

Developing Space Instrumentation for Exploring Airless
Bodies in the Mid-Infrared in Support of NASA's Lunar
Trailblazer and ESA's Comet Interceptor



Nicholas Julian Elkington

Supervised by Professor Neil Bowles

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2024

Abstract

This thesis describes work that supports the development of thermal instruments for two key space missions: NASA's Lunar Trailblazer and ESA's Comet Interceptor. The focus is on the Lunar Thermal Mapper (LTM), designed for accurate thermal mapping of the lunar surface, and the Modular InfraRed Molecules and Ices Sensor (MIRMIS), intended to study the volatile content of a dynamically new comet.

The work presented covers the optimisation of internal temperature measurements, calibration procedures, and overall radiometric performance of each instrument. For LTM, extensive testing and calibration ensure that the instrument meets stringent performance requirements, enabling precise temperature measurements and supporting accurate mineralogical assessments of the lunar surface. Key challenges include accounting for the thermal environment and improving calibration accuracy under flight conditions. For MIRMIS, the work described within this thesis develops operational simulations to guide instrument design and ensure that it meets its scientific objectives during high-velocity flybys of comets. The simulation provides a comprehensive end-to-end model, linking operational parameters to scientific outputs. This simulation is crucial for enhancing the involvement of the scientific community in the development of the MIRMIS instrument, ensuring that its design and operations are fully aligned with the mission's scientific goals.

This research forms the foundation for future instrument operations, with applications covering mission scenarios of complex thermal environments, thus contributing significantly to the scientific success of both space missions.

Acknowledgements

"I can't carry it for you, but I can carry you" – Samwise Gamgee

This thesis may be a record of my work as an individual, but this project would never have been completed but for the support of a great many people.

First and foremost, I want to thank my supervisor, Neil, for taking a chance on me, opening the doors to this project, and being a constant source of support over these four years. His encouragement was always apt - nudging me when I lagged and trusting me when I moved forward independently.

A huge thank you goes to Tris, Katherine, Keith, and Maisie. Tris, your practical advice and endless patience with my attempts to automate anything electronic have been critical. Katherine, thank you for bringing us back to actually thinking about the targets. Keith, for reminding me of the fun in our mad work environment; and Maisie, for setting me on the right track to genuine scientific independence.

I also want to thank everyone else in the lab: Rory, for designing LTM so that I had a mission to support; Greg and Jordan, for their tireless electronics support; Andy, Jon, and Waqas, for building and maintaining all the equipment that makes doing science possible; and Henry, Burtie, Rowan, Eloise, and all the other students who have passed through our group for their company.

Outside of Oxford, I'm grateful to Aurelie, Sara, Geronimo, and the whole MIRMIS Team, as well as Bethany and everyone involved in Lunar Trailblazer. A sincere thanks goes to STFC and UKSA, whose financial support was key to bringing this project to life.

As essential as the academic support has been, the non-academic support has been equally important, and often more meaningful. To my fellow Fellows at Bullingdon College,

particularly the Master and the Dean, thank you for adding such colour to my life and making Oxford truly my home.

To my countless crewmates over the last eight years: Charlie, without whom this project would likely never have started; Ben, for maintaining my sanity during a year of lockdown; and Luke, Chris, Tom, Tobias, Johnny, Saxon, Alex, Felix, Louis, Tavish, and everyone else who has hauled a boat down a river with me, or against me. Those experiences, much like the friendships they have wrought, will last a lifetime.

I'm so grateful to the friends outside of the boat too, particularly those in the DBR who have consistently brought so many great times. Annie and Julia, Augustin and Erin – Oxford simply would not have been the same without you.

I appreciate the support of all the other friends I have who have provided me with so many wonderful distractions along the way. Particularly all my peers from Christ Church who have been patient through my DPhil-induced reclusion.

Lastly, I am profoundly grateful for the constant support of my family. My elder siblings, Annie and Jack, deserve a special mention. Their achievements across academic, social and sporting settings throughout my life have paved the way for my own journey and continue to serve as inspirations. I owe a special thanks to all my grandparents for laying the foundations for the curious and scientific environment I was raised in – especially my grandfather Julian, whose mind has left an indelible mark in so many ways.

Most importantly, I thank my parents for their endless love and unwavering support. This has been most conspicuous in their permanent willingness to indulge in my distractions on the river, without which this DPhil would never have come to pass. It is no surprise to me that one of my earliest memories is a demonstration of Newton's Laws of Motion at the Science Museum in Kensington. I have loved coming to understand the world around me ever since, and that is all because of you.

Table of Contents

ABSTRACT	I
ACKNOWLEDGEMENTS	II
TABLE OF CONTENTS.....	IV
TABLE OF ACRONYMS	VIII
IMPACT OF COVID-19 AND MISSION DELAYS	X
1 INTRODUCTION.....	1
1.1 WATER AND MINERALOGY ON THE MOON	5
1.1.1 <i>Volatiles in Lunar Surfaces</i>	6
1.1.2 <i>Volatiles in PSRs</i>	7
1.1.3 <i>Neutron Spectroscopy</i>	10
1.1.4 <i>M³ Measurements of Lunar Water</i>	12
1.1.5 <i>LCROSS Ejecta Measurements</i>	15
1.1.6 <i>SOFIA Measurements</i>	15
1.1.7 <i>Chang'E 5 Results</i>	16
1.2 VOLATILES IN COMETS AND ASTEROIDS	17
1.2.1 <i>Evolution of Comet Knowledge</i>	18
1.2.2 <i>Isotropic Ratios for Water Origination</i>	21
1.2.3 <i>The Thermal Environment of the Nucleus</i>	25
2 MEASURING THE LUNAR SURFACE	29
2.1 METHODOLOGY OF REMOTE SENSING	29
2.1.1 <i>Remote Sensing of Volatiles</i>	29
2.1.2 <i>Compositional Remote Sensing</i>	32
2.2 PREVIOUS LUNAR MEASUREMENTS	34
2.2.1 <i>The Diviner Lunar Radiometer</i>	34

2.2.2	<i>The Moon Minerology Mapper</i>	37
2.3	LUNAR TRAILBLAZER.....	39
2.3.1	<i>The Trailblazer Mission</i>	39
2.3.2	<i>The High-Resolution Volatiles and Minerals Moon Mapper</i>	41
2.3.3	<i>Lunar Thermal Mapper</i>	43
2.4	MISSION REQUIREMENTS.....	47
2.4.1	<i>Science Requirements of LTM</i>	49
3	LTM CHARACTERISATION AND DESIGN WORK	51
3.1	RADIOMETRIC MODELLING.....	52
3.1.1	<i>Instrument Model</i>	52
3.1.2	<i>A Combined Thermal-Radiometric Model</i>	62
3.2	CHARACTERISING THE LTM OPTICAL PATHWAY.....	73
3.2.1	<i>Spectral Measurements</i>	74
3.2.2	<i>Interference Filters</i>	76
3.2.3	<i>Mesh Filters</i>	82
3.2.4	<i>Instrument Transmission Functions</i>	87
3.2.5	<i>Verification of Instrument Transmission Function</i>	90
3.3	INTERNAL TEMPERATURE MEASUREMENTS.....	92
3.3.1	<i>The PRT Calibration Procedure</i>	93
3.3.2	<i>Housekeeping Board Development</i>	99
3.4	LTM BLACKBODY CALIBRATION	109
3.4.1	<i>Calibration Procedure</i>	110
3.4.2	<i>Blackbody Calibration Results</i>	116
3.5	RADIANCE-TEMPERATURE REFERENCE TABLES	122
4	COMET INTERCEPTOR AND MIRMIS	125
4.1	THE COMET INTERCEPTOR MISSION.....	125
4.2	THE MIRMIS INSTRUMENT.....	127
4.2.1	<i>The Thermal InfraRed Imager</i>	128

4.2.2	<i>Mid-InfraRed and Near-InfraRed Imagers</i>	130
4.3	MIRMIS CONCEPT OF OPERATIONS.....	130
4.3.1	<i>Key Scanning Challenges</i>	133
4.3.2	<i>The Comet Interceptor Nominal Case</i>	135
4.4	INSTRUMENT REQUIREMENTS.....	136
4.4.1	<i>Official Requirements</i>	137
4.4.2	<i>Community Requirements</i>	139
5	MIRMIS DESIGN WORK	141
5.1	DEFINING THE MIRMIS OPERATIONS	142
5.1.1	<i>Thermophysical Modelling</i>	142
5.1.2	<i>Radiometric Performance Modelling</i>	150
5.1.3	<i>Blurring and Other Integration Time Effects</i>	154
5.1.4	<i>Mirror Pointing Modelling</i>	159
5.1.5	<i>The Full TIRI Pointing Plan</i>	166
5.1.6	<i>Alternate Pointing Plans</i>	169
5.2	MIRMIS POWER CONSUMPTION.....	171
5.2.1	<i>Power Modelling</i>	171
5.2.2	<i>Alternative Power Scenarios</i>	176
5.3	TIRI FILTER SELECTION MODELLING	177
5.4	MOTOR COMMANDING.....	182
5.4.1	<i>Analysis Methodology</i>	184
5.4.2	<i>Simulation of the Motor Commanding</i>	186
5.4.3	<i>Commanding Conclusions</i>	198
5.5	HOUSEKEEPING BOARD DEVELOPMENT.....	199
5.5.1	<i>Relevant Requirements</i>	199
5.5.2	<i>Test Procedure</i>	200
5.5.3	<i>Testing Results</i>	204
5.5.4	<i>Housekeeping Board Conclusions</i>	213

6	TIRI OBSERVATION SIMULATION.....	214
6.1	SIMULATION INPUTS	215
6.2	SIMULATION STRUCTURE OVERVIEW	216
6.3	THE FULL FLYBY SIMULATION	218
6.3.1	<i>Instrument Setup</i>	218
6.3.2	<i>Target Parameter Processing</i>	221
6.3.3	<i>Flyby Parameter Processing</i>	223
6.3.4	<i>Instrument Operations Simulation</i>	235
6.3.5	<i>Data-Cube Generations</i>	238
6.4	DATA VISUALISATION	243
6.4.1	<i>Standard Methods</i>	243
6.4.2	<i>Multicolour Images</i>	245
6.4.3	<i>Interactive Tools</i>	246
6.5	VERIFICATION	247
6.6	SIMULATION RESULTS	249
6.6.1	<i>The Hot Comet Case</i>	249
6.6.2	<i>Cold Nucleus Case</i>	251
6.6.3	<i>Composite Material Case</i>	253
6.6.4	<i>Representative Comet Materials</i>	256
6.7	SURFACE PARAMETER RETRIEVALS	263
6.7.1	<i>The Challenge of Retrieval</i>	263
6.7.2	<i>Proposed Temperature-Emissivity Technique</i>	266
6.8	SIMULATION CONCLUSIONS	269
6.9	FUTURE WORK.....	270
7	CONCLUSIONS AND SUMMARY	273
7.1	OVERVIEW OF RESEARCH	273
7.2	KEY CONTRIBUTIONS AND FINDINGS	274
7.3	FUTURE DIRECTIONS OF RESEARCH.....	276

8	BIBLIOGRAPHY	278
9	APPENDICES	292
9.1	TIRI DATA-CUBE DETAILS	292
9.2	PSG SETTINGS FOR REPRESENTATIVE SPECTRA	295

Table of Acronyms

Acronym	Meaning
ADC	Analog-to-Digital Converter
AU	Astronomical Unit
Bb	Blackbody
CMBR	Cosmic Microwave Background Radiation
CA	Closest Approach
CDHU	Command and Data Handling Unit
CI	Comet Interceptor
CSV	Comma Separated Values
DAC	Digital-to-Analog Converter
DAQ	Data Acquisition
DTGS	Deuterated Triglycine Sulphate
EFM	Engineering Flight Model
ESA	European Space Agency
ESATAN	European Space Thermal Analysis Network
FIR	Far-Infrared
FoV	Field of View
FPGA	Field-Programmable Gate Array
HDO	Deuterated Water
HRTS	High-Resolution Temperature Sensor
HVM ³	High-Resolution Volatiles and Minerals Moon Mapper
IML	Infrared Multilayer Laboratory
IMS	Ion Mass Spectrometer
INL	Integral Non-Linearity
INO	Institut National D'Optique
JAXA	Japan Aerospace Exploration Agency

JCMT	James Clerk Maxwell Telescope
JFC	Jupiter-Family Comet
JPL	Jet Propulsion Laboratory
LAMP	Lyman Alpha Mapping Project
LCROSS	Lunar Crater Observation and Sensing Satellite
LEND	Lunar Exploration Neutron Detector
LPNS	Lunar Prospector Neutron Spectrometer
LRO	Lunar Reconnaissance Orbiter
LTM	Lunar Thermal Mapper
LTB	Lunar Trailblazer
Mb	Megabyte
M ³	Moon Mineralogy Mapper
MIR	Mid-Infrared/ Mid-Infrared Radiometer
MIRMIS	Modular InfraRed Molecules and Ices Sensor
MLI	Multi-Layer Insulation
MPA	Movement Pattern Analysis
NASA	National Aeronautics and Space Administration
NEP	Noise Equivalent Power
NEDT	Noise Equivalent Differential Temperature
NIR	Near-Infrared/ Near-Infrared Radiometer
NMS	Neutral Mass Spectrometer
OH	Hydroxyl
PID	Proportional-Integral-Derivative
PRT	Platinum Resistance Thermometer
PSG	Planetary Spectrum Generator
PSR	Permanently Shadowed Regions
RAL	Rutherford Appleton Laboratory
RT-DLaTGS	Deuterated Lanthanum α Alanine doped TriGlycine Sulphate
SETN	Sensor of Epithermal Neutrons
SNR	Signal-to-Noise Ratio
SOFIA	Stratospheric Observatory for Infrared Astronomy
STL	Stereolithography
TIRI	Thermal InfraRed Imager
URF	User Reference Frame
VTK	Visualization Toolkit
VTT	Technical Research Centre of Finland

Impact of COVID-19 and Mission Delays

This DPhil project commenced in October 2020, during a period when the United Kingdom and the University were under significant restrictions due to the COVID-19 pandemic. As the country entered its second winter of COVID restrictions, remote working, limited access to laboratories, and restrictions on large gatherings remained in place within the University until early 2022. These restrictions had a profound impact on the early stages of this project. The constraints on formal teaching and training, combined with reduced opportunities to engage with more experienced colleagues in the laboratory, significantly influenced the direction of my research. Consequently, this thesis exhibits a stronger emphasis on modelling and computational analysis, as these methods were more feasible under the circumstances. The cancellation and postponement of conferences, both domestically and internationally, limited my opportunities to present my work, exchange ideas, and build connections within the planetary physics community until the latter stages of my DPhil.

In addition to the challenges posed by the pandemic, this project was also affected by delays in mission timelines. At the project's inception, the Lunar Trailblazer mission was scheduled to launch in late 2022, with LTM expected to be delivered to NASA earlier that year and data from the lunar surface anticipated in 2023. However, due to delays and overruns beyond the scope of this project, the delivery of LTM was postponed to 2023, and the Trailblazer will not launch before 2025. These delays necessitated a pivot in the focus of this thesis. Consequently, a significant portion of the work presented here relates to the development of MIRMIS. This work applies lessons learned from the late-stage development of LTM, with the goal of ensuring that MIRMIS delivers superior scientific returns and is equipped with a robust set of tools to optimise its challenging in-flight operation.

1 Introduction

Why is there water on Earth? The processes that led to the Earth's habitable environment, rich in the water essential for sustaining carbon-based life, remain a fundamental puzzle in modern astrophysics and planetary science. Solving this puzzle is a foundational question and a basis for continued Solar System exploration.

All the Solar System's rocky planets exhibit evidence of a significant water presence in their past. Isotopic analysis of the Venusian atmosphere implies the loss of an ancient water ocean (Donahue et al., 1982). Geomorphology suggests an aqueous history for Mars (Baker, 2006), with Neutron spectroscopy (Feldman et al., 2004) and seismic data (Wright et al., 2024) suggesting significant sub-surface deposits still exist. The presence of significant terrestrial water deposits on Earth has been well-documented since ancient times (Pyramid Texts, 2400 BCE), and has been recently reconfirmed by the author through direct ocular inspection. Although the presence of water on Earth is well-documented, the exact origins and mechanisms responsible for its abundance remain topics of active research and debate.

Recent studies suggest that the early materials that accreted to form the Earth were dry, and highly reduced (Rubie et al., 2011; Wood et al., 2008). This implies that Earth's water must have been added either during late accretion or after the planet had largely formed. One major weakness of the direct accretion model arises from the isotopic composition of terrestrial water. If Earth's water had been inherited directly from the planetary disk, it would likely exhibit a Deuterium to Hydrogen (D/H) ratio similar to that of the Solar nebula. However, observations reveal that the water on rocky planets, including Earth, has a D/H ratio significantly higher than the Solar value - approximately six times higher for Earth, eight times for Mars, and six times for Vesta (Hallis, 2017). This

suggests that the water across these bodies likely shares a common source that is distinct from the proto-Solar nebula.

Alternative theories for the late addition of water include the oxidation of a hydrogen-rich atmosphere by iron oxide in the terrestrial magma ocean (Genda & Ikoma, 2008) or by the inward migration of water bearing phyllosilicate dust to early planetesimals (Ciesla & Lauretta, 2005). Each of these theories, however, face challenges and contradictory evidence, although they may have still contributed to some degree (O'Brien et al., 2018).

While the material that formed near 1 AU likely did not contain water, water-bearing compounds were certainly formed at greater distances from the Sun, where lower temperatures prevailed. Numerous simulations have shown that late accretion of water-bearing materials from beyond 2.5 AU is possible, provided that the initial disk of embryos and planetesimals extended beyond the asteroid belt region (Morbidelli et al., 2000; Raymond et al., 2004). Advances in our understanding of early Solar System dynamics, particularly the "Grand Tack" hypothesis, have significantly influenced these simulations, leading to a wide range of possible outcomes depending on the parameters used (Hansen, 2009).

The "Grand Tack" hypothesis suggests that Jupiter, followed by Saturn, migrated inward from its formation point at ~3.5 AU to ~1.5 AU, before reversing its course and travelling outward again. This migration scattered asteroids and comets, truncating the planetesimal disk and potentially facilitating the delivery of water to the inner Solar System. Simulations that account for such dynamics, including the effects of accretion, water delivery, and giant impacts, have produced results where Earth's final water levels are consistent with current estimates (O'Brien et al., 2014). Figure 1-1 displays the results of some of these simulations, showing that the proportion of planets formed that originated in the outer Solar System increases with time. Since the primitive planetesimals

Introduction

from the outer Solar System carried water and other volatiles, and given that all of the curves increase with time, this suggests that the water would have arrived late in the accretion process. By this mechanism, and other travelling Jupiter models, early planetesimal scattering may allow for late accretion of water to the rocky planets.

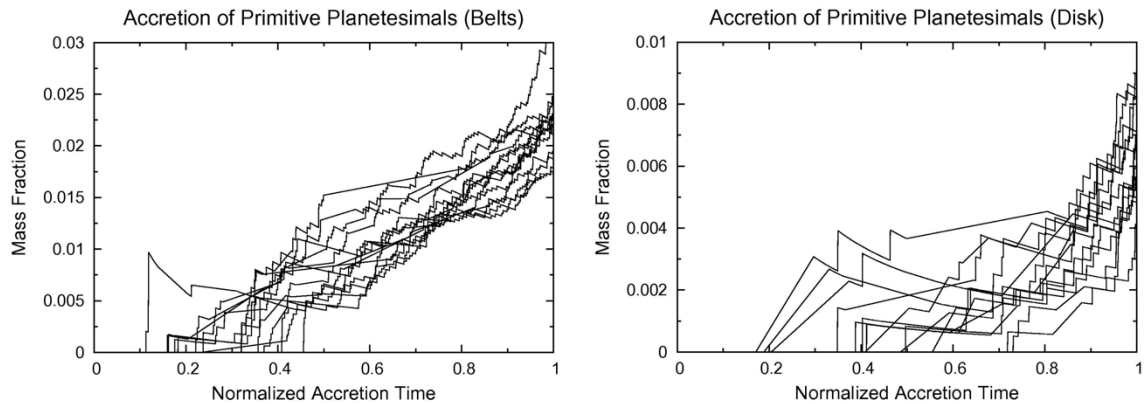


Figure 1-1: The accretion curves for planets larger than 0.75 Earth masses from the Grand Tack simulations of O'Brien et al. The mass fractions of material originally from the belt region (between the orbits of the giant planets) and the disk population (beyond the giant planets) are both plotted. The fractions are plotted against a normalised time value, which is given by the mass of the planet divided by its final mass. (O'Brien et al., 2014)

As a final alternative, comets have long been proposed as potential sources of Earth's water due to their ice-rich composition. One study, based on the lunar impact record and mass scaling laws, concluded that Earth could have accumulated an entirely exogenous ocean between 4.5 and 3.7 billion years ago, provided that 10% of impacts by mass consisted of comets (Chyba, 1987). However, the primary challenge to this hypothesis is the sheer volume of water that would need to be delivered, coupled with the low probability ($\sim 1 \times 10^{-6}$) of any comet from the giant planet region striking Earth. According to the estimates in O'Brien's 2018 review *'The Delivery of Water During Terrestrial Planet Formation'*, if the comets that impacted with Earth were completely ice, they could still only provide 10% of the water present today. Despite these limitations, comets are likely to have contributed some volatiles to Earth's surface in what is known as the "Late Veneer" - a period after the initial formation of the planet when additional material, including volatile-rich comets, may have been added to Earth's crust.

Introduction

It is mostly likely that the presence of terrestrial water originates from a combination of the methods mentioned above. Current scientific interest in the presence of volatile materials in Solar System bodies aims to better constrain and understand these processes. Similar to how the Chyba paper combined lunar impact records with cometary estimates to explore the origins of Earth's water, our approach involves gathering data from a variety of sources and integrating them to develop a more comprehensive understanding of the processes that shaped the Solar System as we observe it today.

This thesis focuses on the development of instrumentation designed to study the thermal and compositional properties of airless bodies, with the goal of understanding how these properties influence the storage and transport of volatiles within the Solar System. Specifically, the work supports the scientific objectives of NASA's Lunar Trailblazer and ESA's Comet Interceptor missions.

Lunar Trailblazer will investigate the presence of volatiles on the lunar surface. By analysing the quantity, distribution, and environmental context of lunar volatiles, the mission aims to shed light on the mechanisms and cycles that govern lunar water, thereby offering insights into the Moon's history. These findings will help us use contemporarily present deposits to better understand the evolution of the lunar surface.

Comet Interceptor, on the other hand, will explore the volatile content, mineralogy, and thermodynamic environment of a dynamically new comet. This first observation of a dynamically new comet will provide a rare opportunity to directly measure materials that have remained unchanged since the early Solar System. The data gathered will offer crucial insights into the structure and formation of comets and early Solar System materials, contributing to our understanding of how water may have been delivered to the inner planets.

In the remainder of this introduction, I will discuss the scientific context and background relevant to these two missions, including the study of water and mineralogy on the lunar surface and the volatile presence in comets.

1.1 Water and Mineralogy on the Moon

Over the past two decades, new research has significantly reshaped our understanding of water and other volatiles on the lunar surface. Earlier results from NASA's Apollo missions and the Soviet Luna robotic spacecraft, which returned samples between 1969 and 1972, led scientists to conclude that the lunar surface was largely devoid of volatiles. This conclusion was based on the absence of minerals containing H₂O and OH in situ and in the returned samples (Haskin & Warren, 1991). However, it had long been hypothesised that water could exist in the subsurface or within permanently shaded regions (PSRs) of the lunar environment, where temperatures remain below 100 K (Watson et al., 1961). These cold traps could accumulate water ice from various sources, including meteor impacts and Solar wind induced production.

In 2009 the Moon Mineralogy Mapper (M³) published results that showed the detection of widespread hydroxyl groups bound to the lunar regolith (Pieters et al., 2009). These results were corroborated by the detection of hydroxyls by the LCROSS spacecraft in the ejecta plume thrown up after the final stage of its launch vehicle was directed into a PSR in the Moon's southern hemisphere (Colaprete et al., 2010). Orbital neutron spectroscopy such as the LEND instrument on LRO provided further evidence for hydrogen in the polar regions of the Moon; however, this was not limited to the low temperature PSRs and thus agreed with the M³ widespread OH measurements. The most recent remote sensing evidence of lunar water was uncovered in October of 2020 when the results of observations of the lunar surface made by the Stratospheric Observatory for Infrared Astronomy (SOFIA) instrument were published. The team analysing the SOFIA data

measured H₂O across the sunlit surface of the Moon, contradicting earlier predictions that lunar water could exclusively be maintained at temperatures below 112 K (Honniball et al., 2020).

With the confirmation that volatiles exist on the lunar surface, key questions remain regarding their abundance, distribution, and the mechanisms that govern their transport across the Moon. Addressing these questions requires advanced orbital remote sensing instrumentation capable of mapping large areas of the lunar surface and providing data from a variety of environments. The Lunar Trailblazer mission aims to tackle these science questions and build on the data gathered by past missions like M³, which was not specifically designed to focus on volatiles. The mission's primary instrument, the High-Resolution Volatiles and Minerals Moon Mapper (HVM³), will measure spectral absorption features of hydroxyl, water, and ice across the 2.6 to 3.6 μm range (Thompson et al., 2020). Complementing HVM³, the Lunar Thermal Mapper (LTM), an infrared radiometer developed at the University of Oxford, will provide simultaneous surface temperature measurements to enable precise thermal correction of the low-energy tail of the Planck function. Additionally, LTM's 11 narrow-bandpass filter channels will facilitate complementary evaluations of the lunar surface's mineral composition (Bowles et al., 2020).

1.1.1 Volatiles in Lunar Surfaces

Numerous attempts were made to measure the hydrated content of lunar samples returned by the Apollo 15 and 17 missions (Friedman et al., 1970; Gibson & Moore, 1973). Early analyses detected very low concentrations of hydrogen (<1 ppm), which were attributed to Solar wind irradiation, while traces of molecular water were dismissed as laboratory contamination. Based on these findings and additional measurements from

other Apollo and Luna samples, it was initially concluded that the lunar surface was composed entirely of dehydrated minerals.

However, since 2008, a series of studies have revisited the question of volatiles in lunar samples using advanced analytical techniques developed since the 1970s. These studies have focused on glass spherules found in Apollo magma samples, revealing the presence of water during their formation. The water content in these glass beads ranges between 260 and 745 ppm, with subsequent research reporting even higher volatile abundances (Saal et al., 2008). These findings suggest that the lunar magma from which these spherules formed contained significant amounts of volatiles, including H₂O.

To estimate the original abundance of volatiles in the lunar magma, researchers have modelled the outgassing processes that occurred during the formation of these spherules. By comparing the distribution of volatiles across the sample spherules, they have inferred that magmatic water, originating from the Moon's interior, has been present since its formation (Friedman et al., 2009; Saal et al., 2008).

1.1.2 Volatiles in PSRs

The size of the Moon and its weak magnetic field mean that over large time scales it is unable to retain a substantial atmosphere. The lunar surface, covered in a fine powder known as regolith, exhibits very low thermal conductivity. This lack of conduction, combined with the absence of atmospheric convection, results in significant temperature gradients across the surface, driven by Solar illumination. These gradients create cold traps - regions where temperatures remain cryogenically low, extending beyond areas of permanent shadow (Paige et al., 2010).

Introduction

The Moon's rotational axis is tilted by only 1.5° relative to the plane of orbit around the Sun, as illustrated in Figure 1-2. This slight tilt causes sunlight to strike the poles at very low angles, leaving certain craters in permanent shadow. These permanently shaded regions (PSRs) have become prime candidates for the accumulation of water ice, as they can maintain temperatures well below 100 K.

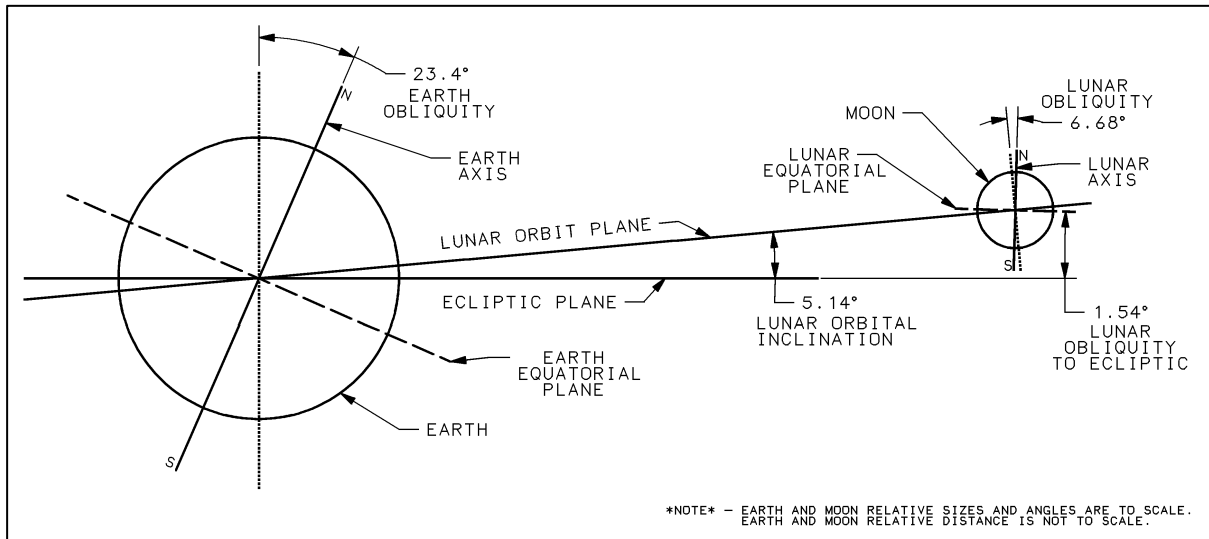


Figure 1-2: The obliquity of the Moon to the elliptic plane is demonstrated here. (Image credit: Peter Sobchak, Wikimedia Commons)

Thermal escape of water molecules occurs when their velocities, described by the Maxwell distribution of molecular kinetic energies, exceed a certain threshold which is defined by the Moon's gravity. The process, known as Jeans escape, involves molecules at the high-energy tail of the distribution reaching velocities greater than the Moon's escape velocity, leading to their loss from the surface. Even at low temperatures, such as those below 100 K, Jeans escape will still occur, though at a greatly reduced rate. For equivalent H_2O densities, the Jeans escape rate at 100 K is 10^{11} times smaller than at 300 K.

Introduction

Calculations using the Jeans escape equations indicate that water ice can be effectively trapped in PSRs, with an escape time of 10^9 years, provided the surface temperature remains below 112 K (Arnold, 1979). Topographical data from the Japanese Kaguya orbiter and NASA's Lunar Reconnaissance Orbiter (LRO) have confirmed the existence of such cold traps in polar craters like Cabeus. By combining these topographical




Figure 1-3: A daytime brightness temperature plot of the lunar south pole as measured by the Diviner Lunar Radiometer. Note the craters which remain in permanent shadow are maintained at temperatures below 112K. Usefully the colouration change to a dark blue coincides with the 112K Jeans Escape limit. The outer circle on the plot represents 80° south latitude (Paige et al., 2010).

data with thermophysical modelling, researchers have predicted that some of these craters maintain temperatures at or below 110 K, making them likely sites for stable water ice deposits (Kozlova & Lazarev, 2010).

Further refinement of these models has aimed to constrain the depth at which stable water could be trapped in the lunar subsurface. An example of this type of model is the Oxford 3D thermophysical model, created by Dr Warren, one of the main advisors to this project, and described in (Warren et al., 2019). This model includes surface shadowing and scattering effects, as well as a discrete subsurface exponential density profile, has been used to identify potential sites where water ice could be accessible within 1 metre of the surface. Figure 3 illustrates the temperature distribution in the southern polar region, as measured by the Diviner Lunar Radiometer, highlighting the PSRs that are potential candidates for water ice deposits. Evidence of water on the Lunar surface comes from neutron spectroscopy, infrared spectroscopy and direct measurements, as will be discussed through the following sections.

1.1.3 Neutron Spectroscopy

Neutron and gamma ray spectrometers are passive instruments that detect neutrons and gamma rays produced by cosmic ray interactions with the lunar regolith. Epithermal neutrons, with energies ranging from 0.4 eV to 1 MeV, are particularly sensitive to hydrogen due to hydrogen's large cross-section at these energies. This sensitivity allows neutron spectrometers to detect hydrogen within the top metre of the lunar subsurface (Hayne et al., 2014).

Neutron spectrometers were deployed in lunar orbit during the Lunar Prospector Mission and later with the Lunar Exploration Neutron Detector (LEND) aboard LRO. The Neutron Spectrometer aboard Lunar Prospector detected increased levels of epithermal neutron emission from the lunar poles, indicating higher hydrogen concentrations -

Introduction

around 150 ppm - compared to 50 ppm in the equatorial regions (Elphic et al., 2000). However, limitations in spectral resolution prevented the identification of specific sources of these neutrons, such as individual craters or PSRs. Figure 1-4 illustrates the hydrogen abundance at the lunar south pole as derived from Lunar Prospector data.




Figure 1-4: Maps of hydrogen abundance for the lunar south pole formed from data from The Neutron Spectrometer on Lunar Prospector. The data shown is for regions poleward of 70° and the plot is taken from (Elphic et al., 2000).

LEND provided global hydrogen mapping at a spatial resolution of 10 km (Sanin et al., 2007). Although LEND does not distinguish between different forms of hydrogen, such as hydroxyls, hydrated minerals, or adsorbed water, its global plots (shown in Figure 1-5) reveal a clear latitudinal variation in hydrogen presence within the lunar subsurface. The epithermal neutron count rate is sharply reduced at the poles, indicating a relative increase in hydrogen, which aligns with the OH presence detected by M³ shown in Figure 1-6.

Neutron spectroscopy is fundamentally limited in spatial resolution due to the low count rates involved. The data recorded by LEND was formed by combining locally

coincidental measurements from multiple orbits to build up a significant signal to noise ratio.

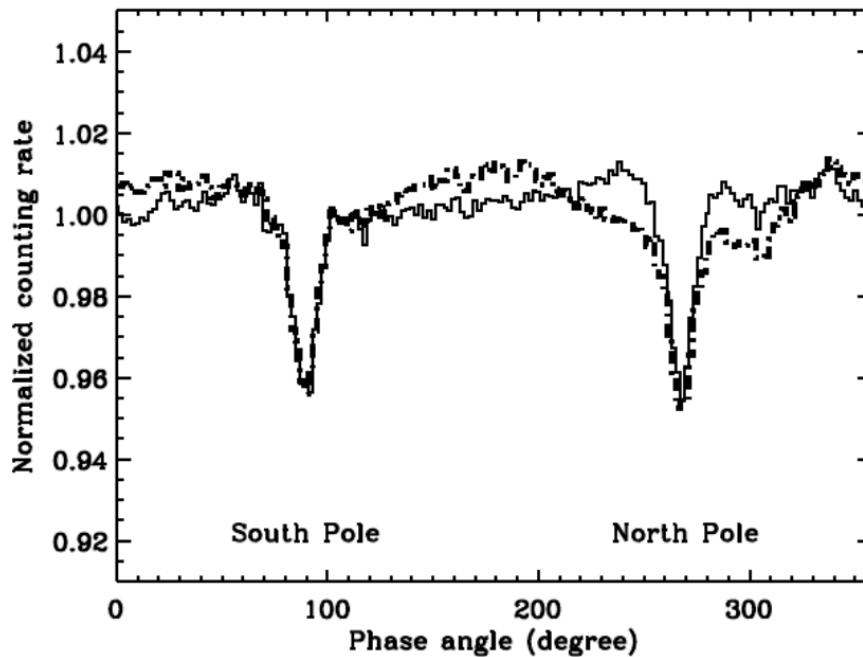


Figure 1-5: Corrected orbital phase curves showing regional latitude variations of epithermal neutron flux across the lunar surface measured by LEND omnidirectional SETN (dotted line) and LPNS (solid line) detectors. The x-axis is the phase angle measured in degrees from 0° to 360° , where 90° corresponds to the south pole and 270° to the north pole. The y-axis represents the normalised counting rate of epithermal neutrons. Figure taken from (Litvak et al., 2012).

1.1.4 M³ Measurements of Lunar Water

The measurement of an H₂O/OH absorption feature over much of the Moon's surface by M³ in 2009 was a surprise (Pieters et al., 2009; Schultz et al., 2010). Although previous research, such as that by Saal and colleagues, had detected water in magmatic glasses, it was not anticipated that hydroxyls would exist on the surface, especially outside of permanently shaded regions (PSRs). This expectation was based on the assumption that any water or hydroxyls exposed to sunlight would reach temperatures high enough to enable escape through thermal processes.

While there was a latitudinal dependence for the presence of the H₂O/OH feature, the molecules being measured are evidently present on the illuminated surface in polar regions rather than being limited to PSRs. A global plot of M³ data is reproduced in Figure

Introduction

1-6. The blue water/hydroxyl feature extends from the poles down to latitudes of 60°. Cross referencing against the Diviner temperature plot in Figure 1-3 reveals that the 80° latitude temperatures were reaching over 200K, well above the Jeans escape limit of 112K. These two observations demonstrate that the form of water/ hydroxyls on the lunar surface is not limited to water ice in PSRs. The measured signature is likely to originate from a very thin (few microns) layer in the upper surface as the OH 3 μm band is very strong.




Figure 1-6: Combination of spectral images taken by M³. Blue regions show strong water/hydroxyl absorption features, red colouring denotes pyroxene features and green shows the brightness of the surface. (Pieters et al., 2009).




Figure 1-7: The imaginary index values showing the absorption features of water at 3 and 6 μm . Plot reproduced from (Lucey, 2017).

The spectral feature observed by M³ was at 3 μm and the absorption features of water are shown in Figure 1-7. Since this was at the long-wavelength limit of the instrument's spectral range. The instrument could not measure the entire width of the feature since it was not designed to look for water. This is because the moon was believed to be dry. Furthermore, this is in the range where the thermal emission of the sunlit regions of the Moon's surface interferes with the measurements. Accurate mapping of absorption features requires precise temperature measurements of the surface as it is being observed, so that thermal emission can be removed from the spectrum. Since M³ had not been optimised for measuring these features it was limited in its capability to correct for the thermal emission in spectrum processing. The method applied to correct for emission led to varying interpretations of the 3 μm feature, with new empirical methods developed to utilise Diviner data to inform the correction (Li & Milliken, 2016). Most recently in 2020 research has been conducted to provide independent verification of thermal correction methodologies to establish which most accurately describe the emission at 3 μm (Tai Udovicic et al., 2020).

The Lunar Trailblazer mission aims to repeat these measurements with a much improved spectral and spatial resolution, full spectral coverage of absorption features, and

a dedicated secondary instrument for measuring the surface temperature enabling accurate thermal correction.

1.1.5 LCROSS Ejecta Measurements

Shortly after the publication of the M³ data NASA launched the Lunar Crater Observation and Sensing Satellite (LCROSS) to explore the presence of water ice in a lunar polar PSR (Schultz et al., 2010). In this experiment the upper stage of the launch vehicle was directed to impact in Cabeus Crater, a permanently shadowed crater near the lunar South Pole. The LCROSS spacecraft then passed through the ejecta plume launched into the exosphere by the impact, where its instruments could measure and monitor the constituents of the ejecta.

The results from the LCROSS mission confirmed the presence of molecular water in the ejecta. Comprehensive analysis of the mission data estimated that the lunar regolith at the impact site contained approximately $5.6 \pm 2.9\%$ water ice by weight, along with smaller quantities of other volatile compounds, such as hydrogen sulphide, carbon dioxide, and ammonia (Gladstone et al., 2010). These findings provided direct evidence of water ice in a PSR, supporting the hypothesis that such regions serve as cold traps capable of preserving volatiles over geologic timescales. However, the σ value for this measurement is at 1.93, and therefore the confidence level for these water-ice measurements is below 95%.

1.1.6 SOFIA Measurements

In 2020, the Stratospheric Observatory for Infrared Astronomy (SOFIA) made a groundbreaking remote measurement of volatiles on the lunar surface. For the first time, observations of a sunlit region at high southern latitudes detected spectral features at 6 μm , a characteristic resonance of H₂O molecules (Honniball et al., 2020). This novel use of

the 6 μm feature for identifying water on the lunar surface revealed water concentrations in the sunlit regolith ranging from 100 to 412 parts per million - higher than previously reported measurements. Identification of volatile specimens across the lunar surface by HVM³ will be able to corroborate this measurement (Thompson et al., 2020).

1.1.7 Chang'E 5 Results

The most recent measurements of lunar surface volatiles come from China's Chang'E-5 mission, which landed in the lunar northeast Oceanus Procellarum region in late 2020. The mission conducted in-situ spectral measurements and returned samples to Earth for laboratory analysis. The in-situ spectral findings indicated an average hydroxyl concentration of 28.5 ppm in the regolith, with returned samples showing a range from 0 to 179 ± 13 ppm (Liu et al., 2022). The corresponding reflectance spectra, shown in Figure 8, clearly display the absorption feature slightly short of 3 μm , which is indicative of hydroxyl presence (further discussed in Section 2.1.1).

Detailed analysis identified the hydroxyl-containing mineral apatite as a significant source of these hydroxyls, leading researchers to conclude that this mineral may contribute to the presence of hydroxyl across young mare regions of the lunar surface. This discovery highlights the relationship between mineralogy and the long-term stability of volatiles on the Moon.




Figure 1-8: The 2700-3100 nm continuum removed absorption features of separate surface regions measured by the on-board spectrometer on Chang'E-5. Plot reproduced from (Liu et al., 2022).

These findings from Chang'E-5 provide a valuable framework for future exploration. They suggest that a more comprehensive dataset, combining mineralogical and volatile surface maps, could significantly enhance our understanding of volatile distribution across the Moon. Producing such a dataset is a primary scientific goal of the Lunar Trailblazer mission.

1.2 Volatiles in Comets and Asteroids

The composition of comets is much less well understood than that of the lunar surface due to the challenges posed by their location in deep space. The high velocities and extreme trajectories of comets make extended missions challenging due to late warning of new comets, and the relative scarcity of these bodies means that missions are typically limited to studying only a few comets or asteroids. This contrasts sharply with lunar missions, where proximity to Earth allows for years of continuous data collection, enabling

comprehensive mapping of the lunar surface. Despite some similarities in the surfaces of these bodies, the study of comets requires a different approach, with missions governed by specific operational constraints and challenges.

As seen in Section 1, ice-rich comets may have played a crucial role in delivering water and other volatiles to the rocky planets of the inner Solar System. These exogenous sources could account for a significant portion of the volatile content observed on Earth and other terrestrial planets. To determine the extent to which comets and asteroids contributed to Earth's water, it is essential to make direct measurements of the volatiles present on these bodies as such measurements enable isotopic analysis, which can inform the origins and history of these volatiles.

Remote measurements of volatile presence, when combined with additional data on shape, composition, and temperature, provide critical insights into the conditions under which volatiles can survive on different Solar System bodies (Guilbert-Lepoutre A. , 2012). These observations allow us to develop models that link the observed volatile quantities to their evolutionary history since the early Solar System. In this section, I will describe the current knowledge of volatiles on the surfaces of small Solar System objects and discuss the additional parameters - such as surface composition, albedo, emissivity, and thermal properties - that are used to characterise their volatile retention and release behaviour.

1.2.1 Evolution of Comet Knowledge

Our understanding of comets has evolved significantly over the centuries. Early observation by the ancient Greeks were explained through different theories. Pythagoras hypothesised that they were wandering planets while Aristotle considered them to be atmospheric phenomena. More scientific investigations began with the observations of Tycho Brahe in the late 16th century, in which he utilised parallax to prove that comets were more distant than the Moon and thus must be celestial bodies travelling far from

Introduction

Earth (Pater & Lissauer, 2015). The hypothesis that comets are solid bodies composed of volatile materials emerged in the 18th century, most notably with Edmond Halley's proposal that the tails of comets are formed by the vaporization of volatile substances as they approach the Sun (Halley, 1705). In the same work, Halley used Newton's gravitational laws to describe the parabolic orbits of 24 comets observed over the preceding 350 years and correctly identified that three of these comets were the same returning body - now known as Halley's Comet, which shall be discussed further later.

The idea of a volatile nucleus gained traction over time, and confirmation came with the first spectroscopic measurements of cometary comae by William Huggins in 1866, which identified spectral lines corresponding to carbon and nitrogen gases (Huggins, 1866). This concept was formalised in the mid-20th century with Fred Whipple's "Dirty Conglomerate" model, which proposed that a comet's nucleus consists of volatile ices mixed with dust and rock. As the comet approaches the Sun, the heat causes these volatiles to sublimate, releasing dust and gases that form the coma and the characteristic tail pointing away from the Sun (Whipple, 1949, 1950). Whipple's model was motivated by a desire to understand the impact of non-gravitational forces on cometary orbits as a method of understanding the deviations seen from perfect ellipses.

The 'Dirty Snowball' model was validated in 1986 when ESA's Giotto spacecraft conducted the first flyby of a comet, directly observing the nucleus of Halley's Comet (Figure 1-9). The spacecraft captured images of bright jets emanating from the nucleus and collected compositional data, identifying water vapour, carbon dioxide, carbon monoxide, and other volatiles in the jets (Eberhardt et al., 1986). While Giotto's instruments could not fully penetrate the nucleus, the mission revealed that Halley's Comet had a much lower albedo than expected, with a reflectance of just 0.04. This was the first comet flyby mission, and its success inspired similar missions to directly measure comets, taking advantage of new technology as it became available.




Figure 1-9: The nucleus of the Comet Halley, as observed by the Halley Multicolour Camera on the Giotto spacecraft. The frame size is 30km. In the centre can be seen two bright jets pointing directly towards the Sun which is to the left of the image. Image from (ESA, 2012).

In 2005, NASA's Deep Impact mission further advanced our understanding by sending an impactor into the nucleus of Comet 9P/Tempel 1, creating a crater and analysing the resulting ejecta. The mission detected a significant presence of water ice, which sublimated rapidly, and observed an increase in organic molecules, suggesting that such organics may exist below the surface, protected by layers of fine dust (A'Hearn et al., 2005). The thermal properties and structural findings from this mission are discussed in more detail in Section 1.2.3.

The most detailed examination of a comet to date was provided by ESA's Rosetta mission, which orbited the Kuiper belt Comet 67P/Churyumov-Gerasimenko from 2014 to

2016. Rosetta returned highly detailed morphological data, revealing a much more complex surface than expected (Thomas et al., 2015). The mission also identified significant diversity in the comet's ice composition, detecting nearly all the molecules known to exist in cometary comae (Le Roy et al., 2015). Notably, Rosetta measured the deuterium-to-hydrogen (D/H) ratio in the comet's water, providing critical insights into the origins of water in the Solar System (Altwegg et al., 2014). The surface of 67P was observed to substantially vary in hardness. While some areas were mechanically hard, the landing point exhibited a softer surface than expected. These findings have modified our expectations of a "Dirty Conglomerate" formed of a denser mixture of rock and ice, which has implications for the thermal conductivity of the material. The data collected by Rosetta has greatly enhanced our understanding of cometary structure and composition, as we will explore further in the following sections.

1.2.2 Isotopic Ratios for Water Origination

As discussed earlier, one of the key methods used to trace the origins of water and other volatiles in the Solar System is the measurement of isotopic ratios, particularly the Deuterium-to-Hydrogen (D/H) ratio. By distinguishing between regular water (H_2O) and deuterated water (HDO), scientists can measure the relative abundance of Deuterium and Hydrogen. The D/H ratio is crucial because it reveals isotopic fractionation - the degree to which different isotopes are separated or concentrated. This ratio is particularly sensitive to the temperature of the environment where the water originally formed, with lower temperatures resulting in higher D/H ratios. As a result, the D/H ratio can provide insights into the temperature and location within the Solar accretion disc where the water was formed. Similar D/H ratios in water found in different regions of the Solar System may imply a common origin.

Introduction

After water is formed, the D/H ratio can still be altered by processes such as evaporation into space, which preferentially removes lighter H₂O molecules and enriches the remaining water in Deuterium. This process is one explanation for the elevated D/H ratio observed in the Venusian atmosphere (Donahue et al., 1982). While other processes, such as hydrothermal, geochemical, and biological activities, can also affect isotopic fractionation, these are generally not relevant for water-ice within most Solar System bodies. The D/H ratio in such environments is relatively stable, with any changes being gradual and driven by identifiable processes.

To date, only two comets have had their volatile content directly measured via mass spectrometry, providing crucial data on H₂O isotopic ratios. The D/H ratio in comet 1P/Halley was characterised using two mass spectrometers aboard the Giotto spacecraft during its 1986 flyby. The Ion Mass Spectrometer (IMS) recorded a D/H ratio of $3.08 - 0.53 + 0.38 \times 10^{-4}$ in H₂O across a wide distance range from the comet nucleus (Balsiger et al., 1995). Independently, the Neutral Mass Spectrometer (NMS) measured the D/H ratio in the coma to be $3.02 \pm 0.22 \times 10^{-4}$. The combined analysis from both instruments suggested a best estimate of $3.16 \pm 0.34 \times 10^{-4}$ for the D/H ratio in Halley's nucleus water and water-ice (Eberhardt et al., 1995). This value, approximately twice that of terrestrial seawater, suggests that Oort cloud comets, like Halley, are unlikely to be the dominant source of Earth's water (Drake & Righter, 2002; Meier & Owen, 1999).

More recently, the Rosetta mission's Rosina mass spectrometer measured the water vapor surrounding comet 67P/Churyumov-Gerasimenko. The D/H ratio obtained was significantly higher, at $5.3 \pm 0.7 \times 10^{-4}$, more than three times the Earth's ratio (Altwegg et al., 2014). This elevated D/H ratio in a Jupiter-family comet (JFC) underscores the isotopic diversity among comets, suggesting that different cometary families may have distinct formation histories.

In addition to in-situ measurements, the D/H ratio in comets has also been studied using spectroscopy with modern ground-based and space-based telescopes, such as the Caltech Submillimeter Observatory (CSO) and the James Clerk Maxwell Telescope (JCMT). These high-resolution spectrometers can detect and differentiate the specific rotational and vibrational transitions of H₂O and HDO in the comet coma, allowing scientists to estimate the relative abundance of each molecule based on the strength of these spectral lines. However, these measurements have lower signal-to-noise compared to direct measurements, and require assumptions about the coma's structure and dynamics, making in-situ measurements critical for verification and refinement. Table 1-1 presents a selection of measured D/H ratios from spectroscopic studies, showing that, with the exception of Comet Hartley 2, most comets exhibit D/H ratios higher than that of Earth.

Comet Target	Telescope Used	D/H Ratio (x10 ⁻⁴)	Reference
C/1996 B2 (Hyakutake)	JCMT	2.9 ± 1	(Bockelée-Morvan et al., 1998)
C/1995 O1 (Hale-Bopp)	JCMT	3.3 ± 0.8	(Meier et al., 1998)
8P/Tuttle	Very Large Telescope	4.09 ± 1.45	(Villaneuva et al., 2009)
103P/Hartley 2	Herschel Space Observatory	1.61 ± 0.24	(Hartogh et al., 2011)
C/2009 P1 (Garradd)	Herschel Space Observatory	2.06 ± 0.22	(Bockelée-Morvan et al., 2012)

Table 1-1: A selection of measured D/H ratios for comets produced by spectroscopic examination. For reference Earth has a D/H ratio of 1.56×10^{-4} .

The full set of retrieved D/H ratio data for the Solar System is shown in Figure 1-10. This expanded dataset reveals that either chondritic meteorites from the asteroid belt or Jupiter-family comets could be a match for terrestrial water, whereas Oort cloud comets consistently exhibit D/H ratios that are too high in deuterium to be considered a primary source of Earth's water.

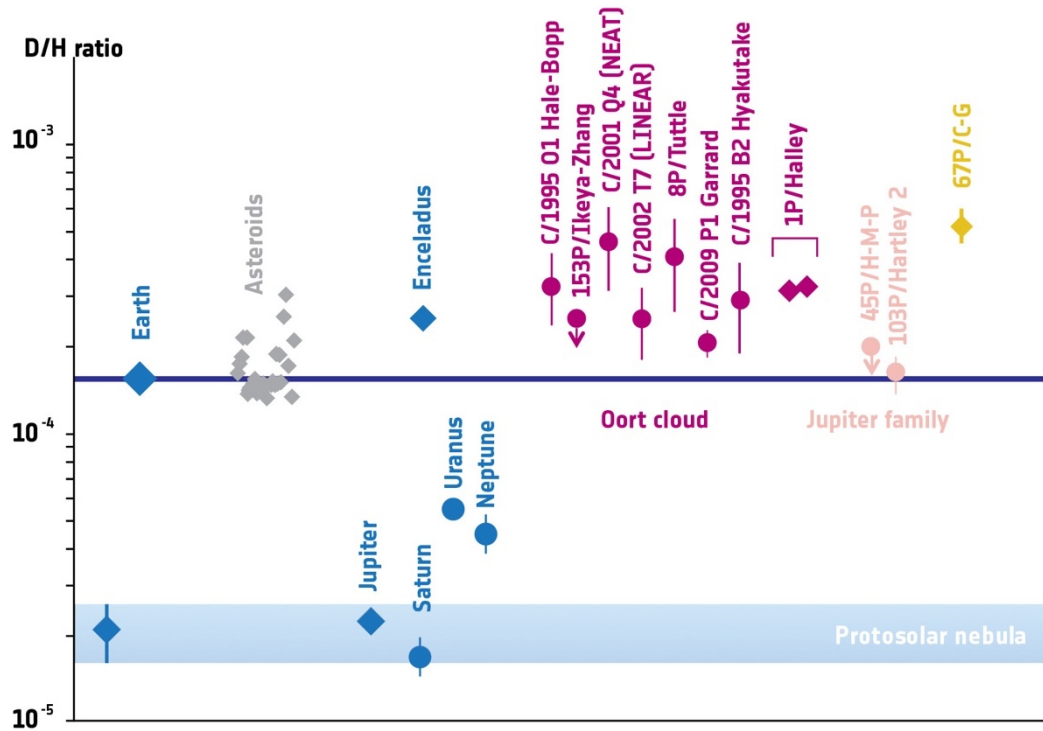


Figure 1-10: The different values of D/H ratios in water measured at different Solar System bodies. The figure is taken from the ESA website (ESA, 2019) produced with the data from (Altwegg et al., 2014) and the references therein. D/H ratios are plotted vertically on a logarithmic axis. The data is grouped by colour with planets and moons in blue, chondritic meteorites from the asteroid belt in grey, comets from the Oort cloud purple and JFCs in pink. Diamonds represent in situ measurements and circles are remotely sensed.

The goal of the Comet Interceptor mission is to visit a dynamically new comet, one making its first transit through the inner Solar System. Such a comet will not have been exposed to the intense heating and radiation effects of proximity to the Sun, thereby serving as a pristine example of material from the early Solar System. By adding a new data point to the D/H ratio dataset illustrated in Figure 1-10, the Comet Interceptor mission will provide critical insights into the extent to which space weathering, along with the heating and cooling cycles experienced during multiple Solar flybys, contributes to the elevated D/H ratios observed in comets studied to date. While it cannot measure isotopic ratios, MIRMIS will provide additional environmental data to support the analysis, for example identifying the quantity of water ice present within the observed nucleus.

1.2.3 The Thermal Environment of the Nucleus

To fully understand the behaviour of volatiles in comets, it is crucial to study their thermal environments, as temperature and heat transfer directly impact sublimation, activity levels, and the stability of the nucleus. In this section, I will discuss the prevailing theories about cometary surfaces based on measurements from previous missions and explore how these insights contribute to our understanding of comet evolution.

Albedo, as a measure of a surface's reflectivity, plays a key role in determining how much sunlight is absorbed or reflected by a comet. Bond albedo measures the fraction of the total Solar energy reflected by a comet over all wavelengths and angles, while geometric albedo quantifies the brightness of the surface at zero phase angle relative to a perfectly diffusing disk. Both parameters are critical in understanding a comet's energy balance and thermal behaviour. A higher albedo in either sense indicates that more sunlight is reflected, which tends to cool the surface and reduce the sublimation rate of volatiles. Conversely, a lower albedo means more absorption of Solar energy, resulting in warmer surface temperatures and enhanced sublimation.

The surface composition of a comet plays a significant role in determining these albedo values. Icy materials generally exhibit high albedo, whereas dust and rock are associated with lower albedo values. Consequently, a rock-ice mixture on a comet may display a wide range of albedo values, leading to diverse thermal behaviours. For instance, Comet Halley, as observed by ESA's Giotto spacecraft, showed a very low geometric albedo of 0.04 due to a surface layer dominated by dust. Similarly, Comet Tempel 1 had a geometric albedo of .04; however, impact observations shed that this low reflectivity was caused by a thin dust layer covering more reflective volatile minerals underneath which rapidly sublimated once exposed (A'Hearn et al., 2005). This suggests that the dust layer acts as a protective barrier, shielding the underlying ice from sublimation until it is

exposed. In the case of Comet 67P/Churyumov-Gerasimenko, studied by the Rosetta mission, a slightly higher geometric albedo of 0.06 was recorded, though it still featured significant dust coverage (Fornasier et al., 2015). It is important to note that both Halley and Tempel 1 have been processed by repeat passages near the Sun, which will have altered their surface properties. One of the key objectives of the Comet Interceptor mission is to assess whether a dynamically new comet, one that remains relatively unaltered, exhibits similar albedo characteristics, shedding light on the relationship between surface composition and cometary thermal processes.

A dark surface layer that absorbs sunlight will heat the comet's surface significantly. However, the effect of this heating depends on the thermal conductivity of the surface material. If the surface material has low thermal conductivity, similar to the lunar regolith discussed in Section 1.1.2, heat transfer into the nucleus will be limited, causing the surface to heat up while the interior remains cool. This localised heating can drive sublimation of volatiles from the surface layer, but the extent of sublimation depends on how much heat is transferred inward. For instance, the thermal map of Tempel 1, shown in Figure 1-11, reveals extreme temperature variations across the comet's illuminated regions, with temperature differences of up to 50 K due to shadows on the surface.

Surface morphology adds further complexity. The irregular shape of 67P/Churyumov-Gerasimenko, for example, creates varied surface textures, ranging from smooth to rough. Rough surfaces can trap heat by reflecting and reabsorbing emitted radiation, while also producing shaded regions. This texture variation can lead to localised "hot spots" where temperatures exceed the surrounding surface, potentially triggering the sublimation of underlying volatiles. Once exposed to sunlight, these volatiles may rapidly sublimate, creating a feedback loop that drives the formation of jets from the comet's nucleus. The Rosetta mission identified such jets on 67P, originating from active regions with smooth surfaces, exposed water ice, and high insolation (Lara et al., 2015). These

factors contributed to higher temperatures, increased activity, and more erosion, producing the anisotropic dust jets observed.




Figure 1-11: This figure displays the temperature maps of the surface of Tempel1 as taken by the infrared spectrometer on Deep Impact. The temperature scale in the centre of the figure is in Kelvin (A'Hearn et al., 2005).

In the absence of direct measurements via sample return, the parameters at the surface of the comets can be defined through a combination of modelling and observation. Surface temperatures maps like the one in Figure 1-11 can be used as 'ground truths' for simulation work. Once a complex simulation of the comet surface temperature which accounts for the real geometry and morphology has been created, surface parameters can be varied to produce results that match the measurement. By this process the thermal inertia of Tempel1 was constrained from the temperatures measured via infrared measurement by Deep Impact (Davidsson et al, 2007). There is further discussion on the specifics of thermal modelling of comet surfaces in Chapter 5 of this thesis. Informing

parameter retrieval modelling drives the requirements on spatial resolution and thermal accuracy of temperature maps for flyby missions, since ground and space-based observations can record bulk temperatures for comet nuclei but lack the relation to shape models.

Understanding the thermal environment of a comet is closely linked to the behaviour of its volatile content, and these two aspects must be studied together (Guilbert-Lepoutre et al., 2015). Capturing accurate thermal maps alongside compositional data that can identify volatiles and their associated minerals is a primary scientific goal for the MIRMIS instrument on Comet Interceptor. The work I have undertaken to support the development of MIRMIS is detailed in Chapters 5 and 6.

2 Measuring the Lunar Surface

In this chapter, I will delve into the process of conducting measurements across the lunar surface via remote sensing, with a particular emphasis on the challenges associated with infrared remote sensing. I will examine two historical examples that have significantly contributed to our current understanding of the lunar surface: the Moon Mineralogy Mapper (M^3) and the Diviner Lunar Radiometer. These missions laid the groundwork for their next-generation counterparts, the High-Resolution Volatiles and Minerals Moon Mapper (HVM³) and the Lunar Thermal Mapper (LTM), both of which are set to fly on NASA's Lunar Trailblazer mission. I will outline the science goals of this mission and the specific roles of each instrument, with a particular focus on the Lunar Thermal Mapper, whose operations this project aims to support. The chapter will conclude by establishing the performance requirements that the LTM must meet and how my work seeks to ensure these requirements are achieved.

2.1 Methodology of Remote Sensing

2.1.1 Remote Sensing of Volatiles

To understand the form and abundance of volatiles across the lunar surface, it is essential to map these volatiles on a large-scale using satellite-based instrumentation, a process known as remote sensing. Water and hydroxyl can be inferred from absorption features observed in reflected near-infrared spectra between 2.6 and 3.6 μm . However, a significant challenge in remote sensing is the lack of knowledge about surface material parameters, such as particle size distribution and porosity, which limits our ability to accurately model radiation scattering.

To address these uncertainties, researchers have employed a combination of radiative transfer modelling and laboratory experiments to identify the parameters that best predict water content as detailed in (Milliken & Mustard, 2005) and (Milliken & Mustard, 2007). These studies have shown that the band parameters correlating the strength of the 3 μm absorption feature with water content increase with particle size, while overall reflectance tends to increase as particle size decreases. Such findings are crucial for the analysis of orbital spectral data.

Despite these advancements, imaging spectroscopy within permanently shadowed regions (PSRs) of the Moon remains challenging due to the limited light available for detection. Previous instruments lacked the sensitivity to capture data based on the minimal reflected light in these areas. However, recent data from ShadowCam, an imager on the Korean Pathfinder Lunar Orbiter, has provided high-resolution imaging of the interiors of lunar PSRs. This data has been used to refine models of secondary illumination within PSRs, thereby enhancing our understanding of their metre-scale morphology (Mahanti et al., 2023). This new dataset and the insights it generates will serve as a crucial reference for analysing the thermal mapping of PSRs that will be produced by Lunar Trailblazer once it becomes operational.

A significant challenge in measuring OH/H₂O features on the Moon is the need for thermal correction. The absorption features for hydroxyl, water, and water-ice overlap in the 3 μm region, as shown in Figure 2-1, and represent several different resonances. The band depth of the OH- feature at 2.8 μm indicates the presence of water and/or hydroxyls. By combining this measurement with the O-H stretch band at 3.0 μm , it is possible to distinguish between water and hydroxyl (Klima & Petro, 2017). However, a key limitation is obtaining accurate reflectance measurements in this wavelength range due to interference from thermal emission. As the surface temperature increases, the blackbody radiance spectrum shifts towards shorter wavelengths, as governed by the Planck curve.




Figure 2-1: The reflectance spectra for hydroxyl, water, and ice, showing their absorption features around the 3000 nm (3 μ m) region. The unshaded section to the left of the 3000 nm mark indicates the limit of the M3 spectrometer's range, while the HVM³ will extend measurements to 3600 nm. (Clark et al., 2016).

This effect is illustrated in Figure 2-2, where radiance spectra between 2 and 3 μ m are plotted for various temperatures. At higher temperatures, the short-wavelength tail of the blackbody emission spectrum can become comparable to the reflected light at 3 μ m, complicating the extraction of the absorption feature's precise shape in regions of direct Solar illumination.

Thermal corrections are possible if the surface temperature is accurately known, as will be discussed further in Section 2.2.2 with respect to M³ data. It is important to note that this spectroscopy technique only probes the upper few millimetres of the surface regolith, and non-uniform heating across the upper regolith surface means that thermal corrections are non-trivial to implement.

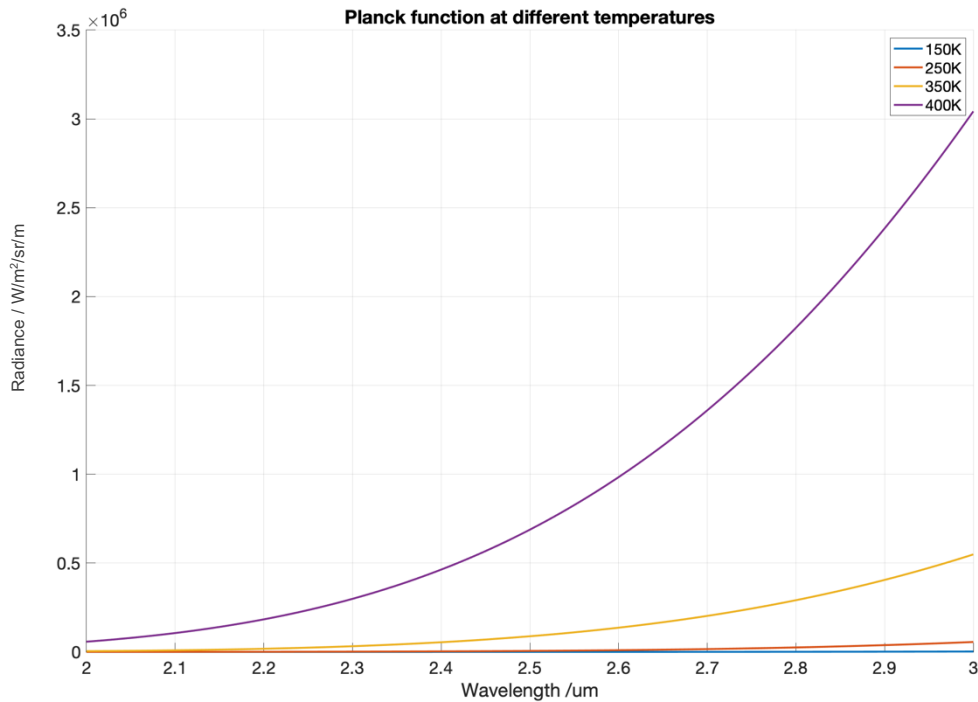


Figure 2-2: A plot of the Planck function between 2 and 3 μm at different temperatures. The low-energy tail of the Planck function increasingly influences measurements within this wavelength range as temperatures rise. Thermal correction algorithms aim to estimate and remove this radiance from instrument measurements.

2.1.2 Compositional Remote Sensing

Remote sensing can also be used to establish the composition of surface materials. A key feature used for compositional analysis in the mid-infrared region (3 to 15 micrometres) is known as the Christiansen Feature (CF). This is a peak in the emissivity of a material, or the nadir in the reflectance spectrum, which occurs in the mid-infrared (5-25 μm). The wavelength position of the CF relates to silicate polymerisation, and it allows an understanding of the bulk composition of a material (Salisbury et al., 1997). The emissivity spectra of two Apollo soil samples have been reproduced in Figure 2-3 to demonstrate the use of the CF position as a tool for compositional analysis.

Analysis of the emissivity spectra reveals that the maximum of the Ap10084 spectrum is found at $8.39 \pm 0.01 \mu\text{m}$ and the maximum of the Ap66031 spectrum is at $8.09 \pm 0.01 \mu\text{m}$. Laboratory analysis of these Apollo samples, which are Lunar soil which has been returned to Earth, has identified that Ap10084 is a mafic rock, an igneous rock with a silica content between 45% and 55%, while Ap66031 is a felsic rock, characterised by a

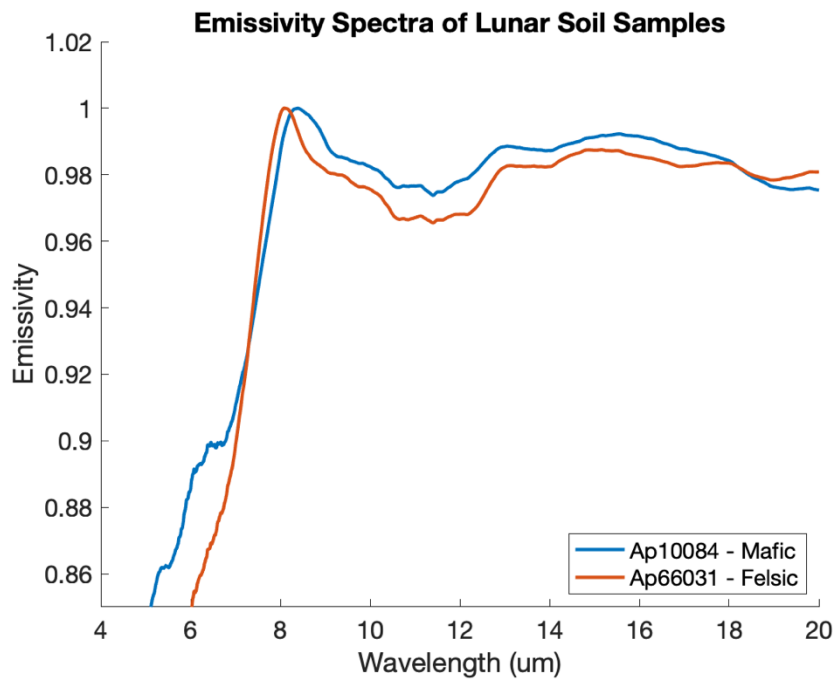


Figure 2-3: Plot of the emissivity spectra of two returned Apollo soil samples. The spectra have each been normalised against their maximum values. The Christiansen Feature is the peak in emissivity observed in both graphs near 8 μm . Data from (Donaldson-Hanna et al., 2014).

silica content exceeding 65%. The correlation between CF position and silica content demonstrates how remote sensing can be used to infer the composition of lunar surface materials from orbit.

The emissivity spectra of lunar soil samples are catalogued in reference libraries, providing a crucial resource for interpreting remote sensing data. Emissivity is typically measured via reflectance measurements in controlled environments and calculated via Kirchoff's law. Future missions can compare their orbital measurements of the CF to these reference spectra, allowing for the identification of the lunar surface's composition remotely. However, an additional challenge in replicating these laboratory measurements is the need to accurately simulate the lunar environment, particularly in achieving realistic thermal gradients in the upper layers of regolith (Donaldson-Hanna et al., 2014).

Remote sensing is also valuable for estimating the exposure age of surface materials on the Moon. The reflectance of the lunar surface changes over time due to space weathering, and this effect varies across different wavelength regions. Specifically, in the 2-

5 μm and 8-14 μm ranges, the lunar surface reflectance decreases with age, while in the 5-8 μm range, the albedo, or reflectivity, increases as the soil ages (Salisbury et al., 1997).

Understanding these spectral changes helps in interpreting the age and weathering history of lunar regolith through remote sensing.

2.2 Previous Lunar Measurements

2.2.1 The Diviner Lunar Radiometer

The Diviner Lunar Radiometer Experiment (shown in Figure 2-4) is an infrared radiometer aboard NASA's Lunar Reconnaissance Orbiter (LRO). In orbit around the Moon since 2009, it is a 'push broom' mapping radiometer which includes 9 filter channels for surface temperature sensing and compositional analysis. The Diviner filter channels are detailed in Table 2-1.




Figure 2-4: The Diviner Lunar Radiometer Experiment with major components labelled. (Paige et al., 2010).

The first two channels measure Solar radiation reflected from the lunar surface. Channels 3, 4, and 5 are compositional channels used to calculate the CF position. The final four channels, 6, 7, 8, and 9, are broadband thermal channels used to map the bolometric surface temperature profile of the entire lunar surface. This bolometric temperature is the product of the Diviner calibration process and represents the brightness temperature of a blackbody that would emit the same amount of radiation as observed in the instrument (assuming a unity emissivity). Black-body brightness temperatures vary across channels since the underlying emissivity is not constant, and the bolometric temperature is the product of a consistent weighted average process across these channel measurements (Paige et al., 2010).

Channel Number	Channel Name	Channel Bandpass (μm)	Purpose
1	A1	0.35 – 2.8	High Sensitivity Solar
2	A2	0.35 – 2.8	Reduced Sensitivity Solar
3	A3	7.55 – 8.05	Compositional
4	A4	8.1 – 8.4	Compositional
5	A5	8.4 – 8.7	Compositional
6	A6	13 – 23	Thermal
7	B1	25 – 41	Thermal
8	B2	50 – 100	Thermal
9	B3	100 – 300	Thermal

Table 2-1: Bandpass details for Diviner. Channels 3-5 are compositional channels for CF mapping and Channels 6-9 are thermal for temperature mapping. (Paige et al., 2010).

In each thermal band, the radiance is recorded, and the corresponding bolometric surface temperature is calculated. The thermal channels are spread across regions in order to be most sensitive at different temperature ranges, with longer wavelength channels being more sensitive to lower temperatures. All the channels operate by recording the power incident on the detector and calibrating this against space views and a well-known blackbody target inside the instrument to derive the radiance emitted from the surface.

To determine the CF value, the emissivity data from the three compositional channels are fitted with a second-order polynomial. The maximum of this parabola is taken as the CF value for that specific location on the lunar surface, as demonstrated in Figure 2-5 (Greenhagen & Paige, 2006, 2009). The CF values derived through this method are good fit for minerals where the CF feature lies within the range of Diviner’s 3 compositional channel, however accuracy falls off as the CF value moves away from this wavelength range. LTM has a broader set of compositional channels and so will maintain accuracy over a wider wavelength range.




Figure 2-5: Diviner compositional three-channel spectra for two apollo samples and a plagioclase from the Inner Rook Mountains. Fitted quadratic are shown to demonstrate the calculation of CF value from the three data points. Plot taken from (Greenhagen et al., 2010).

2.2.2 The Moon Minerology Mapper

The Moon Minerology Mapper (M^3) is an imaging spectrometer which was the first to identify pervasive hydroxyl/water signatures across the sunlit surfaces of the Moon. However, since the instrument was designed to create a mineralogical map of the surface it covered a spectral range of 0.43 - 3 μm . Therefore, the 3 μm position for the H_2O and OH absorption features sits at the limit of the instrument's measurement capability. While the feature could be observed in the data, its full spectral shape could not be mapped or characterised in its entirety as can be observed from the spectra plotted in Figure 2-1, and as discussed in Section 2.1.1.

The relative global band depth of the 3 μm feature, calculated from the available spectral data, is plotted in Figure 2-6. This plot reveals a clear latitudinal dependence, with the band depth increasing significantly from the 45-60° latitude range to the poles. Additionally, there is a noticeable compositional dependence, with the feature being present (purple colouration) in the lighter highland regions but not detected (transparent colouration) in the adjacent lunar mare regions at the 45° latitude, 240° - 300° longitude




Figure 2-6: Global lunar distribution of OH relative band depth as measured by M^3 . The OH relative band depth has been overlaid onto a plot of the 750nm reflectance. Regions with relative band depth close to zero have been made transparent. Band depth shows a marked increase beyond 60° of latitude. (Klima & Petro, 2017).

boundary. At the latitude where OH reflectance can begin to be measured ($\sim 50^\circ$), it is observed only in the highlands and not in the mare (the northern section of Mare Imbrium).

Because the $3\ \mu\text{m}$ feature was at the limit of M³'s spectral range, the band depth could not be accurately assessed, as an anchor point could only be established on the short-wavelength side of the feature. This limitation makes the band depth particularly sensitive to the spectral slope, which would be influenced by any residual thermal component. This sensitivity underscores the need for more precise surface temperature measurements to accompany future spectral missions to improve their accuracy (Klima & Petro, 2017).

Different thermal correction models have been applied to the M³ data, each of which produces slightly different band depths and, consequently, different water maps of the lunar surface. To demonstrate this, variations in results after different thermal corrections are illustrated in Figure 2-7. The uncorrected and M³ Level 2 corrected data

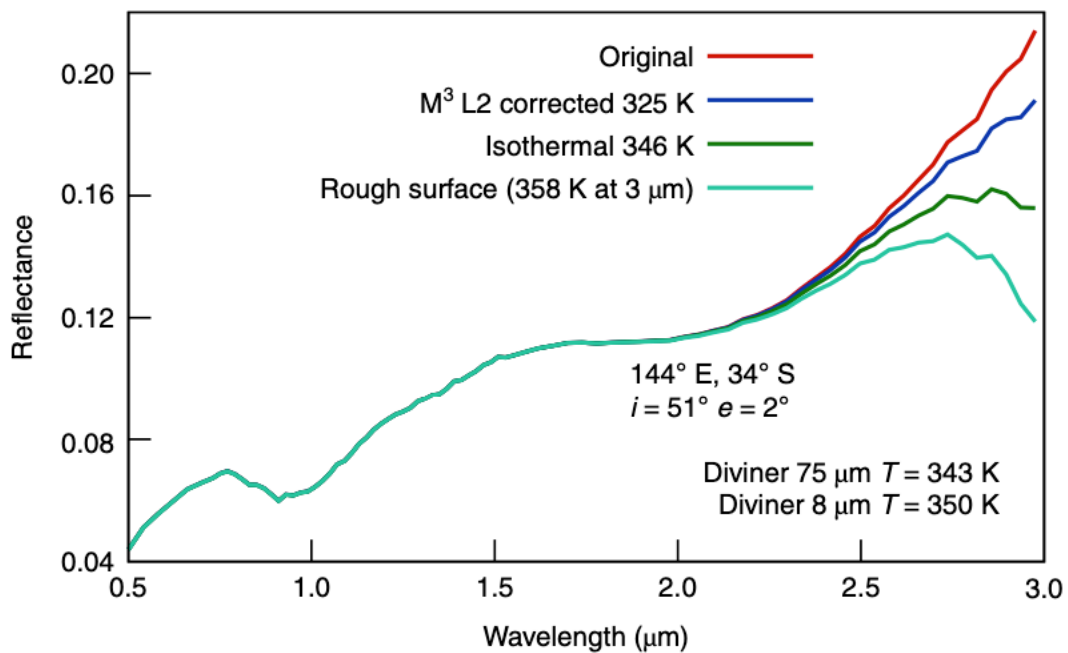


Figure 2-7: An example spectrum showing the effects of three different thermal corrections compared to the original uncorrected data. The uncorrected (red), M³ Level 2 corrected (blue) data are plotted alongside a radiative equilibrium model (green) and a model including surface roughness (cyan). This plot has been reproduced from (Bandfield et al., 2018).

show little-to-no evidence for a 3 μm absorption feature. However, when considering warmer surfaces via radiative equilibrium assumptions or accounting for surface roughness, the thermal contribution is increased significantly and therefore the corrections produce a meaningfully different absorption feature. Since the results of the measurements are so heavily dependent on the model used to analyse them, it is clear we need to be able to better constrain the models used in the analysis. One example of this would be to produce coincident high-accuracy temperature measurements as the spectral measurements are being captured.

Furthermore, while M³ scanned 95% of the lunar surface in its low spectral and spatial resolution operating mode, it achieved relatively few targeted observations in its highest resolution mode due to the early end of the mission. In summary M³ was limited in its spectral range and sensitivity, as well as in its ability to thermally correct at the long-wavelength part of its spectrum.

2.3 Lunar Trailblazer

2.3.1 The Trailblazer Mission

Following the confirmation of volatiles on the lunar surface discussed in Section 1.1, key questions remain regarding their abundance, distribution, and transport across the Moon.

The Lunar Trailblazer mission (LTB) will tackle these science questions and improve on the data gathered by past missions, such as M³, which was not designed to focus on the discovery and analysis of volatiles. The High-Resolution Volatiles and Mineral's Moon Mapper (HVM³) is the primary instrument on Trailblazer and will measure the spectrum across the absorption features for hydroxyl, water and ice between 2.6 and

3.6 μm (Thompson et al., 2020). The Lunar Thermal Mapper (LTM), an infrared radiometer developed at the University of Oxford, will provide HVM³ with coincidental surface temperature measurements to allow for a precise thermal correction of the low energy tail of the Planck function. LTM will also have the capability to evaluate the surface mineralogy composition with 11 narrow-bandpass filter channels (Bowles et al., 2020). These instruments, working in concert, will produce localised maps of volatiles, mineralogy and surface temperature on the Lunar surface that directly confront the limitations present within the M³ and Diviner data.

Lunar Trailblazer's science goals are reproduced here, taken from (Ehlmann et al., 2021):

- Determine the form, abundance and distribution of H₂O and OH on the lunar surface.
- Assess possible time variations of lunar volatiles in the sunlit surface.
- Determine the form, abundance and distribution of water ice in permanently shadowed regions.
- Quantify localised gradients in surface temperature to understand how these affect water and ice concentrations and search for small cold traps.

In order to support these goals LTM's main function is to provide a simultaneous independent estimate of surface temperature, thereby improving the thermal correction for HVM³ in the $\sim 3 \mu\text{m}$ region. LTM will also expand upon the compositional investigation made by the Diviner instrument through its narrow band channels, enabling the presence of volatiles to be related to the mineralogical environment of the local Lunar surface.

The work I will detail in this thesis aims to support the ground testing and in-flight calibration of LTM by providing calibration products for the data pipeline and quantifying the uncertainties therein to understand the overall uncertainty of LTM's final

measurements. This feeds through the instrument's science return by improving the thermal correction for HVM³, thus enabling precise analysis of the volatiles present on the lunar surface.

2.3.2 The High-Resolution Volatiles and Minerals Moon Mapper

The High-Resolution Volatiles and Minerals Moon Mapper (HVM³) builds on the capabilities of the M³ instrument, with significant enhancements to improve the detection of water on the lunar surface. HVM³ is an imaging spectrometer with a spectral scanning range from 0.6 to 3.6 μm , a spectral sampling rate of 10 nm, and a pixel field of view of 0.77 mrad, which will correspond to a ground image spatial resolution below 100 metres in the typical science orbit. The instantaneous field of view (iFOV) is the angle subtended by a single pixel on an instrument. By extending its measurement range to 3.6 μm , HVM³ can distinguish between the spectral signatures of water-ice, hydroxyls, and molecular H₂O, enabling detailed mapping of these species across the lunar environment.




Figure 2-8: The High-Resolution Volatiles and Minerals Moon Mapper (HVM³) instrument, fully assembled, prior to integration with the spacecraft.

The instrument also boasts greater sensitivity, with signal to noise predicted up to 100 from scattered light, which allows it to perform measurements inside permanently shadowed regions (PSRs) by utilizing the illuminated edges of craters. This makes HVM³

the first imaging spectrometer capable of probing spectral features within PSRs, providing unprecedented insights into the distribution of water in these regions.

The expanded spectral range compared to M³ allows for the observation of the continuum on both sides of the 3 μm absorption feature. This capability facilitates more accurate thermal corrections, essential for distinguishing between thermal and compositional effects. These corrections will be further refined by the Lunar Thermal Mapper, which provides simultaneous direct temperature measurements.

Additionally, HVM³ will conduct repeated observations of selected sites at various phase angles, corresponding to different times of the lunar day. Observations across these lunar times will help to constrain the existence of a diurnal water cycle by detecting any variations in the form and abundance of volatiles throughout the lunar day (Thompson et al., 2020).

2.3.3 Lunar Thermal Mapper

The Lunar Thermal Mapper (LTM) is a push-broom multispectral thermal mapper, with its optical design illustrated in Figure 2-9. When the pointing mirror is oriented in the nadir direction, radiation from the lunar surface is reflected through the instrument's optics onto a microbolometer array detector. The incoming light is focused by the first two




Figure 2-9: A diagram of LTM showing the progress of radiation through the instrument optics when the instrument is viewing the Lunar surface. The pointing mechanism can rotate the pointing mirror so that the incident light on the detector is from above (deep space view) or from the blackbody target (Oxford).

mirrors onto a preliminary focal plane where the filter assembly is located. After passing through the filters, the light is refocused by an additional three mirrors to form an image on the detector surface. The pointing mirror can alternate between viewing the lunar surface, deep space, or an internal blackbody for calibration purposes.

LTM's filter design draws on the heritage of the Diviner instrument, incorporating both wide bandpass filters for surface temperature measurements and narrow bandpass filters for compositional analysis. The full filter arrangement is detailed in Table 2.

2.3.3.1 LTM Thermal Measurements

LTM's primary scientific objective is to provide independent and simultaneous surface temperature estimates to support the thermal correction of HVM³'s spectral

measurements. Temperature measurements are formed through the LTM calibration procedure. The microbolometer in LTM measures resistance within each pixel, which rapidly change to reflect the incident radiation. These resistances are transformed into accurate radiance values through a calibration process which compares against the resistances measured when observing the known radiance from two calibration targets. These targets are the view of deep space (radiating the CMBR), and an internal blackbody with a well-characterised spectral output and precisely measured internal temperature. The calibration process, discussed in detail in Section 3.1, operates on an individual pixel level and allows for the calculation of radiance values for each filter channel. Since the surface emissivity is unknown, the brightness temperature within each channel is established. Brightness temperature is the temperature at which a blackbody would have to be to duplicate the observed radiance intensity within a wavelength range. The relative magnitude of brightness temperatures reflects the underlying emissivity values as well as the different depths penetrated by light at different wavelengths. These brightness temperature values can then be used to estimate the surface's kinetic temperature profile (Bowles et al., 2020).

For effective thermal correction of the 3 μm absorption feature, LTM must achieve a radiometric accuracy of ± 1 K through its design and calibration procedures. This is the instrument's top level science requirement.

Channel Name	Channel Bandpass / μm	Purpose
A1	6.875 – 7.125	Compositional
A2	7.125 – 7.375	Compositional
A3	7.375 – 7.625	Compositional
A4	7.625 – 7.875	Compositional
A5	7.875 – 8.125	Compositional
A6	8.125 – 8.375	Compositional
A7	8.375 – 8.625	Compositional
A8	8.627 – 8.875	Compositional
A9	8.875 – 9.125	Compositional
A10	9.375 – 9.625	Compositional
A11	9.875 – 10.125	Compositional
A12	6.25 – 12.5	Thermal
A13	12.5 – 25	Thermal
B1	25 – 50	Thermal
B2	50 – 100	Thermal

Table 2-2: Bandpass ranges for LTM. It ought to be noted that the limits on the bandpasses are not exact since the filters taper at each end of the bandpass. The compositional channel bandpasses (A1-11) have been calculated from the central position of each bandpass and the targeted bandwidth of 0.25 μm . These are the bandpass values used for the idealised filters used in radiometric modelling in Section 3.1.

2.3.3.2 LTM Compositional Analysis

LTM is equipped with 11 compositional channels, designed for precise mapping of the CF position across the lunar surface. These channels span the 7-10 μm range, offering a significant improvement over Diviner's 3-channel system, which covered the 7.5 - 8.7 μm range. This expanded spectral coverage allows LTM to capture the entire shape of the CF, rather than just its position.

Returning to the Apollo emissivity spectra considered in Section 2.1.2, which have been reproduced in Figure 2-10 over the wavelength range examined by LTM, it is apparent that the spectra differ in more than just CF position. The gradients of the spectra on the short wavelength tails of the peaks are different, and the shape of the spectra deviate on the long wavelength side of the feature. This information, which relies on the entire shape of the feature, rather than just the CF position, could not be mapped by the three channel compositional data on Diviner. LTM's enhanced spectral resolution allows for a more detailed mapping of these variations, enabling more sophisticated mineralogical investigations to contribute to the Trailblazer dataset. For instance, the gradient of the short-wavelength tail provides information about particle size, with steeper tails indicating smaller particles. The mid-infrared emission spectrum also contains Reststrahlen bands, which are directly related to the fundamental vibrational modes of minerals. LTM's compositional channels are designed to capture these features, offering valuable data for constraining the mineralogical composition of the lunar surface.

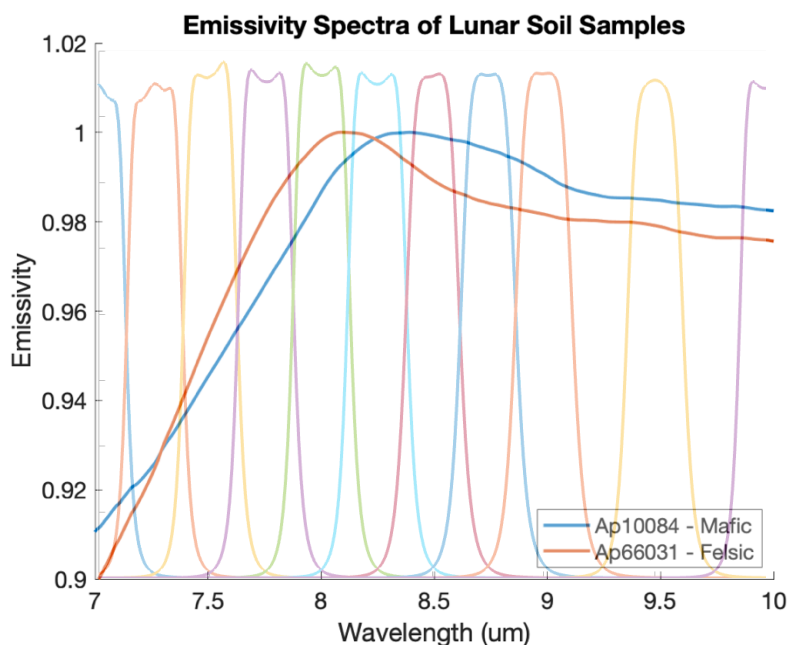


Figure 2-10: Plot of the emissivity spectra of two Apollo soil samples over the wavelength range covered by LTM's compositional filters. The spectra have each been normalised against their maximum values. Data from (Donaldson Hanna et al, 2017). The LTM compositional bandpasses have been overlaid to show their position in comparison to the mineral spectra.

Due to the narrow bandpasses of the compositional channels, less power reaches the detector. Therefore, these channels are the primary driver of the instrument's performance requirements. In order to accurately determine surface composition, the emissivity must be measured with a signal-to-noise ratio providing a precision considerably higher than ± 0.05 , as this is the typical magnitude of emissivity features. This allows observed features to be measurable beyond the magnitude of the combined instrument and calibration noise sources.

2.4 Mission Requirements

Having discussed the scientific context and the challenges associated with the measurements that the Lunar Trailblazer mission aims to produce, I may discuss the performance requirements imposed on the Lunar Thermal Mapper. These requirements are meticulously defined to ensure that the mission achieves its science goals with the necessary precision and accuracy.

Space missions are complex and expensive undertakings, often involving budgets of hundreds of millions of pounds and development timelines spanning many years or even decades. Given the substantial investment in time and resources, the development process for such missions is highly rigorous. Central to this process is the definition and flow-down of requirements – a systematic approach to ensuring that the final data products are capable of addressing the mission's scientific questions.

A requirement flow is a hierarchical decomposition of high-level mission objectives into specific, measurable engineering and operational criteria. These flows begin with the top-level science requirements – such as the need to accurately measure the presence of water on the lunar surface – and break them down into progressively more detailed requirements that govern every aspect of the instrument's design, calibration, and

operation. The goal of this flow is to trace the path from broad scientific goals down to the fine details of engineering specifications, ensuring that each step in the process supports the mission's overall objectives.

To illustrate the concept, consider the following example of a requirement flow for Lunar Trailblazer, focusing on the measurement of water on the lunar surface (these numbers are only representative):

1. **Top-Level Requirement:** The mission must measure the temperature-corrected reflectance of the lunar surface to within 1% accuracy to be able to detect and quantify water and hydroxyl.
2. **Derived Requirement:** To achieve this reflectance accuracy, the LTM must measure the surface temperature with a precision of ± 1 K, since temperature corrections directly affect reflectance measurements in the 3 μm absorption region.
3. **Subsystem Requirement:** Accurate surface temperature measurements via calibration require the radiance output from LTM's on-board blackbody calibration target to be known with an uncertainty of no more than 0.5%.
4. **Component-Level Requirement:** Achieving this radiance accuracy necessitates that the temperature sensors embedded in the blackbody calibration target measure temperature with an accuracy of ± 200 mK.

In practice, requirement flows are far more complex than the example provided. For instance, the requirement to know the calibration target radiance to within 0.5% introduces further requirements related to:

- **Temperature Calibration:** the thermometer must have been calibrated in the laboratory against a known thermometer in a highly controlled thermal environment.

- **Measurement Electronics Accuracy:** the circuitry used to create the temperature measurement must not introduce additional noise above a certain level.
- **Modelled Radiance Response Accuracy:** the relationship between temperature and radiance must be accurate to within a certain uncertainty.
- **Data Handling:** the computation involved in recording, storing and manipulating the target temperature must not introduce noise above a certain level.

Each of these aspects introduces its own set of requirements, creating a complex web of interrelated criteria that must all be satisfied to ensure the overall success of the instrument. By having a requirement on each of these factors at each level, we can know that the next level up will be satisfied. Therefore, once the requirements have been defined, we can satisfy the instrument performance by designing the instrument to satisfy all of the bottom level requirements. The uncertainty on the final measurement is produced by gathering all the lowest level uncertainties and stacking them together to recreate the full instrument.

2.4.1 Science Requirements of LTM

In this section I reproduce the science requirements that are directly relevant to the work in this thesis. These relate to the accuracy of measurements through temperature sensors, thermal stability, calibration accuracy and signal magnitude relative to noise.

The aim of this project is to support the radiometric in-situ calibration of LTM through ground-based laboratory calibration and computational modelling and simulation. Several of LTM's design requirements are directly applicable to this work and have been reproduced here so they can be referred to in the next chapters as points of reference and validation.

Requirement Name	Requirement
RL3- LTM-Absolute calibration	LTM shall incorporate a calibration target to provide knowledge of the radiometric accuracy to $\pm 1K$.
RL3- LTM- Calibration stability	The LTM filter assembly and detector array shall be stable to $\pm 0.01K$ over 300 seconds.
RL3- LTM- Mineral channels	LTM shall implement a minimum of 3 mineral channels to provide cross reference with Diviner with detector NEP low enough to achieve a signal to noise ratio of 100 for thermal emission from a surface at 333 K.
RL3- LTM- Thermal channel wavelength range	LTM thermal channels shall cover the wavelength range 6 to 100 microns.
RL3- LTM- Thermal channel sampling	LTM shall implement a minimum of 3 thermal channels to cover the wavelength range in RL3- LTM- Thermal channel wavelength range with detector NEP $\leq 3.3 \times 10^{-11} W$.
RL3- LTM- Temperature monitoring	LTM shall implement redundant temperature monitoring of the filter assembly and detector assembly.
LTB-LTM-IRD-098	The LTM provided FPA and filter temperature sensors read by the spacecraft shall be 1 kOhm PRTs.

Table 2-3: The LTM Design requirements relevant to the work in this thesis (D-104876 LTM Requirements).

3 LTM Characterisation and Design Work

This chapter presents the work I conducted in support of the development of LTM. The tasks span from optimizing and testing final design decisions to characterizing flight systems and delivering post-calibration data products for use during the instrument's operational phase. A central focus throughout is the radiometric calibration of the instrument, and the underlying physics and functionality of the radiometric models used to evaluate LTM's performance will be thoroughly explained. Due to the commonality with the TIRI module in MIRMIS this work also prepares for the instrument operations and performance discussions in Chapters 5 and 6.

The chapter is structured around several key areas of development. First, I introduce the radiometric models created in-house, which are essential for evaluating various aspects of the instrument's performance. I then describe the work done to measure the optical transmission of the instrument through spectral measurements, providing results that are crucial for constraining its optical throughput. Next, I discuss the development and testing of the electronics used for high-precision onboard temperature measurements, to ensure that LTM can satisfy its science requirements. The chapter also covers the thermal calibration of LTM's internal blackbody calibration target, followed by the analysis and processing of the resulting data. Finally, I detail the generation of reference tables, which relate the measured radiance to surface brightness temperature for each channel of the instrument.

As LTM is a thermal instrument, its performance is highly sensitive to its internal thermal environment, and the work described here is unified by the theme of characterizing and improving this thermal performance. The radiometric models allow us to predict the effects of thermal instability or measurement errors, while the optical transmission measurements enable accurate radiometric modelling. The development of

the housekeeping board and Platinum Resistance Thermometer (PRT) measurements ensure precise onboard temperature monitoring, most crucially for the blackbody target. The calibration of this blackbody enhances the accuracy of temperature measurements, which reduces radiometric uncertainty. By integrating these elements, the full radiometric model quantifies how uncertainties in the calibration target's temperature propagate to uncertainties in the measured radiance, directly influencing final calibration accuracy. Ultimately, the completed radiometric model is used to create in-flight reference tables that link the instrument's calibrated radiance outputs to the surface brightness temperatures being observed.

3.1 Radiometric Modelling

Radiometric modelling of the instrument enables the calculation of performance metrics based on various input parameters, such as temperature fluctuations or detector pixel area. This approach allows for the evaluation of different instrument designs and provides insight into how specific variables affect overall performance. For example, modelling can help determine the level of thermal stability required to maintain the radiometric accuracy within the limits specified in the requirements, as outlined in Section 2.4.1. By simulating these conditions, we can ensure the instrument meets its design goals under real-world operational scenarios.

3.1.1 Instrument Model

3.1.1.1 Measuring Scene Radiance

To build up the radiometric model, we begin by considering the radiance as generated at the surface of the target scene. Next, we consider how this radiance is passed through the instrument, before we finally examine how the radiance is measured at the detector.

The scene spectral radiance in the instrument model is the lunar surface, emitting thermal radiation. In a simplified model the surface can be modelled as a uniform material, at temperature T_s and with characteristic emissivity ϵ_s . Radiation at the instrument from the surface will be a combination of emitted and reflected thermal radiation. However, for scenes not within a crater, the reflected radiation will largely be from a view of empty space, emitting at 2.7K. While the radiance of a body at Solar surface temperatures is large, the very small solid angle subtended by the sun, combined with the low reflectivity of lunar regolith in the thermal infrared, means that the lunar surface emission dominates over Solar radiation. There will be reflection from other lunar highland surfaces, but this will be marginal. This gives a simple spectral radiance of the surface as:

$$L_s = \epsilon_s B_\lambda(T_s) \quad (3.1)$$

Where $B_\lambda(T_s)$ is the spectral radiance of a surface at a temperature T_s as defined by the Planck function:

$$B_\lambda(\lambda, T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{hc/\lambda k_B T} - 1} \quad (3.2)$$

Where h is the Planck constant, c is the speed of light, λ is the wavelength, k_B is the Boltzmann constant, and T is the temperature of the surface.

As this radiance passes through the instrument it is reflected multiple times as the optics first focus the image on the filter plane before focusing the image onto the detector focal plane. For each reflection the signal is reduced, and signal is also lost at the filters due to imperfect in-band transmission. This gives an overall instrument transmission of

$$\tau = r_d R^n T \quad (3.3)$$

wherein n is the number of mirrors, R is the reflectivity of each mirror, r_d is the response function of the detector and T is the transmissivity of the filter. The reflectivity of

the mirrors may depend on wavelength, but it is typically flat over IR wavelengths. A fully characterised instrument transmission function is described in Section 3.2.4, relating the spectral relationship between incident radiance and the radiance at the detector.

The radiant flux incident on a detector pixel is determined by integrating the surface radiance over the solid angle subtended by the pixel and scaled by the pixel's effective area. This is characterised using L_s , the pixel area A_d , the solid angle Ω_d and the instrument transmission as follows:

$$\Phi_s = A_d \Omega_d \int \tau(\lambda) L_s(\lambda) d\lambda \quad (3.4)$$

The total power that is incident on the detector depends on the radiant flux transmitted through the instrument optics from the source Φ_s , emitted flux from components within the instrument optics Φ_l , and the flux emitted by the detector pixel Φ_d .

$$P_d = \Phi_s + \Phi_l - \Phi_d \quad (3.5)$$

As the instrument will not be isothermal it will emit across a range of temperatures, however, for simplicity this model neglects the effect of self-emission. Its effect will be incorporated into other work building on this model, as discussed in Section 3.1.2. Flux emitted from the pixel can be calculated from the temperature of the detector array, however, for a simple calibration it can be included in the instrument radiance and voltage offset as a constant term.

LTM operates using a microbolometer array as its detector. Each pixel in this array is formed from a pad of thermally absorbent Vanadium Oxide. This pad is suspended in front of a reflector and changes in temperature as the thermal radiation incident on the detector varies. The Institut National d'Optique (INO) create these detectors with a black gold coating for improved flat long wavelength performance (INO, 2024). Readout

electronics below the pixel measure the resistance across the pad, which is then converted to a voltage V_s which is proportional to the incident radiation power P_0 :

$$V_s = aP_0 + b \quad (3.6)$$

As the pixel construction is imperfect each pixel will have a unique response, the gain a and offset b must be determined individually for each pixel during calibration.

Setting aside issues of thermal instability in the instrument and focal plane array, calibration ensures that each pixel in the detector array responds accurately to incident radiance, accounting for inherent pixel-to-pixel variations. This is achieved by determining gain $A_{i,j}$ and an offset $B_{i,j}$ for each pixel (located at i,j), from looking at two targets of known radiance, typically black bodies emitting with Planck spectral radiance B_λ .

$$V_s = A_{i,j} \int \tau L_s d\lambda + B_{i,j} \quad (3.7)$$

For the hot and cold blackbody targets, the calibration equations are:

$$V_{i,j}(T_{hot}) = A_{i,j} \int \tau B_\lambda(T_{hot}) d\lambda + B_{i,j} \quad (3.8)$$

$$V_{i,j}(T_{cold}) = A_{i,j} \int \tau B_\lambda(T_{cold}) d\lambda + B_{i,j} \quad (3.9)$$

Solving for $A_{i,j}$ and $B_{i,j}$ gives:

$$A_{i,j} = \frac{V_{i,j}(T_{hot}) - V_{i,j}(T_{cold})}{\int \tau B_\lambda(T_{hot}) d\lambda - \int \tau B_\lambda(T_{cold}) d\lambda} \quad (3.10)$$

$$B_{i,j} = V_{i,j}(T_{hot}) - A_{i,j} \int \tau B_\lambda(T_{hot}) d\lambda \quad (3.11)$$

The scene radiance is calibrated by comparing the data number from the scene with those of known reference points, including a blackbody target and space view. The

data number is the completely uncalibrated numerical output recorded from the detector's pixel. This calibration process transforms the measured data numbers to measurements of physical radiance values via the following calibration equation:

$$B_{Scene} = (B_{Blackbody} - B_{Space}) \times \left(\frac{D_{Scene} - D_{Space}}{D_{Blackbody} - D_{Space}} \right) + B_{Space} \quad (3.12)$$

Where $B_{Blackbody}$ is the known radiance from the blackbody target, B_{Space} is the known radiance from the space view, D_{Scene} is the data number from the scene view, D_{Space} is the data number from the space view and $D_{Blackbody}$ is the data number from the blackbody observation. This calibration equation makes the assumption that the detector response is linear.

The calibrated scene radiance B_{scene} can allow us to calculate the temperature of the observed scene:

$$B_{scene} = \int \tau \epsilon_s B_\lambda(T_s) d\lambda = \frac{V_s - B_{i,j}}{A_{i,j}} \quad (3.13)$$

While the simplified model presented here assumes uniform temperature and emissivity, the actual lunar surface features a complex mixture of materials, particle sizes, and thermal gradients. The emission spectra allow for investigation of the surface by remote sensing, but it creates challenges in simulating retrieved spectra. For a simple noise analysis, while thermal emission models can be used to create accurate emission spectra, simply altering the emissivity spectrum of an isothermal surface can provide an understanding of the effects of sensing different composition surfaces.

3.1.1.2 Evaluating Noise Performance

To quantify the signal-to-noise performance of the instrument, it is essential to understand and evaluate the noise characteristics of the detector. Both LTM and MIRMIS have top-level requirements to measure surface temperature to an accuracy defined in

kelvin. One crucial metric used to assess the instrument's performance is the Noise Equivalent Delta Temperature (NEDT). NEDT corresponds to the smallest change in temperature that would result in a change in the detected signal equal to the noise level. This metric is primarily driven by the performance of the detector array and the Noise Equivalent Power (NEP), which represents the minimum detectable power for the detector.

However, since NEDT depends on the specific measurement conditions and the instrument's optical setup, a more general metric, the specific detectivity, D^* , is often used to evaluate and predict performance under varying conditions. D^* is an inherent property of the detector and provides a consistent framework to compare the instrument's sensitivity across different optical configurations. Numerically, D^* is defined as:

$$D^* = \frac{\sqrt{A_d}}{NEP} \quad (3.14)$$

where A_d is the area of the pixel and NEP is the noise equivalent power. NEP is calculated as:

$$NEP = \frac{A_d \Omega_d}{\sqrt{f}} NER \quad (3.15)$$

where NER is the noise equivalent radiance and f is the frame rate the detector is operating at. The NER is calculated from the derivative of the calibrated radiance with respect to temperature, evaluated at the temperature of the radiance source:

$$NER = NEDT \left[\frac{d(B_{scene})}{dT} \right]_{T_s} \quad (3.16)$$

Combining the equations 3.14, 3.15 and 3.16 gives our final forms for D^* and $NEDT$:

$$D^* = \frac{\sqrt{f}}{\sqrt{A_d} \Omega_d NEDT \left[\frac{d(\int \tau \epsilon_s B_\lambda(T_s) d\lambda)}{dT} \right]_{T_s}} \quad (3.17)$$

$$NETD = \frac{\sqrt{f}}{\sqrt{A_d} \Omega_d D^* \left[\frac{d(\int \tau \epsilon_s B_\lambda(T_s) d\lambda)}{dT} \right]_{T_s}} \quad (3.18)$$

The performance of an optical system is typically measured by the signal-to-noise ratio (*SNR*), this is calculated from ratio between the incident power on the detector and the noise equivalent power:

$$SNR = \frac{B_{scene} A_d \Omega_d \Delta t}{NEP \sqrt{\Delta t}} = \frac{D^* \sqrt{A_d} \Omega_d \int \tau \epsilon_s B_\lambda(T_s) d\lambda}{\sqrt{f}} \quad (3.19)$$

These performance metrics are calculated based on the solid angles of the optical design. However, since optical systems are typically defined in terms of F-number, F_N , this must be converted to the solid angle. The relationship between the two is given below:

$$\Omega_d = 2\pi \left(1 - \cos \left[\arctan \left(\frac{1}{2F_N} \right) \right] \right) \quad (3.20)$$

3.1.1.3 Modelling Outputs

The baseline performance of the instrument was evaluated in MATLAB using equations 3.18 and 3.19, calculating NEDT and SNR values for all LTM filter channels

Modelling Parameters	Values
D*	8.5e8 cm Hz ^{1/2} W ⁻¹
Mirror reflectivity	0.95
Telescope F-number	1.4
Framerate	65 Hz

Table 3-1: Modelling parameters used for the modelling in this section

across a relevant target temperature range for the lunar surface using the model parameters given in Table 3-1.

The SNR curves in Figure 3-1 demonstrate that the 11 compositional channels are considerably more sensitive to error since their SNR is much smaller than the thermal channels. This is no surprise since the narrower bandpasses allow less radiation to be incident on the detector. The exception to this rule is the final thermal filter, B2, which struggles at high temperatures since it is a long-wavelength channel, and at higher temperatures the Planck function is significantly shifted to short wavelengths. However, the important conclusion from this is that since the compositional channels are more sensitive, they will drive the instrument’s SNR requirements more than the thermal channels. This plot demonstrates the ability of the instrument to satisfy LTM’s Mineral Channel requirement of a 100 SNR at 333 K (Table 2-3).

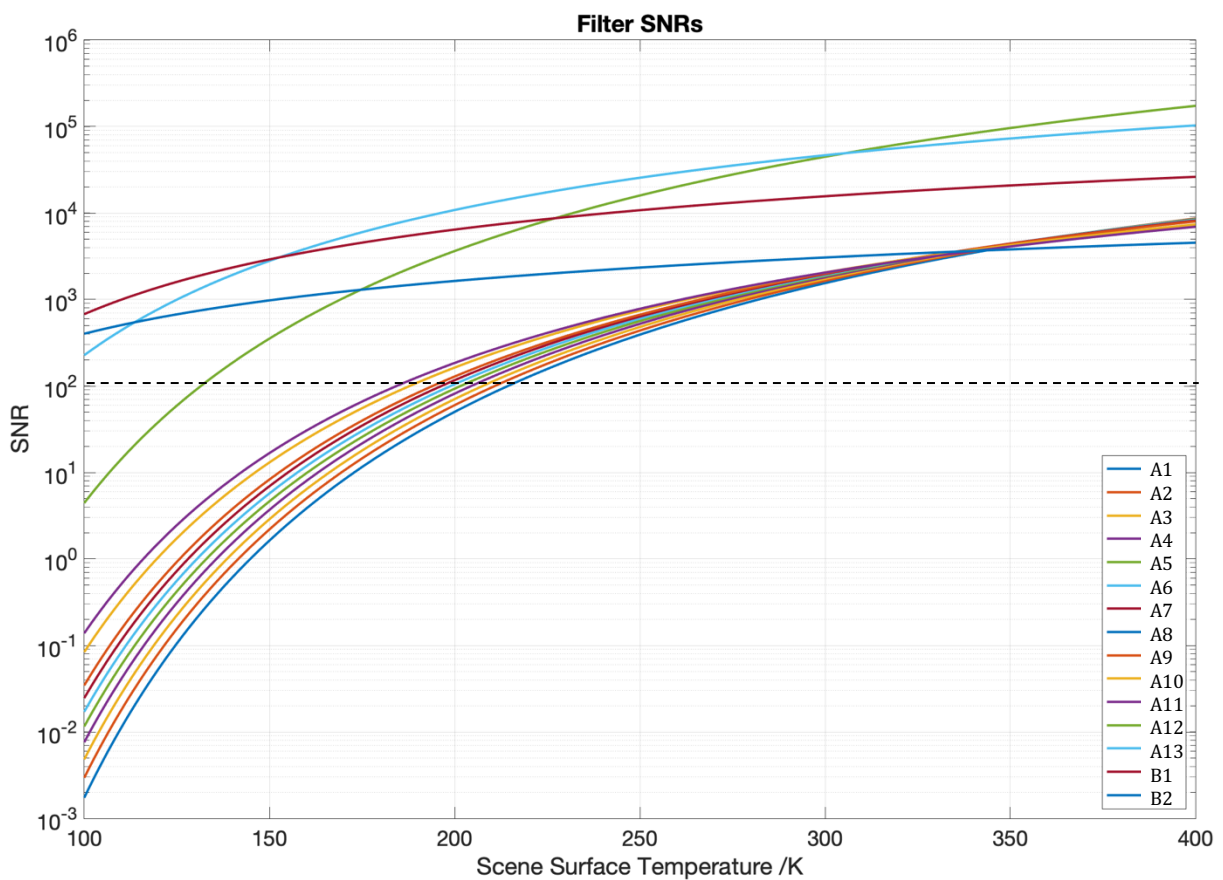


Figure 3-1: Signal to Noise Ratios plotted against temperature for idealised LTM filters. Parameter values are as given in Table 2. Filter bandpasses are as given in Table 4 with a transmissivity of 0.8. the black dotted line represents the mineral channel requirement SNR which must be achieved above a temperature of 333 K.

The significance of the NEDT is that it represents the effective minimum resolvable temperature by the instrument in that channel indicating the best performance the instrument could have based on the detector noise, before other factors. The NEDT curves have been evaluated for each channel in Figure 3-2, over the same temperature range as used for Figure 3-1. This modelling shows that the thermal channels should remain above the radiometric accuracy requirement due to noise to a temperature of 100K, allowing observations inside of PSRs. However, the compositional channels are limited to higher temperatures >180K. Therefore, LTM will be able to measure temperatures inside PSRs in order to support HVM's thermal corrections, but it will not provide the compositional data to accompany that. This is acceptable since most of the lunar surface, aside from PSRs, reaches temperatures above 180K during the day, which will constitute the majority of LTM's observations (Williams et al., 2017).

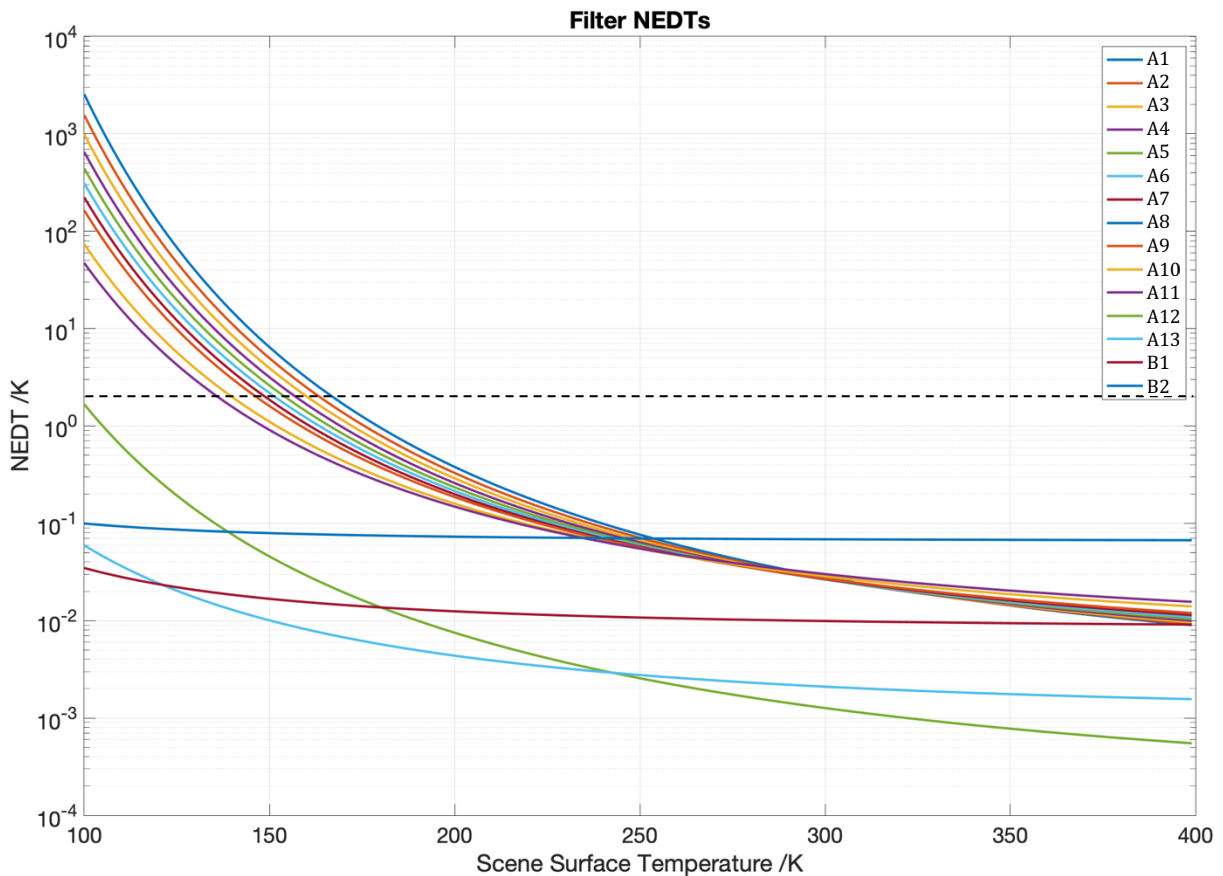


Figure 3-2: NEDT plotted against temperature for idealised LTM filters. Parameter values are as given in Table 2. Filter bandpasses are as given in Table 4 with a transmissivity of 0.8. The black dotted line represents the 1K requirement.

Furthermore, the ability to detect specific features in the emission spectra was examined through a "science-to-noise" ratio. This metric, defined as the ratio between the difference in incident power from two scenes with varying emissivities and the noise equivalent power, is crucial for identifying features in the emission spectrum. The signal-to-noise ratio alone is insufficient in such cases, as resolving features depends on the difference between signals emitted by surfaces with different emissivities. If the signal difference is smaller than the noise, the feature will not be resolved.

$$\text{'science to noise' ratio} = \frac{A_d \Omega_d \sqrt{\Delta t} [B_s(\epsilon_s) - B_s(\epsilon_s - \Delta\epsilon)]}{NEP \sqrt{\Delta t}} \quad (3.21)$$

This ratio has been plotted for each of the compositional channels in Figure 3-3, showing that a 2% emissivity feature can be resolved from noise at surface temperatures ranging from 200 – 230 K, depending on the channel. This is the temperature above which

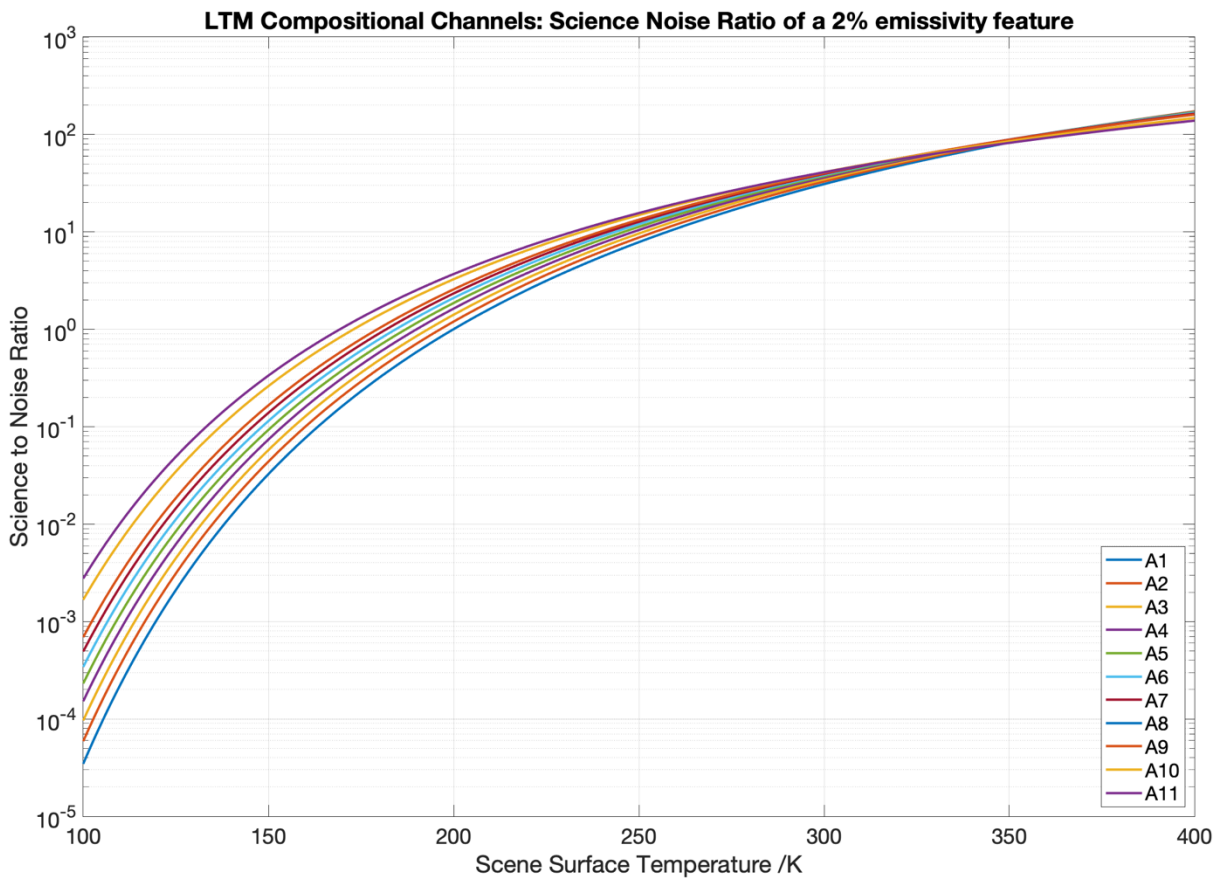


Figure 3-3: 'Science to noise' ratio for a 2% emissivity feature plotted against temperature for idealised LTM filters. Parameter values are as given in Table 3-1. Filter bandpasses are as given in Table 2-2, with a transmissivity of 0.8.

the feature being measured has a spectral signal more than four times the magnitude of the noise. This highlights the challenges of making compositional measurements in cold polar regions and confirms that such measurements will not be feasible inside PSRs. However, these regions are outside the primary requirements for LTM. The ability to differentiate between different mineral types depends significantly on the size of the emissivity features in their spectra. By combining laboratory measurements of Apollo samples and lunar analogues to estimate feature size expectations, along with temperature maps of the Moon's surface, we can identify areas where LTM's compositional measurements will be most effective. It is important to note that mineral differentiation may not be equally achievable across all regions.

3.1.2 A Combined Thermal-Radiometric Model

Having established the basic principles and detector limiting case of the radiometric model, I will now extend its complexity to investigate the instrument's performance under the variable thermal conditions of lunar orbit.

3.1.2.1 *The LTM Calibration Procedure*

As an infrared radiometer utilizing a microbolometer array, LTM detects variations in radiance, relying on a careful calibration process to form surface temperature measurements. This process relies on ground testing data obtained before launch, as well as in-situ calibration views during the mission, to determine offset and gain values necessary for accurate radiometric measurements. The measured scene radiance output is calculated using Equation 3.12.

For in-situ calibration, two known temperature references must be measured in close temporal proximity to each target view: a hot and a cold calibration view. The cold calibration relies on viewing deep space, where the Cosmic Microwave Background

Radiation (CMBR) at a well-characterised temperature of 2.725 K serves as a stable reference, where the instrument is mainly observing self-emission of the instrument. In contrast, the hot calibration view comes from an internal blackbody target that is heated by the instrument itself. Calibration views are taken before and after the instrument views the lunar surface, to account for any drift in the calibration values across the period of observation.

The typical LTM observation sequence is then as follows:

- The pointing mirror rotates to the blackbody view and integrates for 70 seconds.
- The pointing mirror rotates to the space view and integrates for 70 seconds.
- The pointing mirror rotates to the scene view, along the nadir, and starts push broom sampling for approximately 20s.
- The pointing mirror rotates to the space view and integrates for 70 seconds.
- The pointing mirror rotates to the blackbody view and integrates for 70 seconds.

With the known temperature, and therefore emission spectrum, of the blackbody target at an emissivity of unity, and assuming that the temperature of each component in the optical system is constant and uniform, the calibrated scene radiance in each channel is calculated. Since the radiation from space at 2.7 K is effectively 0, the recorded signal from the space view is a measurement of the background emission from the instrument.

The calibration sequence in Equation 3.12 only holds as long as the temperatures of all the optical components remain steady. For a dynamic system with constantly changing temperatures for all components, the calibration sequence relies on estimating

D_{Scene} and $D_{Blackbody}$ from their values during the calibration sequence. This will be nearly perfectly accurate (better than 0.1 K) in the case of a constant linear change in temperature within the instrument. However, for non-linear changes in temperature the calibration sequence will introduce a larger error. A combined radiometric and thermal model of the instrument was created to quantify this error.

3.1.2.2 Integrating Thermal Modelling Results

The radiometric model incorporates the self-emission from all components of the instrument during the calibration sequence. The layout of these components is shown in Figure 3-4, with the direct optical pathway being considered laid out as a block diagram in Figure 3-5 for simplicity. The incident radiance from the scene, as well as from the




Figure 3-4: A diagram illustrating the optical system within LTM. (MIRMIS-OXF-TN-007).

pointing mirror and the first two mirrors, only includes radiation within the filter's bandpass. The incident radiance from the scene is defined as:

$$P_{scene} = A\Omega \int_{\lambda_1}^{\lambda_2} (\epsilon_{surface} B_{scene}(T_{scene}, \lambda) \tau_{scene}) d\lambda \quad (3.22)$$

where $\epsilon_{surface}$ is the emissivity of the scene, T_{scene} is the temperature of the scene, τ_{scene} is the transmissivity of the path from the scene to the detector, and λ_1 and λ_2 define the bandpass of the filter.

The transmissivity of the path from each source to the detector is different for each radiation source. For the scene the path includes the losses at each of the six mirrors and passing through the filter, whereas for the pointing mirror it will be the same but with one fewer mirror since the radiation emitted from the pointing mirror is not reflected by that mirror. Filter transmissivity has been assumed to be flat at 80% across the thermal infrared. There is an analogous equation for each component in the instrument, with $\tau_{component}$ simply the sum of the transmissivities between the component and the detector.




Figure 3-5: The sequence of components on the optical pathway. The detector block sees the self-emission of each component, passed down the remaining pathway.

The radiance on the detector from the filter block is treated as two separate radiances. The first is emitted in the out of bandpass wavelength region, at an emissivity of unity. The second is the emission from the filters, which have an emissivity of 0.2 to match their transmission. The mirror array that provides the final focus is formed of three more low emissivity faces undergoing broadband emission.

The detector is housed in the detector block and baffle, and it is assumed that each pixel in the detector views this with an F-number of unity as a worst case. The assumptions made in the model, such as the detector block F-number and the emissivity of

the filter block are based on estimates of the worst-case scenarios, meaning the estimations of errors will provide an upper bound.

This model was used to simulate calibration errors from two different scenarios. In the first case a singular temperature change in a component between the calibration and the scene views was simulated. This was used to confirm that temperature variation within the filter block and detector block needed to be kept below ± 0.01 K between calibrations in order to keep the error due to calibration drift below the inherent noise of the detector. This mirrors the requirement on temperature stability from Table 2-3.

3.1.2.3 The LTM Thermal Model

The second scenario simulated considers the instrument environment when taking lunar surface observations from orbit. The thermal behaviour of LTM has been modelled throughout its orbit by the LTM industrial partners at the Rutherford Appleton Laboratory




Figure 3-6: Modelled LTM Component temperatures throughout an orbit. Temperatures have been calculated through an ESATAN model of LTM provided by RAL. (LTM-RAL-AN-0001). The grey regions demonstrate the three observations of the surface taken in each orbit. As the pointing mirror observes the surface the internal temperatures can be seen to increase.

(LTM-RAL-AN-0001). The temperatures of relevant components during the orbital period have been reproduced in Figure 3-6. Variations in the instrument temperature are largely governed by shading and illumination. For most of the orbit, LTM remains in the shadow of either the spacecraft or the Moon, with direct sunlight exposure occurring only between 5000 and 6000 seconds. During this period of illumination, the component temperatures rise rapidly, then begin to cool after 6000 seconds as the instrument enters the Moon's shadow, continuing to decay until the next exposure at 5000 seconds in the subsequent orbit.

The model was run in ESATAN-TMS, a thermal analysis software suite used in the aerospace industry for spacecraft thermal control and design. The simulation was performed using LTM's expected orbital parameters: a polar orbit 130 km above the lunar surface. The Moon was modelled as a sphere at 1 AU from the Sun, with a uniform surface temperature of 400 K. LTM was assumed to be thermally coupled to the Trailblazer spacecraft via an interface plate and covered in Multi-Layer Insulation (MLI) to help stabilise the instrument's temperature. The model was iterated over successive orbits until it converged on a stable solution, and the stable temperature cycle is what has been plotted in Figure 3-6.

The grey regions demonstrate the three observations simulated in each orbit as part of the thermal model. This can be seen in the small upward inflections in component temperature, see the detector temperature at 1200 seconds, caused when the pointing mirror rotates to the surface and the radiation from the 400 K surface propagates into the instrument. One important output of this modelling is that the expected filter block temperature varies more (up to 0.06 K) across the 300 second observations than permitted in LTM's design requirements as quoted in Section 2.4.1. The implications of this temperature instability will be considered presently.

3.1.2.4 In-Situ Calibration Error/ Temperature Instability

The temperature data from the thermal model was input into a time-domain version of the radiometric model to simulate the full calibration sequence for each observation. The model included simulations of the four calibration views, as well as the correct scene radiance for the observation. To calibrate the scene view, a linear interpolation between the pre- and post-scene calibration views was applied using the calibration formula (Equation 3.12). The resulting error in surface temperature measured by the instrument is presented in Table 3-2. These results reflect the worst-case scenario due to the assumptions made in the RAL modelling. The results here show a failure to achieve the LTM requirement on radiometric accuracy of ± 1 K, this is a direct result of the failure to achieve the detector block and filter block stability of 10 mK across the calibration sequence.

The method of calibration between the two calibration views is the key source of error. The model linearly interpolates between the first and second calibration view and takes values from this solution to calibrate the scene observations. However, since the error is being introduced by non-linear changes in instrument temperature the radiance does not change linearly across the calibration sequence. The difference between the linear plot between calibration views and the nonlinear radiance change creates the error. A more complex calibration procedure, for example deriving a polynomial curve of the radiances through the two calibration views, informed by instrument temperature measurements, may result in a more accurate estimate of the instrument radiance across the during the scene observation. However, evaluating the true viability of this method would require more example data than was retrieved during the LTM ground campaign. This modelling has shown that deviating from the standard calibration equation may provide an opportunity to improve instrument accuracy, and I hope to further investigate this once we have in-flight measurements of the LTM thermal behaviour .

LTM Channel	Channel Bandpass / μm	Maximum NEDT in Observation 1 /mK	Maximum NEDT in Observation 2 /mK	Maximum NEDT in Observation 3 /mK
A1	6.875 – 7.125	1249	702	1705
A2	7.125 – 7.375	1190	730	1628
A3	7.375 – 7.625	1163	768	1595
A4	7.625 – 7.875	1133	798	1558
A5	7.875 – 8.125	1090	814	1502
A6	8.125 – 8.375	1115	874	1543
A7	8.375 – 8.625	1348	1089	1883
A8	8.627 – 8.875	1332	1122	1866
A9	8.875 – 9.125	1333	1165	1874
A10	9.375 – 9.625	1493	1384	2119
A11	9.875 – 10.125	1611	1572	2304
A12	6.25 – 12.5	65	145	116
A13	12.5 – 25	135	241	34
B1	25 – 50	694	1001	717
B2	50 – 100	5240	7002	6163

Table 3-2: The maximum errors calculated in the thermal-radiative model, measured in each channel for each observation in the orbit. Errors are given as NEDT values in mK in a colour scale where red equates to more error, and green to less.

The first solution to the large accuracy errors given in Table 3-2 is to address their root cause – the temperature instability during calibration. Given the large solid angle between the filter block and the detector, the relatively high emissivity of the filter block (compared to the mirrors) and the high transmissivity between the filter block and the detector, it can be assumed that a large proportion of the instrument radiance is from the filter block. Coupled with the fact that the filter block absorbs a high proportion of the radiation fed through the instrument from the pointing mirror, and thus the filter block exhibits large temperature variations, the filter block is likely to be the dominant source of error introduced through a thermally unstable calibration. To establish the validity of this assumption I adjusted the thermal-radiative model to run as before, but with a limit that the filter block has been limited to a change of 0.01K across the 300s from the beginning of

each observation, matching the requirement stability value. The results of this modelling, with a thermally stable filter block, are shown in Table 3-3.

The stabilisation of the filter block temperature has considerably reduced the calibration error. For the compositional channels it has reduced the error by an average of 67% for Observation 1, 70% for Observation 2 and 67% for Observation 3. This artificially stabilised performance is representative of a physical option of operating LTM.

LTM Channel	Channel Bandpass / μm	Maximum NEDT in Observation 1 /mK	Maximum NEDT in Observation 2 /mK	Maximum NEDT in Observation 3 /mK
A1	6.875 – 7.125	404	210	541
A2	7.125 – 7.375	388	222	519
A3	7.375 – 7.625	381	235	513
A4	7.625 – 7.875	374	245	505
A5	7.875 – 8.125	361	256	490
A6	8.125 – 8.375	372	270	510
A7	8.375 – 8.625	455	326	639
A8	8.627 – 8.875	452	336	638
A9	8.875 – 9.125	455	348	647
A10	9.375 – 9.625	516	407	751
A11	9.875 – 10.125	562	457	833
A12	6.25 – 12.5	57	127	140
A13	12.5 – 25	81	154	134
B1	25 – 50	256	358	234
B2	50 – 100	1651	2008	2706

Table 3-3: The maximum errors calculated in the thermal-radiative model with a filter block stabilised to a 0.01K temperature change, measured in each channel for each observation in the orbit. Errors are given as NEDT values in mK with the same colour scaling as in Table 3-2. The general shift to green shown the error reduction.

The instrument has been designed with two thermo-electric coolers (TECs) which can heat or cool a component by passing a voltage through them. These are mounted to the detector block and filter assembly block respectively. The clear conclusion from these results is that in the case of significant non-linear temperature variation, the use of TECs to stabilise the filter assembly block would provide an improvement in accuracy of up to 70% in the compositional channels (see Table 3-3). This thermal analysis was expanded

upon to define the power required for this TEC and thus choose a suitable model. I bought several of these TECs, which have been implemented into the LTM flight model and flight spare, however time constraints during the instrument ground campaign mean that they have not been tested in the instrument. The use of TECs to stabilise temperature is not without risks. Since the instrument is highly sensitive to temperature changes, using the TEC must create a smoother temperature profile, and without testing it has not yet been verified that this would be the result. Once the spacecraft has been launched, there will be opportunity to test the TEC's in flight and verify their performance relative to the thermal environment of the instrument while in space. Further performance analysis can be verified using the LTM flight spare in the Oxford calibration facility.

This modelling verifies that it may be a necessity to utilise the TECs to achieve the instrument's thermal stability requirement (Table 2-3). Furthermore, this acts as a requirement on the temperature sensing, since achieving that temperature stability of 10mK requires measurement of temperature to better than 10 mK.

3.1.2.5 Relationship between Blackbody PRT uncertainty and radiometric accuracy

An additional important application of the thermal radiometric model is to analyse the relationship between internal temperature measurements and final radiometric accuracy. The model allows us to derive the error in calibrated radiance that arises from a given error in measuring the temperature of the internal blackbody target. This is achieved by simulating the calibration sequence, including each calibration and scene view, while varying the blackbody's temperature. I have simulated the calibrated radiance using a 'true temperature' and a 'measured temperature' with a small error. By adjusting this error and recording the difference in calibrated radiance, the model provides a tool to quantify how temperature measurement uncertainties translate into radiometric inaccuracies. This is crucial for determining the overall radiometric accuracy of the instrument, as specified in the LTM requirements.

The output of this simulation is shown in Figure 3-7. It demonstrates how temperature uncertainty - stemming from measurement accuracy, electronics, onboard readouts, and thermometer calibration - can be related directly to radiometric uncertainty. This is essential for the development of the temperature measurement pipeline, as it informs us of how each uncertainty translates through to the radiometric requirement.

Improving the accuracy of LTM's temperature measurements is a focus of later sections in this chapter. Section 3.3 covers internal temperature measurements, and Section 3.4 addresses the calibration of the blackbody target. These results help contextualise temperature uncertainties within the overall radiometric uncertainty budget.

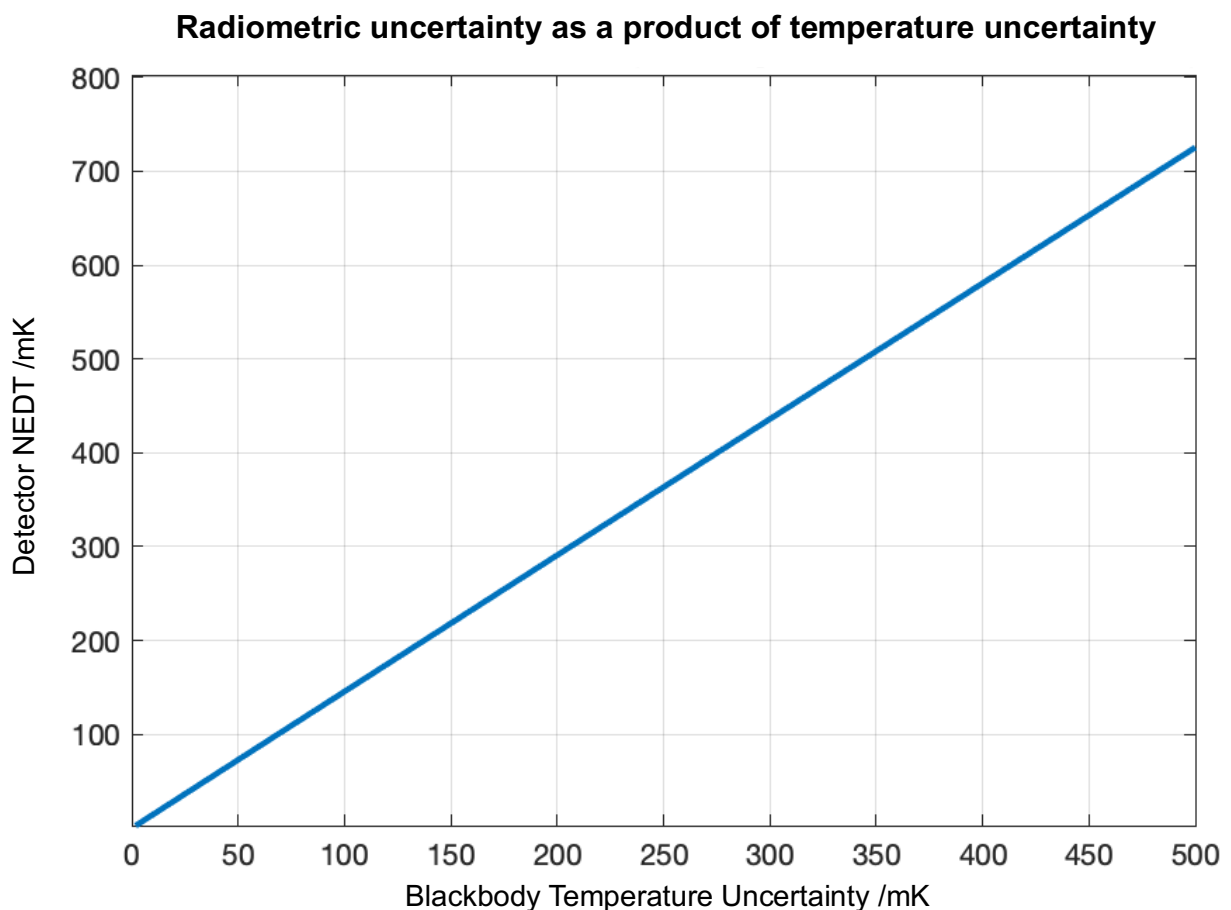


Figure 3-7: A graph relating the blackbody temperature uncertainty to radiometric uncertainty at the detector. The temperature uncertainty can be turned into a radiometric uncertainty post-calibration.

3.2 Characterising the LTM Optical Pathway

Following from the radiometric modelling of LTM in Section 3.1, it is necessary to consider the true transmission function of the instrument. The modelling results seen in the previous section describe the general performance of the instrument using idealised filter bandpasses and flat transmission, reflectivity and emissivity values. While these are useful for informing the development of the instrument and understanding the instrument operations, they are limited in their comparison to the true instrument performance. In order to calibrate the full instrument, we need an accurate model of the instrument. This section describes the formation of the full LTM transmission functions for each filter channel which are then used in the LTM data-pipelines, final calibration report and in future instrument modelling.

The filters can be split into two groupings, the interference filters and the mesh filters. The interference filters are comprised of the narrowband compositional filters and the first two thermal channels (A1 to A13), whereas the mesh filters refer to the final two, thermal broadband filters (B1 and B2). The interference filters were designed and created in-house at the Infrared Multilayer Laboratory (IML) in the University of Oxford and the mesh filters were provided by the Astronomy Instrumentation Group at the University of Cardiff. Mid-infrared (MIR) measurements of the LTM filter transmissions were produced by IML, with further far-infrared (FIR) measurements produced within this project to compare with the modelled FIR wavelengths. The transmission functions of the mesh filters across the full MIR and FIR regions were characterised by the measurements described herein. Further components within the LTM optical chain are also characterised to form a full optical feedthrough transmission function. These additional factors include the gold mirror reflectivity and the detector absorptivity. Due to time constraints direct measurements through the assembled LTM flight filter assembly were not created,

amplifying the importance of this work. There is a possibility of interference effects in the final transmission spectrum, however the full instrument model linearity testing, described in 3.2.5, verifies that the combined throughput function accurately describes the instrument's behaviour.

3.2.1 Spectral Measurements

All transmission measurements presented in this chapter were performed using the same method, outlined here. The measurements were taken with a Bruker VERTEX-70v spectrometer (Figure 3-8) in transmission mode.



Figure 3-8: Bruker VERTEX 70v Spectrometer, used for all transmission measurements detailed in this chapter.

Two detectors were used to cover the full wavelength range of interest. An RT-DLaTGS (Deuterated Lanthanum α Alanine doped TriGlycine Sulphate) detector enabled MIR measurements from less than 1 micron to 25 microns, while an RD-DTGS (Deuterated Triglycine Sulphate) FIR detector covered the FIR range, measuring out to 200 micrometres. Together, these detectors allowed us to measure the full bandpass range for

both the compositional and thermal filters. The source used was a Bruker MIR source, paired with a wide-wavelength-range MIR/FIR beamsplitter.

To measure transmission spectra, a filter sample mount was used, as shown in Figure 3-9. First, signal measurements were taken through the spectrometer with an empty sample mount, under vacuum, to serve as the background spectrum. Then the sample to be measured was then clamped into the sample mount, and the spectrometer was returned to vacuum conditions (measurements were only taken once the pressure had dropped below 6 hPa, where it would stabilise). The signal spectrum was then measured through the filter.




Figure 3-9: Photograph of spectrometer sample chamber. The filter being measured is clamped into the filter sample mount and held at the spectrometer's intermediate focus.

For post processing, the spectrometer data was extracted and turned into a csv format using a custom MATLAB script. The transmission function can be then derived as the ratio of the two measurements using the following formula:

$$Transmission_{sample}(\lambda) = \frac{Signal_{background}(\lambda)}{Signal_{sample}(\lambda)} \quad (3.23)$$

where $Signal_{background}$ and $Signal_{sample}$ are the respective measurements obtained from the spectrometer.

Reflectivity measurements were performed in a similar manner. A reflectivity jig was used to reflect the incident light at either a small angle (for diffuse reflectivity) or a high angle (for specular reflectivity). Measurements were taken through the same jig for both the sample and a reference material with a well-known reflectivity, typically gold reference mirrors due to their high and uniform reflectance in the infrared region. Reflectivity was then calculated using a similar ratio approach to the transmission measurements.

3.2.2 Interference Filters

3.2.2.1 Filter Design

The interference filters in LTM are formed from stacked layers with a Zinc Sulphide substrate. These filters are intended to transmit a narrow band of radiation while blocking radiation outside the specified bandpass. However, they exhibit long-wavelength leakage, which needs to be carefully analysed to ensure that it is minimised. Spectral measurements are combined with IML's predictions to understand how real-world performance deviates from the idealised filter behaviour.




Figure 3-10: The laboratory-measured LTM interference filter transmission functions over a 5-200 micrometre range. These spectra are the product of combining the MIR and FIR spectra. The long wavelength leaking can be observed as everything, except the green spectrum, to the right of 12 microns.

The interference filter bandpasses are given in Table 2-2 and are fully specified in LTM-OXF-SP-044. They have all been measured over a full wavelength range (5 – 200 microns) across two sets of measurements I made in Oxford. These spectra, combined and plotted in Figure 3-10, show that each filter successfully selects a narrow wavelength range, but also displays significant long-wavelength leakage.

3.2.2.2 Long Wavelength Leaking

To limit transmission only to the desired wavelengths, the filters are 'windowed' in the filter assembly using additional blocking materials. Originally, the LTM design included both a Potassium Chloride (KCl) window and a Barium Fluoride (BaF_2) window to mitigate long-wavelength leakage. Transmission spectra for two versions of the KCl window (the flight window would have been parylene coated for protective properties) and the BaF_2 window are shown alongside the compositional filters in Figure 3-11. The BaF_2 window provides a cutoff above 15 microns, while the KCl window blocks radiation above ~ 30 microns (the dataset shows the long wavelength blocking beginning at 25 microns but does not extend into the longer wavelength range. Filter A13 would have only been covered by the KCl window.

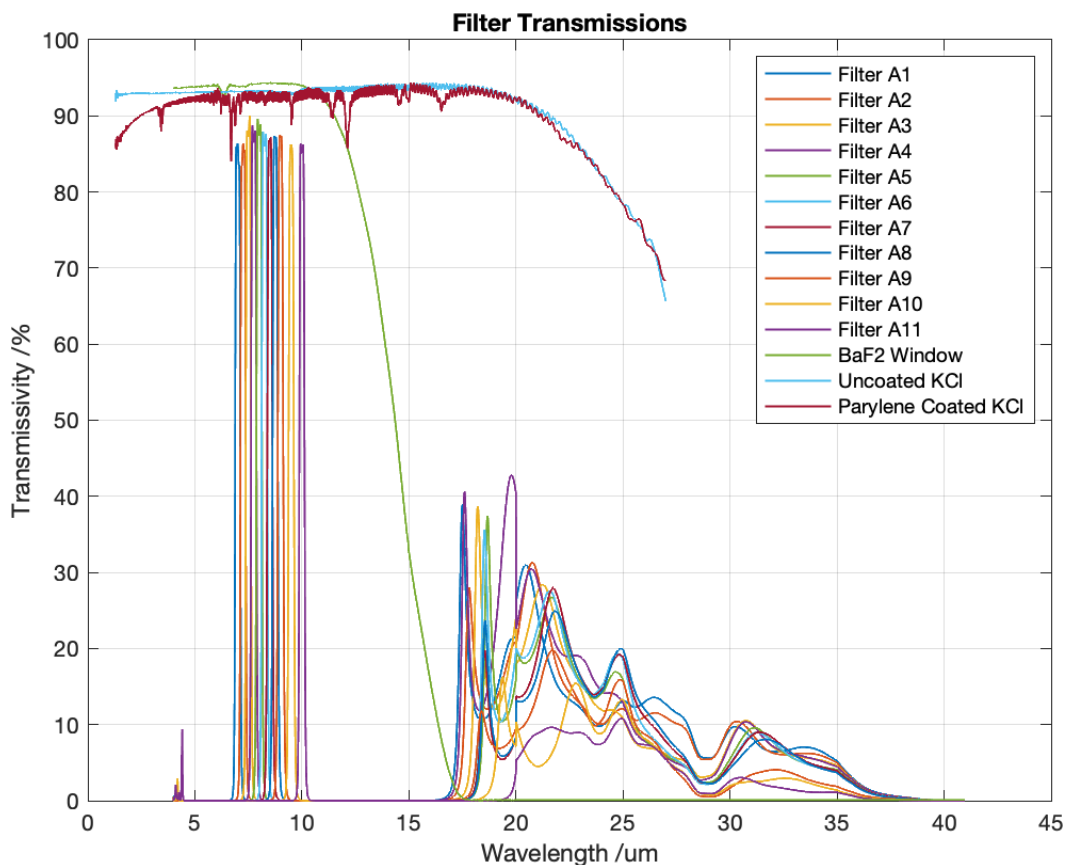


Figure 3-11: The compositional filters with the transmission filters of the windowing materials that are used to prevent long wavelength radiance leaking. The BaF_2 spectrum was measured by me, the KCl spectra were provided by IML for this analysis.

During vibration testing prior to LTM's delivery, it was discovered that the large crystal panes, particularly the KCl window, posed a significant risk. Relative motion during testing caused abrasion, which degraded transmission clarity, and one of the windows cracked. As a result, the design was modified to remove the KCl window entirely, leaving filter A13 without a window, and relying solely on the BaF₂ window for the remaining filters. I conducted an analysis to verify that this new configuration would adequately block long-wavelength leakage in the relevant channels.

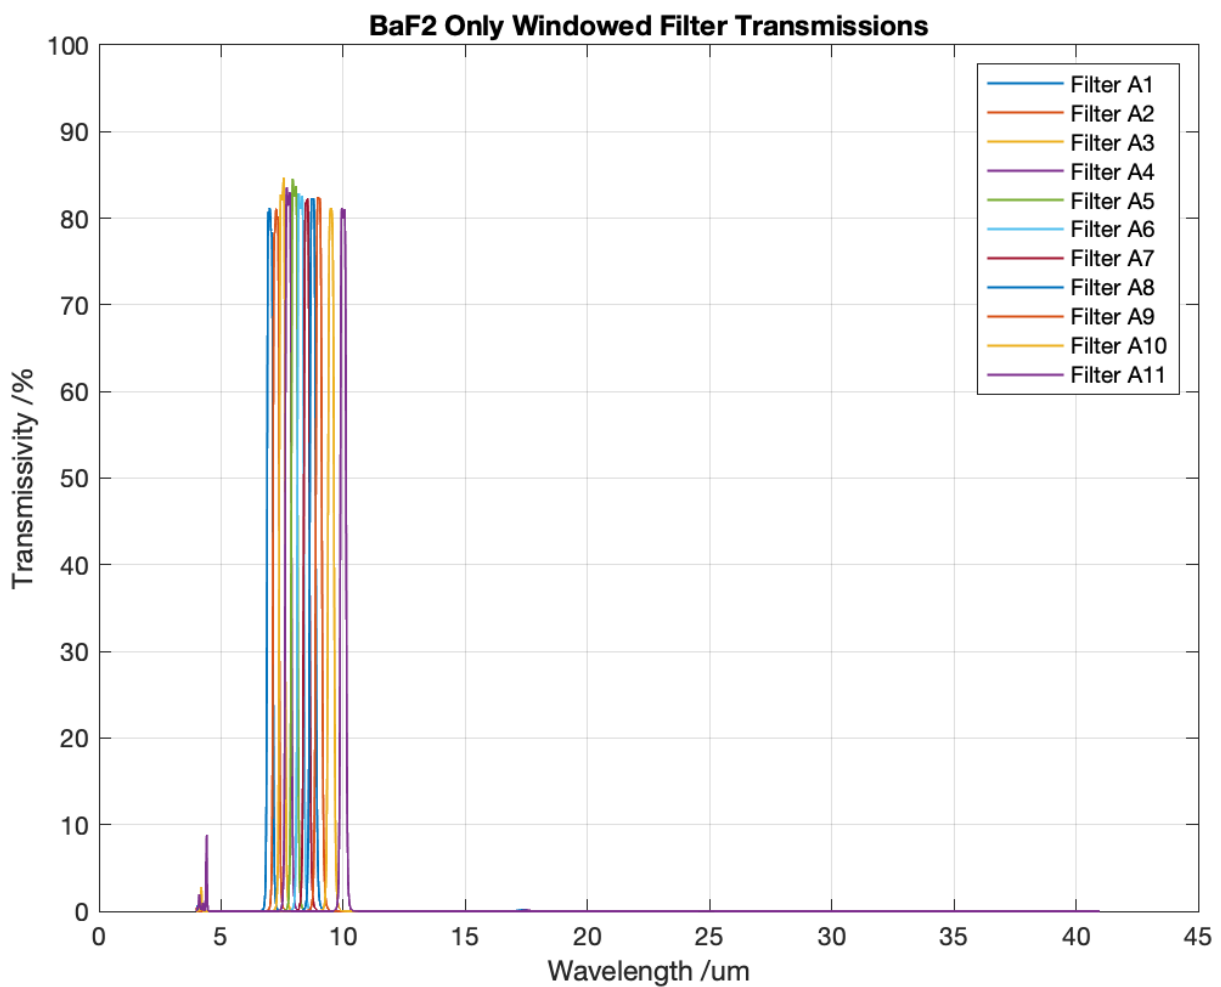


Figure 3-12: The transmission functions of each filter A1-11 when combined with a 50 micrometre Barium-Fluoride window. These transmission functions have been achieved by combining the transmissions, rather than by taking further measurements of combined filter stacks. A small amount of short wavelength leaking is present but accounts for a very small radiance throughput. Slight residual long wavelength leaking can be observed at 17 microns, and the magnitude of this is examined in Table 3-4..

3.2.2.3 Windowed Filter Performance

The resulting transmission functions formed by convolving the filter and window functions are shown in Figure 3-12. This figure clearly illustrates that the BaF₂ window effectively reduces long-wavelength leakage immediately after the filters. Leaking becomes less significant at progressively higher wavelengths due to the shape of the Planck function, and therefore it is most important to minimise leaking in the MIR. The BaF₂ windowed transmission functions have also been plotted in Figure 3-13 to demonstrate the relative scale of the long wavelength leaking. The plot shows that there are leaks in several subsequent areas, with the largest leaks peaking at a transmissivity below 0.3%. the following modelling serves to investigate the impact of the minimal remaining leaking on the total radiance received at the detector:

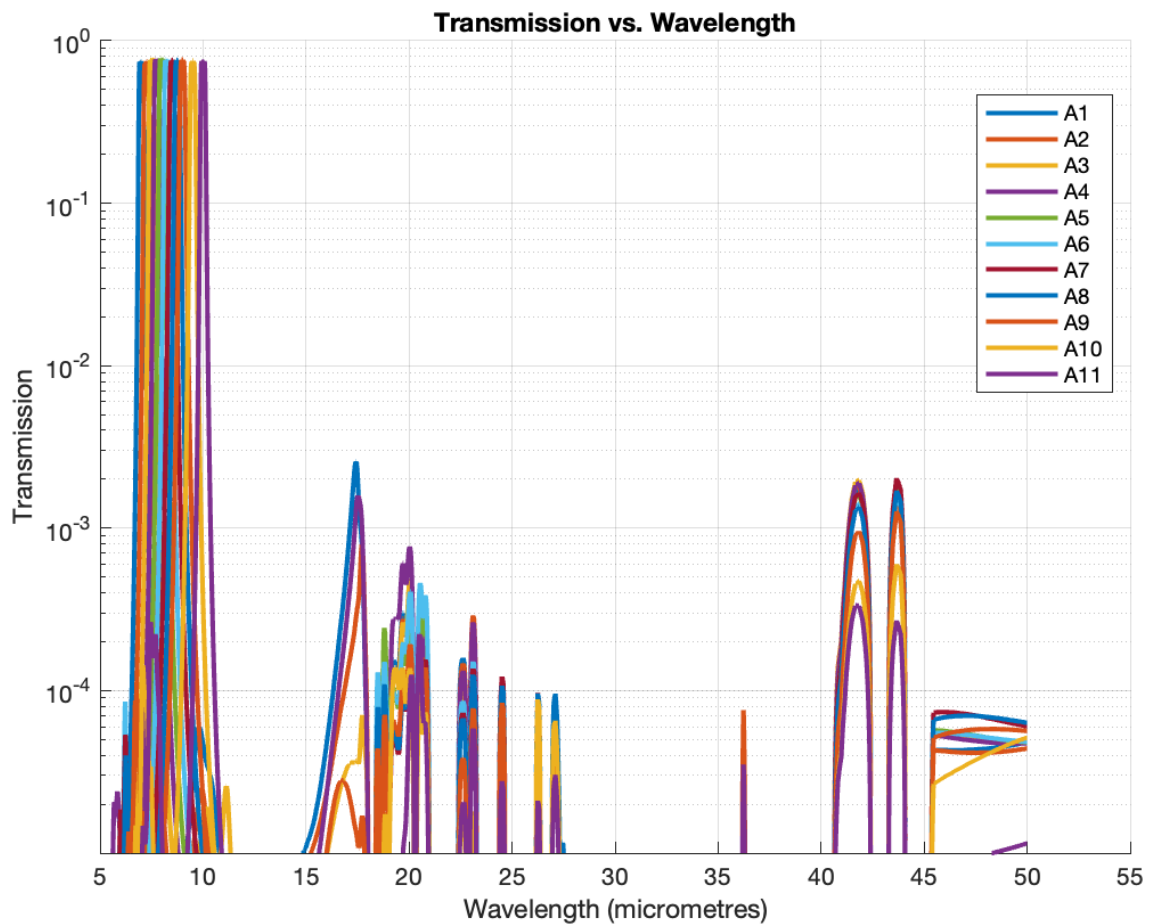


Figure 3-13: The final full-stack LTM transmission functions for each compositional filter, plotted on a logarithmic y-axis to demonstrate the relative transmission scales of the long-wavelength radiance leaks. The impact of the long wavelength leaking is modelled in this section and results are detailed in Table 3-4.

To quantify the impact of the BaF₂ windowing, I used my radiometric modelling tools to evaluate the transmission behaviour and assess the impact of long-wavelength leakage on the instrument's performance. To quantify the effect of the windowing, the ratio between out-of-band power (leakage) to in-band power (desired signal) was calculated. This was done using the following formula, applied over the windowed and unwindowed transmission spectra:

$$\frac{\text{Out - band power}}{\text{In - band power}} = \frac{\int_0^{\lambda_1} \tau(\lambda) \epsilon_s(\lambda) B_\lambda(T_s) d\lambda + \int_{\lambda_2}^{\infty} \tau(\lambda) \epsilon_s(\lambda) B_\lambda(T_s) d\lambda}{\int_{\lambda_1}^{\lambda_2} \tau(\lambda) \epsilon_s(\lambda) B_\lambda(T_s) d\lambda} \quad (3.24)$$

Where λ_1 and λ_2 are the upper and lower wavelengths for the bandpass which is being modelled. A relatively low temperature (270 K) blackbody was used for a worst-case

Channel	Ratio of out-band power to in-band power /%	Ratio of out-band power to in-band power without the blocking filter /%
A1	1.02	584
A2	0.56	495
A3	0.41	441
A4	0.65	528
A5	0.31	322
A6	0.29	317
A7	0.28	275
A8	0.26	257
A9	0.17	209
A10	0.17	143
A11	0.11	101

Table 3-4: Ratio of out-band to in-band powers arriving at the detector. Calculated for a 270K surface temperature observation, with and without the blocking filter.

analysis, since at lower temperatures the Planck function has a higher proportion of radiance in the longer wavelength regions and thus leaking in this area would have a larger effect.

The results of modelling Equation 3.24 for each of the compositional channels are shown in Table 3-4. The power at the detector from the leaking is small (<1%) compared to the power from in-band power so the out-band power will have a marginal effect on the total detector power recorded. The second column shows how the detector would be dominated by out-band power if the blocking window was not present. This result shows that even in a worst case the impact of leaking is at a fraction of a percent for most of the compositional channels. This analysis includes short wavelength leaking which increases the out-band radiance. This further amplifies the evidence from Section 3.1.1.3 that the compositional channels will only be accurate in high temperature regions where the size of emissivity features will be magnified compared to uncertainty sources.

3.2.3 Mesh Filters

The broadband thermal filters B1 and B2 in LTM are constructed from multiple layers of individual filter meshes, which I will refer to as 'meshes'. The individual meshes are stacked together to create the full mesh filter transmission function. Measurements of each mesh in isolation were performed in the Oxford spectroscopy laboratory, and these spectra were used to simulate the transmission functions of the full filter stack. Additionally, a study was conducted to assess the influence of mesh stacking order on the final transmission function.

The stack of specific constituent meshes combine to achieve the desired bandpass. The individual and combined transmission spectra are presented in the following sections. All measurements were taken using the FIR detector, and the specified bandpasses for the mesh filters are given in Table 3-5.

Filter	Bandpass range /μm	$\nu_{\text{min}} / \text{cm}^{-1}$	$\nu_{\text{max}} / \text{cm}^{-1}$
B1	25 - 50	200 ± 20	400 ± 40
B2	50 - 100	100 ± 10	200 ± 20

Table 3-5: Mesh Filter bandpass requirements as defined in LTM-OXF-SP-041.

3.2.3.1 Mesh Band 1

Mesh Band 1 (B1) is designed to cover a 25–50 micron bandpass, and is constructed from four individual meshes, labelled FP3096, FP3097, FP3123, and FP3115. These labels will continue to be used throughout this discussion. The individual transmission spectra for each of these meshes are shown in Figure 3-14. By examining these spectra, we can observe the contribution each mesh makes to the final bandpass of Mesh B1, for example FP3096 clearly provides a short wavelength cutoff, while FP3115 defines the long wavelength end of the spectrum. FP3123 provides short wavelength




Figure 3-14: Mesh Filter B1 individual mesh transmission functions. Spectra generated in the Oxford laboratory using the FIR detector.

leaking coverage. The combination of these individual meshes into a stack forms the overall bandpass necessary to meet the specified design parameters (Section 2.4.1).

3.2.3.2 Mesh Band 2

Mesh Band 2 (B2) is designed to provide a 50–100 micron bandpass. This is achieved by stacking three individual meshes, labelled W3034, FP3098, and FP3099.

Figure 3-15 presents the individual transmission spectra for each of these meshes.




Figure 3-15: Mesh Filter B2 individual mesh transmission functions as measured in the Oxford spectroscopy laboratory.

3.2.3.3 Mesh Filter Stacking

Since the full bandpass of each mesh filter is achieved by stacking multiple meshes, it is important to understand how the stacking order affects the final transmission function. Although mesh filters operate by selectively transmitting certain wavelengths, the distance between the filters is comparable to the wavelength of light passing through and this, combined with the surface properties of the filters can form an interferometer-like effect, shifting the transmission function.




Figure 3-16: The measured transmission functions through Mesh Band 1 for each of the different stack orders.

To investigate this, I measured the transmission spectra of all possible stacking sequences for Mesh Band 1, which consists of four meshes that can be arranged in six different sequences. The results are shown in Figure 3-16, where slight variations between the transmission functions can be observed.

The short-wavelength edge of the bandpass remains consistent across all stack orders, as this is determined primarily by a single mesh. However, the transmission within the bandpass shows clear differences between stacks. For example, Stack 6 shows a local minimum at ~ 70 microns, while Stack 2 exhibits a peak in transmission at the same point. These variations are unlikely to be artifacts of spectrometer drift between measurements, as such drift would produce uniform transmission offsets rather than the localised differences seen here.

Despite these variations, the overall bandpass width and position remain consistent across all stack configurations, and the average transmissivity is comparable. As a result, no significant performance advantage can be gained by choosing a particular stacking order. This suggests that while mesh interaction leads to subtle changes in transmission characteristics, the overall performance of the mesh filter is robust to changes in the stacking sequence.

3.2.4 Instrument Transmission Functions

For most radiometric modelling purposes, we want to relate the radiance spectrum entering the instrument to the radiance spectrum received at the detector. Therefore, an overall 'instrument transmission function, is most useful. The overall transmission function of the LTM instrument is determined by the combined effects of the mirrors' reflectivity, the filters' transmissivity, and the detector's absorptivity. The total instrument transmission function is given by:

$$\tau_{instrument}(\lambda) = \tau_{filter}(\lambda) * R_{Mirror}(\lambda)^6 * A_{Detector}(\lambda) \quad (3.25)$$

Where τ , R and A refer to the transmissivity, reflectance and absorptivity respectively.

3.2.4.1 Other Optical Components

The absorptivity of the INO detector was determined by measuring its specular reflectance. Since the detector is smooth it can be assumed that radiation is predominantly reflected in a specular manner and little radiation is lost to diffuse reflection. In addition, given the opaque surface it was assumed that all non-reflect light is absorbed with no transmittance. Therefore, $absorptivity = 1 - reflectance$, due to the minimal reflective scattering. Under this assumption, the absorptivity of the detector was derived and is plotted in Figure 3-17. The detector exhibits near-unity absorptivity across the compositional channels, though performance slightly degrades at longer wavelengths.




Figure 3-17: The absorption of the INO detector measured through the spectrometer. The absorptivity is near one across the compositional channels, and the performance degrades slightly at longer wavelengths.

The surface of the LTM gold mirrors cannot be measured using the reflectance setup described in Section 3.2.1, since they are not flat and therefore will not provide a clear specular reflectance. Therefore, I used external measurements from ThorLabs to

model the mirrors' reflectivity behaviour. These measurements, shown in Figure 3-18, provide the reflectivity spectrum of the mirrors. While sources such as LBOptics report that gold mirrors maintain 99% reflectivity out to 100 microns, the measured spectra are not provided. For a conservative estimate, I assumed a worst-case reflectivity of 98% beyond the bounds of the ThorLabs data.




Figure 3-18: Gold mirror reflectivity spectrum taken from public ThorLabs data.

3.2.4.2 Aggregate Instrument Transmission Functions

With the detector absorptivity, mirror reflectivity, and filter transmissivity characterised, the complete instrument transmission function for each filter channel was assembled. These transmission functions describe how the radiation incident on the instrument's aperture is transmitted to the detector, accounting for losses at each optical component. Figure 3-19 presents the full instrument transmission functions for all filter channels, combining the effects of all components.




Figure 3-19: The full LTM instrument transmission functions for each filter channel.

These transmission functions are critical for further performance modelling of the instrument and will be used extensively during the data analysis and retrieval processes. They will also play an essential role in post-processing LTM data from the calibration campaign and in preparation for the launch of the Lunar Trailblazer mission. When LTM returns in-situ measurements of the lunar surface, these functions will provide the foundation for accurate data interpretation.

3.2.5 Verification of Instrument Transmission Function

The accuracy of the instruments radiometric model and calibration procedure is assessed during the ground calibration testing by means of a linearity test. During the ground campaign, LTM took measurements of an external blackbody target at various temperatures ranging from 200 K to 350 K. For each calibration, the radiance of the

blackbody target is determined within each channel following the calibration pipeline. Additionally, the radiance is independently derived from the PRT sensor temperatures in the external blackbody, using the LTM instrument model, based on the transmission functions discussed earlier in this section.

By plotting the measured radiance against the model-derived radiance for each channel, we should expect a 1:1 linear relationship for a well-calibrated instrument. This comparison serves as a critical test of the accuracy of the optical model and transmission functions I developed for LTM.

This plot, displayed in Figure 3-20, was created for the LTM Calibration Report by another member of the LTM team using the transmission data I provided (Evans, 2023). The linearity of each channel is evaluated using a metric called integral non-linearity (INL), defined as the maximum fractional error across the full scale. The LTM microbolometer has a requirement for linearity better than 5%, which the plot shows was achieved for every channel when using the LTM model.

The relatively higher INL observed in the compositional channels can be attributed to noise in the measured radiance at lower temperatures, where signal levels are substantially reduced. Despite this, the linearity results demonstrate that the transmission functions I have developed in this chapter provide an accurate representation of LTM's behaviour.




Figure 3-20: LTM radiance linearity test. Channels are listed in their filter assembly order. Red lines indicate the noise equivalent power. Linearity is quoted as INL. Plot is recreated from the LTM calibration report (Evans, 2023).

3.3 Internal Temperature Measurements

As discussed in Section 3.1.2.1 LTM's calibration relies on the radiance emitted by the hot calibration view, which is provided by the internal blackbody target. Therefore, accurate characterization of the blackbody's radiance spectrum is critical to the instrument's performance. LTM's design requirements state that the instrument must incorporate an internal calibration target capable of delivering radiometric accuracy to within ± 1 K (Table 2-3). Therefore, the uncertainties associated with the blackbody target

- including its physical properties and the accuracy of in-flight temperature measurements
- must remain within this limit. Given that the effective emissivity of the target is close to unity due to the Nextel paint used and the blackbody's geometry, the precision of the target radiance is primarily determined by the accuracy of the temperature measurements.

The thermal calibration of the LTM blackbody is discussed in Section 3.4; this section focuses on the electronics used to measure temperatures within LTM. Temperature measurements are carried out using Platinum Resistance Thermometers (PRTs), with the data read out through on-board electronics. In this section, I will describe the methodology of PRT-based temperature sensing and outline the work done to improve and test the LTM's PRT readout electronics.

A full PRT calibration procedure, as well as full data analysis is provided first. This provides an overview for the process and discusses the derivation of calibration constants based on experimental results.

3.3.1 The PRT Calibration Procedure

In this section, I will discuss the use of PRTs for temperature measurement and the procedures used to calibrate them for improved accuracy and precision. This includes an example of a calibration procedure used to evaluate different PRT mounting options.

PRTs are mounted throughout LTM and measure temperature by varying their electrical resistance in response to temperature changes. A PRT is formed from a small piece of platinum wire with a highly linear and repeatable temperature dependence in resistivity. The PRTs used have a 1000 Ω resistance at 0 °C, and since the resistance of the PRT is approximately linear with temperature, the instantaneous resistance is calculated by applying a known current to the PRT and measuring the resulting voltage drop.

Simultaneously, the current is measured by recording the voltage drop across a reference resistor. The PRT's resistance is determined by the following equation:

$$R_{PRT} = R_{ref} \times \frac{V_{PRT}}{V_{ref}} \quad (3.26)$$

Where V and R denote voltage and resistance respectively and the subscripts PRT and ref refer to the values for the PRT and for the reference resistor respectively.

Temperature can be calculated from the R_{PRT} by means of a calibration curve, this is typically a polynomial function which has been fitted to measurements taken at known temperatures.

$$T_{PRT}(R_{PRT}) = \sum_{k=0}^n A_k R_{PRT}^k \quad (3.27)$$

3.3.1.1 PRT Calibration

The method of mounting PRTs is crucial for accurate temperature measurements, as improper coupling will permit a temperature gradient between thermometer and target, and mechanical stresses will alter the calibration coefficients. To ensure optimal performance, two mounting procedures were tested: PRTs glued directly into drilled holes using thermally conductive epoxy, and PRTs mounted into holders bolted onto a testing block with thermal grease (Apiezon N) applied between the mount and the block.

Ten PRTs were used, split evenly between the two methods, and mounted onto a testing block. The block was then subjected to thermal cycling to evaluate the stability and performance of the mounting techniques. The experimental setup is shown in Figure 3-21.

The calibration of the PRTs was performed using a testing block coupled to a temperature-controlled baseplate within a vacuum chamber. The block was housed in a shroud, which was then covered in MLI to create an isothermal environment. The baseplate and shroud were equipped with Rhodium-Iron (RhFe) thermometers with well-known calibration curves, used as reference points with an accuracy of ± 8 mK. The shroud creates a uniform thermal environment, and the PRTs were calibrated against the RhFe thermometers. Therefore, the absolute maximum precision attainable depends on the uncertainty of the RhFe sensors.




Figure 3-21: Photographs of the test block used for thermometer calibration. A shows the test block on the baseplate and the sensor mountings can be observed. The two mounting types, glued into a mounting block or glued directly into the test block can be discerned. In B the thermal shroud has been placed over the top and bolted into the baseplate.

The baseplate temperature was controlled by balancing the cooling power of an attached cryocooler and the thermal input from resistance heaters. The cryocooler operates in a constant mode, and heater powers were set to achieve different temperature set points. Once thermal equilibrium was achieved, the PRT resistance was recorded. This method is used, rather than a PID control loop since it achieves a more stable equilibrium and thus a more precise temperature. This is the preferred method for PRT calibration, but it comes at the cost of time since reaching thermal equilibrium in this manner takes a few days. This process was repeated at various temperatures to gather data for calibrating the PRTs.

Each PRT was supplied with a ~ 1 mA current, supplied in series, and the resulting voltage drop was measured. A 500.9 Ohm reference resistor was inserted in the series resistor chain and the voltage measured across it was used to monitor the current. The circuit setup is shown in Figure 3-22.

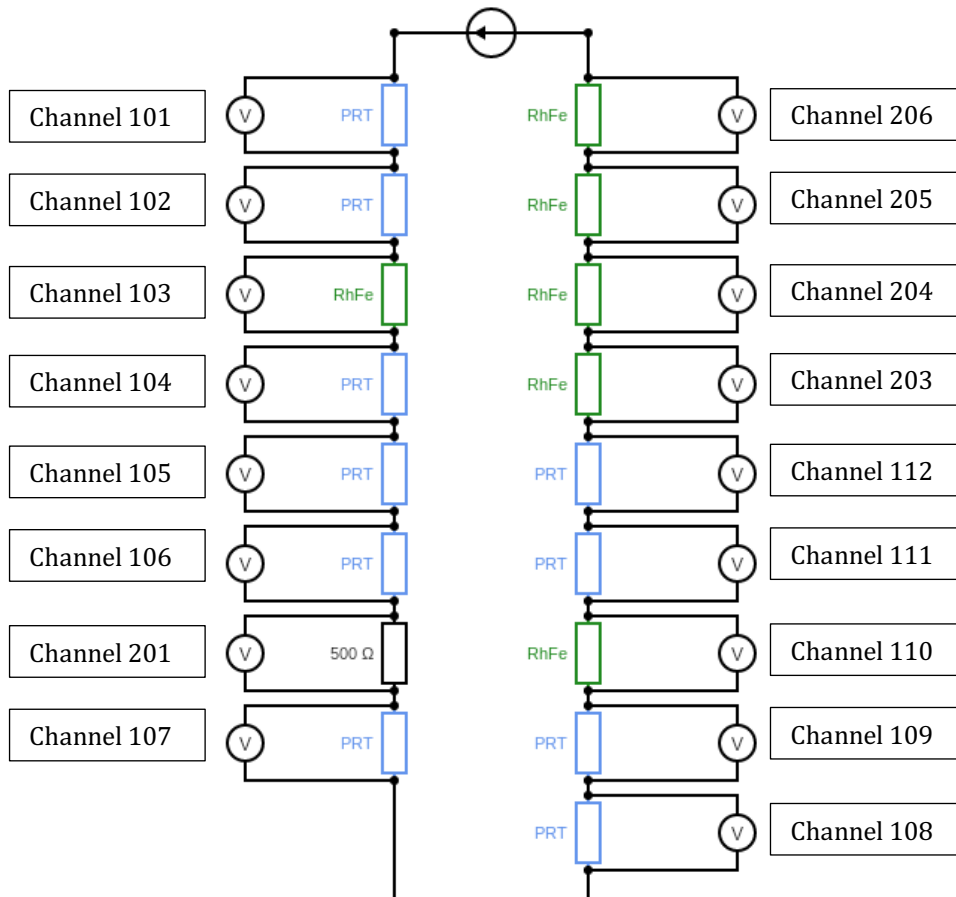


Figure 3-22: Circuit diagram of the testing apparatus. PRTs are coloured blue, RhFe reference thermometers green and the 500 Ω reference resistor black.

The stability of various PRT measurements at their setpoints is shown in Figure 3-23, where the steepest temperature gradient observed was less than 0.006 K/min. The most stable points in each dataset were used for calibration, as these correspond to the smallest thermal gradients within the test setup. These plots immediately reveal that two resistors have a slightly deviant performance with respect to the others as between the plots, as temperature increases, their resistance increases more than the other PRTs. This is why the bottom two plots, the hottest setpoints, have two lines separate from the rest.

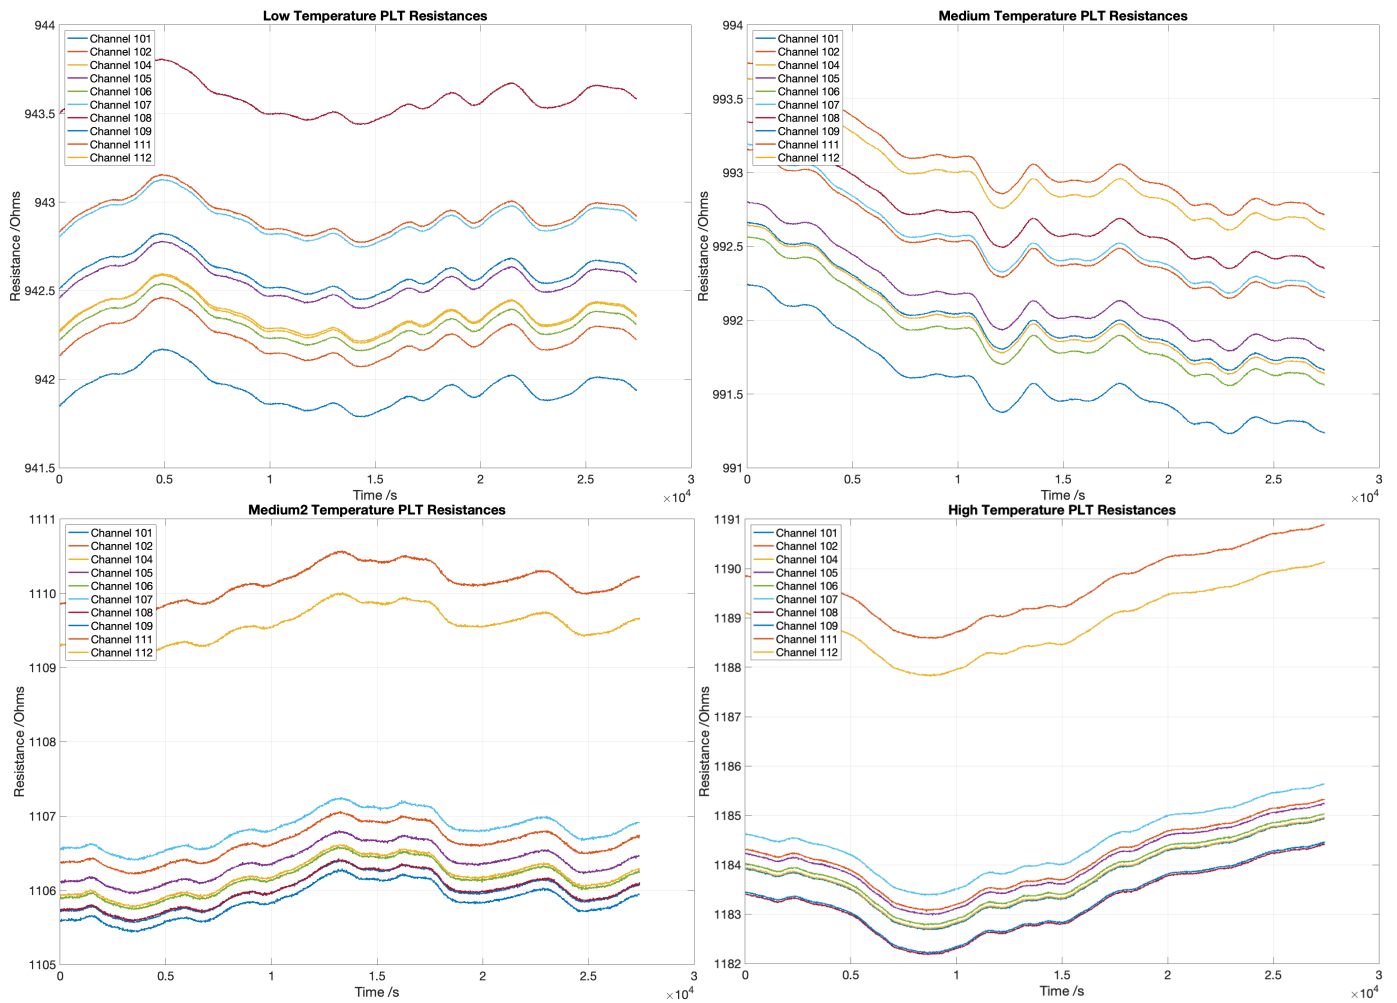


Figure 3-23: Graphs showing the PRT resistances across 8-hour recordings. These are reproduced to show the stability of the measurements. The temperatures maintained for each of these are: Top left – 257 K, Top Right – 272 K, Bottom Left – 276 K and Bottom Right – 323 K.

Calibration coefficients were derived in MATLAB using a four-point polynomial fit (Equation 8.2), with the calibration constants presented in Table 3-6 and the curves plotted in Figure 3-24 according to Equation 3.27. The very small values for A_2 in each channel show the near linear response of the PRT's to temperature. The calibration covers a temperature range of 255–335 K, ensuring accuracy within the expected operational range of LTM.

The accuracy of the temperature sensors is chiefly limited by the calibration process. Considering the fluctuations in the temperature within each reading (<0.35 K per hour) I decided to attempt to calibrate the PRT's to both the most stable and least stable points on the reading to compare the difference in result. The least stable points, for

example 1.2×10^4 s into the Medium Temperature Resistance plot in Figure 3-23, ought to give the least accurate calibration data. Calibrating to different points on the graph resulted in a measured temperature difference of less than 50 mK at 330 K, the expected operating temperature range of the instrument.

Sensor Channel	A_0	A_1	A_2
101	18.5 ± 0.1	0.2412 ± 0.0002	$(1.33 \pm 0.01) \times 10^{-5}$
102	24.7 ± 0.1	0.2342 ± 0.0001	$(1.36 \pm 0.01) \times 10^{-5}$
104	22.1 ± 0.1	0.2383 ± 0.0003	$(1.21 \pm 0.01) \times 10^{-5}$
105	18.2 ± 0.1	0.2412 ± 0.0004	$(1.34 \pm 0.02) \times 10^{-5}$
106	18.2 ± 0.1	0.2413 ± 0.0002	$(1.33 \pm 0.01) \times 10^{-5}$
107	19.4 ± 0.1	0.2387 ± 0.0002	$(1.45 \pm 0.01) \times 10^{-5}$
108	12.5 ± 0.2	0.2497 ± 0.0008	$(1.04 \pm 0.03) \times 10^{-5}$
109	17.1 ± 0.1	0.2424 ± 0.0005	$(1.33 \pm 0.02) \times 10^{-5}$
111	19.4 ± 0.1	0.2384 ± 0.0003	$(1.49 \pm 0.01) \times 10^{-5}$
112	19.7 ± 0.1	0.2384 ± 0.0002	$(1.48 \pm 0.01) \times 10^{-5}$

Table 3-6: The calibration constants for the PRTs as according to the polynomial function in Equation 3.26. Since the coefficients are the product of the fitting algorithm these uncertainties are interdependent.

The full uncertainty of this calibration technique will depend on temperature stability, reference resistor uncertainty and RhFe thermometer accuracy. These will all be considered in Section 3.4 where I discuss the most important thermal calibration I undertook as part of this project; the PRT's mounted into LTM's internal blackbody.

Once the accuracy of the PRTs is known, its impact on the wider calibration procedure can be evaluated using the thermal-radiometric model discussed in Section 3.1.2. Additional sources of temperature uncertainty come from the measurement and readout electronics.




Figure 3-24: Calibration curves for the PRTs prior to any testing. The divergence at high temperature shows some element of the procedure has affected some of the sensors. All of the lines follow the higher line, with the lower line corresponding to two of the PRTs which are glued into drilled holes (Channel 111 and Channel 112).

3.3.2 Housekeeping Board Development

The housekeeping board in LTM is the circuit board which contains the electronics through which internal PRT's are measured. It is essential for producing accurate measurements of internal temperatures, which are managed via the Command and Data Handling Unit (CDHU). A block diagram of the CDHU is provided in Figure 3-25. This unit interfaces with the spacecraft, manages both analogue and digital electronics, and allows the instrument to operate autonomously during its mission. Temperature data from 21 PRTs, distributed across key components such as the internal blackbody, filter assembly, detector assembly, and the instrument housing, is collected through the housekeeping board.

This section discusses the development, testing, and troubleshooting of the housekeeping board, starting with an engineering model and proceeding to the flight models. The testing revealed several issues with the original design, which led to important modifications that were implemented in the final flight hardware.

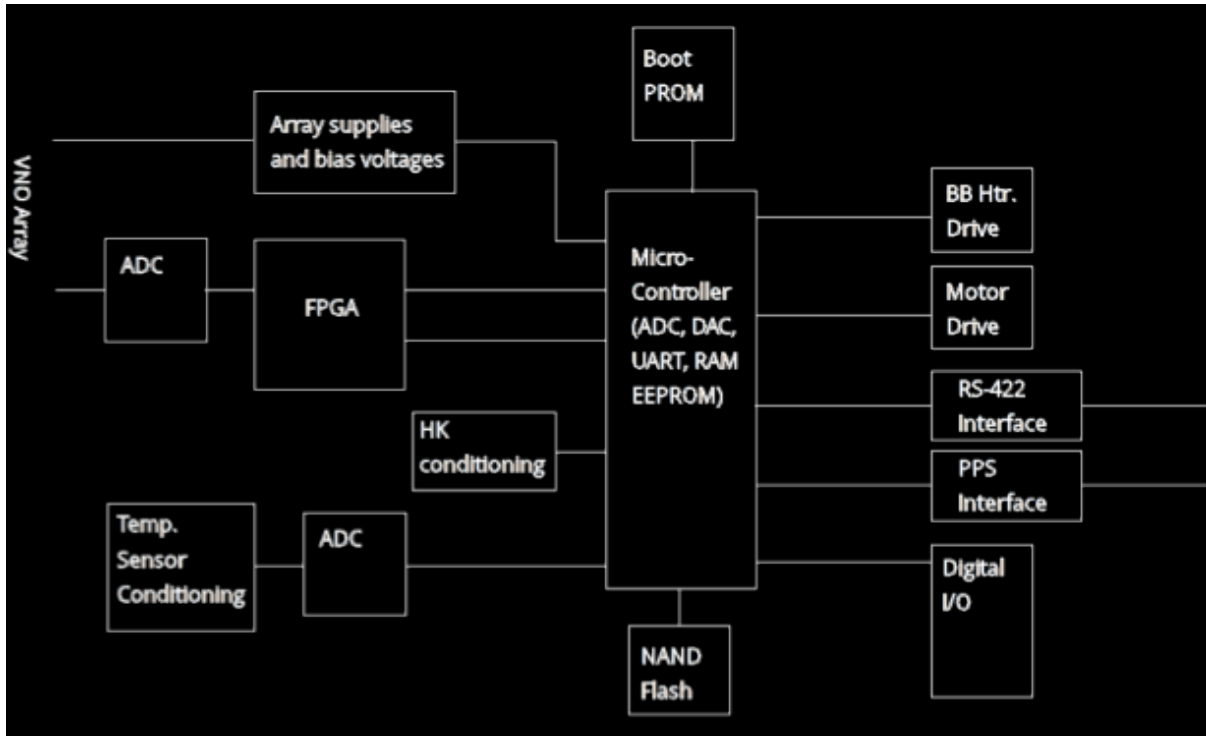


Figure 3-25: The LTM Command and Data Handling Unit (CDHU). The CDHU enables the instrument to operate autonomously from the spacecraft to take, process and store data.

I first tested the engineering model of the housekeeping board and found that errors within the design were preventing proper function. The voltage measurements through the board were lower than expected, and the measurements in some channels would produce the correct resistance, and in others they would not. This led to the improvement of the board design before development of the flight board. I then also tested the flight board to evaluate the level of accuracy for the resistance measurements that could be obtained through the housekeeping electronics to add to the uncertainty on the blackbody calibration.

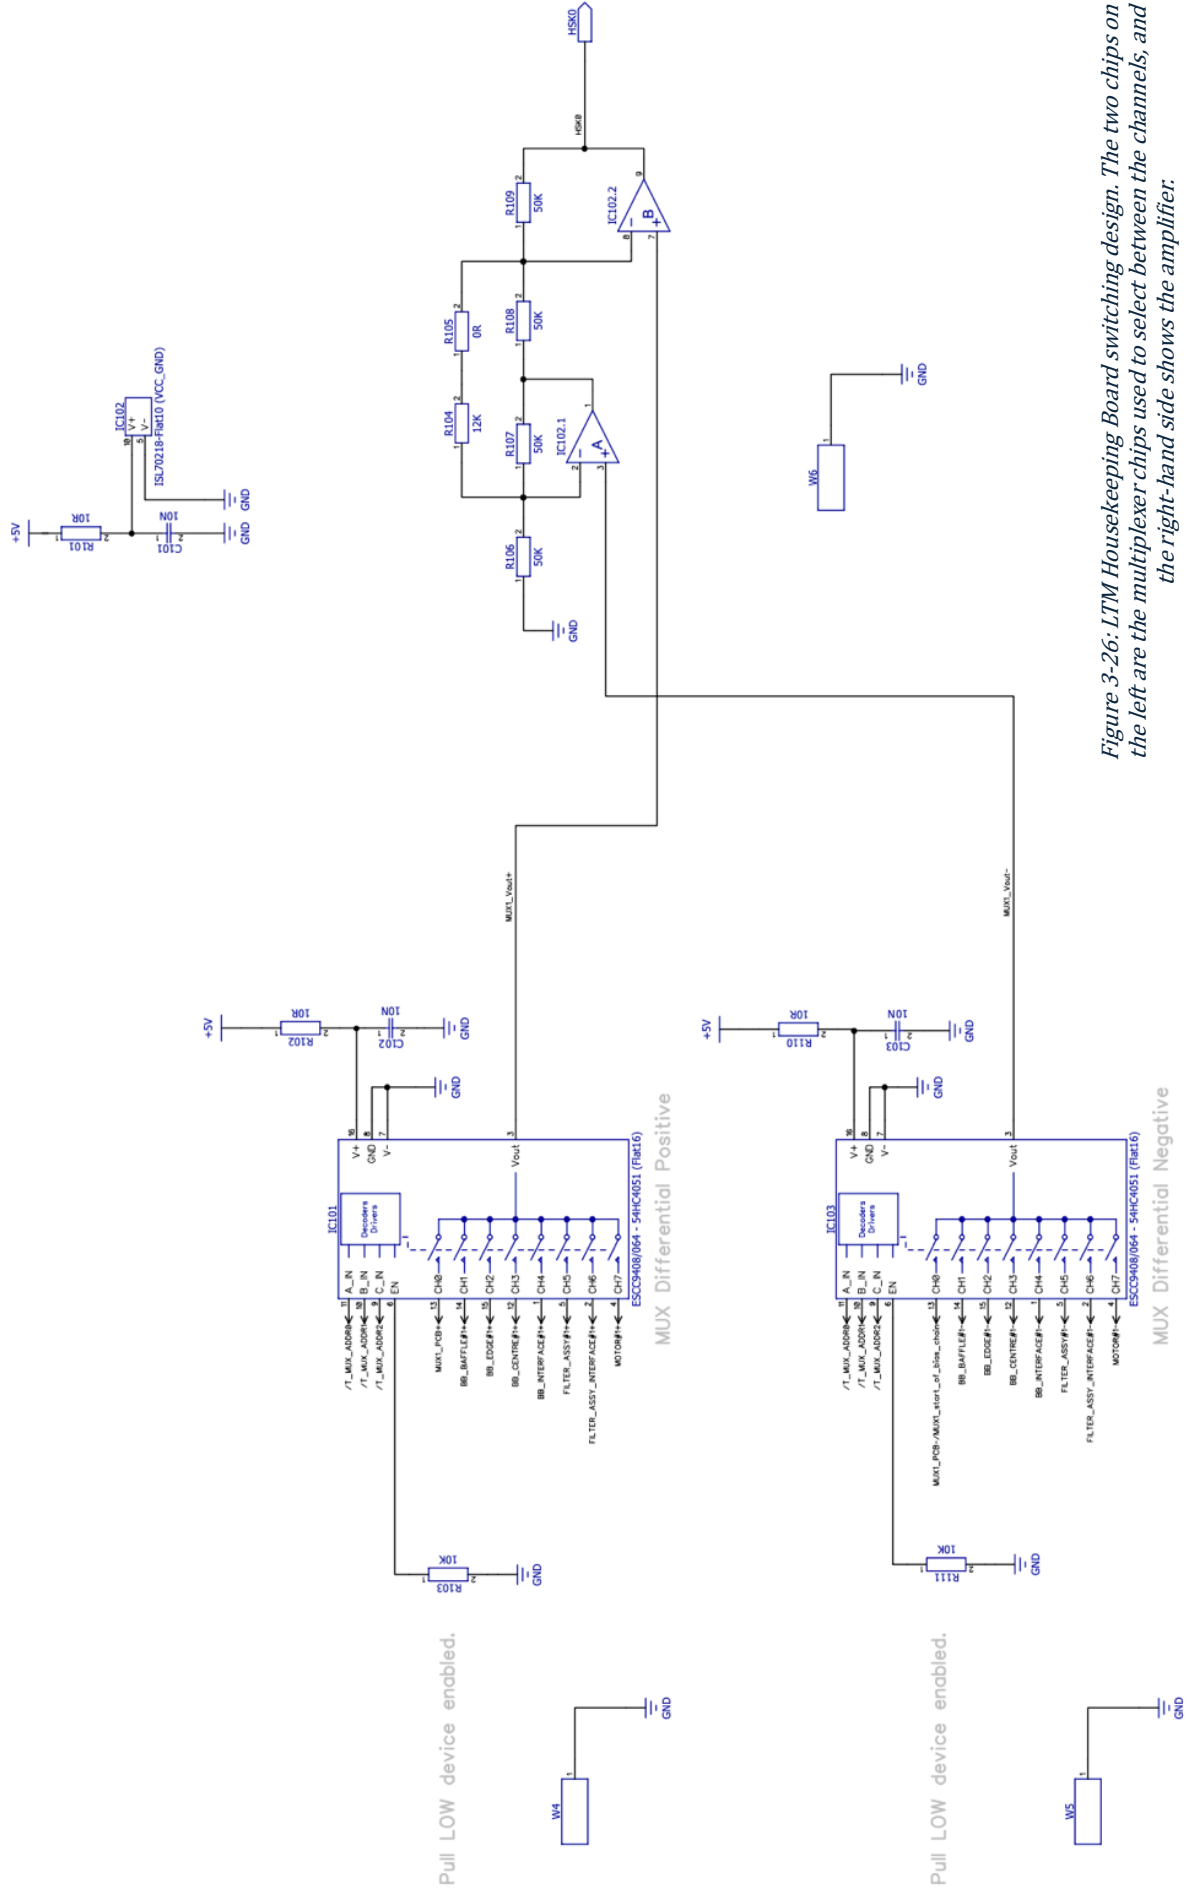


Figure 3-26: LTM Housekeeping Board switching design. The two chips on the left are the multiplexer chips used to select between the channels, and the right-hand side shows the amplifier.

3.3.2.1 *Housekeeping Board Engineering Model Testing*

The original design for the housekeeping board is shown in Figure 3-26, this is a heritage design from a previous space instrument designed at Oxford; the short period Seismometer (SEIS-SP) on the InSight Lander. The component temperature in LTM is measured by a set of 21 four-wire PRTs, measured using three sets of identical circuitries in order to maintain redundancy for key temperature measurements (Table 2-3). In each set, 7 PRTs are connected in a resistor chain between a temperature dependant current source (AD590) and ground. A 953 Ω reference resistor is inserted at the start of the chain and a 1 k Ω resistor between the last PRT and ground.

Two ECSS9408 multiplexer chips are used as a pair to select the PRT being measured. Each multiplexer has 8 channels which correspond to the reference resistor and the 7 PRTs in each chain. Each channel in one multiplexer is attached above each PRT in the chain and each channel in the other multiplexer is attached below. The multiplexers are controlled by the same three control lines, so they always have matching channels open so the voltage between the multiplexer outputs is the voltage across the PRT selected. Therefore, the voltage across each resistor can be measured. In order to increase signal-to-noise, the voltage between the multiplexer outputs is amplified by a constant value (10.3) before being read out at the analogue to digital converter (ADC). The ADC used to measure from the housekeeping board has 16 bits of digital resolution.

When the housekeeping board is in use, the control lines switch the multiplexers through the channels and the voltage across each is recorded. The current is calculated from the voltage drop across the reference resistor V_{ref} , and this value gives an indirect measurement of the board temperature. The resistances of each PRT, R_{PRT} , can then be calculated from the voltage drop across that PRT, V_{PRT} as in Equation 3.26. The common amplification between each voltage and the measured voltage at the ADC cancels here.

$$R_{PRT} = R_{ref} \times \frac{V_{PRT}}{V_{ref}} \quad (3.28)$$

This is effectively the same methodology as in the previous Section, except for the fact that the voltages measured are not the voltages across the resistors, due to the amplification, so when the current is calculated explicitly the voltage needs to be reduced by a factor of 10.3.

I tested these boards using an engineering model, shown in Figure 3-27, which had differing resistors ranging from 845 Ω to 1300 Ω in lieu of PRTs. The control lines were controlled via an Arduino Uno. The three voltage outputs were recorded using a Hewlett-Packard Data Acquisition Unit (DAQ) and the experiments were run via a MATLAB script which controlled the Arduino and recorded data from the DAQ. Power for the board was supplied by benchtop supplies.




Figure 3-27: The LTM housekeeping EM board under testing. The three sets of components can be seen running vertically down the image. Small copper wires show where modifications have been made to this board, which were then implemented in the final design.

Initial testing on the engineering models highlighted two functional errors. The first was that the 5 V line used to power the AD590 current supply was insufficient. This

was a result of doubling the number of PRTs in the resistor chain relative to the heritage design. The matter was solved by updating the engineering model to run off a +12 V line which had been provided to the board but was previously unused.

Multiplexer Channel	Resistor value / Ω	Measured resistance through housekeeping EM board / Ω
0	953	Reference Resistor
1	845	843.96
2	931	930.39
3	1000	999.62
4	1020	1019.39
5	1150	1149.68
6	1210	1133.95
7	1300	739.401

Table 3-7: This table compares the manufacturer resistor values with the resistance that was measured through the housekeeping engineering model board. The board clearly measures the resistances correctly for channels 1 through 5, and then a systematic error is introduced which has resulted in a reduced reading for channels 6 and 7.

Once the current supply was fixed, a second problem arose whereby the first 6 channels would measure correctly through the housekeeping board, but the final two would not. Example measurements of this are demonstrated in Table 3-7. This error was found to be caused by the function of the ISL70218 op-amp used to amplify the voltage drop between the two multiplexers. It performed poorly at amplifying signals at low common-mode voltages. Therefore, it would perform as expected at the top of the resistor chain, but at the bottom of the resistor chain, where the common-mode voltage being supplied to it was small, it failed to provide the proper amplification and therefore the resulting measurement was lower. The solution I devised was to increase the size of the resistor at the bottom of the chain, after the PRTs, to 3.3 k Ω from 1 k Ω , to increase the common mode voltage being fed to the amplifier. Therefore, it was also necessary to run the op-amp off the +12 V line that had been introduced previously to give the chip enough head room. The ISL70218 op-amp is rated for a supply voltage from 3 V to 42 V, however

its performance degrades below a voltage supply of ± 5 V as the amp does not draw enough current for the gain to remain constant across the range of common mode voltages used (intersil, 2020). These changes fixed the performance issues and enabled the board to operate as intended, so they were incorporated into the final design for the flight boards. Figure 3-28 shows a plot of resistance data retrieved through the housekeeping board as part of this testing. In comparison to the values in Table 3-6, the resistances measured at each of these values is very close ($<0.1\%$ error) to the true resistor values. This is within the precision of the resistor manufacture values.




Figure 3-28: A sample of the resistance plot generated by my test setup. The voltage trace has been converted to a resistance. The resistance changes as the housekeeping multiplexer boards are switched through by the Arduino micro-controller. This plot is from the fully modified engineering model board and it can be seen that each of the resistances (the selected Y-coordinates) measured is now correct.

3.3.2.2 Flight Model Housekeeping Testing

Two flight models of the housekeeping board, referred to as Flight 1 and Flight 2, were built and populated by the department's technical staff, one of which is shown in Figure 3-29. I tested both boards using the same setup as described in the previous section

to ensure their functionality and assess the accuracy of the resistance measurements. The test criteria included verifying proper channel switching, ensuring correct control line function, stability of channel measurements, and consistency across multiple readings.




Figure 3-29: One of the LTM flight housekeeping boards ready for testing. Each of the electrical mounting points has a connector so that the test resistors can be plugged and unplugged without touching the flight board's connection points.

One limitation of this test, however, is that the accuracy of the temperature measurements is not solely dependent on the analogue electronics, provided they are functioning correctly. The primary limiting factor lies in how the voltages are read and digitised through the board. The Data Acquisition Unit (DAQ) used for benchtop testing, a Hewlett-Packard 34970A, has an accuracy of $10 \mu\text{V}$, which is significantly higher than the Analog-to-Digital Converter (ADC) that will be used in flight. The resolution of the ADC is a critical consideration, as it defines how finely the voltages can be measured and converted into digital values.

The resolution of the ADC is dependent on the number of bits available, and the range of voltages it is required to measure. According to the LTM Thermal Analysis Report discussed in Section 3.1.2.3, the operating temperature range of the instrument spans from -60°C to 60°C , which corresponds to a current range of $213 \mu\text{A}$ to $333 \mu\text{A}$ and a PRT resistance range of 750Ω to 1230Ω . This produces a dynamic voltage range of 1.65V to

4.22 V, or 2.57 V overall. With a 16-bit ADC, this would provide a voltage resolution of 39 μV . However, testing revealed that the ADC used in the housekeeping board only achieves 14-bit resolution, increasing the voltage resolution to 157 μV .

An uncertainty of 157 μV corresponds to a resistance uncertainty of 0.1 Ω , and consequently a temperature measurement uncertainty of 22 mK. If the housekeeping electronics can demonstrate accuracy better than 0.1 Ω , then the primary error in the temperature measurement will come from the ADC's performance. If the ADC's effective resolution drops below 14 bits, this uncertainty will increase further. The only solution to this noise is to enable super sampling of the ADC in the flight software.




Figure 3-30: LTM housekeeping flight board setup. The PRT simulator is connected to the board with high precision resistors in each channel. The Arduino Uno on the lefthand side is providing signal to the address lines, controlling the multiplexers and data is logged through an external DAQ, controlled via MATLAB.

During the benchtop testing of the housekeeping flight boards, all voltages were measured using a Hewlett-Packard 34970A Data Acquisition Unit (DAQ), which has an accuracy of 10 μV (Keysight Technologies, 2020). Following the same calculations as above

this voltage resolution defines a measurement accuracy of $6.6 \text{ m}\Omega$, corresponding to a temperature accuracy better than 2 mK .

Compared to the engineering model tests, a key difference in the flight board testing is that high-accuracy resistors (0.1%) were used and mounted onto a separate board to ensure identical test conditions for each flight board. The 0.1% uncertainty in the resistors translates to a temperature variation of approximately 289 mK for a PRT at 293 K . Figure 3-30 shows the flight board being tested within the full setup.

Example results of the testing for Flight 1 and Flight 2 are displayed in Table 3-8 and Table 3-9 respectively. The measured resistance through each channel for each of the three multiplexer pairs is given. Each board was tested four times with the same procedure to verify that the function was repeatable and consistent. The resistances measured through the housekeeping boards was always within the uncertainty of the resistors. The standard error on the resistance measurements was at the noise of the DAQ (which corresponds to $\pm 0.007 \Omega$). The maximum error on the resistances of $\pm 0.05 \Omega$, if repeated

Channel Number	Expected Resistance / Ω	Measured Resistance 1 / Ω	Measured Resistance 2 / Ω	Measured Resistance 3 / Ω
0	Reference	Reference	Reference	Reference
1	845 ± 0.845	844.63 ± 0.030	844.49 ± 0.033	844.62 ± 0.034
2	931 ± 0.931	930.95 ± 0.026	930.92 ± 0.023	931.01 ± 0.021
3	1000 ± 1.00	999.52 ± 0.033	999.70 ± 0.034	999.57 ± 0.037
4	1020 ± 1.02	1020.40 ± 0.039	1020.41 ± 0.039	1020.28 ± 0.041
5	1150 ± 1.15	1149.75 ± 0.042	1150.00 ± 0.041	1149.51 ± 0.046
6	1210 ± 1.21	1209.76 ± 0.026	1209.99 ± 0.022	1209.42 ± 0.014
7	1300 ± 1.30	1299.94 ± 0.021	1300.23 ± 0.018	1299.74 ± 0.014

Table 3-8: Test results from one of the test runs for Flight 1. All of the resistances measured here are within the uncertainty of the resistors themselves.

at the blackbody target, corresponds to an uncertainty in the temperature recording of 15 mK according to the radiometric modelling discussed in Section 3.1.2.

Channel Number	Expected Resistance / Ω	Measured Resistance 1 / Ω	Measured Resistance 2 / Ω	Measured Resistance 3 / Ω
0	Reference	Reference	Reference	Reference
1	845 ± 0.845	844.616 ± 0.030	844.591 ± 0.035	844.820 ± 0.040
2	931 ± 0.931	931.039 ± 0.037	930.937 ± 0.041	931.351 ± 0.046
3	1000 ± 1.00	999.697 ± 0.048	999.630 ± 0.050	1000.037 ± 0.059
4	1020 ± 1.02	1020.585 ± 0.049	1020.299 ± 0.053	1020.773 ± 0.060
5	1150 ± 1.15	1150.050 ± 0.055	1149.707 ± 0.057	1150.158 ± 0.059
6	1210 ± 1.21	1210.132 ± 0.029	1209.611 ± 0.026	1210.156 ± 0.023
7	1300 ± 1.30	1300.462 ± 0.025	1299.674 ± 0.031	1300.625 ± 0.037

Table 3-9: Test results from one of the test runs for Flight 2. All of the resistances measured here are within the uncertainty of the resistors themselves.

Therefore, this electronics testing has delivered an improved housekeeping board design, as well as verification that the flight boards themselves are capable of measuring temperature to an accuracy that satisfies LTM's radiometric accuracy requirements.

3.4 LTM Blackbody Calibration

While I was involved in many aspects of the LTM ground testing, my most significant independent work was the calibration of the PRTs mounted in the internal blackbody target. As discussed previously, providing an accurate radiance spectrum from the calibration target is critical to achieving the radiometric accuracy of the instrument at better than 1K. To facilitate this, I calibrated the blackbody with its PRTs mounted, before it was integrated into the instrument.

In this section, I explain the thermal calibration in two parts: first, outlining the procedure undertaken and the challenges encountered; and second, analysing the data measured, leading to the final calibration constants and the uncertainties inherent in these measurements.

3.4.1 Calibration Procedure

The LTM internal blackbody has 8 PRTs mounted in it. There are two each mounted to the baffle, centre and edge of the blackbody, and finally there are two mounted to the blackbody-instrument interface. The layout of these is shown in Figure 3-31. The mounted




Figure 3-31: LTM internal blackbody before coating, viewed from the rear. Locations for embedded PRTs are marked on.

PRTs were screened prior to mounting, meaning that a large number of PRTs purchased from the manufacturer underwent a light calibration procedure to identify and remove any components with sub-standard performance. For each PRT used, a full thermal calibration was undertaken after mounting. This is because the process of mounting can introduce mechanical strains into the body of the PRT, which in turn change the calibration constants. The goal of this calibration is to derive precise calibration constants for each PRT. The factory calibration constants for the PRTs (Honeywell HRTS-5760s) provide a certainty of ± 300 mK at 300 K (Honeywell, 2015), whereas the in-house calibration procedure in a stable environment can calibrate the PRTs to approximately 10 mK. Since the 300 mK factory accuracy would consume too much of the uncertainty budget for the blackbody's radiometric accuracy, the more precise in-house calibration is necessary.




Figure 3-32: Blackbody target mounted inside the calibration chamber. Note the MLI surrounding the thermal plate to reduce thermal load on the heating/cooling system and create an isothermal environment, the copper pipes of the liquid cooling/heating for the thermal plate, and the wiring showing where one of the edge PRTs is mounted.

The LTM blackbody was calibrated over a temperature range of 253 K to 333 K to cover the full possible operating range of the instrument. The blackbody was mounted on a thermal plate within a vacuum chamber to provide temperature control, and radiometric coupling between the blackbody/thermal plate and the interior walls of the vacuum chamber was minimised by encasing the blackbody/thermal plate in MLI. A photograph of the blackbody target mounted inside of the vacuum chamber is included in Figure 3-32.




Figure 3-33: Blackbody calibration experimental setup showing the vacuum chamber and equipment, temperature control equipment and datalogging equipment.

The temperature control of the thermal plate was managed via liquid cooling using a bench-mounted Huber ministat 125, regulated by a Huber Pilot ONE using Huber SilOil thermofluid. This system has a manufacturer quoted temperature range of 248 K to 398 K, and a temperature stability of 20mK (huber, 2014). Though this system provides a lesser

cooling power (50 W at 253 K) than the Oxford Thermal Calibration Chamber and less thermal stability (20 mK as opposed to <10 mK), it has the advantage of being faster to reach that stability (hours rather than days) due to being a liquid controlled system. This speed advantage was important due to scheduling constraints as there was a narrow window for thermal calibration in the LTM assembly schedule. Figure 3-33 is a photograph showcasing the calibration setup during calibration.

The planned calibration procedure involved taking measurements at stabilised setpoints ranging from 253 K to 333 K, increasing in increments of 10 K. However, once the calibration began, an issue with the procedure became apparent. The Huber system's cooling power, rated at 150 W at 263 K and 50 W at 253 K (huber, 2014), proved insufficient to achieve a stable temperature at 253 K. This issue likely arose because the Huber was operating in a small room, causing the ambient temperature to rise, reducing the thermodynamic efficiency of the system.

To address this, I adjusted the calibration procedure by starting at 273 K, allowing the Huber more time to reach stability at this setpoint, and then progressing through the 273 K to 333 K setpoints. The lowest temperature setpoints were reattempted at the end of the process. Given the time constraints, this adjustment ensured that most of the temperature range was measured within the LTM build schedule. By the end of the calibration process, stable readings were successfully obtained at 273, 283, 293, 303, 313, 333, and 263 K. The details of the achieved temperatures are laid out with the results of this measurement in Table 3-10. After adjusting the calibration procedure to account for the Huber's limitations, we proceeded with measuring the PRT resistances.

Huber setpoint / K	Temperature achieved /K	Standard Deviation /K
263.15	267.032 ± 0.013	0.0156
273.15	275.719 ± 0.013	0.0047
283.15	284.629 ± 0.013	0.0042
293.15	294.099 ± 0.013	0.0040
303.15	303.078 ± 0.013	0.0113
313.15	312.321 ± 0.013	0.0145
333.15	329.999 ± 0.013	0.0095

Table 3-10: The average stabilised temperature achieved at the thermal plate according to the RhFe sensor, compared to the Huber setpoint, as well as the standard deviation in temperature across the period used for the measurements.

I undertook the calibration using a procedure based on the calibration discussed in Section 3.3.1. Each of the PRTs were connected in series, with a 500 Ω reference resistor in the chain, and voltage drops were measured across each resistor. Voltages were measured using a Hewlett-Packard 34970A Data Acquisition Unit, which achieves an accuracy of 10 μ V, the order of 0.001% of the readings (Keysight Technologies, 2020). The datalogger was controlled via a MATLAB script I wrote to record all measurements with timestamps and uploaded the data to a cloud server in real-time in order to ensure data would not be lost in the event of a failure in the experimental system. At the time of the calibration only 6 of the 8 PRTs in the LTM design had been mounted. The two missing blackbodies were on the interface which mounts the blackbody to the instrument rather than the blackbody itself, are not used directly in the calibration and are thus less vital. Those sensors have been calibrated from their screening data, and this will be verified in flight once the instrument undergoes space environment checks.

A small current (~ 1 mA to prevent self-heating) was applied across the resistor chain by an external power supply. This current was precisely measured using the voltage across the reference resistor. The four-wire PRTs are connected to the datalogger by means of an adjusted LTM housekeeping board. This board provided the correct connections for the blackbody which already had its flight connectors attached but contained no chips. Where the PRTs would attach to the multiplexers on the housekeeping board, they instead feed through to a D-type connector and pass out of the vacuum chamber. This modified housekeeping board can be seen in Figure 3-34.




Figure 3-34: The modified housekeeping board made to feed the PRT connections from the blackbody flight connector out of the vacuum chamber.

The temperature of the thermal plate was recorded by three Rhodium-Iron (RhFe) temperature sensors. These reference sensors had been calibrated separately and provide a high accuracy (<10 mK) measurement of the temperature. These sensors, also four-wire temperature sensors, are on the same resistor chain as the PRTs, and are measured by the

datalogger in a similar manner. A circuit diagram detailing the resistor chain system, how each of the four-wire sensors is attached to the datalogger, and how this relates to the vacuum chamber, is laid out in Figure 3-35.

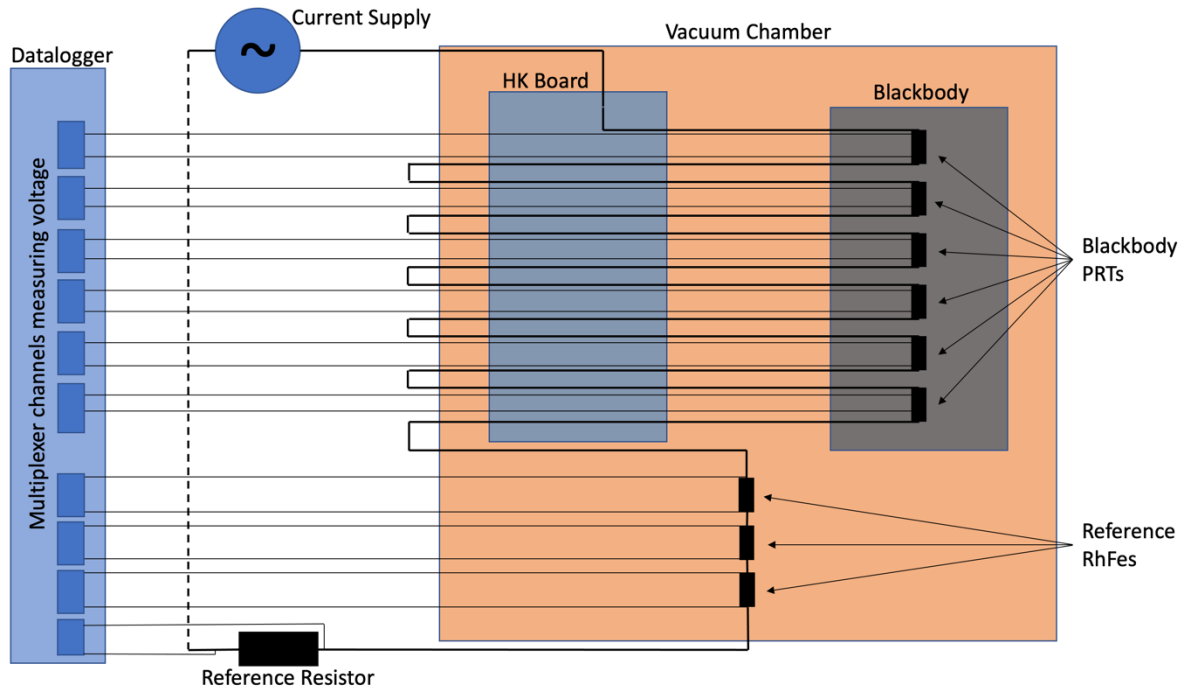


Figure 3-35: Circuit diagram for the PRT calibration electronics. The thick black wires indicate the series resistor chain, the thin black wires indicate the channels where the multiplexer datalogger is measuring each voltage.

3.4.2 Blackbody Calibration Results

The resistances measured for each PRT over the course of the calibration procedure are given in Figure 3-36. One notable outlier is the Blackbody Baffle #2 PRT, which exhibits resistance values that deviate from those of the other PRTs, despite being of the same type and operating at the same temperature. In contrast, the resistance measurements from the remaining PRTs align closely, to the point where their data overlaps with the blue line representing BB Centre #1. Further investigation of the Baffle #2 data from subsequent instrument calibration revealed additional inconsistencies, such as instances where its resistance increased while the resistances of the other PRTs decreased. This suggests that the Baffle #2 sensor may be faulty. Further analysis of this sensor will be made based on data retrieved once the instrument is in operation. If the

behaviour of the PRT is internally consistent it may be possible to calibrate it against the other PRTs from instrument or spacecraft thermal cycling tests.




Figure 3-36 Resistances for each of the 6 PRTs across the duration of the calibration procedure. The initial failed attempt to stabilise at -20°C can be seen, followed by 36 hours stabilising at 273 K and then steps at 12 hour intervals up to 333 K. Finally, there is a 36 hour measurement at 263 K where the huber is not always stable.

Figure 3-37 shows the temperature of the block to which the blackbody was mounted throughout the calibration procedure. This temperature data was generated using a pre-existing calibration curve for the RhFe sensor (Clarendon Laboratory, 2005), and the blackbody PRTs were calibrated against these reference temperatures. The RhFe sensor itself was calibrated in-house at Oxford with an accuracy of ± 8 mK, comparing to a National Physical Laboratory standard.

The noise in the reference temperature was measured at 0.0035% of the signal, corresponding to approximately 10 mK for a 294 K measurement. Since the noise in temperature readings and the inherent accuracy of the RhFe calibration are assumed to be uncorrelated, the total uncertainty in the temperature measurement is derived by

combining these factors in quadrature. This results in an overall uncertainty of 13 mK for the temperature measurements.




Figure 3-37: The plot of temperature for the blackbody against time over the duration of the calibration procedure, as measured by the RhFe reference resistors.

It is assumed that the blackbody reaches thermal equilibrium with the block to which it is mounted after 12 hours of stable temperature. Figure 3-38 illustrates the stability achieved during the calibration, focusing on the 30°C measurement. A bulk temperature gradient of approximately 10 mK per hour is evident across the entire reading. Additionally, oscillations caused by PID control are visible, with a periodicity of about one hour and an amplitude of 30 mK.

I selected and averaged the data from the final 2 hours of each setpoint measurement, specifically the last two oscillations in the PID loop. This subset of data, with a standard deviation of 11.4 mK, represents the most stable period of the measurement. Figure 3-39 and Figure 3-38 also show that while a temperature gradient is present in the PRT readings, the oscillations are significantly dampened by the thermal break between the temperature-controlled thermal plate and the blackbody. The




Figure 3-38: A focused display of the 30°C measurement including measurement uncertainty. The fluctuations from the control system are evident, as well as a slight downward trend across the measurement once it peaks. The bulk gradient shows a drop in temperature of 100mK over a 10 hour measurement, with a standard deviation of 23mK across this whole period.

blackbody is mounted on PTFE struts specifically to achieve this damping. By focusing on the last 2 hours of stable data for each setpoint, the calibration will reflect the most isothermal conditions, minimising any thermal gradients between the reference thermometers and the PRTs.

The temperature values used in the calibration equation for the PRTs will include multiple sources of uncertainty. These include the uncertainty from the RhFe temperature measurement and any fluctuations in temperature observed during each measurement. Similarly, the resistance values will reflect the inherent uncertainties in the measurements taken by the DAQ, as well as variations recorded throughout the measurement period. Additionally, the final uncertainty in the calibrated PRT values will also include an error introduced by the polynomial fitting process, as described in Equation 3.27.




Figure 3-39: The PRT resistances from the same measurement as displayed in Figure 3-36. The bulk gradient is still visible, but the oscillations have a smaller amplitude relative to the size of the change over the full recording. This shows how there is some damping between the thermal plate and the temperature measured at the PRTs.

The results from the calibration procedure are presented in full in the following tables. The temperatures achieved at the thermal plate for each two-hour period were given previously in Table 3-10, along with their respective standard deviations. These

	Average PRT Resistances $\pm 0.010 / \Omega$				
Temperature	Baffle #1	Edge #1	Edge #2	Centre #1	Centre #1
-6.118 ± 0.013	984.987	985.174	985.252	985.487	984.963
2.569 ± 0.013	1014.894	1014.994	1015.088	1015.288	1014.771
11.479 ± 0.013	1045.918	1045.961	1046.062	1046.256	1045.711
20.949 ± 0.013	1079.304	1079.286	1079.402	1079.584	1079.013
29.928 ± 0.013	1109.998	1109.873	1110.003	1110.192	1109.570
39.171 ± 0.013	1143.709	1143.524	1143.681	1143.851	1143.203
56.849 ± 0.013	1213.582	1213.309	1213.527	1213.649	1212.972

Table 3-11: Average resistances for each PRT at each temperature in Ohms.

values represent the stability of the temperature environment during each calibration setpoint. The corresponding PRT resistance data is detailed in Table 3-11, providing the specific resistances measured for each PRT at each temperature setpoint. Table 3-12 contains the standard deviations of the resistance values, which will be factored into the overall uncertainty calculations for the calibration.

	PRT Resistance Standard Deviations /Ω				
Temperature	Baffle #1	Edge #1	Edge #2	Centre #1	Centre #1
-6.118 \pm 0.013	0.0615	0.0620	0.0618	0.0613	0.0618
2.569 \pm 0.013	0.0183	0.0179	0.0182	0.0176	0.0176
11.479 \pm 0.013	0.0076	0.0076	0.0079	0.0077	0.0078
20.949 \pm 0.013	0.0092	0.0089	0.0090	0.0089	0.0086
29.928 \pm 0.013	0.0247	0.0238	0.0243	0.0242	0.0240
39.171 \pm 0.013	0.0573	0.0579	0.0584	0.0581	0.0586
56.849 \pm 0.013	0.0826	0.0847	0.0844	0.0839	0.0849

Table 3-12: Standard deviation for each PRT over the stable period used for calibration. These will be used to derive uncertainties.

The calibration of the PRTs was performed by fitting a third-order polynomial to the seven data points collected during the calibration process. This polynomial, described in Equation 3.26, establishes the relationship between the resistance of each PRT and the temperature, allowing accurate temperature readings to be derived from future resistance measurements. The calibration coefficients for each of the PRTs, which define this polynomial relationship, are listed in Table 3-13. These coefficients are unique to each PRT and have been used to form the final calibration products in the LTM calibration report, and feature in the final LTM calibration pipeline.

The uncertainty associated with temperature measurements made using the calibrated PRTs has also been calculated and are provided in Table 3-14. This uncertainty

	LTM PRT Calibration Coefficients			
	A₁	A₂	A₃	A₄
Baffle #1	732.227	-1.86749	2.06626 x10 ⁻²	-6.59652 x10 ⁻⁷
Edge #1	744.866	-1.90477	2.10218 x10 ⁻²	-6.70973 x10 ⁻⁷
Edge #2	745.203	-1.90559	2.10285 x10 ⁻²	-6.71188 x10 ⁻⁷
Centre #1	735.241	-1.87826	2.07769 x10 ⁻²	-6.63447 x10 ⁻⁷
Centre #2	751.665	-1.92459	2.12143 x10 ⁻²	-6.77135 x10 ⁻⁷

Table 3-13: The calibration constants for the LTM blackbody PRTs.

arises from the combined sources discussed earlier, including measurement noise, fluctuations during the data collection period, and the accuracy of the polynomial fit. The overall uncertainty for the PRTs at room temperature (approximately 293 K) corresponds to 88 mK. This value represents the precision that can be achieved in temperature readings taken using the calibrated PRTs and is well within the acceptable limits for maintaining the radiometric accuracy required by LTM.

PRT Label	Calibrated PRT Uncertainty
Baffle #1	0.028%
Edge #1	0.029%
Edge #2	0.029%
Centre #1	0.029%
Centre #2	0.029%

Table 3-14: The uncertainties for the calibrated LTM blackbody PRTs.

3.5 Radiance-Temperature Reference Tables

For any part of the lunar surface imaged by LTM, the top-level data product will consist of images where each pixel corresponds to a calibrated radiance measurement

across multiple channels. More advanced retrieval algorithms can then process these radiances to infer surface properties, such as temperature and emissivity. As detailed in Section 2.3.2, the immediate data product required by the HVM³ instrument is the brightness temperature. This value is critical for performing thermal corrections on their spectral measurements.

To support this requirement, I developed a look-up table that relates the calibrated radiance measured by the instrument to the corresponding brightness temperature at the viewed surface. This table is the product of a radiometric model similar to those discussed earlier in this chapter, and using the full instrument transmission functions created in Section 3.2. The table spans a temperature range of 90 K to 400 K, with 1K increments. For each temperature, the spectral radiance was calculated using the Planck function at a 10-nanometre resolution across the 5 to 150 micrometre wavelength interval.

The detector radiance was then calculated by convolving the spectral radiance with the instrument transmission functions. Finally, measured radiance was extracted by integrating the radiance at the detector. Since we are considering the brightness temperature of the body being observed, the emissivity of the surface was set to 1. This generated for each filter a list of measured radiances produced by the list of temperatures. Any calibrated radiance value which is returned by the instrument can be related by interpolation to the brightness temperature which would have produced that measurement.

This look-up table, relating the calibrated radiance from LTM to a surface brightness temperature were generated and delivered to our colleagues at JPL along with a document explaining its derivation. For illustrative purposes the data from the tables have been plotted in Figure 3-40. Since these results will be used directly in the calibration of

the HVM³ data, this is an extremely important output of the work developing radiometric models and the optical measurements that have been described in this chapter.




Figure 3-40: The curves from the LTM Calibrated Radiance to Surface Brightness temperature lookup tables plotted. Calibrated radiance is on a logarithmic axis

4 Comet Interceptor and MIRMIS

4.1 The Comet Interceptor Mission

The Comet Interceptor mission, a joint collaboration between ESA and JAXA, aims to conduct the first in-depth study of a dynamically new comet—a comet that has never entered the inner Solar System. Such comets are of particular scientific interest because they remain largely unchanged since the formation of the Solar System, retaining their pristine structure and composition. This provides a rare opportunity to study the early Solar System's building blocks. The primary objective of Comet Interceptor is to observe and capture data from a dynamically new comet before its volatile elements are significantly altered by Solar radiation and heat during its first close passage to the Sun (Snodgrass & Jones, 2019).

Comet Interceptor is scheduled for launch in 2029, where it will initially reside in a holding orbit at the Earth-Sun Lagrange Point 2. This vantage point allows it to remain in a stable orbit while awaiting the opportunity of a suitable target. Once an appropriate comet is detected, the mission will be rapidly deployed to intercept the comet at high speed. During the flyby, Comet Interceptor will split into three separate spacecraft, each equipped with a suite of scientific instruments and following distinct trajectories. This multi-spacecraft configuration ensures comprehensive coverage of the comet from different distances and risk profiles, allowing each spacecraft to gather different data while minimizing the possibility of damage from the comet's dust and debris. The mission will achieve unprecedented multi-point observations, giving scientists detailed views of the comet's nucleus, coma, and the surrounding environment.

The scientific payload on Comet Interceptor includes the Modular InfraRed Molecules and Ices Sensor (MIRMIS), which is designed to analyse the thermal environment and composition of the comet's surface and surrounding coma. MIRMIS, integral to the broader instrument suite, supports the mission's primary scientific goal of investigating the diversity of comets, by mapping the temperature and composition of the nucleus and coma across a wide spectral range (0.9–25 μm) (Bowles et al., 2021). It will provide detailed characterisation of the comet's surface morphology, thermal properties, and volatile composition, allowing a reference point to answer the questions posed around the transfer of volatiles through the Solar System in Section 1.2. By offering the most detailed observations of a comet to date, and specifically applying this capability to a dynamically new comet, Comet Interceptor will provide insights into the composition and behaviour of these pristine objects.

In this chapter, I will provide a comprehensive overview of the Comet Interceptor mission and focus specifically on the MIRMIS instrument. I will discuss MIRMIS's design and how it aims to operate to measure the comet's nucleus, alongside some of the key challenges facing its development. This discussion will culminate in defining the official requirements, as well as the personally derived requirements, that shaped the work I discuss. Chapter 5 will then delve into the design work I have carried out for MIRMIS, including the development of its full operational program and performance verification. Finally, in Chapter 6, I will discuss the end-to-end simulation I developed for the Thermal InfraRed Imager (TIRI), which supports further development of TIRI and provides a framework for the continued refinement of other modules and instruments on the mission.

4.2 The MIRMIS Instrument

MIRMIS, developed through a collaboration between the University of Oxford and the VTT Technical Research Centre of Finland, is a thermal imager consisting of three independent modules, each designed to observe the comet across different spectral ranges.

These modules are:

- The Near InfraRed Hyperspectral Imager (NIR): Developed by VTT Finland, NIR operates in the 0.9–1.7 μm spectral range. It fields a spectral resolution of approximately 30–40 nm, a field of view (FOV) of $6.7^\circ \times 5.4^\circ$, and an instantaneous field of view (iFOV) of 0.18 mrad. The iFOV is the angular resolution of a single pixel.
- The Mid InfraRed Spectrograph (MIR): Also developed by VTT Finland, MIR functions as a single-channel spectrometer within the 2.5–5 μm range. It offers a similar spectral resolution of 30–40 nm and features a circular FOV of 2° .
- The Thermal InfraRed Imager (TIRI): Developed by the Planetary Instrumentation Group at Oxford, TIRI is a multispectral imaging radiometer operating across the 6–25 μm spectral range. It features a FOV of $9^\circ \times 7^\circ$ and an iFOV of 0.26 mrad.




Figure 4-1: The MIRMIS internal structure from the CAD model. TIRI is on the right (with an identical optical design as LTM) and NIR and MIR are mounted on the left.

The three modules share a common mechanical, thermal, and electrical interface to the spacecraft. The combined MIRMIS instrument command and data handling unit (CDHU) is designed to allow the instrument independent operation during the encounter with the Comet Interceptor target body.

4.2.1 The Thermal InfraRed Imager

TIRI's design shares many similarities with LTM, including the optical feedthrough and electrical framework. However, due to the unique operational requirements of the Comet Interceptor mission, and thanks to the opportunity to learn from the development of LTM, TIRI is an updated and improved version of the instrument. The primary difference lies in the operational mode: while LTM functioned as a push-broom radiometer, TIRI images the entire cometary nucleus through each filter independently. These images are then collated to form a multispectral image which we call a data-cube. The data-cube is essentially formed from the stack of different filter channel images to relate the spectral

information to the spatial location. A more detailed discussion of this is contained in Section 4.3.

TIRI is equipped with eight narrow-bandpass compositional filters, covering the 7–25 μm region, along with a single broadband thermal channel for temperature measurements. These filters allow TIRI to capture critical compositional data, while the broadband channel provides precise surface temperature readings. The original filter bandpasses are detailed in Table 4-1. Upon further analysis of the bandpasses (detailed in Section 5.3) the decision was made to slightly alter these, and as such the analysis in Chapter 6 uses a slightly different set of bandpass data. As the instrument is modelled through these sections, it will be made clear which set of bandpasses have been used for the modelling.

Channel	Bandpass / μm
A	7.56 – 7.96
B	8.76 – 9.20
C	10.13 – 10.65
D	11.73 – 12.33
E	13.58 – 14.28
F	15.73 – 16.53
G	18.21 – 19.23
H	20.61 – 22.61
Broadband	7 - 25

Table 4-1: TIRI filter bandpasses from the original instrument design.

TIRI's detector differs from that of LTM, offering a x 480 pixel array with a 17 μm pixel pitch. However, the principle on which it operates, which relies on calibrating against external views and an internal blackbody for comparative radiance measurements, remains the same. The TIRI concept of operations will be discussed in later in this chapter, Section 4.3, while a full operational plan, informed by detailed modelling, is presented in Chapter 5.

4.2.2 Mid-InfraRed and Near-InfraRed Imagers

The MIR and NIR modules operate within shorter wavelength regions than TIRI and serve different scientific objectives. MIR is a single-channel spectrometer covering the 2.5–5 μm range, while NIR is a hyperspectral imager operating within the 0.9–1.7 μm range. These instruments focus on analysing different aspects of the comet's composition.

NIR and MIR share a single pointing mirror to control their viewing. This motor assembly is being designed and produced within Oxford, and as such this project has played a role in that development. This thesis will only address work supporting NIR and MIR in the capacity of developing the pointing assembly and the methods used for its control.

4.3 MIRMIS Concept of Operations

In this section I will give a brief overview of the concept behind how MIRMIS will operate during the encounter, covering how the data will be taken. A special focus is given to TIRI since it has significant complexity in its operations.

The spacecraft platform will not track the comet directly during the closest approach. Instead, it will maintain a fixed orientation relative to the comet nucleus, requiring MIRMIS to compensate for this lack of rotation using pointing mirrors. MIRMIS is mounted on the anti-ram plate of the spacecraft, shielding it from the dust in the comet's coma. As a result, the instrument will only be able to observe the comet from a limited range of angles - starting just before closest approach and continuing as the spacecraft recedes from the comet. Figure 4-2 illustrates the observation windows.

Since the mission timeline depends on velocity and encounter distance, it is most useful to refer to the operational profile in terms of angle. I define the operation angle as the

angle between the spacecraft velocity and the spacecraft-comet vector. During the encounter, MIRMIS will be able to observe the comet within an operation angle range of 85 to 173 degrees.

The primary function of the pointing mirrors is to keep the instruments following the comet nucleus throughout the encounter. For the NIR and MIR modules, the mirrors will continuously track the nucleus, ensuring that it remains centred in the field of view for the entire observation window.




Figure 4-2: This diagram demonstrates the encounter phase for Comet Interceptor spacecraft A, defining the direction of the spacecraft velocity vector \vec{V}_{sc} , with respect to the ram direction. The observation angle is defined as the angle between the nominal velocity vector and the line of sight to the comet nucleus. Coloured regions represent the angular range in which data-cubes must be generated (not to scale).

For TIRI, however, the pointing process is more complex. Unlike NIR and MIR, TIRI needs to scan the entire comet nucleus through each filter individually. This scanning process is essential for generating multispectral data across different wavelengths in each of the instrument's channels. Each filter channel is observed by deviating from the tracking

position, to move the nucleus in the field of view. The filters are laid out in the detector field of view in a pattern as demonstrated in Figure 4-3.

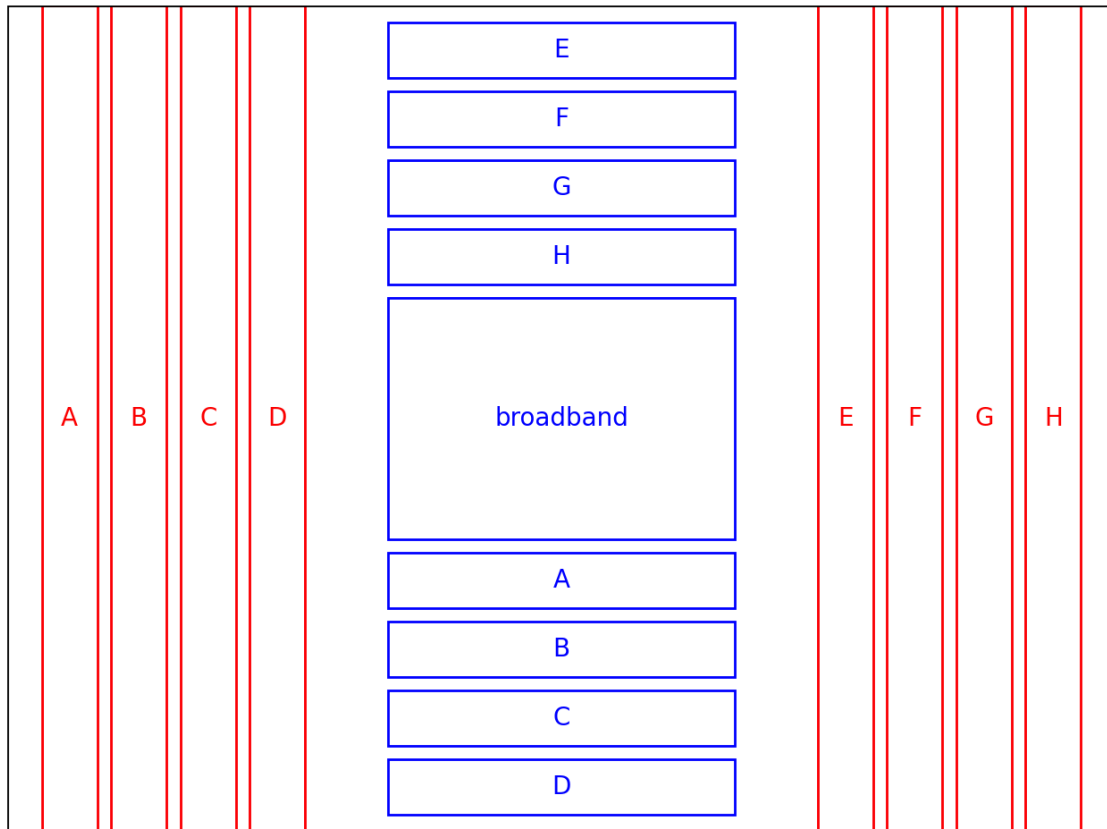


Figure 4-3: The layout of TIRI's filters, showing the central broadband channel and the repeated compositional filters arranged for horizontal and vertical scanning phases.

When TIRI scans the comet nucleus through its multiple filters, the resulting images are collocated and stacked to form a data-cube. Each data-cube is essentially a 3D image of the comet's surface, where every pixel contains the full set of spectral data across TIRI's eight compositional filters. TIRI is required to capture five data-cubes during the encounter, with the angles demonstrating the range within which data-cubes need to be taken described by the shaded regions in Figure 4-2. The specifications and constraints related to these data-cube acquisitions are detailed in Section 4.4.1.

The filter layout is designed to address one particular challenge of following the comet nucleus. As the pointing mirror rotates around the axis which the incoming light is reflected along, that mirror movement rotates the image in the detector's field of view. This is shown in Figure 4-4, and there is a full mathematical description of this vector




Figure 4-4: The rotation of the TIRI fov through operation illustrated by projecting an orange rectangle as the instrument fov as a set of different observation angles. As the pointing mirror moves through 90 degrees, the image rotates from the detector frame of reference.

transformation in Section 6.3.3. This means that during the initial phase of the encounter, at an operation angle near 90° , data-cubes are generated by scanning the nucleus horizontally across the detector, utilising the red filters from Figure 4-3. However, near the end of the encounter, at operation angles nearing 180° , the scanning is shifted to a vertical orientation, using the blue filter from Figure 4-3. There is, however, a period in the middle where the filters cannot access the entire nucleus, limiting the formation of complete data-cubes. This process is described in greater detail in Section 5.2.4.

4.3.1 Key Scanning Challenges

The MIRMIS Concept of Operations faces several challenges due to the high relative velocity between the spacecraft and the comet, as well as the spacecraft's fixed orientation during closest approach. The spacecraft is expected to exhibit a relative velocity of up to 70

km/s, imposing strict constraints on the available timeframe for data collection and placing significant demands on the pointing mechanism and control systems.

The first major challenge is related to pointing and tracking the comet's nucleus. Since the spacecraft itself does not track the comet, the responsibility falls entirely on MIRMIS's pointing mechanisms. For the NIR and MIR modules, the pointing mirrors must continuously follow the comet's nucleus, keeping it centred in the field of view. Any misalignment during this period could result in a significant loss of data.

For TIRI, the challenge is more complex due to the requirement to scan the comet's nucleus multiple times, covering each filter. The first complexity arises from the narrow size of the filters: at closest approach, the nucleus is expected to be approximately 20 pixels wide, and it must be imaged through a 30-pixel-wide filter. Compounding this is the rotation of the field of view, meaning the offset required to position the comet in each filter changes dynamically as a function of both filter and operation angle. This requires precision in both targeting and pointing control, as calculating where the mirror should be aimed is no more trivial than achieving fine motor control accuracy while moving at high velocity.

Defining TIRI's operation within the broad parameter space available is also a non-trivial task. Variables such as detector integration times, filter order, pointing mirror motor speed, and others can all be adjusted, and a strategy must be developed to maximise the scientific return from the measurements. A key example of this is the balance between spatial resolution and filter performance. As was the case with LTM, the narrow compositional channels of TIRI are highly sensitive to temperature and integration periods. To maximise the SNR, longer integration times are desirable. However, since the comet nucleus is being observed from different angles during the flyby, and the field of view is rotating, the object imaged by TIRI is effectively rotating about two axes. Any

increase in integration time introduces more blur, both within individual images and across the entire data-cube, as each image is taken from a slightly different angle.

Another similar juxtaposition exists around waiting for the mirror to stabilise before the detector takes a measurement. If the instrument takes images as rapidly as it can between movements, it may introduce blur due to jitter on the motor while the detector is being integrated. However, pausing to allow the mirror to settle will increase the time between images and therefore improving the individual image comes at the expense of increasing blur across the data-cube (between different images).

Much of the work in Chapters 5 and 6 is dedicated to developing tools and methods that enable informed decisions during the engineering, electronics, and software development phases, optimising the science return from MIRMIS.

4.3.2 The Comet Interceptor Nominal Case

The Comet Interceptor mission presents a wide range of possible flyby and encounter conditions. These parameters include the relative spacecraft-comet trajectory - such as angles, distances, planar offsets, and relative velocities - as well as target-specific variables, including the comet's nucleus size, surface temperature, and rotational axis. Many of these factors will not be well constrained prior to the actual encounter. To ensure that the instrument can meet its scientific goals under these uncertain conditions, a nominal case has been defined.

The nominal case represents an expected scenario for the comet encounter, with some parameters pushed to their worst-case values. This approach allows the instrument design to be validated by ensuring that it can, not only meet, but exceed its performance requirements under these more challenging conditions. By designing MIRMIS to operate successfully in this nominal case, we can be confident in its ability to function during the actual mission.

The nominal case is consistent across the various instruments on board the spacecraft, although the specific worst-case parameters have been selected based on the unique challenges each instrument faces. The table below outlines the key parameters for TIRI's modelling in the nominal case. This does not include the nominal case parameters for dust and gas production rates which govern the dust, fields and plasma experiments on the spacecraft.

Parameter Name	Parameter Value
Distance at Closest Approach	1000 km
Encounter Velocity	70 km/s
Pointing Mirror Rotation Velocity	6 degrees /second
Comet Surface temperature	300 K
Nucleus size	5 km

Table 4-2: The nominal case parameter values for TIRI.

4.4 Instrument Requirements

As with LTM, MIRMIS must meet a set of predefined requirements to ensure the scientific return of the mission aligns with the goals outlined in the Comet Interceptor mission requirements document. Given the complexity of MIRMIS operations and the wide parameter space still open, I have expanded upon these by defining additional scientific requirements. These requirements will guide my work, ensuring alignment with the official parameters, which will be confirmed by ESA.

4.4.1 Official Requirements

The Level 2 science requirements for MIRMIS operations define the operation angles at which the MIRMIS modules must collect data. These angles have a direct impact on the operations plan, as they dictate the necessary pointing mirror angles and rotation rates for NIR, MIR, and TIRI, compensating for the spacecraft's relative velocity vector during the encounter. The specific requirements for each module are summarised below and illustrated in Figure 4-2.

Requirement Identifier	Requirement
R2-P-056	A-MIRMIS-TIRI shall obtain at least 3 thermal images and 1 complete data-cube throughout the Central fly-by region ($80^\circ < \vartheta < 175^\circ$), at $\vartheta = 90^\circ \pm 10^\circ$ with resolution at least 300 m/px and second at $\vartheta = 110^\circ \pm 15^\circ$ with resolution at least 600 m/px.
R2-P-066	A-MIRMIS-TIRI shall obtain at least 2 additional complete (TBC) data-cubes throughout the region of ($125^\circ < \vartheta < 173^\circ$), first at $\vartheta = 140^\circ \pm 5^\circ$ and second at $\vartheta = 165^\circ \pm 5^\circ$.
R2-P-138	MIRMIS-TIRI should obtain at least 2 complete data-cubes of large-scale coma structures at angles $\vartheta > 173^\circ$. Note that, the flux in the coma drops as $1/r$, until molecules get photo-dissociated and dust grains sublime.

Table 4-3: The MIRMIS TIRI requirements defining the positional of the data-cube retrieved, defined relative to the observation angle, here defined as ϑ .

Additional requirements for MIRMIS-TIRI are related to temperature monitoring, calibration, and instrument accuracy, given in Table 4-4. These are similar to the LTM requirements and exist for the same reasons, temperature knowledge through operation is very important for accurate calibration. Note that TIRI does not have the temperature stability requirement that LTM does. This is because the calibration procedure of TIRI is not so well defined. Since TIRI will only view the nucleus in one frame at a time, the rest of the detector (more than 95% of pixels) will be viewing deep space. As such there is a near

constant cold target calibration view being produced and this can be used to track and account for the instrument background radiance.

Requirement Identifier	Requirement
MIRMIS-TIRI-PER-010	MIRMIS shall implement redundant temperature monitoring of the TIRI filter assembly and detector assembly.
MIRMIS-TIRI-PER-070	TIRI shall incorporate a calibration target to provide knowledge of the radiometric accuracy to ± 1 K.
MIRMIS-TIRI-BB-0500	The TIRI calibration blackbody target temperature sensors shall have an accuracy of 0.25 K.
MIRMIS-TIRI-IF-180	TIRI module shall install 2 temperature sensors on both detector and filter assemblies to permit monitoring by spacecraft while TIRI module is powered off.

Table 4-4: The MIRMIS TIRI requirements governing internal temperature measurement. These are similar to the requirements for LTM.

Finally, MIRMIS has a set of power requirements which govern the average and peak power that can be drawn by the modules and the instrument, these are laid out in Table 4-5. For MIRMIS the short period power consumption is dominated by the pointing motors.

Requirement Identifier	Requirement
MIRMIS-NIRMIR-IF-100	The NIR/MIR module shall have an average power of 4.3 W
MIRMIS-IF-370	The average power draw of the MIRMIS instrument shall be less than or equal to 11.2 W
MIRMIS-IF-380	The peak power draw of the MIRMIS instrument shall be less than or equal to 20 W
MIRMIS-IF-1920	The instrument PI shall provide a power profile for all units. Note: Power profiles with respect to the different modes and mission phases is needed by the S/C to size the power subsystem.

Table 4-5: The MIRMIS requirements governing power draw in terms of averages, peak power and the provision of expected power draw profiles.

4.4.2 Community Requirements

Given the broad parameter space in which MIRMIS operates, I developed additional requirements and constraints to guide its development. Variables such as motor velocity, integration times, data retrieval methods, and operational schedules all significantly influence the quality of the data captured during the encounter. This project seeks to establish clear links between these operational parameters and their effect on the instrument's scientific output, allowing for informed trade-offs to be made during the design of MIRMIS's electronics, software, and operational plans.

To assist in defining the key constraints that might limit the utility of TIRI's measurements, I consulted with the MIRMIS instrument science team. The specific questions posed to them included:

- What spatial resolution is required to meet the science goals and given these are dynamic measurements, what level of spatial degradation caused by rotational and linear blurring can be accommodated?
- Are any of the TIRI filters a higher priority in case a data-cube has to be compacted, or if there are opportunities to gather extra data?
- Is there any constraint on compositional order, and are there any relationships between compositional filters which should prioritise making those measurements as coincidental as possible?
- What is the minimum SNR within each filter which will enable mineral retrieval? Should each filter target the same SNR, or do some channels require a higher SNR?
- What degree of post-processing should be done on the retrieved data before it is released to the science users?

The responses from the science team were clear and led to the definition of two key Community Science Requirements that will guide TIRI's operations:

1. The spatial resolution of the mapping produced by TIRI should be maintained below 500m in order to allow observation and differentiation between features.
2. All of the compositional filters should target an SNR of 100.

These two requirements establish clear boundaries within which TIRI must operate. However, achieving these goals requires careful balancing of competing factors. For example, the need to maximise spatial resolution directly conflicts with the goal of achieving high SNR. Increasing integration times can help improve SNR but also introduces blurring and uncertainty due to the dynamic nature of the comet's motion and the rotating field of view of the instrument. Therefore, these limits define an operational envelope, ensuring that TIRI's data measurements remain scientifically valuable while balancing the need for both high spatial resolution and signal clarity.

5 MIRMIS Design Work

In this chapter, I describe the work I have done to support the development and design of the MIRMIS instrument. One of the core challenges in designing MIRMIS for the complex operations outlined in Section 4.3 is the intricate relationship between subsystem design, operational constraints, and the resulting scientific measurements. To address this, I developed a series of models and tools that link instrument design parameters to the scientific performance. This work not only guides the engineering teams towards optimal design solutions but also provides verification to the Comet Interceptor's central management and funding bodies that MIRMIS will exceed its performance requirements, as specified in Section 4.4.

The first section of this chapter details the work I conducted to define the full set of instrument operations during the flyby. This includes thermophysical modelling of cometary surfaces, radiometric performance assessments, and spatial and pointing analysis. Building on the defined operations plan and timeline, I demonstrate that the instrument can perform within power draw requirements, even in a worst-case scenario. I then describe the analysis and trade-study I conducted to evaluate different bandpass options, which led to the selection of a higher-performing set of filters for TIRI.

Next, I address the requirements for the development of control software for the motor pointing system and describe the tools I created to relate the operations timeline to instrument commanding, thus establishing accuracy limits. Finally, I discuss the work I carried out to further refine the performance of the TIRI housekeeping boards, building on the LTM housekeeping testing in Section 3.3.2, and how I developed operational limits that ensure the instrument meets its performance requirements.

5.1 Defining the MIRMIS Operations

As discussed earlier, the operation of MIRMIS is highly complex, necessitating a well-defined strategy to develop an operations plan that meets the requirements outlined in Section 4.4. This section details the modelling and analysis efforts I undertook to support the creation of this plan, with particular focus on optimising signal, noise, spatial resolution, and the power budget during the encounter.

While I have produced a full operations plan for the nominal case, the true value of this work lies in the underlying analysis. The operations plan will certainly need to be revised once the Comet Interceptor target is selected, and detailed information on the trajectory, encounter parameters, and comet characteristics becomes available in the lead-up to the encounter. Consequently, my aim is not to provide a final, rigid operations plan but rather to develop flexible tools and models. These tools will allow for the rapid generation of an optimised operations plan as soon as the mission's final details are confirmed.

In the immediate term, my focus has been on verifying that the instrument requirements can be met under the nominal case scenario, with particular attention to how decisions are made to balance signal-to-noise ratio and spatial resolution, while also providing a model for power consumption throughout the flyby.

5.1.1 Thermophysical Modelling

To effectively use radiometric modelling tools and investigate the signal performance of the TIRI instrument, I first needed to identify the likely surface temperatures that TIRI will observe. Given the complexity of this topic, I conducted an initial investigation myself before consulting colleagues within the MIRMIS science team for a more developed perspective. These discussions led to the development of a new

cometary thermal model, which is currently being worked on by other members of the team. This model will be discussed in more detail in Section 6.3.2 and will be used extensively in Chapter 6. Here, I briefly describe my own modelling techniques, and the conclusions reached in collaboration with my colleagues.

5.1.1.1 A 1-Dimensional Thermophysical Model

I created a multilayer 1D model of the comet surface in python based on the John Spencer model (Spencer et al., 1989). This model is based on an energy balance at the surface, considering the heat flow:

$$Net\ Heat\ Flow_{Surface} = Q_{absorbed} - Q_{emitted} - Q_{conducted\ to\ layer\ 2} \quad (5.1)$$

where:

$$Q_{absorbed} = (1 - bond\ albedo) \times Solar\ Constant \times \cos(\theta) \quad (5.2)$$

$$Q_{emitted} = \epsilon\sigma T^4 \quad (5.3)$$

In these equations, θ represents the Solar incidence angle, ϵ is the emissivity, and σ is the Stefan-Boltzmann constant. Each subsequent layer below the surface has similar terms, though they only account for heat transfer between adjacent layers without radiative terms:

$$Net\ Heat\ Flow_{layer\ n} = -Q_{conducted\ to\ layer\ n-1} - Q_{conducted\ to\ layer\ n+1} \quad (5.4)$$

Heat transfer between layers is governed by Fourier's law of heat conduction, which relates heat transfer Q to the difference in temperature ΔT and the thickness of the layer Δx , via the thermal conductivity k :

$$Q = -k \frac{\Delta T}{\Delta x} \quad (5.5)$$

The simulation calculates these over a set of time periods and updates the temperatures as follows:

$$\Delta T = \frac{\text{Net Heat Flow} \times \Delta t}{m \times c_p} \quad (5.6)$$

where ΔT is the temperature change over the time step Δt , m is the mass of the layer (density \times area \times depth), and c_p is the specific heat capacity of the material at that layer.

This model allows each layer to have distinct properties and thicknesses, though it assumes simplified geometry with no shading scattering or lateral heat flow and neglects sublimation effects.

The model was run for a comet with 10 surface layers, with the temperatures for the top 3 layers plotted as results in Figure 5-1. The model parameters used are all given in Table 5-1. These values were proposed by members of the MIRMIS science team as representative nucleus parameters. Their accuracy was not vital as the primary purpose of this modelling was to understand the sensitivity of the model outputs to parameter values such as layer depth. Initial conditions were derived from the equilibrium temperature at the selected Solar distance, leading to a slight overestimate which was rapidly corrected. The results show a convergence from the initial conditions to a stable diurnal temperature cycle, with decreasing thermal fluctuations in the deeper layers. These results aligned with expectations, confirming the model's functionality. However, this simulation, based on an astronomical body at a constant Solar distance of 1 AU is not representative of a comet's evolution, thus the surface temperature range of 230 to 370 K is not sufficient to define TIRI operations.




Figure 5-1: The temperature across the top three layers, radiative inflows and outflows of power and the cumulative heat flow into the surface layer plotted for a 12-hour rotational period surface which has its normal perpendicular to the axis of rotation.

To generate more useful temperature predictions, I modified the 1D model to allow the Solar constant to vary as a function of distance from the Sun. I then simulated a Jovian comet on an elliptical orbit from 5.2 AU at aphelion to 0.5 AU at perihelion. The distance against time for this segment of the orbit (farthest to closest) is shown in Figure 5-2.

Parameter Name	Parameter Value
Rotation Period	12 hours
Emissivity	0.9
Density	500 kg/m ³
Specific Heat Capacity	840 J/kg/K
Layer Depth	0.1 metre
Thermal Conductivity	0.1 W/m/K
Time Step	120 seconds
Distance to the Sun	1 AU

Table 5-1: The surface parameters used for the thermophysical comet model described in this chapter. These values are based on measured comet values and water-ice.

Running the model over this trajectory, starting with low initial temperatures, yielded the full set of outputs shown in Figure 5-3, with the likely period of a mission encounter (near 1 AU) illustrated independently in Figure 5-4. These modelling results show that as the comet spends a great duration of this journey at a large distance from the sun, the surface will reach a stable diurnal cycle, varying within a very low temperature range of 125-135 K. The starting temperature is increased at model initiation because the initial temperature is calculated as a Solar thermal balance at the initial distance, not counting for a rotating body. Once the model starts the temperature rapidly drops to a




Figure 5-2: The comet-Solar distance plotted against time from aphelion (5.2 AU) to perihelion (1AU) within a single orbit for a Jovian comet.

more correct value. However, after approximately 750 days (as the comet crosses within 2 AU of the Sun), absorbed radiation increases significantly, raising surface temperatures considerably by the time the comet reaches perihelion.




Figure 5-3: The temperature across the top ten layers, radiative inflows and outflows of power plotted for a 12-hour rotational period surface on my Jovian comet (trajectory in Figure 5-2). The plot shows broad coloured bands for each temperature which display the full range of the daily variation in temperature for each layer.

Further analysis of a 7-day period at 1 AU (around 830 days according to the trajectory in Figure 5-2) reveals diurnal variations in surface temperature (220-350 K), that are reduced compared to those in the stable case of Figure 5-1 due to approaching equilibrium from a low temperature. These provide a guide for the temperature ranges that MIRMIS may need to observe and show that given a thin layer of dark and poorly conductive dust, the surface can warm significantly despite the cold starting condition.

There are, however, considerable limitations to this model. The lack of sublimation provides an incomplete boundary layer at the surface, and the model is sensitive to layer depth. Rerunning the simulation with 20 layers, at half the layer depth gives a surface temperature variation of 205-375 K at 1 AU (although the average across the top two layers is similar). The question of how heat is transferred through the comet's surface, as well as accurately constraining the likely dust layer depths and thermal conductivity values remain unresolved. These are key parameters that MIRMIS will help measure, as they are currently poorly understood.




Figure 5-4: The same data from Figure 5-3 with 7 days of data plotted for when the comet is at 1 AU, in order to show the diurnal variation when the mission would visit the comet.

Recognizing the limitations of my model, I consulted members of the instrument science team with expertise in thermophysical modelling of comets. Their insights provided valuable guidance, allowing us access to considerably more complex modelling results to verify the likely temperatures that MIRMIS will be required to measure.

5.1.1.2 External Thermophysical Modelling

With the support of Dr. Guilbert-Lepoutre from the Laboratoire de Géologie de Lyon, I was provided with the results of a more mature 3D thermophysical model for simulating several cometary bodies. This model is a thermal evolution model that accounts for factors such as illumination, shading, lateral heat flux, varying albedo properties, and complex boundary conditions (Guilbert-Lepoutre et al., 2011). The model is crucially able to account for lateral fluxes and heterogeneous boundary conditions that reflect albedo properties. Boundary conditions are represented as spherical harmonics in order to improve numerical efficiency, although for high-resolution models the simulation is still computationally intensive.

Simulations were run using the Comet 67P shape model for two different material profiles. Once for a ‘dust’ rocky comet, and once for an ‘ice’ based comet. In each simulation the comet was simulated throughout a one AU orbit, with a 12-hour rotational

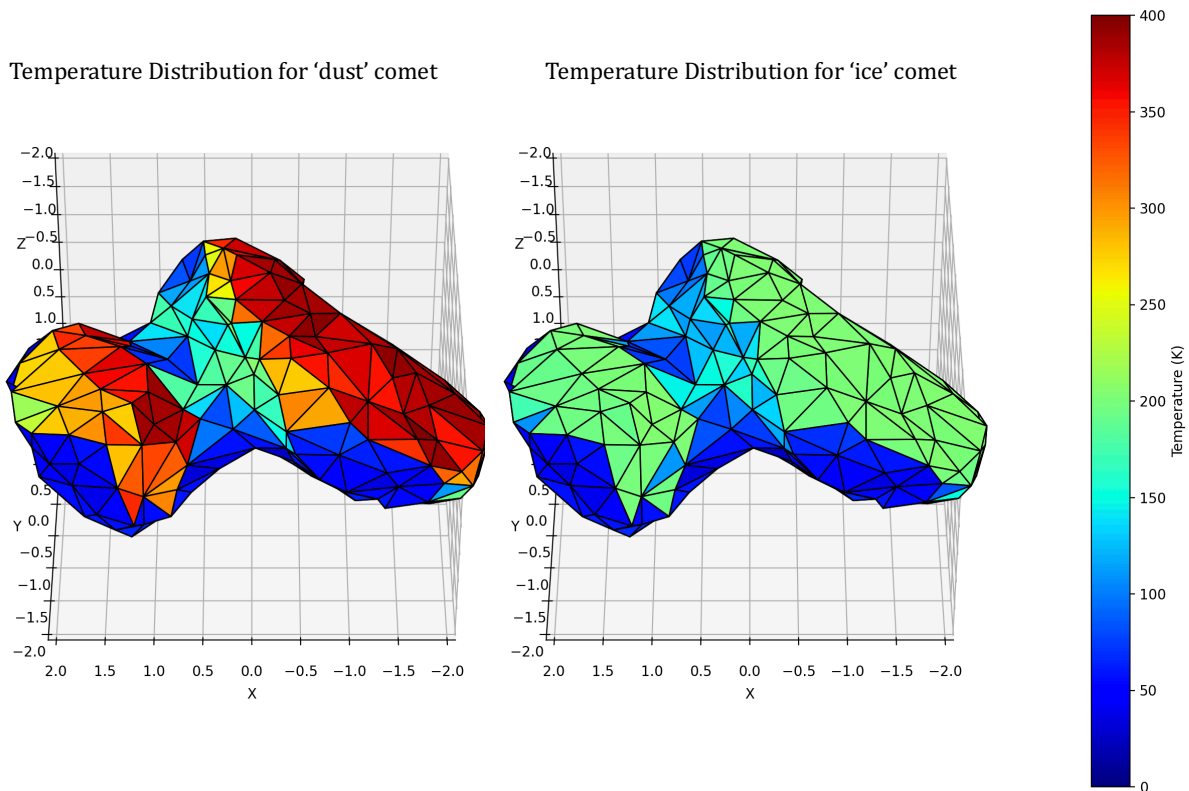


Figure 5-5: Snapshot temperature distributions for two comet simulations modelled at 1 AU. The dust rock-based comet is on the left, the icy comet on the right. The comets are viewed from along their rotational axis, with the Solar illumination coming from the top right corner. Each comet has its temperature plotted on the same scale.

period. The temperature profiles from the end of these simulations are shown in Figure 5-5.

The results of the simulation clearly indicate that an icy comet could exhibit significantly lower temperatures than a rocky, dust-based comet. The maximum temperature in illuminated regions for the icy comet was around 220 K, compared to 370 K for the dust comet. The extremely low temperatures on the night side of the comet are physically unlikely, and possibly represent an error in the model functionality, such as a lack of lateral thermal transfer between facets, or a geometric illumination issue. The dust comet's temperature roughly aligns with observations of Temple 1 (Figure 1-11) which demonstrated a maximum temperature of 330 K, and the Vega measurements of Halley which measured a sunlit temperature between 300 and 400 K (Mendis, 1988).

While the TIRI instrument is capable of measuring temperatures as low as 200 K through its broadband channel, its narrowband compositional channels are optimised for higher temperatures. If the comet TIRI encounters has a temperature profile similar to the icy model shown in Figure 5-5, the only science goals retrieved will be the thermal mapping of the nucleus, since the signal will be too low in the compositional channels. Given the consistency between my own modelling, Dr. Guilbert-Lepoutre's dusty comet simulations, and the Temple 1 and Halley measurements, I concluded that the Temple 1 measurement is likely representative, and a surface temperature of 330 K for the illuminated regions of the target that TIRI will observe is reasonable. I have therefore used this value as a nominal case temperature, and the basis for defining the operations plan in the sections that follow this.

5.1.2 Radiometric Performance Modelling

As discussed in Section 4.4.2 TIRI's compositional channels require a SNR of 100. This would enable measurements capable of determining emissivity features of a few

percent and therefore identification of spectral shape related to mineral and ice composition. This SNR requirement drives the duration over which the detector needs to integrate as, in the absence of instrument drift, longer integrations gather more signal and therefore have an improved SNR. I used a TIRI radiometric model, similar to the LTM model described in Section 3.1.1 to define the required integration time to produce an SNR of 100 in each channel.

The integration time requirement was established by generating SNR products for each TIRI channel using the radiometric model. For the surface temperature, I used a value of 350 K, based on the discussion from Section 5.1.1. Throughout the instrument's operation, from an operation angle of 85° to 173° , the distance covered by a single pixel varies from 0.26 km to 3 km. Since the nominal case comet is 5 km in diameter, I assumed that each pixel is fully filled, meaning that the radiance does not vary with the distance from the nucleus. While this may not be an accurate approximation for the last data-cubes taken at the largest operation angles, the mode of operation is expected to change at those angles, as I will discuss in Section 6.6.1.

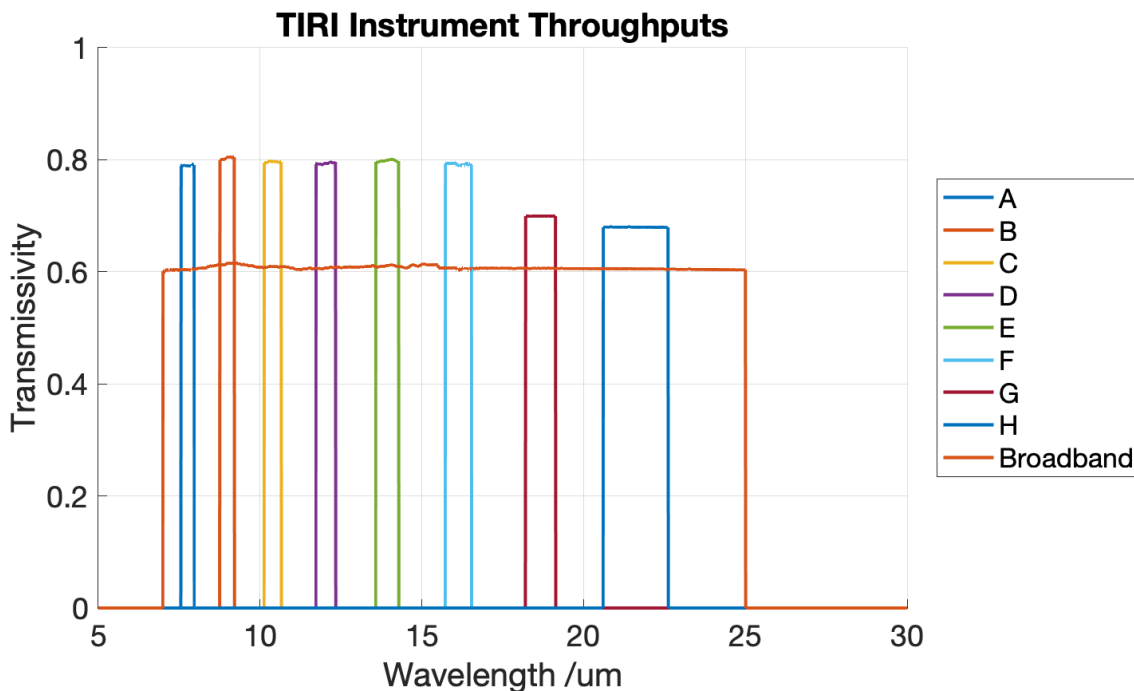


Figure 5-6: The TIRI instrument transmission function for each channel. This has been produced from the bandpass definitions, the filter transmissions of the LTM compositional filters and the reflectivity and absorptivity data used for LTM.

Since I do not have precisely modelled or measured filter transmissions for TIRI, as the filters are currently in manufacture, I have used simulated values based on the on the bandpass definitions in Table 4-1. The instrument transmission functions, plotted in Figure 5-6, are derived from a combination of TIRI's design specifications and measured optical performance from the LTM filter block, as discussed in Section 3.2. The TIRI detector's NEP used in this model, is $1.7 \times 10^{12} \text{ W/Hz}^{1/2}$, based on the detector D^* of 1×10^9 .

The SNR for each channel at a given integration time is calculated using the following equation:

$$SNR = \frac{B_{comet-detector} * A\Omega * \Delta t}{NEP * \sqrt{\Delta t}} \quad (5.7)$$

where $B_{comet-detector}$ is the comet radiance as viewed at the detector, $A\Omega$ is the etendue of the TIRI optical system, Δt is the integration time and NEP , the noise equivalent power of the system.

The structure for the model used to define the integration requirement is detailed in Figure 5-7. This model calculates the required integration time for each filter channel sequentially using an iterative process. Initially, the SNR is calculated for an integration time of 1 second. Based on the returned SNR, the integration time is adjusted, and the SNR

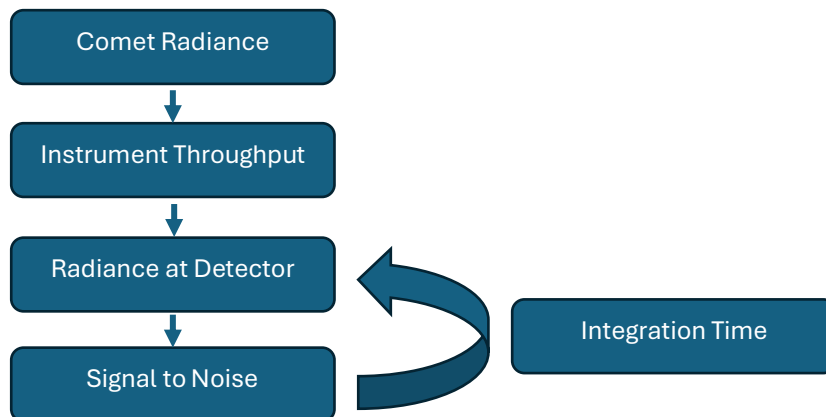


Figure 5-7: The iterative radiometric modelling process used to generate the integration time requirement for each channel.

is recalculated. This process continues until the required SNR of 100 is achieved, with the model converging to an accuracy of 0.1%.

The resulting integration times from this analysis are given in Table 5-2. These values are critical for use in the instrument pointing analysis, as they inform the required dwell times for each image. The dwell time refers to the period during which the pointing motor must hold the comet nucleus steady within the detector's field of view to allow for image capture. These values will be referenced throughout the remainder of this chapter.

Channel	Bandpass / μm	Integration Time Required /s
A	7.56 – 7.96	0.0552
B	8.76 – 9.20	0.0448
C	10.13 – 10.65	0.0400
D	11.73 – 12.33	0.0435
E	13.58 – 14.28	0.0520
F	15.73 – 16.53	0.0743
G	18.21 – 19.23	0.1452
H	20.61 – 22.61	0.0710

Table 5-2: The integration time requirements for each channel to produce an SNR of 100 for a target surface at 350 K.

Additionally, I used this model to generate a graph that shows how SNR changes with integration time, providing a reference for the science team. This graph allows for quick identification of the additional integration time required if a higher SNR is desired. The broadband channel is not included in this analysis, as its wide bandpass results in an integration time below the minimum achievable with our detector and readout electronics, meaning that achieving the required SNR for this channel does not drive the overall integration time.

This analysis demonstrates how extended integration times impact SNR. My work shows that we can meet the defined science requirement, though future work will focus on surpassing these goals and maximizing the instrument's performance throughout the

encounter window. Retaining flexibility in the operations plan is important given the uncertainty regarding the target environment.




Figure 5-8: The relationship between SNR and the integration time required for each of the TIRI channels.

5.1.3 Blurring and Other Integration Time Effects

As discussed in Section 4.3, there is a trade-off between spatial resolution and the signal captured by the instrument. In the previous section, I examined the integration time requirements needed to achieve specific SNR levels. In this section, I will address the other side of this compromise, focusing on the effects of integration time on spatial resolution.

To illustrate this, I have used simulated images of the comet generated from a spacecraft view simulation created one of my colleagues, Maisie Rashman (Rashman, ESA report MIRMIS 030). The simulation is purely illustrative and shows how the comet rotates within the field of view – due to both the motion of the comet itself and the rotation of the field of view. These images represent the geometry of the system, rather than the actual instrument’s view, for example the resolution is far higher than TIRI will achieve, and there is no accounting for positioning the target in the field of view here. The

simulation was rerun with frames generated at 10 ms intervals to allow for the analysis I present in this section.

One cost to increased integration times is spatial blur, which can manifest in two distinct ways. First, there is blur within individual images during the integration period, as the comet's image changes while the detector is being read. Second, there is data-cube blur, which occurs when images from different filters, taken at slightly different times, are stacked to form the data-cube. These blurs are both caused by changes in the detector image across the period of data retrieval and are most significant at closest approach, when the relative angular velocities between the spacecraft and target are highest.

To illustrate this, Figure 5-9 shows an exaggerated example of long integration times, demonstrating the blur of an image created with a longer exposure. Initially, TIRI's intended integration period was 1 second, until the analysis in Section 5.1.2 showed that much shorter times were satisfactory to achieve scientifically useful signal. As the figure




Figure 5-9: An image showing the effect of integration. This image is a zoomed view showing the blur created by integrating across 2.5 seconds of frames.

shows, longer integration periods give a significant loss of spatial resolution, as multiple comet angles are blended together.

Figure 5-10 meanwhile shows two frames from the flyby simulation demonstrating how much the comet will appear to rotate across a three second interval. The object imaged on the right is slightly rotated compared to the image on the left. The surfaces imaged will be positioned differently, as well as being at different angles, substantially complicating retrievals. While some rotational effects can be corrected through the data-cube co-alignment algorithm, the motion caused by the spacecraft's flyby cannot be fully compensated for. This underscores the need to minimise integration times and optimise motor speed, as both factors affect the overall scan procedure duration.




Figure 5-10: Images from the start and end of a 3-second data-cube.

To further visualise the effects of blur, I generated false-colour images by layering three images taken at different times during the flyby, each from a different filter. The images were coloured red, green, and blue, then overlaid to create an analogue to how data-cubes are formed by combining signals from each filter. If the images are perfectly aligned, the false-colour image appears unchanged. Any spatial discrepancies due to rotation or blur will be revealed through shifts in colour. It's important to note that in this simulation, the intensity of the colours represents Solar reflection and is unrelated to TIRI's measurements.




Figure 5-11: This image has been produced by overlaying the images from the first three adjacent filters from the first scan.

Figure 5-11 the false-colour image created from the first three filters shows minimal spatial distortion. Using adjacent filter means that the time between captured images is short. The colours remain aligned, producing a stable image.

In contrast to this, the image in Figure 5-12 was created by layering images from the first, middle, and last filters in the data-cube. Since there is a 2.5-second gap between the first and last images, significant rotational blur is visible. This effect is highlighted by rainbow-coloured banding along regions with steep dark-light gradients. The frames used




Figure 5-12: This image has been produced by overlaying the first, middle and final images from the first data-cube. The vertical rainbow coloured banding shows regions of high intensity gradient moving between frames.

to create Figure 5-10 are from the first and last filters in Figure 5-12, the rotation is now visualised through the addition of colour:

The first two scans are the most susceptible to both types of rotational blur due to the higher angular velocity between the spacecraft and the comet at closest approach, coupled with the rapid rotation of the pointing mirror. These are also the only scans where spatial resolution is crucial, as they are taken at the closest distance to the comet, offering the highest spatial resolution (260 metres per pixel at closest approach). The graph in Figure 5-13 demonstrates how quickly spatial resolution deteriorates as the distance between the spacecraft and the comet increases after closest approach. Within 30 seconds




Figure 5-13: The number of pixels that a 5 km target nucleus will subtend at a certain time after closest approach for a nominal case flyby.

of the encounter, half of the spatial resolution available at closest approach will have been lost due to distance. To meet the spatial mapping requirement of 500 metres per pixel, data must be captured before the comet's nucleus appears smaller than 10 pixels across.

In addition to blur, spatial resolution is also degraded by pointing accuracy and uncertainties during data-cube construction, both of which will be discussed in subsequent sections. Given these constraints, the mapping data required to satisfy our spatial resolution requirements must be taken as early in the encounter. This highlights the pressure to minimise integration times while still meeting the SNR requirements.

5.1.4 Mirror Pointing Modelling

A key component to defining the operations of the instrument is understanding the pointing mirror positioning through the encounter. As discussed previously, TIRI creates its data-cubes by stacking images on the comet nucleus taken through each filter. In order to change the filter being used, the mirror must move a small amount, changing the direction of the instrument pointing vector and thus moving the nucleus in the field of

view. We refer to the process of moving the nucleus through each of the TIRI filters as scanning. In this section I will describe the modelling undertaken to derive the necessary movement patterns throughout the operation, to enable us to produce data-cubes throughout the flyby. The motion of the NIR/MIR mirror is much simpler and is also discussed (shown on the left in Figure 4-1).

For this model, I have assumed that the TIRI motor assembly has a maximum rotation rate of 60 seconds per revolution, or 6 degrees per second. Since the pointing mechanism is driven by a stepper motor with 29,000 steps per rotation, this allows for a stepping control of 0.0124 degrees per step. Given the motor's rotational velocity, the assembly is capable of moving at a rate of 483 steps per second, or 0.483 steps per millisecond.

The dwell time requirements derived previously play an important role in determining the pointing mirror's motion. Initial power modelling for TIRI was based on a worst-case scenario with a flat 1-second dwell time across all filters. However, in the operational plan described here, I have used improved dwell time estimates to define how the mirror moves throughout the encounter, assuming that TIRI is able to achieve the minimum integration times calculated previously. As a result, each scan only utilises a narrow portion of the windows allocated for data-cube acquisition in the requirements.

Throughout the encounter the nucleus is tracked by the pointing mirror. The tracking path through the encounter (for the nominal case) has been plotted in Figure 5-14. For the NIR/ MIR motor this is all the motion required, this angular position will keep the nucleus in the centre of each of those detectors. In the case of TIRI the movement to access each filter is just a small perturbation from this curve, and I will now explain how this deviation is modelled.

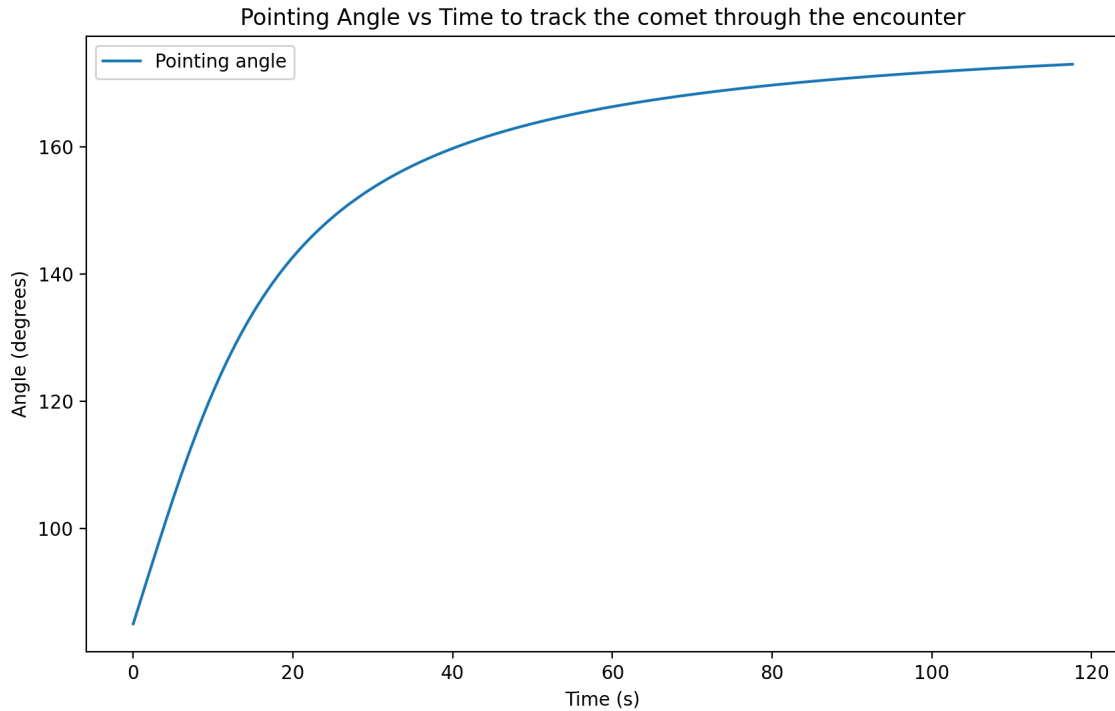


Figure 5-14: The pointing angle for the pointing mirror required through the encounter to keep the comet centred within the centre of the TIRI and NIR/MIR FoV throughout the encounter. The pointing angle is defined as the angle between the spacecraft velocity vector and the spacecraft-comet vector. Time is measured from the start of the encounter.

The positions of each TIRI filter in the detector field of view was displayed in Figure 4-3. The comet is scanned across the field of view using the single rotational axis of the pointing mirror. As there is only a single axis for rotation, we only have control of the comet position along a line within the field of view. Due to the rotation of the field of view with the operation angle (visualised in Figure 4-4), this line is an arc, which I refer to as the 'arc of control'. The arc of control defines the possible positions where the comet can be placed at any given operation angle. The shape and movement of this arc throughout the encounter are illustrated in Figure 5-15. At the beginning of the encounter (operation angles near 90°) turning the mirror moves the comet horizontally in the field of view, at the end of the encounter it moves the comet vertically. This explains the filter layout with each filter repeated, once vertically and once horizontally.

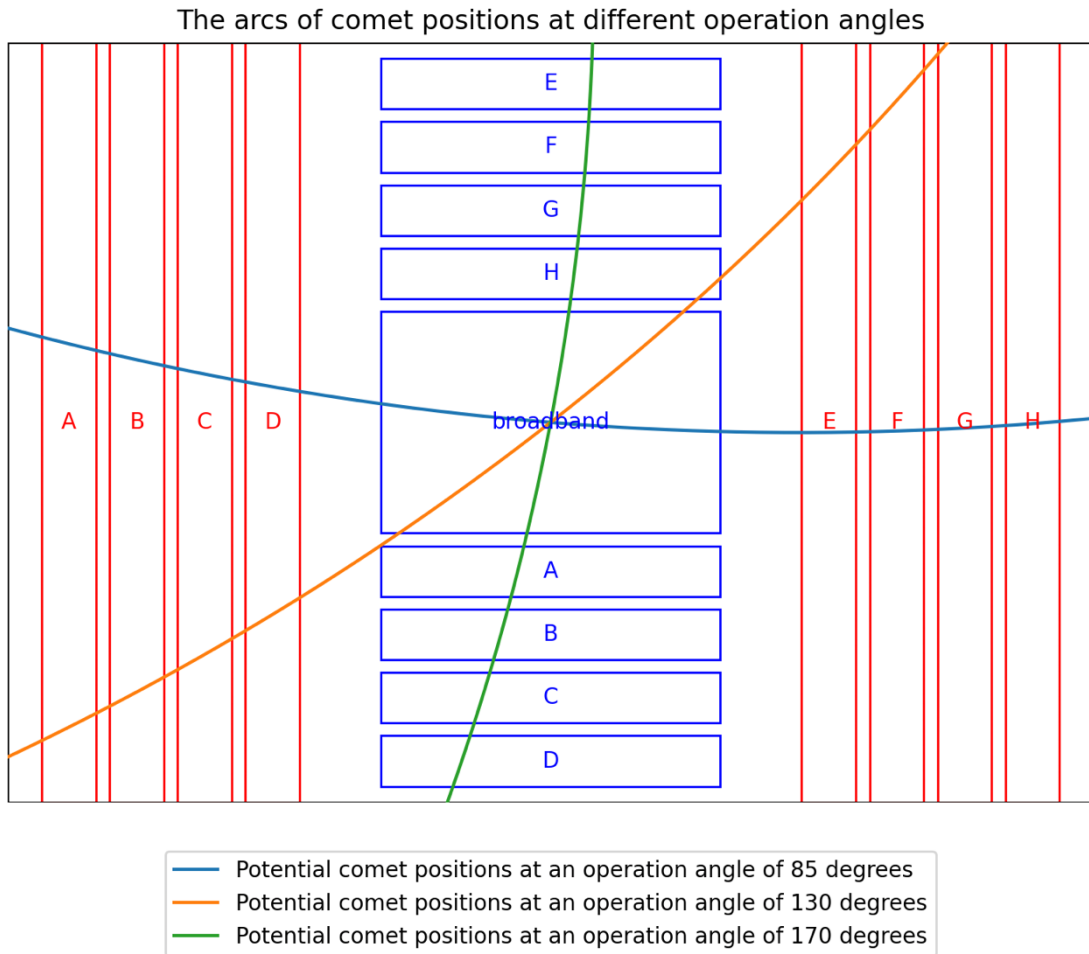


Figure 5-15: The arc of control for 3 different operation angles, overlaid with the position of the filters in the detector field of view.

The arc of control is defined from the geometry of the system, with the curve of the arc based on a circle drawn around a centre point (x_0, y_0) which is defined relative to the size of the detector field of view which is being rendered. The arc is then rotated via a rotation matrix, using the operation angle at that point, to produce the rotated central coordinates (x_{rot}, y_{rot}) .

$$\begin{pmatrix} x_{rot} \\ y_{rot} \end{pmatrix} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} x - x_0 \\ y - y_0 \end{pmatrix} + \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} \quad (5.8)$$

Once the arc is defined as a function of the operation angle, the intersection between the arc and each filter can be calculated at any point during the encounter. Python algorithms I developed automatically constrain these intersections and calculate the angular offset required to position the nucleus at the correct location. This offset is always

relative to the operation angle, which is the pointing angle which would place the nucleus in the centre of the field of view.




Figure 5-16: The relationship between the angular offset required to place the nucleus in the centre of each filter at any given operation angle through the encounter. These are a worst-case scenario since this only considers intersection with the centre of the filter, and there is some benefit to be gained by considering the edged of windows (at cost of risk).

Figure 5-16 illustrates the required offsets for each filter at different operation angles. The lines initially curve upward with increasing operation angles, as the arc of control moves to the diagonal, increasing the distance between the centre of the field of view and the filter intercept. The secondary lines decrease as the operation angle increases for the same reason. The discontinuities indicate regions where the arc of control does not intersect with either filter window, meaning data cannot be retrieved at those angles. For this reason, data-cubes are typically acquired early or late in the flyby, discussed further in Section 5.1.5. The filter layout can now be seen as an effort to minimise the gaps in this plot, subject to limitations on the filter block design, such as providing adequate support for the filters themselves to survive launch.

With computational tools to find the angular offset required to produce a filtered image defined, the next task is to relate these positions to a movement plan for the TIRI motor. I created a time-series dataset based on the flyby geometry and velocity (as shown

in Figure 5-14), and then dynamically layered the targeting onto this timeseries data. The algorithm for this takes an initial position and time, and an order in which to scan through the filters. The algorithm targets each filter by calculating the pointing angle from the operation angle and offset for that filter at that time. The model targets each filter for the integration period for that filter (taken from the requirements defined in Section 5.1.2), and after that time has elapsed it updates its target to the next filter in the list. The mirror position then moves at the maximum motor speed until it coincides with the target angle, at which point the next integration period is begun. By this method the model simulates the full scan, accounting accurately for the full geometry of the flyby and filter assembly, integration time requirements and motor max velocity.

Scanning methods for TIRI are plotted in Figure 5-17. The orange line always shows the pointing target and the blue the pointing position. Where the lines overlap the target is centred in the filter and the images are being produced. Once an image is taken, the target immediately switches to the next filter's required position, and the pointing position moves at maximum speed to meet it. The diagonal gradient that each filter target follows is the curve to track the comet stably in the detector view (Figure 5-14) and the difference between these filters is the offset from Figure 5-16.

The two position plots show two different strategies for scanning across the filters (we almost certainly want to scan across every filter sequentially either forward or backward). In the top plot, the scan occurs in the opposite direction to the comet's motion across the field of view, allowing the motor to move quickly between filters. In contrast, the bottom plot shows the scan moving in the same direction as the comet, resulting in slower filter changes, as the instrument must "overtake" the comet's motion. Despite identical integration times and motor velocities, the top scan takes only 1.4 seconds, compared to 4.5 seconds for the bottom scan. Figure 5-17 demonstrates the impact of operational plan choices on measurements as despite using the same integration time and motor velocity,

the top scan takes 1.4 seconds rather than 4.5 on the bottom scan. As discussed in Section 5.1.3, minimizing the time between images is crucial to reducing blur in the data-cube, making the top scan preferable.




Figure 5-17: The movement pattern to scan a single data-cube. In the top graph the model has been run scanning from Filter A through H in order, for the bottom graph the model was run in the opposite direction, from filter H to filter A.

However, the top scan introduces challenges. The motor must rotate bidirectionally, and the greater velocity differences between tracking the comet and moving between filters may introduce jitter, positional uncertainty, or require additional stabilization time, reducing the overall benefit. These factors will be further investigated through laboratory testing, but for now, I assume the benefits of reduced blur outweigh the potential drawbacks. Where possible, the shortened scan duration will be preferred.

5.1.5 The Full TIRI Pointing Plan

With the tools to define TIRI's motor pointing plan now developed, the full pointing procedure for the encounter can be created. This process involves setting the operational parameters for each data-cube (filter order, start position, and integration times), and combining the outputs into a comprehensive motor movement plan.

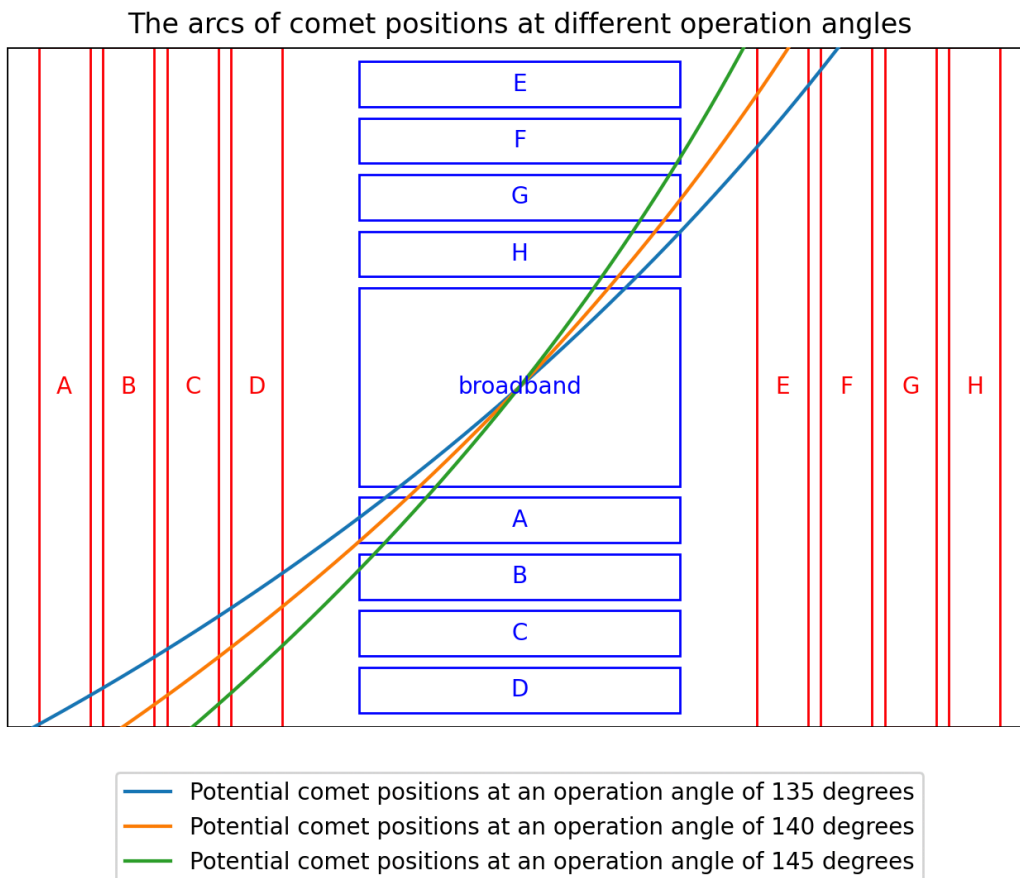



Figure 5-18: The arc of control for an operation angles of 135°, 140° and 145°. This is the range of positions required for the third data-cube.

For the first and second data-cubes, the pointing mirror scans the comet nucleus across the centre of each vertical filter. For the fourth and fifth data-cubes, the mirror scans the nucleus through all eight horizontal filters. However, Figure 5-18 highlights the difficulty in obtaining a full data-cube in the intermediate stages of the encounter. At an operation angle of 140°, the arc of control does not intersect with filter F. Attempting to scan earlier or later to include filter F would exclude either filter G or E, making it impossible to retrieve a full data-cube at this position.

To address this issue, I redefined the third data-cube, since the originally proposed position was physically unretrievable (initial plan in Table 4-3). By moving the third data-cube earlier in the encounter, to an observation angle of 123°, I not only made the scan possible but also improved the spatial resolution of that data-cube set of measurements to provide another map of the nucleus surface.

The full details for the third data-cube are given in Figure 5-19. The filters in this observation have been scanned from H to A on the vertical filters. Priority is given to measuring the comet with filter H at the start of the scan, as H will no longer be accessible later in the observation. If the filters were scanned in the opposite direction, from A to H, the observation would need to start earlier, around 119°, in order to capture H at the end



Parameter	Value
Start Position	123°
Filters Used	Vertical
Filter Order	H - A

Figure 5-19: The movement plan for the third data-cube, the arc of control at the beginning of that scan and the set of parameters used to define that movement pattern.

of the scan. However, given the current motor velocity assumptions (6 degrees per second), the motor would not be able to move fast enough from the end of Data-cube 2's scan to the starting position for the third data-cube. Moving from H to A is slow, regardless

of whether data is being collected or not. Full details for every data-cube, in the style of Figure 5-19, are given in the appendices (Section 9.1).

The complete motor movement plan, as derived from this pointing analysis, is shown in Figure 5-20. It combines the movement required for each data-cube scan and directs the motor to prepare for the next scan as soon as the previous one is completed.




Figure 5-20: The TIRI motor pointing plan across the entire encounter, incorporating the five data-cube scans.

The pointing plan will be used for several further analyses, looking at the motor commanding and development (Section 5.4) and within the TIRI instrument simulations (Chapter 0). One immediate use of the full pointing plan is its integration into the MIRMIS operations timeline, which allows us to verify the instrument's power consumption. The MIRMIS operations timeline this generated is shown graphically in Figure 5-21 and the details are contained in Table 5-3.




Figure 5-21. the MIRMIS mission timeline around the flyby, from instrument power on 24 hours before closest approach through to power-down 24 hours after.

In the nominal flyby scenario, MIRMIS is expected to be powered on for 48 hours, spanning 24 hours on either side of closest approach (CA). This time window allows for communications and instrument checks before CA, and data processing and downlink after CA. It may be possible to reduce these time windows in the future. MIRMIS will perform internal calibration starting 30 minutes before CA and will collect science data for 5 minutes after CA, with an additional 25 minutes allocated for data transmission.

Action	Operation Angle / degrees	Time
Instrument Power-on		24 hours before CA
Comms (checkup)		24 hours before CA
Comms (if needed)		18 hours before CA
Comms (if needed)		15 hours before CA
NIRMIR Cal Procedure		290 seconds before CA
TIRI Calibration Procedure		100 seconds before CA
Observation 1 Start	85	1.25 seconds before CA
Observation 2 Start	106	4 seconds after CA
Observation 3 Start	123	9 seconds after CA
Observation 4 Start	165	53.5 seconds after CA
Observation 5 Start	172	101.6 seconds after CA
TIRI Calibration Procedure		225 seconds after CA
NIRMIR Cal Procedure		300 seconds after CA
Data Download		TBC
Instrument Power-Down		24 hours after CA

Table 5-3: The MIRMIS operations timeline.

5.1.6 Alternate Pointing Plans

At the time of writing, full testing of the MIRMIS motor mechanisms is being conducted at Oxford. One early finding from this testing is that the assumed worst-case

motor maximum velocity may have been a significant underestimate. Preliminary results suggest that the mechanism may allow for a rotational speed exceeding 12 degrees per second, which is more than double the value used in my initial modelling.




Figure 5-22: The pointing position and pointing target for the entire encounter utilising a 12 degree per second maximum motor speed. Time is measured relative to closest approach.

In response to this, I developed an alternate pointing profile to support ongoing motor testing, incorporating this higher motor velocity. The new pointing profile, shown in Figure 5-22 demonstrates how the increased speed allows the third data-cube to be scanned in the same direction as the first and second data-cubes. This adjustment reduces the overall time required to retrieve data for the third cube.

Not only did this alternate profile support motor testing and development, but it also validated the flexibility and modularity of the operations modelling tools I developed and have detailed here. These tools enabled me to rapidly produce an updated operations profile as operational parameters evolved.

5.2 MIRMIS Power Consumption

The requirements contained in Table 4-5 specify MIRMIS's peak power draw of 20W, and average of less than 11.2 W, and necessitate the provision of a detailed power draw profile. To verify that MIRMIS can achieve its objectives without exceeding these limits, and to provide a framework for future power consumption analysis, I developed a Python model to predict the instrument's power profile.

The predicted non-motor rotating power consumption of MIRMIS is 10.3 W; 5.3 W from TIRI and 5 W from NIR/MIR, when no motors or heaters are running and without any margin. The power consumption of a single motor is estimated to be 6 W. With both motors running, the expected instantaneous power draw rises to 22.3 W. In the nominal flyby scenario, TIRI is expected to be active and collecting data for less than 5 minutes, as the comet nucleus will become sub-pixel in size shortly after closest approach. Over this 5-minute period, the estimated average power consumption of MIRMIS varies depending on the time frame. Given that the spacecraft operates at 28 V, the maximum instantaneous current draw for MIRMIS is approximately 0.8 A (22.3 W), without margin.

Since TIRI is required to take measurements across its five data-cubes and remain inactive between scans, the motor alternates between periods of movement and inactivity. In my model, once an observation is completed, the instrument moves immediately to its starting position for the next scan, and then waits. This movement plan follows the pointing analysis discussed in Section 5.1.5, with the motor's true pointing position modelled at a high resolution of 200 μ s.

5.2.1 Power Modelling

To calculate the power draw, I used this high-resolution model to track the TIRI and MIR/NIR motor positions. The model accounts for motor-specific parameters such as

step size, pulse duration, and maximum rotation rate. A sample of the true motor pointing position is shown in Figure 5-23. The stepped movement data allows me to simulate individual motor steps and the resulting instantaneous power draw when moving the mirror.




Figure 5-23: A sample of the discrete motor pointing, representative of the motor's stepped movement. This sample demonstrates the end of the first filter scan in the first observation, and the start of the movement to centre the second filter.

Each motor pulse draws full power for 0.5 ms. The motor velocity is controlled by adjusting the time between pulses. The power profile generated from this model is shown in Figure 5-24, where the velocity control can be seen by regions where the motor is not powered. The constant 6 W sections correspond to the motor moving at maximum speed between filters.




Figure 5-24: Sample power draw of the motor scanning across the first two filters of the first observation. The lower rotation velocity of tracking the comet in the filter is visible with pauses between pulses when compared with the maximum velocity attainable which is used when moving between filters and between observations.

This power dataset can then be binned into a lower resolution to produce a plot for the power required throughout the encounter. The average power across each 0.02 s bin has been calculated in this manner and the binned power draw for TIRI’s motor across the entire encounter is displayed in Figure 5-25. The long periods of inactivity between the third and fourth, and fourth and fifth data-cubes will drop the average power draw considerably.




Figure 5-25: Power draw across the encounter for the TIRI Motor, calculated for 0.02s time buckets.

The power consumption of the NIR/MIR motor was calculated using the same model applied to its position data from Figure 5-14. The power curve produced for this motor is given in Figure 5-26 averaged over 0.25-second intervals. The drop in power consumption later in the observation is due to the lower motor velocity.

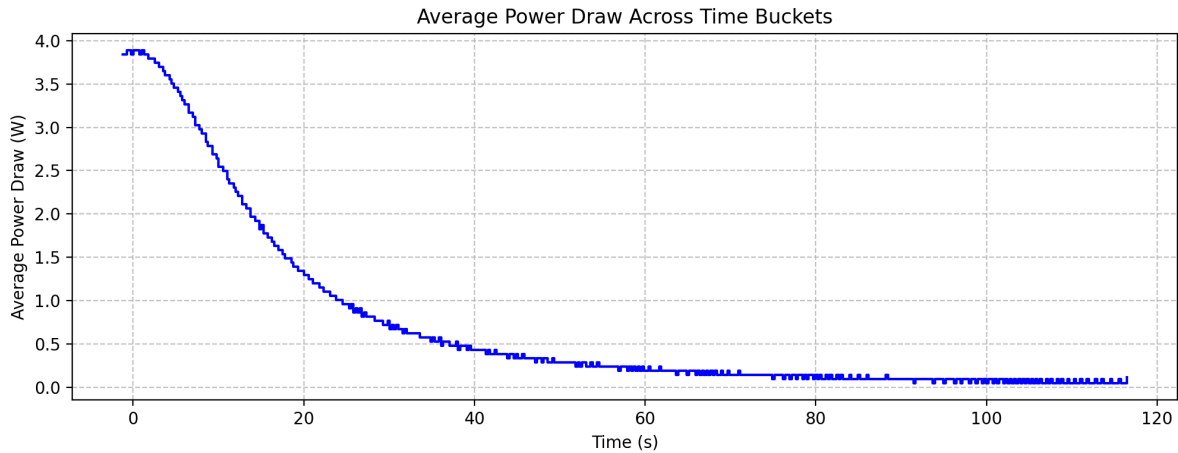


Figure 5-26: The bucketed power draw across the encounter for NIR and MIR, binned every 0.25s.

The individual contributions to the total instrument maximum power draw are detailed in Table 5-4, where these powers are given without any margin. The power draw for the instrument's electronics is well-characterised, as the engineering boards for all MIRMIS modules have been tested extensively. However, the motor power consumption remains the most uncertain aspect, as it depends heavily on the FPGA control software, which is still in development. The current model assumes worst-case power consumption for the motors.

Power Draw Contributor	Maximum Power Draw /Watts
TIRI Motor	6
NIR/MIR Motor	6
TIRI Electronics	5.3
NIR/MIR Electronics	5
Total	22.3

Table 5-4: The individual contributions to the MIRMIS maximum power draw.

To estimate the total power usage over the encounter, I combined the subsystem power profiles with the operational timeline detailed in Table 5-3. The contributions of each MIRMIS subsystem, including motor movements for instrument calibration, are plotted in Figure 5-27. These profiles have been combined to produce the overall mission power profile.




Figure 5-27: A decomposition of each of the power draws within MIRMIS across the period of the encounter. The modules NIR and MIR are displayed separately from TIRI, and each motor is displayed separately. The motor powers are unbinned.

The total instrument power draw is illustrated in Figure 5-28, showing the baseline from each module’s electronics, punctuated by motor movement during calibration and observation.




Figure 5-28: Total power draw for MIRMIS through comet encounter: 0 time corresponds to closest approach. This is unbinned power.

The requirement on MIRMIS stipulates that average power is kept below 12 W across any 5-minute period during the encounter. In order to evaluate this requirement, I have calculated the rolling 300-second average of power across the instrument’s total power profile. The result, displayed in Figure 5-29, demonstrates that the average power

draw across any 5-minute window remains below 12W and thus the requirement is satisfied.




Figure 5-29: 300-second rolling average total power draw for MIRMIS across the encounter, using data taken from the power draw in Figure 5-28. Time is measured relative to closest approach.

5.2.2 Alternative Power Scenarios

Several power-reducing schemes do exist should further developments prove to increase power draw. One potential approach involves switching off the motor encoder and zero position sensor, which together draw about 1 W. Additionally, the current motor function model likely overestimates the power consumption. If the motor rotates faster than modelled, it will spend less time drawing 6 W pulses. Moreover, the model assumes that each motor step draws power for the full duration of the step, whereas in practice, the pulse only needs to accelerate the motor to speed and can be turned off as the step completes. At the time of writing a full investigation into this approach is underway at Oxford, using the final TIRI mechanism design. Reducing pulse duration would not only lower the average power consumption but also allow for staggered operation of the NIR/MIR and TIRI motors, ensuring they are never running simultaneously. This would further reduce peak power draw.

5.3 TIRI Filter Selection Modelling

The goal of maximizing signal within the TIRI filter bandpasses raised the question of whether wider filters could be used (original design specification in Table 4-1). The original compositional filters are narrow enough that it would be possible to more than double the width of most filter positions without causing overlap between channels. This change would roughly double the signal size at the detector, improving the SNR at the expense of reduced spectral resolution. A reduction in spectral resolution would lead to a corresponding loss in the precision of the wavelength position of the recorded signal, and compromise sensitivity to variations in composition. The two sets of potential filter bandpasses are displayed in Figure 5-30.




Figure 5-30: Proposed TIRI filter bandpasses. The narrow blue bandpasses are the originally proposed filters. The wider orange bandpasses are the hypothetical wider channels we are investigating.

This investigation was partly motivated by the thermophysical modelling conducted by the MIRMIS science team and by my own geospatial modelling, which suggested that depending on the flyby trajectory, MIRMIS could spend a significant portion of the encounter observing unilluminated regions of the comet. By the time TIRI acquires the final data-cubes, up to 84% of the visible comet surface could be from the nightside.

Combined with the thermophysical modelling results for an ice-dominated comet surface (Section 5.1.1.2), which predict much lower surface temperatures than TIRI was originally designed to measure (low-temperature performance was traded for improved spatial resolution when adapting LTM into TIRI), there was interest in exploring whether wider filters might be better suited for these lower temperatures. I undertook this investigation, and the results are described here.

To explore the advantages of wider filters, I developed an instrument simulation model in MATLAB to test both filter configurations. The simulations were designed to generate instrument measurements with uncertainty values, which could then be evaluated by the broader science team for their utility. The model aimed to quantify the impact of wider filters on the data product and provide a basis for further discussion.

The simulation used an array of different representative cometary radiance spectra at different surface temperatures, generated through the Planetary Spectrum Generator (PSG) (Villanueva et al., 2022). The spectra utilised were the same set of mineral spectra utilised in the original filter selection investigation (Shirley et al., 2022) and cover six compositional mixtures of crystalline and amorphous silicates combined with six temperature profiles. Full details of the mineral spectra can be retrieved from Shirley et al., 2022. These spectra were fed through the instrument's transmission function, using the same TIRI radiometric model described in Section 5.1.2 to calculate the power incident on the detector. I also modelled noise levels to provide error bars for the measurements. The integration times for each filter channel were based on those defined in Section 5.1.2.

The simulation was run multiple times, testing both the narrow and wide filter bandpasses across a range of surface compositions and temperatures. The results for three cases; 300 K, 250 K, and 200 K on an identical composition (a dusty, rocky nucleus formed of silicates – full details for the PSG composition is given in the appendix in Figure 9-7), are shown in Figure 5-31, Figure 5-32, and Figure 5-33 as examples of the results produced for

the science team. This material is not carbonaceous but was chosen as it exhibits features in the same wavelength region (centred on 8-12 microns) as the features in representative spectra formed from carbonaceous meteorites and OSIRIS-REx (Donaldson Hanna et al., 2021, Bates et al., 2021).

In Figure 5-31 the results for a 300 K comet surface are shown. At this temperature, there is ample signal, even with the narrow filters. The noise is relatively small compared to the differences between data points, meaning uncertainty does not limit our ability to identify spectral features. In this case, the loss in spatial specificity that accompanies the wider filters brings no meaningful increase in spectral performance.

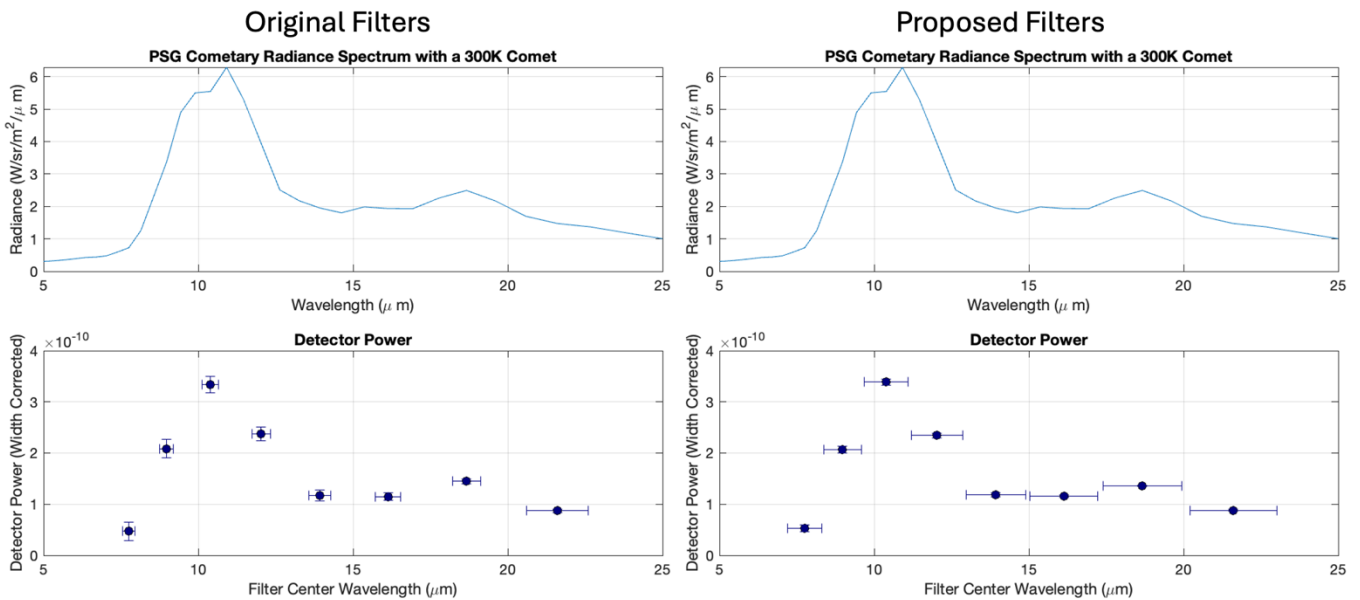


Figure 5-31: The simulated radiance measurement for a 300 K comet, with the narrow filters shown on the left and the wider filters on the right. The top plots show the PSG generated radiance, and the bottom plots show the measured datapoints with noise bars.

Figure 5-32 presents results for a cooler comet with a surface temperature of 250 K. In the original filters with the narrow bandpasses, the more substantial uncertainties now show interference in distinguishing between the filter channels. For example, the uncertainties in channels C and D (at 11 and 12.5 microns respectively) overlap, meaning noise could obscure the difference between their radiances, despite the true radiance values being substantially different. In the proposed filter case, with wider bandpasses, the correct relationship between these data points remains clear despite the noise. Therefore, the information regarding temperature and composition that comes from the relative

radiance magnitudes is preserved. For this 250 K surface, the wider filters have generated a meaningful improvement to the data retrieved.

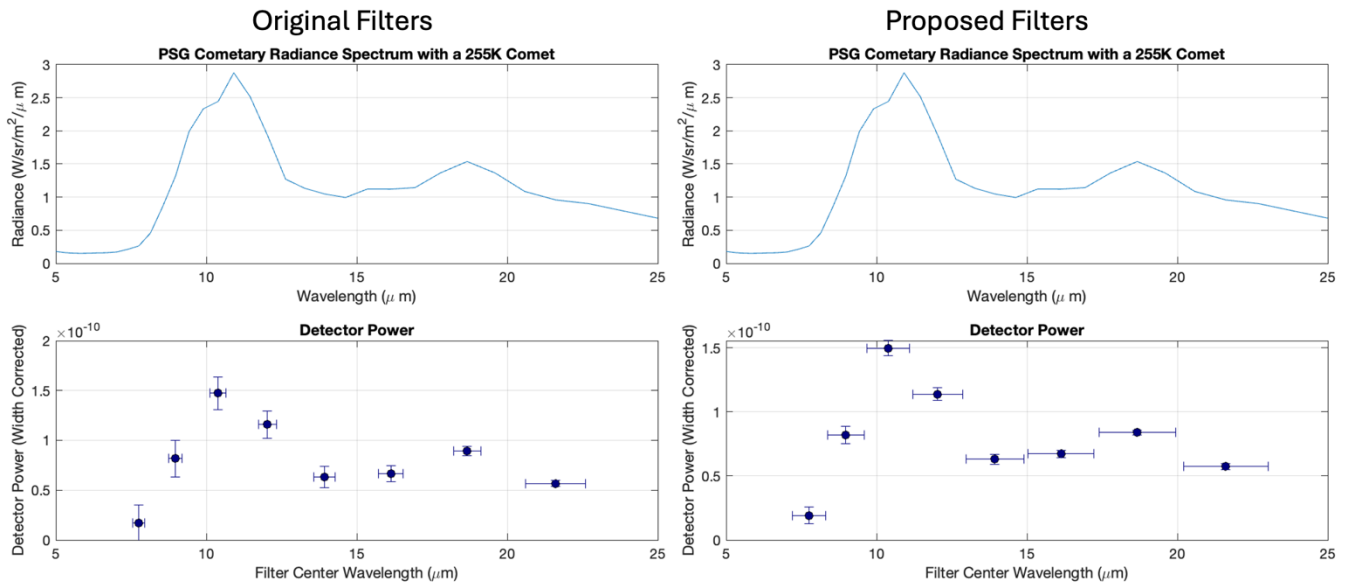


Figure 5-32: The simulated radiance measurement for a 250 K comet, with the narrow filters shown on the left and the wider filters on the right. The top plots show the PSG generated radiance, and the bottom plots show the measured datapoints with noise bars.

Lastly, Figure 5-33 shows the simulated measurements for a 200 K surface. In the narrow bandpasses, the signal is so weak relative to the noise that compositional retrievals are not feasible. This outcome is expected, as TIRI’s compositional channels were not designed to perform well at such low temperatures. While the wider bandpasses perform somewhat better, the SNR is still too low to differentiate between the power levels in each filter, making the improved SNR meaningless in this scenario. TIRI is incapable of probing composition at this temperature without considerably trade-offs such as substantially increased integration times.

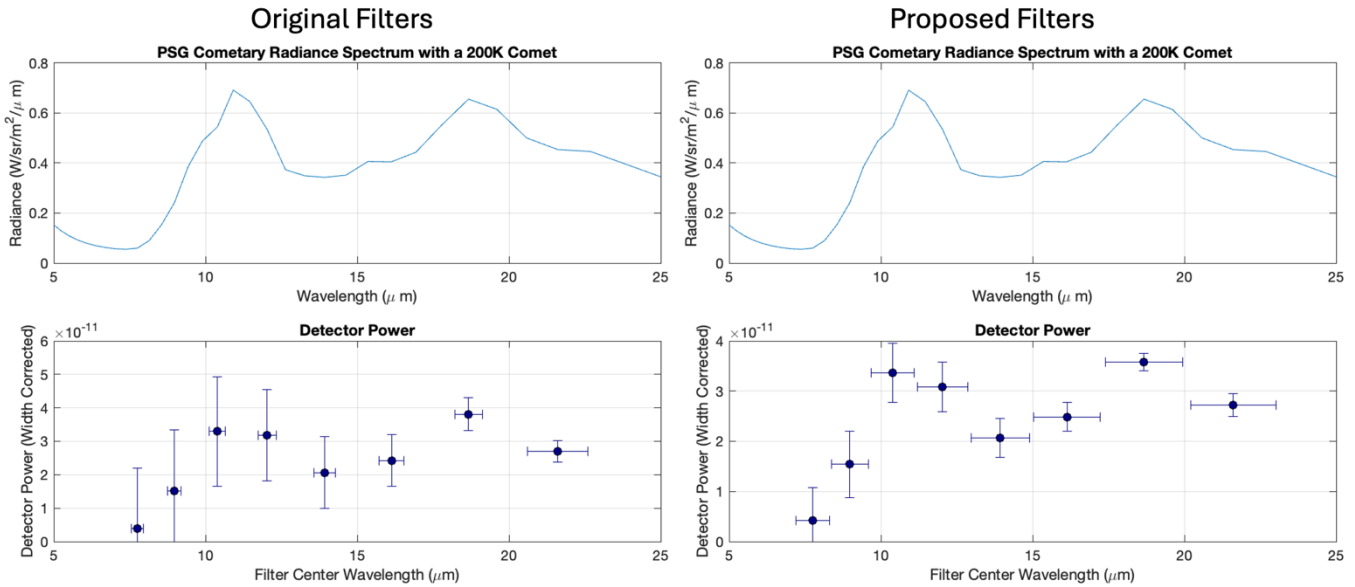


Figure 5-33: The simulated radiance measurement for a 200 K comet, with the narrow filters shown on the left and the wider filters on the right. The top plots show the PSG generated radiance, and the bottom plots show the measured datapoints with noise bars.

This analysis explored the relationship between surface temperature and the instrument's ability to detect spectral features. As the temperature decreases, the instrument's capacity to distinguish between features, and therefore determine composition, diminishes. However, larger spectral features can still be detected at lower temperatures compared to more subtle features. The graphic in Figure 5-34 was produced to visualise the benefits of adopting wider bandpass filters. The plot shows the expected performance of TIRI, with the surface temperature and compositional feature size governing the detectability of features. The curves on the graph represent the boundaries of detectability - if the comet surface falls below these curves, the instrument will be unable to characterise the composition. Broadening the filters lowers these curves, decreasing the likelihood that the comet surface will fall below the detection threshold.

The core challenge in making this trade-off is the uncertainty of the target. Since the specific characteristics of the target comet are unknown, we cannot definitively place it on this plot. To assess the probability of benefiting from wider filters, a comprehensive population analysis of likely targets would be required, modelling their formation, trajectories, and thermophysical properties. This work is currently being undertaken by




Figure 5-34: A qualitative graphic describing the trade-off space between feature size and surface temperature. Any comet surface will exist within these two parameters. The curves separate the regions where spectral features are detectable (above the curve) from regions where spectral features are undiscernible (below). Curves for narrow and wide bandpass filters are shown.

members of the MIRMIS science team. However, due to time constraints, filter design decisions had to be finalised before this analysis could be completed. In the end, a compromise was reached: slightly broader filters were chosen, retaining much of their original narrowband characteristics while offering some improvement in signal.

5.4 Motor Commanding

With the spatial and spectral resolution analysed for SNR, the performance of the pointing motor now needs to be considered. The spatial resolution of TIRI's images is constrained not only by the integration time but also by the pointing accuracy of the instrument. Uncertainty in motor control can lead to errors in data alignment and may introduce blur if the detector image is not stable during integration. This is especially important because the motor will move while the detector is actively integrating. To ensure precise pointing throughout the encounter, I conducted an analysis of the control and commanding of the pointing mirrors, aiming to relate pointing uncertainty to the requirements for the commanding software design.

The on-board flight software for TIRI is being developed by an external supplier, KISPE, who specialising in space and telecommunications software. Given TIRI's unique operational requirements as a scanning radiometer, it is essential to clearly define the software's interactions with the spacecraft and the instrument's CDHU. While the spacecraft and CDHU requirements are well-defined, the operational requirements for TIRI had not been fully constrained prior to the development of the operations plan outlined in Section 5.1. The purpose of this modelling is to provide a framework and tools to evaluate the requirements on the software relating to the instrument commands. The modelling will use the known instrument parameters to translate operational requirements into software requirements. The initial modelling will focus on the motor commanding, as the most operationally vital part of the commanding suite, but it can be further extended to consider additional constraints such as taking images, detector windowing and data extraction.




Figure 5-35: The position of the TIRI pointing mirror motion across encounter for the nominal case as defined by the operations plan. This is the high-resolution dataset which needs to be reduced and recreated by motor commands.

The operations plan defines a profile (reproduced here in Figure 5-35) for the nominal CI encounter which represents the most demanding observation as it is at the upper limit of flyby velocity. The motor position profile output from the planning process is defined at a 1 ms resolution throughout the encounter, based on geometry and maximal motor performance parameters. However, the instrument cannot control motor position at this resolution because uploading such high-resolution time-series data to the FPGA would be too costly in terms of data volume. To address this, my task was to define a command regime that compresses the motor position profile into a computationally efficient set of commands while minimizing data loss and maintaining pointing accuracy.

Discussion with the software development team led to the definition of three questions that need to be established to support and inform their work.

- How many commands are necessary over the flyby?
- How should the motor be commanded to move?
- What are the requirements on the timing uncertainty of commands?

5.4.1 Analysis Methodology

To address these questions, I developed a model to simulate the motor movement and evaluate different command schedules. The movement profile in the operations plan was based on data-cube requirements and a conservative estimate of the motor's maximum velocity (see Section 5.1). I modelled the motor performance for this movement profile to derive an efficient command schedule that minimises the number of commands while maintaining pointing accuracy. Once the command generation was verified, I introduced timing uncertainty into the model and evaluated its impact on pointing accuracy. This analysis provided insight into how timing errors propagate through the system and allowed me to relate time uncertainty to the operational requirement of pointing accuracy.

The model proceeds in steps to evaluate efficient ways to move from the movement profile defined in the operations plan to instrument commands.

- **Movement Position Analysis:** I first analysed the motor position profile to reverse-engineer a suitable command schedule.
- **Command Simulation:** New pointing profiles were generated from the command schedule, and the errors in these commands were compared to the original profile to evaluate their suitability.
- **Time Uncertainty Simulation:** I then introduced timing uncertainty into the command schedule and evaluated the resulting pointing accuracy. By examining the positional residuals, I was able to relate timing errors to pointing accuracy requirements.
- **Further Investigation:** With these tools, I explored different operational modes and their impact on pointing accuracy. The relationship between timing uncertainty and pointing accuracy was analysed to define the command execution timing requirements.

The commands will be evaluated based on the process laid out in Figure 5-36, where the pointing residuals are the ultimate measure of final pointing accuracy and will be compared against the pointing accuracy requirement to define the success of commanding regimens.



Figure 5-36: the pipeline for evaluating TIRI commanding schedules against the pointing sequence defined by the operations plan.

5.4.2 Simulation of the Motor Commanding

In this section, the full breakdown of the modelling undertaken will be provided, displaying the processes used, as well as the assumptions made. This will enable this work to be recreated and improved on in future. This modelling is all executed in Python and structured to be modular to interact with other MIRMIS modelling efforts, such as updated operations plans and feeding the output into instrument simulations.

5.4.2.1 Movement Position Analysis

The first problem we needed solved was to understand how to turn the positional profile plotted in Figure 5-35 into a series of discrete commands. The intended methodology for operating the TIRI motor is to send commands to the FPGA which set a motor velocity which will then be maintained until the next motor command is received. This will have the effect of changing the smooth profile into one formed of straight lines between command executions. The first task was to identify where these commands should be placed. This is tackled by executing the Movement Position Analysis (MPA) code. The MPA code examines the TIRI pointing sequence and begins by plotting the velocity against time and the acceleration against time to identify points of discontinuity as potential commanding locations.

The rotational velocity across the first data-cube has been plotted in Figure 5-37. We can see that the velocity across the timescale considered for each image to be taken is essentially constant (the movement to track the comet is a curve, but the plot shows that across the short timeframes we are considering the curve is imperceptible, true residuals are retrieved and examined later). This supports the hypothesis that we may be able to recreate profile from commands defining velocity. Each of these velocities may be set as a single command, with that velocity being overwritten only when the next command is received. The slight deviations in velocity seen while tracking the comet in filters represent the fact that the velocity is not actually a straight line. The errors against this true line will be demonstrated when we plot residuals between our command simulated sequence and the target sequence.




Figure 5-37: The target position profile velocity plotted across the first TIRI data-cube. Regions of constant positive velocity show the positions where the TIRI motor is keeping the nucleus in the centre of a filter and therefore this velocity is essentially 'tracking' the comet with the pointing mirror. The constant negative velocity shows the motor moving at maximal velocity to move to the next filter. At the end the velocity moves to maximal positive velocity to move the nucleus back to the first filter to commence the second data-cube.

The next task is to consider the acceleration profile, as we want to select peaks of acceleration (changes between the near flat velocities observed in Figure 5-37) as the locations of our commands. The full encounter acceleration has been plotted in Figure 5-38, showing that the critical acceleration trigger in the positional dataset is somewhere about 2500 degrees per second squared. A function was defined to sweep through the acceleration data, recording the index for each incident of extreme acceleration values as a 'significant time' to record the positions for commanding. One shortfall of this method of identifying commands was the relatively small accelerations from a zero velocity in data-cube 3 (a product of scanning in the opposite direction), so an additional filter was used to identify these cases.




Figure 5-38: The acceleration profile across the target sequence, plotted for the entire sequence. Note how the acceleration required reduces across the encounter since the angular velocity to track the nucleus reduces as we receded from the comet. The exception to this is in data-cube 3 (8 to 13 seconds) which has lower acceleration since the motor is scanning the nucleus across the TIRI filters with the comet movement, rather than contra the flyby movement. See the operations plan for filter 3 for more details of how this was defined.

The final product of this analysis is to then use the index values of the ‘significant times’ to identify the positions of commands across the encounter. The result of this is displayed in Figure 5-39. The bottom plot shows that the velocity commands do correspond to the times at which the motor substantially changes velocity. These will be the locations for commands that are used in the next section.




Figure 5-39: The TIRI motor position plotted against time for the entire encounter (top) and only across the first data-cube (bottom). Overlaid in red on these plots are the significant times, which are the identified commanding positions, as described in this section. Each red line is the time at which the command will be executed.

5.4.2.2 Generating Commands

Once the command locations have been identified, the full command dataset needs to be created. This is a relatively simple process based on the outputs shown in Figure

5-39. The velocity for each command is defined by calculating the average velocity between the position when the command is issued and the position when the next command is issued. This data is then saved as a text file of command data, containing the timestamp and velocity defined for each command. An excerpt of the generated command data is shown in Figure 5-40. To make a quick comparison about the data compression of transferring from the high-resolution position time-series data to the command data shown in Figure 5-40 we can look at the relative size of the files. The full position data is a 2.5 MB file, whereas the full command file is 2 KB, so this file takes less than 0.1% of the data volume.

	#	Start Time	Velocity
1			
2	-1.250000	3.925330	
3	-1.194000	-6.200000	
4	-1.134000	-2.291198	
5	-1.133000	3.951622	
6	-1.088000	-6.195053	
7	-1.028000	3.951468	
8	-0.988000	-6.200000	

Figure 5-40: The first 8 lines of the saved Command_Data.txt file generated by the MPA and command generation code.

This command table has been used with the TIRI breadboard motor assembly, used to drive the stepper motor under vacuum through simulated encounter regimens. This enables us to directly measure and validate velocity, acceleration and pointing accuracy on a representative assembly, both in the case of capability measurements and testing encounter-like movement patterns.

5.4.2.3 Validating Commanding Method

Next, we need to understand how much precision has been lost in transforming from the positional data to the command data. To do this we must simulate the TIRI motor movement pattern from the command dataset, and this can then be compared against the original dataset to ascertain the residuals.

The simulated profile is generated by starting with the initial position and time (an assumption is made that the instrument is able to start the encounter sequence from the correct position and at the correct time) and defining the distance moved to the next position by multiplying the velocity set for that command by the time difference to the next command. This is repeated for each command to the end of the command dataset. The generated profile is plotted along with the original position data in Figure 5-41. Small variations in the angular difference would not be visible unless zoomed in considerably, which would reduce our ability to consider the profile across the whole encounter.




Figure 5-41: The position profile simulated from the command data (orange) plotted against the position data as taken from the operations plan (blue). The overall shape is clearly similar, but small deviations need to be examined.

Therefore, I have also interpolated the simulated positions onto the target position timeframe and plotted the residuals between them in Figure 5-42.




Figure 5-42: The residual between the simulated and operations defined position datasets across the entire encounter.

In order to evaluate the meaning of the residuals displayed in Figure 5-42 we need to relate the residuals through to pointing uncertainty. We will do this using the following TIRI instrument parameters:

- TIRI instantaneous Field of View is 0.26 mrad (0.0149 degrees)
- 29,000 step stepper motor gives 0.0124 degrees per step

This allows us to make the comparison that the maximal residual error of 0.0024 degrees corresponds to a pointing error of 16% of a pixel at the detector. Crucially considering this is less than 20% of the angular difference between each step on the stepper motor (which is our fundamental minimal pointing resolution) the resolution of individual steps is going to be a dominant factor when attempting to hold the comet nucleus stationary in the instrument field of view. Therefore, the positional resolution lost

in transforming our position data to a velocity command data is not a dominant source of pointing uncertainty.

5.4.2.4 Introducing Time Uncertainty

Using the tools above, and following the pipeline laid out in Figure 5-36 we can now investigate the effect of timing uncertainty in command execution. The maximal positional error introduced in a single command can be identified directly from the motor parameters, and the relationship between time and position errors is best understood with the aid of the diagram in Figure 5-43. The maximal positional error would come from a timing error in a scenario where we are moving between the maximal velocity in one direction and the maximal velocity in the opposite direction. The position error follows the equation:

$$error_{position} = \frac{1}{2}(v_1 - v_2) * error_{time}^2 \quad (5.9)$$

Where v_1 and v_2 are the velocity before and after the command is executed, so a smaller change in velocity will produce less position error for the same time uncertainty. If moving from the maximum speed of 6 degrees per second to the maximum speed in the opposite direction, with a timing error of 10 milliseconds, then that one command would lead to an error of 0.006 degrees. This is approximately half a pixel, but since these




Figure 5-43: A diagram demonstrating the relationship between an error in the time of a command and the error in position. The blue line represents the motor position, with the dotted line denoting the intended path. However, the command has been called early, so the motor has changed to the next velocity at an earlier time, creating a positional offset (the vertical difference between lines).

position errors cumulatively combine over the encounter; across 120 commands the error could propagate to a dominant size. Individual command shifts are important to understand for co-aligning images within data-cubes, whereas the cumulative error needs to be understood to verify that the image retrieved in the final data-cube will fall within the filter windows. We need to investigate the errors in the context of a full set of encounter movements.




Figure 5-44: A residual plot generated for the command data executed with random time errors distributed about a time uncertainty of 2 milliseconds. Each time a command is generated a pointing error is introduced relative to the position before, and this error is then propagated forward to impact every point after this time. This leaves the possibility that across the entire encounter the position may drift very substantially if repeated errors occur in the same direction.

This is done by generating an array of random timing error values, normally distributed about 0 with a standard deviation equal to half of time uncertainty. The time uncertainty is the constrained temporal accuracy of any given command. The array is the length of the command data. The time errors are added to the time column in the command data. This then gives a new command dataset with each command execution randomly perturbed around its intended position. The outcome of this for a single case has been plotted in Figure 5-44. A single case of random error is not representative, so in the

next section we move to evaluate the propagation of errors across many repeated simulations.

5.4.2.5 Monte Carlo Simulation of Command Errors

The model was adjusted to loop over many cases, with each experiencing an independent random distribution of command timing errors. The errors were again normally distributed about an average of 0, with a set standard deviation, which is half of what here is referred to as the ‘time uncertainty’. This reflects the fact that the quoted uncertainty will typically refer to the two standard deviation limit, the 95th percentile of the noise. The distribution of outputs can then be analysed statistically. At each time frame, the average position error, as well as the standard deviation of the distribution is calculated. The results are then plotted as errors, with the 1st, 2nd and 3rd standard deviation boundaries displayed as in Figure 5-45.




Figure 5-45: The mean pointing error with standard deviation envelopes. In this case the monte-carlo simulation was run 5,000 times, with a time uncertainty of 2 milliseconds. The grey envelopes show the 1st, 2nd and 3rd standard deviations from the mean.

The results in Figure 5-45 show how the error spreads across the encounter. At the first and second data-cubes (<8 seconds) the potential error increases steeply as these are high change in velocity movements. Data-cubes four and five (at 55 seconds and 105 seconds) show slightly less steep gradients for two reasons derived from the tracking velocity. Firstly, since the tracking velocity is lower, the amount of time taken to produce all the measurements for that data-cube is increased as it takes longer to move between filters. Secondly, as described in Section 5.4.2.1, the position error is based on the difference in velocity at the commands, so a reduced tracking velocity means a smaller error increase. The third data-cube (8-15 seconds) is an illuminating example because it is the only case where the filter scan is run in the same direction as the comet movement in the detector FoV. This means that while the individual images are more spaced out in time, the velocity changes are very small compared to the other case, and thus the pointing error introduced in the change of direction is much smaller. There is another important source of pointing accuracy which is not accounted for in the above model, which is 'jitter' as the




Figure 5-46: The mean pointing error with standard deviation envelopes. In this case the monte-carlo simulation was run 5,000 times, with a time uncertainty of 0.5 milliseconds. The grey envelopes show the 1st, 2nd and 3rd standard deviations from the mean.

motor settles between different velocities. While this also needs to be measured in the motor assembly engineering model, this is expected to scale with the magnitude of the velocity changes and so should also be considerably lower in the low acceleration schedule of data-cube 3.

The three standard deviation position error (99.7th percentile limit) at the end of the observation in Figure 5-45 is nearly 0.2 degrees, which corresponds to 13 pixels which is a critical failure since at a filter width of 30 pixels, 13 pixels is enough error to put the nucleus outside of the filter window, meaning we would not view the comet in the final data-cube.

In light of this, the simulation was rerun at a time uncertainty of 0.5 milliseconds. This has generated the pointing error distribution plotted in Figure 5-46. This shows the 3rd standard deviation limit reduced to 0.05 degrees. This error corresponds to 3.3 pixels and thus would enable measurement of the nucleus throughout the encounter.

5.4.2.6 Encoder Position Resetting

Since the motor assembly carries a Gurney encoder to measure the motor position, this can be used to reset position and account for error positions which are introduced. The simplest implementation of this would be to realign the position in between data-cubes where the FPGA resources are in lower demand. This scenario has been simulated by resetting the pointing error to 0 between observing each data-cube. The result of these simulations has been plotted in Figure 5-47. The discussion from Section 5.4.2.4 regarding the error introduced in each data-cube based on the motor velocities in that data-cube is well visualised here. The third data-cube shows the lowest introduced pointing error due to the smaller velocity changes. The largest accumulated errors here in the 3rd standard deviation is below 0.03 degrees, which corresponds to 2 pixels of accumulated pointing

error across a single data-cube. This would dominate the pointing uncertainty across the data-cube and constrain the spatial resolution.




Figure 5-47: The position error across the entire encounter in the case where the position is reset to its intended position, and error is therefore set to 0, in between data-cubes. This simulation was run 5,000 times and the mean error and 1st, 2nd and 3rd standard deviation envelopes have again been plotted.

If the required pointing accuracy cannot be achieved due to the uncertainty introduced in command timing, and in jitter or dropped steps at the high acceleration regime of the first two data-cubes, this uncertainty can be greatly reduced by changing the direction of filter scanning.

5.4.3 Commanding Conclusions

The results of the TIRI pointing commanding simulations produce a requirement of 0.5 millisecond accuracy on the commanding execution time. This analysis has also generated tools to create motor command tables which are currently being used within the full set of MIRMIS motor assembly tests. The algorithms used to generate these tables stand to be reused to allow rapid testing of updated mission operation plans, and by enabling pointing plans to be tested within the motor they will contribute to analysing the

performance of the motor in terms of velocity, acceleration, pointing accuracy, jitter and reliability. The knowledge gained regarding motor performance will feed back into the operational design modelling, ensuring the instrument is run to produce optimal scientific outcomes. The command tables generated for the two operation plans shown in this chapter (Figure 5-20 and Figure 5-22) have both been used directly to command the motor assemblies during testing, and future work combining these command tables with high-resolution encoder outputs will constrain the true pointing accuracy during the encounter.

If the required pointing accuracy cannot be achieved due to the uncertainty introduced in command timing, and in jitter or dropped steps at the high acceleration regime of the first two data-cubes, this uncertainty can be greatly reduced by changing the direction of filter scanning.

5.5 Housekeeping Board Development

In MIRMIS, the internal temperatures are recorded exactly the same way as in LTM, with four pairs of ESCC9408/064 multiplexers where each channel in a multiplexer pair measures the voltage across a PRT1000. Four of these circuits are included in the final design. The multiplexer chips are controlled by a DAC via address lines.

5.5.1 Relevant Requirements

The TIRI module must monitor its internal temperature components in accordance with the MIRMIS TIRI requirements, as outlined in Table 4-4. This is achieved using the Housekeeping Board, which is designed to measure and monitor data from 28 temperature sensors distributed throughout the TIRI module. This satisfies Requirement MIRMIS-TIRI-

PER-010, which mandates redundant temperature data recording to prevent critical point failure.

The housekeeping board's role is also tied to Requirement MIRMIS-TIRI-BB-0500, which stipulates the accuracy of the temperature measurements for the blackbody target used in calibration. As in LTM, meeting this requirement is essential for fulfilling MIRMIS-TIRI-PER-070, which governs the precision needed in calibration procedures.

5.5.2 Test Procedure

The housekeeping board used for testing was an engineering model built using flight-like parts. Each resistor chain was modified by soldering resistors directly across the multiplexer pins, rather than connecting to an external PRT. These resistor values were chosen to represent the span of values expected from the instrument PRTs during flight.

The values are provided in Table 5-5.

Channel Number	Resistance /Ohms
Channel 0	845
Channel 1	931
Channel 2	1000
Channel 3	1020
Channel 4	1150
Channel 5	1210
Channel 6	1300
Channel 7	953 (Reference resistor)

Table 5-5: TIRI housekeeping board PRT substitution resistor values.

The board was then adapted with electrical leads available for the ground, power inputs, address lines and the multiplexer outputs for each of the four internal multiplexer-pair circuits. The power lines were set using external bench power supplies, and the address lines were set via Arduino, controlled by MATLAB scripts. The output of each Op-

Amp was measured either by the Arduino (in preliminary testing) or via a high-resolution NI DAQ. The testing equipment and software is described below:

- **Housekeeping board control signals:** Managed by an Arduino Uno controlled through custom written MATLAB scripts.
- **Data acquisition:** Handled by both the Arduino Uno and National Instruments DAQ (NI USB-6218), with custom MATLAB scripts.
- **Data storage:** Saved online as CSV files to the University's OneDrive.
- **Data analysis:** Performed in Jupyter Notebooks using Python, with relevant libraries (pandas, matplotlib, seaborn) for statistical manipulation and visualization.

This was a two-stage test, consisting of:

1. **Operational checks:** The first stage aimed to verify the correct function of each component, ensuring and probing circuitry performance under various conditions.
2. **Detailed measurements:** The second stage focused on more precise voltage measurements across the board, including high-resolution time-series data collection from different areas of the circuit.

5.5.2.1 Preliminary Investigation

The preliminary investigation aimed to collect baseline data, check the integrity of the board, and ensure correct operating conditions. It also served as an opportunity to test the MATLAB diagnostic software that would manage the system.

The base experimental setup consisted of recording the HSK output of each op-amp at a rate of 30Hz while the address lines rotated through each of the multiplexer truth table values to sweep across all the resistors. The controlling MATLAB script would

identify which channel was activated by the Arduino while the measurements were being taken to save the data in a labelled format.

The initial data analysis used the following pipeline to derive the measured resistance from the voltage readings within each multiplexer circuit:

1. The averages and standard deviations for each channel's voltage measurement is calculated.
2. The current within the resistor chain is calculated by comparing the average voltage in the reference resistor channel to the known resistance (the factor of 10.3 relates the voltage at the amplifier output, which is measured, to the voltage across the resistor, which is before the amplifier):

$$Current_{chain} = \frac{Voltage_{reference}}{10.3 * Resistance_{reference}} \quad (5.10)$$

3. Resistance across each resistor is then calculated from the common current along the chain using Ohm's law:

$$Resistance_i = \frac{Voltage_i}{10.3 * Current_{chain}} \quad (5.11)$$

This allowed for immediate comparison between the measured and expected resistances, providing a check on correct functionality. The measured voltage noise was also analysed to determine the noise uncertainty introduced by the electronics.

5.5.2.2 *Reduced Op-Amp Supply Voltage Test*

One goal of the testing was to measure the noise in the measurements in order to constrain the uncertainty of temperature measurements through the board. By measuring the noise, and experimenting with the board, I wanted to assess ways to reduce noise.

A hypothesis was that the amplification circuit contributed to noise. The board's Op-Amps are powered by a 12 V supply, but it was proposed that a 5 V supply might

reduce noise while it should maintain functionality. The board was tested at both 12 V and 5 V, with additional measurements taken at key points in the circuit using a multimeter to investigate the performance.

5.5.2.3 Noise Investigation

The next stage of testing aimed to identify the sources of noise across the circuit by measuring voltages at different points: across the resistors, the multiplexer outputs, and the amplifier outputs. This allowed for differentiation between noise introduced by the current, multiplexers, or amplifier.

This was enabled by taking high-frequency measurements at different points on the board using additional leads that were connected throughout each circuit. Measurements are taken across the resistor directly, across the multiplexer outputs and at the amplifier circuit output. Each is constrained of, noise in the current, noise through the multiplexers and noise at the output. These measurements were taken across two resistors, the highest and lowest value, and repeated for each of the four internal circuits on the housekeeping board.

Each of these measurements were taken with the high-frequency NI DAQ, taking 2 seconds of data in each channel at 1000 Hz. Since each measurement had to be taken individually, the datasets were coincided using the leading edge after the reference resistor measurement. The data was then analysed by calculating the averages and standard deviations for each of these different measured voltages. Comparison between these statistics describes where the noise has originated.

5.5.2.4 Voltage Supply Investigation

Using the enhanced measurement setup, the impact of the voltage supply (5 V vs. 12 V) on the board's performance was investigated with high-resolution and high-

frequency measurements. These tests further refined the assessment of whether a lower voltage supply could reduce noise levels.

5.5.2.5 Switching Velocity

The high-frequency DAQ also enabled direct measurement of the settling time for the voltage to stabilise when the multiplexers switched channels. This is critical for determining how rapidly the flight software can switch between channels, ideally allowing all temperature measurements to be temporally coincident.

5.5.3 Testing Results

5.5.3.1 Preliminary Data Analysis

Initial testing and data analysis were conducted as outlined in Section 5.5.2. The voltages recorded via the Arduino (shown in Figure 5-48) demonstrate clear and correct switching between channels, consistent across all circuits. Calculated resistances through the calibration process were accurate to within 0.1%, with the current constant between circuits to within 1%, matching the expected current flow.

However, voltage noise was measured with a standard deviation of up to 28 mV, corresponding to a resistance uncertainty of around 8 Ω , which in turn relates to a temperature uncertainty of nearly 2°C. This is higher than the required accuracy of 0.25°C (as per MIRMIS-TIRI-BB-0500) and exceeds the tolerance for noise in temperature measurements. This noise level provides the motivation for the further testing described here. The noise level, though too high for individual measurements, can be reduced using super sampling and averaging in the flight software, however this comes at the cost of on-board computation and longer measurement windows.

Average Voltage Table (in Volts):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Channel_0	2.5752	2.5581	2.5733	2.5665
Channel_1	2.8391	2.8191	2.8372	2.8291
Channel_2	3.0493	3.0301	3.0468	3.0385
Channel_3	3.11	3.0908	3.1075	3.0992
Channel_4	3.5085	3.4891	3.5061	3.4965
Channel_5	3.6929	3.6684	3.6908	3.6747
Channel_6	3.9688	3.9465	3.9707	3.956
Channel_7	2.9054	2.8871	2.9031	2.8958

Standard Deviation Table (in Volts):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Channel_0	0.006075	0.015597	0.015617	0.00837
Channel_1	0.0064645	0.01602	0.016857	0.0094235
Channel_2	0.0082069	0.018411	0.018647	0.012156
Channel_3	0.010278	0.019486	0.019642	0.013278
Channel_4	0.012828	0.0219	0.021867	0.014293
Channel_5	0.016648	0.023684	0.023189	0.019583
Channel_6	0.022101	0.028829	0.028736	0.024511
Channel_7	0.0055057	0.015081	0.015072	0.0071632

Current Table Across Circuits (in Amperes):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Current (A)	0.00029599	0.00029412	0.00029576	0.00029501

Resistance Table Across Channels and Circuits (in Ohms):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Channel_0	844.68	844.41	844.72	844.62
Channel_1	931.25	930.56	931.37	931.05
Channel_2	1000.2	1000.2	1000.2	999.95
Channel_3	1020.1	1020.3	1020.1	1019.9
Channel_4	1150.8	1151.7	1150.9	1150.7
Channel_5	1211.3	1210.9	1211.6	1209.3
Channel_6	1301.8	1302.7	1303.4	1301.9
Channel_7	953	953	953	953

Figure 5-48: Basic housekeeping metrics for each internal circuit as recorded and calculated by my MATLAB script. The average voltage and voltage standard deviation in each channel, resistor chain current and each recovered resistance measurement is displayed.

5.5.3.2 5 V Supply Test

One hypothesis for reducing noise was that lowering the operational amplifier's supply voltage might lead to smoother performance. The amplifier was originally designed to operate at 5 V, but the design had been changed to use 12 V as a product of my testing on the LTM board (Section 3.3.2). Given the availability of 5 V lines, and the fact this was the original intention, using a 5 V supply was tested. Results are shown in Figure 5-49.

Observing the measured and calibrated resistances immediately tells us that this has prevented the board from operating correctly. Switching the supply to 5 V resulted in a malfunction, invalidating the calibrated resistances as the final channel, which measures the reference resistor, produced incorrect results. Since the reference resistor voltage is used to calibrate all other channels, this prevents correct measurements in any channel.

Average Voltage Table (in Volts):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Channel_0	2.576	2.5592	2.5741	2.5671
Channel_1	2.8389	2.8216	2.8393	2.8287
Channel_2	3.0495	3.0309	3.0491	3.0381
Channel_3	3.1108	3.0927	3.1095	3.0988
Channel_4	3.508	3.4891	3.5095	3.4942
Channel_5	3.6915	3.6703	3.6912	3.6749
Channel_6	3.9702	3.9473	3.9692	3.9562
Channel_7	3.1775	3.1154	3.1805	3.1356

Standard Deviation Table (in Volts):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Channel_0	0.003721	0.0094604	0.010882	0.005262
Channel_1	0.0029467	0.0057913	0.0057894	0.0031566
Channel_2	0.0026884	0.0060376	0.0062054	0.0033365
Channel_3	0.0030405	0.0088957	0.0093749	0.0036403
Channel_4	0.0035064	0.011953	0.011915	0.0047864
Channel_5	0.0044263	0.015846	0.016176	0.0061512
Channel_6	0.011803	0.021591	0.021749	0.014644
Channel_7	0.0017164	0.0024467	0.0021936	0.002474

Current Table Across Circuits (in Amperes):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Current (A)	0.00032371	0.00031739	0.00032402	0.00031944

Resistance Table Across Channels and Circuits (in Ohms):

	<u>Circuit_1</u>	<u>Circuit_2</u>	<u>Circuit_3</u>	<u>Circuit_4</u>
Channel_0	772.58	782.86	771.3	780.22
Channel_1	851.43	863.13	850.74	859.74
Channel_2	914.59	927.15	913.63	923.39
Channel_3	932.99	946.04	931.72	941.81
Channel_4	1052.1	1067.3	1051.6	1062
Channel_5	1107.1	1122.7	1106	1116.9
Channel_6	1190.7	1207.5	1189.3	1202.4
Channel_7	953	953	953	953

Figure 5-49: The housekeeping metrics for the 5 V op-amp supply test. The measured resistances (at the bottom) diverge from the resistor values (Table 5-5). On examination the current is wrong (μA values should be temperature in K). When compared with Figure 5-48 it can be seen that the values for Channels 6 and 7 are elevated compared to what they should be. Channel 7 then invalidates all resistance conversions as it is the reference resistor.

The original hypothesis was however somewhat vindicated, with the voltage noise in the measurements reduced by between 30 and 50%. Having observed a similar

behaviour when housekeeping board setups were tested prior to the finalisation of the LTM design, I hypothesised that this inconsistency was a result of non-constant amplification across common-mode voltages. Figure 5-50 shows the expected common-mode voltages on either side of each resistor in the housekeeping resistor chain. It is the voltage across these resistors which is fed into the amplifier circuit.




Figure 5-50: The expected common-mode voltages between each resistor, defined relative to ground. The inputs to the amplifier circuit will be the voltages on either side of any resistor.

This was investigated by setting the housekeeping board to measure different resistors through the amplifier and directly measuring each pin in the amplifier circuit to identify whether there was any failure point. Each of the locations to probe is illustrated in

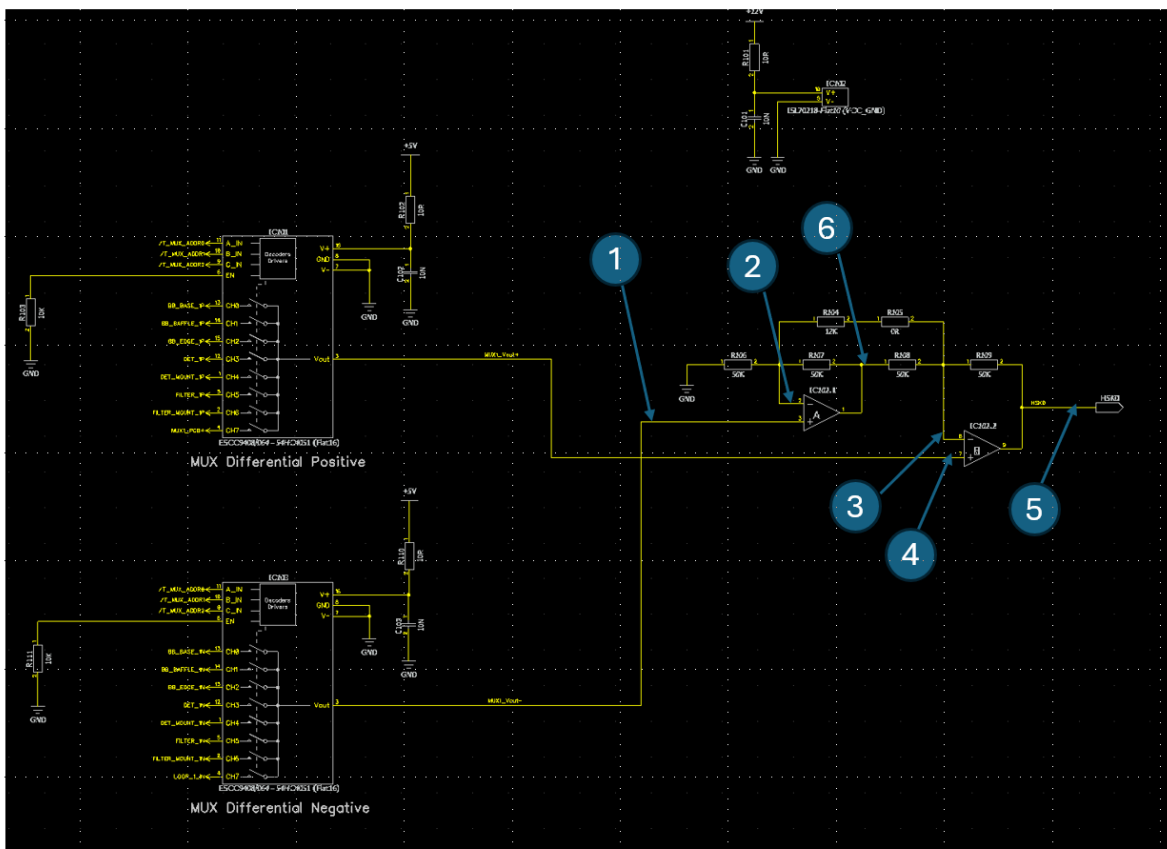


Figure 5-51: Test points on the housekeeping board for additional probing while different channels are selected. This circuit diagram is the same as the housekeeping board layout described more clearly in Figure 3-26.

Figure 5-51, and the results of these measurements for three resistors are listed in Figure 5-52.

The results show that at the end of the resistor chain, the amplifier circuit became saturated as internal voltages needed to surpass 5 V, which was not possible with only a 5 V supply. The conclusion was to continue with the 12 V power supply as originally designed.

Resistor Chain Position	Multiplexer Channel	Test Point Voltages					
		1 (pin 3)	2 (pin 2)	3 (pin 8)	4 (pin 7)	5 (pin 9)	6 (pin 1)
Lowest	0	2.434	2.436	2.688	2.687	2.086	4.328
Ref	7	2.695	2.698	2.981	2.977	2.428	4.705
Highest	6	0.490	0.490	0.867	0.876	3.493	-0.147

Figure 5-52: Housekeeping probe testing results. Each of the locations described the in Figure 5-51 for the reference resistor, as well as the lowest and highest resistors in the resistor chain. The failure point can be observed in the amplifier circuit at the highest resistor.

5.5.3.3 High-Frequency Measurements

The next phase utilised the high-frequency data acquisition unit to investigate performance in greater detail (described in Section 5.5.3.3). Figure 5-53 shows the four measured voltages. The signal-to-noise ratio improved significantly after amplification (transition from the green to the purple line). However, despite using a higher resolution DAQ, the remaining noise is still substantial and requires further investigation.




Figure 5-53: Plots of the four recorded voltage lines measured through the high-frequency NI DAQ. The three lines at the bottom of the graph show the voltage across the 850 Ω and 1300 Ω resistors (constant blue and red lines) with the input line of the amplifier circuit in green switching between the voltage of each resistor as the multiplexers switched through their channels.

5.5.3.4 Noise Investigation

The graph in Figure 5-54 shows the non-amplified voltage through the multiplexers, as well as two directly measured voltages. The size of the noise is comparable to the separation between different voltages, complicating the accurate measurement of resistance.

Python scripts were setup to evaluate the noise in three positions. The noise was evaluated in the direct voltage measurement across the resistor, in the multiplexer output/amplifier input and at the amplifier output. Then this analysis was run for both the case where the housekeeping board was measuring the 845 Ω and the 1300 Ω resistors. The increase in noise as the voltages is passed through the system at each step (selection via multiplexer pair and amplification) is examined and broken down to be viewed. The data for the first circuit on the housekeeping board is given in Figure 5-55. The testing was




Figure 5-54: The direct resistor voltages and the amplifier input voltage, plotted to show the size of the noise relative to the different resistor values. The uncertainty in measurement is clearly comparable to the different between the voltage across the 1210 Ω and 1300 Ω resistors.

repeated for each circuit, and the results were consistent, for space efficiency only the firsts set of results are provided here. Key findings include:

1. Noise at the resistor level increased by $\sim 70\%$ for the 1300 Ω resistor compared to the 845 Ω resistor. Proportionality between resistor size and noise suggests that the noise is primarily due to fluctuations in the current supply, translated to voltage via Ohm's law.
2. Both resistors showed a small increase in noise ($>20\%$) after passing through the multiplexer pairs.
3. Amplification magnified noise significantly (by 65 mV), increasing it by 130% and 80%, respectively. However, since amplification also increases signal strength by a factor of 10.3, the overall signal-to-noise ratio still improved.

The final output noise (14 mV) corresponds to a resistance uncertainty of 4 Ω , translating to a temperature uncertainty of about 1°C. This level of uncertainty is an order of magnitude too high to meet the requirement, suggesting that oversampling and averaging will be needed to reduce noise further. A large portion of the noise originates from the current supply, limiting the improvement that can be gained from optimizing other parts of the housekeeping board.

```
Noise in 845R Voltage while measuring 845R resistor: 0.004002523799209288 Volts
Noise in Input Voltage while measuring 845R resistor: 0.004794931728764913 Volts
Noise in Output Voltage while measuring 845R resistor: 0.011022123207440566 Volts
Percentage increase in noise from 845R to Input: 19.80%
Percentage increase in noise from Input to Output: 129.87%
Absolute noise at resistor: 0.0040 Volts
Additional noise from 845R to Input: 0.0008 Volts
Additional noise from Input to Output: 0.0062 Volts

Noise in 1300R Voltage while measuring 1300R resistor: 0.007104097270925083 Volts
Noise in Input Voltage while measuring 1300R resistor: 0.008243988526494625 Volts
Noise in Output Voltage while measuring 1300R resistor: 0.014752580292772688 Volts
Percentage increase in noise from 1300R to Input: 16.05%
Percentage increase in noise from Input to Output: 78.95%
Absolute noise at resistor: 0.0071 Volts
Additional noise from 1300R to Input: 0.0011 Volts
Additional noise from Input to Output: 0.0065 Volts
```

Figure 5-55: Noise analysis of the first circuit based on the high frequency measurements. The same analysis has been conducted for the 845 Ω and the 1300 Ω resistor.

While these results provide valuable insights, there are some limitations to the testing:

- The experiment measures noise within the test setup, which may not accurately represent conditions aboard the spacecraft.
- The power supplies used in this test are bench supplies, which may differ from spacecraft power lines in terms of noise characteristics.
- The spacecraft's onboard ADC will have a lower resolution than the high-resolution DAQ used in this test, which will introduce additional noise.

To address these limitations, I recommend verification testing through the TIRI electrical functional model (EFM) once available. The EFM is similar to the engineering

board tested here and similarly incorporates dummy resistors in place of PRTs. The internal board voltages will not be able to be probed, however the final measurements will be more representative being measured through the full instrument electronics.

5.5.3.5 Multiplexer Switching Velocity

The final component of testing was to measure the channel switching velocity on the multiplexers. This is done by observing the voltage of the output voltage, as measured through the NI DAQ, in order to view the moment at which the address lines trigger the multiplexers to change measured channels. These magnified views can be observed in Figure 5-56 and in Figure 5-57.

Examination of Figure 5-57 shows there is no visible settling or transition period visible on a voltage measurement run at 1000 Hz, therefore any effects take place on a sub-millisecond timescale. This means the instrument may take measurements within a




Figure 5-57: Upon further magnification, the switching velocity becomes apparent. Since we cannot see any settling process in a 1000 Hz measurement, we can tell that the multiplexer switches on a sub-millisecond timescale.

switch between channels.

millisecond of changing channel, which guides how rapidly the instrument can rotate through the address lines and take a full set of measurements.

5.5.4 Housekeeping Board Conclusions

The testing and analysis of the TIRI housekeeping board have provided valuable insights into its performance and highlighted areas for further testing once the MIRMIS EFM is available for testing.

Initial testing has confirmed the basic functionality of the circuitry with all channels correctly switching and resistances accurately measured to within 0.1%. The necessity of the 12 V power supply for the amplifier circuits has been validated as necessary for reliable operation.

The noise within individual PRT measurements, found to correspond to 1 K temperature errors, has been investigated and decomposed. Amplification has been found to increase noise, but significantly improve SNR, and the source of most noise was traced to the current supply. Further investigations with the EFM, using the spacecraft simulator to provide the instrument power supply will evaluate this performance in the most representative environment. The multiplexer switching has been shown to be fast (quicker than 1 millisecond) guiding the rate at which the channels can be changed. Super-sampling will be tested on the EFM as a way to reduce noise and enhance reliability, and this may be the limiting factor for the coincidence of temperature measurements.

6 TIRI Observation Simulation

This chapter presents a comprehensive simulation of the TIRI operation, building upon the parameters explored in the preceding chapter regarding TIRI's operational design. As outlined in the discussion about defining TIRI's operations, the instrument's performance depends on a multitude of interrelated and nonlinear factors. Therefore, its full capabilities can only be accurately assessed by simulating these factors together in a flight-like scenario. The aim of this simulation is to emulate TIRI's observations during the comet flyby, using a limited set of inputs to generate representative science products that can be evaluated against the mission's scientific objectives.

Realistic science products will provide a vital link between the science team and the instrument team. Enabling specialists in a range of cometary science fields to directly probe simulated science results will allow these scientists to advise the instrument team through instrument and mission design, as well as target selection. This collaborative process is expected to accelerate the development of an optimal flyby strategy. In addition, the simulation tool will allow rapid assessment of potential target comets by evaluating the instrument's performance under different flyby conditions and cometary backgrounds. This will allow prompt analysis and feedback on the relative strengths in science returns of different potential cometary targets and therefore play a role in supporting the central ESA mission team at the decision point for target selection. Furthermore, the TIRI simulation model provides a baseline and includes the majority of the framework necessary to integrate NIR and MIR operational models, to form complete MIRMIS science products.

Beyond supporting flyby planning and target selection, the TIRI simulation will produce example calibrated radiance spectra for each channel, incorporating noise and uncertainty levels. These outputs will be used to verify the instrument's performance

relative to requirements on signal retrieval, temperature accuracy, spatial resolution, and spectral differentiation.

6.1 Simulation Inputs

To simulate the TIRI observation, the parameters needed to define the input space are divided into three groups. These parameters flow through the full end-to-end simulation, as outlined in the simplified data-pipeline shown in Figure 6-1.

- **Target parameters:** These pertain to the characteristics of the comet nucleus being targeted. They include the nucleus shape model, size, temperature distribution, orientation, and surface properties.
- **Flyby parameters:** These describe the mission's general conditions, independent of the instrument, including the flyby velocity, angle, and distance at closest approach.
- **Operational parameters:** These are specific to the MIRMIS instrument and its operations. They include the pointing mirror velocity profile, filter positions, individual filter integration times, and data-cube locations. These inputs are used to define the operations plan (see Section 5.1), which is central to the simulation.

The target and flyby parameters feed into the operations modelling, shaping the operations plan, which in turn drives the simulation. For instance, a faster encounter will necessitate more rapid motor movements, leading to increased pointing uncertainty and spatial blur compared to a slower flyby. Additionally, a warmer target would allow shorter integration times, whereas a cooler target would require longer exposures. For the simulations presented in this chapter, a consistent operations plan has been used to demonstrate how varying inputs affect the outputs. This is the operation plan defined in

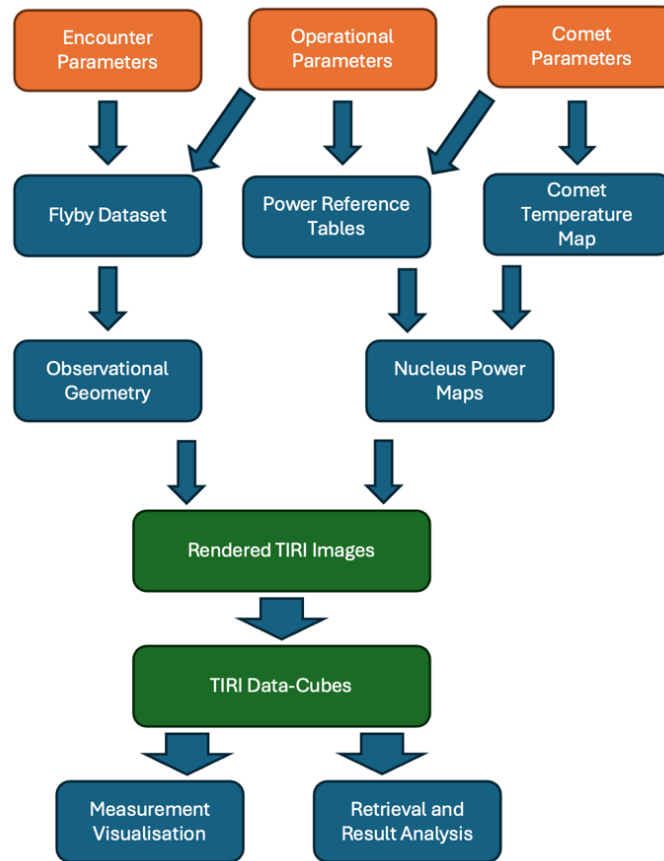


Figure 6-1: Flowchart describing the operation and data pipeline covered in this end-to-end simulation in a simplified manner.

Section 5.1.5, and used for the rest of the analysis in Chapter 5. For the actual mission scenario, the operations will be fine-tuned based on the expected characteristics of the target comet to optimise the scientific retrieval.

6.2 Simulation Structure Overview

The simulation is designed with modular sections, each providing intermediate outputs that can function independently. The key components are as follows.

Instrument Setup: The TIRI instrument's transmission properties are defined within each filter using transmission and reflectivity data. A temperature-to-detector-power reference table is generated, which maps the power incident at the detector for

each filter. If multiple emissive surfaces are present, separate reference tables are created for each emissivity function.

Target Parameter Processing: The nucleus data is imported and processed in Python. The comet's shape model is scaled to the appropriate size and oriented based on the direction of Solar illumination. Temperature and emissivity maps are then used to compute a detector-power map for each filter.

Flyby Parameter Processing: The flyby geometry is calculated using data from the operations plan. A time-series dataset with 1 ms temporal resolution is generated for the entire observation. This dataset includes the spacecraft's position in the sun-comet reference frame, as well as the orientation, pointing direction, and focal point of the TIRI field of view (FoV) at every time step. These parameters allow accurate object rendering in the simulation of the observation.

Instrument Operation Simulation: This module simulates the detector's view by rendering throughout the observation. Based on the data-cube configuration from the operations plan (including positions, integration times, and filter sequences), simulated detector images are produced. For each data-cube, an image is generated for every filter used.

Data-Cube Generation: The simulation uses a co-location algorithm to align the nucleus across all filter images within each data-cube. A four-dimensional stacked NumPy array data-cube is then constructed, in which each pixel contains a full set of values across all filters. This includes information such as colour, detector power, power noise, energy, energy noise, calibrated radiance, radiance noise, brightness temperature, and temperature uncertainty. This fully stacked array represents the final data-cube, serving as the science output of the simulation.

6.3 The Full Flyby Simulation

This section expands upon the previously introduced modules, detailing how each operates within the overall simulation framework.

6.3.1 Instrument Setup

The goal of this module is to derive the power that is received at the detector through each channel, depending on the temperature of the surface. This is done by defining a look-up table relating set temperatures to the detector radiance through the instrument transmission model. Specific powers are then interpolated from this.




Figure 6-2: The simulated instrument transmission spectra for each TIRI filter, as used in the end-to-end simulation.

The TIRI instrument transmission spectra used here are a combination of the TIRI filters which were the output of Section 5.3, with mirror and detector measurements following the description in Section 5.1.2. The instrument transmission spectra that this created are plotted in Figure 6-2. Due to the modular nature of the code these spectra can be easily replaced with higher accuracy simulated or measured filter spectra once it becomes available.

Now the instrument transmission spectra can be used to create a look-up table which relates a surface temperature and emissivity through to measured detector power. A temperature grid was made spanning the expected range in 0.5 K intervals. For each temperature the incident power at the TIRI detector is calculated. First, we define the spectral radiance at that temperature using the Planck function. Then the detector power is found by integrating the spectral radiance received at the end of the optical chain.

$$Detector\ Power_{filter} = A\Omega \int \epsilon_{surface} \tau_{instrument,filter} B_{\lambda}(T_{surface}) d\lambda \quad (6.1)$$

Where A is the pixel area, Ω is the solid angle, $\epsilon_{surface}$ is the emissivity of that surface, $\tau_{instrument,filter}$ is the instrument transmission function through that filter, and $B_{\lambda}(T_{surface})$ is the Planck spectrum for the surface temperature. For computational efficiency the Planck function for each temperature is generated once and saved, and then that data is used in each filter's integration.

An example of the resulting data, created for a blackbody surface, is displayed in Figure 6-3. The power curves differ between filters, reflecting the instrument's spectral sensitivity. For instance:

- Filter A (centred on 7.76 μm) shows the lowest power at low temperatures but experiences the largest inflection, rising to nearly the highest power at higher temperatures.

- Filter H (centred on 21.6 μm), on the other hand, receives the highest power at low temperatures but the lowest at high temperatures.




Figure 6-3: Power incident at the detector through each filter for a given surface temperature. This data has been generated for a perfect blackbody. The broadband channel is also calculated but it omitted from this plot to give clarity of the differing compositional channel curves.

This means that in our simulated data we expect to see a substantially greater contrast between high and low temperature regions in Filter A than in Filter H. this can be used to verify simulated measurements.

For cases where the comet model being observed will have multiple different emissivity present on its surface, multiple reference tables are created, one for each different emissivity profile, these are labelled based on the emissive function and so the correct look-up table can be selected during rendering.

6.3.2 Target Parameter Processing

The thermophysical comet data is loaded together from the output file of a thermophysical model which was created by Duncan Lyster in part to support this work. This model is similar to the external model used in Section 5.1.1.2, and is described in greater detail in (Lyster et al., 2024). The model predicts diurnal temperature variation across airless bodies by means of a three-dimensional illumination model which accounts for morphology, shadowing and sub-surface heating. This dataset includes the shape model (in STL format) and a corresponding CSV file containing temperature values for each facet of the comet's surface, as calculated by the thermophysical model. The shape model is then scaled to match the desired comet size, ensuring proper alignment with mission parameters.

Figure 6-4 shows the temperature map of the comet shape model used for this simulation. The illumination originates from infinity along the $X = Y$ line, making this side the warmest. The comet rotates around its Z-axis with a 12-hour period, resulting in cooler regions at the poles where the surface receives minimal or indirect sunlight. The purpose of using this irregular shape model, rather than a more spherical approximation, is to investigate the effect of complex topography on TIRI measurements. While a spherical comet would present a smooth temperature gradient (hottest at the equator and coldest at the poles), an irregularly shaped nucleus has warmer regions juxtaposed with colder ones, posing a greater challenge for TIRI's measurement capabilities and the interpretation of those measurements.

If a variable emissivity map of the nucleus is going to be used in the simulation, it can be defined at this stage. For instance, in Section 6.6.3, a composite nucleus is used with an emissivity map divided geometrically, where one half of the comet carries a different emissivity label than the other. In Section 6.6.4, a partially hydrated nucleus is modelled

with an emissivity map defined latitudinally, where the cooler polar regions carry hydrated minerals.




Figure 6-4: A visualisation of the temperature map on the surface of a comet from Duncan's open-source comet nucleus model (Lyster et al., 2024). The 67P shape model has been used, and here scaled to be 5 km across for the Comet Interceptor nominal case. The sun is within the X-Y plane at positive infinity in both X and Y.

Once the temperature and emissivity maps are aligned with the comet's shape model, the next step is to generate the detector power maps. These maps represent the power that would be incident on a detector pixel from each facet of the comet through each filter. The power for each facet is derived by interpolating values from the relevant temperature-to-detector-power reference table, based on the facet's temperature. The output of this step is a Power Data file, which contains the power incident at the detector for each facet (rows) across all filters (columns). This file serves as the foundation for rendering the detector's view in the subsequent stages of the simulation.

6.3.3 Flyby Parameter Processing

The flyby parameters are processed to calculate the trajectory and pointing information needed for the observation period, as outlined in the operations plan (Section 5.1). The time-series data from the operations plan provides the operation angle and the pointing mirror position at 1 ms intervals throughout the encounter. The instrument's view is simulated using the Visualisation Toolkit (VTK) in Python, which renders the comet's shape model in a 3D space. The toolkit creates a virtual camera representing the instrument, and its position, focal point, and orientation vector must all be defined in the 3D reference frame. Therefore, before simulating the instrument's view, these parameters need to be calculated for the full duration of the flyby.

The operation angle serves as a central reference for defining observation events such as data-cube positioning and is defined according to the vector relationship in Figure 6-5. This angle is useful because it retains its meaning independently of flyby-specific variables such as closest approach distance, encounter velocity, or flyby angle.




Figure 6-5: A diagram to clarify flyby parameters. The flyby angle is the angle between the comet-sun vector and the Comet Interceptor velocity vector. The operation angle is the angle between the velocity vector and the satellite-comet vector.

The flyby angle (which may range between 45° and 135°) significantly affects the TIRI observations. At lower flyby angles, TIRI will begin acquiring data-cubes while the

nucleus's terminator (the boundary between illuminated and shadowed regions) is still visible, allowing high-resolution mapping of both illuminated and unilluminated regions. As the encounter progresses, later data-cubes will capture a mostly illuminated surface, ensuring strong signals even as spatial resolution decreases due to increasing distance.

Conversely, at higher flyby angles, the initial data-cubes will have similar coverage due to the symmetry of the encounter, but later data-cubes will mostly observe the shadowed side of the nucleus. With much of the surface in shadow, temperatures could be too low for meaningful compositional signal retrieval in TIRI's channels, reducing the scientific utility of the later data-cubes.

Full derivation of the modelling used to generate the necessary geometric parameters for the model's rendering module are now provided.

6.3.3.1 Deriving Position Vector

The spacecraft's position throughout the encounter is derived using the operation angle, based on the geometry of the flyby parameters: closest approach distance and flyby angle.

- 1. Closest Approach Position:** The position at closest approach is simply the distance at closest approach, projected in the direction of Solar illumination, then rotated about the origin by 90° minus the flyby angle.
- 2. Distance from Closest Approach:** The distance travelled from the point of closest approach is computed using the operation angle.
- 3. New Position Calculation:** Knowing the flyby angle (which defines the direction of the spacecraft's velocity vector), the position at closest approach, and the distance travelled, the spacecraft's new coordinates can be calculated at any given operation angle.

Figure 6-6 illustrates the Comet Interceptor’s trajectory in the reference frame. The nucleus is at the origin and Solar illumination is from the $X=Y$ line. The spacecraft's trajectory is perpendicular to this when the flyby angle is 90° . Adjusting the flyby angle rotates the trajectory accordingly.




Figure 6-6: The spacecraft position in the space reference frame through the encounter (X and Y axis units in metres). This simulation is for a scenario where the comet is at the origin, the Solar illumination coming from an infinite point along $X = Y$, and the Comet Interceptor flyby angle is 90° and the distance to closest approach is 1000 km.

6.3.3.2 *Deriving Pointing Vector*

Given the position of the spacecraft and the angular position of the pointing mirror, we need to be able to define the instrument pointing vector so we can define render the detector's field of view. Confirming that the comet is in the FoV and, particularly that our scanning procedure, outlined in the operations plan, will result in the comet being placed in each of the filters correctly, allowing us to create a full data-cube.

This section will describe how the TIRI pointing vector can be calculated from the TIRI pointing mirror angle (Θ), the spacecraft velocity vector (\mathbf{V}), the spacecraft position vector (\mathbf{P}) and the CI target position vector (\mathbf{T}). This will be calculated in two discrete mathematical steps.

Firstly, a Θ dependent rotation matrix will be defined to find the TIRI pointing vector in the instrument reference frame. Secondly, a matrix transformation to move from the instrument reference frame to a space reference frame will be defined.

Consistent understanding of instrument and spacecraft axes is critical to ensure correct final vector definitions. The full definition of the MIRMIS coordinates and axes are contained in Section 3 of the MIRMIS Operations Plan document (COMET-OXF-MIR-PL-090). I repeat here the key information to enable the following analysis.

Figure 6-7 shows the URF axes overlaid onto the TIRI CAD model. The important directions to note are as follows:

- The X axis corresponds to the Spacecraft-Comet vector at closest approach.
- The Z axis corresponds to the negative spacecraft velocity vector, and thus defines the point the Spacecraft-Comet vector asymptotically approaches upon infinite recession of the spacecraft.




Figure 6-7: MIRMIS alignment cube and Universal Reference Frame (URF) definition. The instrument URF axes are coincident with the spacecraft URF axes [COMET-OXF-MIR-PL-090].

Defining key vectors:

Now that we have defined the TIRI reference frame we can define the key pointing vectors within it.

$$\text{TIRI pointing at } \theta = 90^\circ \text{ (closest approach)} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \mathbf{a} \quad (6.2)$$

$$\text{TIRI pointing at } \theta = 180^\circ \text{ (Z axis)} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \mathbf{b} \quad (6.3)$$

$$\text{TIRI pointing at } \theta = 270^\circ \text{ (calibration target)} = \begin{pmatrix} -1 \\ 0 \\ 0 \end{pmatrix} = \mathbf{c} \quad (6.4)$$

Each of these vectors is produced by the reflection of the pointing vector from the TIRI instrument out to the pointing mirror, I will refer to this as the input view.

$$\text{The input view (constant vector from instrument into pointing mirror)} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} = \mathbf{v} \quad (6.5)$$

Defining a Theta dependant rotation matrix:

The pointing mirror is at a constant 45° to the input view and rotates about the Y axis to enable the output view (the pointing vector from the pointing mirror out from the instrument) to take any value in the TIRI X-Z plane.

The pointing vector after the mirror can be described by a rotation matrix $\mathbf{R}(\theta)$ which satisfies the following equations:

$$\mathbf{R}(90)\mathbf{v} = \mathbf{a} \quad (6.6)$$

$$\mathbf{R}(180)\mathbf{v} = \mathbf{b} \quad (6.7)$$

$$\mathbf{R}(270)\mathbf{v} = \mathbf{c} \quad (6.8)$$

Since this only needs to be applicable for \mathbf{v} and this does not need to be generally reproduceable for any vector being reflected in the mirror, $\mathbf{R}(\theta)$ can take a very simple form:

$$\mathbf{R}(\theta) = \begin{pmatrix} 0 & \sin\theta & 0 \\ 0 & 0 & 0 \\ 0 & -\cos\theta & 0 \end{pmatrix} \quad (6.9)$$

Verification of $\mathbf{R}(\theta)$:

I wrote python code which plots the pointing vector before and after reflection at the surface of the mirror, using the newly defined transformation matrix $\mathbf{R}(\theta)$.

Several screenshots of plots from this code have been included in Figure 6-8. They show the expected behaviour in the pointing vector as the motor position is varied. At 90 degrees the TIRI pointing vector is entirely along the X axis. As the angle increases (tracking the position of the comet in this reference frame as its position increases in Z), the pointing vector moves within the X-Z plane, satisfying each of the required key pointing vectors in a smooth and continuous manner.




Figure 6-8: Representation of the pointing vector from the TIRI optics before and after reflection, at a set of different motor pointing angles.

Now that we can define the pointing within the context of the instrument, we need to be able to relate this to a pointing vector within the context of the flyby, and the space reference frame that the simulation is rendered in. The orientation of the spacecraft reference frame within the space reference frame must be defined. Then a transformation matrix can be created to rotate vectors in the spacecraft reference frame into the space reference frame.

Defining Reference Frame, Spacecraft Velocity and Position Vectors

The spacecraft velocity vector is a constant, defined by the flyby angle. The position however is defined along a straight line (the velocity vector) through the position of closest approach. Closest approach is defined as the 0 time, with a timeseries dataset of positions in space defined from that point. In this document the space reference frame being used is centred on the target with the comet nucleus as the origin. The Solar illumination comes in parallel along the $(-1, -1, 0)$ vector. Shifting to another reference frame will be another simple transformation.

Define reference frame orientation

The orientation between the two reference frames is defined in terms of certain constants that we know about the spacecraft, its orientation and trajectory. The Z axis of the spacecraft frame is aligned with the negative of the velocity vector of the spacecraft in the space frame. The X axis of the spacecraft frame is oriented such that the comet nucleus always lies within the X-Z plane. The origin of the space frame (the comet nucleus) takes a positive value in the X axis (the instrument views towards the comet). The y-axis of the spacecraft frame is orthogonal to both the X and Z axes and completes a regular right-handed coordinate system.

Calculating the space frame orientation

The following steps define the methodology used to define the spacecraft reference frame in its correct orientation within the space reference frame.

Determine the z-axis:

- Align the z-axis of the spacecraft frame with the negative of the velocity vector of the spacecraft.
- Normalise the velocity vector to ensure it has a unit length, then negate it to point in the opposite direction.

Mathematically this is:

$$\mathbf{v}_{normalised} = \frac{\mathbf{v}}{\|\mathbf{v}\|} \quad (6.10)$$

$$\mathbf{z}_{spacecraft} = -\mathbf{v}_{normalised} \quad (6.11)$$

Determining the x-axis:

- The x-axis should point towards the space frame origin from the spacecraft's position.
- Compute the vector that points from the spacecraft's position to the origin of the space frame.
- Project this vector onto the plane perpendicular to the z-axis to ensure orthogonality with the z-axis.
- Normalise the resulting vector to make it a unit vector for the x-axis.

Mathematically this is:

$$\mathbf{r}_{spacecraft-to-origin} = -\mathbf{r}_{spacecraft} \quad (6.12)$$

$$\mathbf{x}_{direction} = \mathbf{r}_{spacecraft-to-origin} - (\mathbf{r}_{spacecraft-to-origin} \cdot \mathbf{z}_{spacecraft})\mathbf{z}_{spacecraft} \quad (6.13)$$

$$\mathbf{x}_{spacecraft} = \frac{\mathbf{x}_{direction}}{\|\mathbf{x}_{direction}\|} \quad (6.14)$$

Determining the y-axis:

- The y-axis must be orthogonal to both the x and z axes to maintain a right-handed coordinate system.
- Compute the cross product of the z-axis and x-axis. The cross product of two vectors results in a third vector that is orthogonal to both.
- The resulting vector from the cross product is the y-axis and is already normalised.

Mathematically this is:

$$\mathbf{y}_{spacecraft} = \mathbf{z}_{axis} \times \mathbf{x}_{axis} \quad (6.15)$$

Constructing the transformation matrix

Once the X, Y, and Z axes of the spacecraft frame are determined, they can be assembled into a rotation matrix. This matrix represents the orientation of the spacecraft frame relative to the space frame.

The matrix is constructed by placing the X, Y, and Z unit vectors of the spacecraft frame as the rows of the matrix. Transposing this matrix aligns it correctly for transforming vectors from the spacecraft frame to the space frame.

$$\mathbf{R}_{temp} = \begin{bmatrix} \mathbf{x}_{spacecraft,x} & \mathbf{x}_{spacecraft,y} & \mathbf{x}_{spacecraft,z} \\ \mathbf{y}_{spacecraft,x} & \mathbf{y}_{spacecraft,y} & \mathbf{y}_{spacecraft,z} \\ \mathbf{z}_{spacecraft,x} & \mathbf{z}_{spacecraft,y} & \mathbf{z}_{spacecraft,z} \end{bmatrix} \quad (6.16)$$

$$\mathbf{R}_{transform} = \mathbf{R}_{temp}^T = \begin{bmatrix} \mathbf{x}_{spacecraft,x} & \mathbf{y}_{spacecraft,x} & \mathbf{z}_{spacecraft,x} \\ \mathbf{x}_{spacecraft,y} & \mathbf{y}_{spacecraft,y} & \mathbf{z}_{spacecraft,y} \\ \mathbf{x}_{spacecraft,z} & \mathbf{y}_{spacecraft,z} & \mathbf{z}_{spacecraft,z} \end{bmatrix} \quad (6.17)$$

This transformation matrix $\mathbf{R}_{transform}$ can be used to change a pointing vector defined in the spacecraft reference frame into the space frame. If the position also needs to be transformed, then the position and orientation information can be rotated together by using homogeneous coordinates (4x4) and then applying a transformation matrix formed of $\mathbf{R}_{transform}$ and the translation information. However, this is not currently necessary in this case as we just need to transform the pointing direction between frames.

These equations can be used in simulations to define the instrument pointing vectors in the space frame of reference, from the spacecraft velocity, position and motor angle. The final two components for the rendering are the focal point and the orientation about the pointing vector.

6.3.3.3 Deriving Focal Point

The Python VTK rendering functions require a set focal point, about which the scene will be rendered (to optimise computational efficiency), it cannot directly use a pointing vector. We can compute the focal point of the camera based on the position and pointing vectors by projecting the down the pointing vector, from the position vector by a distance equal to the magnitude of the pointing vector. This will place the focal point on the line of the FoV and close enough to the origin that the nucleus will fall within the scene limits and be rendered.

The focal point is calculated by projecting the normalised pointing vector from the spacecraft's position vector by a distance equal to the magnitude of the position vector:

$$Focal\ Point = Position\ Vector + \|\mathbf{Position\ Vector}\| \cdot \frac{\mathbf{v}}{\|\mathbf{v}\|} \quad (6.18)$$

6.3.3.4 Deriving the Orientation Vector

Once the camera position and pointing vectors have been defined, we still have a free parameter that needs to be constrained before image production can begin. The camera is currently free to rotate about its pointing vector, rotating the target in the field of view, and in the case of the off-centre views which are very important for TIRI, since they govern filtered images, the orientation will rotate the nucleus about the comet centre. An upward vector is defined to constrain the orientation of the image.

As we have noted previously when defining the TIRI operations plan, the instrument FoV rotates as the pointing mirror turns. This is a result of the FoV reflecting at 90 degrees off the pointing mirror as it rotates. Fortunately, this FoV rotation is linear with the mirror pointing angle and can be defined by setting an original upward vector and using a Rodriguez rotation along the pointing vector.

The up vector is initially set to $\begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix}$, representing a base direction at closest approach. The rotation angle is defined as the difference between the motor position and its starting position (at 90 degrees).

Rodriguez rotation allows for the rotation of a vector around a specified axis by a given angle. It involves calculating the components of the rotation matrix using trigonometric functions and the axis components. The resulting matrix is then used to transform the upward vector.

The Rodriguez matrix **Rod** is given by:

$$\mathbf{Rod} = \mathbf{I} \cos(\theta) + (\sin(\theta)) \mathbf{K} + (1 - \cos(\theta))(\mathbf{u} \otimes \mathbf{u}) \quad (6.19)$$

where θ is the rotation angle, \mathbf{I} is the identity matrix, \mathbf{K} is the cross-product matrix of the rotation axis \mathbf{u} , and $\mathbf{u} \otimes \mathbf{u}$ is the outer product of the rotation axis with itself.

For a given axis $\mathbf{u} = \begin{pmatrix} u_x \\ u_y \\ u_z \end{pmatrix}$, the cross-product matrix:

$$\mathbf{K} = \begin{pmatrix} 0 & -u_z & u_y \\ u_z & 0 & -u_x \\ -u_y & u_x & 0 \end{pmatrix} \quad (6.20)$$

With these derivations, the necessary geometric parameters for rendering the TIRI pointing vectors, focal point, and orientation are ready for use in the simulation.

6.3.4 Instrument Operations Simulation

6.3.4.1 *Rendering Individual Detector Views*

The instrument's field of view is rendered using the vtk in python. The rendering process incorporates several key inputs:

- Shape model (as described in Section 6.3.2),
- Power Data (derived from facet temperatures and filter response functions, Section 6.3.2),
- Spacecraft position, instrument focal point, and FoV orientation (Section 6.3.3),
- Filter labels, and
- Camera-specific parameters like FoV, pixel dimensions, and resolution scaling.

All images are initially rendered at high resolution (sub-pixel level), before being down sampled the detector resolution. This allows examination of the images in finer detail to draw out performance evaluation.

Figure 6-9 demonstrates the rendering output for four different TIRI filters. The comet's facets are visualised in grayscale, with each facet's brightness representing the detector power observed through the specific filter, according to the Filter Data dataset defined in Section 6.3.1. For visibility purposes, power values are normalised before rendering, ensuring that every filter reveals a range of facet brightness. This normalization process can be adjusted to reflect dynamic scaling challenges in the detector, while reverse mapping can be used to relate image brightness back to detector power, ensuring conservation of the thermophysical data.




Figure 6-9: Rendered images of the comet nucleus through three of TIRI's compositional filters, A, E and H, as well as the broadband channel. These renders are taken at a perpendicular to the plane of rotation, with the illumination from directly below in this orientation. These are not TIRI images, just renders using the filters to demonstrate the differences in the observed target.

6.3.4.2 Simulating Operation

Once the complete time series dataset has been generated for the flyby (with all geometric and operational parameters calculated at 1 ms intervals), the TIRI operation simulation can be run independently. This dataset, produced and stored separately, can be further analysed as needed.

The simulation utilises the data-cube information outlined in the operations plan (e.g., filter orders, positions, and integration times). For each period of integration (each time step during an integration window), the detector view is rendered for the relevant filter, and the corresponding rendered images are saved into folders.

After rendering, image formation occurs through down-sampling the high-resolution images to the TIRI detector's resolution of 640x480 pixels. This process averages each pixel's "colour" (which relates to the detector power) across the entire rendered facet area, simulating how the real detector pixels integrate power across the viewed surface. Essentially, the process mirrors how the detector records incident power from observed sources.

Once all filter images are created, these individual images are stacked into a multidimensional image array, representing each integration over time. The pixels are then averaged across the layers to form the measured detector image for each filter.

Figure 6-10 illustrates the progression of image processing, from rendering to the final detector measurement. The final image on the right shows some blurring compared to the central image, primarily due to the limitations of keeping the comet nucleus




Figure 6-10: The down sampling and integration process. These three images show the comet at different stages of rendering for the formation of a filter G image in the first data-cube. The left image shows the nucleus as rendered for a 1ms view. The central image shows the image on the left but downsampled to the detector resolution. The right image is the filter G measurement, having averaged a series of images like the one in the centre.

perfectly stationary in the detector's FoV. Blurring predominantly occurs in the vertical direction, as the pointing motor only controls the horizontal axis during the first data-cube. Further investigations into motor position noise and jitter could introduce additional blur along the motor control line, warranting further exploration to assess how motor pointing uncertainty impacts spatial resolution in the measured data.

6.3.5 Data-Cube Generations

Once the simulated detector measurements are produced, the process for generating data-cubes follows the same pipeline as the real TIRI data handling, with some modifications for the simulation. The data-cube generation involves two key steps: firstly, the coalignment of images so that the spectral data from each filter corresponds to the same spatial region of the comet, and secondly, the transformation of pixel colour data into scientifically meaningful outputs, such as calibrated radiance, power, energy, brightness temperature, and emissivity.

6.3.5.1 Measurement Coalignment

Due to TIRI's filter layout, the comet nucleus appears in different positions across the images from each filter. The first data-cube's images, shown in Figure 6-11 illustrate this misalignment. Accurate coalignment is critical for meaningful comparison between the spectral data across the filters, as misalignment would distort the spectral information.




Figure 6-11: All filter images for the first data-cube. As we can see the nucleus inhabits a different region of the detector in each image as the motor has moved the field of view to capture each image in a different filter.

This process is done in Python using an object detection algorithm suitable for the simulated image data. The main goal of this algorithm is to identify and isolate a specific object across the multiple filtered images, align these images based on the object's position, and then crop these images to a uniform size for further analysis. More detail of each step is given below.

Identifying the Comet

Since the rendered images have a dark background (with a colour value set to zero), the comet appears as the only object with non-zero pixel values. The comet is identified by scanning for these non-zero values and defining a bounding box that encloses the comet nucleus. This bounding box is calculated by determining the minimum and maximum coordinates that contain all the non-zero pixels in the grayscale image.

Cropping the Image

Once the bounding box is determined for each image, the images are cropped to a uniform size. The cropped dimensions match the largest bounding box found, and the cropping is centred on the comet nucleus in each filter image. This ensures that the images are equal in size and have the nucleus aligned in the centre. To refine the results, a manual fine-tuning tool was created, allowing the user to adjust the final alignment of the images within the data-cube if necessary. Figure 6-12 shows the cropped objects identified using this algorithm from the images in Figure 6-11.




Figure 6-12: The comet nuclei, selected and cropped by the object identification algorithm, labelled by the filter they have been imaged through. These are the images that are stacked to form the final data-cube.

Stacking and Saving

After all images have been aligned and cropped, they are stacked into a single three-dimensional NumPy array. This array enables multi-dimensional analysis and serves as the foundation for further processing, including the calculation of the science products that will be derived from the measurements.

6.3.5.2 *Populating the data-cube*

Once the co-alignment process is complete, the next step involves extending the stacked array of pixel measurements to calculate the relevant scientific properties of interest. This process transforms the raw pixel values into meaningful data outputs such as incident power, received energy, calibrated radiance, brightness temperature, and emissivity. These calculated properties form a 4D NumPy array, where each pixel's position (X, Y), the filter, and the calculated property are recorded.

By populating the data-cube with these properties, we can derive meaningful scientific insights from the simulated TIRI measurements. Each property includes its own uncertainties, allowing for a comprehensive understanding of the instrument's performance under varying operational conditions. Here is an overview of how each property is derived:

Pixel Value

The pixel value is the initial grayscale intensity or brightness value from the coaligned, rendered images. This value forms the basis of further calculations.

Incident Power

Incident power is determined by using a reverse mapping from the pixel value to power for each filter, based on the Filter Data from the instrument setup. Each pixel's value corresponds to a weighted average of the detected powers, as the TIRI detector views multiple facets of the comet surface at different temperatures within a single pixel. This behaviour simulates the actual instrument operation. The NEP of the TIRI detector is used to calculate the uncertainty in the power measurement.

Received Energy

Received energy is derived by multiplying the incident power by the integration time for that measurement. This calculation allows visualization of how noise levels within the detector affect the ability to distinguish signals over different integration periods. The uncertainty in energy, tied to the NEP, remains constant, but as the energy increases with longer integration, the SNR improves.

Calibrated Radiance

Calibrated radiance is the primary science product, representing the radiance values measured by TIRI, typically returned as the highest-level data product. It is calculated from the power measured by each pixel, converted into radiance using the pixel area and solid angle product ($A\Omega$) in m^2 per steradian:

$$\text{Calibrated Radiance}_{\text{filter,pixel}} = \text{Power}_{\text{filter,pixel}} * A\Omega \quad (6.21)$$

The uncertainty in calibrated radiance is similarly derived from the uncertainty in the power measurement.

Brightness Temperature

Brightness temperature is interpolated from the incident power using a precomputed reference dataset that links detector power to specific temperatures, based on the Planck function. However, since each pixel captures power from multiple temperature sources, the returned brightness temperature is an approximation and does not represent the actual surface temperature. This aspect mirrors the challenges of the real instrument and provides a valuable opportunity to explore methods of improving temperature retrieval accuracy. The uncertainty in brightness temperature is determined by recalculating the temperature with the incident power adjusted by its uncertainty (positive and negative), yielding a temperature range.

Emissivity

Emissivity is calculated by comparing the calibrated radiance measured by the instrument to the radiance expected from a surface at a given brightness temperature. For this model we are using the brightness temperature through the broadband channel as a definitive surface temperature (this is not the true surface temperature however).

$$Emissivity_{pixel,filter} = \frac{Calibrated\ Radiance_{pixel,filter}}{Reference\ Radiance_{filter}(Brightness\ Temperature_{pixel})} \quad (6.22)$$

The reference radiance is derived from the instrument model for a blackbody or known emissive surface. Since brightness temperature is used to approximate the true surface temperature, current emissivity retrievals may exceed physical bounds (i.e., emissivity > 1), but ongoing work on improved retrieval algorithms will refine this process. This is discussed further in Section 6.7.

6.4 Data Visualisation

I developed a set of Python tools to allow the Comet Interceptor science team to visualise and interact with the data generated by the TIRI simulation. This section outlines these tools, highlighting how they support engagement with the data and enable a deeper analysis of the simulation outputs. These tools are then used to explore several simulated cases, demonstrating various aspects of the simulation's performance and utility.

6.4.1 Standard Methods

The simulated data-cube can be examined either by individual layers or at the pixel level. Python functions accompany the data-cube to display all calculated properties - such as power, radiance, and temperature - at any selected pixel and filter. Figure 6-14 shows an example of this pixel-level interrogation.

```

Analyzing Pixel (4, 13) for Filter 'A':
Pixel Value: 39.0 Scaled /255
Power: 4.022047064806825e-12 Watts
Energy: 2.252346356291822e-13 Joules
Calibrated Radiance: 0.03794384023402665 W/m^2/sr
Brightness Temperature: 174.17950536389236 Kelvin
Power Noise: 7.18381116519239e-12 Watts
Energy Noise: 1.2020815280171307e-12 Joules
Radiance Noise: 0.06777180344521123 W/m^2/sr
Temperature Uncertainty: 10.014268293093437 Kelvin
Emissivity: 0.08029664728977483 Unitless
Emissivity Uncertainty: 0.14343491695962116 Unitless
    
```

Figure 6-14: Example data output for a single pixel and filter; where all the properties have been retrieved.

By selecting the same property across different filters at a specific pixel location, the spectrum of that property can be retrieved, as demonstrated in Figure 6-13. For example, in the calibrated radiance spectrum for a high-temperature, high-emissivity surface, uncertainty is minimal. However, as surface temperature decreases, or if the surface exhibits complex topography with varying emissivities, the spectrum will have greater uncertainty and a less distinct curve.




Figure 6-13: Measured calibrated radiance spectrum for Pixel (7, 7). This spectrum is generated for a surface with a brightness temperature measured as 331.8K, modelled with an emissivity of unity.

These functions also allow mapping of specific properties, such as temperature or radiance, across the comet's surface. This is particularly useful for generating images that highlight gradients and contrasts within the selected property across the nucleus.

6.4.2 Multicolour Images

A more advanced method involves creating multicolour composite images from different layers of the data-cube. Three layers - such as the radiance or emissivity in different filters - are assigned distinct colours (e.g., red, green, and blue), and combined into a single image. This technique reveals gradients and correlations between properties across the comet's surface.

Figure 6-15 shows an example of this visualization approach, useful for identifying underlying differences in surface properties. For instance, assigning Filter A's emissivity to red and Filter B's emissivity to blue would result in purple areas indicating high values in both filters, while red areas highlight regions with stronger Filter A emissivity. By adjusting these mappings and colour weightings, specific mineral compositions or temperature variations can be isolated and visualised with greater contrast.




Figure 6-15: A set of multilayer nucleus images. In each of these images the red colour value is set by the Filter A radiance, green is set by the Filter E radiance and blue is set by the Filter H radiance. This imaged nucleus has a unitary emissivity so the underlying colour changes only show the different radiance through the filters due to different temperatures present in each pixel.

6.4.3 Interactive Tools

To enhance exploration of the data-cube, I developed an interactive Python script that dynamically visualises the dataset in real-time. This tool displays a broadband image of the comet alongside a spectral plot for the selected pixel, as shown in Figure 16. The user can toggle between various properties, Radiance, Power, Energy, Brightness Temperature, and Emissivity, via buttons at the bottom of the interface. The spectrum for the selected pixel updates as the cursor is moved over different areas of the image.




Figure 6-16: Screenshot of dynamical data-cube visualisation, showing the simulated measured calibrated radiance spectrum in pixel (10, 8) from the 1st data-cube.

Additionally, functions have been created to quickly plot or retrieve specific spectral data from the stacked array, facilitating rapid pixel-level investigation. This interactive approach encourages familiarity with the dataset and assists in identifying specific regions within the data-cube that require further scrutiny, speeding up the process of data interpretation.

These visualization tools provide a robust framework for analysing TIRI's simulated observations, enabling a direct comparison of how different flyby parameters or

target properties influence the data product. The ability to explore both spectral and spatial properties interactively ensures that the science team can efficiently identify key features and potential areas of interest within the dataset.

6.5 Verification

The process of retrieving calibrated radiance through this simulation involves several steps of abstraction. The temperature and emissivity maps, and instrument transmission functions have been used to define the detector power, and thus an arbitrary colour used to render each facet for the detector-power map. The powers are normalised, and the colours are normalised before the process of rendering starts, and then pixels are down sampled, averaging colour values into sizeable pixels. Finally, the colours are transformed back into radiative measurements using the reverse mappings and normalisations that were used before rendering. I wanted to be able to prove that the results that we are generating are genuinely representative of the surface and that there are no errors in the code which are affecting the results.

To verify the simulation outputs, I compared the results against an independent TIRI model written in MATLAB. I wrote this model to assess the performance of alternative TIRI filters, discussed in Section 5.3. It uses radiance spectra generated by PSG, and evaluates the measurements that would be obtained through the TIRI instrument. Unlike the Python simulation, this MATLAB model does not involve any geometric considerations, focusing solely on radiative performance. Crucially, it has been previously validated against other instrument simulations within the MIRMIS science team, and its consistent results provide a reliable baseline for comparison. The only common feature between the MATLAB model and the Python simulation is the use of TIRI transmission functions, ensuring that any discrepancies are not due to shared code.

I chose the first data-cube for this comparison since it covers the smallest area per pixel, making it the most likely to observe uniform surface regions. I generated a radiance spectrum for this data-cube and compared it to the spectrum produced by the MATLAB model, which used a unitary emissivity and a 320 K surface temperature from the PSG.

Figure 6-17 displays the comparison between the two radiance spectra. Both spectra align closely, with agreement within a 10 percent in most channels. Given that the Python simulation models a comet surface with varied topography and temperature, while the MATLAB model assumes a flat, uniform surface, some discrepancy was expected. However, the small differences between the two demonstrate that the Python simulation is accurate and free from significant errors.




Figure 6-17: Comparison of radiance spectra between the full simulation output (right) and the expected TIRI outputs from a much simpler, and previously verified, TIRI instrument model (left) written in MATLAB. This is an important verification step since the method of arriving at the spectrum through the full optical simulation is substantially different to the MATLAB model.

I have also verified the computation of the all the parameter values in the final data-cube as seen in Figure 6-14. This has been checked by simulating representative radiances through the TIRI transmission functions and manually calculating each parameter value and cross referencing those against the values produced by the simulation.

6.6 Simulation Results

All the images shown throughout Section 6.4 were from an initial simulation run, considering a comet with flat unitary emissivity. In this section I probe the results of that simulation further (6.6.1), as well as considering a cold nucleus (6.6.2), a binary material comet (6.6.3), and finally a representative nucleus, with a mixture of PSG generated rock and ice spectra (6.6.4).

6.6.1 The Hot Comet Case

This simulation utilises the same comet temperature map from throughout Section 6.3, demonstrated in Figure 6-4, which exhibits surface temperatures from 330 to 200 Kelvin. Figure 6-18 presents the brightness-temperature maps for the comet nucleus, as captured by TIRI in each data-cube. Several important conclusions can be drawn from these representations. Firstly, the model's ability to capture the geometric behaviour of the flyby is validated. The comet's anticlockwise rotation in the field of view can be traced by following the movement of the hottest part of the nucleus across the images. Furthermore, as the flyby progresses, the spacecraft's viewing angle clearly shifts. Notably, the third data-cube captures an "end-on" view of the comet, in contrast to the "side-on" perspective in the first data-cube.

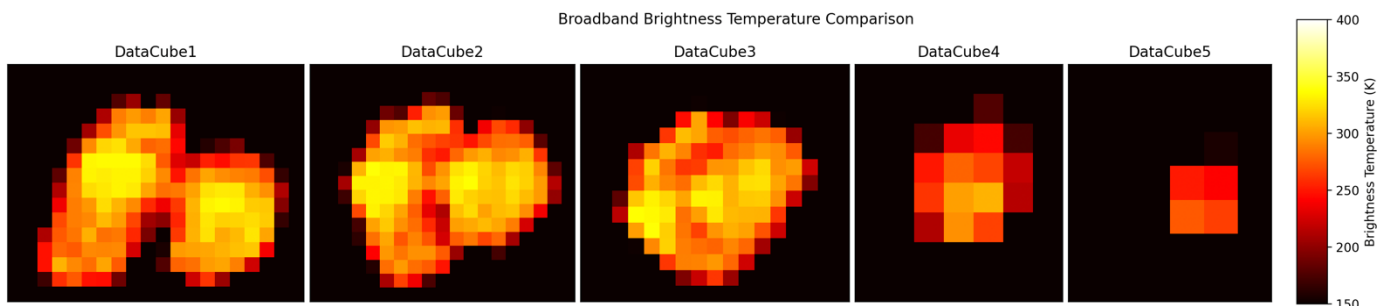


Figure 6-18: Nucleus brightness temperature maps from the broadband channel for each TIRI data-cube.

The relative resolution between data-cubes provides further visual confirmation that detailed nucleus mapping will be constrained to the first few data-cubes. In the final data-cubes, the comet's apparent size is significantly reduced relative to the pixel resolution, making high-resolution mapping impossible. Therefore, I propose different operational modes for each data-cube, tailored to the specific data that can be realistically extracted. For the first three data-cubes, it is essential to balance the demands on spatial resolution and signal generation, as outlined in the operations plan. However, in the fourth and fifth data-cubes, spatial resolution is so diminished that the priority shifts towards maximizing the signal. In these cases, longer integration times per channel, and averaging across pixels, would allow the instrument to focus on retrieving a "bulk compositional spectrum" with enhanced SNR. One confounding factor to manage with this is the observation of multiple temperatures within a single pixel.

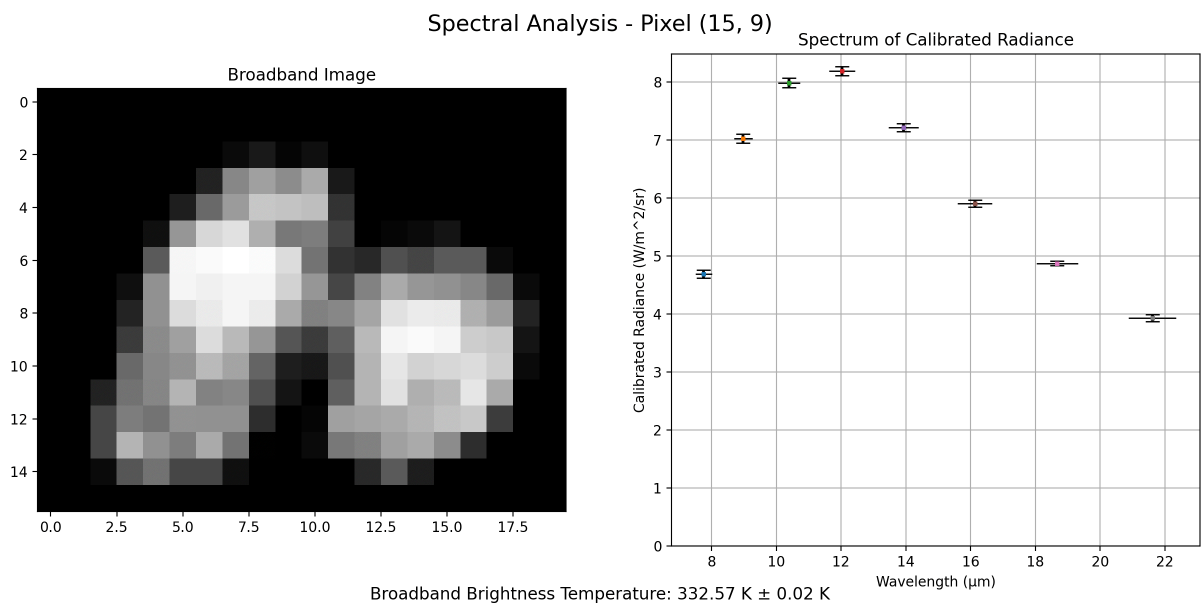


Figure 6-19: The calibrated radiance for the warmest pixel on the first data-cube. The signal-to-noise on the radiance measurements are very high since the surface is so warm.

In the case of this simulated hot comet, with surface temperatures reaching up to 350 K in the warmest regions, the retrieval of accurate radiance spectra presents minimal challenges. As shown in Figure 6-19, the error bars on the calibrated radiance for the hottest pixel in the first data-cube are very small, confirming the high SNR (> 90) of the radiance measurements at these temperatures. However, this performance will degrade as

emissivity drops or in areas with lower surface temperatures, where noise becomes more significant.

6.6.2 Cold Nucleus Case

To investigate how nucleus temperature impacts the measurement results, I conducted a simulation of a colder comet. Building on the thermophysical modelling results in Figure 5-5, which suggest that an ice dominant comet could be much colder, the temperature map from the previous section was used, but each facet's temperature was reduced by 100 degrees, resulting in a maximum temperature below 250K. This modified temperature map is shown in Figure 6-20. Aside from the temperature change, the flyby and operational parameters remain identical to the hot comet case, allowing for direct comparison between the two scenarios.

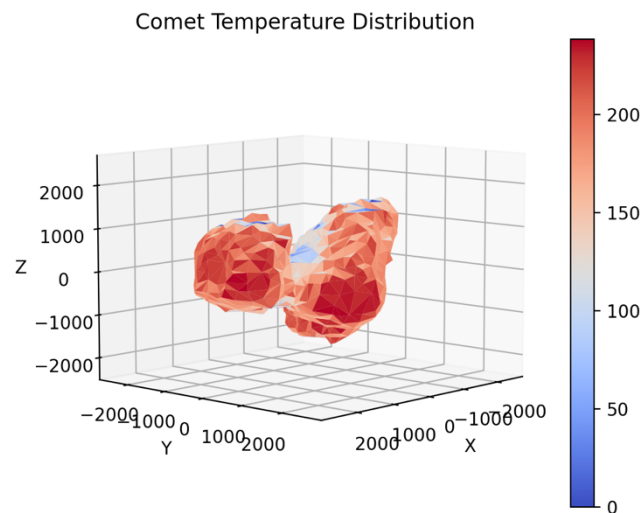


Figure 6-20: The reduced temperature map used in this cold comet analysis. The comet temperatures range from 250 K in the illuminated equatorial region, to below 100 K in the polar and shaded regions.

The cold nucleus presents significant challenges due to the lower radiance reaching the detector. The signal incident on the detector is weak, and as outlined in the operations plan, the scientific requirements demand a high SNR. This SNR is not for detecting the signal itself, but for distinguishing between the emissivities of different

bands. The challenge arises when noise levels approach the magnitude of the differences between band signals.




Figure 6-21: Calibrated radiance measurements for the hottest pixel in this data-cube from the simulation of the cold comet scenario.

By examining the radiance spectrum in Figure 6-21, we see the noise levels are comparable to the differences between the channels. For instance, in channels D, E, F and G (at approximately 12, 14, 16, and 19 microns respectively), the uncertainties overlap considerably. If we cannot differentiate the radiance between these channels, it will be difficult to accurately retrieve surface temperature or discern emissivity relationships in these specific bands, preventing reliable identification of mineral spectra. The challenges or retrieving parameters from radiance is discussed further in Section 6.7. This contrasts with the hot comet case in Figure 6-19, where the noise is comparatively lower, allowing each channel to remain distinct. Consequently, the cold case simulation demonstrates that the instrument will struggle to retrieve compositional data when observing lower-temperature regions, whether on a warm comet or in a uniformly cold environment. This highlights the importance of selecting higher-temperature targets to optimise science returns.

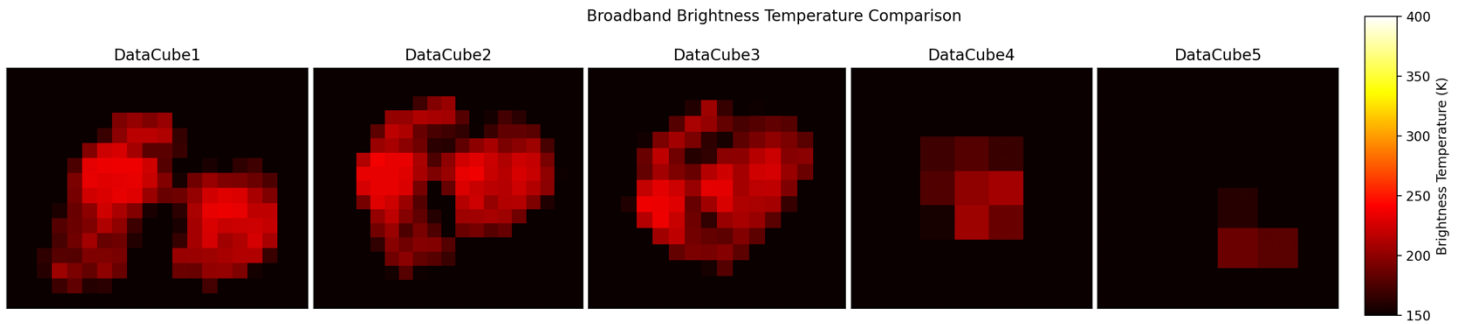


Figure 6-22: The brightness temperature images of the nucleus produced for each data-cube across the encounter. These have been plotted on the same temperature range as the plots in Figure 6-18. Despite the width of the broadband channel, the signal at short integration times will still be low in cases of low temperatures.

Additionally, temperature mapping becomes more challenging when the signal is reduced, as is the case with this cold comet. The full set of brightness-temperature maps for the encounter is shown in Figure 6-22. For real-world low-temperature targets, the instrument's operational plan should be adapted accordingly. Prioritizing spatial resolution for compositional maps is not practical if the composition cannot be reliably retrieved. A more effective strategy might be a hybrid operational mode, where broadband images are captured with short integration times, similar to the original MIRMIS operations plan, while the compositional filters use extended integration times, sacrificing spatial accuracy for more robust spectral retrievals. Further data processing techniques, such as averaging across pixels to improve signal strength, could also enhance results. These alternative operational strategies can be evaluated using the TIRI simulation tool.

6.6.3 Composite Material Case

To simulate a comet with variable compositions, an emissivity map was created for the 67P comet shape model, representing a nucleus formed from two distinct materials. Each lobe of the comet model was assigned a different emissivity value. To simplify the analysis and enhance visibility, two extreme simulated emissivity spectra were used. In the first spectrum, referred to as Material A, the emissivity is set to 1 for wavelengths below 14 microns and 0.4 for wavelengths above. In contrast, Material B has the opposite characteristics, with an emissivity of 0.4 below 14 microns and 1 above.

The emissivity map was constructed by identifying a plane that separates the two lobes of the comet and assigning the appropriate material to each facet based on its location relative to that plane. The resulting material distribution is shown in Figure 6-23. Aside from the differing emissivity distributions, the same shape model, temperature map, and flyby parameters were used as in the hot comet simulation to ensure comparability.




Figure 6-23: Material distribution across the shape model for the variable emissivity simulation. Material A is coloured in red, and material B in blue.

The individual filter images, shown in Figure 6-24 , display the radiance observed through each filter, effectively capturing the emissivity differences across the comet's surface. This is an extreme case designed to validate whether the simulation images can accurately represent emissivity variations across the nucleus. The results clearly illustrate the differing radiance magnitudes caused by the material distribution, confirming the model's ability to capture this information.




Figure 6-24 : The variable emissivity comet nucleus imaged through each filter. In filters A, B, C, and D, Material A appears much brighter; making the other lobe difficult to see due to colour normalization. In filters F, G, and H, the opposite occurs, with Material B dominating. Filters E and the broadband filter span the 14-micron change point, rendering both materials at similar brightness levels.

The variability in signal levels enables the use of the three-color images discussed in Section 6.4.2, as seen in Figure 6-25. In this image, we can discern the material distribution across the comet's surface. The left-hand side of the nucleus is predominantly blue due to the stronger signal in filter H, indicating a higher emissivity at longer wavelengths, corresponding to Material B. Conversely, the right-hand lobe appears red, representing Material A. The regions where green dominates (from filter E) over red reflect the influence of surface temperature variations. As the temperature changes, the blackbody radiation peak shifts within the wavelength region observed, causing the relative radiance levels in each filter to fluctuate, as previously observed in Figure 6-3.

Three Color Image
R: A, G: E, B: H




Figure 6-25: A three-colour image of the nucleus formed from the filter images in Figure 6-24. Filter A signal is shown in red. Filter E signal is imaged in green and filter H signal is imaged in blue.

Retrieving surface temperatures is relatively straightforward when the emissivity is flat across the 7 to 25-micron region, as the Planck function's peak can be easily identified from the measured radiance. However, when emissivity is non-constant, retrieval becomes more challenging. Additionally, each pixel observes regions at different temperatures and with varying emissivity spectra, further complicating the retrieval process. This simulation tool proves invaluable for generating controlled spectral datasets,

allowing for the testing and verification of retrieval algorithms under well-defined surface conditions.

6.6.4 Representative Comet Materials

The previous analyses have focused on simulations using unrealistic, simplified cometary spectra, either with uniform emissivity across the surface or highly polarised emission values not typically observed in real cometary materials. These simulations have been useful in evaluating the functionality of the model and identifying the limitations of the instrument's data products, such as difficulties in compositional retrieval at low temperatures. However, to further explore instrument performance in a more realistic scenario, I have developed a simulation of a comet with representative surface spectra, highlighting the key challenges we might face.

For this simulation, I used the same shape model and temperature map as before (shown in Figure 6-4) but generated more representative emissivity spectra using the PSG. Two distinct surface spectra were considered. The first was a "dry" rocky mixture composed of three rock minerals, a fayalite, enstatite and pyroxene in equal proportions. The full details of the minerals used are given in the Appendix Section 9.2. The second was a mixture, 60% of the same rocky material combined with 40% water ice. The emissivity map was generated by assigning the dry rock to all surfaces above 220 K, while the rock-ice mixture was applied to colder regions (below 220 K), primarily in the comet's polar regions and shaded areas with lower average temperatures. This setup reflects the logic that a comet with a rotational axis perpendicular to the orbital plane would experience high temperatures at the equator (up to nearly 350 K), reducing the likelihood of surface ice presence, while the polar regions, staying below 200 K, would retain higher ice concentrations. This emissivity map is shown in Figure 6-26.

The primary objective of this simulation is to determine whether the rock composition can be identified and whether the presence of water ice can be detected based on TIRI measurements under these conditions.




Figure 6-26: The compositional mapping of the comet nucleus used for the representative spectrum simulation, viewed from above to display one polar region. The red regions represent the facets labelled with the pure rock mixture emissivity. The blue facets are the cold regions selected by a temperature cutoff, which will exhibit the rock and ice mixture emissivity.

The PSG-generated surface spectra are displayed in Figure 6-27. From this, we can identify specific filters where the differences between the rock and rock-ice spectra are most pronounced. For instance, in filters B, C, F, and G, the emissivity of the pure rock surface is higher than that of the rock-ice mixture. A reduction in emissivity across these channels may suggest the presence of water ice. Conversely, filter E shows higher emissivity for the rock-ice mixture, making it a potential indicator for ice detection. This multi-filter approach allows us to isolate compositional differences from temperature variations, enhancing our ability to accurately interpret the data.




Figure 6-27: The dry rock, and rock and ice mixture spectra as generated through PSG, the exact PSG inputs are provided in the appendix. The emissivity spectra have been overlaid on the TIRI filter measurements so we can see which filter channels will enable differentiation between these two spectra.

The simulation was executed with a spacecraft trajectory at a 45-degree flyby angle, ensuring the first data-cube captured the polar region where ice is expected. A second, identical simulation was run for a purely rocky surface to serve as a point of comparison, allowing us to differentiate the ice signals.

Based on the plots in Figure 6-27, filters C, E, and G were chosen for analysis in the emissivity plots, with filter C mapped to red, E to green, and G to blue. Since the emissivity in the rock-only regions is higher in filters C and G, these areas will appear more purple (high red and blue values), while regions with ice should appear green due to elevated emissivity in filter E. When plotted as a radiance map (see Figure 6-28), the dominant factor is temperature, making it difficult to extract compositional information. The goal of

post-measurement data analysis is to separate temperature effects and isolate the compositional features, which I will demonstrate.

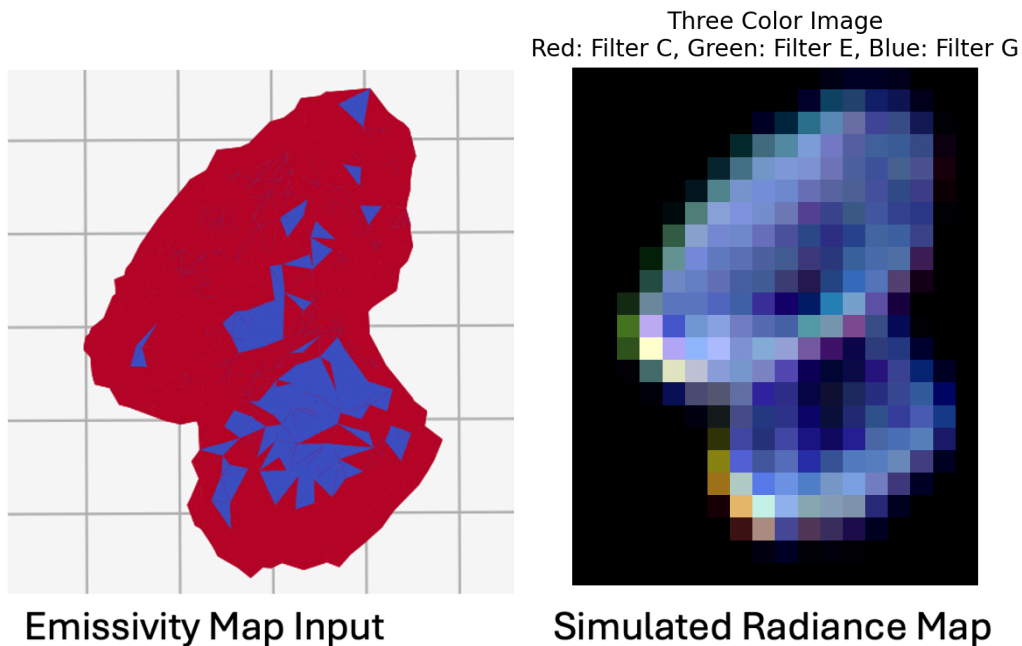


Figure 6-28: The three-colour radiance image of the comet plotted next to the emissivity map. Since this is a radiance map the colour variation is governed by the surface temperature rather than any emissivity variations.

The initial three-color emissivity plot is not informative due to issues at the edges of the nucleus, where pixels capture both the comet and background space. These edge cases generate extreme values, which damp other results in the image normalisation process. Improved co-alignment algorithms during data-cube generation could mitigate these issues. For now, a series of filters was applied to handle extreme values, masking anomalous emissivities and normalizing the remaining pixels, as shown in Figure 6-29. This significantly improves the image, allowing for a clearer analysis of the comet's surface.

For the refined image in Figure 6-29 we expected to see a predominantly purple comet, with green regions indicating the presence of ice. Initially, the simulation showed the comet as green, with purple areas where ice was expected. Upon investigation, a user error was discovered where the emissivity values had been misapplied, causing the polar

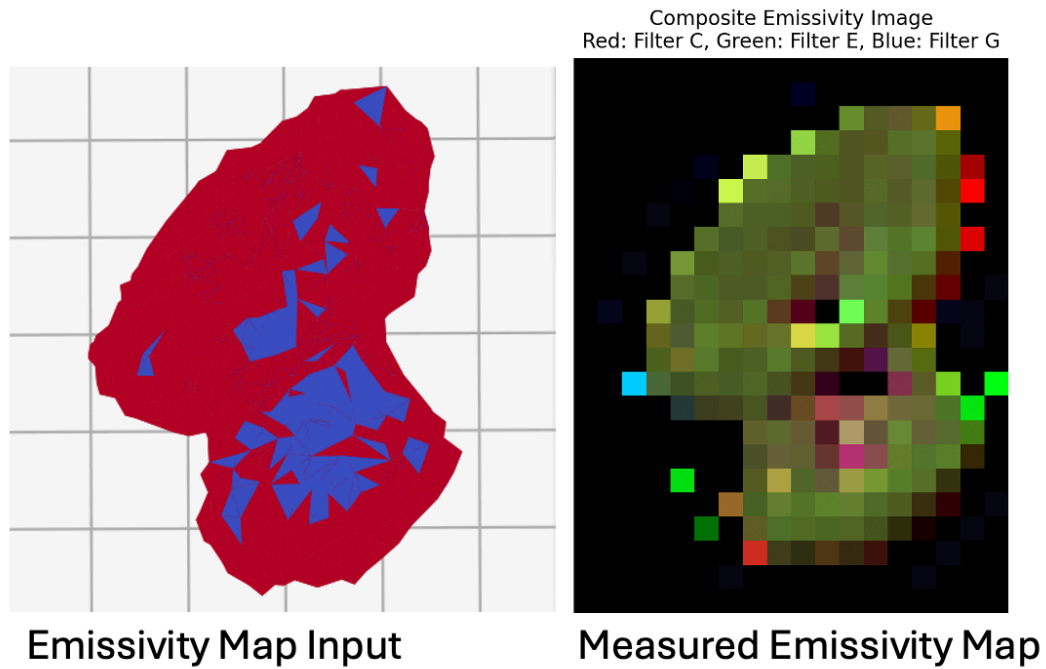


Figure 6-29: A well filtered three-colour emissivity image for filters C, E and G, shown next to the emissivity map used to generate the simulation. The image has been formed from the Figure 6-26 emissivity map by using masks to remove the anomalously high and low emissivity pixels, setting them to 0, and then renormalising the remaining pixels.

regions to be simulated as pure rock. Remarkably, the TIRI measurements correctly identified the erroneous ice signals, demonstrating the instrument's capability to detect water ice under the proper conditions. After correcting the mapping, the comet's surface, as shown in Figure 6-30, exhibits the expected composition - a brownish, rocky nucleus with distinct green regions denoting water ice. The smaller areas of ice in the upper regions are also clearly visible, indicated by subtle green shading due to elevated emissivity in filter E.

Composite Emissivity Image
 Red: Filter C, Green: Filter E, Blue: Filter G




Figure 6-30: The comet nucleus re-modelled with the correct mapping for the rock and rock-ice mixture regions: A well filtered three-colour emissivity image for filters C, E and G, shown next to the emissivity map used to generate the simulation. The image has been formed with the same masking and filtering as the previous image.

To directly compare the emissivity spectra, Figure 6-32 shows the measurements from the same pixel (located at 9,18) for both the rock and the rock-ice simulation. The rock-ice mixture shows a lower emissivity in filter C and a higher emissivity in filter E, exactly as predicted by the spectra in Figure 6-27. However, the uncertainties for filters C and E overlap, meaning actual measurements could make distinguishing between features more challenging. The size of the emissivity uncertainty bars is comparable to the differences in emissivity values, necessitating further testing of retrieval techniques against measured and simulated noise to assess their robustness (discussed in Section 6.7). Fortunately, the tools developed here can generate well-labelled, representative surface data for testing these techniques.




Figure 6-32: The measured emissivity spectra for the same pixel (9,18) on the 1st data-cube of the rock and ice simulation as well as the pure rock simulation. The graph on the left shown the emission measured in the presence of water ice, and the graph on the right shows the retrieved emissivity spectrum for a rock only simulation. The emissivity values are so large due to the error in surface temperature retrieved. Both set of values were retrieved by the same methodology, so the errors are consistent.

To fully evaluate the impact of the emissivity uncertainties we need to observe noise within the measurement, I created a script to inject random noise into the data, using a Gaussian distribution with a standard deviation based on half the uncertainty for each filter and pixel (which is defined based on the detector noise). The resulting emissivity maps, shown in Figure 6-31 , demonstrate that even with noise, the ice-related emissivity




Figure 6-31: Three emissivity plots for the mixed material comet, where random noise has been injected into the data-cube according to the uncertainty on each filter measurement in each pixel. The image has been reproduced multiple times to show the random nature of the noise injected, and to demonstrate that the ice produced emissivity feature (the green region), is typically still visible beyond the noise.

feature (green regions) remains visible, confirming the instrument's ability to detect this water-ice despite the noise in the measurement.

6.7 Surface Parameter Retrievals

Analysis of the measurements produced through the simulation has highlighted the importance of the TIRI retrieval algorithms. These are the tools which use the calibrated radiance measurements produced by the instrument and derives surface property parameters from them through a parameter optimisation process. The simulation has provided a framework for an instrument data pipeline, and it provides the opportunity to create data to probe and refine this process.

6.7.1 The Challenge of Retrieval

One of the principal challenges in TIRI retrievals surrounds the ability to produce accurate temperature measurements from the data retrieved. As discussed, the instrument records radiance in reference to calibration targets, and these are accurate to within the requirements contingent on them. However, these radiances need to be processed in order to produce temperature measurements. The brightness temperatures defined in 6.3.5.2 are a method for this, but the assumptions of unity emissivity mean the brightness temperature will be increasingly different to the surface temperature as emissivities decrease. Prior measurements of comet surfaces on 67P and Halley imply emissivity values above 0.8 (Emerich, 1987; Marshall et al., 2018), but a dynamically new comet may be lower due to the increased ice present at the surface. An illustrative higher albedo body is Europa, which exhibits an emissivity of 0.75 in its ice-rich surface (Trumbo et al., 2018).

On LTM (as on Diviner), this problem is solved by using the four broadband thermal channels in concert. These discrete channels measure from 7 to 100 micrometres, and therefore the shape of the blackbody curve can be constrained, which allows for reliable inference of the observed temperature. However, since TIRI only has a single broadband channel across 7-25 microns, the shape of the Planck function cannot be constrained. The shape of the radiance measurements across the compositional channels is dominated by the emissivity spectrum, but accurate retrieval of that spectrum is contingent on accurate removal of the surface temperature.




Figure 6-33: The radiance measurements and associated derived emissivity values for a first data-cube pixel in a simulation of a nucleus using the rock only emissivity spectrum from Figure 6-27. Note that many of the emissivity values are well above 1.

This becomes apparent when attempting to derive emissivities from the measured spectra using the simplified retrieval method in Section 6.3.5.2 (which just uses the brightness temperature in the broadband channel). I use measurements from a simulation of the nucleus using the emissivity of the rock-only mixture from Figure 6-27. Since the average emissivity across the broadband channel is well below the average emissivity in most of the channels, the brightness temperature is thus also lower, and therefore when emissivity is retrieved via Equation 6.22, many channels show an emissivity substantially larger than 1.

The radiance values measured, and emissivity retrieved for one such pixel is shown in Figure 6-33. One of the advantages of the simulation is retaining the true data to compare with the measured data, as well as simulated measurements errors. In this case the brightness temperature measured through the broadband channel was 184 K compared to the true surface temperature of ~220 K. As such the emissivities are distorted. Shape comparison can be done by comparing between normalised emissivities, as is shown for this data in Figure 6-34. The distortion observed, with overestimates at low wavelength, and underestimates at longer wavelengths, are a product of the incorrect temperature, with the Planck function that has been removed by the retrieval shifted to the right. The error in shape, here exhibited most clearly by a factor of two error on the last channel, would not permit accurate compositional analysis.




Figure 6-34: The normalised emissivity measurements plotted alongside the normalised emissivity spectrum assigned to the surface in the simulation.

I have devised an algorithmic method for improving this retrieval process, explained in the following section. This model is under further development external to this thesis.

6.7.2 Proposed Temperature-Emissivity Technique

A core issue with the retrieval process is its inherent degeneracy. From any pixel the instrument produces 9 measurements, a calibrated radiance in each of its filter channels. However, the radiances are defined by more than 9 parameters, most clearly there is the emissivity across each of the 9 channels, but also temperature. This is before we account for further aggravating factors such as a non-uniform temperature across the pixel and consider how to account for uncertainty in measurements.

I have proposed a method to simplify this regime, under certain assumptions, which reduces degeneracy in the system and thus permits the use of optimisation to derive surface parameters, given a set of TIRI radiance measurements.

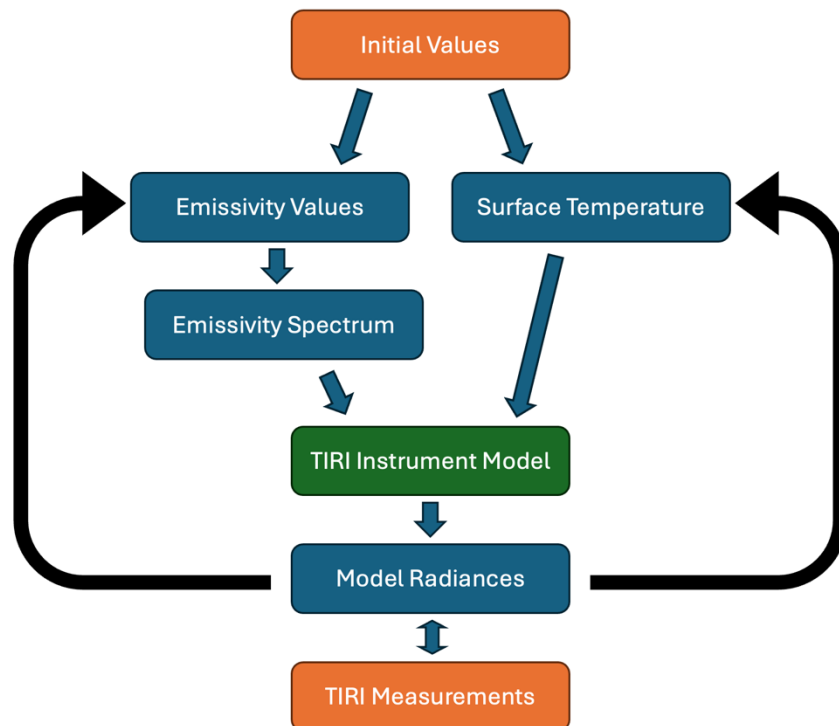


Figure 6-35: A schematic detailing the structure of the optimisable TIRI retrieval algorithm. Initial values for surface parameters are fed into the model, from which it derives radiance measurements which are compared to the provided TIRI measurements. The surface parameters are then iteratively tuned to minimise the difference between the model's output radiances and the measured TIRI radiances.

This technique uses the TIRI measurements, along with an initial guess at the surface parameters, to derive the most suitable surface parameters (structural overview in

Figure 6-35). The problem of degeneracy is simplified through the handling of the emissivities. The model considers 8 emissivity values (for the compositional filters) and then draws an emissivity spectrum between them. There is an assumption that the surface has a single temperature. A TIRI instrument model is defined which takes the emissivity spectrum and the surface temperature, and calculates model radiances in each filter channel via:

$$Radianc_{e_{filter}} = \int \epsilon(\lambda) * B(T, \lambda) * \tau_{filter}(\lambda) d\lambda \quad (6.23)$$

This as a whole structure can be considered a function with 9 inputs (8 emissivities and temperature) and 9 outputs (radiances in each channel). Since this is no longer degenerate, it should be possible to constrain the set of surface parameters which produce the closest match with the measured radiance values. An inverse method such as this should avoid the systematic error observed in Figure 6-34 since increasing the surface temperature will lower short wavelength emissivities and increase long wavelength ones, reducing the overall error. The inclusion of measurement uncertainty does reintroduce degeneracy, and I have outlined a research pathway to establish the best method to manage this uncertainty using methods such as Bayesian inference and regularisation.

Using this optimisation-based approach for such a non-linear function as the instrument model creates the opportunity for substantial further complexity, using different optimisation algorithms, feature structures and initial assumptions. To verify the method, I created this algorithm in Python utilising standard optimisation algorithms from the scipy library and using the outputs of the previous retrieval algorithm as my initial parameter values. The results of testing this on a 1-dimensional, uniform temperature surface are displayed in Figure 6-36. The accuracy of retrieved results is greatly improved in shape, and the true parameter values have been retrieved including temperatures, without requiring normalisation.




Figure 6-36: The retrieved emissivity values for a set of simulated TIRI measurements produced on a flat, constant temperature surface.

To further examine the efficacy of this method, I applied the same retrieval technique to the radiance values shown in Figure 6-33. These values are taken from my simulation and thus are the measurements from a pixel observing many facets with varied temperatures, and likely including slight pointing error in the colocation meaning each filter is not viewing the same set of facets. This is the most challenging case to produce retrievals for. These results are reproduced in Figure 6-36. The algorithm has retrieved accurate emissivity values for most filters and shows a marked improvement over the results in Figure 6-34 since the shape is captured without normalisation. The errors in the

longer wavelengths may be a product of fitting a single temperature to data generated by a temperature spread, a limitation in the optimisation algorithm or an outcome of alignment error.

I think that this shows a potential in my technique for generating surface parameter retrievals, and with further investigation I hope to see a further improvement in performance in comparison with other retrieval techniques. The TIRI simulation can generate a useful testing ground, where measurements can be compared to the underlying geometry, temperature and surface properties viewed by that pixel to understand where the algorithms struggle.




Figure 6-37: The emissivities retrieved through the optimisation method, plotted alongside the underlying rock emissivity used for this simulation. The measurements were produced for a pixel observing a range of temperatures.

6.8 Simulation Conclusions

The TIRI simulation tool has been successfully developed and is capable of generating data-cube measurements for in-depth analysis by the Comet Interceptor science community. The tool has been written up in technical documents, and the code will

be made available for use by the MIRMIS team. This tool offers flexibility within a broad parameter space, allowing for the investigation of diverse comet topologies, temperature distributions, and emissivity profiles. Additionally, it facilitates the examination of various flyby scenarios and the performance evaluation of different instrument operational methods, which is critical for optimizing mission design.

The tool's utility extends beyond individual simulations. It provides the capability to produce large volumes of well-characterised, simulated data, which can be used to refine and test surface parameter retrieval techniques from the calibrated radiance measurements. This is particularly important for verifying the accuracy of compositional, temperature, and emissivity retrievals under different mission conditions.

Preliminary investigations have validated the robustness of the simulation in handling geometric, radiometric, and data analysis components. The results demonstrate that retrieving the desired scientific data in data-cube format is feasible, and temperature uncertainty remains well within the required limits at this stage. However, it is important to note that the full accuracy of the final measurements may be impacted by factors not currently simulated, such as internal calibration uncertainties due to temperature variation in-flight. However, this simulation tool will provide a valuable method to probe the effect of each type of uncertainty and systematic error on the instruments produced measurements.

6.9 Future Work

The simulation described in this chapter establishes a robust framework for evaluating the performance of the TIRI instrument, as well as beginning the process of developing the instrument data-pipeline. It also highlights the need for further analysis to fully understand the instrument's capabilities and limitations under various mission

conditions. The results thus far demonstrate that the instrument can retrieve spectral data, identify compositional changes, and measure surface temperatures. Furthermore, the instrument has the potential to exceed its performance requirements, even under the most challenging scenarios, such as the high-velocity (70 km/s) flyby. Despite these successes, several areas require further investigation.

Scenarios with Low-Temperature Comets:

Thermophysical models suggest that some comets with a majority surficial ice content may exhibit average temperatures below 200 K. TIRI encounters such a cold comet, it will be crucial to adjust the operational parameters to maximise scientific return. For example, higher-frequency thermal broadband maps might be prioritised during the encounter. This simulation tool can now be used to explore how the instrument might perform under these conditions and identify optimal strategies for such encounters. Extreme flyby angles also warrant further study, as they could lead to observations dominated by unshaded regions of the nucleus, further challenging data retrieval.

Surface Structure and Composition Interrelation:

While the current simulation uses the Oxford thermophysical model to generate temperature maps from shape models, the version employed only accounts for a uniform surface composition. In real comets, surface features like ice would not only influence emissivity but also affect temperature due to differences in thermal conductivity and albedo. Future work should focus on generating simulations which incorporate compositional data earlier in the simulation pipeline, ensuring that the temperature maps reflect the underlying material structure. This will help ensure that nominal cases for operational decisions are relevant to real comet surfaces. Promising developments are being made on the Oxford thermophysical model, driven in part by accuracy requirements generated through this project.

Pointing and Data-Cube Co-Location Analysis:

Further work is required to refine the co-location of data-cube images utilising pointing data, rather than image analysis, to provide accuracy at a sub-pixel scale. This will involve accounting for the pointing mirror's motor positions, rotations, and spacecraft movements, based on onboard encoder data. The current work on the pointing vector analysis provides the foundation for this which can be extended to achieve precise alignment and stacking of the images into data-cubes.

Retrieval Algorithms for Surface Parameters:

The development and verification of retrieval algorithms for surface temperature and emissivity spectra from calibrated radiance measurements are critical next steps. While lessons can be drawn from similar instruments, such as the Diviner Radiometer, TIRI's retrieval techniques must be tailored to its specific mission requirements. I have developed a verification plan utilising retrieved LTM data to test temperature retrieval on real data. In addition, the simulation tool developed here can generate well-labelled simulated datasets, providing an ideal platform to test and refine retrieval algorithms. With the wide range of potential compositions expected in real cometary data, more complex retrieval challenges will arise, beyond the relatively simple ice-rock mixture used in these initial tests.

Expansion to NIR and MIR:

Finally, the work in this report provides a strong foundation for developing equivalent simulations for the NIR and MIR modules. These simulations will be critical for generating an end-to-end model for MIRMIS, capable of simulating the full set of calibrated measurements across all instruments for any given comet shape model and surface properties. This integrated model will serve as an invaluable tool for both target analysis and post-encounter data interpretation, ensuring the mission achieves its science goals.

7 Conclusions and Summary

7.1 Overview of Research

This thesis has outlined my contributions to the development of instrumentation at the University of Oxford for two planetary science missions: NASA's Lunar Trailblazer and ESA's Comet Interceptor. Through the design, testing and calibration of the LTM, and the development of operation plans and simulation tools for MIRMIS, the work presented here aims to contribute to wider goals of understanding the wider planetary science questions regarding the transport of volatiles through the Solar System.

The bulk of the research focused on the rigorous design and testing of LTM has seen efforts directed toward utilising ground-based calibration, laboratory measurements and computational modelling to characterise the performance of the instrument and ensure satisfaction of the performance requirements.

The instrument knowledge has been expanded on within the development of the TIRI instrument, informing the creation of simulation tools to accurately model the instrument behaviour through the encounter. These tools have supported the definition of TIRI operation plans throughout the flyby, and enabled evaluation of instrument performance based on these plans. This has been used to confirm the instrument's ability to meet its science requirements and has laid the groundwork for future refinement of operational strategies.

7.2 Key Contributions and Findings

LTM Calibration and Characterisation

The radiometric performance of LTM has been examined through the development of radiometric models, with thermal-radiometric models created to probe the effect of thermal instability and temperature inaccuracy on the instrument's calibrated data products. These models have been made more accurate through a thorough characterisation of the LTM optical pathway, determining the spectral transmission of each LTM filter, as well as mirror and detector performance through a full set of laboratory spectroscopic measurements in the mid- and far-infrared. The accuracy of the final instrument transmission functions has been verified through successful linearity testing on the instrument radiance response. The culmination of this modelling has been the production of radiance-temperature look-up tables which have been delivered to our mission partners at NASA to form a part of the initial spacecraft data-analysis pipeline.

Significant effort has been made to optimise the internal temperature measurements made by LTM, which are essential for providing accurate calibrated radiances from the lunar surface. The housekeeping electronics have been tested and improved to ensure their proper performance. The instrument's internal blackbody target has been thermally calibrated in the laboratory to provide temperature measurements to an accuracy of 88 mK. This is well within the requirements to calibrate the instrument to its requirement radiance accuracy of ± 1 K, which is crucial for the instrument's science goals. The calibration constants created here will be used every time LTM takes a measurement, enabling coincidental temperature measurements to enable thermal corrections to the HVM³ measurements.

MIRMIS and TIRI Operational Design

For the MIRMIS instrument, the TIRI module forms a central part of the scientific payload. Work contained in this thesis has played a key role in defining the TIRI's operations. This has been done by using internal and external thermophysical models to define likely limitations of target temperature and updated radiometric models to define integration times. These have then been incorporated into dynamic pointing modelling tools to define the instrument operations plan and identify how best to form TIRI's data-cubes. Motor and pointing modelling tools have been developed to further probe the operations plans, defining commanding requirements and predicting the instrument power draw. Following from the LTM housekeeping board investigation, the TIRI electronics has been further investigated to define its performance and guide the temperature measurement in the instrument from the CDHU.

The result of this modelling has been to verify that the instrument can perform in excess of requirements in terms of SNR, spatial resolution, power draw and radiometric accuracy for a worst-case nominal flyby scenario. This ensures TIRI's ability to measure surface temperature, emissivity, and compositional data during the Comet Interceptor's flyby, regardless of the parameters of the mission target. Furthermore, the modular nature of the analysis allows for rapid development of new operations plans in future, in response to a variety of parameter changes.

Simulation of TIRI Operations

The final component of the work contained within this Thesis is the development of an end-to-end simulation of TIRI's operations. This tool provides a way to relate engineering and operational design considerations to the instrument data product, allowing members of the MIRMIS science team with diverse comet science expertise to guide to the instrument development. The simulation has enabled the creation of

simulated data which reflect the entire spectrum of potential flyby parameters. Furthermore, in developing this simulation I have begun the process of developing the TIRI data-pipeline, which has created many avenues for future work to optimise the science retrieval from the TIRI dataset. Finally, this simulation provides a framework for similar developments for the other MIRMIS module, as well as other instruments across Comet Interceptor, to provide a full suite of tools to enable the ESA target selection team to evaluate the scientific potential of the alternative targets they will have to consider, enabling the mission to ensure it selects the most scientifically rewarding target.

7.3 Future Directions of Research

The research presented in this thesis establishes a strong foundation for further studies aimed at enhancing the understanding and performance of thermal instruments like LTM and MIRMIS. Several key areas warrant continued investigation.

Further developments in the LTM will be driven by in-orbit data once Lunar Trailblazer is launched. In particular, the thermal environment of the instrument will be assessed in space, emphasizing the significance of the thermal stability analysis discussed in Section 3.1.2. This will highlight the impact of temperature fluctuations on calibration accuracy. Additionally, the instrument's thermal environment can be rigorously tested using the flight spare in the Oxford calibration chamber to quantify calibration uncertainties. Precision improvements in internal temperature measurements, facilitated by updated software, will be evaluated, and the thermal-radiometric model can be employed to refine calibration accuracy. Furthermore, once surface data is returned from LTM, attention will turn to using the characterisation of the instrument, such as transmission functions, to optimise retrievals of surface parameters.

For MIRMIS, the operations plan will continue to evolve based on results from engineering model tests, incorporating updated information regarding motor speeds, power consumption, thermal performance, and filter measurements. Enhanced thermophysical modelling will help better define the operational requirements of the instrument. Testing the operations plan on these models will provide critical insights into the instrument's performance, for example defining the pointing accuracy achieved by the motor throughout simulated scans.

The MIRMIS simulation may be developed as described in Section 6.9. In addition, the pointing analysis underlying the model may be used to define a colocation algorithm capable of aligning images within data-cubes at a sub-pixel level, based on motor position data. Such an algorithm would be invaluable for post-flyby data analysis. Furthermore, significant work remains to develop and refine methods for retrieving surface parameters. The simulation tool described in this thesis is uniquely positioned to aid in this research, as it can generate high-resolution, labelled data, describing the underlying surface which is viewed within each detector pixel.

TIRI has already benefited from lessons learned during LTM's development, such as modifications to the housekeeping board that were carried over between instruments. This cross-instrument learning will continue, with improved ground-based calibration for TIRI resulting from LTM's test campaign. As LTM progresses through its mission cycle, being tested in space, producing measurements and being analysed, these lessons will continue to inform and enhance MIRMIS's operational design.

8 Bibliography

A'Hearn, M. F., Belton, M. J., Delamere, W. A., Kissel, J., Klaasen, K. P., McFadden, L. A., . . .

Collins, S. (2005). Deep Impact: Excavating Comet Tempel 1. *Science, Vol 310*, 258-264.

Altwegg, K., Balsiger, H., Bar-Nun, A., Berthelier, J., Bieler, A., Bochslers, P., . . . LeRoy, L.

(2014). 67P/Churyumov-Gerasimenko, a Jupiter family comet with a high D/H ratio. *Science, Vol 347*, I6620.

Arnold, J. (1979). Ice in the lunar polar regions. *Journal of Geophysical Research: Solid Earth, Vol 87*, 5659-5668.

Baker, V. (2006). Geomorphological Evidence for Water on Mars. *Elements Vol 2*, 139-143.

Balsiger, H., Altwegg, K., & Geiss, J. (1995). D/H and 18O/16O ratio in the hydronium ion and in neutral water from in situ ion measurements in comet Halley. *Journal of Geophysical Research: Space Physics, Vol 100*, 5827-5834.

Bandfield, J. L., Poston, M. J., Klima, R. L., & Edwards, C. S. (2018). Widespread distribution of OH/H₂O on the lunar surface inferred from spectral data. *Nature Geoscience, Vol 11*, 173-177.

Bates, H., Donaldson Hanna, K., King, A., Bowles, N., & Russell, S. (2021). A Spectral Investigation of Aqueously and Thermally Altered CM, CM-An, and CY Chondrites Under Simulated Asteroid Conditions for Comparison With OSIRIS-REx and Hayabusa2 Observations. *JGR Planets, Vol 126*, e2021JE006827.

Bockelée-Morvan, D., Biver, N., Swinyard, B., de Val-Borro, M., Crovisier, J., Hartogh, P., . . .

Davis, G. (2012). Herschel measurements of the D/H and 16O/18O ratios in water

Bibliography

- in the Oort-cloud comet C/2009 P1 (Garradd). *Astronomy and Astrophysics, Vol 544*, L15.
- Bockelée-Morvan, D., Gautier, D., Lis, D., Young, K., Keene, J., Phillips, T., . . . Wootten, A. (1998). Deuterated Water in Comet C/1996 B2 (Hyakutake) and Its Implications for the Origin of Comets. *Icarus, Vol 133*, 147-162.
- Bowles, E., Näsilä, A., Kohout, T., Villanueva, G., Howe, C., Jones, G. H., . . . Penttila, A. (2021). THE MODULAR INFRARED MOLECULES AND ICES SENSOR FOR ESA'S COMET INTERCEPTOR F-CLASS MISSION. *51st Lunar and Planetary Science Conference* (p. 2097). Houston: Lunar and Planetary Science Conference.
- Bowles, N., Elhmann, B., Klima, R., Blaney, D., Calcutt, S., Dickson, J., . . . T. (2020). The Lunar Thermal Mapper Instrument for the Lunar Trailblazer Mission. *51st Lunar and Planetary Science Conference* (p. 1380). Houston: Lunar and Planetary Science Conference.
- Cameron, A. G., & Ward, W. R. (1976). The Origin of the Moon. *Abstracts of the Lunar and Planetary Science Conference, Vol 7* (p. 120). Houston: Lunar and Planetary Science Conference.
- Chyba, C. (1987). The cometary contribution to the oceans of primitive Earth. *Nature 330*, 632-635.
- Ciesla, F., & Lauretta, D. (2005). Radial migration and dehydration of phyllosilicates in the solar nebula. *Earth and Planetary Science Letters, Vol 321*, 1-8.
- Clarendon Laboratory. (2005). *Calibration of a Sensor for Temperature Measurement: Calibration Number 10248*. Oxford: Clarendon Laboratory.

Bibliography

- Clark, P., Malphrus, B., Reuter, D., Hurford, T., MacDowall, R., Petro, N., . . . Cha. (2016). Broadband infrared compact high-resolution exploration spectrometer: lunar volatile dynamics for the lunar ice cube mission. *SPIE Optics and Photonics*. San Diego, California.
- Colaprete, A., Schultz, P., Heldmann, J., Wooden, D., Shirley, M., Ennico, K., . . . Sollitt, L. (2010). Detection of water in the LCROSS Ejecta Plume. *Science, Vol 330*, 463-468.
- Davidsson, B., Gutierrez, P., & Rickman, H. (2007). Nucleus properties of Comet 9P/Tempel 1 estimated from non-gravitational force modeling. *Icarus, Vol 191*, 547-561.
- Donahue, T., Hoffman, J., Hodges, R., & Watson, J. (1982). Venus Was Wet: A Measurement of the Ratio of Deuterium to Hydrogen. *Science, Vol 216*, 630-633.
- Donaldson Hanna, K. L., Pieters, C. M., Thomas, I., Bowles, N. E., & Greenhagen, B. T. (2014). Characterization of Apollo Bulk Soil Samples Under Simulated Lunar Conditions. *Lunar and Planetary Science Conference*, (p. 2345). Houston.
- Donaldson Hanna, K., Bowles, N., Warren, T., Hamilton, V., Schrader, D., McCoy, T., . . . Lauretta, D. (2021). Spectral Characterization of Bennu Analogs Using PASCALE: A New Experimental Set-Up for Simulating the Near-Surface Conditions of Airless Bodies. *JGR Planets, Vol 176*, e2020JE006624.
- Drake, M. J. (2005). Origin of water in the terrestrial planets. *Meteoritics & Planetary Science, Vol 40*, 519-527.
- Drake, M. J., & Righter, K. (2002). Determining the composition of the earth. *Nature 416*, 39-44.
- Eberhardt, P., Krankowsky, D., Schulte, W., Dolder, U., Lammerzahl, P., Berthelier, J., . . . Illiano, J. (1986). ON THE CO AND N2 ABUNDANCE IN COMET HALLEY. *In ESA*

Bibliography

- Proceedings of the 20th ESLAB Symposium on the Exploration of Halley's Comet.*
ESA.
- Eberhardt, P., Reber, M., Krankowsky, D., & Hodges, R. (1995). The D/H and 180/160 ratios in water from comet P/Halley. *Astronomy and Astrophysics*, 302, 301-316.
- Ehlmann, B. L., Klima, R. L., Bennett, C. L., Blaney, D., Bowles, N., Calcutt, S., . . . Wood, J. (2021). LUNAR TRAILBLAZER: A PIONEERING SMALLSAT FOR LUNAR WATER AND LUNAR GEOLOGY. *Lunar and Planetary Science Conference*, (p. 1325).
Houston.
- Elphic, R. C., Lawrence, D. J., Feldman, W. C., Barraclough, B. L., Maurice, S., Binder, A. B., & Lucey, P. G. (2000). Lunar rare earth element distribution and ramifications for FeO and TiO₂: Lunar Prospector neutron spectrometer observations. *Journal of Geophysical Research; Planets*, Vol 105, 20,333-20,345.
- Emerich, C. G. (1987). Temperature of the nucleus of Comet Halley. *Proceedings of the International Symposium on the Diversity and Similarity of Comets.* ESA.
- ESA. (2012, November 27). *THE NUCLEUS OF COMET HALLEY*. Retrieved from ESA Science: <https://sci.esa.int/web/giotto/-/51183-the-nucleus-of-comet-halley>
- ESA. (2019, September). *DEUTERIUM-TO-HYDROGEN RATIO IN THE SOLAR SYSTEM*. Retrieved from European Spcae Agency: <https://sci.esa.int/web/rosetta/-/55118-deuterium-to-hydrogen-ratio-in-the-solar-system>
- Evans, R. (2023). *Performance and Calibration Report*. Oxford: University of Oxford.
- Feldman, W., Prettyman, T., Maurice, S., Plaut, J., Bish, D., Vaniman, D., . . . Metzger, A. (2004). Global distribution of near-surface hydrogen on Mars. *Journal of Geophysical Research: Planets*, Vol 109, E09006.

Bibliography

- Fornasier, S., Hasselmann, P. H., Barucci, M. A., Feller, C., Besse, S., Leyrat, C., . . . Agarwal, J. (2015). Spectrophotometric properties of the nucleus of comet 67P/Churyumov-Gerasimenko from the OSIRIS instrument onboard the ROSETTA spacecraft. *Astronomy and Astrophysics, Vol 583*, A30.
- Friedman, B., Saal, A., Hauri, H., van Orman, J., & Rutherford, M. (2009). The Volatile Content of the Apollo 15 Picritic Glasses. *40th Lunar and Planetary Conference* (p. 2444). Houston: Lunar and Planetary Conference.
- Friedman, I., Gleason, J., & Hardcastle, K. (1970). Water, hydrogen, deuterium, carbon and C13 content of selected lunar material. *Proceedings of the Apollo 11 Lunar Science Conference* (pp. 1103-1109). Houston: Geochimica et Cosmochimica Acta Supplement, Volume 1.
- Genda, H., & Ikoma, M. (2008). Origin of the ocean on the Earth: Early evolution of water D/H in a hydrogen-rich atmosphere. *Icarus, Vol 194*, 42-52.
- Gibson, E., & Moore, G. (1973). Volatile-Rich Lunar Soil: Evidence of Possible Cometary Impact. *Science, Vol 179*, 69-71.
- Gladstone, G., Hurley, D., Retherford, K., Feldman, P., Pryor, W., Chaufray, J., . . . Stern, S. (2010). LRO-LAMP Observations of the LCROSS Impact Plume. *Science, Vol 330*, 472-476.
- Greenhagen, B., & Paige, D. (2006). Mapping Lunar Surface Petrology Using the Mid-Infrared Emissivity Maximum with the LRO Diviner Radiometer. *37th Annual Lunar and Planetary Science Conference* (p. 2406). Houston: Lunar and Planetary Science Conference.

Bibliography

- Greenhagen, B., & Paige, D. (2009). Overview of the 2009 LRO Diviner Lunar Radiometer Compositional Investigation. *40th Lunar and Planetary Science Conference* (p. 2255). Houston: Lunar and Planetary Science Conference.
- Greenhagen, B., Lucey, P., Wyatt, M., Glotch, T., Allen, C., Arnold, J., . . . Paige, D. (2010). Global Silicate Mineralogy of the Moon from the Diviner Lunar Radiometer. *Science, Vol 329*, 1507-1509.
- Guilbert-Lepoutre, A. (2012). Survival of Amorphous Water Ice on Centaurs. *The Astronomical Journal, Vol 144*, A97.
- Guilbert-Lepoutre, A., Besse, S., Mousis, O., Ali-Dib, M., Höfner, S., Koschny, D., & Hager, P. (2015). On the Evolution of Comets. *Space Science Review, Vol 179*, 271-296.
- Guilbert-Lepoutre, A., Lasue, J., Federico, C., Coradini, A., Orosei, R., & Rosenberg, E. (2011). New 3D thermal evolution model for icy bodies application to trans-Neptunian objects. *Astronomy and Astrophysics, Vol 529*, A71.
- Halley, E. (1705). *Synopsis of the Astronomy of Comets*. London: John Senex.
- Hallis, L. J. (2017). D/H ratios of the inner Solar System. *Phil. Trans. R. Soc, A375*.
- Hansen, B. (2009). Formation of the Terrestrial Planets from a Narrow Annulus. *The Astrophysical Journal, Vol 703*, 1131-1140.
- Hartmann, W., & Davis, D. (1975). Satellite-Sized Planetesimals and Lunar Origin. *Icarus, Vol 24*, 504-515.
- Hartogh, P., Dariusz, L., Bockelée-Morvan, D., Val-Borro, M., Biver, N., Kuipers, M., . . . Blake, G. (2011). Ocean-like water in the Jupiter-family comet 103P/Hartley 2. *Nature 478*, 218-220.

Bibliography

- Haskin, L., & Warren, P. (1991). Lunar Chemistry. In *Lunar Sourcebook*. Cambridge University Press.
- Hayne, P. O., Ingersoll, A. P., & Paige, D. A. (2014). *New Approaches to Lunar Ice Detection and Mapping*. Keck Institute for Space Studies.
- Honeywell. (2015). *Thin Film Platinum RTDs, HRTS series*.
- Honniball, C., Lucey, P., Shenoy, S., Orlando, T., Hibbitts, C., Hurley, D., & Farrell, W. (2020). Molecular water detected on the sunlit Moon by Sofia. *Nature Astronomy* 5, 121-127.
- huber. (2014). *ministat 125*. Retrieved from huber-online: https://www.huber-online.com/daisy_website_files/datenblaetter/en/2014.0011.01.PDF
- Huggins, W. (1866). *On the Results of Spectrum Analysis Applied to the Heavenly Bodies: A Discourse Delivered at Nottingham, Before the British Association*. Nottingham: Ladd, W.
- INO. (2024, 10 08). *INO*. Retrieved from INO website: https://inostorage.blob.core.windows.net/media/1234/datsheet-384x288-fpa_2019.pdf
- intersil. (2020). *ISL70218SEH Datasheet*. Retrieved from Renesas: <https://www.renesas.com/en/document/dst/isl70218seh-isl70218srh-datasheet?srsltid=AfmBOoqkZjZC4XjvZFPo05YSrBjdlFaGleUd7VsfLDCbQdHf9B1aWKhd>
- Keller, H., Arpigny, C., Barbieri, C., Bonnet, R., Cazes, S., Coradini, M., . . . Wilhelm, K. (1986). First Halley Multicolour Camera imaging results from Giotto. *Nature* 321, 320-326.

Bibliography

- Keysight Technologies. (2020). *34970A Data Acquisition/ Switch Unit Family technical Overview*.
- Klima, R., & Petro, N. (2017). Remotely distinguishing and mapping endogenic water on the Moon. *Phil. Trans. R. Soc*, A375.
- Kozlova, E., & Lazarev, E. (2010). Crater Cabeus as Possible Cold Trap for Volatiles Near South Pole of the Moon. *41st Lunar and Planetary Science Conference* (p. 1533). Houston: Lunar and Planetary Science Conference.
- Lara, L. M., Lowry, S., Vincent, J.-B., Gutiérrez, P. J., Rožek, A., La Forgia, F., . . . Bertaux. (2015). Large-scale dust jets in the coma of 67P/Churyumov-Gerasimenko as seen by the OSIRIS instrument onboard Rosetta. *Astronomy and Astrophysics, Vol 583*, A9.
- Le Roy, L., Altwegg, K., Balsiger, H., Berthelier, J.-J., Bieler, A., Briois, C., . . . Tzou, C.-Y. (2015). Inventory of the volatiles on comet 67P/Churyumov-Gerasimenko from Rosetta/ROSINA. *Astronomy and Astrophysics, Vol 583*, A1.
- Li, S., & Milliken, R. (2016). An empirical thermal correction model for Moon Mineralogy Mapper data constrained by laboratory spectra and Diviner temperatures. *Journal of Geophysical Research: Planets, Vol 121*, 2081-2107.
- Litvak, M. L., Mitrofanov, I. G., Sanin, A., Malakhov, A., Boynton, W. V., Chin, G., . . . Zube. (2012). Global maps of lunar neutron fluxes from the LEND instrument. *Journal of Geophysical Research, Vol 117*, E12.
- Liu, J., Liu, B., Ren, X., Li, C., Shu, R., Guo, L., . . . Ouyang, Z. (2022). Evidence of water on the lunar surface from Chang'E-5 in-situ spectra and returned samples. *Nature Communications 13*, 3119.

Bibliography

LTM Requirements - Internal Document. (2020). *D-104876*. Planetary Instruments Group.

Lucey, P. G. (2017). Water Absorption at 6 Microns: A New Tool for Remote Measurements of Lunar Surface Water Abundance and Variation. *Lunar and Planetary Science XLVIII*, (p. 1314). Houston.

Lyster, D., Howett, C., & Penn, J. (2024). Predicting Surface Temperatures on Airless Bodies: An Open-Source Python Tool. *EPSC2024-1121*. Berlin: Copernicus Meetings.

Mahanti, P., Robinson, M., Wagner, R., Estes, N., & Humm, D. (2023). First Look, First Results - Comparing Secondary Illumination at Lunar Permanently Shadowed Regions From the First Shadowcam Image and Topography Based Simulation. *2023 IEEE International Geoscience and Remote Sensing Symposium* (pp. 4162-4165). Pasadena: IGARSS.

Marshall, D., Groussin, O., Vincent, J., Brouet, Y., Kappel, D., Arnold, G., . . . Thomas, N. (2018). Thermal inertia and roughness of the nucleus of comet 67P/Churyumov-Gerasimenko from MIRO and VIRTIS observations. *Astronomy and Astrophysics, Vol 616*, A122.

Meier, R., & Owen, T. (1999). Cometary Deuterium. *Space Science Reviews*, 33-43.

Meier, R., Owen, T., Matthews, H., Jewitt, D., Bocklee-Morvan, D., Biver, N., . . . Gautier, D. (1998). A determination of the HDO/H₂O ratio in comet C/1995 O1 (Hale-Bopp). *Science, Vol 279*, 482-484.

Mendis, D. (1988). A Postencounter View of Comets. *Annual Review of Astronomy and Astrophysics, Vol 26*, 11-49.

Bibliography

- Milliken, R., & Mustard, J. (2005). Quantifying absolute water content of minerals using near-infrared reflectance spectroscopy. *Journal of Geophysical Research*, Vol 110, E12001.
- Milliken, R., & Mustard, J. (2007). Estimating the water content of hydrated minerals using reflectance spectroscopy. *Icarus*, Vol 189, 574-588.
- Morbidelli, A., Chambers, J., Lunine, J. I., Petit, J. M., Robert, F., Valsecchi, G. B., & Cyr, K. E. (2000). Source regions and timescales for the delivery of water to the Earth. *Meteoritics & Planetary Science*, Vol 35, 1309-1320.
- Morbidelli, A., Chambers, J., Lunine, J., Petit, J., Robert, F., Valsecchi, G., & Cyr, G. (2000). Source regions and timescales for the delivery of water to the Earth. *Meteoritics and Planetary Science*, Vol 35, 1309-1320.
- O'Brien, D., Izidoro, A., Jacobson, S., Raymond, S., & Rubie, D. (2018). The Delivery of Water During Terrestrial Planet Formation. *Space Science Reviews*, Vol 214, A47.
- O'Brien, D., Walsh, J., Morbidelli, A., Raymond, S., & Mandell, A. (2014). Water delivery and giant impacts in the 'Grand Tack' scenario. *Icarus*, Vol 239, 74-84.
- Paige, D., Siegler, M., Zhang, J., Hayne, P., Foote, E., Bennett, K., . . . Ba. (2010). Diviner Lunar Radiometer Observations of Cold Traps in the Moon's South Polar Region. *Science*, Vol 330, 479-482.
- Pater, D., & Lissauer. (2015). *Planetary Sciences*. Cambridge University Press.
- Pieters, C., Goswami, J., Clark, R., Annadurai, M., Boardman, J., Buratti, B., . . . Petro, N. (2009). Character and Spatial Distribution of OH/H₂O on the Surface of the Moon Seen by M3 on Chandrayaan-1. *Science*, Vol 326, 568-572.
- RAL Space. (2021). *LTM Thermal Analysis Report*. RAL Space.

Bibliography

- Raymond, S., Quinn, T., & Lunine, J. (2004). Making other earths: dynamical simulations of terrestrial planet formation and water delivery. *Icarus*, *Vol 168*, 1-17.
- Reinhard, R. (1986). The Giotto encounter with comet Halley. *Nature* *321*, 313-318.
- Ringwood, A. E. (1989). Flaws in the giant impact hypothesis of lunar origin. *Earth and Planetary Science Letters*, *Vol 95*, 208-214.
- Rubie, D., Frost, D., Mann, U., Asahara, Y., Nimmo, F., Tsuno, K., . . . Palme, H. (2011). Heterogeneous accretion, composition and core–mantle differentiation of the Earth. *Earth and Planetary Science Letters*, *Vol 301*, 31-42.
- Saal, A., Hauri, E., Cascio, M., Orman, J., Rutherford, M., & Cooper, R. (2008). Volatile content of lunar volcanic glasses and the presence of water in the Moon's interior. *Nature* *454*, 192-195.
- Salisbury, J., Basu, A., & Fischer, E. (1997). Thermal Infrared Spectra of Lunar Soils. *Icarus*, *Vol 130*, 125-139.
- Sanin, A. B., Boynton, W., Evans, L., Harshman, K., Kozyrev, A., Litvak, M., . . . Vostrukhin, A. (2007). LUNAR EXPLORATION NEUTRON DETECTOR (LEND) FOR NASA LUNAR RECONNAISSANCE ORBITER. *Lunar and Planetary Science XXXVIII* (p. 1648). Houston: Lunar and Planetary Institute.
- Schultz, P., Hermalyn, B., Colaprete, A., Ennico, K., Shirley, M., & Marshal, W. (2010). The LCROSS Cratering Experiment. *Science*, *Vol 330*, 468-472.
- Shirley, K., Warren, T., Faggi, S., Villanueva, G., Protopapa, S., Hanna, K. D., . . . Thirumangalath, S. (2022). Optimizing Filter Bandpass Selection for the Thermal Infrared Imager on ESA's Comet Interceptor Mission. *EPSC2022-486*. Granada: Copernicus Meetings.

Bibliography

Snodgrass, C., & Jones, G. (2019). The European Space Agency's Comet Interceptor lies in wait. *Nature Communications* 10, A5418.

Spencer, J., Lebofsky, L., & Sykes, M. (1989). Systematic biases in radiometric diameter determinations. *Icarus*, Vol 78, 337-354.

Tai Udovicic, C. J., Bandfield, J. L., Ghent, R. R., Farrand, W. H., Edwards, C. S., & North Arizona University. (2020). Testing Lunar Near-Infrared Corrections with Diviner Observations. *51st Lunar and Planetary Science Conference* (p. 2632). Houston: Lunar and Planetary Science Conference.

Thomas, N., Sierks, H., Barbieri, C., Lamy, P. L., Rodrigo, R., Rickman, H., . . . Deppo, D. (2015). The morphological diversity of comet 67P/Churyumov-Gerasimenko. *Science*, Vol 347, 16220.

Thompson, D., Green, R., Ehlmann, B., Klima, R., Pieters, C., Blaney, D., . . . Sci. (2020). THE HIGH-RESOLUTION VOLATILES AND MINERALS MOON MAPPER (HVM3) ON THE LUNAR TRAILBLAZER MISSION. *51st Lunar and Planetary Science Conference* (p. 2052). Houston: Lunar and Planetary Science Conference.

Tissot, J., Robert, P., Durand, A., Tinnes, S., Bercier, E., & Crastes, A. (2013). Status of Uncooled Infrared Detector Technology as ULIS, France. *Defence Science Journal*, Vol 63, 545-549.

Trumbo, S., Brown, M., & Butler, B. (2018). ALMA Thermal Observations of Europa. *The Astronomical Journal*, Vol 156, 156-161.

Unas, Teti, I, P., Merenre, & II, P. (2400 BCE). *Pyramid Texts*. Samuel A. B. Mercer.

Bibliography

- Villaneuva, G., Mumma, M., Bonev, B., DiSanti, M., Gibb, E., Böhnhardt, H., & Lippi, M. (2009). A Sensitive Search for Deuterated Water in Comet 8p/Tuttle. *The Astrophysical Journal Letters*, Vol 690, L5.
- Villanueva, G., Liuzzi, G., Faggi, S., Protopapa, S., Kofman, V., Fauchez, T., . . . Mandell, A. (2022). *Fundamentals of the Planetary Spectrum Generator*. Maryland: <https://psg.gsfc.nasa.gov/help.php>.
- Warren, T., King, O., Bowles, N., Sefton-Nash, E., Fisackerly, R., & Trautner, R. (2019). The Oxford 3D Thermophysical Model with application to the Lunar PROSPECT Mission. *51st Lunar and Planetary Science Conference* (p. 2004). Houston: Lunar and Planetary Science Conference.
- Watson, K., Murray, B., & Brown, H. (1961). The Behaviour of Volatiles on the Lunar Surface. *Journal of Geophysical Research*, Vol 66, 3033-3045.
- Whipple, F. (1949). A Comet Model. I. The Acceleration of Comet Encke. *Astrophysical Journal*, Vol 111, 375-394.
- Whipple, F. (1950). A Comet Model. II. Physical Relations for Comets and Meteors. *Astrophysical Journal*, Vol 113, 464-474.
- Williams, J. P., Paige, D. A., Greenhagen, B. T., & Sefton-Nah, E. (2017). The global surface temperatures of the moon as measured by the diviner lunar radiometer experiment. *Icarus*, Vol 283, 300-325.
- Wood, B., Wade, J., & Kilburn, M. (2008). Core formation and the oxidation state of the Earth: Additional constraints from Nb, V and Cr partitioning. *Geochimica et Cosmochimica Acta*, Vol 72, 1415-1426.

Bibliography

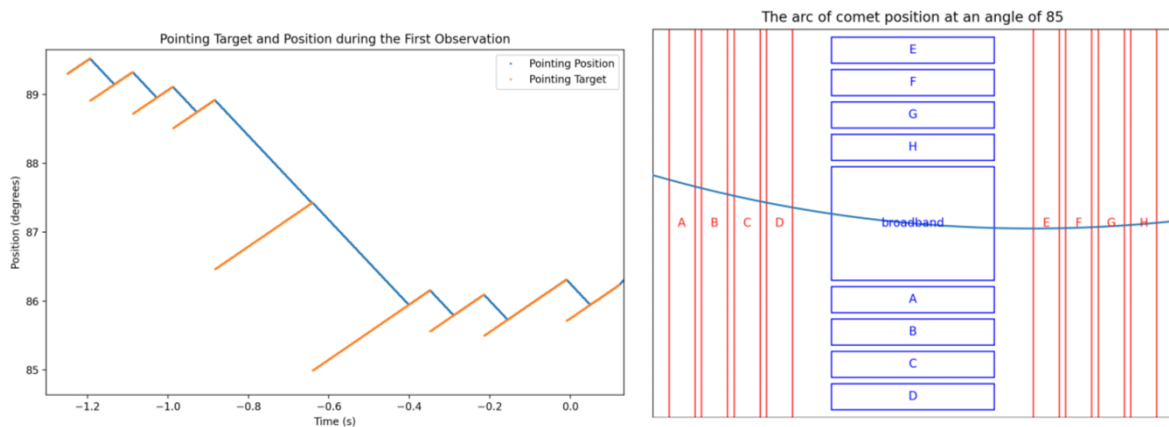
Wright, V., Morzfeld, M., & Manga, M. (2024). Liquid water in the Martian mid-crust.

Proceedings of the National Academy of Sciences, Vol 121, No. 35.

9 Appendices

9.1 TIRI Data-Cube Details


This appendix lists the operation angle position, filter order, and movement pattern for each TIRI data-cube.



Parameter	Value
Start Position	85°
Filters Used	Vertical
Filter Order	A - H


Figure 9-1: The movement plan for the first data-cube, the arc of control at the beginning of that scan and the set of parameters used to define that movement pattern.

Appendices



Parameter	Value
Start Position	106°
Filters Used	Vertical
Filter Order	A - H

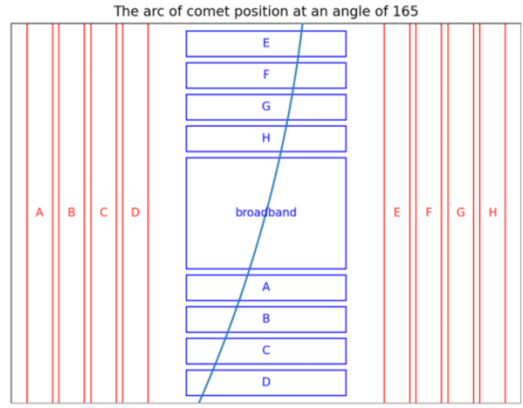

Figure 9-3: The movement plan for the second data-cube, the arc of control at the beginning of that scan and the set of parameters used to define that movement pattern.



Parameter	Value
Start Position	123°
Filters Used	Vertical
Filter Order	H - A



Figure 9-2: The movement plan for the third data-cube, the arc of control at the beginning of that scan and the set of parameters used to define that movement pattern.

Appendices



Parameter	Value
Start Position	165°
Filters Used	Horizontal
Filter Order	D – A, H - E

Figure 9-5: The movement plan for the fourth data-cube, the arc of control at the beginning of that scan and the set of parameters used to define that movement pattern.



Parameter	Value
Start Position	172°
Filters Used	Horizontal
Filter Order	D – A, H - E

Figure 9-4: The movement plan for the fifth data-cube, the arc of control at the beginning of that scan and the set of parameters used to define that movement pattern.

9.2 PSG Settings for Representative Spectra

Here are the compositional values that were used through PSG to generate the rock and rock-ice mixture emissivity spectra used in Chapter 7.



Figure 9-6: The minerals and abundances used in PSG to generate the emission spectrum for the rock and ice mixture material surface.



Figure 9-7: The minerals and abundances used in PSG to generate the emission spectrum for the purely rock material surface.