

Measuring the infrared sky background from the high Antarctic plateau

Author:

Freeman, Matthew

Publication Date:

2018

DOI:

<https://doi.org/10.26190/unsworks/20985>

License:

<https://creativecommons.org/licenses/by-nc-nd/3.0/au/>

Link to license to see what you are allowed to do with this resource.

Downloaded from <http://hdl.handle.net/1959.4/61302> in <https://unsworks.unsw.edu.au> on 2024-04-24

Measuring the infrared sky background from the high Antarctic plateau

Matthew Stanley Ryan Freeman



A thesis in fulfilment of the requirements for the degree of

Doctor of Philosophy

School of Physics

Faculty of Science

University of New South Wales

December 2018



FOR OFFICE USE ONLY Date of completion of requirements for Award:

Declaration

Originality Statement

‘I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or substantial proportions of material which have been accepted for the award of any other degree or diploma at UNSW or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by others, with whom I have worked at UNSW or elsewhere, is explicitly acknowledged in the thesis. I also declare that the intellectual content of this thesis is the product of my own work, except to the extent that assistance from others in the project’s design and conception or in style, presentation and linguistic expression is acknowledged.’

Signed:

Date:

17/08/2018

Copyright Statement

‘I hereby grant the University of New South Wales or its agents the right to archive and to make available my thesis or dissertation in whole or part in the University libraries in all forms of media, now or here after known, subject to the provisions of the Copyright Act 1968. I retain all proprietary rights, such as patent rights. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation. I also authorise University Microfilms to use the 350 word abstract of my thesis in Dissertation Abstract International (this is applicable to doctoral theses only). I have either used no substantial portions of copyright material in my thesis or I have obtained permission to use copyright material; where permission has not been granted I have applied/will apply for a partial restriction of the digital copy of my thesis or dissertation.’

Signed: _____

Date: 17/08/2018

Authenticity Statement

‘I certify that the Library deposit digital copy is a direct equivalent of the final officially approved version of my thesis. No emendation of content has occurred and if there are any minor variations in formatting, they are the result of the conversion to digital format.’

Signed: _____

Date: 17/08/2018

INCLUSION OF PUBLICATIONS STATEMENT

UNSW is supportive of candidates publishing their research results during their candidature as detailed in the UNSW Thesis Examination Procedure.

Publications can be used in their thesis in lieu of a Chapter if:

- The student contributed greater than 50% of the content in the publication and is the “primary author”, ie. the student was responsible primarily for the planning, execution and preparation of the work for publication
- The student has approval to include the publication in their thesis in lieu of a Chapter from their supervisor and Postgraduate Coordinator.
- The publication is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in the thesis

Please indicate whether this thesis contains published material or not.



This thesis contains no publications, either published or submitted for publication



Some of the work described in this thesis has been published and it has been documented in the relevant Chapters with acknowledgement



This thesis has publications (either published or submitted for publication) incorporated into it in lieu of a chapter and the details are presented below

CANDIDATE'S DECLARATION

I declare that:

- I have complied with the Thesis Examination Procedure
- where I have used a publication in lieu of a Chapter, the listed publication(s) below meet(s) the requirements to be included in the thesis.

Name	Signature	Date (dd/mm/yy)
Matthew Freeman		17/08/18

Postgraduate Coordinator's Declaration

I declare that:

- the information below is accurate
- where listed publication(s) have been used in lieu of Chapter(s), their use complies with the Thesis Examination Procedure
- the minimum requirements for the format of the thesis have been met.

PGC's Name	PGC's Signature	Date (dd/mm/yy)

Acknowledgements

I would like to thank:

My supervisor Michael Ashley, for the amazing opportunities he's given me. I never thought I'd get the chance to travel to Antarctica, and do science in the coldest place on Earth.

My co-supervisor Michael Burton, for his productive discussions and suggestions on my project every week.

My MSc supervisor Phil Yock, for suggesting I apply to UNSW. I'd say it turned out pretty well.

My wonderful partner Sally, for all her support, and also for submitting her thesis first so I knew what to expect.

All the friends I made at UNSW: Shaila, Domenico, Burcu, Anita, Nigel, Graeme, Anant, John, Daniel, Tiger, and everyone else.

And finally, Mum and Dad, for always encouraging me in everything I've done.

Abstract

The high plateau of Antarctica is known to have the best observing conditions on Earth for infrared astronomy. The extremely low levels of water vapour there open up new atmospheric windows, making ground-based infrared and terahertz observations possible. In particular, at the site Ridge A in Antarctica the atmospheric water vapour level regularly drops below 0.1 mm, lower than any other site on Earth. The low temperatures, stable atmosphere, and clear skies also contribute to the superb observing conditions at this site.

For most near-infrared observations the limiting factor is the background emission, which primarily comes from the atmosphere. To conclusively demonstrate the quality of a site, it is necessary to take measurements of the background in-situ. This has been performed for the first time away from the South Pole with the installation of the Near Infrared Sky Monitor (NISM) at Ridge A in 2015. NISM has taken measurements of the spectral radiance of the atmosphere in the Kdark band, at 2.4 microns. These measurements can determine how much infrared noise will be present in astronomical observations, and hence the quality of Ridge A as an observatory site.

In this thesis, the data sent back from NISM are analysed to determine the zenith spectral radiance of the atmosphere at Ridge A. This analysis posed a challenge, as a number of problems occurred with the instrument, including problems with a motor used to set the beam altitude, and ice forming on the black body calibrator. As NISM is completely inaccessible for most of the year, these problems had to be diagnosed remotely, and dealt with in software.

Ultimately, it was determined that there may be a long-wavelength leak in NISM's Kdark filter, and that it would be necessary to recover the instrument in order to characterise the leak. The team was not able to reach the site in the 2017/2018 season due to the poor condition of the skiway. A definitive measurement of the infrared sky background will have to wait until NISM is recovered in a future mission, when inspection of the Kdark filter will allow the data collected so far to be calibrated.

Contents

1	Introduction	14
1.1	Choosing an observatory site	16
1.2	Site comparisons	17
1.3	Antarctic astronomy	20
1.3.1	Boundary layer	22
1.3.2	Wind speed	22
1.3.3	Cloud cover	23
1.3.4	Aurora	23
1.3.5	Aerosols	24
1.3.6	Airglow	24
1.3.7	Precipitable water vapour	24
1.3.8	Atmospheric thermal emission	28
1.3.9	Surface temperature	28
1.3.10	Free atmosphere seeing	29
1.4	Antarctic infrared site testing	30
1.5	Terahertz astronomy	34
1.5.1	Terahertz Surveys	36
2	Ridge A	39
2.1	The Ridge A environment	39
2.1.1	Wind speed	43
2.1.2	Temperature	43
2.1.3	Water vapour	44
2.1.4	Access	45

2.2	HEAT	48
2.3	PLATO-R	51
2.3.1	Power	52
2.3.2	Control	53
3	NISM	55
3.1	Detector	58
3.2	Optical configuration	59
3.3	Beam	62
3.4	Optical modelling	62
3.5	Passband	68
3.6	Cooler	68
3.7	Motor	70
3.8	Electronics	71
3.8.1	ADC board	71
3.8.2	Microcontroller board	74
3.9	Winterisation	76
3.10	Black body	76
3.11	Installation	79
4	Atmospheric modelling	82
4.1	VSTAR	83
4.1.1	DISORT	84
4.1.2	HITRAN	85
4.2	Ridge A atmospheric model	86
4.3	VSTAR results	90
4.3.1	NISM's passband	91
4.3.2	VSTAR prediction	94
4.4	Contribution of different gases	95
5	Data analysis	98
5.1	Binary packets	98

5.2	Baseline subtraction	104
5.3	Motor slip analysis	105
5.3.1	Lowess smoothing	110
5.3.2	The Fibonacci search algorithm	111
5.4	Correcting angle of elevation	115
5.4.1	Aligning to previous block	117
5.4.2	Aligning to a reference block	117
5.4.3	Fitting a skydip	119
5.4.4	Lmfit	121
5.4.5	Katana cluster	122
6	Calibration	126
6.1	Calibration using the Milky Way	127
6.1.1	Atmospheric attenuation	131
6.1.2	Subtracting the Galactic flux	133
6.2	Calibration using the black body	136
7	Black body investigation	139
7.1	Dynamic time warping	143
7.1.1	Accumulated cost matrix	143
7.1.2	Optimal warping path	146
7.2	Shape dynamic time warping	147
7.2.1	Custom step sizes	148
7.2.2	Custom weights	149
7.3	Results of DTW	150
7.3.1	Investigating the jump	153
8	Aurora	157
8.1	NOAA satellites	157
8.2	Aurora measurements	158
8.3	Comparison	161

9	Red leak	163
9.1	Filter	163
9.2	Recovery	166
10	Lessons	168
10.1	Instrument housing	168
10.2	Optics	169
10.3	Detector	169
10.4	Monitoring	170
10.5	Logistics	170
11	Conclusion	171

Authors Note

Investigations into the astronomical potential of Antarctica have been going on for a number of years at UNSW before work on this thesis began.

The establishment of the observatory at Ridge A described in Chapter 2 was the work of Craig Kulesa, David Lesser, and Abram Young from the University of Arizona, along with Michael Ashley, John Storey, Nic Bingham, Daniel Luong-Van, Colin Bonner, Jon Everett, Geoff Sims, Yael Augarten, Campbell McLaren, and Luke Bycroft from UNSW.

The development of the instrument NISM described in Chapter 3 was also completed before this thesis. The electronics were designed by Michael Ashley and Nic Bingham, and the optical design was based on the original NISM by John Storey, Max Boccas, Jon Lawrence, and Michael Ashley. The AAO was largely responsible for the mechanical design, with contributions from UNSW.

The analysis in Chapter 4 onwards is the key part of this thesis, and is the work of the author.

Chapter 1

Introduction

Astronomy is a science that is constantly advancing. New observatories are continually under development, and new advances in telescope design allow us to see wider and deeper, revealing more about the universe we live in.

There are two factors that determine how good the observations through a telescope will be. One is the quality of the telescope itself, including the diameter and smoothness of the primary mirror, the design and construction of the telescope, and the type of detectors installed.

The other factor is the site at which the telescope is located. There are a number of aspects of a site whose levels and stability all contribute to the quality of observations, including the elevation, cloud cover, turbulence, light pollution, and levels of absorbing molecules such as water vapour, CO_2 , and CH_4 .

In order to take the best observations possible, you want the best telescope at the best site. Once telescopes became sufficiently advanced, choosing a suitable site became of equal importance to having a high quality telescope. All the effort designing and building a telescope would be wasted by placing it at a poor location.

In this thesis, we investigate the high plateau of Antarctica as a possible site for future observatories. As will be explained, the conditions at the Antarctic plateau make it the best location on Earth for almost all kinds of astronomy, and it could become the home of the world's major observatories in the future. As such, it is very important to take in-situ measurements of the site's quality, to conclusively demonstrate its merits.

The core subject of this thesis is the processing of data from NISM, a small instrument which has measured the infrared sky background in Antarctica. These measurements should in principle have been able to demonstrate the quality of the site, but have proved a challenge to acquire and analyse.

This thesis is organised as follows:

Chapter 1 gives a review of the field of Antarctic astronomy, covering the environmental conditions, infrared site testing experiments, and the existing terahertz surveys.

Chapter 2 explores the site Ridge A, on the high plateau of Antarctica. This site is likely to have the best, or close to the best, conditions in Antarctica, and is home to several exploratory instruments, including HEAT, PLATO-R, and NISM.

Chapter 3 describes the instrument NISM (the Near Infrared Sky Monitor), which was installed at Ridge A to measure the infrared sky background.

Chapter 4 models the atmosphere at Ridge A using the program VSTAR, which was written by Jeremy Bailey, to produce a prediction of the measurements we would expect to see with NISM.

Chapter 5 covers the initial processing of the data packets returned by NISM, and the corrections that were made to solve some of the problems encountered.

Chapter 6 attempts to correctly calibrate the NISM data, using the on-board black body reference, as well as the brightness of the Milky Way.

Chapter 7 investigates NISM's black body reference source, which gave anomalous brightness readings.

Chapter 8 looks at the Aurora Australis as a possible source of the infrared emission measured by NISM.

Finally, Chapter 9 explores a possible problem with the instrument: that the infrared filter may have a long-wavelength leak, causing extra light to leak through and increasing the measured brightness.

1.1. CHOOSING AN OBSERVATORY SITE

Table 1.1: Altitudes of various astronomical sites.

Altitude	Site
2228 m	Mt Kosciusko; the highest mountain in Australia
2400 m	San Pedro de Atacama
2635 m	VLT, Chile
2800 m	Hale Pohaku, Maunakea
2835 m	South Pole
3500 m	South Pole barometric altitude (50% of people develop symptoms of Acute Mountain Sickness)
3656 m	Lhasa, Tibet
4087 m	Dome A
4145 m	Keck Telescope, Maunakea
4500 m	Ali (Shiquanhe) , Tibet
4600 m	Dome A barometric altitude
5059 m	ALMA Observatory, Chile
5100 m	Ali Observatory, main site
5150 m	Everest North Base Camp
5310 m	Highest point at Ali Observatory
5612 m	Cerro Chajnantor, Chile

1.1 Choosing an observatory site

When choosing a site for an observatory, elevation is one of the first considerations. Since the atmosphere is responsible for the majority of obstructions, a straightforward answer is to build the telescope somewhere high up, to get above as much of the atmosphere as possible. This is why virtually all observatories are built at the tops of mountains. This also gives the telescope a good view of the sky, unobstructed by nearby terrain. Table 1.1 lists the heights of various astronomical sites for comparison.

Cloud cover is another obvious factor. Optical telescopes can not see through cloud, so a site with nice clear skies is required. These include particularly dry

areas, and high elevation sites that are above the cloud layer.

Light pollution is a factor that is of increasing importance as populated areas encroach on observatories. The light produced by cities scatters off the atmosphere, significantly increasing the background sky brightness. As a result, modern telescopes need to be built far away from city lights. A number of major historical observatories, such as Mt Wilson Observatory near Pasadena and the Royal Observatory in Greenwich, were built close to populated areas, and are no longer prominent observing sites due to the increased light pollution.

Once major considerations like elevation and cloud cover are checked off, the next step is to look for a site with weak atmospheric turbulence. Turbulence causes the light from stars to refract by differing amounts as it travels through the atmosphere, and is responsible for the prominent twinkling of stars. This scintillation blurs the image of a star, limiting the resolution of observations. To avoid this, astronomers look for sites with stable atmospheres. These are commonly found near the coast on the windward side, as the lower layers of air flowing over the flat ocean form a relatively laminar flow with weak turbulence.

Another factor that can have an effect is the level of water vapour in the atmosphere. Water vapour is transparent at visible wavelengths, but can absorb and emit light at infrared and submillimetre wavelengths, making astronomy at these wavelengths difficult to impossible. This is one of the factors that has made visible and radio the two major types of astronomy: other wavelengths are blocked by the atmosphere.

So, astronomers are looking for high locations with clear skies to build telescopes. And as the capabilities of telescopes improve, it becomes more important to find an excellent location to take advantage of them.

1.2 Site comparisons

The search for observing sites has led astronomers to some extreme locations on the planet. They have built airborne telescopes mounted on balloons and aircraft, and have launched telescopes into space to avoid the atmosphere entirely.

1.2. SITE COMPARISONS

Two of the premier ground-based astronomical sites are Maunakea in Hawaii, and the Atacama desert in Chile. The Maunakea observatories are located at the peak of the tallest volcano on the Big Island of Hawaii, at an elevation of 4145 m. At this elevation the observatories are above $\sim 40\%$ of the atmosphere. Being located in the Pacific Ocean, the winds passing over the mountain tops of Hawaii are relatively laminar, so there is little low level turbulence to disrupt observations. Although the high altitude jet stream does still have an effect.

The Atacama desert in northern Chile is a very dry location due to the dehumidifying effect of the cold ocean currents off the west coast. It is home to a number of telescopes built in the Andes, some at elevations over 5000 m. Its proximity to the ocean provides a relatively laminar airflow over the site, reducing the low level atmospheric turbulence.

These two sites are generally regarded as the premier observing sites on Earth. With one in the northern hemisphere and one in the southern hemisphere, they provide full-sky coverage, and they are home to some of the largest optical telescopes, including Keck and the VLT.

Balloon-borne and aircraft-borne telescopes can observe from much higher up in the atmosphere, so have improved seeing conditions over ground based telescopes. However, they do come with some downsides: they are very expensive to launch and operate, they are limited in the size of telescope they can carry, and they can only observe for a limited duration before landing again. The airborne SOFIA has an operating cost of $\sim \$100\,000$ per hour (Sloan, 2016). Since they are not firmly mounted on the ground, there is increased motion and vibration of the telescope to deal with.

Space-based telescopes are the ultimate step in avoiding the atmosphere. By operating in orbit, space telescopes have almost no atmospheric issues to worry about—no turbulence, no water vapour, just a clear view of the sky, with resolution limited only by the size of the mirror. Telescopes in low Earth orbit do have to deal with the albedo of the Earth, which can increase the background when observing in certain directions, but can still achieve superb seeing with suitable pointing.

However, like airborne telescopes, this excellent seeing comes with trade offs. The size of the mirror is limited, both by the weight, and by what can be folded to fit in a cylindrical rocket. Additionally, space telescopes are enormously expensive. Simply launching the mission can cost \$20 000 per kg, in addition to the design and construction of the telescope itself. The infrared Herschel Space Observatory cost €1100 million over its three year lifetime¹, which gives it an operating cost of \sim €50 000 per hour.

Space telescopes are not easily accessible once launched, and may be impossible to service in some orbits. The *Hubble Space Telescope* is in low Earth orbit, so could be reached by the Space Shuttle. However its replacement, the James Webb telescope, will be at L2, far beyond the reach of any realistic repair mission. If there is any problem with the telescope on launch, it will not be possible to fix it.

Even without any errors, space missions have a limited lifetime. Once components degrade or the cryogen for the detector runs out, the mission is over. Failing reaction wheels led to the end of the *Kepler* space telescope's primary mission. Additionally, the technology in a space telescope can quickly become obsolete. Once the design phase of the telescope is complete it is not possible to make changes, so components cannot be upgraded as technology improves.

Space telescopes provide fantastic capabilities, but at a high cost. For a similar budget, a larger ground-based telescope could be operated for a much longer duration. Such a telescope could perform larger surveys, collect more light, and be upgraded with the latest detectors as technology improves. Atmospheric turbulence will have some effect, but adaptive optics can mitigate some of this. Modern adaptive optics systems can achieve diffraction limited resolution in the infrared for some fraction of the field of view.

For the above reasons, ground-based telescopes have an important role to play in astronomy. So it is in the interest of astronomers to find the best sites in the world to construct these telescopes. There are observatories located all over the Earth, but there is still one place that is very promising, but as yet has few optical/infrared telescopes—Antarctica.

¹esamultimedia.esa.int/docs/herschel/Herschel-Factsheet.pdf

The centre of Antarctica has excellent conditions for astronomy, but due to the remoteness and extreme conditions it has been relatively underutilised. It is a difficult site to operate in, but it is far cheaper to transport equipment to Antarctica than it is to launch it to space. The cost to operate the HEAT telescope is approximately \$100 per hour (Craig Kulea, private communication). This is remarkably cheap, compared to the $\sim\text{€}50\,000$ for Herschel and the $\sim\$100\,000$ for SOFIA. It is also possible to access the site regularly, to perform maintenance and install new equipment.

The high Antarctic plateau has some of the best conditions on the planet, and it could become the premier site for infrared and terahertz/sub-millimetre observatories. Such observatories would have a large advantage over traditional temperate latitude sites for projects such as mapping of interstellar molecular clouds in the terahertz, and observations of star forming regions and distant galaxies in the infrared.

1.3 Antarctic astronomy

Antarctica is the coldest, highest, and driest continent. While the weather around the coast can be stormy, in the centre of the continent the winds are low and the atmosphere is dry and stable. These conditions make Antarctica potentially an excellent site for ground based astronomy. Figure 1.1 shows the locations of some existing Antarctic stations. The primary locations for astronomy are those in the interior of the continent, i.e., the South Pole, Concordia (Dome C), Kunlun (Dome A), and Dome Fuji (Dome F).

Antarctic astronomy is a relatively recent field. The first optical analysis of the astronomical conditions at the interior of the continent was only carried out in 1964, with the installation of a small telescope at the South Pole (Indermuehle et al., 2005). This analysis suggested that the seeing would be excellent, which lead to the first optical astronomical research being carried out in 1979, observing solar oscillations from the South Pole (Grec et al., 1980).

Other interior sites have only become available more recently. Concordia

Station was established at Dome C in 1996, and was upgraded to allow winter-overs in 2005 (Dolci et al., 2016). The Dome A site at the centre of the high plateau is even more recent, with the first humans visiting the site in 2005, and Kunlun station being established in 2009 (Hu et al., 2014). The last site to be reached was Ridge A, with the first humans and instruments only arriving in 2012 (Kulesa et al., 2013a; Ashley, 2013).

Saunders et al. (2009) reported predictions and measurements of the conditions at a number of sites in Antarctica, specifically the South Pole, Dome A (Kunlun), Dome B, Dome C (Concordia), Dome F (Fuji), as well as Ridge A and Ridge B. They looked at a range of environmental factors relevant to astronomy, including boundary layer thickness, wind speed, cloud cover, aurorae, airglow, precipitable water vapour (PWV), atmospheric thermal emission, surface temperature and free atmosphere seeing.

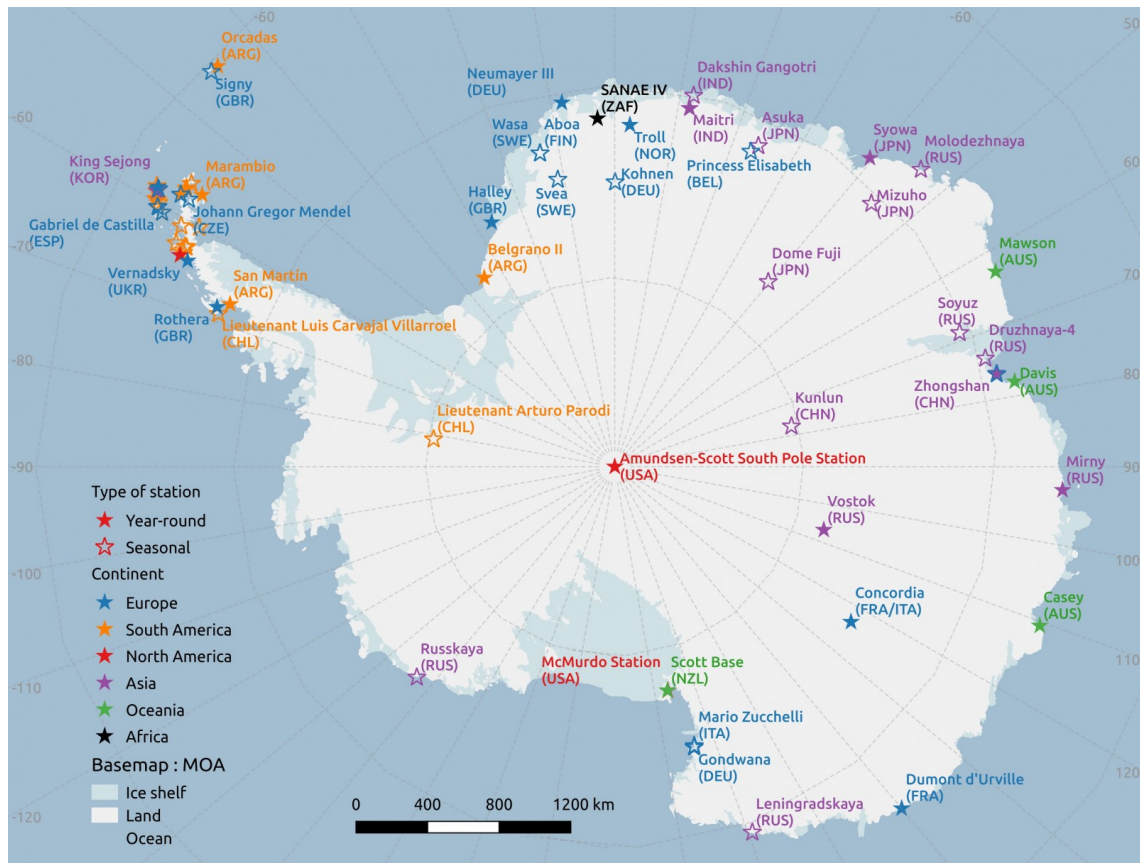


Figure 1.1: A map of Antarctica showing permanent and seasonal bases [Credit: Sophie Berger].²

²<https://blogs.egu.eu/divisions/cr/2016/08/>

1.3.1 Boundary layer

As there is no jet stream over Antarctica, the turbulence in the atmosphere is almost entirely contained to a turbulent boundary layer close to the surface. Swain & Gallée (2006b) modelled the thickness of the boundary layer, defining it as the height where the boundary seeing contribution is 0.1 arcseconds or better for 50% of the time. They predicted the boundary layer thickness to be 19 m at Dome F, 22 m at Dome A, < 24 m at Ridge B, and 28 m at Dome C. Bonner et al. (2010) measured the thickness of the boundary layer at Dome A using SODAR, and found it to be 14 m, while Lawrence et al. (2004b) measured the boundary layer at Dome C to be less than 30 m. The ground level seeing in Antarctica is worse than at temperate sites, but the thin boundary layer offers the possibility of raising a telescope up on a tower. There it would achieve free air seeing, which has been measured at ~ 0.23 arcseconds at Dome C (Lawrence et al., 2004b).

Additionally, having the turbulence confined to a thin boundary layer is the ideal situation for implementing adaptive optics. The lower altitude turbulence has a lower spatial frequency, which gives a larger isoplanatic angle and a longer coherence time. These factors allow a low-order adaptive optics system to correct for the seeing over a large angle (Marks, 2002).

1.3.2 Wind speed

Swain & Gallée (2006a) presented further simulation results, including wind speeds. They found wind speeds to be below 5 ms^{-1} at all the sites. The wind speed results showed a similar pattern to the boundary layer results, with the lowest speeds at Dome F, increasing speeds at Dome A and Ridge B, then Dome C, and then highest at the South Pole.

These match up well with field measurements. The average wind speed at the South Pole was measured to be 4.2 ms^{-1} (Hudson & Brandt, 2005), and the wind speed at Dome C was measured to be 2.9 ms^{-1} (Aristidi et al., 2005). A weather tower was only installed recently at Dome A, allowing the wind speed at a height of 2 m to be measured as 1.5 ms^{-1} Hu et al. (2014).

1.3.3 Cloud cover

Cloud cover data in Antarctica have come from passive satellite measurements. Unfortunately, these are notoriously unreliable, as satellites have difficulty distinguishing between a cloud layer and the ice surface. Saunders et al. (2009) includes data from the Aqua MODerate-resolution Imaging Spectroradiometer (MODIS) and the Ice, Cloud, and Elevation Satellite Geoscience Laser Altimeter System (GLAS). They found that all the sites had similar winter fractional cloud cover of ~ 0.2 . There was only a 4% difference in the cloud cover between the best site (Dome C) and the worst (Dome F).

1.3.4 Aurora

Auroral activity tends to form a ring at 60° to 75° geomagnetic latitude. The south magnetic pole was at 64.28° South, 136.59° East in 2015, placing it off the coast of Antarctica near Dumont d’Urville station in Figure 1.1. Hence, there is greater activity at the South Pole and Dome F, and lower activity at Dome A and Dome C. Measurements at the South Pole show aurora contributing $\sim 22 \text{ mag arcsec}^{-2}$ in the B and V bands (Dempsey et al., 2005). Saunders et al. (2009) extrapolated from this, finding that Dome A, Dome B, and Dome C are likely to have auroral contributions of $\sim 23 \text{ mag arcsec}^{-2}$ in the B and V bands. They state that this sky brightness is worse than at good temperate sites by a factor of 2 in B , and by 20-30% in V .

Measurements of the aurora were taken at Dome A in 2009, using the Nigel and Gattini instruments (Sims et al., 2012). The median sky brightnesses in the B , V , and R bands were found to be 22.9, 23.4, and $23.0 \text{ mag arcsec}^{-2}$ respectively, which is in good agreement with the predictions from Saunders et al. (2009).

For comparison, at Maunakea the zodiacal light at the zenith is $23.1\text{--}22.1 \text{ mag arcsec}^{-2}$ (Kenyon & Storey, 2006). Due to Maunakea’s position near the equator, its zenith is closer to the ecliptic, and hence the zodiacal light is brighter. So the optical zenith sky brightness at Dome A is always lower than at Maunakea, despite the aurora.

1.3.5 Aerosols

The concentration of aerosols over the high plateau is very low (Ashley, 2012). Measurements of the aerosol optical depth (AOD) at 500 nm wavelength over the South Pole find a median AOD of 0.015 (Tomasi et al., 2007). For comparison, the AOD(500 nm) over a populated city can be ~ 0.3 . The low value over Antarctica is a consequence of the fact that there is very little dust, pollen, or pollution. Some scattering that is observed is suspected to be due to fine ice crystals, also called diamond dust, with dimensions on the order of $\sim 10 \mu\text{m}$ (Hidas et al., 2000).

1.3.6 Airglow

Airglow is emission from excited molecules in the upper atmosphere. Air molecules are photoionized by sunlight during the day, and undergo chemical reactions that produce excited states which radiate at a range of optical wavelengths. Airglow is primarily composed of OI lines at 557.7 nm and 630 nm, NO₂ continuum from 500–650 nm and OH continuum from 700–2300 nm (Benn & Ellison, 1998; Oliva et al., 2013; Liu et al., 2008). Beyond 2300 nm ($2.3 \mu\text{m}$) airglow drops off rapidly. OI airglow is predicted to be strong in Antarctica, but can be filtered out by narrow band filters. OH airglow is also quite strong, but there is an “OH hole” that forms over Antarctica in the summer, which could allow observations in the *J* and *H* bands. Unfortunately, this is the period when many sites have continuous daylight. In order to take advantage of this feature, a telescope would need to be placed as far north as possible on the Antarctic Plateau. Dome C is the furthest north of the interior sites, having roughly 3.5 hours of usable darkness per day at the equinox (Saunders et al., 2009).

1.3.7 Precipitable water vapour

Precipitable Water Vapour, or PWV, is a measure of the total column of water vapour in the atmosphere. It gives the depth of water that would be created if all the water vapour in a column the atmosphere were to be condensed into a liquid, and is usually given in millimetres. Water vapour has strong absorption bands in




Figure 1.2: Predictions of the 2004 winter average precipitable water vapour levels across Antarctica. Taken from Swain & Gallée (2006a).

the infrared and sub-mm, so the level of water vapour in the atmosphere is one of the primary obstacles to IR/sub-mm astronomy.

Simulations done in 2006 (Swain & Gallée, 2006a) showed that Dome C and the South Pole have good PWV levels, Dome A and Dome F have even better PWV's, but the best PWV is found roughly 100 km from Dome A, at Ridge A. Here the predicted winter median PWV is 0.118 mm. See Table 1.2 and Figure 1.2 for more details.

Due to the importance of the precipitable water vapour level for terahertz astronomy, the instrument Pre-HEAT was deployed to Dome A in 2008 to make measurements of the PWV, in advance of the HEAT terahertz telescope (Yang et al., 2010). Pre-HEAT measured the atmospheric transmission at 661 GHz. This measurement was used to validate the transmission model shown in Figure 1.3.

The 661 GHz transmission data were combined with microwave soundings from




Figure 1.3: Comparison of the atmospheric transmission as a function of frequency at Dome A/Ridge A, Chajnantor, and Maunakea. Dome A has superior transmission, particularly at higher frequencies. Taken from Yang et al. (2010).

Table 1.2: PWV quantiles in μm for the South Pole, Dome C, Dome A, Ridge A, Dome B and Dome F, during the 2008 season. Measurements taken with the MHS sensor on the NOAA-18 satellite. Taken from (Saunders et al., 2009).

	SP	DC	DA	RA	DB	DF
Annual median	437	342	233	210	274	279
Winter median	324	235	141	118	163	163
Winter 25%	258	146	103	77	115	114
Winter 10%	203	113	71	45	83	90
Winter σ	133	122	65	64	67	98

the MHS instrument on the NOAA-18 satellite. Combining the two data sets allowed the transmission to be converted to a measurement of PWV (Saunders et al., 2009), shown in Table 1.2. The best winter quartile PWV at Dome A was found to be 0.1 mm, the lowest PWV measured from anywhere on the ground. For comparison, Chajnantor’s best quartile is 0.35 mm, and Maunakea’s is 1.0 mm (Yang et al., 2010).

Importantly, the low PWV at Dome A was also very stable, with the standard deviation of the transmission being 9%, which is equal to 0.07 mm PWV. The stability of the PWV can be even more critical than the absolute level. If the PWV is stable, it is possible to account for a higher level by integrating for a longer time. But an unstable PWV reduces the advantage from long integrations, and so prevents deep terahertz observations.

The combination of Pre-HEAT and NOAA-18 data predicted that the lowest PWV was to be found at Ridge A, where the best winter quartile PWV was 0.08 mm. This is at the centre of the contours in Figure 1.2. Due to this extremely low PWV, Ridge A was chosen to be the site of the terahertz telescope HEAT (Kulesa et al., 2013b)

With the installation of the HEAT telescope at Ridge A in 2012 (see Section 2.2), it has been possible to measure the PWV at Ridge A. These measurements confirmed the predictions, finding that the PWV was nearly always below 0.25 mm, and regularly dropped below 0.1 mm (Burton et al., 2015). With these

ground-based measurements taken, it has been possible to calibrate satellite measurements, so it is now possible to obtain PWV measurements from almost anywhere in Antarctica (Craig Kulesa, private communication).

1.3.8 Atmospheric thermal emission

Atmospheric thermal emission depends on the airmass, the temperature profile of the atmosphere, and the density of molecules that emit, such as H_2O , CO_2 , and CH_4 . The coldest air is found between the South Pole and Dome A. Saunders et al. (2009) used a line-by-line radiative transfer code to model the sky brightness in the infrared over Antarctica. They ranked the sites from best to worst as Ridge A, Dome A, Ridge B, Dome F, Dome C, and finally the South Pole. They found Ridge A to be ~ 3 times better than the South Pole in the K , L , and M band infrared thermal windows (2.0–2.4, 3.0–4.0, and 4.6–5.0 μm), and 1.5 times better in the optically thick bands outside these windows.

The South Pole is the only site where near-infrared measurements have been taken prior to the work described in this thesis (Ashley et al., 1996; Nguyen et al., 1996; Phillips et al., 1999). The sky brightness in the K_{dark} band was found to be $80 - 200 \mu\text{J arcsec}^{-2}$. There have been no measurements of the K_{dark} background at other sites, but since they are higher and drier than the south pole, Saunders et al. (2009) predicted that they would have darker skies.

1.3.9 Surface temperature

A lower surface temperature results in lower thermal emission from the telescope itself. Swain & Gallée (2006a) included surface temperatures in their models; they found the coldest region in Antarctica runs from Dome F to Dome A to Ridge B. There is another cold minimum at Dome C. But the temperatures at these sites are all within a few degrees of each other, so there is not expected to be much difference in the effects on any telescopes. The average winter temperature at these sites is around -70°C .

Hu et al. (2014) reported temperature measurements from Dome A, finding a year-round average of -54°C at a height of 2 m. This can be separated into a daily average of -35°C in January, an average of -60°C in April, and an average of -70°C in the middle of winter.

Satellite measurements have found pockets on the high plateau that reach -98.6°C during the winter, which is the coldest temperature ever recorded on Earth (Scambos et al., 2018). These are located along the ridge between Dome A and Dome F, as expected. However, since they are recorded by satellites they do not count as official weather records.

1.3.10 Free atmosphere seeing

Direct measurements of the free atmosphere seeing have been taken from the South Pole (Marks et al., 1999), Dome C (Lawrence et al., 2004b), and Dome F (Okita et al., 2013). These measurements are difficult to obtain, as they have to be taken above the boundary layer.

Marks et al. (1999) used balloon-borne microthermal probes to measure the mean free atmosphere seeing above the South Pole. They found a seeing of 0.37 ± 0.07 arcseconds.

To measure the free atmosphere seeing above Dome C, Lawrence et al. (2004a) installed the AASTINO module at Concordia Station. This module included a Multi-Aperture Scintillation Sensor (MASS), which observed the scintillation of a single star to calculate the seeing above 500 m, and a SODAR (Sonic Detection And Ranging) to measure the seeing from 30–500 m. The median free atmosphere seeing above Dome C was measured to be 0.23 arcseconds (Lawrence et al., 2004b)(note the value in the abstract of this paper is the mean, not the median).

Okita et al. (2013) took measurements at Dome F with a Differential Image Motion Monitor (DIMM) from the top of an 11 m tower, and found the best seeing to be between 0.2 and 0.3 arcseconds. They presumed that these measurements were from times when the boundary layer was below the 11 m height of the tower, and thus were measurements of the free atmosphere seeing.

Predicting the seeing using meteorological data is difficult, because the

turbulent layers that cause poor seeing are thinner than the vertical resolution of satellites. Saunders et al. (2009) used wind speed models to make rough predictions of the seeing at various sites. They found the best seeing to be at the South Pole and Ridge A, then Dome A, Dome F, Dome B, and finally Dome C had the worst. However, this disagrees with the data from Marks et al. (1999), which shows worse free-atmosphere seeing than has been measured at Dome A, Dome C, and Dome F.

1.4 Antarctic infrared site testing

Infrared astronomy unlocks many interesting observations of the universe. Infrared wavelengths can penetrate through dense interstellar dust clouds which would block visible light. This allows observations of stellar formation and galactic nuclei. Infrared astronomy is also well suited to observing the early universe, using the red-shifted light from distant galaxies.

However, infrared astronomy is difficult to do from the ground. The main limiting factor for infrared observations is the high amount of background light from airglow, atmospheric thermal emission, and zodiacal light (light emitted and scattered from dust grains in the solar system). Zodiacal light is pretty much unavoidable, but fortunately is faint, being measured at $6 \mu\text{Jy arcsec}^{-2}$ at the south ecliptic pole (Hauser et al., 1991). Airglow does vary in strength at different sites, but is still present all across the Earth. However, different sites on Earth can have significantly different levels of thermal emission.

Due to the superb conditions in Antarctica, it was quickly identified as a potential site for infrared telescopes. Early measurements in the millimetre and sub-millimetre at the South Pole indicated a very low level of water vapour. Harper (1989) used these measurements to predict the infrared background emission, and calculated that there would be 220 times less emission than at Maunakea in the $2.27\text{--}2.45 \mu\text{m}$ band. This leads to a telescope having a 15 times better signal-to-noise ratio. In very good conditions, Harper predicted that the background could approach the limit set by zodiacal light.

These promising predictions demanded further investigation. The first infrared

site test was performed at the South Pole in the winter of 1994 (Ashley et al., 1996; Phillips et al., 1999). They made measurements of the infrared sky background over a range of wavelengths from 1–5 μm , using the Infrared Photometer Spectrometer (IRPS), which had previously been an instrument on the Anglo-Australian Telescope. These measurements are shown in Figure 1.4.

IRPS found the darkest window to be the long end of the K_{dark} band, from 2.35 μm to 2.45 μm , where they observed a sky flux of 80–200 $\mu\text{Jy arcsec}^{-2}$ (the dip at the top right of Figure 1.4). This K_{dark} flux is 20–100 times lower compared to mid-latitude sites like Siding Spring or Maunakea. They also found that aurora have a negligible effect on the background in the K_{dark} window.

The reason for the low emission in the K_{dark} band is that it falls neatly in between the thermal emission and airglow spectra. Thermal emission is much higher at longer wavelengths, and airglow is higher at shorter wavelengths. As mentioned previously in section 1.3.6, OH airglow drops off around 2.3 μm . So in the K_{dark} band, both thermal emission and airglow have a comparable low effect.

Also in 1994, the SPIREX 60 cm telescope was installed at the South Pole (Nguyen et al., 1996). SPIREX was equipped with the GRIM spectrometer, which was sensitive over a range of 1–2.5 μm . They also made measurements of the zenith brightness in the K_{dark} window. They measured a brightness in the 2.29–2.43 μm wavelength range of $162 \pm 67 \mu\text{Jy arcsec}^{-2}$ during the winter.

The SPIREX telescope was later used to observe Polycyclic Aromatic Hydrocarbons (PAH's), which are tracers of photodissociation regions in molecular clouds (Rathborne et al., 2002), and to observe the NGC612 star forming region in the infrared (Storey et al., 2000).

Smith & Harper (1997) measured the mid-infrared sky brightness at the South Pole in the summer of 1997. They used the NIMPOL instrument to observe in the 8–13 μm N band, and found a sky brightness of 484 $\mu\text{Jy arcsec}^{-2}$. This is ~ 8 times better than Maunakea, but not as dark as the near-IR K_{dark} band. The near-IR is on the Wien side of the black body curve, so the lowered temperature at Antarctic sites produces a significantly lower near-IR emission.

These initial instruments had a design feature that limited their use: they were




Figure 1.4: Sky emission from the South Pole as a function of wavelength. The K_{dark} window is visible as a dip around $2.4 \mu\text{m}$. Taken from Phillips et al. (1999).

cooled with cryogenics. SPIREX was cooled to $\sim 75 \text{ K}$ with liquid nitrogen, IRPS was cooled to $\sim 65 \text{ K}$ by solid nitrogen, and NIMPOL was cooled to 18 K with liquid helium (Nguyen et al., 1996; Ashley et al., 1996; Smith & Harper, 1997). Cryogen boils away over time, so needs to be refilled periodically. This makes these instruments unsuitable for testing remote sites on the high plateau, where they would need to operate for long periods with no human contact.

Thus, a different cooling method was required. The next generation of site testing began with the installation of AASTO at the South Pole in 1997 (Storey et al., 1996; Storey et al., 2000). AASTO (Automated Astronomical Site Testing Observatory) was a module designed to house and power a number of site testing instruments, covering UV to sub-millimetre wavelengths. AASTO used miniature Stirling-cycle coolers to cool the instruments to their operating temperatures. These coolers have a small power budget and replace liquid cryogenics, so are suitable for use at remote field sites. AASTO generated power using a propane-fuelled catalytic oxidiser, which generated 2.5 kW of heat. A

thermoelectric generator then used this heat to generate 50 W of electrical power.

This independent power generation enabled AASTO to be used in remote locations on the high plateau, although it was set up at the South Pole for initial testing. The thermoelectric generator proved to have a notable failure mode: the thermosyphon could leak freon into the AASTO enclosure. This would then react with the catalytic oxidiser to produce hydrochloric and hydrofluoric acid, which caused “devastating corrosion to the electronics, computers, and exposed parts of the instruments” (Storey et al., 2000). Remote plateau observatories would obviously require a more reliable power supply.

AASTO contained the first incarnation of NISM, the Near Infrared Sky Monitor. A later version of this instrument became the one used in this thesis. NISM observed the sky spectral brightness in the longer end of the K_{dark} window, from 2.27–2.50 μm . In this window there are only four strong OH emission lines, and these are much weaker than the OH lines shorter than 2.3 μm . This instrument suffered from a $\sim 0.01\%$ long-wavelength leak from 2.5–5 μm , which had to be accounted for in the data analysis. The optical layout of the original NISM is shown in Figure 1.5. It uses a chopper wheel to chop between the sky and a black body reference.

Lawrence et al. (2002b) reported NISM measurements from the 2001 season. After rejecting data that were affected by cloud cover or the Moon passing through the beam, they measured the median sky spectral brightness to be 120 $\mu\text{Jy arcsec}^{-2}$, consistent with other observations. For comparison, they estimated the sky flux at Maunakea in the same band to be $\sim 1000\text{--}2000 \mu\text{Jy arcsec}^{-2}$.

The infrared sky flux at the centre of the plateau is expected to be even lower than 120 μJy , due to the lower temperatures and drier atmosphere. But, prior to this thesis, no infrared measurements have been taken there.

Dong et al. (2018) describes the development of a multi-band near infrared sky brightness monitor (MNISBM), to be installed at an Antarctic site at a future date, most likely Dome A. MNISBM has some similarities to NISM: it uses an InSb detector, a thermoacoustic cooler, a long-wavelength blocking filter, and a chopping mirror to measure the infrared background. It has filters to observe in

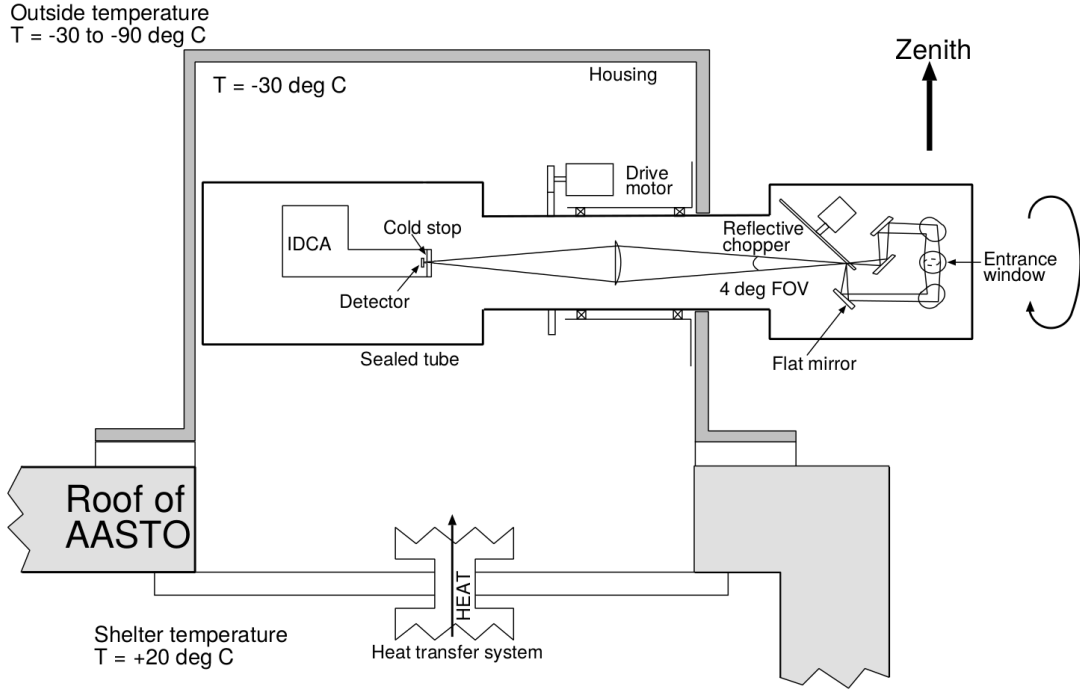


Figure 1.5: Optical layout of the original NISM instrument. (Figure created by modifying the image of MISM in Storey et al. (1999)).

the J , H , and K bands, at 1.15–1.35, 1.48–1.8, and 1.98–2.32 μm respectively. The K band filter of MNISBM covers shorter wavelengths than NISM’s K_{dark} filter, which extends to 2.5 μm .

1.5 Terahertz astronomy

Along with infrared, the other atmospheric window opened up by the superb conditions in Antarctica is in the terahertz/sub-millimetre band, covering the range $0.3\text{--}3\text{ THz} = 0.1\text{--}1\text{ mm}$. Observing in the terahertz portion of the spectrum is very difficult, for a number of reasons. The primary one is the opacity of the atmosphere: terahertz radiation is strongly absorbed by water vapour. The high levels of water vapour in our atmosphere make it almost completely opaque to terahertz radiation. Thus, terahertz astronomy is virtually impossible from almost all ground locations on Earth (Kulesa et al., 2013b).

The only way to get around this is to find extremely dry sites, and observe at high elevation to get above as much of the atmosphere as possible. Telescopes such

as JCMT and ALMA are located on top of high mountains in Hawaii and the Atacama desert in Chile respectively, some of the driest sites on Earth. The telescope SOFIA is mounted in an aircraft, to take it up above most of the atmosphere. Balloons offer another way to reach very high altitudes, as used in long duration balloon experiments like the Stratospheric THz Observatory (Walker et al., 2010). Finally, space based telescopes like *Herschel* are outside the atmosphere entirely, so have no problems with water vapour.

Terahertz radiation is also difficult to detect with standard detectors. Terahertz frequencies are generally too high for conventional electrical circuits, as used in radio telescopes. But the energy of terahertz radiation is too low to be picked up by CCD detectors, as used in optical and infrared telescopes. This places terahertz astronomy in an awkward middle ground.

There are two main methods for detecting terahertz radiation: bolometers and heterodyne systems. These come at the problem from techniques used at different ends of the electromagnetic spectrum.

Bolometers absorb the terahertz radiation falling on them and measure the tiny increase in heat in the detector. Bolometers have wide bandwidths, so have a higher sensitivity than heterodyne systems, making them suitable for detecting continuum emission from faint objects. However, this gives them less frequency resolution, making them unsuitable for collecting velocity information from emission lines. The SCUBA-2 instrument on JCMT is a bolometer.

Coming from the radio end of the spectrum are heterodyne detectors. These instruments generate a local signal at a frequency close to the frequency being observed. The observed signal is then combined with the local signal, and the ‘beat’ frequency measured. This heterodyne frequency is the difference between the two input frequencies, and is normally set to a few gigahertz, which is low enough to be digitised directly in electronic circuits. Heterodyne systems can detect the frequency of the signal, and hence calculate its velocity from the Doppler shift. ALMA uses heterodyne detectors.

Ground based sites such as ALMA and JCMT still have to deal with atmospheric water vapour, so can only observe in good conditions. Airborne

observatories like SOFIA and space observatories such as *Herschel* can all but ignore water vapour (SOFIA commonly reaches precipitable water vapour levels of $10\ \mu\text{m}$, (Young et al., 2012)), but are expensive and have limited observation time.

One of the primary uses of terahertz astronomy is to observe the interstellar medium (ISM). Dense molecular clouds in the ISM contain a large fraction of the total mass, and consist primarily of H_2 molecules (Bally, 1986). These molecules do not have any easily excited emission lines at low temperatures, so instead tracer molecules are used to detect molecular clouds.

The brightest emission line in a typical galaxy is the $[\text{CII}]$ line at $158\ \mu\text{m}$, which is in fact the primary cooling method for the diffuse ISM (Pineda et al., 2013). This can be seen in Figure 2 of De Bernardis et al. (2010), which shows the spectrum of the M82 galaxy in the submillimetre. The CII line (from C^+ ions) clearly stands out as the brightest on the logarithmic scale.

Given the CII line's location in the terahertz spectrum, it is difficult to observe from the ground, so terahertz surveys of the ISM have to go to some lengths to get their observations.

1.5.1 Terahertz Surveys

Some of the earliest terahertz measurements were taken in 1985 using NASA's airborne Lear Jet observatory (Stacey et al., 1983). This project observed the $158\ \mu\text{m}$ $[\text{CII}]$ forbidden transition line from ionized carbon. Again, we see that the primary concern when observing in the terahertz is trying to avoid the water vapour in the atmosphere. This experiment measured the CII flux from five lines of sight in the Galactic Plane by chopping between the galactic plane and the Moon, which has negligible $158\ \mu\text{m}$ emission. This CII emission line is the primary cooling method for the diffuse ISM (Pineda et al., 2013), so it is clearly an important measurement tool for studying how molecular clouds form.

One of the first submillimetre surveys of the entire sky was carried out with the COBE satellite, which operated from 1989 to 1993 (Bennett et al., 1994). This survey mapped the Galaxy using the FIRAS spectrometer, which covered wavelengths from $0.1\text{--}10\ \text{mm}$. This wavelength range allows the detection of

emission lines from CII, CI, NII and CO. However, this instrument had a beam size of 7 degrees, which gives the final survey a very low resolution. It also had a spectral resolution of 100 kms^{-1} , which is insufficient to completely resolve these spectral lines. They found that CII is the brightest emission line in the submillimetre, and that it comes primarily from the cold neutral medium.

Another way to get above the water vapour in the atmosphere is to use a high altitude balloon. The BICE experiment was one of these, using a balloon launched up to an altitude of 37 km to look at the centre of the Galaxy in [CII] (Nakagawa et al., 1998). There were two BICE flights launched from Texas in 1991, with the second one collecting data on the Galactic Plane for 6 hours (the observations from the first flight suffered from high levels of noise). The survey covered the centre of the Galaxy, from $l = 350^\circ$ to $l = 25^\circ$, with an angular resolution of 15 arcmin and a spectral resolution of 175 kms^{-1} .

SOFIA (Stratospheric Observatory For Infrared Astronomy) is a modern successor to the Kuiper Airborne Observatory and NASA's Lear Jet observatory (Young et al., 2012). This modified Boeing 747 carries a 2.7 m mirror up to altitudes of 13.7 km. Precipitable Water Vapour levels (PWV) at this altitude are roughly $10 \mu\text{m}$, excellent for terahertz observations. SOFIA first flew in 2010, and aims to achieve 960 hours of observation per year, split up over 120 flights. However, SOFIA has a maximum flight time of 10 hours, and combined with the requirement for it to return to its base, it is limited to about 4 hours per night on any one source.

One of the highest resolution terahertz surveys has been carried out by the *Herschel* Space Observatory. *Herschel* was launched in 2009, and operated until 2013. This spacecraft had a 3.5 m primary mirror, and observed the Galaxy in infrared and submillimetre wavelengths from its position at L2 (Pilbratt et al., 2010). *Herschel* used a liquid helium cryostat to cool its instruments to a sufficiently low temperature, and this was the primary time constraint on the mission—observations ended in 2013 when the helium ran out.

Ground based sub-millimetre telescopes need to be situated on high mountain tops, above much of the atmosphere. The James Clerk Maxwell Telescope is a

15 m single dish telescope on top of Maunakea in Hawaii at 4092 m elevation, where atmospheric windows are available to observe bands at 345 and 860 GHz (Holland et al., 1999). JCMT uses the SCUBA-2 bolometer array to observe continuum emission from cold clouds of dust in star forming regions (Holland et al., 2013). SCUBA-2 consists of an array of 10,000 pixels, cryogenically cooled to 100 mK. It observes with a beam size of 7.9 and 13.0 arcseconds, at the 450 μm and 850 μm bands respectively.

The largest terahertz telescope in operation is the Atacama Large Millimetre Array (ALMA). This telescope array consists of 66 dishes at a site in the Atacama Desert in Chile, and has been operating since March 2013. This region is very dry, and the telescope is up at 5050 m elevation, so PWV levels are around 1 mm (Wootten & Thompson, 2009). This allows ALMA to observe in its instrumental range of 31–950 GHz (0.3–10 mm). The 66 dishes give ALMA very high resolution and sensitivity. The total collecting area is over 6600 m², and a maximum baseline of 15 km gives ALMA a resolution of 0.015 arcseconds (Wootten & Thompson, 2009). This means that ALMA can resolve much smaller structures than any other terahertz telescope. However, its high resolution and small beam size make it less suited to large scale surveys.

As mentioned earlier, the HEAT terahertz telescope was installed at Ridge A in Antarctica in January 2012, to take advantage of the superb conditions for infrared and terahertz astronomy there (Kulesa et al., 2013b). HEAT mapped the Milky Way using the Ci line, and is covered in further detail in Section 2.2.

Chapter 2

Ridge A

As we've seen in the previous chapter, the high plateau of Antarctica, by virtue of its height, dryness, and cold atmosphere, is the best location on Earth for infrared astronomy. But this prediction is based on a combination of satellite measurements, theoretical models, and a handful of ground-based measurements. In order to conclusively show the advantages of this site, more measurements need to be taken in-situ.

This poses a challenge. The high plateau is one of the most extreme locations on Earth. It is remote, difficult to access, at a very high altitude, and experiences extremely low temperatures. These all present difficulties, for both personnel and equipment. Taking measurements of the sky brightness requires a carefully designed instrument, and a carefully planned expedition.

This chapter covers the remote observatory site at Ridge A: the conditions at the site, the instruments installed there, and the supporting equipment that keeps the site operational.

2.1 The Ridge A environment

Ridge A is a site at the high plateau of Antarctica, approximately 144 km from the peak of Dome A. It is located on the ridge that runs from Dome A to Dome F, hence was given the name Ridge A. The coordinates are $-81:40:25$ South and $72:42:58$ East. Ridge A is at a physical elevation of 4040 m, but due to the cold

2.1. THE RIDGE A ENVIRONMENT

conditions in Antarctica the atmosphere contracts down, and the pressure altitude (which is the relevant number for comparison with temperate latitude sites) is closer to 4400 m Hu et al. (2014).

The site at Ridge A was first established in January 2012, with the installation of the HEAT terahertz telescope. This project was a collaboration between the University of Arizona and the University of New South Wales. UNSW designed and built the PLATO-R module to power the site, and UoA designed and built the HEAT telescope. The United States Antarctic Program (USAP) provided the logistics for this project with assistance from a grant from the Australian Antarctic Division for the Australian component.

Figure 2.1 shows a topographic map of Antarctica, with the locations of McMurdo, the South Pole, Ridge A, and Dome A marked. The topographic data were collected by the Bedmap-2 survey (Fretwell et al., 2013), and are available at their website¹. The highest point of the Antarctic plateau is at Dome A, with an elevation of 4090 m. This is the location of the Chinese Kunlun Station.

Most of this elevation is due to the thickness of the East Antarctic ice sheet. Figure 2.2 shows the topography of the land underneath the ice, also taken from the Bedmap-2 survey. We can see that Ridge A and Dome A are above a mountain range, known as the Gamburtsevs. This mountain range is approximately the size of the European Alps, and has peaks up to 2600 m tall, but is completely buried by kilometres of ice.

At Ridge A the Gamburtsevs are 1860 m high. This is then covered by 2180 m of ice, bringing the total elevation up to 4040 m. It is an interesting fact that if you are standing at Ridge A, there is a mountain range hidden 2 km below your feet.

Another interesting fact shown by Figure 2.2 is that much of the bedrock of Antarctica is close to sea level, so in many places the high elevation is entirely due to the thickness of the ice sheet (although if the ice were removed the land would rise due to isostatic rebound).

¹<https://www.bas.ac.uk/project/bedmap-2/>

²<http://www.earthpoint.us/Convert.aspx>

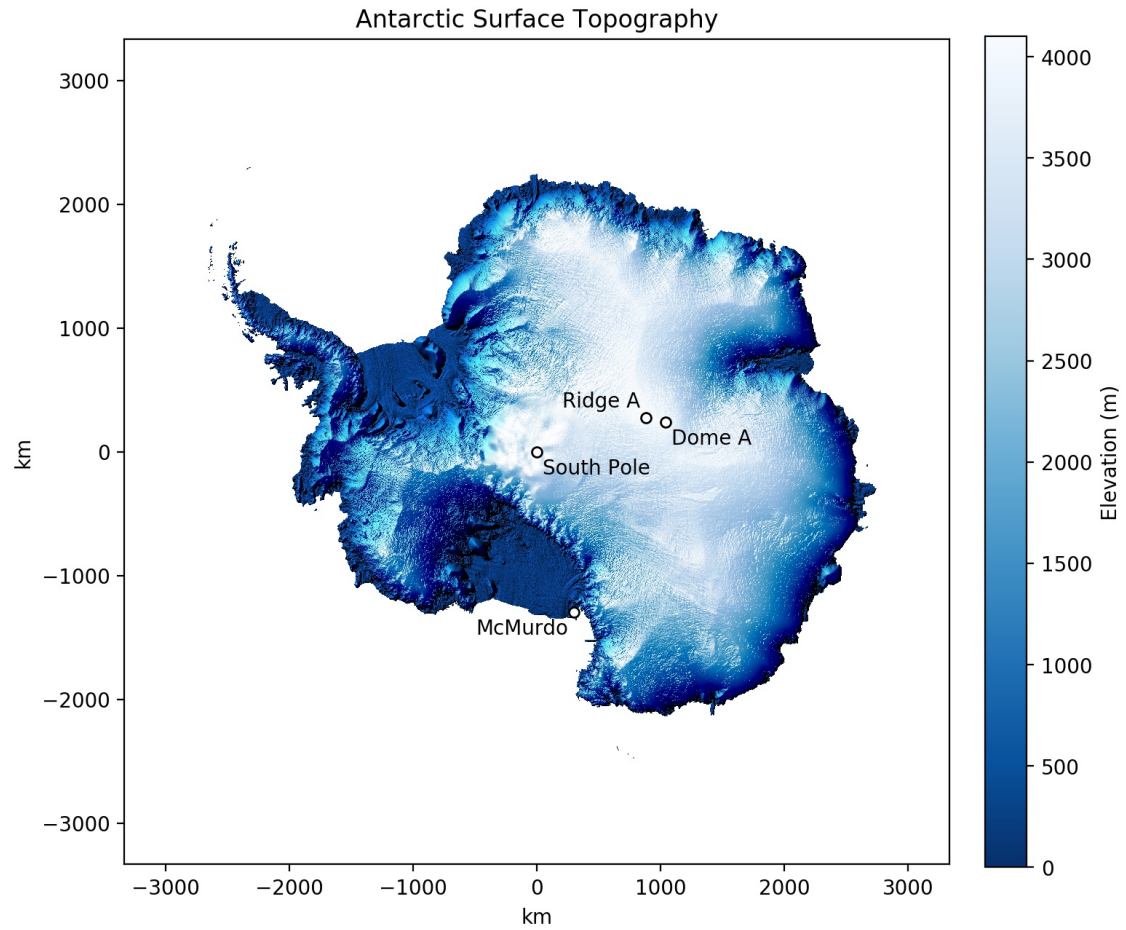


Figure 2.1: The surface topography of Antarctica, with a few stations marked. The fuzzy circle around the South Pole is due to the lack of satellite coverage that far south. Data taken from the Bedmap-2 survey (Fretwell et al., 2013). Coordinates converted using an online calculator².

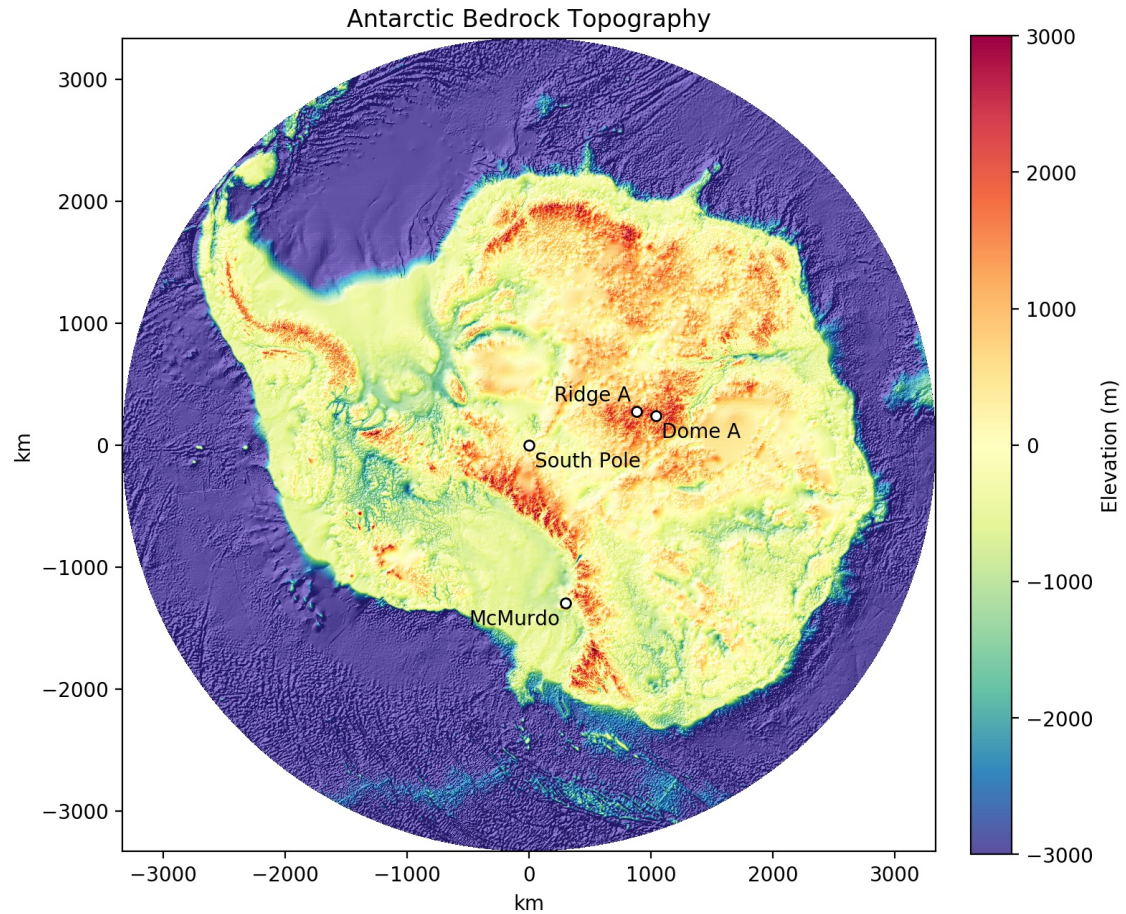


Figure 2.2: The topography of the Antarctic bedrock underneath the ice. This map does not take into account isostatic rebound - the fact that if the weight of ice was removed, the land would rebound up. Data taken from the Bedmap-2 survey (Fretwell et al., 2013). Coordinates converted using an online calculator²

2.1.1 Wind speed

At Ridge A the wind is very calm, with an average 2.5 ms^{-1} wind speed³.

Antarctica has katabatic winds: cold dense air rolls downhill off the high plateau towards the coast. This produces fierce winds near the coast, but very little wind at the centre of the plateau itself. This is a benefit for astronomy, as a stable atmosphere with calm winds provides better seeing.

This low wind speed affects the development of sastrugi: the ‘waves’ in the snow surface (See Figure 2.3). These are caused by the blowing wind in the same manner as sand dunes. In most places in Antarctica, the katabatic winds have a particular direction, and the sastrugi will be aligned in this direction. At Ridge A, the wind speed is very low and has no preferred direction, so the sastrugi tend to be lumpy and disorganised. This presents a slight difficulty when trying to land a plane at the site. With regular sastrugi it is possible to land along their direction for a smoother touchdown. At Ridge A, the lumpy sastrugi make landing difficult from any direction, so more effort has to be made in setting up and maintaining a skiway. This fact will have a major bearing on the work in this thesis, as will be explained in section 9.2.

2.1.2 Temperature

Ridge A is extremely cold. In summer the maximum temperature is around -35°C , and in winter the temperature regularly drops below -75°C . Measurements from satellites have shown the surface temperature falling below -90°C , and some sites near Ridge A have dropped below -98°C (Scambos et al., 2018). These are the coldest temperatures ever recorded on Earth. However, since they are satellite measurements it is not an official record, as this would require a ground based measurement.

The low temperature at Ridge A is due to a number of factors. Firstly, Ridge A is at a high altitude. Secondly, everywhere around Ridge A is also at a high altitude, so there is no warm air nearby. The low wind speed means that the cold

³<http://soral.as.arizona.edu/heat/>



Figure 2.3: Sastrugi (the ripples in the snow) at AGAP. Formed by the blowing wind, in the same manner as sand dunes. [Credit: Matthew Freeman, 2015]

air can stay where it is, rather than being blown away and mixing with warmer air at a different location. The extreme temperatures below -98°C were recorded in shallow depressions, where the cold air can gather.

Additionally, the atmosphere is extremely transparent to infrared radiation. Not only is this good for infrared astronomy, it also allows the ground to radiate its heat away to space, allowing it to cool to extremely low temperatures.

2.1.3 Water vapour

Ridge A was selected by the UoA/UNSW collaboration as one of the most promising sites for infrared and terahertz astronomy. It is located right at the centre of the lowest water vapour contours in the PWV simulations in Swain & Gallée (2006a) (See Figure 1.2). Thus, it was predicted to have extremely low PWV; possibly the lowest of anywhere on Earth. This fact could open up new atmospheric windows, and make it worth the effort of setting up a telescope in the extreme Antarctic conditions.




Figure 2.4: Level of water vapour at the Chajnantor plateau in the Atacama desert (green, measured by the APEX telescope radiometer), compared to the PWV measured at Ridge A in Antarctica by the HEAT telescope. Taken from Burton et al. (2015)

These predictions were verified with the installation of the HEAT telescope, which allowed measurements of the PWV. Figure 2.4 compares the water vapour measurements at Ridge A (blue) with water vapour measurements from the Chajnantor plateau in the Atacama desert; the home of the ALMA telescope, and one of the best astronomical sites on the planet (green). The Atacama regularly has water vapour levels below 1 mm, but Ridge A regularly has PWV levels below 0.1 mm, with the worst Ridge A levels being comparable to the best Atacama levels.

2.1.4 Access

Ridge A is rather difficult to access. It is located ~ 950 km from the South Pole, and ~ 1200 km from the coast. There are two main methods of accessing inland

2.1. THE RIDGE A ENVIRONMENT

sites in Antarctica: an overland traverse, or flying.

A traverse can take several weeks, as a train of tractors drive across the snow. This is very slow, and requires a large team of people and vehicles. It is also affected by terrain. Mountains need to be navigated around, and hidden crevasses are a significant danger. However, it is the most efficient method for transporting heavy cargo.

A flight is much faster, taking only a few hours to reach almost any location. However, flights have a limited cargo capacity, and burn significant amounts of fuel. Flights also require substantial logistical support. A supply chain needs to be in place to bring fuel and crew down to Antarctica and to maintain an aerodrome. Depending on the type of plane, a skiway may need to be prepared at the site. Planes also have a limited range, so it may be necessary to set up fuel caches at sites around Antarctica beforehand.

Generally, planes are preferred for light cargo and fast trips, whereas traverses are used for heavy cargo, or where the capacity for launching aircraft is not available.

The Polar Research Institute of China reaches Kunlun Station at Dome A via a traverse from Zhongshan station on the coast (See Figure 1.1). The United States Antarctic Program accesses their South Pole Station by flying from McMurdo Station. Since 2005 the US have also been using a traverse to transport fuel from McMurdo.

For our experiment at Ridge A, since we are only transporting smaller instruments (and are working with the USAP), we access Ridge A by flying from the South Pole station. The ~ 950 km distance makes it the most remote location that is operated out of the South Pole Station.

The USAP uses a number of DHC-6 Twin Otter aircraft for transporting field teams and cargo, in contract with Kenn Borek Air Ltd. (See Figure 2.5). Twin Otters are small transport planes powered by two turboprop engines, and are fitted with skis for operating on snow. They come equipped with 19 passenger seats which can be removed to transport cargo, although actual capacity will be limited by weight constraints. Twin Otters are purpose-built for short-takeoff-and-landing



Figure 2.5: A Twin Otter light aircraft at the South Pole, used for transport to Ridge A. [Credit: Matthew Freeman, 2014]

on snow or gravel, so are ideally suited to Antarctic conditions.

The Twin Otters have a maximum range of 1435 km with zero payload⁴. The return trip to Ridge A is farther than this, and the Twin Otter is carrying a full load of cargo, so it is necessary to refuel on the way. For our mission we refuel at a site called AGAP, roughly half way between the Pole and Ridge A. This site has a cache of fuel barrels that were airdropped by C-17s during the 2008–2009 season. The Twin Otters stop here on the outbound and return journeys to refuel. This site was established for the Antarctica’s Gamburtsev Province (AGAP) project, which mapped the topography of the Gamburtsev mountains under the ice.

Transporting cargo and personnel to Ridge A is a complicated logistical task, and depends heavily on the training and expertise of the Twin Otter pilots and mechanics.

If Ridge A becomes a successful observatory site, it will be necessary to transport heavier loads in the future, to construct a station and larger telescopes.

⁴<https://www.vikingair.com/viking-aircraft/dhc-6-twin-otter>

One possibility for this is to use the USAP’s LC130 Hercules planes. These can carry much more cargo than a Twin Otter, but would require a large and well maintained skiway at Ridge A in order to land and take off again. Other options that have been considered include airdropping cargo by parachute, or doing a “combat offload”, where the plane touches down, the cargo is pushed out the back ramp, and the plane takes off again without stopping (Craig Kulesa, private communication). However, both these options are probably not suitable for delicate astronomical equipment.

2.2 HEAT

The Ridge A site was first established in January of 2012, with the installation of the HEAT telescope and the PLATO-R observatory modules. HEAT (the High Elevation Antarctic Telescope) is a small 61 cm telescope that observes in the terahertz band. This telescope was established as a pathfinder, to test the atmospheric transparency at Ridge A. It has been observing the plane of the Milky Way, slowly building up a map of our Galaxy at terahertz frequencies.

Figure 2.6 shows the HEAT telescope with its protective cover, on site at Ridge A. The grey window is made of high density polyethylene (HDPE) that is transparent to terahertz radiation, so it does not affect the observations. The body of the telescope is on legs to raise it up above the snow. The purpose of this design is to stop snow piling up against the telescope. In Antarctica the wind blows the snow around, and whenever the wind hits an obstruction it stops and drops its snow. So any object left on the ground will accumulate snow, and eventually be buried. HEAT’s platform is designed to allow the wind to pass underneath without disruption, to avoid snow accumulating. It is not perfect, and some snow does build up, but the rate is greatly reduced in comparison to a solid structure. The Amundsen-Scott South Pole Station is up on legs for the same reason.

HEAT uses Schottky diode detectors and heterodyne receivers to measure the CI lines at 492 and 809 GHz, as well as the NII line at 1461 GHz . Two of these receivers are on site at a time, and they are swapped around each year when the



Figure 2.6: The HEAT terahertz telescope with its protective cover, located at Ridge A. [Credit: Matthew Freeman, 2015]

site is accessed. Future plans involve adding new receivers to observe the 1900 GHz C+ line. The receivers have a 1 GHz bandpass, and a 1 km s^{-1} spectral resolution Burton et al. (2015).

The receivers are mounted inside a cryostat to keep them at the required 50 K. This is visible in the centre of Figure 2.7. Since HEAT operates remotely for a year at a time, it is not possible to use liquid cryogenics to cool the detectors. Instead, they are cooled with a Stirling microcooler, and the vacuum in the cryostat insulates them from the environment. This cooling system uses only 80 W of power. Including all the electronics, HEAT only requires a power budget of 150 W.

HEAT uses polished aluminium mirrors in an off-axis Gregorian layout fed by a siderostat, which can be seen in Figure 2.7. Light reflects from the flat mirror (left) to the 61 cm parabolic primary mirror (right), then to the small ellipsoidal secondary, before entering the cryostat. The flat mirror has one axis of movement: it can rotate to adjust the elevation of HEAT's beam. This means HEAT only has two moving parts (the cooler and mirror), maximising reliability. This layout gives



Figure 2.7: The HEAT telescope without its cover, showing the mirrors and cryostat. [Credit: Craig Kulesa, 2012]




Figure 2.8: HEAT’s drift scan pattern across the Milky Way. The white lines are simulated HEAT scans, superimposed on a map of the Galaxy. Taken from the HEAT website.⁵

HEAT a diffraction limited beam size of 2 arcseconds (Burton et al., 2015).

With only one axis of movement, HEAT observes using drift scanning (i.e., the telescope scans by letting the sky drift through its beam due to the Earth’s rotation). Since Ridge A is relatively close to the South Pole, the motion of the stars through the sky is mostly horizontal. HEAT observes at a particular elevation, and lets the Milky Way drift across the view. Once HEAT has observed a slice of the Milky Way, it adjusts its elevation, and lets another slice of the Galaxy drift through its field of view. This pattern can be seen in Figure 2.8.

2.3 PLATO-R

At such a remote location, powering and operating a telescope is a difficult task. It is similar to designing a satellite: the instruments need to have a self contained power source, and must be operated remotely without any access.

The HEAT telescope is powered and controlled by PLATO-R (the PLATeau Observatory for Ridge A). PLATO-R consists of a series of modules installed at Ridge A. These include the Instrument Module (IM) and the Engine Module (EM), as well as a cube of solar panels. Each module is roughly a 1 m cube. This is just small enough to fit through the doors on the Twin Otter (1.27×1.42 m), which are currently the limiting factor for the size of equipment we can bring.

The cargo doors of the Twin Otter present another limit on the cargo that can be carried: they are over a metre off the ground. Currently at Ridge A, any cargo must be lifted up and down by hand, or slid up or down a steep ramp, without mechanical assistance. This restricts the weight of any single piece of cargo to that

⁵<http://soral.as.arizona.edu/heat/>

which can be lifted by the crew.

2.3.1 Power

Power comes from two sources. During the summer months when the Sun is above the horizon, the cube of solar panels provides 1 kW of power. These can be seen on the right in Figure 2.9

For the dark winter months, power comes from the two small HATZ 1B-30 diesel engines contained in the Engine Module (the green box in Figure 2.10). These engines are connected to fuel-cooled brushless alternators with NdFeB magnets. Each engine provides 1 kW of power. They run on AN8 kerosene, the same fuel that is used to power all of the US equipment, including the base's generators, the aircraft, traverse tractors, snow mobiles, and any other vehicles. Using a single fuel greatly simplifies logistics, and the addition of racing oil allows it to be used in PLATO-R's diesel engines. The engine module is located ~ 60 m away from HEAT, so that exhaust or vibrations do not affect the observations.

The two diesel engines are mounted in the top half of the EM, while the bottom half contains an 800 L fuel tank. The module is constructed from fibreglass, and the gaps between the upper and lower sections are insulated with wool, which has excellent properties at very low temperatures. Temperature management is an important issue at Ridge A. The engines need to be sufficiently warm to start at the beginning of winter, so heaters are mounted inside the module to warm it. However, once the engines are running, cooling becomes the important factor. Due to the low air pressure, the air is a much less efficient conductor of heat, and so is less able to cool the engines. Thus, fans are installed to keep air flowing through the module, keeping the engines cool while running.

Only one engine is run at a time. The module contains the second engine as a backup, to be switched on when the first engine fails. Diesel engines require regular maintenance, but these ones are run in extremely harsh conditions for months at a time with no access, so they usually fail before the Sun rises again at end of winter. Improving the reliability of the engines is an important step to establishing future telescopes at this site.

Wind power has been considered as another option for power during winter, to try to reduce the weight of fuel that needs to be brought in. It would certainly be beneficial to use renewable sources, but due to the very low winds speeds at Ridge A, it may not be practical. The low temperatures and icing may also be issues. Work is being done at the University of New South Wales into possible wind turbine designs that could work in these conditions.

2.3.2 Control

The Instrument Module is shown in Figure 2.9. This insulated yellow box contains the supervisor computer that controls power and communications to HEAT and the other equipment, as well as a bank of LiFePO_4 batteries to store and regulate power. The batteries power HEAT and the other equipment, and the solar panels and engines are used to charge the batteries when they are low. The battery pack can store up to 20 kWhr of energy.

An Iridium satellite connection allows remote control of PLATO-R and HEAT, and transmits data back. The bandwidth is limited, so the raw data are processed in the IM, and the reduced products are transmitted. The Iridium antenna is visible in Figure 2.9 as a short stick on the top right of the IM box.

The instrument module also has a number of instruments attached to it. These include 2 anemometers for measuring wind speed, a sky camera, a cluster of webcams for monitoring the other modules, and, since 2015, the infrared sky monitor NISM, which is the key experiment for this thesis.




Figure 2.9: The PLATO-R Instrument Module at Ridge A, with HEAT and the solar panels in the background. [Credit: Nic Bingham, 2015]



Figure 2.10: The PLATO-R Engine Module. The snow piles up and has to be dug away each season, resulting in the crater around the module.[Credit: Nic Bingham, 2015]

Chapter 3

NISM

Section 1.4 described measurements of the infrared sky background that were successfully taken from the South Pole with the original instrument NISM (Lawrence et al., 2002a). Given the success of these measurements, the next step is to take measurements from the high plateau, where the conditions are superior.

The cryostat module from NISM (including the dewar, cooler, and detector) was re-purposed, and incorporated into a new instrument for deployment at the high plateau. This instrument was, somewhat confusingly, also called NISM. This chapter describes the new instrument, covering its specifications and installation. From this point on, all references to NISM refer to the new version of the instrument.

The new NISM (Near Infrared Sky Monitor) is a small instrument designed to measure the infrared sky background at the high plateau of Antarctica. It was installed at Ridge A in the 2014/2015 summer, during a maintenance mission for HEAT and PLATO-R. NISM was designed and constructed before I joined the project, so my personal contribution started with the installation of NISM, described in Section 3.11. The total construction costs of NISM were approximately \$200 000 (Michael Ashley, private communication).

Figures 3.1, 3.2, and 3.3 show the NISM instrument. The anodised red housing was constructed by the Australian Astronomical Observatory (AAO). The box on the left houses the cryostat module containing the detector. The box on the right houses the stepper motor. A fold mirror is located in the middle, between the two




Figure 3.1: The Near Infrared Sky Monitor (NISM) [Credit: AAO, 2014]

sections, and is attached to the stepper motor. The cylinder on the back of the instrument is a black body, which serves as a reference brightness for calibration.

The goal of NISM is to determine the infrared emission from the atmosphere by measuring the spectral radiance at the zenith in the K_{dark} band. The previous version of NISM operated by ‘chopping’ between the target and a reference position. However, the current NISM operates by rotating continuously, completing one rotation in ~ 10 seconds. This time is short enough that all the sources of infrared radiation should be constant, and so chopping is not required.

The stepper motor rotates the mirror to scan NISM’s beam in elevation, dividing one 360 degree rotation into 8000 steps. NISM records ~ 16.7 measurements per second. Since NISM rotates once per ~ 10 seconds, and measures with both channels of the analogue-digital converter before advancing, there are roughly 90 motor steps between each measurement pair.

One of the problems discovered with the original NISM is that it had a




Figure 3.2: NISM installed on top of the IM at Ridge A. The detector is housed in the section on the left. The mirror is visible in the centre of the instrument, and the section on the right houses the motor that rotates the mirror. [Credit: Nic Bingham, 2015]

long-wavelength leak (Lawrence et al., 2002a). This allowed extra light in at longer wavelengths, which caused it to measure an anomalously high brightness. Given the extremely rapid rise in the atmospheric thermal emission towards longer wavelengths, even a tiny leak can overwhelm the K_{dark} emission. This error required significant correction, so the new version of NISM included an additional long-wavelength blocking filter, to prevent the extra light from leaking through.

The parameters of NISM are detailed in the following sections. Unfortunately, since various pieces of NISM were constructed as long ago as 1996, some documentation is no longer available. Additionally, the new instrument was only completed immediately before it was due to be shipped down to Antarctica, so there was not time to do a detailed analysis in the lab. Without complete documentation, there is some uncertainty to the exact specifications of the instrument. The data below are the best we have available, based on planning documents, manufacturer’s specifications, email correspondence, and optical




Figure 3.3: Cross section of NISM, showing the cryostat, black body, mirror, and motor. [Credit: AAO, 2014]

modelling with the software Zemax. Once NISM can be recovered from Antarctica, proper measurements of the instrument specifications can be taken. These would include careful measurements of the dimensions of the instrument, paying particular attention to the optical path. The thickness and radius of curvature of each optical element needs to be measured to model the final beam size, which should also be measured experimentally. A measurement of the final passband including the short-pass filter should also be done.

3.1 Detector

The detector of NISM is a single Indium Antimonide diode, with a diameter of 2.032 mm (“final optical configuration” sketch from John Storey, 10 July 1996). Note that $2.032 \text{ mm} = 0.08 \text{ inches}$. This detector was supplied by Cincinnati Electronics, and its parameters are described in Table 3.1.

The detector is housed inside an integrated detector-dewar cooler assembly. This small vacuum dewar keeps the detector at an operating temperature of 77 K (see Section 3.6).

Table 3.1: NISM InSb diode properties

Quantum Efficiency	$> 70\%$
R_0A	$2 \times 10^6 \Omega\text{cm}^2$
Dark Current	$5.1 \text{ fA}/\sqrt{\text{Hz}}$
Johnson Noise	$12.7 \text{ fA}/\sqrt{\text{Hz}}$

3.2 Optical configuration

Inside the dewar, 6.50 mm away from the detector, is a cold stop with a 3 mm diameter that blocks all thermal radiation from outside the field of view (See Figure 3.5 for the optical layout model). The 3 mm diameter gives the centre of the detector a 26° field of view (“final optical configuration” sketch from John Storey, 10 July 1996), and the field of view from one edge of the detector to the other is 42° (Cincinnati quote, 6 May 1996).

Attached to the cold stop is a filter that limits observations to the K_{dark} band. Lawrence et al. (2002b) states that the central wavelength of the filter is $2.379 \mu\text{m}$, with a full width half maximum of $0.226 \mu\text{m}$. The documentation from Cincinnati Electronics states that the filter is $2.30\text{--}2.45 \mu\text{m}$, with a peak transmission of $\geq 85\%$, and an average blocking of $\leq 0.1\%$ from $0.1 \mu\text{m}$ to the cut-on, and from the cut-off to $6.0 \mu\text{m}$. (Cincinnati quote, 6 May 1996). The NISM design document states that the window is germanium, has a diameter of 10 mm, a thickness of 1 mm, and a 60% minimum peak transmission (NISM design document, 23 May 1996).

10 mm away from the detector is the window of the cryostat (measured from an unknown 1:1 drawing by Nic Bingham or Michael Ashley). The dewar window is anti-reflection coated germanium (Cincinnati quote, 6 May 1996), has a diameter of 12 mm, is 1 mm thick, and has a transmission of 95% (NISM design document, 23 May 1996). Everything inside the cryostat is cooled to 77 K.

Unfortunately, the cryostat was constructed in 1996, and proper documentation is not available. We have detailed all the information we can, but the exact

3.2. OPTICAL CONFIGURATION

arrangement of components inside the cryostat still has some uncertainties.

Outside the cryostat, and built into the housing of NISM, is a long wavelength blocking filter made of borofloat glass (Jon Lawrence’s request for quote email, 5 April 2002). The gap between the cryostat window and the blocking filter was difficult to measure, as the cryostat mount has some freedom of movement. Using photographs, it was estimated at 8.3 mm (Michael Ashley) or 8.6 mm (Matthew Freeman). The blocking filter is 1.75 mm thick, and has a physical diameter of 32.0 mm. The clear aperture is 29.0 mm, defined by the AAO lens spacer.

The blocking filter was manufactured by NDC Infrared Engineering, and ordered by Jon Lawrence from UNSW in Oct/Nov 2002 as “FILTER 2268 NM CENTRE WAVELENGTH 32 mm diam, Block to $5.5\text{ }\mu\text{m}$ ”. It is a bandpass filter, with a multi-layer dielectric, blocking $< 0.1\%$ out to $5.5\text{ }\mu\text{m}$ (from the manufacturer’s website¹). It has a 75% transmission from ~ 2.1 to $2.5\text{ }\mu\text{m}$ (Jon Lawrence email, 23 Dec 2002, presumably from looking at the filter scan data that would have come with the filter). Borofloat 33 glass² has a transmission of $\sim 85\%$ at $2.4\text{ }\mu\text{m}$, dropping rapidly to $\sim 10\%$ at $2.7\text{ }\mu\text{m}$, peaking again at $\sim 50\%$ at $3.2\text{ }\mu\text{m}$, and then $< 5\%$ out to $5\text{ }\mu\text{m}$.

The blocking filter was added to stop the red leak discovered in the previous version of NISM. It was originally intended to be placed on the outer ‘sky’ side of the lens, in the 4 degree field of view. In the final design it was placed on the inside of the lens, in the possibly 42 degree field of view of the detector. However, this should not be a problem, as the design is not using the full f/1 potential of the lens.

The side of the filter facing away from the detector was noted to be purple, which is presumably the dielectric coating (Nic Bingham , email 17 April 2014).

Next to the long wavelength blocking filter is the lens, which focuses light onto the detector of NISM, giving it an approximately 4 degree field of view on the sky. This lens was also re-used from the original NISM instrument, where it brought the beam to a focus at the chopper wheel (Figure 1.5). The AAO lens separator

¹<https://web.archive.org/web/20020617200702/http://www.ndc.com/products/newtfod/narrow/narrow.html>

²https://www.us.schott.com/d/borofloat/1ca1519d-b5b3-469d-b3eb-9b80235b1783/1.2/borofloat33_opt_usa_web.pdf

design gives a 4.25 mm distance between the filter and the rim of the lens. Due to the curvature of the lens, the distance between the filter and the centre of the lens is 6 mm. The clear aperture of the lens is defined by the lens holder and o-ring, which is a Parker 2-028 with an internal diameter of 34.65 mm.

The lens as ordered was an f/1 ZnSe AR coated meniscus lens with a diameter of 38.1 mm (Max Boccas email, 28 Mar 2013). The focal ratio f/1 means that paraxial rays (close to the optical axis) focus at a point 38.1 mm behind the lens. Measurements of the lens shows that it has an edge thickness of 2.75 mm, and the radius of curvature of the concave surface is 104 mm (Nic Bingham, email of 17 April 2014).

The radius of curvature of the convex side is more difficult to measure. Estimates put it at 29 mm, with a central lens thickness of 7 mm (Nic Bingham, email of 17 April 2014). A more accurate way to calculate the convex radius of curvature is to make an optical model of the lens, using the focal length and the concave radius of curvature. The focal length does depend somewhat on the wavelength, but ZnSe lenses are normally used for mid-infrared wavelengths, and the refractive index does not vary greatly from the mid to the near infrared, so the reported f/1 focal length should be accurate. Modelling with Zemax determined that the centre thickness of the lens is 6.35 mm, and the radius of curvature of the convex side is 36.6 mm. This radius of curvature is somewhat different to that measured physically, but is probably more accurate given the difficulty of the measurement.

This lens appears to be similar to the model ZE-PM-1526-38.1 lens³, which has an aperture of 38.1 mm, an edge thickness of 3 mm, a centre thickness of 6.60 mm, and a focal length of 38.1 mm.

The final optical element is the fold mirror. This flat mirror is mounted at a 45 degree angle, to direct the beam 90 degrees towards the sky. It is an ellipse, with a semi-minor axis of 37.5 mm, and the centre of the mirror is 44.9 mm from the lens.

³https://www.latticeoptics.com/web/cboard/main.php?act=view&mc=3&mid=3&arrange=input_date&direct=desc&mc=3&main_no=20&page=1&rtm_page=main.php&ttn=15&cn=1

The stepper motor rotates the mirror to scan the field of view in elevation, in a similar manner to HEAT, which allows NISM to eventually observe the entire sky by drift scanning.

3.3 Beam

NISM’s optical design should give it a beam size that is approximately 4 degrees across. Measurements of the beam shape can be seen in Figure 3.4. To obtain these calibration measurements, the reading from NISM was recorded when looking at a hot target (a soldering iron tip) at various positions 4.87 m away in front of the detector. A smooth 2D surface was fitted to these measurements to determine the location of the centre of the beam, and the xy positions of the measurements were converted into radial distances from the centre. These radial distance measurements are then plotted against the signal from the NISM detector.

The beam is larger than expected, with ‘wings’ extending out to 7.5 degrees, so the optics of NISM were carefully modelled using Zemax to understand the cause. The large beam diameter means NISM has trouble resolving small sources, like stars and planets, due to the increased background signal from the sky.

3.4 Optical modelling

To better understand the optics of NISM, the optical model was recreated in the software Zemax (release 13, service pack 2). This model is shown in Figure 3.5, and the parameters are listed in Table 3.2. We constructed this model based on the best information we have about the layout of NISM, as described in Section 3.2. Zemax models each optical surface in NISM and simulates a large number of light rays passing into the instrument, to determine how many fall on the detector.

The model uses “non-sequential components” in Zemax, in order to incorporate both sources and multiple reflections. All surfaces are anti-reflection coated, and “split rays” are used to model reflections. The wavelength of light used was $2.4\,\mu\text{m}$. To recreate the soldering iron measurements, a point source was placed




Figure 3.4: The beam shape of NISM (mirrored). Measured by observing a hot target (a soldering iron tip) at a number of angles from the beam centre. The y axis is the ADU readout of the detector.




Figure 3.5: The Zemax optical layout of NISM. Left to right, the components are the InSb detector, the cold stop, the cryostat window, the blocking filter, and the curved lens. The vertical lines extending above and below the cold stop and cryostat window are “absorbing stops”, added to ensure that rays have to travel through the optics to reach the detector.

Table 3.2: Zemax optical layout for NISM. The thickness parameter is the distance from the current surface element to the next one. All surfaces are AR coated. Units are in mm.

Surface	Radius of Curvature	Thickness	Material	Radius (physical)	Radius (clear)
1 lens front	36.6	6.35	ZnSe	19.05	17.330
2 lens back	104.0	6.00		19.05	17.330
3 short-pass filter front		1.75	Borofloat	16.00	14.500
4 short-pass filter back		8.40		16.00	14.500
5 cryostat window front		1.00	Germanium		5.500
6 cryostat window back		1.50			5.500
7 K_{dark} filter front		1.00	Germanium		1.500
8 K_{dark} filter back		6.50			1.500
9 Detector			InSb		1.016




Figure 3.6: Zemax optical modelling of NISM. The blue rays are from a point source on axis, 5 m away. All rays land on the detector.

5 m away from NISM. The source emits a cone of rays with a spread of ± 10 degrees, with rays uniformly distributed in the cone.

An example of the model in action can be seen in Figures 3.6, 3.7, and 3.8. The point source is 5 m away to the right, emitting rays in a 10 degree cone. These rays pass through the optics and may fall on the detector. In Figure 3.6 the source is on axis, and all the rays fall on the detector. In Figure 3.8 the source is moved to 6.5 degrees off axis, and all the rays miss the detector.

From Figures 3.6, 3.7, and 3.8, it is immediately apparent that there is a problem with the optical layout: the detector is clearly not at the focal point of the lens. We suspect that the AAO optical design used a rough measurement of the radius of curvature for the lens. As mentioned in Section 3.2, this is a difficult physical measurement to make. We think that it is more accurate to use our Zemax model of the lens, which assumes that the focal length is exactly $f/1$.

With this problem of the detector not being at the focal point of the lens, we are not using the full lens diameter. Figures 3.6, 3.7, and 3.8 show that only light rays passing through the centre of the lens will reach the detector, and that the edge of the lens is unused. This decreases the light falling on the detector,




Figure 3.7: Zemax optical modelling of NISM. The blue rays are from a point source 2 degrees off axis, 5 m away. Some of the rays miss the detector.




Figure 3.8: Zemax optical modelling of NISM. The blue rays are from a point source 6.5 degrees off axis, 5 m away. All of the rays miss the detector.




Figure 3.9: The beam shape of NISM as calculated by the Zemax optical model. The measured beam shape is overlaid in blue. There is a significant discrepancy, with the measured beam having a sharper peak and ‘wings’ out to a large angle.

reducing our sensitivity.

The Zemax model was run with a point source at a number of angles off axis. For each run the number of rays falling on the detector was counted, allowing the (relative) beam shape to be calculated. The results of these runs are shown in Figure 3.9. It shows the results from the Zemax modelling, overlaid with the experimentally measured beam shape. The Zemax ray count is scaled to the same units as the experimental measurement.

We can see that there is a noticeable discrepancy. The Zemax model has a fairly smooth drop off, while the experimental measurements show a sharp central peak with extended ‘wings’. In the original design of NISM it was intended to have a beam that was approximately 4 degrees across. This would correspond with the central peak of the measured beam, without the extended wings. It is clear that the beam shape is not matching what was intended, but it is not clear why. The

Zemax optical model is as accurate as we can make it, given the limited information on the instrument, but does not match the observed beam. Altering the model to account for uncertainties, such as the exact position of the cryostat, has not been able to account for the measured beam. A more detailed exploration of the reasonable parameter space for the optical model could not match the measured beam either. At this stage, the extended beam size is a mystery, and will have to wait until the instrument is recovered to be resolved.

3.5 Passband

To limit observations to the K_{dark} band, a bandpass filter is placed in front of NISM's detector, attached to the cold stop inside the cryostat. The wavelength response of this filter can be seen in Figure 3.10. These calibration measurements show the transmission window at $\sim 2.4 \mu\text{m}$.

The extra peak at $\sim 4.5 \mu\text{m}$ is unwanted, so an additional blocking filter was installed to block out any longer wavelength radiation. The filter blocks to $< 0.1\%$ out to $5.5 \mu\text{m}$. Refer back to Section 3.2 for more information. This filter is an important component: the thermal emission from the atmosphere is much greater at wavelengths longer than $\sim 2.4 \mu\text{m}$, so a small amount leaking through the filter could produce a significant reading.

3.6 Cooler

Since NISM operates at a remote site for long periods of time, liquid cryogenics would not be suitable, as they would evaporate. Instead, a Stirling-cycle cooler is used to cool the cryostat. The cooler is a Model K508 Stirling engine manufactured by Ricor, which operates on a maximum of 10 W of power (typically 2–3 W). It can provide 0.5 W of cooling, and has an expected lifetime of 8000 hours. The cooler is attached to the dewar, and the whole cryostat module weighs 500 g. Figure 3.11 shows the cryostat module.

The cooler was found to be a source of noise to the system. Electrical noise




Figure 3.10: NISM’s measured passband, showing the window in the K_{dark} band at $\sim 2.4 \mu\text{m}$. The y axis is the electrical signal measured from NISM’s diode. Subsequent to this measurement, an additional short-pass filter was added to block the peak at $\sim 4.5 \mu\text{m}$ (Lawrence, J., private communication).




Figure 3.11: The cryostat/Stirling cooler module, with the microprocessor and ADC boards attached. Taken from Bingham & Ashley (2014)

from the motor, thermal cycling, and vibration can all introduce noise into the diode, which can affect the precise measurements. Bingham & Ashley (2014) goes into detail on the modifications that were made to reduce the noise. The most important modification was to synchronize the ADC to the cooler piston.

3.7 Motor

To control the elevation of NISM's beam, the mirror is attached to a S17-3-E3-01 Silvermax stepper motor, supplied by Quicksilver Controls Inc. It is visible on the right in Figure 3.3. The motor rotates the mirror, scanning the view in elevation. The motor includes an internal optical encoder with 8000 steps. The position reported by the optical encoder gives the position of the shaft of the motor, and hence the elevation of any particular observation.

The S17 range of motors⁴ have an operating temperature range of -10°C to $+80^{\circ}\text{C}$. Ridge A is always colder than this, so it is necessary to warm the motor to its operating temperature.

3.8 Electronics

NISM is DC coupled, meaning that the current from the detector is measured directly. It is typical in circuits to have capacitors to damp out unwanted frequencies. If an instrument is chopping at 100 Hz, then capacitors can be used to select only the signal at 100 Hz. However, with NISM we want to measure all the signal from the detector, so do not have any intermediate capacitors.

The signal from the InSb diode goes directly to the analogue-digital converter, where it is digitised. The digitised signal is then passed to a microcontroller board, which then communicates with the PLATO-R supervisor computer using RS485.

NISM samples approximately 1000 times per minute, equivalent to a sampling rate of $\sim 16.7\text{ Hz}$, or a sample duration of 0.06 seconds.

3.8.1 ADC board

The electronics diagram for NISM's ADC converter board is shown in Figure 3.12, and the board itself is shown in Figure 3.13. The main feature on the left of Figure 3.12 is the electronics layout for the DDC112 Analogue Digital Converter. This is the chip that converts the analogue signal from the InSb diode into a digital signal for processing. These components are also described in Bingham & Ashley (2014).

The main two inputs are labelled INSB_A and INSB_K, and are visible on the right of the chip. These are the inputs from the anode and the cathode of the InSb diode respectively. The current from the detector is fed directly into the ADC chip. To minimise noise, these two connections are “air wires”: they are lifted up above the printed circuit board, to isolate them from stray currents. The wires are visible on the right of Figure 3.13. The two terminals are placed as close as is practical to

⁴<http://www.quicksilvercontrols.com/SP/OB/QCI-SilverMax17.pdf>

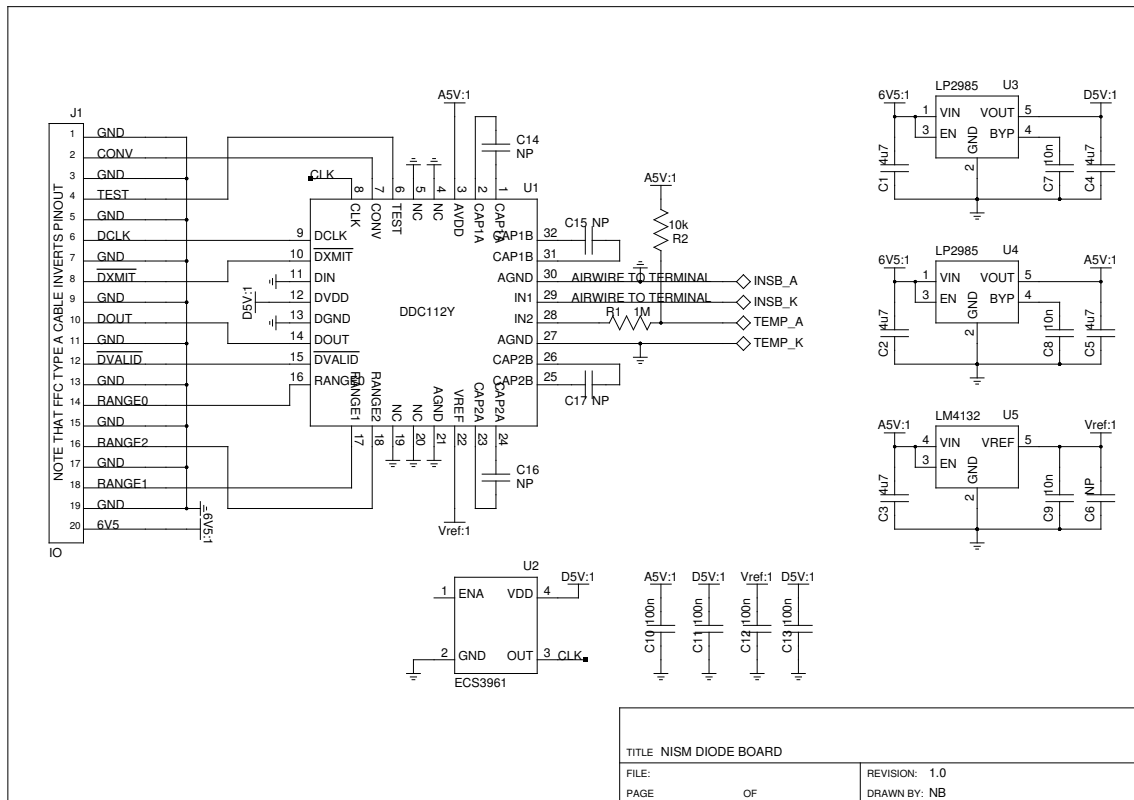


Figure 3.12: Electronics diagram for NISM's DDC112 based analogue digital converter board.

the ADC and the detector, to minimise the area of the circuit. A larger area would result in a larger current being generated by any stray magnetic fields enclosed.

The DDC112 chip has two internal channels. The analogue-to-digital conversion takes some time, so the DDC112 integrates on one channel while converting on the other. This allows continuous integration to be performed. However, the two channels have slightly different gains and offsets, which must be accounted for when processing the data.

TEMP_A and TEMP_K are the other two outputs from the cryostat. These are the anode and cathode of the temperature sensor, which is a silicon diode that returns a current based on the temperature inside the cryostat.

The large block on the left of the chip is the output connector, visible on the left in Figure 3.13. This sends the information from the ADC to the microcontroller board, and includes a 6.5 V input that powers the ADC.

The three diagrams on the right of Figure 3.12 are voltage regulators. They convert DC from the 6.5 V input into a different, lower voltage. In order they are a




Figure 3.13: NISM’s DDC112 analogue digital converter. The two terminals on the right with visible “air wires” are the inputs from the InSb detector [Credit: Nic Bingham, 2014].

3.8. ELECTRONICS

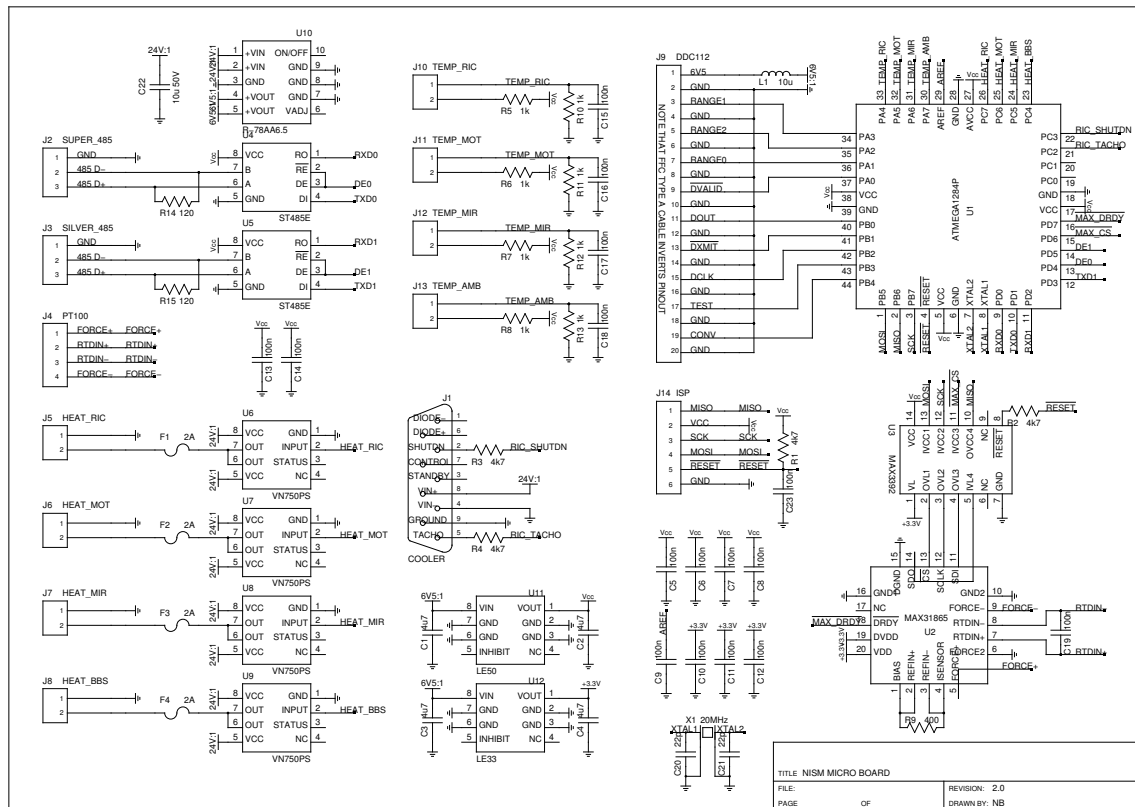


Figure 3.14: Electronics diagram for NISM’s microcontroller board. This board communicates with the PLATO-R supervisor computer using RS485.

digital 5 V source, an analogue 5 V source, and a reference voltage. One is attached on each side of the ADC chip. They are labelled A5V:1, D5V:1, and Vref:1 in Figure 3.12, and are visible as black chips in Figure 3.13 (on the top, bottom, and lower left). These are placed as far apart as possible, to ensure that the different voltages are stable and unaffected by noise (the digital D5V:1 is particularly noisy, much more so than the analogue A5V:1).

On the bottom right of Figure 3.12 are four capacitors that damp out noise. The final component on the bottom left of figure 3.12 is the clock, which keeps time for the circuit. It is visible as a shiny rectangular component on the upper left of Figure 3.13

3.8.2 Microcontroller board

The ADC board is connected with a ribbon cable to an Amtel AVR 8-bit microcontroller board, which governs NISM and communicates with the PLATO-R

supervisor computer using RS485. The electronics diagram for the microcontroller board is shown in Figure 3.14.

On the upper right of Figure 3.14 is the main chip, labelled ATMEGA1284P. The connections at the top of the chip lead to the four temperature sensors and the four heaters. On the top right of the chip are the connections to the Ricor shutdown and the Ricor tachometer, which is a hall sensor added to measure the position of the Ricor cooler motor (Bingham & Ashley, 2014).

On the left of the chip is the connection to the ADC board, via a ribbon cable. This allows the microcontroller to receive data from the ADC board, and give it commands.

At the top middle of Figure 3.14 are four A590 platinum resistance thermometers. They measure the temperatures of the Ricor motor, the stepper motor, the mirror, and the ambient temperature.

At the top left of Figure 3.14, SUPER_485 is the connection to the PLATO-R supervisor computer. Below that, SILVER_485 is the connection to the Silvermax stepper motor that rotates the mirror in elevation. Next, PT100 is the temperature sensor attached to the black body. This temperature is much more important to record, so has a more accurate temperature sensor. Finally on the left, there are the four heater circuits, which warm the Ricor motor, stepper motor, mirror, and black body to their operating temperatures. They are relays connected to 24 V supplies. The small current from the chip triggers them to send a large current to the heater.

In the centre of Figure 3.14, the diagram labelled COOLER is the connection to the Ricor cooler motor. This allows the microcontroller to command the Ricor cooler, to adjust the temperature inside the cryostat.

To the right of COOLER, J14ISP is a plug that allows the microcontroller to be programmed. Below this are a number of capacitors to damp out noise. Finally, the circuits on the lower right connect the chip to the black body resistance temperature detector.

3.9 Winterisation

NISM has to survive the extreme temperature at Ridge A, which often drops below -75°C during winter, requiring a number of steps to be taken to protect it from the cold. The detector and motor are both in housings insulated with phenolic foam, and have heaters to keep them above -20°C . They are thermally insulated from the metal frame of NISM by a fibreglass baseplate. There is a small heater on the back of the mirror to keep it a few degrees above ambient temperature, which prevents ice forming. The cables connecting the different sections of NISM are coated with Teflon, as a standard plastic coating would go brittle and crack in the cold temperatures.

3.10 Black body

On the back of the instrument is a black body, to act as a reference brightness source (the red cylinder in Figure 3.1). It consists of a cylinder of copper with a cone machined into the end, inside an insulated housing. The copper was coated with a layer of copper oxide using commercial Ultra-Blak 420, to make it as black as possible in the infrared.

The exposed surface of the cone is facing the mirror. Heaters warm the copper body to $\sim 30^{\circ}\text{C}$ above ambient, and platinum resistance thermometers closely monitor the temperature.

Figure 3.15 shows a simulation of the temperature of the black body while it is operating. The outer housing is close to ambient temperature, while the inner cone is a uniform $\sim 30^{\circ}\text{C}$ warmer. The insulation and heating are carefully designed so that the cone shows minimal variation in temperature across its face. This should give it a uniform brightness, making it an excellent black body reference. The simulated black body temperature model in Figure 3.16 shows that the temperature variation across the exposed face of the black body is $\sim 0.05^{\circ}\text{C}$, making it extremely uniform.




Figure 3.15: Temperature analysis of NISM's black body. [Credit: Nicholas Staszak, ANSYS, 2014]

3.10. BLACK BODY




Figure 3.16: Temperature gradient across the face of the black body. [Credit: Nicholas Staszak, ANSYS, 2014]

3.11 Installation

NISM was installed in the 2014/2015 maintenance mission to Ridge A, which I was fortunate to be a part of. The HEAT/PLATO-R group tries to send a team every year to do maintenance. We replace the broken engines, refuel them, upgrade the electronics, swap out the cryostat for HEAT, and install new instruments.

The 2014/2015 mission team consisted of 2 people from UNSW and 2 people from UoA. The missions are supported by the USAP, so we travel to Antarctica via a flight on a Hercules, from Christchurch to McMurdo Station. More information about this trip can be found in my blog⁵.

At McMurdo Station we met up with a field guide, who looks after the science team in the field. It is their responsibility to cook food, melt snow for drinking water, keep in radio contact with the station, and keep an eye on the health of the team. At McMurdo we also collected our camping gear and supplies, and undertook a field training course to learn how to live at a remote field camp.

We travelled from McMurdo Station to the South Pole on another Hercules flight. At the South Pole we spent some time preparing for our mission to Ridge A. We worked on pieces of equipment, charged batteries, planned flights and cargo, and pumped the cryostat down to a vacuum. The time at the South Pole is also crucial for acclimatisation to the high altitude. Due to poor weather and mechanical problems with the aircraft, we were delayed for several weeks.

We eventually managed to secure a Twin Otter flight out to Ridge A on the 1st of January. The first flight out is a reconnaissance flight, to check the conditions at Ridge A. The pilots fly several fuel barrels and other cargo in to the site, and fly out empty. This minimises the take-off weight of the plane, in case the skiway is in poor condition. For our trip the skiway was usable, so the second flight brought us in on the 7th of January, along with some cargo. For safety, the pilots stayed until we had set up tents, made some water by melting snow on a stove, and radioed the station.

We spent four ‘nights’ at Ridge A. The first day was spent setting up camp.

⁵<https://coldsdkies.wordpress.com/>

3.11. INSTALLATION

We each pitched our own tent in the snow, and set up a kitchen tent where our field guide could make water and cook meals. The next four days were spent working on HEAT and PLATO-R. There are a lot of tasks that need to be done with a small team. The first one was digging snow away from the instruments. The wind-blown snow piles up against everything (see Figure 2.9), so a lot of effort has to be spent clearing it away.

For PLATO-R Engine Module there are the tasks of removing the old engines, installing the new ones, and refuelling the tanks. New electronics need to be installed in the Instrument Module, and new instruments attached to the outside, such as NISM. On HEAT, the old cryostat must be removed, and the new one installed.

After any components have been changed, testing needs to be done to check that everything is working. We only have this one week in the field to do maintenance, so everything has to be completely operational by the time we leave.

On the fourth day, a third flight brought in some more fuel barrels, and took back the old engines. On the fifth day, the fourth and final flight picked the team up and brought us back to the South Pole.

Ridge A is a very difficult environment to work in. At ~ 4600 m pressure altitude, the air pressure is $\sim 50\%$ of that at sea level. Physical exertion is exhausting, and Ridge A requires shovelling snow away from instruments, lifting engines, and pushing fuel barrels around. This is made even more difficult by the weight of clothing that needs to be worn.

We travel there at the warmest time of year, but the temperature is still below -35°C . Extreme cold weather clothing is required to be able to work in this environment. Additionally, we do not really have any other sources of heat. Food is cooked over a camp stove, and small items such as laptops and bread can be placed in the IM to warm up, but the team has to rely on their clothing to keep warm. It is necessary to keep your water bottle inside your jacket to stop it from freezing. The tents warm up slightly in the sun, but it is still below freezing inside, and heavy duty sleeping bags are required. The high altitude also makes sleeping difficult.

The five day length of the mission is the shortest that gives us enough time to complete work on PLATO and HEAT, so that we minimised the time spent in these extreme conditions. In future, if there are more telescopes installed at Ridge A, it would be practical to set up some sort of heated cabin or hut that the team can stay in, to temporarily get out of the cold. This would allow the mission time to be extended, so more work could be done.

NISM was installed successfully and operated over the 2015 winter, collecting measurements of the sky brightness. In the next chapter we simulate the atmosphere over Ridge A, to predict what values NISM might observe. Then Chapter 5 begins the analysis of the data that NISM sent back.

Chapter 4

Atmospheric modelling

Models of the water vapour over Antarctica and satellite measurements predict that Ridge A should have the lowest levels of PWV of anywhere in the world (Saunders et al., 2009). This should in turn give Ridge A the lowest atmospheric infrared background emission of anywhere in the world. We can predict the infrared background we would expect to see by modelling the atmosphere, and solving the radiative transfer equation.

In this chapter we use the program VSTAR (Bailey & Kedziora-Chudczer, 2012) to model the atmosphere at Ridge A. By solving the radiative transfer equation, VSTAR can calculate a prediction for the brightness we would expect to observe with NISM.

The radiative transfer equation models the behaviour of radiation passing through a medium. It includes the effects of the three kinds of interactions that can take place in the medium: absorption, emission, and scattering. A complete discussion of radiative transfer can be found in texts such as Liou (2002).

One form of the equation of radiative transfer is given here:

$$I_{\lambda}(s_1) = I_{\lambda}(0)e^{-\tau_{\lambda}(s_1,0)} + \int_0^{s_1} J_{\lambda}(s)e^{-\tau_{\lambda}(s_1,0)}k_{\lambda}\rho ds \quad (4.1)$$

Where λ is the wavelength, s is the distance along a ray of light, I is the intensity, τ is the optical depth, J is the source function, k is the mass extinction cross section, and ρ is the density of the medium.

Picture a beam of radiation with wavelength λ passing through a medium (in

our case, this will be the atmosphere). The beam has an initial intensity $I_\lambda(0)$. We would like to know $I_\lambda(s_1)$, the intensity when it has travelled a distance s_1 .

The first half of Equation 4.1 shows this initial intensity being exponentially attenuated as it passes through the medium. The rate of attenuation is determined by the optical depth parameter $\tau_\lambda(s_1, 0)$. This optical depth is determined by the density of the medium ρ , the mass extinction cross section of the medium k_λ , and the distance travelled through the medium, from 0 to s_1 . τ_λ is defined as:

$$\tau_\lambda(s_1, s_0) = \int_{s_0}^{s_1} k_\lambda \rho ds \quad (4.2)$$

In addition to attenuation, the beam of radiation can be brightened due to emission and multiple scattering. The second half of Equation 4.1 covers this. $J_\lambda(s)$ is the source function. This is the increase in intensity of the radiation due to emission and scattering at the position s . Each layer s acts as its own source of radiation, so each emission $J_\lambda(s)$ is exponentially attenuated as it travels to s_1 , just as in the first half of the equation.

This emission is then integrated over the path, from 0 to s_1 , to find the total contribution. Note that $k_\lambda \rho ds = d\tau$, so that we are effectively integrating over the optical depth rather than the distance, as it is the relevant parameter for this situation.

4.1 VSTAR

VSTAR is a piece of software that solves the radiative transfer equation through an atmosphere. It was written by Professor Jeremy Bailey at UNSW for the purposes of modelling exoplanet atmospheres. It models the atmosphere as a number of plane parallel layers, and calculates the molecular absorption for each molecular species present, as well as scattering from clouds, aerosols, and Rayleigh scattering. A more complete explanation of the software can be found in Bailey & Kedziora-Chudczer (2012), but a brief overview is given here.

VSTAR was designed for ground-based observations of planetary atmospheres in the solar system. Its main purpose is to model a target planet's atmosphere, but

it is also necessary for it to be able to model the Earth’s atmosphere, to account for the emission and absorption from terrestrial atmospheric species that will be present in ground-based observations. As such, VSTAR was designed to be very versatile. It can model a range of atmospheres at a wide range of temperatures, from exoplanets to brown dwarfs, and works at wavelengths from microwave to ultraviolet.

Comparisons with other models and with observations show that VSTAR accurately models atmospheric emission (Bailey & Kedziora-Chudczer, 2012), hence it can provide us with an accurate estimate for the infrared emission at Ridge A, given a suitable model of the atmosphere.

4.1.1 DISORT

VSTAR implements the DISORT package to solve the radiative transfer equation. DISORT is an implementation of the DIScrete Ordinate Radiative Transfer method (Stanmes et al., 1998). This method approximates the atmosphere as a number of discrete layers.

Rather than performing the integration of the radiative transfer function over all angles, DISORT uses a Gaussian quadrature approximation. It performs the calculation for a number of specific quadrature angles (streams), and sums them with appropriate weights. By carefully choosing the positions and weights, the Gaussian quadrature rule can produce a very accurate approximation for a function that is close to polynomial.

By default, VSTAR uses 8 streams: 4 up and 4 down. This reduces computation time while still producing an accurate model. If more accuracy is required, the number of streams can be changed. VSTAR will then select that number of angles according to the Gaussian quadrature approximation. The default 8 stream model was sufficient for our needs.

To run the DISORT method on an atmosphere, VSTAR creates a two dimensional model. The first dimension is the height of the atmosphere, and the second dimension is the wavelength of radiation that will be transferring through the atmosphere. To generate this model two inputs are needed. Firstly, a vertical

profile of the atmosphere, with values for the height, pressure, temperature, and gas volume mixing ratios of each layer. Secondly, a range of wavelengths that the radiative transfer function will be solved for, including a grid step size.

Along with the atmospheric profile, VSTAR needs to be passed a model of the albedo of the surface of the planet. This is simply a file with the surface albedo given for a range of wavelengths.

4.1.2 HITRAN

VSTAR can read in molecular line data from a range of sources. The primary catalogues used are HITRAN and GEISA. These have extensive data on molecular emission and absorption data for species in the Earth’s atmosphere. For other atmospheres at higher temperatures, such as brown dwarfs, VSTAR can use additional catalogues, such as HITEMP. A full list of the catalogues used can be found in Bailey & Kedziora-Chudczer (2012). In our case we are simply modelling the Earth’s atmosphere at low temperatures, so HITRAN contains the lines we need.

VSTAR is a modular program, which allows it to be adapted for a range of atmospheres. These are the primary modules:

- MOD—generates the two dimensional atmospheric model.
- LIN—calculates molecular line emission and absorption using a catalogue such as HITRAN.
- RAY—implements Rayleigh scattering from gas molecules.
- PART—implements particle scattering from clouds and aerosols.
- RT—solves the radiative transfer equation using information from the previous packages. Outputs the results.

The amplitude of Rayleigh scattering is inversely proportional to the square of the wavelength. Since we are looking at (relatively long) infrared wavelengths in our particular model, Rayleigh scattering is negligible. Thus, we do not need to

implement the RAY module . Additionally, the atmosphere in Antarctica is extremely clear with very few aerosols (Tomasi et al., 2007; Ashley, 2012), so we do not need to use the PART module either.

The RT package can output three files. Two of them contain the fluxes and radiances observed; one file for top of the atmosphere, and one file for at the surface. The other file has the optical depth of the atmosphere (τ), and the total transmission through it ($e^{-\tau}$).

4.2 Ridge A atmospheric model

To use VSTAR to model the expected emission at Ridge A, we need a model file containing the atmospheric profile. This file contains a two dimensional array, with each row being one layer of the atmosphere. The columns contain the data for that layer: height, pressure, temperature, and the fractional volume mixing ratios of the species present in the atmosphere. For our Earth atmosphere model, we include H₂O, CO₂, O₃, N₂O, CO, CH₄ and O₂.

However, there have not been any measurements taken of the atmospheric layers above Ridge A. Thus, it is necessary to create a model based on other measurements that are available. Fortunately, a number of measurements have been taken from sites around Antarctica, including the South Pole and Dome A. Additionally, VSTAR has a default atmospheric model for high latitudes, labelled Sub Arctic Winter.

Shi et al. (2016) constructed a model of the atmosphere above Dome A, shown in Figure 4 of their paper. Dome A is only 144km from Ridge A, so should have a very similar atmospheric profile. Their model is mostly based on satellite measurements from the NASA MERRA reanalysis (Modern Era Retrospective-analysis for Research and Applications). However, these measurements do not have sufficient resolution to resolve the temperature inversion close to the ground. Thus, they combined the MERRA data with their own ground-based measurements of the inversion layer. Their model shows the temperature and H₂O mixing ratio as a function of air pressure.

However, the model from Shi et al. (2016) does not include the altitude, which is the required parameter for VSTAR. Since the atmosphere above Antarctica is so cold, it contracts down, and the pressure at the high plateau is lower than at equivalent temperate sites. Thus, we cannot use the standard pressure-altitude model of the atmosphere.

Fortunately, there are numerous measurements of the atmospheric profile taken from the South Pole. NOAA (the US National Oceanographic and Atmospheric Administration) regularly launches radiosonde balloons from the South Pole. These balloons rise up through the atmosphere, recording measurements of the pressure as they go. Data are available from their website¹.

To generate our model of the Atmosphere above Ridge A, we used pressure data from the Ozone Sonde, launched from the South Pole on 01 August 2016. We used this pressure/height relationship to determine the temperature and H₂O mixing ratio as functions of height, based on the Shi et al. (2016) model. These parameters are shown in Figure 4.1

VSTAR also requires mixing ratios for other molecular species in the atmosphere. For these, we used the default settings in VSTAR's Sub Arctic Winter profile. These other molecular species are evenly distributed across the atmosphere, so Ridge A would not be expected to have any differing molecular ratios. We did make one change to the default: since the development of the software, the concentration of CO₂ in the atmosphere has risen to 400 ppm.

Other measurements of the atmosphere in Antarctica are available. Tomasi et al. (2011) presents measurements of the atmosphere above Dome C, taken between 2005 and 2009. However, Dome A is really the only site that is close enough to Ridge A to have similar temperature and water vapour profiles. Hence, we based our model mainly on the Dome A measurements from Shi et al. (2016).

The parameters for the model we used can be seen in Figure 4.2. The temperature inversion is clearly visible at ~ 500 m. This inversion is associated with an increase in the concentration of water vapour at that height. Since the air is warmer, the amount of water vapour it can hold also increases.

¹<https://www.esrl.noaa.gov/gmd/dv/data/index.php?site=spo&type=Balloon>




Figure 4.1: Temperature and water vapour profiles in the Ridge A atmospheric model used for VSTAR. The temperature (and thus emission) are highest in the lowest layers, where there are few data points.

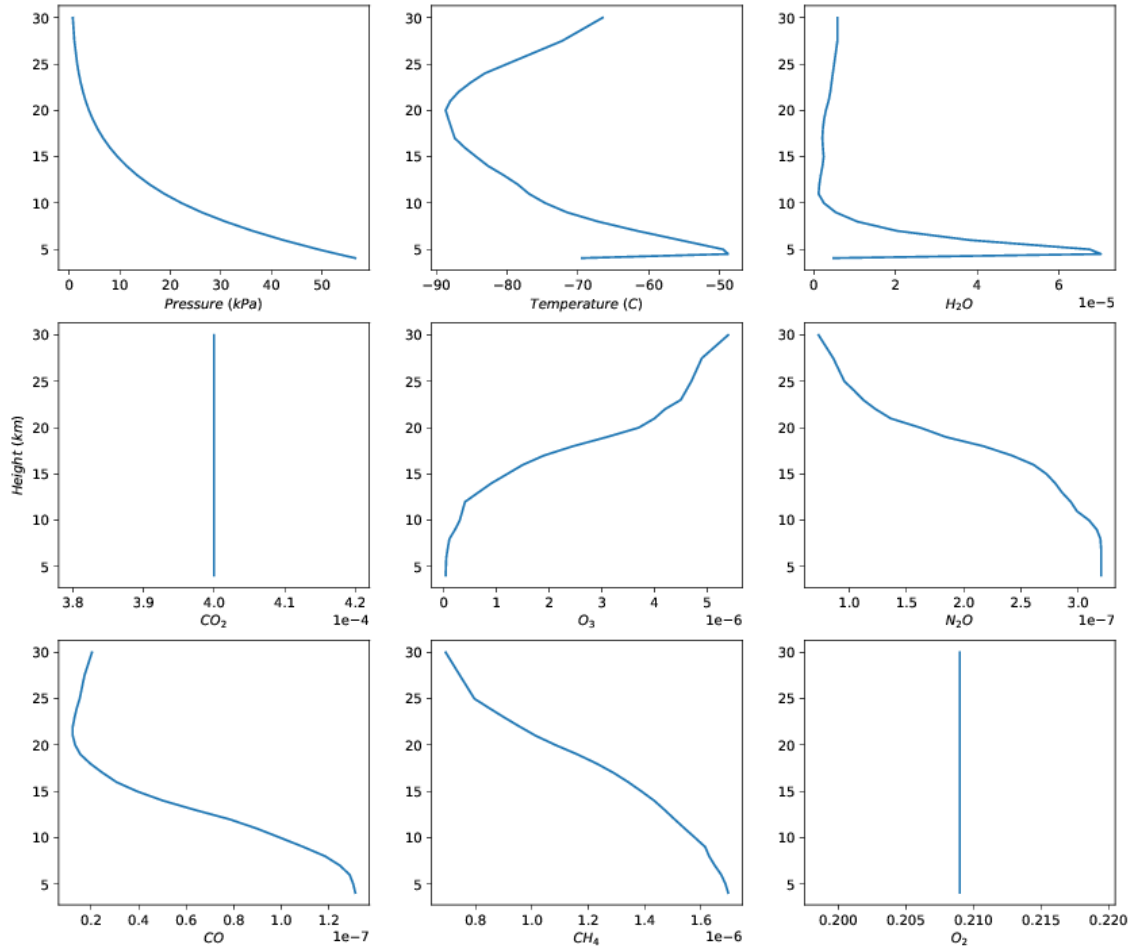


Figure 4.2: All parameters in the Ridge A atmospheric model used for VSTAR: pressure, temperature and atmospheric mixing ratios.




Figure 4.3: Atmospheric spectral radiance as modelled by VSTAR at Ridge A, at a zenith distance of 22° , using the input atmosphere in Figure 4.2

The final input required by VSTAR is the surface albedo. In our particular case the albedo has very little effect. The surface can emit and reflect radiation, but due to the negligible scattering in the atmosphere, none of it will be scattered back down to the instrument. As such, we set the albedo to 1 at all wavelengths.

4.3 VSTAR results

With the input atmospheric profile defined, it is a simple matter to run the VSTAR program to obtain a model of the emission. The output of VSTAR is shown in Figure 4.3. This plot shows the spectral radiance of the atmosphere over the range of wavelengths that NISM observes.

VSTAR outputs the radiance at eight angles, from the eight stream approximation, and we use the four positive ones. The output angles are labelled

as 0.93, 0.66, 0.33, and 0.06. This number is the cosine of the angle from the zenith, so the four streams are at zenith angles of 21.6° , 48.7° , 70.7° , and 85.5° . The 0.93 stream is the closest to the zenith, at 21.57° away, so that is the stream we used. Since the cosine of this angle is 0.93, the emission at the zenith should be $\sim 93\%$ of that at this angle (based on path length through the atmosphere, and assuming that the lines are not saturated).

The main feature visible in Figure 4.3 is the significant increase in radiance towards longer wavelengths. This is due to the black body emission from the atmosphere peaking at a longer wavelength. The wavelength of maximum emission for a black body is given by Wien's displacement law: $\lambda_{max} = b/T$, where $b = 2900 \mu\text{mK}$. For an atmosphere at $\sim -70^\circ\text{C}$, this corresponds to a peak in emission at $\sim 14.5 \mu\text{m}$.

However, this is clearly not a smooth black body curve. The forest of lines are due to molecular emission and absorption. At wavelengths where a line is present, the atmosphere has a brightness similar to a black body. At wavelengths where there are not lines the atmosphere is close to transparent, and the emission is very low.

VSTAR also outputs the transmission through the atmosphere at each wavelength. This is shown in Figure 4.4. The average transmission over the NISM wavelengths is 97%.

4.3.1 NISM's passband

NISM has a filter that limits its view to the K_{dark} window at $2.4 \mu\text{m}$. The passband of this filter is shown in Figure 3.10. This plot does not include the effect of the blocking filter, which has approximately 75% transmission between 2.1 and $2.5 \mu\text{m}$. It is likely that this filter cuts off some of the radiance longer than $2.5 \mu\text{m}$. The exact passband of the blocking filter is not known, so it is not included in the calculations in this section.

Multiplying the atmospheric emission in Figure 4.3 by NISM's passband produces the weighted radiance plot in Figure 4.5. This shows the radiance of the atmosphere viewed through NISM's filter, at an angle of 21.6° from the zenith.




Figure 4.4: Atmospheric transmission as modelled by VSTAR for $\cos(z)$ of 0.93, using the input atmosphere in Figure 4.2. The average transmission over this range is 97%.




Figure 4.5: The modelled atmospheric spectral radiance at Ridge A, at a zenith distance of 22° , weighted by NISM's passband. Modelled by VSTAR

4.3. VSTAR RESULTS

Table 4.1: Ridge A atmospheric spectral radiance through NISM’s passband, as calculated by VSTAR.

Zenith angle (Degrees)	Radiance ($\mu\text{Jy arcsec}^{-2}$)
22	3.8
49	4.8
71	8.0
86	22.0

Integrating over the weighted radiance in Figure 4.5 gives the total radiance that will be measured by NISM. Since VSTAR outputs four different angles, there are four total radiances. These are displayed in Table 4.1.

As expected, the radiance is lower at angles closer to the zenith, as this is the shortest path length through the atmosphere. Close to the horizon, the airmass is much greater, and hence the emission is much greater.

4.3.2 VSTAR prediction

As we have seen in the last section, the VSTAR model predicts that the zenith emission that NISM would measure is $3.8 \mu\text{Jy arcsec}^{-2}$. This is an extremely low number. Measurements from the South Pole find values around $\sim 100 \mu\text{Jy arcsec}^{-2}$ (Phillips et al., 1999). Ridge A is expected to be lower, but not this low.

It is likely that this model is over-estimating how good Ridge A is. One atmospheric effect that is not accounted for is diamond dust: small ice crystals suspended in the air. These could be a source of scattering, although they are expected to be negligible. Smith & Harper (1997) took measurements of the mid-infrared sky brightness at the South Pole, and noted that the ‘ice haze’ increased the brightness in the N band ($8\text{--}13 \mu\text{m}$) by $\sim 16\%$, which is a relatively minor effect.

The atmospheric model does not have much resolution for the lowest layer of the atmosphere. This layer is the densest, and has the highest temperature (at the inversion layer). Interpolation between the layers in the model may be giving

incorrect values for the atmospheric parameters.

Perhaps more significantly, this model does not include airglow. Emission from excited OH molecules in the upper atmosphere is the primary source of brightness at wavelengths short of 2.4 microns (Saunders et al., 2009). In the K_{dark} band, airglow and thermal emission are expected to have similar magnitudes. By excluding airglow in our model, we may be missing out on the major source of brightness.

It is also possible that auroral emission has some effect. The aurora mostly emits in visible wavelengths, but it may have some effect in the infrared too. This topic is discussed further in Chapter 8.

Overall, it is unlikely that the atmospheric emission at Ridge A is as low as $3.8 \mu\text{Jy arcsec}^{-2}$. However, this result should be taken to indicate that at least the thermal emission from the atmosphere is very low, and that Ridge A should be an excellent site. This makes an actual measurement from NISM all the more important.

4.4 Contribution of different gases

VSTAR allows us to investigate the contributions the various atmospheric molecules make to the radiation emitted. By setting the mixing ratio for a gas to zero, that gas will not be used in the model, and we can see the emission from the remaining molecules in that atmosphere.

Figure 4.6 shows the atmospheric emission when the level of water vapour in the atmosphere is set to zero. By comparing this with Figure 4.3, we can see a massive decrease in emission shorter than ~ 2.65 microns. This is essentially repeating the core idea of this experiment: water vapour is a primary contributor to the opacity of the atmosphere at infrared wavelengths, and a site with very low water vapour could become a good infrared observatory.

Figure 4.7 shows the atmospheric emission when CO_2 is removed. This spectrum is very similar to the original spectrum in Figure 4.3, but we no longer see the dip in emission just before 2.7 microns. This would suggest that CO_2 is

4.4. CONTRIBUTION OF DIFFERENT GASES

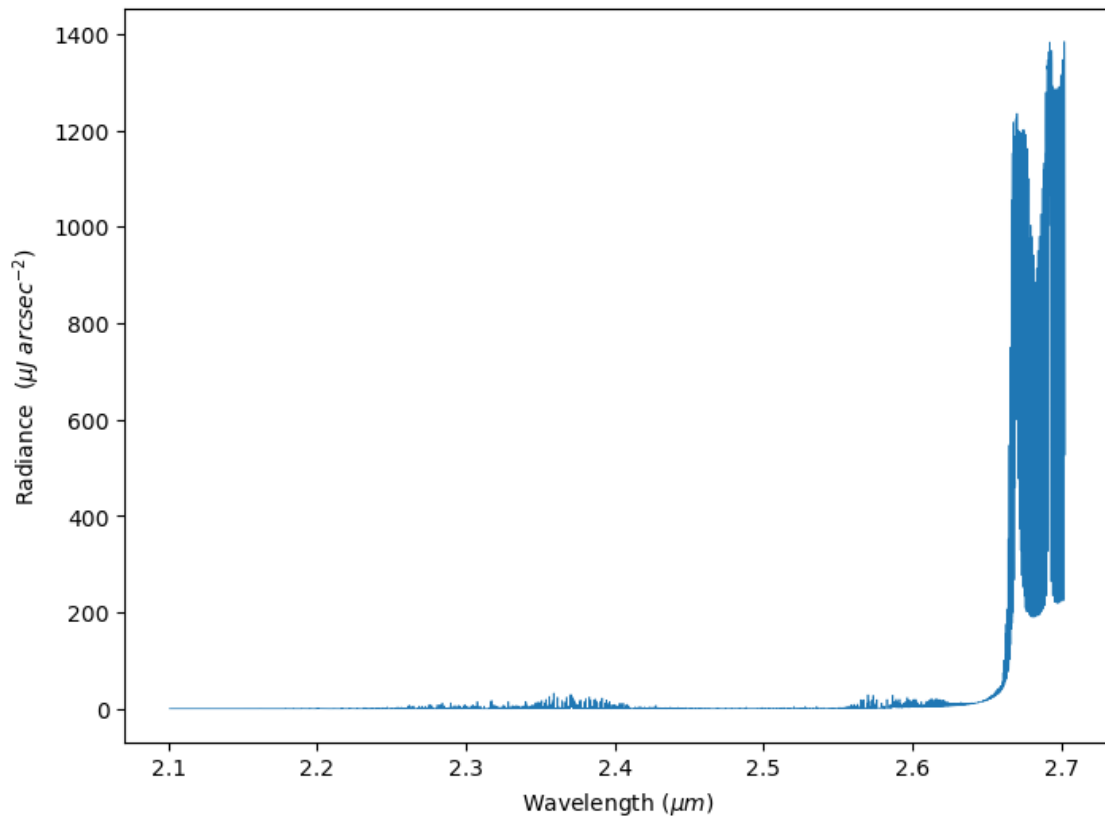


Figure 4.6: The modelled atmospheric spectral radiance at Ridge A, at a zenith distance of 22° , with no H_2O in the atmosphere. Modelled by VSTAR.

absorbing or scattering emission at this wavelength. However, since CO_2 is pretty uniformly distributed across the atmosphere, we would not expect Ridge A to have any significant difference to other sites with respect to the CO_2 emission.




Figure 4.7: The modelled atmospheric spectral radiance at Ridge A, at a zenith distance of 22° , with no CO_2 in the atmosphere. Modelled by VSTAR.

Chapter 5

Data analysis

The purpose of NISM is to measure the zenith spectral radiance during the winter at Ridge A. As we have seen in Chapter 4, this radiance is expected to be very low, making Ridge A an excellent location for infrared astronomy.

NISM was installed at Ridge A in January of 2015 (see Section 3.11), and took its first set of measurements during that winter. Unfortunately, there was a problem with the modem during the 2016 season, so no data could be collected that year. A repair mission could not be sent over the 2016/2017 summer, so no data could be collected during 2017 either.

The raw data from NISM were returned to UNSW via an Iridium satellite connection. The data were in a very raw state, with simple lists of detector readings. This chapter covers the processing of the raw NISM data into physical units, and the challenges that were encountered during this processing.

5.1 Binary packets

Data are returned from NISM in binary files, each containing data from one day. These files can contain “short packets”, “long packets”, and ASCII text. The Perl script `proc-ns.pl` (written by Michael Ashley) extracts the short and long packets from a binary file and writes them out as ASCII text. The raw binary files are referred to as Stage 1 in my scripts, and the extracted ASCII files are referred to as Stage 2.

Table 5.1: The contents of the short packets

Index	Parameter	Description
0	time	UNIX epoch in seconds when the data packet was received
1	short_time	Time -1428000000.0
2	sequence	Integration sequence number, should increment by one
3	motor	Actual motor position, in steps
4	dt	Length of the last integration; units $64.0/20e6$ seconds
5	insb	Last integration on the detector, in ADU
6	temp	Last integration of the detector temperature, in ADU

The short packets contain the primary data from NISM, and the long packets contain additional information on the state of the instrument. The parameters in the short and long packets are listed in Tables 5.1 and 5.2 respectively.

The primary data are the parameters ‘dt’, ‘insb’, and ‘motor’, returned in the short packets. These are the integration time length, integration value, and motor position respectively. The motor position, modulo the number of steps per revolution, gives the direction that NISM’s beam is pointing, and the measurement of the sky brightness is found by dividing the integration value by the integration time. This returns a measurement of the flux hitting the diode in ADU per second (ADU = analogue digital unit, the digitised readout from the diode). This value can be converted into physical units by using some suitable calibration source.

The signal measured by the diode will be the sum of the infrared emission from a number of sources, including astronomical sources, emission from the atmosphere, as well as emission from the mirror, from the filter optics, and other internal thermal emission. In fact, the majority of the signal seen will be from instrumental sources, rather than the desired atmospheric emission.

An example of the raw data can be seen in Figure 5.1. It shows the measured signal varying over the course of a day. Most of this variation is due to the changing ambient temperature. This changes the temperature of the mirror and optical surfaces, and thus changes their thermal emission.

Zooming in on a section of the raw data in Figure 5.2, we can see the desired

5.1. BINARY PACKETS

Table 5.2: The contents of the long packets

Index	Parameter	Description
0	time	UNIX epoch when the data packet was received
1	short_time	Time -1428000000.0
2	enable bits	Indicates which devices are enabled: bit 0 Ricor cooler (0 = off, 1 = on) bit 1 Silvermax motor polling bit 2 DDC112 (the ADC used to measure the detector) bit 3 MAX31865 (ADC used for black-body temp)
3	ddc_test	0 for normal mode, 1 for test mode
4	ddc_range	The ADC range; 7 is most sensitive
5	ddc_period	If non-zero it is the manually-set period of the ADC
6	ricor_heater	Ricor heater setting; 0 = off, 65535 = fully on
7	motor_heater	Motor heater setting
8	mirror_heater	Mirror heater setting
9	bb_heater	Black-body heater setting
10	ricor_temp	Ricor cooler temperature; degrees C = $\text{temp} * 1100.0 / 1023 - 273.15$
11	mirror_temp	Mirror temperature
12	internal_temp	Temperature inside the detector housing
13	external_temp	Ambient air temperature
14	avr_time	Microprocessor time; 16-bit counter; units are $(0x10000 * 64) / 20e6$ seconds
15	bb_temp	Black-body temperature; degrees C = $\text{ADUs} / 32.0 - 256 - 273.15$
16	motor_target	Motor target position (8000 steps per revolution); 32 bit
17	motor	Motor reported position
18	sequence	Integration sequence number, should increment by one
19	dt	Length of the last integration; units $64.0 / 20e6$ seconds
20	insb	Last integration on the detector, in ADU
21	temp	Last integration of the detector temperature, in ADU




Figure 5.1: Raw data from NISM, from one day. Blue is the data, red is a baseline fitted to the data. The variation over the day is due to the change in temperature. The Sun is up at ~ 8 hours, producing the large range of brightnesses at that time.




Figure 5.2: A zoomed view of the raw data from Figure 5.1, showing the change in emission as the motor rotates. At this point in time the Sun was up, so there is a very high signal from the horizon. The zig-zag nature of the baseline is caused by the two channels of the ADC having different gains (See Section 5.2).

signal superimposed on the thermal variation. As the mirror rotates, it scans NISM’s field of view through 360° of elevation. It scans across the sky, starting from the black body reference source located at one horizon, going up to the zenith, and then continuing down to other horizon. It then scans across the ground, seeing the cold snow, and parts of NISM’s frame and mount. A measurement of the atmospheric opacity like this is referred to as a skydip, as the observations dip from a high elevation angle to a low elevation angle.

The most noticeable features in the data will be the horizon and the black body. The black body is roughly 30 K above ambient temperature, so emits much more thermal emission than the sky. Looking at the other horizon, we are looking through a large airmass, and so see a high signal from the atmospheric emission.

As NISM scans across the sky, the measured emission follows an approximately

$\sec \theta$ function ($\sec \theta = 1/\cos \theta$). This is due to the path length through the atmosphere: if we model the atmosphere as a horizontal slab, then the path length through the atmosphere at an angle θ away from the zenith is $\sec \theta$ times the vertical height.

This model assumes that the atmosphere is optically thin: that light can easily pass through it, and we see emission from all parts of the atmosphere along the line of sight. Realistically, the atmosphere will have some optical depth: nearer sections along the line of sight block the emission from more distant sections, so the measured emission will be lower than the $\sec \theta$ model. Also, the $\sec \theta$ function goes to infinity at the horizon (because we model the atmosphere as an infinite horizontal slab) which is clearly unphysical.

When looking at the ground, we expect the emission to be very low. The wintertime ambient temperature at Ridge A is $\sim -70^\circ\text{C}$. Since there is a temperature inversion, the ground is usually colder than the air, so snow below -70°C should give off very little emission in the infrared.

The exact emission is difficult to calculate. It depends on the snow temperature, as well as the grain size. The emissivity of snow is found to decrease as the grain size increases (Salisbury et al., 1994; Carlson et al., 1992). The emission from snow is the result of a combination of surface and volume scattering, and the relative effect of these two factors is dependent on grain size (Salisbury et al., 1994).

The grain size at Ridge A is difficult to determine, because it depends on a number of factors. In Antarctica, small snow grains settle on the surface. As more layers settle, pressure welds them together into larger grains, resulting in a stratified range of grain sizes (Brucker et al., 2010). The observed grain size then depends on how deep your wavelength penetrates into the snow.

Combined with the lack of snow measurements from the high plateau, and possible reflections from the snow surface, the precise emission cannot really be calculated. However, we can use an ideal black body as an upper limit. Based on the temperature, the emission must be very low to the point of being insignificant when compared to emission from the atmosphere. So it is a reasonably safe

assumption that the coldest observed signal from NISM when looking at the snow is approximately zero. This is further covered in Chapter 6.

There are also some sections of NISM’s frame and mounting visible when looking downwards. These sections have uncertain temperatures and emissivities, and may be highly reflective, giving them unpredictable brightnesses. We ignore these regions when doing the data analysis.

5.2 Baseline subtraction

The sky signal in the NISM data is superimposed on the signal from the ambient temperature variation. To separate these two signals we fit a baseline following the temperature variation, and subtract it from the data. This is done using the Python script `subtract_baseline.py`.

As described in Section 3.8, the analogue digital converter (ADC) in NISM has two channels, A and B. NISM alternates between using channels A and B, which allows continuous sampling of the diode’s signal. While one channel is digitising, the other can be sampling. As a result, all the data packets with an even sequence number are taken using one channel, and all the data packets with an odd sequence number are taken using the other channel. The two channels have slightly different offsets and gains, so the even and odd data packets are processed separately when subtracting a baseline.

The script `subtract_baseline.py` reads in the short and long data packets from the Stage 2 ASCII files. A copy of the long packets is stripped down to the primary data and combined with the short packets, so that we have a set of all the primary data points. Some initial data processing is then done. Any points where the motor position is greater than 10^9 are masked out, using a Numpy masked array. The motor position is an incremental count of the motor steps taken, increasing by 8000 every revolution. It normally increases from 0 to $\sim 10^7$ during one 4-hour observing cycle. At the end of a cycle, the reported number jumps a few orders of magnitude to some very large value. This appears to be some integer overflow effect, and does not represent the real motor position, so these points

should be masked out.

The script also masks out any points where the motor is not moving. If the motor is not moving, then there is no sky dip to be seen; the measured signal is purely thermal variation in the detector. This contains no useful data, and so should be masked out. The script does this by comparing the position of each measurement to the position of the previous one. If it has not changed, then the motor is not moving and the data are masked.

Next, the script finds each rotation of the motor. It does this by looking at the motor position mod 8000. The motor has 8000 steps per revolution, so this calculation gives the motor's actual angular position. The motor advances ~ 90 steps between measurements. Whenever the motor position (mod 8000) decreases by ~ 8000 from one measurement to the next, it must have gone past 8000 and wrapped around to 0, indicating a rotation has been completed. The variable step size introduces some uncertainty in this value, so we set a broad limit of > 6000 , to catch all completed rotations.

Finally, we calculate the median time and the median measured value for each rotation of the motor, and create a new point at this average. The script then linearly interpolates between these points to create a baseline function. This baseline is subtracted from the data to give us the result that we want: data that shows how the signal changes with elevation angle, independent of the thermal background.

As there are two channels of the detector with slightly different gains and offsets, we create and subtract a baseline twice, once for the even points and once for the odd points. The even and odd points can then be combined together as a single baseline-subtracted set. The data in this format are referred to as Stage 3. A section of this baselined data can be seen in Figure 5.3.

5.3 Motor slip analysis

Now that we have the baseline subtracted from the data, we should be able to simply fold the data on a period of 8000 motor steps to get the skydip signal.




Figure 5.3: A full day of NISM data after subtracting a baseline.




Figure 5.4: Folding the NISM data in Figure 5.3 on its motor position mod 8000 does not give us the smooth skydip as expected. Instead we see a smeared out blur.

Unfortunately, this did not work. An example of the problem can be seen in Figure 5.4. This shows the measurements from one day (from Figure 5.3), folded on a period of 8000 steps. We should expect to see the characteristic $\sec \theta$ skydip shape, with some slight dispersion in the measured signal due to changing conditions during the day. However, the plot is smeared out, with no clear skydip visible. There are some peaks visible near 5000 and 1000, which are presumably the black body and the far horizon, but they should be clear and distinct, not swamped in a sea of points.

The apparent smearing of the data can be explained by errors in the calculation of the motor position. This calculation is only based on the motor having a period of 8000, therefore the motor is not taking 8000 steps in one revolution. This is a fairly major problem. We rely on the motor position to determine the elevation of NISM's beam. If we do not know the elevations of any of the measurements, then we can not calculate the brightness of the zenith.

5.3. MOTOR SLIP ANALYSIS

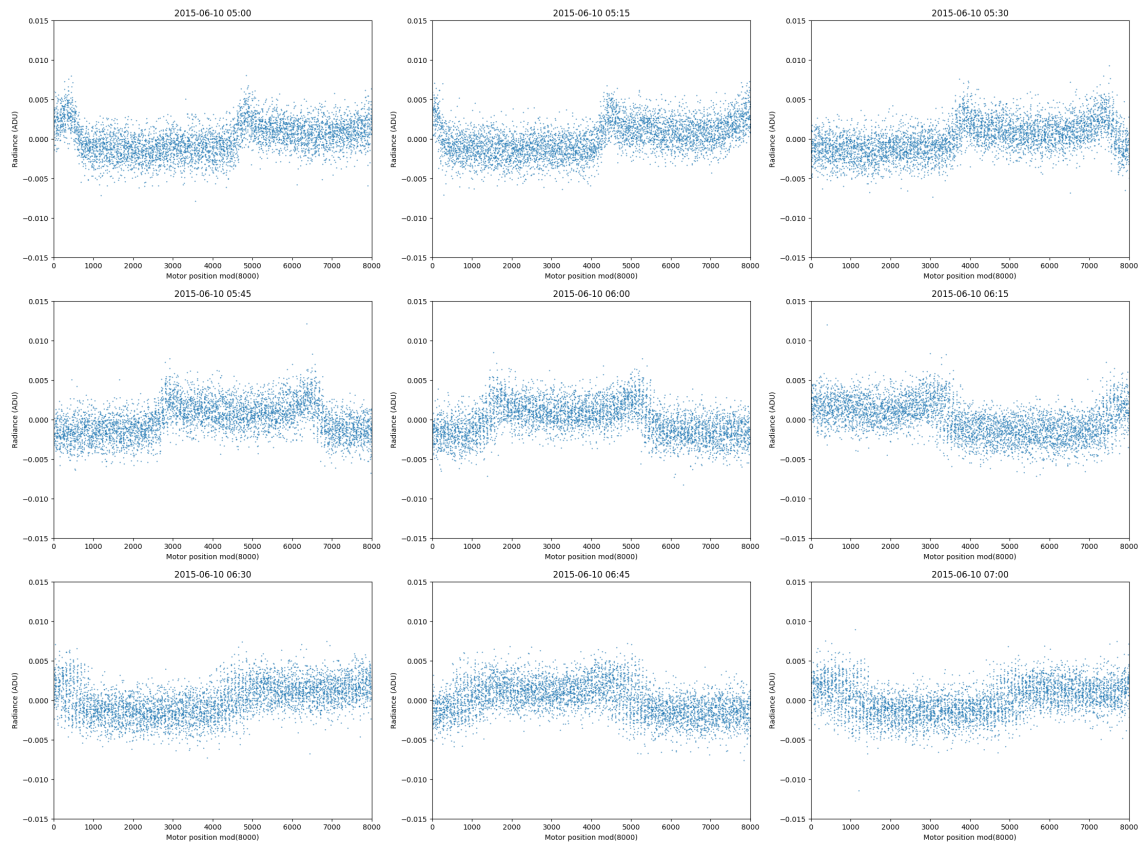


Figure 5.5: Five minute blocks selected from the data in Figure 5.4. The skydip is initially visible, but breaks up and flows away after some time. This indicates that the motor is no longer making 8000 steps per rotation.

More information on the problem is revealed in Figure 5.5. This shows several 5-minute blocks taken from the data in Figure 5.4 (timestamps are in the figure titles). In the initial blocks we can see the expected skydip shape. However, after some time the skydip shape begins sliding sideways and smearing out. This indicates that initially the motor works correctly, but after some time it begins slipping positions. Instead of counting 8000 steps in a rotation, it misses some, and will only have recorded, say, 7999 steps. This means that the period of rotation is now 7999. When we fold the data on 8000 steps, we are using the wrong period, and instead of seeing the regular skydip we just see aliasing effects.

The motor has an optical encoder with 8000 steps per revolution, which is used to determine the motor position. In the cold Antarctic conditions, we are running the motor outside its specified temperature range. It is possible that the optical encoder is therefore marginal, and occasionally misses steps.

The motor rotates once every 10 seconds. If it is only slipping one step per rotation (7999 rather than 8000), then after four hours that is 1440 steps, or 64 degrees away from the reported position. A small number of steps missed can have a very large effect.

Fortunately, this should be recoverable. The measurements are still sequentially in order, and the motor seems to be missing only a small number of steps out of the 8000, so it should be possible to determine the period of the motor by looking for features in the data.

There are a number of methods for finding a periodic signal in a dataset. These are often used in other areas of astronomy. For example, searching for transiting exoplanets by looking for a periodic dip in a light curve. Different methods can be used depending on the shape of the features being looked for. In the case of NISM, the features can not be described by a simple function. There will be a $\sec \theta$ curve from the atmosphere, a peak at one horizon from the black body, cold ice, NISM's metal frame, possible clouds, and a signal from bright astronomical sources, such as the centre of the Milky Way. Additionally, the data are not uniformly sampled in time. Therefore, simple period-finding functions such as a Fourier transform will not work well.

To search for this periodic signal, we use a method called Phase Dispersion Minimisation. In this method, first the data are folded on some period, and a smooth curve is fitted to the data. The χ^2 residual between the curve and the data points is then calculated. If the χ^2 is low, then the chosen period is a good fit for the data. The reasoning for this can be understood by comparing the curves in Figure 5.5 to the blur in Figure 5.4. When the period is correct, the data points follow a smooth curve, and the residual will be low. When the period is incorrect, the data points are smeared out, deviating away from a curve and producing a higher residual. By searching a range of possible periods and testing each one, it is possible to identify the correct period of the motor, and recover the NISM data.

5.3.1 Lowess smoothing

To generate a smooth curve to an unusually shaped dataset, we use a lowess (local weighted regression) fitting function. Our data does not show simple features, such as a sinusoidal variation or Gaussian peaks, so we require a fitting function that can generate a smooth curve for arbitrarily shaped data. A lowess function uses linear regression to fit a local area around each data point, and repeats this over all data points.

Since the amount the motor slips by varies throughout the day, we need to break the day up into sections, and find a lowess fit for each one. The sections need to be long enough to have a sufficient number of data points to be able to generate a good fit, but we also want them short enough to retain a good time resolution for the variation in the motor slip. Thus, we have broken each day's data into 5-minute blocks, and we fit a lowess curve to each one. This is a short enough time that we can assume that the motor slip is roughly constant over the 5 minutes. Each block has ~ 5000 data points, so we can fit an accurate curve. This then gives us an estimate of the motor slip every 5 minutes.

In Lowess fitting, a linear fit is generated for each point. First, a data point n is selected. A window of points is chosen around point n , by selecting a certain number of points either side. For our method we used a window containing 5% of the data. For ~ 5000 points, this is ~ 250 points, or ~ 125 either side of point n .

Since the x-axis of our data is an angle, it wraps around back to 0 at motor position 8000. If point n is close to the start or end of the range, the window selected wraps around to include points at the other end.

A linear fit is then generated to the points in the window, using a standard least squares polynomial fit provided by the NumPy function `polyfit`. This linear fit is then evaluated at the x position of the initial point n . The value of the fit at the x position of n now becomes the smoothed version of n ; it is a representative value of the adjacent points.

The lowess fitting procedure is then repeated for every point; creating a window and finding a linear fit each time. The final result is an array of smoothed points that follows the trend in the data. The advantages of the lowess function are that it returns a smoothed point for every input data point, and it can be used on arbitrarily shaped data.

5.3.2 The Fibonacci search algorithm

To search over a range of possible periods we use a Fibonacci search algorithm. A Fibonacci search is a recursive algorithm that efficiently finds the minimum or maximum of a function in some given range. It only searches integer values, and requires that the function be unimodal; that is, the function has a single minimum or maximum, and it increases or decreases monotonically on either side.

This is an appropriate method to use because the motor position is recorded as an integer value. Some other search methods look at the derivative of the function, but they require that arbitrarily small steps can be taken. Since the measured motor position is always an integer, these methods would not work.

The steps of the Fibonacci algorithm are:

1. Select two test points inside the search range, at positions of the two largest Fibonacci numbers that will fit. This splits the range into 3 regions.
2. Evaluate the function at these two test points. This is shown in Figure 5.6. Select the point with the lower value.

3. Since the function is unimodal, the minimum must be in one of the regions next to the point with the lower value. It could not be in the third region, because then the function would have low points on either side of a high point, and would not be unimodal. This is demonstrated in Figures 5.7 and 5.8.
4. This third region can be excluded, and the algorithm repeated on the new smaller range.
5. Eventually there is a single element left in the range, which will be the minimum of the function.

Since we are using Fibonacci numbers, one of the test points in the previous range will be carried over to the new range, and so does not need to be evaluated again. We only need to evaluate one new point at each step. This makes the Fibonacci search very efficient, especially when the function is slow to evaluate.

If the function were being evaluated over real numbers rather than integers, this algorithm can still be used, with the Golden Ratio replacing the Fibonacci numbers.

To use this algorithm, we assume that the number of steps being missed is less than 100, i.e., the period is the range 7900 to 8000. This range covers all the observed motor slips. We also assume that the function of residual vs period is unimodal over this range: that is, it has one minimum at the correct period, and is monotonic either side. Using an incorrect period smears the data out, so the further from the correct period we are, the more smear we expect to see. Therefore, the residual will increase monotonically as we move away from the correct period. If we were to extend the range to some factor of 8000, say, 4000, then we would expect to see the fit improve again due to aliasing effects. But this effect happens far away from our range of 7900–8000.

With these assumptions, we can apply the Fibonacci search algorithm to find a minimum residual. The fact that the Fibonacci search minimises the number of points evaluated is very useful for us, as finding the residual for each possible period is a slow process.




Figure 5.6: A Fibonacci search algorithm being carried out on some function $f(x)$. Two test points are chosen at x_1 and x_2 , such that F_1 and F_2 are the two largest Fibonacci numbers that will fit in the range. This divides the range into three regions a, b, and c. Since $f(x_1) < f(x_2)$, the minimum of $f(x)$ must be in region a or b. These possibilities are illustrated in Figures 5.7 and 5.8. The minimum cannot be in region c, so c can be excluded, and the algorithm repeated on the new smaller range




Figure 5.7: The minimum may be in region a.




Figure 5.8: Or the minimum may be in region b. But it cannot be in c.

We select two test points inside the range, at values of the two largest Fibonacci numbers that will fit. Initially, these will be at $7900 + 55$ and $7900 + 89$. These two test points break the range up into three regions. In the first stage, these regions will be 7900–7955, 7955–7989, and 7989–8000. We then use a lowess function to generate a smooth curve to the data folded on these two periods, and calculate the residuals.

Since the residual vs period function is unimodal, we know that the region between the period with the higher residual and the range edge does NOT contain the minimum. This region can be discarded from the search range.

This process can now be repeated with the new smaller range. Selecting two points at Fibonacci numbers that fit, and then removing the end near the higher residual. After repeating this several times, you will eventually be left with a single point inside the range, which is the correct period.

We ran this algorithm on all NISM data, to create a corrected data set with the right period of rotation. An example of a corrected 5-minute block is shown in Figure 5.9. The black body is visible as the bright peak on the left.

In the original design of NISM the black body completely fills the beam, so we would expect to see a flat top to the signal as the beam scans across the constant brightness. The measured beam is larger than expected, so instead we see a sharper peak: the black body only fills the beam when it is directly in the centre of view.

5.4 Correcting angle of elevation

Using the Fibonacci search algorithm and lowess fitting function, we can find the correct period of each 5 minute block of data. However, by folding the data on a period other than 8000, we no longer know which angle of elevation each motor position corresponds to. It was expected that 0 would be straight down, 2000 would be one horizon, 4000 would be the zenith, and 6000 would be the other horizon, and that we could simply take the motor position $\text{mod}(8000)$ to find the elevation. If we use the corrected period, calculating motor position $\text{mod}(\text{period})$




Figure 5.9: A 5-minute block of NISM data with the correct period. The green line is the lowess fit to the data. The black body is visible as the peak on the left, and the other horizon is visible on the right.

will give the data in the right phase, but not in the right locations. This is because the motor period is changing over time. If it was constant, then motor position $\text{mod}(\text{period})$ would work fine. However, the period changes while the motor is operating, so $\text{mod}(\text{period})$ is effectively folding all the data before the current 5-minute block by the wrong period, and hence the current block has an unknown starting elevation. Figure 5.10 shows the elevation angle drifting over one day.

To solve this, we can align each 5-minute block by searching for features in the data, and aligning based on those. We considered a range of possible methods for this alignment.

5.4.1 Aligning to previous block

One method is to align each 5 minute block to the block immediately before. By ‘sliding’ the data left and right, we can compare the lowess curves and minimise the difference. Each of our data files contains data from one day, broken up into 5-minute blocks. The first block of the day is usually aligned correctly, because the motor has just reset its position and the period should be correct.

However, this method has some problems. The first block of the day may not be aligned correctly, which then causes the later blocks to be misaligned. The other problem is that the alignment is not perfect. The lowess fitting is not exact: it is based on noisy data, and assumes that the motor period is constant in the 5 minute duration, which we know to be only an approximation. This means that each block can be slightly out of alignment with the previous one. Over the day, the alignment of the blocks can drift, and end up far away from the correct position.

5.4.2 Aligning to a reference block

We would prefer to have some fixed reference, so that every block is aligned in the same way. We considered two simple ways to implement this method. The first was to choose one block as a reference block, and align every 5-minute block to that one. The data are slid sideways until the difference between the lowess curves in

5.4. CORRECTING ANGLE OF ELEVATION

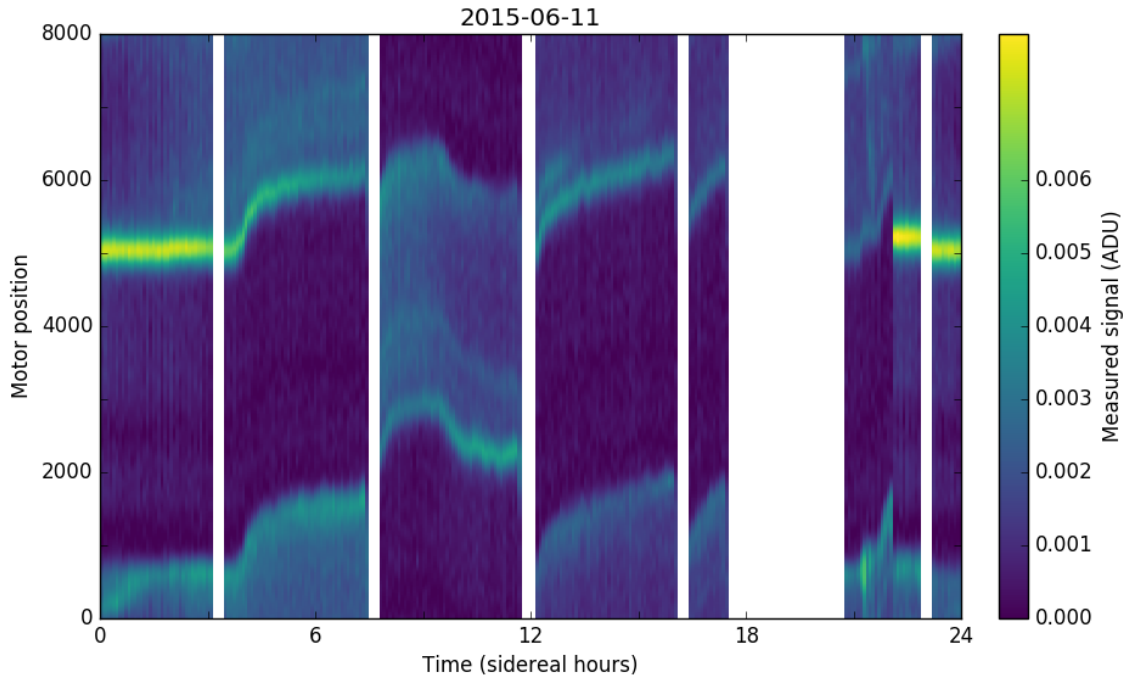


Figure 5.10: NISM data over one day before the alignment is corrected. Each 5-minute block becomes one vertical slice. White areas are times when the motor is not moving. The brightest region is the black body.

minimised, as before.

The second method was to look for the two bright peaks in the data, from the black body and the far horizon. Two windows are selected slightly less than 180° apart, and slid across the data till the values inside the windows are maximised.

Both of these had a similar problem: there can be significant variation in the shape of the signal, depending on the conditions. The elevation of the Sun and Moon, the temperature of the air and black body, and the position of clouds and the Milky Way all change throughout the year. So the particular block we are looking at could have a different shape to the reference block, or could have bright sections that are not exactly at the two horizons. These two methods provided inconsistent results, but the method of aligning each block individually is still preferred.

5.4.3 Fitting a skydip

The method we settled on was an improved version of looking for the two peaks. Instead of just looking at two windows, we generate a model of the skydip, and fit it to the observed data. This is more complicated, but can fit a wide range of possible skydips, and also returns some useful parameters.

The skydip model we used is based on the models featured in Phillips et al. (1999). Specifically equation (3), the model for airglow-dominated emission. In the K_{dark} band we expect airglow and thermal emission to both contribute, but based on the VSTAR modelling, airglow is expected to be stronger. The airglow-only model is also simpler, and can still provide a good fit to the observed skydip data. The equation we used for our models is:

$$F = A + f \sec \theta e^{-\tau \sec \theta} + G \quad (5.1)$$

In this equation, F is the sky spectral radiance, A is the constant background signal, f is the spectral radiance at the zenith, assuming no absorption, θ is the angle away from the zenith, τ is the optical depth parameter at the zenith, and G is the astronomical signal from the Milky Way.

This model of the atmosphere assumes that the observed emission is airglow that comes from a flat layer above the atmosphere. The emission from a section of this layer is proportional to the path length through the layer, which is given by $\sec \theta$.

This signal is then attenuated by its path through the atmosphere. This attenuation is given by $e^{-\tau \sec \theta}$. τ is the optical depth parameter of the atmosphere at the zenith, which is then multiplied by the path length through the atmosphere, $\sec \theta$. The atmosphere exponentially attenuates the signal as it passes through, resulting in the attenuation factor $e^{-\tau \sec \theta}$.

A is some constant offset due to the thermal background. In our baseline subtraction process (Section 5.2), we generate a baseline from the median points and subtract this from the data. This means that the median point in the data is now at zero.

We expect the lowest point in the data to be cold snow, which has close to zero emission in the K_{dark} band (see Chapter 6). The parameter A shifts the data so that the lowest point is at 0, rather than the median being at 0.

Our equation contains two modification from Equation (3) in Phillips et al. (1999). The first is the inclusion of an astronomical signal, given by G in Equation 5.1. NISM should be able to see bright astronomical sources, such as the Milky Way. This will contribute an extra signal when the Milky Way is in NISM's field of view. To find the expected signal from the Milky Way, we use infrared data from the DIRBE instrument on the COBE satellite (Silverberg et al., 1993).

Briefly, DIRBE measured the entire sky at infrared wavelengths. These data can be used to create a model G of the expected astronomical signal observed by NISM at any particular time. A more complete explanation of the DIRBE data can be found in Section 6.1

The other modification from Phillips et al. (1999) is the inclusion of the black body. NISM has a black body reference at one horizon, so when looking in this direction the sky signal is obscured. In this region, we replace the model of the atmospheric emission with a model of the expected black body emission. The black body has a uniform temperature, so has a uniform brightness across its face (see Figure 3.15). We model the black body in two dimensions, as a circle with a uniform brightness. We also include the mounting structure below the black body as an attached rectangle with a smoothly changing temperature gradient.

This model covers emission from the sky. For emission from the ground, we do not try to fit a model, and just use a constant value at the lowest observed point. When looking downwards NISM can see the snow, the Instrument Module, and sections of its own frame and mount. These all have uncertain positions, emissivities, and reflectivities. Since we do not know the expected emission from this region, we do not try to fit a model to it.

Our model does not include clouds, the Moon, or the Sun. Clouds are difficult to model exactly. They can have uneven distributions across the sky, and have an unknown brightness. However, the effect of cloud cover should be similar to a change in the optical depth of the atmosphere, so our model should still find a

good fit.

The Sun and the Moon are both bright, and produce a large signal, often saturating the detector. However, at Ridge A's latitude they will always be near the horizon, so will blend in with the skydip signal. More importantly, we are looking for dark skies when the Sun and Moon are down. Having a poor model fit when they are up is not a major problem.

5.4.4 Lmfit

To fit our model to the observed data points, we use the python package `lmfit`¹. This package implements the Levenberg-Marquardt algorithm, which solves non-linear least squares problems by following the steepest slope.

The `lmfit` package makes this algorithm very easy to implement. All the variables in the model can be passed to `lmfit` as `Parameter` objects. These objects contain a starting value for the variable, along with any other constraints, such as some values being limited to ≥ 0 .

These parameters can then be passed to the `lmfit` method, along with the model and the measured data points. `lmfit` will then iteratively generate models and compare them to the data points. It will take small steps in parameter space and follow the gradient, until it finds a location where the gradient is zero. This location will be a minimum, and should be a good fit to the data.

A disadvantage of the `lmfit` method is that it finds a local minimum near to the starting parameters. If there are several minima it can fail to find the global minimum. Fortunately, our model seems to be well behaved, and `lmfit` usually finds a good fit.

`Lmfit` does have some trouble finding a fit when the Sun or Moon is up. Since the Sun or Moon has a low elevation, we see a much brighter signal close to the horizon, as well as reflected off some of NISM's components. In these cases, `lmfit` usually fits this with a very bright atmospheric signal, with the $\sec \theta$ function giving a very high value near the horizon.

Occasionally `lmfit` will find a very poor fit to the data. These are usually times

¹<https://lmfit.github.io/lmfit-py/>

where the Sun or Moon is up, or the motor has stopped rotating, or some other problem has occurred. Since these times do not give us useful data, it is safe to exclude any times when the lmfit solution is poor.

We considered using a Markov-Chain Monte Carlo method (MCMC) to find a minimum, but eventually decided on lmfit. MCMC explores the parameter space to find a global minimum, but it takes a longer time than a simple gradient method like lmfit. Since we are fitting a large number of curves and have a fairly simple model, the faster lmfit should be better for our purposes.

5.4.5 Katana cluster

Our data processing takes some time. Generating fits for the Fibonacci search and running lmfit are fairly slow processes. To complete this processing in a reasonable time we made use of the Katana computing cluster at the UNSW Faculty of Science.

Katana consists of a cluster of 150 computational nodes, which contain a total of 2160 Intel CPU cores. Jobs can be sent to the cluster, and will be automatically allocated to nodes based on availability.

Our data are well suited for cluster computing. The data returned from NISM are broken up into 1-day files, which are each processed individually. We therefore write each job to process one 1-day file, and submitted all of these jobs to Katana. Katana then distributes these jobs as nodes become available, runs each job, and collects the results.

To process data on the Katana cluster we used the python script `fit_motor_model_align_katana.py`. This script takes the baselined data and implements the methods discussed in this chapter to correct for the motor slipping, to produce usable data.

It first calculates the motor slip by breaking the data up into 5-minute blocks, and running a Fibonacci search on each block. The Fibonacci search generates lowess smoothed curves for a range of different slip amounts, and searches in this range to find the lowest residual.

Then, it generates a model of the sky, and uses lmfit to vary the parameters to

find a best fit, in order to find the location of the zenith, and correctly align each 5-minute block.

The end result of this process is correctly aligned measurements from NISM—what we would have expected had the motor been running correctly. There are still some issues: the correct period is not always found, and the blocks still have slight misalignments, but the majority of the data are usable. An example of a day of data is shown in Figure 5.11. In this figure, the x axis is sidereal time, and the y axis is motor position, which translates to the elevation angle. Each 5-minute block of data becomes one vertical slice in this plot.

A number of features are visible in this plot. The bright horizontal band at motor position 2000 is the black body. It is promising to note that the observed brightness of the black body correlates with the temperature of the black body in the second subplot. The bright horizontal band at motor position 6000 is the opposite horizon. The region between these shows the sky dip, and the dark region outside these bands is the cold snow. The white vertical strips are times when NISM is not recording. The small regular gaps are the scheduled times when the motor resets its position to zero. The larger white gap around 18 hours is a time when the motor stopped rotating.

There is a noticeable discontinuity in the data at ~ 22 hours. This is a consequence of converting the plot to sidereal time. Each day of data recorded by NISM starts at 00:00 GMT. However, Figure 5.11 is plotted in sidereal time, so that astronomical features such as the Milky Way are always in the same place. On this particular day, 00:00 GMT corresponds to $\sim 22:00$ sidereal. So the discontinuity at 22 hours is in fact the start of the dataset for that day. A small dashed line in subplot 3 marks this time.

Possibly the most important feature in Figure 5.11 is the curve visible between 6 and 12 hours. This is the Milky Way, stretching across the sky. With NISM's fixed azimuth, only one part of the Milky Way is visible at a time. Over a day, as the Earth rotates, different sections of the arc of the Milky Way come into view, and it appears to rise and set. It then appears on the opposite horizon, and is visible as another arc around 18 hours, although much of it is obscured by the

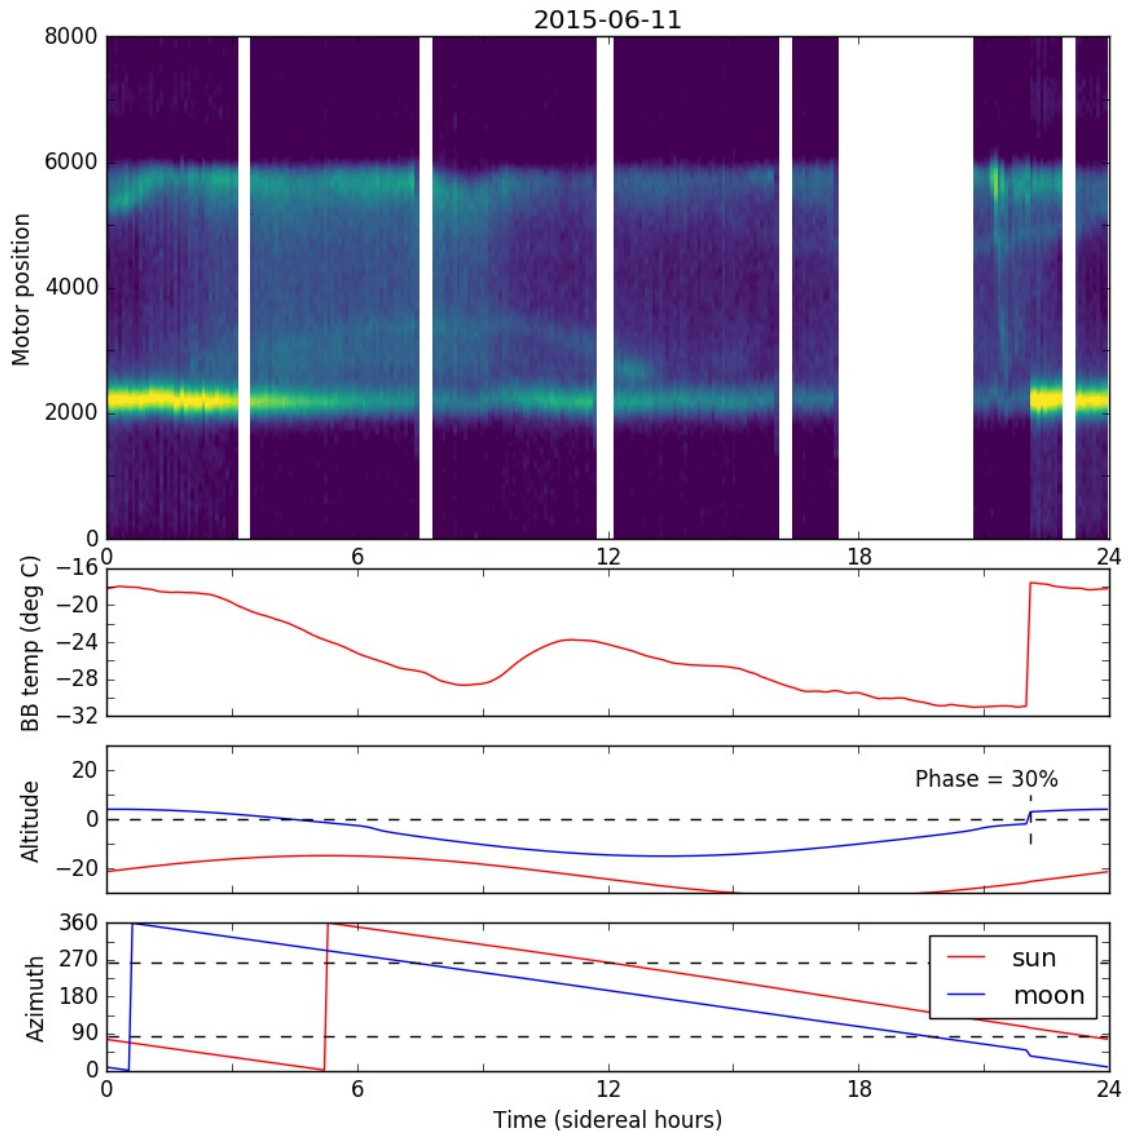


Figure 5.11: NISM data from 1 day after being corrected and aligned. Subplot 2 shows the temperature of the black body, and subplots 3 and 4 show the altitude and azimuth of the Sun and Moon.

motor stopping at that time.

Of the 88 days of data collected during 2015, 35 days proved to have usable data. On these days the Sun and Moon were below the horizon, the motor was operating for a significant fraction of the day, and the motor slip correction procedure worked well. Over these 35 days the mean temperature was -58.8°C . The mean temperature over the full 88 days was -59.0°C , so the 35 days do not show a bias terms of ambient temperature.

The data are returned from NISM as ADU (Analogue Digital Units). The next step is to use some calibration source to convert these into actual measurements of the sky brightness. There are two main possible calibration sources that NISM can see: the Milky Way, and the black body attached to NISM. We investigate using both of these as calibrators in the next chapter.

Chapter 6

Calibration

After processing the data using techniques from the last chapter, we have accounted for NISM's motor slipping, and have collected the measurements of the zenith sky brightness. NISM reports its measurements in Analogue Digital Units per millisecond, found by dividing the ADU readout on the detector by the time over which it has been integrating measurements.

To obtain a calibrated sky brightness measurement, we need to convert the ADU/ms measurements into units of spectral radiance. The SI units for spectral radiance are $\text{W sr}^{-1} \text{m}^{-2} \text{Hz}^{-1}$, i.e., watts per steradian of sky, per metre squared of collecting area, per unit frequency of the radiation. By convention, the units used when looking at the infrared sky background are $\mu\text{Jy arcsec}^{-2}$, i.e., microjanskys per square arcsecond of sky. Janskys are the commonly used units for radio astronomy. One jansky is defined as $10^{-26} \text{W m}^{-2} \text{Hz}^{-1}$.

Arcseconds are a unit of angle. 1 degree = 60 arcminutes = 3600 arcseconds, and 1 steradian = $4.25 \times 10^{10} \text{arcsec}^2$. We can easily convert from $\text{W sr}^{-1} \text{m}^{-2} \text{Hz}^{-1}$ to $\mu\text{Jy arcsec}^{-2}$ by dividing by $10^{-6} \times 10^{-26} \times 4.25 \times 10^{10}$ (i.e., by 4.25×10^{-22}).

To be able to convert our ADU measurement into physical $\mu\text{Jy arcsec}^{-2}$, we need a suitable absolute calibration source. There are two possible sources in the NISM data: the attached black body, and the Milky Way. The brightness of the Milky Way is known from previous measurements, and the brightness of the black body can be calculated from its temperature. Both of these features are clearly visible in the data.

There are two significant differences between using the Milky Way and using the black body as calibrators. The first is that they have different angular sizes. The Milky Way has a small angular size and does not fill the beam, hence we are very sensitive to the extended wings of NISM's beam. The black body fills the beam, so the exact beam shape is much less important. The second is that the Milky Way and black body have different spectra, which will affect how much light gets through NISM's filter at different wavelengths.

Additionally, we make the assumption that the coldest measured direction is at zero emission. This will be looking down at the snow, which is extremely cold. The air temperature during winter is $\sim -70^\circ\text{C}$, and the ice surface can be even colder. This is a consequence of the fact that the surface has a higher emissivity than the atmosphere (Hudson & Brandt, 2005).

Ice at $\sim -70^\circ\text{C}$ has a negligible thermal emission compared with the black body ($\sim -40^\circ\text{C}$) and the Milky Way. When viewed through NISM's filter, a black body at -40°C is 40 times brighter than a black body at -70°C . Therefore the emission from the ground will be very low, such that we can approximate it as zero.

If our assumption is incorrect, then the zero point would be even lower. This increases the measured brightness of both the zenith and the black body, giving a higher final zenith brightness. Therefore, our method calculates a lower bound for the zenith brightness.

6.1 Calibration using the Milky Way

The arc of the Milky Way is clearly visible in the data (See Figure 5.11). Of the 88 days of data we processed, 35 proved to be usable, and the Milky Way is clearly visible in 16 of them. The mean temperature over these 16 days was -55.8°C , so $\sim 3^\circ\text{C}$ higher than usual. The days where the Milky Way is visible may be biased towards having clear skies and higher air pressures, which would increase the path length through the atmosphere, which would make the observed milky way brightness dimmer than expected. The calibration factor would then be higher, which would give a higher zenith measurement. However, any variations in the air

pressure should have a very small effect on the attenuation. It could be worth monitoring the air pressure in any future missions to account for this.

Using the data from these 16 days, the Milky way should be an excellent calibration source. But to be able to use it as a calibrator, we need some absolute measurements of its spectral radiance.

For this, we use data from the DIRBE instrument, which was mounted on the COBE satellite that was launched in 1989 (Silverberg et al., 1993; Hauser et al., 1998). DIRBE (the Diffuse InfraRed Background Experiment) measured the spectral radiance of the entire sky at a range of wavelengths from 1.25 microns to 240 microns. The short end of this wavelength range nicely covers the K_{dark} band measured by NISM. The data are available on their website¹.

The DIRBE wavelengths that are closest to NISM's K_{dark} passband are at 1.25 μm , 2.2 μm , 3.5 μm , and 4.9 μm . DIRBE's measurements of the brightest part of the Galactic Centre at each wavelength are plotted in Figure 6.1. This figure includes a simple polynomial fit to the points. This polynomial shows that there is only a gradual change in the Galaxy's brightness with wavelength (in this region at least). The brightness at NISM's passband (2.4 microns = 1.24×10^{14} Hz) is close to that at the DIRBE 2.2 micron band. It is marked on Figure 6.1 as a cross. This simple polynomial interpolation suggests that the brightness of the Galaxy in the K_{dark} band is 96% of that at 2.2 microns. Therefore, we should be able to use the 2.2 micron measurement to compare with NISM's observations, after scaling it by 0.96.

Figures 6.2, 6.3, 6.4, and 6.5 show equirectangular projections of the full sky DIRBE data sets, at 1.25 μm , 2.2 μm , 3.5 μm , and 4.9 μm respectively. The Milky Way is clearly the dominant feature in the sky at these wavelengths. There are also numerous stars that are bright in the infrared. Of interest is the dark band cutting across the Galactic Centre that is present in the 1.25 μm plot, but absent in the others. This is due to absorption by interstellar dust grains in the plane of the Milky Way. These dust grains absorb shorter wavelengths of light (closer to visible), but are transparent to longer infrared wavelengths. This is a nice visual

¹https://lambda.gsfc.nasa.gov/product/cobe/dirbe_aasm_data_get.cfm




Figure 6.1: The brightness of the Galactic Centre as observed by DIRBE. Points are the observed wavelengths at 4.9, 3.5, 2.2, and 1.25 μm . The curve is a simple second order polynomial fit, for the purposes of interpolation. The cross is the position of the K_{dark} band at 2.4 μm .




Figure 6.2: DIRBE all sky infrared measurements at 1.25 μm wavelength.




Figure 6.3: DIRBE all sky infrared measurements at 2.2 μm wavelength.




Figure 6.4: DIRBE all sky infrared measurements at 3.5 μm wavelength.




Figure 6.5: DIRBE all sky infrared measurements at 4.9 μm wavelength.

example of one of the benefits of infrared astronomy—we can observe regions in the centre of the Galaxy that are obscured at visible wavelengths.

Figure 6.3 shows the DIRBE 2.2 micron data, which is closest to NISM’s observing wavelength. However, NISM has a very low resolution, so sees a much blurrier view of the Galaxy (See Figure 3.4 for NISM’s beam shape). To calculate the brightness of the Galaxy as would be observed by NISM, we convolve the DIRBE data with NISM’s beam shape.

The result of this convolution can be seen in Figure 6.6. NISM’s low resolution is very obvious. Notably, many of the bright stars are no longer visible, as they provide too little total brightness to be detectable in NISM’s large beam. They are effectively smeared out across a larger area of sky.

6.1.1 Atmospheric attenuation

From the convolution of the DIRBE data with NISM’s beam shape, we have the brightness of the Galaxy that NISM would see. However, this light from the Galaxy will be slightly attenuated by the atmosphere before it reaches the ground, so the brightness recorded by NISM will be slightly lower than the true brightness.

We can calculate the attenuation of the atmosphere from its optical depth,




Figure 6.6: DIRBE 2.2 μm observations convolved with NISM's beam shape.

which we know from the skydip shape. The greater the optical depth of the atmosphere, the greater the deviation from a pure $\sec(z)$ skydip. If the optical depth is zero, then emission is seen from all sections of the atmosphere along the line of sight, and the measured brightness is proportional to the airmass, giving a pure $\sec(z)$ skydip. If the optical depth is very high, then the atmosphere is effectively very foggy, and looks the same in all directions. Emission is only seen from the nearest section of the line of sight, and the skydip is instead flat, with no dependence on elevation. Phillips et al. (1999) cover this in Figure 6 of their paper.

The optical depth parameter of the atmosphere is found by the lmfit model fitting procedure, described in Section 5.2. We find the median value of the optical depth of our fitted skydips to be 0.2. This can be used to calculate the ground-level flux F_{ground} using the following equation.

$$F_{\text{ground}} = F_{\text{space}} e^{-\tau \sec z} \quad (6.1)$$

F_{space} is the flux seen above the atmosphere by DIRBE (Figure 6.6), τ is the optical depth parameter (0.2), and z is the angle from the zenith.

Applying this scaling factor to the DIRBE data converts them into the brightness that would be seen through the atmosphere at ground level.




Figure 6.7: NISM data from one day.

Ideally, we would like to have well-sampled measurements of the air pressure at Ridge A. The air pressure affects the total column of air along the beam, and thus the attenuation. However, no continuous air pressure measurements have been taken at Ridge A, so instead we perform all calculations relative to the mean airmass, with its optical depth of 0.2. Any variations in pressure should have a relatively small effect.

6.1.2 Subtracting the Galactic flux

We now have a model of the brightness of the Milky Way as seen by NISM. We can compare this model to the signal that NISM measures from the Milky Way, in order to find an ADU-to- $\mu\text{Jy arcsec}^{-2}$ conversion factor and calibrate the NISM data.

To do this, we subtract the DIRBE data from the NISM data using a range of different conversion factors. When we have found the correct conversion factor, the Milky Way in the DIRBE data will match the Milky Way in the NISM data, and subtracting them will leave solely the atmospheric signal behind.

The first step is to align the two data sets. The DIRBE data are already in




Figure 6.8: NISM data from one day after subtracting the Milky Way, using the correct conversion factor. Compare with Figure 6.7.

units of RA and DEC. The NISM data are recorded as a motor position (which corresponds to an elevation), and a timestamp.

NISM has a fixed azimuth. After some trial and error of subtracting the DIRBE data, we determined that the azimuth values where NISM is looking at the horizon are at 83 and 263 degrees. We also found that NISM is not quite vertical: it is tilted 5 degrees away from the zenith. This should have no effect on our measurements, but it is necessary to include in order to align the NISM and DIRBE data.

After aligning the two data sets, we subtract the DIRBE data using a range of possible conversion factors, and manually inspect the plots to find the point where the Milky Way disappears.

Manual inspection of the data is the best method to use here. The data are fairly noisy, and there may be slight misalignments, and it would be difficult to automate a system that can understand what is happening in each case. Manually looking at the plot allows us to intelligently determine when the signal matches, taking into account noise and alignment mismatches.




Figure 6.9: Histogram of Ridge A zenith spectral radiance measurements from the usable 35 days of data from the 2015 season. Calibrated using the observed brightness of the Milky Way.

We ran this process on the 16 days where the Milky Way was clearly visible. There is some variation in the results we find, as they can be affected by clouds, the motor slipping, alignment with DIRBE, and other noise. But over the 16 days, we find an average conversion factor of $7 \pm 0.5 \times 10^{-6}$. This can be seen in Figures 6.7 and 6.8: the Galaxy disappears once the correct conversion factor is found.

We can then use this calibration factor to determine the zenith spectral radiance. Dividing the zenith ADU values from the usable 35 days of data by the conversion factor of 7×10^{-6} (and scaling by 96% for the wavelength difference) gives the desired zenith spectral radiance values, shown in the histogram in Figure 6.9 and in Table 6.1.

We find a 50th percentile value of $294 \mu\text{Jy arcsec}^{-2}$. This is significantly higher than expected. Measurements from the South Pole find a zenith spectral radiance

Table 6.1: Ridge A zenith spectral radiance measurements from the usable 35 days of data from the 2015 season (from Figure 6.9). Calibrated using the observed brightness of the Milky Way.

	$(\mu\text{Jy arcsec}^{-2})$
Mean	306
Median	294
25th percentile	241
10th percentile	207

of $\sim 120 \mu\text{Jy arcsec}^{-2}$ Lawrence et al. (2002b). Ridge A is higher, drier, and colder than the South Pole, and so should have a lower sky brightness. The modelling with VSTAR also agrees that the brightness should be well below $100 \mu\text{Jy arcsec}^{-2}$. Therefore, we suspect that this result is incorrect, and that there is a problem somewhere with the analysis or the instrument.

6.2 Calibration using the black body

There is another option for calibration: NISM's inbuilt black body. This cone of copper is heated to ~ 30 K above ambient temperature. Since we know the temperature of the black body, we can calculate its radiance from Planck's Law,

$$B_\nu(\nu, T) = \frac{2h\nu^3}{c^2} \frac{1}{e^{\frac{h\nu}{k_B T}} - 1} \quad (6.2)$$

NISM was designed so that the black body fills the beam. Measurements of the beam shape show that it is larger than expected, but it is still filled by the black body. Since the black body has a very uniform temperature across its face (see Figure 3.16), we do not need to worry about the exact beam shape. We can simply use the black body temperature and Planck's Law to calculate the appropriate black body spectrum. We then multiply this spectrum by NISM's passband response (Figure 3.10) and integrate to find the spectral radiance that should be measured by NISM.

By comparing this calculated brightness to the observed brightness we can

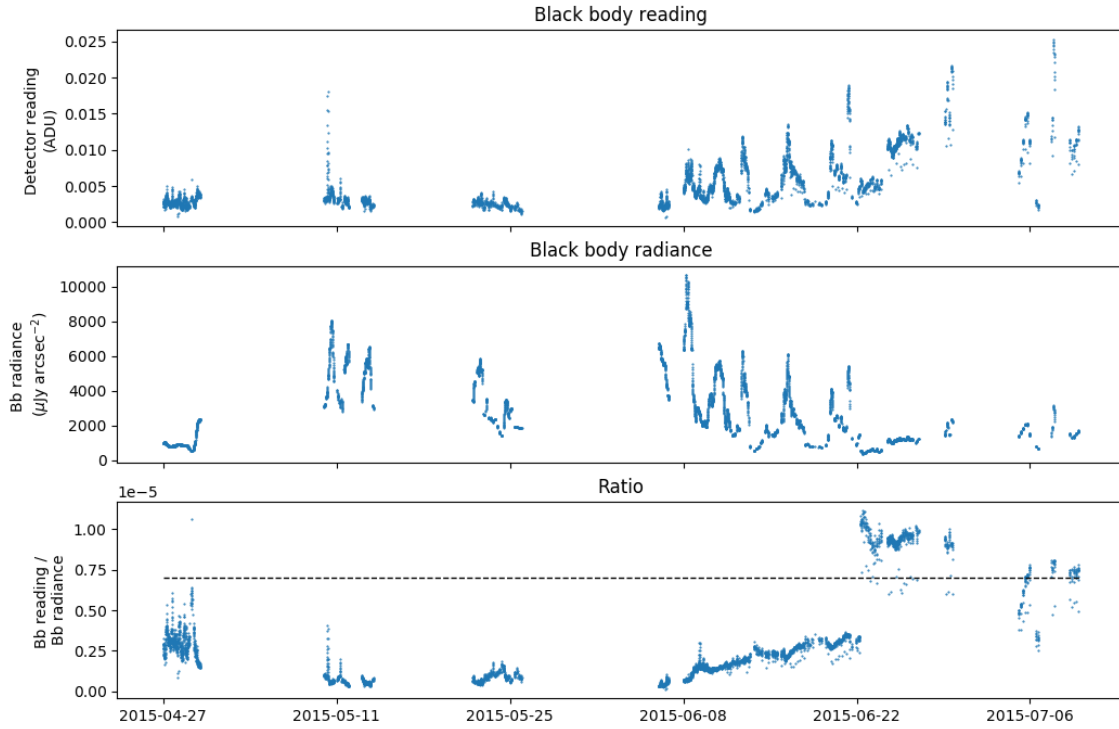


Figure 6.10: Measurements of the black body brightness over the year. The first subplot shows the apparent brightness of the black body as measured by NISM’s detector, the second subplot shows the calculated radiance based on the black body’s temperature, and the third subplot shows the ratio between the first two. The dashed line is the ratio calculated using the Milky Way in Section 6.1.

determine the conversion factor. The observed and calculated brightnesses of the black body are shown in Figure 6.10. This plot shows data from the black body over the 2015 observing season.

The first subplot of Figure 6.10 shows the brightness of the black body in ADU, as measured by NISM. To obtain this measurement we select the brightest part of the signal when NISM is looking in the direction of the black body, to account for any misalignment from the motor slip fitting process.

The second subplot shows the calculated brightness of the black body in $\mu\text{Jy arcsec}^{-2}$, based on its temperature. Heaters keep this temperature ~ 30 K above ambient, so changes in the ambient temperature cause the changes in the black body temperature. The third subplot shows the ratio between the first two.

We would expect this third subplot to show a constant value, equal to the calibration factor. However, it clearly does not. The ratio between the observed

and calculated brightnesses varies throughout the year, which indicates some problem.

The good news is that we are definitely recording some signal from the black body. In the time between 2015-06-08 and 2015-06-22 there are significant variations in the temperature of the black body, which produce the noticeable spikes in the calculated signal in subplot 2. These spikes match up well with spikes in the observed brightness in the first subplot. The third subplot shows how closely these two measurements follow each other. During this time the brightness ratio is increasing fairly smoothly, despite the temperature variations. This strong correlation indicates that we can see the black body, at least partially.

The dashed line on the third subplot in Figure 6.10 is the conversion factor found using the Milky Way as a calibration source. We can see that the black body calibration factor varies above and below it at different times. It is not immediately clear if either of these calibration sources is correct.

The black body should in theory be the most reliable calibration source, yet measurements show its apparent brightness changing unpredictably. The next chapter investigates this phenomenon in detail, with the aim of correcting for this error and using the black body as an accurate calibration source.

Chapter 7

Black body investigation

In the last chapter, we discovered that there is some problem with the measurements of the black body. NISM measures the brightness of the black body, and temperature sensors on the black body allow us to theoretically calculate its brightness. We would expect there to be some constant ratio between these two numbers, which is the conversion factor from NISM's ADU measurements into $\mu\text{Jy arcsec}^{-2}$. However, this ratio is definitely not constant. It varies throughout the year, although the variations seem to be relatively smooth.

What could cause of this discrepancy? At the fundamental level, we are comparing two measurements: the temperature of the black body reported by sensors attached to the copper cone, and the signal reported by the InSb diode when NISM is looking in the direction of the black body. When things are working correctly, these measurements should be directly linked: the signal recorded by NISM should be solely radiation emitted from the black body.

We can see from the third subplot in Figure 6.10 that the system is working at least somewhat correctly. The ratio between the calculated and measured signals is, while not constant, at least strongly correlated. The spikes in measured brightness between 2015-06-08 and 2015-06-22 closely match the spikes in calculated brightness. This tells us that we are observing the black body, and that the signal NISM records is at least partially due to thermal emission from the black body.

The chain of events for a measurement is as follows:

-
1. The black body emits radiation according to its temperature and emissivity.
 2. This travels to NISM's mirror, where it is reflected towards the detector housing.
 3. It passes through a lens and filter, before entering the detector and being recorded by the InSb diode

The close correlation between the measured and calculated signals indicates that the temperature sensors and detector are both working. The emissivity of the black body should be constant, and the relatively high temperature of the surface should prevent any ice from growing. In the dry Antarctic atmosphere, it is sufficient to heat surfaces just a few K above the ambient temperature to sublime away any ice. Additionally, we can discount any thermal emission from ice on the mirror or lens. This would have a constant effect on the measured signal which would not vary with elevation angle, and so would be subtracted out as part of the baseline emission.

We think that the most likely cause of this variation is that the apparent brightness of the black body is affected by ice partially covering the front of the black body cavity. While the temperature of the black body should prevent any ice growing on its surface, it is possible that ice or snow could collect on the outside of the housing, or on the frame of NISM. If enough ice is present, it could block the line of sight towards the black body, and thus block some of the radiation. This would mean that the black body fills less of the beam of NISM. As a result, the signal recorded by NISM will be lower, and thus the ratio of measured/calculated signal will be lower too. It is also possible that, with the larger than expected beam, we are seeing changing ice cover on reflective surfaces near the black body, which would also alter the observed signal.

Changing ice coverage does fit with the ratio seen in the third subplot of Figure 6.10. The initially low ratio could be due to ice blocking the view of the black body. Later in the season the ice may be sublimating or blowing away, revealing more of the black body and increasing the measured brightness. The sudden jump at 2015-06-22 could be a large chunk of ice falling off. This is supported by looking




Figure 7.1: A comparison of 5-minute blocks from before and after the jump. The black body peak is wider after the jump, suggesting that ice fell away exposing a greater surface area.

at the width of the black body peak before and after the jump. Figure 7.1 shows two 5-minute blocks of data taken before and after the jump. The black body peak is noticeably wider after the jump, which is consistent with ice falling off and exposing a greater surface area.

NISM definitely does accumulate ice, as can be seen in Figure 7.2.

Unfortunately, we do not have a view of the front of the black body, so cannot confirm that there is ice blocking the aperture.

However, it may be possible to infer the presence of ice on the black body. The black body has a heater attached which is set to a constant level. This keeps the black body roughly 30 degrees above the ambient temperature. As the ambient temperature goes up and down, the black body temperature goes up and down too, but, importantly, with some time lag. The black body has some thermal mass, and is insulated from the environment, so it takes some time to respond to variations in the air temperature.

If the black body is covered in ice, then we would expect this to further insulate the black body from the air, and thus we would expect there to be a longer delay between the air temperature changing and the black body temperature changing.

NISM has a sensor that records the external air temperature. What we want to do is compare the external and black body temperatures, and look for a time delay in the two signals. This delay should hopefully correlate with the amount of ice on the black body. We expect this delay to vary over time with the amount of ice on



Figure 7.2: Ice on the back of NISM early in the season, on 2015-03-30 (From the Ridge A webcam).

NISM.

A function like cross-correlation may not be ideal here. Since we expect the offset to vary over time, we would not find a single maximum at some global shift value. Instead we need a function that looks for a time-varying offset between two signals. For this purpose, we investigated using dynamic time warping algorithms.

7.1 Dynamic time warping

Dynamic time warping is an algorithm that ‘warps’ one signal in time in order to match it to another signal. The amount of warping can be different at different times, i.e., it can vary dynamically. This seems to be an appropriate algorithm for our purpose. We have two sequences that we expect to show a similar shape, but with some dynamically varying offset between them.

Dynamic time warping was originally developed for speech recognition. The speeds at which words are spoken can vary from speaker to speaker, so it is necessary to dynamically time warp a recorded phrase in order to match it to a library of known reference words.

The full dynamic time warping algorithm involves two main steps. The first is to create an accumulated cost matrix from the two sequences, and the second is to find an optimal warping path through this accumulated cost matrix.

The dynamic time warping method described in this section is based on the method in Rabiner & Juang (1993) and Müller (2007).

7.1.1 Accumulated cost matrix

The first step in dynamic time warping is creating a cost matrix from the two input sequences. For this, the sequences must be evenly sampled in time. Our measurements of the external and black body temperature were not evenly sampled, so we linearly interpolated between them at uniform time steps.

The two sequences are placed along the x and y axes of a matrix. At each cell of the matrix (n,m) , the n th value of one sequence is compared with the m th value of the other sequence. The difference between these two values becomes the entry

in that cell of the matrix. An example cost matrix is shown in Figure 7.3. We can see the two sequences placed along the x and y axes of the matrix. The values in the cells of the matrix show the difference between the two functions at each location: dark areas are where the sequences have similar values, and light areas are where the sequences have differing values.




Figure 7.3: An example of a cost matrix used in dynamic time warping. Taken from page 70 of Müller (2007)

The idea of dynamic time warping is fairly simple. We would like to shift the time values of some elements in a sequence, so that they match up with similar elements in a second sequence. This can be shown intuitively on the cost matrix in Figure 7.3. We would like to find some path through this matrix (from the lower left corner to upper right corner) that follows the dark regions, thus minimising the differences between the two signals. This path will match up elements from the two sequences such that the difference between them is minimised. An example of one such warping path is shown on the left in Figure 7.4.

There are several constraints applied to the path. Firstly, it must start at the bottom left corner of the cost matrix, and finish at the upper right corner. This ensures that all points in both input sequences are matched up with corresponding

points in the other sequence. Secondly, we cannot step backwards in time. Each point the path steps to must either have the same time value, or a larger time value. This makes sure that the points in the input sequences remain in the right order. Thirdly, specific step conditions are specified for moving from one cell to the next. In the standard DTW algorithm, 3 possible steps are allowed: up, right, and diagonally up and right (i.e., step sizes $(1,0)$, $(0,1)$ and $(1,1)$). This step condition makes sure that no values in either input sequence are missed out. It also ensures that the second condition is met: if only steps up and right are allowed, then no backwards steps can occur.

To find this optimal warping path, we use the cost matrix to create an accumulated cost matrix. Each cell in the cost matrix shows the difference between the two sequences at that point. Each cell in the accumulated cost matrix shows the total value of an optimal warping path that has reached that particular cell. The total value is the sum of all the cost matrix cells that that optimal warping path has gone through. An example of an accumulated cost matrix is shown on the right in Figure 7.3.




Figure 7.4: An example of a cost matrix and an accumulated cost matrix used in dynamic time warping. The white line is the optimal warping path. Taken from page 74 of Müller (2007)

For DTW we want to find a path through the cost matrix that follows the dark regions, and ends up having the lowest total difference between the two signals. This is what the accumulated cost matrix shows: the total cost we would have accumulated when getting to a cell. Specifically, it shows the lowest possible value, which is generated by taking the best path.

The accumulated cost matrix is calculated in a recursive fashion. To find the value in a particular cell of the accumulated cost matrix, we look at the values in the cells that we could have previously stepped from. In standard DTW, these will be at steps $(-1,0)$, $(0,-1)$, and $(-1,-1)$. The length of the shortest possible path to the current cell must be the length of the shortest possible path to the previous cell, plus the cost value of the current cell. So to calculate the accumulated cost, all we need to do is look at the accumulated cost values in the three possible previous cells, pick the lowest, and add the cost of the current cell.

When calculating values in row and column 0, this algorithm will look for values in cells outside the matrix and crash (here we are using the 0 indexing convention where 0 is the first row). To avoid this, an extra row and column of cells are added to the left and bottom of the accumulated cost matrix. The values in these cells are then set to infinity, so that they are never selected as the shortest path length. Additionally, we need a starting value, so the corner of these two new rows (position $(-1,-1)$) is set to zero. This ensures that the accumulated path length of the first cell $(0,0)$ is just the cost value of that cell.

With these adjustments, the algorithm can be run for every cell in the matrix, producing an accumulated cost matrix like the one in Figure 7.4. It immediately becomes clear where the optimal warping path lies: along the dark region in the accumulated cost matrix.

7.1.2 Optimal warping path

Once the accumulated cost matrix has been calculated, finding the optimal warping path is simple. Again, a recursive algorithm is used, but now we look at the path in reverse, starting at the top right corner, and working back to the origin. We simply look at the three possible cells we could have stepped from, select the one with the lowest accumulated cost, and step there. We only have three possible options for where we could have stepped from, and since the one we select has the lowest accumulated cost, it must be on the path that has the lowest possible accumulated cost.

We then repeat this process, stepping back one cell each time, until we reach

the origin. The path we have followed will be the one with the minimum possible accumulated cost, and is therefore the optimal warping path.

This path allows us to align elements of the first signal so that they best match the elements in the second signal. Or, equivalently, to align the second signal to the first. The amount of time warping can be found by calculating the distance of the optimal warping path from the diagonal (as the diagonal represents 0 time warping between the two signals).

7.2 Shape dynamic time warping

The previous sections describe the simplest DTW algorithm, but there are more complex ones that can be used for different data sets. For our data, we used a method called shape dynamic time warping.

In standard dynamic time warping, the two sequences are compared directly. The value put in a cell of the cost matrix is simply the value of one sequence minus the value of the other sequence. This is useful for cases where the two signals should have the same value. However, in our case we expect the two sequences to have different values; the temperature of the black body should be ~ 30 degrees higher than the ambient temperature.

Instead of comparing the temperature values, we would like to compare the trends: both measurements should increase and decrease in sync. Shape dynamic time warping does this by selecting a window around each point, calculating the derivatives of the sequences in this window, and comparing those. This lets us compare the time derivative of the temperature, rather than the temperature itself.

In our case, we fit a second-order polynomial to a window of points around each value being compared. We then take the derivative of this polynomial at these points. When generating the cost matrix, instead of subtracting the two values being compared, we find the Euclidean distance between the two derivatives. That is, we subtract the derivatives at each point in the window, square each result, add them all together, and take the square root of the sum. This gives a single value which is entered into the cost matrix.

7.2.1 Custom step sizes

Another change that can be made to the DTW algorithm is to customise the way that steps are taken through the cost matrix to generate the warping path. In the default algorithm, three possible steps are allowed: vertically up, horizontally to the right, and diagonally up and right. That is, to reach cell (i,j) , we can step from one of the cells $(i,j-1)$, $(i-1,j)$, or $(i-1,j-1)$. These three possible steps are shown in the left side of Figure 7.5.

This is the simplest step constraint, but it does have a property that can be undesirable. At each point in the path, it is allowed to move horizontally, vertically, or diagonally. This means that it can take an unlimited number of horizontal or vertical steps in a row. If this happens, the algorithm is matching a large number of values from one sequence with a single value in the other sequence. This is unlikely to be physical. The values in the sequences may match, but we would not expect such a rapid time shift between the sequences.

So, we would like to use a different step constraint that prevents too many horizontal or vertical steps being taken in a row. One possible constraint is shown on the right of Figure 7.5. This constraint only allows two vertical or horizontal steps to be taken in a row, before a diagonal step must be taken. This prevents the warping path from wandering too far away from the diagonal (which would correspond to a very large warp).

This different step constraint is implemented when the accumulated cost matrix is calculated. In the standard condition, we have three possible starting locations. For each one, we take that value of the accumulated cost matrix, and add the value of the cost matrix in the cell we are stepping to.

When using the modified constraint, we have five possible starting locations (Figure 7.5). For each one, we take the value of the accumulated cost matrix at that starting location, and add in the values in the cost matrix at each point along the path. This gives a total path length for each of the five possible paths. We then select the smallest of these, and put that value in the accumulated cost matrix.

The second step of the algorithm can be carried out as normal. Start at the top right corner, look at the three possible steps, and select the one with the lowest




Figure 7.5: Possible step options for DTW. The dots are cells in the cost matrix. On the left is the standard condition. The condition on the right restricts repeated horizontal and vertical steps.

accumulated cost. Because of the change to the way the accumulated cost matrix is created, we will never end up selecting more than two horizontal or vertical steps in a row.

7.2.2 Custom weights

Another possible change is to give the steps different weights. In the default algorithm, the three step possibilities all have equal weighting. This means that there is no preference for a vertical, horizontal, or diagonal step. However, this ‘even’ weighting ends up preferring diagonal steps. One diagonal step is equivalent to taking one horizontal plus one vertical step. This would involve adding the cost of two squares rather than one, and so would add a larger cost to the accumulated cost matrix. As such, diagonal steps are cheaper, and thus preferred by the optimal path.

A solution to this is to weight the different steps. By multiplying the cost of a diagonal step by 2, its cost is brought into line with the horizontal and vertical steps. This makes the path equally likely to choose a vertical, horizontal or

diagonal step each time.

However, care should be taken to make sure that this equally weighted result is the desired one. Diagonal steps indicate no warping between the two sequences. If the two sequences are expected to be similar to one another we would expect little warping, and therefore expect diagonal steps to be the most common. The default weighting gives more prominence to diagonal steps, exactly as would be expected for similar sequences. The weighting used should reflect the expected correspondence between the sequences, and the preferred shape of the warping path.

For our case, we expect the two sequences to be similar and so used the default weighting, where all five step possibilities have equal weight. This slightly favours diagonal steps, as we would expect to see between similar sequences.

7.3 Results of DTW

We implemented the shape DTW algorithm on the temperature by modifying the python package pydtw. This package can implement the standard DTW algorithm in 1D or 2D. We modified the 2D section of this program to use the shape DTW algorithm and the modified step condition.

Figure 7.6 shows the result of running the DTW algorithm on the external and black body temperatures. The first subplot shows the two temperature sequences (with the black body temperature shifted down by 32.3 degrees to visually match). The second subplot shows the result of aligning the external temperature to the black body temperate (although little change is visible at this scale). The third subplot shows the time shift: the amount of time that the external temperature was shifted by to match the black body temperature, according to the DTW algorithm.




Figure 7.6: Dynamic time warping of the external temperature to the black body temperature.

In general, we can say that there is not an obvious pattern in the time shift. It seems to vary randomly throughout the observation period. Thus, we cannot use it to model the amount of ice on NISM. There are likely many other factors that also affect the temperature of the black body, such as the speed and direction of the wind.

However, an interesting feature was found when exploring the dynamic time warping algorithm.

Figure 7.7 shows the result of the dynamic time warping algorithm again, but this time with the simple step condition, rather than the modified one (the left side of Figure 7.5, rather than the right). With this simple step condition, the

7.3. RESULTS OF DTW

DTW algorithm is more free to explore the warping possibilities.




Figure 7.7: Dynamic time warping of the external temperature to the black body temperature, using the simple step condition.

The third subplot now shows significant warping at several points. The most extreme one is the large spike between 2015-04-24 and 2015-05-08 (starting on 2015-05-01). This corresponds with the time when the black body heater setting was increased 50% (from 26200 to 39300 arbitrary units). This caused the black body to heat up to reach a new equilibrium temperature by the next day. This jump in black body temperature is visible in subplots 1 and 2 of Figure 7.7. This sudden rise in the temperature of the black body is very different to the temperature trend of the external sensor, so the DTW algorithm tries to shift the

time by a large amount to match it, resulting in the distinctive spike. The simple step condition allows this sort of major shift.

The other major jump is between 2015-06-19 and 2015-07-03 (on 2015-19-22). This coincides with the jump in the black body reading which was noticed back in Figure 6.10.

So while the DTW algorithm may not give us a nice model of the insulation of the black body, it can highlight areas where the two temperature plots are behaving in noticeably different manners.

7.3.1 Investigating the jump

Figure 7.8 shows the observed brightness of the black body, the radiance of the black body calculated from its temperature, and the ratio of these two values. These are the same measurements as in Figure 6.10, but Figure 7.8 is limited to the times around 5 am GST on 2015-06-22, when the ratio of brightness to radiance suddenly jumps up. It is clear that at this time, the observed brightness of the black body suddenly increases, while the radiance/temperature does not.

Figure 7.9 shows the result of the DTW algorithm over the same period of time. The subplots in Figure 7.9 show the same measurements as in figure 7.7, just restricted to a time window around the jump at 2015-06-22. Just after this jump, the third subplot shows a significant amount of warping produced by the DTW algorithm. This is also clearly noticeable in the full plots in Figure 7.7

The large amount of warping indicates that the black body temperature and the external temperature are behaving very differently. This would be consistent with large pieces of ice falling off the black body, exposing it to the cold air and causing a sudden drop in its temperature.

From this result, we believe that there was ice covering the aperture of the black body at various times throughout the season. This ice would have partially blocked NISM's view of the black body, so it would not fill the beam, and hence the radiation detected by NISM would have been lower than would be calculated based on the black body's temperature. At the time of the jump on 2015-06-22, a large section of ice fell or was blown off. This exposed more of the face of the black




Figure 7.8: Black body reading, radiance and ratio at the time of the jump (magnified from figure 6.10).

body to NISM’s detector, increasing the measured brightness (Figure7.8). It also exposed the black body to more of the air, causing a rapid change in the temperature, which was picked up by the DTW algorithm.

The design of NISM was intended to prevent ice forming on the black body. Its temperature of 30 degrees above ambient should be more than enough to stop any ice forming on the surface. However, it is possible that ice built up on part of the frame of NISM. The ice present in Figure 7.2 clearly indicated that ice builds up on some surfaces. Perhaps enough ice piled up on the frame in front of the black body to obstruct the view. It is also possible that ice crystals formed on the outside of the casing, and grew to extend over the aperture.

In future, it may be useful to set up a webcam and LED lamp facing NISM. If ice is suspected during the winter, these could be used to determine how much is there. Also, changes to the design could be made, perhaps avoiding any shelf in front of the black body where snow could accumulate.

In any case, the presence of ice obscuring the black body means that it is not




Figure 7.9: Dynamic time warping from Figure 7.7 at the time of the jump.

suitable as a calibration source. Since we do not know the percentage of the black body covered at any time, we cannot calculate the radiance that would be observed by NISM at any time, and so cannot get a ADU/radiance conversion factor. It seems that the Galaxy may be the better calibration source.

However, this measurement is still suspect. The measured zenith signal when using the Milky Way as a calibration source is much higher than expected. There is also the problem that the black body calibration factor varied above and below the Milky Way calibration factor (See subplot 3 in Figure 6.10). If ice blocking the black body was the only problem, then the black body calibration factor should only ever be lower than the Galaxy calibration factor.

Clearly more investigation is required. The next chapter investigates the Aurora Australis as a possible source of infrared emission causing the high zenith brightness.

Chapter 8

Aurora

The measurements from NISM indicate a much higher zenith brightness than was expected. Predictions indicated a value below $100 \mu\text{Jy arcsec}^{-2}$, but we found a median value of $294 \mu\text{Jy arcsec}^{-2}$ (Figure 6.9). While looking for other possible sources of emission, we investigated whether the Aurora Australis could be a cause. It is definitely a major source of emission at visible wavelengths. It is not expected to be a major source at infrared wavelengths, but it could still have some emission that contributes to the signal seen by NISM.

The aurora is caused by charged particles from the Sun interacting with the Earth's atmosphere. These particles are directed along the Earth's magnetic field lines towards the poles, where they excite various molecular species in the atmosphere. The brightest emission is from the [OI] line at 556.6 nm, which gives the aurora its characteristic green glow. There are also a host of other lines from N_2 and O_2 (Gattinger & Jones, 1974; Sims et al., 2012).

If NISM is detecting light from the aurora, then it should be fairly simple to find a correlation between the zenith brightness measured by NISM, and the auroral activity at that time.

8.1 NOAA satellites

To find information on the intensity of the Aurora Australis, we used satellite data from the National Oceanographic and Atmospheric Administration (NOAA). They

maintain a number of satellites that measure space weather and the conditions in the upper atmosphere. For our test, we used data from the NOAA19 satellite.

This satellite was launched in 2009, and is in a polar orbit at ~ 850 km altitude¹.

NOAA19 is the latest in NOAA's series of Polar Orbiting Environmental Satellites (POES). These satellites measure the flux of charged particles in near Earth orbit. They have two instruments for doing so: TED and MEPED. TED measures the lower energy particles, which are responsible for causing the aurora. MEPED detects the higher energy particles, which are more of a hazard to satellites.

TED has detectors that separately measure protons and electrons. It has two telescopes, one aimed at the zenith, and one 30° away. It also has two different energy bands: 50 eV to 1 keV, and 1 keV to 20 keV. Therefore, TED contains a total of 8 electrostatic analysers to report all combinations of these measurements ($2 \times 2 \times 2$).

Data are available on the NOAA website² in both the raw form (detector counts) and the processed form (physical units). For our overall measurement of the aurora, we used the parameter `ted_total_eflux_atmo`. This is the processed total energy flux in the atmosphere at 120 km elevation, measured by TED in mW m^{-2} . It includes both protons and electrons, and covers the full energy range of 50 eV to 20 keV.

These charged particles streaming past the satellite will go on to interact with the Earth's atmosphere, where they will produce an aurora. The greater the flux of charged particles, the greater the brightness of the aurora, so we can use this measurement as a proxy for aurora brightness.

8.2 Aurora measurements

The NOAA satellites are in polar orbits, so they pass over Antarctica. However, they do not regularly pass directly over Ridge A. Thus, we have to choose some

¹<https://www.wmo-sat.info/oscar/satellites/view/341>

²<https://satdat.ngdc.noaa.gov/sem/poes/data/processed/ngdc/uncorrected/full/2015/noaa19/>

spatial range around Ridge A to take auroral measurements from. For this experiment we made the assumption that auroral activity over Antarctica is spatially correlated. That is, when the auroral emission is strong over one part of Antarctica, the emission should be strong in other parts of Antarctica too. Thus, we can take auroral measurements from anywhere in Antarctica, and expect them to be correlated with the auroral emission at Ridge A.

Limiting the selected auroral measurements to a closer radius around Ridge A would give a stronger correlation to the Auroral activity happening at Ridge A, but would significantly reduce the data available. The satellites make 14 orbits per day, and their closest approach to Ridge A changes as the Earth rotates. Only allowing measurements very close to Ridge A could leave too few data points to investigate any correlation.

Auroral activity is confined to an annulus around the south magnetic pole. To determine the strength of the aurora, it is important to select measurements in this region. Selecting a small sampling area could give worse results as it may miss the annulus entirely.

We selected a radius of 3000 km around Ridge A. This includes most of the continent of Antarctica. A plot of the orbital coverage can be seen in Figure 8.1. This shows the foot-of-field of NOAA19 over one day³. Ridge A is marked with a green dot, and all the orbital locations within 3000 km of Ridge A are highlighted in red.

³The foot-of-field is the location on Earth where a field line passing through the satellite intersects with the ground.




Figure 8.1: The orbit of NOAA19 over one day (2015-05-13). Sections within 3000 km of Ridge A are highlighted in red.

Figure 8.2 shows the particle flux detected by NOAA19 as it orbits during one day. For this plot, a cut-off was set to $> 0.5 \text{ mW m}^{-2}$. We can see the auroral activity occurring in a roughly annular shape over Antarctica. On this particular day, little activity was detected over East Antarctica.




Figure 8.2: Flux detected by NOAA19 during 2015/05/13.

8.3 Comparison

Figure 8.3 shows the NOAA19 measurements of the aurora during 2015 (red), plotted on the same axes as the NISM zenith brightness measurements (blue). This includes all measurement within the radius marked in Figure 8.1. At first glance there does not appear to be much correlation between them. This is confirmed by plotting one against the other, as shown in Figure 8.4. This figure shows no correlation. There are a large number of times when there is zero detected aurora, but even when there is detected particle flux, it shows no relationship with the infrared zenith brightness.

8.3. COMPARISON




Figure 8.3: NOAA aurora measurements (red lines) and NISM zenith brightness measurements (blue dots) plotted together.




Figure 8.4: Correlation of the aurora and zenith brightness measurements in Figure 8.3.

Therefore we can conclude, as was expected, that auroral activity has no detectable effect on the zenith brightness as measured by NISM.

Chapter 9

Red leak

The last chapter indicated that the Aurora is not causing a high signal to be measured by NISM. This then leaves us with the problem that the zenith brightness according to NISM is $\sim 294 \mu\text{Jy arcsec}^{-2}$, far higher than expected.

There is also the problem that the calibration factors found by using the black body and the Milky Way are inconsistent. The third subplot in Figure 6.10 shows how the calibration factor calculated using the black body varies above and below the calibration factor calculated using the Milky Way. If the sole problem with the measurements was ice blocking the black body, then it would only be dimmer than expected and the calibration factor would only be lower than the one found using the Galaxy.

The last source of error we can suppose is that there is a problem with the filter on NISM, and that it is letting in too much light at longer wavelengths. This could definitely be a source of extra zenith emission. It would also affect the calibration factors found using the black body and Milky Way, as they are at different temperatures.

9.1 Filter

InSb diodes are sensitive out to $\sim 5.5 \mu\text{m}$. This is visible in Figure 3.10, which shows initial testing of NISM's passband. There is the desired window around $2.4 \mu\text{m}$, but there is also an unwanted peak around $4.5 \mu\text{m}$. An additional blocking




Figure 9.1: Ridge A atmospheric spectral radiance as modelled by VSTAR, extending out to 5.5 microns

filter was added to NISM to block this extra long wavelength emission. However, if the filter is not perfectly blocking 100% of the emission, then extra light at longer wavelengths (‘redder’) could be leaking through.

This extra light could have a major effect on the total detected signal, since the atmosphere is much brighter at longer wavelengths. This was discussed back in Section 4.3 and shown in Figure 4.3.

We can see the effect at wavelengths out to 5.5 μm by re-running VSTAR with a larger wavelength range. This is shown in Figure 9.1. We can see that the emission continues to rise sharply as the wavelength increases. Comparing the vertical scale of Figure 4.3 with Figure 9.1 shows an increase of many orders of magnitude.

This effect is made even more dramatic when we include a slight red leak in NISM’s filter out to 5.5 μm. Figure 9.2 shows the emission from the atmosphere




Figure 9.2: Ridge A atmospheric spectral radiance as modelled by VSTAR. Weighed by NISM's bandpass, including a red leak out to 5.5 microns.

multiplied by NISM's passband when a small red leak is included (in the same manner as Figure 4.5). For this test, we have extrapolated from the bandpass measurement in Figure 3.10. The signal observed at $4.25\ \mu\text{m}$ was $0.221\ \text{V}$, so we have assumed that the bandpass leak has a value of $0.221\ \text{V}$ out to $5.5\ \mu\text{m}$. A leak of this magnitude is only 0.02% of the maximum transmission, and would be lost in the noise in the bandpass characterisation in Figure 3.10.

But here we can see that this tiny leak has a huge effect. The light coming in through the intended window at $2.4\ \mu\text{m}$ is the small bump on the left of Figure 9.2. The light coming in through the leak has a far greater effect on the total signal measured. The brightness at the different zenith angles is listed in Table 9.1. Including the leak changes the predicted zenith emission from 3.8 to $108\ \mu\text{Jy arcsec}^{-2}$, an increase of 25 times.

So a leak of this magnitude could be a very significant problem. A small leak

Table 9.1: Atmospheric spectral radiance through NISM’s passband, including a red leak, as calculated by VSTAR.

Angle (Degrees)	Radiance ($\mu\text{Jy arcsec}^{-2}$)
22	108.7
49	126.5
71	175.1
86	343.6

that would have been barely detectable in the original filter characterisation could account for $\sim 96\%$ of the observed zenith emission.

This leak would also affect the calibration factor. The brightness of the Milky Way (as shown in Figure 6.1) is not increasing towards longer wavelengths, so its apparent brightness would not be increased in the same manner as the sky and the black body. This could account for the differing calibration factors between the black body and the Milky Way.

9.2 Recovery

Unfortunately, documentation on the particular filter used is scarce. It is also not possible to characterise the filter based on the data we have collected; there are no carefully calibrated spectral sources visible at Ridge A. So it was decided that the best option was to recover NISM from Ridge A at the end of the project, and bring it back to a lab in Sydney for more detailed analysis of the filter. This was scheduled to take place in the 2017/2018 mission to Ridge A.

This was the last scheduled mission to the site for this project. As part of the Antarctic Treaty, equipment is not allowed to be left in Antarctica to become rubbish. Instead, everything must be brought back to a station, or out of Antarctica. So bringing NISM back was a natural part of this mission.

Unfortunately, the team was not able to access Ridge A. In the two years since the last mission, snow had blow over the skiway, forming sastrugi. The Twin Otter

pilots attempted to touch down, but they determined that the skiway was too rough, and that there was a risk of breaking the undercarriage. So they made the decision to return to the South Pole. It was not possible to organise another mission for that season, so the only option is to wait till next season and try again. By then the skiway may have smoothed out through natural process. Or the mission could include a groom team to travel to the site beforehand and set up a fresh skiway.

With the failure to reach Ridge A, it was not possible to recover NISM, and it is therefore not possible to carefully characterise the filter. Reconstructing the optical system and testing that as a proxy would not be sufficient, as there are a number of uncertainties as to the exact design of the instrument which would leave any reconstruction inaccurate. Hopefully the mission next summer will be successful, and the instrument can be safely returned.

With the return of the instrument, the filter could be carefully characterised. In a lab setting, an infrared spectrophotometer could be used to determine NISM's transmission as a function of wavelength. The original bandpass measurement was done using an infrared grating monochromator with a temperature controlled black body, so this is another option for testing, although a spectrophotometer would be preferable.

An accurate measurement of NISM's bandpass would allow us to calculate the fraction of light coming through at longer wavelengths, like the 96% figure calculated with VSTAR. This factor could then be subtracted from the Ridge A measurements, allowing a true measurement of the near infrared sky background to be made.

Chapter 10

Lessons

Throughout this thesis, a number of problems have been found with the NISM instrument. In light of these, there are a number of changes that could have been made to NISM. These problems could also affect future missions, so should be carefully examined for future reference.

10.1 Instrument housing

The design of NISM has a weakness that was uncovered in Chapter 7: there is a space for snow to settle in front of the black body. This snow accumulation is presumed to have led to the changing brightness of the black body. The two most likely spots are the inner rim of the black body housing (visible as the light grey strip inside the black body in Figure 3.1), and the section of baseplate in front of the black body (the anodised red piece visible in Figure 3.3). The black body itself is sufficiently hot to prevent ice forming, but these two other surfaces are not.

Future designs should take this into account, by designing the instrument in such a way that there are no cold surfaces at the front of the black body where snow could build up. Adding heaters to these sections could prevent ice forming, but would be a large energy drain and could result in unwanted infrared emission.

Another idea that could be worth investigating is a hydrophobic coating, which could theoretically prevent ice from forming on the cold surfaces. Investigation would need to be done to determine whether this is an effective or practical

method in the extremely cold and dry Antarctic conditions.

10.2 Optics

It was determined by optical modelling that NISM is not making use of the full size of its lens. As can be seen from Figure 3.8, rays from slightly off axis do not land on the detector, so rays from the edge of the lens will not be focused to the detector either. This is clearly a waste of light, and was not intended in the original design. It is possible that the wrong curvature of the lens was used when designing the optical layout.

Another possible change is to move the blocking filter to inside the cryostat. This would prevent light scattering off the filter and getting to the detector. This change would not have been possible with NISM, as the cryostat was a sealed module that could not easily be modified. However, it could be useful for future instrument designs.

10.3 Detector

The InSb diode used in NISM has a weakness: it is sensitive to radiation out to 5 microns. The atmosphere is far brighter at $5\text{ }\mu\text{m}$ than it is at $2.4\text{ }\mu\text{m}$, so this extra noise can overwhelm the desired signal. This fact necessitated adding the long-wavelength blocking filter. A different chip could avoid this problem. InAs diodes, for example¹, have a sharp wavelength cut off at $3.5\text{ }\mu\text{m}$, so light from $5\text{ }\mu\text{m}$ will not be observed.

Another possibility is to change the detection system entirely. Instead of a diode with a filter, a wavelength depended detector such as a spectrograph could be used. This would give a very clear measure of the sky brightness as a function of wavelength, without the worry of light leaking through the filter.

¹https://www.hamamatsu.com/resources/pdf/ssd/p10090-01_etc_kird1099e.pdf

10.4 Monitoring

Many of the problems with NISM were difficult to diagnose. We had to process the data before the problems with the motor or icing could be detected. These could have been detected much earlier if more monitoring tools were in place. The PLATO-R module has a set of webcams on it, which take pictures every 15 minutes. They are used to monitor the status of the HEAT telescope, cloud cover, and wind speed (by observing the flag). Another camera aimed at NISM could have made it much easier to detect any ice on the black body.

A camera could also have diagnosed the problem with the motor not turning, although this may have been more difficult. Given the limited data uplink, photos can only be taken periodically, and videos are probably not possible. The mirror rotates fairly rapidly, and it may not be easy to see if it is at a slightly different angle to what is reported by the encoder. But more information is always helpful. It may have been much easier to spot the motor stopping entirely.

10.5 Logistics

Inconveniently, a lot of information about NISM is unknown, such as its bandpass and exact optical layout. Assembly of the instrument was finished shortly before it was to be transported down to Antarctica, so there was not time to properly characterise it. This highlights an important consideration for Antarctic instruments: logistics in Antarctica usually require that there is a hard deadline for shipping cargo down at the start of a summer season. Missing this deadline means waiting another year for the next season. Of course, all this information is well known, but the example of NISM further highlights the requirement to have everything ready to go on schedule. Future missions need to be well planned, so that all relevant information on the instrument can be measured and recorded before it is installed on the plateau.

Chapter 11

Conclusion

The instrument NISM (the Near Infrared Sky Monitor) was installed in January 2015 at Ridge A, a site on the high Antarctic plateau. Ridge A has superb conditions for infrared astronomy, including clear skies, extremely low temperatures, and an extremely dry atmosphere, with PWV levels regularly dropping below 0.1 mm. These conditions give ridge A a very low infrared sky background, giving it the best observing conditions on Earth for infrared and terahertz astronomy. NISM took measurements of the infrared sky background in the $2.4\text{ }\mu\text{m}$ K_{dark} band during the 2015 winter season, in order to definitively determine the quality of the site.

Modelling of the atmosphere above the site supported the excellent quality of Ridge A, predicting very low thermal emission from the sky.

Unfortunately, a number of problems were discovered on processing the data. NISM has a stepper motor that rotates the beam in elevation. The stepper motor has an optical encoder that keeps track of its position, to determine the angle of elevation of the beam. It was discovered that the optical encoder was slipping positions as it rotated, so the reported position could be incorrect. Without the position data, it is not possible to tell which data points are from the zenith. This problem was corrected by looking at the signal in the data, and using it to determine the period that the motor was rotating on using phase dispersion minimisation. Future experiments should make it a priority to monitor and maintain the motor's operating temperature, to avoid this problem.

With the motor slip corrected, the next step was to calibrate the data, using the two possible calibration sources: the Milky Way, and the attached black body reference. By comparing the observed signal from NISM with measurements of the Milky Way from the DIRBE instrument on the COBE satellite, it was possible to determine a conversion factor. However, using this conversion factor resulted in a measurement of the median zenith spectral radiance of $294 \mu\text{Jy arcsec}^{-2}$. This measurement is far higher than expected. Measurements from the south pole found a zenith spectral radiance of $\sim 120 \mu\text{Jy arcsec}^{-2}$, and Ridge A is colder, higher, and drier, so by all expectations should have a zenith spectral radiance well below $100 \mu\text{Jy arcsec}^{-2}$.

Using the black body as a calibration reference also proved challenging. The brightness of the black body should be solely dependent on its temperature, but it was found that this relationship varied over the course of the year. After investigation of the black body temperature behaviour using dynamic time warping, it was determined that the most likely cause of the discrepancy was ice covering the aperture of the black body, reducing its beam filling factor and hence its brightness. This ice cover varied throughout the year, making the black body an unreliable reference. Future experiments should take this icing into account when designing their instruments, making sure that surfaces where ice could grow or settle are not close to the line of sight.

The Milky Way was chosen as a more reliable calibration reference, but it still resulted in an anomalously high reading. It is possible that there is some other source of light that caused these high readings. The aurora was investigated as a source, but auroral activity showed no correlation with the observed brightness of the sky in the infrared.

The last explanation for the high readings is that there is a red leak in the filter, letting in extra light at longer wavelengths. The original NISM instrument suffered from a red leak, so an additional short-pass filter was added to block this. However, filters only block to some high percentage, not usually 100%. The atmosphere is far brighter at wavelengths longer than K_{dark} , so it is possible that a very small leak caused the higher than expected readings. Future experiments

should take care to make the long wavelength blocking as complete as possible, or to use some other detector that has a higher wavelength resolution, such as a spectrograph.

It is not possible to characterise the leak using the data returned from NISM, as there are no spectrally calibrated sources visible. The only solution is to recover the instrument and return it to a lab in Sydney, where a proper analysis of the filter can be carried out. With the leak characterised, it should be possible to calibrate the 2015 data from Ridge A, and conclusively determine the quality of the site.

A recovery mission was planned in the 2017/2018 summer season, but was unable to reach the site. During the previous years, sastrugi had built up over the Ridge A skiway. On attempting to land, the pilots determined that it was too bumpy, and that there was a risk of breaking the undercarriage.

The next possible mission is in the 2018/2019 summer. By then the sastrugi may have smoothed out, or a groom team could be sent to prepare the skiway. A definitive measurement of the infrared sky background will have to wait until then.

Bibliography

- Aristidi, E., Agabi, K., Azouit, M., et al. 2005, An analysis of temperatures and wind speeds above Dome C, Antarctica, *Astronomy and Astrophysics*, 430, 739
- Ashley, M. C. B. 2012, Site characteristics of the high Antarctic plateau, *Proceedings of the International Astronomical Union*, 8, 15
- . 2013, Care and Feeding of an Antarctic Telescope, *Physics Today*, 66, doi:<https://doi.org/10.1063/PT.3.1987>
- Ashley, M. C. B., Burton, M. G., Storey, J. W. V., et al. 1996, South Pole Observations of the Near-Infrared Sky Brightness, *Publications of the Astronomical Society of the Pacific*, 108, 721
- Bailey, J., & Kedziora-Chudczer, L. 2012, Modelling the spectra of planets, brown dwarfs and stars using vstar, *Monthly Notices of the Royal Astronomical Society*, 419, 1913
- Bally, J. 1986, Interstellar molecular clouds., *Science*, 232, 185
- Benn, C. R., & Ellison, S. L. 1998, Brightness of the night sky over La Palma, *New Astronomy Reviews*, 42, 503
- Bennett, C. L., Fixsen, D. J., Hinshaw, G., et al. 1994, Morphology of the interstellar cooling lines detected by COBE., *The Astrophysical Journal*, 434, 587
- Bingham, N. R., & Ashley, M. C. B. 2014, Reducing noise from a Stirling micro cooler used with an InSb diode, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 9154, 91541V

- Bonner, C. S., Ashley, M. C. B., Cui, X., et al. 2010, Thickness of the Atmospheric Boundary Layer Above Dome A, Antarctica, during 2009, Publications of the Astronomical Society of the Pacific, 122, 1122
- Brucker, L., Picard, G., & Fily, M. 2010, Snow grain-size profiles deduced from microwave snow emissivities in Antarctica, Journal of Glaciology, 56, 514
- Burton, M. G., Yang, J., & Ichikawa, T. 2015, Astronomy from the High Antarctic Plateau, Publications of the Korean Astronomical Society, 30, 611
- Burton, M. G., Ashley, M. C. B., Braiding, C., et al. 2015, Extended Carbon Line Emission in the Galaxy: Searching for Dark Molecular Gas along the G328 Sightline, The Astrophysical Journal, 811, 13
- Carlson, R. W., Arakelian, T., & Smythe, W. D. 1992, Spectral reflectance of Antarctic snow: Ground truth and spacecraft measurements, Antarctic Journal of the US, 5, 296
- De Bernardis, P., Bagliani, D., Bardi, A., et al. 2010, in On recent developments in theoretical and experimental general relativity, astrophysics and relativistic field theories. Proceedings, 12th Marcel Grossmann Meeting on General Relativity, Paris, France, July 12-18, 2009. Vol. 1-3, 2133–2144
- Dempsey, J. T., Storey, J. W. V., & Phillips, A. 2005, Auroral Contribution to Sky Brightness for Optical Astronomy on the Antarctic Plateau, Publications of the Astronomical Society of Australia, 22, 91
- Dolci, M., Valentini, A., Tavagnacco, D., Di Cianno, A., & Straniero, O. 2016, Operational challenges for astronomical instrumentation in Antarctica: results from five years of environmental monitoring of AMICA at Dome C, 99089D
- Dong, S.-c., Wang, J., Tang, Q.-j., et al. 2018, Design of a multiband near-infrared sky brightness monitor using an InSb detector, Review of Scientific Instruments, 89, 023107
- Fretwell, P., Pritchard, H. D., Vaughan, D. G., et al. 2013, Bedmap2: improved ice bed, surface and thickness datasets for Antarctica, The Cryosphere, 7, 375

- Gattinger, R. L., & Jones, V. 1974, Quantitative Spectroscopy of the Aurora. II. The Spectrum of Medium Intensity Aurora Between 4500 and 8900 Å, Canadian Journal of Physics, 52, 2343
- Grec, G., Fossat, E., & Pomerantz, M. 1980, Solar oscillations: full disk observations from the geographic South Pole, Nature, 288, 541
- Harper, D. 1989, Infrared Astronomy in Antarctica, AIP Conference Proceedings, 198, 123
- Hauser, M. G., Kelsall, T., Moseley, Jr., S. H., et al. 1991, in American Institute of Physics Conference Series, Vol. 222, After the first three minutes, ed. S. S. Holt, C. L. Bennett, & V. Trimble, 161–178
- Hauser, M. G., Arendt, R. G., Kelsall, T., et al. 1998, The COBE Diffuse Infrared Background Experiment Search for the Cosmic Infrared Background. I. Limits and Detections, The Astrophysical Journal, 508, 25
- Hidas, M. G., Burton, M. G., Chamberlain, M. A., & Storey, J. W. V. 2000, Infrared and Submillimetre Observing Conditions on the Antarctic Plateau, 260
- Holland, W., Robson, E., Gear, W., et al. 1999, SCUBA: a common-user submillimetre camera operating on the James Clerk Maxwell Telescope., Monthly Notices of the Royal Astronomical Society, 303, 659
- Holland, W. S., Bintley, D., Chapin, E. L., et al. 2013, SCUBA-2: the 10 000 pixel bolometer camera on the James Clerk Maxwell Telescope., Monthly Notices of the Royal Astronomical Society), :, 430
- Hu, Y., Shang, Z., Ashley, M. C. B., et al. 2014, Meteorological data for the astronomical site at Dome A, Antarctica, 126, 32
- Hudson, S. R., & Brandt, R. E. 2005, A look at the surface-based temperature inversion on the Antarctic Plateau, Journal of Climate, 18, 1673

- Indermuehle, B. T., Burton, M. G., & Maddison, S. T. 2005, The History of Astrophysics in Antarctica, Publications of the Astronomical Society of Australia, 22, 73
- Kenyon, S. L., & Storey, J. 2006, A Review of Optical Sky Brightness and Extinction at Dome C, Antarctica, Publications of the Astronomical Society of the Pacific, 118, 489
- Kulesa, C. a., Ashley, M. C., Augarten, Y., et al. 2013a, Opportunities for Terahertz Facilities on the High Plateau, Proceedings of the International Astronomical Union, 8, 256
- Kulesa, C. a., Ashley, M. C. B., Augarten, Y., et al. 2013b, Opportunities for Terahertz Facilities on the High Plateau., Proceedings of the International Astronomical Union, S288, 256
- Lawrence, J., Ashley, M., Kenyon, S., et al. 2004a, A robotic instrument for measuring high altitude atmospheric turbulence from Dome C, Antarctica, Proceedings of the SPIE, Volume 5489, pp. 174-179 (2004)., 5489, 174
- Lawrence, J. S., Ashley, M. C. B., Burton, M. G., et al. 2002a, Operation of the Near Infrared Sky Monitor at the South Pole., Publications of the Astronomical Society of Australia, 19, 328
- Lawrence, J. S., Ashley, M. C. B., Burton, M. G., & Storey, J. W. V. 2002b, Observations of the Antarctic infrared sky spectral brightness, SPIE, 4836, 176
- Lawrence, J. S., Ashley, M. C. B., Tokovinin, A., & Travouillon, T. 2004b, Exceptional astronomical seeing conditions above Dome C in Antarctica, Nature, 431, 278
- Liou, K. N. 2002, An introduction to Atmospheric Radiation (Academic Press)
- Liu, G., Shepherd, G. G., & Roble, R. G. 2008, Seasonal variations of the nighttime O(1S) and OH airglow emission rates at mid to high latitudes in the context of the large scale circulation, Journal of Geophysical Research: Space Physics, 113, <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2007JA012854>

- Marks, R. D. 2002, Astronomical seeing from the summits of the Antarctic plateau, *Astronomy and Astrophysics*, 385, 328
- Marks, R. D., Vernin, J., Azouit, M., Manigault, J. F., & Clevelin, C. 1999, Measurement of optical seeing on the high antarctic plateau., *Astronomy and Astrophysics Supplement Series*, 134, 161
- Müller, M. 2007, *Information Retrieval for Music and Motion* (Springer-Verlag Berlin Heidelberg)
- Nakagawa, T., Yui, Y. Y., Doi, Y., et al. 1998, Far-infrared [CII] line survey observations of the galactic plane., *The Astrophysical Journal*, 115
- Nguyen, H., Rauscher, B. J., Severson, S., et al. 1996, The South Pole near-infrared sky brightness, *Publications of the Astronomical Society of the Pacific*, 108, 718
- Okita, H., Ichikawa, T., Ashley, M. C. B., Takato, N., & Motoyama, H. 2013, Excellent daytime seeing at Dome Fuji on the Antarctic plateau, *Astronomy and Astrophysics*, 554, L5
- Oliva, E., Origlia, L., Maiolino, R., et al. 2013, A GIANO-TNG high resolution IR spectrum of the airglow emission, *a&a*, 1
- Phillips, A., Burton, M. G., Ashley, M. C. B., et al. 1999, The Near-Infrared Sky Emission at the South Pole in Winter, *The Astrophysical Journal*, 527, 1009
- Pilbratt, G. L., Riedinger, J. R., Passvogel, T., et al. 2010, Herschel Space Observatory, An ESA facility for far-infrared and submillimetre astronomy., *Astronomy and Astrophysics*, 518
- Pineda, J., Langer, W., Valusamy, T., & Goldsmith, P. 2013, A Herschel [CII] Galactic plane survey., *Astronomy and Astrophysics*, 554
- Rabiner, L., & Juang, B.-H. 1993, *Fundamentals of Speech Recognition* (Prentice-Hall, Inc)

- Rathborne, J. M., Burton, M. G., Brooks, K. J., et al. 2002, Photodissociation regions and star formation in the Carina nebula, *Monthly Notices of the Royal Astronomical Society*, 331, 85
- Salisbury, J. W., D'Aria, D. M., & Wald, A. 1994, Measurements of thermal infrared spectral reflectance of frost, snow, and ice, *J. Geophys. Res.*, 99, 24235
- Saunders, W., Lawrence, J. S., Storey, J. W. V., et al. 2009, Where is the best site on Earth? Domes A, B, C, and F, and Ridges A and B., *Publications of the Astronomical Society of the Pacific*, 121, 976
- Scambos, T. A., Campbell, G. G., Pope, A., et al. 2018, Ultralow Surface Temperatures in East Antarctica From Satellite Thermal Infrared Mapping: The Coldest Places on Earth, *Geophysical Research Letters*, 0, <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2018GL078133>
- Shi, S.-C., Paine, S., Yao, Q.-J., et al. 2016, Terahertz and far-infrared windows opened at Dome A in Antarctica, *Nature Astronomy*, 1, 0001
- Silverberg, R. F., Hauser, M. G., Boggess, N. W., et al. 1993, Design of the diffuse infrared background experiment (DIRBE) on COBE, *Proceedings of the SPIE*, 2019, 10
- Sims, G., Ashley, M. C. B., Cui, X. Q., et al. 2012, Airglow and Aurorae at Dome A, Antarctica, *Publications of the Astronomical Society of the Pacific*, 124, 637
- Sloan, C. 2016, NASA's super 747 SOFIA: The world's biggest flying observatory, *CNN Travel*
- Smith, C. H., & Harper, D. A. 1997, Mid-Infrared Sky Brightness Site Testing at the South Pole, *Publications of the Astronomical Society of the Pacific*, 110, 747
- Stacey, G., Smyers, S., Kurtz, N., & Harwit, M. 1983, The Galaxy's 157 micron [CII] emission: Observations by means of a spectroscopic lunar occultation technique., *The Astrophysical Journal*, 268

- Stanmes, K., Tsay, C., Wiscombe, W., & Jayaweera, K. 1998, Numerically stable algorithm for discrete-ordinate-method radiative transfer in multiple scattering and emitting layered media., *Applied Optics*, 27, 2502
- Storey, J., Ashley, M., Boccas, M., Phillips, M., & Schinckel, A. 1999, Infrared Sky Brightness Monitors for Antarctica, *Publications of the Astronomical Society of the Pacific*, 111, 765
- Storey, J. W., Ashley, M. C., & Burton, M. G. 2000, in *Optical and IR Telescope Instrumentation and Detectors*, Vol. 4008, 1376–1382
- Storey, J. W. V., Ashley, M. C. B., & Burton, M. G. 1996, An Automated Astrophysical Observatory for Antarctica, *Publications of the Astronomical Society of Australia*, 13, 35
- Swain, M., & Gallée, H. 2006a, A comparison of possible Antarctic telescope locations., *Proceedings of SPIE*, 6267
- . 2006b, Antarctic Boundary Layer Seeing., *Publications of the Astronomical Society of the Pacific*, 118, 1190
- Tomasi, C., Petkov, B., Benedetti, E., Valenziano, L., & Vitale, V. 2011, Analysis of a 4 year radiosonde data set at Dome C for characterizing temperature and moisture conditions of the Antarctic atmosphere, *Journal of Geophysical Research Atmospheres*, 116, 1
- Tomasi, C., Vitale, V., Lupi, A., et al. 2007, Aerosols in polar regions: A historical overview based on optical depth and in situ observations, *Journal of Geophysical Research: Atmospheres*, 112,
<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2007JD008432>
- Walker, C., Kulesa, C., Bernasconi, P., et al. 2010, The Stratospheric THz Observatory (STO), *Proceedings of the SPIE*, 7733, 77330N
- Wootten, A., & Thompson, R. 2009, The Atacama Large Millimeter/submillimeter Array, *Proceedings of the IEEE*, 97, 1463

- Yang, H., Feng, L., Lawrence, J. S., et al. 2010, Exceptional Terahertz Transparency and Stability above Dome A, Antarctica., Publications of the Astronomical Society of the Pacific, 890, 490
- Young, E. T., Becklin, E. E., Marcum, P. M., et al. 2012, Early science with SOFIA, the Stratospheric Observatory For Infrared Astronomy., The Astrophysical Journal, 749, L17