

Gravitational and Gravity Waves: the EBEX and PMC Turbo
Balloon Experiments

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Christopher Peter Geach

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Shaul Hanany, Adviser

December 2020

ACKNOWLEDGEMENTS

First and foremost, I'd like to thank my advisor, Shaul Hanany, for pushing me when I needed pushing, and for knowing when to stop pushing. All of my best qualities as a scientist are those that I have learned from him. I have always known that he has my best interests at heart; at times, I've felt sure of little else, and that knowledge has been a great comfort.

To the Cosmolab-ers: Qi Wen and Karl Young, Kate Raach, Jeff Klein, Kyle Zilic and Asad Aboobaker, Franky Aubin and Chaoyun Bao, and to the PMC Turbo team: Bjorn Kjellstrand, Bernd Kaifler and Michele Limon, Biff Williams, Natalie Kaifler, Dave Fritts, Glenn Jones and Amber Miller – it was a pleasure to work with and learn from all of you.

To Frank Kelso: your limitless patience and generosity of spirit in explaining basic concepts of mechanical design (and, later on, less basic concepts) prevented PMC Turbo from ending up flattened at the bottom of a crater in northern Canada. You have my enduring gratitude.

To everyone at CSBF, especially Chris Schwantes and Justin Marsh: your help in preparing and launching PMC Turbo went above and beyond any reasonable expectation. Working with you all, in Palestine and in Sweden, is something I will cherish.

To Courtney Antone, my long-suffering girlfriend (and now fiancée!), whose loving companionship has buoyed up my spirits and kept me sane, and who has kept me in touch with the whole wide world that exists outside of physics. I'm sure it felt at times that this chapter of our lives would never end. Thank you for riding it out with

me.

To my friends and family: it's possible that without you to distract me, I might have finished sooner. It's much more likely that I wouldn't have finished at all. Thank you for the distractions.

DEDICATION

To my mother. I miss you.

ABSTRACT

I describe my contributions to two balloon-borne experiments in two different fields: the E & B EXperiment (EBEX), a mm-wave polarimeter that flew from Antarctica in 2013 in search of a gravitational wave signature in the cosmic microwave background (CMB); and the Polar Mesospheric Cloud Turbulence (PMC Turbo) experiment, an atmospheric science payload that flew from Sweden to Canada in 2018 in search of a signature of gravity waves in images of polar mesospheric clouds. Gravitational waves in the primordial Universe are predicted to have imprinted a specific pattern in the polarization of the CMB; a detection of this signature would constrain models of inflation. Gravity waves in our atmosphere transport momentum vertically and are an important driver for the circulations and thermal structure of the upper atmosphere.

My contributions to EBEX came after flight; I finalized aspects of the spectral and polarization calibration through a combination of analytical derivations and experimental measurements. For PMC Turbo, my role was more encompassing: I was responsible for the mechanical and thermal design of the payload; I developed the post-flight image processing algorithm to calibrate the images and to remove spurious signals; and I developed analysis techniques to identify and characterize gravity waves in the images. This work culminated with a detailed analysis of a particular event in which strong gravity wave breaking was captured in the PMC Turbo images, enabling quantification of the momentum deposition and energy dissipation associated with this event.

Contents

List of Tables	ix
List of Figures	x
1 Gravitational and Gravity Waves: Experiments in Scientific Ballooning	1
1.1 Gravitational Waves in the Cosmic Microwave Background	1
1.1.1 The E & B EXperiment	3
1.2 Gravity Wave Breaking at the Mesopause	4
1.2.1 The Polar Mesospheric Cloud Turbulence Experiment	8
2 EBEX Calibration	10
2.1 Spectral Calibration	10
2.2 Polarization Calibration	18
2.2.1 Analytical Polarization Calibration	18
2.2.2 Experimental Polarization Calibration	24
2.3 Polarized Thermal Emission from the Mirrors	26
3 PMC Turbo Instrumentation	31
3.1 Camera Systems	31
3.2 Rayleigh Lidar	34

3.3	Power System	35
3.3.1	Power Budget at Float	36
3.3.2	Power Budget during Ascent	40
3.3.3	Power System Performance	44
3.4	Mechanical, Thermal and Optical Design	46
3.4.1	Mechanical Design	46
3.4.2	Thermal Design	54
3.4.3	Optical Design	63
4	PMC Turbo 2018 Flight and Data	71
4.1	PMC Turbo 2018 Dataset	73
4.2	Image Calibration	73
4.2.1	Stray Light Contamination	79
4.3	Image Projection	80
5	PMC Turbo Data Analysis	83
5.1	Linear Gravity Wave Theory	84
5.2	Mechanisms for Brightness Perturbations in Camera Images of PMCs	88
5.2.1	Horizontal Convergence at the PMC Layer	89
5.2.2	Temperature Fluctuations	92
5.3	Identifying a Gravity Wave in Image and Lidar Data	95
5.3.1	Spanwise Averaging of Camera Data	95
5.3.2	Spatial Power Spectrum of Image Data	98
5.3.3	Identifying Gravity Waves in Lidar Data	99
5.4	Lidar and Camera Correlations	100
5.5	PMC Layer Wind Estimates	103
5.5.1	Wind Estimates from <i>Trackpy</i>	103

CONTENTS	vii
5.5.2	<i>Trackpy</i> Analysis Pipeline 112
5.5.3	Wind Estimates from NAVGEM 113
6	Case Study: Gravity Wave Breaking and Vortex Ring Formation 114
6.1	Vortex Rings 116
6.2	Characterization of Gravity Wave and Background Environment . . . 117
6.2.1	Gravity Wave Characterization 117
6.2.2	Power Spectrum Analysis 118
6.2.3	Confirmation of Gravity Wave Identification in Lidar Data . . 120
6.2.4	Local and Background Wind Estimates 121
6.3	Vortex Ring Characterization 123
6.4	Interpretation and Discussion 124
6.4.1	Momentum Deposition by the Gravity Wave 125
6.4.2	Comparison to Numerical Simulation 127
6.4.3	Vertical Structure of the Vortex Rings 129
References	130
A	PMC Turbo Stress Report 138
A.1	Gondola Drawings 140
A.2	Margins of Safety Tables 142
A.3	Suspension System Hand Calculations 149
A.3.1	Load Cases 150
A.4	Ballast Attachment Hand Calculations 159
A.5	Suspension System Finite Element Analysis 161
A.6	Rigid Rod Analysis 167
A.6.1	Buckling Analysis 167
A.6.2	Tension Analysis 168

A.7 Gondola Cage Finite Element Analysis	171
A.8 Ballast Attachments Finite Element Analysis	181
A.9 Middle Floor Finite Element Analysis	192
A.10 Middle Floor Welds Finite Element Analysis	198
A.11 Battery Floor Finite Element Analysis	204
A.12 Radiator Finite Element Analysis	209
A.13 Radiator Welds Finite Element Analysis	215
A.14 Solar Array Finite Element Analysis	221
A.15 Camera Mounts Finite Element Analysis	227

List of Tables

2.1	Uncertainties of polarization calibration parameters	24
3.1	PMC Turbo power consumption	39
3.2	Albedo and Earth IR flux values	58
3.3	Allowable, expected and measured temperatures	60
6.1	Gravity wave and background environment parameters	126
A.1	Margin of safety table for PMC Turbo subsystems.	147
A.2	Margin of safety table for PMC Turbo COTS components.	148

List of Figures

1.1	<i>Planck</i> 2015 temperature power spectrum	3
1.2	Example PMC images	6
2.1	Schematic of the EBEX optical system	11
2.2	Schematic of the Ebert-Fastie spectrometer	12
2.3	Calibration scans of the Ebert-Fastie spectrometer	16
2.4	Measured and predicted EBEX bands	17
2.5	Idealized polarimeter schematic	18
2.6	Achromatic HWP phase	23
2.7	Polarization calibration results	25
2.8	Polarization of thermal emission	28
2.9	Polarization of thermal emission	29
3.1	PMC Turbo gondola	32
3.2	PMC Turbo FOV diagram	32
3.3	Spatial resolution schematic	34
3.4	PMC Turbo power system schematic	36
3.5	Solar radiation and solar cell efficiency spectra	38
3.6	Ascent power and pointing plots	45
3.7	PMC Turbo suspension system	50
3.8	Pre-flight predictions for cold- and hot-case temperatures	59

3.9	In-flight instrument temperatures	62
3.10	Stitched difference images from the wide-field cameras	66
3.11	Camera FOV and cone of admitted light	67
3.12	Updated optical baffle design	68
3.13	Piggyback difference images	69
4.1	Ground track, altitude, and pointing plots for PMC Turbo	72
4.2	Sky brightness model curvature bias	75
4.3	Sky brightness correction for flat-field images	78
4.4	Image processing pipeline	81
5.1	Spatial GW perturbations	90
5.2	GW brightness perturbations	91
5.3	GW brightness perturbations due to horizontal convergence	91
5.4	GW brightness perturbations due to temperature perturbation	94
5.5	Example spanwise average binning	96
5.6	Simulated spanwise average demonstration	97
5.7	Simulated spanwise average results	97
5.8	Power spectrum of simulated data	99
5.9	Lidar vs. camera correlation grids	102
5.10	Lidar vs. camera timestream correlation	103
5.11	Recovered <i>Trackpy</i> velocities from manipulated flight data	105
5.12	Recovered <i>Trackpy</i> velocities from simulated data	106
5.13	Size dependence of the <i>Trackpy</i> algorithm	108
5.14	Memory dependence of the <i>Trackpy</i> algorithm	109
5.15	Mass dependence of the <i>Trackpy</i> algorithm	110
5.16	Occurrence dependence of the <i>Trackpy</i> algorithm	111

5.17	Example <i>Trackpy</i> results	112
6.1	Vortex rings at 2:45 UT on July 10, 2018	115
6.2	Amplitude of best-fit sinusoid vs. bin orientation angle	118
6.3	Spanwise average of flight data with best-fit peaks	118
6.4	Colorplot of spanwise average	119
6.5	Power spectrum plot of camera data at 2:45 UT on July 10, 2018 .	119
6.6	Lidar backscatter data	120
6.7	Lidar altitude and GW brightness perturbation	121
6.8	<i>Trackpy</i> velocity estimates for 2:40 to 3:00 UT	122
6.9	<i>Trackpy</i> velocity estimates for 2:40 to 3:00 UT	123
6.10	Zoom of vortex ring evolution	124
6.11	Vortex rings in lidar data and images	128
A.1	PMC Turbo gondola with annotations	140
A.2	PMC Turbo gondola with dimensions	140
A.3	Stripped PMC Turbo gondola with dimensions	141
A.4	Schematic of the suspension system	149
A.5	Schematic of the suspension system 2	151
A.6	Schematic of the suspension system 3	152
A.7	Schematic of the rotation of the suspension system	154
A.8	Geometrical constraints of the suspension system	155
A.9	Schematic of the suspension system 4	157
A.10	Schematic of the setup of the suspension system FEA	162
A.11	Suspension system FEA results 1	163
A.12	Suspension system FEA results 2	163
A.13	Suspension system FEA results 3	164

A.14	Suspension system FEA results 4	164
A.15	Suspension system FEA results 5	165
A.16	Suspension system FEA results 6	165
A.17	Suspension system FEA results 7	166
A.18	Suspension system FEA results 8	166
A.19	Rigid rod FEA buckling analysis	167
A.20	Rigid rod FEA setup schematic	168
A.21	Rigid rod FEA results 1	169
A.22	Rigid rod FEA results 2	170
A.23	Rigid rod FEA results 3	170
A.24	Gondola cage FEA analysis setup	171
A.25	Gondola cage FEA results 1	172
A.26	Gondola cage FEA results 2	173
A.27	Gondola cage FEA results 3	174
A.28	Gondola cage FEA results 4	175
A.29	Gondola cage FEA results 5	176
A.30	Gondola cage FEA results 6	177
A.31	Gondola cage FEA results 7	178
A.32	Gondola cage FEA results 8	179
A.33	Gondola cage FEA results 9	180
A.34	Setup of the ballast attachment FEA	181
A.35	Ballast attachment FEA results 1	182
A.36	Ballast attachment FEA results 2	183
A.37	Ballast attachment FEA results 3	183
A.38	Ballast attachment FEA results 4	184
A.39	Ballast attachment FEA results 5	184

A.40	Ballast attachment FEA results 6	185
A.41	Ballast attachment FEA results 7	185
A.42	Ballast attachment FEA results 8	186
A.43	Ballast attachment FEA results 9	186
A.44	Ballast attachment FEA results 10	187
A.45	Ballast attachment FEA results 11	187
A.46	Ballast attachment FEA results 12	188
A.47	Ballast attachment FEA results 13	188
A.48	Ballast attachment FEA results 14	189
A.49	Ballast attachment FEA results 15	189
A.50	Ballast attachment FEA results 16	190
A.51	Ballast attachment FEA results 17	190
A.52	Ballast attachment FEA results 18	191
A.53	Ballast attachment FEA results 19	191
A.54	Setup of the middle floor FEA	192
A.55	Middle floor FEA results 1	193
A.56	Middle floor FEA results 2	193
A.57	Middle floor FEA results 3	194
A.58	Middle floor FEA results 4	194
A.59	Middle floor FEA results 5	195
A.60	Middle floor FEA results 6	195
A.61	Middle floor FEA results 7	196
A.62	Middle floor FEA results 8	196
A.63	Middle floor FEA results 9	197
A.64	Setup of the middle floor welded parts FEA	198
A.65	Middle floor FEA results 10	199

A.66	Middle floor FEA results 11	199
A.67	Middle floor FEA results 12	200
A.68	Middle floor FEA results 13	200
A.69	Middle floor FEA results 14	201
A.70	Middle floor FEA results 15	201
A.71	Middle floor FEA results 16	202
A.72	Middle floor FEA results 17	202
A.73	Middle floor FEA results 18	203
A.74	Setup of the battery floor FEA	204
A.75	Battery floor FEA results 1	205
A.76	Battery floor FEA results 2	205
A.77	Battery floor FEA results 3	206
A.78	Battery floor FEA results 4	206
A.79	Battery floor FEA results 5	207
A.80	Battery floor FEA results 6	207
A.81	Battery floor FEA results 7	208
A.82	Setup of the radiator FEA	209
A.83	Radiator FEA results 1	210
A.84	Radiator FEA results 2	210
A.85	Radiator FEA results 3	211
A.86	Radiator FEA results 4	211
A.87	Radiator FEA results 5	212
A.88	Radiator FEA results 6	212
A.89	Radiator FEA results 7	213
A.90	Radiator FEA results 8	213
A.91	Radiator FEA results 9	214

A.92	Setup of the radiator welded parts FEA	215
A.93	Radiator FEA results 10	216
A.94	Radiator FEA results 11	216
A.95	Radiator FEA results 12	217
A.96	Radiator FEA results 13	217
A.97	Radiator FEA results 14	218
A.98	Radiator FEA results 15	218
A.99	Radiator FEA results 16	219
A.100	Radiator FEA results 17	219
A.101	Radiator FEA results 18	220
A.102	Setup of the solar array FEA	221
A.103	Solar array FEA results 1	222
A.104	Solar array FEA results 2	222
A.105	Solar array FEA results 3	223
A.106	Solar array FEA results 4	223
A.107	Solar array FEA results 5	224
A.108	Solar array FEA results 6	224
A.109	Solar array FEA results 7	225
A.110	Solar array FEA results 8	225
A.111	Solar array FEA results 9	226
A.112	Setup of the solar array FEA	227
A.113	Camera mounts FEA results 1	228
A.114	Camera mounts FEA results 2	229
A.115	Camera mounts FEA results 3	230
A.116	Camera mounts FEA results 4	231
A.117	Camera mounts FEA results 5	232

A.118	Camera mounts FEA results 6	233
A.119	Camera mounts FEA results 7	234
A.120	Camera mounts FEA results 8	235
A.121	Camera mounts FEA results 9	236

Chapter 1

Gravitational and Gravity Waves: Experiments in Scientific Ballooning

Since the early 1990s, NASA has conducted hundreds of launches of long-duration balloons addressing a wide range of science questions. I had the fortune of working on two such balloon missions, the E & B EXperiment (EBEX) and the Polar Mesospheric Cloud Turbulence (PMC Turbo) experiment. These two experiments had deceptively different goals: the detection of the signature of gravitational waves and the detection of the signature of atmospheric gravity waves, respectively.

1.1 Gravitational Waves in the Cosmic Microwave Background

The cosmic microwave background (CMB) provides a look back into the early history of the Universe. The light that now reaches us was last scattered in the epoch of recombination, $\sim 380,000$ years after the Big Bang. The existence of the CMB, and in particular its isotropy to the level of 10^{-5} , indicate that all the visible Universe was once in causal contact, providing compelling evidence for the paradigm of inflation

[1, 2]. Beyond that, the anisotropies that remain contain a wealth of information about the Universe, as phenomena of all sorts have had a chance to imprint their signature on this ubiquitous signal.

The primary tool in studying the anisotropies of the CMB is the angular power spectrum, in which fluctuations are plotted according to their angular scale; see an example in Figure 1.1. The temperature anisotropies were first measured by the COBE satellite in the 1990's [3], revealing the dipole in the temperature of the CMB that arises from our motion relative to the CMB frame. Measurements of higher-order multipoles by COBE demonstrated the existence of angular structure on smaller angular scales. This marked the beginning of the modern era of precision cosmology. As measurements with ever-increasing sensitivity and resolution have been made by instruments such as BOOMERANG [4], MAXIMA [5], WMAP [6] and *Planck* [7], among many others, the CMB has quantified the density parameters of baryonic matter, dark matter, and dark energy, demonstrated the flatness of the Universe, and provided constraints on a wide range of cosmological parameters, from the Hubble constant to the optical depth to reionization. Figure 1.1 shows the temperature power spectrum as measured by *Planck*, together with the best-fit Λ CDM-model, demonstrating exquisite agreement.

Today, the anisotropies in the polarization maps of the CMB are of particular interest. These maps can be decomposed into 'E-mode' (curl-free) and 'B-mode' (divergence-free) patterns, each of which have distinct causes. The only source of primordial B-modes was tensor perturbations associated with gravitational waves in the early Universe, while E-modes were primarily caused by scalar perturbations associated with density fluctuations. E-modes have been observed by a variety of experiments (e.g. [4, 5, 8]), as have B-modes formed by gravitational lensing of primordial E-modes [9]. Primordial B-modes remain elusive, though ever-decreasing




Figure 1.1: (Top) *Planck* 2015 temperature power spectrum (blue dots) showing 1σ error bars and the best-fit Λ CDM spectrum (red line). (Bottom) Residuals with respect to the best-fit spectrum, showing exquisite agreement. Figure is taken from [7].

upper-limits have been imposed on the magnitude of that signal (most recently by [10]). A detection of primordial B-modes would constrain the energy scale of inflation and provide insight into the first 10^{-36} s after the Big Bang.

Detection of primordial B-Modes is complicated by the presence of galactic foregrounds: synchrotron emission and polarized dust. Incomplete knowledge of these foregrounds requires CMB observations over multiple frequency bands to leverage the different spectral signature of these signals.

1.1.1 The E & B EXperiment

The E & B EXperiment (EBEX) was designed, built, and launched with this goal in mind, with 2,000 detectors spread across three frequency bands centered around 150, 250, and 410 GHz. Polarimetry was achieved through the use of a continuously rotating half-wave plate (HWP) and polarizing grid. EBEX launched on December 29,

2012 from McMurdo Station in Antarctica and operated for 11 days. The payload, flight, detector characterization and data analysis are described in [11–16] among many others. I describe my contributions to the calibration of the instrument after flight in Chapter 2.

1.2 Gravity Wave Breaking at the Mesopause

Fast-forwarding through 13.8 billion years and zooming in from the scale of the Universe to our little blue planet, we turn our attention to another type of wave. Gravity waves (GWs), or buoyancy waves, are waves for which gravity provides the restoring force. Ripples on the surface of a pond into which a pebble has been thrown are an example of such a wave (though in that case, surface tension of the pond also plays an important role). GWs are also present, indeed ubiquitous, in the air around us. GWs exist at all levels of the atmosphere, but are of particular interest in the upper atmosphere. Upward propagating GWs grow exponentially in amplitude in response to decreasing density, transporting momentum and energy vertically in the process. This process leads to systematic forcing on global scales that shapes the thermal structure and circulation patterns at levels from the troposphere into the thermosphere ([17], and references there-in).

GWs arise from a variety of causes, primarily in the troposphere. Orographic GWs occur when surface winds encounter topographical features (mountains, for example) that disrupt the flow and cause GWs to propagate upwards; the topographical features exert a drag on the mean flow, and GWs transport this drag vertically before breaking at higher altitudes. Non-orographic sources include convective cells in the troposphere, such as those associated with thunderstorms, wherein the induced upwelling generates GWs at a range of scales. Additionally, secondary GWs can arise

due to momentum deposition associated with GW breaking in the upper atmosphere, resulting in upward- and downward-propagating waves.

The GW scales that are thought to contribute most to the momentum flux have horizontal wavelengths of less than 100 km [18]. This is below the resolution of global circulation models (GCMs) that attempt to describe climate and weather patterns of the Earth; thus, GCMs must rely on parameterizations of this effect. These parameterizations are poorly constrained by sparse observations and rely on simplifying assumptions that fail to describe known GW behaviour. Further observations of GWs and in particular of GW breaking are needed to further constrain their influences on the circulation and structure of the atmosphere.

A challenge in characterizing GW drag in GCMs is the spatial variability and temporal intermittency of the phenomenon. A wide array of experiments have measured a range of GW parameters over different regions of the GW spectrum and at different altitudes, but none offer a complete picture. Rayleigh lidars measure backscatter along a (typically) vertical profile, revealing GW-induced density perturbations. Other lidars are tuned to specific resonances and measure the vertical structure of corresponding layers of the atmosphere - for example, the band of neutral sodium that exists between altitudes of 80 and 100 km. Ground-based imagers also take advantage of sodium and other atmospheric constituents that make up the airglow. Since these phenomena are temperature-dependent, they are sensitive to the temperature perturbations associated with GWs; due to the thickness of the airglow layer, airglow imagers have limited sensitivity to GWs with vertical wavelengths less than ~ 15 km. Satellite-based instruments, such as the CIPS instrument aboard the AIM satellite, provide more global coverage; CIPS has recently demonstrated sensitivity to GWs with horizontal wavelengths from 15 to 600 km and vertical wavelengths greater than 15 km, so long as the GWs occur at altitudes of 50 to 55 km.




Figure 1.2: Example PMC images taken from the ground (left, from [17]) and from the International Space Station (right, [21]).

In addition to uncertainty regarding the global sources and spectra of GWs, parameterizations at present are also limited in their description of GW breaking. Improving these descriptions represents a separate observational challenge, requiring high spatial and temporal resolution to characterize the transition from coherent wave structure at horizontal wavelengths of 10's of km to turbulence at the inner scale.

A promising window into these GW breaking dynamics, used by lidars and imagers alike, makes use of polar mesospheric clouds (PMCs, also known as noctilucent clouds). PMCs are clouds of ice particles that form at the temperature minimum of the polar summer mesopause, at an altitude of approximately 83 km, and at high latitudes (typically poleward of 55°). The temperature at these altitudes can reach below 120 K [19]; despite the low concentration of water vapor, these low temperatures allow for the growth of ice crystals. Ice crystals are thought to nucleate on meteoric dust particles, though the possibility of homogenous nucleation has also been suggested [20]. The PMC cloud layer can be as thin as 100 m, meaning that small-scale features with characteristic scales down to ~ 10 m can be resolved ¹.

¹Though high-resolution images of the airglow can also be captured, small-scale structure is highly unlikely to be coherent over the several km thickness of the airglow layer and will therefore be suppressed

Figure 1.2 shows example images of PMCs, taken from the ground and from space. The structure visible in the images is due to modulation of the PMC field by GWs and their associated instability structures. Horizontal and vertical wind perturbations associated with GWs distort the PMC layer, while temperature perturbations associated with these vertical motions result in enhanced growth or sublimation of the PMC particles; it is through these mechanisms that GWs imprint their signature onto the PMC layer. These mechanisms are discussed in greater detail in Section 5.2.

PMCs have been used to study GWs for 60 years [22], despite significant experimental challenges in capturing the PMC signal. Observations from satellites are limited by the large viewing distance and high orbital speed: the state-of-the-art PMC-observing imagers aboard the AIM satellite, for example, have spatial resolution of >2 km per pixel. Ground-based imaging of PMCs has its own challenges. PMCs are optically thin, and the signal is dominated by Rayleigh scattering from the intervening atmosphere when the sun is above the horizon. They are visible only when the sun is $>6^\circ$ below the horizon, but close enough to the horizon that the clouds themselves are still illuminated. Since PMCs form during a 2-month interval in each hemisphere's summer, the prevalence of these conditions is rare, and observations are correspondingly sporadic. Furthermore, such observations occur at large off-zenith angles which increases viewing distance and reduces spatial resolution of the images and apparent brightness of the clouds. A new observational window into these processes was discovered serendipitously by EBEX in its 2012-2013 flight.

Pointing reconstruction for EBEX was achieved through the use of two star cameras on the payload. These consisted of commercial CCD-based cameras with narrow field-of-view (FOV) lenses housed in cylindrical pressure vessels. Images taken in flight showed fields of stars that were compared to catalogs of known star locations to solve for the pointing of the telescope with high precision. When these images

were inspected after flight, it was discovered that many of them contained cloud-like features that were identified as polar mesospheric clouds .

The PMC images taken by EBEX represented a new observational approach. Long-duration balloon payloads, like EBEX, float at an altitude of 35-40 km, above $\sim 99\%$ of the Earth's atmosphere. This implies a corresponding reduction in sky brightness. The result, as shown by EBEX, is that PMCs can be observed even with the sun above the horizon. These observations can occur at much smaller off-zenith angles (limited only by the region obstructed by the balloon itself, which extends to $<30^\circ$) with reduced viewing distance and increased spatial resolution. The EBEX images had spatial resolution of 3 m per pixel when projected to the PMC layer; these are the highest resolution images of PMCs ever taken. The ability to image the PMCs with such high resolution is significant, because the inner scale of turbulence (i.e. the scale at which viscous forces begin to dominate over inertial forces) is approximately 10 m at PMC altitudes. Thus, balloon-borne imagers can provide sensitivity to the full range of scales involved in GW breaking, from coherent GWs with horizontal wavelengths on the order of 100 km to turbulence at the inner scale. The EBEX star cameras had two shortcomings with respect to PMC observations: the total FOV was 3x4 km when projected to the PMC layer, and the interval between consecutive images was 30 s. The resulting images gave no information on large-scale dynamics, and little information on temporal evolution since the same features were rarely captured multiple times.

1.2.1 The Polar Mesospheric Cloud Turbulence Experiment

The Polar Mesospheric Cloud Turbulence experiment (PMC Turbo) was proposed as a dedicated PMC-observing mission optimized for capturing the evolution of GW dynamics over four spatial orders of magnitude at high cadence using a suite of seven

16 MP CCD-based cameras; a Rayleigh lidar provided complementary information on the vertical structure of the PMC layer. I describe my role in the design and development of the PMC Turbo instrument in Chapter 3; the 2018 flight of PMC Turbo, the data thereby obtained and the algorithm I developed to calibrate the raw data are described in Chapter 4; analysis techniques that I developed to aid in the interpretation of that data are described in Chapter 5; and the application of those analysis techniques to a segment of flight of particular interest is described in Chapter 6. Appendix A contains a report on the stress analysis I performed to demonstrate the compatibility of the PMC Turbo gondola with the NASA guidelines for structural integrity of balloon-borne payloads.

Chapter 2

EBEX Calibration

The E and B Experiment (EBEX) was a balloon-borne mm-wave polarimeter that was designed to measure the polarization of the CMB. It launched from McMurdo Station in December of 2012, flew for 25 days and collected 11 days of science data. A schematic of the optical system is shown in Figure 2.1. An ambient temperature Gregorian-Dragone reflecting telescope was coupled to a cryogenic receiver that housed a continuously-rotating, achromatic half-wave plate (AHWP) to modulate the polarization signal and 1,500 bolometric transition edge sensors in three frequency bands, centered at 150 GHz, 250 GHz, and 410 GHz. A polarizing wire grid split the incoming light into two polarization states, each of which was coupled to a separate focal plane.

2.1 Spectral Calibration

Before flight, other team members had performed a spectral calibration of the instrument in order to characterize the spectral bands. Spectral calibration was important because various frequency-dependent effects appear in the science data. For example, the measured signal in flight was composed of contributions from the CMB and from galactic foregrounds, among others. The relative strength of these signals was




Figure 2.1: Ray trace of the EBEX optical system (left) and a zoom of the cold optics (right) with components labelled. Figure taken from [11].

frequency-dependent. Similarly, the AHWP introduced frequency-dependent rotations to the polarization signal (see Section 2.2). Properly accounting for these and other systematic effects required accurate knowledge of the relative spectral response of each frequency band. The overall calibration of the EBEX spectral bands has been described by [13]; I describe here the small component of this calibration for which I was responsible.

The relative spectral calibration was performed using an Ebert Fastie (EF) monochromator. The EF, shown in Figure 2.2, consisted of a 1200 K blackbody source at the entrance aperture, a spherical mirror, and a diffraction grating that selected the output wavelength according to the grating equation:

$$m\lambda = d(\sin \alpha + \sin \beta) \quad (2.1)$$

where m is the diffraction order, λ is the output wavelength, d is the spacing between adjacent grooves of the grating, and α and β are the incident and reflected angles off




Figure 2.2: Schematic of the Ebert-Fastie spectrometer, showing the relevant geometry. Figure adapted from [23].

the grating (measured relative to the grating normal). The EBEX EF was designed to use 2nd order diffraction via diffraction gratings with blaze angles of 30° and groove spacings of 4.0, 2.4, and 1.5 mm for the 150, 250, and 410 GHz bands, respectively.

Given the geometry of the EF design, the angles are constrained such that

$$\alpha - \beta = 2\theta, \quad (2.2)$$

where $\theta = 8.93^\circ$ is as shown in Figure 2.2. Furthermore, we have

$$\beta = \phi - \theta, \quad (2.3)$$

where ϕ is the grating angle, measured relative to the plane containing the entrance and exit apertures. We can rearrange (2.1) to find the frequency at the exit aperture

as a function of ϕ :

$$\nu = \frac{mc}{2d \sin \phi \cos \theta}. \quad (2.4)$$

Thus, by rotating the grating one changes the output frequency of the EF. The spectral resolution of the EF is the quadrature sum of two terms: the resolving power of the EF and the spectral broadening due to the finite width of the entrance and exit apertures [24]:

$$\frac{\Delta\nu}{\nu} = \sqrt{\left(\frac{1}{mN}\right)^2 + \left(\frac{[1 + (l/F)^2] \cos(\phi - \theta) \Delta l}{2 \cos \theta \sin \phi F}\right)^2}, \quad (2.5)$$

where N is the number of grooves on the grating (37.5, 62.5, and 100 for 150, 250 and 410 GHz, respectively), $l = 110$ mm is the distance between the exit aperture and the symmetry axis of the EF, $F = 700$ mm is the focal length of the spherical mirror, and $\Delta l = 12$ mm is the width of the exit aperture. For the EF design used here, the second term dominates at all diffraction orders, and

$$\frac{\Delta\nu}{\nu} \simeq \frac{[1 + (l/F)^2] \cos(\phi - \theta) \Delta l}{2 \cos \theta \sin \phi F}. \quad (2.6)$$

For the EF design used here, we have $\frac{\Delta\nu}{\nu} \simeq 0.035$, independent of frequency band and diffraction order.

In the pre-flight calibration of the EBEX receiver, the EF was coupled to the receiver on a translational stage to determine the relative spectral response of each detector. The EBEX spectral bands were in the Rayleigh-Jeans limit of the blackbody source, so the relative signal strength across each band was known, independent of precise knowledge of the blackbody temperature. But the spectral response of the EF itself, in particular the diffraction gratings that were used, was not measured

prior to flight. Therefore, I was tasked with performing the spectral calibration of the EF. Another graduate student at this time worked to model the expected spectral response of the gratings using PCGrateTM [23, 25].

For this measurement, I used a small cryostat with a single bolometric detector coupled to a Winston cone, with a spectral response that we assume to be flat across the EBEX frequency bands. This cryostat was placed directly at the exit aperture of the EF. A mechanical chopper between the thermal source and the entrance aperture modulated the incoming signal, and a lock-in amplifier increased the signal-to-noise of the final measurement by reading out at the chop frequency. High- and low-pass filters for each frequency band were placed at the exit aperture of the EF; these were the same band-defining filters that were used in the EBEX receiver. I measured the frequency response of each grating by scanning across the frequency range for each band in steps of 1° (corresponding to approximately 10 GHz), integrating for 2-5 minutes at each step. The experimental setup I inherited lacked repeatability, so I made the following improvements:

- I mounted the cryostat on a 2D translational stage. This enabled me to maximize the signal by optimizing the alignment; furthermore, this alignment was repeatable between trials.
- I re-designed the brackets that were used to mount the diffraction gratings. The previous brackets had play in the attachment mechanism, so the alignment would change each time a grating was removed and replaced. The new brackets also included set screws to adjust the tilt of the gratings, improving alignment and increasing signal.
- I re-designed the baffles to minimize stray light from entering the optical path.
- I identified a time-dependence to the noise level; measurements taken outside

normal working hours had lower noise by a factor of 2-3.

With these changes, the EF gave consistent, repeatable results, but the results showed much stronger frequency-dependence than expected (relative to other gratings of similar design [26] and relative to the simulated performance).

I identified the cause of this anomalous behavior to be the leakage of high-frequency light into our measurement bands through higher-order diffraction. For a given orientation of the grating angle, corresponding to an output frequency of ν_0 , the EF will also output light with frequencies $\frac{3}{2}\nu_0$, $2\nu_0$, $\frac{5}{2}\nu_0$, etc. In particular, at the blaze angle of the grating, the efficiency remains high for higher order diffraction – as the wavelength of light becomes small relative to the groove spacing, each groove of the grating acts as a mirror oriented parallel to the plane containing the entrance and exit aperture. Additionally, the black body source has an intensity that increases with frequency, up to a peak frequency of ~ 7 THz.

When the EF was used to calibrate the EBEX receiver, multiple thermal filters prevented out-of-band light from reaching the focal planes, but in the more simplified experimental setup used to calibrate the EF, the band-defining filters provided insufficient rejection to out-of-band light. I added additional low-pass filters with known in-band response to reduce this leakage. The calibration of the 250 GHz grating, for example, originally used the one high- and two low-pass filters that defined the 250 GHz band, with cut-offs at approximately 220, 285, and 335 GHz, respectively. I added the low-pass filters from the 410 GHz band with cut-offs at approximately 445 and 560 GHz to eliminate contributions from higher-order diffraction. To test for further out-of-band leakage, I replaced the 250 GHz high-pass filter with the high-pass filter from the 410 GHz band (365 GHz cutoff) and repeated a scan across the range of grating angles. This showed a constant, angle-independent offset, which I corrected for. Figure 2.3 shows the final results for the 250 GHz grating. This mea-

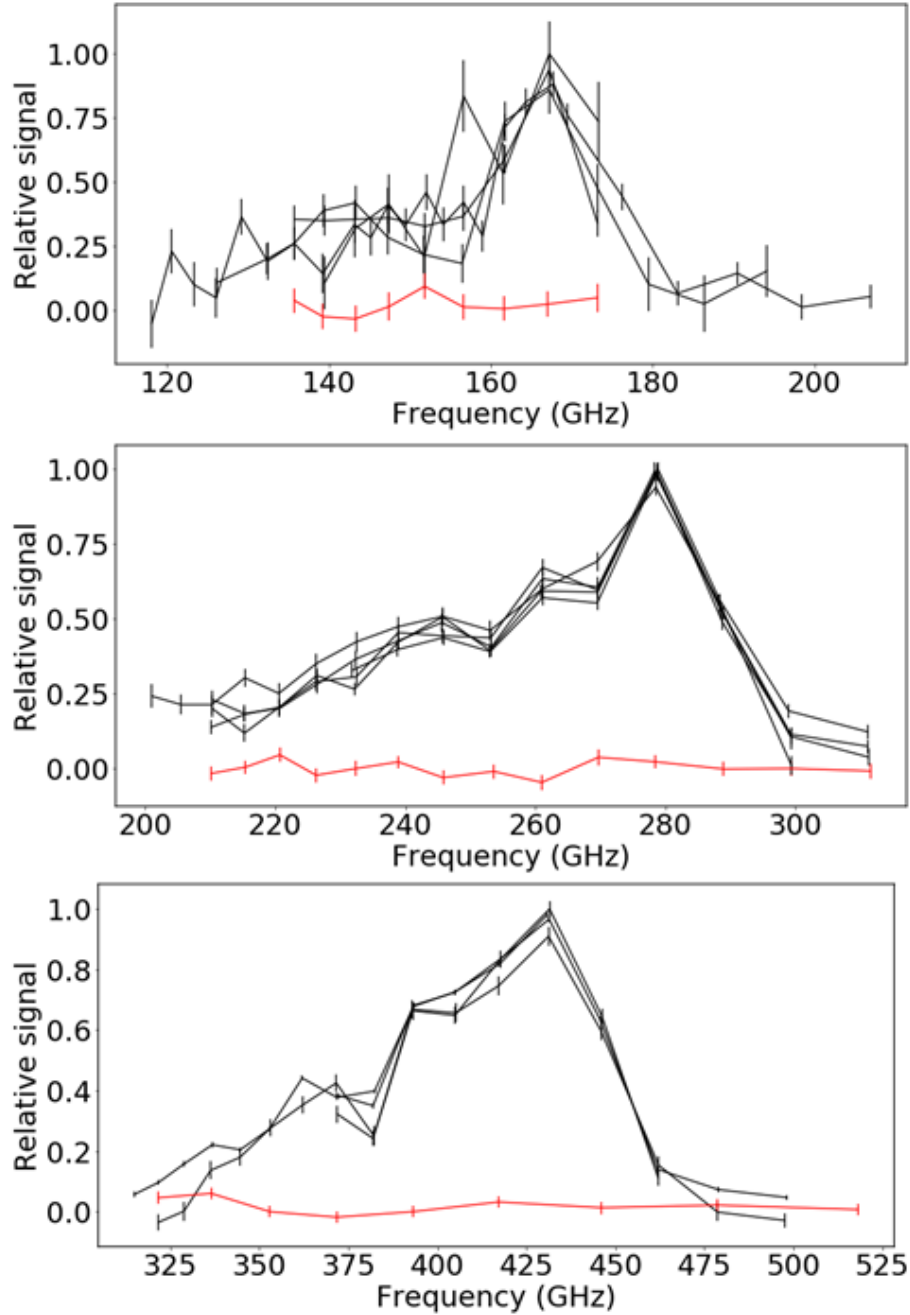


Figure 2.3: Final calibration results showing the relative spectral response for the Ebert-Fastie with the 150 GHz grating (top), 250 GHz grating (middle), and 410 GHz grating (bottom). Each plot shows a series of scans taken across multiple days in black, and a reference scan with the signal blocked in red.

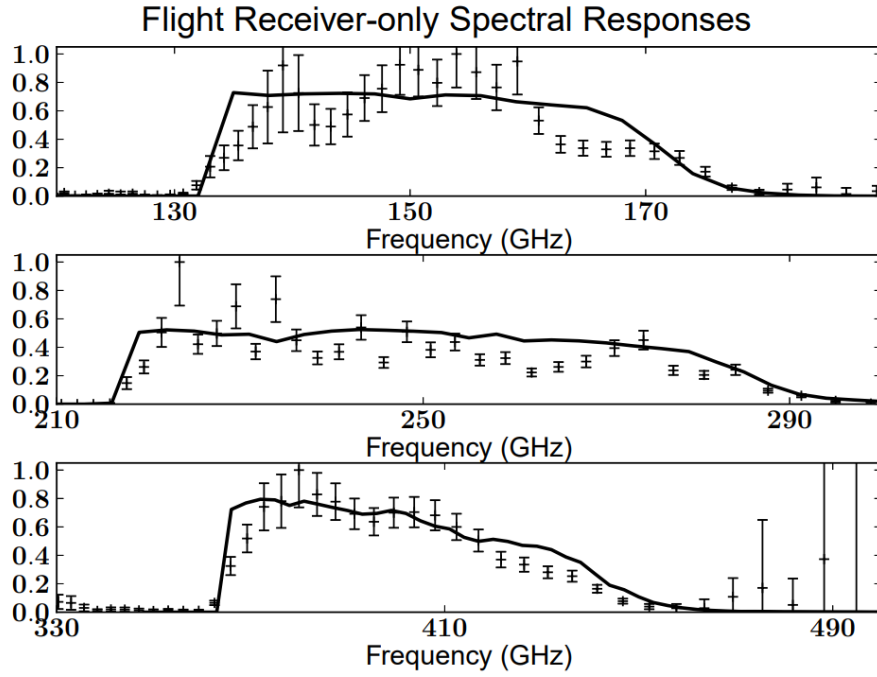


Figure 2.4: Final calibration results showing the measured and predicted EBEX spectral bands: 150 GHz (top), 250 GHz (middle), and 410 GHz (bottom). Figure taken from [13].

surement enabled a final determination of the EBEX spectral bands; Figure 2.4 shows the measured and predicted bands as detailed in [13].




Figure 2.5: Schematic of an idealized polarimeter containing a continuously rotating HWP and polarizing grid.

2.2 Polarization Calibration

Polarization calibration is the determination of the conversion of polarization signal measured at the focal plane to incoming polarization on the sky. This can be done empirically or analytically, though any analytical calibration will require empirical measurements of physical parameters and any empirical calibration will require an analytical understanding to extend the results beyond the precise conditions of the calibration measurement. For EBEX, a series of empirical calibration measurements were made, as described by [14]. My responsibility was to construct an analytical framework as a cross-check on this empirical calibration. This analytical framework was then used to adjust the empirical results to account for changes in the experimental setup between calibration and in-flight measurements.

2.2.1 Analytical Polarization Calibration

Figure 2.5 shows an idealized schematic of the basic requirements for polarimetry as implemented by EBEX. Shown is the coordinate system I will adopt, with $+z$ pointing towards the sky – thus a positive rotation will be counter-clockwise, as

viewed from above. Consider monochromatic incident radiation with Stokes vector parameters I , Q , U , and V , and let $\rho(t)$ denote the orientation of the ordinary axis of the HWP, ϵ denote the polarization modulation efficiency of the HWP, and β denote the orientation of the transmission axis of the polarizing grid. I will also refer to α_{in} , the polarization angle of incident radiation, given by $\alpha_{in} = \frac{1}{2} \arctan(\frac{U}{Q})$. Angles are measured relative to the $+x$ -axis, and Q and U are defined such that $+Q$ lies along the x -axis and $+U$ lies at $= 45^\circ$. The Mueller matrices for the HWP and the polarizing grid are M_{HWP} and M_{Grid} , respectively:

$$M_{HWP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \epsilon \cos(4\rho(t)) & \epsilon \sin(4\rho(t)) & 0 \\ 0 & \epsilon \sin(4\rho(t)) & -\epsilon \cos(4\rho(t)) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (2.7)$$

$$M_{Grid} = \frac{1}{2} \begin{pmatrix} 1 & \cos(2\beta) & \sin(2\beta) & 0 \\ \cos(2\beta) & \cos^2(2\beta) & \sin(2\beta) \cos(2\beta) & 0 \\ \sin(2\beta) & \sin(2\beta) \cos(2\beta) & \sin^2(2\beta) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (2.8)$$

The HWP and polarizing grid give a total Mueller matrix for the instrument of

$$M_{Total} = \frac{1}{2} \begin{pmatrix} 1 & \epsilon \cos[4\rho(t) - 2\beta] & \epsilon \sin[4\rho(t) - 2\beta] & 0 \\ \cos(2\beta) & \epsilon \cos[4\rho(t) - 2\beta] \cos(2\beta) & \epsilon \sin[4\rho(t) - 2\beta] \cos(2\beta) & 0 \\ \sin(2\beta) & \epsilon \cos[4\rho(t) - 2\beta] \sin(2\beta) & \epsilon \sin[4\rho(t) - 2\beta] \sin(2\beta) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (2.9)$$

The EBEX bolometers are sensitive only to intensity of incident radiation; the

signal detected at the focal plane in this idealized model is

$$S(t) = \frac{1}{2} (I + \epsilon Q \cos[4\rho(t) - 2\beta] + \epsilon U \sin[4\rho(t) - 2\beta]) \quad (2.10)$$

In general, we expect deviations from this idealized model that contribute overall rotations to the measured signal. I encode these deviations in an overall offset Φ such that

$$S(t) = \frac{1}{2} (I + \epsilon Q \cos[4\rho(t) - \Phi] + \epsilon U \sin[4\rho(t) - \Phi]) \quad (2.11)$$

The polarization signal is extracted from the timestreams by demodulating $S(t)$ with $\cos(4\rho(t))$ and $\sin(4\rho(t))$ and keeping only the DC terms:

$$Q' \equiv S(t) \cos(4\rho(t)) = \frac{\epsilon}{4} [Q \cos(\Phi) - \epsilon U \sin(\Phi)] \quad (2.12)$$

$$U' \equiv S(t) \sin(4\rho(t)) = \frac{\epsilon}{4} [Q \sin(\Phi) + \epsilon U \cos(\Phi)] \quad (2.13)$$

Noting that $U = Q \tan(2\alpha_{in})$, these can be rewritten as

$$Q' = \frac{\epsilon Q \cos(\Phi + 2\alpha_{in})}{4 \cos(2\alpha_{in})} \quad (2.14)$$

$$U' = \frac{\epsilon U \sin(\Phi + 2\alpha_{in})}{4 \sin(2\alpha_{in})} \quad (2.15)$$

Finally, these terms can be combined to solve for α_{in} , assuming that Φ is known:

$$\alpha_{in} = \frac{1}{2} \arctan \left(\frac{U'}{Q'} \right) - \frac{\Phi}{2} \quad (2.16)$$

The overall rotation Φ is therefore required to convert from the bolometer timestreams to the incident polarization angle.

For EBEX, the deviations from the idealized model described above that must be accounted for are :

- non-zero orientation of the polarizing grid,
- measurement conventions of the HWP orientation,
- frequency-dependence of the HWP effective ordinary axis, and
- bias due to differential readout times,

each of which is discussed in detail below.

Orientation of the EBEX Polarizing Grid

Measuring β in the reference frame of the instrument was done using a coordinate measuring machine (CMM). Tooling balls at the edges of the polarizing grid were used as reference points to define its transmission axis relative to tooling balls at the aperture stop; a separate set of measurements defined the location of the aperture stop tooling balls relative to another set of tooling balls at the top of the LN2 stage; and a third set of measurements defined the the location of these LN2 tooling balls relative to the tooling balls on the top plate of the cryostat that define the instrument reference frame. I combined these various measurements in SolidWorks to create a 3D model containing the tooling balls from each stage and measured the orientation of the grid transmission axis in the instrument frame.

By comparing (2.10) to (2.11), we see that Φ contains a contributions 2β due to the arbitrary rotation of the grid transmission axis.

Half-wave Plate Orientation Angle

The ordinary axis of the HWP was measured using a Fourier Transform Spectrometer over a range of 0-600 GHz and was indicated on the edge of the HWP with a black

arrow. Uncertainty in the spectrum used in that measurement resulted in an ambiguity of $\pm 3^\circ$ in the interpretation of the black arrow (and therefore all subsequent HWP orientation measurements). Note that this is a potential overall offset rather than statistical uncertainty.

The orientation of the HWP was determined using an LED and photo diode modulated by a slotted ring around the circumference of the HWP (see [14]). The absolute orientation of the HWP was encoded by a single slot of double width. The orientation of the HWP was measured between these two reference points: the fixed location of the LED/photo diode, and the rotating location of the double slot. Thus, the orientation of the HWP was a combination of three terms: $\rho(t) = \tilde{\rho}(t) + \gamma + \psi$, where $\tilde{\rho}(t)$ is the position data that was written to disk, γ is the angle between the double slot and the black arrow marking the ordinary axis of the HWP, and ψ is the orientation of the LED/photo diode relative to the instrument frame. Both γ and ψ were measured using the CMM.

Adding these terms to (2.10) gives

$$S(t) = \frac{1}{2} \left(I + \epsilon Q \cos[4\tilde{\rho}(t) + 4\gamma + 4\psi - 2\beta] + \epsilon U \sin[4\tilde{\rho}(t) + 4\gamma + 4\psi - 2\beta] \right) \quad (2.17)$$

Frequency-dependence of the HWP effective ordinary axis

The HWP was comprised of 5 sapphire plates of varying thicknesses, designed to provide high polarization modulation efficiency over a broad range of frequencies. However, its performance was not uniform across that frequency range. The effective ordinary axis of the HWP was frequency dependent (see Figure 2.6, and we must introduce a frequency-dependent term, $\Delta_\phi(\nu)$, into the polarization calibration to account for this effect. Note that this term depends not only on frequency band but also on the input spectrum of light within each band. This effect is discussed in more




Figure 2.6: Phase offset of the achromatic HWP as a function of frequency. Phase offset is a measurement of the orientation of the effective ordinary axis of the 5-stack HWP relative to the ordinary axis of the topmost sapphire plate of the HWP. This figure is from [16].

detail in [16].

Including this effect gives

$$S(t) = \frac{1}{2} \left(I + \epsilon Q \cos[4\tilde{\rho}(t) + 4\gamma + 4\psi - 4\Delta_\phi(\nu) - 2\beta] + \epsilon U \sin[4\tilde{\rho}(t) + 4\gamma + 4\psi - 4\Delta_\phi(\nu) - 2\beta] \right) \quad (2.18)$$

Readout Delay Correction

Due to arcane details of the EBEX readout system, there was a $t_d = 10.69$ ms delay in the readout of the HWP orientation angle relative to the bolometer readout. Therefore, the bolometer data recorded at a time t corresponds to HWP angle data recorded at $t + t_d$. Let $\hat{\rho}(t)$ denote the actual orientation of the HWP at a time t , as opposed to the recorded orientation at that time, with $\tilde{\rho}(t) = \hat{\rho}(t + t_d)$. This

Parameter	Uncertainty
γ	0.1°
ψ	0.1°
$\delta_{readout}$	0.1°
$\Delta_\phi(\nu)$	0.7°
β	0.3°
Φ	2.9°

Table 2.1: Uncertainty on each parameter in the analytical polarization calibration prediction. The quoted uncertainty on Φ is a function of the other uncertainties using standard error propagation. This table is adapted from [11]

introduces a phase shift of $\delta_{readout} \equiv \omega t_d$ where ω is the rotation rate of the HWP:

$$S(t) = \frac{1}{2} \left(I + \epsilon Q \cos[4\tilde{\rho}(t) + 4\gamma + 4\psi + 4\delta_{readout} - 4\Delta_\phi(\nu) - 2\beta] \right. \\ \left. + \epsilon U \sin[4\tilde{\rho}(t) + 4\gamma + 4\psi + 4\delta_{readout} - 4\Delta_\phi(\nu) - 2\beta] \right) \quad (2.19)$$

Taken all together, we have

$$\Phi = -4\gamma - 4\psi - 4\delta_{readout} + 4\Delta_\phi(\nu) + 2\beta, \quad (2.20)$$

with uncertainties on each of these values given in Table 2.1. The total uncertainty on Φ is 2.9° , which translates to a 1.5° contribution to the uncertainty on α_{in} .

2.2.2 Experimental Polarization Calibration

The polarization calibration was measured in a series of tests: (1) the EF was used to measure the polarization rotation as a function of frequency within each frequency band. This test was performed with no absolute reference for the incident polarization, so the results are used to verify our model of the frequency response of the HWP within each frequency band rather than to fix the absolute polarization rotation; (2)




Figure 2.7: Polarization calibration measurements for the receiver-only test (red dots), the end-to-end test (blue dots), and the Ebert-Fastie test (gray dots). Note that the Ebert-Fastie only measured relative phase shift as a function of frequency within each band and have been adjusted vertically to fit the other measurements. Also shown is the analytical prediction (blue dashed line and shaded region). Figure is adapted from [11]

a receiver-only test was performed with an external polarizing grid, the orientation of which was measured relative to the receiver frame; and (3) an end-to-end test was performed for the whole telescope, scanning a polarized source in the far-field. For each test, Φ was determined within each frequency band. The results, together with the analytical prediction, is shown in Figure 2.7. The relative phase between frequency bands is well captured by the analytical prediction, but there is an overall offset of 12° between the prediction and the measurements. One possible explanation for this offset is the ambiguity in the meaning of the black arrow used to designate the nominal ordinary axis of the HWP discussed above - a misinterpretation of the meaning of this mark by 3° would result in the required 12° offset in Φ .

2.3 Polarized Thermal Emission from the Mirrors

The discussion above is idealized, in that it omitted any mention of spurious polarization signals. Spurious polarization signals that originate sky-side of the HWP are modulated by the HWP and result in large swings in intensity at the focal plane. In EBEX, the size of this modulated signal, or ‘template,’ was sufficient to drive the detectors out of the regime in which they respond linearly to incoming light, complicating the data analysis process. Two primary sources of spurious polarized signal were identified: instrumental polarization, primarily by the field lens, and polarized thermal emission from the warm mirrors. A discussion of these effects and their estimated contributions in comparison to the total measured signal in flight appears in [27]. Here, I discuss the second of these two effects and describe how I quantified its expected magnitude. This effect was identified by Andrei Korotkov; the calculations that follow are my own.

Light in the optical path of the telescope has angles of incidence of up to 45° . Correspondingly, light emitted from the primary mirror into the optical path of the telescope is emitted at angles of up to 45° relative to the surface normal. The associated emissivity depends on the angle of emission, θ , and on the polarization of the emitted light.

The reflectivity for light that is polarized in the plane of incidence (R_{\parallel}) and light that is perpendicular to the plane of incidence (R_{\perp}) is

$$R_{\parallel} = 1 - \frac{\epsilon_0}{\cos \theta} \tag{2.21}$$

$$R_{\perp} = 1 - \epsilon_0 \cos \theta \tag{2.22}$$

where ϵ_0 is emissivity of the surface at normal incidence [28]. The reflectivity and

absorbtivity (or emissivity) of the mirrors sum to unity, thus

$$\epsilon_{\parallel} = \frac{\epsilon_0}{\cos \theta} \quad (2.23)$$

$$\epsilon_{\perp} = \epsilon_0 \cos \theta \quad (2.24)$$

where ϵ_{\parallel} and ϵ_{\perp} are the emissivities of in-plane and out-of-plane polarization, respectively. The degree of polarization, p , of the emission at a given θ is given by

$$p \equiv \frac{\epsilon_{\perp} - \epsilon_{\parallel}}{\epsilon_{\perp} + \epsilon_{\parallel}} = \frac{\sin^2 \theta}{1 + \cos^2 \theta}. \quad (2.25)$$

Thus, polarized thermal emission from the EBEX primary mirror will have p from 0.04 to 0.33, with additional contributions from the secondary mirror; the emission that reaches the focal plane will represent an average across that range of values. Thermal emission from the mirrors is one of the dominant sources of radiative loading [12], so polarized thermal emission represents a significant source of the polarized light.

The expected contribution of this effect to the template corresponds to the average of the Stokes vectors of each ray across the field of view. Let E_p denote the Stokes vectors of the bundle of emitted rays from the primary mirror that enter the optical path of the central pixel, and let $\langle E_p \rangle$ denote an average across the bundle. Rather than propagating the emitted rays forward and computing $\langle E_p \rangle$ at the focal plane, I instead determined hypothetical incident rays, I_p that would produce reflected rays with the same Stokes vectors as the actual emitted rays (ignoring the differential reflectivity mentioned above). It is not clear how to calculate $\langle E_p \rangle$, as all of the emitted rays have different propagation directions, but calculating $\langle I_p \rangle$ is straightforward, as the hypothetical incident rays are all parallel. More importantly, the hypothetical contribution from I_p to the signal detected at the focal plane is identical to the actual contribution from the E_p .




Figure 2.8: Degree of polarization (left) and polarization angle (right) of thermal emission from the primary mirror as a function of position within the beam of the central pixel of the focal plane.

Figure 2.8 shows the degree of polarization and polarization angle of thermal emission from the primary mirror across the beam of the central pixel of either focal plane. The coordinate system used is shown in Figure 2.1, with the origin at the center (chief ray) of the beam. The degree of polarization depends almost wholly on x and increases with decreasing x (as is apparent from Figure 2.1 the range of incidence angles shown in Figure 2.1). The polarization angle depends mostly on y but is symmetric across the xz -plane and thus averages to zero. The resulting polarization is therefore entirely in the x -direction, or $+Q$ in our coordinate system. I determined equivalent incident rays on a 1 cm x 1 cm grid across the 1.05 m entrance aperture; the resulting average gave a $p = 15.5\%$.

A simpler approach to estimating this degree of polarization, used for example by [29], is to ignore the curvature of the primary mirror - thus all emitted rays have equal propagation directions and emission angles. Employing that approach for EBEX, I found the plane tangent to the primary mirror at the location of the chief ray. This gave an emission angle of 32° , corresponding to $p = 16.3\%$, a difference of 5% relative to the more sophisticated approach laid out here (with the caveat that




Figure 2.9: Degree of polarization (left) and polarization angle (right) of thermal emission from the secondary mirror as a function of position within the beam of the central pixel of the focal plane.

there is degeneracy in the simpler approach as to how to determine the orientation of the best-fit planar mirror).

I repeated this process for the thermal emission from the secondary emission, with some added complication. To determine the range of emission angles from the secondary mirror, I start with the bundle of parallel rays that make up the field of view of the central pixel and propagate it forward through its reflection off the primary mirror and then the secondary mirror. For each ray, I determine the point and angle of incidence on the secondary mirror. The set of rays that are reflected from the secondary mirror correspond to the set of emitted rays that are in the optical path of the central pixel. Given this set of emitted rays (E_s), with polarization angles and degrees again determined by their angle of emission, let I_s denote the bundle of rays that produce E_s upon reflection from the secondary mirror. Further, let I'_s denote the bundle of rays that produce I_s upon reflection from the primary mirror. Once again, I'_s consists of parallel rays, and computing $\langle I'_s \rangle$ is again straightforward.

Figure 2.9 is analogous to Figure 2.8 for emission from the secondary mirror. Again, as expected by the symmetry of the optical system, the polarization angle of

the thermal emission is symmetric across the xz -plane and $\langle I'_s \rangle$ is entirely in $+Q$. The range of emission angles is smaller for the secondary mirror than for the primary mirror, and the degree of polarization is correspondingly smaller. I found $\langle I'_s \rangle$ to have an average degree of polarization of 6.0%.

The thermal emission from the mirrors are aligned and thus add coherently. The total thermal emission from the mirrors is polarized at 10.8% for the central pixel of the focal plane. The total field of view for EBEX spans just $\pm 3.3^\circ$ each in azimuth and elevation, so the range of emission angles varies only slightly across the focal plane. Due to geometrical complexities of applying this technique individually for each position on the focal plane, I take this value to be representative of the whole focal plane.

The significance of this contribution to the template is discussed in detail in [27]. The conclusion presented there is that polarized emission from the mirrors in EBEX was subdominant to instrumental polarization induced by the field lens, but only because the measured loads in-flight were significantly in excess of the predicted loads, probably due to spillover of the beam onto warm surfaces around the mirrors. If EBEX had performed as predicted, thermal emission from the mirrors would have been the dominant source of polarized signal to be modulated by the HWP.

Chapter 3

PMC Turbo Instrumentation

In this chapter I describe the The PMC Turbo payload consisted of 7 cameras, a Rayleigh lidar, and associated power and telemetry hardware (see Figure 3.1). It was housed aboard an aluminum frame, referred to here as the gondola. A NASA-supplied rotator provided stable pointing, such that the solar panels received constant illumination and the FOVs of the cameras was centered on the anti-sun direction. Multiple antennae provided redundant uplink and downlink during flight.

3.1 Camera Systems

The PMC Turbo cameras were modeled after the EBEX cameras, with two important modifications: much larger FOVs and increased storage capacity. The cameras were Allied Vision Prosilica GT 4907 with Kodak KAI-16070 CCDs with 16 megapixels (3,232 x 4,864 pixels). Each camera was housed in a cylindrical pressure vessel that also contained a control computer and four 8 Tb hard drives. Four of the cameras, referred to here as the wide-field cameras, had Canon f/1.2 50 mm lenses with FOVs of $27 \times 40^\circ$, spanning a combined FOV of 150×100 km, with an average spatial resolution of 20 m per pixel (see discussion below). The other three cameras (the “narrow-field cameras”) had Canon f/2 135 mm lenses, with FOVs of $10 \times 15^\circ$, were positioned in the




Figure 3.1: Schematic of the PMC Turbo gondola, with the major components labelled.

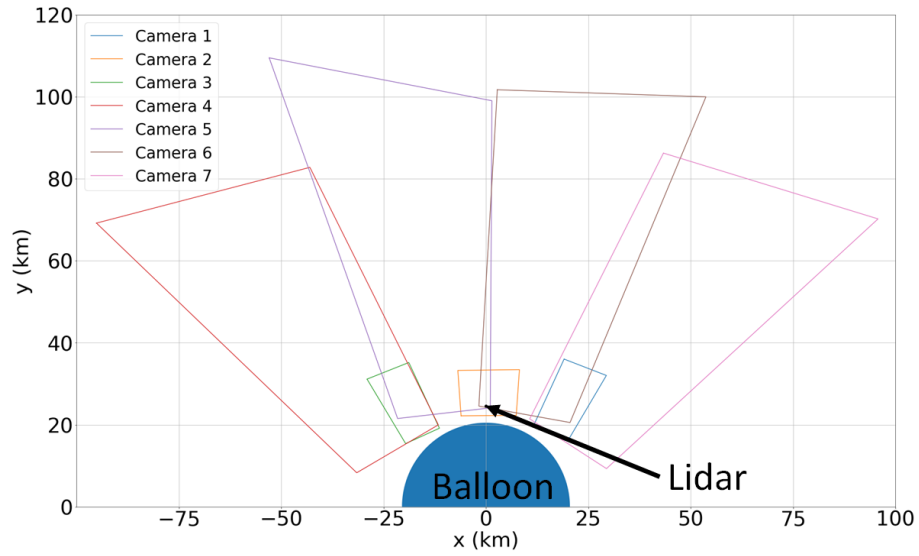


Figure 3.2: Fields of view of the PMC Turbo cameras and lidar, as projected onto the PMC layer. In the coordinate system used here, the y-axis corresponds to distance to the north, and the x-axis to distance to the east. The FOVs rotate to point anti-sun at all times; the orientation shown here corresponds to local noon, with the sun due south. The blue-shaded area is the region of sky that is obscured by the balloon.

areas of overlap of the wide-field cameras (see Figure 3.2), and had spatial resolution of 3-4 m per pixel. Images were taken in bursts of 2-4 at a cadence of 2 s per burst. Since the FOV of the cameras was projected onto the cloud layer (approximated by a tangent plane at an altitude of 83 km), the spatial resolution of the cameras varied across the FOV. Two factors contributed: (1) pixels at larger off-zenith angles imaged regions of the sky that were farther away, increasing their projected size; and (2) the PMC layer was imaged at non-normal incidence, so each (square) pixel projected to a rectangular region on the sky (i.e. the FOV of a single pixel was stretched more in the off-zenith direction than in the azimuthal direction); see Figure 3.3. Given a pixel with a FOV given by $d\theta$ in each direction oriented at an angle ϕ off-zenith, the distance to the PMC layer is given by $\frac{h}{\cos\phi}$, where h is the vertical distance to the PMC layer (approximately 45 km). The FOV of this pixel projected onto the PMC layer has dimensions of $ds_1 = \frac{hd\theta}{\cos\phi}$ in the azimuthal direction and $ds_2 = \frac{hd\theta}{\cos^2\phi}$ in the off-zenith direction. For the wide-field cameras, with $d\theta = 30$ arcseconds, the spatial resolution varied from 7x8 m per pixel at small off-zenith angles to 17x45 m per pixel at large off-zenith angles. For the narrow-field cameras, with $d\theta = 11$ arcseconds, the effect was less pronounced: the spatial resolution varied from 2.7x3.0 m per pixel to 3.2x4.3 m per pixel. Note that this assumes that the long axis of the cameras was aligned along the off-zenith direction, whereas the wide-field cameras were rotated by 12-15° relative to that alignment to create a continuous combined FOV. Furthermore, this only gives the native pixel resolution at the PMC layer and neglects the effect of motion blur due to relative motion between the PMCs and the gondola which acts to degrade the native resolution. Effective resolution during flight varied with differential wind speed, rotation rate of the gondola, and exposure time of the images.




Figure 3.3: Schematic demonstrating the geometry of the projection of a pixel onto the PMC layer, illustrating the resolution in the azimuthal direction (left panel) and in the off-zenith direction (right panel).

3.2 Rayleigh Lidar

The Rayleigh lidar was provided by the German Aerospace Center (DLR); a more detailed description appears in Kaifler et al. [30]. It consisted of co-aligned laser and telescope oriented at 28° off-zenith, in the region of overlap between the central two wide-field cameras. It had an average power output of 4.2 W, a pulse repetition rate of 100 Hz, an output wavelength of 532 nm, and a nominal power consumption of 300 W. A radiator and liquid cooling loop provided active thermal control. The telescope collected backscattered photons and recorded the arrival time relative to the most recent laser pulse, thereby measuring the backscatter coefficient as a function of altitude. Below the PMC layer, backscatter was due to air particles; measurement of the backscatter was used to reconstruct a vertical profile of temperatures. Within the PMC layer, backscatter was dominated by the PMC particles, so the lidar measured the altitude, depth, and vertical distribution of the PMC layer.

A balloon-borne lidar has many of the same advantages as balloon-borne cameras: the float altitude of the payload gives a reduction in sky brightness by two orders of magnitude and a factor of 2 reduction in viewing distance relative to ground-based lidars. The coincident nature of the lidar profiling and imaging data is another advan-

tage of a balloon-borne payload. The backscatter signal measured by a lidar decreases as distance squared (ignoring the exponential decrease in atmospheric density); thus, ground-based lidars used in PMC observations are oriented at small off-zenith angles in order to minimize the distance to the PMCs. As discussed above, ground-based imaging of PMCs occurs at large off-zenith angles. Coincident measurements must therefore be made by instruments at different locations, increasing the probability of unfavorable tropospheric weather conditions and increasing complexity. Co-location and coincident pointing of the PMC Turbo cameras and lidar provided coincident image and lidar data throughout flight, whenever PMCs were present.

3.3 Power System

The payload was powered by 15 solar panels from SunCat Solar coupled to Valence¹ U1-24RT Lithium Iron Phosphate rechargeable batteries through a Morningstar TS-MPPT-60A charge controller. The science instruments were divided into two separate circuits to reduce the possibility of single-point failures: five cameras were connected to a circuit with 6 solar panels, and the other two cameras and the lidar were connected to a separate circuit with 9 solar panels. I will refer to these two circuits as A and B, respectively. Each circuit contained two batteries with nominal voltage of 25.6 V and capacity of 20 Ah each, wired in parallel. A battery box for each circuit contained the batteries and charge controller, and acted as the interface between the solar arrays and the science instruments. The output voltage of the solar panels depended on illumination but was typically 15-20 V for a single panel. The panels were connected in strings of three (with each string connected in parallel to the other strings) to ensure sufficiently high voltage to charge the batteries. A schematic of Circuit B is shown in Figure 3.4.

¹Valence has since been purchased by LithiumWerks.

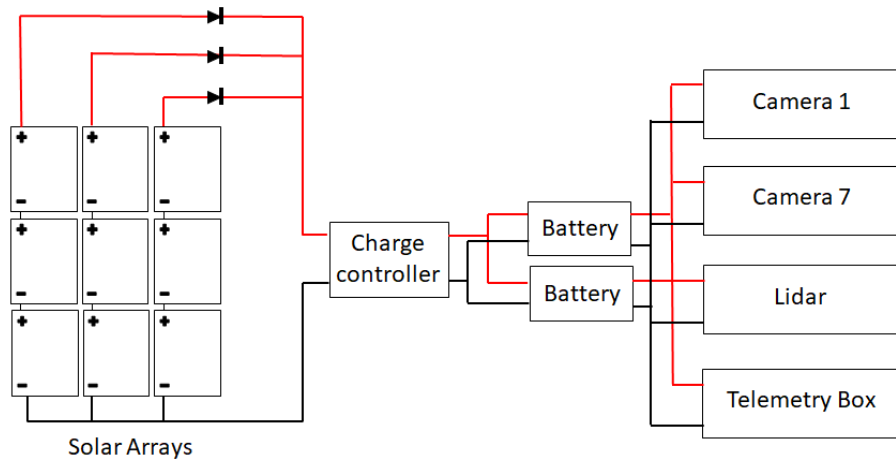


Figure 3.4: Schematic of the power system for Circuit B; red lines indicate high voltage (60 V between solar arrays and charge controller, 24 V elsewhere) and black lines indicate electrical ground. Diodes on each string of solar panels prevent backwards flow of current in the event of an un-illuminated or differentially illuminated solar array.

The two principal requirements of the power system were (1) to provide sufficient power at float, and (2) to provide sufficient power during ascent. These were distinct requirements, because the former was driven by the power output of the solar arrays while the latter was driven primarily by battery capacity.

3.3.1 Power Budget at Float

The solar cells were quoted to have an efficiency of 23% under standard test conditions (STC) – i.e. 1000 W m^{-2} insolation at normal incidence from an AM 1.5 Global solar spectrum at 25°C , with a 0.35% reduction in efficiency for each $^\circ\text{C}$ above that temperature (to clarify, this effect is multiplicative to the 23% efficiency, not additive: a temperature increase of 1°C would result in an efficiency of $99.65\% \times 23\% = 22.92\%$). To estimate the power supplied in flight, we need to account for several factors: the increased temperature due to constant illumination, the increased insolation due to reduced air mass, the additional insolation due to albedo, and the non-normal

incidence of the illumination. We quantify these effects as follows:

- Prior balloon flights suggest, and thermal modelling for PMC Turbo confirm (see Section X.X), a typical solar array temperature of 100°C, which give a 26% decrease in efficiency.
- The total insolation at the top of the atmosphere is 1310 W m⁻² during the PMC Turbo launch and flight window [31]. Naively, we would expect the power supplied by the solar arrays to increase by 31% relative to the ground. But in fact, this extra radiation predominantly occurs at infrared frequencies that the solar cells are not sensitive to. Three spectra are needed here: the AM 1.5 G solar spectrum, the extra-terrestrial solar spectrum, and the solar cell efficiency spectrum. The first two are provided by the NREL [32], while the latter was provided by the manufacturer of the solar panels. All three spectra are shown in Figure 3.5. To a good approximation, the quantum efficiency of the solar cells acts as a boxcar filter between 400 and 1100 nm. By integrating over each of the two solar spectra over this band, I find that 75.7% of the AM 1.5 G spectrum falls in the useful range of wavelengths, while only 67.0% of the extraterrestrial spectrum does. The result is that the additional insolation at the top of the atmosphere provides a 16.8% increase in output power relative to STC.
- The incidence angle of sunlight on the solar arrays is dictated by the relative pointing in azimuth and in elevation:
 - The NASA rotator was quoted to have an accuracy of $\pm 2^\circ$ (1σ) in azimuthal pointing – i.e. the rotator was expected to keep the solar arrays pointed within 2° of the sun in azimuth. The in-flight performance of the rotator exceeded this expectation, with a 1σ variation of 0.8° . A 2° deviation from normal incidence results in a 0.1% decrease in incident power,

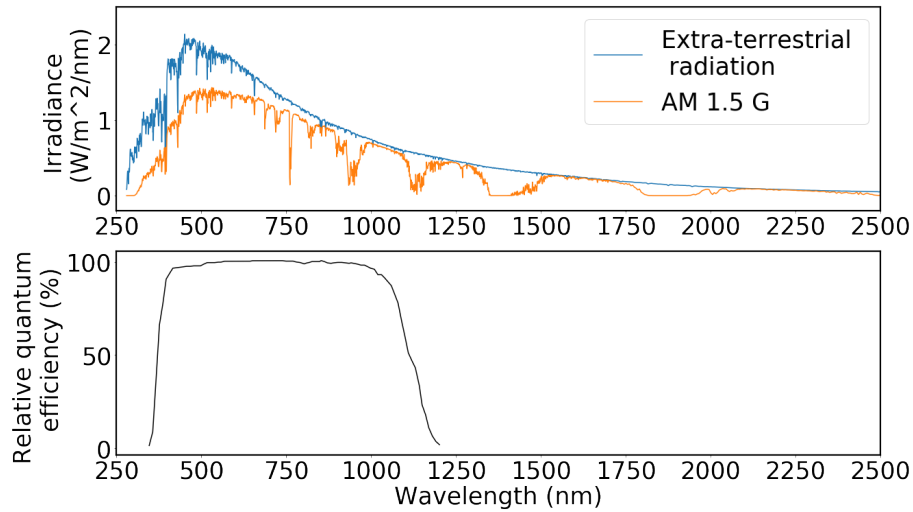


Figure 3.5: (Top) Irradiance spectra for extra-terrestrial radiation (blue) and for AM 1.5 G conditions (orange). (Bottom) Relative quantum efficiency of the solar cells as a function of incident wavelength.

so my pre-flight calculations of power budget neglected this deviation, and the in-flight performance validated that assumption.

- An analysis of previous flights from Esrange determined that our trajectory would have a lower bound of 66°N . Our launch window was dictated by the stability of the stratospheric jet stream and the presence of PMCs at our expected latitudes: we planned for a launch window between June 21 and July 21. With an expected flight duration of up to 6 days, the range of possible solar elevation angles was between -5° and 48° . Note that the horizon is at -6° at balloon float altitudes, so the sun was not expected to go below the horizon during flight. The solar arrays were tilted such that the normal to the arrays had an elevation angle of 23° . I assumed a maximum angle of incidence of 28° , which gave a 20% reduction in incident power relative to normal incidence.

Each panel was made up of 30 individual solar cells, with an area of 0.0153 m^2

Circuit A	
Cameras (5)	275 W
Telemetry box	10 W
Miscellaneous losses (diodes, etc)	40 W
Internal battery heating	15 W
Total:	340 W
Circuit B	
Cameras (25)	110 W
Lidar	450 W
Telemetry box	10 W
Miscellaneous losses (diodes, etc)	40 W
Internal battery heating	30 W
Total:	640 W

Table 3.1: Power consumption breakdown between PMC Turbo instruments on each circuit.

per cell. Therefore, we expected a minimum power output of 73 W per panel, giving 440 W for Circuit A and 660 W for Circuit B, compared to power demands of 340 and 640 W, respectively (see Table 3.1). This power budget excluded the effects of albedo, which contributes additional incident power on the solar panels and increases the power supplied to the instrument. Albedo is generally quoted as a single number averaged across the solar spectrum, with little information available on the spectral dependence of that number, so this effect was not quantified. Neglecting this effect, the power supply was sufficient to meet the expected demands during flight. The calculated margin for Circuit B is small, but since the estimate of solar power was conservative (using the extreme value of the off-normal solar incidence), the risk was determined to be acceptable.

3.3.2 Power Budget during Ascent

In an ideal world, the power budget for ascent would be determined well in advance of integration, and ample battery capacity would be provided. This was not the case for PMC Turbo; these calculations were not carefully done until integration was well underway and the flight hardware was fixed. I reproduce the considerations that affect the power budget during ascent and encourage future balloonists to think about this early and not skimp on batteries. What follows represents our state of knowledge before flight. The power supply and demands are not constant during the pre-launch and ascent periods, so I instead consider energy supply and demand.

The energy supply is made up of the battery capacity and the solar power that is generated during ascent. The nominal battery capacity is 0.5 kWh per battery, or 1.0 kWh per battery box, but we did not plan to use the nominal capacity for two reasons: (1) the batteries had a safety mechanism that prevented total discharge. Triggering this safety mechanism would prevent the instrument from being powered and would not be resolvable in flight; and (2) the battery capacity was temperature dependent, with smaller capacity at low temperatures. We planned to use no more than 60% of the nominal battery capacity, with plans in place to power down instruments if that limit was reached, to avoid catastrophic failure. The usable battery capacity per circuit was therefore 0.6 kWh.

Prior to launch, the instrument is powered by a diesel generator on the launch vehicle, but safety considerations dictate that the instrument be unplugged prior to the start of balloon inflation. This occurs at least 45-60 min prior to launch. (The preference of the launch crew is that the instrument be disconnected from ground power before the balloon is removed from its crate, 2 h prior to launch; this was the subject of ongoing negotiations, even until the day of launch itself. A less-anxiety inducing solution would be to provide battery capacity for a full 2 h pre-launch

period.) It is possible that the layout of the flight train and the direction of the sun will cooperate such that the solar arrays are fully illuminated prior to launch, but there is no guarantee – for example, it was often overcast in Esrange, which almost totally eliminated power output from the solar arrays.

The total time from launch to float altitude depends on atmospheric conditions, the weight of the payload, and real-time decisions regarding ballast made by the launch crew. Past flight profiles suggested a total time to float of 3-4 h. During the majority of the ascent, the air density is too high for the NASA rotator to provide stable pointing, so the payload will rotate freely – past flights suggested a typical period of several minutes. The rotator starts working at an altitude of 90,000 ft, roughly an hour before reaching float altitude. I conservatively budgeted 4 h for the total time between being disconnected from ground power and being stably pointed at 90,000 ft.

While the payload spins during ascent, the solar arrays provide power whenever they are illuminated, with power proportional to $\cos \theta$, where θ is the angle between the vector pointing to the sun and the solar array normal. We define $\sigma = \cos \theta$, and we let ϕ_A be the elevation of the solar array normal, and ϕ_S and θ_S be the elevation of the sun and the azimuth of the sun relative to the solar arrays, respectively, then we have

$$\sigma = \cos \phi_A \cos \phi_S \cos \theta_S + \sin \phi_A \sin \phi_S. \quad (3.1)$$

The average of σ over one revolution gives the factor by which the incident power on the arrays is reduced relative to constant normal incidence. The contribution from

the front side of the arrays, $\langle\sigma_f\rangle$, is given by

$$\langle\sigma_f\rangle = \frac{1}{2\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos\phi_A \cos\phi_S \cos\theta_S + \sin\phi_A \sin\phi_S d\theta_S, \quad (3.2)$$

where the integration is over only half of a revolution since the front side of the arrays is shaded for the other half of each rotation.

The backsides of the solar panels are somewhat transparent, so the solar arrays continue to produce power when illuminated from behind. We measured the relative efficiency, ϵ_{rel} , of the solar panels to be 0.3. In this case, we can still use (3.1), with $-\phi_A$ substituted for ϕ_A (since the solar array normal is tilted *down* when viewed from behind). This gives an additional contribution $\langle\sigma_b\rangle$, given by:

$$\langle\sigma_b\rangle = \frac{\epsilon_{rel}}{4\pi} \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \cos\phi_A \cos\phi_S \cos\theta_S - \sin\phi_A \sin\phi_S d\theta_S \quad (3.3)$$

where an additional factor of 1/2 has been added because the backside of the solar arrays are shaded by the payload for approximately one quarter of each revolution. The total efficiency factor of the panels due to rotation is the combination of these two factors, and can be written:

$$\langle\sigma\rangle = \frac{2 + \epsilon_{rel}}{2\pi} (\cos\phi_A \cos\phi_S) + \frac{2 - \epsilon_{rel}}{4} (\sin\phi_A \sin\phi_S) \quad (3.4)$$

As discussed above, $\phi_A = 23^\circ$, and ϕ_S was constrained only by season and geographic location, with an upper bound of 48° . Power concerns precluded a launch with the sun below the horizon, but this condition had minimal impact on the launch window: the sun set for the first time in Esrange on 15 July 2018, and by 21 July 2018, Esrange experienced less than 3 h of night. Given these constraints, $\langle\sigma\rangle$ can take values between 0.34 and 0.38.

The lowest temperatures encountered in flight are experienced during ascent, particularly at the tropopause, with temperatures as low as -50°C . Given the upward motion of the payload and the relatively dense atmosphere, convective cooling is significant. Previous flights suggested an average solar panel temperature of 0°C during ascent, giving an increase in efficiency of 9% relative to STC.

Taken together with the AM 1.5 G spectrum, these estimates suggested an average power of 40 W per panel for the 3 h period between launch and stable pointing. This translates to additional energy of 0.7 kWh and 1.1 kWh for Circuits A and B, respectively, giving a total energy supply of 1.3 and 1.7 kWh, respectively.

The low temperatures during ascent also placed more stringent demands on the energy supply. To avoid instruments from getting too cold, additional heaters were installed inside the battery boxes and on the radiator. Each battery box had 40 W of heater on a thermal relay, and the radiator had additional heaters that increased the total power draw of the lidar to 700 W. Estimates of the convective coupling of the radiator were poorly constrained, but we estimated that up to an hour of additional heating would be required for the radiator and batteries. To conserve power prior to launch, we planned to reduce the lidar power consumption to 190 W. The total expected energy demand was 1.3 and 2.4 kWh.

Thus, in the pessimistic but plausible scenario outlined above, we expected a significant shortage of energy during launch. For Circuit A, we could power down a single camera and stay within the estimated supply, but Circuit B posed a bigger challenge. We decided to power down the pressure vessels on Circuit B, and limit power consumption by the lidar as much as possible. We added an external, non-rechargeable Saft LiSO₂ battery that provided 100 W of additional heating to the radiator for up to 10 h (the battery was current-limited, so this power output could not be exceeded). Together, these brought the expected power demands down to

1.65 kWh, reducing but not eliminating the potential shortfall.

3.3.3 Power System Performance

When PMC Turbo launched on 8 July 2018, the layout of the balloon was such that the solar panels were fully illuminated before launch, and the batteries were still fully charged up until launch. Ascent was faster than expected: the total time from launch to float altitude was 3.0 h, and the rotator turned on 2.0 h after launch. We had a comfortable surplus of power throughout ascent; see Figure 3.6, which shows the battery state of charge throughout ascent. This was as much a result of luck as proper planning and I reiterate the need to do a proper accounting of the power demands prior to finalization of the experimental design.

The voltage and current of the solar array output were measured at the input to each battery box. When the batteries were fully charged, the output from the solar panels was limited to meet the demand of the instrument, with the rest of the energy being dissipated on the panels. Over the period during which the gondola was rotating and the batteries were less than fully charged, the average power output from the solar arrays was 60 W per panel, compared to an expected value of 30 W. Adjusting the above calculations During ascent, the sun had an average elevation angle of 26.5° , which increases our expected power output to 44 W per panel, still significantly less than the values measured in flight.

An examination of Figure 3.6 shows that our expectation of uniform, relatively rapid rotation of the gondola during ascent did not occur. Whereas the longest rotation during ascent for EBEX2013 was 6 minutes, the longest rotations for PMC Turbo were 6, 14, and 34 minutes. The gondola rotated slowly and spent a disproportionately large fraction of time (60%) with the solar arrays facing the sun, another fortunate accident. When the gondola did rotate, the power output per panel varied

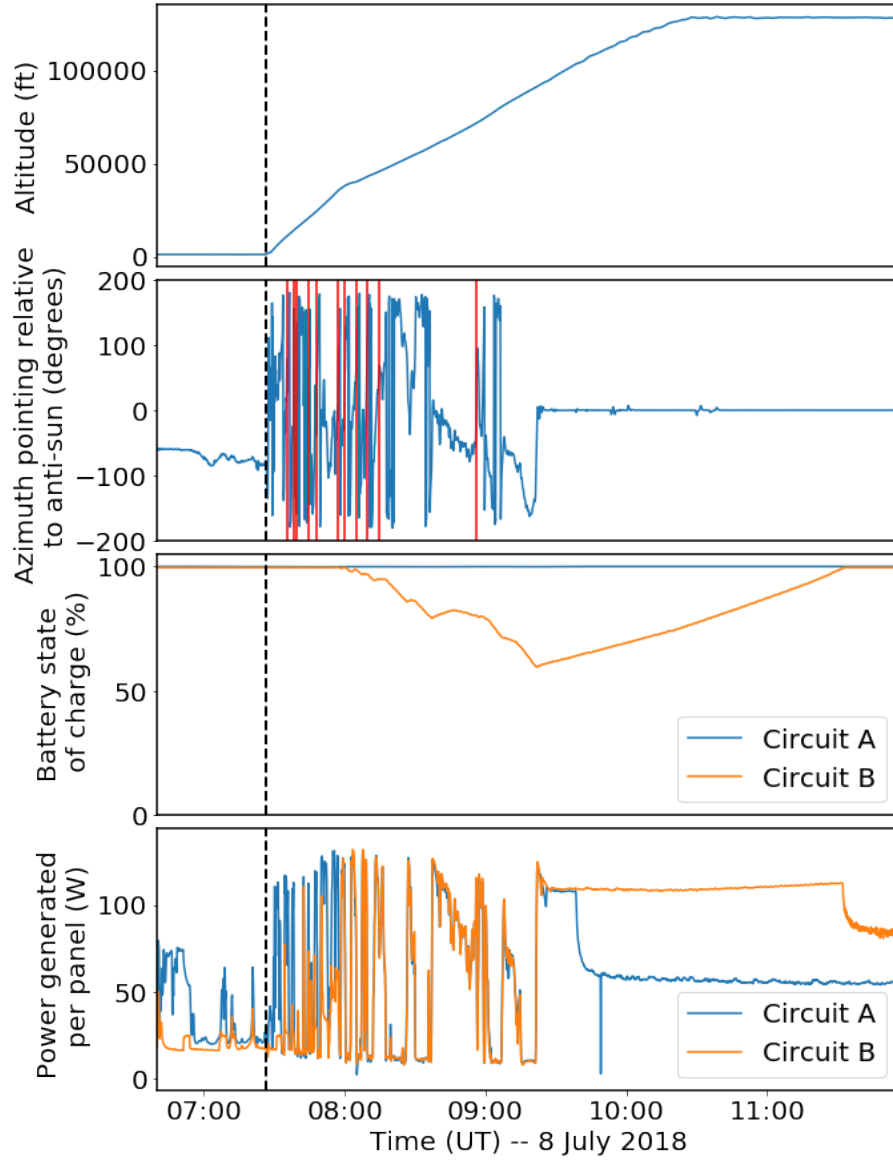


Figure 3.6: Altitude, azimuthal pointing relative to the sun, state of charge of the batteries, and power generated per panel during ascent and the first couple hours of float. The black dashed line in all plots shows demarks the time of launch. The red lines in the pointing plot show times that the gondola completes a full rotation.

between 15 and 90 W for the front and back sides of the arrays. The calculations shown above give an average expected value of 12 W when the backside of the arrays are illuminated and 94 W when the front side is illuminated. Thus the majority of the discrepancy between the expected and actual power output of the solar arrays was due to the duty cycle of the rotation rather than the performance of the panels. After reaching float altitude, the batteries quickly charged to full capacity and remained fully charged thereafter flight. As a lower limit, this means that the solar arrays produced sufficient power throughout flight. The only portion of flight that reveals the full capacity of the solar arrays is the brief period after reaching float as the batteries recharged. Given the elevation angle of the sun (which had risen to an average value of 34.5° over this period), the stable pointing of the gondola such that the solar panels faced the sun, and the measured solar panel temperatures of 60°C , the expected output of the solar arrays is 106 W per panel; we observed 110 W per panel.

3.4 Mechanical, Thermal and Optical Design

In addition to its numerous advantages, a balloon-borne payload also comes with significant challenges that are relevant to ground-based experiments. Of these challenges, I focus here on the mechanical and thermal design; I also discuss the baffles that were used to reduce stray light in the images from flight, the shortcomings of those baffles and the re-design of the baffles that I produced after flight.

3.4.1 Mechanical Design

In the following, I describe the design drivers and decisions for the PMC Turbo gondola. I took the lead responsibility for the following design, initially under the super-

vision of Michele Limon. Frank Kelso also provided invaluable guidance throughout the design process. Without both of their contributions, I have little confidence that this process would have converged.

Design Drivers

The basic design considerations for the PMC Turbo gondola were: (1) to provide a rigid, balanced structure to mount the instruments described above; (2) to provide the cameras and lidar with an unobstructed view of the sky; (3) to keep aforementioned instruments securely attached to the gondola throughout flight; (4) to provide protection to the instruments on landing; (5) to provide mounting points that place the antennae as the topmost element of the payload; and (6) to do all of the above using standard structural members whenever possible. Of these, (3) is the most important: instruments falling from the payload during flight represent a safety risk, and NASA will not approve a payload for launch until the mechanical design has been vetted against that possibility. There are also constraints on the size and weight of the payload, but PMC Turbo was well below these limits, so these constraints did not drive the design.

The CSBF guidelines [33] for the structural design of balloon gondolas has been updated since the PMC Turbo flight; the guidelines presented here are similar but not identical to the currently applicable set. The largest stresses during flight are encountered at termination, when the connection between the payload and the balloon is severed. For several seconds, the payload falls freely towards the Earth, until the air becomes sufficiently dense to inflate the parachute. By the time the parachute opens, the payload can be falling at close to 50 m/s, and the parachute causes a swift deceleration. The NASA guidelines dictated that the gondola design be able to withstand:

- a 10g load applied vertically
- a 5g load applied at 45° from the vertical in each of four orthogonal directions
- a 5g load applied horizontally

For the first two requirements, the load was applied at the interface between the suspension system and the NASA rotator. The third requirement applied only to the instruments attached to the gondola, and not to the suspension system.

The primary driver for these design requirements is safety rather than preservation of the instrument. The parachute slows the payload as it falls, but it still hits the ground with appreciable speed (10-15 m/s). Cardboard crush pads are added at the bottom of the payload in the hopes of softening the landing, but there is no guarantee of their effectiveness (if the payload lands with sideways velocity, the crush pads can be stripped from the bottom of the payload before they have a chance to cushion the impact – CSBF recovery personnel have commented on payloads they’ve recovered that the only things to survive impact were the crush pads). In addition to the initial impact, the NASA rotator (a.k.a. the “wrecking ball”) weighs 200 lb and is centered above the rest of the payload. In addition to meeting the flight requirements given to us by NASA, we wanted to minimize the risk of damage to our flight hardware at landing.

With few exceptions, the gondola was constructed from 6061-T6 aluminum, because it is relatively cheap, widely available, and has good strength-to-weight ratio.

Gondola Cage

The design of the PMC Turbo gondola consisted of the “cage” – a robust aluminum structure that housed the PMC-Turbo instruments – and the suspension system, which interfaced between the cage and the NASA rotator. The cage consisted of

three levels made from C-channels connected at the corners by vertical L-brackets. Note that in this design, the bolts between the vertical and horizontal members are oriented horizontally; they provide a clamping force between the members such that any vertical loading is opposed by static friction at the joint rather than the shear strength of the bolt itself. This is an important aspect of mechanical design that I failed to appreciate in initial iterations. Due to time constraints, the design of the cage was finalized and fabrication began before the design of the suspension system was completed. This created fixed geometrical constraints and added to the challenge of designing the suspension system, the difficulty of which I did not appreciate at the time.

The size of the cage was determined by the instruments that it contained. The heaviest elements of the payload were the SIP and the lidar, so these were centered in the gondola cage with the other instruments arranged around them. Three I-beams were placed along the short axis of the gondola at the middle level of the cage; the SIP hung below these I-beams, and the lidar mounted above them. The cameras were positioned at the top level of the cage and were arranged symmetrically, with a bias towards the front of the gondola to prevent obstructions to their FOV. With the exception of the lidar radiator and solar arrays, all other instruments were contained within the cage to provide protection at landing. The size of the radiator and solar arrays and the uncertainty in the orientation of the payload relative to its motion at landing precluded any effort to protect those components.

Suspension System

The suspension systems for balloon payloads typically fall into one of two categories: rigid or flexible. Rigid suspension systems provide protection from the rotator and other elements of the flight train at landing, and protect the payload if the payload




Figure 3.7: Suspension system of the PMC Turbo gondola, showing the 4 rigid members and the 4 tensioning cables. The purple arrows show the attachment point to the flight line, where the loads described above are applied.

rolls over. The challenge of designing a rigid suspension system is that significant bending moments will occur at the attachment points, especially in the “5g at 45° from vertical” load cases and as a result rigid suspension systems are heavier than their flexible counterparts. Flexible suspension systems, on the other hand, provide no protection at landing. Furthermore, at termination, in the “5g at 45° from vertical” load case, flexible suspension members (braided steel cabled, generally) will translate significantly from their nominal configuration at float; this can damage other flight components, especially given a compact layout such as that of PMC Turbo. Finally, given the requirement that the antennae be positioned above all other gondola components, a flexible suspension system would require a separate, rigid structure to mount the antennae sufficiently high. This last point is due to practical considerations of the pre-launch procedure: it is entirely conceivable that a payload be designed along the same lines as PMC Turbo, with antennae in the same position but a purely flexible

suspension system supporting the cage below. On the launch vehicle and in flight, this configuration is feasible. But prior to launch, the payload rests on the ground, and a flexible suspension system cannot support the antenna boom. Therefore, the suspension system would have to be remounted prior to each launch attempt, and cables between the antennae and SIP would need to be re-attached and re-tested each time.

Due to these considerations, we opted for a hybrid approach, shown in Figure 3.7. The NASA rotator attached to a spreader bar that also carried the antenna boom (not shown; see Figure 3.1). The principal suspension members were 4 steel rods, with spherical ball joints at either end to eliminate bending moments. These rigid members mounted at either end of the spreader bar and provided the necessary strength and stiffness in the vertical load case and in the 5g at 45° load case when that load was applied toward the front or the back of the gondola. But due to the spherical ball joints, these rigid members provided no stiffness in the lateral (left-right) direction; to add lateral stiffness, we added steel cables that ran to each corner of the cage. These cables attached via hoist rings, and turnbuckles were used to tension the cables and precisely adjust of the cable length. The position of the spreader bar was therefore fully constrained relative to the gondola cage.

To adequately size the suspension members and attachment points, I calculated the expected loads on each suspension member in each load case analytically (see Appendix A). For all such calculations, I used an estimated total suspended weight of 3500 lb, which included a 200 lb margin over our estimated total weight. This is a sensible precaution: pre-integration weight estimates inevitably fail to account for everything, and the final weight is likely to be higher than expected. In our case, the final weight of PMC Turbo was within 50 lb of my design criterion.

Stress Analysis

I performed stress analysis of the PMC Turbo gondola iteratively as the design progressed to identify areas that needed reinforcement. I used the Finite Element Analysis (FEA) module of SolidWorks 2015 to conduct this stress analysis, verified by hand calculations when possible. Appendix A describes this stress analysis in detail; it is an updated version of the stress report I submitted to NASA in order to certify the PMC Turbo gondola for flight. The analysis described here includes all components up to and including the adapter fitting that interfaces between the gondola and the NASA rotator. All components above that adapter fitting have been designed and validated by NASA engineers for payloads much heavier than ours.

Since the gondola design was performed in SolidWorks, using the integrated FEA module was a streamlined and efficient process. In principal, one could simply apply the appropriate boundary conditions and load cases to the entire gondola and calculate the resulting stresses everywhere. In practice, this is too computationally expensive to do. Furthermore, I was unable to properly simulate the degrees of freedom present in the suspension system, which resulted in unrealistically large bending moments at the attachment points. Therefore, I analyzed different components of the gondola separately, as follows:

- the suspension system; loads in each suspension member were calculated by hand. All OTC components (hoist rings, spherical ball joints, clevis brackets and pins, etc.) were omitted in the analysis; the manufacturers provided the relevant load limits.
- The rigid members of the suspension system were analyzed separately for expected tensile and compressive loads.
- the gondola cage; simplified mass models of all subsystems (cameras, lidar,

radiator, solar arrays, battery boxes, etc) were used to make the analysis less computationally intensive. Again, OTC components of the suspension hardware were omitted in the analysis.

- individual subsystems, analyzed in full detail. These include:
 - camera mounting brackets and attachment points
 - lidar mounting brackets and supports
 - solar arrays
 - radiator
 - battery box mounts
 - ballast attachment points

In each case, the selection of appropriate boundary conditions was a crucial step in performing the stress analysis. Since the payload is under free-fall in each of the load cases, there are no fixed geometry constraints that can be realistically applied. For example, consider the analysis of the gondola cage. As discussed above, the gondola cage was analyzed separately from the suspension system, with resulting loads in each of the suspension members calculated manually. A conservative approach would be to apply fixed geometry constraints to the base of the gondola cage while applying the calculated loads at the attachment points of each suspension member. This results in large apparent stresses at the base of the gondola, as the entirety of the upward forces must be transmitted through the structure to the fixed constraints. In reality, the base of the gondola cage sees almost no stresses, as each level of the gondola structure is only responsible for accelerating those components beneath it. More accurate results can be obtained by applying "inertial relief" which is intended for exactly this scenario: a model with unbalanced external loads. By definition, the

simulation is dynamic – a net external force will result in an overall acceleration – but the inertial relief option determines the equivalent static simulation by applying inertial loads based on the masses of individual finite elements.

For subsystems, it was straightforward to apply fixed geometry constraints at some distance from the subsystem of interest, then apply a 10g or 5g load to all components. Unrealistic stress concentrations would result in the vicinity of the fixed constraint, but these concentrations were outside the region of interest.

Examples of areas where the final design was informed by stress analysis include: (1) the vertical T-bars that provided vertical support to the upper level of the gondola cage where the cameras were mounted; (2) the steel or aluminum plates that reinforced the attachment points of the various suspension members; and (3) the rigid members of the suspension system. In previous versions of the stress analysis, I had failed to consider the possibility of buckling of these rigid members when they were placed under compressive loads. A rigid column has a critical compressive load at which it becomes elastically unstable; a compressive load above this limit will cause buckling. This placed more stringent requirements on the stiffness of these rigid members than a naïve stress analysis that considers perfectly axial loading. This highlights the necessity of considering all possible failure modes when validating a mechanical design.

3.4.2 Thermal Design

Each of the instruments flown on PMC Turbo had an expected range of operating temperatures, beyond which they would cease to function as intended. The purpose of the thermal modelling described here was to gain confidence that each of the instruments would remain in their respective operating temperature ranges. Much like the stress analysis described in the previous section, the thermal modelling was done iteratively, as the thermal design evolved in response to the results of the thermal

analysis. The output from the thermal model is the range of temperatures that are likely to be reached by each instrument. Testing was then performed at the thermal-vacuum chamber during integration in Palestine, TX to ensure proper operation of the instruments over their possible temperature ranges.

Much like the power system (Section 3.1.3), the thermal design of a balloon payload has two primary drivers: maintaining appropriate temperatures during ascent, and maintaining appropriate temperatures at float. These are somewhat opposing: during ascent, the payload passes through the tropopause, with ambient temperatures of -50°C or lower and significant convective cooling due to the upward motion of the gondola; whereas at float, the payload is subjected to constant solar illumination with relatively little convective cooling to prevent overheating. The convective cooling during ascent is difficult to quantify, but, with the exception of the lidar, our instruments were closely modelled after previous flight hardware (especially EBEX2013), so we had high confidence that they would survive ascent. Therefore, the main driver for the thermal design was the thermal environment at float.

The thermal environment at float is dominated by radiation and power dissipation. The radiative environment consists of contributions from the ever-present sun, albedo from the surface of the Earth, and infrared thermal emission from the Earth. The payload, in turn, re-radiates thermal energy outward in all directions. Temperatures are therefore dictated by the optical properties of the gondola, quantified by the solar absorptivity (α) and the infrared emissivity (ϵ). To illustrate the importance of these two parameters, consider an aluminum sheet floating in space at 1 AU (i.e. the average distance from the Earth to the sun), oriented perpendicularly to incident solar radiation. If it is unpainted and polished, it has $\alpha=0.14$ and $\epsilon=0.03$; if it is painted with GSFC White Paint MS74, it has $\alpha=0.17$ and $\epsilon=0.91$. In the first case, the steady-state temperature of the sheet is 214°C ; in the second, it is -56°C , a

temperature swing of 270°C. Thus, the surface coating of the various components is a principal control we have over their temperatures in flight.

The thermal analysis was performed in Thermal Desktop, an AutoCAD-based geometric thermal modeler. Before performing thermal analysis of the instrument, I verified that the various features of Thermal Desktop worked as expected. I simulated a single sheet floating in space, as described above, and verified that the Thermal Desktop results agreed with hand-calculated temperatures, with appropriate dependence on α and ϵ . I simulated many parallel sheets, each one heated by re-radiation from neighboring sheets. I converted the sheet into a 3-D object and verified that the thermal gradient across the object was consistent with the input thermal conductivity. I added heat loads to simulate heat dissipation by the instrument, and included Earth IR flux and albedo from the surface of the Earth. In particular, Thermal Desktop allowed for a temperature-dependent heat load, which is an important characteristic of the Valence batteries. In all cases, the results from Thermal Desktop agreed with hand calculations.

Having gained confidence in the functionality of Thermal Desktop, I reproduced a simplified version of the PMC Turbo gondola in AutoCAD and input the relevant surface and conductive properties. For the camera pressure vessels, no detail was included as to the internal components; instead, 55W of power was dissipated over the surface of the pressure vessel. Given that the vessels remain pressurized and have internal fans to circulate the air, the temperature of the pressure vessel wall determines the internal temperatures. Laboratory testing had indicated that the motherboard would reach a temperature 15°C hotter than the pressure vessel wall, and that this temperature differential held for a range of wall temperatures – thus the problem was reduced to finding the expected range of pressure vessel wall temperatures.

The lidar had an active cooling system that transferred heat to the radiator;

design of this system involved dedicated thermal analysis by the team at DLR. In particular, the liquid cooling loop required functionality that was not available in Thermal Desktop. Their analysis indicated that: (1) the thermal link between the lidar and gondola was such that no more than 80 W would transfer between gondola and lidar, in either direction, (2) the average temperature of the radiator was 4-11.5°C colder than the radiator outlet temperature, and (3) the radiator outlet temperature must remain between 16 and 20°C in order for the lidar to operate, and above 5°C for the lidar to survive. Additionally, the steady-state power consumption was measured to be 298 W. This power consumption occurs within the lidar pressure vessel, but the waste heat is transferred to the radiator. A more accurate accounting would have deducted the 4.2 W laser output power but such deduction was omitted from my analysis. The open questions to be solved by Thermal Desktop were: in the hot-case, what is the maximum expected temperature of the radiator, and in the cold-case, what heat load is required to keep the radiator in the appropriate temperature range.

The solar arrays were modelled after the description in Soler et al. [34] who used Thermal Desktop to model the BLASTPol instrument. The solar arrays are an important aspect of the thermal design, due to their large surface area and the requirement that they be placed in direct solar illumination. A preliminary, more compact design of PMC Turbo had the solar arrays directly behind the instrument, providing shading to the instrument from the sun. Thermal modelling indicated that such a configuration would result in pressure vessel temperatures of 70°C, well in excess of their maximum operating temperature. This prompted the final design, in which the solar arrays are spaced out to the sides of the gondola and the space between was filled with aluminized mylar to provide shading to the instrument.

My approach was to find the steady-state hot- and cold-case temperatures for the various components of the instrument, assuming that the instrument came to thermal

	Albedo (dimensionless)	Earth IR flux (W/m²)
Minimum	0.30	148
Mean	0.53	177
Maximum	0.73	215

Table 3.2: Albedo and Earth IR flux values based on flight bounded by June 21 and July 21 in time, 66°N and 77°N in latitude and 130°W and 25°E in longitude.

equilibrium given the conditions in those respective scenarios. The advantage of this approach is that it narrows the parameter space of simulations and obviates the need for information on the heat capacity of the instrument; the disadvantage is that the result is necessarily conservative: the conditions of the hot-case scenario, for example, do not necessarily persist long enough for the instrument to come to thermal equilibrium, and therefore the steady-state temperatures will be warmer than any encountered in flight. Given the wide operating temperature ranges of the PMC Turbo instruments, this conservative approach did not cause a significant obstacle.

The parameters for the hot- and cold-case scenarios were provided by Doug Ferguson at CSBF and were based on the launch window and range of possible trajectories. Table 3.2 shows the range of expected albedo and Earth IR flux values assuming a flight between June 21 and July 21 with trajectory bounded by 66°N and 76°N in latitude and 130°W and 25°E in longitude. Albedo and Earth IR flux are roughly anti-correlated (see Figure X); therefore, it is overly conservative to assume extreme values of both of these parameters will occur simultaneously. Instead, I ran two simulations each for the hot- and cold-case: for the hot-case, I used the mean albedo and the maximum expected Earth IR flux, then I used the maximum albedo and the mean Earth IR flux. Similarly, for the cold-case, I used the mean albedo and the minimum Earth IR flux, then used the minimum albedo and the mean Earth IR flux. In both cases, I found that the instrument had greater sensitivity to the IR flux than to the albedo. Representative examples of cold-case and hot-case model outputs are shown




Figure 3.8: Pre-flight predictions for cold-case (left) and hot-case (right) temperatures for the PMC Turbo gondola, assuming a stable thermal environment at float.

in Figure 3.8. Note that there is a significant temperature gradient from the front (active) side of the solar panels to the back; the predicted solar panel temperatures were based on the front side of the panels, whereas only the backside of the arrays are visible in Figure 3.8.

The final variable dictating the thermal environment was the elevation angle of the sun. Given flight dates between June 21 and July 21 and an expected southern-most latitude of 66°N , we anticipated solar elevation angles between -3.5° and 47.5° , though in principal the hot- and cold-case scenarios can occur at any elevation angle in that range. Note that at float altitudes, the horizon is roughly at -6° , so over the expected range of flight dates and trajectories there was no danger of the sun going beneath the horizon (which would have resulted in much larger temperature swings). The aluminized mylar sun shades were designed such that all instruments were shaded over the expected range of solar elevation angles, assuming that the NASA rotator maintained stable pointing to within $\pm 5^\circ$ in azimuth. I ran each hot- and cold-case scenario at increments of about 5° across the expected range of

	Minimum allowable/ expected/measured temperature ($^{\circ}\text{C}$)	Maximum allowable/ expected/measured temperature ($^{\circ}\text{C}$)
Camera pressure vessels	-40/-15/-10	55/37/30
Solar panels	-/70/47	110/99/87
Batteries	0/19/9	50/37/21
Telemetry box	-70/-10/-	110/40/-

Table 3.3: Allowable, expected and measured temperature ranges for PMC Turbo components during flight. The measured temperatures are limited to temperatures encountered at float rather than during ascent.

solar elevation angles; the hottest and coldest resulting temperature across that set of simulations determined the upper and lower limits, respectively, for the expected range of temperatures for each instrument in flight (excluding during ascent). The allowable, expected and measured temperature ranges are given in Table 3.3.

Evident in Table 3.3, and shown more clearly in Figure 3.9, is a bias in the model prediction of the temperatures of the solar panels. Furthermore, whereas the thermal model predicted maximum temperature swings of 30°C , the measured swings in temperature were up to 40°C . Three possible factors explain these effects:

- The predicted temperatures constitute an average over the solar array, whereas the measurement was made on the CSBF solar panels, which were the farthest panels from the gondola. Figure 3.8 shows a thermal gradient of $\sim 10^{\circ}\text{C}$ across the solar arrays, which explains part of this effect.
- The thermal model assumes no convective coupling between the gondola and the ambient environment. Any such coupling would act to cool the gondola.

Given the large surface area, low thermal mass, and position of the solar areas at the exterior of the payload, the solar panels would be particularly affected by convective cooling.

- An error in the thermal model of the solar arrays. This is of course the least imaginative explanation, but is certainly plausible. The model of the solar arrays is based on the description in Soler et al. [34]; the measured solar array temperatures reported in that paper demonstrate a similar overall bias towards lower temperatures to what we observe here.

Given that there is no minimum allowable temperature for the solar arrays, the bias towards lower temperatures is not of immediate concern. However, the solar arrays are thermally coupled to the rest of the gondola, so this bias towards warmer model solar array temperatures can result in warmer predicted temperatures for the other instruments as well.

The batteries also experienced lower temperatures in flight than the predicted cold-case scenario. This is probably due to an underestimation of the conductive path between the batteries and the outer walls of battery box. Our primary concern prior to flight was that the batteries might overheat; this concern was reflected in the thermal model, which was aimed at conservatively estimating the hot-case scenario. An overall bias towards colder temperatures is therefore unsurprising.

The cameras remained comfortably within the predicted and allowable temperature ranges. Furthermore, the cameras near the rear of the gondola reached temperatures $\sim 20^{\circ}\text{C}$ higher than those near the front, which agrees with the model predictions. Furthermore, each camera had temperature swings of $\sim 20^{\circ}\text{C}$, also in accordance with the model predictions.

Each of the instruments was tested over the expected temperature range at a pressure of 5 mbar during integration in Palestine, TX, with the exception of the

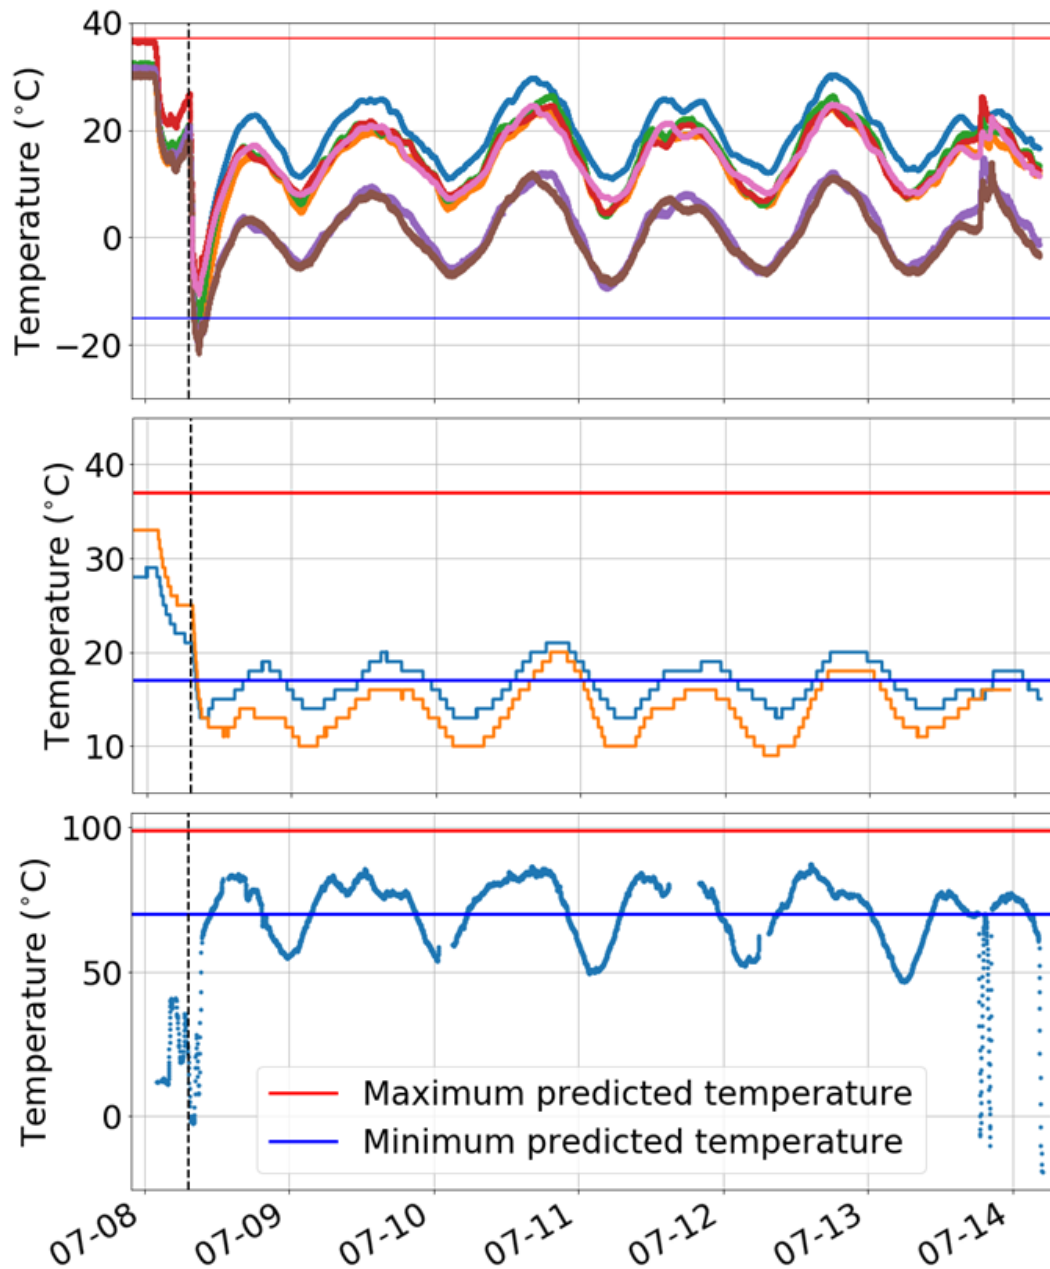


Figure 3.9: In flight temperatures of the cameras (top, each camera marked with a distinct color), batteries (middle, with batteries from each battery box shown), and solar panels (bottom). Launch is marked with a dashed black line in each plot, and the red (blue) lines indicate the maximum (minimum) predicted temperature for each component. Note that predictions were made for the thermal environment at float; no attempt was made to model temperatures during ascent.

solar panels. SunCat solar panels have flown on many previous balloon missions and were therefore deemed safe for flight without the need for thermal-vacuum testing of individual panels. One of the pressure vessels showed a minor malfunction at low temperatures: below -15°C , the focus step of the lens would slip and would need to be re-focused. Other than that, the instruments performed as expected across the full temperature range.

3.4.3 Optical Design

The design of the PMC Turbo pressure vessels was modeled after the EBEX2013 star cameras and is described in detail in Kjellstrand et al. [35]. Here I describe the design of the camera mounts that dictate the alignment of the cameras, the process by which alignment was determined, and the design of the optical baffles.

Camera Mounts

The alignment of the cameras is dictated by the science goals of the experiment: to image GW breaking processes from scales of 100 km down to the resolution of the cameras. The orientations of the wide-FOV cameras were optimized to provide the maximum contiguous sky coverage. Regions at small off-zenith angles give improved spatial resolution due to the decreased viewing distance and are therefore prioritized. The narrow-FOV cameras were pointed such that they covered the areas of overlap between the wide-FOV cameras at small off-zenith angles to make maximal use of their increased spatial resolution. The limit at small off-zenith angles is imposed by the presence of the balloon, which subtends an opening angle of 23.8° from the perspective of the payload. I included a 2° margin to allow for fabrication errors in the camera mounts, misalignment between the optical axis of the cameras and the symmetry axis of the pressure vessels, and horizontal displacement between the

payload and balloon. Figure 3.2 shows the alignment of the cameras as measured in flight, projected onto the PMC layer at 83 km.

The design of the camera mounts included two degrees of freedom: the pressure vessels can rotate freely about their symmetry axis, and the mounts can pivot $\pm 5^\circ$ in “cross-elevation” – i.e. the direction perpendicular to the off-zenith direction. Note that this is similar but not identical to the azimuthal direction. The camera mounts were measured using a coordinate measuring machine while mounted to the gondola and were determined to be within 0.4° of the nominal angle in every case, in elevation and in azimuth. We taped strips of paper with lines spaced at $0.07''$ intervals to the outer wall of the pressure vessels – given the $8''$ OD, this corresponds to 1° increments. This allowed for repeatable alignment of the “roll” of the pressure vessels and for specific incremental adjustments.

Optical Alignment

The alignment of the pressure vessels was measured by taking pictures of the night sky and solving for the resulting star fields – this was, after all, the primary function of the EBEX2013 star cameras. This was performed in Palestine, TX during the initial integration of the payload, and again in Esrangle, Sweden prior to flight. Due to the summertime launch and northern latitude of Esrangle, stars were not visible from the ground in the weeks prior to launch. The need to measure the alignment again after re-integrating the payload dictated the early start of the Esrangle campaign. We arrived on May 1, 2018 and quickly re-assembled the payload. We confirmed that the alignment was close to nominal on the night of May 7. Already, civil twilight (which requires the sun to be more than 6° below the horizon) had ended for the season, but it was still barely possible to find star solutions for images with sufficiently short exposure times not to cause star tracks. The star solutions were found using

astrometry.net [36].

Optical Baffling Design (and Redesign)

In order to minimize the amount of stray light being scattered into the optical path, optical baffles were placed at the front end of the camera pressure vessels. These were coarsely modeled after those used with the star cameras on EBEX2013, as described in Chapman [37]. Prior to flight, we had unfounded and untested confidence that the cameras would demonstrate strong rejection of stray light, and as a result no great effort was put into the design of the optical baffles. This resulted in inadequate baffling that led to significant degradation of the science data in flight. Therefore, this section is written with the benefit of hindsight and describes the considerations that should have been made prior to flight rather than the considerations that were actually accounted for. The shortcomings of the baffles that were flown serve to illustrate the importance of these considerations.

Figure 3.10 shows three sets of images made by differencing pairs of images taken at 20 minute intervals for each of the four wide-field cameras. The first two sets of images were taken at times with little PMC activity to isolate the contributions from stray light, and the third set was taken at a time of particularly intense PMC activity; all three sets have the same color scale applied. These images demonstrate a range of characteristics: (1) the magnitude of the stray light contribution is large, comparable to or larger than the PMC contribution, (2) the stray light contribution is neither uniform between cameras nor constant in time, and (3) cameras 4 and 7 (the leftmost and rightmost cameras) are particularly badly afflicted.

The goal of the baffles is to minimize the amount of spurious light that enters the optical path. The first-order design of the cameras and mounts should be such that no bright objects are within the FOV of the cameras. The pressure windows

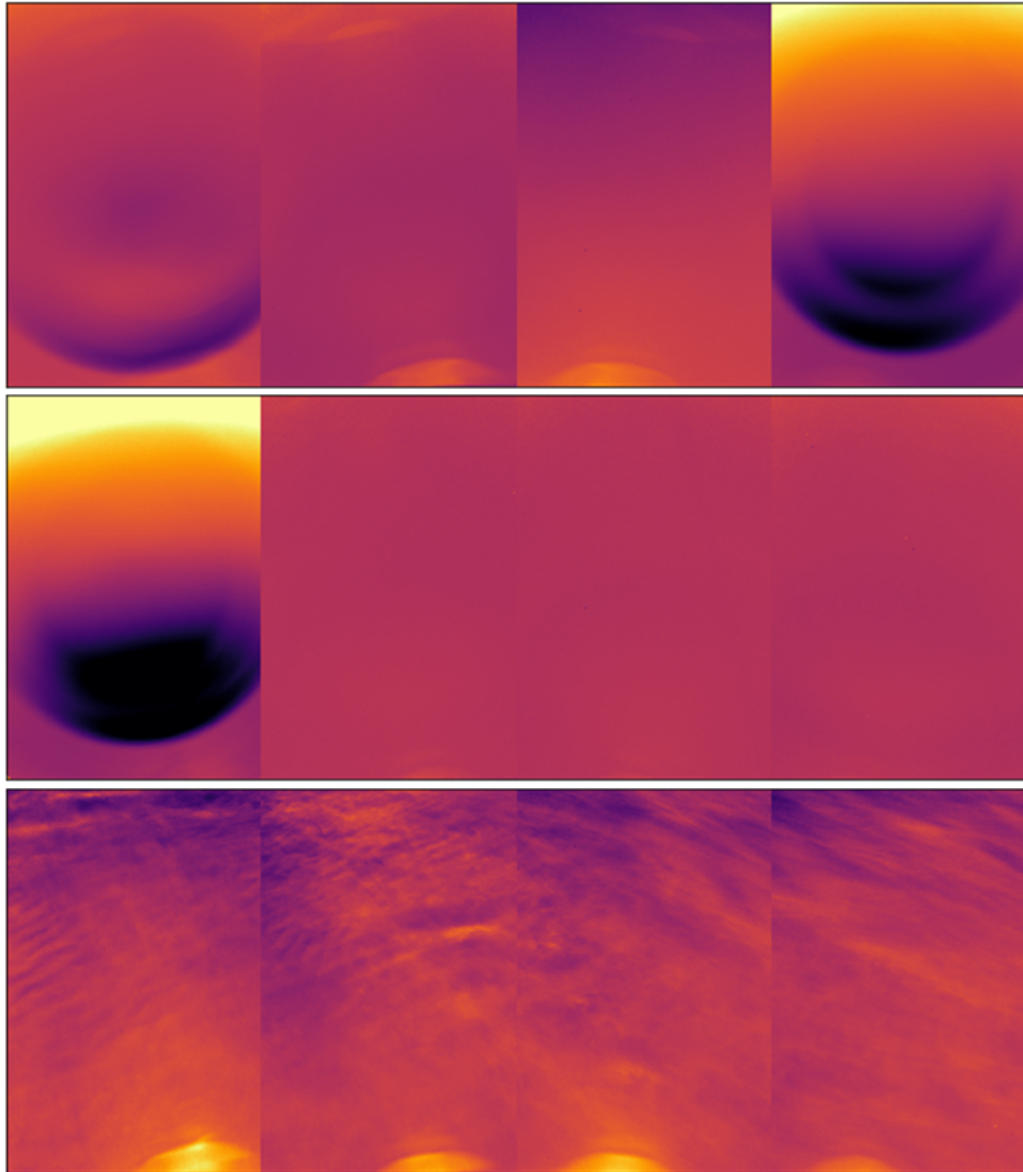


Figure 3.10: Stacked difference images from the four wide-field cameras at three different times during flight (top, middle, and bottom). Each image is found by taking the difference between a pair of images spaced at 20 minutes. The first two sets of images were taken during a time with little PMC activity to demonstrate the stray light artifacts in isolation; the third was taken at a time of particularly intense PMC activity to demonstrate the relative magnitude of the scattered light effect. All images are shown with the same color scale.

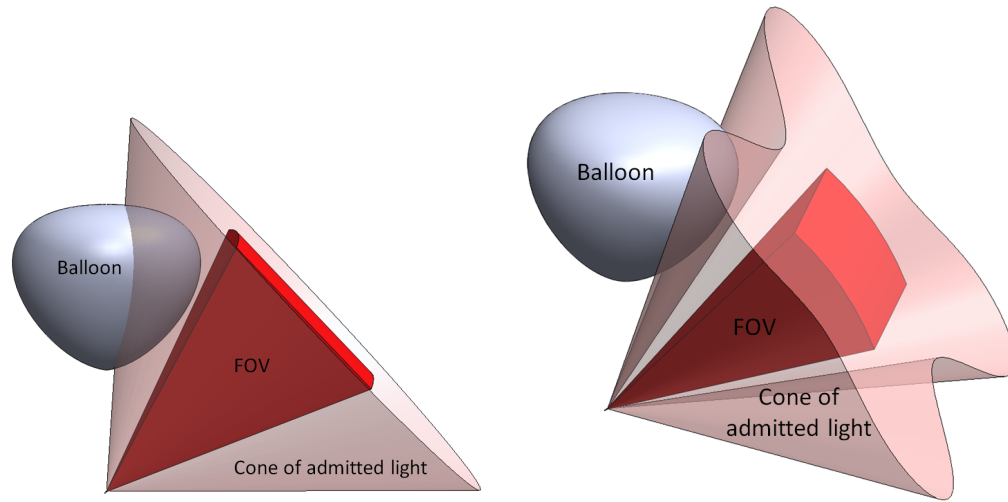


Figure 3.11: (Left) Wide-field camera FOV and cone of admitted light for the 2018 flight of PMC Turbo. (Right) Same as on left, but for the updated optical baffle design.

are imperfectly transmissive and therefore scatter incident light into the optical path; therefore, the design of the baffles should be such that no bright objects are permitted to shine onto these windows. The EBEX2013 design included internal vanes to meet an even stricter demand: in order for external light sources to reach the window, multiple reflections within the baffle were required; see Chapman [37].

A key distinction that must be made is between the FOV and the “cone of admitted light”, i.e. the region of space that can directly enter the pressure vessel window. For the PMC Turbo wide-FOV cameras, I determined the longest cylinder (with 10” diameter to match the OD of the camera pressure vessels) that would not encroach on the FOV. The FOV of the PMC Turbo cameras was much wider than that of the EBEX2013 star cameras, and the baffles were accordingly shorter: 7” for the wide-FOV cameras. This gives a 1.7° margin at the corners of the FOV, but the cone of admitted light has an opening angle of 43° . Given the small margin between the edge of the FOV and the balloon in flight, the balloon is well within the cone of admitted light, as shown in the left panel of Figure 3.11.




Figure 3.12: PMC Turbo camera pressure vessel with updated optical baffle design.

After the PMC Turbo flight from Sweden, one of the wide-FOV cameras flew as a piggyback on the SuperTiger-2 payload in December of 2019, launching from McMurdo Station in Antarctica. This presented an opportunity to re-design the optical baffle and camera mount to avoid the mistakes that were made in the original design. Specifically, we adjusted the elevation angle of the camera mount to leave a larger margin between the FOV and the balloon, we rotated the camera so that the short axis of the FOV was aligned with the off-zenith direction (whereas the PMC Turbo wide-FOV cameras had the long axis of the FOVs roughly aligned with the off-zenith direction), and I adapted the baffle design to limit the cone of admitted light. Rather than the original axially-symmetric design, the adapted design included lobes on the four sides of the FOV to provide a constant margin between the FOV and the edge of the baffle. The result is shown in Figure 3.12 and unintentionally mimics the lens hood of a commercial DSLR: form follows function. An updated version of the FOV schematic with the cone of admitted light and the balloon is shown in the right panel of Figure 3.11 for the camera and baffle in their nominal configuration.

The upper edge of the FOV in this updated configuration was at 36.5° off-zenith, more than 10° lower than the upper edge of the FOVs during the flight from Sweden. This results in a 10% and 20% increase in the projected pixel size, in the azimuthal

and off-zenith directions, respectively. For a single camera flying as a piggyback on another payload, it is not feasible to include sufficient baffling to image up to such small off-zenith angles while still excluding the balloon from the cone of admitted light. For future PMC-observing missions, providing shading from the balloon should be a design driver – for example, a horizontal shade placed 2 m above the cameras could provide unobstructed view of the sky up to off-zenith angles of 27° while blocking any light from the balloon out to an angle of 26° . Alternatively, or additionally, future missions may be able to use a flightline extensions to increase the distance from the payload to the balloon, thereby reducing the apparent size of the balloon and expanding the area of available sky.




Figure 3.13: Difference images from the piggyback camera at four different times during flight, two when clouds were present (top) and two when no clouds were present (bottom). Each image is found by taking the difference between a pair of images spaced at 20 minutes. The same colorscale is used as in Figure 3.10.

Unfortunately, there was a mistake in the design or fabrication of the updated baffle that resulted in the baffle being rotated by 15° relative to the camera during the SuperTiger-2 flight. As a result, the top and bottom lobes of the baffle impinged upon two corners of the camera's FOV, oversaturating these regions and contributing spurious signal across the FOV. In addition to the mis-alignment of the baffle, the SuperTiger-2 payload offered less shading to the camera than the PMC Turbo payload, in which all cameras were completely shaded from all direct solar illumination. Between these two factors, the magnitude of the excess scattered light was higher for the piggyback flight than it had been for the 2018 flight.

The excess stray light observed during the 2018 flight was problematic because of its magnitude and because of the timescales on which it evolved - often 10 minutes or less. These are precisely the timescales of interest to our studies. Though larger in magnitude, the excess scattered light observed in the piggyback flight was more constant in time, evolving on a timescale of ~ 1 hour; as a result, it is more easily separable from the PMC signal we hoped to observe. See Figure 3.13, which shows difference images from the piggyback flight, with and without clouds present. No contribution from the excess scattered light is detectable.

Chapter 4

PMC Turbo 2018 Flight and Data

PMC Turbo launched on July 8, 2018 at 7:27 UT and flew for 5.9 days before landing in Nunavut, Canada. The trajectory and altitude track of the payload are shown in the top and middle plots of Figure 4.1. Performance of the instrument was nominal, aside from the excess brightness issue that has previously been commented. The Rayleigh lidar also functioned as expected, though it too experienced higher than expected background brightness, again thought to be caused by scattered light from the balloon entering the optical path.

The suspension system survived termination, to my eternal relief. We made no measurements of the acceleration of the payload at termination, nor of the stresses that occurred within the suspension system, so a more detailed validation of the mechanical design cannot be made. The payload also remained upright upon landing, and the rigid members of the suspension system provided protection from the rotator and spreader bar (though two of the four rigid members were slightly damaged in the process). The CSBF solar panels and ballast hoppers, which hung below the payload, sustained significant damage; the damage to the payload itself was minimal.

The NASA-provided rotator kept the payload pointed such that the center of the cameras' combined FOV was within 2° of anti-sun for $>96\%$ of flight; see the bottom plot of Figure 4.1. The rotator sporadically lost its lock on the sun, resulting

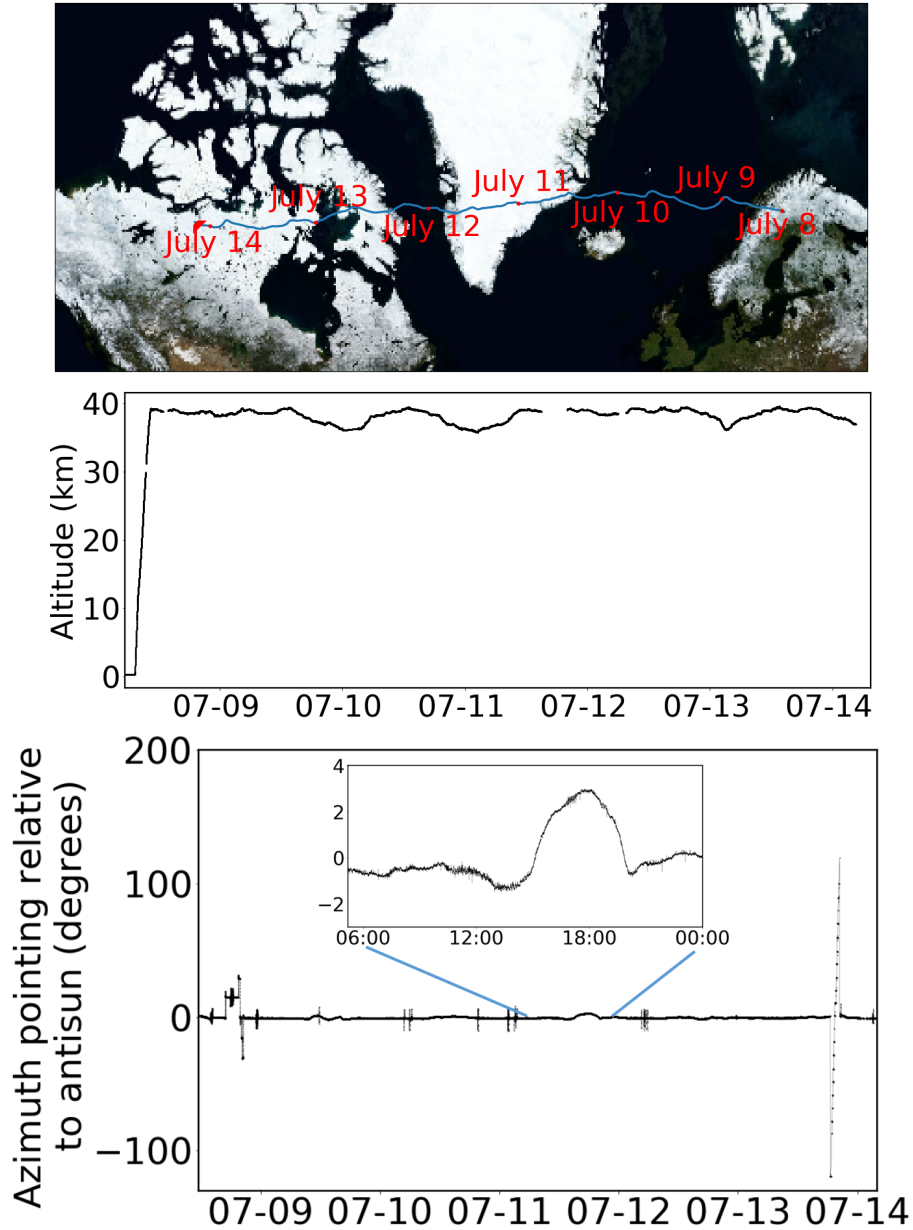


Figure 4.1: (Top) Ground track of the 2018 flight of PMC Turbo from Erange, Sweden to Nunavut, Canada. (Middle) Altitude as a function of time for the PMC Turbo payload during its 2018 flight. (Bottom) Azimuth pointing of the center of the FOV relative to the antisun direction. Inset plot shows a zoom of an 18-hour section of July 11 showing short-timescale oscillations and longer-timescale drifts in the pointing of the instrument.

in large swings away from its nominal position, though these departures were short-lived. Note that the large swings in pointing towards the end of flight were intentional and were intended to provide additional information about the performance of the cameras in flight.

4.1 PMC Turbo 2018 Dataset

PMC Turbo captured over 6 million images, comprising 120 Tb of data. Approximately 60% of these images contained PMCs. The cameras were operated in two modes, in each for approximately half of flight. In one mode, the exposure time was set to 100 ms to prevent smearing from motion blur. In the other mode, a custom auto-exposure algorithm determined the appropriate exposure time, usually between 200-300 ms, to maximize signal-to-noise of each image. The exposure time, and other metadata for each image, was also recorded. Additionally, the Rayleigh lidar determined volume backscatter, giving atmospheric temperatures up to an altitude of ~ 80 km and PMC density when clouds are present. Finally, in addition to the science data described above, various sensors logged the latitude, longitude, altitude, and azimuthal pointing of the payload.

4.2 Image Calibration

The PMC Turbo signal of interest is the cloud brightness, which we take as a proxy for the integrated column density of PMC particles along a given line-of-sight. This signal is faint, in particular when compared to Rayleigh scattering off the intervening atmosphere. Therefore, flat-fielding is required to make the PMC signal visible and remove other contributions to the raw data. Some of these contributions are additive, whereas others are multiplicative; the first step is to construct an analytical model of

the data. Let R denote the raw data – the image that is written to disk. The most significant contribution to R is the sky brightness, S . When clouds are present, they contribute an additional signal, C . Ideally, these would be the only components of the signal, and image processing would consist only in determining a way to model the sky brightness. Instead, there are several deviations from this idealized model:

1. Due to the voltage bias applied to the camera CCDs and thermal noise within the CCDs, each camera has a characteristic “dark image” – i.e. the image that result when no external light is incident on the CCD.
2. Vignetting by the optical system, natural illumination falloff, and variations in pixel sensitivity across the focal plane all contribute multiplicative factors to the recorded data. These factors are combined into a single “flat-field,” F , representing the sensitivity of each pixel in the focal plane.
3. Images during flight were taken at a range of exposure times, t_{exp} . In order that inter-comparison among the images be meaningful, this variable exposure time must be taken into account. The aperture was also adjustable, but all images that are used for science purposes had fully open apertures.
4. All of the contributions discussed above have some associated noise. This noise is dominated by photon noise and is unavoidable – therefore I will omit this factor from the analytical model discussed below. This noise can be reduced by longer exposure times, but this comes at the expense of increased motion blur. Alternatively, the images can be binned 2x2, 4x4, or larger ‘super-pixels’ to increase signal to noise. The latter is preferred, as it results in a symmetric reduction in spatial resolution, whereas motion blur decreases resolution only in the direction of relative motion between the cameras and the clouds.




Figure 4.2: Percent error in the sky brightness model that results from neglecting the curvature of the Earth when computing the air mass as a function of off-zenith angle.

Taken together, these contributions suggest the following model for the flight data:

$$R = F(S + C) + D \quad (4.1)$$

Note that multiplication here is performed elementwise rather than standard matrix multiplication. To isolate and calibrate the cloud contribution, we define:

$$I = \frac{R - D}{F t_{exp}} - \frac{S}{t_{exp}} \quad (4.2)$$

In order to make use of this formula, each of the variables must be determined. These are discussed below.

- Sky brightness:** I constructed a simplistic model of sky brightness that accounted for the primary sources of variability over flight: the air mass along a given pixel's line-of-sight and the scattering angle between the pixel's line-of-sight and solar illumination of the air column. Consider a pixel viewing at an off-zenith angle ϕ and at an azimuth angle θ , at a time when the sun is at off-zenith angle ϕ_S and at an azimuth angle θ_S . To a good approximation, the air mass along the pixel's line-of-sight is proportional to $\frac{1}{\cos \phi}$; the curvature of the Earth (and the Earth's atmosphere) prevents this expression from being exact, but over the range of angles observed by PMC Turbo, this approximation

is good to better than 0.8%; see Figure 4.2. Furthermore, the error associated with this approximation smoothly varies across the FOV, so the effect on small-scale structures is negligible. Next, I assume that the sky brightness is due to Rayleigh scattering, with a $1 + \cos^2 \theta_{scatter}$ dependence on $\theta_{scatter}$, where $\theta_{scatter}$ is the angle between the pixel viewing angle and solar radiation and has:

$$\cos \theta_{scatter} = \cos \theta_S \sin \phi \sin \phi_S + \cos \phi \cos \phi_S \quad (4.3)$$

The un-scaled sky brightness model is thus:

$$S_{unscaled}(\theta, \phi, \theta_S, \phi_S) = \frac{1}{\cos \phi} 1 + (\cos(\theta - \theta_S) \sin \phi \sin \phi_S + \cos \phi \cos \phi_S)^2 \quad (4.4)$$

As discussed in Section 3.4.3, the images themselves can be used to determine the viewing angle of each pixel by comparing the imaged star fields to known catalogs of stars. GPS sensors on the payload gave the location of the payload throughout flight, indirectly giving the position of the sun as a function of time. Thus, for each image in flight, $S_{unscaled}$ could be determined according to the above formula.

For the overall scaling, I made the approximation that the raw image is dominated by this contribution. Let $\langle R - D \rangle$ denote the average across the raw data with the associated dark image subtracted, and $\langle S_{unscaled} \rangle$ denote the average across the unscaled sky brightness model. The scaled sky brightness model is therefore given by:

$$S = \frac{\langle R - D \rangle}{\langle S_{unscaled} \rangle} S_{unscaled} \quad (4.5)$$

A valid criticism of this approach is that increased cloud brightness will in-

crease the amplitude of the estimated sky brightness. But even for bright cloud displays, the cloud brightness contributes $\sim 2\%$ of the total signal, so this is a small effect. Furthermore, the resulting bias associated with overestimating the amplitude of the sky brightness has a smooth gradient across the image, which has a negligible effect on subsequent data analysis.

- **Dark image:** D has a constant term due to the voltage bias applied to the CCD and a term that is proportional to exposure time due to thermal noise. In order to characterize D for the full range of expected exposure times, we took images before flight for each camera with the entrance aperture fully blocked. For each camera, we took images with exposure times of $60\ \mu\text{s}$ and $1\ \text{s}$; for a given exposure time, I interpolate between these two images.
- **Flat-field:** Ideally, each camera would be used to image a uniformly-illuminated field. Normalizing this image (after subtracting D) would give F directly. If a uniformly-illuminated field is not available, a field with known illumination can also be used to calculate F – dividing the image by the known illumination pattern (again, after subtracting D) and then normalizing yields F . This latter solution is the one I used. I used images taken at dusk during the final verification of our optical alignment and assumed with the sky brightness model described above accounted for the variation in illumination across the images, with one further complication. The images were taken after the sun had set (this was necessary in order that stars be visible and therefore pointing reconstruction be possible). As a result, and unlike the situation encountered in flight, the air mass imaged by the cameras was not fully illuminated by the sun. Furthermore, this effect was not uniform across the image: pixels with line-of-sight closer to zenith encountered illuminated air at lower altitudes than pixels that viewed farther off zenith. Given the exponential decrease in atmospheric

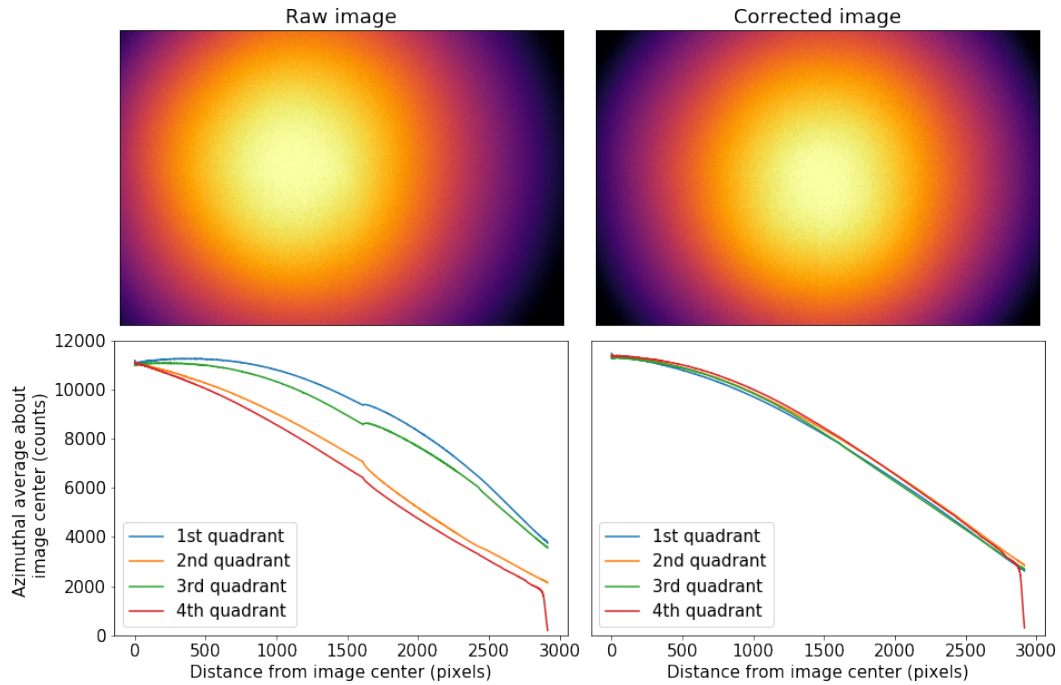


Figure 4.3: (Left, top) A raw image taken on the ground during the optical alignment tests prior to launch in Esrangle. (Left, bottom) Azimuthal average of each of the quadrants of the above image plotted as a function of distance from the center of the image. (Right, top and bottom) Same as on left, but for an image that has been corrected for the varying sky brightness in the raw image. The corrected image is approximately azimuthally symmetric about the center of the image, as expected.

density as a function of altitude, the integrated brightness above a certain altitude is proportional to the density at that altitude. Thus, for each pixel within each image, I determined the lowest point of illumination numerically, given the sun elevation angle at the time the image was taken. I then added a factor to the sky brightness model that was proportional to the atmospheric density at the lowest point of illumination (assuming an atmospheric scale height of 8 km).

To verify that this successfully removed the gradient in sky brightness, I took the azimuthal average of each quadrant of an image and over-plotted them as a function of distance from the center of the image; see Figure 4.3. In the raw image, the brightest point of the image is off-center, and the azimuthal

averages correspondingly vary. For the corrected image, the azimuthal averages are in much closer agreement. Note that the edge of the sensor had a cluster of unresponsive pixels, causing a sharp drop in the curve for quadrant 4.

Each camera had several hundred images taken under appropriate conditions. I constructed corrected images for each image individually (to account for changing sky conditions), co-added all the images for each camera, and normalized the result to construct the final flat-field image for each camera.

4.2.1 Stray Light Contamination

As discussed in Section 3.4.3, insufficient optical baffling between the camera windows and the balloon resulted in stray light in each camera’s FOV that varied on a time scale of minutes to hours. In order to emphasize the cloud signal, we perform an additional step in the image pre-processing: we subtract a moving average, usually of 10 minutes. Note that this approach also corrects for any potential errors in the sky brightness model described above. I co-added flat-fielded images in 1-minute intervals throughout flight so that generating final cleaned images required minimal processing; for a given image, the corresponding moving average image was generated by interpolating over the pre-generated co-added images for the desired interval. For a standard 10-minute moving average, this required loading 11 images, as opposed to the 1200 images required to produce the moving average from scratch.

Subtraction of this moving average results in high-contrast images of small-scale instability structure; a consequence is that it can suppress long-period motions – i.e. low-frequency GWs can be reduced or eliminated. Furthermore, this step eliminates any information about the absolute brightness of the PMCs. The final, pre-processed data product is therefore relative brightness integrated along a sloped path through the cloud layer in un-calibrated units. In the discussion that follows, I will use “im-

ages” to refer to the final, processed images rather than the raw images, except where specified otherwise.

4.3 Image Projection

To convert features within the images to physical scales on the sky, we approximated the PMC layer as a flat plane, parallel to the Earth’s surface at the location of the payload and at an altitude of 83 km. Note that a better approximation would have been to assume the PMCs formed a layer at constant altitude; the difference between these two assumptions results in a 2.5% discrepancy in interpreted spatial scales at the bottom of the FOV and a 0.3% discrepancy in the center of the FOV. We again used pointing solutions from star fields in flight to determine the mapping from each pixel to the corresponding point on the PMC layer. The cameras were rigidly mounted to the gondola frame, so the relative alignment of the cameras was fixed throughout flight. Therefore, we took a single pointing solution for each camera to be representative for the whole flight, up to an overall rotation in azimuth.

In addition to determining spatial scales, it is also of interest to be quantify translation of features within the images between frames. This apparent motion was a combination of the actual motion of the features, the translational motion of the gondola relative to the Earth, and the rotation of the gondola. Thus, for each period of interest, we designated a coordinate system that is fixed relative to the Earth with standard NESW coordinates. We used the pointing solution discussed above together with the geographic location of the gondola and the azimuthal pointing logs to determine a mapping from pixel coordinates to this fixed coordinate system, thereby incorporating the translation and rotation of the gondola.

To illustrate the entire pre-processing pipeline, Figure 4.4 shows a set of raw im-




Figure 4.4: (Top) A set of 4 raw images, each taken at 13:30 UT on July 12, 2018. (Bottom) The final pre-processed data product: these same 4 images with the flat-field algorithm applied and projected onto a plane that is tangent at an altitude of 83 km, directly above the location of the gondola at this time.

ages (top panel) all taken at the same time, then the same four images (bottom panel) having been flat-fielded according to the algorithm described above and projected onto the PMC layer. The coordinate system is such that the origin is directly above the gondola at the time the images were taken, and the +x-axis (+y-axis) is oriented in the east (north) direction. I use ‘meridional’ to refer to winds in the north-south direction and ‘zonal’ to refer to winds in the east-west direction, with the convention that a positive wind refers to a wind blowing *from* the south and west, respectively.

Chapter 5

PMC Turbo Data Analysis

The stated goals of the PMC Turbo mission were, as quoted from the original proposal[38]:

SG1: Identify the dominant GW, instability, and turbulence dynamics, scales, and intensities that define the character and scales of GW dissipation events at PMC altitudes, and

SG2: Assess the magnitudes and scales of GW momentum fluxes and momentum deposition, their statistics, and their implications for MLT forcing and secondary GW radiation.

Accomplishing these goals requires the ability to identify and characterize GWs that were present in the PMC Turbo images. In this chapter, I layout some relevant relations from linear GW theory (for which I claim no responsibility). I then describe a series of tools that I developed in order to quantify various GW parameters. In the next chapter, I apply these tools to a specific event during flight in which strong GW breaking was evident in camera and lidar data. The analysis methods presented here appear in a recently published article [39] of which I am first author.

5.1 Linear Gravity Wave Theory

The linear theory of GWs offers an approximation of the behaviour and propagation of GWs through the atmosphere in the absence of wave breaking, which is inherently a non-linear phenomenon. This linear theory is derived from the Navier-Stokes equations under the Boussinesq approximation, with the further assumptions that the background atmosphere is in hydrostatic equilibrium and that the GW presents small perturbations to the background state. Following the derivation by Nappo [40], I define a coordinate system such that the GW propagates in the xz -plane and is uniform in the y -direction. Note that I will use the terms ‘streamwise’ and ‘spanwise’ to indicate the direction of GW propagation and the perpendicular to that direction in the horizontal plane. Further, I will use ‘upstream’ (‘downstream’) to refer to motion or distance in the positive (negative) streamwise direction. I let $w_1(x, z, t)$ denote the vertical component of the GW velocity perturbation. Furthermore, I define $\tilde{w}(z)$ to carry the z -dependence of $w_1(x, z, t)$ and assume plane-wave solutions in x and t :

$$w_1(x, z, t) = \tilde{w}(z)e^{i(kx - \Omega t)} \quad (5.1)$$

where k is the horizontal wavenumber and Ω is the intrinsic frequency of the wave - this is related to the apparent frequency of the wave, ω via $\Omega = \omega - u_0k$, where u_0 is the background horizontal velocity in the direction of wave propagation. Finally, I define $\hat{w}(z) = \tilde{w}(z)e^{-(z - z_0)/2H_s}$, where z_0 is a reference altitude and H_s is the scale height of the atmosphere:

$$H_s = \frac{RT}{g} \quad (5.2)$$

with R as the universal gas constant, T the temperature and g the acceleration due to gravity. $H_s \simeq 4.4$ km at PMC altitudes. The Navier-Stokes equations can then be combined into a single equation in terms of \hat{w} ; the result is the Taylor-Goldstein equation, derived separately by Taylor [41] and Goldstein [42]:

$$\frac{d^2\hat{w}}{dz^2} + \left[\frac{N^2}{(c-u_0)^2} + \frac{u_0''}{(c-u_0)} - \frac{1}{H_s} \frac{u_0'}{(c-u_0)} - \frac{1}{4H_s^2} - k^2 \right] \hat{w} = 0 \quad (5.3)$$

Here, N is the Brunt-Väisälä, or buoyancy, frequency; typical values at the PMC layer are $N \sim 0.017$ rad s⁻¹. The apparent horizontal phase speed of the wave is given by c , and primes indicate derivatives with respect to z . Note that the phase speed is related to the apparent frequency via

$$c = \frac{\omega}{k} \quad (5.4)$$

and to the intrinsic phase speed c_i via

$$c = c_i + u_0. \quad (5.5)$$

Solutions to (5.3) give the full dynamical account of the GW by way of the “polarization equations”:

$$i\Omega\tilde{u} - \tilde{w} \frac{du_0}{dz} = \frac{i}{\rho_0} k\tilde{p} \quad (5.6)$$

$$i\Omega\tilde{w} = \frac{1}{\rho_0} \frac{d\tilde{p}}{dz} + \frac{\tilde{\rho}}{\rho_0} g \quad (5.7)$$

$$ik\tilde{u} + \frac{d\tilde{w}}{dz} = 0 \quad (5.8)$$

$$i\Omega\tilde{\rho} + \tilde{w} \frac{\rho_0}{g} N^2 = 0 \quad (5.9)$$

where quantities with subscript 0 indicate horizontally uniform background quantities and quantities with tildes indicate GW perturbations: \tilde{u} is the horizontal velocity perturbation, \tilde{p} is the pressure perturbation, $\tilde{\rho}$ is the density perturbation.

If I make the further simplifying assumption that the GW is propagating in an environment with no background wind shear (i.e. u_0'' and u_0' are both zero) and consider a GW with a vertical wavelength λ_z that is small relative to H_s^1 , (5.3) simplifies further to

$$\frac{d^2\hat{w}}{dz^2} + \left[\left(\frac{Nk}{\Omega} \right)^2 - k^2 \right] \hat{w} = 0 \quad (5.10)$$

where I have used $c - u_0 = c_i = \frac{\Omega}{k}$. This yields general solutions

$$\hat{w}(z) = Ae^{imz} + Be^{-imz}. \quad (5.11)$$

Here, the vertical wavenumber m satisfies

$$m^2 = k^2 \left(\frac{N^2}{\Omega^2} - 1 \right) \quad (5.12)$$

I keep only the positive branch of (5.11) with a negative value of m and positive value of k - this corresponds to a GW propagating in the positive x -direction with upward group velocity but downward phase propagation. The complete solution for $w_1(x, z, t)$ is thus

$$w_1(x, z, t) = |w_1| e^{z - z_0/2H_s} e^{i(kx + mz - \Omega t)}. \quad (5.13)$$

where $|w_1|$ denotes the magnitude of the vertical velocity perturbations at the reference altitude. Together with the polarization equations above ((5.6)-(5.9)), this

¹More explicitly, the requirement is that $m \gg \frac{1}{2H_s}$, i.e. $\lambda_z \ll 4\pi H_s$

provides a complete description of the GW perturbations.

Additional Relations from Linear Gravity Wave Theory

Given a wave with horizontal and vertical wavenumbers k and m , respectively, the wave propagates with k-vector

$$\vec{k} = \begin{pmatrix} k \\ m \end{pmatrix} \quad (5.14)$$

while lines of constant phase are along $kx + mz = 0$, i.e. $z = -\frac{k}{m}x$. Such a wave propagates at an angle

$$\beta = \arctan\left(\frac{m}{k}\right) \quad (5.15)$$

relative to the horizontal plane. The intrinsic frequency of the wave depends wholly on this angle of propagation: noting that $\cos \beta = \frac{k}{\sqrt{k^2 + m^2}}$ and solving (5.12) for Ω gives

$$\Omega = \frac{Nk}{\sqrt{k^2 + m^2}} = N \cos \beta. \quad (5.16)$$

I define the intrinsic GW period to be $T_i = \frac{2\pi}{\Omega}$, and the buoyancy period to be $T_b = \frac{2\pi}{N}$.

Finally, I define the final, dimensionless GW amplitude

$$a = \frac{|u_1|}{c} = \frac{|w_1|}{c_z} \quad (5.17)$$

where $c_z = \frac{\Omega}{m}$ is the vertical phase speed and an amplitude $a = 1$ is defined as the overturning amplitude, at which point the GW is convectively unstable. By integrating (5.13) with respect to time, we see that a GW with amplitude a will cause vertical

displacements with magnitude

$$\delta_z = \frac{|w_1|}{\Omega} = \frac{|w_1| c_z}{c_z \Omega} = a \frac{1}{m} = a \frac{\lambda_z}{2\pi}. \quad (5.18)$$

5.2 Mechanisms for Brightness Perturbations in Camera Images of PMCs

An advantage of studying GW breaking at the PMC layer is that GWs are the dominant forcing at these altitudes, and perturbations to the cloud brightness can be attributed to GWs with high confidence. An unforeseen disadvantage is that the signature of GWs in the PMC layer is not as strong as was expected, due to the thinness of the layer, whereas small-scale instabilities are bright and ubiquitous in comparison.

Airglow, another common tracer used to study GWs (see, for example Taylor [43]), responds immediately to temperature perturbations associated with GWs, so GWs feature prominently in airglow images, but the mechanism by which GWs imprint their signature on PMCs is more complicated. Before attempting to identify GWs in the PMC Turbo images, it is worthwhile to consider why we expect GWs to leave a signature in the PMC brightness in the first place.

There are several competing mechanisms by which a GW can perturb the apparent brightness of the PMC layer. In the discussion that follows, it is important to keep in mind that a camera is sensitive only to the integrated column brightness along each pixel's line of sight. For the purposes discussed here, I will assume a camera that is viewing directly overhead so that this column brightness corresponds to a purely vertical integration.

5.2.1 Horizontal Convergence at the PMC Layer

As a GW perturbs the PMC layer, there are areas of horizontal convergence and divergence that create perturbations in the apparent PMC brightness. To quantify this effect, I assumed the GW described by (5.13) with a horizontal wavelength of 15 km and a vertical wavelength of 5 km, propagating in a background atmosphere with no background wind (in these conditions, there is no difference between apparent and intrinsic parameters of the wave). I assumed an initially flat and unperturbed 1 km thick PMC layer centered at an altitude of 82.5 km, and sampled this cloud layer via a lattice of points spaced by 100 m in x and by 10 m in z . Each of these lattice points represents a PMC ice particle being advected by the GW as it propagated through the layer; it is these motions that result in perturbations in the apparent brightness of the PMC layer. I let the GW amplitude gradually ramp up over a period of 1,000 s then determined the column density of the original lattice points in their perturbed locations. The apparent brightness detected by a camera is proportional to this column density. Note that I assume that the PMC particles do not undergo any sublimation or growth as they are advected to and fro; this mechanism will be discussed in Section 5.2.2.

Figure 5.1 shows a horizontal slice of the cloud layer at intervals of 400 s; the cloud layer is distorted into a sawtooth pattern that propagates from left to right. At the vertical portion of the sawtooth pattern, the horizontal convergence of the PMC layer results in higher column density and therefore apparent brightness.

I ran the above simulation for final GW amplitudes of 0.1, 0.5, 0.7, and 0.9. I show the resulting spatial perturbations and associated brightness. The region of greatest horizontal convergence and greatest brightness occurs downstream from the greatest vertical perturbation; at small GW amplitudes, the phase shift between these quantities is $\sim 90^\circ$ but at larger amplitudes this phase shift shrinks. The




Figure 5.1: Spatial GW perturbations to an initially flat layer at (from top to bottom) 0 s, 400 s, 800 s, and 1200 s. The points are color-coded by their initial location in the x -direction.

magnitude of this effect increases with GW amplitude, but even at an amplitude of $a = 0.1$, horizontal convergence due to spatial GW perturbations will result in a $\sim 10\%$ perturbation; see Figure 5.3, which shows the magnitude and phase shift associated with this effect as a function of GW amplitude.

As is evident from (5.9), there are density perturbations associated with the GW that effect the PMC brightness - but these density perturbations are a result of the spatially-varying velocity perturbations rather than a separate effect. These density perturbations contribute to the brightness variations shown in Figure 5.2, but they are a small contribution to the overall effect. From (5.9), the fractional perturbation in density is given by

$$\frac{\tilde{\rho}}{\rho_0} = i \frac{N^2}{\omega g} \tilde{w}. \tag{5.19}$$

As the GW amplitude approaches overturning, \tilde{w} approaches c_z , the vertical phase speed, which is given by $\frac{\omega}{m}$. Thus, the largest possible amplitude that the density

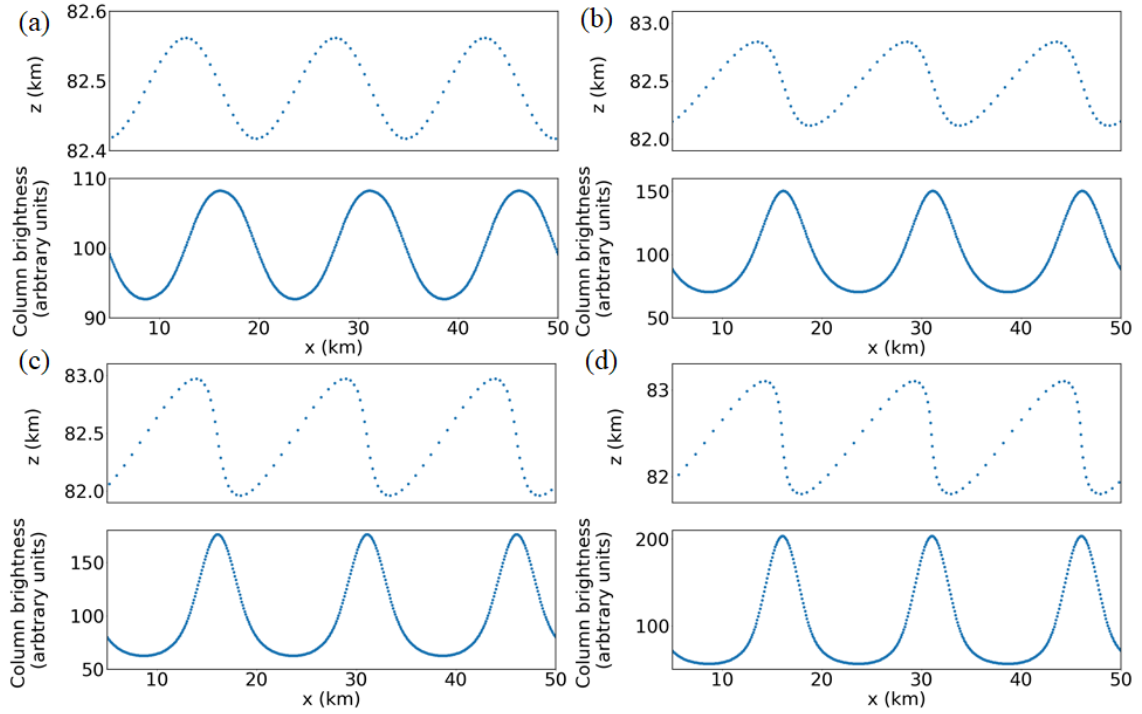


Figure 5.2: (a, top panel) Perturbed GW layer at 2000 s showing horizontal- and vertical-perturbations for a GW amplitude of $a = 0.1$. (a, bottom panel) Integrated column brightness (density) in arbitrary units of the GW layer illustrated in the top panel. (b-c) Same as (a) for GW amplitudes of 0.5, 0.7 and 0.9, respectively.

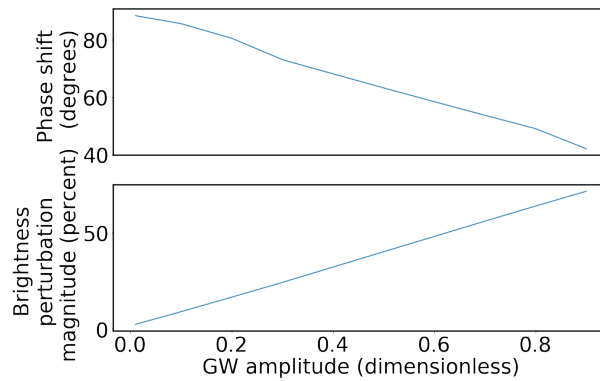


Figure 5.3: (Top panel) Phase shift between the region of greatest vertical displacement and highest brightness as a function of the dimensionless GW amplitude a for the horizontal convergence mechanism of brightness modulation. (Bottom) Magnitude of the brightness perturbations as a function of a .

perturbations can attain is

$$\left| \frac{\tilde{\rho}}{\rho_0} \right| = \frac{N^2}{mg}. \quad (5.20)$$

For the values used here, this limits the fractional density perturbations to less than 2.5%. Thus, the majority of the horizontal convergence effect is due to the folding of the PMC layer rather than perturbations to the PMC layer density.

5.2.2 Temperature Fluctuations

In the analysis described above, the PMC particles were treated as immutable objects, subject only to spatial translations. In reality, vertical motion imparted to the PMC particles causes temperature perturbations that can drive particle sublimation or growth. Upward motions are accompanied by expansion and decreased temperatures, while downward motions result in compression and warming.

The particle growth and sublimation rate of PMCs has been estimated by a range of experiments and models: Chandran et al. [44] estimated the PMC particle growth rate to be no more than 10 nm hr⁻¹ based on model output; Zasetsky et al. [45] measured rates up to 30 nm hr⁻¹; Gadsden and Schröder [46] calculated growth and sublimation rates up to 35 nm hr⁻¹; and Rapp et al. [47] presented model results that suggest growth rates up to 8 nm hr⁻¹ but sublimation rates that exceed that by almost an order of magnitude. It is important to note that both Chandran et al. [44] and Rapp et al. [47] base their results on models of specific GW parameters, namely waves with horizontal wavelength $\lambda_h \sim 600$ km and periods of many hours. These are not representative of waves that PMC Turbo was sensitive to, and it is certainly plausible that higher-frequency waves, with higher associated vertical velocity perturbations, result in higher growth and sublimation rates. On the other hand, Rapp et al. [47]

also suggest that the effect more generally is asymmetric in growth and sublimation, with sublimation occurring much more rapidly than growth. This implies that in as much as temperature may be a mechanism responsible for modulating the PMC particle size, the net effect will be an overall depletion of the PMC particles.

Given the need to subtract a moving average of 10 or tens of minutes, PMC Turbo was primarily sensitive to dynamics occurring on shorter timescales than that. If we consider phenomena with a timescale of 5 minutes, the estimates quoted above suggest changes in particle size of no more than 3 nm. For a PMC particle radius of 55 nm [48], this corresponds to a 5.5% change in particle size.

The PMC Turbo cameras were sensitive to wavelengths above 600 nm, with quantum efficiency of the CCD decreasing to zero around 1000 nm. Given this range of observing frequencies, Savigny and Hoffman [49] suggest a dependence of scattering cross section on particle radius, r , that scales with r^k with $5.1 \leq k \leq 5.9$. Let B denote the cloud brightness. Taking the lower range of possible k , we have

$$B \propto r^{5.1} \tag{5.21}$$

and therefore

$$\frac{\partial B}{\partial t} \propto 5.1 r^{4.1} \frac{\partial r}{\partial t} \tag{5.22}$$

which we can re-arrange to find

$$\frac{1}{B} \frac{\partial B}{\partial t} = 5.1 \frac{1}{r} \frac{\partial r}{\partial t} \tag{5.23}$$

i.e. the fractional change in brightness is a factor of 5.1 larger than the fractional change in particle radius. The 5.5% change in radius discussed above corresponds




Figure 5.4: Phase shift between the region of greatest vertical displacement and highest brightness as a function of the dimensionless GW amplitude a for the temperature perturbation mechanism of brightness modulation.

to a 28% change in brightness. This magnitude of variability is detectable in the PMC Turbo data for even modest PMC displays, so this remains a plausible (though unlikely) mechanism for PMC brightness modulation.

Qualitatively, temperature modulation of PMC brightness will cause a decrease in brightness for PMC particles that have been advected downwards and an increase in brightness for PMC particles that have been advected upwards. Assuming a PMC layer initially at equilibrium, PMC particles will grow as they are advected upward, and continue to grow as they are advected back down, until they reach their starting position. Thus, the greatest change in brightness will occur upstream from the area of greatest vertical displacement; the magnitude of this phase shift is $\sim 90^\circ$ at small GW amplitude, but increases as the amplitude increases; see Figure 5.4. Recall that the phase shift for the horizontal convergence mechanism is in the opposite direction for all GW amplitudes. The phase shift between the region of greatest PMC brightness and the region of greatest vertical displacement is therefore a reliable indicator of which mechanism is responsible for perturbing PMC brightness.

5.3 Identifying a Gravity Wave in Image and Lidar Data

Given these two plausible mechanisms by which a GW may imprint its signature on images of the PMC layer, there is motivation to search for GWs in the camera data. I performed this search with two independent methods: in the spatial and in the spatial frequency domains. The lidar data provides independent validation to the identification of a GW in the camera data by quantifying the associated vertical displacements.

5.3.1 Spanwise Averaging of Camera Data

GWs rarely feature prominently in the PMC Turbo images; more typically, small-scale instabilities exhibit much greater contrast in brightness, obscuring the larger-scale patterns. But the GW contribution is expected to be coherent over large scales; in particular, the GW signal is expected to be uniform over lines of constant phase. Averaging over such lines will thus amplify the GW signal relative to the small-scale structure.

We have no *a priori* knowledge of the propagation direction of any hypothetical GW. Therefore, I perform the averaging over all possible propagation directions, in 5° intervals. In each case, I divide the projected combined FOV of the cameras into parallel 1 km wide bins and average across each bin; see an example in Figure 5.5. When the orientation of these bins lies in the spanwise direction (i.e. perpendicular to the GW propagation direction, or streamwise direction), I expect the average to exhibit periodic structure; at other orientations, each bin will cross multiple phases of the wave, and the periodic structure will be suppressed. These bins each contain on average 5,900 points; this averaging will therefore result in a factor of ~ 77 in-




Figure 5.5: Example binning for an assumed GW propagation direction (streamwise direction, blue arrow) for the central two wide-field cameras. Points are color-coded by bin; all points within a given bin are averaged together.

crease in signal-to-noise and a similar reduction in the contributions from small-scale instabilities.

For each orientation of the bins, I plot the average across each bin as a function of distance in the assumed streamwise direction. I then find the amplitude of the best-fit sinusoid to each average. I take the angle at which this amplitude is maximized to be the angle that corresponds to true spanwise averaging.

Figure 5.6 demonstrates this technique for simulated data with low-amplitude periodic structure and a relatively high level of white noise. The top panel on the right shows the spanwise average as a function of distance in the streamwise direction, demonstrating that the sinusoidal structure is recovered with good signal-to-noise. The bottom panel shows binned averaging at an orientation that is rotated by 20° relative to the spanwise direction; the periodic structure is no longer apparent. Figure 5.7 shows the amplitude of the best-fit sinusoid for each of the bin orientation angles. As expected, a sharp peak occurs at the orientation that aligns with the spanwise direction of the simulated GW.




Figure 5.6: (Left) Simulated data for two wide-field cameras' projected FOV, showing a faint periodic signal dominated by noise. (Right) Binned average of the data plotted as a function of distance in the perpendicular direction, with the bins aligned with the spanwise direction (top) and rotated by 20° relative to the spanwise direction.




Figure 5.7: Amplitude of the best-fit sinusoid for binned averaging of the simulated data shown in Figure 5.6, as a function of bin orientation angle (black dots). The blue vertical line corresponds to the spanwise direction. The sharp peak in the best-fit amplitude demonstrates that this technique correctly identifies the GW orientation angle.

An advantage of this approach is that once I have identified the orientation of a GW in the images, I can repeat the spanwise averaging on subsequent images. The periodic structure advances across the image as the wave propagates and is advected through the FOV. In order to quantify this motion, I assume that each peak and trough in the spanwise averaged data has a Gaussian profile. At intervals of 10 s, I find the best-fit location of all peaks and troughs in the spanwise averaged data, letting the location, vertical offset, amplitude and width of the Gaussian vary freely. Over the time segment of interest, I obtain location as a function of time for each peak and trough. The average separation of neighboring peaks (or troughs) gives an estimate of the horizontal wavelength. I find the best-fit line to the locations of each peak and trough, and I take the average of the slopes of these lines to be the apparent phase speed of the wave. This technique is demonstrated on flight data in Figure 6.3.

5.3.2 Spatial Power Spectrum of Image Data

Given the presence of faint but coherent periodic structure, another method of quantifying that structure is through a Fourier transform from the spatial to the spatial frequency domain. A slight complication is that the data is irregularly sampled on the cloud layer (and irregularly periodic in pixel coordinates) - I can interpolate onto a regularly sampled grid on the cloud layer, but some loss of information occurs in the process. Additionally, given the size of the FOV ($\sim 100 \times 100$ km), the maximum resolution in the spatial frequency domain is 10^{-2} km^{-1} , corresponding to steps of 4 to 9 km at scales of 20 to 30 km.

Nevertheless, I use this method to validate the results found in the spanwise averaging technique described above. I rotate the coordinate system to make the camera FOV fit more compactly in a rectangular domain, such that the y - and x -axes lie in the anti-sun and perpendicular to anti-sun directions, respectively. Figure 5.8




Figure 5.8: Power spectrum of the simulated data shown in Figure 5.7. The white dashed line corresponds to the orientation of the GW. The coordinate system used here has k_{\parallel} as the anti-sun direction and k_{\perp} as the perpendicular to anti-sun direction.

shows a portion of the power spectrum of the simulated data shown in Figure 5.7 with a white dashed line that corresponds to the orientation of the GW. Unsurprisingly, the power spectrum shows a peak along that line. I fit a 2D Gaussian to the power spectrum to determine the location of that peak and take the uncertainty in each coordinate to be half a pixel. For an input GW with λ_h of 21 km with a propagation direction 125° clockwise from the anti-sun, I found $\lambda_h = 20 \pm 3$ km at $122^\circ \pm 7^\circ$ from anti-sun.

5.3.3 Identifying Gravity Waves in Lidar Data

As demonstrated in Figure 5.1, GWs cause vertical displacements to an initially flat PMC layer. The lidar is sensitive to the vertical structure of the PMC layer and can therefore measure the altitude of the PMC layer at the location of the lidar beam as a function of time; in the presence of a GW of sufficiently large amplitude, the lidar will resolve vertical oscillations in the PMC layer altitude at that point.

The processed data product that I use for this analysis consists of volume backscatter coefficients as a function of time and altitude, with resolution of 5 s and 60 m, respectively. I estimate the PMC layer altitude by finding the altitude with the greatest volume backscatter coefficient at each time. To identify the dominant GW

perturbation over an interval of interest, I fit a sinusoid to the altitude as a function of time. Note that the period of any apparent GW is neither the intrinsic period of the wave (measured in a frame that propagates with the wave) or the apparent period (measured in a fixed coordinate system). Still, the lidar data can provide (or fail to provide) validation to the identification of a GW in the image data; the data can reveal the phase shift between the GW brightness perturbation and vertical displacement maxima; and the data can quantify the magnitude of the vertical displacements δ_z , which gives a lower bound on the vertical wavelength via (5.18).

5.4 Lidar and Camera Correlations

The coincident lidar profiles provide information on the vertical structure of the PMC layer that is absent from the camera images and are valuable in aiding the interpretation of the dynamics observed in flight. In order to make an accurate comparison between the lidar and camera data, I first determined the point at which the lidar beam intersected with the PMC layer within the camera images. I did this by correlating two datasets: the pixel-by-pixel timestreams of brightness in the final processed images in the known vicinity of the lidar beam, and the vertically integrated brightness through the PMC layer as measured by the lidar. The lidar data used in this comparison was sampled at a 10 s cadence, so I binned the camera timestream data temporally in 10 s intervals, averaging the data within each bin. The final processed images used in this comparison had a 10 minute moving average subtracted, so I removed a 10 minute moving average from the lidar timestream as well. I normalized each of the resulting timestreams to lie between 0 and 1; this aids in visually comparing the two datasets, though it does not affect the resulting correlation analysis. I performed the following analysis on a segment of data from 2:40 - 3:00 UT on July 10,

2018, a period in which bright small-scale structure was visible in both the camera images and the lidar timestream. The significance of this period will be discussed further in subsequent sections.

I iteratively zeroed in on the precise location of the lidar beam: as a first pass, I took a 300x300 pixel region of the images from camera 5 (corresponding to approximately $2.5 \times 2.5^\circ$) and binned it in 4x4 pixels to obtain a 75x75 pixel grid. For each (binned) pixel timestream, I computed the Pearson p-value of correlation between it and the lidar timestream. Where the correlation was strongest (lowest p-value), the likelihood of overlap between the lidar and that particular pixel was highest. I fit a 2D Gaussian profile to the logarithm of Pearson p-value and took the center of that Gaussian fit to be the best estimate of the location of overlap. Figure 5.9 shows the p-value correlation data and best-fit Gaussian for this and subsequent steps.

I repeated the above process on a 150x150 pixel grid centered around the solution in the above step, binned in 2x2 pixels, then again on an unbinned 75x75 pixel grid centered on the solution found in the second step. I took the center of the 2D Gaussian fit to the unbinned data to be the best estimate of the overlap location. I performed a χ^2 analysis on the fitting process to determine the uncertainty on the result; I found an uncertainty of 3 pixels in each direction, corresponding to 20 m when projected to the PMC layer. The timestreams for the lidar and for the camera pixel at the determined point of overlap are shown in Figure 5.10 and demonstrate excellent correlation (as evidenced by the Pearson p-value of 10^{-14} determined from the final Gaussian fit).

I repeated the above analysis for a range of other periods of high PMC signal to noise and found the average of the pixel coordinates of the best-fit location of overlap. This point was within 3 pixels in each direction of the pixel determined above, which is consistent with the uncertainty determined above.




Figure 5.9: (Left column) Logarithm (base 10) of the Pearson p-value of the correlation between lidar and pixel timestreams. (Right column) Best-fit 2D Gaussian profiles to those data. From top to bottom, the rows contain data for 4x4 binning, 2x2 binning, and unbinned data.




Figure 5.10: Integrated PMC brightness as measured by the lidar (orange) and one of the cameras (blue) at the determined point of overlap between these instruments. Both datasets have been normalized to lie between 0 and 1.

5.5 PMC Layer Wind Estimates

In order to differentiate between the apparent and intrinsic properties of any GW that we detect in the PMC Turbo data, it is crucial to constrain the background wind in which the GW propagates. I used two independent methods to estimate this background wind. The first, from the PMC Turbo data itself, was through feature tracking of objects within the images between consecutive frames using an open-source Python package *Trackpy* [50]. The second method was from a global reanalysis using the US Navy Global Environmental Model (NAVGEM, Eckermann et al. [51]) that provided wind and temperature estimates for altitudes up to 100 km.

5.5.1 Wind Estimates from *Trackpy*

The PMC Turbo dataset consists of images with a cadence of >0.5 Hz. Whenever clouds were present, small-scale structures that evolved on timescales of minutes were ubiquitous within the images. These characteristics make feature tracking a suitable analysis technique.

Trackpy is an open-source package that was written in IDL with wrappers that allow it to be integrated easily into Python, which was used for all the data analysis I performed. It automatically identifies features (regions within the image that ex-

hibit sharp contrast in brightness over a user-selected spatial range) and links these features between consecutive images to form trajectories. GWs at different scales can contribute phase-synchronous velocity perturbations to these trajectories: GWs with $\lambda_h < 100$ km will have their contributions average to zero over the full FOV, and GWs with larger λ_h will have a slowly-varying contribution that effectively contributes to the background wind for the purposes of the scales we are sensitive to. Thus, while individual trajectories may be biased by the presence of GWs, the average of all trajectories across the FOV will represent an accurate estimate of the background wind.

Optimization and Validation of *Trackpy*

Before applying *Trackpy* to flight data, I conducted a series of tests to gain confidence in the output of the *Trackpy* algorithm and select optimal input parameters.

First, I generated a series of input images based on a single images from flight. Each frame was a subsection of the complete image, with the window shifting by a fixed amount between each frame to simulate motion. I added a unique noise realization to each frame, with a noise level estimated from the standard deviation across small portions of the original image. This is conservative, for two reasons: (1) the original image already contains noise, so adding a unique noise realization to each frame will result in an artificially high noise level, and (2) the method I used to estimate the noise level makes the implicit assumption that there is no contribution from cloud brightness on small scales. I chose regions of the image that were relatively devoid of clouds, but this still results in an over-estimate of the noise level. I generated a series of images for each of a range of input velocities and ran the *Trackpy* algorithm on each set of images. For each input velocity, *Trackpy* found a list of feature trajectories within the series of images. I found the average velocity of each




Figure 5.11: Input velocities (black dots) and recovered *Trackpy* velocity estimates (blue dots and error bars) for images with motion simulated by taking a moving window across a single image from flight. Each panel represents a single input velocity and *Trackpy* estimate.

trajectory, and I took the average of all the trajectories to be the best estimate of the input trajectory. I used the standard deviation among the trajectories to estimate the uncertainty.

The result is plotted in Figure 5.11. In each case with non-zero input velocity, the estimated velocity is smaller in magnitude than the input velocity; the average magnitude of this bias is 3%. This bias is the inevitable result of spurious feature-linking between consecutive images: each trajectory consists of a series of locations that *Trackpy* has deemed to correspond to the same feature. Where *Trackpy* has correctly identified repeated features, the average inferred velocity will match the input velocity. But when *Trackpy* incorrectly identifies features as being repeated, or finds a particle trajectory that is entirely an artifact of the noise realizations,




Figure 5.12: Input velocities (black dots) and recovered *Trackpy* velocity estimates in the presence of a GW (blue dots and error bars) and in the absence of a GW (red dots and error bars) for simulated images. Each panel represents a single input velocity.

the inferred velocities will average to 0. The overall velocity estimate represents an average of these two types of trajectory identifications, so the result will tend to approximate the input velocity with a bias towards values that are closer to zero. Note that this bias is well within the calculated uncertainty

Next, I simulated images with three components: a random realization of small-scale structures modelled after those observed in the PMC Turbo data, a sinusoidal variation in the brightness to simulate the passage of a GW, and a unique noise realization, with the amplitude of each of these components based on estimates from the image data. I input a range of input velocities for the small-scale structure and let the GW field propagate separately. Further, for each input velocity, I ran the *Trackpy* algorithm with and without GW propagation to determine if the GW motion biased the results. The result is shown in Figure 5.12 for each input velocity. Again, there is a slight bias in the retrieved velocities towards smaller magnitude values, though in this case the bias is $<1\%$. Interestingly, the simulations with a propagating GW have smaller bias and smaller uncertainties than the simulations with a static GW.

In all cases, the retrieved velocity estimates are well within the uncertainties of the input values.

Finally, I tested the robustness of *Trackpy* across a range of input parameters. These parameters consists of

- Feature size: the size of feature that *Trackpy* looks for in the images.
- Memory: the number of frames that *Trackpy* will look through when linking features between consecutive images. A longer memory allows features to disappear and reappear while still being linked into coherent trajectories.
- Mass: *Trackpy* quantifies the magnitude of each given feature and allows a user-input cutoff limiting the minimum magnitude of feature that is included in the analysis.
- Occurrences: the number of times a given feature must make an appearance before it is included in the analysis.
- Search range: the distance, in pixel counts, that *Trackpy* will search over when linking features between consecutive frames.

The last parameter can be determined *a priori* - wind speeds above 100 m s^{-1} are rare at these altitudes, and this velocity limit can be converted into a search range depending on the bin size used in the images and the time interval between images. A larger-than-necessary search range allows for greater possibility that *Trackpy* falsely identifies features with each other.

For the rest of the parameters listed above, I tested *Trackpy* over a range of reasonable input values to determine the sensitivity to and the optimal value for each parameter. I used a series of images from flight binned by 8x8 pixels spaced at 2 s over an interval of 2 minutes. I used *Trackpy* to find trajectories in pixel coordinates, then




Figure 5.13: (Top) Estimated zonal velocity as a function of feature size. (Middle) Same as top, for meridional velocity. (Bottom) Number of trajectories found as a function of feature size.

converted those trajectories to coordinates on the sky to retrieve velocity estimates in physical units.

Figure 5.13 shows the results of this analysis over a range of input feature sizes. The error bars in the top two panels indicate the standard deviation of the across the FOV. Clearly, the results are not sensitive to this parameter. I chose a feature size of 7 pixels to maximize the number of features (and therefore make the coverage of the FOV more uniform). Note that this is counting binned pixels and corresponds to 56 pixels in native pixel units. The features being found are therefore ~ 1 km in size.

Figure 5.14 shows the results over a range of input memory values. Again, the results are highly robust across the full range of input values. As the length of memory increases, there is a slight increase in the number of features found. This is accompanied by an increase in the error bar on the resulting meridional velocity value - suggesting that these additional trajectories may be spurious. I therefore chose to




Figure 5.14: (Top) Estimated zonal velocity as a function of length of memory. (Middle) Same as top, for meridional velocity. (Bottom) Number of trajectories found as a function of length of memory.

use a memory value of 2.

Figure 5.15 shows the results over a range of input minimum mass values. Again, the algorithm is robust over the full range that I tested. There is a relatively steep fall-off in the number of trajectories that were found as the minimum mass increased, and low-mass features do not appear to bias the retrieved velocities. All masses are positive, by definition; I chose a minimum mass of zero, meaning no trajectories were rejected on the basis of mass.

Note that by default *Trackpy* identifies features that are bright relative to the background, but there is the option to invert the image to find features that are dark relative to the background in the original images. To maximize the number of features found and obtain greater coverage of the FOV, I found all trajectories in the original and in the inverted images.

Figure 5.16 shows the results over a range of minimum occurrence values. Yet




Figure 5.15: (Top) Estimated zonal velocity as a function of feature mass. (Middle) Same as top, for meridional velocity. (Bottom) Number of trajectories found as a function of feature mass.

again, the retrieved velocities show little dependence on this parameter, though there is an increase in the calculated uncertainties at small values of the occurrences limit, which indicates a higher rate of spurious trajectories. I chose a lower limit of 10 frames to reduce this effect.

A final input parameter that is independent of *Trackpy* is the filter that I apply to the images before running the *Trackpy* algorithm. This filter is necessary because stars in the raw images will appear as long, dark streaks in the final, processed images due to the subtraction of a moving average. These streaks, which are several hundred pixels in length but less than 10 pixels in width, are reliably identified as features due to their sharp contrast. The apparent motion of these features is due solely to the translation of the gondola (which changes the mapping of pixel to sky coordinates) and slow rotation of the gondola (which shifts the stars in pixel coordinates). This motion is completely independent of the winds at the PMC layer and systematically




Figure 5.16: (Top) Estimated zonal velocity as a function of minimum number of occurrences. (Middle) Same as top, for meridional velocity. (Bottom) Number of trajectories found as a function of minimum number of occurrences.

biases the results. Thus, I apply a 2D Gaussian high-pass filter to the images in the spatial domain to eliminate these small, sharp features. I quantified the effectiveness of each filter by counting the number of trajectories that were clustered around the drift speed of the gondola (see, for example Figure 5.17, which shows the individual trajectory velocities for a series of unfiltered images). I found that a filter width of 31 pixels eliminated all trajectories associated with stars.

Given a set of trajectories found by *Trackpy*, I filtered out those trajectories that were not well approximated by a constant velocity. I quantified this by assuming a single pixel uncertainty in the feature locations. I fit the x position and y position separately as a function of time and found the χ^2 value for each fit. I kept only those trajectories with probabilities to exceed of less than 99.5% in both coordinates.




Figure 5.17: Average velocities for trajectories found by *Trackpy* using unfiltered images. Each dot represents the average velocity for a single feature trajectory. The black errorbars indicate the average and standard deviation across all trajectories. The cluster of trajectories at the top-middle of the figure (black circle) correspond to stars that are tracked between images.

5.5.2 *Trackpy* Analysis Pipeline

Overall, the results of this analysis demonstrate an algorithm that is robust over a wide range of input parameters. As such, it is a valuable tool to estimate the local wind velocities within the PMC Turbo FOV whenever clouds are present in the images. The final analysis procedure I followed was as follows:

- Generate a series of 60 images at 2 s intervals.
- Apply a 2D gaussian high-pass filter to the images, with a filter width of 31 pixels.
- Bin the images into 8x8 pixel super-pixels.
- Find all trajectories in the images, with the following parameters:
 - Images inverted: True and false, separately
 - Feature size: 7 pixels
 - *Trackpy* memory: 2 frames
 - Minimum feature mass: No limit.

- Minimum number of feature occurrences: 10 occurrences
- Search range: 3 pixels
- Find best-fit velocities for each pixel in zonal and meridional directions independently.
- Filter trajectories using the χ^2 metric to remove trajectories with non-linear trajectories.
- Average the best-fit zonal and meridional velocities across all remaining trajectories.

5.5.3 Wind Estimates from NAVGEM

Independent wind estimates were also available from a NAVGEM reanalysis that incorporated meteor-radar wind measurements from 25 sites around the globe (note that the inclusion of these measurements is a new development since the most recent NAVGEM publications). The native resolution of the NAVGEM output was 1 km vertically and 100 km horizontally on a time scale of 1 hour. In addition to meridional and zonal winds, NAVGEM also supplied temperature profiles which gave higher-confidence constraints on H_s (which is a function of temperature) and N (which is a function of temperature gradient). Uncertainties in the NAVGEM output were assessed by comparing it to radar measurements at an additional site over a period of ~ 2.5 months, as described by Eckermann et al. [51].

Chapter 6

Case Study: Gravity Wave Breaking and Vortex Ring Formation

From 2:40 to 3:00 UT on July 10, 2018, bright and ubiquitous ring-shaped instability structures were apparent across the entire combined FOV of the central two wide-field cameras. Stray light in the other two wide-field cameras was particularly bad during this period, so the analysis presented here focuses solely on the central two cameras. This event was the subject of a recent article [39] of which I was first author; much of the work presented in this chapter appears in that paper. Unless otherwise noted, I am responsible for the work I describe here.

At the time of this observation, the gondola was at an altitude of 36.3 km above 69.3°N and 17.4°W. The FOV of the cameras was oriented towards the south-southwest, and the gondola was drifting towards the west-southwest at 8 m s⁻¹.




Figure 6.1: Combined and projected FOV of the central two wide-field cameras at 2:45 UT on July 10, 2018. The vortex rings appear as dark circular or slightly elliptical cores with diameters of 2-5 km surrounded by a partial or complete bright rings.

6.1 Vortex Rings

These instability structures, shown in Figure 6.1 are of particular interest because of their resemblance to a type of convective instability known as vortex rings. Numerical simulations of gravity wave breaking have demonstrated that gravity waves that grow in amplitude and become convectively unstable (rather than dissipating through dynamical instabilities or undergoing critical layer filtering) will reliably produce vortex rings [52, 53], but there have been few empirical confirmations of vortex rings formation.

In particular, Fritts et al. [54] has demonstrated the formation of vortex rings in response to breaking GWs at a range of inertial frequencies. They have extracted quantitative relationships between the diameter (D) and separation (d) of the vortex rings and the vertical wavelength and inertial frequency of the underlying GW:

- $D \simeq \frac{\lambda_z}{2}$
- $d \simeq D$ for a GW with $\omega_i \geq N/3$

where both D and d refer to the values of those parameters prior to vortex ring dissipation. Vortex rings often split into smaller vortex rings (see Figure 6.10), so the comparison is only meaningful for the largest observed vortex rings. Fritts et al. [54] also predicts a typical time scale of $\sim 1T_b$ for the evolution from coherent vortex rings to turbulence.

Empirical validation of these relationships requires observation of vortex rings formation in response to GW breaking and independent characterization of the parameters of the underlying GW parameters.

6.2 Characterization of Gravity Wave and Background Environment

In this section, I apply the techniques that are laid out in Chapter 5 to characterize the propagation direction, horizontal wavelength, apparent phase speed, and vertical displacements associated with the GW that was responsible for the observed vortex rings, and the background environment (and in particular winds) in which this event occurred.

6.2.1 Gravity Wave Characterization

Propagation Direction, Horizontal Wavelength, Apparent Phase Speed and Vertical Displacements

As described in Section 5.3, there are two distinct methods for identifying and characterizing a GW in the image data: spanwise averaging in the spatial domain, and a power spectrum in the spatial frequency domain. I applied each of these techniques to the image data during this event.

Calculating binned averages at a range of bin orientation angles determined a GW propagating at an angle 125° clockwise from north (see Figure 6.2) - hence I will refer to binned averaging at an orientation perpendicular to that direction as spanwise averaging. I calculated spanwise averages at 10 s intervals over the 20 minute duration of the event, and found the best fit location of the peaks and troughs at each time step. Figure 6.3 shows an example spanwise average with peaks and troughs marked in red and blue, respectively, and Figure 6.4 shows the series of spanwise averages as a color plot with a red (blue) line marking the linear fit to the locations of each peak (trough). The average separation between neighboring peaks gives a horizontal wavelength of $\lambda_h = 18 \pm 0.5$ km. The average slope of the linear fits to the peak/trough




Figure 6.2: Amplitude of the best-fit sinusoid to the binned average of the image data as a function of orientation angle of the bins, using data from the central two wide-field cameras at 2:45 UT on July 10, 2018.




Figure 6.3: Spanwise average of flight data at 2:45 UT on July 8, 2018. Best-fit Gaussian profiles to each peak (red) and trough (blue) are shown. Image and caption are taken from Geach et al. [39].

locations gives an apparent phase speed of $c = 21 \pm 6 \text{ m s}^{-1}$. In each case, the quoted uncertainty represents the standard deviation across all measurements.

6.2.2 Power Spectrum Analysis

As described in Section 5.3.2, I interpolated the camera data onto a regular grid and found the power spectral density in two dimensions. The data, shown in Figure 6.5, exhibit a sharp peak corresponding to the GW I identified in the previous section.




Figure 6.4: Spanwise average of camera data as a function of time, from 2:40 to 3:00 UT on July 8, 2018. Red (blue) lines indicate linear fits to the locations of peaks (troughs) in the spanwise average data. Image and caption are taken from Geach et al. [39].




Figure 6.5: Portion of the spatial frequency power spectrum of the camera data from 2:45 UT on July 10, 2018, exhibiting a sharp peak that corresponds to the GW identified in the previous section. The white dashed line corresponds to the orientation of the GW found previously. Image are taken from Geach et al. [39].

I fit a 2D Gaussian profile to determine the location of this peak. Again taking the uncertainty in the location of the peak to be half a pixel, this analysis gives $\lambda_h = 17 \pm 1.7$ km with a propagation direction $123^\circ \pm 6^\circ$ clockwise from north. Both these values are consistent with the spanwise average results reported above, albeit with larger uncertainties.




Figure 6.6: Colorplot of the backscatter coefficient measured by the lidar, with the altitude of the brightest bin at each time interval (black dashed line).

6.2.3 Confirmation of Gravity Wave Identification in Lidar Data

To check for confirmation of this GW identification, I analyzed the altitude of the PMC layer as a function of time. Figure 6.6 shows the backscatter coefficient in the PMC layer as measured by the lidar in this interval. It also shows the PMC layer altitude, as determined from the altitude of the brightest bin in the backscatter profile at each time interval. Qualitatively, the lidar data shows a PMC layer with an average altitude of 82.3 km and a vertical oscillation in the PMC layer with a period of approximately 10 minutes.

Using the analysis described in Section 5.4, I found the location of the lidar beam at each time interval in the lidar data and determined the spanwise average of the image data at each of those locations. I removed an overall linear trend from this data, which is shown in Figure 6.7. Also shown is the lidar altitude and a best-fit sinusoidal to the lidar altitude data, with an amplitude of vertical displacements of $\delta_z = 0.3 \pm 0.1$ km. The comparison of these two measurements demonstrates that the periodic motion detected by the lidar correlates well with the brightness perturbations detected by the cameras over this interval. Noting that the lidar beam scans in the upstream direction and recalling the possible mechanisms for brightness perturbation (Section 5.2), Figure 6.7 demonstrates that the brightness perturbation




Figure 6.7: PMC layer altitude (red, left ordinate) and best-fit sinusoid (black dashed line), and spanwise-averaged brightness at the streamwise location of the lidar beam (blue, right ordinate) as a function of time. Image and caption are taken from Geach et al. [39].

and vertical displacement associated with the GW are closely aligned in phase, with the brightness perturbations slightly downstream from the vertical displacement. This suggests that the brightness perturbations are primarily caused by the GW-induced horizontal convergence at the PMC layer; furthermore, the small apparent phase shift is indicative of a large-amplitude GW.

6.2.4 Local and Background Wind Estimates

Applying the analysis pipeline described in Section 5.5.2 gives zonal and meridional winds averaged across the combined FOV of the cameras over the duration of the event. Figure 6.8 shows the *Trackpy* zonal and meridional winds at intervals of 2 minutes over this period, with error bars denoting the standard deviation within each interval. The average velocity over the entire interval is $U_{Trackpy} = \begin{pmatrix} -10 \pm 7 \\ -49 \pm 8 \end{pmatrix}$ where the uncertainty is taken as the quadrature sum of the standard deviation across all intervals and the average standard deviation within each interval.

As described in Section 5.5.3, we have independent wind estimates from NAVGEM at 2 and 3 UT, as shown in Figure 6.9. Interpolating between those profiles to 2:50 UT (the center of the interval studied here), the NAVGEM output gives wind estimates




Figure 6.8: *Trackpy* zonal (top) and meridional (bottom) velocity estimates for 2 minute intervals from 2:40 to 3:00 UT.

$$\text{of } U_{NAV\text{GEM}} = \begin{pmatrix} -28 \pm 9 \\ -28 \pm 4 \end{pmatrix}.$$

The difference between the *Trackpy*- and NAVGEM-derived wind estimates is $\begin{pmatrix} 18 \pm 11 \\ -21 \pm 9 \end{pmatrix}$, an apparent discrepancy of $> 2\sigma$. Given that we expect the *Trackpy* wind estimates to be influenced by local perturbations by the GW identified above, I take the $U_{NAV\text{GEM}}$ value to be the best-estimate for the background winds. I discuss the reason for this apparent discrepancy in greater detail in Section 6.4.1. From here on, I will refer to $U_{NAV\text{GEM}}$ as U_0 .

Temperature estimates from NAVGEM also allow calculation of the following parameters: $H_s = 4.4 \pm 0.1$ km, $N = 0.020 \pm 0.003$ rad s⁻¹, and $T_b = 310 \pm 50$ s. I estimate the uncertainty on these values from the change in temperature over a range of 2 km centered on the PMC altitude (using standard error propagation to convert uncertainty on temperature to uncertainty on these parameters).




Figure 6.9: NAVGEM profiles of zonal winds, meridional winds, temperature, and buoyancy frequency averaged horizontally over a 200 km radius around the location of the gondola during the GW breaking event, at 2 UT (red) and at 3 UT (blue). The gray horizontal bar indicates the approximate altitude of the PMC layer. Image and caption are taken from Geach et al. [39].

6.3 Vortex Ring Characterization

The relevant parameters to characterize the vortex rings are the diameter D of the rings prior to dissipation, the separation d between the rings, and the breakdown time of the vortex rings. A representative series of images showing the evolution of a pair of vortex rings appears in Figure 6.10.

I estimated D by manually fitting circles to a selection of vortex rings that I chose across the FOV; I found a range of vortex ring diameters from 2-5 km. Obviously, this approach is prone to selection effects and is likely biased towards larger vortex rings. Since the parameter of interest is the vortex ring diameter prior to breakdown, it is precisely the largest vortex rings present in the images that I want to characterize; therefore this selection bias is not problematic.

The series of images in Figure 6.10 shows the many vortex rings; in almost every case, the vortex rings are immediately adjacent to neighboring rings. Thus, $d \simeq D$ for the rings observed here.




Figure 6.10: Zoom of a section of a series of projected images showing the evolution of a pair of vortex rings oriented roughly N-S (white dotted circles). The image in panel (a) was taken at 2:41 UT, with subsequent images spaced at 80 s intervals. Vortex ring formation is evident in panel (b); bright spots at the north and south edges of the vortex ring are evident in panels (e) and (f), indicating trailing vortices linking adjacent rings. By panel (j), significant loss of definition has occurred in both rings, and the lower ring has diffused entirely. Also apparent in panel (a) is a single vortex ring (white dashed circle) that over the next four panels splits into several smaller rings. Image and caption are taken from Geach et al. [39].

The vortex rings shown in Figure 6.10 form by panel (b), start to lose definition in panel (f), and have lost all definition by panel (j). The frames are spaced at 80 s intervals, giving a total vortex ring lifetime of ~ 640 s and a breakdown time of ~ 320 s.

6.4 Interpretation and Discussion

The GW parameters derived in Section 6.2, the wind estimates described in Section 6.2.4, and the GW relationships derived in Section 5.1 allow for a complete characterization of the GW and background environment. The relevant parameters

and their provenance are listed in Table 6.1.

6.4.1 Momentum Deposition by the Gravity Wave

I interpret the apparent discrepancy between the *Trackpy* and NAVGEM-derived winds as being the result of momentum deposition by the GW as it breaks. The scale of the GW is well below the resolution of the NAVGEM analysis, so NAVGEM is not sensitive to perturbations on these scales. This implies a total momentum deposition per unit mass of

$$\Delta\vec{U} = U_{Trackpy} - U_0 = \begin{pmatrix} 18 \pm 11 \\ -21 \pm 9 \end{pmatrix} m s^{-1}. \quad (6.1)$$

The vertical momentum flux associated with a GW is given by

$$\tau = -\rho_0 \overline{u_1 w_1} \simeq \frac{1}{2} \rho_0 \frac{k}{m} u_1^2 \quad (6.2)$$

where the bar indicates a horizontal average over a full wavelength and I have used a simplifying assumption (borne out by the parameters in Table 6.1) that $\lambda_z \ll 4\pi H_s$.

At low amplitudes, this momentum flux is constant: the amplitude grows exponentially as the GW rises and density decreases. But when the GW becomes unstable and starts to break, momentum is no longer conserved. Instead, the GW deposits the momentum it carries into the mean flow at the altitude at which it breaks. Numerical simulation of GW breaking has suggested a typical timescale for GW breaking of T_i [55], and I approximate the vertical layer over which the momentum is deposited to be H_s . Thus we expect a net change in the local momentum per unit mass of

$$\Delta\vec{U} \simeq \frac{T_i}{H_s} \frac{\tau}{\rho_0} \simeq \frac{1}{2} \frac{k u_1^2 T_i}{m H_s} \quad (6.3)$$

Gravity Wave or Background Parameter	Symbol	Value	Source
Mean wind	\vec{U}_0	$\begin{pmatrix} -28 \pm 9 \\ -28 \pm 4 \end{pmatrix} \text{ m s}^{-1}$	NAVGEN
Buoyancy frequency	N	$0.020 \pm 0.003 \text{ rad s}^{-1}$	NAVGEN
Buoyancy period	T_b	$310 \pm 50 \text{ s}$	NAVGEN
Scale height	H_s	$4.4 \pm 0.1 \text{ km}$	NAVGEN
Propagation direction	\hat{k}_h	$\begin{pmatrix} 0.82 \pm 0.03 \\ -0.57 \pm 0.02 \end{pmatrix}$	PMC Turbo
Horizontal wavelength	λ_h	$18 \pm 0.5 \text{ km}$	PMC Turbo
Vertical wavelength	λ_z	$10 \pm 5 \text{ km}$	Linear theory
Apparent phase speed	c	$21 \pm 6 \text{ m s}^{-1}$	PMC Turbo
Intrinsic phase speed	c_i	$28 \pm 10 \text{ m s}^{-1}$	Linear theory
Intrinsic period	T_i	$650 \pm 230 \text{ s}$	Linear theory
Intrinsic frequency	Ω	$0.01 \pm 0.003 \text{ rad s}^{-1}$	Linear theory

Table 6.1: Inferred estimates of relevant GW and background parameters and their source. Values derived from linear theory are calculated from NAVGEN estimates and PMC Turbo observations.

Given that the GW is breaking, we can assume that it was close to its overturning amplitude. Thus, $u_1 \simeq c_i$ and $u_1 T_i \simeq \lambda_h$, and we have

$$\Delta \vec{U} \simeq \frac{1}{2} \frac{k \lambda_h}{m} \frac{c_i}{H_s} \simeq \frac{\lambda_z}{2 H_s} c_i. \quad (6.4)$$

Thus, a GW with the parameters listed in Table 6.1 can be expected to cause a net change in local winds of $\Delta \vec{U} \simeq \begin{pmatrix} 25 \pm 15 \\ -18 \pm 11 \end{pmatrix} \text{ m s}^{-1}$, which is consistent with the measured discrepancy in (6.1).

I tested this hypothesis by comparing the NAVGEN and *Trackpy* winds at 2 UT on July 10, 2018, prior to the onset of GW breaking. In this case, I found $U_{Trackpy} = \begin{pmatrix} -20 \pm 9 \\ -35 \pm 12 \end{pmatrix} \text{ m s}^{-1}$ and $U_{NAVGEN} = \begin{pmatrix} -13 \pm 9 \\ -35 \pm 4 \end{pmatrix} \text{ m s}^{-1}$. In this case, the two measurements are consistent; this supports the hypothesis that the discrepancy at

later times is due to local momentum deposition rather than some systematic error in the *Trackpy* or NAVGEM output.

The total body forcing due to GW breaking at the mesopause required to maintain the observed thermal structure is $100 \text{ m s}^{-1} \text{ day}^{-1}$ [56]. The direction of this forcing is eastward, in general agreement with the momentum deposition reported here. Taking only the eastward component of the inferred momentum deposition, this event constitutes $\sim 30\%$ of the expected daily total in the span of 20 minutes. Clearly this rate of momentum deposition is not representative; this event was analyzed specifically because there was evidence of strong GW breaking at these times. The magnitude of momentum deposition is significant, though, in that it suggests that the total body forcing may be made up of relatively few large events each day rather than a more continuous smaller amplitude forcing.

The hypothesis that the GW has undergone significant breaking is also consistent with the vertical displacements measured by the lidar, $\delta_z = 0.3 \pm 0.1 \text{ km}$. Recalling (5.18), this implies a GW amplitude of $a = 0.2 \pm 0.1$. Given that GW breaking is not expected as such low GW amplitudes, this implies that the GW had undergone significant amplitude reduction since the onset of vortex ring formation.

6.4.2 Comparison to Numerical Simulation

Given the vortex rings characteristics reported in Section 6.3, the numerical simulations by Fritts et al. [54] suggest a GW with $\lambda_z \simeq 10 \text{ km}$ and $\Omega \geq N/3$. In both cases, this is in good agreement with the GW parameters listed in Table 6.1. Furthermore, the time scale for vortex breakdown, as shown in Figure 6.10, is $\sim 320 \text{ s}$, in good agreement with the $\sim 1T_b$ prediction made by Fritts et al. [54].

Beyond these numerical relationships, Fritts et al. [54] also quantified the energy dissipation rate associated with GW breaking and vortex ring formation via $\langle \epsilon \rangle$ (the




Figure 6.11: (Top, left to right) Zoomed sections of the camera field of view, starting at 2:45:40 UT (left) and proceeding in 40 s intervals to the right, with the location of the lidar beam within the field of view (white dot). The images track the passage of a vortex ring through the lidar beam (over-sized white circles). (Bottom) The lidar profile of the GW breaking event, with blue arrows to indicate the corresponding times of each of the images above. Image and caption are taken from Geach et al. [39].

mean energy dissipation rate) and ϵ_{high} (the 99th percentile value over the domain). The results they report are dimensionless and need to be scaled by a parameter $C = \lambda_z^2/T_b^3$. Given the parameters listed in Table 6.1, this event has $C = 3.5_{-2.7}^{+5.9}$, yielding $\langle \epsilon \rangle$ between 0.03 and 0.73 W kg⁻¹ and ϵ_{high} between 0.3 and 7.3 W kg⁻¹. These values are similar to or higher than values that have been reported by past experiments that made measurements of the polar summer mesopause: Lübken et al. [57] reported values up to 2.4 W kg⁻¹, and Strelnikov et al. [58] reported values up to 2.0 W kg⁻¹. In both cases, measurements were made in situ by way of sounding rockets; as such, the measurements were sparse and highly localized. As pointed out by Fritts et al. [54], such experiments are unlikely to sample the regions in which the highest energy dissipation rates occur.

6.4.3 Vertical Structure of the Vortex Rings

This is the first instance of coincident lidar and image data for an observation of vortex rings at the PMC layer, allowing insight into the 3D structure of the vortex rings. In the lidar timestream, the vortex rings appear as sharp downward excursions of the lower edge of the PMC layer; see Figure 6.11. By correlating the lidar data to the images, we see that the downward excursions correspond to the bright edges of the rings. From 2:48 UT onward, the PMC layer is relatively undisturbed in the vicinity of the lidar beam, with sparse and faint vortex rings. In that section of data, the PMC layer has a more uniform thickness of ~ 1 km. This suggests that the appearance of vortex rings in the PMCs is primarily due to the motions at the core of the ring rather than the brightly visible fringe. The numerical simulations of Fritts et al. [54] suggest rings that are inclined at $\sim 45^\circ$ relative to the horizontal and have associated plunging motions along the axis of the vortex ring cores. Such motions would act to deplete the PMC layer at those locations, both by advecting PMC particles away from the cores and by accelerating sublimation as the particles are advected downward.

References

- [1] A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Physical Review D*, 23:347–356, 1981.
- [2] A. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389–393, 1981.
- [3] J. C. Mather, Fixsen D. J., R. A. Shafer, et al. Calibrator design for the COBE Far Infrared Absolute Spectrophotometer (FIRAS). *The Astrophysical Journal*, 512(2), 1999.
- [4] T. E. Montroy, P. A. R. Ade, J. J. Bock, et al. A measurement of the CMB $\langle EE \rangle$ spectrum from the 2003 flight of BOOMERANG. *Astrophysical Journal*, 647(2), 2006.
- [5] J. H. P. Wu, J. Zuntz, M. E. Abroe, et al. MAXIPOL: Data analysis and results. *Astrophysical Journal*, 665(1), 2007.
- [6] C. L. Bennett, D. Larson, J. L. Weiland, et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Final maps and results. *The Astrophysical Journal Supplement Series*, 208(2), 2013.
- [7] *Planck* collaboration. *Planck* 2015 results XIII: Cosmological parameters. *Astronomy & Astrophysics*, 594, 2016.

- [8] E. M Leitch, J. M. Kovac, N. W. Halverson, et al. Degree angular scale interferometer 3 year cosmic microwave background polarization results. *Astrophysical Journal*, 624(1), 2005.
- [9] The POLARBEAR collaboration. A measurement of the cosmic microwave background B-mode polarization power spectrum at sub-degree scales with POLARBEAR. *Astrophysical Journal*, 794(171), 2014.
- [10] Keck Array and BICEP2 collaborations. Constraints on primordial gravitational waves using *Planck*, WMAP, and new BICEP2/*Keck* observations through the 2015 season. *Physical Review Letters*, 121, 2018.
- [11] The EBEX collaboration. The EBEX balloon-borne experiment — Optics, receiver, and polarimetry. *The Astrophysical Journal Supplement Series*, 239, 2018.
- [12] The EBEX collaboration. The EBEX balloon-borne experiment — detectors and readout. *The Astrophysical Journal Supplement Series*, 239, 2018.
- [13] Kyle Zilic. *Calibration and Design of the E and B Experiment (EBEX) Cryogenic Receiver*. PhD thesis, University of Minnesota, Twin-Cities, 8 2014.
- [14] Jeffrey Klein. *Design, Implementation, and Calibration of Half-Wave Plate Polarimetry for the E and B Experiment*. PhD thesis, University of Minnesota, Twin-Cities, 10 2014.
- [15] Catherine Raach. *Characterization, Optimization, and Operation of TES Bolometers for EBEX*. PhD thesis, University of Minnesota, Twin-Cities, 4 2018.
- [16] Chaoyun Bao. *Foreground Cleaning for Cosmic Microwave Background Po-*

- larimeters in the Presence of Instrumental Effects*. PhD thesis, University of Minnesota, Twin-Cities, 6 2015.
- [17] D. C. Fritts and M. J. Alexander. Gravity dynamics and effects in the middle atmosphere. *Reviews of Geophysics*, 41, 2003.
- [18] M. J. Taylor, P.-D. Pautet, D. C. Fritts, et al. Large-amplitude mountain waves in the mesosphere observed on 21 june 2014 during DEEPWAVE: 1. wave development, scales, momentum fluxes, and environmental sensitivity. *Journal Geophysical Research - Atmospheres*, 124, 2019.
- [19] F.-J. Lübken, J. Lautenbach, J. Höffner, et al. First continuous temperature measurements within polar mesosphere summer echoes. *Journal of Atmospheric and Solar-Terrestrial Physics*, 71:453–463, 2009.
- [20] B. J. Murray and Jensen E. J. Homogeneous nucleation of amorphous solid water particles in the upper mesosphere. *Journal of Atmospheric and Solar-Terrestrial Physics*, 72:51–61, 2010.
- [21] NASA. Polar mesospheric clouds, south pacific ocean. <https://earthobservatory.nasa.gov/images/80217/polar-mesospheric-clouds-south-pacific-ocean>, 2013. Accessed: 2020-12-04.
- [22] C. O. Hines. Internal atmospheric gravity waves at ionospheric heights. *Canadian Journal of Physics*, 38:1441–1481, 1960.
- [23] Aaron Smith. Ebert-fastie spectral response measurements and simulation. Master’s thesis, University of Minnesota, Twin-Cities, 5 2014.
- [24] Rowland William Cobb. A high-resolution infra-red ebert spectrometer. Master’s thesis, University of British Columbia, 8 1961.

- [25] PCGrate. S6.1, 2013.
- [26] E. G. Loewen, M. Nevière, and D. Maystre. Grating efficiency theory as it applies to blazed and holographic gratings. *Applied Optics*, 16(10), 1977.
- [27] J. Didier, A. D Miller, D. Araujo, et al. Intensity-coupled polarization in instruments with a continuously rotating half-wave plate. *The Astrophysical Journal*, 876, 2019.
- [28] T. Renbarger, J. L. Dotson, and G. Novak. Measurements of submillimeter polarization induced by oblique reflection from aluminum alloy. *Applied Optics*, 37, 1998.
- [29] S. Takakura, M. Aguilar, Y. Akiba, et al. Performance of a continuously rotating half-wave plate on the POLARBEAR telescope. *Journal of Cosmology and Astroparticle Physics*, 2017, 2017.
- [30] B. Kaifler, D. Rempel, P. Roßi, et al. A technical description of the balloon lidar experiment BOLIDE. *Atmospheric Measurement Techniques*, 13(10), 2020.
- [31] Hongfei Zheng. *Solar Energy Desalination Technology*. Elsevier, 2017.
- [32] NREL. Reference Air Mass 1.5 G Spectra. <https://www.nrel.gov/grid/solar-resource/spectra-am1.5.html>, 2020. Accessed: 2020-10-23.
- [33] CSBF. Gondola Structural Design Requirements. <https://www.csbf.nasa.gov/documents/gondola/820-PG-8700.0.1%20Gondola%20Structural%20Design%20Requirements.pdf>, 2019. Accessed: 2020-10-23.
- [34] J. D. Soler, P. A. R. Ade, F. E. Angilè, et al. Thermal design and performance of the balloon-borne large aperture submillimeter telescope for polarimetry BLAST-

- Pol. In L. M. Stepp, R. Gilmozzi, and H. J. Hall, editors, *Ground-based and Airborne Telescopes V*, volume 9145 of *Proceedings of SPIE*, page 34, 2014.
- [35] C. B. Kjellstrand, G. Jones, C. Geach, et al. The PMC Turbo balloon mission to measure gravitywaves and turbulence in polar mesospheric clouds: Camera, telemetry, and software performance. *Earth and Space Science*, 7(8), 2020. doi: <https://doi-org.ezp3.lib.umn.edu/10.1029/2020EA001238>.
- [36] D. Lang, D. Hogg, K. Mierle, et al. Astrometry.net: Blind astrometric calibration of arbitrary astronomical images. *The Astronomical Journal*, 139(5):1782–1800, 2010. doi: <https://doi-org.ezp3.lib.umn.edu/10.1088/0004-6256/139/5/1782>.
- [37] Daniel Chapman. *EBEX: A Balloon-Borne Telescope for Measuring Cosmic Microwave Background Polarization*. PhD thesis, Columbia University, 2 2015.
- [38] D. C. Fritts et al. Nasa 2015 proposal for the PMC Turbo mission. LCAS proposal, 2015.
- [39] C. Geach, S. Hanany, D. Fritts, et al. Gravity wave breaking and vortex ring formation observed by PMC Turbo. *Journal Geophysical Research - Atmospheres*, 2020. Accepted but unpublished as of November 23, 2020.
- [40] C. Nappo. *An Introduction to Atmospheric Gravity Waves*. Elsevier, 2002.
- [41] G. I. Taylor. Effect of variation in density on the stability of superposed streams of fluid. *Proceedings of the Royal Society*, pages 499–523, 1931.
- [42] S. Goldstein. On the stability of superposed streams of fluids of different densities. *Proceedings of the Royal Society*, pages 524–548, 1931.
- [43] M.J Taylor. A review of advances in imaging techniques for measuring short period gravity waves in the mesosphere and lower thermosphere. *Advances in Space*

- Research*, 19(4), 1997. ISSN 0273-1177. doi: [https://doi.org/10.1016/S0273-1177\(97\)00161-0](https://doi.org/10.1016/S0273-1177(97)00161-0).
- [44] A. Chandran, D.W. Rusch, G.E. Thomas, et al. Atmospheric gravity wave effects on polar mesospheric clouds: A comparison of numerical simulations from CARMA 2D with AIM observations. *Journal Geophysical Research - Atmospheres*, 117(D20), 2012. doi: <http://dx.doi.org/10.1029/2012JD017794>.
- [45] A.Y. Zasetsky, S.V. Petelina, R. Remorev, et al. Ice particle growth in the polar summer mesosphere: Formation time and equilibrium size. *Geophysical Research Letters*, 36, 2009. doi: <https://doi.org/10.1029/2009GL038727>.
- [46] M. Gadsden and W. Schröder. *Noctilucent clouds*. Springer-Verlag, 1989.
- [47] M. Rapp, F.-J. Lübken, A. Müllemann, et al. Small-scale temperature variations in the vicinity of NLC: Experimental and model results. *Journal Geophysical Research - Atmospheres*, 107(D19), 2002.
- [48] O.S. Ugolnikov, A.A. Galkin, S.V. Pilgaev, et al. Noctilucent cloud particle size determination based on multi-wavelength all-sky analysis. *Planetary and Space Science*, 146(15), 2017.
- [49] C. Savigny and C. Hoffman. Issues related to the retrieval of stratospheric-aerosol particle size information based on optical measurements. *Atmospheric Measurement Techniques*, 13(4), 2020.
- [50] Trackpy. v04.1, 2019. <https://doi.org/10.5281/zenodo.3492186>.
- [51] S. Eckermann, J. Ma, K. Hoppel, et al. High-altitude (0–100 km) global atmospheric reanalysis system: Description and application to the 2014 austral winter

- of the Deep Propagating Gravity Wave Experiment (DEEPWAVE). *Atmospheric Measurement Techniques*, 13(4), 2018.
- [52] Ø Andreassen, P. Ø. Hvidsten, D. C. Fritts, et al. Vorticity dynamics in a breaking internal gravity wave: 1. initial instability evolution. *Journal of Fluid Dynamics*, 367:27–46, 1998.
- [53] D. C. Fritts, S. Arendt, Ø Andreassen, et al. Vorticity dynamics in a breaking internal gravity wave: 2. vortex interactions and transition to turbulence. *Journal of Fluid Dynamics*, 367:46–65, 1998.
- [54] D. C. Fritts, L. Wang, G. Baumgarten, et al. High-resolution observations and modeling of turbulence sources, structures, and intensities in the upper mesosphere. *Journal of Atmospheric and Solar-Terrestrial Physics*, 162:57–78, 2017.
- [55] D. C. Fritts, L. Wang, J. Werne, et al. Gravity wave instability dynamics at high reynolds numbers: 1. wavefield evolution at large amplitudes and high frequencies. *Journal of Atmospheric Science*, 66(5), 2009.
- [56] R. Vincent. The dynamics of the mesosphere and lower thermosphere: A brief review. *Progress in Earth and Planetary Science*, 2(4), 2015.
- [57] F.-J. Lübken, Rapp. M, and P Hoffman. Neutral air turbulence and temperatures in the vicinity of polar mesospheric summer echoes. *Journal of Geophysical Research: Atmospheres*, 107(D15), 2002.
- [58] B. Strelnikov, M. Rapp, I. Strelnikova, et al. Small-scale structures in neutrals and charged aerosol particles as observed during the ECOMA/MASS rocket campaign. *Annales Geophysicae*, 27:1,449–1,456, 2009.

- [59] W. Young and R. Budynas. *Roark's Formulas for Stress and Strain*. McGraw-Hill, 2002.

Appendix A

PMC Turbo Stress Report

Here I reproduce the stress report I submitted to CSBF to demonstrate that the PMC Turbo gondola met the design criteria to be certified for flight. Professor Frank Kelso provided guidance, and Charlie Houtz performed an independent analysis to verify the results. I found an oversimplification in the assumptions I used to model the suspension system, which I have corrected for here. The margins of safety have changed as a result, but the conclusion of the stress report, namely that the PMC Turbo gondola was sufficiently strong to satisfy the CSBF design criteria, did not change. The stress report has also been reformatted to fit the style guidelines of this document.

The following subsystems of the PMC Turbo gondola were analyzed:

- The suspension system, starting with the NASA-supplied lower triplate (D27505 - Lower Rotator Adapter Fitting). This is rigidly attached to a spreader bar, from which the gondola hangs via 4 rigid members and 4 steel wire cables. The 4 rigid members are analyzed separately for failure in buckling and in tension.
- The gondola cage. Simplified mass models are used for the more complicated elements of the gondola.
- Ballast attachment points.

- Middle floor.
- Battery floor.
- Radiator.
- Solar array.
- Two cameras on a C-channel.

The total weight of the PMC Turbo gondola is 1860 lbs., excluding ballast, bolts (estimated conservatively at 40 lbs.), antennae, the suspension system cables, and some minimal mylar shielding. To take a conservative approach, I assume 1200 lbs of ballast and add 300 lbs. of additional margin to cover late modifications. I take the total suspended weight to be 3500 lbs.

All the principal joints are bolted. A few mounting brackets are welded aluminum - these parts are analyzed separately, with a de-rating of their ultimate strength due to the welding.

A.1 Gondola Drawings




Figure A.1: PMC Turbo gondola with annotations.




Figure A.2: PMC Turbo gondola with dimensions.




Figure A.3: PMC Turbo gondola without radiator and solar arrays, with dimensions.

A.2 Margins of Safety Tables

Subsystem Name	Load Case	Material Type	Ultimate Tensile Strength	Maximum Stress	Margin of Safety
Suspension System	10g	Al 6061-T6	37 ksi	33.5 ksi	0.10
Suspension System	5g at 45° backwards ¹	Al 6061-T6	42 ksi ²	24.8 ksi	0.69
Suspension System	5g at 45° left/right	4140 steel	60 ksi	35.9 ksi	0.67
Suspension System	5g at 45° left/right	Al 6061-T6	37 ksi	30.6 ksi	0.21
Gondola Cage	10g	Al 6061-T6	37 ksi	20.1 ksi	0.84
Gondola Cage	5g at 45° forwards	Al 6061-T6	37 ksi	18.3 ksi	1.0
Gondola Cage	5g at 45° backwards	Al 6061-T6	37 ksi	24.2 ksi	0.53
Gondola Cage	5g at 45° left	Al 6061-T6	37 ksi	31.0 ksi	0.19
Gondola Cage	5g at 45° right	Al 6061-T6	37 ksi	27.2 ksi	0.36
Gondola Cage	5g at 90° forwards	Al 6061-T6	37 ksi	27.7 ksi	0.34
Gondola Cage	5g at 90° backwards	Al 6061-T6	37 ksi	27.7 ksi	0.34
Gondola Cage	5g at 90° left	Al 6061-T6	37 ksi	32.8 ksi	0.13
Gondola Cage	5g at 90° right	Al 6061-T6	37 ksi	32.8 ksi	0.13
Ballast Attachment	10g	4140 Steel	60 ksi	30.0 ksi	1.0
Ballast Attachment	10g	Al 6061-T6	37 ksi	16.2 ksi	1.3
Ballast Attachment	5g at 45° backwards ¹	4140 Steel	60 ksi	37.8 ksi ³	0.59

Subsystem Name	Load Case	Material Type	Ultimate Tensile Strength	Maximum Stress	Margin of Safety
Ballast Attachment	5g at 45° backwards ¹	Al 6061-T6	37 ksi	16.5 ksi	1.2
Ballast Attachment	5g at 45° left ¹	4140 Steel	60 ksi	37.8 ksi ³	0.59
Ballast Attachment	5g at 45° left ¹	Al 6061-T6	37 ksi	13.0 ksi	1.8
Ballast Attachment	5g at 90° backwards ¹	4140 Steel	60 ksi	47.1 ksi ³	0.27
Ballast Attachment	5g at 90° backwards ¹	Al 6061-T6	37 ksi	19.4 ksi	0.91
Ballast Attachment	5g at 90° left ¹	4140 Steel	60 ksi	30.0 ksi ³	1.0
Ballast Attachment	5g at 90° left ¹	Al 6061-T6	37 ksi	16.9 ksi	1.2
Middle Floor	10g	Al 6061-T6	37 ksi	12.1 ksi	2.1
Middle Floor	5g at 45° forwards	Al 6061-T6	37 ksi	5.5 ksi	5.7
Middle Floor	5g at 45° backwards	Al 6061-T6	37 ksi	7.9 ksi	3.7
Middle Floor	5g at 45° left	Al 6061-T6	37 ksi	11.2 ksi	2.3
Middle Floor	5g at 45° right	Al 6061-T6	37 ksi	9.9 ksi	2.7
Middle Floor	5g at 90° forwards	Al 6061-T6	37 ksi	5.4 ksi	5.9
Middle Floor	5g at 90° backwards	Al 6061-T6	37 ksi	5.4 ksi	5.9
Middle Floor	5g at 90° left	Al 6061-T6	37 ksi	12.0 ksi	2.1
Middle Floor	5g at 90° right	Al 6061-T6	37 ksi	12.0 ksi	2.1
Middle Floor (Welded parts)	10g	Al 6061-T6	12 ksi	8.5 ksi	0.41

Subsystem Name	Load Case	Material Type	Ultimate Tensile Strength	Maximum Stress	Margin of Safety
Middle Floor (Welded parts)	5g at 45° forwards	Al 6061-T6	12 ksi	5.4 ksi	1.2
Middle Floor (Welded parts)	5g at 45° backwards	Al 6061-T6	12 ksi	5.4 ksi	1.2
Middle Floor (Welded parts)	5g at 45° left	Al 6061-T6	12 ksi	9.1 ksi	0.32
Middle Floor (Welded parts)	5g at 45° right	Al 6061-T6	12 ksi	8.3 ksi	0.45
Middle Floor (Welded parts)	5g at 90° forwards	Al 6061-T6	12 ksi	5.4 ksi	1.2
Middle Floor (Welded parts)	5g at 90° backwards	Al 6061-T6	12 ksi	5.4 ksi	1.2
Middle Floor (Welded parts)	5g at 90° left	Al 6061-T6	12 ksi	7.2 ksi	0.67
Middle Floor (Welded parts)	5g at 90° right	Al 6061-T6	12 ksi	7.2 ksi	0.67
Battery Floor	10g	Al 6061-T6	37 ksi	27.5 ksi	0.34
Battery Floor	5g at 45° forwards	Al 6061-T6	37 ksi	13.2 ksi	1.8
Battery Floor	5g at 45° backwards	Al 6061-T6	37 ksi	13.4 ksi	1.8
Battery Floor	5g at 45° left	Al 6061-T6	37 ksi	12.7 ksi	1.9
Battery Floor	5g at 45° right	Al 6061-T6	37 ksi	7.6 ksi	3.9
Battery Floor	5g at 90° forwards	Al 6061-T6	37 ksi	8.8 ksi	3.2
Battery Floor	5g at 90° backwards	Al 6061-T6	37 ksi	8.8 ksi	3.2
Battery Floor	5g at 90° left	Al 6061-T6	37 ksi	4.3 ksi	7.6
Battery Floor	5g at 90° right	Al 6061-T6	37 ksi	4.3 ksi	7.6
Radiator	10g	Al 6061-T6	37 ksi	8.1 ksi	3.6

Subsystem Name	Load Case	Material Type	Ultimate Tensile Strength	Maximum Stress	Margin of Safety
Radiator	5g at 45° forwards	Al 6061-T6	37 ksi	5.7 ksi	5.5
Radiator	5g at 45° backwards	Al 6061-T6	37 ksi	2.1 ksi	17
Radiator	5g at 45° left	Al 6061-T6	37 ksi	6.7 ksi	4.5
Radiator	5g at 45° right	Al 6061-T6	37 ksi	6.7 ksi	4.5
Radiator	5g at 90° forwards	Al 6061-T6	37 ksi	4.0 ksi	8.3
Radiator	5g at 90° backwards	Al 6061-T6	37 ksi	4.0 ksi	8.3
Radiator	5g at 90° left	Al 6061-T6	37 ksi	9.5 ksi	2.9
Radiator	5g at 90° right	Al 6061-T6	37 ksi	9.5 ksi	2.9
Radiator (Welded parts)	10g	Al 6061-T6	12 ksi	4.2 ksi	1.9
Radiator (Welded parts)	5g at 45° forwards	Al 6061-T6	12 ksi	1.2 ksi	9.0
Radiator (Welded parts)	5g at 45° backwards	Al 6061-T6	12 ksi	2.5 ksi	3.8
Radiator (Welded parts)	5g at 45° left	Al 6061-T6	12 ksi	2.3 ksi	4.2
Radiator (Welded parts)	5g at 45° right	Al 6061-T6	12 ksi	2.1 ksi	4.7
Radiator (Welded parts)	5g at 90° forwards	Al 6061-T6	12 ksi	1.7 ksi	6.1
Radiator (Welded parts)	5g at 90° backwards	Al 6061-T6	12 ksi	1.7 ksi	6.1
Radiator (Welded parts)	5g at 90° left	Al 6061-T6	12 ksi	1.4 ksi	7.6
Radiator (Welded parts)	5g at 90° right	Al 6061-T6	12 ksi	1.4 ksi	7.6

Subsystem Name	Load Case	Material Type	Ultimate Tensile Strength	Maximum Stress	Margin of Safety
Solar Array	10g	Al 6061-T6	37 ksi	31.1 ksi	0.18
Solar Array	5g at 45° forwards	Al 6061-T6	37 ksi	24.0 ksi	0.54
Solar Array	5g at 45° backwards	Al 6061-T6	37 ksi	24.2 ksi	0.53
Solar Array	5g at 45° left	Al 6061-T6	37 ksi	9.1 ksi	3.1
Solar Array	5g at 45° right	Al 6061-T6	37 ksi	8.8 ksi	3.2
Solar Array	5g at 90° forwards	Al 6061-T6	37 ksi	33.5 ksi	0.10
Solar Array	5g at 90° backwards	Al 6061-T6	37 ksi	33.5 ksi	0.10
Solar Array	5g at 90° left	Al 6061-T6	37 ksi	3.0 ksi	11.3
Solar Array	5g at 90° right	Al 6061-T6	37 ksi	3.0 ksi	11.3
Two Cameras and Mounts	10g	Al 6061-T6	37 ksi	9.2 ksi	3.0
Two Cameras and Mounts	5g at 45° forwards	Al 6061-T6	37 ksi	2.3 ksi	15
Two Cameras and Mounts	5g at 45° backwards	Al 6061-T6	37 ksi	2.3 ksi	15
Two Cameras and Mounts	5g at 45° left	Al 6061-T6	37 ksi	1.6 ksi	22
Two Cameras and Mounts	5g at 45° right	Al 6061-T6	37 ksi	3.0 ksi	11
Two Cameras and Mounts	5g at 90° forwards	Al 6061-T6	37 ksi	2.7 ksi	13
Two Cameras and Mounts	5g at 90° backwards	Al 6061-T6	37 ksi	2.7 ksi	13
Two Cameras and Mounts	5g at 90° left	Al 6061-T6	37 ksi	1.9 ksi	18

Subsystem Name	Load Case	Material Type	Ultimate Tensile Strength	Maximum Stress	Margin of Safety
Two Cameras and Mounts	5g at 90° right	Al 6061-T6	37 ksi	1.9 ksi	18
Rod Legs (Tension)	10g	1018 Steel	55 ksi	33.2 ksi	0.66
Rod Legs (Tension)	5g at 45° forwards ¹	1018 Steel	55 ksi	39.5 ksi	0.39
Rod Legs (Tension)	5g at 45° left ¹	1018 Steel	55 ksi	0.8 ksi	67.8
Rod Legs (Compression)	10g	1018 Steel	6,260 lbs ⁴	N/A	N/A
Rod Legs (Compression)	5g at 45° forwards ¹	1018 Steel	6,260 lbs	4,550 lbs	0.38
Rod Legs (Compression)	5g at 45° left ¹	1018 Steel	6,260 lbs	1,750 lbs	2.6

Table A.1: Minimum margin of safety for each subsystem of the PMC Turbo gondola.

Hardware Component	Load Case	Ultimate Load	Maximum Load	Margin of Safety
Clevis Bracket ⁵	10g	54,400 lbs	8,840 lbs	5.2
Clevis Bracket	5g at 45° forwards/backwards	54,400 lbs	10,520 lbs	4.2
Clevis Bracket	5g at 45° left/right	54,400 lbs	1,750 lbs	30
Clevis Pin ⁶	10g	77,200 lbs	8,840 lbs	7.7
Clevis Pin	5g at 45° forwards/backwards	77,200 lbs	10,520 lbs	6.3

¹These test were only performed in backwards and left directions – forwards and right are equivalent, by symmetry.

²42 ksi is used instead of 37 ksi because the maximum stress occurs in the NASA triplate, which has higher ultimate strength.

³A stress singularity was observed. The values here are taken within 0.1” of the singularity. Some minimal plastic deformation may occur.

⁴For the rod legs under compression, we determine the buckling load rather than the ultimate strength.

Hardware Component	Load Case	Ultimate Load	Maximum Load	Margin of Safety
Clevis Pin	5g at 45° left/right	77,200 lbs	1,750 lbs	43
Spherical Ball Joint ⁷	10g	28,050 lbs	8,840 lbs	2.2
Spherical Ball Joint	5g at 45° forwards/backwards	28,050 lbs	10,520 lbs	1.7
Spherical Ball Joint	5g at 45° left/right	28,050 lbs	1,750 lbs	15
Small Hoist Ring ⁸	5g at 45° left/right	20,000 lbs	9,600 lbs	1.1
Large Hoist Ring ⁹	5g at 45° left/right	35,000 lbs	18,340 lbs	0.91
Turnbuckle ¹⁰	5g at 45° left/right	26,000 lbs	9,600 lbs	1.7
Eye Bolt ¹¹	10g	7,800 lbs	4,250 lbs	0.84
Eye Bolt	5g at 45° forwards/backwards	7,800 lbs	3,730 lbs	1.1
Eye Bolt	5g at 45° left/right	7,800 lbs	3,250 lbs	1.4

Table A.2: Minimum margin of safety for the COTS components of the PMC Turbo gondola. For each hardware component, only load cases for which that component experiences an applied load are shown.

⁵Miller forged steel clevis bracket, part number 0960160075

⁶Miller pivot pin, part number 0683690000

⁷Springfix Linkages rod end ball joint, part number HIM750

⁸Crosby UNC swivel hoist ring, part number HR-125-1016924

⁹Crosby UNC swivel hoist ring, part number HR-125-1016946

¹⁰Crosby eye & eye turnbuckle, part number HG-226-1031430

¹¹Crosby UNC shoulder type machinery eye bolt, part number S-279-9900226




Figure A.4: Suspension system of the PMC Turbo gondola, showing the 4 rigid members and the 4 tensioning cables. The red arrow shows the attachment point to the flight line, which I set as the origin of the coordinate system used here.

A.3 Suspension System Hand Calculations

The suspension system consists of 4 rigid members with ball joints which support the gondola under vertical loading and front-back loading, and 4 cables that support the gondola under left-right loading. We calculate the resulting loads in each load case by hand. The coordinate system is shown in the above figure: the x-axis is to the right, the y-axis is into the page, and the z-axis is vertical. The rigid members have unit vectors $\begin{pmatrix} 0 \\ \pm 0.393 \\ 0.919 \end{pmatrix}$ and the cables have unit vectors $\begin{pmatrix} \pm 0.599 \\ \pm 0.293 \\ 0.745 \end{pmatrix}$. If we set the origin of our coordinate system at the center of the attachment point to flight train (see red arrow in the above figure), then the load from each rigid member

is applied at $\begin{pmatrix} \pm 27.5 \\ \pm 1.75 \\ -15.38 \end{pmatrix}$, and the load from each cable is applied at $\begin{pmatrix} \pm 4.5 \\ 0 \\ -14.35 \end{pmatrix}$, with all coordinates given in inches. These coordinates will become important when calculating the net torque applied to the suspension system.

A.3.1 Load Cases

The load cases tested for here are 10g applied vertically, and 5g applied at 45° to vertical in the 4 principal directions of the gondola. These loads are applied at the attachment point to the flight train (see red arrow in the above figure) and are transmitted to the gondola through the suspension members. In calculating the resulting loads in each of the suspension members, we take into account the weight above the suspension system: the spreader bar, the upper hoist rings and clevis brackets, the antenna boom and antennae (not shown). The total mass of the gondola with ballast is no more than 3,500 lbs; the mass above the suspension members is 250 lbs. Thus, in the 10g case, a load of 35,000 lbs is applied vertically at the attachment point to the flight train. The effective load transmitted to the gondola is 32,500 lbs. Similarly, in the 5g cases, the load applied at the attachment point to the flight train is 17,500 lbs, with an effective load of 16,250 lbs transmitted to the gondola. In the calculations below, we calculate the resulting loads in the suspension members from the applied loads at the attachment point to the flight train. In these calculations we use the effective load (L_{eff}) rather than the total applied load.

10g Applied Vertically

In this case, the rigid rods come under tension. By symmetry, the four rods are tensioned equally and no net torque is applied to the suspension system. Thus, we




Figure A.5: Suspension system of the PMC Turbo gondola, illustrating the ‘10g applied vertically’ load case.

have

$$L_{eff} = 4T_z = 4(0.919T). \quad (\text{A.1})$$

In the 10g case, $L_{eff} = 32,500$ lbs and thus $T = 8,840$ lbs.

5g Applied at 45° Forwards/Backwards

In this case, the front rigid rods come under tension, and the back rods come under compression. The suspension system is free to rotate about its axis (normal to the page in the above figure); therefore, I find the rotation of the spreader bar and the




Figure A.6: Suspension system of the PMC Turbo gondola, illustrating the ‘5g at 45° forwards/backwards’ load case.

resulting loads at that orientation. I do this calculation iteratively: assuming no rotation, I calculate the resulting loads in each of the rigid members. This results in a net torque on the suspension system. I calculate the rotation of the suspension system that eliminates this torque. This results in a new orientation of the rigid members, with new implied loads in each of the rigid members. I repeat this process until of successive iterations result in changes of less than 10 lbs in the resulting loads.

In this case, $L_{eff} = 16,250$ lbs, and $L_y = L_z = 11,490$ lbs. To calculate T and C given L , we first need the orientation of the front and back rigid members with respect to the vertical (θ_T and θ_C respectively). Initially, we have $\theta_T = \theta_C = 23.22^\circ$.

I decompose the forces into their y - and z -components,

$$L_y = 2C_y + 2T_y = 2C \sin \theta_C + 2T \sin \theta_T \quad (\text{A.2})$$

$$L_z = 2T_z - 2C_z = 2T \cos \theta_T - 2C \cos \theta_C, \quad (\text{A.3})$$

and solve this system of equations for C and T :

$$C = \left(\frac{1 - \tan \theta_T}{\sin \theta_C + \cos \theta_C \tan \theta_T} \right) \frac{L_y}{2} \quad (\text{A.4})$$

$$T = \left(\frac{\sin \theta_C + \cos \theta_C}{\cos \theta_T - \sin \theta_T} \right) C \quad (\text{A.5})$$

The first iteration gives $C = 4,160$ lbs and $T = 10,410$ lbs.

To calculate the implied rotation of the suspension system, we consider the torque applied by the rigid members about the attachment point to the flight train. We define d to be the distance between the attachment points and this pivot; τ_T and τ_C to be the torque applied by the front and back rods respectively; $\theta_0 = 6.49^\circ$ to be the angle between the initial lever arm at which torque is applied and the vertical (see figure below); and $\Delta\theta$ to be the rotation of the suspension system.

These torques can be written

$$\tau_T = Td \sin(\theta_T - \theta_0 - \Delta\theta) \quad (\text{A.6})$$

$$\tau_C = Cd \sin(\theta_C - \theta_0 + \Delta\theta) \quad (\text{A.7})$$

The suspension system will rotate such that the net torque is zero, i.e. $\tau_T = -\tau_C$.

This is satisfied when

$$\Delta\theta = \arctan \left[\frac{T \sin(\theta_T - \theta_0) - C \sin(\theta_C - \theta_0)}{T \cos(\theta_T - \theta_0) + C \cos(\theta_C - \theta_0)} \right] \quad (\text{A.8})$$




Figure A.7: Suspension system of the PMC Turbo gondola, illustrating the rotation imparted to the spreader bar under the ‘5g at 45° forwards/backwards’ load case.

In a coordinate system fixed to the gondola, the lower attachment points of each pair of rigid members are fixed a distance l apart, the length of each rod, r , is fixed, and the distance between the upper attachment points, d , is fixed (see figure below). The quadrilateral formed by these four sides is not totally constrained; the rigid members have freedom to rotate. This rotation is limited by the geometry of the attachment points to be approximately $-90^\circ \leq \Delta\theta \leq 90^\circ$.

Given $\Delta\theta$ from the calculation above, the other angles are constrained. We define a coordinate system with its origin at the base of the back rigid member (at B_1 in Figure A.8), with the x -axis to the right and the y -axis upwards. The other points (T_1 at the top of the rear rigid member, T_2 and B_2 at the top and bottom, respectively, of the front rigid member) have the following coordinates:

$$B_2 = (l, 0) \tag{A.9}$$




Figure A.8: Schematic of the suspension system showing the fixed geometrical constraints affecting rotation of the spreader bar under the ‘5g at 45° forwards/backwards’ load case.

$$T_1 = (r \sin \theta_C, r \cos \theta_C) \quad (\text{A.10})$$

$$T_2 = (l - r \sin \theta_T, r \cos \theta_T) \quad (\text{A.11})$$

Given the fixed separation between T_1 and T_2 , we also have

$$T_2 = T_1 + (d \cos \Delta\theta, d \sin \Delta\theta). \quad (\text{A.12})$$

This gives the following system of equations

$$r \sin \theta_C + d \cos \Delta\theta = l - r \sin \theta_T \quad (\text{A.13})$$

$$r \cos \theta_C + d \sin \Delta\theta = r \cos \theta_T \quad (\text{A.14})$$

which can be solved for θ_C and θ_T .

$$\sin \theta_C + \sin \theta_T = \frac{l - d \cos \Delta\theta}{r} = 2 \sin \frac{\theta_C + \theta_T}{2} \cos \frac{\theta_C - \theta_T}{2} \quad (\text{A.15})$$

$$\cos \theta_C - \cos \theta_T = \frac{d \sin \Delta\theta}{r} = 2 \sin \frac{\theta_C + \theta_T}{2} \sin \frac{\theta_C - \theta_T}{2} \quad (\text{A.16})$$

I now define θ' as

$$\theta' \equiv \frac{\theta_C - \theta_T}{2} = \arctan \left(\frac{d \sin \Delta\theta}{l - d \cos \Delta\theta} \right), \quad (\text{A.17})$$

and we can write

$$\theta_T = \arcsin \left(\frac{d \sin \Delta\theta}{2r \sin \theta'} \right) - \theta' \quad (\text{A.18})$$

$$\theta_C = \theta_T + 2\theta' \quad (\text{A.19})$$

I then use these new values for θ_T and θ_C in (A.4) and (A.5) to obtain a new approximation for C and T . Iterating the above process gives a self-consistent solutions for all the relevant variables, with final values for the resulting loads of $T = 10,520$ lbs and $C = 4,550$ lbs, applied at angles of 20.9° and 26.1° , respectively, with respect to the vertical. This corresponds to a rotation of the suspension system by 35.0° .

5g Applied at 45° to the Left/Right

In Figure A.9 and in the equations below, I assume that the rigid members come under tension (this does not impose any constraints on the results: a negative tension implies compression). The cables in the opposite direction to the applied load will also come under tension. In this case, the suspension system does not have freedom of rotation; the net torque and force on the system will therefore be zero. By front-to-back symmetry, the tension and compression in the front and back members will be equal, and all loads will therefore be lie in the same plane, parallel to the plane of the page in Figure A.9. As in the previous case, $L_{eff} = 16,250$ lbs, with $L_x = L_z = 11,490$ lbs. I define T_1 , T_2 , and T_3 to be the net forces at each attachment point as shown in the




Figure A.9: Suspension system of the PMC Turbo gondola, illustrating the ‘5g at 45° to the left/right’ load case.

figure above. In terms of the tension in the individual suspension members,

$$T_1 = 1.91T_C \quad (\text{A.20})$$

$$T_2 = 1.84T_{R_l} \quad (\text{A.21})$$

$$T_3 = 1.84T_{R_r} \quad (\text{A.22})$$

where T_C is the tension in each cable, and T_{R_l} and T_{R_r} are the tension (or compression) in the individual left and right rods, respectively.

I decompose the loads into their x - and z -components:

$$L_x = T_1 \sin \theta_1 \quad (\text{A.23})$$

$$L_z = T_1 \cos \theta_1 + T_2 + T_3 \quad (\text{A.24})$$

where $\theta_1=38.8^\circ$ is the angle T_1 makes with the vertical. I then calculate the torque

about the attachment point to the flight train. The cables apply a net force T_1 at a distance $r_C=15.0''$ from the pivot, with an angle $\theta_C=21.4^\circ$ between the applied force and the lever arm. The rods apply a net force $T_2 - T_3$ at a distance $r_R = 31.5''$ from the pivot, with an angle $\theta_R=60.8^\circ$ between the applied force and the lever arm. Defining a positive torque to be counter-clockwise in Figure A.9, the net torque is

$$\Sigma\tau = T_1 r_C \sin \theta_C + (T_2 - T_3) r_R \sin \theta_R. \quad (\text{A.25})$$

The system of equations, (A.23), (A.24), and (A.25), can be solved for T_1 , T_2 , and T_3 , and therefore for T_C , T_{R_l} , and T_{R_r} , giving $T_C = 9,600$ lbs, $T_{R_l} = -1,750$ lbs, and $T_{R_r} = 230$ lbs. Thus, the rods on the left in Figure A.9 will be under compression, while the rods on the right will be under minimal tension.

To test the assumption made in this analysis that the suspension system will remain stationary under these applied loads, I modelled the cables as being $l = 54''$ in length and $0.375''$ in diameter (this is conservative, as much of the length is comprised by the turnbuckle, which is less elastic). Thus, the cables have spring constants

$$k = \frac{EA}{l} \quad (\text{A.26})$$

where $E = 29,000$ ksi is the modulus of elasticity of steel, $A = 0.11$ in² is the cross-sectional area of the cable. Thus, $k = 59,300$ lbs in⁻¹. Given the calculated tension in the cables of 9,600 lbs, the cables will stretch by $0.16''$. This equates to a horizontal translation of the spreader bar by $0.28''$. This represents a negligible perturbation to the geometry of the suspension system and does not meaningfully impact the results. It also indicates that none of the components of the gondola are in danger of impeding the motion of the spreader bar at termination.

A.4 Ballast Attachment Hand Calculations

The ballast hopper will be suspended from eyebolts at the 4 corners of the gondola cage. We assume that each cable will make a 45° angle with the vertical, and that the ballast will weigh 1,200 lbs. I analyze each load configuration independently below.

10g Applied Vertically

The cables have unit vectors $\begin{pmatrix} \pm 0.60 \\ \pm 0.38 \\ 0.71 \end{pmatrix}$, using the same coordinate system as above.

The load from the ballast is vertical, and has a magnitude of $L = 12,000$ lbs, where

$$L = 4T_z = 4(0.71T) \quad (\text{A.27})$$

which gives $T = 4,250$ lbs in each cable, at an angle of 45° with the vertical.

5g Applied at 45° to the Right

The cables on the right will swing over until they lie in the plane of the applied load, whereas the cables on the left will go slack. The cables on the right will have unit vectors $\begin{pmatrix} -0.65 \\ \pm 0.38 \\ 0.65 \end{pmatrix}$. I decompose the applied loads into x - and z -components and find

$$L_x = 2T_x = 2(0.65T) \quad (\text{A.28})$$

$$L_z = 2T_z = 2(0.65T) \quad (\text{A.29})$$

Here, $L = 6,000$ lbs. Thus, $L_x = L_z = 4,250$ lbs, and $T = 3,260$ lbs. Each of the two cables on the right make an angle of 49° with the vertical.

By symmetry, in the ‘5g applied at 45° to the left’ load case, the resulting tensions

are identical.

5g Applied at 45° Forwards

The cables in front will swing backward until they lie in the plane of the applied load, while the cables in back will go slack. The cables in front will have unit vectors $\begin{pmatrix} \pm 0.60 \\ -0.57 \\ 0.57 \end{pmatrix}$. I decompose the applied loads into y - and z -components and find

$$L_y = 2T_y = 2(0.57T) \quad (\text{A.30})$$

$$L_z = 2T_z = 2(0.57T) \quad (\text{A.31})$$

Here, $L = 6,000$ lbs. Thus, $L_y = L_z = 4,250$ lbs, and $T = 3,720$ lbs. Each of the two cables on the front make an angle of 55° with the vertical.

By symmetry, in the ‘5g applied at 45° backwards’ load case, the resulting tensions are identical.

5g Applied at 90° to the Right

The cables on the right will swing over until they lie in the plane of the applied load, whereas the cables on the left will go slack. The cables on the right will have unit vectors $\begin{pmatrix} -0.92 \\ \pm 0.38 \\ 0 \end{pmatrix}$. The applied loads is entirely in the x -direction; therefore

$$L_x = 2T_x = 2(0.92T) \quad (\text{A.32})$$

Here, $L = L_x = 6,000$ lbs, and $T = 3,250$ lbs. Each of the two cables on the right lie in the horizontal plane.

By symmetry, in the ‘5g applied at 90° to the left’ load case, the resulting tensions

are identical.

5g Applied at 90° Forwards

The cables in front will swing backward until they lie in the plane of the applied load, while the cables in back will go slack. The cables in front will have unit vectors $\begin{pmatrix} \pm 0.60 \\ -0.8 \\ 0 \end{pmatrix}$. The applied load is entirely in the y -direction; therefore

$$L_y = 2T_y = 2(0.8T) \tag{A.33}$$

Here, $L = L_y = 6,000$ lbs, and $T = 3,740$ lbs. Each of the two cables on the front lie in the horizontal plane.

By symmetry, in the ‘5g applied at 90° backwards’ load case, the resulting tensions are identical.

A.5 Suspension System Finite Element Analysis

Material is Aluminum 6061-T6 for most components. Ultimate tensile strength is 37.0 ksi. The NASA triplate is also 6061-T6, but has an ultimate tensile strength of 42.0 ksi due to being machined from thicker stock. Hoist ring reinforcement pieces are made from 4140 steel, with ultimate tensile strength of 60.0 ksi.

Suspension system overview: the NASA triplate is rigidly attached to a spreader bar. The gondola hangs from the spreader bar by way of 4 rigid rods with ball joints at either end, and 4 steel wire cables. The attachment points are clevis brackets for the rods and hoist rings for the cables. We also use turnbuckles to adjust the length of cables.

The constraints were performed in two ways: First, a fixed geometry constraint




Figure A.10: Example model setup, showing fixed geometry at the attachment point between the triplate at the universal joint. Loads are applied at the attachment points of the suspension system. This example is the 10g applied vertically case.

is applied where the triplate interfaces with the universal joint, and the loads are applied at the attachment points to the lower gondola. Second, no constraints are applied – instead, the appropriate load is applied upward at the top of the triplate and at the attachment points to the lower gondola. Inertial relief is used to stabilize the model. We treat the maximum load of the two analysis methods as the true maximum.

Hoist rings and clevis brackets are included in model for a place to apply loads. Parts are hidden in stress results; load capacities are given by their manufacturers in excess of the expected loads – see Table A.2.

Results (demonstrated in the following subsections): For the configuration in which a 10g load is applied vertically, the maximum stress is 33.5 ksi. For the 5g load applied at 45° towards the left or right, the maximum stress is 30.6 ksi in aluminum components and 35.9 ksi in steel components. For the 5g load applied at 45° towards the front or back, the maximum stress is 24.8 ksi.

10g Applied Vertically




Figure A.11: FEA results for the suspension system with 10g applied vertically in the fixed triplate configuration, showing a maximum stress of 33.3 ksi.




Figure A.12: FEA results for the suspension system with 10g applied vertically in the inertial relief configuration, showing a maximum stress of 33.5 ksi.

5g Applied at 45° Forwards

(By symmetry, 5g at 45° to vertical in backwards direction is equivalent)




Figure A.13: FEA results for the suspension system with 5g applied at 45° towards the front in the fixed triplate configuration, showing a maximum stress of 20.3 ksi - note this maximum occurs in the triplate.




Figure A.14: FEA results for the suspension system with 5g applied at 45° towards the front in the inertial relief configuration, showing a maximum stress of 24.8 ksi.

5g Applied at 45° to the Right

(By symmetry, 5g at 45° to vertical to the left is equivalent)




Figure A.15: FEA results for aluminum components of the suspension system with 5g applied at 45° towards the right in the fixed triplate configuration, showing a maximum stress of 30.6 ksi.




Figure A.16: FEA results for steel components of the suspension system with 5g applied at 45° towards the right in the fixed triplate configuration, showing a maximum stress of 35.9 ksi.




Figure A.17: FEA results for aluminum components of the suspension system with 5g applied at 45° towards the right in the inertial relief configuration, showing a maximum stress of 30.3 ksi.




Figure A.18: FEA results for steel components of the suspension system with 5g applied at 45° towards the right in the inertial relief configuration, showing a maximum stress of 35.9 ksi.

A.6 Rigid Rod Analysis

The previous analysis of the suspension system shows that we expect to see compressive forces of up to 4,550 lbs and tensile forces up to 10,520 lbs on the rigid rods that make up the suspension system. Here I demonstrate hand calculations to show that the rods will not fail under such applied loads. I also perform FEA analysis to validate our hand calculations.




Figure A.19: FEA results for the buckling analysis of the rigid rods of the suspension system, showing a critical load of 6,308 lbs.

A.6.1 Buckling Analysis

The rods are 1018 steel, 1" hex stock with each end threaded to fit a ball joint. The rods have the following properties: a modulus of elasticity $E = 29,700$ ksi, an area moment of inertia of $I = 0.06$ in⁴, and an effective length¹² of $L_e = 53$ ".

The load at which such a rod will buckle is given by

$$P_{cr} = \frac{\pi^2 EI}{L_e^2} \quad (\text{A.34})$$

which evaluates to $P_{cr} = 6,260$ lbs for the values gives above.

¹²Since both ends of the rods are attached to spherical ball joints and are therefore free to rotate, the effective length of the rods for the purposes of calculating buckling conditions is simply the length of the rods

This result was confirmed by an FEA buckling simulation. As shown Figure A.19, I applied pinned constraints at either end of the rod and applied a 1 lb test load to one end of the rod. I found a load factor of 6,308, giving a critical load of 6,308 lbs, in good agreement with the hand calculation given above.



Figure A.20: Schematic of the FEA analysis of the rigid rods under tension.

A.6.2 Tension Analysis

The maximum tensile stress will occur where the rod has been turned down to 0.75", due to the stress concentration and reduced area. If we assume a stress concentration factor $K_t = 1.53$ [59], then our expected maximum stress will be given by

$$\sigma_{max} = K_t \frac{L}{A} \quad (\text{A.35})$$

where L is the applied load, and $A = 0.44 \text{ in}^2$ is the cross-sectional area of the turned-down portion of the rod.

I also performed an FEA analysis to validate this analytical results. As shown in

Figure A.20, I fixed one end of the rod and applied a tensile load at the other end. In every case, the maximum resulting stress is less than the ultimate tensile strength of 1018 steel, which is 55 ksi.

10g Applied Vertically

In this case, the rods each come under a tensile load of 8,845 lbs. Analytically, we expect a maximum stress of 30.2 ksi. The FEA analysis shown in Figure A.21 gives a maximum stress of 33.2 ksi.




Figure A.21: FEA results for the rigid rods with 10g applied vertically.

5g Applied at 45° Forwards

(By symmetry, 5g at 45° to vertical in backwards direction is equivalent)

In this case, two of the rods come under a tensile load of 10,520 lbs. Analytically, we expect a maximum stress of 35.9 ksi. The FEA analysis shown in Figure A.22 gives a maximum stress of 39.5 ksi.




Figure A.22: FEA results for the rigid rods with 5g applied at 45° forwards or backwards.

5g Applied at 45° to the Right

(By symmetry, 5g at 45° to vertical to the left is equivalent)

In this case, two of the rods come under a tensile load of 230 lbs. Analytically, we expect a maximum load of 0.8 ksi. The FEA analysis shown below gives a max stress of 0.9 ksi.




Figure A.23: FEA results for the rigid rods with 5g applied at 45° to the left or right.

A.7 Gondola Cage Finite Element Analysis

Material is Aluminum 6061-T6 for all components, with an ultimate tensile strength of 37.0 ksi.

Simple mass models are used for lidar pressure vessel (210 lbs), lidar telescope (87 lbs), SIP crate (350 lbs), cameras (40 lbs each), solar arrays (60 lbs each), and ballast (340 lbs each). The ballast weight is exaggerated to be conservative.

Fixed constraint is placed where suspension members attach to gondola cage. Gravity is then applied as 10g vertically, 5g at 45°, and 5g at 90°.

Result: Maximum stress is 32.8 ksi. No yielding is expected to occur.




Figure A.24: Schematic of the setup used in the FEA analysis of the gondola cage, with fixed geometry constraints (green arrows) applied at the attachment points to the suspension system, and augmented gravity (red arrow) applied in the appropriate direction. This figure shows the '10g applied vertically' load case.

10g Applied Vertically




Figure A.25: FEA results for the gondola cage with 10g applied vertically, showing a maximum stress of 20.1 ksi.


5g Applied at 45° Backwards

Figure A.26: FEA results for the gondola cage with 5g applied at 45° in the backwards direction, showing a maximum stress of 18.3 ksi.


5g Applied at 45° Forwards

Figure A.27: FEA results for the gondola cage with 5g applied at 45° in the forwards direction, showing a maximum stress of 24.2 ksi.

5g Applied at 45° to the Left




Figure A.28: FEA results for the gondola cage with 5g applied at 45° to the left, showing a maximum stress of 22.0 ksi.


5g Applied at 45° to the Right

Figure A.29: FEA results for the gondola cage with 5g applied at 45° to the right, showing a maximum stress of 31.0 ksi.


5g Applied at 90° Backwards

Figure A.30: FEA results for the gondola cage with 5g applied at 90° in the backwards direction, showing a maximum stress of 27.7 ksi.


5g Applied at 90° Forwards

Figure A.31: FEA results for the gondola cage with 5g applied at 90° in the forwards direction, showing a maximum stress of 27.7 ksi.

5g Applied at 90° to the Left




Figure A.32: FEA results for the gondola cage with 5g applied at 90° to the left, showing a maximum stress of 32.8 ksi.

5g Applied at 90° to the Right




Figure A.33: FEA results for the gondola cage with 5g applied at 90° to the right, showing a maximum stress of 32.8 ksi.

A.8 Ballast Attachments Finite Element Analysis

Material: Most parts made out of Aluminum 6061-T6 (ultimate tensile strength 37 ksi). Steel reinforcement plates at each eyebolt are 4140 steel (ultimate tensile strength ≥ 60 ksi). We do not have material certification for the tensile strength of the steel, so we take a conservative value.

Fixed geometry constraints are applied where lower floor of gondola interfaces with the rest of the gondola. Loads at each eyebolt are calculated by hand for each scenario; see Appendix A.3. I assume 1,200 lbs of ballast.

Eyebolts are included in the model for a place to apply loads. They are hidden in stress results; load capacities are given by their manufacturers and are in excess of the expected loads; see Table A.2.

Results: There are stress concentrations exceeding the ultimate tensile strength in the steel reinforcement plates in every loading configuration. I show that these stress concentrations (1) are an artifact of the FEA software, and (2) affect only a minimal volume. I conclude that at worst, we will see minimal plastic deformation with no risk of failure. In all other parts, stresses are below ultimate tensile strength.



Figure A.34: Setup of the ballast attachment FEA, with fixed geometry constraints applied at points where lower floor of gondola (pictured) attached to the rest of the gondola. Loads are applied at each corner towards the center (purple arrows) – I assumed the ballast would be hung such that the cables make a 45° angle with the vertical. This corresponds to the ‘10g applied vertically’ load case.

Diverging singularity in Finite Element Analysis Results

The following figures show the stresses in the steel reinforcement plate for decreasing mesh sizes, demonstrating that the maximum stress diverges. They further demonstrate that the regions where the stress exceeds the ultimate tensile strength of the material are minimal. These results are for the ‘5g applied horizontally to the left’ load case, but they are substantially similar to all other cases.




Figure A.35: FEA results for the ballast attachment points with 5g applied at 90° to the left, with a mesh size of 0.5” (left) and 0.25” (right), showing a maximum stress of 36.2 ksi and 39.9 ksi, respectively.




Figure A.36: Same as Figure A.35, but with a mesh size of 0.15" (left) and 0.1" (right), showing a maximum stress of 65.3 ksi and 94.8 ksi, respectively.




Figure A.37: Same as Figure A.35, but with a mesh size of 0.05", showing a maximum stress of 125 ksi.

Clearly, this is an artifact of the FEA software. But even if we trust the FEA results at a mesh size of 0.05", this will cause slight plastic deformation, not ultimate failure. Below, we show contour plots showing the regions that have higher stress than 60 ksi, demonstrating how small the affected volume is. Some plastic deformation may occur in this region, but it will not lead to ultimate failure of the plate.




Figure A.38: Side-view (left) and top-view (right) of the steel reinforcement plate of the ballast attachment point, highlighting (red) the volume where the stress exceeds 60 ksi (the ultimate tensile strength of the material).

10g Applied Vertically




Figure A.39: FEA results for the ballast attachment points with 10g applied vertically, showing a maximum stress of 115 ksi. This is due to a diverging singularity as demonstrated above; only a small region is affected.




Figure A.40: FEA results for the ballast attachment points with 10g applied vertically, showing where the maximum stresses occur. At distances less than 0.1” from these stress concentrations, the calculated stresses are significantly below the ultimate tensile strength.




Figure A.41: FEA results for the aluminum components of the ballast attachment points with 10g applied vertically, showing a maximum stress of 16.2 ksi.

5g Applied at 45° to the Right

(By symmetry, 5g at 45° to vertical to the left is equivalent)




Figure A.42: FEA results for the ballast attachment points with 5g applied at 45° to the right, showing a maximum stress of 89.6 ksi. This is due to a diverging singularity as demonstrated above; only a small region is affected.




Figure A.43: FEA results for the ballast attachment points with 5g applied at 45° to the right, showing where the maximum stresses occur. At distances less than 0.1” from these stress concentrations, the calculated stresses are significantly below the ultimate tensile strength.



Figure A.44: FEA results for the aluminum components of the ballast attachment points with 5g applied at 45° to the right, showing a maximum stress of 13.0 ksi.

5g Applied at 45° Forwards

(By symmetry, 5g at 45° to vertical in backwards direction is equivalent)




Figure A.45: FEA results for the ballast attachment points with 5g applied at 45° towards the front, showing a maximum stress of 77.6 ksi. This is due to a diverging singularity as demonstrated above; only a small region is affected.

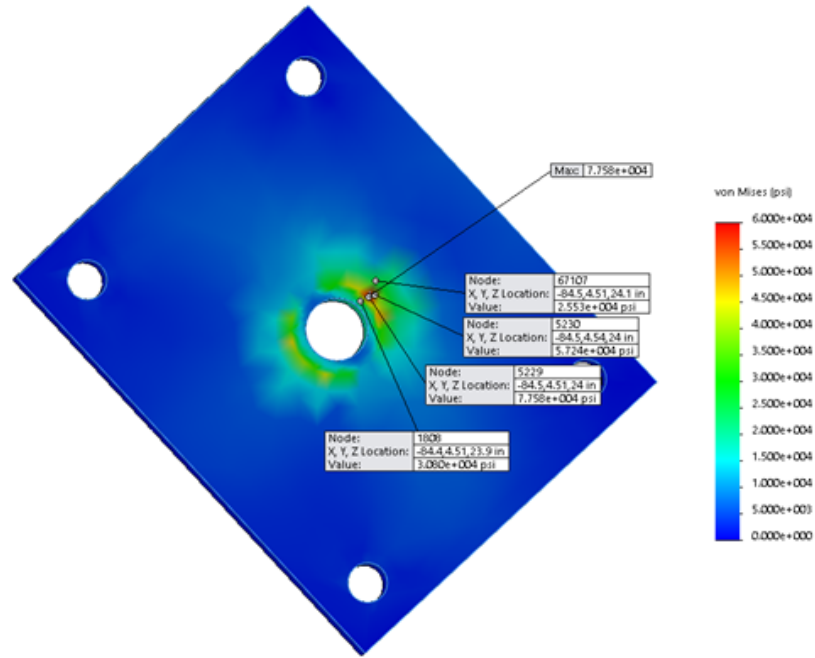


Figure A.46: FEA results for the ballast attachment points with 5g applied at 45° towards the front, showing where the maximum stresses occur. At distances less than 0.1” from these stress concentrations, the calculated stresses are significantly below the ultimate tensile strength.



Figure A.47: FEA results for the aluminum components of the ballast attachment points with 5g applied at 45° towards the front, showing a maximum stress of 16.5 ksi.

5g Applied at 90° to the Right

(By symmetry, 5g at 90° to vertical to the left is equivalent)




Figure A.48: FEA results for the ballast attachment points with 5g applied at 90° to the right, showing a maximum stress of 125 ksi. This is due to a diverging singularity as demonstrated above; only a small region is affected.

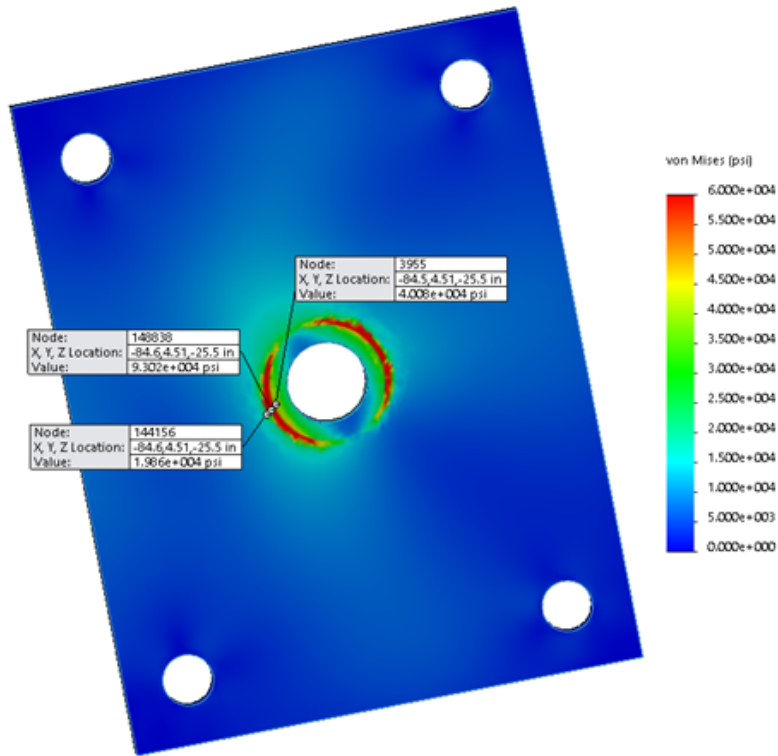


Figure A.49: FEA results for the ballast attachment points with 5g applied at 90° to the right, showing where the maximum stresses occur. At distances less than 0.1” from these stress concentrations, the calculated stresses are significantly below the ultimate tensile strength.




Figure A.50: FEA results for the aluminum components of the ballast attachment points with 5g applied at 90° to the right, showing a maximum stress of 16.9 ksi.

5g Applied at 90° Forwards

(By symmetry, 5g at 90° to vertical in backwards direction is equivalent)




Figure A.51: FEA results for the ballast attachment points with 5g applied at 90° towards the front, showing a maximum stress of 131 ksi. This is due to a diverging singularity as demonstrated above; only a small region is affected.




Figure A.52: FEA results for the ballast attachment points with 5g applied at 90° towards the front, showing where the maximum stresses occur. At distances less than 0.1” from these stress concentrations, the calculated stresses are significantly below the ultimate tensile strength.




Figure A.53: FEA results for the aluminum components of the ballast attachment points with 5g applied at 90° towards the front, showing a maximum stress of 19.4 ksi.

A.9 Middle Floor Finite Element Analysis

Material: All parts made out of Aluminum 6061-T6, with an ultimate tensile strength 37.0 ksi.

Simple mass models are used for the lidar pressure vessel (left, 210 lbs), lidar telescope (right, 87 lbs) and SIP crate (350 lbs).

Fixed geometry constraints are applied where middle floor interfaces with the rest of the gondola.

Augmented gravity is applied in the appropriate direction in each load case.

Results: Maximum stress is 12.1 ksi.




Figure A.54: Setup of the FEA for the middle floor, showing fixed geometry constraints where the middle floor attaches to the rest of the gondola (green arrows). Augmented gravity (red arrow) is applied in the appropriate direction in each load case - the case shown here corresponds to the '5g applied at 45° to the left' load case.

10g Applied Vertically




Figure A.55: FEA results for the middle floor with 10g applied vertically, showing a maximum stress of 12.1 ksi.

5g Applied at 45° to the Right




Figure A.56: FEA results for the middle floor with 5g applied at 45° to the right, showing a maximum stress of 9.9 ksi.

5g Applied at 45° to the Left




Figure A.57: FEA results for the middle floor with 5g applied at 45° to the left, showing a maximum stress of 11.2 ksi.

5g Applied at 45° Backwards




Figure A.58: FEA results for the middle floor with 5g applied at 45° towards the back, showing a maximum stress of 7.9 ksi.

5g Applied at 45° Forwards




Figure A.59: FEA results for the middle floor with 5g applied at 45° towards the front, showing a maximum stress of 5.4 ksi.

5g Applied at 90° to the Right




Figure A.60: FEA results for the middle floor with 5g applied at 90° towards the right, showing a maximum stress of 12.0 ksi.

5g Applied at 90° to the Left




Figure A.61: FEA results for the middle floor with 5g applied at 90° towards the left, showing a maximum stress of 12.0 ksi.

5g Applied at 90° Backwards




Figure A.62: FEA results for the middle floor with 5g applied at 90° towards the back, showing a maximum stress of 5.4 ksi.


5g Applied at 90° Forwards

Figure A.63: FEA results for the middle floor with 5g applied at 90° towards the front, showing a maximum stress of 5.4 ksi.

A.10 Middle Floor Welds Finite Element Analysis

Material: All parts made out of Aluminum 6061-T6. Since these parts contain welds, I assume an ultimate tensile strength 12.0 ksi.

Simple mass models are used for the lidar pressure vessel (left, 210 lbs), lidar telescope (right, 87 lbs) and SIP crate (350 lbs). All parts not containing welds are excluded from the FEA results to demonstrate the highest stresses in welded components.

Fixed geometry constraints are applied where middle floor interfaces with the rest of the gondola.

Augmented gravity is applied in the appropriate direction in each load case.

Results: Maximum stress is 9.1 ksi. Though this stress occurs in a component that contains welds, the stress occurs far from the weld itself. The maximum stress that occurs in the vicinity of a weld is 5.4 ksi.




Figure A.64: Setup of the FEA for the middle floor welded parts, showing fixed geometry constraints where the middle floor attaches to the rest of the gondola (green arrows). Augmented gravity (red arrow) is applied in the appropriate direction in each load case - the case shown here corresponds to the '5g applied at 45° to the left' load case.


10g Applied Vertically

Figure A.65: FEA results for the welded components of the middle floor with 10g applied vertically, showing a maximum stress of 8.5 ksi.


5g Applied at 45° to the Right

Figure A.66: FEA results for the welded components of the middle floor with 5g applied at 45° to the right, showing a maximum stress of 9.1 ksi.


5g Applied at 45° to the Left

Figure A.67: FEA results for the welded components of the middle floor with 5g applied at 45° to the left, showing a maximum stress of 8.3 ksi.


5g Applied at 45° Backwards

Figure A.68: FEA results for the welded components of the middle floor with 5g applied at 45° towards the back, showing a maximum stress of 5.4 ksi.


5g Applied at 45° Forwards

Figure A.69: FEA results for the welded components of the middle floor with 5g applied at 45° towards the front, showing a maximum stress of 5.4 ksi.


5g Applied at 90° to the Right

Figure A.70: FEA results for the welded components of the middle floor with 5g applied at 90° towards the right, showing a maximum stress of 7.2 ksi.


5g Applied at 90° to the Left

Figure A.71: FEA results for the welded components of the middle floor with 5g applied at 90° towards the left, showing a maximum stress of 7.2 ksi.


5g Applied at 90° Backwards

Figure A.72: FEA results for the welded components of the middle floor with 5g applied at 90° towards the back, showing a maximum stress of 5.4 ksi.


5g Applied at 90° Forwards

Figure A.73: FEA results for the welded components of the middle floor with 5g applied at 90° towards the front, showing a maximum stress of 5.4 ksi.

A.11 Battery Floor Finite Element Analysis

Material: All parts made out of Aluminum 6061-T6. Since these parts contain welds, I assume an ultimate tensile strength 12.0 ksi.

Simple mass models are used for the battery box (60 lbs) and telemetry box (20 lbs).

Fixed geometry constraints are applied where battery floor interfaces with the rest of the gondola.

Augmented gravity is applied in the appropriate direction in each load case.

Results: Maximum stress is 27.5 ksi.




Figure A.74: Setup of the FEA for the battery floor, showing fixed geometry constraints where the battery floor attaches to the rest of the gondola (green arrows). Augmented gravity (red arrow) is applied in the appropriate direction in each load case - the case shown here corresponds to the '5g applied vertically' load case. Note that the view here and in the following figures is from the side rather than from the front.

10g Applied Vertically




Figure A.75: FEA results for the battery floor with 10g applied vertically, showing a maximum stress of 27.5 ksi.

5g Applied at 45° to the Right

(By symmetry, 5g at 45° to vertical to the left is equivalent)




Figure A.76: FEA results for the battery floor with 5g applied at 45° to the right, showing a maximum stress of 13.2 ksi.

5g Applied at 45° Backwards




Figure A.77: FEA results for the battery floor with 5g applied at 45° towards the back, showing a maximum stress of 12.7 ksi.

5g Applied at 45° Forwards




Figure A.78: FEA results for the battery floor with 5g applied at 45° towards the front, showing a maximum stress of 7.6 ksi.

5g Applied at 90° to the Right

(By symmetry, 5g at 90° to vertical to the left is equivalent)




Figure A.79: FEA results for the battery floor with 5g applied at 90° to the vertical towards the right, showing a maximum stress of 8.8 ksi.

5g Applied at 90° Backwards




Figure A.80: FEA results for the battery floor with 5g applied at 90° to the vertical towards the back, showing a maximum stress of 4.3 ksi.


5g Applied at 90° Forwards

Figure A.81: FEA results for the battery floor with 5g applied at 90° to the vertical towards the front, showing a maximum stress of 4.3 ksi.

A.12 Radiator Finite Element Analysis

Material: All parts made out of Aluminum 6061-T6, with an ultimate tensile strength 37.0 ksi.

A simple mass model is used for the radiator (40 lbs).

Fixed geometry constraints are applied where radiator interfaces with the rest of the gondola.

Augmented gravity is applied in the appropriate direction in each load case.

Results: Maximum stress is 12.1 ksi.




Figure A.82: Setup of the FEA for the radiator, showing fixed geometry constraints where the radiator mounts attach to the rest of the gondola (green arrows). Augmented gravity (red arrow) is applied in the appropriate direction in each load case - the case shown here corresponds to the '10g applied vertically' load case.

10g Applied Vertically




Figure A.83: FEA results for the radiator with 10g applied vertically, showing a maximum stress of 8.1 ksi.

5g Applied at 45° to the Right




Figure A.84: FEA results for the radiator with 5g applied at 45° to the right, showing a maximum stress of 6.7 ksi.

5g Applied at 45° to the Left




Figure A.85: FEA results for the radiator with 5g applied at 45° to the left, showing a maximum stress of 6.7 ksi.

5g Applied at 45° Backwards




Figure A.86: FEA results for the radiator with 5g applied at 45° towards the back, showing a maximum stress of 2.1 ksi.


5g Applied at 45° Forwards

Figure A.87: FEA results for the radiator with 5g applied at 45° towards the front, showing a maximum stress of 5.7 ksi.


5g Applied at 90° to the Right

Figure A.88: FEA results for the radiator with 5g applied at 90° towards the right, showing a maximum stress of 9.5 ksi.

5g Applied at 90° to the Left




Figure A.89: FEA results for the radiator with 5g applied at 90° towards the left, showing a maximum stress of 9.5 ksi.

5g Applied at 90° Backwards




Figure A.90: FEA results for the radiator with 5g applied at 90° towards the back, showing a maximum stress of 4.0 ksi.


5g Applied at 90° Forwards

Figure A.91: FEA results for the radiator with 5g applied at 90° towards the front, showing a maximum stress of 4.0 ksi.

A.13 Radiator Welds Finite Element Analysis

Material: All parts made out of Aluminum 6061-T6. The two bottom braces of the radiator mounts contain welds; I assume a de-rating of the ultimate tensile strength to 12.0 ksi.

A simple mass model is used for the radiator (40 lbs).

Fixed geometry constraints are applied where radiator interfaces with the rest of the gondola.

Augmented gravity is applied in the appropriate direction in each load case.

Results: Maximum stress is 4.2 ksi.




Figure A.92: Setup of the FEA for the radiator welded parts, showing fixed geometry constraints where the radiator mounts attach to the rest of the gondola (green arrows). Augmented gravity (red arrow) is applied in the appropriate direction in each load case - the case shown here corresponds to the '10g applied vertically' load case.

10g Applied Vertically




Figure A.93: FEA results for the welded components of the radiator with 10g applied vertically, showing a maximum stress of 4.2 ksi.

5g Applied at 45° to the Right

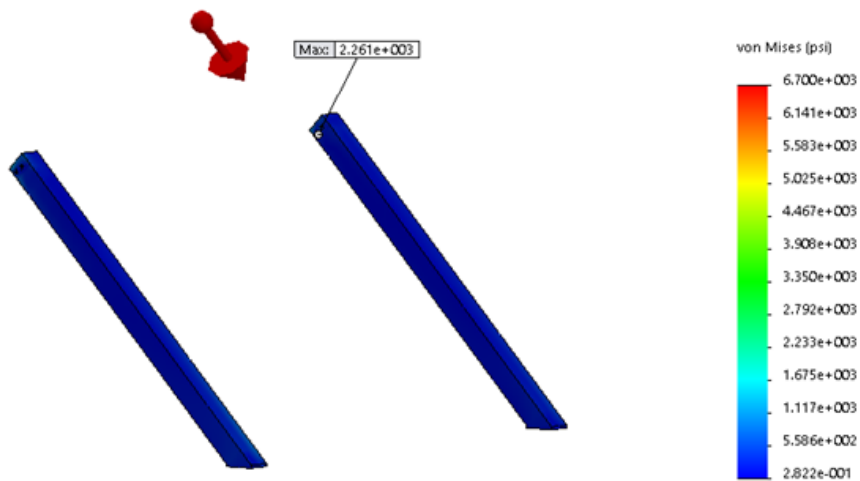


Figure A.94: FEA results for the welded components of the radiator with 5g applied at 45° to the right, showing a maximum stress of 2.3 ksi.

5g Applied at 45° to the Left




Figure A.95: FEA results for the welded components of the radiator with 5g applied at 45° to the left, showing a maximum stress of 2.1 ksi.

5g Applied at 45° Backwards




Figure A.96: FEA results for the welded components of the radiator with 5g applied at 45° towards the back, showing a maximum stress of 1.2 ksi.

5g Applied at 45° Forwards




Figure A.97: FEA results for the welded components of the radiator with 5g applied at 45° towards the front, showing a maximum stress of 2.5 ksi.

5g Applied at 90° to the Right




Figure A.98: FEA results for the welded components of the radiator with 5g applied at 90° towards the right, showing a maximum stress of 1.4 ksi.

5g Applied at 90° to the Left

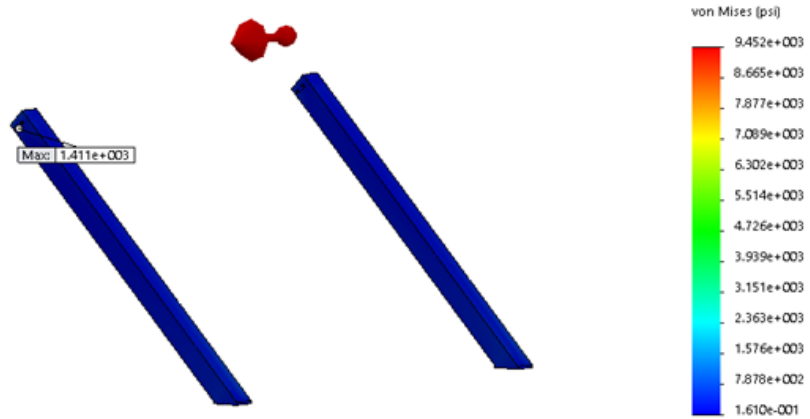


Figure A.99: FEA results for the welded components of the radiator with 5g applied at 90° towards the left, showing a maximum stress of 1.4 ksi.

5g Applied at 90° Backwards




Figure A.100: FEA results for the welded components of the radiator with 5g applied at 90° towards the back, showing a maximum stress of 1.7 ksi.


5g Applied at 90° Forwards

Figure A.101: FEA results for the welded components of the radiator with 5g applied at 90° towards the front, showing a maximum stress of 1.7 ksi.

A.14 Solar Array Finite Element Analysis

Material: All parts made out of Aluminum 6061-T6, with an ultimate tensile strength 37.0 ksi.

A simple mass model is used for each solar panel (1.75 lbs).

Fixed geometry constraints are applied where solar array interfaces with the rest of the gondola. By symmetry, the midpoint of the two beams connecting the two solar arrays is constrained to the symmetry plane, so a fixed plane constraint is applied at those points.

Augmented gravity is applied in the appropriate direction in each load case.

Results: Maximum stress is 31.1 ksi.




Figure A.102: Setup of the FEA for the solar array, showing fixed geometry constraints where the solar array attaches to the rest of the gondola (green arrows) and fixed plane constraints applied where indicated. Augmented gravity (red arrow) is applied in the appropriate direction in each load case - the case shown here corresponds to the '5g applied at 45° forwards' load case.

10g Applied Vertically




Figure A.103: FEA results for the solar array with 10g applied vertically, showing a maximum stress of 31.1 ksi.

5g Applied at 45° to the Right




Figure A.104: FEA results for the solar array with 5g applied at 45° to the right, showing a maximum stress of 8.8 ksi.

5g Applied at 45° to the Left




Figure A.105: FEA results for the solar array with 5g applied at 45° to the left, showing a maximum stress of 9.1 ksi.

5g Applied at 45° Backwards




Figure A.106: FEA results for the solar array with 5g applied at 45° towards the back, showing a maximum stress of 24.2 ksi.

5g Applied at 45° Forwards




Figure A.107: FEA results for the solar array with 5g applied at 45° towards the front, showing a maximum stress of 24.0 ksi.

5g Applied at 90° to the Right




Figure A.108: FEA results for the solar array with 5g applied at 90° towards the right, showing a maximum stress of 3.0 ksi.

5g Applied at 90° to the Left




Figure A.109: FEA results for the solar array with 5g applied at 90° towards the left, showing a maximum stress of 3.0 ksi.

5g Applied at 90° Backwards




Figure A.110: FEA results for the solar array with 5g applied at 90° towards the back, showing a maximum stress of 33.5 ksi.


5g Applied at 90° Forwards

Figure A.111: FEA results for the solar array with 5g applied at 90° towards the front, showing a maximum stress of 33.5 ksi.

A.15 Camera Mounts Finite Element Analysis

Material: All parts made out of Aluminum 6061-T6, with an ultimate tensile strength 37.0 ksi.

A simple mass model is used for each camera (35 lbs).

Fixed geometry constraints are applied where solar array interfaces with the rest of the gondola. By symmetry, the midpoint of the two beams connecting the two solar arrays is constrained to the symmetry plane, so a fixed plane constraint is applied at those points.

Augmented gravity is applied in the appropriate direction in each load case.

Results: Maximum stress is 9.2 ksi.




Figure A.112: Setup of the FEA for the camera mounts, showing fixed geometry constraints where the structural members that support the cameras attach to the rest of the gondola (green arrows). Augmented gravity (red arrow) is applied in the appropriate direction in each load case - the case shown here corresponds to the '5g applied vertically' load case.

10g Applied Vertically




Figure A.113: FEA results for the camera mounts with 10g applied vertically, showing a maximum stress of 9.2 ksi.

5g Applied at 45° to the Right




Figure A.114: FEA results for the camera mounts with 5g applied at 45° to the right, showing a maximum stress of 3.0 ksi.

5g Applied at 45° to the Left




Figure A.115: FEA results for the camera mounts with 5g applied at 45° to the left, showing a maximum stress of 1.6 ksi.


5g Applied at 45° Backwards

Figure A.116: FEA results for the camera mounts with 5g applied at 45° towards the back, showing a maximum stress of 2.3 ksi.


5g Applied at 45° Forwards

Figure A.117: FEA results for the camera mounts with 5g applied at 45° towards the front, showing a maximum stress of 2.3 ksi.


5g Applied at 90° to the Right

Figure A.118: FEA results for the camera mounts with 5g applied at 90° towards the right, showing a maximum stress of 1.9 ksi.

5g Applied at 90° to the Left




Figure A.119: FEA results for the camera mounts with 5g applied at 90° towards the left, showing a maximum stress of 1.9 ksi.

5g Applied at 90° Backwards




Figure A.120: FEA results for the camera mounts with 5g applied at 90° towards the back, showing a maximum stress of 2.7 ksi.


5g Applied at 90° Forwards

Figure A.121: FEA results for the camera mounts with 5g applied at 90° towards the front, showing a maximum stress of 2.7 ksi.