Performance Evaluation of an Uncooled Infrared Array Camera

Luke Alexander Pillans

UCL

-

A Thesis Submitted for the Ph.D Examination

In the Field of Instrumentation

-



UMI Number: U591775

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U591775 Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author. Microform Edition © ProQuest LLC. All rights reserved. This work is protected against unauthorized copying under Title 17, United States Code.



ProQuest LLC 789 East Eisenhower Parkway P.O. Box 1346 Ann Arbor, MI 48106-1346 I, Luke Alexander Pillans, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis

Luke Alexander Pillans

•

Abstract

This thesis describes work carried out on an uncooled pyroelectric infrared array camera with the aim of improving performance and increasing its value in commercial markets. The image processing circuitry of the camera was bypassed and replaced by a purpose built 14 bit digitiser and processing algorithms running on a PC. The constructed digitiser was shown to meet the performance needs of the detector. A model was developed for the camera's performance, taking into account the nature of the chopped pyroelectric detector, and the wavelength passband of the camera. The model suggested that placing a temperature sensor close to the chopper blade of the camera would allow radiometric measurements to be made with the camera. Experimental results verified the predicted camera behaviour and radiometric performance was found to be accurate to within ± 1.5 K when imaging flat fields in a stable thermal environment. Significant distortion and radiometric errors were found when imaging high contrast scenes; an algorithm was written to correct this distortion. The algorithm was shown to perform well, drastically reducing distortion and improving radiometric accuracy in all scenes tested. The source of the distortion was not identified, but it is thought to be unrelated to the physical behaviour of the pyroelectric array. The performance of the modified camera is discussed in relation to the current state of the art, and in relation to the performance needs of existing and emerging infrared imaging markets.

Table of Contents

ABSTRACT	3
Table of Contents	4
Table of Figures	
ACKNOWLEDGEMENTS	13
STRUCTURE OF THIS THESIS	14
CHAPTER 1 BACKGROUND TO INFRARED	15
1.1 INFRARED RADIATION & INFRARED WAVE BANDS	
1.1.1 The infrared spectrum	
1.1.2 Near visible infrared, NIR	
1.1.3 Short wave infrared, SWIR	
1.1.4 Medium wave infrared, MWIR	
1.1.5 Long wave infrared, LWIR	
1.2 RADIOMETRY	20
1.2.1 Physical Laws Governing Radiation from Black Bodies	
1.2.1.1 Planck's Law of Black Body Radiation	21
1.2.1.2 Wien's Displacement Law	
1.2.1.3 Stefan-Boltzmann Law	
1.2.2 Emissivity and Related Surface Properties	
1.2.3 Optics	25
1.2.4 Atmospheric Effects	
1.2.4.1 Effects independent of weather and obscurants	
1.2.4.2 Adverse weather conditions and other obscurants	27
1.2.5 The advantages of controlling spectral band pass	
1.2.6 The emissive properties of materials in MWIR versus LWIR	
1.3 Performance measurement	32
1.3.1 How is the performance of IR cameras measured?	
1.3.2 Noise Equivalent Temperature Difference, NETD	
1.3.2.1 Ways of expressing NETD	
1.3.2.2 Temporal NETD, tNETD	
1.3.2.3 Spatial NETD, sNETD	
1.3.3 Minimum Resolvable Temperature Difference, MRTD	
1.3.4 Radiometric performance measurement	
1.3.5 Spatial resolution and Modulation Transfer Function, MTF	
CHAPTER 2, INFRARED TECHNOLOGY AND ITS	
APPLICATIONS	37
2.1 IR CAMERA TECHNOLOGY OVERVIEW	37
2.2 QUANTUM DETECTORS AND ASSOCIATED TECHNOLOGY	38
2.2.1 Quantum detectors	
2.2.2 Cooling technologies	
2.2.2.1 Overview of cooling technologies	
2.2.2.2 Joule-Thomson	
2.2.2.3 Stirling cycle	
2.2.2.4 Peltier coolers	
2.2.2.5. Choosing a cooling technology	
2.2.3 Scanning systems	
2.2.3.1 Conventional scanning	
2 2 UNCOOLED ID INACING	

2.3.1 The rise of uncooled IR imaging	46
2.3.2 Thermal detectors explained	46
2.3.2.1 The concept of a thermal detector	46
2.3.2.2 The time constant of a thermal detector examined through an electr	rical
analogue	47
2.3.2.3 Types of thermal detector	50
2.3.4 Microbolometers	50
2.3.5 Pyroelectric Detectors	54
2.3.5.1 The principles of pyroelectric detectors	54
2.3.5.2 The Design of pyroelectric detectors	55
2.3.5.2.1 Readout chips	55
2.3.5.2.2 The interface between detector and readout	55
2.3.5.2.3 Hybrid technology; bump bonding	56
2.3.5.2.4 Hybrid Technology, adding reticulation	56
2.3.5.2.5 Advanced bump bonding technologies.	57
2.3.5.2.6 Integrated Technology	58
2.3.6 Microcantilevers	58
2.4 INFRARED OPTICS	61
2.4.1 Traditional IR lens in Germanium	61
2.4.2 Moulded Aspheric IR Glass lenses	62
2.4.3 Moulded plastic Fresnel lenses	63
2.5 APPLICATIONS FOR INFRARED CAMERAS	65
2.5.1 Military applications	65
2.5.1.1 The military as a driving force in IR development	65
2.5.1.2 IR vision systems	65
2.5.1.3 Reconnaissance systems	65
2.5.1.4 Automated guidance systems and fuses	66
2.5.2 Law enforcement services	
2.5.3 Emergency Services	66
2.4.4 Civil & Industrial	
2.5.4.1 Surveillance	66
2.5.4.2 Condition monitoring	66
2.5.4.3 Other civilian uses	67
2.5.5 Scientific	68
CHAPTER 3, REPLACING THE EXISTING PROCESSING	
WITH A 14-BIT DIGITISER	69
3.1 ANALYSING THE DESIGN OF THE 100 X 100 PIXEL PYROELECTRIC CAMER	A69
3.1.1 History of the camera	69
3.1.2 Chopper design	
3.1.3 Processing within the camera	
3.1.4 The need for analogue offset removal	
3.1.5 Analysing the weaknesses of the existing processing system	73
3.2 DESIGN AND CONSTRUCTION OF A NEW DIGITISER	74
3.2.1 Selecting an Analogue to Digital Converter chip	
3.2.2 Selection of construction method	
3.2.3 Circuit design of the digitiser	
3.2.3.1 Input stage	77
3.2.3.2 ADC Stage	
3.2.3.3 Output Stage	82
3.2.3.4 Earthing	82
3.2.4 The frame grabber	
3.3 EXPERIMENTAL ANALYSIS OF DIGITISER PERFORMANCE	85

3.3.1 Sources of noise in the digitiser	
3.3.2 Noise in the analogue circuit	
3.3.2.1 Power supply	85
3.3.2.2 Ground noise	85
3.3.2.3 Interference noise	
3.3.2.4 Electronics within the digitiser	
3.3.2.5 Other analogue noise	
3.3.2.6 Analogue noise summary	
3.3.3 Digital noise	
3.3.4 Investigating digitiser noise	
3.3.4.1 Experimental design	
3.3.4.2 Data analysis	
3.3.4.3 Standard deviation	
3.3.4.4 Fourier transform	
3.3.4.5 Inspecting raw samples	
3.3.4.6 Histogram	
3.3.5 Conclusions on digitiser performance	
CHAPTER 4, MODELLING AND MEASUREMENT OF TH	IE
CAMERA'S RADIOMETRIC PERFORMANCE	93
4.1 MODEL OF THE DETECTOR OPERATION	93
4.1.1 Pyroelectric detectors in detail	
4.1.1.1 Time constant and signal	
4.1.1.2 Linearity of thermal impedance	
4.1.2 Thermal environment of the detector element	
4.1.3 Contributions to the signal and the purpose of measuring chopper	
temperature	
4.1.4 Modelling the response of the detector to a black body	
4.2 Experimental Methods	105
4.3 VERIFYING DIGITISER PERFORMANCE AND ASSESSING ARRAY UNIFO	RMITY
	107
4.3.1 Using linear fitting to compare pixel sensitivity	
4.3.1.1 Profiling the distribution of detector sensitivity	108
4.3.1.2 Inspecting the behaviour of pixels with highest sensitivity	109
4.5.1.5 Inspecting the behaviour of pixels with low output gradient	111
4.5.2 Inspecting the benaviour of pixels with the highest raw voltage	112
4.4 ANALYSIS OF DETECTOR PERFORMANCE FOR THE PURPOSES OF	116
4.4.1 Response to IR	116
4.4.2 Behaviour at different camera temperatures	110
4.4.3 Radiometric calibration	
4 4 4 Radiometric performance measurement	121
4 4 4 1 tNFTD	122
4.3.2 sNETD image uniformity	124
4.3.3 Long term accuracy of radiometric calibration	
4.3.4 Effect of changing camera temperature on radiometric performance.	
CHAPTER 5. SCENE CONTENTS CORRECTION	
5.1 EXPERIMENTAL SETUP	129
5.2 Analysis of patterned images	
5.2.1 Vertical divisions	
5.2.2 Horizontal divisions	
5.2.3 Corner occlusions	
5.2.4 Circular apertures	
•	

5.2.5 Multiple edges	142
5.2.6 Analysing the nature of the distortion	144
5.3 CONSTRUCTING A PROCEDURE TO CORRECT FOR SCENE CONTENTS	147
5.3.1 DEFINING THE GOALS OF THE ALGORITHM	147
5.3.1.1 Algorithm flow chart	147
5.3.2 Describing the evolution of voltage offset ΔV mathematically	152
5.3.3 Investigating the relationship between input level and the magnitude of	ΔV153
5.3.4 Finding ΔV for intermediate pixels	155
5.3.5 Relating this to the non-flat field case	: 156
5.3.6 Scene contents correction algorithm summarised	158
5.3.6.1 Summary of Parameters:	159
5.4 PERFORMANCE OF THE SCENE CORRECTION ALGORITHM	161
5.4.1 Performance on flat field data	161
5.4.2 Performance on non-flat field data	162
5.4.2.1 Vertical divisions	163
5.4.2.2 Horizontal divisions	165
5.4.2.3 Corner Occlusions	168
5.4.2.4 Multiple Edges	170
5.4.2.5 Circular apertures	173
5.4.3 Effect of scene contrast on image correction algorithm	175
5.5 Computational requirements	178
5.6 Conclusions	180
CONCLUSION	182
FURTHER WORK	185
REFERENCES	189
LIST OF ACRONYMS	196

.

Table of Figures

Figure 1.1, Atmospheric absorption of infrared radiation over 1 km 16
Figure 1.2, Black body radiation curves
Figure 1.3, Atmospheric absorption of radiation over 10 m
Figure 1.4, Sources of radiation received by a thermal imaging camera
Figure 2.1, Common detector cooling technologies 40
Figure 2.2, Electrical analogue of a thermal detector
Figure 2.3, Example bolometer readout circuit with compensating resistor and deferential output
Figure 2.4, Microbridge Structure
Figure 2.5, Hybrid construction, the bump bonding process
Figure 2.6, Reticulation by laser photochemical etching 57
Figure 2.7, Bump bonding technologies 58
Figure 2.8, Micro-cantilever detector technology
Figure 3.1, Photo of camera
Figure 3.2, Diagram of chopper
Figure 3.3, Overview of processing within the camera
Figure 3.4, Example image from scope73
Figure 3.5, Overview of digitizer78
Figure 3.6, Schematic circuit diagram of digitiser
Figure 3.7, Scope trace showing dead pixels
Figure 3.8, Fourier transform of battery voltage
Figure 3.9, Digitizer measurements of battery voltage, showing sequential and random
readings
Figure 3.10, Digitizer measurements of battery voltage, sequential and non-sequential readings
Figure 3.11, Histogram of battery samples
Figure 4.1, Thermal behavior of pixel during steady state operation

Figure 4.2, Thermal behavior of pixel during dynamic use	
Figure 4.3, Diagram of detector chamber, showing position of chopper, lens and PRT relative to the pyroelectric array	
Figure 4.4, Photograph of PRT sensor sited inside detector cavity 101	
Figure 4.5, Atmospheric transmission of radiation over 1 m, (32°C and 75% humidity)	
Figure 4.6, Black body radiation power within defined passbands	
Figure 4.7, Response of one pixel with linear fit used to measure relative sensitivity . 107	
Figure 4.8, Histogram of linear fit gradients108	
Figure 4.9, Image of sensitivity (white high)109	
Figure 4.10, Binary image of sensitivity, highest in black	
Figure 4.11, Response curves for the 64 Pixels with highest gain (image difference) 111	
Figure 4.12, Response curves for the 64 Pixels with lowest gain (image difference) 112	
Figure 4.13, Binary image of responsivity, lowest performing pixels in black 112	
Figure 4.14, Response curves for the 64 pixels with highest raw voltage readings 114	
Figure 4.15, Pixels contributing to the 100 highest raw readings 115	
Figure 4.16, Image of raw detector output (white high)115	
Figure 4.17, Linearity of detector response to black body when black body temperature is converted radiation intensity using different models for the detector's spectral response. Linearity is quantified by calculating correlation coefficients	
Fig 4.18, Pixel response to black body temperature with theoretical 8 - 12 µm power curve overlaid	
Figure 4.19, Detector response at different camera temperatures	
Figure 4.20, Graphs of mean tNETD vs. Target temperature, for all camera temperatures, error bars indicate standard deviation across all pixels	
temperatures	
Figure 4.22, Image of tNETD 124	

Figure 4.23, Graph of mean sNETD vs. ΔT , all camera temperatures, with 2nd order
polynomial line of best fit125
Figure 4.24, Calibration stability over twenty months, error bars cover the range of
errors recorded on each
Figure 4.25, Radiometric error when varying camera temperature rapidly
Figure 4.26, Radiometric error when varying camera temperature slowly
Figure 5.1, Images of vertical divisions
Figure 5.2, Graphs of vertical sections though vertical division images in Figure 5.1.132
Figure 5.3, Graphs of horizontal sections though vertical division images in Figure 5.1
Figure 5.4 Images of horizontal divisions 135
Tigure 5.4, mages of nonzontal divisions
Figure 5.5, Graphs of vertical sections though horizontal division images in Figure 5.4
Figure 5.6, Graphs of horizontal sections though horizontal division images in Figure
5.4
Figure 5.7, Images of corner occlusions137
Figure 5.8, Sections though corner occlusion image 5.7 D
Figure 5.9, Images of circular apertures
Figure 5.10, Sections though circular apertures in figure 5.9141
Figure 5.11, Images containing multiple edges
Figure 5.12, Sections though multiple edge images in figure 5.11
Figure 5.13, Graph of 1500 consecutive pixels, starting above the horizontal division in
image 5.4C
Figure 5.14, Image difference level vs. Row number (vertical section) at three black body temperatures
Figure 5.15 Raw signal in open and closed fields with pedestal removed (vertical
section)
Figure 5.16, Logarithmic fits to voltage decay at start of frame vs. Row number. The

Figure 5.17, Swing vs. Mean Vm. Swing is calculated from the logarithmic fits in figure
5.16 and Vm is measured for each combination of black body and camera temperature.
Figure 5.18, Vm and Vi vs. Row number for a single flat field frame
Figure 5.19, Swing vs. Masked input level, the masked input level is a weighted average of the values in the previous frame
Figure 5.20, ΔVt vs. Mean Vi, ΔVt is the "target" value towards which ΔV decays, a line of best fit is calculated which is later used to calculate ΔVt for individual pixels in
an arbitrary scene
Figure 5.21, Block diagram of scene contents correction algorithm
Figure 5.22, Vertical sections through flat fields with and without scene contents correction. As intended the correction has very little effect on flat field data
Figure 5.23, Images of flat fields calibrated with and without scene contents correction
Figure 5.24, Images of vertical divisions with and without scene contents correction 164
Figure 5.25, Scene correction multi-graphs, vertical sections through image 5.24 B 165
Figure 5.26, Images of horizontal divisions with and without scene contents correction
Figure 5.27, Scene correction multi-graphs, vertical sections through images from figure 5.26
Figure 5.28, Images of corner occlusions with and without scene contents correction
Figure 5.29, Scene correction multi-graphs, vertical sections through image 5.28 B 170
Figure 5.30, Images of multiple edges with and without scene contents correction 171
Figure 5.31, Scene correction multi-graphs, vertical sections through image 5.30 A 172
Figure 5.32, Images of circular apertures with and without scene contents correction173
Figure 5.33, Scene correction multi-graphs, vertical sections through image 5.32 A 174
Figure 5.34, Scene correction multi-graph, vertical section through image 5.32 B, column 47
Figure 5.35, Images of angled plate with and without scene contents correction 176

Figure 5	5.36, 8	Scene	correction	multi-graphs,	vertical	sections	through	images	from	figure
5.35				•••••••••••••••••••••••						177

,

•

Acknowledgements

The work presented here was made possible by the funding of Sira Ltd, QinetiQ, the EPSRC and the postgraduate training partnership. I would like to thank Dr Ian Furniss of UCL and Dr Gordon Hopkinson of Sira for their help and supervision throughout this work. Dr Mark Skipper of Sira, for his help in the design of the digitiser. Also Dr Roger Young and Dr Nicholas Pillans for their help in proof reading.

Structure of this thesis

This thesis contains two chapters covering background theory and current technological trends, and individual chapters covering the three areas of practical research which were pursued during the course of this study.

Chapter 1 covers the principles of infrared radiation, the science of radiometry and the performance parameters generally used to characterise infrared imagers. Chapter 2 covers the theory of operation of different infrared detector technologies, and discusses the current state of the art in uncooled detector technology. Established and new markets for infrared imagers are also discussed in Chapter 2.

Chapter 3 describes the design, construction and testing of the digitiser built by the author to replace the existing processing electronics within the camera. Chapter 4 establishes a model for the camera and how it may be used for radiometry; this model is tested experimentally and found to accurately predict the camera performance. Chapter 5 describes a previously undocumented form of distortion found in the output of the detector and explains an algorithm which was developed in the course of this research to remove this distortion.

The Conclusion and Further Work sections of the thesis draw together the results of all three experimental sections and discuss their relevance to the infrared applications and markets described in Chapter 2.

Chapter 1 Background to Infrared <u>1.1 Infrared radiation & infrared wave bands</u>

1.1.1 The infrared spectrum

The term infrared (IR) does not refer to a strictly defined band in the electromagnetic (EM) spectrum, the demarcation boundaries are somewhat fuzzy. The word infrared can be used to describe electromagnetic radiation in a wavelength range starting at the upper edge of visible light, at around 700nm, out to the edge of the microwave band, which is generally taken to be about 1mm. Earthbound IR imaging and radiometry is concerned almost entirely with the lower end of that spectrum, below 15 μ m.

Atmospheric absorption defines three generally accepted bands within the range 700nm – 15 μ m. The graph in figure 1.1 shows typical percentage atmospheric transmission through 1km, across the infrared spectrum^{R1}. There are two major absorption features considered to be of particular importance: the first from 2.5 - 3 μ m and the second from 5.5 - 8 μ m. The first is caused by absorption by carbon dioxide and water molecules, the second by water alone. These absorption features are commonly used to define four major passbands (or 'windows') with the infrared spectrum. These are termed Near visible InfraRed, NIR (700nm - 1.0 μ m), Short Wave InfraRed, SWIR (1.0 - 2.0 μ m), Medium Wave InfraRed, MWIR (3 - 5 μ m), and Long Wave InfraRed, LWIR (8 - 12 μ m). Infrared radiation extends to much longer wavelengths than are covered by this definition of LWIR and there is imaging equipment designed to operate in this region, sometimes termed far infrared or sub-mm band. However, these definitions are standard within terrestrial infrared imaging and will be used here. In addition to the three bands given, there are many more fine distinctions made within specialised areas of use, but they are beyond the scope of this thesis.

Atmospheric transmission is only one of several factors that lead to the partitioning of the infrared spectrum, of comparable importance is the behaviour of the detector materials used. Many detectors only respond to radiation within a certain band, lens materials and coatings can also be wavelength specific. When these factors are combined it is found that very few configurations will allow an infrared camera to operate over more than one of the atmospheric passbands, as a result the bands segregate infrared technology. Again there are exceptions to this rule but they are not of specific relevance to this work.



Figure 1.1, Atmospheric absorption of infrared radiation over 1 km

Radiation within each band interacts differently with materials it encounters in the environment (including the atmosphere) and is emitted at different levels by different bodies, according to Planck's Law and the wavelength dependent emissivity of the material. Images of the same scene in different wavebands may appear markedly different. Certain features may be prominent in one band and less so or perhaps completely invisible in another. Consequently, cameras operating in the different bands are suited to different applications, this discussed further in Chapter 2.

1.1.2 Near visible infrared, NIR

It has been established that the infrared wave band extends directly from the edge of visible light. The proximity of NIR to visible light, means that in some cases the same optical materials and detectors used for visible light may be used for NIR, with little or no modification. Many silicon based CMOS or CCD detectors designed for visible light function acceptably in NIR and are used in the field. As a result there are more comparatively cheap NIR systems available on the market when compared to the other IR bands. NIR performance is sometimes built into widely available video cameras aimed at the consumer market; it is marketed as a form of night vision.

In the NIR band radiation interacts with many materials in the same way as visible light, more so than in the other IR bands. Generally, the same non-metallic materials exhibit high transmission or high reflection in NIR as visible light, this is less true in the other infrared bands. When imaging in NIR most of the radiation received from an object will be radiation reflected from its surface. There is very little thermal emission of radiation in the NIR band by objects at ambient temperature, which is somewhere between -40 to +40 °C across most of the globe, consequently, in darkness some form of illumination is needed, (thermal emission is discussed in more detail in section 1.2).

The small infrared emitters used (often LED's) need not produce any visible light, so NIR systems can be used in covert surveillance and security systems but they are not as attractive in military use because the infrared source can be detected by anyone else with appropriate equipment. Many NIR systems have more in common with visible light cameras than medium and long wave infrared cameras. As such they are not of much interest when assessing the performance of a long wave IR camera, like the one used for this project. Throughout the rest of this thesis, the term infrared cameras will be used to refer to medium and long wave cameras and not NIR cameras, unless specifically mentioned.

1.1.3 Short wave infrared, SWIR

SWIR imaging is dominated by Indium Gallium Arsenide (InGaAs) detector technology. InGaAs is a semiconductor with a band gap suitable for making photodetectors sensitive to wavelengths up to, and in certain cases someway past, 2 μ m, depending on technology used. SWIR detectors are commonly available in high resolutions (VGA and above), they offer wide dynamic range for day and night use, and relatively low noise when compared to longer wavelength imaging. SWIR optics utilise largely the same materials as visible light optics but lens designs and coatings will tend to be dedicated to SWIR use.

Thermal emission in the SWIR band is low at environmental temperatures; in general purpose imaging objects will be identified by reflected rather than emitted light. In daylight few objects will be seen as strong emitters, a surface temperature of hundreds of Celsius is required to provide enough contrast against daylight. At night star light and sky glow can provide enough illumination for clear imaging of the environment. The natural environment and human bodies will still not be detected by emitted radiation, but many types of operating machinery will be seen by their emission; hot exhausts make land vehicles and planes easy to detect. Such sources also provide significant illumination of surrounding objects. Passive SWIR is used by the military for covert observation; the higher resolutions available can offer advantages over MWIR and LWIR.

SWIR sources are not detected by NIR silicon imagers or normal light amplification, and SWIR lasers are considered "eye-safe" (they will not harm the eyes under normal exposure). As such SWIR provides possibilities for illuminated night vision and laser target illumination offering a generational tactical advantage over earlier night-vision technologies.

1.1.4 Medium wave infrared, MWIR

When compared to SWIR, MWIR wavelengths are much closer to the peak in black body thermal emission at ambient temperatures (see Section 1.2). When looking at the world through a medium wave camera, much of the radiation picked up from objects in the scene will be the result of thermal emission by those objects, rather than radiation reflected from their surfaces. Consequently, a scene may appear much the same in the middle of the night, as in middle of the day. Medium wave cameras can be used for night vision in environments with complete (visible) darkness, without any form of illumination, thus making MWIR ideal for covert imaging. Atmospheric absorption of MWIR is very low, so it is possible to see over long distances (compared to LWIR). The spectrum of radiation emitted by the sun contains much more energy in the MWIR band than in LWIR, as a result much more solar MWIR radiation reaches the Earth's surface. This solar radiation is reflected by low emissivity surfaces so large amounts sun glare can arise from water and metallic objects. This can be a disadvantage of MWIR compared to LWIR in such applications and particularly in radiometry.

1.1.5 Long wave infrared, LWIR

In typical LWIR usage, most the radiation detected by a camera originates from thermal emission by the environment, even in daylight. Consequently, differentiating between hot and cold objects is often very easy, warm blooded animals and operating machinery stand out against the environment. This makes LWIR of particular use in military and law enforcement applications, it also makes long wave infrared well suited for performing radiometry.

Another advantageous property of long wave infrared is its reduced interaction with fine particles in the atmosphere. The interaction of EM radiation with small particles is dependent on the size of the particle relative to the wavelength of the radiation. When the size of a particle is smaller than the wavelength of radiation, their interaction is limited, thus LWIR will be transmitted through a mist of particles roughly twice as large as those which obscure MWIR. It so happens that this allows LWIR to see through many forms of smoke (and to some extent mist) which MWIR will not. Conversely, one of the major drawbacks when operating in long wave is the absorption of radiation by moisture in the atmosphere. This reduces the maximum range of a long wave camera to only a few kilometres in humid environments.

•

•

1.2 Radiometry

Radiometry is defined as the measurement of the intensity of electromagnetic radiation received from a body. Within the field of thermal imaging, the term is closely associated with the measurement of temperature. The amount of radiation emitted by a surface is dependent on its temperature and a surface property known as emissivity. A black body is a hypothetical object, a perfect absorber of radiation and thus also a perfect radiator with an emissivity of one. The radiance of a black body is given by Planck's Law of black body radiation, it is dependent on the body's temperature alone. Therefore measuring the radiation emitted by a black body allows its temperature to be calculated.

With an ideal camera, the thermal radiation intensity incident upon a detector is directly proportional to the radiant intensity leaving the surface of an object in the image, regardless of the distance to the object. (The caveat 'with an ideal camera' may suggest that in real world case this statement may be far from the truth but for most cameras it holds very well, limitations are discussed later.) It follows that the surface temperature of a black body can be measured by an imaging detector which is suitably calibrated. These concepts are covered in more detail in the rest of this section.

1.2.1 Physical Laws Governing Radiation from Black Bodies

1.2.1.1 Planck's Law of Black Body Radiation

Planck's Law expresses the spectral radiance of a black body as a function of temperature. The radiation emitted by a black body is not distributed evenly acros's the electromagnetic spectrum; the spectral profile is dependent on temperature and explains the relevance of the infrared waveband to radiometric temperature measurement.

Planck's Law^{R2}:

$$I(\lambda,T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{\frac{hc}{\lambda kT}} - 1}$$

Where:

 $I = Spectral Radiance, Js^{-1}m^{-2}sr^{-1}m^{-1}$

T = Temperature of the black body, K

 λ = Wavelength, m

c, speed of light = $2.998 \times 10^8 \text{ m s}^{-1}$

- e, base of the natural logarithm = 2.718
 - h, Planck's constant = 6.626×10^{-34} J s
 - k, Boltzmann's constant = 1.381×10^{-23} J K⁻¹

Planck's Law produces characteristic curves of spectral luminance for bodies of different temperatures. A set of these curves covering the range from 250 to 400 Kelvin is given in figure 1.2. The curves all peak between 6 and 12 μ m, showing that IR is the ideal waveband in which to make radiometric temperature measurements of the environment at terrestrial temperatures. As temperature increases, the peak wavelength of the black body radiation decreases, at temperatures of around 700K it becomes possible to make radiometric temperature measurements with visible light. The amount of IR radiation emitted by a black body continues to increase even as the peak output moves out of the IR band, so infrared radiometry can still be used for very hot objects.



Figure 1.2, Black body radiation curves

1.2.1.2 Wien's Displacement Law

As black body temperature increases the emission spectrum moves towards higher photon energies, so there is an inverse relationship between T and the peak emission wavelength. The wavelength of peak emission of a black body is given as a function of temperature by Wien's Displacement Law^{R3}:

$$\lambda_{\max} = \frac{b}{T}$$

Where b is Wien's displacement constant = $2.898 \times 10^{-3} \text{ m K}$

1.2.1.3 Stefan-Boltzmann Law

Total radiated power of a black body across the entire EM spectrum is given by the Stefan-Boltzmann Law^{R4}:

$$j^* = \sigma T$$

Where:

 j^* = total radiated power per unit area, W m⁻²

T = temperature, K

σ, Stefan-Boltzmann constant = $5.670 \ge 10^{-8} \ge 10^{-8} \le 10^{$

There is a fourth order relationship between the total radiated power and black body temperature, but it is important to realise that no detector will collect the full spectrum radiated by the black body. Absorption by the atmosphere and optics combined with the spectral sensitivity of the detector mean that only a certain percentage of the energy emitted by a black body will be detected, the exact percentage will vary according to black body temperature. As a result the fourth order relationship breaks down and the detected signal does not rise as quickly with temperature as the Stefan-Boltzmann Law would imply. This effect is discussed in more detail later in this chapter.

1.2.2 Emissivity and Related Surface Properties

At a basic level any surface has three principle optical properties; emissivity, reflectivity and transmittance, these are dimensionless quantities which can take any value between zero and one. Integrated over all angles the sum of these parameters equals one; this is a statement of conservation of energy. Therefore a perfect black body with an emissivity of 1 has zero reflectivity and zero transmission.

There follows a list of terms and definitions commonly used to describe surface properties^{R5}:

Emittance, ε ; the ratio of the radiant power emitted by a body as a consequence of its temperature alone, to the corresponding radiant power of a black body at the same temperature. Therefore emittance is dependent on the physical characteristics of the body in question, it is not directly a property of the surface material of the body.

Absorptance α ; the ratio of the radiant energy absorbed by a body to that incident upon it. A black body has an absorptance of 1. Absorptance is equal to emittance.

Reflectance ϱ ; the ratio of the radiant energy reflected by a body to that incident upon it. A black body has a reflectance of zero.

Transmittance τ ; the ratio of the radiant energy transmitted through a body to that incident upon it. A black body has a transmittance of zero.

Emissivity ε ; this is one of the most widely quoted surface properties in the field of infrared imaging. Emissivity is sometimes defined as the emittance of a body of material which has an optically smooth surface and is sufficiently thick as to be opaque.^{R5} By this definition emissivity is a fundamental property of a material, any surface treatment (e.g. textured surface) or shaping (e.g. cavity emitter) may increase an

object's emittance, but cannot increase the emissivity of the surface material. However, common usage of the word emissivity often does not match this definition; literature will often refer to the emittance of an object or material as its emissivity. Examples include radiating sources specially shaped to achieve high emittance, and the behaviour of bulk materials such as vegetation. Some sources use the terms "equivalent emissivity" or "nominal emissivity". Pragmatically one has to accept that the words emittance and emissivity are to some extent interchangeable. Within this thesis emissivity is used to describe both material properties and object properties.

Reflectivity (ϱ) and **absorptivity** (α) are to reflectance and absorptance as emissivity is to emittance. Transmittance has no equivalent property in this context because a sample of material thick enough to be opaque has no transmission by definition.

All of these parameters are wavelength dependent; it is after all the wavelength dependent reflectivity of objects which allows us to see them in colour. Emissivity, reflectivity and transmittance must still add up to one at every wavelength.

In many cases emissivity is assumed to be equal across a waveband of interest, this is termed a grey body assumption. It is often very practical to use a value of emissivity averaged over the passband of the detector in question, but care should be taken about reusing the same figure when making calculations for a different system. Also, an averaged value of emissivity may be effective for calculating thermal emission, but when it comes to calculating absorption or reflection the spectrum of the incoming radiation becomes important. Using the average value of emissivity for an object illuminated with monochromatic radiation could give rise to large errors in calculation, even within a narrow waveband.

Emissivity is clearly of paramount importance when making radiometric temperature measurements; as emissivity drops below one the amount emitted radiation falls in line with it, but the total radiation emanating from an object may not. As reflectivity rises more reflected radiation emanates from the object. If the surroundings are hotter than the object more radiation will leave its surface than would a black body's surface. So, to an infrared imager, an object with emissivity lower than one may look either hotter or colder than its true temperature.

To make a radiometric temperature measurement of an object with emissivity lower than one the amount of reflected radiation must be calculated or approximated. A common method of approximation is to assume that the reflected radiation will be that of a black body at the environmental temperature multiplied by the reflectivity of the object being measured. Then environmental temperature must then be calculated, this is sometimes done by simply taking a measurement of ambient temperature in the vicinity of the detector. This calculation assumes that the detector is seeing a diffuse reflection of the environment in the surface of the object, however if the object exhibits specula reflection the reflected light could be very different, reflected sunlight being an obvious example.

In general, the lower an object's emissivity the more difficult temperature measurement will be. At the lower end of the market many single point infrared thermometers are calibrated with a fixed correction assuming an emissivity of 0.95, which approximates many high emissivity materials commonly found in the environment. This works well in many applications, care must be taken when choosing which surface to measure, but that is a fundamental rule in radiometric temperature measurement. It should be noted that the IR reflectance of materials does not necessarily correlated to the visible band reflectance; it is dangerous for a user to make assumptions about reflectance. To make accurate temperature measurements of objects with lower emissivity it is important not only to know the emissivity but also to make the measurements in a well controlled, or at least well understood, environment.

1.2.3 Optics

The Lagrange invariant of an imaging lens system relates the size of an object in the image plane to its real size and its distance from the lens. It states that if an object moves twice as far away its linear size in the image will halve^{R6}. The result is that as a target moves further away the reduced amount for radiation reaching the detector is exactly matched by reduction in image size and the flux incident on the detector in the section of image occupied by the target remains constant. To use an analogy from visible light photography, exposure does not have to be compensated for the distance to the subject.

As with any camera the only way to increase the flux at the detector is to decrease the fnumber of the lens, aside from the general requirement for larger and more expensive lens elements this has the disadvantage of giving a shallow depth of field. However, most infrared imagers are of comparatively low resolution compared to visible light cameras and correspondingly the effective depth of field is higher. An image will not appear out of focus until the resolution of the image projected onto the detector approaches the pixel pitch.

There are many choices when it comes to lens materials for infrared optics, the most common are covered in detail in Chapter 2, but whatever material is chosen it will have its own spectral response. Percentage transmission through a lens is will vary with wavelength; this variation is determined by the optical properties of the lens material, the antireflective coating used and the lens design. Refractive index is also a function of wavelength, this means that there are limits to the resolution that can be achieved when the entire passband of the lens is used, it is impossible to focus the entire spectrum radiated by a target to the same point on the detector. Not only does this limit resolution, it also has implications for radiometry, particularly with smaller targets. Depending on which wavelength is most in focus, hot and cold objects will be differentially blurred, if the object size is comparable to the radius of the blur, linearity will be affected. These problems can be reduced if the passband of the system is limited using a wavelength dependent filter. The filter can be placed either in front of the lens or between the lens and detector. The wavelengths reaching the detector can be restricted to a band for which the lens has a consistent refractive index, the defocusing is avoided and resolution can be improved. Constraining the passband can have other benefits for radiometry, total signal may be decreased but the adverse effects of atmosphere, and reflected radiation can be reduced. This is discussed later in the chapter. It is possible to produce low pass, high pass and band pass filters by adding appropriate doping elements to infrared optical materials. This is equivalent to making colour filters for visible wavelengths.

1.2.4 Atmospheric Effects

1.2.4.1 Effects independent of weather and obscurants

Figure 1.1 shows the atmospheric transmission of infrared over a distance of one kilometre and illustrates how absorption attenuates transmitted radiation over long distances. Figure 1.3 gives a graph of atmospheric transmission over a distance of one meter, it shows that percentage transmission is much higher at shorter distances, this is to be expected. If at any given wavelength the fractional transmission through one meter of atmosphere is T, the transmission over a distance d will be equal to T^d. One kilometre of atmosphere can be thought of as 1000 successive one meter thicknesses of atmosphere.



Figure 1.3, Atmospheric absorption of radiation over 10 m

From the discussion of emittance and absorptance in section 1.2 it is evident that as an absorber of radiation the atmosphere must also be an emitter of radiation. The atmosphere between the target and the detector will attenuate both emitted and reflected radiation and contribute a third source of radiation. The various sources of radiation arriving at a camera are illustrated in figure 1.4 ^{R7}.



Figure 1.4, Sources of radiation received by a thermal imaging camera

The amount of radiation absorbed by the atmosphere and the amount radiated by it are both dependent on the distance between camera and target. As far as imaging is concerned this means that the contrast between a target and its surroundings will be reduced when viewed from a distance. As distance increases there will come a point at which the target is no longer distinguishable from its background because of sensitivity limits of the camera, (that is assuming the target does not shrink below the spatial resolution of the camera before then).

1.2.4.2 Adverse weather conditions and other obscurants

Atmospheric transmission introduces a distance dependence to radiometry, not only that, also a dependence on weather; a large part of atmospheric absorption is due to water vapour in the atmosphere so radiometry will be affected by humidity even without precipitation. Any precipitation or other particles in the air will affect IR visibility, but not in exactly the same way as they do in the visible band.

Weather conditions such as rain, fog, and dust can affect performance at short ranges. Over shorter distances radiometric accuracy will be diminished and over larger distances image quality will be degraded. The obscuring effects of rain and fog at visible wavelengths will be familiar to everyone, but in the MWIR and LWIR these effects are reduced, thanks largely to the longer wavelengths. This is one of the major benefits of IR imaging in both military and civilian roles. (See section 2.3.) Dust storms and sand storms are weather phenomena with which many in the western world are unfamiliar, but they occur regularly in many environments across the globe and can have drastic effects on visibility. Dust and sand can have significant effects on imaging in the IR, cutting visibility significantly. Generally, MWIR is affected more strongly than LWIR, the severity of the image degradation is linked closely to the size of the dust particles ^{R8}. Fine dust will have only a small effect but sand with large grain sizes will affect IR imaging almost as badly as it does visible imaging.

In addition to these naturally occurring phenomena there are many man made local environmental conditions which can affect visibility, smoke from fires is a common example. Most types of smoke affect IR visibility to a much lower extent than visible light. This is one of the benefits of IR imaging in fire fighting, and also why LWIR is often preferred in that application.

In military environments smoke is often deliberately generated to obtain tactical advantage; in the military field the products used are referred to as obscurants. Many different methods are used to produce visible wavelength obscurants, from burning heavy oils, known as fog oils, to burning phosphorus. All of these obscurants have less effect on IR imaging than on visible wavelengths; yielding a very significant tactical advantage of IR imaging over light amplification.

It is in the nature of military technology that as soon as a new technology becomes available, work starts on ways to circumvent or negate its advantages. Infrared vision is no exception. Obscurants have been developed explicitly for the infrared which may be used alone or in conjunction with visible light obscurants. Already in use are brass flakes and graphite flakes, and other types are being developed. The principle is similar to smoke based obscurants, generating a suspension of reflecting or absorbing particles in the air. In the infrared, particles must be larger in size to interact with the longer wavelengths. The problem is that larger particles tend to precipitate out of the air too quickly, hence the use of very fine flakes which have a large surface area in relation to their mass. Despites this, IR obscurants still dissipate more quickly than smokes and tactics must be changed slightly to compensate. It is also worth noting that the most widely used IR obscurant, graphite flakes, are toxic if inhaled and use of respirators is advised whenever they are encountered.

1.2.5 The advantages of controlling spectral band pass

It has been shown that all elements involved in forming an IR image have some spectral dependency, the emission from a target, transmission through the atmosphere, transmission through the lens and the sensitivity of the detector itself are all dependent on wavelength. For a given arrangement of camera and target these factors can be combined to give a single graph of sensitivity vs. wavelength. Taking the product of this curve and the black body emission spectrum for the target's temperature and integrating across it will yield a value for the detected radiation power as a proportion of the ideal case. (The ideal case being a black body radiator, no atmosphere, optics with 100% transmission and a black body absorber for the detector.) However, this factor will vary with target temperature, as the black body spectral profile changes the total efficiency of the system changes. The response of the camera to target temperature will not follow the 4th order relationship of Wien's Displacement Law, the exact relationship will depend on all the factors listed above. This is not necessarily an issue if the overall goal is simply infrared imaging, but for radiometry this relationship must be known. This can be achieved in two ways: one is to measure the spectral response and perform the integration described above, and the other is to simply perform an empirical calibration of the system as a whole. Each approach has its merits and these will be discussed later.

Once the relationship between target temperature and signal is determined a problem remains, the system is sensitive to changes in any of the elements contributing to the total passband, including changes in the atmosphere. Since the spectral transmission of the atmosphere is dependent on distance, any calibration will change when the distance to the target changes. This is clearly far from ideal and is not something that can be entirely avoided but the situation can be significantly improved by limiting the spectral response of the detector. Consider the graph of atmospheric transmission in Figure 1.1; one of the most noticeable features is the near total absorption between 5 μ m and 8 μ m. If a broadband detector is fitted with a filter which blocks all wavelengths below 8 µm the band where the atmosphere absorbs and emits most radiation is cut out. The detector will receive less radiation from the target, particularly at short ranges, so the signal is reduced, but the effect of distance to target on the signal is vastly reduced. This is clearly of great benefit for radiometry, but is also beneficial for general imaging in many applications; this is because as well as absorbing radiation, the atmosphere is also a source of radiation. Viewed at close range a scene will have a certain level of contrast between hot and cold areas, if the same scene is viewed from a larger distance the signal contribution from the atmosphere reduces the contrast. Adding the filter reduces the change in contrast with distance, it does not improve the sensitivity at long distances, rather it reduces the sensitivity at short distances. Since it doesn't improve the sensitivity, the filter cannot improve the temperature resolution of distant objects, but it does allow elements of the scene in the foreground and elements of the scene in the distance to be imaged with appropriate contrast at the same time. Without the filter, if the display contrast is increased to obtain good temperature resolution on the distant objects the foreground objects may well be clipped at both the top and the bottom ends of the display's dynamic range.

The case of a high pass filter at 8 μ m is by no means the only useful option but it is an example of particular relevance to making radiometric measurements with a board band detector.

There are additional benefits to wavelength filters. The problem of a lens' refractive index varying with wavelength was discussed earlier. If the passband of the detector is reduced then only the lens' performance over that band is important, so it becomes easier for the designer to choose a material which has uniform refractive index over the band in question. Hot objects and cold objects will be brought into good focus and spatial resolution can be improved.

Choosing the passband of the camera also allows the designer to avoid wavelengths where reflectivity may cause particular problems for radiometry. The most general example being the solar radiation that is present in MWIR, the 8 μ m low pass filter described above will remove the problem of reflected solar radiation disturbing radiometric temperature measurements. The same principle can be applied in specialised cases, for example an application which requires accurate radiometry of a

specific material. Filters may be chosen to constrain the detector's sensitivity to a narrow band within which the material in question has high emissivity. Such specialised radiometric systems are more often found in the field of single point infrared temperature measurement than they are in infrared imaging.

1.2.6 The emissive properties of materials in MWIR versus LWIR

A vast amount of work has been done measuring the emissive properties of different materials. Figures will not be listed here but one catalogue of such data can be found in this reference ^{R9}. The conclusion reached from studying these data is that many of materials which will be observed with infrared cameras have higher emissivity in the LWIR band than the MWIR band. These materials include: Soils, vegetation, asphalt, concrete, wood, and most interior and exterior paints. In many cases emissivity in LWIR may be well over 90% while it is under 50% in MWIR. From this it can be concluded that the LWIR band is most suitable for performing general purpose radiometry for temperature measurement. This is not to say that LWIR imagers are superior to MWIR imagers, it is simply a matter of fitness for purpose.

The lower emissivity of materials in MWIR offers advantages in some applications; military studies have found that whilst target detection ranges are greater in LWIR, target identification ranges are significantly greater in MWIR^{R10}. Target detection range is defined as the range at which a target's presence can just be detected, without identifying it ^{R10}. Higher emissivity gives warm targets greater contrast against the environment in LWIR accounting for its longer detection range. Target identification range is defined as the range at which the nature of a target can be determined ^{R10}. For example, is the target seen a tank, a jeep, or a tent? It has been found that the lower emissivity of materials aids target identification in MWIR; light reflected from the surface of objects picks out surface details and increases contrast across the surface of an object, improving the chances of identification. (This comparison assumes similar spatial resolution available for the imagers in each band; in many situations the choice may not be so simple. Spatial resolution is paramount to detection and recognition, so in real world cases the best tool for the job may be determined by resolution rather than waveband.)

<u>1.3 Performance measurement</u>

1.3.1 How is the performance of IR cameras measured?

The are many ways in which to classify the performance of imaging systems. Since the scope of this work is limited to one camera, many of these performance variables are fixed, for that reason certain parameters were not investigated. The performance of infrared imagers is often limited by noise levels and there are several industry standard measurements which are described here. Where possible and useful these measurements were calculated for the camera used in this project, the results are presented in later chapters.

1.3.2 Noise Equivalent Temperature Difference, NETD

1.3.2.1 Ways of expressing NETD

Noise equivalent temperature difference (NETD) is one of the most widely used parameters for the comparison of infrared imaging systems. As the name suggests the NETD is a measurement of the noise on the signal expressed as a value on a temperature scale^{R11}. The matter is complicated by the fact that the noise on a video stream can be expressed in several ways each of which gives useful but different information about the quality of output, and NETD acts as a scale upon which to express noise measurements. When comparing values of NETD between different imaging systems it is important to understand how each figure is calculated to ensure values are compared like for like.

In later chapters several different NETD figures are given and many of these are calculated several times for different camera conditions. The different measurements of NETD and their merits are discussed here to avoid confusion later. Whilst these different measurements of NETD are calculated independently they are related, random noise on pixels will contribute to all measurements of NETD so these measurements should be examined with reference to each other.

1.3.2.2 Temporal NETD, tNETD

Temporal NETD (tNETD) is the noise measured on the changing value of one pixel, over a series of sequential frames. It serves as a measure of the stability of the image over time and is often the dominant form of noise in the output of an imager. In a single frame of output from an IR camera, tNETD may be the limiting factor in target recognition. Of the various measurements of NETD this most closely relates to the random errors on individual samples because it does not take into account any measurement of fixed pattern noise.

Temporal NETD is commonly expressed as an average value across the detector. This is the mean of the tNETD values calculated independently for each pixel. This figure is the most commonly used measurement of NETD and is often given alone and described as the NETD of a detector.

In some situations the tNETD on the mean value of a group of pixels can be a useful measurement, for example when comparing the performance of detectors of different resolutions and predicting a detectors performance for specific tasks. Taking the mean output of many pixels will usually yield a lower value of tNETD so whenever such measurements are used it is important to state clearly how they are calculated.

1.3.2.3 Spatial NETD, sNETD

sNETD is the noise measured across all the pixels of one frame, so it must be measured by imaging a flat field. It is useful as a gauge of how noisy a single frame will look when viewed as a static image. The sNETD will be made up of the random errors on each pixel (as measured by their tNETD) plus any fixed pattern noise on the detector. Averaging many consecutive fames of video can reduce the contribution of tNETD to the image and allow fixed pattern noise to be measured alone. Large amounts of fixed pattern noise can be particularly objectionable or even confusing to the operator. Fixed pattern noise can, hardware permitting, be measured and calibrated, (or re-calibrated), out of the image. In some types of detector, fixed pattern noise can be induced by image burn-in, in such cases sNETD is one method of quantifying the effect.

sNETD may sometimes be measured in only one axis, horizontal or vertical, the calculation follows the same principal but includes only the pixels in one row or one column. In an ideal camera one might expect the horizontal and vertical measurements to provide similar results, but in reality practical elements of the detector and camera design, (for example scanning systems), often result in significant differences between horizontal and vertical sNETD measurements.

1.3.3 Minimum Resolvable Temperature Difference, MRTD

MRTD is another figure often quoted for infrared imaging systems; it is an expression of the minimum temperature difference between two areas of the image which can be distinguished by the human eye. It is a useful figure because it allows a very direct comparison of how different systems will perform under the same circumstances.

MRTD is always used as a measurement of a complete system including the camera and the display hardware, measured subjectively by a human observer. This subjectivity causes problems in the comparison of MRTD measurements made by different people and at different times. Work as been done to make the measurement of MRTD as objective as possible, but the measurement is still dependent on the observations of a user and is subject to certain errors as a result ^{R12}.

MRTD is often measured as a function of spatial frequency, generally by performing the same test with bar patterns of different sizes; this makes spatial resolution a factor in the result. There is no globally accepted standard for MRTD tests so it is difficult to compare results for different cameras unless the test conditions are known to be consistent.

Since MRTD is a measurement of an entire system including display hardware it is not a useful figure in the context of this work. The aim of this project has not been to produce a complete imaging system but rather to analyse the performance of the detector used.

1.3.4 Radiometric performance measurement

The most common ways of expressing radiometric performance are as a percentage error or as an error on the temperature scale such as plus or minus half a degree. However there are a number of different ways in which a radiometric value can be calculated, with one pixel or many, from one frame, or integrated over time, thus it is difficult to predict what level of accuracy will be obtained without knowing the exact situation and the nature of the camera.

For example, in certain scientific applications such as satellite imaging it may be important to calculate a temperature value for each pixel in a single frame. This is the most extreme case where the calculated temperature is dependent on a single measurement from a single pixel. In many situations this type of measurement is not particularly useful, it is also difficult to display comprehensively in real time. In many cases what is required is a measurement of the temperature of an object within the scene. This can be achieved by comparing the colour or shade of the object in the image with a calibrated scale to the side of the image. This offers great flexibility but also introduces an additional source of error, the user's ability to compare the colour in the image to the colour on the scale. In practice making precise measurements usually requires the use of a colour scale with a discontinuity in it, sometimes called an "isotherm". The user will adjust the scale until the isotherm lies over the object they are concerned with; at this point the isotherm temperature represents the temperature of the object.

It may be more practical for the camera to calculate a numerical value of temperature of the object of concern and report that value to the user, avoiding the need for the user to adjust the scale and possibly making for quicker and more accurate measurements. The limitation being that the area covered by the object must be defined. Computational systems exist for the purpose of automatically defining and tracking objects within a video sequence and these could be used for the purpose of radiometry but that is beyond the scope of this project. Two other basic methods of defining the object of interest are discussed below.

One method is to predefine an area of the image and then to align the object with that area of the image in order to measure its temperature. This technique is often most suitable for portable imaging systems since it allows very quick sequential measurements of different objects. The calculated temperature can be reported constantly to the user or can be recorded to a file in the camera at the touch of a button.

The second method is to fix the field of view of the camera and then manually define an area within the image as the region of interest (ROI) and calculate a mean temperature within that area. Obviously some interface must be built into the system to allow the user to define the ROI, this often involves the user manipulating a cursor to draw a box within an object. An advantage of this arrangement is that it can be possible for the user to define an arbitrary shape as the ROI. A user defined region can be made to fit the object more precisely than a simple box or circle but this does involve a more complex interface and will take the user longer to define.

A user defined ROI system can allow the user to define more than one ROI within the image and measurements for each ROI can be calculated simultaneously. This can be very useful combined with a fixed camera position, since it allows continuous
comparison of multiple ROI's. With a suitable control system, it could be configured to report automatically when an ROI passes certain threshold temperatures.

An advantage to defining an ROI, rather than selecting an individual pixel, is that readings from all pixels within the ROI can be combined to give a mean value. This can increase the accuracy of the measurement by reducing the effect of random and fixed pattern noise between the pixels. Temporal averaging can also reduce tNETD induced errors on the measurement. The camera used in this project produces 50 frames per second, for many radiometric applications this kind of temporal resolution is not needed. Often taking a radiometric measurement over a period of a second or more would be perfectly acceptable and would yield a considerable reduction in random noise.

1.3.5 Spatial resolution and Modulation Transfer Function, MTF

Spatial resolution is related to an imager's ability to resolve fine detail in an image. In digital systems an ultimately limiting factor is the pixel resolution of the image produced but other factors of the optics or detector can reduce resolving power below what would be expected on the basis of the pixel count alone. Poor focus and other optical aberrations can lead to smearing of the image. Thermal bleeding between pixels or smearing in a readout CCD can also reduce the available resolution.

The modulation transfer function (MTF) is a measurement which attempts to describe the combined effect of all these factors. Calculating MTF involves taking a target of bars or edges and measuring the intensity profile of these shapes in the output image. A poorly performing detector will soften the edges of a bar pattern, and at higher spatial frequencies, the measured difference in signal between the hot bars and cold bars will drop. As the spatial frequency increases approaching the pitch of the pixel grid, aliasing can occur. A form of Moiré distortion occurs when beating between the pixel frequency and bar frequency produces patterns in the output which are not present in the input. MTF measures the effect of increasing spatial frequency on detector output and provides an insight to the cameras ability to resolve spatial detail.

Chapter 2, Infrared technology and its applications

2.1 IR camera technology overview

As of the early 1990's a typical high performance infrared camera might have consisted of the following components:

- A detector consisting of one element or a small array of elements, cryogenically cooled.
- A reflecting and/or refracting scanning system.
- A wide aperture $(\sim f/1)$ aspheric germanium lens.
- A system of analogue electronics to process the output of the detector into a video signal.

15 years later a typical IR camera is significantly different, the market has grown and prices for entry level cameras have dropped. Technical advances have made this possible, not all of these advances centre on delivering greater performance; many are made with the goal of functional performance at lower cost. To this end the biggest change has been the large scale adoption of uncooled focal plane arrays. By removing the need for cooling and scanning mechanisms these arrays have reduced the cost of the core detector system, in turn this has produced a pressure to reduce the cost of the peripheral components including the lens and the processing systems. The high performance end of the market is still dominated by cooled detectors.

The fundamental difference between cooled and uncooled IR detectors is that uncooled detectors are all thermal detectors and cooled detectors are almost exclusively quantum detectors. The first sections of this chapter examine the operation of both types of detector and the peripheral elements needed to drive them.

2.2 Quantum detectors and associated technology

2.2.1 Quantum detectors

In a quantum detector individual photons strike the semiconductor material of the detector and are absorbed, triggering an electronic transition and forming an electron-hole pair^{R13}. When enough such pairs are formed the effect can be detected electrically. Detectors operate either as photoconductors or photovoltaics. In a photoconductor the electron-hole pairs formed increase the conductivity of the detector element, the effect is measured by passing a bias current through the detector. In a photovoltaic detector the electronic transitions produce a voltage which can be amplified and measured directly.

An IR photon can trigger the formation of electron-hole pairs in a semiconductor, provided the band gap is smaller than the energy of the photon. The silicon photodetectors used for visible light cannot function beyond NIR for this reason. As the temperature of a quantum detector rises, phonons, (thermal energy), in the semiconductor cause the spontaneous generation of electron-hole pairs. This produces a signal which is indistinguishable from that produced by incoming photons, this signal is called the dark current. As a detector gets hotter the dark current increases and the quantum efficiency of the detector will fall, so the detector becomes less sensitive. (Quantum efficiency is the likelihood than an individual photon will trigger a transition.) At higher temperatures the detector will become completely saturated by the dark current; there are no longer any bound electrons to be excited by the incoming photons. At ambient temperature, MWIR and LWIR quantum detectors tend to be entirely saturated by dark current and so must be cooled.

The level of cooling required depends on the type and design of detector. In many cases it may come down to a balance between absolute performance and the practicalities of the cooling solution. As a rule MWIR quantum detectors are more likely to function at higher temperatures since MWIR photon energy is that much higher.

Quantum detectors have very short response times^{R13}, changes in radiation input can be measured almost immediately at the output. It is this property which makes it possible to use a scanning system to build a very high resolution image from a small detector and still achieve good frame rates. The high speed also makes focal plane array quantum detectors especially suitable for applications which require very fast update rates, particularly in automated systems which can react much more quickly than a human operator.

Quantum detectors offer excellent sensitivity and very high signal to noise ratios, but tend to have rather narrow spectral bands; sensitivity peaks near the band gap of the detector, photons with much higher energy are less likely to trigger transitions^{R13}. As explained in chapter one, a narrow spectral sensitivity range can be beneficial, but on the downside it does mean there are less photons available. It is the broadband response of thermal detectors which allows them to perform on the same scale as quantum detectors.

2.2.2 Cooling technologies

Adding a cooling system to the design of a camera can significantly increase the complexity of the system as a whole and depending on usage requirements, can have serious impacts on usability: These are some inherent disadvantages of a cooled system irrespective of which cooling technology is used:

- When the camera is turned on there is an inevitable delay before the detector reaches operating temperature.
- The energy required to cool the detector must come from somewhere, usually either high pressure bottled gas or electricity.
- If a system is required to be portable, the additional energy requirement represents another resource which must be carried with the camera

2.2.2.1 Overview of cooling technologies

The cooling systems used in infrared cameras take many forms, from very large systems with high thermal inertia, designed to operate continuously for long periods of time, to compact systems designed to cool quickly and run for only minutes or seconds. Despite the range of applications, the systems can be broken down into three widely used cooling technologies. Joule Thomson coolers, which tend to use bottled gas, Stirling cycle coolers, which are a form of heat engine, and Peltier coolers, which are a solid state technology with no moving parts. An example of each technology is given in Figure 2.1, the pros and cons of these systems are discussed below.





2.2.2.2 Joule-Thomson

A Joule-Thomson cooler uses compressed gas, which is released through a nozzle, as the gas expands it cools the area around the nozzle. In many systems a plug of condensed gas forms around the nozzle and limits the gas flow, making the design self regulating. Joule-Thomson coolers have been developed that can cool a detector to 100K in 2.3 seconds and maintain that temperature as long as gas is available ^{R14}. Gas must be provided by a compressor or, for portable units, stored in bottles. Efficiency depends on how well insulated the detector is. Often the detector will be housed inside a vacuum dewar to provide effective insulation. Typically a 1kg bottle (inc. gas and bottle) could provide a camera with 1 to 2 hours of continuous operation. However, a lot of gas is needed to cool the detector initially, so the bottle may only be able to cool the detector to operating temperature a few times.

2.2.2.3 Stirling cycle

A Stirling-cycle engine uses a motor to drive two pistons, compressing a gas and allowing it to expand. The engines can be made small and efficient, weighing as little as 380g, and consuming less than 5W of electrical power ^{R10}. The overall efficiency is greater than an integrated Joule-Thomson cooler and compressor. However, the time required to reach operating temperature is generally several minutes^{R10}. In portable applications the energy needed to drive the Stirling cycle cooler must be carried, usually in the form of an electrical battery.

Classically a Stirling engine would use a rotary motor to drive pistons, converting the rotary motion into linear motion through a cam or crank shaft mechanism. A newer generation of Stirling coolers use linear electromagnetic motors (like a solenoid) rather than rotary motors to produce much smaller piston movements at a higher frequency. The moving parts in such a system can be mounted on elastic restraints rather than moving freely, and the bearings associated with a rotary system aren't needed so such systems can be made more rugged and reliable than rotary driven Stirling engines ^{R15}. This makes them attractive in military applications.

2.2.2.4 Peltier coolers

A Peltier cooler uses the thermoelectric effect to generate a temperature gradient across a cooling element. By maintaining the hot surface of the element close to ambient temperature, allowing it to dissipate heat to the environment, the cool side can be taken well below ambient temperature. A single element thermoelectric cooler based on the Peltier effect can achieve a temperature differential of about 40 degrees. Multiple elements can be stacked, the temperature gradient of each adding, to increase the total gradient across the system. Each element produces heat as well as a thermal gradient, so the more stages used, the less efficient the device becomes. The more heat that is pumped, the more must be transferred from the hot side of the cooler to the environment, so large heatsinks are needed. The practical limit is found to be around four or five elements. Such a device can cool a detector to about 100 degrees below ambient temperature in less than a minute with a power consumption of a few watts.

2.2.2.5. Choosing a cooling technology

When designing a portable camera the choice of cooling system obviously has a very significant effect on how the camera can be used. Stirling cycle coolers are sometimes not a practical choice because of the long cool down time. Joule-Thomson coolers are a common choice in portable and military applications. If bottled gas is used, the cooling system requires no electrical power. The remaining problem is that of the gas supply needed for continued operation, either a large supply must be carried with the camera or the operator must have some means of replenishing it. However, there are military applications where prolonged use is not an issue, for example in missile guidance systems where only enough gas is needed to last the flight time of the missile. Thermoelectric coolers can be easily integrated into portable units. The power consumption is not so great as to make such designs totally impractical, but the level of cooling provided is smaller than other cooling technologies.

2.2.3 Scanning systems

2.2.3.1 Conventional scanning

When a cooled detector is used it has to be enclosed within a vacuum chamber, or some other form of effective thermal insulation. At the time infrared cameras were first developed, it was not possible to produce an array of very small elements, using a large number of elements would have resulted in a very large detector. A larger detector implies a larger vacuum container and hence more weight and more cooling power. In addition most early quantum detectors were photoconductive devices, so measuring the signal involved dissipating significant energy in the detector, consequently more elements would require additional cooling power. Wiring up a detector within a vacuum container poses practical problems; there must be a break in the wall of the chamber for the cabling, and this cabling conducts heat into the chamber. If a large number of elements must be wired up, this problem becomes larger. The leakage can be minimised through various techniques, such as using very fine stainless steel conductors, or using etched conductive traces on a flexible substrate. Advances in electronic design now make it possible to fit multiplexing electronics within the vacuum container reducing the cabling requirements. However, any electronics added inside the chamber introduce an additional source of heat which must be removed by the cooling system. These are the limitations which initially led to the introduction of scanning systems.

A scanning system works by mapping a small section of the total field of view (a window) onto the detector at one time and scanning this window across the total field of view, building up a two dimensional image. The principle is similar to the way in which a cathode ray tube scans its beam across a screen, generating an image. The detector consequently only looks at each part of the scene for a small fraction of the time taken to scan each frame, thus reducing integration time. However, since only a few elements are needed, they can be made comparatively large, and the increased radiation gathering area makes up for the short exposure. Scanning systems are still used in IR cameras to provide higher spatial resolution than is otherwise possible. The output image resolution of the best scanning detectors surpasses the highest resolution 'staring arrays' (a term used to describe detectors which do not scan).

2.2.3.2 Microscanning

A process known as microscanning is sometimes used with high resolution arrays. A scanning system is used to shift the image on the detector between frames. The shift is achieved either with part of the optical system, (such as using angled windows in a chopper blade), or sometimes by shifting the detector itself. The image is usually shifted by half a pixel (or an integer plus a half) between frames, the image may return to its original position for the next frame or may cycle through a number of different shifted positions with a period of 4 or more frames. The pixels from each frame are positioned appropriately on an output grid of increased resolution and either integrated to produce a single image of higher resolution, or displayed in quick succession. The effective increase in detector resolution is termed "super resolution". Since many detectors have

less than 100% fill factor (they are "undersampled"), super resolution can often reduce moiré effects and stair-step edge artefacts, significantly improving image fidelity in some scenes ^{R16}.

Various microscanning sequences are used, some examples are:

- A simple half pixel displacement in the horizontal or vertical direction.
- 4 sequential half pixel shifts in a 2x2 grid, improving both vertical and horizontal resolution over a 4 frame cycle.
- A diagonal shift which can improve vertical and horizontal resolution in a cycle of 2 frames.

With any of these scanning regimes the microscanning distance can be changed from half a pixel to 1.5 pixels (or more), which improves the performance of the camera by sampling adjacent pixels in the compounded image with different elements within the detector, reducing non-uniformity. Quantitative and qualitative research has shown that image quality from undersampled detectors is significantly improved by 2 frame microscanning; 4 frame microscanning improves performance further, but higher order microscanning gives rapidly diminishing returns ^{R16}.

Microscanning has its drawbacks, the scanning must operate at such a rate that each position is sampled at least 25 Hz in order to avoid visible jittering in the image, preferably higher. Take the example of a 4 frame microscanning system with a 100Hz detector, each sub-frame is sampled at 25 Hz. If the display connected to the camera is capable of rendering 100 frames per second the sub-frames can simply be shown in succession shifted by the appropriate amount, the eye and brain of the user integrate the output acceptably. However, if the connected display is only capable of operating at 50Hz then two of the four sub-frames have to be combined into a single image and displayed at once. With a static scene this is not a problem, but if the camera is moving, or if objects are moving within the scene, there is the chance that the same object is displayed at two points in the same frame. When this happens it can be easily spotted by the eye, and it looks to the user like there are two objects instead of one, even at a rate of 50 Hz. The severity of the effect is somewhat subjective and can range from mildly objectionable to genuinely confusing depending on both the operator and the application. The effect can be described as a breakdown in temporal correctness. PAL and NTSC video formats produce 50 and 60 (interlaced) fields per second respectively, and so could not display the required 100 Hz signal. Processing within the camera could be designed to display 100 frames per second, but the LCD display panels commonly used in portable cameras do not normally have a fast enough response time

to be driven effectively at 100 Hz. The requirements are well within the range of CRT's designed as computer monitors, but these cannot generally be integrated into portable units.

•

•

2.3 Uncooled IR imaging

2.3.1 The rise of uncooled IR imaging

Arguably the most significant development in IR imaging technology in the last decade has been the improvements in uncooled imaging. In the past, uncooled IR imaging was seen very much as the poor cousin of cooled imaging and as a niche market, an option considered where quality was not of paramount concern. Uncooled imaging is now seen as the choice that will bring IR to the mass market and as such, a very profitable field for manufacturers. What has changed to bring this about?

Uncooled IR cameras rely on thermal detectors, rather than quantum detectors, the limitation of early thermal detectors was often the thermal capacity of the detector element^{R13}. The large detector elements used had high thermal capacity and therefore slow response and poor sensitivity. Over the last few decades advances in semiconductor manufacturing and the emergence of micromachining has allowed tiny intricate structures to be produced in large arrays at low cost; it is this technology that drives today's thermal detectors.

2.3.2 Thermal detectors explained

2.3.2.1 The concept of a thermal detector

Consider an object sitting isolated in the vacuum of space, there is no material link between the object and its environment. What temperature will this object reach? There is no path for thermal conduction into or out of the object, the only method of thermal transfer is through radiation, the object emits black body radiation and it absorbs radiation from the environment. Over time the object will reach equilibrium with the environment, it reaches a temperature where the power of its black body radiation matches the power of the absorbed radiation. In deep space the equilibrium temperature will be 2.7 K, this is the black body temperature of the universe, known as the cosmic microwave background (CMB). The time taken to reach equilibrium will depend on the emissivity of the body, its surface area and its thermal capacity.

Now consider the case of an object hanging unsupported in front of a substrate; the object is our thermal detector element, the substrate is the silicon chip on which the detector is made. Assume the detector substrate has a high thermal capacity and is at room temperature, so it can act as a heatsink without changing temperature. For now, assume there is still no path for thermal conduction between the detector and its

environment, all transfer is by radiation. If the detector faces a hot scene, for example the inside of an oven, the element is exposed to radiation which will be absorbed and raise its temperature. It will reach equilibrium at a temperature somewhere between that of the substrate and that of the oven. If we can measure the temperature of the detector element we have a thermal infrared detector.

In the real world the detector element must be supported in some way, if we add legs between the element and the substrate we introduce a path for thermal conduction. Heat will flow from the detector element into the substrate, but providing the legs have some thermal impedance the detector element will still settle at a temperature above that of the substrate. The lower the thermal impedance, the colder the detector element will become. Consequently thermal detector elements are often elevated platforms supported above the substrate on narrow legs, the structures created using silicon micromachining techniques.

If a gas is introduced between the detector element and the substrate another thermal conduction path is introduced and more signal is lost. Since conduction through the legs and through the gas are parallel paths, excessive loss through the gas cannot be compensated for by longer legs. Consequently thermal detectors are often packaged in either a vacuum, or a gas with low thermal conductivity, typically noble gasses such as Argon or Krypton.

2.3.2.2 The time constant of a thermal detector examined through an electrical analogue

The system described above can be represented by an electrical analogue^{R17}, Figure 2.2. In this circuit the capacitance represents the thermal mass of the detector. The resistance is the thermal impedance between detector and substrate (the combined impedance of the structural supports and gas in parallel). Voltage represents temperature, with the substrate considered as ground. Net radiation input is represented by a current source at the top of the circuit. As radiation is absorbed by the detector a voltage or temperature differential forms between detector and substrate, this is the signal voltage.

It is clear from the circuit that the system has a time constant created by the thermal mass of the detector element, the more massive the detector element the slower the detector. The time constant is given by^{R17}:

$$r = Z_C \cdot C$$

Where:

r = time constant, s

- Z_{c} = conductive thermal impedance between detector and substrate, °C W⁻¹
- C = heat capacity of the detector element, J ° C⁻¹

The time constant is of particular importance to video applications, a video camera must achieve a frame rate of at least 25 frames per second in order to produce apparently smooth motion, many military and other high motion applications demand at least 50 frames per second. A slow time constant compared to the frame rate will cause ghosting and trails, where features of the previous frame persist in the current frame. (This is most obvious when hot objects move across the field of view; they will appear to leave a decaying trail of light behind them.) To operate effectively when viewing moving scenes a thermal camera must have a time constant shorter than the read out period, this implies a time constant of the order of tens of milliseconds at most.

So the design of a thermal detector is a compromise between sensitivity and time constant, increasing the thermal isolation of the detector element will make it more sensitive, but at the expense of speed^{R17}. The only way to improve speed without sacrificing sensitivity is to reduce the mass of the element; this is a challenge for semiconductor fabrication technology.



Figure 2.2, Electrical analogue of a thermal detector

There is a way to circumvent the problem by adding a mechanical chopper to the system. If the input is modulated, it is possible read the detector element at a higher frequency compared to its time constant, and still get acceptable performance. Measuring the difference between the closed scene looking at the chopper and the open scene looking at the world is a way of measuring the gradient of the temperature vs. time curve of the detector. Since the time constant of the detector does not vary, the gradient measured is proportional to the radiation power absorbed^{R18}. (Wien's Law suggests the time constant will vary with detector temperature, since heat loss through radiation is not directly proportional to temperature. However, with most detectors thermal exchange is dominated by conduction so this nonlinearity is small, this is shown by calculation in chapter 4.)

A chopper can also reduce 1/f noise, but it also has drawbacks, since radiation is blocked during the closed half of the cycle less power is incident on the detector, so there is less thermal input to measure^{R19}. Indeed when chopped faster than the thermal time constant the temperature changes in the element are lower than they would be without the chopper. Choppers tend to be used only when the nature of the detector



demands it, as with pyroelectrics^{R18}. Detectors which suffer from significant 1/f noise or are particularly sensitive to ambient temperature changes are often used with shutters which operate periodically to produce a dark frame for offset correction. How often the shutter operates depends on the detector and application, but this system is fundamentally different to the AC nature of a chopped detector.

2.3.2.3 Types of thermal detector

In theory any physical property which changes with temperature can be used as a thermal detector, but in the field of infrared imaging there are two main types of detector used in uncooled cameras; pyroelectric arrays and microbolometer arrays. Both types of detector perform well and each has its advantages in different situations, both are used in cameras that are in commercial production. However, in recent years interest in pyroelectrics has waned in favour of microbolometers, this can be explained in part by the better compatibility of microbolometers with standard CMOS fabrication technologies. This improves the prospect of scaling to progressively larger array sizes, and of cost reductions through economies of scale. Also, microbolometer technology can take advantage of process improvements developed through other applications of MEMS technology.

The camera used in this project, and all the cameras produced by QinetiQ as part of the same product line, are based on pyroelectric materials. For this reason the work here concentrates on the technology behind these detectors, and draws comparisons with microbolometer technology where appropriate. There are other types of thermal detector used in uncooled cameras, including a recent technology called microcantilevers. This technology is still considered experimental and has little direct relevance to the work in this project, so it is not covered in extensive detail, they are nevertheless interesting because they are a thermal detector measuring a different physical property; thermal expansion.

2.3.4 Microbolometers

Bolometers have been used as infrared detectors for many decades, their use in single point infrared temperature sensors predates their use in large imaging arrays. The term microbolometer is not a strict definition, but it is used to refer to large arrays of bolometers fabricated through lithographic processes. Modern fabrication techniques involve micromachining, (also referred to as Micro Electro-Mechanical Systems or MEMS), to produce the freestanding structures needed for high sensitivity. A bolometer is a resistor with a high temperature coefficient of resistance (TCR); the resistor forms the element of a thermal detector, it is exposed to infrared radiation and the temperature of the element is read by passing a bias voltage trough the resistor. There are a wide variety of conductors to choose from when making a bolometer but the market is currently dominated by Vanadium Oxide^{R20}, commonly abbreviated to VOx, (which achieves a TCR of around 4% /°C), and Amorphous Silicon.

An obvious challenge when measuring bolometers is sensitivity to ambient temperatures, the detector may be expected to operate at temperatures over a range of a few degrees to tens of degrees depending on application, whereas signals may be in the order of microkelvin. Over this ambient temperature range bolometer resistance may change by tens of percent, and signals are tiny by comparison. This poses a problem for the sensitivity of the signal amplifiers and the resolution of the ADC used.

To read the bolometer a potential divider must be made with a second resistor such that the impedance change is manifested as a voltage. Figure 2.3 gives one example of a bolometer read-out circuit, other configurations are possible^{R21}. One method of handling ambient temperature changes is to use a reference resistor made of the same high TCR material as the detector, sometimes called a compensating resistor. The compensating resistor is not sensitive to IR, instead it is kept at the substrate temperature. As ambient temperature changes the impedance of both resistors changes by the same fraction and the ratiometric nature of the circuit means signal voltage remains constant. Impedance mismatching between the active and the compensating resistor will introduce voltage offsets between different voltages in an array. Nevertheless, this form of compensation can be enough to keep the output voltages of an array within the range of the range of a high resolution ADC whilst maintaining enough voltage resolution to measure the IR signal.



Figure 2.3, Example bolometer readout circuit with compensating resistor and deferential output

Whilst simple stored offset values might be enough to produce an infrared image from such a system, other factors combine to introduce non uniformity problems. Mismatch in the TCR values of the active and compensation resistors causes a residual sensitivity to ambient temperatures. Temperature gradients on the detector die can lead to different temperatures in the region of the active and compensation resistors; this will be manifested as a spurious signal. In addition, some of the resistor materials used may show significant changes in absolute impedance or TCR over time.

Cameras are designed to keep the detector temperature as stable and even as possible, often using a low power cooler or heater to allow detector temperature to be kept constant even when ambient temperature changes. Even with very careful camera design, changing pedestal noise remains a problem with bolometer detectors. The most common method of dealing with this is to use a uniform shutter to periodically check the current offset values of all pixels,^{R23} these systems are usually controlled automatically; the camera will monitor the detector temperature and trigger the shutter process whenever it detects a change in ambient temperature which warrants it. The shuttering can take from a fraction of a second to several seconds depending on the camera. Typically the shutter will trigger most often immediately after the camera is turned on, as the detector reaches its operating temperature. After that, if the camera is

operating in a stable environment, there may be several minutes between shutter operations.

This system can correct for changing pedestal noise very well and allows microbolometers to be used as radiometric imagers, but the shutter mechanism does cause problems for certain applications; firstly, it adds a moving part to the system, which is a possible point of failure and can make the camera less rugged. Secondly, the operation of the shutter does mean a momentary loss of image, in many applications that may be seen as a critical flaw, military, fire fighting and other search and rescue type applications are obvious examples.

Microbolometers are commonly made with structures called microbridges, a name derived from the physical structure of the pixel, it consists of a thin layer of detector material, supported above a silicon substrate by raised electrical contacts, producing a vacuum gap below the detector and providing excellent thermal isolation. Each pixel is joined to the contacts by thin legs which reduce the thermal connection with the substrate still further. A typical microbridge structure is shown in figure 2.4.

This structure is achieved by starting with the silicon readout circuit, on the surface of which lie the electrical contacts for the detector elements. These contacts are surrounded by a layer of sacrificial material. A structural layer and then the bolometer material are deposited on top of the sacrificial layer. Later the sacrificial layer is removed, (usually by chemical etching), leaving the free standing bridge structure. This kind of MEMS construction process, depositing and removing sacrificial material, is termed surface micromachining, it differs from bulk micromachining which involves etching down into the bulk silicon wafer, rather than deposited layers.



Figure 2.4, Microbridge Structure

Surface micromachining allows the bridge structure be made very thin, at this microscopic scale the strength of the structure is not generally a limiting factor, the precision and uniformity of the fabrication processes determines how thin the microbridges can be made. Thinner bridges have smaller mass which allows the design to use high thermal isolation whilst maintaining an acceptably short time constant, thus yielding more signal. Below a certain thickness the level of absorption in the bridge can fall, leading to loss of signal. Various techniques are employed to maintain absorption whilst making the bridge thinner; these include tuning the height of the gap under the bridge, and adding special absorbing layers.

At present microbolometer arrays of 640x480 pixels are common but represent the high end of the market, 320x240 arrays are significantly cheaper. High quality arrays can achieve an NETD of <30mK^{R20}. The weight and power consumption of these arrays is low enough to provide handheld imagers with good battery life; one manufacturer has demonstrated a 640x480 array weighing only 126g including the driving electronics and consuming less than 3W of power during continuous operation ^{R22} The performance of microbolometer detectors has improved remarkably in recent years but some theoretical research suggests they are reaching the fundamental limits of their performance ^{R21}. If this is the case the next decade may not see similar gains. However, it will certainly see large increases in array resolution and large reductions in cost.

2.3.5 Pyroelectric Detectors

2.3.5.1 The principles of pyroelectric detectors

Pyroelectrics are materials which generate an electrical charge in response to changing temperature. In crystalline materials with large internal dipoles, changes in the lattice temperature cause the dipoles to move and lead to spontaneous generation of charge. This charge can be measured electrically, so if these crystals are formed into appropriate elements and exposed to thermal radiation, a thermal detector can be made^{R18}. The charge generated by a pyroelectric element normally dissipates rapidly through the movement of free charges within the crystal, so it is not possible to make DC measurements using this kind of detector, however, their sensitivity to changing temperature is extremely good^{R18}. By providing a modulated source of radiation, alternating between radiation from the scene and a constant "reference" radiation input, the delta between the two inputs is presented as an AC signal which can be measured. Provided the reference input is kept constant, even very low frequency

changes in the scene input can be measured absolutely. Typically, the method of modulation is a rotary chopper and the reference input is the blackened surface of the chopper blade.^{R19}

Modulation frequency for pyroelectrics is flexible^{R18}, the lower limit is determined by the thermal time constant of the detector, since once the temperature stops changing signal rapidly disappears. At high frequencies response falls off because the charge does not have time to build up. (Although the movement of free charge starts to cancel the polarisation the moment it forms, the position of the dipoles creating it is a function of absolute temperature, so the temperature must change by a certain amount before peak polarisation is reached.) The drop off in performance at higher frequencies is determined by both material and configuration, but many pyroelectric detectors can be driven to tens of hertz before serious signal degradation, and some much higher^{R18}. When made into a starring array they can easily meet the requirements of general purpose video infrared imaging.

2.3.5.2 The Design of pyroelectric detectors

This section covers the manufacturing techniques used in the development of the pyroelectric detector used in this project and subsequent chips in the same family.

2.3.5.2.1 Readout chips

A pyroelectric array detector can be split into two main components, the pyroelectric material and some form of readout chip. The readout chip must measure the charge produced by each element of the detector and provide some way of multiplexing the output of all pixels. There are two common solutions, one is to use a CCD, the other is to use a CMOS chip consisting of an array of MOSFET (Metal On Silicon Field Effect Transistor) amplifiers (one for each pixel) and multiplexing circuitry.

Of the two designs CCD's offer a generally simpler design and generally greater uniformity but a MOSFET array avoids some of the problems of a CCD. These include smearing caused by readout on the CCD and a tendency towards higher noise ^{R24} R25. The detector used in this project and other detectors in the same family use MOSFET output.

2.3.5.2.2 The interface between detector and readout

In the complete detector the pyroelectric material must be joined to the readout chip in some way that is both electrically conductive and thermally isolating. These two criteria are slightly contradictory and careful design is needed to get maximum performance from the detector. The production techniques used to fulfil these criteria was classed by QinetiQ as a hybrid technology since it involves two production processes on separate dies which are later joined together.

2.3.5.2.3 Hybrid technology; bump bonding

Steps in the process are shown in Figure 2.5. First a layer of pyroelectric is sliced from a block and polished. A common electrode is deposited on one surface and individual electrodes for each pixel, on the other. On top of each pixel's electrode, a non-wettable pad (a material that repels molten solder), is formed. A smaller wettable pad (which attracts solder), is formed on top of that. Solder is plated over these pads producing the structure seen in Figure 2.5a. The chip is then heated melting the solder, which is drawn up into droplets, or bumps, sitting on the wettable pads, a process know as reflowing, which is shown in Figure 2.5b. The chip is aligned and brought into contact with the readout chip whilst the solder is still molten, bonding the two halves together to produce the final chip. One of the most difficult parts of the process is bringing the two halves together properly aligned. The whole process is known as flip chip solder bonding or bump bonding.



Figure 2.5, Hybrid construction, the bump bonding process

A limiting factor in the simplest bump bonded chips is the presence of a single, continuous layer of pyroelectric material, which extends across the whole chip, making the thermal isolation between pixels very poor, another factor is thermal conduction through the solder bond. Solder has high thermal conductivity and effectively leaches signal away from the detector layer. Systems exist to alleviate both of these problems.

2.3.5.2.4 Hybrid Technology, adding reticulation

To reduce the thermal conduction between pixels, the detector material can be reticulated, the process involves cutting groves between the pixels, in order to isolate them. Reticulation can be carried out by two different methods, laser photochemical etching, or ion beam milling.

Laser photochemical etching involves immersing the detector in an environment of potassium hydroxide and using a laser to cut groves in the detector material. The laser beam drives the chemical reaction between the potassium hydroxide and the pyroelectric. The process is depicted in Figure 2.6. This technique requires each cut to be made separately. This can take a long time and becomes impractical for producing large arrays in high volume.

Ion Beam Milling is an alternative technology which allows reticulation to be carried out more easily on large detectors. Rather than a laser beam, a wide stream of argon ions is used to drive the etching process. The beam covers the whole area of the detector, so all of the reticulation is cut in one operation. To protect the areas of the chip which will become the detector elements, a mask is deposited on the surface before etching begins. The mask is also etched away by the ion beam but the thickness of the mask is chosen such that etching in the masked areas reaches the top surface of the detector at the same time that the etching in the groves reaches the end point. Layers know as "etch stops" are used to slow down the rate of etching once it reaches the desired level, they are shown in Figure 2.6.



Figure 2.6, Reticulation by laser photochemical etching

2.3.5.2.5 Advanced bump bonding technologies.

Techniques exist to reduce thermal conduction through solder bumps. A simple solution is to reduce the size of the solder bond. However, this reduces the electrical contact and can lead to greater variations between pixels and more dead pixels. Alternatively, a layer of polymer material with poor thermal conductivity can be included between the solder and the detector, this is shaped so as to form a mesa like structure. The technique is known as mesa isolation. Cross-sections of finished detector elements, using these isolation techniques, are shown in Figure 2.7.

Hybrid production techniques require complex construction procedures, with close tolerances in each phase of the process. The large number of steps involved gives rise to high levels of variation between pixels, and also a high incidence of dead pixels. Making these chips is slow and expensive, but it does work successfully and the techniques and quality of output improve over time.



Figure 2.7, Bump bonding technologies

2.3.5.2.6 Integrated Technology

This term was used to describe an alternative to hybrid technology, a technique where the entire chip would be built up on one die, avoiding the difficult step of flip chip bonding. It was hoped that microbridge structures could be used, similar to those in microbolometers and offering many of the same advantages: thinner detector elements, better thermal isolation and reduced costs. Unfortunately, process compatibility issues prevented the QinetiQ project from reaching this goal and have hampered further pyroelectric array development.

2.3.6 Microcantilevers

The microcantilever is a form of thermal infrared detector which relies on measuring the thermal expansion of materials within the detector element. It is a kind of detector which would not have been possible without recent advances in MEMS technology^{R26}. The structure of a microcantilever pixel element is in some ways similar to a microbridge in that an absorbing plate is suspended above the detector substrate on long arms, but the operation is very different. A microcantilever pixel is depicted in Figure 2.8. Sections of the legs are formed from bi-material strips which, like a bi-metallic strip, are formed of two materials with different coefficients of thermal expansion. Changing temperature causes these legs to bend, lifting the tip of the cantilever up and down.



Figure 2.8, Micro-cantilever detector technology

Microcantilevers differ from pyroelectric and bolometric detectors in that the thermally dependent property being measured is not directly measurable electrically. Some method is needed to measure the movement of the cantilever, at least two methods have been tested successfully. One involves using electrically conductive arms and measuring the capacitance between the IR absorbing end of the cantilever and another plate built into the detector substrate. As the cantilever moves up and down the two plates of the capacitor move closer or further apart changing their capacitance. The other method that has been reported is to measure the position of the cantilevers optically, either with a laser or a visible band imaging array. This offers the interesting prospect of an infrared detector electrically isolated from its readout circuit.

The performance of microcantilevers to date has not matched that of microbolometers or pyroelectrics but some theoretical studies have suggested that the fundamental limits of performance for this technology exceed those of microbolometers ^{R26}. That is why the technology has produced significant interest despite being in the early stages of development. It has the possibility of becoming at least as cheap as microbolometer technology and may provide advantages in very high resolution arrays. Moving the

readout system off the detector chip removes the chance of an electrical defect during manufacture rendering the entire detector chip unusable, and since the pixels require no electrical connection, the occurrence of dead pixels may be reduced. It also removes heat sources from the array reducing the chance of thermal gradients in the substrate.^{R26}

٠

•

2.4 Infrared optics

Most of the plastics and glasses used to make lenses and windows for visible light cameras are completely opaque at MWIR and LWIR wavelengths, so different materials have to be used. There are many materials that have been used to make optics for medium and long wave infrared cameras; Germanium, Zinc Selenide, Zinc Sulphide, Cadmium Telluride and Silicon are all practical solutions.^{R27} Various alkali halide materials can be used but have hygroscopic properties and will degrade in humid environments, making them impractical for many applications.^{R27} These materials have been used for IR optics for many decades, and are considered established technology. From those materials germanium has emerged as the most common choice for mobile imaging cameras. Over recent years the drive for lower cost thermal imaging has pushed forward the development of two alternative technologies, moulded IR glasses and moulded plastics.^{R28} In this section the properties of a typical germanium lens system will be examined and the emerging technologies will be compared to this standard.

2.4.1 Traditional IR lens in Germanium

Compared to the materials used in visible light optics, germanium has a high refractive index, around 4.0,^{R29} which makes the challenges of lens design rather different.

The performance of infrared cameras is often limited by noise levels, and for this reason an increase in the level of radiation incident at the detector will often lead to a significant improvement in image quality. Desirable characteristics of an IR lens include wide aperture, (high F number, often f/1 or greater), and high percentage transmission through the lens system. The percentage transmission through a germanium lens element will be lower than is typical in visible light optics, partly due to poorer transmission through the material itself, and partly due to greater surface reflection caused by the high refractive index, this promotes the use of a small number of lens elements. Whereas a normal 35mm camera lens may have upwards of 6 elements, often as many as 15, many IR lenses will only have two or three. Zoom systems are very rare in IR imaging, if different focal lengths are needed it will generally involve swapping lenses.

Fortunately the high refractive index of germanium makes it practical to build lenses with high f numbers so long as aspheric surfaces are used. In visible light optics aspheric lenses are generally produced by high precision moulding processes, but germanium cannot be moulded in this way and aspheric lenses must be made by lathing a lens from a solid block of germanium. Diamond is used as the cutting material and the process is often referred to as diamond turning or single point diamond turning: SPDT (referring to the fact that a single cutting element is used). SPDT lenses are more expensive than moulded lenses because the process of cutting is slow, requires more work to keep equipment perfectly calibrated and production cannot easily be scaled up to high volumes. Since IR wavelengths are longer than visible light, the surface tolerances in the lens manufacture do not have to be quite as stringent as for visible light lenses, so although diamond turning is expensive, it has not been prohibitively so. Diamond turning has remained viable because of the high cost of other elements in the camera system, making it just one expensive component in a list of expensive components. As uncooled detectors have become a commercial technology, production costs have dropped, as has the cost of the electronics to drive them, consequently the pressure to find a cheaper alternative to germanium has increased.

Whist germanium is a rugged material and would not need coating on the grounds of durability, its high refractive index makes surface coating a necessity for good transmission levels.^{R30} Diamond has an appropriate refractive index for a first layer coating. Diamond coated lens surfaces are very resistant to scratching, so sometimes single layer diamond coatings are used for the external faces of lenses for military applications, however in most application multicoating is used to maximise transmission. With good coating, transmission of well over 95% can be achieved across the passband of a single germanium element.^{R30} Germanium works well in both the MWIR and LWIR bands, and compared to other materials has very low dispersion, which is the variation of refractive index with wavelength. Dispersion gives rise to chromatic aberration, even with monochrome output chromatic aberration degrades the quality of an image, lowering resolution. The refractive index of germanium is sensitive to temperature, giving rise to thermal defocus. This can be partially compensated for by refocusing the lens, but a germanium lens will have an ideal operating temperature which should be maintained for peak optical performance.

2.4.2 Moulded Aspheric IR Glass lenses

Perhaps the most promising new technology for high volume thermal imagers is that of moulded IR glass, called chalcogenide glass.^{R36} Chalcogenide glasses are a family of

semiconducting glasses, containing one or more of the chalcogen elements (group VI elements S, Se, Te) as alloy elements. Umicore IR Glass is one company specialising in these materials, currently their two main products sell under the names GASIR1 and GASIR2 which have compositions of Ge22As20Se58 and Ge20Sb15Se65 respectively (and slightly different passbands). As is evident from their compositions, both contain roughly 20% germanium, this is significant because germanium is expensive in itself. Whilst these lenses have the potential to become far cheaper than SPDT germanium when mass produced, the amount by which the price can fall is limited by the materials cost.

The dispersion in these glasses is higher than that of germanium.^{R36} In order to obtain similar performance the optics must be designed to compensate for this effect, a process known as achromatisation. This is achieved by using diffractive or aspherodiffractive surfaces, the use of these complex surface curvatures is made possible by the moulding technology. Once the high precision mould has been produced the added cost of making lenses with these surfaces is comparatively low.

Like germanium lenses, chalcogenide glass lenses are coated to increase transmission from about 70% for the untreated glass up to 98% (7.5 – 11.5 um) for the highest performance coating.^{R36} Janos Technology also offer a "durable" coating designed to withstand abrasion and the harsh treatment that can be expected in a military environment. This offers lower transmission of 95% but would typically only be used on the front surface of the front element of a lens.^{R36}

Umicore now produce standard lenses using these materials including a two element 25 mm f/1.2 lens available from the manufacturer as a coated but unmounted lens set for as little as 325 euros at the time of writing. Performance of the best chalcogenide lenses has been shown to be close to that of germanium lenses and the technology is eminently suitable for medium to high volume production and use in the medium to high cost IR imaging market.

2.4.3 Moulded plastic Fresnel lenses

Crude polyethylene plastic lenses have, for a long time, been used with non-imaging, single element IR sensors, such as pyroelectric motion sensing intruder alarms. Historically such plastics have not been seen as suitable for imaging applications because the transmittance of the material is so low that total transmission through the lens would be unacceptably poor. A company called Fresnel Technologies Inc. have

tried to overcome the poor transmittance of plastics and produce very low cost imaging lenses by using Fresnel lens technology^{R28}.

Fresnel lenses are often used in visible light applications where very thin lenses are required. The shape of a conventional lens is broken into a number of concentric sections known as Fresnel zones. For each zone the lens thickness is reduced without changing the surface curvature but instead creating surface discontinuities between each zone. Fresnel lenses are more often used in lighting applications than in imaging applications because their performance is limited. (One of the most commonly seen examples of a Fresnel lens is in the back lit writing surface of overhead projectors.)

The reduced thickness of the Fresnel lens makes the lenses viable for basic IR imaging, going some way to compensating for the poor transmittance of the material. When tested, a two element 50mm f/0.8 lens produced by Fresnel Technologies Inc. performed significantly worse than the other lens technologies discussed here. The poorer MTF led to reduced contrast and produced a blurred appearance to the image R²⁸. However, the low refractive index of the plastic (~1.54) allows it to be used uncoated, giving the very real prospect of an imaging IR lens for only a few dollars.

The plastic used is injection moulded and tough enough for mounting flanges to be integrated into the lens mould, reducing total manufacturing cost for a camera. The low material cost makes wide apertures comparatively cheap, which would provide plenty of radiation input for cheap, low sensitivity detectors. These lenses may well find a market in very low cost IR imagers.

2.5 Applications for infrared cameras

2.5.1 Military applications

2.5.1.1 The military as a driving force in IR development

During the development of infrared imaging in the 20th century, national defence became one of the largest markets for infrared imagers. Various military interests around the world have also been responsible for the development of many existing systems. The applications of IR technology for offence and defence cover a vast range, from personal imagers, to satellite imagers, to computer controlled guidance systems.

2.5.1.2 IR vision systems

Infrared vision systems range from personal handheld devices to integrated surveillance systems in all kinds of land, sea and airborne vehicles. These systems offer two very distinct benefits to the user. Firstly, improved target detection; anything warm, people, operating vehicles, occupied buildings, stand out well against the environment in most situations. Thermal camouflage is generally more difficult to implement than visual camouflage. Secondly, IR vision enables passive observation in complete darkness and improved vision in some adverse environmental conditions. Passive observation means no illumination is required, the user can observe an enemy, with reduced chance of detection and at greater distances than is otherwise possible.

2.5.1.3 Reconnaissance systems

Infrared has for a long time been used in satellite earth-imaging and high altitude imaging from planes. Applications include battlefield reconnaissance and tactical intelligence regarding an enemies military and industrial activities. For example it is easy to tell whether an industrial facility is active, based on the heat signature. In some cases it is possible to detect where a military encampment has been, even after the people and equipment have left. Operating generators and running engines can leave heat signatures in the soil, which persist for hours after the heat source has been removed. In this type of application real time imaging is not always necessary, data may be stored and examined later, and still images may supplant video data with equal information value. In some cases passive scanning is used, where a one dimensional detector is used and the passage of the carrier moves the target area across the landscape, building up a two dimensional image.

2.5.1.4 Automated guidance systems and fuses

Heat seeking guidance and fuse systems are examples of computer controlled IR imaging systems, data is captured and interpreted in real time by some form of computer which controls the guidance or detonation of a missile weapon. The output need never be seen by a human operator. Such systems will not always capture an image in the form of a two dimensional grid of data, in some cases a number of single element sensors may be used.

2.5.2 Law enforcement services

Police forces have similar requirements to the military in terms of surveillance and night vision systems. The police generally require these systems in much smaller quantities than the military. However, drug enforcement, and customs services have similar needs. The combined market becomes quite large.

2.5.3 Emergency Services

Fire fighting services can put portable long wave cameras to great use in finding casualties and fighting fires in smoke filled buildings. Most fire services in the United States are now equipped with such imagers, this represents a large market.^{R31} Search and rescue services can use IR imaging to more easily find people lost or stranded in inhospitable environments or dangerous situations.

2.4.4 Civil & Industrial

2.5.4.1 Surveillance

Infrared cameras have obvious applications in civilian security and surveillance systems, the uses and benefits are very similar to those in comparable military and law enforcement applications. A limiting factor in these markets has been the cost of installing and running MWIR or LWIR imaging systems. This has pushed the market towards cheaper alternatives such as SWIR systems with active illumination. The falling price of uncooled cameras may start to change that.

2.5.4.2 Condition monitoring

A major civilian use is condition monitoring through radiometry. This can be extremely useful for non-destructive, non-invasive and non-interruptive testing, fault detection and many similar applications. Many modes of failure in electronic and electrical systems result in an increase in component temperature before subsequent failure. Often temperature changes can be measured before performance has been adversely affected. Thermal monitoring can allow faults to be detected and remedied without loss of functionality. Radiometric imaging has advantages over other forms of thermal monitoring, such as integrated temperature sensors. One imager can be used to inspect many components or systems at distance, as such it can be a cheaper and more effective solution than direct temperature measurement.

Radiometry can also be used to monitor structural integrity, materials under varying or modulated strain are subject to frictional heating. Poor structural integrity often results in strain concentrated around points rather than distributed evenly, the concentrated strain can lead to inelastic distortion or catastrophic failure. When the system as a whole is subjected to stress, the points of highest strain will show the greatest increases in temperature, thermal imaging can identify these points. Systems commonly tested include structural supports, pressure vessels and piping. In situations where a vessel or pipe carries a hot or cold load, flaws may become visible without modulated strain. Thin points or partial fractures offer a path of reduced thermal resistance between the interior and exterior, so they will have a different surface temperature to the rest of the surface. These engineering practices can be used to monitor the condition of existing installations and also as a tool in the design of new systems.

2.5.4.3 Other civilian uses

Other civilian uses which have been suggested include night time and poor weather driver aids in motor cars, and night sights for bird watching and game hunting. Infrared night sights, like any equipment of military origin, will find a market among enthusiasts as prices become affordable.

2.5.5 Scientific

Throughout the history of infrared imaging, scientific research has been a major application for the newest technology. IR is used in a vast range of scientific fields from medical imaging to astronomical observation. Scientific applications often demand the very best fidelity available, either in terms of spatial resolution or temperature resolution.

.

Chapter 3, Replacing the existing processing with a 14-bit digitiser

<u>3.1 Analysing the design of the 100 x 100 pixel</u> pyroelectric camera

3.1.1 History of the camera

The camera used in this project is one of a series of cameras produced by QinetiQ (previously DERA) over a long period of time, with constant improvements and revisions to design. Several newer generations of camera were been produced after this one, they were all similar in operation, the only major changes being to the detector and the packaging. However, development of this line of pyroelectric cameras has now stopped completely and they are no longer in production at QinetiQ. The highest resolution detector that QinetiQ had produced as of 1999 was 380 x 240 pixels ^{R32}. Some models went into commercial production including compact handheld units which were sold to several military forces around the world ^{R33}.

The camera used in this project is a prototype design. The detector is a reticulated 100 x 100 pixel pyroelectric detector. The camera is designed to operate as a self contained system. It can run on internal batteries and has a built in eyepiece display. The camera is shown in Figure 3.1 next to a screen displaying the output of the camera and the desktop power supply used to power the camera.



Figure 3.1, Photo of camera

3.1.2 Chopper design

Being a pyroelectric device the camera has a chopper, this takes the form of a blackened aluminium disk with spiral blades. The aim of using spiral blades is to make the velocity of the chopper blade across the surface of the detector as nearly constant as possible, so that each pixel can be sampled at the same period of time after the chopper blade has passed minimising difference in response. The spiral shape is shown in Figure 3.2. The chopper is driven by a small electric motor wired to a phase locked loop to control the speed. The readout and control circuitry is slaved to the chopper speed so that the camera can operate even when the chopper speed varies slightly ^{R34}.



Figure 3.2, Diagram of chopper

3.1.3 Processing within the camera

The processing within the camera involves a complex series of steps which in the end produce image differenced output from the detector. Many of these steps are designed to make up for the limited resolution of the high speed ADC's available when the system was developed. At the end of processing the data produced is a measurement of the signal voltage at each pixel, with gain adjusted by certain parameters measured during construction of the camera.

When seeking to improve the performance of the camera, examining the operation of the processing system was an obvious first step, particularly since this element of the design had changed little in the newer generations of camera under development when this project was undertaken. A schematic of the internal operation of the camera is shown in Figure 3.3.



Figure 3.3, Overview of processing within the camera

The output of the detector undergoes analogue offset correction before being digitised at 8 bits. The offset removal is a method of removing the variations in mean output voltage between the pixels (sometimes known as pedestal noise). The need for offset removal is discussed in the next section, 3.1.4.

Once the offset is removed and the signal digitised, gain corrections are applied, from more values stored in ROM. Image difference processing (IDP) is applied to the gain corrected signal immediately before video output. The IDP takes the signal from when the chopper last opened and subtracts the signal generated by the chopper shutting. The subtraction removes remaining pedestal noise not dealt with by the analogue offset removal, as well as effects of output voltage drift. (Image difference processing and voltage drift are discussed further in chapter 4.) It also improves signal to noise levels by increasing the size of the signal (compared to monitoring one phase only).
3.1.4 The need for analogue offset removal

A sample trace of the detector output can be seen in Figure 3.4. The lower trace (1) shows the pixel clock, the upper (2) is the analogue output of the detector. (The pixel clock is a square wave signal with a period equal to that of one pixel, it is used to synchronise the processing within the camera.) This trace was captured on a Hewlett Packard digital oscilloscope with 200 MHz input bandwidth, a tool which was used extensively to analyse the detector output. The output of six pixels is visible on the oscilloscope trace. These represent six consecutive pixels in one row of the detector grid. The multiplexer switches between the output of the MOSFET amplifiers behind each pixel. The output level of each pixel is seen as a plateau in voltage. As the voltage changes quickly between pixels, the start of the voltage transition coincides with the rising edge of the pixel clock. This trace was captured with the camera pointed at a static scene at room temperature. The oscilloscope trigger was locked to the field identifier signal, consequently the trace only shows the open field of the camera output. The voltages seen here represent the typical output voltages for those pixels when the camera input is small (scene and chopper temperatures are similar). The output level of pixels in this trace only varies over about 0.1V but across the detector as a whole the variation is much greater, of the order of 2.5V for working pixels.

Analogue offset correction uses stored values to equalise the output voltage of each pixel before digitisation, a process made necessary by the limited resolution of the ADC in the camera. The mean output voltage of individual pixels falls between 0.0V and 2.5V, whereas the signal, the change in output voltage of a single pixel between an open and closed frame, is much smaller, of the order of 1.4 mV per Kelvin. If the full output range output of the detector (2.5V) were sampled directly at 8 bits per pixel, 256 discrete levels input levels could be recorded, and they would be separated by intervals of about 10 mV. The smallest difference in temperature the camera could detect would be around 7 K. The camera would be practically useless for imaging the natural environment.



Figure 3.4, Example image from scope

3.1.5 Analysing the weaknesses of the existing processing system

When the system was first examined for the project, the permanent offset correction values stored in ROM appeared to be a weak point in the design; the addition of a digital to analogue conversion process and the combination of two analogue signals, stood out as stages which could add significant levels of noise to the system. Analogue circuitry is prone to performance variations linked to temperature and calibration drift over time. Removing this stage of processing would benefit radiometric measurement by removing possible sources of error.

Considering the system as a whole, the resolution of the ADC stood out as the greatest limiting factor; it is this converter which necessitates the analogue offset removal. Even assuming the analogue offset removal functioned perfectly, 8 bits of resolution did not appear to be enough to extract all useful information from the signal. If the analogue processing were perfect, the 256 discrete levels of measurement would still limit either the resolution, or the dynamic range which the camera could handle. If the output of the array could be digitised directly with enough resolution to match the performance of the detector, then the circuit can be simplified and dynamic range improved, without any significant disadvantages.

3.2 Design and construction of a new digitiser

3.2.1 Selecting an Analogue to Digital Converter chip

As outlined earlier, the DERA pyroelectric camera was originally built with an analogue offset correction and 8 bit ADC to capture the signal from the detector. In order to improve the performance of the camera it was decided to bypass all the original processing circuitry. Since the time this camera was built the quality of ADC's available has increased remarkably. The pixel clock of this camera operates at 1 MHz. At the start of the project in 1999 several integrated ADC's capable of running at 1MHz were available, with resolutions of up to 24 bits. With this resolution there would be no need to use the analogue offset correction system used in the original design. The raw data from the detector could just be fed straight into the ADC and all processing could be done by digitally. This became the overall aim of the project, to devise a system to capture the raw output of the camera on a computer and to investigate processing techniques that could provide new function to the camera, or improve the performance of the camera.

After looking at the ADC's available it was decided to use a 14-bit ADC from Analog Devices, the AD9240, which is a three stage pipeline convertor. The pipeline converter operates on the following principle; a sample and hold circuit take a nearly instantaneous measurement of the input signal, this signal is used by three separate lower resolution ADC's in turn, each with higher gain than the last. The output of the first ADC is used to provide an offset for the input range of the second, this ensures that the signal voltage falls within the narrower input range which results from the higher gain of the second ADC. In turn the second ADC is used to generate the offset for the third. Digital logic processes the output of all three ADC's and calculates a 14bit result from them.

At the time, higher resolution, 16-bit and even 24-bit devices were available but all those examined were of the Sigma-Delta type, designed for monitoring a continuous wave form, not a switching (multiplexed) output such as that from this camera. A sigma-delta ADC is an increasingly common type of ADC now dominating the market in many applications. In its simplest form a sigma-delta converter consists of a single comparator, and feedback modulator circuit which serves to keep the input to the comparator close to the trigger voltage. Without going into unnecessary detail, the design of the feedback loop means that the density of ones vs. zeros coming out of the comparator is proportional to the voltage input to the sigma-delta. The modulator is in effect a tuned dither generator. A digital decimation filter sums the very high speed stream of single bit results to provide higher bit depth output values at lower frequency.

By running the digital filter for longer between outputs very high resolution can be obtained, often as high as 24 bit. The integrating nature of the system means it is naturally suited to continuous waveforms; discontinuities in the input waveform will reveal a settling time in the digital output. Typically 2 to 3 samples must be discarded after a step change in input before good data appears on the output. In order to sample the switched output of the multiplexer a sigma-delta ADC would have to run at least 3 times faster than the pixel clock rate of the ADC with intermediate samples being discarded.

Sigma-delta converters are often capable of running at very high clock speeds, but as the ADC runs faster the integration time decreases and the noise level increases. At the highest speeds the effective resolution will not match the bit depth of the output. Expert opinion was that the best sigma-delta ADC's available at the time would not match the performance of the AD9240 in this application.

Since work began on the digitizer, higher resolution ADC's have been released but the benefit they would provide to the performance of this camera has been shown to be minimal. With the input amplifier of the ADC adjusted to cover the range of the all functional pixels in the detector, the interval between digital levels represents a difference in input of less than 100 mK. This is shown to be below the noise level of the camera in chapter 4.

3.2.2 Selection of construction method

The choice of the AD9240 placed certain limitations on construction methods. The chip itself has quite a large number of contacts, 44 in total and is only available in surface mount packages. (In the context of integrated circuits (IC's) the term package refers to the plastic body of the device within which the silicon chip is housed and into which the contacts pins are mounted. Often a single component will be available in several package formats, which have different pin spacing (pitch), some formats allow socket mounting, others are intended for surface mounting. The pin pitch on surface mount chips is too fine to allow a circuit to be constructed on a strip board, something which is often used for prototyping circuits.

The simplest method of construction would be to design a printed circuit board (PCB) and have this made professionally. This possibility was considered. It appeared that the PCB could be included in a production job being undertaken by Sira Electro-Optics (SEO). The small circuit board might have been added to the same sheet of substrate as the SEO project, allowing professional production at little extra cost. However, during the design process of the digitiser the SEO project was delayed and the layout of their circuit was being reconsidered. At this point there was no exact production date and it was not even certain whether it would be possible to fit the ADC circuit on the substrate, so the decision was made to construct the circuit entirely by hand.

The circuit was constructed by gluing the chips upside-down onto a copper clad board, connections were made between components using insulated wires and the copper surface of the board served as a ground plane for the entire circuit. This method of construction is sometimes referred to as "hanging gardens" because of the way wires are suspended above the board. The upside-down mounting of IC's is often termed "dead bug" construction, the exposed legs of the chips sticking into the air, bare a resemblance to deceased arthropods resting on their backs. This method of construction was a painstaking manual process but offered certain advantages over a PCB. The circuit could be more easily modified and adjusted after completion, particularly if additional components were required.

Whilst a manually constructed board could not be expected to perform as well as a carefully designed multilayer PCB, with careful construction should it should be possible to reach acceptable performance levels even with a high resolution high speed convertor such as this. The raised position of the contact pins allows surface mount ceramic capacitors to be attached directly to the pins of the chips and directly to the copper ground plane. These capacitors serve a vital role in noise suppression, and mounting the capacitors in direct contact with the pins and the ground plane provides the lowest possible resistance between the pin and ground giving the best noise suppression possible. The continuous ground plane also provides protection from electronic noise induced by EM radiation.

3.2.3 Circuit design of the digitiser

The design of the digitizer board can be split into three main sections; the input stage, the ADC itself and the output stage. A simple diagrammatic overview of the digitiser is given in Figure 3.5, this shows the larger electronic components and external

connections, this diagram is roughly to scale. A full schematic circuit diagram is given in Figure 3.6, it shows all active and passive components in the finalised circuit, this figure will be referred to throughout the rest of this chapter. The only connections not shown in Figure 3.6 are those between the line drivers in the output circuit and the frame grabber, this cabling is specific to the frame grabber used. No electronic components sit between the line drivers and the frame grabber, so the digitiser could be reproduced from this figure.

3.2.3.1 Input stage

The input stage is essentially a preamplifier for the ADC chip. It is required to take the output voltage of the camera and produce a signal with the appropriate voltage range for the ADC chip. The input stage must also protect the ADC from signals that are too high or too low. This protection is very important because the output of the camera is not at all well behaved. There are several pixels in the detector which, through a failure at some point in the manufacturing process, do not function. These are termed dead pixels. The onboard processing of the camera has a record of these pixels in ROM, when these pixels are read during output the previous reading is substituted in their place; this system is bypassed by the new ADC. A significant problem with dead pixels arises when using the new ADC, because in the analogue detector output, rather than registering no signal, these pixels have a large negative voltage, as much as -10 volts, which would be enough to destroy an unprotected ADC chip.



Figure 3.5, Overview of digitizer





A sample of the detector output, including two dead pixels, can be seen in the Figure 3.7. The image is a screen capture from the digital oscilloscope used during testing. The upper trace (input 1) is the pixel clock from the camera, a 5 volt square wave. The trace in the middle of the image (input 2) shows the output of the camera over a period of about 20 pixels ($20 \mu s$). The scale is given on the key at the bottom of the image, 5 volts per division (divisions are represented by the grid, the finer markings are subdivisions), with an offset of zero volts, so the line across the middle of the height represents zero volts. Each pixel covers about 3 subdivisions of the x axis. The voltages of most pixels lie around 2V but in this trace two dead pixels can be seen. The first is third from the left, the pixel has an output voltage of about -1V. The second dead pixel appears in the right hand half of the trace, the 14th from the left, the output of this pixel is almost - 10V.





In order to protect the ADC a 'clamping' op-amp was used in the input stage. This is an operational amplifier which has high and low 'voltage clamp' inputs. The desired maximum and minimum output voltages are applied to these terminals, the amplifier will not produce an output voltage outside of that range, no matter what the input. The amplifier chosen was a CLC502. This amplifier is characterised by fast input settling, and rapid overload recovery. Input settling describes how rapidly the output voltage stabilises after a change in input voltage. It is important that the settling time should be low when monitoring a discontinuous waveform such as the output of this detector. Overload recovery is the process of returning to normal output after overloaded input has forced the amplifier into clamping mode.

Since the amplifier is expected to go into overload several times each frame, every time it encounters a dead pixel, it is important that the recovery should be quick, so that the subsequent pixels can be measured accurately. The CLC502 specifications give a typical settling time of 25 ns, (32 ns max) and typical overload recovery of 8 ns, (15 ns max). These parameters exceed the requirements of the experiment by a healthy margin, with the 1 MHz pixel clock, the period of each sample is 1000 ns.

The input to the CLC502 is protected from negative load by the diode D1. The high voltage clamp (V high) is set by the potential divider formed between R4 and R5, it is set to 2 Volts which is the upper limit of the input range of the AD9240. The low voltage clamp is set to zero volts by shorting V low to ground. Tuned CR (capacitance – resistance) circuits on the 5 Volt power input (+VCC), (R4,C1,C2) and V high (R6,C3,C4) protect the amplifier from high frequency noise on the voltage input. Such noise would be echoed on the output. The high capacitance polar electrolytic capacitors are used in parallel with much lower value ceramic capacitors because the ceramic capacitors have a superior high frequency response. Surface mount ceramic capacitors were connected directly between the pins of the amplifier and the copper ground plane to ensure absolute minimum resistance in the path, and so greatest suppression of high frequency noise.

The input stage also includes adjustable gain and offset controls, so that the input to the ADC can be adjusted to the ideal range. These take the form of two potentiometers, shown as POT1 and POT2 in Figure 3.6. POT1 controls the offset and POT2 the gain.

In common with other clamping amplifiers the CLC502 exhibits some degree of nonlinearity within 200 mV of the clamping voltages. Accordingly the gain and offset were adjusted such that all functioning pixels fell within the range of good linearity. Pixels with very high or very low offset voltages could be pushed into the regions of nonlinearity by very high or low IR input (accordingly). To ensure this could not happen would require a significant reduction in gain, compromising the resolution of the vast majority of pixels which fall nowhere near these non-linearities. As for the handful of pixels which will be affected to some extent, since the signal level is small compared to the offsets the most significant effect will be a reduced gain for these pixels, this will be removed by the calibration regime described in chapter 4. These pixels will also suffer from a degree of non-linearity in their response curves but the calibration is still somewhat robust to the problem, the nature of the high order polynomial fit allows the calibration curve to follow arbitrary non-linearities which fall within the calibrated range.

3.2.3.2 ADC Stage

The ADC stage of the digitizer consists of the ADC chip and its power supply, reference voltages and buffering circuits. There are capacitors connected to ground at the power and signal inputs to decrease noise which might adversely affect performance. Again, when high capacitance is needed small surface mount ceramic capacitors are used in parallel with electrolytic capacitors, to damp high frequency noise. For the analogue power (AVDD) and digital power (DVDD) inputs, separate capacitor sets (C6+C9 and C7+C8) separated by a low value resistor (R8) help to prevent any transfer of switching noise from the digital circuitry back into the analogue circuitry.

The output of the CLC502 preamp (V out) runs through two low value resistors (R11 & R12) to the signal input of the AD9240. The diode D3 protects the ADC from negative input. C18 is a very small (22 pF) capacitor between the input and ground. This serves to suppress noise at frequencies above the clock frequency. For example this will smooth the small spikes seen in the middle of each pixel's output in Figure 3.4.

One element of the ADC stage caused problems during testing. In the original design C19 was not fitted, and output of the digitiser was garbled, a digital signal was found but no usable data could be extracted. A lot of testing was carried out to little avail, until at one point, the sense pin was accidentally left unconnected. In this state the digitiser would work perfectly for a period of time, before encountering the same problem as before. This effect was traced to build up of stray charge on the sense pin triggering the AD9240 to flip into 5V input range. From this behaviour it was deduced that the original problem was caused by electronic noise on the sense pin causing the AD9240 to flip between 2 V and 5 V input modes. The addition of the larger capacitor C19 in parallel to the existing C11 damped the noise and left the digitiser working perfectly.

3.2.3.3 Output Stage

The output stage consists of five DS34C87T 'Quad TRISTATE Differential Line Driver' amplifiers made by National Semiconductor. These are amplifiers designed for sending high bandwidth digital signals along medium length cables. Each chip has four channels, between them the five chips carry the 14 data bits, the pixel clock, the field identifier, and the frame sync. The term differential refers to the fact that two cables are used to carry each channel of data. When the signal is logic 0, one cable carries a positive voltage and the other a negative voltage, when the signal is logic 1, the voltages reverse. The cables are twisted around each other as a pair. The equal and opposite currents in each cable help to reduce the amount of electromagnetic noise radiated by the conductors and so reduce cross talk in cable bundles which carry many data channels in parallel.

The digital stage of this circuit is undemanding as far as signal integrity goes. There should be very little chance of bit errors. Any similarly specified line driver could be used in place of the DS34C87T without modifying the rest of the digitiser.

3.2.3.4 Earthing

Careful grounding and elimination of earth loops is essential to achieving good performance from a high resolution convertor such as this. Fortunately in this case it was not difficult to eliminate earth loops in circuit. The power supplies for the camera and the ADC board were independent and both were fully isolated "floating voltages". The entire system could be earthed through the shielding of the frame grabber cable, to the chassis of the PC. Fourier transform measurements showed good rejection of mains frequency noise and no significant frequency spikes in the noise spectrum which would expose the resonances which can be found in earth loops.

3.2.4 The frame grabber

A digital frame grabber was used to capture the output of the digitiser for computer processing. Being a 100% digital device, so long as the signal falls within the capabilities of the frame grabber, it should work without loosing any data and performance parameters are uncritical. However it will be discussed in some detail here since problems were encountered and specific workarounds used, which may be of interest to anyone working with similar equipment.

A Meteor II Digital frame grabber made by Matrox Ltd was chosen because it was a well specified product. However, with hindsight it was not a particularly good choice

and led to problems, a peculiarity of the camera design proved incompatible with the frame grabber driver software. The camera operates without a line-sync signal. There are pixel and frame sync signals to represent the start of each pixel and new frame, but not a synchronisation signal for each line. Matrox frame grabbers are designed to operate with a line-sync. This caused quite severe problems with the frame grabber software. Despite a lot of experimentation and many attempts to get help from Matrox, it was never possible to make the frame grabber synchronise properly at the start of a grab. Whatever way it was configured, the first half frame's worth of pixels were corrupted. An acceptable workaround was developed, which involved making a single very large grab into an image buffer, equivalent to 202 frames. The frame grabber and software treats this as a single very large frame. The first frame pair, including the corrupted data, were discarded. The rest are then reformatted appropriately. This produced a continuous sample of two seconds, long enough to allow any continuous processing that was needed. If longer video clips were required only two frames were lost in every 202. Using this technique rules out any possibility of real time processing, since the data in the grab buffer is not accessible until the grab finishes. i.e. A two second delay. This was not a problem, since there was no need to carry out processing in real time during experimentation.

There was additional quirk of the digitiser system as a whole, caused by the fact that the ADC produces a three cycle lag in the passage of the data signal. Since the frame-sync and field identifier do not pass through the ADC they are not lagged. This offset had to be corrected for when processing the data.

The field identifier signal has been mentioned in this chapter, this is a signal from the camera which indicates whether the current frame was produced by the opening or the closing of the chopper, (i.e. a square wave with period 0.02 s, equivalent to 2 frames). This signal had to be fed into the frame grabber somehow but the only time most video systems use a field identifier is when the signal is interlaced. That is not the case with this camera but providing a field identifier signal to the frame grabber caused the software to attempt to de-interlace the signal automatically, reversing this process would have been unnecessarily complicated. The practical solution used was to tell the frame grabber to capture at 15 bits per pixel and feed the field identifier into the frame grabber as the most significant bit. That way the captured signal consists of light frames followed by dark frames, so it is possible to tell what field a pixel comes from, simply by its magnitude. This made formatting and recording data with the frame grabber easy

and quick. It also made it possible to check that data was synchronised correctly in the frame grabber software without performing IDP. All processing of the data was carried out off-line in IDL.

.

3.3 Experimental analysis of digitiser performance

Getting the digitiser to a fully functional state, where data could be captured from the camera reliably, took experimentation and some trial and error. This can be considered part of the design process as discussed earlier in this chapter. Once the digitiser was shown to be functional, by capturing images from the camera, it was necessary to measure its performance. Aside from the characteristics of the ADC chip itself, which are fixed, the performance of the digitiser comes down to one factor, noise level. In this section the various possible sources of noise are discussed. This is followed by a description of the experimental methods used to gauge the level of noise contributed by the digitiser.

3.3.1 Sources of noise in the digitiser

Any noise generated by the digitiser board will affect all readings taken. The noise errors introduced by the digitiser can be split into two groups, noise in the analogue circuit and noise in the digital circuit.

3.3.2 Noise in the analogue circuit

This is noise that gets added to the camera signal by the electronics of the digitiser. There are several possible sources for this noise which are discussed individually in the following sub-sections.

3.3.2.1 Power supply

The power supply used for the digitiser may contribute noise. The power input is buffered with capacitors to reduce this contribution but this will not remove all noise from the power source.

3.3.2.2 Ground noise

Ground noise is another contributor. The power supply for the digitiser is not connected to mains ground, it is in effect a "floating voltage", but the ground plane of the digitiser is connected to ground plane of the camera. The camera itself is also disconnected from mains ground but when linked to the computer all three components are connected to mains ground level through the digital signal cable. This means that any noise that the camera or computer introduce to the ground level could contribute to noise in the digitiser. With a true zero resistance path to ground any such noise should be small but in a real world case it may become significant. The cables used between computer digitiser and camera have a finite resistance, they also have capacitive and reactive components producing an LCR circuit between the components the impedance which is dependent on frequency. When dealing with high resolution, high speed digitisers ground noise can become significant.

3.3.2.3 Interference noise

The digitiser is reasonably screened from electromagnetic interference but with a computer, several power supplies, three CRT's and fluorescent lights operating within a few metres of the digitiser, this could make a significant contribution to noise on the signal.

3.3.2.4 Electronics within the digitiser

In addition to the external sources of noise there are possible sources within the digitiser itself. The signal from the camera is not passed to the ADC directly; a clamping operational amplifier (op-amp) is used to provide the correct signal level to the ADC and to protect it from voltage outside its operating parameters, a condition known as overloading, which can damage components. All operational amplifiers produce noise, that in a well designed circuit, should be small in comparison to other sources. The op-amp should not reach its upper clamping voltage during normal operation but it does reach the lower clamping limit several times during each camera frame, at each dead pixel. The dead pixels produce a voltages as low as -12V on the camera output. As the op-amp hits its clamping voltage, noise is produced on the output, again this should be small but being aware of it as a possibility could have been useful if problems arose. The op-amp used was chosen specifically for its low noise characteristics this helped to ensure it did not become a source of problematic noise.

3.3.2.5 Other analogue noise

Microphonic pickup is a problem that can affect any electrical circuit not just audio electronics. Coaxial cables have a significant capacitance. Flexing or stretching these cables will change their capacitance; this can introduce an electrical signal into the circuit they are connected to. All readings were taken in a quiet lab away from any machinery so the equipment was not exposed to any large vibrations during operation. Each reading took only a few seconds and they were all were taken manually. If any large incidental vibrations had occurred, doors slamming and the like, they could have been noted and the reading retaken. This was not needed as no vibrations were detected.

The ADC itself may induce small amounts of noise into the circuit. Most commonly switching noise from the digital circuitry can feed back into the analogue circuitry. ADC chips are designed to minimise this effect internally. So the most common method of transference is through ground or power connections. That is why the AD9240 has separate analogue and digital pins for ground and power. The capacitors mounted between the analogue and digital inputs should help to minimise any noise transfer but finite amount of noise will be present and even a small error in assembly, one poor electrical connection, could increase the transfer significantly.

Along with all these sources there will be noise introduced by the passive components, noise in the resistors and the solder joints of the signal path, in well constructed circuit these effects should be immeasurably small. Care was taken throughout construction to ensure good connections but it remains a possible source of problems and is another reason why a digitiser's performance should be tested.

3.3.2.6 Analogue noise summary

All of the above sources should introduce only small levels of noise and in a perfect digitiser each should not add greatly to the overall noise in the system. Added together, however, the combined level could become enough to show up on the output at a low level.

It is possible that the total noise introduced by a digitiser into the analogue stage of a circuit is much smaller than the quantisation level of the ADC, in this case analogue noise in the digitiser might have no significant contribution to the noise level in the digital output. However the noise level will have significant effects before the noise reaches the scale of one digital unit. Take the situation where there is a constant input level, if this level happens to be close to the boundary between two digital levels then even a small amount of noise originating from the digitiser could cause the digital output to vary randomly from one of the two levels to the other even if the input signal itself is very clean. This is something to bear in mind in situations where one needs high accuracy down to the least significant bit of the ADC scale.

Ironically, a small level of noise operating in this manner can improve the accuracy of measurements taken with a digitiser. The noise causes the output of the ADC to flip between two adjacent levels. Over a large number of samples the proportion of high readings to low reading will be proportional to the true input voltage. So taking an average of a large number of samples will produce a value closer to the true input voltage than the case where there is no noise at all. In some systems a controlled amount of noise is deliberately added to the input of an ADC, a technique known as dithering. Dithering does not reduce the total noise power in the digital output; rather it shifts the noise energy further up the frequency spectrum. If high frequencies close to and above the Nyquist frequency are later removed from the signal the total noise remaining is reduced. Since dithering is more relevant to the recording of continuous waveforms than to the sampling of multiplexed data, as in the case of this digitiser, it was not considered during the design process.

3.3.3 Digital noise

There is also noise generated in the digital circuit. Noise here should not be enough to introduce an inversion in the digital bit stream but this would be possible if there was a flaw in the construction of the circuit. So long as the circuit is functioning correctly the only noise in the digital signal should arise from non linearity of the ADC. This may seem to be disassociated from other problems that may be present in the digital circuitry but produces a source of error on the output scale and is not present in the analogue circuit, it is therefore a form of digital fixed pattern noise.

Nothing in the circuit can affect this problem, it is down to the design of the ADC. The ADC used was specifically chosen for its good linearity figures, this will help to reduce the effect of non-linearity on the quality of output.

Analysing all of these noise sources independently would have required a lot of work and would not necessarily have detected problems that could arise only when the entire circuit was fully assembled and connected to the camera. Instead the approach used was to consider all of these noise sources during the design of the digitiser circuit and to examine the total noise contribution of the digitiser after construction. This noise was found to be well below the noise level of the camera signal itself so the digitiser was considered satisfactory for the job.

3.3.4 Investigating digitiser noise

3.3.4.1 Experimental design

An experiment was devised where the analogue data signal from the camera was replaced with a continuous voltage provided by an alkaline cell, to measure the noise that would be present in normal operation, the circuit was kept as close as possible to that used during normal operation; the same interconnects were used where possible and the same equipment was turned on in the room. The camera was turned on and was used to provide the timing signals needed to operate the frame grabber in exactly the same manner as when the camera was being used normally. The 1.5V nominal voltage of the alkaline cell produced a signal in the middle range of the input scale of the ADC. A single AAA cell was soldered to a short length of wire connected to a miniature BNC connector, this allowed the battery to be connected to the digitiser by the same length of coaxial cable used to connect to the camera.

3.3.4.2 Data analysis

The captured data was processed in a very similar way to data from the camera. It was split into frames in the same way and individual samples treated as pixels. Since the control signals from the camera were connected to the digitiser board, there was the possibility of noise with a frequency component linked to the operation of the camera, the frame clock, the line clock or something else similar, formatting the data in the same way might make any such noise easier to spot.

500,000 sequential battery voltage samples were taken. These were loaded into IDL and various approaches used to analyse the noise present within them.

3.3.4.3 Standard deviation

The standard deviation on the battery signal, using all samples, was calculated to be 2.26 ADC units.

3.3.4.4 Fourier transform

A Fourier transform method was used to inspect the data for frequency components. The plot can be seen in Figure 3.8. The transform yields a flat horizontal band with random variations of equal amplitude along its whole length. This is a good indication that the noise on the input is truly random with no dominant frequency components.



Figure 3.8, Fourier transform of battery voltage

3.3.4.5 Inspecting raw samples

Plotting sequential readings from within these data provides an easy way of inspecting the output for any strangeness in the distribution of samples which might not show up in either the standard deviation or the Fourier Transform. The graphs show the majority of readings fall within 2 units of the mean value. The largest deviations fall about 10 units either side of the mean. An example is shown in Figure 3.9. The plot in black is of sequential samples, the plot in red is of the same number of samples drawn randomly from the full set of data. The randomly selected samples look very similar to the sequential ones, with no discernable difference in pattern between them. This demonstrates that the signal is stable over the course of the whole data set. That is to say there are no very low frequency components to the signal.

Several more plots were made to look for any patterns that might be associated with timing signals from the camera. Figure 3.10 has a number of lines plotted representing different intervals between samples in the set. Details of each line are given in the key. Lines show samples separated by one line, one field and one fame (two fields) none of them appear to show any significant pattern.

3.3.4.6 Histogram

The frequency distribution of values recorded for battery voltage was inspected by producing a histogram of the sample distribution. The histogram is given in Figure 3.11. The histogram shows the distribution of values forms a bell shaped curve. The plot is made rather angular by the digital nature of the input signal, there appears to be a hump in one side of the curve, this could be an effect caused by the discrete nature of the data values or it could even be an effect of very small scale non-linearity in the ADC. It does not appear to be something which would have a significant effect on the measurements made with the camera attached.



Figure 3.9, Digitizer measurements of battery voltage, showing sequential and random readings



Figure 3.10, Digitizer measurements of battery voltage, sequential and non-sequential readings

91



Figure 3.11, Histogram of battery samples

3.3.5 Conclusions on digitiser performance

No proof of excessive noise was found by looking at readings taken using the battery but this is what was hoped for. There is no evidence that any noise is introduced by the control signals of the camera. There is also no evidence of frequency dependent noise arising from the ground plane connected to the computer and the camera or though electromagnetic interference. All of these things suggest that the digitiser was well constructed and was operating at near optimum noise levels.

There is no evidence of noisy bits on the digital stage of the circuit; these would produce a distribution with jumps of discrete numbers of units between each reading.

There is a level of background noise on the readings with a standard deviation of 2.26 units. This is not perfect; the spec sheet for the AD9240 claims a noise level of 0.36 units under ideal conditions, but because of the higher noise level on the camera output it proves to be of little significance. There is nothing to suggest that taking the mean value of several samples could not provide a measurement of a clean voltage input to an accuracy of 1 ADC unit.

What has been described here is not an exhaustive set of performance tests for a digitiser, but enough to show it is capable of matching the performance of the detector in both resolution and noise.

Chapter 4, Modelling and measurement of the camera's radiometric performance

Having built and tested the digitiser it was necessary to assess the performance of the camera when used with the digitiser. There were two goals to this experimentation, the first being to ensure the digitiser was functioning optimally and determine whether it would be a limiting factor in the performance of the camera. The second goal was to ascertain the suitability of this camera for making radiometric temperature measurements. To facilitate radiometric calibration a platinum resistance thermometer (PRT) was mounted within the camera body, close to the chopper blade. A temperature calibrated high emissivity "black body" source was used to provide signal for calibration. This chapter covers the theory of operation of the pyroelectric detector and how an additional temperature sensor can allow radiometric temperature measurement. Theory is used to predict the response profile of the camera and experimental results are compared to this model. Typical radiometric performance camera performance criteria are calculated for this camera. (Within this work the term "detector element" or just "element" will be used to mean the active area of an individual pixel within the array. That is, the pyroelectric material, the absorber and electrodes, but not the bump bond, or the substrate, the CMOS readout chip.)

4.1 Model of the detector operation

4.1.1 Pyroelectric detectors in detail

4.1.1.1 Time constant and signal

Chapter 2 described how a pyroelectric detector is sensitive to changing temperature. To be used as a thermal infrared detector, a pyroelectric sensor such as this must be presented with a chopped input, which is achieved, in this case, by using a rotating chopper wheel close to the plane of the detector. The design of the chopper wheel was described in Chapter 3.

Figure 2.2 gave an electrical analogue of a thermal detector; remembering that the current source represents incident radiation, and the voltage output represents a temperature difference between the detector element and substrate, it is clear that a change in incident radiation will set the detector radiation moving towards a new

equilibrium with a time constant as described in section 2.3.2.2. Since the time constant is fixed, the rate of temperature change at any point in time, dT/dt, is directly proportional to the difference between the current detector temperature and the equilibrium temperature towards which it is heading.^{R18}

If the thermal capacity is constant and the thermal impedance is linear then equilibrium temperature is directly proportional to incident radiation power. Therefore when exposed to a sudden change in incident radiation, such as during the opening of the chopper, dT/dt is directly proportional to the change in incident radiation power. If the electrical response of the pyroelectric to dT/dt is linear over the small ranges generated by changes in incident infrared radiation then the chopped system can be considered linear with respect to radiation power.

4.1.1.2 Linearity of thermal impedance

Linearity of the thermal impedance was mentioned above because there is a good reason that it may not be linear. The total thermal impedance is a combination of all heat-loss paths from the detector element; this includes thermal conduction as well as heat loss by black body emission. Thermal conductivity is a linear phenomenon at the temperature scales dealt with in this system, but, as we have seen, black body emission is not, it has a fourth order relationship to temperature, so the total thermal impedance is not completely linear. However, within this system, thermal conductivity is the dominant heat-loss path, this can be demonstrated mathematically.

Consider a single pixel within a pyroelectric array, in equilibrium at a temperature of 300 K, the pixel is $100 \mu m$ square, and has a solder bump bond on the rear surface, consider the solder bump as cylinder of solder 50 μm in diameter and 50 μm high. (These dimensions approximate those in this camera.) If the pixel is exposed to increased radiation input such that its new equilibrium temperature is 301 K, the additional heat-loss through radiation and conduction can be calculated as follows:

Loss through radiation:

Assume the front surface to be a perfect black body radiator, since it is coated with an absorber. The sides and back are not coated and so have low emissivity, in addition, the sides are small because the element is thin compared to its area. Radiation from the sides and back will be ignored.

The additional power radiated by the front surface will be equal to the black body radiation at 301 K minus that at 300 K. From the Stefan-Boltzmann Law, and the area of the surface:

Black body radiation at 300 K:

$$P_{300} = \sigma T^{4} (100 \times 10^{-6})^{2}$$

= 5.67 \times 10^{-8} \times 300^{4} \times (100 \times 10^{-6})^{2} = 4.59 \times 10^{-6} W

And at 301 K:

 $P_{301} = 5.67 \times 10^{-8} \times 301^{4} \times (100 \times 10^{-6})^{2} = 4.65 \times 10^{-6} W$

The difference:

$$\Delta P = 6 \times 10^{-8} \quad W$$

Loss through conduction:

Additional loss through conduction can be calculated from the thermal conductivity of the solder bond and its dimensions. The thermal conductivity of 63% Sn, 37% Pb solder is 50 W/m K.^{R37}

Heat-loss by conduction, 1 K gradient, 50 µm diameter by 50 µm deep bond:

$$P = \frac{CA}{L} = \frac{50\pi(25\times10^{-6})^2}{50\times10^{-6}} = 1.96\times10^{-3} W$$

When the element's temperature rises by 1 K, over 30,000 times more of the additional heat-loss is by conduction than by radiation. Even if smaller bonds and mesa isolation were to reduce conductive loss by two orders of magnitude; the loss through radiation would remain virtually insignificant. So, black body radiation from the elements will not significantly affect the linearity of the detector.

From this we can also see that to raise the detector element's equilibrium temperature by one degree will require around 2 mW of radiation power, equivalent to a black body at over 1600 K filling the field in front of the detector. To achieve a target temperature resolution of better than one kelvin the detector will have to be sensitive to temperature changes of the order of microkelvin.

4.1.2 Thermal environment of the detector element

Whilst this camera is operating with a steady input and steady ambient temperature the detector elements will have, in effect, two thermal equilibrium temperatures, one for the closed field and one for the open field. The element will normally not reach either equilibrium point, but it will start to move towards the new equilibrium point after each chopper transition. Averaged over an entire chopper cycle, open and closed fields, the detector element will have a mean temperature which will be constant during steady state input. The thermal behaviour of the detector element over time during steady state input is shown graphically in figure 4.1. The top graph shows the temperature of the element and the points of transition between open and closed fields. The lower graph uses the same time base and shows dT/dt, the input to which the detector is sensitive. As described in chapter 2, the signal from the element does not appear instantly after the change in dT/dt, because it takes some time for the charge to build, the point of read-out is delayed after the point of transition accordingly.



Figure 4.1, Thermal behavior of pixel during steady state operation

To get a useful signal from the detector the pixel value of each open field is taken and the value of the previous closed field is subtracted from it, yielding a signal which is proportional to dT/dt. This process is termed image difference processing (IDP), as we have seen; depending on the behaviour of the pyroelectric material this output should be linear with respect to radiation power.

A change in input results in a new mean radiation power integrated over the chopper cycle, and so will cause a change in mean temperature; this is depicted graphically in figure 4.2. The new equilibrium mean temperature is not reached instantly, the mean temperature integrated over one cycle, will move towards the new equilibrium according to the time constant of the detector. This effect will introduce an error in the signal; looking at the graph, taking $o_1 - c_1$ for IDP produces an artificially large result because c_1 was recorded at an earlier point in time than o_1 and the "mean" temperature has changed between the two readings. Changes in ambient temperature will result in a changing substrate temperature and will have a very similar effect.



Figure 4.2, Thermal behavior of pixel during dynamic use

The erroneous signal fades as the detector mean temperature progresses to wards a new equilibrium. However during the unstable period the erroneous signal can be largely compensated for using a technique called multipoint IDP^{R35}, where, instead of subtracting the preceding closed field from the open field, the average of the preceding and the next closed field is used. This is illustrated on the graph.

Since the progression towards a new average temperature is not linear this correction algorithm is not perfect, but given the size of the time constant it is a good approximation. Naturally, nothing comes for free and the payback in this case is that output is delayed by one field, (half a frame), and some temporal resolution is lost. With the field rate of 100 Hz the delay is 10 ms and will usually be imperceptible to a user. However, it could be significant in computer controlled applications.

During the performance analysis later in this chapter most readings are taken in a steady state where multipoint IDP is of little benefit, however some experiments with changing ambient temperature are described at the end and multipoint IDP is used accordingly.

4.1.3 Contributions to the signal and the purpose of measuring chopper temperature.

During the entire chopper period there are three primary sources of radiation incident upon the detector, these are; radiation through the lens, the chopper blade and the interior surfaces of the detector chamber (including the lens body). A schematic diagram of the detector chamber is given in figure 4.3, approximate values are given for relevant dimensions. During the closed field, the proximity of the chopper to the detector surface is such that it will fill almost the entire 180° field of view of the detector element, additional oblique radiation will arrive from emission from the side walls of the detector chamber, and some radiation from the lens may reflect off surfaces within the chamber and eventually arrive at the detector. However, all internal surfaces are blacked with high emissivity paint and radiation received from the lens during the closed period would arrive after several reflections within the chamber so the contribution will be small. In addition, the oblique nature of this radiation means much of it will be reflected from the germanium window forming the front of the detector package. Radiation internal to the detector package, including that reflected back to the detector off the inner surface of the window, does not vary with the chopped cycle and so does not contribute to the signal.



Figure 4.3, Diagram of detector chamber, showing position of chopper, lens and PRT relative to the pyroelectric array

During the open field the detector receives radiation from the lens, which subtends approximately 110° at the detector surface; as such it represents the primary source of radiation during the open cycle. Parts of the chopper wheel still fill some of the field of view; radiation from these has not changed with the opening of the chopper. The rest of the field of view will now be occupied by areas of the front surface of the detector chamber; these will contribute to the radiation arriving at the detector. If these surfaces do not have the same emissivity as the chopper wheel, or are at a different temperature, a signal will result. In this camera the same paint is used on both surfaces, but if they are at different temperatures there will be a signal. The motion of the chopper wheel produces some stirring effect on the air in the chamber, which will help to maintain thermal equilibrium between chopper and chamber. Since the chamber body is aluminium, as is the chopper, their thermal conductivity is good and whilst the camera is in thermal equilibrium temperature differences between the chopper and the chamber front wall should not have a major contribution to the signal. If the environmental temperature is changing this will not be the case, and an erroneous signal offset will remain after multipoint IDP. Temperature gradients, induced within the camera body by changing environmental conditions, can affect the performance of any infrared camera, particularly radiometric performance. Such temperature effects are sometimes referred to as "thermal shock", the term comes from the tendency for the problem to arise quickly after sudden changes in environment, and the way it recovers gradually once conditions are stable.

It is clear that in order to perform radiometric measurements with this camera it is essential to know the chopper temperature, to this end, a PRT was added to the detector chamber in a position illustrated in figure 4.3, it is also shown in a photograph in figure 4.4. The PRT was of the "film" type, which have lower thermal capacity, and hence quicker response, than the wire-wound type. The position of the sensor was chosen to give as good an indication of the chopper temperature as possible. The intention was to place the PRT such that it intersected the volume of moving air created by the spinning chopper, and hence produced a good thermal link with the chopper temperature. The obvious way to get a more accurate measurement of the chopper temperature would be to use a single point infrared temperature sensor, suitable thermopile sensors sell for a few dollars or more depending on specifications. Such a solution would probably have to be designed into the camera; there was no practical way to add this to the existing camera without making major changes to the chassis.



Figure 4.4, Photograph of PRT sensor sited inside detector cavity

It is worth making a brief comparison between radiometry with this pyroelectric detector, and with the bolometer detectors which now dominate the market. With a bolometer, it is vitally important to know the substrate temperature of the detector, since the signal manifests itself as a steady temperature difference between the detector element and the substrate. The AC nature of the pyroelectric detector removes the dependence on the substrate temperature, explaining the lack of a temperature sensor built into the detector package. There remains the possibility that the sensitivity of the pyroelectric material varies with temperature, and this was investigated by repeating experiments with the camera at different temperatures. To do this a temperature control system was made by pumping fluid through a large copper heat exchanger, which was designed to cover the entire front of the detector chamber.

4.1.4 Modelling the response of the detector to a black body

A temperature controlled black body source was used to measure the response of the camera, this consisted of a flat radiating plate coated with a high emissivity textured paint, specified as having a mean emissivity of over 0.97 across the $3 - 14 \mu m$ range,.(Sira Electro Optics internal data). Chapter 1 described how the spectrum of radiation arriving at the detector from a black body can be modified by a number of factors, including the atmosphere and optics, and how this can affect radiometry. The

camera was known to have some form of LWIR band filter in the detector package but the exact specifications were not known, the spectral response of the camera could have been measured with a monochromater, but there was not one available for this project. Since the passband will affect the linearity of the detector response vs. target temperature, having a good idea of the passband of the camera is important in interpreting the response to the black body. To this end, the power of black body radiation within various passbands was modelled, so that the curves could be compared with the camera's response, to confirm the nominal 8-12 µm specification of the camera. Black body radiation within each band was calculated by dividing the band into 100 intervals and using Planck's Law to calculate the spectral radiance at the midpoint of each interval. This value is multiplied by the interval size to approximate the power within that interval. The power contributed by all 100 intervals was summed to give the total power within the waveband.

Transmission through the atmosphere had to be considered in order to make a fair comparison between a the nominal 8 – 12 μ m figure and what might be expected if there were no high pass filter in the camera. During testing the separation between camera and black body was fixed at 1 m, figure 4.5 gives a plot of infrared transmission through one meter of atmosphere.^{R38} Black body power in the 3 – 12 μ m range as viewed through 1 m of atmosphere was simulated by dividing the full range into three sections: 3 – 5 μ m, 5 – 8 μ m and 8 – 12 μ m. In accordance with the transmission graph, the power from the 5 – 8 μ m range was reduced by 20%, this value was summed with the 3 – 5 μ m and 8 – 12 μ m values to simulate transmission through 1 m of atmosphere. These curves are shown in the first graph of figure 4.6, a point to note is the differing levels of curvature seen between passbands, this is illustrated more clearly by plotting intensity rise over the 0 – 100 °C range normalised to one, which is given in the second graph. The camera's response profile is expected to most closely match the curve for the 8 -12 μ m band.





Figure 4.6, Black body radiation power within defined passbands

4.2 Experimental Methods

The following sections analyse the response of the camera to a black body at different temperatures, this section describes the testing methods and those elements of the hardware which have not yet been covered. The broader elements of this methodology also apply to the readings taken in a dynamic temperature environment, discussed later in this chapter, and the scene readings covered in chapter 5, any changes are described where relevant.

With the camera temperature kept constant, the black body source was placed in front of the camera such that it filled the field of view at a distance of 1 m. The black body was varied between room temperature and 100 °C and the output of the camera captured at each point.

The camera temperature was controlled by pumping fluid through the heat exchanger fitted to the front of the camera; the fluid temperature was controlled by a calibrated fluid bath with both heating and cooling capability. The camera body was clad in foam insulation and suspended above the bench on polystyrene foam blocks, in order to reduce thermal gradients within the detector enclosure. The camera was powered on along with the fluid bath, and the system was left running for 120 minutes to equilibrate before any readings were taken. Observing the rate of change in temperature of the PRT sensor in the camera during early capture sessions suggested that two hours was considerably longer than necessary for the temperature to stabilize, but the two hour period happened to fit in with the laboratory work schedule, so it was kept for subsequent capture sessions.

A temperature controlled black body device was used as the target. The black body had a flat target area in the form of a metal plate coated with Nextel brand high emissivity textured black paint. The back of the plate was heated with a network of four high current resistors with integral heat spreaders to provide good contact to the plate. A companion unit provided a power supply and thermostatic control for the unit. For improved temperature measurement a fine hole was drilled towards the center of the target plate, from the side edge. A PRT sensor was fed into the hole, to provide temperature measurements from the centre of the plate. The black body was left for 10 minutes to reach equilibrium at each temperature level. Shutters were drawn to prevent sunlight affecting results. The two PRT's used, one in the camera and one in the black body, were wired with twin runs of cable to allow 4 wire resistance measurements for improved accuracy. Once wired up the two PRT sensors were calibrated against a third PRT which had been calibrated in its wired state by SIRA Testing & Certification. Calibrated 4-wire meters were used to read the PRT's at each point.

4.3 Verifying Digitiser Performance and Assessing Array Uniformity

The performance of the digitiser itself had been tested and the results were covered in Chapter 3. These tests verified that the digitiser was functioning with the framegrabber and not producing excessive noise. Various tests were carried out to ensure the digitiser's proper functioning with the output of the camera. These tests looked for evidence of clipping by the digitiser, or distorted response caused by the analogue electronics, including the clamping op-amp. The distribution of certain pixel parameters across the detector array was investigated; this might reveal problems in the multiplexer, bump bonding or timing of the chopper.

4.3.1 Using linear fitting to compare pixel sensitivity

The first tool used to analyse the response of pixels was a mathematically calculated linear fit. The IDL LINFIT function was used to achieve this. LINFIT fits data to a linear model by minimizing the chi-square error statistic. An example fit to a typical pixel response curve is shown in Figure 4.7. The non-linear nature of black body radiation means that the linear fit does not accurately describe the response to temperature, but it is a valuable tool for analyzing the relative sensitivity of pixels, the gradient of the fit gives a good measure of the relative sensitivity across the array and naturally reduces the effect of inaccuracy in any individual reading.



Figure 4.7, Response of one pixel with linear fit used to measure relative sensitivity
4.3.1.1 Profiling the distribution of detector sensitivity

Figure 4.8 shows a histogram of the gradient values over all pixels, which are distributed around a large peak at the high end of the range. The around 95% of pixels fall within the shape of the peak, but there is also a distinct tail to the right hand side of the peak, representing a group of pixels with higher than average signal, this feature was noted for later investigation. There a scattering of pixels with lower gradients, outside of the bell shaped peak, these were taken to be dead or "nearly dead" pixels, their performance was investigated later to determine how best to deal with them.



Figure 4.8, Histogram of linear fit gradients

Figure 4.9 shows an image of the gradient values across the detector. This can be thought of as an image of the sensitivity of pixels within the array. The bright pixels represent those with the steepest gradient. The image shows an overall gradation of response from top to bottom, with the detector most responsive at the bottom. Six dead pixels can be seen as completely black dots. There are also a number of dark blotches, of approximately 10 pixels in diameter, clustered towards the upper right of the image. A certain amount of horizontal striation over alternate lines is also visible.



Figure 4.9, Image of sensitivity (white high)

It seems likely that these different structures in the response image are caused by different aspects of the detectors design and construction. The 'blotching' may be caused by some kind of flaw in the solder bump bonding process or in the manufacture of the detector or readout layer. The striations might be traced to some aspect of the readout system. On the other hand, it was thought that the general shading around the bottom of the image could be caused by some controllable element of the design, for example, the timing of the chopper. The nature of this phenomenon was investigated further and the results of that work are included in Chapter 5.

4.3.1.2 Inspecting the behaviour of pixels with highest sensitivity

The histograms were used to select groups of pixels according to certain characteristics, to examine their performance and distribution across the detector. Reverse lookup was used to identify those pixels which form the tail at the high end of the gradient histogram. The distribution of those pixels is shown in the binary image in Figure 4.10, which shows that these pixels all lie within to bottom six rows of the array and form the vertical gradation noted in Figure 4.9.



Figure 4.10, Binary image of sensitivity, highest in black

The response curves of the pixels with highest gradient were plotted so their shape could be inspected. In this section sets of large numbers of graphs are presented, laid out in grids. (In each grid the axes and scale are the same for each graph and are not shown on each graph to save space.) Response curves for four pixels are given on each graph. When the curves are similar in shape it is not possible to make out each line, this is not a problem because the function of the graphs is to identify pixels which do not conform to the general trend. This method of graphing was used during the experimentation to allow the response curves of very many pixels to be inspected with relative ease. They are presented here in the same manner as a demonstration of that method. The sets of graphs are ordered starting in the upper left corner of the grid and progressing in rows of four, left to right. Subsequent rows are positioned below the previous row.

Graphs of the 64 pixels with the highest gain can be seen in Figure 4.11, (sixteen sets of four). The graphs show that over this range of 64 pixels the gradient decreases very smoothly from the maximum without any discontinuities. All pixels exhibit a similar curve profile. The shape of the curves, the steady increase in gradient all the way up to the top of the curve, suggests the output of these pixels is not being distorted by non-linearities in the analogue electronics within the detector.



Figure 4.11, Response curves for the 64 Pixels with highest gain (image difference)

4.3.1.3 Inspecting the behaviour of pixels with low output gradient

Figure 4.12 gives another 16 graphs in the same format as Figure 4.11; these represent the 64 pixels with lowest response. In the first two graphs 6 dead pixels are visible as horizontal lines at an output level of zero. The following 58 pixels show a steady increase in gradient. There are visible jumps in gradient over the first twenty four pixels, after that the gradient increases so regularly that individual lines cannot be made out on the graphs. Looking at these there seems to be no reason why these 58 least sensitive pixels should be differently to the rest of the array. The distribution of these pixels with lowest output is shown by the binary image in Figure 4.13. It shows a scattering of pixels across the centre of the image and a slight grouping towards the top right corner and down the right hand side.



with lowest gain in top left graph and progressing left to right, then in rows. Within each graph, colours cycle from black (lowest) through red, green and blue (highest).

Figure 4.12, Response curves for the 64 Pixels with lowest gain (image difference)



Figure 4.13, Binary image of responsivity, lowest performing pixels in black

4.3.2 Inspecting the behaviour of pixels with the highest raw voltage

Inspecting the pixels with highest sensitivity was a useful test for saturation due the limitations of the detector or the electronics within the camera, but it will not necessarily pick up cases of saturation in the digitiser electronics, these could be in the pre-amplifier of the digitiser board, or actually involve overloading of the ADC itself. This form of saturation would be caused by those pixels producing the highest voltage, the highest raw ADC values before image differencing. These are not be the same pixels that produce the highest image difference signal, which were examined above.

The entire data set was searched for the highest raw pixel values. The highest value found was 15663 (mean value over 100 fields). This reading is well below the full scale value of the 14 bit ADC, which is 16383. The mean reading for the negative fields of the same pixel, at the same time, was 14359, giving a signal level of around 1300 ADC units. These were recorded with a ΔT of around 80K and camera temperature of 20 °C. These figures yield an average ADC resolution of 61mK in the 20 – 100 °C temperature range. As will be demonstrated later, this resolution is significantly better than the tNETD of the camera, showing digitiser to be adequate and properly adjusted for taking measurements in the 0 – 100°C range.

The fact that the ADC was not saturated does not guarantee that the signal could not be clipped by the input amplifier, this was investigated by looking at output graphs for those pixels with the highest raw readings. Figure 4.14 shows a series of response plots for the 64 highest raw pixel readings, again there are four pixels plotted on each graph. These pixels are arranged in descending order of raw value left to right across the page. The graph contains curves of varying gradient but they are towards the upper end of the range of gradients found earlier. All graphs show the same upwardly curving response found in all the other working pixels, there is no sign of a downward curve or kink at the top of any of the graphs which would indicate saturation at some stage in the sampling process.

Looking at the graphs in Figure 4.14 it is apparent that the curves do not all come from the set of readings with the highest sensitivity, the pedestal level of the raw signal dominates over the signal generated by the thermal input, in fact, many of these 64 highest readings come from the same pixel on different occasions. This is illustrated in Figure 4.15 which shows all the pixels that appeared in these 64 highest readings, came from fifteen individual pixels clustered around the lower right corner of the image. All fifteen pixels are on even numbered rows (counting from zero at the bottom). Figure 4.16 shows the raw output of the entire detector. Comparing this with Figure 4.15 shows a correlation between the two, the brightest region of Figure 4.16 is in the lower right corner, and striations can be seen, with the even numbered rows brighter than the odd numbered rows, matching the distribution in 4.15.



Response curves for the 64 pixels with highest raw voltage readings. Four pixels on each graph, starting with highest raw voltage in top left graph and progressing left to right, then in rows. Within each graph, colours cycle from black (highest) through red, green and blue (lowest).

Figure 4.14, Response curves for the 64 pixels with highest raw voltage readings

Comparing this with the responsivity image in Figure 4.9 reveals certain similarities. The dead pixels appear in both images. The horizontal striations also appear in both images but the polarity is reversed, those rows on the detector that have the lower raw values have higher responsivity. Also the vertical banding appears in both images and again the polarity is reversed between the two. The conclusions that can be drawn from these observations are limited. It is reasonable to suggest that the features that are shared by both the raw image and the processed image could be caused by non-uniformities in the readout amplifiers and electronics, rather than any non-uniformity in the detector or the bump bonding, since voltage offset is more likely to be caused by the associated electronics than the structure of the detector material. Without deeper investigation it is not possible to say for sure.



Figure 4.15, Pixels contributing to the 100 highest raw readings



Figure 4.16, Image of raw detector output (white high)

<u>4.4 Analysis of detector performance for the purposes</u> of radiometry

4.4.1 Response to IR

Section 4.1 formulated a model for the detector behaviour in which response is linear with respect to radiation power absorbed by the detector; the camera's response to the black body was used to validate this model. Since it was not possible to measure the spectral response of the detector directly, the approach taken was to model the relationship between black body temperature and the power incident at the detector for the nominal passband quoted for the camera; $8 - 12 \mu m$. This is model is compared to that which would be expected through 1m of atmosphere if the camera had no high-pass wavelength filter, and also that for the full black body spectrum. Figure 4.17, is a graph of the results. What follows is a detailed explanation of how these plots were generated.

Using the method described in section 4.2, data was gathered for the response of the camera to the black body. The camera was kept at a stable temperature of 23 °C, the response to the black body was measured at 12 points between room temperature, (~24 °C at the time), and 100°C. The IDP was used on the data and the mean IDP response of each pixel over 100 frames calculated. From these values a mean response for the whole array was calculated at each black body temperature. The resulting curve is shown in figure 4.18, (the significance of the other curve on this graph is discussed later). A 2nd order polynomial curve was fitted to the measured response of the camera, and this was used to extrapolate a response for the camera at black body temperatures of exactly 40 °C and 90 °C. The purpose of the polynomial fit was to reduce the effect of any slight experimental errors at any individual measurement points. (Such errors could be caused by fluctuations in room temperature and suchlike.)

Next, the radiation intensity emitted by the black body within each band was calculated, as described in section 4.1.4, for black body temperatures of 40 ° C and 90° C. These values were used to calculate gain and offset values for converting between the experimentally measured black body temperature and theoretically calculated radiation intensity for each passband. (This process assumes the detector response to be linear.) These gain and offset values were used to calculate were used to calculate radiation intensity for the black body temperature at each experimentally measured point; this was done for each of the three passbands.



Figure 4.17, Linearity of detector response to black body when black body temperature is converted radiation intensity using different models for the detector's spectral response. Linearity is quantified by calculating correlation coefficients.

This procedure is a method of converting the measured black body temperatures to a scale of radiation power; this was done for the three passband models, so that goodness of fit could be measured for each model using the same experimental data, with the hope that the results would support the quoted $8 - 12 \mu m$ passband figure. (The computed data is kept in the "in-band emitted intensity" scale, rather than converting to power received at the detector, this is because converting to received power would involve making assumptions about the efficiency of the lens and fill factor of the detector, the two scales are directly proportional so the linearity measurement is not affected.)

Linear fits and correlation coefficients (R) were calculated for each of the three models these are given in figure 4.17. Since the data has very little noise the correlation coefficient is high for all three models but it is measurably higher for the $8 - 12 \mu m$ model; R = 0.999878 vs. R = 0.999497 for the simulated transmission through 1 m of atmosphere. The value of R for the full spectrum lies between the two at 0.999873, however since this is known to be false, based on the atmospheric transmission data, it should not affect the conclusion; it is included for comparison only.

Whist the graphs in figure 4.17 are computationally valid they do not provide a particularly good visual demonstration of the fit since the errors are so small. Figure 4.18 gives a curve of measured detector response vs. black body temperature with a theoretical curve for the $8 - 12 \mu m$ band scaled to fit, this demonstrates the quality of the fit in a more visual manner.



Fig 4.18, Pixel response to black body temperature with theoretical 8 - 12 µm power curve overlaid

The results here support the model of an $8 - 12 \,\mu\text{m}$ passband and the detector having a linear response to radiation power. Since neither the spectral response profile, or the linearity of the detector could be measured directly, neither can be proved individually but the combined evidence is strong and, without doubt, the model provides an effective way of linearising the response of the camera, as evidenced by the R value.

4.4.2 Behaviour at different camera temperatures

Theory discussed in section 4.1 states that as camera temperature rises, the target temperature which will give an IDP output value of zero will rise with it, but differential response to radiation input will remain constant, so long as the sensitivity of the pyroelectric itself does not change. To test this theory response curves were measured for different camera temperatures, starting from 16 °C and going up to 40 °C. These are plotted in figure 4.19a, they show the vertical offset predicted by theory. Next, the curves for camera temperatures higher than 16 °C were offset by an amount calculated from the new camera temperature and the response at 16 °C. This was done by finding the point on the 16 °C response curve were the black body temperature equals the new camera temperature, the IDP response value at that point was used to offset the readings at the new camera temperature by subtracting it from each value. This computation shifts all the other curves up towards the 16 °C curve effectively removing the reduced IDP signal caused by the additional radiation coming from the chopper when operating at the higher temperature. This represents exactly the kind of correction that could be used in a function radiometric calibration routine.

The shifted curves are shown figure 4.19b, the curves now lie very close to each other. The gradients are all very similar, showing that there is no significant change in the pyroelectric sensitivity of linearity over this temperature range. However, small offsets remain between each curve, this cannot be explained by a change in the response of the pyroelectric material because the IDP removes any DC response from the detector. The implication is that the input to the detector is somehow not matching the model. This is likely to be caused by the way camera temperature was being controlled through the front plate of the detector chamber only. Slight imbalances between the chopper temperature and the temperature of the inside front surface of the detector chamber could produce this kind of offset, but this cannot be confirmed as the cause without further work.



Figure 4.19, Detector response at different camera temperatures

The size of the offsets is enough to cause a 4 °C difference in radiometric temperature measurement between the highest and lowest camera temperatures, (which span a range of 24 °C). The repeatability of measurements over time proved to be better than that, implying that the cause is systematic rather than random error. This was easily corrected for by factoring the offset change into the calibration routine; that solves the problem but does not explain its presence. This effect was not investigated any further because it was felt that a certain answer would not be obtained without considerable work, if at all.

4.4.3 Radiometric calibration

The experiments carried out showed that the detector functions as predicted by the model, the camera can be easily calibrated for radiometry using the data that was recorded. The limiting factors will be repeatability, and especially repeatability in less tightly controlled environmental conditions, this was investigated to some extent, and is documented at the end of this chapter. The controlled IR passband of the camera suggests the radiometric temperature measurements should not be strongly affected by the distance to the target over short to medium ranges, not more so than other LWIR thermal imagers. Inevitably, longer ranges will affect radiometric performance; the reasons for this are explained in chapter 1. All that is needed for calibration is some mathematical routine to calculate target temperature from IDP output and measured camera temperature, using the data already recorded. This section will cover such methods and explain the one chosen to make the radiometric measurements used later in this chapter and the next.

These statements about radiometric performance do not take into account the effect of high contrast levels in the scene, in the idealised theoretical model no such effects should exist, in most real arrays some amount of thermal cross-talk between adjacent pixels is inevitable, but should not extend for a radius of more than a few pixels. Unfortunately, in the case of this camera high contrast levels were found to have extreme effects on radiometric performance, the evidence does not suggest this is down to any physical behaviour of the detector, but is more likely due to some faulty or inadequate element in the analogue electronics, whether in the readout, or the power supply to the detector. This effect and a method to correct for it are discussed in detail in chapter 5. The correction acts independently of the radiometric computation method and so is left until the next chapter.

As stated, what is needed for radiometric calculation is some method of describing the relationship between target temperature and IDP output at the measured points, and a method of interpolating between the measured points. There are a large number of possible techniques; from one giant two dimensional look up table, to a fully parameterised system. A method may, involve converting output to temperature via radiation power, doing so allows the calibration to be broken down into detector behaviour, and behaviour of the other elements in the system, such as the lens. This would be of benefit for a camera used with multiple lenses. Without being able to

measure the spectral response of the camera it was not possible to develop such a segmented radiometric model.

The radiometric model used for this project was simply to fit 4^{th} order polynomial curves, to the response of each pixel. The 4^{th} order curves match the radiation power in the 8 -12 μ m band well. To calculate target temperature, the polynomial equations are solved for the current pixel value, the coarse offset is corrected using the reading from the PRT sensor in the camera and the technique described in section 4.4.2. The residual offset shown in figure 4.19b, was handled by repeating the process for the two calibrated camera temperatures closest to the current camera temperature, and interpolating between the two results. This method is computationally intensive and, as such, unlikely to be used in a production system, but it is conceptually simple and direct enough to avoid introducing any artefacts which might have affected the results in later sections, it was shown to have good repeatability and stability across all the pixels in the array.

4.4.4 Radiometric performance measurement

4.4.4.1 tNETD

The concept of NETD and the various ways of measuring it are covered in detail in chapter 1. tNETD is defined as the standard deviation on the output of one pixel over time, expressed on the temperature scale. NETD is expected to reduce as target temperature increases, as a result of the increasing gradient of the power vs. black body temperature curve given in figure 4.6.

tNETD was calculated for each camera temperature that was tested, results are plotted in figure 4.20. tNETD reduces with increasing target temperature, and appears to stay fairly stable across different camera temperatures. To confirm the relationship between tNETD and the gradient of the black body power curve, the gradient of the 4th order polynomial fit was calculated and plotted on the tNETD graph. This is shown for a camera temperature of 16 °C in figure 4.21, the fit is remarkably good.



Figure 4.20, Graphs of mean tNETD vs. Target temperature, for all camera temperatures, error bars indicate standard deviation across all pixels



Figure 4.21, Graphs of mean tNETD vs. Target temperature, for all camera temperatures

tNETD is routinely quoted for a target temperature of 30 °C, by this measure the camera would have an average tNETD across the array of almost exactly 0.8 K. The range of tNETD values for different pixels runs from 0.94 K to 0.67 K. At a target temperature of 100 °C the mean value drops to around 0.52 K. Values of tNETD across the array were made into an image to check for any peculiarities on distribution across the array, this is given in figure 4.22. No distinct features are visible apart from slightly higher values towards the bottom edge of the array, which could be related to the effect discussed in chapter 5.



Figure 4.22, Image of tNETD

4.3.2 sNETD, image uniformity

At each calibration point the radiometric output over 50 consecutive frames was used to obtain a mean frame with values on the temperature scale, the standard deviation across the frame was calculated to give the sNETD.

sNETD is plotted vs. target temperature in Figure 4.23, sNETD decreases slightly with increasing target temperature, but not as distinctly as tNETD. The values cover a range from 10.1mK up to 313mK. sNETD was measured on all the steady state camera readings taken on different dates, before and after, the calibration date. The mean sNETD of all readings other than those using for calibration was 265mK, showing that image uniformity did not change significantly over time. These values of sNETD are all well below the tNETD of the detector, even at high target temperatures, as a result sNETD is unlikely to have a noticeable effect on the image quality of the camera.



Figure 4.23, Graph of mean sNETD vs. ΔT , all camera temperatures, with 2nd order polynomial line of best fit

4.3.3 Long term accuracy of radiometric calibration

The long term stability of the calibration was tested by capturing data under similar conditions on 3 dates over a period of 21 months. On each occasion three consecutive measurements were made with the camera and target operating in stable conditions. The measurements were made 30 minutes apart. The mean value of 400 pixels in the centre of the detector, integrated over 25 frames was used to calculate the temperature of the target; the errors on these radiometric measurements are plotted in figure 4.24.

The error bar plotted at each date spans the range of the 3 measurements taken on each date. These results show that short term stability is greater than long term stability; the widest spread of temperatures on one day was recorded in January 2002, these covered a range of 0.14K. Whereas the range covered by all readings was 1.7K. It appears from the graph that the calibration had drifted most significantly between November 2002 and October 2003, but since only 3 dates were tested, it is not possible to draw conclusions from that. The results suggest that, with regular recalibration and a stable operating environment, it should be possible to make radiometric measurements to an accuracy of better than 1K, but without recalibration larger errors may occur.



Figure 4.24, Calibration stability over twenty months, error bars cover the range of errors recorded on each

4.3.4 Effect of changing camera temperature on radiometric performance

An attempt was made to assess the camera's susceptibility to changing environment. The black body target was set to a constant temperature and the camera temperature was changed using the water bath. In the first experiment the black body was set to 25 °C and the camera temperature was to 20 °C, the temperature control of the water bath was then increased by 5 °C. The system was left for 5 minutes and the next reading taken, at this point the camera had not yet reached its new target temperature. The thermostat was then increased another 5 °C and the cycle repeated. The error on each reading (the gap between radiometrically measured temperature and PRT measured temperature), is plotted in Figure 4.25. Later the process was repeated leaving 10 minutes between consecutive readings; these results are plotted in Figure 4.26. On each graph error bars show the standard deviation in the response of the 200 pixels used to calculate the radiometric values.



Figure 4.25, Radiometric error when varying camera temperature rapidly





With a simple interpretation of the physical system one would expect the measured error values to be consistent within each test because the rate of change of temperature ought to be similar at each measurement point. One might also expect the errors to be greater when the temperature is changed more rapidly, but the results don't satisfy either of these predictions. Both tests show positive and negative errors, and the size of the errors is similar in both cases. The worst errors are just over 1 K. The only firm

conclusion that can be drawn is that in order to keep radiometric error below 1 K the camera should be left longer than 10 minutes to reach equilibrium with its environment before readings are taken. To learn more about the behaviour of the camera in dynamic environments a more advanced temperature control system would be needed, along with additional PRT's mounted within the camera to give some idea of the temperature gradients which are set up in the process.

Chapter 5, Scene contents correction 5.1 Experimental setup

Experimentation has shown the camera can be calibrated for radiometric measurement with a reasonable degree of accuracy, when the image is filled by a flat field. The next stage of performance assessment was to investigate the performance of the camera when aimed at a scene containing areas of differing intensity.

These experiments were carried out using the same black body used for flat field calibration, in addition to the black body, an aluminium plate was used in the scene to occlude areas of the black body. The aluminium plate was painted on one side with the same highly emissive textured black paint as the black body, (Nextel 101-C10). The other side of the plate was left unpainted. The plate was arranged with the painted surface facing the camera and the unpainted surface facing the black body. This arrangement minimised the effect of radiation from the black body on the temperature of the plate, as well as minimising the reflection of IR radiation from the plate towards the camera.

A retort stand was used to hold the aluminium plate in a number of positions, chosen to allow systematic investigation into the relationships governing effects of the scene content on the cameras response. The patterns chosen included vertical divisions down height of the image, with the aluminium plate occluding different proportions of the image from both the left and right sides. Similar arrangements were used with horizontal edges. Measurements were made with one corner of the image occluded by the plate, and a number of observations were taken using a diagonal division between the hot and cold areas of the image. Additional readings were taken with a division made by the plate and an additional area of the black body occluded by objects of other shapes. A metal plate with a circular aperture was used to produce images with a cool background and a hot circle of varying sizes in the middle of the image.

Data was captured from the camera in the same manner as earlier experiments. Once the raw data was saved to disk, image differencing and further processing was carried out in IDL. The calibration system described in Chapter 4 was applied and the radiometric output was inspected using imaging and graphing techniques. In this chapter most target temperature measurements are expressed as a ΔT value, which is the difference between the target temperature and camera temperature, since the distortions seen here are related to absolute voltage value rather than black body temperature, ΔT is the more useful parameter to graph, because it is the cause of changing voltage on the output.

5.2 Analysis of patterned images

5.2.1 Vertical divisions

Figure 5.1 shows a sequence of images with vertical temperature divisions. These images are given in false colour, to allow finer visual inspection of the intensity variations on the printed page. The colour scale is shown with the image.



Figure 5.1, Images of vertical divisions

The most obvious flaw in these images is a change in intensity towards the bottom edge, both the hot portions of the image become visibly cooler in the bottom 20% of the images and the cool areas become warmer in same region. In the images where the division between hot and cold is closer to the edge of the image, the effect is stronger in the smaller area than in the larger area, see images A and D.

Figure 5.2 shows graphs plotted of vertical sections through these images. Each graph shows the measured temperature in two vertical sections through the image, one in the hot area and one in the cold area. The unprocessed image difference values are over plotted as unlinked crosses, these data are scaled by a factor of 0.1 to fit them on the same graph. Each graph is numbered to match the corresponding images in Figure 5.1. (In this chapter all graphs of vertical sections are plotted in order from the bottom to the top of the image, and horizontal sections are plotted left to right; this matches the order in which the pixels are sampled from the detector.)



Figure 5.2, Graphs of vertical sections though vertical division images in Figure 5.1

These graphs show numerically the variation in temperature towards the bottom of the image, which was noted in Figure 5.1. In the most extreme examples the cool sections of the image become over 20 °C warmer at the bottom edge of the image than in the top half of the image and the hot sections become over 15 °C cooler. Horizontal sections of the images show that each row in the image exhibits apparently even response, despite the large variations seen in the vertical direction. Some example plots are given in Figure 5.3 to illustrate this.



Figure 5.3, Graphs of horizontal sections though vertical division images in Figure 5.1

5.2.2 Horizontal divisions

Figure 5.4 shows images with horizontal temperature divisions. The images are presented in the same false colour scale as those in Figure 5.1. These images show distortions in temperature towards the bottom of the image and on the upper side of the temperature boundary. In those images which are hotter below the boundary, the bottom edge of the image appears cooler than the rest of the hot region. In the cold region at the top of the image, the area just above the boundary appears cooler than the top of the image. In those images where the area below the boundary is cold, the trend is reversed, the bottom edge of the image appears warmer, and the area just above the boundary also appears warmer than it should.

Figure 5.5 shows graphs of vertical sections taken through the images in Figure 5.4. These graphs show distortions of up to 9 K in the cooler regions and up to 6 K in the hot regions. The shape of the temperature distortions appears similar to those noted in section 5.2.1. The shape of the curves at the start of the frame and following the transitions appears to follow a logarithmic form, decaying from an initial level towards a

stable level, with a stable time constant. Since the effect is believed to be related to the behaviour of analogue circuitry within the camera or detector, a logarithmic model is reasonable if the behaviour is linear, the validity is demonstrated by the results later in the chapter.

As in Figure 5.2, these graphs are over plotted with the unprocessed image difference values. One feature stands out as a difference between the processed and unprocessed data. In the scenes which are hot below the boundary and cold above, the processed data shows temperature decreasing towards the bottom edge of the image but the unprocessed data shows output increasing towards the bottom of the image. The relationship correlates with the increased response of the detector towards the bottom of the image noted in Chapter 4. This feature is consistent with all relationships examined so far and is specifically noted since it becomes important later.

Figure 5.6 shows a horizontal section through one of the images of horizontal divisions. As in Figure 5.3 the horizontal section offers little insight into the nature of the distortions seen in the vertical sections.







Figure 5.5, Graphs of vertical sections though horizontal division images in Figure 5.4



Figure 5.6, Graphs of horizontal sections though horizontal division images in Figure 5.4

5.2.3 Corner occlusions

Figure 5.7 gives images with a hot background and one corner occluded by the cool aluminium plate. These images show features in common with both the vertical boundaries examined in section 5.2.1 and the horizontal boundaries examined in section 5.2.2. Around the bottom of the image, the cool area, occupied by the aluminium plate, appears warmer than expected and the hot area, the black body, appears cooler than expected. At the horizontal boundaries, the area immediately above the boundary appears warmer than expected when hot and cooler than expected when cool. These effects all seem stronger in the images where the plate covers more of the total area of the image, Figure 5.7 b and d. (The circular dot seen in the corner of the plate is a mounting hole drilled in the plate, its presence is incidental.)



Figure 5.7, Images of corner occlusions

Figure 5.8 shows graphs of vertical sections through images shown in Figure 5.7. Each graph shows two vertical sections, one through column 4 and another through column 95. In both graphs column 4 falls across the horizontal boundary, and column 95 lies entirely in the hot portion of the image.

As stated the sections through the horizontal boundary exhibit a similar form to those in section 5.2.2, additional insight is gained by comparing the profiles of the sections through column 95. These curves show cooler than expected readings towards the bottom edge of the image, they also show a change in temperature in line with the horizontal boundary in column 4. The change in temperature is not a sudden discontinuity, rather once the boundary is hit, the temperature starts to drift towards a new equilibrium, again with an apparently logarithmic form.

In graph 5.8a the corner occlusion lies in the lower left corner of the image consequently column 95 starts at a higher temperature than column 4. After the boundary is hit at around row 75, the curve for column 4 rises quickly, to a level close to that of column 95, then over the subsequent rows, the level of both curves drops towards a new equilibrium. In Graph b, the two curves start at a similar level, at the boundary the level in column 4 drops rapidly, then both columns rise in level, following the same logarithmic profile.



Figure 5.8b, Graphs of vertical sections though corner occlusion image 5.7 D Figure 5.8, Sections though corner occlusion image 5.7 D

Looking at these curves, it seems as though the horizontal boundary causes an offset in temperature values in subsequent rows. The change is an offset, rather than change in gain, since small and large signals are shifted by the same amount. The graph in Figure 5.8 b shows this most clearly. After row 25 both columns are at very different levels but both drift upwards by the same amount.

5.2.4 Circular apertures

A number of readings were taken with a cold background and a circular aperture through which the hot surface of the back body could be seen. For these readings a metal plate with a multi-bladed adjustable aperture was used, the shape of the aperture very closely approximated a circle at each size. The plate in which the aperture was mounted was not large enough to fill the entire detector whilst keeping the edge of the aperture in focus. To avoid the hot surface of the black body appearing in the edges of the image the black body was moved further from the camera so that the plate would exclude the entire surface. Around the edge of the plate, the image is filled by the wall of the lab, which appears in the image at a similar but not identical intensity to the plate. The outline of the plate can be seen, as can a small section of the integral bar used to mount the plate, this appears as a small dark rectangle to the left hand edge of the images. Images of 3 different sizes of aperture are given in Figure 5.9.



Figure 5.9, Images of circular apertures

The largest aperture appears in Figure 5.9 a. Tracking up this image from the bottom the temperature of the cold plate appears to drop once the hot circle enters the image. As the circle starts to occupy more of the length of each row, the apparent temperature of the background continues to drop. Moving above the aperture, the temperature of the plate rises back towards its initial level. The medium aperture in image b shows the same effect to a smaller extent. In the image with the smallest aperture, image c, the effect cannot be seen. (The slight darkening seen in a circular area around the aperture, is caused by the blades of the adjustable aperture, which have a slightly different emissivity than the plate.)

Graphs of vertical sections through these images are given in Figure 5.10. Each graph has one line representing a section through column 19 and one through column 46. Column 19 is outside the aperture whilst column 46 passes through the aperture in each case. The graphs show levels consistent with the theory of an offset postulated in section 5.2.3. In graph a, the temperature in column 46 rises sharply at the edge of the aperture, progressing up the height of the aperture, the temperature dips slightly from the peak at the bottom edge. Passing the top edge of the aperture, the graph drops and seems to overshoot the level it was at prior to the aperture. Towards the top edge of the image, the graph rises up towards its original level. Column 19 does not contain the peak corresponding to the aperture but the level drops for the rows in which the aperture is present in other columns. The size of the drop matches the size of the drop across the top of the plateau seen in column 46. Above the aperture, the level of column 19 closely matches that of column 46, drifting upwards at the same rate. The same effect is seen to a smaller extent in Graph b but is not detectable in Graph C.

The presence of the hot aperture in the image appears to offset the temperature readings downwards, in the whole of the rows in which it appears. The offset builds up gradually and persists above the aperture, gradually decaying towards the top of the image. The size of the effect seems to be proportional to the area of the aperture in the image.



Figure 5.10, Sections though circular apertures in figure 5.9

5.2.5 Multiple edges

Examining the preceding scene arrangements offered significant insight into the effects of scene contents on the response of the camera, and enabled the formulation of a theory as to what kind of relationship might describe the effects. To investigate the effect further and to allow some quantification, a new scene was devised. A vertical division was used, to since this is known to disrupt the response through the entire height of the frame. An additional rectangular notch was taken out of the hot area of the image, using a metal rod. The aim was to measure the effect of this rod on the measured temperature of both cool and hot columns in the image. The images generated are shown in Figure 5.11





Similar distortions are seen in Figure 5.11 A and B. Image A will be discussed here for convenience, the same observations apply to image B only mirrored around the vertical axis. At the bottom of the image, the hot region is cooler and the cold region warmer, in the same way as noted in section 5.2.1. Moving up the image, as the rod enters the image, at around row 40, the hot section through the centre of the width starts to rise in temperature. The solid cold band at the right hand edge of the image also starts to become hotter. The metal bar on the left of the image appears coldest at its bottom edge, the apparent temperature rises moving upwards through the vertical section of the bar. These relationships are shown numerically in the graph of Figure 5.12.



Figure 5.12a, Graphs of vertical sections though multiple edge image 5.11 A



Figure 5.12b, Graphs of vertical sections though multiple edge image 5.11 B Figure 5.12, Sections though multiple edge images in figure 5.11

Figure 5.12 shows the changes in temperature reported between rows 45 and 56 are similar in magnitude and profile for all three columns. These curves confirm the earlier observations, that a reduction in the mean temperature in each row, causes an increase in the measured temperature of all pixels in subsequent rows. This change is an offset of similar magnitude, whatever the actual temperature in that area of the scene. The change does not start with a discontinuity, rather it takes the from of a logarithmic progression towards a new stable level. When the mean temperature in a row drops, the
offset reverses, the measured temperature in subsequent rows drops, understanding the nature of this shift in response makes it possible to correct for it.

5.2.6 Analysing the nature of the distortion

The effects noted above, seem to take effect over a number of rows, there is little variation in each row. However, given the sequential nature of the read out of the detector, it is possible that the effect occurs on a pixel by pixel basis rather than line by line, the effect being too small to be readily apparent in one line of the image. Figure 5.13 shows graphs of measured temperature vs. pixel number, over a period of 15 rows of the image in Figure 5.4 c, starting from row 37. The smooth nature of this curve confirms that the effect does grow one pixel at a time.



Figure 5.13, Graph of 1500 consecutive pixels, starting above the horizontal division in image 5.4C

It was postulated that this offset effect could be caused by the increased voltage, and hence charge, needed to report higher temperatures in the analogue circuitry of the camera or detector. As the output voltage rises, some capacitive element of the circuitry may lose charge and voltage, causing the output voltage to drop. Conversely as the output voltage drops, the capacitive element would recharge towards its previous value. This theory explains the behaviour noted in the central regions of the images and also the logarithmic nature of the changes but it does not initially explain the behaviour at the bottom edge of the image, where a vertical division causes hot areas to appear cooler and cold areas to appear warmer, considering the nature of the raw data prior to image differencing offers an explanation. Fields are considered positive and negative for the purpose of image difference processing, since negative fields are subtracted from positive fields to calculate the total signal. However, both positive and negative fields appear as positive voltage on the output of the detector. Looking at a warm scene (scene hotter than chopper blade), the average voltage of a negative field is lower than that of a positive field. This implies that, in accordance with the proposed theory, capacitive elements of the readout circuit would be more charged at the end of a negative field than at the end of a positive field. The transition from negative to positive occurs in the time space of two lines, empty of data, at the end of the field. Following this transition, one would expect the voltage measured at the start of a positive field, to have a boost in voltage relative to the end of the field. Conversely, the start of a negative field would be at a low voltage, relative to the end. After image differencing, these changes would lead to much greater responsivity at the start of a frame than at the end. This effect would be larger the hotter the scene and would diminish as the temperature of the scene approaches that of the chopper. When the temperature of the scene was lower than the temperature of the camera, the effect would be reversed.

Going back to the flat field data analysed in Chapter 4, the predictions of the theory above are fulfilled. Figure 5.14 shows plots of the image difference signal level looking at flat fields of 3 different temperatures. This demonstrates that the magnitude of the effect increases with increasing signal strength. Figure 5.15 shows plots of the raw values of pixels when staring at a hot black body, minus each pixel's mean value when staring at a black body close to chopper temperature. (This subtraction must be made to show the trends clearly without pedestal noise.) These graphs are given for both positive and negative fields. They demonstrate that the effect is present with opposing polarity in both positive and negative fields. The model also implies that features in the scene near the top of the image will affect response at the bottom of the image; this is sometimes termed a "wrap around" effect, and is observed in later examples.



Figure 5.14, Image difference level vs. Row number (vertical section) at three black body temperatures





5.3 Constructing a procedure to correct for scene contents

5.3.1 Defining the goals of the algorithm

The challenge was to create a algorithm, through which, the value of each pixels in the frame could be corrected, to compensate for the effects of changing output voltage and hence remove the adverse effects of scene contents on the radiometric performance. Two things were needed to do this, firstly a method of determining the real value from the measured value in at least one pixel of the frame, secondly a method of tracing the changing value of the offset between measured and real values throughout the rest of the frame. From these calculated offsets the real value of every pixel could be calculated.

Before developing an algorithm, the meaning of "real value" must be defined. After some initial investigation, a decision was taken to define the real value of a pixel as the voltage level (in ADC units) that would be measured for that pixel, if all preceding pixels had produced zero output for a long period of time. This value was termed Vi for "Ideal Voltage". The measured voltage was termed Vm and the offset, Vi – Vm, was termed ΔV , the relationship is given in equation 5.1. This definition of Vi was chosen to simplify calculating the development of ΔV throughout the frame, the reasons will become apparent later in the chapter.

 $\Delta V = Vi - Vm \qquad (5.1)$

Once an algorithm was developed to determine Vi from the measured voltage Vm, it would be possible to apply the algorithm to flat field data and recalibrate for radiometric measurement. It would then be possible to apply the algorithm to data from scenes containing any image and use the new calibration data to convert the processed voltages into temperature measurements.

5.3.1.1 Algorithm flow chart

The algorithm described in this section (5.3) has considerable complexity and was conceived through reasoned investigation into the finer details of the distortion effects examined in section 5.20. The stages in the creation of the algorithm and the mathematical reasoning behind each step are given in detail throughout the rest of this section. It was felt that this derivation might be hard to follow without a more simple understanding of the operation of the algorithm as a whole. To aid understanding an

overview of the algorithm and its derivation is given over the following pages in the form of a flow chart. Necessarily the flow chart lacks many of the finer details which are covered in the following subsections. A summary of all the mathematical stages of the algorithm is presented graphically in section 5.3.6 (Figure 5.21).





a.

Distortion is noted in frames with high contrast, for example, in this frame featuring a horizontal division between hot and cold areas, a brighter than expected band is seen at the bottom of the frame and another just above the division, at the bottom of the hot region (around row 40).

Looking at raw data from flat fields reveals similar peaks in output at the start of the frame. This graph shows image difference output in one column of the detector at three different target temperatures. In each case the frame starts with a high voltage which drops noticeably over the first 20 rows. The size of the distortion appears to be related to the magnitude of the input (size of Δ T). (When this data was first examined in Chapter 4 this effect was believed to be natural variation in gain across the detector.)

The characteristics of the distortion were examined by comparing voltage levels in the first 30 lines at the start of the frame to the mean voltage in the rest of the frame. Logarithmic curves were fitted to each case, these are shown in the graph. In all cases the decay constant of the curves was found to be similar, the mean value was taken and assumed to be a constant for the purposes of the algorithm. As noted above the size of the effect appeared to be related to the input level.

d.

The size of the voltage peak was measured by taking the y-intercept of the calculated logarithmic curves shown above. These values were plotted against the mean image difference voltage level (proportional to ΔT) and a straight line was obtained, which is shown in this graph. When viewing an arbitrary scene the value of ΔT is unknown, so this graph could not be used to calculate the size of distortion in the scene. Another model was suggested.



e. A weighted average of the previous frame is used to approximate the size of the distortion at the start of the next frame. Each pixel is weighted by the same exponential form as the fitted curves in part c above, only reversed. So pixels towards the end of the previous frame have a large contribution to the calculation and those at the start have a very small contribution. This weighted average was plotted against the size of the distortion as measured in part d above, again a straight line resulted. This relationship was used in the algorithm to calculate the size of voltage swing at the start of an arbitrary frame.

f.

The voltage swing identified in part e was used to determine the amount of correction to be applied to the first pixel. The level of correction for subsequent pixels was calculated by taking the value of the current pixel and using the same linear relationship to determine what value the signal would reach if all subsequent pixels received the same input. From one pixel to the next, the level of correction moves from its previous value towards the calculated target value by a fixed factor; a factor determined by the logarithmic curves in part c. This process applies to all pixels in the frame. The algorithm was applied to all flat field data, one case is plotted in this graph. In the corrected (green) line the voltage peak at the start of the frame is removed. Radiometric calibration was repeated with the corrected data.







This sequence shows how the algorithm acts on image difference levels to remove the distortion at the start of the frame and the peak that follows a sharp transition. (Fixed pattern noise is removed from these graphs for clarity.) However, notice from graph i, the pixel correction level, ΔV is around zero ADC units at the point just after the transition (around row 37), and yet in graph h, which uses a temperature scale, the uncorrected level (black line) is higher than the true PRT measured temperature (dashed blue line). This may seem anomalous but is explained by the fact that the radiometric scale was recalibrated after applying the algorithm to the flat field results. (After applying the algorithm exactly the same calibration process described in Chapter 4 was used.)



k.

The images to the left show how the correction algorithm removes the undesirable distortion from the image. Throughout this example the same scene has been used, a horizontal division, the algorithm has also been shown to work on many different scenes, examples of which are given later in the chapter.

5.3.2 Describing the evolution of voltage offset ΔV mathematically

It was believed the decreasing voltage observed at the start of flat field frames was caused by a changing value of ΔV , and that this change was logarithmic in nature (discussed in section 5.2.6). To determine whether a logarithmic algorithm would be suitable to describe this effect, the COMFIT (/LOGISTIC) procedure in IDL was used to fit log curves to the data.

Log curves were calculated for all flat field data used previously for radiometric calibration in Chapter 4. To allow simple comparison, each curve was computed as though decaying towards zero voltage. To achieve this, the mean output of the last 40 lines of the frame was calculated, this figure was subtracted from the mean value of each of the first 60 lines of the frame. (It was decided to fit curves to the mean values of each line rather than all pixels, to speed computation.) Best fit logarithmic curves were calculated for the resulting data points. Figure 5.16 shows a graph of these curves for ten different target temperatures. Each target temperature is plotted in a different colour, the calculated log curve is plotted as a solid line and the camera data, to which these were fitted, are plotted as crosses.

All the logarithmic curves in Figure 5.16 show excellent goodness of fit, the largest differences between the measured data and the calculated curves are in the order of 50 ADC units. That occurs on the curve with highest target temperature, whereas the Y-intercept of this curve reaches over 550 units. The points with the largest deviations fall in the same detector row for each curve. This is caused by the variation in gain between rows in the detector, which was noted in Chapter 4. These variations will be compensated for later, in the calibration of each pixel, and are no cause for concern at this point in the process. Most data points in the graph fall much closer to the calculated curve than 50 ADC units, the majority being closer than 10 ADC units. Considering the quality of fit obtained in these graphs, and the suspected cause of the effect, it was decided a logarithmic model would be suitable to describe the evolution of ΔV , as long as the initial value and the value towards which ΔV decays, is an important concept in the formulation of this algorithm, it was termed ΔVt .



Figure 5.16, Logarithmic fits to voltage decay at start of frame vs. Row number. The full range of black body temperatures are plotted for one camera temperature.

5.3.3 Investigating the relationship between input level and the magnitude of ΔV

As a first step towards calculating ΔV , the relationship between the magnitude of ΔV at the start of the frame and input level was investigated. At this point, ΔV could not be measured, since Vi had not been calculated. However, Vi is believed to be constant throughout a flat field, and we have defined ΔV as equal to Vi - Vm (equation 5.1). Therefore, the change in Vm between the start and the end of the frame, must be equal and opposite to the change in ΔV over the same period. (Subscripts are used with the newly created variables to indicate their values at different points in the frame. 0 representing the value at the first pixel, i the value at pixel i, and f the value at the last pixel in the frame.)

$\Delta V_0 = Vi_0 - Vm_0$ $\Delta V_f = Vi_0 - Vm_f$ $\Rightarrow \Delta V_0 - \Delta V_f = - (Vm_0 - Vm_f)$ $S = Vm_0 - Vm_f$	(5.2)
	(5.3)
	(5.4)
	(5.5)

The difference between Vm at the start of the frame and the end of the frame was termed voltage swing (S).

Values of voltage swing, S were calculated from the y-intercepts of the logarithmic curves, discussed in 5.3.2. These were plotted against the mean value of Vm over the

whole frame and appear to lie along a straight line, as shown in Figure 5.17, which includes results from all camera temperatures. The relationship between signal level and voltage swing is linear. As discussed earlier, this voltage swing is contributed to by equal and opposite effects in both the positive and negative fields of the frame. Furthermore, it is believed that a positive voltage swing seen in the positive field is caused as a direct result of the lower output voltage in the preceding negative field and vice versa. From this, it follows that if all preceding pixels had been of zero signal level, the voltage swing at the first pixel in the first frame with non zero signal, would be half of that measured in the flat field case.



Figure 5.17, Swing vs. Mean Vm. Swing is calculated from the logarithmic fits in figure 5.16 and Vm is measured for each combination of black body and camera temperature.

Vi was defined as the voltage that would be generated if all preceding pixels had given zero signal, therefore for pixel 0 in the flat field case, the measured voltage, minus half the voltage swing is equal to the ideal voltage:

$$Vi_0 = Vm_0 - S/2$$
 (5.6)

From 5.1:

$$\Delta V_0 = V m_0 - S/2 - V m_0$$
$$\Delta V_0 = -S/2 \qquad (5.7)$$

Still considering a flat field, it was shown that over the course of one frame the initial measured voltage, Vm_0 decays towards a final value of Vm_0 minus S, from equation 5.5.

From 5.3 and 5.6:

$$Vi_0 = Vm_0 - S/2$$

And:

$$Vm_f = Vm_0 - S$$

From 5.5:

 $\Delta V_{f} = Vi_{0} - Vm_{f} = Vm_{0} - S/2 - Vm_{0} + S$ $\Delta V_{f} = S/2 \qquad (5.8)$

Both ΔV_0 and ΔV_F have been related to S, which is directly proportional to input level, shown in Figure 5.17. For the flat field case it is believed that:

$$\Delta V_{f} \approx \Delta V t$$
$$\Delta V t = S/2 \tag{5.9}$$

5.3.4 Finding ΔV for intermediate pixels

If ΔV can be calculated from ΔV_0 and $\Delta V t$, then Vi can be calculated for every pixel in the frame. Since the decay of ΔV_0 towards $\Delta V t$ is logarithmic, we may define F as a factor by which ΔV decays towards $\Delta V t$ in the space of one pixel.

$$\Delta V_{(i+1)} = \Delta V_i - (\Delta V_i - \Delta V_i) * F$$
(5.10)

and

$$\mathbf{F} = \ln(2)/\mathbf{T} \tag{5.11}$$

Where T is the half life of the logistic decay curves calculated in section 5.3.2. A mean value of T was calculated and found to be 488 Pixels, which yields a value of F, 1.42×10^{-3} .

For flat field data, S can be calculated from the graph in Figure 5.17, both ΔV_0 and ΔVt can be calculated from S. Equation 5.10 can be used to calculate ΔV for every pixel in the frame and hence Vi. This procedure was carried out for flat field data. Figure 5.18 shows both Vm and Vi for column 50 of a flat field. The calculated values of Vi show a similar form to the measured voltages but the characteristic peak of voltage at the start of the frame has been removed as intended.



Figure 5.18, Vm and Vi vs. Row number for a single flat field frame

5.3.5 Relating this to the non-flat field case

Equation 5.9 made it possible to calculate the Vi value of each pixel in a flat field, the only information needed in addition to the measured data was S, which was calculated from the mean value of Vm. In the case of a non-flat field, S could not be calculated from the mean value of Vm. The voltage swing is caused by the level of output in the preceding field but is more heavily dependent on those pixels closer to the end of the previous field, less so on those near the beginning. This was not an issue for the flat field case because the input level was the same throughout the frame and the signal in the current frame was the same as the signal in the previous frame. However, if S were calculated for a non-flat field in the same way, signals near the start (bottom) of the frame might have an undue influence on the calculation of S. To provide a measure of the output history of the detector at the start of a frame, a new measurement, a weighted average of the previous frame, was used. This weighted average gives most significance to the last pixel in the previous frame; the level of significance of earlier pixels drops with the same half life calculated in section 5.3.4, towards zero significance at the start of the frame. The weighted average method was used to calculate the output history of frames in the flat field data, these values were plotted against S in the same way as before, once again a straight line resulted, Figure 5.19.

Using a weighted average allows us to calculate S for any non-flat scene. However, when equation 5.10 was applied to a flat field, a constant value of ΔVt was used, calculated from the value of S at the start of the frame. This is no longer valid since the basis of this correction algorithm is that ΔV is influenced by the contents of the scene,

so a method of calculating ΔVt for each pixel must be used. ΔVt is known to be proportional to the mean value of Vi in a flat field, so it was postulated ΔVt may calculated for each pixel, from the Vi of that pixel. The relationship was found by taking the mean value of Vi in flat fields (calculated in section 5.3.4), and plotting them against the calculated ΔVt for the flat fields as shown in Figure 5.20. The line of best fit was calculated for these data, see equation 5.12 below, this was then used to calculate ΔVt at each pixel in non-flat scenes:

 $\Delta Vt = 0.241Vi + 0.739 \tag{5.12}$

With equation 5.12 it is possible to apply equation 5.10 to pixels in a non-flat scene. During processing each frame of video is analysed to calculate the value of S for the next frame. S is used to calculate ΔV of the first pixel of the frame. ΔV is used to calculate Vi. Vi is used to calculate ΔVt . ΔV and ΔVt are used to calculate the new value of ΔV for the next pixel. This process is carried out in turn for each pixel of the frame. The calculated response curve is each pixel is used to calculate ΔT from Vi. Adding the measured camera temperature to ΔT produces radiometric output. These processing steps are expressed algebraically below:

For the first pixel of the frame:

$$\Delta V_0 = -S/2$$

$$Vi_0 = Vm_0 + \Delta V_0$$

$$\Delta Vt_0 = 0.241 Vi_0 + 0.739$$

$$\Delta Vi_1 = \Delta V_0 - (\Delta V_0 - \Delta Vt_0) \ge 1.42 \ge 10^{-3}$$

For all subsequent Pixels, i:

$$Vi_{i} = Vm_{i} + \Delta V_{i}$$
$$\Delta Vt_{i} = 0.241 Vi_{i} + 0.739$$
$$\Delta V_{(i+1)} = \Delta V_{i} - (\Delta V_{i} - \Delta Vt_{i}) \ge 1.4 \ge 10^{-3}$$

The numeric constants here were calculated from the behaviour observed earlier in the chapter, as such they should be considered camera specific, if this routine were to be applied to a different camera the values would probably have to be recalculated.

Since each pixel value has no dependence on future pixel values, there is no inherent delay involved in this processing. With purpose designed hardware, a radiometric value could be calculated for each pixel before the next pixel has been read from the detector,

and correction made in real time.



Figure 5.19, Swing vs. Masked input level, the masked input level is a weighted average of the values in the previous frame.



Figure 5.20, ΔVt vs. Mean Vi, ΔVt is the "target" value towards which ΔV decays, a line of best fit is calculated which is later used to calculate ΔVt for individual pixels in an arbitrary scene.

5.3.6 Scene contents correction algorithm summarised

The preceding sections have explained in detail the reasoning behind the conception of the scene contents correction algorithm and the mathematical derivation of the computations used. The operation of the algorithm as a whole has been summarised in the block diagram in Figure 5.21, this diagram is intended to serve as an aide memoire, which may be referred back to whist reading the discussions later in this chapter. It also serves as clarified illustration of exactly how much computation is involved in the application of the algorithm; every stage of computation involved is given, as are all the parameters used. The upper part of the diagram includes the stages of computation which occur once per frame at the start of the frame. The lower part of the diagram shows the computations which occur for each pixel of the frame. The per frame computation yields the corrected value of the first pixel, Vi₀, the insertion point of this value into the per pixel computation loop is illustrated

5.3.6.1 Summary of Parameters:

- Vi Ideal Voltage, the result of the algorithm
- Vm Measured Voltage, the output of the camera after IDP
- $\Delta V = Vi Vm$, The difference between measured and ideal voltages
- $\Delta V t$ the target ΔV value towards which ΔV decays

Subscripts 0, f and i are used to denote values for the first, the final and the ith pixels.

 $S = Vm_0 - Vm_f$, voltage swing, the difference between the start and end of the frame

F The factor by which ΔV decays towards ΔVt during the period of one pixel



Figure 5.21, Block diagram of scene contents correction algorithm

5.4 Performance of the scene correction algorithm

The algorithm generated in section 5.3.5 was applied to flat field data not used in the calibration process and to the non-flat field data used in section 5.2. The results are discussed here.

5.4.1 Performance on flat field data

The algorithm was not designed to have any effect on flat field data but it would still operate on flat field data in normal use, it was tested to ensure it has no adverse effect on the accuracy or uniformity of flat field results.

When flat field results with and without the algorithm were compared the values with the scene correction algorithm were not significantly better or worse than those obtained previously. The values with the algorithm are closer to the values calculated without the algorithm than they are to the values measured by the PRT's. This suggests the errors in the two sets of calculated values arise from the same source, whether that is an error in the measurement of the target or camera temperature or a drift in camera calibration over time cannot be determined. Non-uniformity measured by sNETD showed no significant difference either. As a final test, graphs were plotted of vertical sections through the flat field response data, and high contrast images of the flat field output with and without the algorithm compared. These are given in Figures 5.22 and 5.23. No peculiarities of response are seen in either the graphs, or the images. It was concluded that the scene correction algorithm does not adversely affect flat field radiometric performance.







Figure 5.23, Images of flat fields calibrated with and without scene contents correction

5.4.2 Performance on non-flat field data

The algorithm was designed to correct for irregularities of performance noted in section 5.2. In this section the algorithm is tested on the same scenes to asses its performance.

5.4.2.1 Vertical divisions

In section 5.2.1 camera output was tested on scenes with vertical divisions between hot and cold areas of the image. Figure 5.24 shows radiometric images of the same scenes, calculated with and without scene correction. All these images are coloured on the same scale which is given in the figure. The scene correction vastly improves the uniformity of the images, the gradation seen at the bottom of the uncorrected images is virtually eliminated.

Vertical sections through some of these images are given in Figure 5.25. Several parameters are plotted on each graph. This same set of parameters and line colouring is used throughout section 5.4.2. The black line corresponds to the radiometric values calculated without scene correction, the orange line shows values with scene correction. The green line gives Vm/20, and the green crosses are Vi/20. Δ V/10 is plotted in red, and Δ Vt/10 is plotted in blue. This set of lines allows all parameters relevant to scene correction to be compared on one graph, the corrected and uncorrected data are given along with the size of the correction.



Figure 5.24, Images of vertical divisions with and without scene contents correction



Figure 5.25, Scene correction multi-graphs, vertical sections through image 5.24 B

5.4.2.2 Horizontal divisions

Corrected and uncorrected images of the horizontal divisions seen in section 5.2.2 are given in Figure 5.26. In the corrected images both the distortions at the bottom of the image and those immediately above the division are reduced, now needing close inspection to be detected at all.



Figure 5.26, Images of horizontal divisions with and without scene contents correction

Graphs of vertical sections through these scenes are given in Figure 5.27, Graph 5.27a shows the non-uniformity in the hot region of the image is reduced from 5.1K to 1.9K. The uncorrected signal starts lower, at the bottom of the image, the corrected and uncorrected signals converge towards the same ΔT , approximately 30.5K at row 40. This level is approximately 2K below the PRT measured temperature of the black body. Looking at the raw data and ΔV curves, ΔV starts at almost zero, this is expected since in top half of the frame ΔT is close to zero. The level of ΔT increases from row zero to row 40, in response to the high signal level, the rate of change is logarithmic.



Figure 5.27, Scene correction multi-graphs, vertical section through image 5.20 C, column 50 Figure 5.27, Scene correction multi-graphs, vertical sections through images from figure 5.26

In the cold region at the top of the frame, the non-uniformity in uncorrected values is greater, 8.3K, this is reduced by over 80% to 1.5K in the corrected values. At the boundary between hot and cold, the value of ΔVt changes rapidly since it is determined by Vi, whereas the size of ΔV once again changes logarithmically.

The values of corrected and uncorrected data in Figure 5.27b have a form similar to those in graph a, with the y-axis inverted. However the form of the ΔV curve is different, it starts at a negative value, and drops towards zero before the division, and becomes positive after the division.

5.4.2.3 Corner Occlusions

The effect of corner occlusions was first examined in section 5.2.3. In Figure 5.28 images are shown of corner occlusions before and after correction. The most notable effect of the occlusions was distortion in the measured temperature of the hot background, with a strong horizontal boundary in background temperature at the same height as the horizontal edge of the cold plate. In the corrected images this effect can still just be made out but is only easily visible in those images with the greatest level of occlusion.

The magnitude of the remaining distortion was examined using graphs, two of these are given in Figure 5.29, both representing vertical sections through the same scene, Figure 5.28b. The first represents a section through column 20; this column covers the cool metal plate and crosses into the hot background at the top of the image. The second is a section through column 95; this column lies in the hot background along its entire height. The uncorrected curve in Graph a shows a similar profile to that in Graph 5.27b. Measured temperature is too high at the bottom of the frame and above the boundary. The level of correction achieved is slightly higher than that in Graph 5.27b. Once corrected the temperature variation in the cool area is only around 2 K, and less than 1 K in the hot region at the top of the image.



Figure 5.28, Images of corner occlusions with and without scene contents correction

The graph in Figure 5.29b also shows good correction. A variation of almost 10 K is reduced to 1.5 K. The red curve of the correction level shows how the initial correction decays to a steady value, which then becomes more negative above the boundary, at the top of the image. After correction, the boundary produces barely any distortion of the temperatures in column 95.



Figure 5.29b, Scene correction multi-graph, vertical section through image 5.28 B, column 95 Figure 5.29, Scene correction multi-graphs, vertical sections through image 5.28 B

5.4.2.4 Multiple Edges

Section 5.2.5 described a type of scene which was specifically designed to produce distortion in several areas of the image. Corrected and uncorrected images of these scenes are given in Figure 5.30. In both image the correction algorithm improves the image fidelity considerably. Banding in the hot regions is reduced to the extent that it appears invisible in one image and is only just discernable in the other. Within the cool region running up one side of the scene, banding is very prominent in the uncorrected image and reduced to the limits of visibility in the corrected image.



Figure 5.30, Images of multiple edges with and without scene contents correction

Figure 5.31 shows plots of vertical sections through those images in Figure 5.30a. The first graph runs through column 20, it shows the algorithm improves the uniformity in both hot and cold regions of the image. The two hot sections still exhibit slightly different temperatures relative to each other, this appears to be caused by a slight mismatch at the start of the frame, a mismatch between the time constant of the decay in the uncorrected data and the evolution of ΔV . It leads to a variation of approximately 0.6 K in the lower hot region. There are several possible causes of this mismatch, this effect was not investigated any further but is discussed in more detail in the conclusions of this chapter. Figure 5.31b is a plot through the hot band which runs up the centre of the image. Distortions of 7.5 K are reduced to 1.8 K, which is reduction by a factor of 4.2. The mean value of the corrected curve is very close to the measured black body temperature, whereas the uncorrected curve lies above this temperature for most of its length. Figure 5.31c is a plot through the cool region in the right of the image. The uncorrected curve shows the characteristic hump in the centre, in line with the bar on the left of the image. This effect is almost undecernable in the corrected curve. Aside from the hump mentioned, the uncorrected curve shows a general trend downwards, this effect remains in the corrected data but is significantly reduced in magnitude, from 13 K to 4.2 K.



Figure 5.31c, Scene correction multi-graph, vertical section through image 5.30 A, column 80 Figure 5.31, Scene correction multi-graphs, vertical sections through image 5.30 A

20

correction (ADC units / 10)

ΔVt / 10, target level of correction (ADC units / 10)

Section through

image 5.30 A:

oð

AT (K)

80

Row Number

5.4.2.5 Circular apertures

Figure 5.32 shows images of the circular apertures discussed in section 5.2.4, before and after correction. The noted distortions in these images were dark banding of the cool background, level with the height of the aperture. This banding is no longer visible in the corrected images. Two sections through Figure 5.32a are given in Figure 5.33. The section in the first graph passes through the aperture. What is apparent from this graph but not the image, is that the temperature measured in the aperture is significantly different in the corrected and uncorrected data. Over the aperture the uncorrected curve varies over a range of 4.5 K with a mean of 73 K. The corrected curve varies by only 1 K and has a mean of 63 K, much closer to the measured black body temperature of 64 K.



Figure 5.32, Images of circular apertures with and without scene contents correction

Figure 5.33b shows a section passing up the height of the metal plate, here a non-uniformity of 10 K is reduced to 2.1 K, the large dip at the height of the aperture, (seen as a dark band in the uncorrected image), is almost totally removed. These results suggest that this correction algorithm would vastly increase the effectiveness of the camera in taking radiometric measurements of small, hot targets against cool backgrounds. This observation is backed up further by the graph in Figure 5.34, which is a vertical section through the smaller aperture in Figure 5.32b. Of note is the height of the peak in the corrected and uncorrected curves and the measured black body temperature, given by the dashed blue line. Applying correction, the temperature



Figure 5.33, Scene correction multi-graphs, vertical sections through image 5.32 A



Figure 5.34, Scene correction multi-graph, vertical section through image 5.32 B, column 47

5.4.3 Effect of scene contrast on image correction algorithm

Figure 5.35 shows images of a scene not previously discussed, it includes an angled edge. The same scene is repeated with the black body at different temperatures. The uncorrected images show slight temperature distortion in the hot region and more significant distortion in the cool region. The distortion is in the form of a cooler band at the height of the black body. In the corrected images these distortions disappear almost completely. The level of distortion in the uncorrected images increases with black body temperature. At each temperature, the level of correction applied by the algorithm is appropriate to remove the distortion.



Figure 5.35, Images of angled plate with and without scene contents correction

Figure 5.36 gives three graphs through the same section through different images in Figure 5.35. The column is in the centre of the image, crossing the angled edge about half way up the height of the image. In graph a, the only distortion visible on the uncorrected curve is a dip in temperature after the transition from hot to cold, it is of the order of 1 K and is not present in the corrected curve. Graph b shows a slight gradient in the uncorrected temperature passing up the height of the hot black body; the same dip noted in graph a is also seen here, only now the magnitude is larger. Both of these effects are corrected by the algorithm, reducing the distortion in each area to the order of 1 K. In both graphs, correction brings the reported temperature of the black body closer to the PRT measured temperature. In each case, although the corrected level is not a perfect match, the degree of correction seems appropriate to the size of the distortion present.





5.5 Computational requirements

The algorithm developed is not computationally complex; it is, in essence, a type of digital filter, and could be executed on a low cost DSP. The computational requirement can be estimated by consulting the block diagram in figure 5.21. Considering first the per frame section; the weighted average looks complex on paper, but in fact the weighting factors can be stored in an array constant in memory, (ten thousand 16 floating point numbers will take 20 kB of memory). Calculating the weighted average is reduced to one multiply-accumulate operation per pixel, 10,000 per frame. These can be computed as the readings come in; there is no need to wait until the end of the frame before starting the calculation, so no latency issues arise from this. Multiplyaccumulate is a common DSP function which multiplies two variables and adds the result to a third variable, an accumulator. Multiply-accumulate can commonly be carried out with a single instruction which can be executed every clock cycle, adding to the same accumulator. The second block in the per frame section can be carried out with a single multiplication followed by addition, this needs two 16-bit floating point constants stored in memory. The third block is not computational, and the fourth block is a single addition. So the per frame section involves a total of 10,003 instructions, all but three of which can be computed at a steady pace during the previous frame.

Moving to the per pixel section; the first block is non-computational, the second involves a single addition, and the third requires one multiplication and one addition, and another two floating point constants. The fourth block is the final computational block, it requires one addition followed by a multiplication and then another addition, and also one more floating point variable. The total for the per pixel section is 6 instructions, or 60,000 per frame.

In total the algorithm requires 70,003 instructions per frame, which is 3.5 million instructions per second at 50 Hz frame rate, (since the algorithm runs after IDP it only happens every frame rather than every field). It requires 10,005 floating point constants, taking just over 20 kB of memory. Since the algorithm contains only multiplication and addition, and has no branches, it can be expected to run at close to 100 % efficiency; one instruction per clock on a floating point DSP. Meaning a DSP running at less than 10 MHz would meet the needs.

Alternatively, the algorithm could be run on a integer DSP using floating point emulation, this may be attractive because integer only DSP's tend to be cheaper, run at higher clock speeds and consume less power than their floating point counterparts. For example, the Analog Devices Blackfin family of DSP's are available at speeds of up to 750 MHz. At the time of writing a 300 MHz part sells for less than \$10 USD wholesale. At 300 MHz the Blackfin would have more than enough overhead for floating point emulation. By comparison, an Analog Devices Sharc floating point DSP costs around \$30 USD for a 300 MHz part.

These estimates do not make any use of single instruction, multiple data (SIMD) processing, which is available on most modern DSP's. The sequential nature of the algorithm would make the acceleration available through SIMD negligible. However, radiometric calculations would lend themselves well to SIMD acceleration, because the values are independent for each pixel but the operations required are the same. What becomes apparent is that it may well be possible to run both the scene correction algorithm and radiometric calculations on the same DSP. The number of instructions needed to calculate radiometric output from the IDP data will vary widely depending on how the radiometric model is implemented so no attempt will be made to quantify it here.
5.6 Conclusions

It was found that when bright patterns are introduced to scenes captured with this camera, various characteristic distortions occur. Particularly noticeable are the effects of hot objects on the temperature recorded in the cooler areas of the image. These effects were investigated by creating a logical series of patterns and examining the resulting output both visually and numerically. It was postulated that these distortions are caused by the output gain of the detector falling when the output voltage is high for extended periods of time. A mathematical model of the observed behaviour was conceived and from this an algorithm was created to calculate and correct for the level of distortion in any image. The parameters used in the algorithm were derived by measuring properties of the flat field output, not by adjusting them to fit the patterned scenes.

When the algorithm was applied to patterned scenes the results were very positive. Image uniformity was improved in all scenes tested. Distortion in what should be flat areas is usually reduced by a factor of 4 or better. In no cases was the corrected image found to be worse than the uncorrected image. In nearly all cases the remaining distortion was of a similar form to the original distortion, except smaller in magnitude. An obvious but misguided conclusion would be to think the level of correction should simply be increased, however increasing the level of correction would lead to overcorrection in other areas of the image. It seems likely that some distortion remains because relationships assumed to be linear are in fact slightly non-linear. It was assumed that the time constant describing the logarithmic evolution of the correction factor is constant across the image but since some areas are well corrected and others less so, this may not be the case. Specifically, the term F in equation 5.1 may have some slight dependence on Vm or ΔV ; it is even conceivable that it varies with in a fixed pattern for each pixel of the detector.

Investigating these possibilities would be a long process and may not prove fruitful. It can be considered as possible future work. It would be more useful to repeate the same experiments with other detectors of the same model, and with higher resolution detectors which use the same readout technology. Such an investigation would have to be carried out before the value of this algorithm to any commercial camera design can be predicted. Since this range of pyroelectric detectors is no longer being developed or manufactured, it seems unlikely that a direct commercial application for the algorithm with be found. Another area for further investigation would be to examine the design of the detector and its readout circuitry, to see if the whole problem might be removed at source, this would clearly be a preferable solution if it were possible. Unfortunately, due to the evolution of DERA into QinetiQ during the process of this research, it was not possible to examine the finer details of the readout circuitry in this way.

As well as reduced distortion, the radiometric accuracy of the corrected images was also improved, in some cases by large margins. Of particular note are scenes where small hot shapes appear against a cool background. In the test scenes used radiometric error was reduced from around 10 K to less than 1 K. This would offer great benefit to many practical radiometric applications, where the object whose temperature must be known is hotter than the local environment and cannot easily be made to fill the image.

Conclusion

In this work the internal signal processing electronics of an uncooled pyroelectric LWIR camera were replaced with a purpose built digitiser. The new digitiser was of a simplified design; all signal processing was moved from dedicated hardware to software analysis on a PC. A model was conceived which predicted that the response of the detector to radiation power would be linear, and consequently, that the relationship between target temperature and IDP output would follow the form of the integral of Planck's Law of black body radiation over the $8 - 12 \,\mu$ m waveband, this was found to be the case. The model also predicted that knowing the chopper temperature would permit radiometric measurements to be made with this camera. A platinum resistance thermometer was installed within the detector chamber for this purpose, radiometric calibration was successful, and radiometric performance was shown to be accurate within 1.5 K, and stable when imaging flat fields or low contrast scenes. When imaging high contrast scenes significant temperature distortion was found in the output, an algorithm was devised to correct for this distortion. The correction algorithm improved radiometric performance and temperature uniformity in all scenes tested.

The original signal processing system combined analogue offset removal with a low resolution analogue to digital converter. This system was replaced by an ADC with 14-bit resolution and no offset removal. The new digitiser was shown to have a noise level of 2.3 ADC units (standard deviation on the digital signal). This cannot be considered perfect performance, but in comparison to a noise level on the output of the detector, which was five times higher, it proved adequate. The new digitiser reduces the number of components needed in the analogue circuit and removes one stage of detector calibration. In the original design offset voltages must be measured and stored in ROM, in the new system they are simply ignored, being cancelled out later by the image difference processing routine. Both of these changes would reduce manufacturing costs.

The camera had not previously been calibrated for radiometric measurement. One of the main aims of this research was to attempt radiometric calibration of the camera, to improve its value to commercial markets. A calibrated PRT was installed within the camera to measure the temperature of the chopper blade, so it was as close to the chopper blade as possible. Detector performance fitted the model for a thermal detector, as described above. Measured response data were used to calibrate the camera, on a per-pixel basis, for radiometric temperature measurement, over a range of target temperatures from 20 °C up to 100 °C and a range of camera temperatures from 16 °C to 40 °C. One calibrated the camera was found to have a tNETD of 0.8 K at a scene temperature of 30 °C. Having calibrated the camera once, the largest radiometric error found when measuring the temperature of flat fields at various times over an 18 month period was less than 1.5 K. In all cases images good uniformity was maintained; the sNETD remained smaller than the tNETD by a at least a factor of two, in all cases. This implies that comparative radiometric measurement between different black body sources within the same frame would be accurate to the level of the tNETD. Unfortunately, this was not the case, tt was found that when imaging high contrast scenes the detector's output was severely distorted, the presence of hot or cold regions within an image causing the detector's gain to vary throughout the image, and producing radiometric errors as large as 15 K.

By examining graphs of cross-sections through various scenes a model was proposed for this behaviour, and an algorithm was created to correct the distortion. Measurements on flat field data were used to calculate the parameters used in the algorithm. When applied to high contrast scenes, the algorithm produced remarkable improvements in image quality and radiometric accuracy. The algorithm reduced the peak radiometric error to less than 5 K, and had no adverse effects on low contrast and flat field scenes. Whilst the algorithm has been shown to function very well, the source of the distortion it corrects remains unknown. Evidence suggests that this effect may well be a quirk of the individual camera tested, a flaw of some kind in the analogue electronics of the detector or camera. It is possible that the distortion could be removed at source by repairing or redesigning elements of the camera or the detector, this must be considered a more desirable solution.

The level of radiometric accuracy achieved is not exceptional but is good enough to be of value in many possible markets, accuracy might be improved by further work, particularly by placing more temperature sensors within the detector chamber and producing a more extensive model of the thermal environment within the chamber. However, since this camera architecture is now considered obsolete, such work would likely have no commercial value.

The temperature range over which the camera could be calibrated was limited by the hardware available. Nevertheless, a distinct advantage of the simplified digitiser system is that the camera could be calibrated over a very large range of target and camera temperatures, with no loss of thermal resolution. The voltage input range of the ADC covers the entire output range of the detector (excluding dead pixels), the size of the signal produced by each pixel being small in comparison to the offset voltages. As the input level increases no change in analogue gain is needed. By contrast, if analogue offset removal is used the input range to the ADC must be chosen to match the expected signal level. Large input signals may saturate the ADC if the input is outside the expected range. Analogue offset removal will always allow finer voltage resolution with the same ADC, but when the resolution of the ADC can be made finer than the noise level of the detector without offset removal (as in this case), offset removal becomes redundant. It will introduce an extra source of error and introduces a risk of signal saturation, a risk which need not be present.

Whilst all processing in this work was carried out off-line after saving data to a hard disk, the algorithms designed do not preclude real time use, including the scene contents correction algorithm. Since the correction applied to a pixel is dependent on previous pixel readings but not subsequent ones, a radiometric value could be calculated for each pixel before the subsequent pixel voltage has even been read from the detector. This distinction is relevant to some markets, perhaps most obvious being the field of machine vision. A computer can respond much faster than a human and in some applications it may be important that latency should be reduced to an absolute minimum, the algorithms devised allow that. Some human operated systems also benefit from low latency, for example, imaging systems used for navigation and IR targeting scopes for weapons systems.

Further Work

Considering the project as a whole, a discussion of further work must include how the systems developed here could be integrated into commercial products. It was established in chapter 2 that the growth area in infrared imaging is the field of low cost, high volume infrared imagers. The three areas of work in this project, the new digitiser, the radiometric calibration and the scene contents correction algorithm could all be of value in the development of low cost cameras, but since other detector technologies have proved to be more cost effective, the areas of work specific to this architecture are unlikely to be applied directly. A limitation of the research carried out in this project is that it has all been conducted with only one detector, a detector which was not of the latest type when the project was started in 1999.

The simplified digitiser offers the prospect of reduced materials cost and a cheaper calibration process. However, it would have to be developed further to be applied to new detectors. The higher pixel count of current detectors must be considered, the ADC used is capable of running at up to 10MHz. The detector in this project produced 1,020,000 pixels per second running at its frame rate of 100Hz, this means the ADC could operate with a detector of up roughly ten times the pixel count with its current frame rate. That is not enough to support a 640x480 (VGA) resolution detector but is enough to drive a 320x240 (QVGA) sensor at 100Hz (which requires a 7.7MHz pixel clock). At the time of writing VGA uncooled arrays have become widely available at the high end of the uncooled imaging market and QVGA arrays are quite common in the middle of the price range, although the cheapest imagers still run at lower resolutions. The ADC used could match the clock speed requirements of the mid-range market today and could also drive low resolution sensors at very high clock speeds. The digitiser design might be suited to reading the output of a microbolometer array with little or no modification.

After clock speed the second limitation of the ADC is its 14-bit resolution. As discussed in Chapters 3 and 4, the resolution of the ADC limits the ultimate thermal sensitivity of the imager it is used in. In the absence of analogue offset removal this limitation becomes more significant than it would otherwise be; with the output range of this detector the 14 bit resolution would limit thermal sensitivity to 80 mK at low input levels. (Resolution with this detector would be slightly finer at high signal levels, a point discussed in Chapter 4.) If this digitiser was used with another detector the

limiting resolution would depend on the characteristics of that detector. If the total voltage output range of the detector were reduced in relation to the signal then the ADC could reach finer resolutions. This could be achieved in two ways, smaller offsets between pixel voltages or greater signal response; either of which may be applicable for newer detectors. It was found that the performance of the digitiser did not match its ultimate resolution; noise on the output was found to have a standard deviation of 2.26 units which compares a figure of 0.36 quoted in the specification sheet. The noise level on the detector signal was found to be several times larger, so the digitiser noise was not considered a problem in this case, but it could limit performance if this digitiser were used with another detector. To maximise performance from this ADC chip the digitiser would have to be rebuilt, probably starting by having the same design built by machine on a multilayer PCB.

To match the best detectors available today a new ADC chip would have to be used. Fortunately in the time since this digitiser was built, higher resolution, higher speed ADC chips have become available. The fundamental design of the digitiser could remain largely unchanged, the clamping op-amp and the line drivers used would work with a more highly specified ADC. Any plan to use such a digitiser would have to start by identifying the detector to be coupled, analysing its output and choosing an ADC which would match its performance requirements.

The radiometric calibration of the camera relied on the placement of a temperature sensor inside the detector cavity. This was found to allow reasonable radiometric performance so long as the camera was kept in a stable thermal environment. In order to improve radiometric performance in the less stable thermal environments which might commonly be encountered in real world applications, further work would be required. The placement of the single PRT added to the detector chamber used was restricted by the design of the camera. The single sensor was used to correct for the temperature distribution within the entire inside of the detector cavity, and the temperature of the chopper blade, this worked acceptably because the temperature relationship between those components was kept stable by the stable environment. To make significant progress in this area would probably require at least 3 sensors in the detector cavity, one for the detector, one for the wall of the cavity and one for the chopper blade. Since the chopper blade is spinning a single element radiometric sensor might be suitable for that task. Once these were added, the performance of the camera in dynamic thermal environments would have to be calibrated, a task of some complexity which would require a thermally controlled chamber large enough to house the whole camera.

The radiometric performance under dynamic thermal environments might be further improved, and calibration simplified, by making some changes to the camera design within the detector cavity; the addition of a device called an optical shield has been shown to benefit performance in similar cameras^{R39}. An optical shield is a form of hood or cowling which fits around the detector, it is made of a metal with high thermal conductivity and given a black coating. The purpose of the shield is to block stray thermal radiation from other parts of the detector cavity from reaching the detector. Since the shield is thermally conducting and maintains a uniform surface temperature, the amount of radiation coming from it can be accurately predicted with a single temperature sensor within the fabric of the shield. By putting the shield in thermal contact with the detector chip, the relationship between detector temperature and shield temperature might be made constant, possibly reducing the number of temperature calibration parameters from at least 3 to 2 (the shield/detector and the chopper).

It would be useful to extend the calibrated range of the camera to cover both hotter and colder target temperatures. As discussed in chapter 4 one advantage of abandoning analogue offset removal is that the calibrated temperature range could be increased massively without loosing temperature resolution, the range covered in this project was limited by the black body hardware available. The calibration in this project was all carried out with one lens at f/1 aperture; obviously the calibration would have to be adjusted if a different lens were used. The model developed for the camera would allow for a correction factor when using the same lens with a smaller aperture, but to extend this to other types of lens would probably require the use of a monochromater to measure the true spectral response of the detector and the lenses. This is because the different lenses would have different transmission spectra and as such would change the shape of the response curve. The use of a monochromater would allow a more complete modelling of the detectors performance. Changing lenses would also introduce the problem of lens related non-uniformity, any such effect was effectively "calibrated out" in this case, but if multiple lenses were used it might be necessary to calculate and apply corrections for each lens. That said; inspecting the distribution of response and noise levels across the detector in chapter 4 did not show any clear

evidence for lens related non-uniformity, darkening towards the corners of the array would have been the most obvious sign.

Further work on the scene correction algorithm could go in two directions; the algorithm might be developed further to improve performance with the same camera, or work might be carried out to see if the algorithm would be applicable to other kinds of detectors. A first line of investigation to improve the performance of the algorithm for this detector might be to investigate further the nature of the logarithmic evolution of the correction level ΔVi throughout the image, this possibility and some of the problems associated with it were discussed in the conclusion of Chapter 5. There seems to be a reasonable possibility that a modified version of the algorithm could be applicable to other detectors, perhaps the biggest challenge here would be identifying which detectors to test. Exposing candidate detectors to high contrast patterns similar to those used in Chapter 5 and observing the results would quickly give some idea as to whether it was worth taking the process further. Similar distortion could occur in a camera which does not use image difference processing; indeed the suggested cause of the effect given in Chapter 5 has nothing to do with image difference processing. However, the parameters used in the calibration were found by measuring the size of effects at the start of the frame which would probably not be present in a camera which did not use IDP. The relationship between ΔVt and Vi, given in the graph in figure 5.20 would have to be obtained in some other way, such as measuring the size of distortion caused by horizontal divisions with different contrast ratios. The same principle could equally apply to contrast in visible light imagers. It has already been mentioned, but should be repeated here, that it may be possible to remove the distortion at source, this would be preferable to using the algorithm, but the algorithm could still prove useful in situations where a direct solution is not possible.

References

R1

Principles of Thermal Imaging A P Davis, A H Lettington Publication: Applications of thermal imaging, p.1-34, S G Burnay, T L Williams, C H Jones; Eds. Publisher: Published under Adam Hilger, Bristol and Philadelphia, by IOP Publishing Ltd. Publication Date: 1988 Page: 3, fig 1.2

R2

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 2.3, Page 9, Eqn 2-8

R3

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 2.3, Page 10, Eqn 2-24

R4

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 2.3, Page 10, Eqn 2-21

R5

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 20, Page 784-5, Sec 20.1, Surface Property Definitions

R6

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 9, Page 373, Sec 9.2.1.3 Lagrangian Invariant

Industrial and Civil Applications of Thermography T L Williams Publication: Applications of thermal imaging, p.35-72, S G Burnay, T L Williams, C H Jones; Eds. Publisher: Published under Adam Hilger, Bristol and Philadelphia, by IOP Publishing Ltd. Publication Date: 1988 Page: 39, fig 2.1

R8

Analysis of emissivity effects on target detection through smokes/obscurants W. Michael Farmer Optical Engineering -- November 1991 -- Volume 30, No.11, p 1701 Publication Date: Nov 1991

R9

Sources of Radiation, (Particularly: Chapter 3, 'Natural Sources' by David Kryskowski, Gwynn H. Suits, p. 139-316) Publication: The Infrared & Electro-Optical Systems Handbook, Volume 1: Sources of Radiation, Ed. George J. Zissis Co-published by: Infrared Information Analysis Centre Environmental Research Institute of Michigan & SPIE Optical Engineering Press Publication Date: 1993

R10

Challenges for third-generation cooled imagers Stuart Horn, Paul Norton, T. Cincotta, Andrew J. Stoltz, Jr., J. D. Benson, Philip Perconti, J. Campbell Publication: Proc. SPIE Vol. 5074, p. 44-51, Infrared Technology and Applications XXIX; Bjorn F. Andresen, Gabor F. Fulop; Eds. Publication Date: Oct 2003

R11

Theory and Practice of Radiation Thermometry D P DeWitt, Gene D Nutter Publisher: A Wiley-Interscience Publication, John Wiley & Sons, INC. Published:1988 Chapter 4, Sec. 4.4.4, Page 281

R12

Alternative measurement techniques for infrared sensor performance Arie N. de Jong, Eric M. Franken, and Hans Winkel TNO Physics and Electronics Laboratory, P.O. Box 96864, 2509 JG The Hague, The Netherlands Optical Engineering -- March 2003 -- Volume 42, Issue 3, pp. 712-724 Publication Date: March 2003

Theory and Practice of Radiation Thermometry D P DeWitt, Gene D Nutter Publisher: A Wiley-Interscience Publication, John Wiley & Sons, INC. Published:1988 Chapter 4, Sec. 4.4.2, Page 275

R14

Fast-cool-down dual gas spray-cooler for pivoted IR detectors Uwe G. Hingst Publication: Proc. SPIE Vol. 5074, p. 620-629, Infrared Technology and Applications XXIX; Bjorn F. Andresen, Gabor F. Fulop; Eds. Publication Date: Oct 2003

R15

Piccolo: a high-performance IR detector optimized for handheld applications
Tuvy Markovitz, Fabian Schapiro, D. Alfiya, S. Hasson, A. Magner, Ofer Nesher
Publication: Proc. SPIE Vol. 5406, p. 239-248, Infrared Technology and
Applications XXX; Bjorn F. Andresen, Gabor F. Fulop; Eds.
Publication Date: Aug 2004

R16

Benefits of microscan for staring infrared imagers
John L. Miller, John M. Wiltse
Publication: Proc. SPIE Vol. 5407, p. 127-138, Infrared Imaging Systems: Design,
Analysis, Modeling, and Testing XV; Gerald C. Holst; Ed.
Publication Date: Aug 2004

R17

Theory and Practice of Radiation Thermometry D P DeWitt, Gene D Nutter Publisher: A Wiley-Interscience Publication, John Wiley & Sons, INC. Published:1988 Chapter 7, Sec. 7.4.3, Page 518, Time Constants

R18

Theory and Practice of Radiation Thermometry D P DeWitt, Gene D Nutter Publisher: A Wiley-Interscience Publication, John Wiley & Sons, INC. Published:1988 Chapter 7, Sec. 7.4.3.3, Page 522, Pyroelectric detectors

R19

Theory and Practice of Radiation Thermometry D P DeWitt, Gene D Nutter Publisher: A Wiley-Interscience Publication, John Wiley & Sons, INC. Published:1988 Chapter 7, Sec. 7.2.6.1, Page 507, Chopped radiometers.

Resolution and sensitivity improvements for VOx microbolometer FPAs Daniel F. Murphy, Adam Kennedy, Michael Ray, Richard Wyles, Jessica Wyles, James F. Asbrock, C. Hewitt, David Van Lue, T. Sessler, John S. Anderson, Daryl Bradley, Richard Chin, H. Gonzales, C. Le Pere, Thomas Kostrzewa Publication: Proc. SPIE Vol. 5074, p. 402-413, Infrared Technology and Applications XXIX; Bjorn F. Andresen, Gabor F. Fulop; Eds. Publication Date: Oct 2003

R21

Theory and Practice of Radiation Thermometry D P DeWitt, Gene D Nutter Publisher: A Wiley-Interscience Publication, John Wiley & Sons, INC. Published:1988 Chapter 7, Sec. 7.4.1.3, Page 515, Thermistor Bolometers

R22

Commercialization of uncooled infrared technology Peter W. Norton, Margaret Kohin, Mathew Dovidio, Brian S. Backer Publication: Proc. SPIE Vol. 5563, p. 55-61, Infrared Systems and Photoelectronic Technology; Eustace L. Dereniak, Robert E. Sampson, C. Bruce Johnson; Eds. Publication Date: Oct 2004

R23

Performance of 320x240 uncooled bolometer-type infrared focal plane arrays Yutaka Tanaka, Akio Tanaka, Kiyoshi Iida, Tokuhiro Sasaki, Shigeru Tohyama, Akira Ajisawa, Akihiro Kawahara, Seiji Kurashina, Tsutomu Endoh, Katsuya Kawano, Kuniyuki Okuyama, Kazuyuki Egashira, Hideo Aoki, Naoki Oda Publication: Proc. SPIE Vol. 5074, p. 414-424, Infrared Technology and Applications XXIX; Bjorn F. Andresen, Gabor F. Fulop; Eds. Publication Date: Oct 2003

R24

The design of low-noise MOSFETS for pyroelectric array readout, (LAMPAR) R. Watton, Paul A. Manning Publication: Proc. SPIE Vol. 807, p. 98-105, Passive Infrared Systems and Technology Publication Date: 1987 Page 98

R25

The design of low-noise MOSFETS for pyroelectric array readout, (LAMPAR) R. Watton, Paul A. Manning Publication: Proc. SPIE Vol. 807, p. 98-105, Passive Infrared Systems and Technology Publication Date: 1987 Page 98

High-sensitivity uncooled microcantilever infrared imaging arrays
Scott R. Hunter, Robert A. Amantea, Lawrence A. Goodman, David B. Kharas, Sergey
Gershtein, James R. Matey, Steven N. Perna, Young Yu, Nagi Maley, Lawrence K.
White
Publication: Proc. SPIE Vol. 5074, p. 469-480, Infrared Technology and
Applications XXIX; Bjorn F. Andresen, Gabor F. Fulop; Eds.
Publication Date: Oct 2003

R27

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 8, Page 315-326, Optical Materials

R28

Polymer imaging optics for the thermal infrared Nelson E. Claytor, Richard N. Claytor Publication: Proc. SPIE Vol. 5406, p. 107-113, Infrared Technology and Applications XXX; Bjorn F. Andresen, Gabor F. Fulop; Eds. Publication Date: Aug 2004

R29

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 8, Page 326-333, Optical Material Properties

R30

Handbook of Military Infrared Technology Editor: William L Wolfe Publisher: Office of Naval Research Department of the Navy Washington D.C. Published: 1965 Chapter 8, Page 347-358, Reflection & Transmission, Optical Surface Coatings

R31

Evaluation of thermal imaging cameras used in fire fighting applications Francine Amon, Nelson Bryner, Anthony Hamins Publication: Proc. SPIE Vol. 5407, p. 44-53, Infrared Imaging Systems: Design, Analysis, Modelling, and Testing XV; Gerald C. Holst; Ed. Publication Date: Aug 2004

European uncooled thermal imaging sensors Kennedy R. McEwen, Paul A. Manning Publication: Proc. SPIE Vol. 3698, p. 322-337, Infrared Technology and Applications XXV; Bjorn F. Andresen, Marija Strojnik; Eds. Publication Date: Jul 1999 Page 323

R33

European uncooled thermal imaging sensors Kennedy R. McEwen, Paul A. Manning Publication: Proc. SPIE Vol. 3698, p. 322-337, Infrared Technology and Applications XXV; Bjorn F. Andresen, Marija Strojnik; Eds. Publication Date: Jul 1999 Page 333

R34

An uncooled solid state thermal imager McEwen, R.K.; MacLeod, A.D.; Manning, P.A.; Perkins, M.C.J. Publication: IEE Conf. Publication No 321, p. 93-99, Advanced Infrared Detectors and Systems, 1990., Fourth International Conference on Publication Date: 5-7 Jun 1990 Page 94

R35

European uncooled thermal imaging sensors Kennedy R. McEwen, Paul A. Manning Publication: Proc. SPIE Vol. 3698, p. 322-337, Infrared Technology and Applications XXV; Bjorn F. Andresen, Marija Strojnik; Eds. Publication Date: Jul 1999 Page 324

R36

Extensive IR lens and coating performance data was available from: Janos Technology, The Monroe Group, 400 Skokie Blvd., Suite 560, Northbrook, IL 60062 (http://www.janostech.com/)

R37

Thermal conductivity data from NIST Boulder Laboratories NIST Boulder, 325 Broadway, Mailcode 107.00 Boulder, Colo. 80305-3328 (materials data was available at www.boulder.nist.gov at the time of writing.)

Technische Temperaturmessungen-Strahlungsthermometrie Publisher: VDI/VDE Richtlinie. Publication: Technische Temperaturmessungen-Strahlungsthermometrie. VDI 3511 Page 4. Published: January, 1995. (Found in literature from Raytek Corporation: Raytek Corporation, 1201 Shaffer Road, Santa Cruz, CA 95061-1820 USA)

R39

320x240 uncooled microbolometer 2D array for radiometric and process control applications Jean-Luc Tissot, Jean-Pierre Chatard, Sebastien Tinnes, Bruno Fieque Publication: Proc. SPIE Vol. 5074, p. 396-401, Infrared Technology and Applications XXIX; Bjorn F. Andresen, Gabor F. Fulop; Eds. Publication Date: Oct 2003

List of acronyms

Analogue to Digital Converter
Charge Coupled Device
Complimentary Metal Oxide Semi-conductor
Cathode Ray Tube
Defence Evaluation Research Agency (now part of QinetiQ)
Electromagnetic
Integrated Circuit
Interactive Data Language
Image Difference Processing
Infrared
Long Wave Infrared
Metal On Silicon Field Effect Transistor
Minimum Resolvable Temperature Difference
Modulation Transfer Function
Medium Wave Infrared
Noise Equivalent Temperature Difference
Printed Circuit Board
Platinum Resistance Thermometer
Region Of Interest
Sira Electro Optics
Spatial Noise Equivalent Temperature Difference
Short Wave Infrared
Temporal Noise Equivalent Temperature Difference

