

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]

UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and Physical Science
Department of Aeronautics, Astronautics and Computational Engineering
Astronautics Research Group

Attitude Characterisation of Space Objects using Optical Light Curves

by

Laurence David James Blacketer

MSc

ORCID: [0000-0002-5941-2312](https://orcid.org/0000-0002-5941-2312)

*A thesis for the degree of
Doctor of Philosophy*

March 2022

University of Southampton

Abstract

Faculty of Engineering and Physical Science
Department of Aeronautics, Astronautics and Computational Engineering

Doctor of Philosophy

Attitude Characterisation of Space Objects using Optical Light Curves

by Laurence David James Blacketer

This PhD thesis focuses on the relationship between space object attitude state and time-varying brightness signals. The large and increasing number of objects in Earth orbit has motivated improvements to the global Space Situational Awareness (SSA) capability. Radar installations are invaluable for observing space objects as they can operate 24-hours per day regardless of weather and, in the cases of phased array radars, can observe large numbers of objects simultaneously. However, these radar installations are extremely costly to build and operate. Optical telescopes can be built and operated at substantially reduced cost and therefore present an attractive alternative. For this reason, there is considerable interest in the global SSA community to maximise the potential of optical installations.

One important area of interest for exploiting optical installations is to understand the relationship between a space object's brightness and its attitude state. The brightness of an object is, in part, a function of its orientation with respect to the source of illumination and the observer. It is therefore possible to determine information about a space object's attitude state by inverting light curve data. In this context, a light curve is a high cadence series of brightness measurements, typically in the visual band on the order of a few minutes in duration. This has been performed extensively in the literature for natural objects such as asteroids, and has also been applied to artificial space objects such as active or inactive spacecraft and rocket bodies.

The aim of this PhD research was to develop novel methods to derive information on space objects from their light curves, and so augment present day SSA capabilities. The selected approach was to develop a synthetic light curve generation model which was then used to improve understanding of optical light curve data. This synthetic light curve model was applied to a wide range of simulated and real-world light curve examples. First, changes in standard deviation of light curves were used to differentiate active spacecraft from inactive spacecraft and identify possible failure dates. Second, the sensitivity of the synthetic light curve generation model to its various inputs was examined. Third, the relationship between illumination conditions

and attitude state was determined to identify those conditions in which attitude states were most distinguishable. Fourth, a technique for determination of the angular velocity vector for rocket bodies was developed. Finally, the synthetic light curve model was used to used to determine reasons for the very bright appearance of Starlink satellites in the night sky and the role of the solar panels in this.

These novel applications and results can be used to further improve the value of space object light curve data and their integration into present day SSA capabilities.

Contents

List of Figures	ix
List of Tables	xvii
Declaration of Authorship	xix
Acknowledgements	xxi
Definitions and Abbreviations	xxiii
1 Introduction	1
1.1 The Orbital Environment	1
1.1.1 Overview	1
1.1.2 Detection of Space Objects	2
1.1.2.1 Radar	3
1.1.2.2 Optical	4
1.1.2.3 In-Situ	6
1.1.3 Space Object Catalogue	8
1.1.3.1 SGP and the TLE format	8
1.1.4 The Space Object Population	10
1.1.4.1 Tracked Objects	10
1.1.4.2 Untracked Objects	11
1.2 Debris Modelling	12
1.2.1 History of Debris Modelling	13
1.2.1.1 EVOLVE	14
1.2.1.2 ORDEM	14
1.2.1.3 LEGEND	15
1.2.1.4 DAMAGE	15
1.2.1.5 Other Debris Models	16
1.2.2 Modelling Results	16
1.3 Debris Management	18
1.3.1 The IADC Space Debris Mitigation Guidelines	18
1.3.2 Active Debris Removal	19
1.4 The Next Generation of Space Situational Awareness	20
1.4.1 Optical Light Curves	21
1.4.1.1 History of Light Curve Analysis	23
1.4.1.2 Light Curve Analysis of Spacecraft	23
1.5 Research Objectives	25

2	Synthetic Light Curve Model Development	27
2.1	Motivations	27
2.2	Model Requirements	28
2.3	Code Operation	28
2.3.1	Model Inputs	29
2.3.2	Pre-Light Curve Generation	29
2.3.2.1	Initialisation Module	29
2.3.2.2	Coordinate Transformation Module	31
2.3.3	Light Curve Generation	31
2.3.4	Bidirectional Reflectance Distribution Function Calculation	33
2.3.4.1	Cook-Torrance BRDF	34
2.3.5	Self-Shadowing Calculation	37
2.3.6	Measuring the difference between light curves	37
2.3.6.1	Root Mean Squared Error	37
2.3.6.2	Hilbert distance calculation	38
2.4	Code Validation	40
2.4.1	Position Data Uncertainty	40
2.4.2	Position Data Validation	41
2.4.3	Coordinate Transformation Validation	42
2.4.4	Brightness Calculation Validation	43
3	Investigating Simple Light Curve Properties	49
3.1	Introduction	49
3.2	Methodology	50
3.3	Results	51
3.3.1	Real Light Curve Data	52
3.3.2	Synthetic Light Curve Data	56
3.4	Discussion	61
3.5	Conclusions	62
4	Investigation of Variable Rocket Body Light Curve Periods	63
4.1	Introduction	63
4.2	Methodology	64
4.3	Results and Discussion	65
4.3.1	Analysis of Period Extraction	70
4.3.1.1	Orderly Category Period Extraction	71
4.3.1.2	Complex Category Period Extraction	77
4.3.1.3	Summary of Period Extraction Results	80
4.3.2	Causes of Changes in Periodicity	81
4.3.2.1	CZ Rocket Bodies	83
4.3.2.2	DELTA Rocket Bodies	85
4.3.2.3	ATLAS Rocket Bodies	85
4.3.2.4	Summary of Causes of Changes in Periodicity Results	87
4.4	Conclusions	88
5	Synthetic Light Curve Model Sensitivity Study	91
5.1	Introduction	91

5.2	Methodology	91
5.3	Results	93
5.3.1	Sensitivity to Phase Angle	93
5.3.2	Sensitivity to BRDF Parameters	96
5.3.3	Sensitivity to Attitude State	99
5.3.4	Sensitivity to Rotation Rate	100
5.3.5	Summary of RMSE Differences	101
5.4	Conclusions	103
6	Investigation of the Reflection Characteristics of Primitive Geometries	105
6.1	Introduction	105
6.2	Methodology	105
6.3	Results and Discussion	107
6.3.1	Cube with diffuse reflection	108
6.3.2	Cube with specular reflection	110
6.3.3	Icosahedron with diffuse reflection	111
6.3.4	Icosahedron with specular reflection	113
6.3.5	Polyhedron with diffuse reflection	114
6.3.6	Polyhedron with specular reflection	114
6.3.7	Synthetic Light Curve Demonstration	118
6.4	Conclusions	118
7	Angular Velocity Vector Determination of Rocket Bodies	123
7.1	Introduction	123
7.2	Methodology	124
7.2.1	Real Target Light Curve	124
7.2.2	Synthetic Target Light Curves	125
7.2.3	Determining Best Fitting Light Curve Inputs	127
7.3	Results	130
7.4	Discussion	133
7.4.1	Polar Plot Discussion	133
7.4.2	Light Curve Plot Discussion	134
7.5	Conclusions	135
8	Analysis of the Brightness of Starlink Satellites	137
8.1	Introduction	137
8.2	Methodology	142
8.2.1	Starlink Light Curve Data	142
8.2.1.1	Synthetic Light Curve Inputs	144
8.3	Results	144
8.3.1	Synthetic Light Curves	144
8.3.2	Real Light Curves	147
8.4	Discussion	149
8.5	Conclusions	151
9	Conclusions	153
9.1	Introduction	153
9.2	Summary of Results Chapters	153

9.2.1	Chapter 3: Investigating Simple Light Curve Properties	153
9.2.2	Chapter 4: Investigation of Variable Rocket Body Light Curve Pe- riods	154
9.2.3	Chapter 5: Synthetic Light Curve Model Sensitivity Study	155
9.2.4	Chapter 6: Investigation of the Reflection Characteristics of Prim- itive Geometries	156
9.2.5	Chapter 7: Angular Velocity Vector Determination of Rocket Bodies	156
9.2.6	Chapter 8: Analysis of the Brightness of Starlink Satellites	157
9.3	Thesis Conclusions	158
Appendix A Quantifying the Difference Between Light Curves using Hilbert Space		161
Appendix A.1	Introduction	161
Appendix A.2	Recap: Hilbert Space Distance Calculation	161
Appendix A.3	Example 1	162
Appendix A.4	Example 2	164
Appendix A.5	Conclusion	165
References		171

List of Figures

1.1	Monthly mass of objects in the United States Space Surveillance Network (SSN) catalogue (Liou et al., 2020).	3
1.2	The Haystack radar complex (NASA, 2008).	4
1.3	The RAF Fylingdales phased array radar. From: https://en.wikipedia.org/wiki/RAF_Fylingdales	5
1.4	One of the three GEODSS sites (NASA, 2008).	6
1.5	The MODEST (front centre) and the SMARTS (left) telescopes at CTIO (NASA, 2008).	6
1.6	A picture of the LDEF platform on-orbit taken from the space shuttle (NASA, 2008).	7
1.7	Monthly number of objects in the United States Space Surveillance Network (SSN) catalogue (Liou et al., 2020).	11
1.8	Spatial distribution with altitude of objects in the TLE catalogue (Liou, 2017).	12
1.9	Combined radar and in-situ measurements of debris flux (Schildknecht, 2007)	13
2.1	The operation of the first two modules of the synthetic light curve generation software.	30
2.2	The RSW satellite-centred reference frame used for light curve generation.	32
2.3	The simplified geometry of a box-wing satellite. The central cube represents the main spacecraft bus and thin rectangles represent solar panels, attached via booms.	32
2.4	Flow chart showing the calculation loops in the Light Curve Generation Module.	33
2.5	The reflection geometry. \mathbf{V} is the observation unit vector, \mathbf{L} is the illumination unit vector and \mathbf{H} is the angular bisector of these two unit vectors. \mathbf{N} is the surface normal unit vector and A is the surface area.	35
2.6	The difference in phase angle (left) and distance (right) between the MMT data and the light curve simulation model, for a light curve of a CZ-2D rocket body (Norad ID: 41858) collected at 03:24:05 UTC on 07/01/2017. The percentage differences are 0.14 % (left) and 0.01 % (right).	42
2.7	The difference in position of the Observation vector, Illumination vector and Rotation vector, from the Earth, Velocity and Orbital Plane Normal vectors, between the orbital and body reference frames.	43
2.8	The difference in position of the Z body-fixed axis from the Earth, Velocity and Orbital Plane Normal vectors, between the orbital and body reference frames.	44

2.9	a) Actual light curve of an Atlas-Centaur II rocket body collected by the Advanced Electro-Optical System (AEOS) telescope on Maui Hawaii. b) Synthetic light curve generated by Watterer and Jah (Watterer and Jah, 2009).	44
2.10	The geometry model used to generate the validation example light curve.	45
2.11	A comparison synthetic light curve generated with the synthetic light curve model described in this chapter.	46
3.1	Simple geometry model used to represent a Globalstar satellite.	51
3.2	Average magnitudes and standard deviations of all light curves in the MMT database for the first group of 4 Globalstar satellites launched on the 8th of February 2000.	53
3.3	Average magnitudes and standard deviations of all light curves in the MMT database for the second group of 4 Globalstar satellites launched on the 29th May 2007.	54
3.4	Average magnitudes and standard deviations of all light curves in the MMT database for the third group of 4 Globalstar satellites launched on the 20th October 2007.	55
3.5	Average magnitudes and standard deviations of all light curves in the MMT database for all three groups of Globalstar satellites.	56
3.6	Average magnitudes and standard deviations of the real and synthetic light curves of the first group of Globalstar satellites. The inactive attitude state was simulated using a random orientation and a rotation period of 10 s.	57
3.7	Average magnitudes and standard deviations of the real and synthetic light curves of the first group of Globalstar satellites. The inactive attitude state was simulated using only a random orientation without rotation.	58
3.8	Average magnitudes and standard deviations of the real and synthetic light curves of the second group of Globalstar satellites.	59
3.9	Average magnitudes and standard deviations of the real and synthetic light curves of the third group of Globalstar satellites.	59
3.10	Average magnitudes and standard deviations of the real and synthetic light curves of all of the considered Globalstar satellites.	60
4.1	Rotation period, provided by the MMT database, against modified Julian date (MJD) for the selection of rocket bodies in Table 4.1.	66
4.2	The sum of the squared residuals of the least-squares fit between the polynomial and the rotation period data against number of degrees in the polynomial, for the orderly category.	67
4.3	The best-fitting polynomials for the objects in the orderly category.	68
4.4	The best-fitting polynomials for the 'fast de-spin' subcategory of the orderly category.	68
4.5	The best-fitting polynomials for the 'slow de-spin' subcategory of the orderly category.	69
4.6	The sum of the squared residuals of the least-squares fit between the polynomial and the rotation period data against number of degrees in the polynomial, for the complex category.	69
4.7	The best-fitting polynomials for the objects in the complex category.	70

4.8	The dispersion in the best-fitting period of the light curves in the ‘fast de-spin’ subcategory.	71
4.9	The PDM result for the light curve of object CZ-2D (41858) with the lowest dispersion.	73
4.10	The PDM result for the light curve of object CZ-2D (41858) with the highest dispersion.	73
4.11	An erroneous PDM result for the light curve of object CZ-4C (40879) with the lowest dispersion.	74
4.12	A corrected PDM result for the light curve of object CZ-4C (40879) with the lowest dispersion.	74
4.13	The PDM result for the light curve of object CZ-4C (40879) with the highest dispersion.	75
4.14	The dispersion in the best-fitting period of the light curves in the ‘slow de-spin’ subcategory.	76
4.15	The PDM result for the light curve of object ATLAS-3B-CENTAUR (39475) with the lowest dispersion.	76
4.16	The PDM result for the light curve of object ATLAS-3B-CENTAUR (39475) with the highest dispersion.	77
4.17	The PDM result for the light curve of object CZ-4 (25942) with the lowest dispersion.	77
4.18	The PDM result for the light curve of object CZ-4 (25942) with the highest dispersion.	78
4.19	The dispersion in the best-fitting period of the light curves in the ‘complex’ category.	78
4.20	The PDM result for the light curve of object CZ-3B (38253) with the lowest dispersion.	79
4.21	An erroneous PDM result for the light curve of object CZ-3B (38253) with the lowest dispersion.	79
4.22	A corrected PDM result for the light curve of object CZ-3B (38253) with the lowest dispersion.	80
4.23	The perigee altitude, apogee altitude and eccentricity of all objects.	82
4.24	The perigee altitude, apogee altitude and eccentricity of CZ rocket bodies in the MMT database.	84
4.25	The perigee altitude, apogee altitude and eccentricity of DELTA rocket bodies in the MMT database.	86
4.26	The perigee altitude, apogee altitude and eccentricity of ATLAS rocket bodies in the MMT database.	87
5.1	The five geometry models used in this chapter.	92
5.2	Illumination geometry used to quantify the change in a light curve that results from increasing the phase angle θ	93
5.3	The change in a synthetic light curve, quantified using RMSE, that resulted from increasing the phase angle between the illumination and observation vectors. A entirely diffuse reflection model was used.	94
5.4	The change in a synthetic light curve, quantified using RMSE, that resulted from increasing the phase angle between the illumination and observation vectors. A entirely specular reflection model was used.	95
5.5	The change in a synthetic light curve, quantified using RMSE, that resulted from decreasing the BRDF parameter omega (diffuse albedo).	96

5.6	The change in a synthetic light curve, quantified using RMSE, that resulted from decreasing the BRDF parameter d (ratio of diffuse reflection).	97
5.7	The change in a synthetic light curve, quantified using RMSE, that resulted from decreasing the BRDF roughness parameter m .	98
5.8	The change in a synthetic light curve, using entirely specular reflection and quantified using RMSE, that resulted from increasing the BRDF ω (diffuse albedo).	98
5.9	The change in a synthetic light curve resulting from changes in attitude state under entirely diffuse reflection.	99
5.10	The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.	100
5.11	The change in a synthetic light curve resulting from changes in attitude under diffuse reflection.	101
5.12	The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.	102
6.1	The three object geometries considered in this paper. A 'b' subscript denotes a body-fixed axis.	106
6.2	A figure showing the reflection geometry and object geometry. The observation vector, \mathbf{V} , is aligned with the X-axis of the system and the position of the illumination vector, \mathbf{L} , is defined by the angles ψ and ϕ .	107
6.3	An example polar plot for the cube geometry with a diffuse reflection and a rotational motion about the Y body-fixed axis.	108
6.4	Results for a diffuse reflection model and a cubic object geometry. RMSE is measured from a no-rotation no illumination offset baseline.	110
6.5	The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.4. The two rotation cases in question are labelled above each polar plot. The reflection model is diffuse and the object geometry is a cube.	111
6.6	Results for a specular reflection model and a cubic object geometry. RMSE is measured from a no-rotation no illumination offset baseline.	112
6.7	The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.6. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a cube.	112
6.8	Results for a diffuse reflection model and a icosahedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.	113
6.9	The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.8. The two rotation cases in question are labelled above each polar plot. The reflection model is diffuse and the object geometry is a icosahedron.	113
6.10	Results for a specular reflection model and a icosahedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.	115
6.11	The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.10. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a icosahedron.	115

6.12	Results for a diffuse reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.	116
6.13	The result of performing a subtraction between two of the plots from Figure 6.12. The two rotation cases in question are labelled above each polar plot. The reflection model is diffuse and the object geometry is a 5-frequency subdivided polyhedron.	116
6.14	Results for a specular reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.	117
6.15	The result of performing a subtraction between two of the plots from Figure 6.14. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a 5-frequency subdivided polyhedron	117
6.16	The initial and final illumination conditions used in the first light curve scenario. \mathbf{L} is the illumination vector and \mathbf{V} is the observation vector. The curved dotted line shows the outline of the Earth.	119
6.17	Light curves for the four rotation states considered in this paper, using the illumination geometry presented in Figure 6.16. All light curves, with the exception of the 'z-rot' case are superimposed.	119
6.18	The initial and final illumination conditions used in the second light curve scenario. \mathbf{L} is the illumination vector and \mathbf{V} is the observation vector. The curved dotted line shows the outline of the Earth.	119
6.19	Light curves for the four rotation states considered in this paper, using the illumination geometry presented in Figure 6.18.	120
7.1	Six examples of real rocket body light curves collected using MMT. Red 'A', 'B' and 'C' letters are used to highlight repeating sequences of reflections at three distinct brightness levels where applicable.	125
7.2	A CZ-2D launch vehicle with the upper stage that remains on-orbit indicated by a red box.	126
7.3	Object geometry used to represent a rocket body. The red, blue and green facets indicate the three zones of the geometry, which are given different reflection characteristics.	128
7.4	The orbital reference frame used to define the angular velocity vector. The angles ϕ and ψ describe the direction of the angular velocity vector with respect to the Earth zenith direction.	128
7.5	Four synthetic target light curves.	129
7.6	The RMSE measured between the synthetic target light curves and new synthetic light curves generated with a angular velocity vector direction given by the polar axes.	131
7.7	The RMSE measured between the real light curve from the MMT and the best-fitting synthetic light curve generated with a new angular velocity vector given by the polar axes.	131
7.8	Comparison between the real MMT light curve and the best-fitting synthetic light curve. This plot shows the first two rotation periods.	132
7.9	Top: Comparison between the real MMT light curve and the best-fitting synthetic light curve. This plot shows the full light curve. Bottom: Absolute residuals.	132

8.1	A render of a Starlink satellite on-orbit, generated by SpaceX and available at: www.starlink.com .	139
8.2	A render of the two geometry configurations used by Starlink satellites. Available at: https://www.spacex.com/updates/starlink-update-04-28-2020/ .	139
8.3	A render showing how the ‘DarkSat’ experiment aimed to reduce the brightness of Starlink satellites in a shark-fin configuration. Available at: https://www.spacex.com/updates/starlink-update-04-28-2020/ .	141
8.4	A render showing how the ‘VisorSat’ experiment aimed to reduce the brightness of Starlink satellites in a shark-fin configuration. Available at: https://www.spacex.com/updates/starlink-update-04-28-2020/ .	141
8.5	Apparent magnitude and orbital altitude, against date, for the first 9 batches of Starlink spacecraft.	143
8.6	Apparent magnitude and orbital altitude, against date, for the first 9 batches of Starlink spacecraft. Data has been filtered to show only the data for spacecraft whilst on-station.	143
8.7	The two geometries used to represent a Starlink system. On the left is the full spacecraft, including solar panel, and on the right is the payload-only geometry.	145
8.8	Minimum, maximum and average synthetic light curve brightness against solar panel involvement.	145
8.9	Minimum, maximum and average synthetic light curve brightness against solar panel shadowing.	147
8.10	Minimum, maximum and average real light curve brightness against solar panel involvement predicted in the synthetic data.	148
8.11	Minimum, maximum and average real light curve brightness against solar panel shadowing predicted in the synthetic data.	148
8.12	Average light curve brightness of all available data in the MMT database for the first 9 batches of operational Starlink satellites.	149
Appendix A.1	The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.	163
Appendix A.2	The baseline light curve from Figure A.1 (black) and the approximated signal generated from the Fourier coefficients determined for the baseline light curve (red).	164
Appendix A.3	A light curve from Figure A.1 with a rotation about the body-fixed x -axis of 0.2 RPM (black), and the approximated signal generated from the Fourier coefficients determined for this light curve (red).	165
Appendix A.4	The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.	166
Appendix A.5	<i>Results for a specular reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.</i>	167
Appendix A.6	<i>The result of performing a subtraction between two of the plots from Figure 6.14. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a 5-frequency subdivided polyhedron</i>	167

Appendix A.7 Results for a specular reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.	168
Appendix A.8 The result of performing a subtraction between two of the plots from Figure 6.14. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a 5-frequency subdivided polyhedron	169

List of Tables

2.1	The input parameters used to generate the synthetic light curve of Figure 2.11.	47
3.1	Table of data on three groups of Globalstar spacecraft	52
3.2	Globalstar spacecraft state according to either the MMT database or Celestrak catalogue.	52
4.1	The initial set of rocket bodies analysed in this chapter.	65
4.2	The categorisation statistics of the CZ rocket bodies plotted in Figure 4.24	83
4.3	The categorisation statistics of the DELTA rocket bodies plotted in Figure 4.25	85
4.4	The categorisation statistics of the ATLAS rocket bodies plotted in Figure 4.26	86
5.1	The RMSE results for each of the plots shown in this section.	103
7.1	The properties of the real target light curve.	125
7.2	Table of randomly generated inputs used to generate the synthetic target light curves of Figure 7.5.	129
7.3	Best-fitting inputs to fit a synthetic light curve to the real target light curve.	130
8.1	Data of the Starlink batches considered in this chapter	142

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission

Signed:.....

Date:.....

Acknowledgements

I would like to express my gratitude to my supervisors Prof. Hugh Lewis and Dr. Hodei Urrutxua. The value of the mentorship and supervision of Prof. Lewis, throughout this PhD and before, cannot be quantified and it is no exaggeration to say I would not be here today without his support. Thank you Hugh. The mentorship and support of Dr. Urrutxua during this PhD has also been invaluable and I cannot imagine doing it without him. Thank you Hodei. I could not have asked for a better supervisory team and I am extremely grateful to you both.

I would also like to thank all of the people that I have met at conferences and meetings throughout working on this PhD for their kindness and encouragement. I am extremely proud and honoured to be associated with this incredible group of people.

Finally, but certainly not least, I would like to thank my parents. To say that none of this would have been possible without their support is an understatement.

Definitions and Abbreviations

AAS	American Astronomical Society
ACI	Astrodynamics Community of Interest
ADR	Active Debris Removal
AEOS	Advanced Electro-Optical System
AIAA	American Institute of Aeronautics and Astronautics
AIUB	Astronomical Institute of the University of Bern
ASAT	Anti-Satellite
BMEWS	Ballistic Missile Early Warning System
BNSC	British National Space Centre
BRDF	Bidirectional Reflectance Distribution Function
CGI	Computer Generated Imagery
CNSA	China National Space Administration
COPUOS	Committee on the Peaceful Uses of Outer Space
CRS	Commercial Resupply Services
CTIO	Cerro Toledo Inter-American Observatory
CZ	Chang Zheng (Long March)
DAMAGE	Debris Analysis and Monitoring Architecture to the Geosynchronous Environment
DSN	Deep Space Network
DSTL	Defence Science and Technology Laboratory
EOL	End of Life
EOP	Earth Orientation Parameter
ESA	European Space Agency
EURECA	European Retrievable Carrier
GCRF	Geocentric Celestial Reference Frame
GEO	Geostationary Earth Orbit
GEODSS	Ground-based Electro-Optical Deep Space Surveillance System
GNC	Guidance, Navigation, and Control
GNSS	Global Navigation Satellite System
GSSR	Goldstone Solar System Radar
HAX	Haystack Auxiliary Radar
HST	Hubble Space Telescope

HTG	Hyperschall Technologie Gottingen GmbH
IADC	Inter-Agency Debris Coordination Committee
ICRF	International Celestial Reference Frame
IERS	International Earth Rotation and Reference Systems Service
ITRF	International Terrestrial Reference Frame
iOTA	In-Orbit Tumbling Analysis
ISAR	Inverse Synthetic Aperture Radar
ISS	International Space Station
ITU	International Telecommunication Union
JAXA	Japan Aerospace Exploration Agency
JIT	Just in Time
JPL	Jet Propulsion Laboratory
LDEF	Long Duration Exposure Facility
LEGEND	LEO-GEO environment debris model
LEO	Low Earth Orbit
LISP	Laser Impulse Space Propulsion
LUCA	Long-term Utility for Collision Analysis
MC	Monte Carlo
MCAT	Meter Class Autonomous Telescope
MEDEE	Modelling the Evolution of Debris in the Earth Environment
MEO	Medium Earth Orbit
MJD	Modified Julian Date
MMT	Mini-Mega TORTOTA wide-field monitoring system
MODEST	Michigan Orbital DEbris Survey Telescope
NASA	National Aeronautics and Space Administration
NORAD	North American Aerospace Defense Command
ORDEM	Orbital Debris Engineering Model
PDM	Phase Dispersion Minimisation
PMD	Post Mission Disposal
RCS	Radar Cross Section
RMSE	Root Mean Squared Error
RPM	Rotations per minute
RPO	Rendezvous and Proximity Operations
RSO	Resident Space Object
SDMG	Space Debris Mitigation Guidelines
SDP	Simple Deep-space Perturbations
SDS	Space Debris Sensor
SFU	Space Flyer Unit
SGP	Simple General Perturbations
SLR	Satellite Laser Ranging
SMARTS	Small and Moderate Aperture Research Telescope System

SOFA	Standards of Fundamental Astronomy
SRP	Solar Radiation Pressure
SSA	Space Situational Awareness
SSC	Surrey Space Centre
SSN	Space Surveillance Network
SST	Space Surveillance and Tracking
SSTL	Surrey Satellite Technology Limited
TDB	Barycentric Dynamical Time (Temps Dynamique Barycentrique)
TEME	True Equator Mean Equinox
TLE	Two-Line Element
TT	Terrestrial Time
UKF	Unscented Kalman filter
UKSA	United Kingdom Space Agency
UN	United Nations
UT1	Universal Time 1
UTC	Universal Time Coordinated

Chapter 1

Introduction

1.1 The Orbital Environment

1.1.1 Overview

Reaching Earth orbit requires a huge expenditure of energy. In the early days of spaceflight, rockets and satellites carried only enough fuel to complete their objective and so did not have the capability to remove themselves from Earth orbit once they reached the end of their lifetime. Historically, objects were made use of for as long as possible, then left to drift on-orbit. However, more recently, space systems use the last of their propellant to either de-orbit or move into shorter lifetime orbits.

As a result of this, approximately 2000 rocket bodies, 2000 bits of mission related debris and 5000 spacecraft (both functional and non-functional) were in Earth orbit at the start of 2020 ([Liou et al., 2020](#)).

The primary mechanism by which objects are naturally removed from the Earth orbital environment is atmospheric drag. However, because atmospheric drag decreases with altitude, this is only applicable to objects in Low Earth Orbit (LEO). For Medium Earth Orbit (MEO) and upwards, an object must modify its orbit at End of Life (EOL), either to increase the atmospheric drag or to move into a less congested region of space. Alternatively, new strategies can be devised such as using the orbital perturbations of Solar Radiation Pressure (SRP) and the J_2 term to deorbit objects in MEO ([Lücking et al., 2012](#); [Rossi et al., 2017](#)).

In the last 60 years, the rate at which objects have been introduced into the environment has outpaced these removal mechanisms. Therefore, the total mass of artificial objects in Earth orbit has, on average, increased continually since the launch of the first artificial objects in the late 1950s. Figure 1.1 shows the mass on orbit, per month, for four different types of object: spacecraft, rocket bodies, fragmentation

debris and mission-related debris. This figure shows the imbalance between the rate at which new objects have been introduced and the rate at which they have been removed by drag, which has resulted in a net increase in the total orbital mass. This figure also shows that the vast majority of the mass of artificial objects in Earth orbit comprises intact spacecraft and rocket bodies.

In addition to the intact spacecraft and rocket bodies, the Earth orbital environment also contains fragmentation debris, (objects that were once part of a larger object), and mission-related debris, (objects such as adaptor rings and payload shrouds that were intentionally released from a larger object as part of nominal operations).

In the early days of spaceflight, there was little concern paid to the risk to on-going space operations posed by discarded space objects. This was because the small number of objects occupied a huge volume. However, as this number of objects increases, it becomes increasingly likely that these may eventually collide with one another. The relative velocities between objects are often measured in kilometres-per-second and such hyper-velocity collisions can result in catastrophic destruction. Even a collision with a relatively small (eg. tennis-ball-sized) piece of fragmentation debris could lead to the destruction of an expensive space mission.

Safeguarding spacecraft against collisions has become a normal aspect of day-to-day space operations. Achieving this requires a good understanding of the positions, and good predictions of the future positions, of all space objects in Earth orbit. Detecting and tracking the space object population is known as Space Surveillance and Tracking (SST).

1.1.2 Detection of Space Objects

Our understanding of the positions and motions of space objects is generated through applying mathematical models to collections of observations. Hence, performing observations of space objects is one of the most crucial aspects of modern day-to-day space operations. The detection and measurement of space objects is primarily achieved in three ways; observations with optical telescopes, observations with radar installations and in-situ measurements, with each having specific object sizes and altitude ranges for which they are most useful.

Although many nations have some means of performing SST, the most capable by far is the United States of America (USA). It's Space Surveillance Network (SSN) comprises a number of optical telescopes and radar installations. The primary objectives of the SSN are to maintain a space object catalogue, including the search for new as yet untracked objects, and to identify, characterise and monitor active space systems (Sridharan and Pensa, 1998).




FIGURE 1.1: Monthly mass of objects in the United States Space Surveillance Network (SSN) catalogue (Liou et al., 2020).

1.1.2.1 Radar

Radar installations form an invaluable component of the SSN. Radars have a number of characteristics that make them well suited for observing space objects. Firstly, radar is unaffected by illumination conditions and so can operate in both daytime and nighttime. Secondly, radar is unaffected by weather and so does not require a clear sky. When compared to an optical system, a radar system therefore has a considerably higher uptime. A further advantage provided by radar is the availability of a beam-park operation mode. In this mode the radar energy is emitted into a fixed direction and objects that intersect the beam are detected. Although this does not allow for determination of a predicted orbit, it does allow surveying of the debris populations down to smaller object sizes, which can then be used to derive statistics for the debris population (NASA, 2008).

A number of radar techniques are employed, they generally fall into one of two broad categories. The first is a mechanical tracking radar, which uses a large movable parabolic dish to track a particular object as it moves across the sky. One example of a mechanical tracking radar used within the SSN is the Haystack radar, a 37 m steerable parabolic dish located in Massachusetts, USA. Haystack is capable of detecting objects as small as 5 mm, at ranges that equate to an orbit of approximately 400 km altitude. An image of Haystack is shown in Figure 1.2 Another example of a major mechanical tracking radar component of the SSN is the Goldstone Solar System Radar (GSSR), which has a 70 m antenna. GSSR is the most sensitive component of the SSN and is

able to detect objects as small as 2 - 3 mm at ranges that equate to an orbit of approximately 400 km altitude. However, because GSSR is part of the Deep Space Network (DSN) the time it spends providing observations to the SSN is limited (NASA, 1996; Stokely et al., 2006).



FIGURE 1.2: The Haystack radar complex (NASA, 2008).

The second broad category of radar is a phased-array. Unlike mechanical tracking techniques, phased-array radars have no moving parts and are instead steered electronically. This type of system provides a key advantage over a mechanical tracker: without the limitation of having to perform any mechanical movements, the system can scan large areas of the sky in fractions of a second and so is capable of maintaining tracks on multiple objects simultaneously. However, phased-array radars are much more expensive due to the large number of complex antenna elements that comprise the sensor array. An example of a phased-array radar is the Royal Air Force (RAF) Fylingdales radar in the North Yorkshire Moors in the north of England. Although Fylingdales' primary function is as a Ballistic Missile Early Warning System (BMEWS), it is also used within the SSN for tracking space objects (NASA, 2008).

For the reasons that have been mentioned above, radars are ideal components within an SST framework. However, they have an important downside: they are extremely costly, both in terms of construction and operation. Although mechanical tracking radar systems have a number of areas in which they outperform optical telescopes, they are substantially more expensive. The additional costs are larger still for phased-array radar systems.

1.1.2.2 Optical

Alongside radar, the other major components of the SSN are optical telescopes. Unlike radar, optical systems are subject to more constraints in their operation. Most



FIGURE 1.3: The RAF Fylingdales phased array radar. From: https://en.wikipedia.org/wiki/RAF_Fylingdales.

obviously, the telescope can only be used at night-time and when the weather is clear. Additionally, the object must be illuminated by the Sun. Another downside of an optical system is that its performance is highly dependent on the light pollution in the surrounding environment and so they are typically located in remote locations away from large population centres, which makes operation more inconvenient. The reason why optical systems are included within the SSN is because at large ranges, they outperform radar in terms of sensitivity. As a result, observations of objects in Medium Earth Orbit (MEO) up to Geostationary Earth Orbit (GEO) are primarily carried out using optical telescopes.

The first major optical components of the SSN were the three sets of three 1 m telescopes, located at each of the following sites; Socorro, New Mexico; Maui, Hawaii; and Diego Garcia, British Indian Ocean Territories, giving a total of nine telescopes. One of these sites is shown in Figure 1.4. Together, this network is the Ground-based Electro-Optical Deep Space Surveillance System (GEODSS). The GEODSS network is responsible for tracking the majority of so-called deep space objects, which have orbital periods longer than 225 mins (NASA, 2008), and can track objects of approximately 1 m and larger in GEO.

Another optical member of the SSN is the Michigan Orbital DEbris Survey Telescope (MODEST), a 0.61 m optical system at the Cerro Tololo Inter-American Observatory (CTIO) in Chile. MODEST is tasked with statistical mapping of GEO and is able to detect (but not track) GEO objects smaller than 20 cm in diameter (Seitzer et al., 2004). The 0.9 m Small & Moderate Aperture Research Telescope System (SMARTS), also located at CTIO, assists MODEST with follow-up observations for orbit determination. Both these telescopes are shown in Figure 1.5. One final example, that is not part of the SSN, is the Meter Class Autonomous Telescope (MCAT), which is a 1.3 m system located on Ascension Island. MCAT is regularly used for observing



FIGURE 1.4: One of the three GEODSS sites (NASA, 2008).

artificial space objects and is able to observe objects as small as 13.5 cm at GEO ranges (Lederer et al., 2013).



FIGURE 1.5: The MODEST (front centre) and the SMARTS (left) telescopes at CTIO (NASA, 2008).

1.1.2.3 In-Situ

In-situ measurements, although not a part of the SSN, have been an invaluable technique for understanding the Earth's orbiting population of space objects. This is because these measurements give the only insight into the numbers of the smallest objects that cannot be reliably detected from the ground. However, the handful of in-situ data sources pales in comparison to the huge number of observations that have been carried out globally using optical and radar systems.

An in-situ measurement is performed by exposing some surface to the space environment for a period of time. By examining the surface for impacts, and taking into consideration the orbit and attitude that the surface was in, estimations of particle flux can be derived. In the search for micro-meteoroids, in-situ measurements were first carried out early in the space age with acoustic sensors and penetration cells being flown on the Explorer-1 and Pioneer-1 missions, the first satellites flown by the

USA and NASA respectively. Subsequently, any other surfaces returned from space, such as the Gemini capsules, were also analysed to measure meteoroid fluxes.

Once the space shuttle entered operation it became much easier to return surfaces from space, such as the solar panels from the Hubble Space Telescope (HST), all of which were examined for impacts. Later, the space shuttle surfaces themselves were also used to make measurements of the spatial density of small objects.

The first major in-situ measurement platform was the Long-Duration Exposure Facility (LDEF) which was launched in 1984 and collected by the space shuttle in 1990 after 6 years on-orbit. Due to the size, fixed orientations of the various surfaces and duration spent on-orbit, the LDEF has been an invaluable resource for the measurements of small debris objects and micro-meteoroids. The LDEF platform is shown in Figure 1.6.



FIGURE 1.6: A picture of the LDEF platform on-orbit taken from the space shuttle (NASA, 2008).

Following the success of LDEF, ESA and the Japan Aerospace Exploration Agency (JAXA) conducted the European Retrievable Carrier (EURECA) and Space Flyer Unit (SFU) missions respectively. EURECA spent a year on-orbit between 1992 and 1993, and the SFU was on-orbit between 1995 and 1996 (NASA, 2008).

The in-situ experiments discussed above are well over 20 years old. More recently, the International Space Station (ISS) has been used as a platform for debris measurement. On December 4th 2017, NASA launched the Space Debris Sensor (SDS) to the ISS on a cargo resupply mission. It was intended for the SDS to stay on-orbit for 2 - 3 years, measuring small debris impacts using a combination of dual-layer films, acoustic

sensors and resistive grid sensors (Hamilton, 2017). Unfortunately the SDS suffered an unrecoverable anomaly after 26 days aboard the ISS (Anz-Meador et al., 2019).

1.1.3 Space Object Catalogue

The various observations made by the SSN are used to generate a space object catalogue. This catalogue stores information such as position, velocity, Radar Cross Section (RCS, where available) and ballistic coefficient of all objects that are detected by the SSN. Using orbital propagation, these observations can then be used to predict future states of the catalogued objects. However, even when using high accuracy orbital propagation techniques, errors will accumulate due to unknown parameters, such as the future atmospheric density at each point of the object's orbit. An object's orbit can therefore be propagated only so far into the future before the differences between the real and propagated orbits are so large that the object can no longer be found when pointing a sensor towards its expected position. The SSN is therefore required to perform catalogue maintenance, where all objects in the catalogue are observed with sufficient frequently so that their positions are updated regularly enough that they do not get lost to the network.

1.1.3.1 SGP and the TLE format

When the SSN was still in its infancy, a method had to be devised for efficiently storing and sharing observation data and performing orbital propagation. However, this had to be achieved using the computation, data storage and data transfer hardware of the 1960s. The system that the USA designed for this task was the Simple General Perturbations (SGP) theory and the Two-Line-Element (TLE) format. A TLE is two lines of 69 characters, which describe the orbit of a particular object at a particular time.

Finding the orbit that best-fits a series of observations (typically optical observations are right ascension and declination, and typical radar observations are elevation, azimuth, range and range rate) is a process called orbit determination. The USA performs this process in a way that exploits all of the available accuracy and precision of the SSN measurement data. However, what gets released publicly is a TLE, which is generated by finding the best-fitting orbit generated using the SGP theory. Unlike a special perturbations theory, which performs the perturbation calculations numerically at each time step, SGP theory uses analytic perturbation theories and truncated force models. Furthermore, the final orbital information can only use a maximum of two lines of 69 characters and so the number of decimal places with which each parameter can be presented is limited. As a result, SGP theory is less accurate than special perturbations techniques, but the calculations are significantly

less expensive computationally, which was a critical consideration at the time of the system development.

Although the accuracy of a TLE is reasonable at the time it is generated (typically 1 km uncertainty at the generation epoch, although this greatly depends on the frequency and quality of the underlying observation data), the positional accuracy with which the object's future position can be predicted using SGP deteriorates with time separation from the TLE creation epoch. This degradation in positional accuracy is caused by the inherent inaccuracies of the SGP model, in addition to unknown parameters. One such parameter is atmospheric density, which is highly dependent on solar activity that cannot be accurately forecast ([Emmert et al., 2014](#)). Hence new TLEs must be generated continually.

An important property of the SGP / TLE system is that the orbit-fitting process provides mean orbital elements, where periodic effects are averaged out, rather than the true instantaneous values. As a result, to convert a TLE back into a state vector, the same version of the SGP model must be used, to ensure that the periodic effects are reinstated correctly. However, because SGP and TLEs are used operationally by the US military, it is usually uncertain exactly which code was used to generate a TLE as this information is not publicly released. For this reason, and to improve compatibility between users of TLE data, the source code of the SGP models was made public as part of Space-Track Report No. 3, which was released in 1980. Space-Track is the online database through which TLE data is provided to the public. The report contained source code for five models, SGP, SGP4, SGP8, SDP4 and SDP8. The SDP (Simple Deep-space Perturbations) codes are used for objects in orbits with periods longer than 225 minutes, with SGP4 used for anything less. The SGP8 and SDP8 codes were developed to better handle re-entering objects. However, there has been no evidence to suggest that they have ever been used in the generation of a catalogue TLE ([Hoots and Roehrich, 1980](#)).

In the years following the 1980 public release of SGP codes, it was suspected that changes were made to the actual operational code used by the USA to generate TLEs. It was expected that these changes introduced additional uncertainties into data generated using the public version of the code. In 1996, a newer version was released by NASA after they were provided the code by Space-Track in support of a climate science mission. This confirmed that changes had been made and that the public and operational SGP codes were inconsistent. These were the motivations behind the 2006 revisiting of the Space-Track Report by [Vallado et al. \(2006\)](#), which was the first comprehensive report, complete with analysis and results, on SGP codes compatible with the TLE database.

1.1.4 The Space Object Population

The population of space objects is a combination of two categories: tracked objects whose position is known well enough to be catalogued, and objects that cannot be tracked and therefore cannot be catalogued. The limiting factor for whether or not an object can be tracked is the sensitivity of the SSN, which is only capable of forming tracks for objects larger than ~ 10 cm. Objects smaller than this may be detectable to the SSN (and for some their orbit can be determined depending on the object's characteristics), but mostly no best-fitting orbital parameters can be resolved. Generally only estimates of the numbers of these objects can be made. The best estimates of the total population are generated by combining data on tracked objects with the statistical measurements obtained using surveying observations and in-situ measurements (NASA, 2008).

1.1.4.1 Tracked Objects

The graph in Figure 1.7 shows the average monthly number of objects in the public TLE catalogue by year. The figure shows four categories of object; rocket bodies; intact spacecraft, mission-related debris and fragmentation debris. This plot shows that there is an increasing trend in the number of objects. When compared to Figure 1.1, which showed that the majority of mass on-orbit is contained within rocket bodies and spacecraft, Figure 1.7 shows that most objects on-orbit are fragmentation debris. It should be noted that this plot does not show when each new object entered the catalogue. Instead, it has been edited by NASA to show when the object was expected to have entered the environment. For example, when a fragmentation occurred, all the resulting objects are shown to have appeared on the day of the fragmentation, rather than being spread across the days/months over which they were gradually detected, tracked and catalogued by the SSN.

In Figure 1.7 two large jumps in the number of fragments occurred in 2007 and 2009. The first of these is due to a Chinese Anti-Satellite (ASAT) test conducted on January 11, 2007 in which a 750 kg weather satellite, Fengyun-1C, was intentionally impacted by a ground-launched kinetic projectile. As a result of this intentional collision, more than 3300 tracked fragments were generated. This action has been universally condemned by the international community as highly irresponsible, in particular because the high altitude of the target object means that many of these fragments will remain on-orbit for many years (Johnson et al., 2008). The second of these two events is a February 2009 collision between a defunct Russian communications satellite, Cosmos 2251 (900 kg), and an operational commercial communications satellite, Iridium 33 (556 kg). This was the first recorded collision between two intact spacecraft and resulted in approximately 1850 newly-tracked fragments (Wang, 2010).




FIGURE 1.7: Monthly number of objects in the United States Space Surveillance Network (SSN) catalogue (Liou et al., 2020).

Figure 1.8 shows the spatial distribution of objects in the LEO environment. This figure plots the number of objects in a 20 kilometre altitude bin against altitude. The blue line plots the population as it was in 2007 and the red line plots the population in late 2010, after the two collision events described above. This Figure shows that the spatial density peaks between 700 and 900 km. These are primarily sun-synchronous orbits, which are highly valuable for Earth observation and surveillance systems. The plot also shows the change in spatial density that resulted from the Fengyun ASAT test and the Iridium-Cosmos collision. The figure shows that these two major fragmentation events occurred in the most congested region of space and caused large increases in the peak spatial density (Liou, 2017).

1.1.4.2 Untracked Objects

Untracked objects are those objects that are either not observable to the SSN, or are not observed reliably enough by the SSN to generate predicted orbits and so do not become members of the space object catalogue. Typically, this is because the object is too small for the network to track. Figure 1.9 shows the various sources of information that are used to estimate the numbers of untracked space objects. The graph plots the logarithm of the object diameter in centimetres against the logarithm of cross-sectional flux in number per square meter per year. In Figure 1.9, the data for objects with sizes of 10 cm and larger is from the TLE catalogue. The data from 10 cm down to 0.1 cm is generated from radar detections, primarily from Goldstone and Haystack. Finally, the




FIGURE 1.8: Spatial distribution with altitude of objects in the TLE catalogue (Liou, 2017).

data below 0.1 cm is generated from in-situ surfaces including the LDEF and EURECA experiments. This Figure shows that flux increases as size decreases and that there are orders of magnitude more untracked objects than there are objects in the catalogue.

In the not too distant future, the number of tracked objects will increase dramatically as the Space Fence System comes online. The Space Fence is a very large and powerful phased array radar system being developed by the USA that may extend the smallest trackable object size down to a few centimetres (Haimenl and Fonder, 2015). The Space Fence, as of 2020, is operational but the objects being detected have not yet entered the public catalogue.

1.2 Debris Modelling

Observations and in-situ measurements of the space environment show that there are large numbers of space objects in Earth orbit. Many of these objects can be tracked and the future position predicted; this is approximately 20,000 objects in total (excluding classified objects, that do not appear in the catalogues), of which 2000 are rocket bodies, 2000 is mission related debris, 5000 are spacecraft and 11,000 is fragmentation debris. However, the vast majority cannot be reliably tracked. Understanding this space object environment and the risks it may pose to space operations and to the

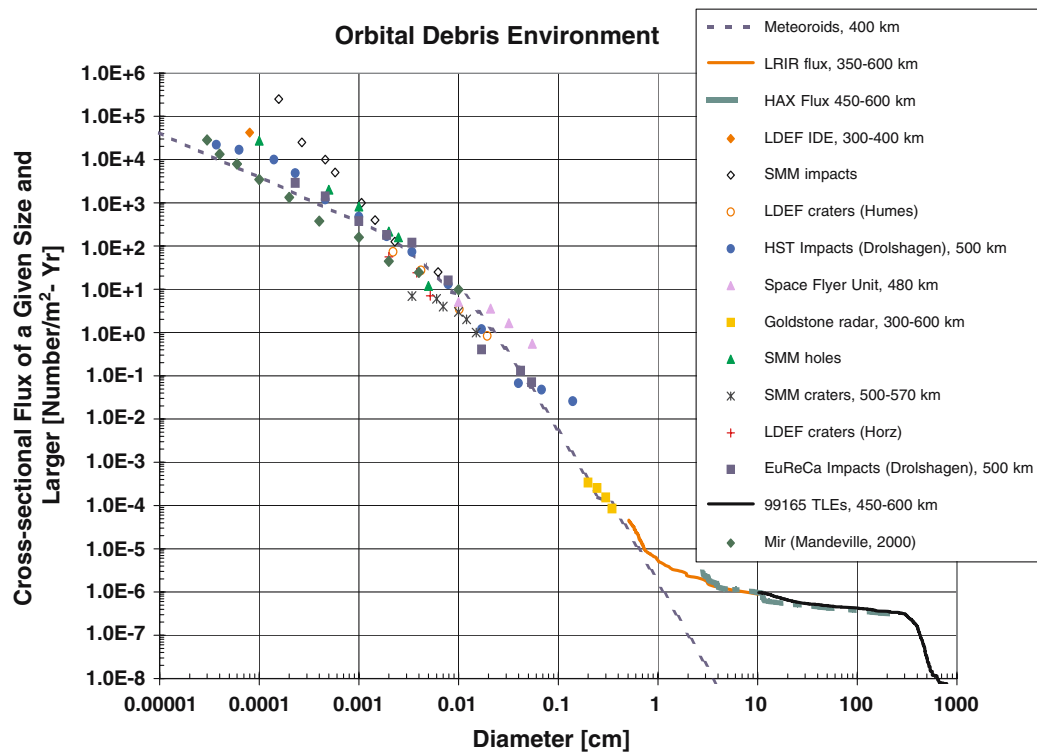


FIGURE 1.9: Combined radar and in-situ measurements of debris flux (Schildknecht, 2007)

sustainability of the orbital environment is the motivation behind space object population modelling.

The term ‘space debris’ is used to refer to artificial objects that are no longer performing their mission and are no longer being controlled. Space debris can therefore be used to refer to all of the artificial objects on orbit with the exception of the approximately 2500 active satellites.

1.2.1 History of Debris Modelling

Space debris models have two, closely connected, primary uses. The first and most obvious is to better understand the space debris environment. The results of these studies form the foundation for the development of space operations ‘best-practice’, such as the space debris mitigation guidelines introduced by the Inter-Agency Debris Coordination Committee (IADC). The second, is to provide a predictive capability to better understand how the space debris environment may evolve into the future. This capability can be used to identify actions, to be taken in the present, that may ensure the short term (10 - 100 years) and long term (100+ years) sustainable use of the space environment.

As an additional application, space debris models can be used to make safety assessments for specific space missions operating within the environment. The results of these analyses feed back into spacecraft design and orbit selection. The models can also be used by regulators to assess the risks to a proposed mission and to existing missions, and so inform a licencing decision.

NASA's first debris model was developed in the late 1970s and was known as AVERAGE. This was the first unification of a simple break up model, a simple orbital propagator and an intact LEO population of objects, together with averaged explosion and collision rates. The capabilities demonstrated by AVERAGE motivated a commitment by NASA to develop the more robust EVOLVE model (NASA, 2008).

Starting with EVOLVE, a subset of example debris models are described over the subsequent subsections.

1.2.1.1 EVOLVE

The objective for the EVOLVE model, first developed in the late 1980s, was to increase the fidelity of AVERAGE. Additionally, EVOLVE was required to make estimates of future fragmentations, which was achieved through implementation of a Monte Carlo (MC) simulation approach. The development of EVOLVE continued throughout the 1990s incorporating additional data from observations and ground-based experimentation as it became available. In particular, by comparing propagations to real data, the orbital propagator was adapted to include higher order gravitational terms, luni-solar and solar radiation pressure perturbations, an oblate Earth and rotating atmosphere. Further advances were made to the breakup model to include delta-velocity distributions of generated fragments. Due in part to limitations in computation time, EVOLVE was a 1-dimensional model in which collisions were assigned to spherical altitude shells between 200 - 2000 km (NASA, 2008).

1.2.1.2 ORDEM

An engineering debris model is used by spacecraft operators as a relatively fast and user-friendly way to calculate the debris flux on spacecraft and instrument surfaces. Engineering models have very different requirements to evolutionary models such as EVOLVE and AVERAGE. Additionally, different operators have varying requirements based on their specific mission. A space debris engineering model must accurately include debris distributions as functions of altitude, latitude and debris size. Engineering models like ORDEM are used by spacecraft operators and designers to perform initial analyses on the debris environments that may be experienced by a spacecraft in a particular orbit. ORDEM 3.0 for example calculates debris flux from 10

cm to 1 m in GEO and from 10 μm up to 1 m elsewhere, from 2010 through to 2035 (Krisko, 2014).

1.2.1.3 LEGEND

In 2001 NASA initiated work on the replacement to the EVOLVE model. This was motivated by two factors; firstly, the 1-dimensional approach of EVOLVE was deemed to be insufficient to adequately characterise the spatial density of debris; secondly, a full-scale debris model (one that describes the near-Earth orbital environment from LEO past GEO) was required in order to accurately describe the growing debris population. It was decided that a high-fidelity 3-dimensional model was needed.

In 2003 the LEO-GEO environment debris model (LEGEND) took over as the standard NASA long-term space debris evolution model. LEGEND was designed to provide an accurate representation of the historical debris population from 1957 to 2001, and 200 year predictions using Monte Carlo simulations. This was done using a mostly deterministic approach aimed at mimicking the real population, with stochastic collision and fragmentation/breakup algorithms. The code was validated by reproduction of the present epoch debris population, through propagations from 1957 using realistic launch rates and historical break-up data.

The future predictions of LEGEND are made by averaging the results of at least 100 MC runs. The numbers and distributions of debris fragments generated from break-ups are modelled using the NASA Standard Break-up Model, which was developed using data collected on actual breakups and collisions together with ground-based high velocity impact testing (Liou et al., 2004).

1.2.1.4 DAMAGE

In addition to efforts taking place within NASA, other bodies were also working on developing evolutionary tools for studying the debris population. The University of Southampton's Debris Analysis and Monitoring Architecture to the Geosynchronous Environment (DAMAGE) software is another Monte Carlo simulation based debris model.

DAMAGE has similarities to LEGEND, in that it uses a 'Cube' method for determining collision probabilities and the NASA Standard Break-up Model for modelling breakup fragments, but also differences such as the choice of propagator and atmospheric model (Lewis et al., 2011).

1.2.1.5 Other Debris Models

In addition to the two discussed above, a number of other evolutionary debris models exist. Examples include: the Modelling the Evolution of Debris in the Earth Environment (MEDEE) model of the French space agency (Dolado-Perez et al., 2012); the Long Term Utility for Collision Analysis (LUCA) of the German space agency (Radtke et al., 2017); and the Debris Environment Long Term Analysis (DELTA) model of the European Space Agency (Walker et al., 2001). Although all of these models feature slightly different implementations, the fundamental concept is the same.

1.2.2 Modelling Results

The first major result of LEGEND was a 2003 examination of the debris population under a “no new launches” scenario, i.e. all launch activity stopped in the year 2003. It was found that, over 200 years, the number of objects generated by collisions was predicted to outpace the number of objects being removed by atmospheric drag. This is caused by collision activity between 900 - 1000 km, resulting from a high spatial density of objects (Liou and Johnson, 2006, 2008).

A subsequent sensitivity study was performed using LEGEND in order to determine the effects of removing massive debris objects on the debris population growth, a strategy known as Active Debris Removal (ADR). A future launch rate, based on the 1999 - 2000 launch cycle, was assumed. It was additionally assumed that the capability to perform ADR could be implemented by 2020. The results show that, under these assumptions and on average across all MC simulation runs, an ADR strategy of between five and twenty removals per year may prevent the net growth of debris objects (Liou and Johnson, 2009).

Two further sets of LEGEND simulations studied the necessity of ADR in more detail. The simulations were run with the same launch rate but with 90 % Post Mission Disposal (PMD) success. The results of the first simulation showed that ADR may still reduce or negate the LEO debris growth rate (Liou et al., 2010). In the second, two different years of implementation was tested; 2020 and 2060. ADR strategies that were implemented on either date were shown to be effective at reducing debris growth. The simulation results indicate that moving from 2020 to 2060 may result in seven more collisions and 2000 additional debris objects. However, it was also concluded that further, more detailed, trade-off studies should be performed to identify a reasonable time-scale for ADR implementation (Liou, 2011).

The University of Southampton’s DAMAGE model was also used to examine the evolution of the debris population under a number of test cases, including the “no new launches” and “no ADR” scenario examined using LEGEND. An additional “no

collision” baseline case was also considered. Projecting forwards by 100 years, the results of the “no new launches” and “no ADR” scenario predicted an increase in the number of objects, which is in agreement with LEGEND. This was not the case for the “no collision” scenario, indicating that the increase is indeed due to collisional activity. An ADR strategy of two or five removals, selected by mass and collision probability, was tested. The results suggest that five removals per year can prevent a population growth by 2109 and that two removals would be insufficient.

DAMAGE was also used to study the effect of introducing a payload and rocket body as part of each ADR mission. In this scenario, twice as many objects were introduced into the environment than were removed through ADR. It was assumed that both new objects would comply with the IADC 25-year PMD rule and it was shown that five removals per year may still prevent the population growth. This highlighted the effectiveness of the PMD guidelines at managing population growth due to launch traffic (Lewis et al., 2009). A further study was performed in which the relationship between ADR removals and the PMD rule was examined. The results showed that decreasing the PMD rule can reduce the required number of ADR removals needed to negate population growth. It was suggested that a 5.7 year reduction in the PMD rule is on average equivalent to the removal of a single, high-risk, object (Lewis et al., 2012). Moreover, adapting the targeted number of removals using up-to-date predictions was shown to increase the probability of preventing population growth (White and Lewis, 2014a). Hence it is suggested that the best way to manage the LEO population, if the goal is to achieve a non-positive growth rate, may be a combined strategy of debris mitigation guidelines and ADR.

It is important to note that the results from these two models do not provide definitive solutions. Long-term future modelling of the debris environment relies heavily on a number of assumptions, such as future launch traffic, solar activity, PMD compliance rates and explosion rates. Although sensible estimates can be made of these values from historical data and predictions of the future directions of the global space industry, the uncertainty in their future values is high. The models can address this uncertainty by using a Monte Carlo approach. DAMAGE was used to show that varying these parameters greatly increased the spread in the MC simulations of the final debris populations. In some of the MC runs, an ADR strategy of 10 removals per year was insufficient to prevent the growth. In others there was no growth even without ADR (White and Lewis, 2014b).

Another important assumption that was used in many of these studies was that the removal of selected objects is instantaneous. This is a major misrepresentation of the technological, financial and political challenges that have yet to be overcome before the successful removal of an object from orbit. Also, this neglects the possibility that the ADR mission itself is involved in any collisions, possibly nullifying all potential benefits gained from the mission (Lidtke et al., 2017). Due to the statistical nature of

this research, the reality may be better or worse than that suggested by the average of a number of MC simulations.

Despite the uncertainties present in these results, it is generally agreed that near-Earth space is a vital resource that must be safeguarded. For this reason many organisations around the world are working to ensure that it is used in a sustainable manner. This includes having the capability to perform ADR if necessary.

1.3 Debris Management

In 1986, an Ariane-1 second stage exploded in LEO. At the time, this was the largest breakup event in the history of space use, creating 492 tracked debris objects. Subsequently, space debris was mentioned for the first time in US national space policy on January 5th, 1988, in a presidential directive by President Ronald Reagan. As a policy applying to, and binding on, all US national security and civil space sectors, it stated that “all space sectors will seek to minimise the creation of space debris” ([Executive Office of the President, 1988](#)).

By 1992, NASA was holding bilateral orbital debris coordination meetings with JAXA, ESA and the Soviet Union. These meetings were consolidated as the Inter-Agency Space Debris Coordination Committee (IADC) in 1993. The purpose of the IADC is to; “review all ongoing cooperative debris research activities between member organisations”; “identify, evaluate, and approve new opportunities for cooperation”; and “serve as the primary means for exchanging information and plans concerning orbital debris research activities”. Over the following years the major space-faring nations of the world joined the IADC. This included the British National Space Centre (BNSC), (now the United Kingdom Space Agency (UKSA)), which joined in 1996 ([NASA, 2008](#)).

1.3.1 The IADC Space Debris Mitigation Guidelines

One of the major successes of the IADC was the creation of the Space Debris Mitigation Guidelines. Based on the research and modelling studies performed by the IADC member nations, these guidelines were designed to represent best practice for management of space assets, in order to safeguard the use of the space environment for future generations. Most notable is the IADC’s “25-year rule” which states that an object should manoeuvre such that the orbital lifetime, at end of life, is reduced. It is suggested, based on studies performed by the IADC as well as other entities, that 25 years is a reasonable and appropriate lifetime limit ([IADC, 2007](#)). The IADC 25-year rule has been widely adopted across the space industry.

In 2007 the United Nations (UN) Committee on the Peaceful Uses of Outer Space (COPUOS) formally endorsed the IADC's mitigation guidelines and encouraged member states to promote adherence to the guidelines through national-level mechanisms. The seven Space Debris Mitigation Guidelines, as laid out by COPUOS, are ([United Nations Office For Outer Space Affairs, 2010](#)):

1. Limit debris released during normal operations.
2. Minimise the potential for break-ups during operational phases.
3. Limit the probability of accidental collision in orbit.
4. Avoid intentional destruction and other harmful activities.
5. Minimise the potential for post-mission break-ups resulting from stored energy.
6. Limit the long-term presence of spacecraft and launch vehicle orbital stages in the LEO region after the end of the mission.
7. Limit the long-term interference of spacecraft and launch vehicle orbital stages with the GEO region after the end of their mission.

1.3.2 Active Debris Removal

Based on the results of debris models, there is a consensus that the required research and development should be undertaken so that, unless alternative less challenging and less costly solutions can be found, an ADR strategy can be implemented.

Currently, the most developed ADR proposals require rendezvous and interface with an object, such as with a robotic arm. Rendezvous and Proximity Operations (RPO) and interfacing with an object require very precise knowledge of its orbit and attitude state. For this reason, it has never been performed on a non-cooperative target. Many potential interface techniques are being explored, including; multiple flexible grabbing arms, single robotic arms, multiple robotic arms, nets and harpoons ([Shan et al., 2016](#)). Development of ADR capability therefore relies on the simultaneous development of techniques to accurately and precisely characterise the state of the target.

Other, more uncertain, techniques for ADR are also being developed. Although these are typically at a lower level of technological readiness, and possibly more speculative in terms of effectiveness or feasibility, they can potentially provide improvements or cost savings over RPO type implementations. One example is an extremely powerful ground or space-based laser system tasked with de-orbiting small debris fragments using Laser Impulse Space Propulsion (LISP). Such a system would face political issues resulting from its possible militaristic applications, but it is suggested that it

would be theoretically capable of removing large amounts of debris over relatively short time-scales. However a number of technical challenges, such as accurate tracking of small objects, would need to be overcome before this system could be realised. Once again a key requirement for this proposal would be accurate attitude information and it is noted that a dedicated observation network would have to be set up for this purpose (Phipps et al., 1996). Another approach is a “just in time” (JIT) system tasked with targeted deflection or removal of an object once a significant risk of it being involved in a future collision has been identified. Since debris population growth has been shown to be closely tied to collision activity, preventing collisions may be just as effective as debris removal. An example would be a system, launched on a ballistic trajectory, which would externally modify the orbit of the target object such as through release of a gas to artificially increase atmospheric drag. A JIT strategy would again require a robust attitude determination infrastructure, which would be needed for accurately predicting collision events.

Numerous organisations are performing research and development in this area. One example is the Surrey Space Centre, in conjunction with a number of partners including Surrey Satellite Technology Ltd. (SSTL) and Airbus. They launched the RemoveDebris mission to the ISS in April 2018 aboard a Falcon-9 rocket as part of Crew Resupply Mission 14 (CRS-14) (Forshaw et al., 2016). In June 2018, RemoveDebris was released from the ISS into a 405 km orbit. The mission objective of RemoveDebris was to test a number of key ADR and debris mitigation technologies such as vision-based navigation, capture of objects using a net or harpoon and a drag sail. These experiments were successfully carried out throughout 2018 and 2019, however the drag sail failed to deploy. The RemoveDebris mission garnered worldwide media attention and brought increased awareness to the issue of space debris (Aglietti et al., 2020).

Research into ADR implementation has many areas of overlap with in-orbit servicing research, which is also concerned with the rendezvous and capture of a potentially non-cooperative target.

1.4 The Next Generation of Space Situational Awareness

The space object cataloguing infrastructure has changed very little since it was first introduced in the 1960s. Despite the inaccuracies and uncertainties that result from a system designed to run on the computers of the time, the cataloguing process was sufficient for keeping track of the relatively small number of objects. However with modern day computing power, storage space and data transfer rates, many of the limiting factors that influenced the design of this system are no longer limitations. Additionally, there has been a large increase in the number of space objects since the

1960s, and new processes such as conjunction event screening benefit from the highest accuracy tracking information. For these reasons there is high interest in the next generation of SSA that is no longer constrained by computational and technical limitations of the past.

There are two key ways that the present SSA system can be improved. The first is to increase data availability. At present, with the large number of objects that are currently catalogued, the SSN is required to devise strategies to prioritise sensor time. As a result, objects that are of less immediate interest, for example debris, are observed less frequently than higher interest targets such as active systems. It will always be necessary to prioritise sensor time, but increasing the volume of observations will decrease the mean age of an observation and improve the accuracy of the catalogue. While this could be achieved by building additional powerful phased array radar installations, it would be extremely costly. Therefore, new novel techniques for increasing the number of observations of space objects are desirable. This has created openings for commercialisation of space observation data, as demonstrated by ExoAnalytic Solutions - a company that sells optical GEO observations, and LEO Labs - a company selling LEO radar data and data analysis.

The second key way that the system can be improved is to incorporate new information. The SGP4 code does not take into consideration the object's shape and attitude motion. It instead uses a 'cannonball' approach, in which the object is treated as a sphere, and the parameter that encapsulates the atmospheric drag, the 'B-star' term (which incorporates the ballistic coefficient), is used as a fitting parameter. Atmospheric drag force is highly dependent on the object shape and attitude, and is the largest perturbing force in Low Earth Orbits (LEO) up to 500 - 600 km. Therefore, incorporation of object shape and attitude information would provide substantial improvements in the accuracy of predicted positions of objects in this region. As discussed, just tracking and cataloguing the positions of the large numbers objects in Earth orbit is a challenge. Hence also acquiring shape and attitude information on all of these objects is a huge task that would require major developments or expansions in hardware and software capabilities. Even so, widespread availability of attitude information has the potential to provide large improvements over present day SSA capabilities. Therefore existing techniques and technologies for determining attitude should be developed and research into novel ones should be performed.

1.4.1 Optical Light Curves

As identified in the previous section there are two key ways that the present SSA capability can be improved: 1) increase the quantity of available data, in a cost-effective manner; and 2) introduce new information that has not previously been available for wide-scale application in SSA, such as attitude.

Although they were historically expensive pieces of equipment, rapid developments in camera technology over the last 20 years has made high quality optical telescope systems much more affordable. As a result, the widespread deployment of low-cost optical telescope systems may be a cost-effective way of augmenting the observation capabilities of the SSN.

A particular advantage of optical telescopes is that when observing an object for catalogue maintenance purposes, the brightness of the object against time can be collected simultaneously. The time-varying measurement of the brightness of an object is known as a light curve. When discussing light curves of artificial space objects in low-Earth orbits, it is typically a high-cadence series of measurements, made in the visual band, with total durations on the order of minutes. Because the brightness of an object is, in part, a function of its orientation, then through analysis of the time-varying brightness signal it is possible to reveal information on the object's attitude state. Attitude determination from optical data has been considered and examined for many years. However, it is only more recently that it has become feasible to apply these techniques to a significant proportion of the space object population.

Hence, both of the key areas identified in the previous section can potentially be addressed using optical telescope systems. Their economy means that many systems can be deployed, leading to a large increase in data availability; and light curve analysis may reveal information about an object's attitude state across the space object population.

In recent years, these types of optical systems have started to be deployed. One example is the commercial company ExoAnalytic Solutions, who use a number of optical observatories around the globe to take large numbers of observations of the geostationary belt. In addition to generating tracking data, analysis of brightness variations are used for anomaly detection. Another example is the Mini-Mega TORTOTA (MMT) wide-field monitoring system. The MMT system is a fully autonomous array of 9 optical telescopes (with 1 additional redundant telescope for a total of 10), installed in pairs on 5 mounts, located in the Russian Caucasus and operated by the Kazan Federal University (Biryukov et al., 2014). In addition to its primary scientific function, the MMT system detects 200 - 500 satellite tracks nightly (Karpov et al., 2016b). The MMT system performs real-time autonomous data processing to extract tracking and brightness data from each observation (Karpov et al., 2016a). Commendably, the data collected by the MMT system is published to the publicly accessible MMT database, available at <http://mmt.favor2.info/satellites>. This database, which contains 310,429 light curves across 7828 LEO objects (as of December 2020), is used extensively throughout this thesis.

1.4.1.1 History of Light Curve Analysis

Using brightness analysis of light curves of artificial space objects, builds on a long heritage of using light curve data to determine properties of astronomical objects such as asteroids and other solar system bodies. For asteroids, optical light curves have been used to determine various properties, including rotation period (Goguen et al., 1976) and shape (Kaasalainen et al., 1992b,a). Complete light curve inversion has also been attempted on asteroids, in which the attitude and the shape are solved for simultaneously (Kaasalainen and Torppa, 2001; Kaasalainen et al., 2001; Torppa et al., 2003). In all cases, some form of a priori information was required.

1.4.1.2 Light Curve Analysis of Spacecraft

Light curves have also been used to determine various properties of spacecraft. To reduce the number of unknowns, early light curve analysis research constrained the spacecraft shape to simple geometries such as cylinders and spheres. One example is a 1963 paper by R. H. Giese, which presented a process of attitude determination for cylindrical objects, based on techniques developed by Davis, Wells, and Whipple in 1957. The result was a mathematical description for the brightness of a tumbling cylinder, resulting from a combination of specular and diffuse reflection (Giese, 1963). This research was then developed into methods of estimating the direction of the rotation axis of a cylindrical object (Giese, 1967; Williams, 1967; Bent, 1967), which were later applied to real light curve data for rocket bodies. The conclusions were that the accuracy of the imaging techniques has a strong effect on the accuracy of attitude determination. For an optical observation with a typical uncertainty of 0.05 magnitude, this technique was shown to be capable of measuring the angular velocity unit vector with an error of 1 - 10 degrees (Santoni et al., 2013). One of these methods was later adapted to make use of simpler input data: the ratio of maximum and minimum observed brightnesses (Williams, 1979a). This method was then applied to actual light curves of an Earth-orbiting cylindrical satellite. Although the solution acquired was mostly consistent, the results were limited by insufficiently frequent observations (Williams, 1979b). Simultaneously, research was being carried out to devise a method for estimating reflection characteristics of spherical satellites, (for which dependence on attitude state is much lower), such as the ratio of diffuse to specular reflection (Smith, 1967).

Another common form of spacecraft geometry are 'Box-Wings'. These consist of a cuboid body and two symmetrically placed solar panels. Research has also been performed on box-wing spacecraft where object shape was constrained to a simple box-like representation of the spacecraft geometry. Examples include observing and determining rotational periods of 20 uncontrolled GEO satellites using Fast Fourier

Transforms [Papushev et al. \(2009\)](#) and determining rotation rate and changes in rotation rate of 4 uncontrolled GEO satellites [Earl and Wade \(2015\)](#).

These techniques were generalised to any space object, in a paper by Hinks et al. in 2013. Here, the object was treated as a geometrical shape containing a number of facets. The brightness of each facet was modelled using a bidirectional reflectance distribution function (BRDF) and the total object brightness was calculated as the sum over all facets. This model was used to explore the amount of attitude information that could be acquired from a single light curve, in contrast to previous studies that correlate across multiple light curves. The results showed that the amount of attitude information available has a strong dependence on the attitude itself and the particular brightness sensitivity to attitude of that object ([Hinks et al., 2013](#)). The observability of attitude information in light curve data was also examined by Friedman et. al. in 2019, where it was shown that observation strategies can be devised to most efficiently generate datasets most suitable for performing light curve inversion to determine attitude state of an Atlas V rocket body ([Friedman et al., 2019](#)).

BRDF modelling has also been used to generate synthetic light curves. This technique was used by Wetterer and Jar in 2009 to test the performance of an unscented Kalman filter (UKF) at determining attitude from light curve data. Kalman filtering, has a long history of use to determine attitude from onboard sensor data ([Farrell, 1970](#); [Lefferts et al., 1982](#); [Crassidis and Markley, 2003](#)), but in this example it was used to determine attitude from observational data. The results showed that UKF techniques can provide estimates with errors for rotation pole direction and surface reflection properties. However, in the absence of real-world comparison data, it was concluded that further research was required to explore this technique ([Wetterer and Jah, 2009](#)).

Research was also carried out by Linares et al. in 2014 to solve for both object shape and attitude motion simultaneously. In this work, light curves were simulated for two GEO objects and a UKF was used for attitude estimation. Using a multiple-model adaptive estimation algorithm, together with a bank of shape hypotheses, the most likely geometry was determined. When the true object shape was contained within the hypotheses, this technique successfully identified the correct geometry. When the correct shape was not within the bank of hypotheses, the technique identified the best approximate geometry ([Linares et al., 2014](#)).

Another important development in attitude determination and propagation is the in-Orbit Tumbling Analysis (iOTA) tool. Funded by ESA, the iOTA model is being developed by Hyperschall Technologie Gottingen GmbH (HTG) in partnership with the Astronomical Institute of the University of Bern (AIUB). The objective is to provide a full six degrees-of-freedom simulation of an object, taking all relevant forces and torques into account, in order to predict attitude over daily and yearly timescales. The aim is that the high accuracy predictions of iOTA will synergise with ESA's ADR

initiatives (Kanzler et al., 2015). One particular capability of iOTA is generating synthetic ground measurements such as light curves. The first results of iOTA showed it to be capable of reproducing the orbit and attitude characteristics of ENVISAT, an ESA Earth observation mission that failed in 2012 (Kärräng et al., 2016).

The initial validation of iOTA is being performed by comparing the iOTA simulated attitude results to real light curves and Inverse Synthetic Aperture Radar (ISAR) observations. Five objects were examined, which included ENVISAT, for which there is also Satellite Laser Ranging (SLR) data. There was agreement between the real and synthetic measurements, particularly for SLR. Development of iOTA is ongoing.

1.5 Research Objectives

This literature review has indicated that as a result of increasing demand for observation data, and the reducing cost of optical systems, the utilisation of optical data for tracking and characterisation of space objects will increase. It was also shown that there is a strong academic precedent for using optical data to derive information on both natural and artificial space objects.

The aim of this PhD research was to assess how optical light curve data can be used to derive information on space objects, and potentially augment present day SSA/SST capabilities. Throughout the remainder of this thesis, a “light curve” is defined as a series of visual-band brightness measurements, of an object in LEO, with a time step of 1 s or less and a total duration between approximately 30 seconds to 5 minutes. Within this research aim, there was a particular focus on understanding the relationship between light curve data and spacecraft attitude state.

The selected approach was to develop a synthetic light curve generation model that can be used to improve understanding of optical light curve data. Using such a model to analyse light curve data has been done before, most notably by: Piergentili et. al. in 2017 where a ‘virtual reality model’ was examined for attitude determination of a box-wing satellite (Piergentili et al., 2017); and Wetterer and Jah in 2009 for attitude determination of a rocket body (Wetterer and Jah, 2009). In this research, a synthetic light curve model is examined in greater detail and further, more general, applications are demonstrated. The synthetic light curve model used in the thesis is different from the iOTA model in development by HTG and AIUB in that it does not perform any simulation of on-orbit torques. Instead, it is focused on calculating expected brightness signals of predefined attitude states.

A key assumption used in this work is that if the synthetic light curve model can reproduce real light curve data, then the simulated attitude state is equal to the real attitude state. This is a common technique for understanding reality - if a physical

system can be modelled perfectly then the parameters of the model will be exactly equal to the parameters in reality. However, it is impossible to create a perfect model of reality and hence it is always necessary to introduce simplifications into the model. The goal therefore is to determine what modelling simplifications can be made that do not compromise the capability of the model to determine information about the reality of the system. Considering how the simplifications made to the employed synthetic light curve model compromise its ability to model real light curve data was examined throughout this thesis. An important limitation of this technique is the uniqueness of the solution. In the cases where the synthetic light curve model was used to determine attitude parameters, all possible attitude parameters were tested to ensure that the determined solution was unique.

The research aim was broken down into a number of key research objectives. These are:

1. Develop a synthetic light curve model that can be used to replicate real optical light curve data.
2. Improve the understanding of real optical light curve data by applying the synthetic light curve model.
3. Evaluate the performance of the synthetic light curve model characteristics that are most important to attitude determination applications.
4. Evaluate how a synthetic light curve model can be used for attitude determination.
5. Demonstrate how a synthetic light curve model can exploit the relationship between attitude and brightness for purposes other than attitude determination.

Throughout this thesis a number of content chapters are presented that address various aspects of the research aim. In the final chapter, a conclusion is presented that describes how each of the research objectives were achieved.

Chapter 2

Synthetic Light Curve Model Development

2.1 Motivations

The objective of this research was to develop techniques that can be used to characterise space debris objects using optical light curves. The primary focus was on attitude determination, with particular emphasis on ADR targets such as the orbital stages of rockets. Available for this research are light curves of upper stages collected by ESA, NASA, CNSA and JAXA, provided through the IADC. Additionally, large numbers of light curves of various object classes are available through the MMT database.

The fundamental assumption used in this work was that if light curve data can be reproduced using a synthetic light curve model, then the simulated attitude state is equal to the real attitude state. In addition to attitude determination, a synthetic light curve model could be applied in other ways that further the understanding and usage of light curve data, such as anomaly detection, analysis of rotation rate dynamics or analysis of how particular object geometry and attitude state combinations manifest themselves in the shape of a light curve.

Light curve modelling can be a highly complex and computationally expensive task when all of the intricacies of object shape and reflection characteristics of the different surfaces are taken into account. Presently, there is no evidence to suggest that building a synthetic light curve model that incorporates all of this complexity is necessary to achieve the objectives of this thesis. For this reason, developing a synthetic light curve generation model should be only as complex as necessary. Additional complexity should be added as the requirement for it is identified.

The first versions of the synthetic light curve model developed in this work used a simple calculation in which the brightness of each face of a geometry model were summed, assuming each face to be a Lambertian reflector. A Lambertian reflector is an ideal matte surface that reflects light entirely diffusely. Early comparisons of this brightness calculation against real light curve data showed reasonable agreement, but the brightest reflections were missing from the synthetic light curves. This was because of the lack of inclusion the specular component of the reflection and so it was determined that this should be added to the model. The technique of Wetterer and Jah (Wetterer and Jah, 2009) was adopted to facilitate comparison and validation against existing results. In the work of Wetterer and Jah, a Bidirectional Reflectance Distribution Function (BRDF) was applied to each face of a geometry model.

2.2 Model Requirements

From the motivations described above, the synthetic light curve model was designed with two key requirements. These were:

Light Curve Replication The model should be capable of reproducing real world light curve data. This would allow comparisons to be made to real-light curves and so assess the performance of the model.

Simplicity The model should not introduce additional complexity until the need was identified. This would avoid introduction of additional unnecessary sources of uncertainty and computational burden in the modelling and parametrisation of these additional components.

To achieve the first of the model requirements, accurate positions of the Sun, the Earth, the observation site and the RSO need to be computed at the time of some real light curve. This requirement, therefore, dictates that the light curve model must be capable of ingesting the position data for these objects and synchronising the data in time and reference frame.

In conjunction with these two requirements, the aim was to make the model as modular as possible. This facilitated the implementation of improvements or changes to the code operation in the future.

2.3 Code Operation

The operation of the synthetic light curve generation model, which is coded primarily in Python, is shown in the flow charts of Figure 2.1 and Figure 2.4. The model can be

broken down into two key modules, pre-light curve generation (Figure 2.1) and the light curve generation module itself (Figure 2.4). This section provides an overview of the operation of these two modules and the flow of inputs and outputs between them.

2.3.1 Model Inputs

The model inputs are as follows. The first is a NORAD catalogue ID, which is used to identify the required object in the Space-Track catalogue. The second input is the time data, which includes a start time, a simulation time step and a total number of steps. The third input is the site data, consisting of a latitude, a longitude and an altitude of the observation location. The final two inputs, which are only required for the final light curve generation step, are a faceted geometry model and BRDF parameters, where the BRDF parameters define the reflection characteristics of each of the geometry facets.

2.3.2 Pre-Light Curve Generation

The pre-light curve generation module can be further broken down into the two modules shown in the flow chart of Figure 2.1, the initialisation module and the coordinate transformation module.

2.3.2.1 Initialisation Module

The initialisation module is responsible for parsing the input data and using this information to generate all the inputs that are required by subsequent modules. This includes conversion to the required time systems and interfacing with external databases to generate the Sun, Earth and RSO positions. The top half of Figure 2.1 shows the flow of data through this module and the three external databases that need to be accessed. The first of these is the International Earth Rotation and Reference Systems Service (IERS) database¹, which is used to acquire the Earth Orientation Parameters (EOP). The various EOPs are necessary for converting between different time systems and between celestial and terrestrial reference frames. Within the initialisation module, the initial Universal Coordinated Time (UTC) input is converted to the Barycentric Dynamical Time (TDB), Terrestrial Time (TT) and Universal Time 1 (UT1) specifications that are used throughout the model. The second database is the Space-Track Two Line Element (TLE) database². A TLE contains data that provides a description of a satellite's orbit at a particular instant in time. The final database is the

¹<https://www.iers.org/IERS/EN/DataProducts/data.html>

²<https://www.space-track.org>

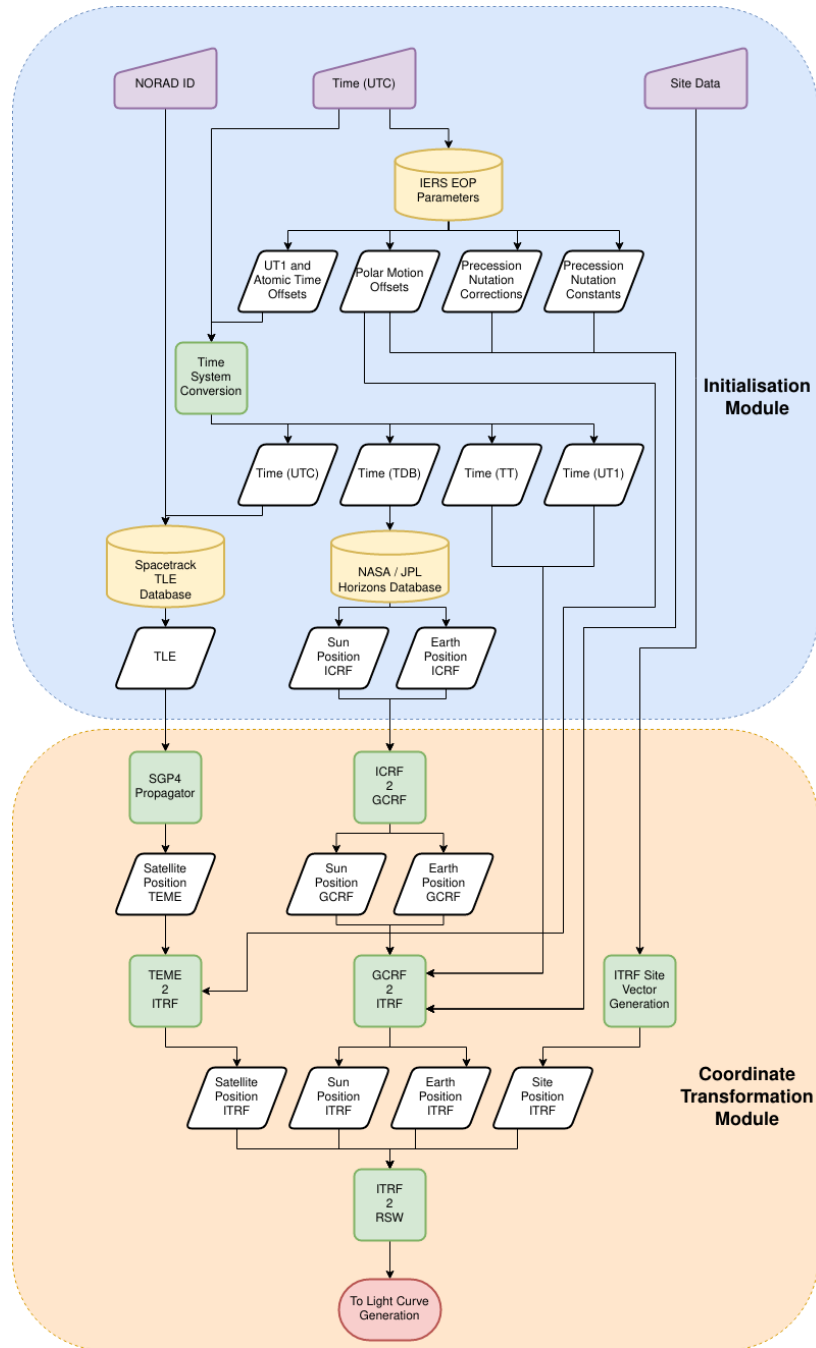


FIGURE 2.1: The operation of the first two modules of the synthetic light curve generation software.

Horizons database³, which is used to acquire ephemeris files for the Sun and the Earth. The Sun and Earth data is requested using TDB time and is received in the International Celestial Reference Frame (ICRF).

All aspects of the initialisation module are coded in Python and all data is stored as comma-separated variable files.

³<https://ssd.jpl.nasa.gov/horizons.cgi>

2.3.2.2 Coordinate Transformation Module

When the initialisation module has finished running, the coordinate transformation module is run. Within this module, a satellite ephemeris is generated by using the TLE acquired in the initialisation module with the SGP4 orbital propagator. The particular SGP4 implementation used is the one provided by Vallado as part of the Revisiting Spacetrack Report No.3. This code is available from <https://celestrak.com/publications/AIAA/2006-6753/> in a number of programming languages. The C++ implementation was used in order to maximise speed and a wrapper was used to allow it to function in a Python environment. The resulting object position data is in the True Equator Mean Equinox (TEME) reference frame.

The Sun and Earth ephemeris data acquired from Horizons is converted from the ICRF to the Geocentric Celestial Reference Frame (GCRF). This is achieved by simply translating the data so that the Earth is at the origin of the coordinate system. Then the data is transformed into the International Terrestrial Reference Frame (ITRF). This is a sequence of transformations that corrects for the complex motion of the Earth's crust with respect to the celestial sphere. The process requires the use of the EOPs provided by the IERS database. Due to the complexity of these transformations and in order to promote consistency, IERS publishes standardised conventions and provides a library of software routines, known as the Standards of Fundamental Astronomy (SOFA). In this work the SOFA routines that implement the IERS 2010 conventions were used. Because the SOFA routines are written in FORTRAN77, they were integrated into the model using a wrapper.

For the light curve generation, the data is required to be in a satellite-centred reference frame. Hence, the final step of the coordinate transformation module is to rotate into the 'RSW' frame. In the RSW frame the coordinate origin is coincident with the centre-of-mass of the satellite; the R-vector is the radial vector, which points directly away from the centre-of-mass of the Earth; the W-vector points normal to the orbital plane, which is obtained from the cross product of the radial vector with the satellite velocity vector. The S-vector is the vector that is normal to both the R and W vectors and completes the set of three orthogonal axes. For a circular orbit, the S-vector coincides with the velocity vector (Vallado et al., 2006). This reference frame is illustrated in Figure 2.2.

2.3.3 Light Curve Generation

Following completion of the Initialisation and Coordinate Transformation modules, all necessary data is available for synthetic light curve generation. This is achieved by performing a brightness calculation for each facet of the object geometry model, and




FIGURE 2.2: The RSW satellite-centred reference frame used for light curve generation.

summing each of the individual brightnesses at each time step. The mathematics of the brightness calculation is presented in the following subsection, 2.3.4.

Figure 2.3 shows an example of a faceted object geometry model that could be used to represent a simple “box-wing” spacecraft. A box-wing spacecraft is defined by a cuboid shaped body to which two solar panels are attached on opposite sides.

A flow chart showing the light curve calculation process is shown in Figure 2.4. Looping through each time-step, then through each facet, if the facet is illuminated and observable then a reflected brightness is calculated. The total reflected brightness at each time-step is then the sum of all the individual brightness calculations.




FIGURE 2.3: The simplified geometry of a box-wing satellite. The central cube represents the main spacecraft bus and thin rectangles represent solar panels, attached via booms.

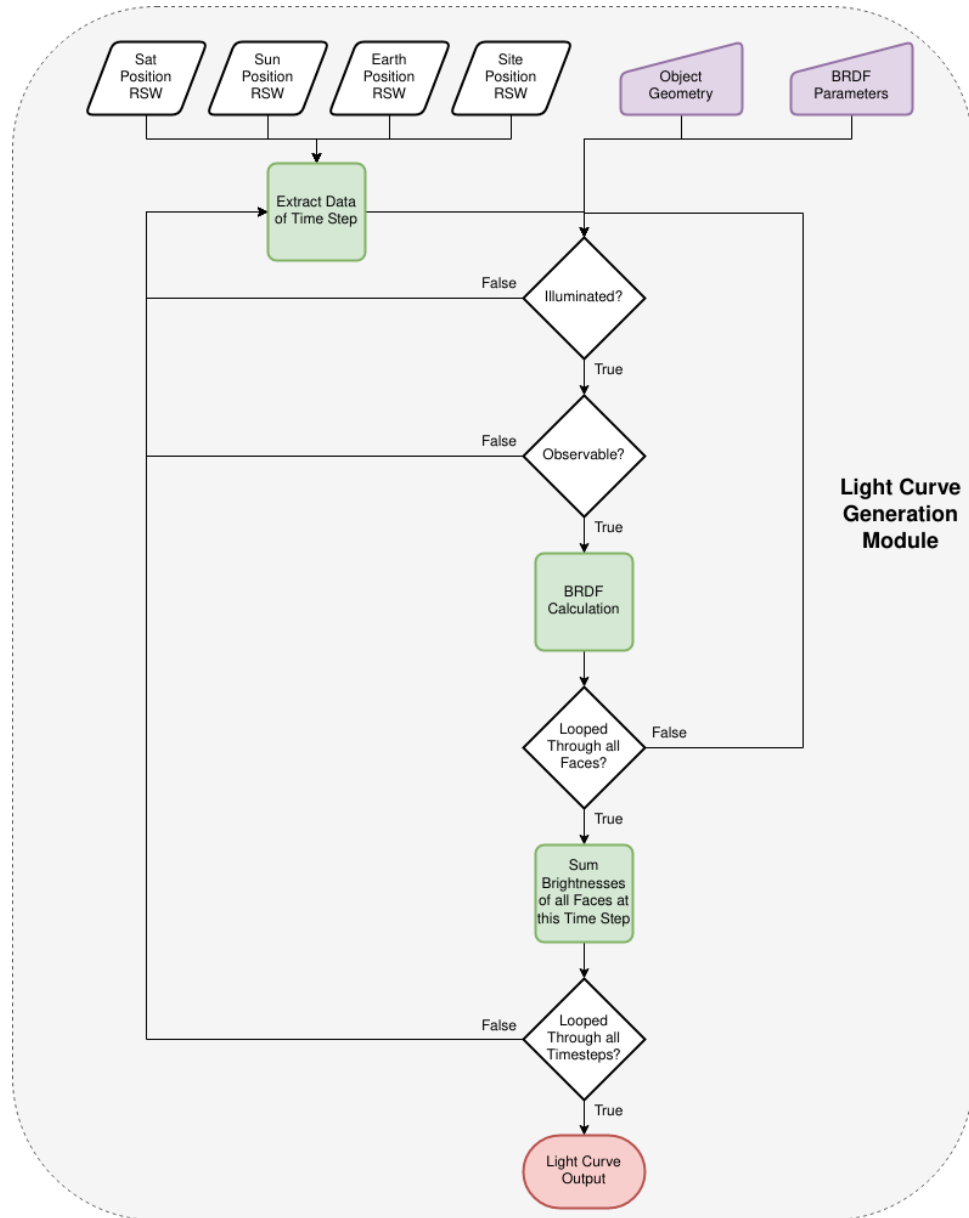


FIGURE 2.4: Flow chart showing the calculation loops in the Light Curve Generation Module.

2.3.4 Bidirectional Reflectance Distribution Function Calculation

Calculating the properties of light reflected from a surface into a particular direction, due to illumination from another direction, is the purpose of a Bidirectional Reflectance Distribution Function (BRDF). Because of their extensive use in Computer Generated Imagery (CGI), many different BRDF models of varying levels of complexity exist (Montes and Ureña, 2012).

The Cook-Torrance BRDF was selected for use in this work for three reasons. The first reason was because it was the same model used by Wetterer and Jah (Wetterer and Jah, 2009), and so provides a point of comparison for validation. Secondly, it provides

a sufficient level of realism for modelling satellite light curves that include specular reflection, because it is well-suited to modelling metallic surfaces. Thirdly, it is computationally inexpensive compared to other models with similar capabilities. To avoid unnecessary increases to the complexity of the light curve model, it was decided that a more complex model of reflection, such as ray-tracing, would only be implemented if the necessity could be clearly established (Montes and Ureña, 2012; Cook and Torrance, 1982).

2.3.4.1 Cook-Torrance BRDF

The Cook-Torrance BRDF is built upon the Torrance-Sparrow BRDF, which is a geometric optics theory-based model. The Torrance-Sparrow BRDF calculates reflected brightness as a linear combination of diffuse and specular reflection components. The diffuse component is modelled as reflection from an ideally diffuse (Lambertian) surface. The specular reflection component is dependent on the roughness of the surface, which is modelled as a collection of microfacets of equal area with random orientations. The microfacet distribution is controlled through user-defined parameters. The specular component is further modified by the Fresnel factor, which gives the fraction of reflected light from the surface. Finally, the specular component is modified by the Geometric attenuation factor, which gives the fraction of reflected light that is not occluded by other facets. The performance of the Torrance-Sparrow BRDF is validated using ray-tracing-like simulations.

The Cook-Torrance BRDF is very similar to the Torrance-Sparrow BRDF, but now only those microfacets that are aligned with the normal of a hypothetical surface that reflects light specularly from the incidence vector to the reflection vector contribute to the calculated brightness.

The Cook-Torrance BRDF is given by

$$B_i = (sR_s + dR_d)\pi A_i(\mathbf{N}_i \cdot \mathbf{L})(\mathbf{N}_i \cdot \mathbf{V}) \quad (2.1)$$

where B_i is the ratio of incident to reflected flux at the i^{th} facet, A_i is the area of the i^{th} facet, R_s and R_d are the specular and diffuse bidirectional reflectances and s and d are the specular and diffuse coefficients, with $s + d = 1$. The vectors, \mathbf{V} , \mathbf{L} and \mathbf{N} are defined in the following paragraph.

Figure 2.5 illustrates a simplified reflection geometry showing the surface area, A , surface normal unit vector, \mathbf{N} , illumination unit vector, \mathbf{L} , observation unit vector, \mathbf{V} , and the angular bisector of these two vectors, \mathbf{H} . The angle between the \mathbf{V} and \mathbf{L} unit vectors is 2θ and the angle between the \mathbf{H} and \mathbf{N} unit vectors is α . Because \mathbf{H} is the angular bisector of the illumination and observation vectors, it is therefore the surface




FIGURE 2.5: The reflection geometry. \mathbf{V} is the observation unit vector, \mathbf{L} is the illumination unit vector and \mathbf{H} is the angular bisector of these two unit vectors. \mathbf{N} is the surface normal unit vector and A is the surface area.

normal of a hypothetical surface that would reflect light from the illumination vector specularly into the direction of the observation vector.

The diffuse reflectance is given by the equation

$$R_d = \frac{\omega}{\pi}, \quad (2.2)$$

with ω being the diffuse albedo of the surface. The specular reflectance is

$$R_s = \frac{F}{\pi} \frac{DG}{(\mathbf{N} \cdot \mathbf{L})(\mathbf{N} \cdot \mathbf{V})}, \quad (2.3)$$

where F is the surface reflectance, G is the geometric attenuation factor and D is the facet slope function.

The geometric attenuation factor, G , which accounts for the masking and shadowing of facets by one another, is

$$G = \min \left\{ 1, \frac{2(\mathbf{N} \cdot \mathbf{H})(\mathbf{N} \cdot \mathbf{V})}{\mathbf{V} \cdot \mathbf{H}}, \frac{2(\mathbf{N} \cdot \mathbf{H})(\mathbf{N} \cdot \mathbf{L})}{\mathbf{V} \cdot \mathbf{H}} \right\}, \quad (2.4)$$

which returns the minimum of the three comma separated numbers. The angular bisector, \mathbf{H} , is calculated from

$$\mathbf{H} = \frac{\mathbf{V} + \mathbf{L}}{|\mathbf{V} + \mathbf{L}|}. \quad (2.5)$$

The facet slope function, D , gives the fraction of the facets that are orientated in the direction \mathbf{H} . A number of facet slope functions have been formulated. In this thesis, the Beckmann distribution function is used as it is applicable to a wide range of

surface materials and values for roughness. The Beckmann distribution function is given by

$$D = \frac{1}{m^2 \cos^4(\alpha)} \exp \left(- \left[\frac{\tan(\alpha)}{m} \right]^2 \right) \quad (2.6)$$

which, aside from the angle α , is a function of only the 'slope term', m , which parametrises the surface roughness. Higher values for m correspond to a rougher surface and hence a broader reflected brightness distribution. In this work, m will be referred to as the 'roughness parameter'.

The reflectance, F , of a perfectly smooth, mirror-like surface, with a zero extinction coefficient, can be obtained from the Fresnel Equation:

$$F = \frac{1}{2} \frac{(g - c)^2}{(g + c)^2} \left\{ 1 + \frac{[c(g + c) - 1]^2}{[c(g - c) + 1]^2} \right\} \quad (2.7)$$

where:

$$c = \cos(\theta) = \mathbf{V} \cdot \mathbf{H} \quad (2.8)$$

and:

$$g^2 = n^2 + c^2 - 1. \quad (2.9)$$

At normal incidence, ($\theta = 0$, therefore: $c = 1$ and $g = n$), this equation becomes

$$F_0 = \left\{ \frac{n - 1}{n + 1} \right\}^2 \quad (2.10)$$

and this can be rearranged to obtain the refractive index, n

$$n = \frac{1 + \sqrt{F_0}}{1 - \sqrt{F_0}}. \quad (2.11)$$

For simplicity, the reflectance at normal incidence, F_0 , is taken as being equal to the diffuse albedo, ω (Wetterer and Jah, 2009).

The reflected brightness versus time data is then converted to an apparent magnitude signal using the equation (Kutner, 2003):

$$m_{app} = -26.74 - 2.5 \log_{10} \left(\sum_{i=1}^{N_{facets}} \frac{B_i}{4\pi r^2} \right) \quad (2.12)$$

where -26.74 is the apparent magnitude of the Sun, B_i is the ratio of incident to reflected flux at the i^{th} facet, N_{facets} is the total number of reflecting facets and r is the distance from the surface to the observer.

Therefore, the reflection characteristics of each facet of the object geometry are parametrised using three numbers: the diffuse reflection coefficient, d (where the specular coefficient is: $s = 1 - d$); the diffuse albedo, ω ; and the roughness parameter, m .

2.3.5 Self-Shadowing Calculation

Depending on the object geometry, it may be necessary to calculate the shadowing of facets by one another. This is performed in the light curve generation model by first projecting the object geometry onto a two-dimensional plane perpendicular to the vector of illumination. Each face of the geometry is then stacked in this plane and assigned an “illumination fraction” based on how much of each face is obscured by another. The same process is then performed with respect to the observation vector and each face is assigned an “observation fraction”. The brightness contribution from each face is obtained by multiplying the calculated brightness by the illumination and observation fractions. This shadowing calculation method relies on the assumption that the observable fraction of a face is contained within the illuminated fraction.

2.3.6 Measuring the difference between light curves

The key assumption in this thesis is that when a simulated light curve reproduces a real world light curve, the simulated attitude state is equal to the real attitude state. Hence, a method for determining whether a simulated and real-world light curve are the same, or otherwise quantifying the differences between them, was required. Two approaches were formulated.

2.3.6.1 Root Mean Squared Error

The first approach for quantifying the difference between light curves was to measure the Root Mean Squared Error (RMSE) as per the equation

$$RMSE = \sqrt{\frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{n}}, \quad (2.13)$$

where \hat{y}_i and y_i are the i -th data points of two light curves that each contain an integer n number of data points. Using this method, a synthetic and real-world light curve, each with n time steps, are compared at each time step and the RMSE quantifies the overall degree of fit.

2.3.6.2 Hilbert distance calculation

A second, more complex approach was also used. In this approach the light curves are expressed as mathematical functions, then ‘distance’ between these functions is measured in Hilbert space, where a Hilbert space is an infinite-dimensional space in which each point is a function. The position of a function in this space can therefore be described by analogy to vector spaces, so that the coordinates of this vector represent the projection of the given function (i.e. light curve) onto a chosen basis of orthogonal functions. With this mathematical construction the distance between two functions in the space can be calculated using standard vector mathematics. The distance between two functions in Hilbert space will be zero if and only if they are the same function.

Although mathematically more complex and computationally more expensive than an RMSE calculation, the Hilbert space distance technique provides two key advantages. First, it provides additional dimensions of information when compared to the one-dimensional property of RMSE. Secondly, this technique may highlight differences between the light curves not revealed by other techniques, depending on the choice of basis functions.

Please note that the notation used in this section is isolated from what is used in other sections, despite the reuse of particular letters. Mathematically, a Hilbert space H , is defined as the set of all functions that are square integrable in a set g . Any square integrable function will be called a *point* of H , and the *origin* of H is the function that is zero almost everywhere in g .

If $\{\varphi_n\}$ is a complete sequence of orthonormal functions in a set g , and f is any square integrable function in g , then the sequence $\{f_n\}$, which is given by

$$f_n = a_1\varphi_1 + a_2\varphi_2 + \dots + a_n\varphi_n \quad (2.14)$$

where

$$a_n = \int_g \varphi_n f dt \quad (2.15)$$

will converge in the mean to a function f , i.e.

$$\lim_{n \rightarrow \infty} \int_g |f - f_n|^2 dt = 0. \quad (2.16)$$

The values generated through Eq. (2.15) are the *coefficients* of f with respect to the system φ_n and provide the coordinates of the function f in the Hilbert space H .

Hence, the position of the point f in H is given by

$$f \equiv (a_1, a_2, \dots, a_n), \quad (2.17)$$

where

$$f \sim \sum_{n=1}^{\infty} a_n \varphi_n. \quad (2.18)$$

If two functions in H share the same coefficients, then they are the same point of H .

If f_1 and f_2 are two distinct points of H , where

$$f_1 \equiv (a_1, a_2, \dots, a_n) \quad (2.19)$$

and

$$f_2 \equiv (b_1, b_2, \dots, b_n) \quad (2.20)$$

then the distance between f_1 and f_2 in H is

$$d^2 = \sum_{n=1}^{\infty} (a_n - b_n)^2. \quad (2.21)$$

The complete system of orthonormal functions, $\{\varphi_n\}$, selected for use in this work is the trigonometric system

$$\frac{1}{\sqrt{2\pi}}, \frac{\cos nx}{\sqrt{\pi}}, \frac{\sin nx}{\sqrt{\pi}} \quad n = 1, 2, \dots, \infty. \quad (2.22)$$

These were chosen due to their suitability for approximating periodic functions.

For a square integrable, L -periodic, function $f(x)$, the *Fourier coefficients* are given by the equations

$$a_0 = \frac{1}{L} \int_{-L}^L f(x) dx, \quad (2.23)$$

$$a_n = \frac{1}{L} \int_{-L}^L f(x) \cos nx dx \quad (2.24)$$

and

$$b_n = \frac{1}{L} \int_{-L}^L f(x) \sin nx dx. \quad (2.25)$$

The *Fourier series* of $f(x)$ is given by

$$f(x) \sim \frac{1}{2}a_0 + \sum_{n=1}^{\infty} (a_n \cos nx + b_n \sin nx). \quad (2.26)$$

If the function $f(x)$ satisfies the condition

$$f(x) = f(-x) \quad (2.27)$$

which is to say it is an even function, then the Fourier series of $f(x)$ reduces to only cosine terms, whereby

$$f(x) \sim \frac{1}{2}a_0 + \sum_{n=1}^{\infty} a_n \cos nx \quad (2.28)$$

with the Fourier coefficients given by

$$a_n = \frac{2}{L} \int_{-L}^L f(x) \cos nx dx. \quad (2.29)$$

Therefore the process for calculating the difference between two light curves, $f_1(x)$ and $f_2(x)$, is to first satisfy Eq. (2.27) by mirroring the light curve in the negative x domain. Subsequently the Fourier cosine coefficients of the two light curves can be determined from Eq. (2.29). Inserting the two sets of Fourier coefficients into Eq. (2.21), and setting the set g to be the interval $[-L, L]$, the distance between the two light curves in Hilbert space can be calculated. For the light curves examined throughout this thesis, the Fourier series was truncated to 1000 terms ($n = 1000$), which provided a sufficiently accurate approximation (Sansone, 1959).

2.4 Code Validation

2.4.1 Position Data Uncertainty

The sources of uncertainty in the modules that lead to the light curve generation are those in the position data acquired from Horizons, the position of the object as calculated from SGP4 and the uncertainties introduced by the transformations between reference frames.

Considering Horizons, the positional accuracy of the Sun and Earth data is determined by the underlying ‘Development Ephemeris’ theory used by NASA/JPL (National Aeronautics and Space Administration / Jet Propulsion Laboratory) to model the solar system. For the Sun and the inner planets, the positional error amounts to approximately $0.001'' - 0.002''$ or approximately 1 km at typical inner planet interplanetary distances (Standish and Williams, 2003).

The uncertainty in an object position calculated from SGP4 using TLE data is much more difficult to quantify. The first reason is that the precise version of SGP4 that was used to generate a TLE must be used to convert it back to a state vector to retain maximum accuracy. However, due to SGP4 being software used operationally by the

US military, and the numerous changes and modifications that have been made to it over the years, the version of SGP4 used to generate any particular TLE is unknown to the public. In addition, the SGP4/TLE system was designed and operationalised in the 1970s, and so was designed to run on hardware and Internet infrastructure that, by today's standards, is extremely low performance. This is why the orbital data is encoded in two lines of only 69 characters, which leads to further loss of accuracy. Therefore, there is a disparity between the state vector obtained by running SGP4 on a TLE and the true state vector. This is due to two factors: firstly, and most importantly, is the error in the initial position at TLE epoch that results from converting observation data into the TLE format; secondly, the predicted future position will have errors introduced by the deficiencies of the SGP4 propagation model. The precise positional uncertainty will depend on many factors, including properties of the underlying observation data that does not get released publicly, and will be specific to particular objects and orbits. It therefore cannot be quantified for any specific TLE. This positional uncertainty increases as the object is propagated into the future, which is a result of the relatively simple force modelling used in SGP4, together with the usual sources of orbital propagation uncertainty (such as unpredictable atmospheric density, unknown or unmodelled attitude state, etc). The rate at which this positional uncertainty increases will again be dependent upon the particular object and orbit as well as external factors such as solar activity/space weather.

A further problem with the SGP4 and TLE system is that it is only designed to propagate using natural forces and so cannot handle manoeuvring objects. This has become a particular issue with the increasing number of spacecraft that use electric propulsion, which apply low levels of thrust over extended time periods.

The final step that may introduce positional uncertainty is the transformation from the celestial frame to the terrestrial frame performed by the SOFA software libraries. The SOFA routines use a number of mathematical theories to perform transformations that correct for effects such as the slowly evolving position of Earth's rotation axis with respect to the celestial sphere and movement of Earth's crust. To achieve this, very precise measurements are made of the Earth's orientation. However, the total error in these transformations is orders of magnitude less than those that result from the TLE and SGP4 system (Petit and Luzum, 2010; Bizouard et al., 2017).

2.4.2 Position Data Validation

The position data used within the light curve model was validated by comparing against the data in the MMT database. Every light curve published by the MMT contains five columns of information: time, standard magnitude, apparent magnitude, range and phase. The accuracy of the position data was validated by comparing the range and phase outputs of the model to those in the MMT database. Figure 2.6 shows

an example plot that compares range and phase extracted from the light curve generation model to the range and phase data provided by the MMT database, for a light curve of a CZ-2B rocket body. One possible source of disagreement between the two datasets is the actual accuracy of the MMT measurements. Although there will always be some difference between a measured object position and the truth data, which result from the accuracy limits of observation hardware and processing pipeline, it is reasonable to assume that the observed positions from the MMT database provide higher accuracy than those obtained from propagating TLE data using SGP4. Therefore, for the purposes of this validation, it was assumed that MMT data provides the true object positions. Hence, there are two reasons why the range and phase data would disagree between the MMT data and the light curve model outputs. Firstly, the TLE from which the model generates the position data is old or of poor quality. Secondly, the object has performed a manoeuvre since the TLE was generated. Therefore, in all light curves generated in this work where the objective was to reproduce real data, if the MMT and TLE data disagreed by more than 1%, then the simulated light curve was flagged as erroneous and was discarded.

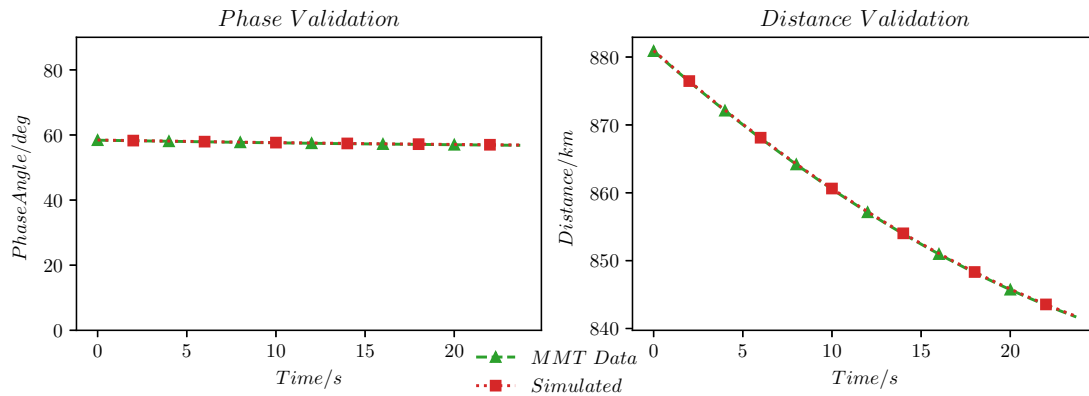


FIGURE 2.6: The difference in phase angle (left) and distance (right) between the MMT data and the light curve simulation model, for a light curve of a CZ-2D rocket body (Norad ID: 41858) collected at 03:24:05 UTC on 07/01/2017. The percentage differences are 0.14 % (left) and 0.01 % (right).

2.4.3 Coordinate Transformation Validation

The rotational motion is defined as a rotation rate about a defined rotation axis. The application of this rotational motion was validated by rotating a geometry model about a range of rotation vector directions and visually inspecting the geometry to ensure it rotated about the defined axis.

The light curve generation is performed in a satellite-centred body-fixed reference frame, which requires transformation into this frame from a satellite-centred orbital reference frame. To validate this step, it is ensured that the angles between various vectors remain the same in both reference frames. In Figure 2.7, the angles between

the observation vector, illumination vector and rotation vector, and each of the Earth, velocity and orbital plane normal vectors, are plotted for both the orbital and body-fixed reference frames. This plot shows that the angles between each of the vectors are the same in both reference frames. The next plot in Figure 2.7, shows the angular difference between the Z body-fixed axis and the Earth, velocity and orbital plane normal vectors. This plot validates that the attitude motion is being correctly applied in the orbital and body-fixed frames. Similar to the position validation, this validation was performed on every light curve that was generated. If ever the angles between the two frames were different by more than 1%, then the simulated light curve was flagged as erroneous. Because this validation does not rely on external data, all synthetic light curves passed this validation step.




FIGURE 2.7: The difference in position of the Observation vector, Illumination vector and Rotation vector, from the Earth, Velocity and Orbital Plane Normal vectors, between the orbital and body reference frames.

2.4.4 Brightness Calculation Validation

The final validation step was applied to the brightness calculation. This was done by comparing a synthetic light curve to the synthetic light curve example provided in Wetterer and Jah (Wetterer and Jah, 2009). Figure 2.9 shows two light curves. On the left is a real light curve of an Atlas-Centaur II rocket body (NORAD 12069) collected by the Advanced Electro-Optical System (AEOS) telescope on Maui Hawaii (Latitude 20.708°N, Longitude 156.257°W, Altitude 3075 m) at 14:04 on 24/03/2003. On the right is a synthetic light curve generated by Wetterer and Jah to replicate this real light curve. To generate this synthetic light curve, the following properties were required as inputs: a faceted geometry model; BRDF parameters for each facet; and attitude state. All other parameters, such as illumination and observation vectors, are defined by the position of the rocket body with respect to the observation site and the Sun at the time the light curve was collected.

The geometry model used in the generation of this synthetic light curve was described as: “a cylinder (height = 9 m, radius = 1.5 m) with 100 facets around its diameter and

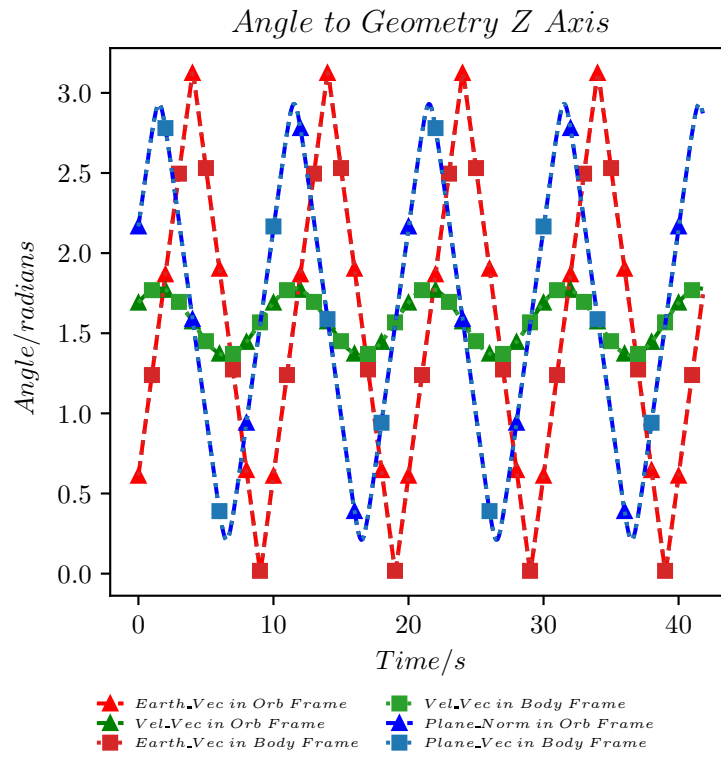


FIGURE 2.8: The difference in position of the Z body-fixed axis from the Earth, Velocity and Orbital Plane Normal vectors, between the orbital and body reference frames.

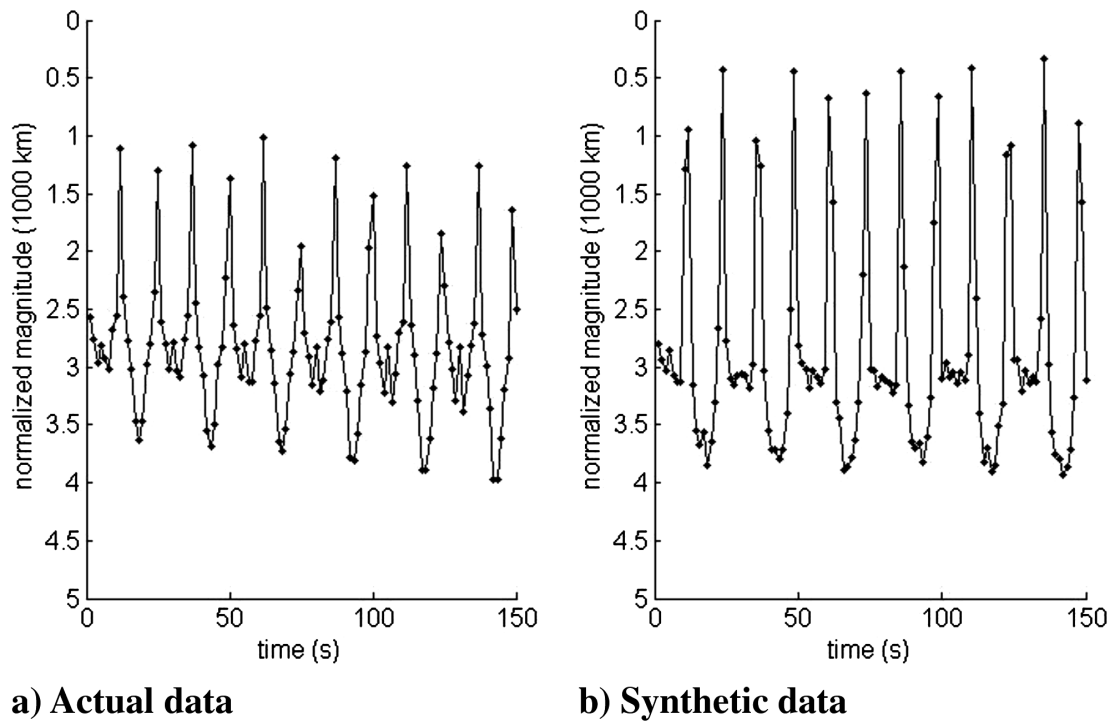


FIGURE 2.9: a) Actual light curve of an Atlas-Centaur II rocket body collected by the Advanced Electro-Optical System (AEOS) telescope on Maui Hawaii. b) Synthetic light curve generated by Watterer and Jah (Wetterer and Jah, 2009).

two 100-facet hemispherical end caps flattened to extend only 1 m beyond the end of the cylinder. The object's y axis is along the axis of the cylinder and the z axis is along the radius. The object is modelled as spinning solely about its z axis". Using this description, the geometry model shown in Figure 2.10 was generated.




FIGURE 2.10: The geometry model used to generate the validation example light curve.

In the example light curve generated by Wetterer and Jah, the cylindrical sides of the rocket body and the two end caps were given individual BRDF parameters for the diffuse albedo, ω , and ratio of diffuse to specular reflection, d . These three geometry zones are coloured blue, red and green in Figure 2.10. The roughness parameter, m , was constant across the geometry model.

It is not clearly stated in the paper what the values for the BRDF parameters were and what attitude state was used to generate the synthetic light curve shown in Figure 2.9. Therefore, these values were determined manually to generate a best-fitting light curve using the synthetic light curve model described in this chapter. Figure 2.11 shows this synthetic light curve and Table 2.1 shows the input parameters used in its generation.




FIGURE 2.11: A comparison synthetic light curve generated with the synthetic light curve model described in this chapter.

Comparing the synthetic light curve generated by the model described in this chapter (Figure 2.11) to the synthetic light curve generated by Wetterer and Jah (Figure 2.9 b), there is good agreement between the two light curves. Both exhibit a peak brightness between magnitude 0.5 - 1.0, and both exhibit two brightness minima, one at approximately magnitude 3.0 and a second at approximately magnitude 4.0. It should be noted that the synthetic light curve generated by Wetterer and Jah contains synthetic noise, whereas the light curve of Figure 2.11 does not. This is the reason why the light curve of Figure 2.11 is more uniform. Additionally, the objective of the fitting process was to show that this model can produce a light curve with a similar shape and brightness as the one shown by Wetterer and Jah. As a result, a much closer match may be possible by further tweaking the input parameters. However, this was deemed to be unnecessary.

Property	Value Used	Units
Attitude x -Euler Angle	0.691	Radians
Attitude y -Euler Angle	0.440	Radians
Attitude z -Euler Angle	3.204	Radians
Rotation Rate	0.252	Radians per Second
m	0.27	-
ω_{side}	0.25	-
d_{side}	0.3	-
ω_{top}	0.6	-
d_{top}	0.8	-
ω_{bottom}	0.47	-
d_{bottom}	0.8	-

TABLE 2.1: The input parameters used to generate the synthetic light curve of Figure 2.11.

Chapter 3

Investigating Simple Light Curve Properties

3.1 Introduction

The content of this chapter was motivated by a request from the Astrodynamics Community of Interest (ACI) to develop techniques to characterise the operational state of spacecraft from observation data. The ACI is an ongoing series of workshops organised by the Defence Science and Technology Laboratory (DSTL), a research and development branch of the United Kingdom Ministry of Defence. The purpose of the ACI meetings is to promote cooperation between the military, commercial and academic sectors of the UK space industry.

On 14-15th of March 2018, DSTL held a joint meeting with the US Air Force Office of Scientific Research at the European Office of Aerospace Research and Development in London. The focus of this meeting was on characterising space objects, and attendees were invited to present novel approaches to characterising objects as ‘active’ or ‘inactive’, with a high degree of confidence. The work in this chapter was derived from a proof-of-concept pilot study that was presented at this meeting.

The hypothesis of this pilot work was that an object with a controlled attitude state will have a lower standard deviation in light curve brightness than an object with an uncontrolled attitude state. The assumption was that an active satellite will maintain a particular attitude state with respect to a ground-based observer and so will always be observed in the same orientation. As a result, light curves of such an object will exhibit low standard deviations around a mean brightness value. An inactive satellite on the other hand will not maintain a particular attitude state and will begin to tumble. As a result, an inactive satellite will be observed in a range of different orientations, which will introduce larger standard deviations.

To test this hypothesis, the light curves of a number of active and inactive satellites were examined. The hypothesis was tested by comparing the real light curves to synthetic light curves generated with either controlled or uncontrolled attitude states.

3.2 Methodology

The method used in this chapter was to apply thresholds to light curve standard deviation data to differentiate active satellites from inactive satellites. Hence, two sets of data were required: 1) a large number of light curves of a range of active and inactive satellites; 2) the operational status of the same range of satellites. The objects selected for analysis were Globalstar satellites. These were selected for two reasons. Firstly, a large number of light curves are available in the MMT database. Secondly, the satellite shape, a box-like bus with solar panels attached to opposite sides (commonly referred to as a 'box-wing'), is a common shape amongst satellites. Although other types of satellites were not considered in this pilot work, it was hoped that by selecting satellites with this shape, the probability that the techniques can be applied more generally to other box-wing spacecraft would be improved.

The operational status of the considered satellites was acquired from two sources. The first was the MMT database, which assigns an 'ACTIVE', 'INACTIVE' or 'UNCERTAIN' flag to each object. The second source was Celestrak (<https://celestrak.com>), which assigns each object in their database a status of 'Operational' or 'Nonoperational'.

For each real light curve considered in this work, a synthetic light curve was also generated. This was to confirm that, in synthetic data, uncontrolled and controlled attitude states results in the hypothesised differences in light curve standard deviation. All inputs used to generate these synthetic light curves were defined by the properties of the real light curves, with the exception of spacecraft geometry, attitude state and BRDF parameters. Figure 3.1 shows the geometry model that was used to approximate a Globalstar satellite. This geometry was generated using measurements available online for the Globalstar platform. In Figure 3.1, the faces that are coloured red are those that contain the antennas, which are required to be Earth-facing. Therefore, an active attitude state was simulated by constructing an attitude profile whereby these faces maintained an Earth-pointing condition. For an inactive attitude state, the geometry was given a random initial orientation at the start of each light curve and a rotation about the z-body-fixed axis with a period of 10s.

The last remaining inputs required for synthetic light curve generation were the BRDF parameters. The purpose of the synthetic light curves was to show the relative change in standard deviation that results from a change from controlled to uncontrolled attitude state. Therefore, it was not necessary fit the synthetic data to the real data.

Additionally, this work was not concerned with the effect that the reflection characteristics may have on the synthetic light curve standard deviations. For these reasons, a selection of realistic BRDF values were used. These were: a diffuse albedo of 0.4; a roughness parameter of 0.2; a ratio of diffuse-to-specular reflection of 0.8.

3.3 Results




FIGURE 3.1: Simple geometry model used to represent a Globalstar satellite.

Three groups of Globalstar satellites were examined in this pilot study. The first was a group of 4 satellites, launched on the 8th of February 2000 aboard the same Delta-7420-10C rocket, which, when this work was carried out, contained a mix of active and inactive satellites. The second was a group of 4 satellites launched on 29th May 2007 aboard a Soyuz-FG Fregat rocket. All four of these satellites were active at the time of the analysis. The final group was 4 satellites launched on 20th October 2007 also aboard a Soyuz-FG Fregat rocket. This was the final batch of the first generation of Globalstar satellites and, at the time of this analysis, contained both active and inactive satellites. This information is summarised in Table 3.1, which shows the satellites that belong to each group, and Table 3.2 which shows their operational statuses.

Group	Intl. Desig.	Spacecraft ID	NORAD ID	Launch Date
1	2000-008A	M063	26081	2000/02/08
1	2000-008B	M062	26082	2000/02/08
1	2000-008C	M060	26083	2000/02/08
1	2000-008D	M064	26084	2000/02/08
2	2007-020A	M065	31571	2007/05/29
2	2007-020C	M069	31573	2007/05/29
2	2007-020D	M072	31574	2007/05/29
2	2007-020F	M071	31576	2007/05/29
3	2007-048A	M067	32263	2007/10/20
3	2007-048B	M070	32264	2007/10/20
3	2007-048C	M066	32265	2007/10/20
3	2007-048D	M068	32266	2007/10/20

TABLE 3.1: Table of data on three groups of Globalstar spacecraft

Group	Spacecraft ID	MMT Status	Celestrak Status
1	M063	UNCERTAIN	ACTIVE
1	M062	INACTIVE	INACTIVE
1	M060	INACTIVE	INACTIVE
1	M064	INACTIVE	INACTIVE
2	M065	ACTIVE	ACTIVE
2	M069	ACTIVE	ACTIVE
2	M072	ACTIVE	ACTIVE
2	M071	ACTIVE	ACTIVE
3	M067	INACTIVE	ACTIVE
3	M070	ACTIVE	ACTIVE
3	M066	ACTIVE	ACTIVE
3	M068	UNCERTAIN	INACTIVE

TABLE 3.2: Globalstar spacecraft state according to either the MMT database or Celestrak catalogue.

3.3.1 Real Light Curve Data

Figures 3.2 to 3.5 show the real light curve data for the Globalstar satellites. When both the MMT database and Celestrak catalogue list an object as inactive, the data is coloured red. When the MMT database and Celestrak catalogue list the object as active, then the data is coloured green. If the two sources state different statuses, then the data is coloured using a unique colour so that this object may be easily identified.

In Figure 3.2, the top plot shows the average brightness of all the light curves available in the MMT light curve database for group 1. The bottom plot in Figure 3.2 shows the standard deviations of these light curves. With each data set, a 40-point moving average line is also plotted. The focus of this analysis is the averaged values and not the point-to-point variations. For this reason, the data points are transparent so that the moving average lines can be more easily seen. This figure contains three satellites

that both sources predict to be inactive (red), and a single satellite (M063 - 26081) that Celestrak flags as active but the MMT database flags as uncertain. This object is coloured blue.

Considering firstly the moving average light curve brightness lines at the top of Figure 3.2, there is no average magnitude threshold value that can separate the object with uncertain operational status (blue), from the data for the inactive objects (red). The moving averages of all four objects exhibit magnitudes of 6.0 ± 0.75 .

Now considering the standard deviations data at the bottom of the figure, a black line has been drawn that shows a threshold between the inactive and uncertain status satellites. This line is drawn at a standard deviation of 0.27. With the exception of the moving average line of M064, which has a small number of points below the threshold (approximately 10% of the total), all the moving average lines of inactive spacecraft are above the threshold. For the satellite with uncertain status (M063) all points on the moving average line are below the threshold until early August 2018, after which all points on the line are above the threshold.

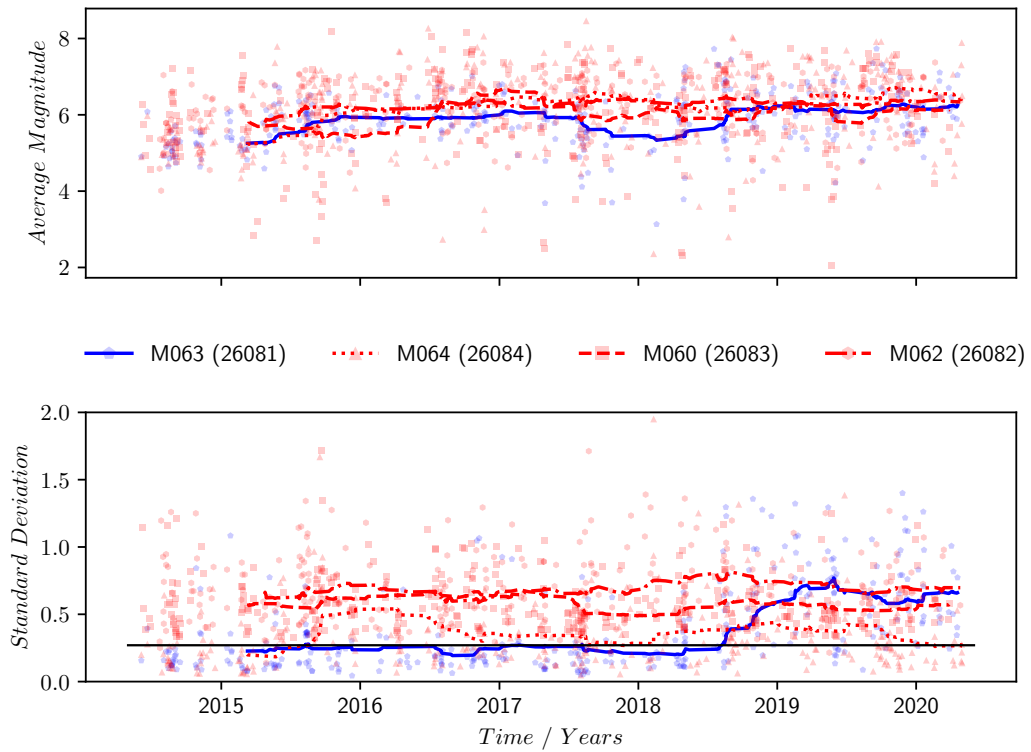


FIGURE 3.2: Average magnitudes and standard deviations of all light curves in the MMT database for the first group of 4 Globalstar satellites launched on the 8th of February 2000.

Figure 3.3 shows the data for group 2. Both sources predict that the satellites in this group are active. In the average brightness plot at the top of the figure, the moving averages of all satellites exhibit an average magnitude of 5.75 ± 0.5 . In the standard deviations plot at the bottom of the figure, the moving averages of all satellites have

standard deviations of 0.2 ± 0.06 . All points on the lines are below the 0.27 threshold value that was used in the data for group 1.




FIGURE 3.3: Average magnitudes and standard deviations of all light curves in the MMT database for the second group of 4 Globalstar satellites launched on the 29th May 2007.

Figure 3.4 shows the data for the final group of Globalstar satellites. This group contains two satellites that both sources classify as active (M070 - 32264 and M066 - 32265) and two satellites for which the two sources contradict on their status (M068 - 32266 and M067 - 32263). These two objects are coloured cyan and magenta. The moving average lines of the light curve magnitudes for this group show a larger distribution than the previous two groups. The moving average magnitude values are 6.0 ± 1.0 . Additionally, it is possible to set a threshold value that separates sections of the moving average lines for the two satellites with uncertain status from the two active satellites. However, a single threshold value cannot be used to separate the entirety of the data from the two types of satellite and, in general, these objects exhibit similar average brightnesses.

In the standard deviation plot at the bottom of Figure 3.4, the threshold value of 0.27 has once again been draw. The overwhelming majority of the moving average standard deviation data for the two active satellites appears below the threshold value, with only a small number of exceptions (20% of data points for M070 and 8% of data points for M066). For object M068 (cyan), the moving averages of standard deviation are entirely above the threshold. For object M067 (magenta) the moving

averages of standard deviation are initially below the threshold, then increase above the threshold in approximately mid-2018, before returning below in late-2020.

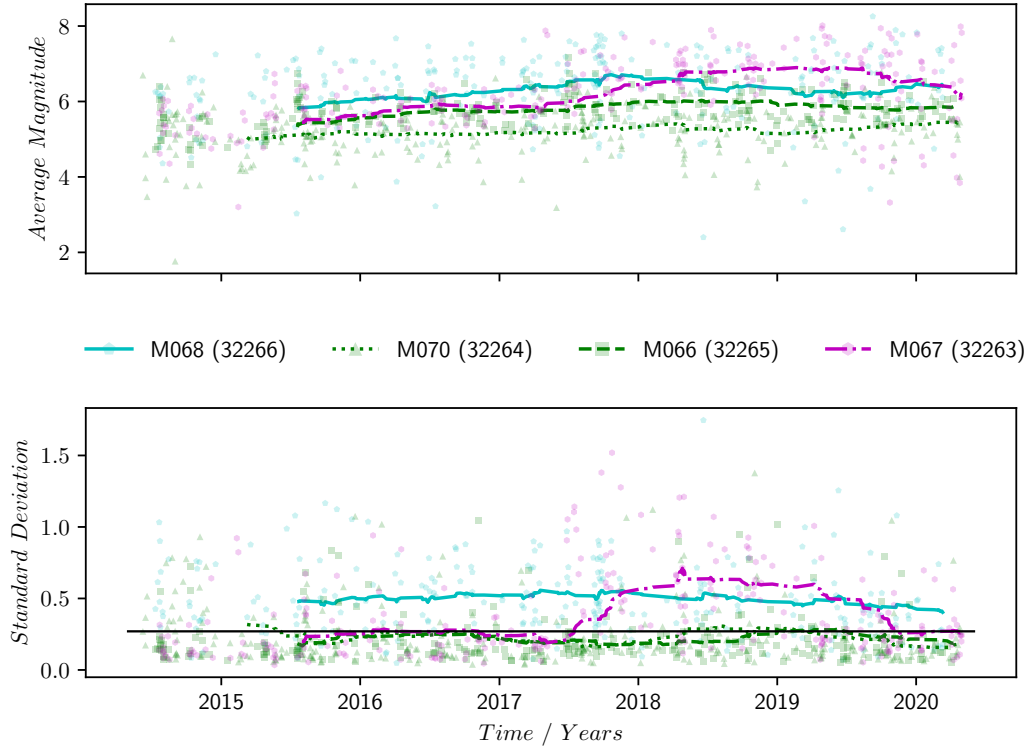


FIGURE 3.4: Average magnitudes and standard deviations of all light curves in the MMT database for the third group of 4 Globalstar satellites launched on the 20th October 2007.

The final figure of this section is presented as Figure 3.5. In this figure, all real light curve data for the three groups of Globalstar satellites has been plotted. Considering firstly the average brightness data at the top of the figure, there is no single threshold value that can separate the three types of object - active, inactive and uncertain operational status. All objects exhibit moving average magnitudes of approximately 5.75 ± 0.5 .

In the standard deviation data for the active spacecraft, the overwhelming majority of the moving average lines are below the threshold value of 0.27, with only a small number of exceptions. The inverse is the case for the three inactive satellites - the overwhelming majority of data points on the moving average lines are above the threshold. The only exception in this case is a very small number of moving average data points for object M064 (red dotted line).

It is important to note that the position of the threshold line is dictated entirely by the data of M064, which has the lowest moving average of standard deviation of all of the inactive satellites. This object is plotted as a red dotted line in the figure. If this object is excluded, then the difference in moving average standard deviation between the lowest value for an inactive satellite, and the maximum value of an active satellite, is

0.1705. As a result, the threshold could be placed anywhere in this region and 100% of the active standard deviation moving averages would be below this value and 100% of the inactive standard deviation moving averages would be above.

In the moving averages of standard deviation for the three objects with uncertain statuses, M068 (cyan) is exclusively above the threshold value, whereas M063 (blue) and M067 (magenta) are initially below the threshold, before then increasing above it. In the case of M067 (magenta), the standard deviation moving average values decrease back below the threshold at the end of the available data.

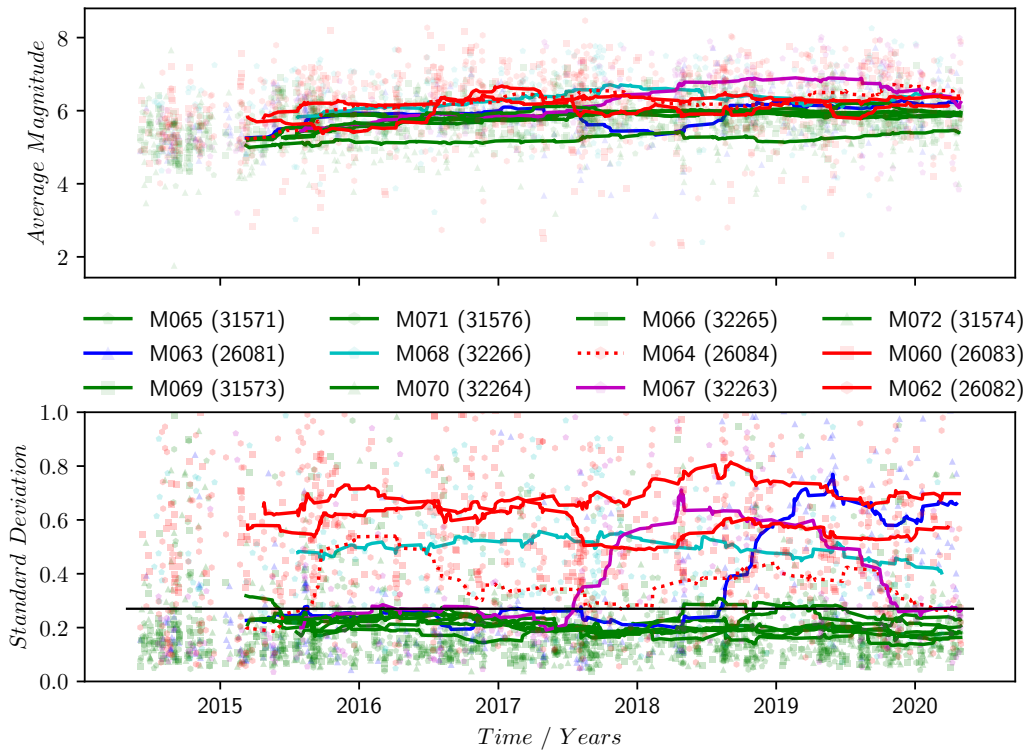


FIGURE 3.5: Average magnitudes and standard deviations of all light curves in the MMT database for all three groups of Globalstar satellites.

3.3.2 Synthetic Light Curve Data

The real light curve data presented throughout the previous section has shown that a threshold can be applied to the standard deviation moving averages of Globalstar satellites to differentiate active satellites from inactive satellites.

To test the hypothesis that the differences in standard deviations between active and inactive satellites are a result of controlled and uncontrolled attitude states, each light curve was replicated using the synthetic light curve model. For objects listed as active, the geometry maintained an attitude state where the red-coloured faces of Figure 3.1 were Earth-pointing. For the objects listed as inactive, a random orientation was

generated for each light curve, and a rotation was applied about the z-body-fixed axis with a period of 10 s.



FIGURE 3.6: Average magnitudes and standard deviations of the real and synthetic light curves of the first group of Globalstar satellites. The inactive attitude state was simulated using a random orientation and a rotation period of 10 s.

Figure 3.6 shows the real and simulated light curves for the first group of objects. All of these light curves were simulated with an inactive attitude state. The synthetic data is coloured in a darker red and blue than the real data. In this and all subsequent plots, the individual data points have been excluded so that the moving average lines are more visible.

Considering the moving average brightnesses at the top of the figure, the synthetic data is slightly brighter than the real data. The moving averages of the real data have magnitudes of 6.0 ± 0.75 , whereas the moving averages of the synthetic data have magnitudes of 7.0 ± 1.0 . This is not unexpected, as no attempt was made to find best-fitting BRDF parameters. This fit could be improved through tuning these parameters, but because the focus of this work is on the relative change in the synthetic light curve data that results from attitude state, this was not required.

The darker lines at the bottom of Figure 3.6 show the standard deviation moving averages for the synthetic data. These standard deviations are higher than the real data, with values of 1.1 ± 0.3 . All data points in the synthetic data are well above the threshold value.

Again, the agreement between the real and synthetic data could be improved through modifying the simulation inputs. This is illustrated by Figure 3.7, which shows the same results as Figure 3.6, but the rotation has been removed, i.e. each light curve was generated with a random but stationary orientation. In this figure, the standard deviations of the synthetic light curves have decreased dramatically such that they are now mostly below the threshold value. This shows that increasing the rotation period from 10 s may decrease the standard deviations and improve the fit to the real data. However, as has been mentioned, this was not required due to the focus of this work being the relative change to standard deviations resulting from active and inactive attitude states.

Figure 3.7 also provides an additional result. This plot shows that in the synthetic data, it is the rotation and not the random initial orientation that is the primary cause of the increase in standard deviation.




FIGURE 3.7: Average magnitudes and standard deviations of the real and synthetic light curves of the first group of Globalstar satellites. The inactive attitude state was simulated using only a random orientation without rotation.

Figure 3.8 shows the real and simulated data for the second group of objects. The most important property of this plot is that the moving average standard deviations are entirely below the threshold value with the exception of a small number of data points in the synthetic data of one object.

Figure 3.9 shows the real and simulated data for the third and final group of objects. Object M068 was simulated using an inactive attitude state and all other objects were

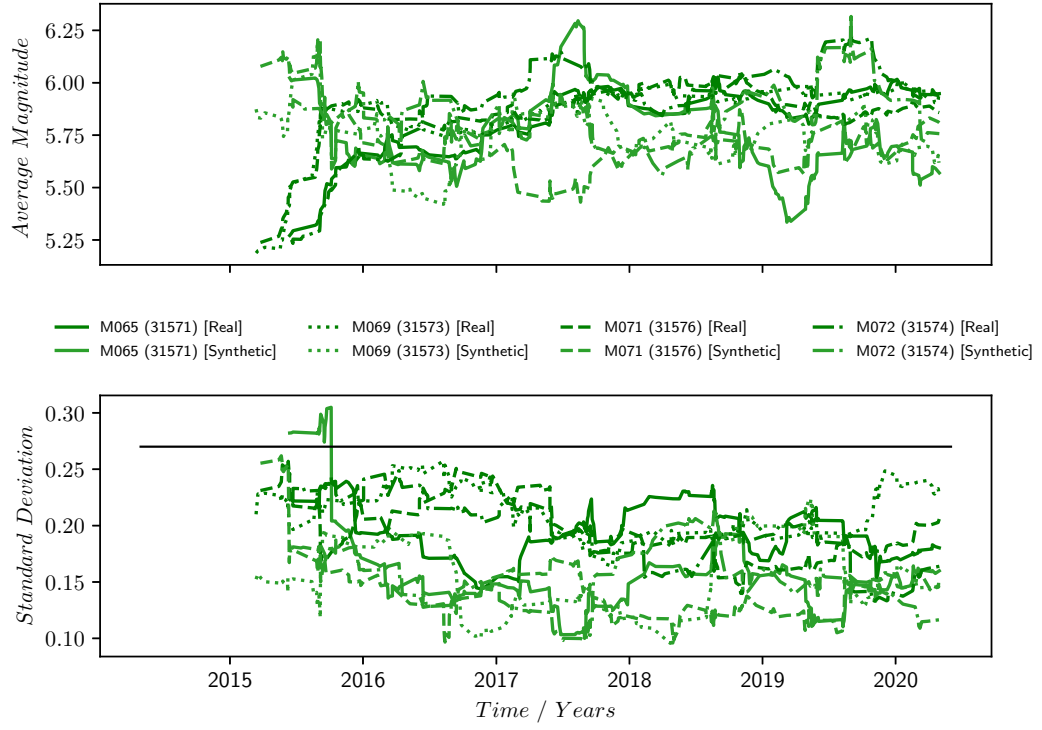


FIGURE 3.8: Average magnitudes and standard deviations of the real and synthetic light curves of the second group of Globalstar satellites.




FIGURE 3.9: Average magnitudes and standard deviations of the real and synthetic light curves of the third group of Globalstar satellites.

simulated as active. All of the points on the standard deviation moving average lines for the active objects are close to or below the threshold value. The standard deviation moving average line for M068 however, is entirely above the threshold. In the real data of this object, the standard deviation moving averages are approximately 0.5, whereas in the synthetic data they are approximately 1.0. The agreement between these two standard deviation moving average lines may be improved through increasing the rotation period from 10 s.

The final plot, Figure 3.10, shows all of the real and synthetic data for the three groups of Globalstar satellites within a single figure. There is a lot of information presented in this figure, but the two most important properties are as follows. Firstly, the standard deviation moving averages of all real and synthetic data for active satellites are at or below the threshold value of 0.27. Secondly, the standard deviation moving averages of all real and synthetic data for inactive satellites are at or above this threshold value. Additionally, if object M064 is excluded (red dotted line), there is a minimum standard deviation separation of 0.17 between all active and inactive data, resulting in a clear divide between the standard deviation moving averages of active and inactive satellites.




FIGURE 3.10: Average magnitudes and standard deviations of the real and synthetic light curves of all of the considered Globalstar satellites.

3.4 Discussion

There were three objects considered in this paper for which the two sources of operational status information were in not in agreement. These were: M063 (coloured blue), which Celestrak categorises as active, but the MMT database flags as ‘uncertain’; M068 (cyan), which is categorised as inactive by Celestrak but ‘uncertain’ by the MMT database; and M067 (magenta), which Celestrak categorises as active, but the MMT flags as inactive.

Using the results of this pilot work, the operational statuses of these objects can now be derived. Considering the data for M063 (blue), the standard deviation moving averages of this object were typical of a satellite with a controlled attitude state up until approximately mid-2018. Subsequently, the standard deviation moving averages increased to values typical of a satellite with an uncontrolled attitude state. This indicates that this satellite may have experienced a failure in mid-2018 and that the Celestrak database is yet to change its operational status from active to inactive. This may be explained by the fact that commercial satellite operators commonly do not publicly announce the failure of their spacecraft.

The standard deviation moving average data for M068 (cyan) consistently exhibited values that are typical of an inactive satellite. This suggests that Celestrak is correct in their inactive categorisation and that the MMT flag of ‘uncertain’ can be finalised as inactive.

The data for M067 (magenta) is more unusual. Similar to M063, the standard deviation moving averages of this object transitioned from values typical of an active system, to values typical of an inactive system in approximately mid-2017. This may indicate a possible failure date. However, the standard deviation moving averages then decreased in late-2019 and are now again exhibiting values that are typical of an active system. This may indicate that the satellite has since recovered from this failure. However, as was shown and discussed in Figure 3.7, the decrease in standard deviation may be due to an increase in rotation period rather than a return to a controlled operational attitude state.

An important aspect of this analysis is that it is performed by examining average trends in a large number of light curves. The point-to-point variability in the standard deviation data is very high. Therefore this technique is limited in application to large data-sets from which long-term trends can be extracted. Another important aspect of this analysis is that it relies heavily on the assumption that the higher standard deviations were caused by uncontrolled attitude motion, which is caused by the failure of the spacecraft.

3.5 Conclusions

It has been shown in this work that a threshold can be applied to the moving averages of standard deviation for the considered Globalstar satellites that differentiates the light curves of active and inactive spacecraft. If Globalstar M064 is included, this threshold value is 0.27, whereas if M064 is excluded, this value can be anywhere in the region of 0.27 to 0.44.

It was hypothesised that this difference in standard deviation between active and inactive satellites was caused by the differences in their attitude states - an active satellite maintains a particular attitude state and so is observed in the same orientation for each light curve, whereas an inactive satellite may be seen in a range of different orientation for each light curve as a result of tumbling motion.

This hypothesis was tested by generating synthetic light curves using either a constant attitude state, or a tumbling attitude state. In the synthetic data, a constant attitude state generated low standard deviations that correlated well with the low standard deviations observed in the real light curves of active systems. Switching to tumbling attitude motion resulted in dramatic increases in the standard deviations. However, when using only a random orientation and no tumbling motion, this increase was removed. This shows that it is the tumbling motion and not the random orientation that is the cause of the standard deviation increase in the synthetic light curves. These results support the hypothesis that the increase in standard deviation observed in the real data of inactive satellites may be caused by uncontrolled tumbling attitude states.

By making use of the results in this work, the operational status of three Globalstar satellites, for which there was disagreement in their status between the MMT and Celestrak databases, was refined.

Although the results of this proof-of-concept study are very promising, there is further work that should be undertaken to test the robustness of this technique. Only Globalstar satellites were examined, which were selected due to their 'box-wing' geometries being a common type of spacecraft design. However, additional spacecraft should be considered to confirm that the results of this chapter are not specific to Globalstar spacecraft. Furthermore, expanding the analysis to consider a larger number of satellites will increase the confidence in these results.

Chapter 4

Investigation of Variable Rocket Body Light Curve Periods

4.1 Introduction

In 2014, the IADC conducted an optical observation campaign on defunct rocket bodies on-orbit. The objective was to investigate what could be characterised about their attitude state, using optical data. This work could potentially be used to support ADR missions. The first step in target selection for an ADR mission is to identify high risk objects, such as those with high mass and high collision probability, whose removal would provide the greatest benefit. Many rocket bodies are massive objects in congested orbits and so are good ADR targets. Subsequently, for an ADR mission that uses a mechanical grabber, such as a robotic arm, rotation rates must be surveyed. If an object's rotation rate is beyond the capability of the ADR spacecraft to de-spin, then it is not a viable target. To ensure that an ADR mission does not arrive at a target and find that the rotation rate is too fast for the mission to succeed, it is important to understand its rotation rate dynamics in advance. Therefore, techniques for characterising this motion, and its evolution, from observation data are required.

To this end, light curve data was collected by IADC members on a range of rocket bodies in LEO and MEO. Estimates were then made of the rotation period of these rocket bodies. However, following the collection and initial analysis of these light curves, it was found that many of the rocket bodies exhibited unusual secular variations in their determined rotation rates. These results generated interest in the causes of the unusual rotation period variability and were part of the initial motivations for this PhD.

In 2017, IADC members were first made aware of the MMT light curve database. In 2018, a JAXA representative presented an examination of the MMT light curves of the

objects observed by IADC members and confirmed that the same rotation period variability was also present in the MMT data.

This chapter examines MMT light curves of the rocket bodies that were observed by IADC members. The variability of the determined rocket body rotation rates was investigated with an objective to determine the causes.

4.2 Methodology

This chapter makes use of phase dispersion minimisation (PDM) for determining the period of light curves. This technique was primarily selected to be consistent with the techniques used for the MMT database. It is intended for period determination in non-sinusoidal signals with sparse and irregularly spaced measurements, so it is also well suited for this application.

In PDM, the variance, σ^2 , of a total of N measurements, x , is given by

$$\sigma^2 = \frac{\sum (x_i - \bar{x})^2}{N - 1}, \quad (4.1)$$

where \bar{x} is the mean and is given by

$$\bar{x} = \frac{\sum x_i}{N}. \quad (4.2)$$

The variance of a subset of x is defined as s^2 and can be calculated using Equation 4.1. If the measurements are divided into M distinct samples, each with n_j data points and variances of s_j^2 , where $j = 1, M$, then the overall variance of the samples is

$$s^2 = \frac{\sum (n_j - 1) s_j^2}{\sum n_j - M}. \quad (4.3)$$

To perform PDM, the measurements x are divided into M samples. A test period value, T is then selected and x is ‘folded’ such that the measurements for each successive period are overlaid. Equation 4.1 is then used to measure the variance of each sample around the *mean curve*, which is given by the means of x_i in each sample. A new test period value T is then selected and the process is repeated.

To evaluate which value of T provides the best solution, the statistic

$$\theta = \frac{s^2}{\sigma^2}. \quad (4.4)$$

NORAD ID	Intl. Desig.	Name
10675	1978-018B	N-1 R/B
15392	1984-115B	DELTA 1 R/B(1)
25637	1999-008D	DELTA 2 R/B
25942	1999-057C	CZ-4 R/B
26906	2001-040B	ATLAS 2AS
27706	2003-010C	DELTA 2 R/B(2)
28476	2004-045C	DELTA 2 R/B(2)
28538	2005-004B	ATLAS 3B
31702	2007-027B	ATLAS 5
32712	2008-012B	DELTA 2 R/B(2)
34662	2009-014B	DELTA 2 R/B(2)
37151	2010-042B	CZ-3B R/B
37229	2010-062H	HAPS R/B STPSat
37365	2011-006B	MINOTAUR R/B
38253	2012-018D	CZ-3B R/B
39198	2013-033B	PEGASUS R/B
39475	2013-072P	ATLAS 5 CENTAUR R/B
40698	2015-028B	AVUM R/B
40720	2015-032F	PSLV R/B
40879	2015-040B	CZ-4C R/B
41847	2016-066G	CZ-11 R/B
41858	2016-068B	CZ-2D R/B
43161	2018-008G	CZ-11 R/B
43278	2018-034D	CZ-4C R/B

TABLE 4.1: The initial set of rocket bodies analysed in this chapter.

is defined. For incorrect period solutions θ will be approximately equal to 1.0, whereas for the correct solution, θ will be minimised. Hereafter, the property θ is referred to as the ‘dispersion’ of the fitted period (Stellingwerf, 1978).

4.3 Results and Discussion

Table 4.1 presents the rocket bodies identified by JAXA as interesting examples of rotation rate variability in the light curves collected by the IADC, which are also available in the MMT database. These objects form the initial sample set of rocket bodies that are considered in this chapter.

The MMT database supplies rotation rates calculated using PDM. Figure 4.1 plots the periods of the light curves, taken from the MMT database, for the objects in Table 4.1. Period data is no longer supplied once the analysis performed for the MMT database fails to extract a period, at which point the light curve is flagged as aperiodic. Therefore, each of the lines in Figure 4.1 ends once the light curve was determined to cease periodic behaviour. This is due to one of three reasons: 1) all subsequent light

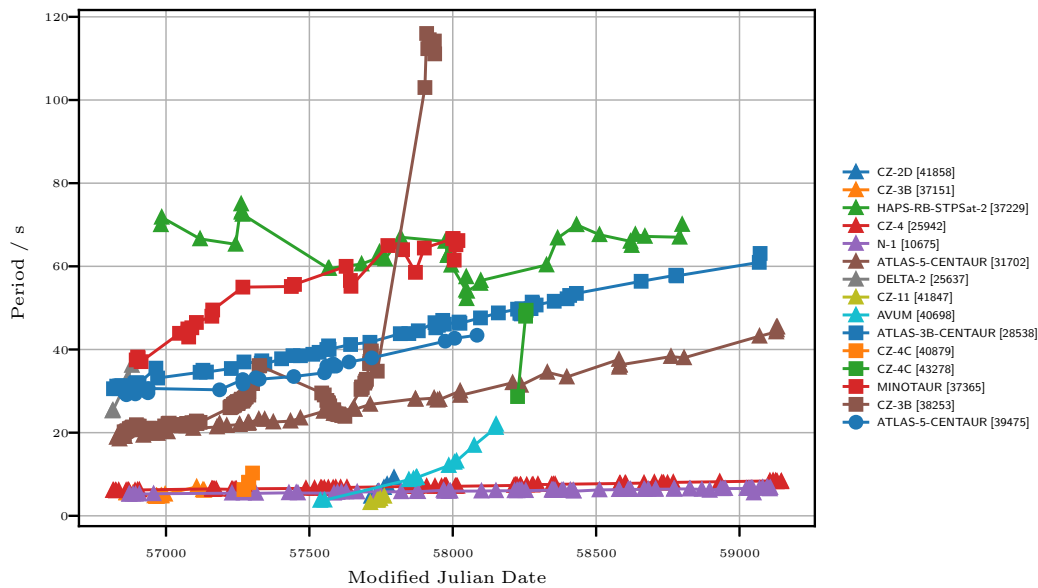


FIGURE 4.1: Rotation period, provided by the MMT database, against modified Julian date (MJD) for the selection of rocket bodies in Table 4.1.

curves were of insufficient quality for period extraction; 2) the light curves were no longer periodic, 3) the light curve periods were longer than the observation span.

The light curve period trends in this figure can be split into two categories. The first category exhibits orderly increases in rotation period. This is exemplified by the Atlas 3B Centaur rocket body (28538), which is plotted in blue with square markers. The second category exhibits complex changes in rotation rate, which is exemplified by the HAPS rocket body (37229) and is plotted in green with triangular markers.

To differentiate these two categories, the rotation periods of each object were fitted using a polynomial. The ‘orderly’ category was defined as those objects with increasing rotation periods for which the best-fitting number of polynomial degrees was either 1 or 2. The ‘complex’ category was all objects that do not meet this definition. The way in which the best-fitting polynomial number of degrees was determined is discussed below for each of the two categories individually.

Figure 4.2 shows the sum of the squared residuals of the least-squares fits, between the polynomial and the rotation period data, against number of degrees. Only those objects for which the best fitting number of degrees was less than or equal to 2 are included - the orderly category. The determined best fitting number of degrees is shown using a black cross. This value was determined in one of two ways. The first way was to use the degrees value where the sum of the least-squares residuals was below a threshold value used to define a satisfactory fit. This value was set at 0.025, and is indicated in the figure as a horizontal solid black line. The majority of the objects in this figure were fit in this way. The second way in which the best-fitting




FIGURE 4.2: The sum of the squared residuals of the least-squares fit between the polynomial and the rotation period data against number of degrees in the polynomial, for the orderly category.

number of degrees was determined was applied to those objects where the residuals sum was above the threshold when the degrees was set to 1. In these cases, the number of degrees was increased until one of two conditions was met. The first condition was that the least-squares residuals sum decreases below the 0.025 threshold. Objects Atlas-5 Centaur (39475, blue with circle markers) and CZ-11 (43161, pink with square markers) were fit in this way. The second condition was that two further increases in the number of degrees did not decrease the least-squares residuals sum by more than 10%. Objects N-1 (10675, purple with triangular markers), Atlas-5 Centaur (31702, brown with triangular markers) and Atlas-3B Centaur (28538, blue with square markers) were fit this way.

Figure 4.3 shows the period data for the orderly category with the best-fitting polynomials plotted as dotted lines. In this plot, the x-axis has been modified to show the days since the first periodic light curve and the y-axis has been modified to show normalised period. The normalised period is calculated by dividing all periods by the period of the first light curve. Hence, this plot shows the change in light curve period relative to its initial value.

Examination of Figure 4.3 shows that the objects in the orderly category can be further divided into two subcategories based on gradient. Treating the period data of each object as straight lines, these two subcategories can be differentiated by setting a gradient threshold of 0.018. Figure 4.4 shows only those objects with gradients greater than the threshold. This is referred to as the ‘fast de-spin’ subcategory. Figure 4.5

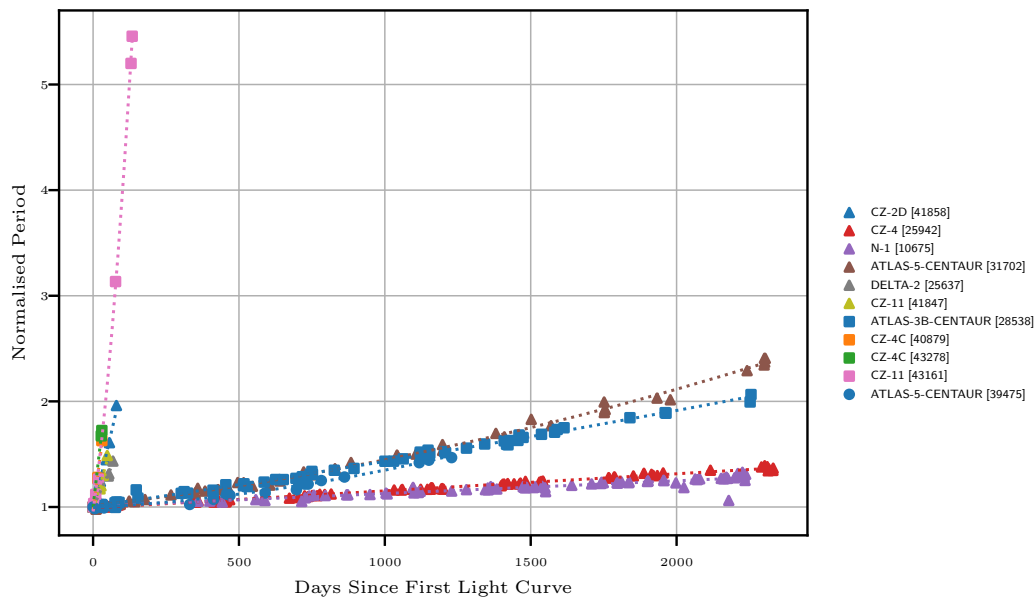


FIGURE 4.3: The best-fitting polynomials for the objects in the orderly category.

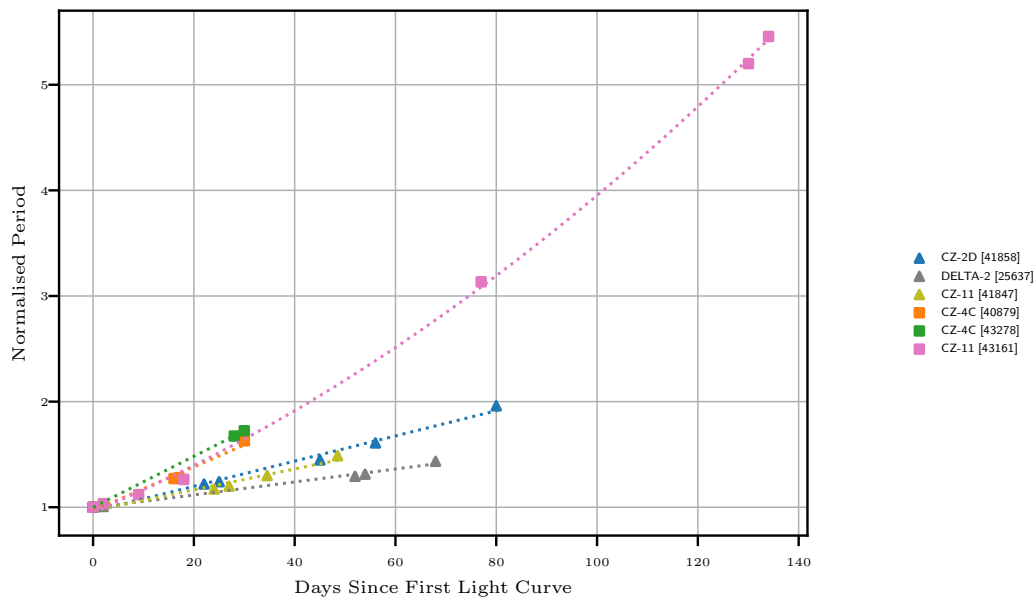


FIGURE 4.4: The best-fitting polynomials for the ‘fast de-spin’ subcategory of the orderly category.

shows only those objects with gradients less than the threshold. This is referred to as the ‘slow de-spin’ subcategory.

Next, the polynomial fitting process of the complex category is considered. Figure 4.6 shows the sum of the least squared residuals, against number of degrees, for the complex category. This category is defined by a best-fitting polynomial of order 3 or

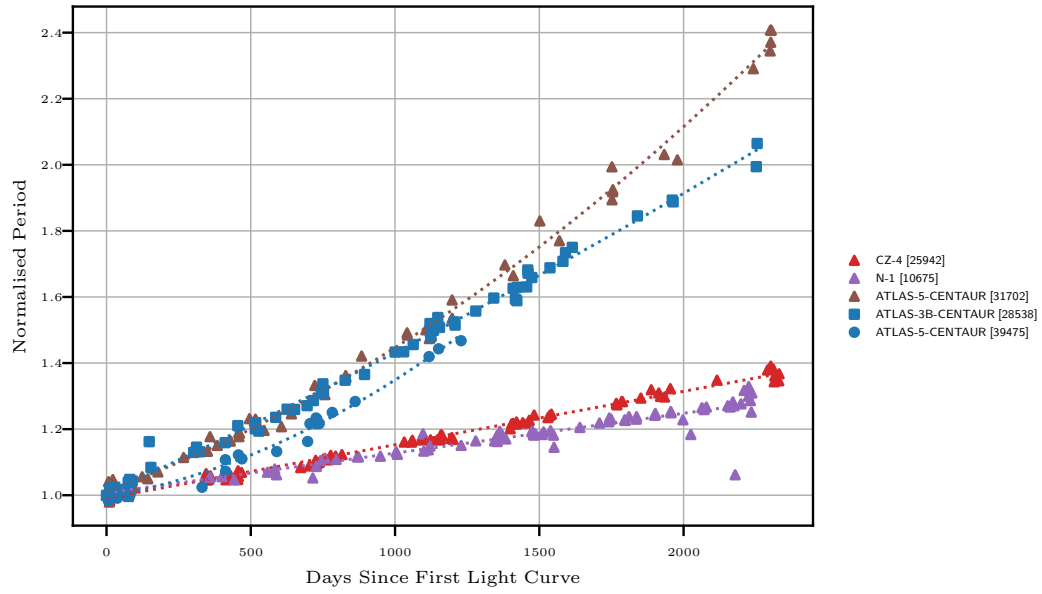


FIGURE 4.5: The best-fitting polynomials for the ‘slow de-spin’ subcategory of the orderly category.




FIGURE 4.6: The sum of the squared residuals of the least-squares fit between the polynomial and the rotation period data against number of degrees in the polynomial, for the complex category.

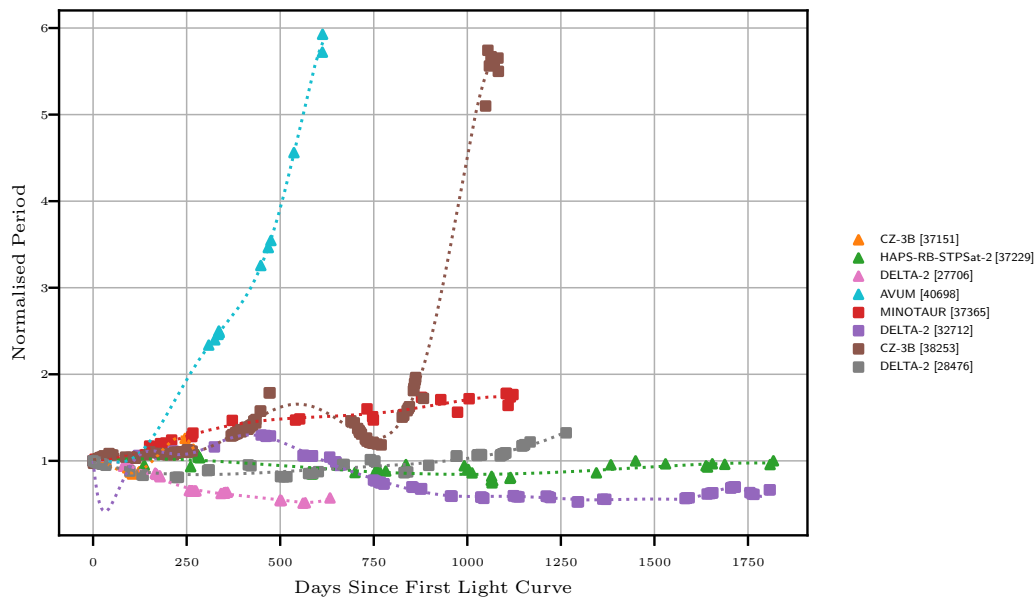


FIGURE 4.7: The best-fitting polynomials for the objects in the complex category.

higher. The best-fitting number of degrees was determined in the same two ways as for the orderly category: the first number of degrees for which the sum of the residuals was below the 0.025 threshold, or the number of degrees where two further increases did not decrease the sum of the residuals by more than 10%. It is probable that some of the best-fitting polynomials in this category have been overfit. However, the only objective in this fitting was to differentiate between the two categories and so this was not a major concern.

Figure 4.7 shows the period data for the complex category together with the best-fitting polynomials.

4.3.1 Analysis of Period Extraction

Before analysis of period variations could take place, the accuracy of the extracted periods needed to be assessed. It was important to understand whether the changes to light curve period observed in the above figures were the result of true changes in the periodicity of the data, and not caused by the period extraction process.

To answer the question, the PDM technique was again applied to the data. This was done so that the quality of the extracted periods could be examined in more detail than that provided in the PDM results from the MMT database. In each case, these new PDM results were compared to the PDM results from the MMT database to ensure agreement between extracted period values.




FIGURE 4.8: The dispersion in the best-fitting period of the light curves in the ‘fast de-spin’ subcategory.

4.3.1.1 Orderly Category Period Extraction

Fast De-Spin Subcategory

Figure 4.8 plots the dispersions of the best-fitting period for the light curves in the fast de-spin subcategory of the orderly category. In this section, objects CZ-2D (41858, blue line with triangular markers) and CZ-4C (40879, orange with square markers) are examined in closer detail to determine the causes of the large differences in minimum dispersion of the best-fitting period. These two objects were selected as they exhibit the largest and smallest dispersion values.

Figure 4.9 shows the PDM results for the light curve of object CZ-2D that had the lowest dispersion in the extracted period. On the left, the original light curve is plotted. In the centre, dispersion is plotted against period. The period value that minimises the dispersion is shown with a vertical dotted green line and the period value provided by the MMT database is shown with a vertical dashed-dotted red line. On the right, the apparent magnitude is split at the end of each period. Plotting the apparent magnitude against the decimal period then results in each period overlaying onto the last, providing a folded light curve. In this case, the best-fitting period was 6.89 seconds. Therefore, the decimal period indicates multiples of this value. For comparison, Figure 4.10 shows the PDM results for the light curve of object CZ-2D that had the highest dispersion in the extracted period.

The folded light curves in these two examples exhibit the same shape. In the folded light curve of Figure 4.9: there is a primary brightness peak, when the decimal period is 0.0 (labelled with a red “1”); a primary minimum, at approximately 0.25 (labelled with a “2”); and another primary brightness maximum, at approximately 0.5 (labelled with a “3”). The light curve then has a secondary minimum at approximately 0.6 (labelled with a “4”), a secondary maximum at approximately 0.75 (labelled with a “5”) and another secondary minimum at approximately 0.9 (labelled with a “6”). The same sequence is present in the folded light curve of Figure 4.10, but with a different starting point. In this case, there is a primary minimum at a decimal period of 0.0 (labelled with a red “1”) followed by a primary maximum at 0.3 (labelled with a “2”). Then there is a secondary minimum at 0.4 (labelled with a “3”), a secondary maximum at 0.55 (labelled with a “4”) and another secondary minimum at 0.7 (labelled with a “5”). Another primary maximum then occurs at 0.8 (labelled with a “6”). This folded light curve shape was confirmed to be present in all the other results for this object. Changes in light curve period are caused by a ‘stretching’ of this shape. This shows that the changes in light curve period were caused by actual changes in signal periodicity and were not caused by erroneous period extraction results.

A possible reason for this shape can be provided by considering the surfaces of a rocket body from which light can be reflected. At the most basic level, a rocket body is an approximately cylindrical object. Therefore, if the rocket body is tumbling, a repeating sequence of reflections would be expected from the curved sides of the cylinder and the two end caps. This sequence may manifest as: a reflection from the side, followed by a reflection from one end cap, followed by another reflection from the side, followed by a reflection from the opposite end cap, then another reflection from the side. This concept is explored in greater detail in Chapter 7, which is concerned with attitude determination of rocket bodies.

In both of these examples, the dispersion minimum is unique. The only comparable minima appear at multiples of either 0.5 or 1.5 of the extracted period, which are caused by ‘over-folding’ or ‘under-folding’ of the light curve. This plot shows that there are no other period values that can generate folded light curves with comparably low dispersions.

Next, the PDM results of CZ-4C (40879) were considered. The PDM results for this object exhibit the highest dispersions in the extracted periods of any of the objects in Figure 4.8 and so are an example of a possibly lower quality result. Figure 4.11 shows the PDM result with the lowest dispersion for this object. In this plot an inaccurate period has been extracted due to folding at half of the true period, which resulted in a lower dispersion. This result shows that additional processing of the PDM results is performed for the MMT database in order to extract an accurate period. This can be done by reevaluating an extracted period based on the periods extracted from other light curves of this object to ensure consistency in the solutions. If the range of periods




FIGURE 4.9: The PDM result for the light curve of object CZ-2D (41858) with the lowest dispersion.




FIGURE 4.10: The PDM result for the light curve of object CZ-2D (41858) with the highest dispersion.

that are tested are reduced to exclude multiples of 0.5 and 1.5 the MMT period, then the PDM results of Figure 4.12 are obtained. Because this work is concerned with only the change in signal period, ‘over-folding’ or ‘under-folding’ the light curve is not a problem, as long as it is consistent across all light curves. This can be tested by confirming that the folded light curves exhibit a consistent shape across all period extraction results.

The reason why the dispersion in this result is higher than the dispersions in the results of object CZ-2D (41858) is due to noise in the light curve. Considering the apparent magnitudes of the light curves of object CZ-2D (41858), the minimum brightness is approximately 3.5. In these results for object CZ-4C (40879), the minimum brightness is an apparent magnitude of approximately 7.0. This is, therefore, a much dimmer light curve and so more susceptible to noise, resulting in a higher dispersion in the folded light curve.

Figure 4.13 shows the PDM result with the highest dispersion for this object. In this




FIGURE 4.11: An erroneous PDM result for the light curve of object CZ-4C (40879) with the lowest dispersion.




FIGURE 4.12: A corrected PDM result for the light curve of object CZ-4C (40879) with the lowest dispersion.

case, the dispersion has been inflated by a change in average brightness throughout the start of the light curve. This shows that the PDM technique is not well suited to signals with a changing average value. Therefore, the value of the dispersion is not necessarily indicative of a good or bad period extraction solution.

Comparing the shapes of the two folded light curves in Figures 4.12 and 4.13, both exhibit two primary maxima and two primary minima. Once again, changes in light curve period are caused by a ‘stretching’ of this shape. This consistency in shape indicates that the period variability has been accurately extracted.

These results have shown that the minimum dispersions are not indicative of the accuracy of the PDM solution. Instead, consistency in the shape of the folded light curves can be used to provide confidence that the period trends have been extracted correctly. The light curves of all objects in the fast de-spin subcategory were examined to ensure that this was the case.




FIGURE 4.13: The PDM result for the light curve of object CZ-4C (40879) with the highest dispersion.

The primary reason why a period result may be incorrect is due to fitting multiples of 0.5 or 1.5 the correct solution. This does not occur in the period data provided by the MMT database, indicating that some pre-processing, post-processing, filtering, or application of some additional processes beyond simple PDM analysis was applied to the results to ensure consistency.

Slow De-Spin Subcategory

Next the PDM results for the slow de-spin subcategory were considered. Figure 4.14 shows the dispersions in the best-fitting period values for this subcategory. The object in this plot that exhibits the lowest dispersions in the extracted periods is ATLAS-3B-CENTAUR (39475), plotted in blue with circle markers. The PDM result with lowest dispersion for this object is shown in Figure 4.15. This plot shows that a low dispersion in the extracted period does not necessarily indicate an accurate solution. In this example, the total length of the light curve, 70 s, is equal to 1.5 times the extracted period of 45 s. There was, therefore, only a small amount of overlap in the folded light curve from which dispersion could be measured. Hence, although the dispersion is low, the lack of data results in a higher uncertainty.

Figure 4.16 shows the PDM result for ATLAS-3B-CENTAUR (39475) with the highest dispersion. This result shows, once again, that the dispersion can be inflated due to changes in average brightness throughout a light curve, causing successive periods to not perfectly overlay when the light curve is folded. Despite the larger dispersion, this light curve contains a larger number of periods than the previous figure and so the uncertainty in this solution is lower. In both this and the previous figure, the folded light curves contain a consistent number of maxima and minima. Each PDM result for this object was examined to ensure that this was true in all cases.




FIGURE 4.14: The dispersion in the best-fitting period of the light curves in the ‘slow de-spin’ subcategory.




FIGURE 4.15: The PDM result for the light curve of object ATLAS-3B-CENTAUR (39475) with the lowest dispersion.

Next, the object in the slow de-spin subcategory that exhibited the highest dispersions (CZ-4, 25942, plotted in red with triangular markers), was examined. Figure 4.17 shows the PDM result with the lowest dispersion for this object. Similar to a previous result, the dispersion in this folded light curve is low due to the short time span of the original light curve. Figure 4.18 shows the PDM result with the highest dispersion. The dispersion in this case has been inflated for two reasons: firstly, because of changes in average brightness; secondly, because of increased noise due to a lower brightness. In both Figure 4.17 and Figure 4.18, the folded light curve shape is consistent with two maxima and two minima. Consistency in the number of maxima and minima was confirmed to be the case throughout all the light curves in the slow




FIGURE 4.16: The PDM result for the light curve of object ATLAS-3B-CENTAUR (39475) with the highest dispersion.




FIGURE 4.17: The PDM result for the light curve of object CZ-4 (25942) with the lowest dispersion.

de-spin subcategory.

Examination of the light curves in the slow de-spin subcategory confirmed that the light curve shape remained consistent and that changes in light curve period were the result of a ‘stretching’ of the shape.

4.3.1.2 Complex Category Period Extraction

Finally, period extraction in the complex category was examined. The minimum dispersions of this category are presented in 4.19. For this category, object CZ-3B (38253, plotted in brown with square markers) is considered in greater detail, as it shows the largest spread in dispersions throughout the light curve data. Figure 4.20 shows the PDM results with the lowest dispersion for this object. This is another example of how a low dispersion does not necessarily indicate a period extraction solution with low uncertainty. The short length of the light curve compared to the




FIGURE 4.18: The PDM result for the light curve of object CZ-4 (25942) with the highest dispersion.




FIGURE 4.19: The dispersion in the best-fitting period of the light curves in the 'complex' category.

determined period results in a low number of folds, causing a low dispersion to be measured. The shape of the folded light curve exhibits two maxima and two minima. This defined the folded light curve shape that was required to remain constant in all folded light curves of this object, to be confident that changes in extracted period were the result of changes in the periodicity of this signal, and not due to errors in the extraction.

Figure 4.21 shows the PDM results with the highest dispersion for this object. This shows another example of how the PDM analysis performed for the MMT database involves further processing, rather than simply finding the period value that




FIGURE 4.20: The PDM result for the light curve of object CZ-3B (38253) with the lowest dispersion.




FIGURE 4.21: An erroneous PDM result for the light curve of object CZ-3B (38253) with the lowest dispersion.

minimised the dispersion. In this PDM example, high noise and a change in brightness at the end of the light curve caused an erroneous global minimum in the dispersion.

For Figure 4.22, the period was constrained to within 3 seconds of the period provided by the MMT, and only the first 400 seconds of the light curve were used. The folded light curve in this figure now shows the expected shape - two maxima and two minima. The extracted period in this figure is approximately 30 s, whereas the extracted period from Figure 4.20 was approximately 115 s, where the difference between these two periods is caused by a stretching of the light curve shape.

The PDM results for all other objects in the complex category were examined to confirm that the extracted folded period shapes were consistent.




FIGURE 4.22: A corrected PDM result for the light curve of object CZ-3B (38253) with the lowest dispersion.

4.3.1.3 Summary of Period Extraction Results

Throughout this section, a selection of PDM results have been presented showing the process by which light curve periods were extracted. It was shown that the value of the dispersion is not indicative of the quality of the solution. For example, if the total length of the light curve is similar to the extracted period, then the light curve can only be folded once. This results in a low folded light curve dispersion, but the lack of data results in a lower quality extracted period. As another example, if the average magnitude of the light curve changes through time, then each successive fold of the light curve will be different from previous folds. This causes a higher dispersion in a period extraction result even though the large number of folds causes the quality of the result to be high.

It is possible for these deficiencies in the application of the PDM technique to be mitigated. For example, a change in average magnitude throughout the light curve was one cause of high dispersion. De-trending the light curves before application of the PDM could be used to prevent these higher dispersions in what would otherwise be high quality period extraction results. Another cause of higher dispersions was noise in the light curve. Application of a filter to the light curves to remove the higher frequency noise could be used to again prevent unnecessarily high dispersions. Applying these two mitigation techniques would result in a substantial reduction in dispersions across all light curves and would allow for a more direct comparison between the dispersions in different PDM results. However, these changes do not address the low dispersion values received from a lower number of folds. It would remain the case that light curves containing a larger number of periods, which is preferable for producing high quality period determination results, would exhibit larger dispersions than light curves containing only a single period. Therefore, development of a metric that could be used to quantify the quality of a particular

period extraction result would require both de-trending and filtering of the light curves to receive the most accurate dispersions, in addition to analysis of total number of periods contained in the light curve. However, the objective of this section was simply to ensure that the period results provided by the PDM technique are true representations of the periodicity of the considered light curves, which was instead achieved by ensuring consistency in the folded light curve shape.

Using the assumption was that a light curve of a tumbling rocket body will exhibit a repeating sequence of brightness levels as the top, bottom or sides of the rocket body become the dominant source of reflected light. This will produce a light curve with a repeating sequence of brightness maxima and minima. Hence, if the rotation period changes, these maxima and minima will still be present but the time taken to complete one cycle will be different, resulting in a 'stretching' of the shape. Working under this assumption, the light curves of all objects were examined to ensure that the number of maxima and minima was consistent.

4.3.2 Causes of Changes in Periodicity

The results of the previous section have indicated that the observed changes in light curve period were caused by changes in the periodicity of the signal and were not caused by the period extraction technique. The next step was to explore the reasons why the period data could be categorised into the orderly and complex categories.

The first possibility explored was that the changes in period were a result of the specific response to on-orbit torques from particular designs of rocket body. However, this was disproved by the fact that CZ and DELTA-2 rocket bodies appear in both the orderly and complex categories. This shows that the presence of a rocket body in a particular category is independent of the rocket body design.

The next possibility was that the presence of a rocket body in a particular category was caused by the characteristics of its orbit. To explore this, Figure 4.23 was generated, which shows the perigee altitude, apogee altitude and eccentricity of the rocket bodies over the time period that the periodic light curves were collected. This information was obtained from the Space-Track database and the objects have been coloured by category. In each plot, the average value of the orderly data is plotted with a dashed line and the average value of the complex data is plotted with a dotted line. It should be noted that 3 objects from the complex category and 2 objects from the orderly category are missing from this figure. This was because no orbital information was available for these objects in the Space-Track database, probably because they were used to launch classified payloads.

Figure 4.23 shows a clear difference between the orbital data of the two categories. In the perigee altitude plot at the top of the figure, the perigees of the objects in the




FIGURE 4.23: The perigee altitude, apogee altitude and eccentricity of all objects.

complex category are all approximately 200 km. The perigee altitudes of the orderly category are approximately 500 - 1000 km. The middle plot shows that the apogee altitudes of the complex category all begin above 4000 km. Additionally, these apogee altitudes all decrease over time as a result of the high atmospheric drag caused by a 200 km perigee. The apogee altitudes of the orderly category remain relatively constant at approximately 1000 km. The eccentricity plot at the bottom of the figure presents this same information in a different format. It shows that objects in the complex category are in eccentric orbits, with atmospheric drag causing eccentricity to decrease, whereas objects in the orderly category are in stable low-eccentricity orbits.

This suggests that the difference in rotation period variability is caused by the

Category	Total	Correct	Incorrect	Percentage Correct
Complex	48	46	2	95.8%
Orderly	19	8	11	42.1%

TABLE 4.2: The categorisation statistics of the CZ rocket bodies plotted in Figure 4.24

characteristics of the orbits. All objects in the orderly category have low eccentricities and hence the magnitudes of the experienced torques will remain approximately constant throughout the orbit. This may explain why all objects in this category exhibit mostly linear changes in rotational motion. Conversely, all objects in the complex category have minimum orbital altitudes in the region of approximately 200 km. This is a very low altitude for an orbiting object and so it will experience a significant force due to the higher atmospheric density at this altitude. Depending on the object's attitude and the position of its centre of mass, this will result in relatively large torques being applied to the object. However, because all of these objects have maximum orbital altitudes on the order of 1000s of km, these atmospheric torques will reduce to effectively zero for the remainder of the orbit. Additionally, at these high altitudes and in the absence of atmospheric torques, other torque sources such as solar radiation pressure and eddy currents will begin to dominate. The cyclical application of these torques may result in attitude-orbit coupling and introduce resonances into the objects spin rate. This may be a possible reason for the complex attitude motion behaviour.

There is one object that does not fit the pattern, AVUM (40698), which has been coloured blue. The period data of this object was best-fit with a 5-degree polynomial, causing it to be a member of the complex category. However, the orbital data shows that this object has a perigee altitude of approximately 400 km and an apogee altitude of approximately 1000 km, which is a lower eccentricity orbit relative to the objects presented in this plot. This object is, therefore, more similar to the orderly category in terms of its orbital characteristics.

To further test this relationship between rotation period variability and orbital data, the analysis was expanded to include all CZ, ATLAS and DELTA rocket bodies available in the MMT database.

4.3.2.1 CZ Rocket Bodies

Figure 4.24 shows the orbital data for CZ rocket bodies. The period data for these objects was again fit using a polynomial, and the objects were categorised based on the number of degrees used in the best fit. A dashed and a dotted line have again been provided to show the average values of the orderly and complex categories respectively.

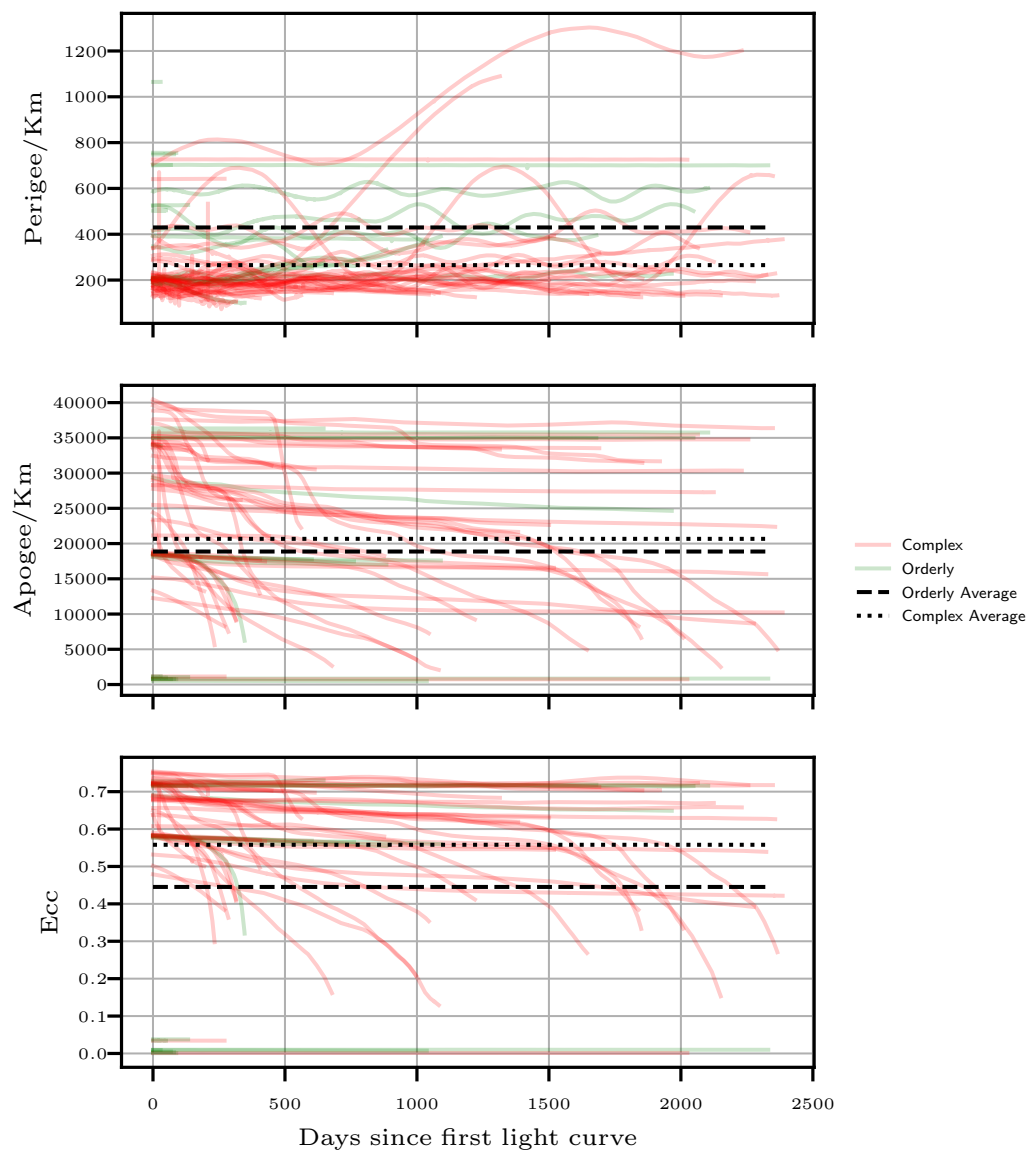


FIGURE 4.24: The perigee altitude, apogee altitude and eccentricity of CZ rocket bodies in the MMT database.

Based on the results of Figure 4.23, the objects with a complex categorisation (coloured red) were expected to have initial eccentricities greater than 0.2 and objects with an orderly categorisation (coloured green) were expected to have initial eccentricities less than 0.2. Table 4.2 presents the number of objects in this figure that were in each category and the number of objects in each category that exhibited the expected eccentricity. This table shows that of the 48 objects that were categorised as complex, 46 also exhibited the expected high eccentricity. However, of the 19 objects categorised as orderly, only 8 showed the expected low eccentricity.

Considering the perigee altitude data at the top of Figure 4.24, there is a clear concentration of objects that are categorised as complex that have perigee altitudes

Category	Total	Correct	Incorrect	Percentage Correct
Complex	16	12	4	75.0%
Orderly	8	3	5	37.5%

TABLE 4.3: The categorisation statistics of the DELTA rocket bodies plotted in Figure 4.25

close to 200 km. This is reflected by the fact that the average complex perigee altitude is approximately 250 km. As expected from Figure 4.23, the perigee altitudes of the orderly category are higher on average.

Considering the apogee altitude data in the middle of Figure 4.24, the majority of objects in the complex category have initial apogees greater than 10,000 km, leading to the expected high eccentricity. However, many of the objects categorised as orderly also exhibit apogees greater than 10,000 km, contrary to the expectation from Figure 4.23.

This is reflected in the information presented in the eccentricity plot at the bottom of Figure 4.24. A large majority of the objects in the complex category exhibit the expected high eccentricity (95.8%), with a small number of outliers. Some objects categorised as orderly exhibit the expected low eccentricity (42.1%). However, more than half exhibit the high eccentricity that is typical of objects in the complex category.

4.3.2.2 DELTA Rocket Bodies

Next, the orbital data of DELTA rocket bodies was considered. The orbital data is shown in Figure 4.25 and the categorisation statistics are shown in Table 4.3. This data shows a similar results distribution to the previous figure. The majority of objects in the complex category have the expected high eccentricity (75.0% have an initial eccentricity greater than 0.2), but only a minority of the objects in the orderly category have the expected low eccentricity (37.5% have initial eccentricities less than 0.2).

4.3.2.3 ATLAS Rocket Bodies

Finally, the orbital data of the ATLAS rocket bodies was considered. The orbital data is shown in Figure 4.26 and the categorisation statistics are shown in Table 4.4. This data shows that objects in the complex category have the expected high eccentricity at a similar level as the results for CZ and DELTA rocket bodies (87.0%). However, in this case, of the 32 objects categorised as orderly, zero exhibited the expected low eccentricity.

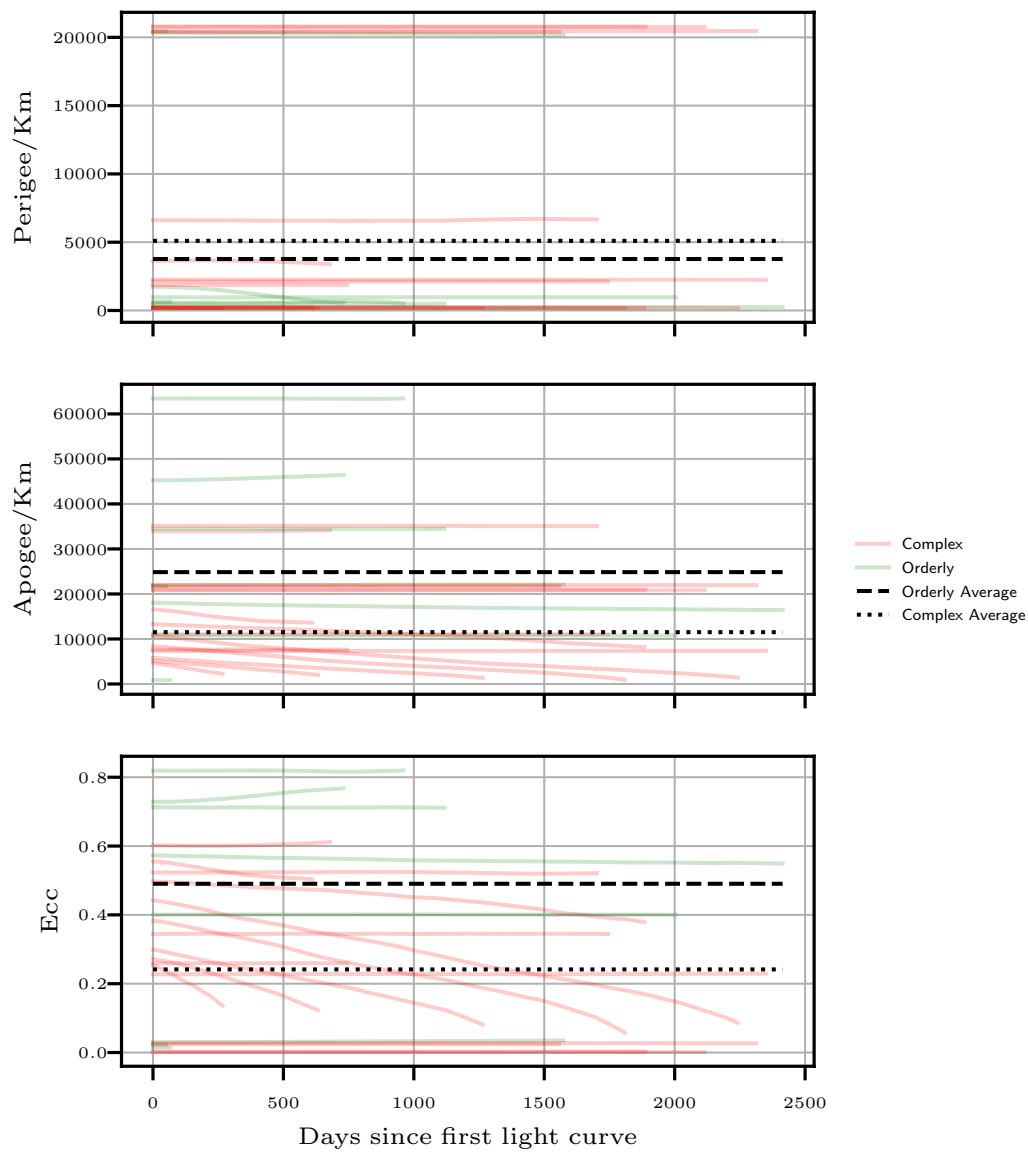


FIGURE 4.25: The perigee altitude, apogee altitude and eccentricity of DELTA rocket bodies in the MMT database.

Category	Total	Correct	Incorrect	Percentage Correct
Complex	46	40	6	87.0%
Orderly	32	0	32	0.0%

TABLE 4.4: The categorisation statistics of the ATLAS rocket bodies plotted in Figure 4.26

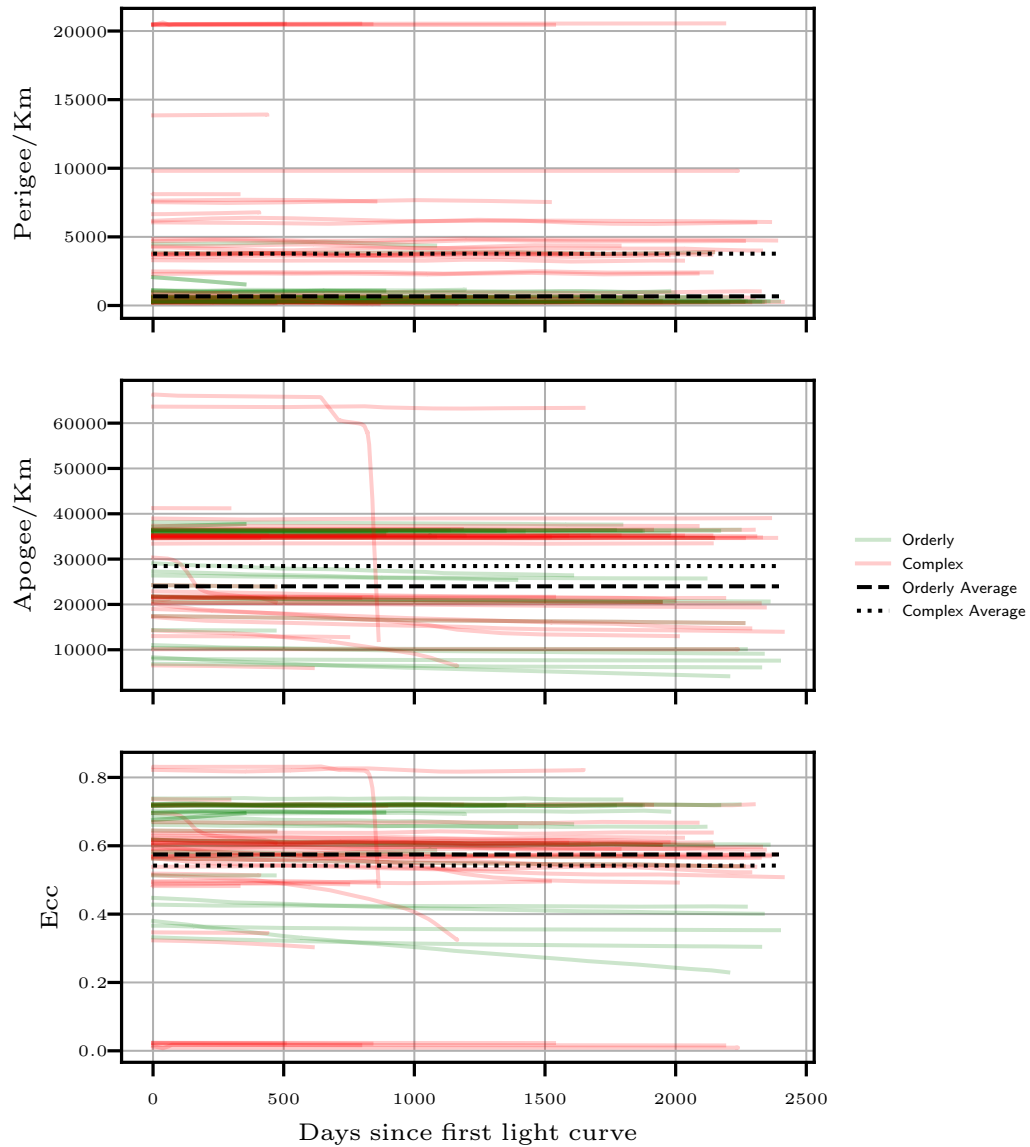


FIGURE 4.26: The perigee altitude, apogee altitude and eccentricity of ATLAS rocket bodies in the MMT database.

4.3.2.4 Summary of Causes of Changes in Periodicity Results

In the initial examination of the orbital data for the original set of objects that were considered in this chapter (Figure 4.23), the results suggested that the characteristics of an object's orbit may be the cause of the differences in period variability. This was because, with a single exception, all objects in the orderly category were in orbits with eccentricities less than 0.2, and all objects in the complex category exhibited eccentricities greater than 0.2. However, this was not the case when the number of objects was expanded to include all CZ, DELTA and ATLAS rocket bodies in the MMT database. The percentage of objects in the complex category with the expected high

eccentricity remained high and constant for all three rocket body types at $85 \pm 10\%$. However, in the orderly category, the percentage of objects with low eccentricity was always less than 50% and in one case was 0%.

The additional data considered for these three rocket body types was not considered in the same detail as the data that was used throughout this chapter. It is therefore possible that there was a greater occurrence of anomalous results, or a breakdown in the categorisation methodology.

4.4 Conclusions

In this chapter, the periodicities of rocket body light curves have been examined. In a subset of rocket bodies that were presented at the IADC as having unusual periodicity variability, it was found that the objects could be categorised as having either 'orderly' or 'complex' period data.

Orderly period data was defined by best-fitting polynomial functions that used 1 or 2 degrees. The periods of these objects increased at a relatively constant rate over time. This orderly category could be further broken down into two subcategories, depending on gradient. These were termed the 'slow de-spin' and 'fast de-spin' subcategories.

Complex period data was defined by best-fitting polynomial functions of greater than 2 degrees. These objects exhibited non-uniform increases and decreases in rotation period.

Before analysis could proceed on the possible causes of these period variations, it was important to establish confidence that the periods provided by the MMT database were true reflections of a change in signal periodicity, and not caused by the period extraction process. This was achieved by examining the folded light curves of each object to ensure that the shape, in terms of number of maxima and minima, remained constant. This indicated that the observed changes in period were the result of real changes in light curve periodicity.

Analysis of the original orbital data showed that the objects in the orderly and complex categories shared the same orbital characteristics. Objects in the complex category were in high-eccentricity orbits and objects in the orderly category were in low-eccentricity orbits. There are a number of mechanisms that may cause the 'spin-up' 'spin-down' behaviour observed in the complex category. These include eddy current and radiation pressure torques, which can have a more significant effect on higher eccentricity orbits due to the increased amount of time they spend away from lower orbital altitudes where atmospheric torques will dominate. This may lead to orbit-attitude coupling which introduces a resonance into the rotational motion.

This categorisation was less clear once all CZ, DELTA and ATLAS rocket bodies in the MMT database were included. There are three possible reasons. Firstly, the categorisation technique developed for use with the original data set loses efficacy when applied to large numbers of objects. As a result, objects may be erroneously categorised as orderly, when in fact their period data was more similar to objects in the complex category. Secondly, there may be additional categories and subcategories that were not exemplified in the limited set of objects from which the categorisation was defined. Both of these points could be improved by reevaluating the categorisation technique. Thirdly, there may be further factors that cause an object to be present in a category other than orbital characteristics. These three reasons can be addressed in future work by continuing and expanding this analysis.

Chapter 5

Synthetic Light Curve Model Sensitivity Study

5.1 Introduction

If using a synthetic light curve model to develop understanding and interpretation of real light curve data, then it is important to understand the operation of that model. Of particular interest is the sensitivity of the model to the model inputs. This knowledge is important for improving the efficiency and quality of the fitting process, as additional focus can be given to those parameters that the synthetic light curve model is most sensitive to. Additionally, this knowledge will improve understanding of the strengths and limitations of this relatively complex model, which will allow for better interpretation of its results.

To achieve this, a sensitivity study was performed on the synthetic light curve model. The objective was to quantify the change in a simulated light curve that results from changes to the inputs and so identify the parameters that have the largest effect on the light curve brightness.

5.2 Methodology

To perform this sensitivity study, a light curve scenario was constructed that allowed for each parameter to be varied and the resulting difference to the synthetic light curve to be quantified. Some light curve generation parameters were kept constant across all of the sensitivity simulations. These were light curve duration, which was set to 180 s, and distance from the object to the observer, which was set to 1000 km. Duration has no effect on the differences between light curve brightness, and distance affects all light curves by an equal amount.

In each simulation, the five geometry models in Figure 5.1 were analysed. The first model at the left of the figure is a 20-sided regular icosahedron where all vertices are 1 m from the geometric centre. In each subsequent geometry model the facets of the 20-sided icosahedron were divided into an increasing number of equal-area triangles and their positions were altered such that all vertices were a distance of 1 m from the geometric centre. For the second model, each facet was divided into 4 facets, giving a total of 80 facets. For the third model the facets were divided into 9 facets giving a total of 180. For the fourth model each facet was divided into 16 facets giving a total of 320. For the final geometry model each facet was divided into 25 facets giving a total of 500. This range of geometry models was selected as they represent a transition from an approximately spherical geometry, for which brightness should be independent of attitude state, to a geometry with a smaller number of larger-area facets, for which each individual facet has a much larger effect on light curve brightness.

The remaining inputs that were varied were the following: the angle between the illumination and observation vectors (phase angle); the three BRDF parameters associated with each facet; the initial attitude state; and the angular rotation rate. In the cases where the attitude of the object was modified (either as a change in initial attitude state or by introduction of a rotational motion), the rotation was applied to one of three body-fixed orthogonal axis, where the body-fixed x-axis is initially set to be coincident with the vector through which the object is observed. This coordinate system is shown in Figure 5.2. The simulations were performed in this way because it was not feasible to test all possible combinations of initial attitude state and rotational motions in terms of computational cost and results communication. For each simulation the difference that resulted from varying each input was quantified by measuring an RMSE between the light curve and a baseline light curve. The baseline light curve is described in the discussion of each simulation result. All of these results were generated twice, firstly using an entirely diffuse reflection model and again using an entirely specular reflection model.




FIGURE 5.1: The five geometry models used in this chapter.

5.3 Results

5.3.1 Sensitivity to Phase Angle

The first result examines the sensitivity of the light curve to phase angle. To simplify the possible combinations of illumination and observation vector directions, the observation vector remained constant and aligned with the x-axis of Figure 5.2 and the illumination vector was offset into the direction of the y-axis by the angle θ . In this experiment, the reflection was entirely diffuse ($d = 1.0$) and the diffuse albedo was set to 0.9. The value of the roughness parameter, m , was irrelevant as it only affects the specular component of reflection. The object had rotation rate of zero and so remained in the orientation shown in Figure 5.2 throughout the light curve.




FIGURE 5.2: Illumination geometry used to quantify the change in a light curve that results from increasing the phase angle θ .

Figure 5.3 shows the RMSE (Eq. 2.13) plotted on the y-axis, against the number of facets in the geometry model plotted on the x-axis, as the angle between the

observation and illumination vectors (phase angle) is increased. The reflection model was entirely diffuse ($d = 1$). The phase angle that corresponds to each coloured line, expressed as multiples of $\pi/2$ radians, is shown in the legend. The black star denotes the position of the baseline light curve from which all RMSE values are measured. The baseline in this case was the light curve for the 500-facet polyhedron, with a zero phase angle. This plot therefore shows two pieces of information. First, by following the lines from right to left, the change in the light curve that resulted from the transition from an approximately spherical object geometry to the 20-sided icosahedron is shown. Second, by observing the vertical change between the different lines, the change in the light curve that resulted from the increasing phase angle is shown.

Figure 5.3 shows that phase angle has a larger effect on synthetic light curve brightness than the accuracy of the spherical approximation. This is indicated by the fact that the maximum difference between the different lines is larger than the change in a line that resulted from decreasing the number of facets. It is further shown that the difference introduced by increasing the phase angle becomes larger as phase angle becomes larger. This is indicated by the increased spacing between the lines as phase angle is increased. The change to the light curve caused by decreasing the number of faces is the same across all phase angles, and only begins to have a noticeable effect once the number of facets drops below 150. The maximum difference that resulted from increasing phase angle was approximately 1.3 and the maximum difference caused by reducing the number of facets was approximately 0.3. All RMSE difference values presented throughout these sections are summarised in a table at the end of this chapter, so that they can be compared more easily.




FIGURE 5.3: The change in a synthetic light curve, quantified using RMSE, that resulted from increasing the phase angle between the illumination and observation vectors. A entirely diffuse reflection model was used.

Figure 5.4 shows the result of a simulation with the same inputs as the previous case, but now using an entirely specular reflection model. This figure shows that switching to a specular model of reflection changes the response of a simulated light curve to phase angle. In this plot, decreasing the number of facets, rather than increasing phase angle, has the largest effect on the light curve. The small differences between the light curves of the 500-facet polyhedron at all phase angles is because the brightness is primarily due to specular reflection from the small number of facets orientated to reflect light in the direction of the observation vector. Increasing the phase angle only serves to change the specific facets from which this reflection is received. As the number of facets are reduced, there are fewer viewing conditions that can reflect specularly in the direction of the observation vector, but the brightness of this reflection is increased as a result of the larger facet area. This is why larger differences are measured between the light curves at different phase angles once the number of facets has been reduced. The irregularities of the lines in this plot, relative to the cleaner distribution of the lines in the previous case, is caused by specular reflection resulting from particular illumination conditions. These particular conditions are such that strong specular reflections are detected, that can be orders of magnitude brighter than reflections caused by slightly different illumination conditions. The maximum difference that resulted from increasing phase angle was approximately 3.5 and the maximum difference caused by reducing the number of facets was approximately 3.75.




FIGURE 5.4: The change in a synthetic light curve, quantified using RMSE, that resulted from increasing the phase angle between the illumination and observation vectors. A entirely specular reflection model was used.

5.3.2 Sensitivity to BRDF Parameters

In all subsequent results, the phase angle remained constant at $(0.5 * \pi/2)$ radians, or 45° . In the next set of results, the BRDF parameters ω , d and m are examined.

Figure 5.5 shows the change in the simulated light curve that resulted from decreasing the diffuse albedo, under an entirely diffuse reflection model. The baseline was the light curve for the 500-facet polyhedron with a diffuse albedo of 0.9. This result is very similar to that for increasing phase angle with diffuse reflection, which was presented in Figure 5.3. First, the difference is greater when reducing the diffuse albedo, than from reducing the number of facets. Second, reducing the number of facets had only a very small effect on the light curve until number of facets decreased below approximately 150 facets. Finally, the change introduced into the light curve by decreasing the diffuse albedo became larger when further from the baseline light curve. The maximum difference that resulted from decreasing the diffuse albedo was approximately 2.4 and the maximum difference caused by reducing the number of facets was approximately 0.4.




FIGURE 5.5: The change in a synthetic light curve, quantified using RMSE, that resulted from decreasing the BRDF parameter omega (diffuse albedo).

Figure 5.6 plots the change in a simulated light curve as the ratio of diffuse to specular reflection, d , is decreased. The parameter m was set to 0.2, ω was set to 0.9 and for the baseline light curve, $d = 1.0$. This graph shows the transition from purely diffuse, to purely specular reflection. The plot shows that the largest change to the simulated light curve is caused by the initial introduction of the specular component (between $d = 1.0$ and $d = 0.9$). As the light curve becomes more and more specular, the increase in the difference from the baseline light curve becomes smaller. As the reflection becomes increasingly specular, decreasing the number of facets becomes a larger cause

of difference in the light curve. Additionally, the RMSE values are much larger in this plot than in all other figures thus far. In the previous cases, the largest measured difference was approximately 3.75, which was caused by increasing the phase angle under purely specular reflection conditions, (Figure 5.4). In Figure 5.6, differences larger than 8.0 were measured, more than twice the previous maximum. The maximum difference caused by reducing only the number of facets was approximately 1.5.




FIGURE 5.6: The change in a synthetic light curve, quantified using RMSE, that resulted from decreasing the BRDF parameter d (ratio of diffuse reflection).

Figure 5.7 shows the differences that result from increasing the roughness parameter, m . Lower values of m describe a lower roughness, and hence a more tightly directional reflection. Because the roughness parameter only affects the specular component of reflection, the reflection model for this result was entirely specular. The baseline in Figure 5.7 used $m = 0.2$ and $\omega = 0.9$. This plot therefore shows the difference that results from decreasing m , i.e. decreasing the surface roughness. In the right-most datapoints for $N_{\text{faces}} = 500$, there are only very small differences between the values for m , until m drops below approximately 0.1. As the number of facets is decreased, the lines for the different values of m begin to diverge. The maximum difference resulting from a decrease in m was approximately 3.0 and the maximum difference resulting from a decrease in the number of facets was approximately 2.7.

Next, the response of the light curve model to changes in the diffuse albedo, under entirely specular reflection conditions, was examined. The diffuse albedo affects the specular component of the reflection due to a simplification used in the BRDF calculations: the reflectance at normal incidence, F_0 , is taken to be equal to the diffuse albedo ω . Figure 5.8 shows the result of increasing the diffuse albedo under specular reflection conditions. The parameter m was set to 0.2 and the baseline value of ω was




FIGURE 5.7: The change in a synthetic light curve, quantified using RMSE, that resulted from decreasing the BRDF roughness parameter m .

set to 0.1. This plot shows that increasing the diffuse albedo has a larger effect on the light curve than reducing the number of facets. The maximum difference that resulted from increasing the diffuse albedo under specular reflection conditions was approximately 7.9 and the maximum difference caused by reducing the number of facets was approximately 1.2.




FIGURE 5.8: The change in a synthetic light curve, using entirely specular reflection and quantified using RMSE, that resulted from increasing the BRDF ω (diffuse albedo).

5.3.3 Sensitivity to Attitude State

The three plots in Figure 5.9 show the difference between light curves as attitude is changed under diffuse reflection conditions. Plot 5.9A shows the result from applying a rotation to the x-body-fixed-axis of Figure 5.2, plot 5.9B shows the result of applying this rotation to the y-body-fixed-axis and plot 5.9C shows the result of applying this rotation to the z-body-fixed-axis. In all three of these plots, for geometries with more than 150 facets, only very small differences were measured between light curves as the rotation angle was increased. This is an expected result, as the brightness of reflection from a sphere should be independent of attitude and these geometries has a sufficient number of facets to exhibit sphere-like behaviour. Below approximately 150 facets, the light curves of the different attitudes become measurably different. This is the point at which the geometry becomes sufficiently non-spherical and the attitudes can be distinguished. In all three plots, the lines are grouped into pairs. For example, the line corresponding to an attitude angle of 0.0π (blue line with circular markers) is always coincident with the line that corresponds to an attitude angle of 1.0π (purple line with '+' markers). This is a result of the symmetry of the geometry models, hence a rotation of π radians (180°) produces the same light curve as for no rotation. Figure 5.9 shows that attitude has a very small effect on the light curves relative to previous results. The maximum difference that resulted from increasing the attitude angle was approximately 0.05 whereas the maximum difference caused by reducing the number of facets was approximately 0.35. Additionally, the results for applying the rotation to any of the three axes are the same.

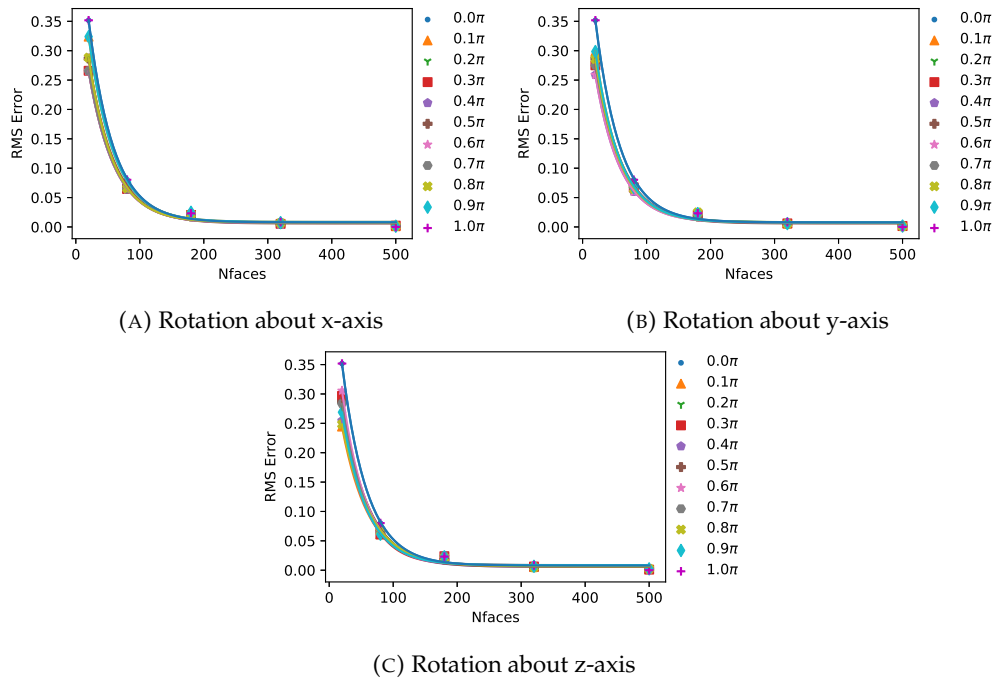


FIGURE 5.9: The change in a synthetic light curve resulting from changes in attitude state under entirely diffuse reflection.

The three plots in Figure 5.10 show the difference that results from changing attitude under specular reflection conditions. Other than reflection model, all inputs are the same as the previous case. In all three plots, little or no difference is measured when the number of facets is greater than approximately 150. This shows that even under specular reflection, these geometries are sufficiently spherical that light curve brightness is independent of attitude. However, for geometries with less than 150 facets, large differences are measured between the light curves with different attitudes. This is caused by strong specular reflections from the larger area facets of these geometries, which only occur at particular rotation angles. The maximum difference that resulted from increasing the attitude angle was approximately 12.0 and the maximum difference caused by reducing the number of facets was also approximately 12.0. This is the largest measured difference thus far. Furthermore, a larger difference is caused by a rotation about the z-body-fixed axis than rotations about the x or y-body-fixed axis.

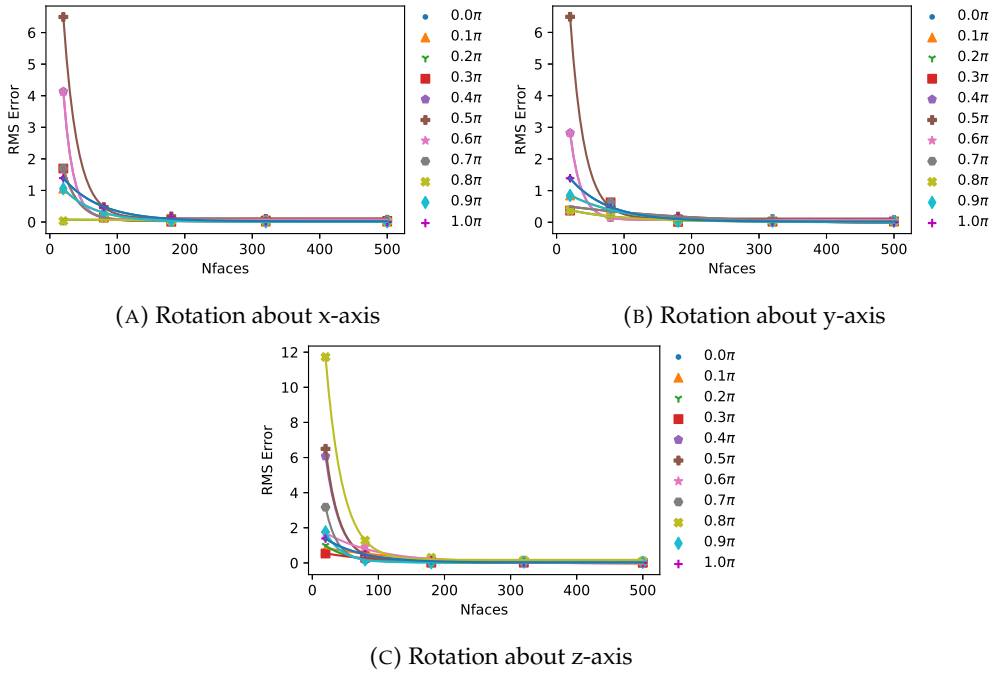


FIGURE 5.10: The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.

5.3.4 Sensitivity to Rotation Rate

Figure 5.11 shows the penultimate result. In this figure the change to a light curve that results from increasing rotation rate, using a diffuse reflection model, is shown. The baseline light curve was a stationary object and each line corresponds to increasing rates of rotation up to a maximum of 2 Rotations Per Minute (RPM). ω was set to 0.9. In all three plots, only very small differences were measured between the light curves with increasing rotation rates. This is because increasing the rotation rate only serves

to increase the frequency of the oscillations observed in the light curve. This does not cause an increase in measured difference because, regardless of rotation rate, the oscillations occur around the flat baseline light curve and so do not change the measured RMSE. This highlights a limitation of the RMSE technique for quantifying the difference between light curves. The maximum difference that resulted from increasing the rotation rate was approximately 0.04 and the maximum difference caused by reducing the number of facets was approximately 0.34.

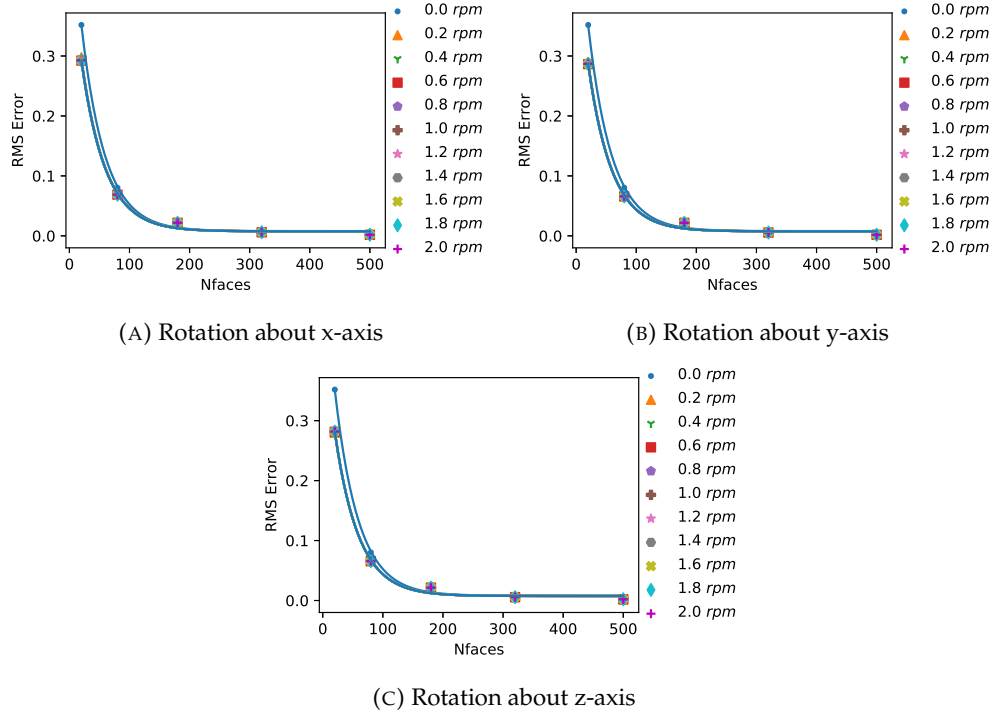


FIGURE 5.11: The change in a synthetic light curve resulting from changes in attitude under diffuse reflection.

Figure 5.12 shows the final result. This plot is the same as the preceding case, but now with a specular reflection model. Like the results of the preceding case, little or no difference is observed between the light curves with different rotation rates. Again this is due to increased oscillation frequency about a flat baseline not being measurable using RMSE. However, decreasing the number of facets results in large differences from the baseline light curve. This resulted in an increase in RMSE of approximately 2.5 in the y-rotation case, 2.75 in the x-rotation case and 4.4 in the z-rotation case.

5.3.5 Summary of RMSE Differences

Table 5.1 summarises the RMSE differences of each of the presented results. The $RMSE_{Parameter}$ column shows the maximum RMSE that resulted from changes to the input parameter. The $RMSE_{Facets}$ column shows the maximum RMSE that resulted from decreasing the number of facets in the geometry model. The parameter d is

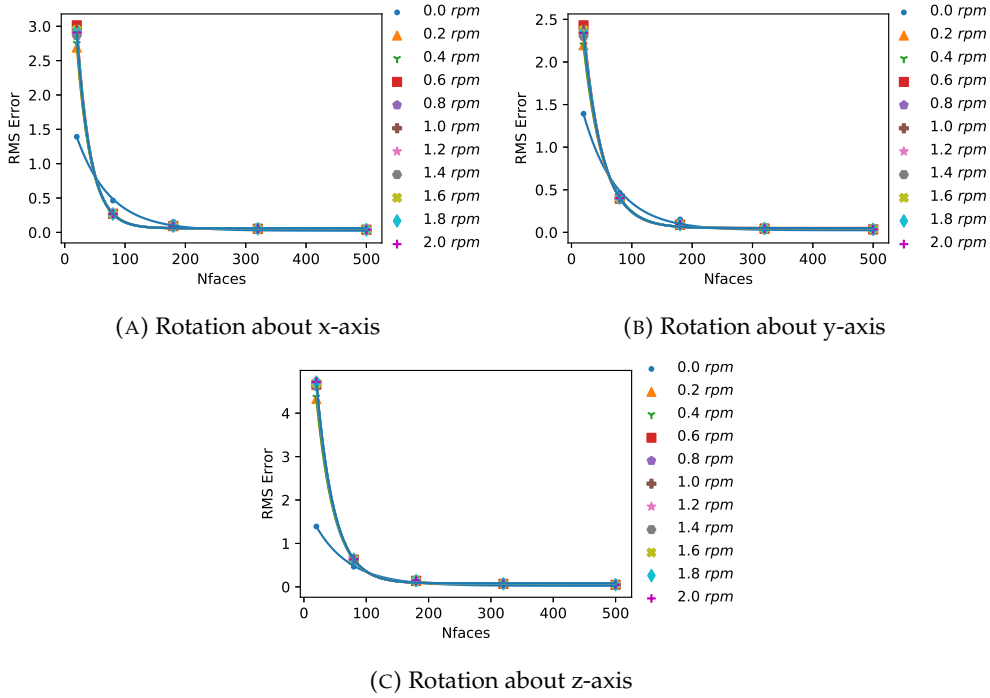


FIGURE 5.12: The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.

separated from other results as it shows the difference caused by changing from a diffuse to a specular reflection model.

Considering the column of diffuse results. Changes to θ and ω caused much larger differences than reducing the number of facets. For attitude and rotation rate, which produce the smallest differences, the opposite is true. These results show that the diffuse reflection component of the light curve model is most sensitive to the diffuse albedo, which causes large differences in light curve brightness. Under diffuse reflection conditions the model is comparatively insensitive to attitude for the considered geometry models, which is a result of the isotropic nature of diffuse reflections. For the specular to diffuse ratio d , the difference caused by changing this parameter was much larger than the difference caused by reducing the number of facets. Changes to d also cause larger differences than changes to all other parameters discussed thus far. This shows that the model is more sensitive to the handling of the specular reflection component than to any aspect of the diffuse reflection.

Considering the column of specular results, all RMSE differences are larger than the equivalent diffuse result. This shows that not only is the model highly sensitive to the specular component of reflection, but that as the reflection becomes more specular, the sensitivity to the other parameters is increased. For attitude and rotation rate, the differences caused by changes to the parameter are now larger than the differences resulting from changing the number of facets. These results show that, under specular

Parameter	Diffuse		Specular	
	$RMSE_{Parameter}$	$RMSE_{Facets}$	$RMSE_{Parameter}$	$RMSE_{Facets}$
θ	1.4	0.3	3.5	3.75
ω	2.5	0.4	7.9	1.2
m	N/A	N/A	3.0	2.7
Attitude	0.05	0.35	12.0	12.0
Attitude Motion	0.04	0.34	4.4	0.4
d	8.0	1.5	N/A	N/A

TABLE 5.1: The RMSE results for each of the plots shown in this section.

reflection conditions, the model is significantly more sensitive to attitude than under diffuse reflection conditions.

5.4 Conclusions

The results of this chapter have quantified the change in a synthetic light curve due to changes in the four key model inputs: illumination conditions, object geometry; attitude state; the BRDF parameters. The change in a light curve was quantified by measuring the RMSE between two light curves. The results of this chapter have highlighted certain circumstances where the performance of using RMSE to quantify differences was decreased. This was when comparing oscillating signals of varying frequency to a common baseline signal that maintained a constant value, which was demonstrated in Figures 5.11 and 5.12. In both results, light curves for increasing rates of rotation were compared to a constant baseline light curve generated with no rotation. The result of an increasing rotation rate is an oscillating light curve with increasing frequency. Hence, as the oscillation frequency increases, the RMSE measured from the baseline light curve remains unchanged. Therefore, although the primary objective of this chapter to quantify the sensitivity of the synthetic light curve model has been achieved, future work could reconsider these particular cases using a technique more suited to quantifying the differences between periodic signals, such as frequency analysis.

Under diffuse reflection conditions, changing the phase angle and the diffuse albedo produced the largest differences. Changes to attitude or rotation rate produced very small differences between light curves.

One of the largest single changes to the synthetic light curves was caused by introducing a specular component to the reflection. The specular reflection component also increased the sensitivity of the synthetic light curve model to the input parameters. Under specular reflection conditions, the differences introduced by changing the attitude or rotation rate are much higher and can be the cause of some of the largest measurable differences between two synthetic light curves.

In conclusion, when the reflection is primarily diffuse, the differences between light curves with different attitude states of the objects and reflection geometries considered in this chapter are very small. Hence the phase angle and diffuse albedo, which cause the largest differences, must be modelled accurately. As the reflection becomes more specular, the model becomes more sensitive to all input parameters. This is particularly the case for attitude state. Hence modelling of the specular reflections of an object can be highly beneficial when performing attitude determination from light curves, as these can be the cause of the largest differences between light curves of objects with different attitude states.

It should be noted that this chapter has not considered if a change in an input parameter would actually produce a detectable change in real light curve data. In cases where the change introduced into a light curve due to a change in an input parameter was small, then the presence of noise in real data may result in this change not being detected. A likely example where this would be true is in the sensitivity to attitude state under diffuse reflection case from Figure 5.9. The results of this chapter are valuable for understanding the fundamental sensitivities of the employed synthetic light curve model and the relative changes introduced into a light curve from the various model inputs, which was the primary objective. However, for the results to have value in understanding if the same changes can be detected in real light curve data, then the method would need to be adapted so to include simulated noise in the synthetic measurements. Further analysis could then be performed to assess if changes to the input parameters produce a sufficiently large change to the light curve that they could be unambiguously detected.

Chapter 6

Investigation of the Reflection Characteristics of Primitive Geometries

6.1 Introduction

The objective of this chapter was to identify illumination conditions under which the light curves of an object with different attitude states are most different from one another. This will help to devise observation strategies that maximise the likelihood that an object is observed in a way that its rotational motion is detectable and distinguishable from similar rotational motions. Additionally, it will help to identify and avoid illumination conditions that are not well suited for detecting or distinguishing rotational motion.

To achieve this objective, synthetic light curves were generated for a selection of object geometries and rotational motions and the illumination conditions under which they were most different were identified. It is not possible to test all combinations of object geometry, attitude state and illumination conditions. These three properties were therefore constrained to a smaller range of representative examples.

6.2 Methodology

Three object geometries were selected. These were: a cube, a 20-sided icosahedron and a 500-sided polyhedron. The cube geometry was selected as it is a good analogue of the typical 'box-like' shape of a satellite. The icosahedron was selected because it has facets at a range of different angles, which is representative of satellites that have less




FIGURE 6.1: The three object geometries considered in this paper. A 'b' subscript denotes a body-fixed axis.

uniform surfaces. Finally, the polyhedron geometry was used to be representative of a sphere. Although spherical objects are much rarer in Earth orbit, in theory their rotational motions should be the least detectable and distinguishable. The polyhedron was therefore included as a control.

The rotational motions considered in this chapter were constrained to rotations about an axis that remained fixed in inertial space. In order to measure differences between different rotations, three rotation vectors were selected. These were rotations about each of the three orthogonal body-fixed axes shown in Figure 6.1. In addition, a 'no-rotation case' was included as a control.

Finally, a range of illumination conditions were required. This was achieved by fixing the direction of the observation vector to be coincident with one of the three body-fixed orthogonal axes, then varying the direction of the illumination vector. Differences between light curves were quantified using Root Mean Square Error (RMSE). Light curves generated with various combinations of object geometry, rotational motion and illumination conditions were compared first using a purely diffuse reflection model, then again using a purely specular model of reflection.

The three object geometries are shown in Figure 6.1. These are: a cube, a 20-sided regular icosahedron, and a polyhedron where each facet of the icosahedron has been subdivided into 25 facets, such that the centroid of each facet is equidistant from the centre of the geometry. The cube has a side-length of 2 m and the icosahedron and polyhedron each have a radius of 1 m.

Figure 6.2 shows the reflection configuration used in the generation of each synthetic light curve with respect to the prescribed reference frame XYZ. The vector \mathbf{V} is the observation vector and always remains aligned with the X axis. The vector \mathbf{L} is the illumination vector, the direction of which is varied to analyse its influence on the generated light curve. The illumination vector is offset from the X axis with a direction described by the two angles ψ and ϕ , where ψ is the angle between the XZ plane and the plane containing the X axis and \mathbf{L} vector, and ϕ is the angle formed between the X axis and \mathbf{L} vector. These two angles therefore describe the offset of the illumination vector, \mathbf{L} , with respect to the observation vector, \mathbf{V} . ϕ is the angle between the illumination and observation vectors, and is therefore the phase angle.




FIGURE 6.2: A figure showing the reflection geometry and object geometry. The observation vector, \mathbf{V} , is aligned with the X -axis of the system and the position of the illumination vector, \mathbf{L} , is defined by the angles ψ and ϕ .

All synthetic light curves were generated with the object placed at 1000 km from the observer. The position of the object was fixed and was only allowed to rotate about its geometric centre. Each light curve was 240 s long with a 0.1 s time step.

6.3 Results and Discussion

The results for this work are RMSE measurements between a baseline light curve, which will be discussed later, and a second light curve generated with a different illumination vector direction. Because the direction of the illumination vector is described using the two angles ψ and ϕ , the results are presented in the form of polar plots. An example polar plot is provided as Figure 6.3, where ψ is measured clockwise along the periphery of the plot and ϕ is measured radially. These polar plots describe the direction of an illumination vector in a hemisphere centred on the direction of the observation vector. The angle ϕ is therefore equivalent to the phase angle. The colour of each pixel shows the RMSE between a light curve generated with an illumination vector direction given by the position on the plot and a baseline light curve. When the colour is white, such as the innermost ring of pixels in Figure 6.3, the RMSE is zero. These plots were generated for four different rotational motions, either no-rotation, or rotation about one of three body-fixed axes: x_b , y_b and z_b . Initially, the observation vector and the body-fixed x_b axis were coincident with the X axis of Figure 6.2.

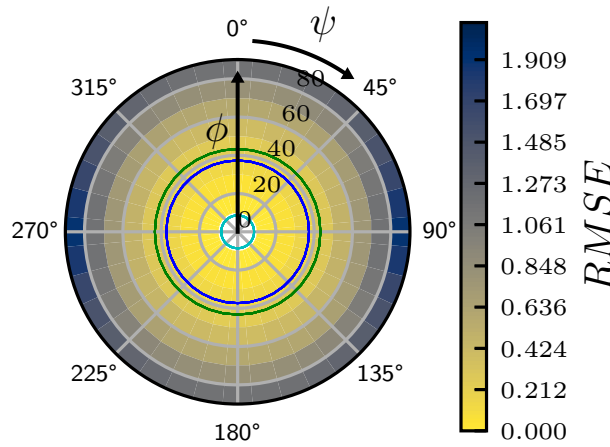


FIGURE 6.3: An example polar plot for the cube geometry with a diffuse reflection and a rotational motion about the Y body-fixed axis.

Therefore, a rotation about the x_b axis is the same as a rotation about the observation vector. This set of four plots were then generated for each of the three object geometries. This entire process was performed twice, firstly using a purely diffuse reflection model, then again using a purely specular reflection model.

The reference frame in Figure 6.2 is used to represent a satellite-centred orbital reference frame. In such a frame, and assuming a circular orbit, the Z axis is normal to the orbital plane, the Y axis is aligned with the velocity vector and the X axis points away from the centre of the Earth. Each plot also includes three concentric circles, which indicate the angular region within which the theoretical position of the Earth would obscure these illumination vectors. These three circles correspond to, in order of increasing value of ϕ , orbital altitudes of 35,786 km, 2000 km and 1000 km. In Figure 6.3 and all subsequent Figures, these three circles are coloured cyan, blue and green respectively. In this work the objects have been set at 1000 km, but these circles illustrate how some illumination vectors would become available as orbital altitude increases.

6.3.1 Cube with diffuse reflection

The first result is presented as Figure 6.4, for the cube geometry and a diffuse reflection model. The cube was selected to represent the typical ‘box-like’ shape of satellites. The baseline light curve, indicated at the centre of the no-rotation plot using a green ‘x’, corresponds to a non-rotating object with $\phi = \psi = 0$. Within this figure, are the results for the four different rotational motions that were considered. For the three rotating cases, the direction of the rotation axis is indicated using a red ‘+’. The pixelation of the plots is caused by using 10° bins for ϕ and ψ . Because the geometry being considered is a cube, an illumination vector offset of 90° ($\phi = 90^\circ$) results in no

light being reflected in the direction of the observation vector. For this reason, the outermost ring of pixels in each of these polar plots is excluded.

In these plots the RMSE measured against the baseline case increases as phase angle ϕ is increased. This is caused by an overall reduction in light curve brightness that results from an increased phase angle. Considering only the no-rotation and x-rotation cases, the results are identical and independent of ψ . This is because only a single face of the cube contributes to the reflection, which remains perpendicular to the observation vector throughout the light curve timespan. In the y-rotation and z-rotation cases, the largest RMSEs were measured when the illumination vector was closely aligned with the rotation vector (red '+'s). In all four polar plots, zero RMSE was measured in the innermost ring of pixels that corresponded to $\phi = 0$. This indicates that under such illumination conditions, rotating and non-rotating objects were indistinguishable from each other.

Figure 6.5 shows the result of performing a pixel-to-pixel subtraction between the polar plots of each of the rotating cases. To correct for negative values, the absolute value of this subtraction was used. These plots remove the effect of brightness reduction due to increased phase angle, allowing the differences between the individual rotational motions to be seen more easily. These plots do not measure the RMSE between the two rotational motions, but rather the difference in the RMSE measured from the baseline for each state. Therefore, zero-valued pixels in this plot do not necessarily show that the light curves of the two different rotational motions are identical, just that the two light curves are different from the baseline by an exactly equal amount. The rotation vectors of the two rotational motions being compared are again indicated using red '+'s. These plots show that the direction of the illumination vector can have a large effect on the difference in RMSE measured between the two light curves and the baseline. Depending upon which two rotational motions are being compared, the largest RMSE differences are recorded either when the illumination vector is aligned with the rotation vector (z-y case), or when the ψ angle of the illumination vector is offset from the ψ angle of the rotation vector by 90° (y-x and z-x cases). In all cases, the largest differences are always measured when the illumination vector is offset from the observation vector by 90° .

These results for the cube geometry highlight the illumination conditions that are most ideal for distinguishing rotational motions. In all cases, a large offset between the observation and illumination vectors is desired. The ideal direction of the illumination with respect to the rotation vector is dependent upon the two rotational motions to be distinguished. In Figure 6.5, the band of white pixels that are horizontal through the origin in the y-x plot (middle) and vertical in the z-x plot (right) shows that these attitude motions are indistinguishable if the observation vector, the illumination vector and the rotation vector are all in the same plane.

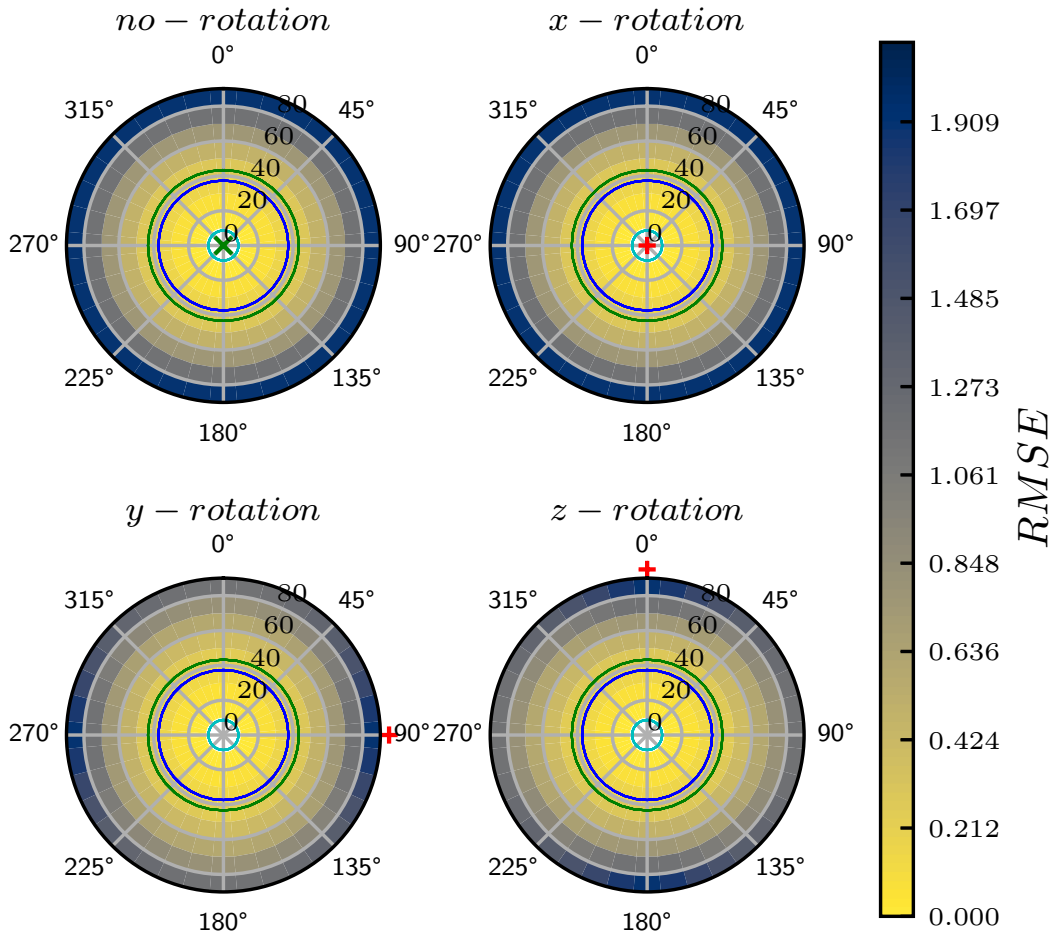


FIGURE 6.4: Results for a diffuse reflection model and a cubic object geometry. RMSE is measured from a no-rotation no illumination offset baseline.

Note that the y-x plot (middle) is exactly the same as the z-x plot (right), but rotated through 90° . This is the case for all of the subtraction results presented in this chapter. For this reason, the z-x polar plot is excluded from all subsequent figures.

6.3.2 Cube with specular reflection

The next result is for the cube geometry with a specular reflection model. All model inputs and parameters are identical to the diffuse results, with the only exception being the reflection model. Considering Figure 6.6, the RMSE measurements are larger than for the diffuse results, but the positions of the largest RMSE values with respect to the illumination and rotation vectors are the same. The most notable difference is that the innermost ring of pixels for $\phi = 0$ in the y-rotation and z-rotation cases no longer correspond with zero RMSE. This shows that with specular reflection these rotational motions can be distinguished, which was impossible under diffuse reflection conditions.




FIGURE 6.5: The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.4. The two rotation cases in question are labelled above each polar plot. The reflection model is diffuse and the object geometry is a cube.

Figure 6.7 shows the result of the pixel-to-pixel subtraction. The $z - y$ case is the same as for the diffuse reflection model. However, the $y - x$ and $z - x$ cases show an inverted behaviour. The largest RMSE now occurs when the illumination vector is aligned with the rotation vector, rather than when they are offset by 90°.

With specular reflection, the most ideal conditions for distinguishing rotational motions are again with maximum illumination vector offset. The ideal direction for the illumination vector is to be coincident with the rotation axis.

6.3.3 Icosahedron with diffuse reflection

The next two figures show the results for the icosahedron geometry with a diffuse reflection model. The icosahedron geometry was selected due to it having a more varied range of facet angles. In Figure 6.8 the y -rotation and z -rotation cases have been excluded because there was no discernible difference between those two plots and the two plots that have been provided. Figure 6.8 shows that the largest RMSEs were caused by increasing the phase angle, ϕ . All of the results appear to be independent of ψ angle. The four plots for the four rotational states cannot be distinguished from this figure.

In the pixel-to-pixel subtraction plots of Figure 6.9, the plots for the rotational motions are now distinguishable. In Figure 6.9, the change in pixel values as ψ is varied shows that ψ does have an effect on the difference in RMSE. However, the magnitude of this difference is very small, which is shown by the scale bar having a maximum value of 0.05.

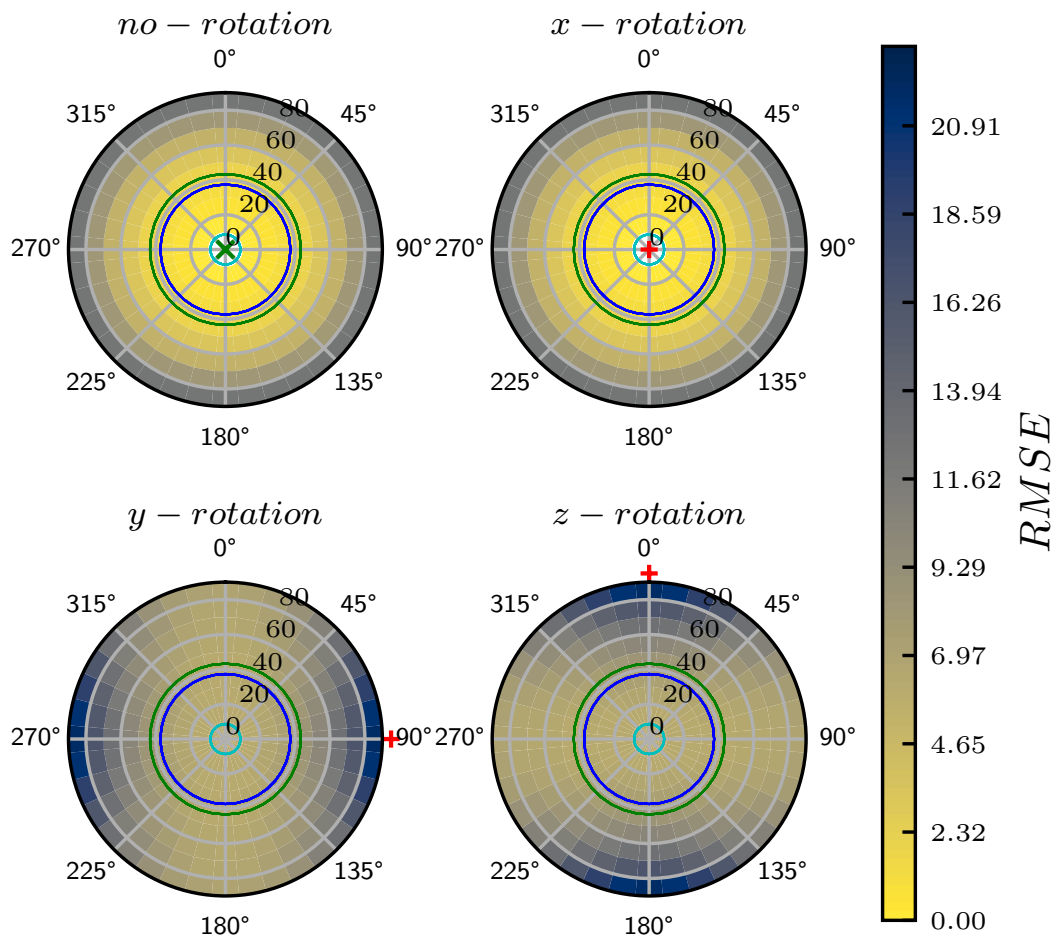


FIGURE 6.6: Results for a specular reflection model and a cubic object geometry. RMSE is measured from a no-rotation no illumination offset baseline.




FIGURE 6.7: The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.6. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a cube.

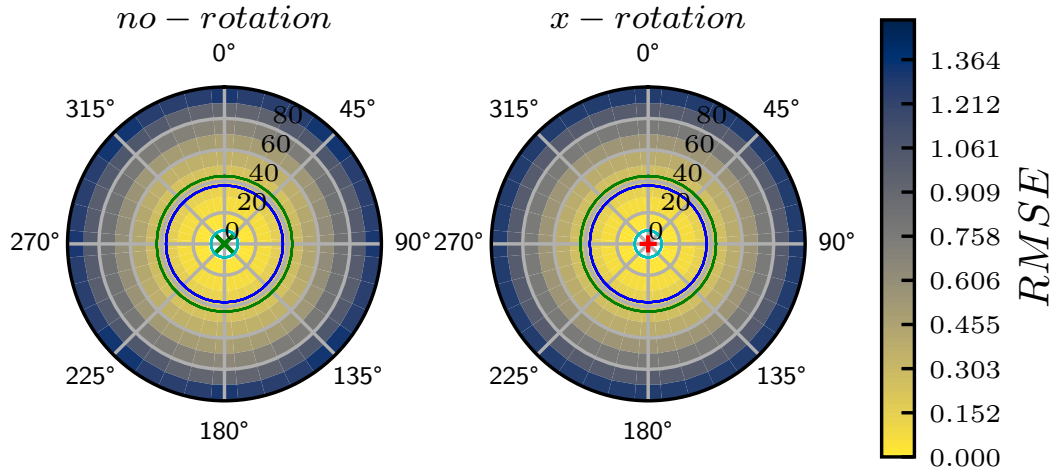


FIGURE 6.8: Results for a diffuse reflection model and a icosahedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.

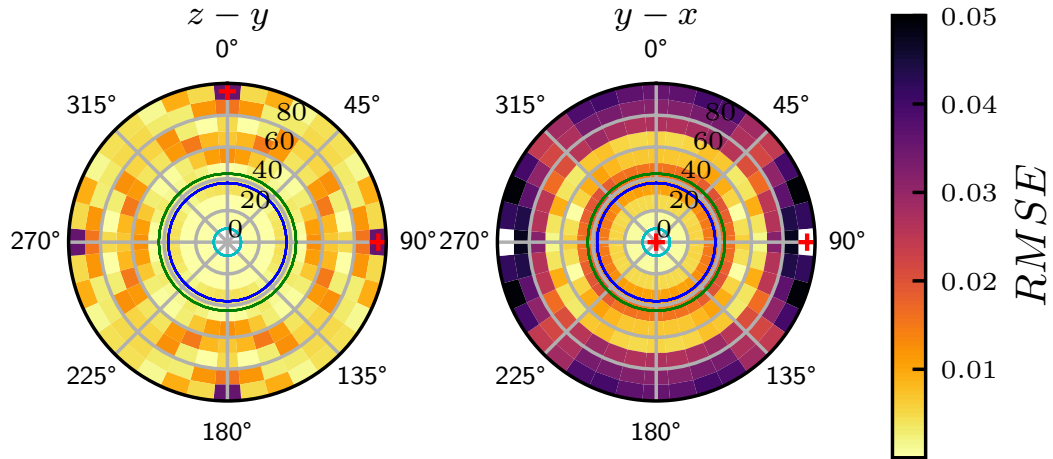


FIGURE 6.9: The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.8. The two rotation cases in question are labelled above each polar plot. The reflection model is diffuse and the object geometry is a icosahedron.

The much smaller RMSE measurements between the light curves of different rotational motions shows that the icosahedron with diffuse reflection has adopted a more 'sphere-like' behaviour. However as was the case for the previous results, the ideal conditions to distinguish these rotational motions is with maximum phase angle and an illumination vector closely aligned with the rotation vector.

6.3.4 Icosahedron with specular reflection

The next two figures show the results for the icosahedron with a specular reflection model.

In Figure 6.10, increasing the phase angle is no longer the cause of the largest RMSE measurements. Instead, it is particular illumination vector directions that lead to these largest values, which is due to the high sensitivity of specular reflections to specific illumination conditions. In the no-rotation case, the large RMSE measurements at approximately $\phi = 60^\circ$ and $\psi = 0^\circ$ or 180° are caused by very bright specular reflections from a single facet of the icosahedron. In the rotating cases, the largest RMSE measurements were made when the illumination vector was placed such that bright specular glints were introduced into the light curve, that were not present in the baseline.

Figure 6.11 shows the pixel-to-pixel subtractions. The largest differences in the RMSE measurements shown in this plot result from illumination vector directions where specular glints are introduced into the light curves of one rotation state, but not the other.

These results exhibit a departure from the behaviour observed in previous figures. With specular reflection, the icosahedron geometry no longer demonstrates ‘sphere-like’ behaviour, as the measurable differences between rotational motions have greatly increased.

6.3.5 Polyhedron with diffuse reflection

The next two results are for the polyhedron, which was used to represent a sphere. Figure 6.12, once again, only shows the no-rotation and x-rotation cases because the y-rotation and z-rotation cases are visually identical. In this figure, any changes in RMSE measurements due to the different rotation motions were negligible compared to brightness reduction due to increasing phase angle. This is the same result as the icosahedron geometry with a diffuse reflection model of Figure 6.8. This highlights the ‘sphere-like’ behaviour of icosahedron geometry under diffuse reflection conditions.

The subtraction plots are also very similar to those of the icosahedron with diffuse reflection, but the magnitude of the difference decreased by two orders of magnitude. The fact that differences are still measured between the different rotational motions is likely to be caused by the discretisation used to generate the object geometry. This again highlights the ‘sphere-like’ behaviour of the icosahedron with diffuse reflection.

6.3.6 Polyhedron with specular reflection

The next two figures present the results for the polyhedron with a specular reflection model. The polar plots of Figure 6.14 show all the same phenomena as the previous results. The only noticeable difference is a very slight ϕ angle dependence in the no-rotation case, which indicates that the polyhedron geometry is not perfectly

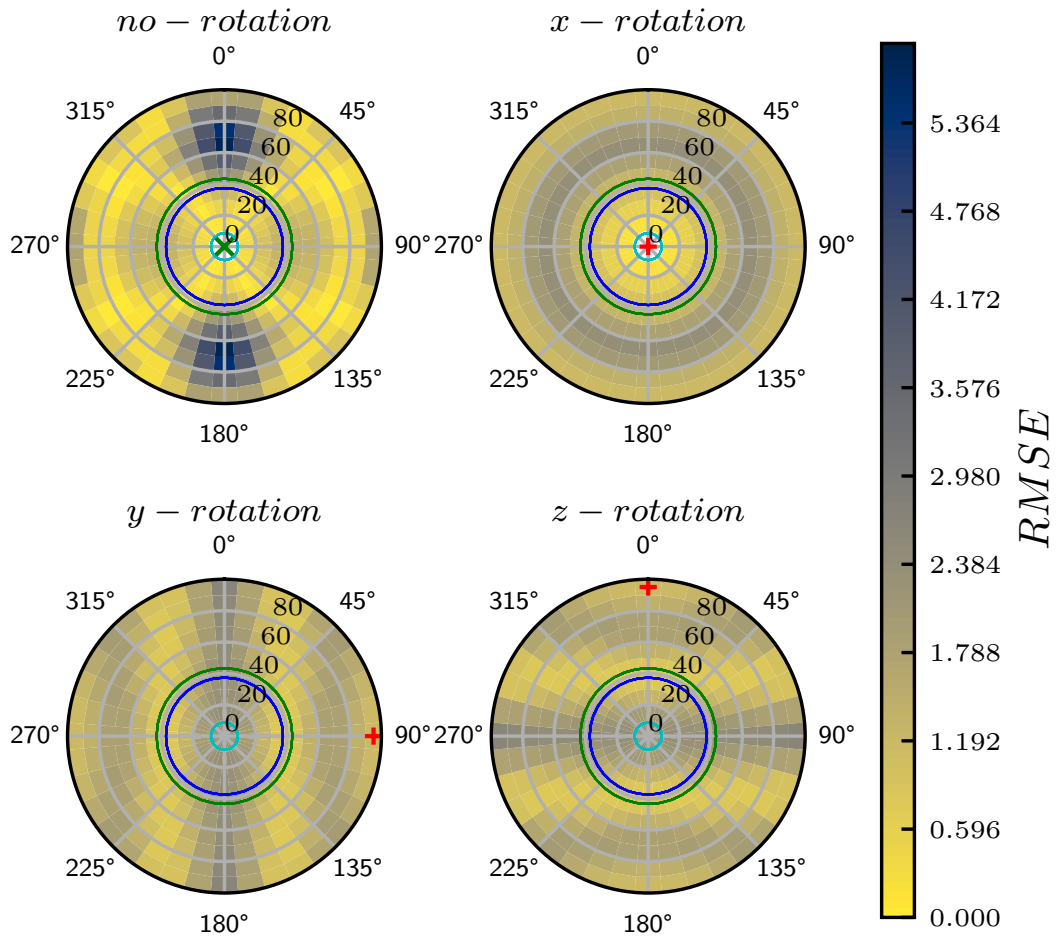


FIGURE 6.10: Results for a specular reflection model and a icosahedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.

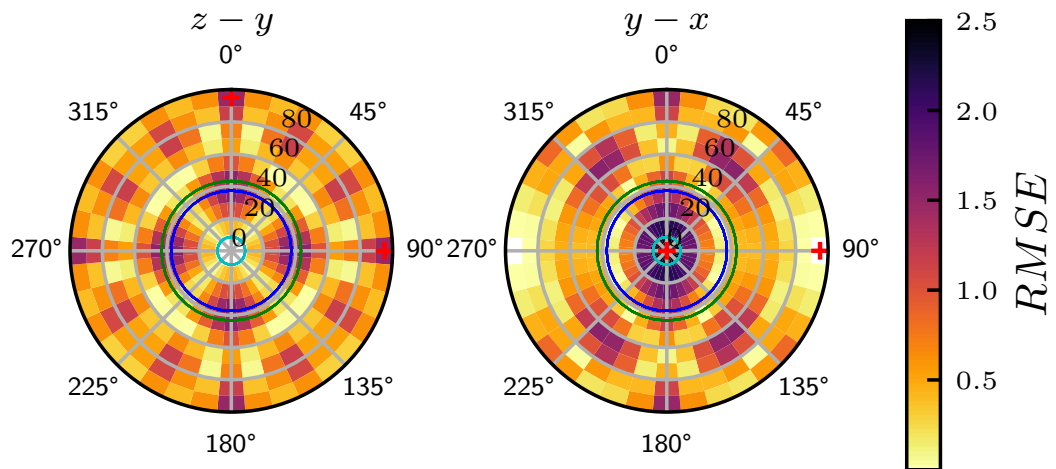


FIGURE 6.11: The result of performing a pixel-to-pixel subtraction between two of the plots from Figure 6.10. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a icosahedron.

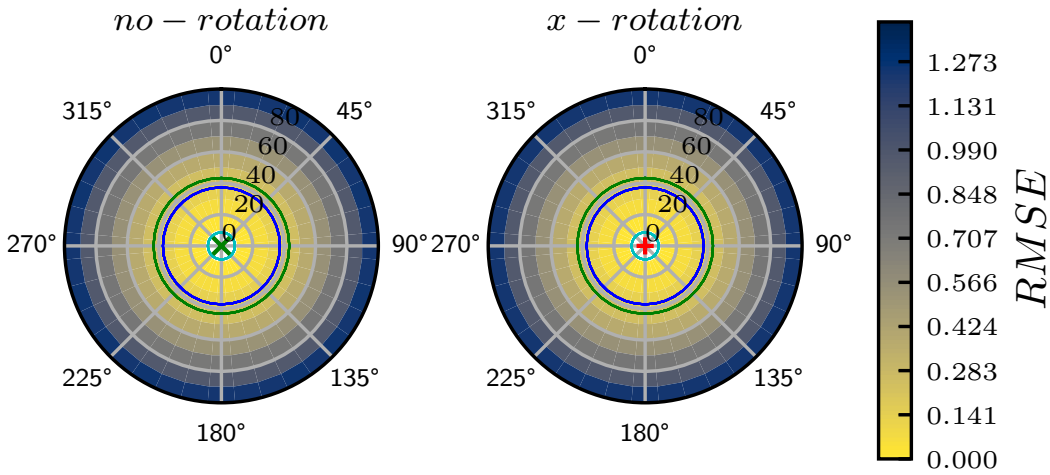


FIGURE 6.12: Results for a diffuse reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.

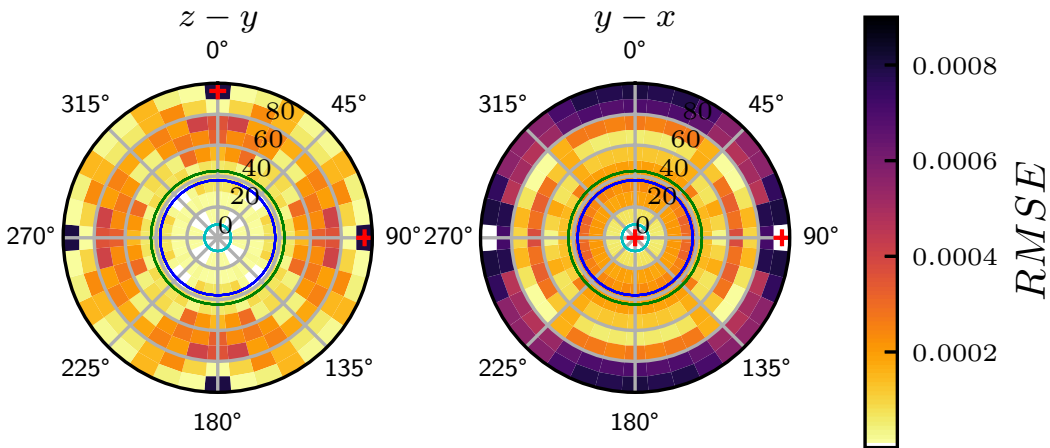


FIGURE 6.13: The result of performing a subtraction between two of the plots from Figure 6.12. The two rotation cases in question are labelled above each polar plot. The reflection model is diffuse and the object geometry is a 5-frequency subdivided polyhedron.

isotropic. The fact that this result is so similar to the diffuse results shows that the polyhedron geometry is sufficiently spherical that strong specular reflections from single facets no longer have an effect on the light curves. This is in contrast to the icosahedron, which exhibited a breakdown in ‘sphere-like’ behaviour after switching to specular reflection.

The final subtraction plot is presented in Figure 6.15. The differences measured, which are larger than those for the diffuse case, are likely due to the object geometry discretisation and exacerbated by the specular reflection model.

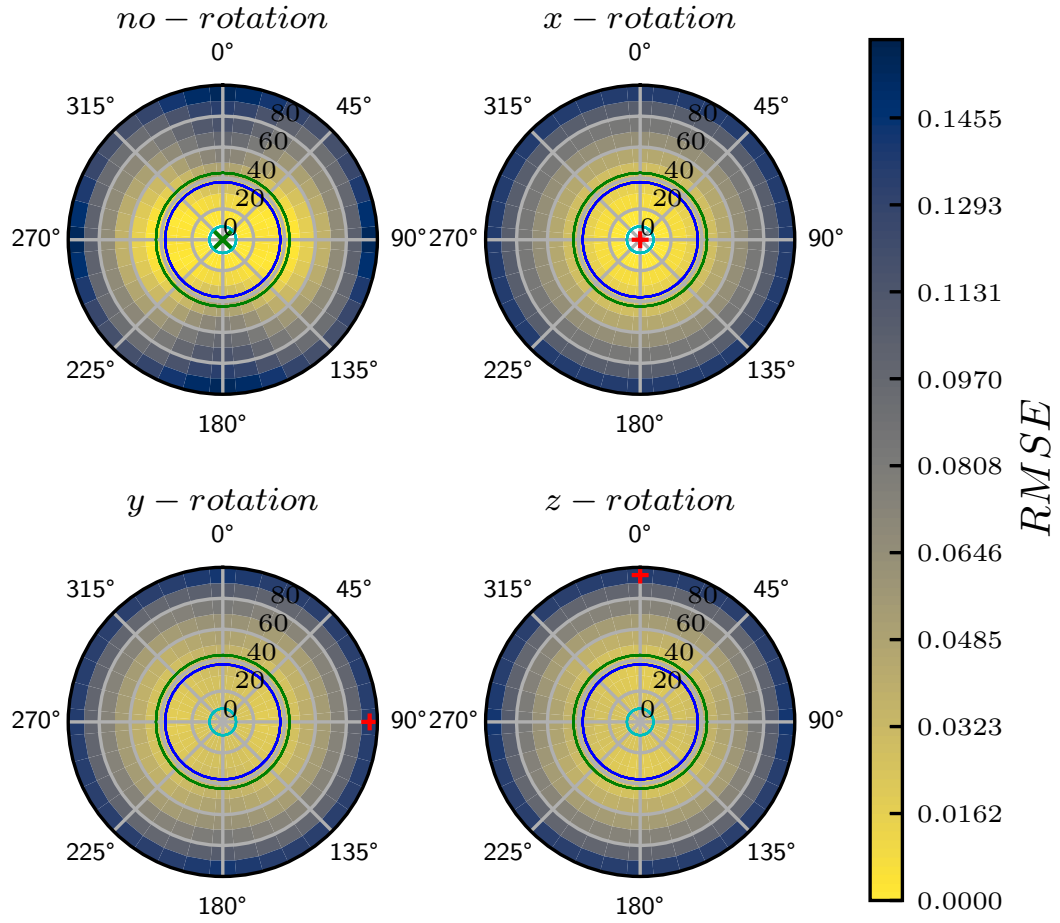


FIGURE 6.14: Results for a specular reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.

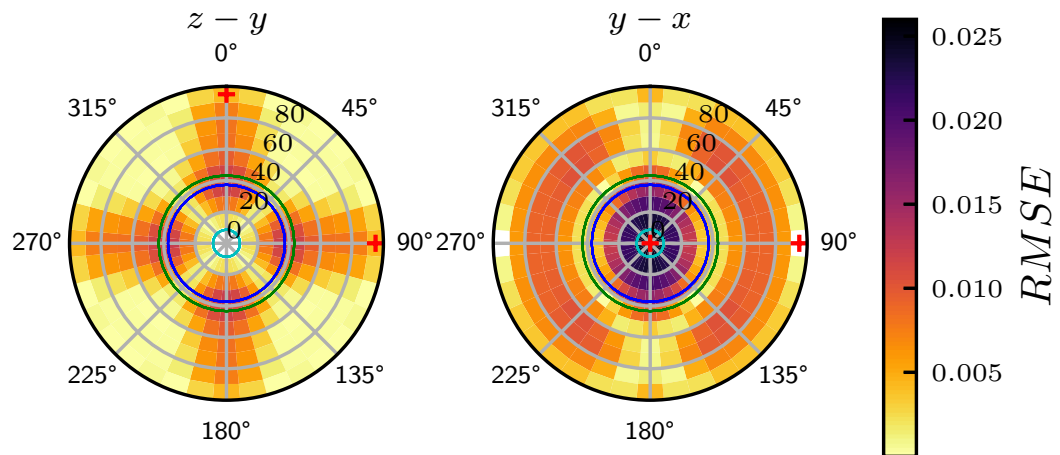


FIGURE 6.15: The result of performing a subtraction between two of the plots from Figure 6.14. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a 5-frequency subdivided polyhedron

6.3.7 Synthetic Light Curve Demonstration

To illustrate how the results presented in this paper may affect real-world light curve characterisation, a realistic demonstration light curve scenario was constructed. Like all previous synthetic light curves, the light curve was 240 s long and the object was placed at a fixed distance from the observer. This scenario used the cube geometry and a diffuse reflection model. Figure 6.16 shows the light curve illumination geometry at $t = 0$ and at $t = 240$. In this figure, all vectors lie in the plane of the x_b and y_b body-fixed rotation vectors, with the z_b vector pointing out of the page. This reference frame represents a satellite-centred orbital reference system for a circular orbit, where the X-axis points to the centre of the Earth, the Y-axis is coincident with the velocity vector and the Z-axis (out of the page) is the normal to the orbital plane. Throughout the light curve, the direction of the observation vector moved from the initial position shown on the left, to the final position shown on the right. This change in observation vector direction is equivalent to the observation vector motion that results from a satellite passing over a ground-based observation site. As a result of this motion, the phase angle increases from 40° at the start of the light curve to 80° at the end.

Figure 6.17 shows the light curves of the four considered rotational motions for this scenario. Two results are demonstrated in this plot. Firstly, only the z-rotation case shows a difference from the no-rotating case. The no-rotation, x-rotation and y-rotation cases are indistinguishable under these illumination conditions. These three light curves are superimposed in Figure 6.17. The second result is that the measurable difference between the z-rotation case and no-rotation case increased as phase angle increased. These light curves demonstrate the results of Figure 6.5, which showed that only a z-rotation is detectable if the observation vector, illumination vector and rotation vector are in the same plane.

A second light curve scenario was generated using a different object attitude state such that none of the three body-fixed geometry axes were in the same plane as the observation and illumination vectors. This illumination geometry is shown in Figure 6.18. The only change from the previous scenario is the object's attitude state. Considering the light curves in Figure 6.19, the four rotational motions can now be distinguished, as the rotation axes are no-longer in the same plane as the observation and illumination vectors.

6.4 Conclusions

The results presented in this work have demonstrated that the direction of the illumination vector with respect to the observation and rotation vectors can determine




FIGURE 6.16: The initial and final illumination conditions used in the first light curve scenario. \mathbf{L} is the illumination vector and \mathbf{V} is the observation vector. The curved dotted line shows the outline of the Earth.




FIGURE 6.17: Light curves for the four rotation states considered in this paper, using the illumination geometry presented in Figure 6.16. All light curves, with the exception of the 'z-rot' case are superimposed.




FIGURE 6.18: The initial and final illumination conditions used in the second light curve scenario. \mathbf{L} is the illumination vector and \mathbf{V} is the observation vector. The curved dotted line shows the outline of the Earth.




FIGURE 6.19: Light curves for the four rotation states considered in this paper, using the illumination geometry presented in Figure 6.18.

whether or not differences can be measured between light curves with different rotational motions.

In the results that used a diffuse reflection model, the most important parameter was phase angle and large phase angles were best able to distinguish between rotational motions. At low phase angles, the light curves of different rotational motions were either very similar or identical. Fortunately, if using ground-based observations, very low phase angles are made impossible by the position of the Earth. However, this would be an important consideration for space-based observations as, typically, low phase angles are desired to maximise signal. When phase angle was sufficiently high, the direction of the illumination vector with respect to the axis of rotation became the most important parameter determining the measurable difference between rotational motions. No single illumination vector direction with respect to rotation vector was ideal for distinguishing all of the different rotational motions for each of the four object geometries. However, in all cases the ideal direction was either closely aligned with the rotation axis, or offset from it by 90° .

In the results using a specular reflection model, the direction of the illumination vector with respect to the object geometry had a much larger affect on the differences measured between rotational motions. This is particularly true for object geometries that have facets at a range of different angles such as the icosahedron. However, this result breaks down as the object becomes more spherical causing the specular results to become more similar to those for diffuse reflection.

Regardless of object geometry, attitude state and reflection model, the largest differences between rotations were measured with an illumination vector either closely aligned with the rotation axis, or offset from the rotation axis by 90° .

The results in this chapter have identified illumination conditions that produce the largest measurable differences between different rotational motions. Certain

properties can be generalised across the considered geometries and rotational motions. For example, a large phase angle typically increased the measurable differences. However, there was no single set of illumination conditions that provided large measurable differences between all the considered rotation motions for all geometries. Rather, the ideal conditions were specific to the object's shape, the reflection characteristics of its surfaces and the rotational motion to be detected. The results also showed particular illumination conditions that resulted in only very small measurable differences between rotation motions. This agrees with the results of Hinks et. al., which found that the detectability of attitude state in a single measurement is dependent on the particular attitude state and the sensitivity of the light curve to changes in that state (Hinks et al., 2013).

This work indicates that if in a favourable illumination geometry, characterising the rotational motion of an object from its brightness data is possible from a single light curve. However, it is best performed using multiple observations under a range of different illumination geometries. This improves the likelihood of receiving favourable geometries for detecting the object's rotational motion and also decreases the likelihood of incorrect characterisation due to components of the rotational motion being undetectable under the considered illumination geometries. Hence, these observations should be constructed in such a way that maximises the difference in the illumination geometry. One way of achieving this would be by using observations from multiple observatories with a diverse range of geographical locations. This would maximise the measurable differences between rotation motions and minimise the possibility of incorrect characterisation.

Chapter 7

Angular Velocity Vector Determination of Rocket Bodies

7.1 Introduction

In this chapter a technique is presented for determining the angular velocity vector for a tumbling rocket body. This has been attempted before by Williams ([Williams, 1967](#)), where the difference between minimum and maximum light curve brightness was used to determine the vector. In this work, the angular velocity vector was determined by generating a match between a real target light curve and a synthetic light curve for which the angular velocity vector was known.

In order to be confident that the simulated angular velocity vector is the same as the real one, it must be the only angular velocity vector for which a matching synthetic light curve can be generated. If multiple synthetic light curves with different angular velocity vectors can be matched to the real data, then further analysis would be required to identify the angular velocity vector solution that is the best match to the real data.

Because the angular velocity vectors of real rocket bodies are typically unknown, there is no truth data against which to compare the results. Therefore, the effectiveness of this technique was first evaluated on four synthetic target light curves, for which all inputs were known. It was therefore necessary to establish that the rocket body light curves generated by the synthetic light curve model are good representations of real light curves of rocket bodies, which was done by comparing them to real light curves in the MMT database.

This technique was then applied to a real target light curve. The performance of the technique was assessed by comparing the angular velocity vector determination results between the real and synthetic cases. In this chapter, the term ‘target light

curve' is used to refer to the light curve from which the angular velocity vector is to be determined. The target light curve may be either real or synthetic.

7.2 Methodology

7.2.1 Real Target Light Curve

Rocket bodies at the most basic level have a common shape: cylindrical with a rocket nozzle at one end. The observed brightness of a cylindrical rocket body is therefore due to reflected sunlight from some combination of three zones: the end with the rocket nozzle; the end without the rocket nozzle; and the curved sides of the main body that houses the propellant tanks. Assuming that a rocket body is in a tumbling motion, and is not in an orientation where an equal amount of light is reflected from two of the zones simultaneously, the light curve would be expected to contain a repeating sequence of reflections at three distinct brightness levels. This would be caused by each of the three zones becoming the dominant source of reflected light as the object tumbles. In those cases where simultaneous reflections from multiple zones is possible, these three brightness levels may merge and become less distinguishable.

To test if the light curves of tumbling rocket bodies exhibit these three hypothesised brightness levels, Figure 7.1 presents six examples of real light curves collected by the MMT on a range of tumbling LEO rocket bodies. Where possible, the repeating sequence of three brightness levels is annotated using the letters 'A', 'B' and 'C'. Using Figure 7.1e as an example, brightness level 'A' presents at a standard magnitude of approximately 4.0, brightness level 'B' presents at approximately 9.0 and brightness level 'C' presents at approximately 6.0. Hence the reflection sequence is [A, B, C, B, A, B, C, B, A, ...]. This sequence has also been labelled for cases 7.1a, 7.1b and 7.1c.

Cases 7.1d and 7.1f do not exhibit three clearly identifiable brightness levels. This may be due to reflections from multiple zones simultaneously as described above or, alternatively, reflections from two of the zones are at very similar brightnesses.

For this chapter, the Long March (abbreviated in Chinese to 'CZ') 2D rocket body light curve in Figure 7.1a was selected as the real target light curve. The properties of this light curve are given in Table 7.1. This light curve was selected due to the low noise and the clear repeating sequence of reflections from the three zones of the geometry, making it a high quality example of a typical rocket body light curve. A picture of a CZ - 2D rocket is shown in Figure 7.2, and the section of the rocket that remains on-orbit has been outlined in red.

For reference, MMT magnitude measurements have a typical error on the order of 0.1 - 0.2 magnitude in the visual band for an object size of approximately 5 square meters at

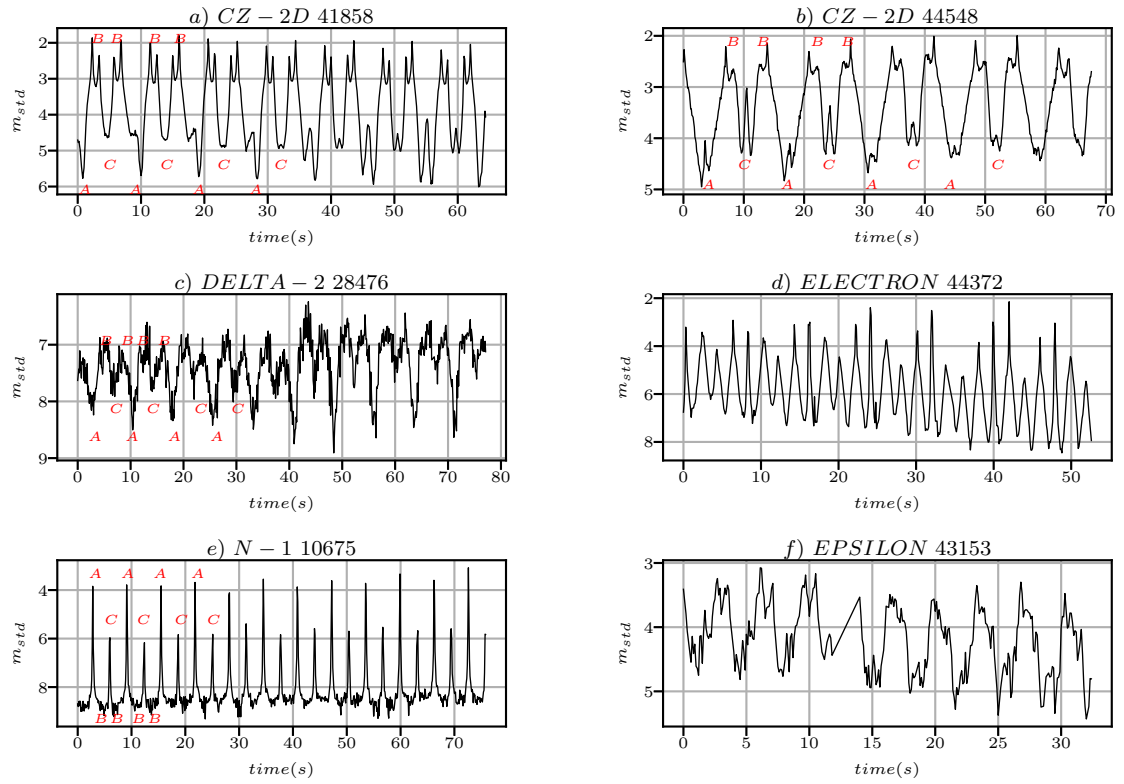


FIGURE 7.1: Six examples of real rocket body light curves collected using MMT. Red 'A', 'B' and 'C' letters are used to highlight repeating sequences of reflections at three distinct brightness levels where applicable.

TABLE 7.1: The properties of the real target light curve.

Property	Value
Name	CZ-2D R/B
NORAD ID	41858
Start Date	11/02/17
Start Time	03:02:59
Duration	64.4 s

ranges of approximately 550 - 750 km (Mallama, 2021). The MMT database provides a rotational period of 9.2 s for this light curve, calculated using period folding.

7.2.2 Synthetic Target Light Curves

Because angular velocity vector truth data is not available, the capability to determine the angular velocity vector from light curve matching was first tested using synthetic light curves for which the angular velocity vector was known. However, the ability of the synthetic light curve model to generate realistic rocket body light curves must first



FIGURE 7.2: A CZ-2D launch vehicle with the upper stage that remains on-orbit indicated by a red box.

be considered. This was performed by generating synthetic light curves and comparing them to the real light curves from Figure 7.1.

A range of example synthetic light curves were generated using position data defined by the real light curve data from Table 7.1. The following additional inputs were required: an object geometry with associated BRDF parameters for each facet; an angular velocity vector in the body-fixed frame; and the initial offset between the body-fixed and orbital reference frames.

The objective in the selection of the geometry was to capture the primary cylindrical feature of a rocket body, whilst minimising the level of complexity. Additional complexity can be added if it is found that this less complex representation is insufficient. Figure 7.3 shows this geometry, which is a 3.3 m by 9.5 m cylinder, capped with a cone that extends by a further 1 m from each end. The cones at each end were added for two reasons. Firstly, so that the ends of the geometry were not flat

surfaces at 90° from the sides of the cylinder, which could cause large disparities between the shapes of real and synthetic light curves. Secondly, because a real rocket body will have features that extend from the ends, such as a rocket nozzle, and these will introduce additional reflective surfaces. Cones were selected over more complex shapes to reduce complexity and to remove the need to calculate shadowing from concave features.

The geometry was broken into three zones, with each being given individual sets of BRDF parameters. These three zones are indicated by the three colours in Figure 7.3 and correspond to the top, bottom and side. The geometry was also assumed to be of constant density, the centre of mass is therefore the same as the geometric centre. This was done, again, to minimise complexity.

The angular velocity vector in the body-fixed frame was set to be coincident with the geometry x-axis, a ‘flat spin’ attitude state. This is a sensible assumption for the rotation state of a tumbling rocket body as it presents a rotation about the object’s principle axis of inertia. Although this will not be object’s initial rotation state, external torques such as gravity gradient and eddy current, in addition to internal dissipative forces such as fuel sloshing, will cause the object to adopt a stable minimum energy state - a rotation about its minimum axis of inertia (Praly et al., 2012). The rotation angular rate was derived from the 9.2 s rotation period provided by the MMT database for this light curve.

Figure 7.4 shows the orbital reference frame. The angular velocity vector in this frame is given as an offset from the Earth zenith vector, using the two angles ϕ and ψ . As shown in Figure 7.4, ψ is the angle between the projection of the angular velocity vector in the Y-Z plane and the Z-axis, and ϕ is the angle between the angular velocity vector and the X-axis. The angular velocity vector is assumed to remain constant in the orbital frame. As a result of the angular velocity vector being aligned with the X-axis of the body-fixed frame, the position of the angular velocity vector in the orbital frame is the same as the offset between the orbital and body-fixed frames.

Figure 7.5 presents four synthetic target light curves, generated using the randomly generated values for the BRDF parameters and angular velocity vector offset angles presented in Table 7.2.

7.2.3 Determining Best Fitting Light Curve Inputs

To find the synthetic light curve that best fits the real target light curve the three sets of BRDF parameters and two angular velocity vector offset angles must be determined. This was achieved by first generating an initial fit, where the two angular velocity vector offset angles were adjusted until the phase of the synthetic and real light curves approximately matched. The next step was to determine the three sets of best fitting




FIGURE 7.3: Object geometry used to represent a rocket body. The red, blue and green facets indicate the three zones of the geometry, which are given different reflection characteristics.




FIGURE 7.4: The orbital reference frame used to define the angular velocity vector. The angles ϕ and ψ describe the direction of the angular velocity vector with respect to the Earth zenith direction.

BRDF parameters for each zone of the geometry. To do this, the synthetic light curve was filtered to only show those reflections from a single zone. All combinations of BRDF parameters were then tested for this zone using the ranges shown in Table 7.3 in steps of 0.1 (i.e. $[d_{top} = 0.0, m_{top} = 0.1, \omega_{top} = 0.1]$, followed by $[d_{top} = 0.0, m_{top} = 0.1, \omega_{top} = 0.2]$ etc.). The best fitting BRDF parameters were those that minimised the RMSE between the synthetic light curve and the corresponding sections of the real light curves. This process was repeated for each of the three geometry zones to determine the three sets of best-fitting BRDF parameters.

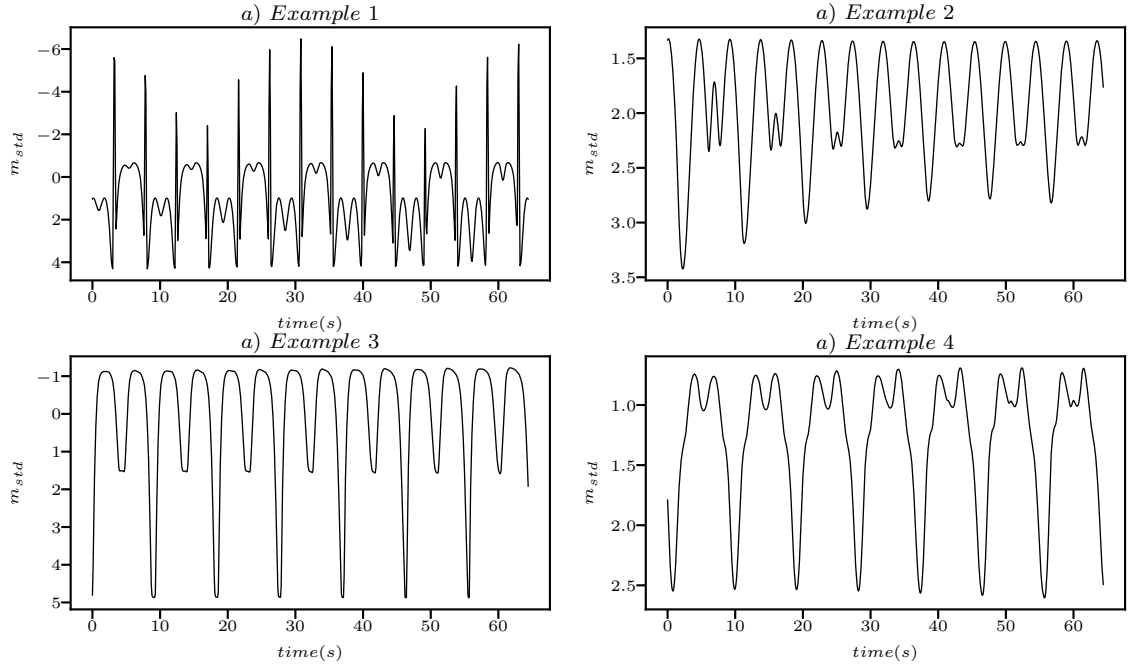


FIGURE 7.5: Four synthetic target light curves.

TABLE 7.2: Table of randomly generated inputs used to generate the synthetic target light curves of Figure 7.5.

Property	Generation Range	Example 1	Example 2	Example 3	Example 4
ψ	0 - 359	267	114	160	324
ϕ	0 - 90	65	15	89	9
ω_{top}	0.00 - 1.00	0.79	0.66	0.51	0.42
m_{top}	0.00 - 1.00	0.34	0.20	0.52	0.61
d_{top}	0.00 - 1.00	0.91	0.11	0.98	0.49
ω_{bottom}	0.00 - 1.00	0.77	0.08	0.73	0.63
m_{bottom}	0.00 - 1.00	0.55	0.63	0.65	0.91
d_{bottom}	0.00 - 1.00	0.26	0.65	0.84	0.09
ω_{side}	0.00 - 1.00	0.83	0.61	0.78	0.35
m_{side}	0.00 - 1.00	0.04	0.39	0.82	0.66
d_{side}	0.00 - 1.00	0.48	0.94	0.62	0.27

TABLE 7.3: Best-fitting inputs to fit a synthetic light curve to the real target light curve.

Property	Value	Testing Range	Step
ψ	350	0 - 350	10
ϕ	60	0 - 90	10
ω_{top}	0.20	0.1 - 0.9	0.1
m_{top}	0.80	0.1 - 0.9	0.1
d_{top}	0.40	0.0 - 1.0	0.1
ω_{bottom}	0.40	0.1 - 0.9	0.1
m_{bottom}	0.40	0.1 - 0.9	0.1
d_{bottom}	0.70	0.0 - 1.0	0.1
ω_{side}	0.20	0.1 - 0.9	0.1
m_{side}	0.40	0.1 - 0.9	0.1
d_{side}	0.60	0.0 - 1.0	0.1

7.3 Results

The two primary results are the polar plots shown in Figs. 7.6 and 7.7. Each pixel in these plots shows the RMSE, calculated using Eq. 2.13, between the target light curve and a synthetic light curve generated with angular velocity vector offset angles given by the position on the polar axes. This range of synthetic light curves generated by varying the angular velocity vector offset angles are referred to as the ‘synthetic fitting light curves’. In these polar plots the radial-axis plots the angle ϕ (0° - 90°) and the rotational-axis plots the angle ψ (0° - 360°). Comparing with Figure 7.4, the origin of the polar plots corresponds to the X-axis, the coordinate $[\phi = 90^\circ, \psi = 0^\circ]$ corresponds to the Z-axis and the coordinate $[\phi = 90^\circ, \psi = 90^\circ]$ corresponds to the Y-axis. The RMSE in each plot was normalised against the maximum recorded RMSE for that plot.

In Figure 7.6, the target light curves are the synthetic light curves of Figure 7.5. These plots measure the RMSE between these synthetic target light curves and the synthetic fitting light curves generated with angular velocity vector offset angles given by the position on the plot. The RMSE was therefore due only to the difference in angular velocity vector, as all other inputs remained the same.

In Figure 7.7 the target light curve is the real light curve of the CZ-2D, and so this figure measures the RMSE between the real target light curve and the range of synthetic fitting light curves with varying angular velocity vector offset angles defined previously. Hence the RMSE shown in this plot was due to the difference in angular velocity vector in addition to errors in the modelling. These included: errors in the synthetic light curve model, such as disparities in the reflection calculated by the BRDF and reality; errors in the synthetic light curve inputs, such as differences between the object geometry of Figure 7.3 and the real rocket body geometry.

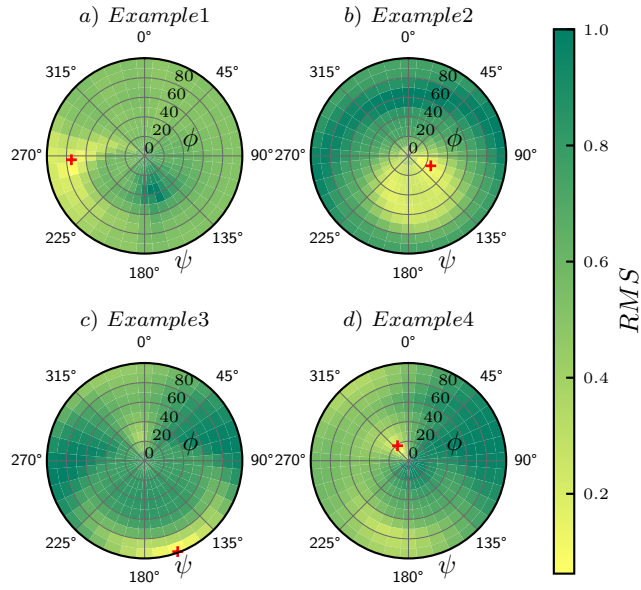


FIGURE 7.6: The RMSE measured between the synthetic target light curves and new synthetic light curves generated with a angular velocity vector direction given by the polar axes.




FIGURE 7.7: The RMSE measured between the real light curve from the MMT and the best-fitting synthetic light curve generated with a new angular velocity vector given by the polar axes.




FIGURE 7.8: Comparison between the real MMT light curve and the best-fitting synthetic light curve. This plot shows the first two rotation periods.




FIGURE 7.9: Top: Comparison between the real MMT light curve and the best-fitting synthetic light curve. This plot shows the full light curve. Bottom: Absolute residuals.

Figures 7.8 and 7.9 show the best-fitting synthetic light curve overlaid onto the real target light curve. Figure 7.8 shows only the first two full rotation periods, whereas Figure 7.9 shows the full light curve.

7.4 Discussion

7.4.1 Polar Plot Discussion

In the four plots of Figure 7.6, the only differences between the synthetic target light curves and the synthetic fitting light curves are the angular velocity vector offset angles. The purpose of these results was to demonstrate the operation of this angular velocity vector determination technique in cases where the true angular velocity vector is known. In all four plots, the pixel with the lowest RMSE is coincident with the actual offset angles of the angular velocity vector that were used to generate the synthetic target light curves, which are marked with red '+'s. The reason why this pixel does not display zero RMSE is due to the 10° bin size of the angular velocity vector offset angles. For example, the true angular velocity vector offset angles in the top left polar plot were $\psi = 267$ and $\phi = 65$, whereas the closest bin was $\psi = 270$ and $\phi = 70$, hence a non-zero RMSE is measured between these light curves. If the bin size was decreased to 1° , then there would be a single pixel that measured exactly zero RMSE.

Due to the normalisation of the RMSE values, these plots also provide a visual indication of the uncertainty in the solution. The plots show that the pixels that surround the correct angular velocity vector solution have values within 10 - 20% of the correct solution and the remainder of the pixels in the plot indicate larger RMSE values - i.e., the solution is unique. This result defines the expected ideal performance of this technique, where the system has been modelled perfectly, and provides a baseline against which to compare.

In Figure 7.7, the target light curve is the real light curve of the CZ-2D rocket body, hence the measured RMSE values are now due to differences in angular velocity vector offset angles combined with errors in the modelling. Because the target light curve is now a real light curve, the true angular velocity vector offset angles are unknown.

This plot is very similar to the results in Figure 7.6: the plot contains a clearly defined RMSE minimum (marked with a green '+') and the RMSE increases as distance from the minimum increases. However, unlike the synthetic results of Figure 7.6, this figure also encapsulates all of the modelling errors between the real target and synthetic fitting light curves. The most substantial modelling error was in the object geometry. The real rocket body has a much more complex structure than the cylindrical geometry used in the synthetic light curve generation, particularly on the end with the rocket nozzle. Furthermore, the attitude state used in the synthetic light curve was simplified and the angular velocity vector was assumed to remain constant in the orbital frame. In reality, the attitude motion may be more complex. Most of these

errors will be captured within the determined values for the BRDF inputs due to their use as fitting parameters.

Despite these differences, the polar plot of Figure 7.7 remains similar to the synthetic polar plots, for which there were no differences between inputs apart from angular velocity vector position. This indicates that this angular velocity vector determination technique is insensitive to errors in object geometry and BRDF modelling.

The technique was applied to a rocket body light curve that the real light curves of Figure 7.1 showed to be typical of rocket body light curves across a range of different rocket body types (CZ-2D, Delta, N-1). Furthermore, the rocket body geometry used was generic and not tailored to specifically represent a CZ-2D rocket body. This suggests that this angular velocity vector determination technique can be generalised for application to any rocket body light curve similar to those presented in Figure 7.1.

7.4.2 Light Curve Plot Discussion

The precise differences between the best-fitting synthetic light curve and the real MMT target light curve are shown in the final two plots, Figures. 7.8 and 7.9.

Firstly, the plot containing only the first two rotation periods, Figure 7.8, is considered. A visual inspection of this figure shows a number of features that are similar between the real and synthetic light curves. These include the phase of the two curves, the positions of the peaks and troughs and the fact that the brightest reflection has two peaks, with one being slightly larger than the other. However, there are also dissimilarities between the two light curves. Most notably, in the twin-peaked brightest reflection the two curves do not precisely overlap and in the two dimmest reflections the shape of the synthetic curve is not the same as in the real curve. The two dimmest reflections in this light curve correspond to reflections from the two ends of the geometry (the red and blue coloured facets of Figure 7.3). The ends of the rocket body will have the largest differences between the geometry used in the synthetic light curve generation and reality, due to features such as a rocket nozzle not being considered. This may be the cause of the differences between the synthetic and real reflections from the ends of object.

Now considering the full light curve shown in Figure 7.9. This figure additionally shows the absolute residuals with a trend line calculated using linear regression. The residuals in the plot show that the difference between the two curves increases with time. Because the synthetic light curve was only fitted to the first rotation period of the real light curve, the residuals plot shows that the closeness of the fit slowly decreases as the light curve progresses. This indicates that there were properties, other than the viewing conditions, that changed throughout the real light curve but remained constant in the synthetic light curve. This may be, for example, because the

angular velocity vector does not remain constant in the orbital frame, which was one of the assumptions used in the synthetic light curve generation. Other reasons for why this may be the case include a small inaccuracy in the rotation period value used, a rotation period which is slowly changing over time, additional rotational motion components beyond the ‘flat spin’ assumption, or small modelling errors resulting from the difference between synodic and sidereal rotation.

7.5 Conclusions

The results in this chapter suggest that fitting a synthetic light curve to a real light curve can be used to determine the approximate direction of the angular velocity vector, for a tumbling rocket body in a “flat spin” attitude state. In both the synthetic and real examples, varying the angular velocity vector resulted in similar relative changes in the RMSE. This indicates that the method was not substantially affected by errors present in the object geometry and that the BRDF parameters can be used for fitting. A unique angular velocity vector solution was obtained for the real light curve example. However, without truth data against which to compare, the accuracy of this solution can only be inferred from the similarities to synthetic results. Although only a single real light curve example was examined, the angular velocity vector determination was performed in a generalised way that was not specific to any particular light curve or rocket body type.

It should be noted that although using RMSE performed well in this case, there are circumstances where performance may be decreased when comparing real and synthetic light curve data. One such case is when the real light curve data contains specular glints which are unmodelled in the synthetic data. This circumstance becomes more likely when using simplified object geometry models, such as using a cylinder to represent a rocket body. In this case, small subsurface components with highly reflective surfaces, from which strong specular reflections could be received, are not present in the model. As a result, inflated RMSE measurements would be received from a synthetic light curve that would otherwise be a high quality fit to the real data. Where necessary, this issue can be mitigated by applying some form of outlier removal to the real data.

In the literature, application of a similar technique to a box-wing satellite in a flat-spin attitude state was used to determine the angular velocity vector. It was found that, similar to this chapter, a relatively simple box wing geometry could be used to match real light curve data and generate an attitude state solution (Piergentili et al., 2017). This provides confidence that the results of this chapter could be applied to objects other than rocket bodies.

Because the first steps of the fitting process were performed manually, all possible combinations of input parameters have not been tested. As a result, it is possible that additional best-fitting angular velocity vectors exist. Furthermore, only a single real light curve was considered, which exhibited a particularly clear structure. This was required in order to facilitate the manual fitting. However, it was shown that these light curves are common for tumbling rocket bodies. Although it is expected that the model would be capable of determining the angular velocity vector in light curves with less clear structure, this would need to be tested. The first steps of the future work will address these issues by automating the fitting process, such as through implementation of a genetic algorithm. This will allow a more thorough examination of the solution space in order to investigate the uniqueness of the solution. Additionally, this will allow a wider range of real rocket body light curves to be tested and so confirm that this technique is applicable to rocket bodies more generally.

Chapter 8

Analysis of the Brightness of Starlink Satellites

8.1 Introduction

In recent years, the global spaceflight sector has shifted from being the exclusive arena of governments and state-actors to a market for commercial exploitation. This has been facilitated by lower satellite manufacture and launch costs. One of the most popular developments resulting from this lower cost is to provide global high-speed, low-latency, broadband internet using a large constellation of satellites operating in LEO. According to the International Telecommunication Union (ITU), 51.2 % of the Earth's population did not have internet access at the start of 2019, with the majority of these people living in less-developed areas that are extremely difficult to serve using conventional infrastructure ([International Telecommunications Union, 2018](#)). With access to reliable internet becoming an increasingly important aspect of life in the 21st century, having the capability to 'beam' internet to anywhere on Earth could be revolutionary for how many people live their lives.

An important consideration for space-based internet infrastructure is latency. To achieve a low latency, internet-providing satellites will operate at relatively low altitudes. However, lower orbital altitudes lead to an increase in the number of satellites required to maintain coverage over the Earth's surface. Coverage considerations, combined with the number of users that a single satellite can support, dictate the required number of satellites in the constellation. Companies that are vying to provide space-based internet address this trade-off with a range of different altitude and satellite number solutions. OneWeb, for example, intends to operate 648 satellites at an altitude of 1200 km. Other examples include Boeing with 2956 satellites at 1200 km, Amazon with 3236 satellites at between 590 - 630 km and Samsung with 4600 satellites at an, as yet, unspecified orbital altitude less than 1400 km ([Reid et al., 2016](#)).

Of these companies, OneWeb is the only one to have begun deploying satellites on-orbit.

Another example is the SpaceX Starlink constellation, for which approval was granted by the Federal Communication Commission in 2017 for 12,000 satellites between 340 - 1150 km, with efforts underway to extend the number to as many as 42,000 satellites. The first deployment phase for Starlink was for the launch of 1600 satellites into 550 km orbits (Foreman et al., 2017). Although these precise numbers and altitudes have been revised several times, SpaceX certainly intends to operate thousands of satellites in LEO.

On 24th May 2019 a Falcon 9 rocket left the Kennedy Space Centre carrying the first set of 60 v0.9 satellites of SpaceX's Starlink constellation. These satellites, known as 'batch-0', were the final prototype systems. The first operational members of the constellation, v1.0, began with the 'batch-1' launch on 11th Nov 2019. As of 12th Nov 2020, 16 launches of v1.0 satellites have been completed giving a total of 833 launched satellites.

Almost immediately upon deployment of the first 60 v0.9 Starlink satellites, it was reported that they were clearly visible to the naked eye as a string of light sources in the night sky. Concerns amongst the world-wide astronomy community quickly grew that large constellations could have a detrimental effect on ground-based astronomy, both optical and radio. For optical astronomy, this can range from degradation of image quality due to streaks caused by satellites passing through the field of view during an exposure, to overexposing and damaging a sensitive instrument.

As part of the publicity surrounding the Starlink deployment, SpaceX released a number of computer-generated images of Starlink satellites operating on-orbit. An example image is shown in Figure 8.1. It was hypothesised that the unusual 'flat-pack' design (i.e., the satellite can be folded flat for improved volume efficiency) may be the reason that Starlink satellites were so bright. In particular, the large solar array may provide a large surface area from which reflected sunlight can be observed.

SpaceX has acknowledged that the brightness of the Starlink satellites is an unexpected problem and has stated that it is committed to ensuring that the Starlink constellation does not interfere with astronomical research. In April 2020, SpaceX released a statement discussing the brightness of Starlink satellites and the steps that were being taken to reduce it. This statement showed that the spacecraft change their configuration, depending on their operational state, to reduce atmospheric drag. Figure 8.2 shows these two configurations. On the left is the operational configuration, in which the angle between the solar panel and spacecraft bus is 90° and the spacecraft bus is orientated such that the large-area-face that contains the antennas are pointing towards Earth. SpaceX refers to this as the 'shark-fin' configuration. On the right is the non-operational configuration, in which the solar panel is folded flat, such




FIGURE 8.1: A render of a Starlink satellite on-orbit, generated by SpaceX and available at: www.starlink.com.

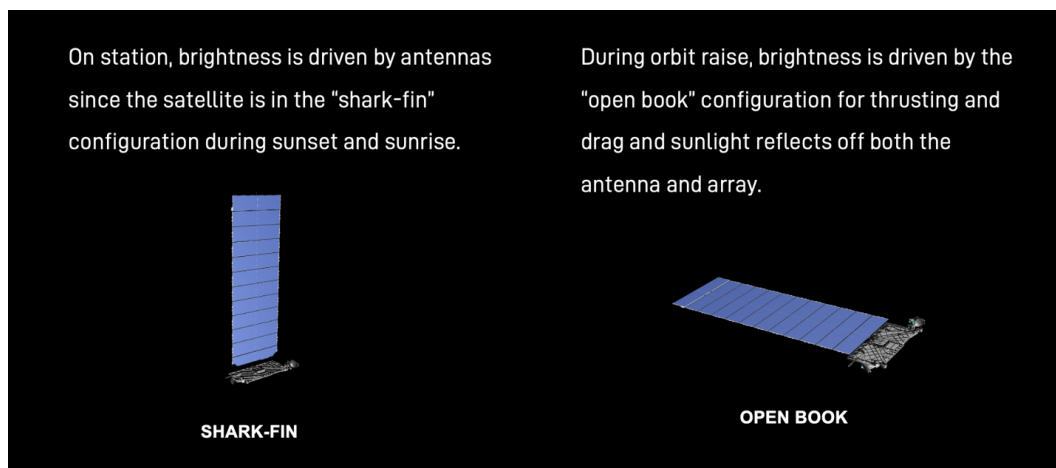


FIGURE 8.2: A render of the two geometry configurations used by Starlink satellites. Available at: <https://www.spacex.com/updates/starlink-update-04-28-2020/>.

that only the edge of the array is presented to drag. SpaceX refers to this as the ‘open book’ configuration.

The open book configuration is only used when the satellite is increasing its orbital altitude to reach the operational altitude of 550 km. Therefore, the system will spend the vast majority of its operational lifetime in the ‘shark-fin’ configuration. This is a potential explanation for why the Starlink spacecraft were so unexpectedly bright when first deployed on-orbit. In the open book configuration the largest possible area is presented to an observer and there is no possibility for the solar panel to be hidden from view behind the spacecraft bus. However, this configuration is also potentially

the easiest to mitigate brightness with procedural changes rather than hardware redesigns. For example, at those times in the orbit where the brightest reflections happen, the spacecraft's attitude can be changed such that it is 'edge-on' to an observer. This has already been considered and implemented by SpaceX.

Mitigating reflected brightness in the operational shark-fin configuration is less straightforward, as there are two reflective surfaces (the spacecraft bus and the solar array) with a 90° offset between them. Furthermore, these surfaces are required to point in particular directions to achieve the communications mission objective. SpaceX's first attempt to reduce the brightness of Starlink spacecraft in a shark-fin configuration was to reduce the albedo of the most reflective elements of the spacecraft bus: the antennas. This was trialled on Starlink-1130, launched as part of the 2nd operational batch of satellites on the 7th January 2020. Figure 8.3 shows that this was achieved by changing the colour of the antenna surfaces from white to black. Although this technique may potentially reduce the brightness of reflections, it will also result in an increase in absorption of solar thermal radiation, possibly causing the components to overheat.

Analysis of observations performed on Starlink-1130 (AKA "DarkSat") by Tregloan-Reed et. al., using a 0.6m telescope in northern Chile, showed a factor of two reduction in brightness of DarkSat when compared to an unmodified spacecraft from the same launch (Tregloan-Reed et al., 2020). There is a single light curve available for DarkSat in the MMT database. Although this single light curve exhibits brightnesses greater than visual magnitude 2.5, it was collected when the spacecraft was orbit raising and so is likely to be inflated by solar panel involvement. Since the launch of DarkSat on the 7th January 2020, no further spacecraft have had this darkening applied to the antenna surfaces. The reason for this is not public, but may be due to unsatisfactory brightness reduction performance, or due to undesirable changes to the thermal properties of the spacecraft that compromised the operation of the antennas.

The most recent change to the Starlink design, aimed at reducing brightness, is to use a deployable visor to shade the antenna surfaces from sunlight. Figure 8.4 shows how the deployment of this visor casts a shadow over the most reflective elements on the Earth-facing side of the spacecraft bus. A single so-called 'VisorSat' was tested in batch-7, launched on the 4th June 2020. In batch-9, launched on the 7th August 2020, every Starlink spacecraft had a sun visor. SpaceX has stated that it now intends for all subsequent Starlink satellites to be equipped with a visor.

This chapter uses the synthetic light curve model to investigate the causes of the high brightness of Starlink satellites whilst in an operational configuration. It was hypothesised that the large solar array was at least partially responsible for the high brightness. This was tested by comparing real light curves of Starlink satellites from the MMT database to synthetic light curves generated using geometries without solar




FIGURE 8.3: A render showing how the 'DarkSat' experiment aimed to reduce the brightness of Starlink satellites in a shark-fin configuration. Available at: <https://www.spacex.com/updates/starlink-update-04-28-2020/>.




FIGURE 8.4: A render showing how the 'VisorSat' experiment aimed to reduce the brightness of Starlink satellites in a shark-fin configuration. Available at: <https://www.spacex.com/updates/starlink-update-04-28-2020/>.

Batch	International Designator	# of Spacecraft	Launch Date
1	2019-074	60	11/11/2019
2	2020-001	60	07/01/2020
3	2020-006	60	29/01/2020
4	2020-012	60	17/02/2020
5	2020-019	60	18/03/2020
6	2020-025	60	22/04/2020
7	2020-035	60	04/06/2020
8	2020-038	58	13/06/2020
9	2020-055	57	07/08/2020

TABLE 8.1: Data of the Starlink batches considered in this chapter

panels and with solar panels in a shark-fin configuration. The difference between these two synthetic light curves would indicate the expected level of solar panel involvement in the reflected brightness. In addition, synthetic light curves were generated using a geometry that included a solar panel with and without self-shadowing (i.e., the proportion of facets that are shadowed by other facets and so do not contribute to the observed brightness). The difference between these two synthetic light curves was used to measure an expected level of solar panel shadowing by the spacecraft bus.

8.2 Methodology

8.2.1 Starlink Light Curve Data

Table 8.1 shows the 9 batches of Starlink spacecraft that were considered for this chapter and the graph at the top of Figure 8.5 shows all light curves in the MMT database for these satellites, which present as vertical lines due to the scale of the x-axis. This amounts to 2461 light curves from 490 objects. The graph at the bottom of Figure 8.5 shows their orbital altitude, from launch to the start of November 2020, calculated by converting all available TLEs into state vectors using SGP4. The first step in this analysis was to extract from the data of Figure 8.5 a sub-set of light curves that were collected while each object was in its final mission orbit. This was done so that the spacecraft could be assumed to be in a ‘shark-fin’ configuration. It was achieved by extracting only those light curves that were collected after the spacecraft had a minimum of three TLEs with altitudes within 2 km of the operational altitude of 550 km. Figure 8.6 shows the result of this filtering and contains 2051 light curves from 439 objects.

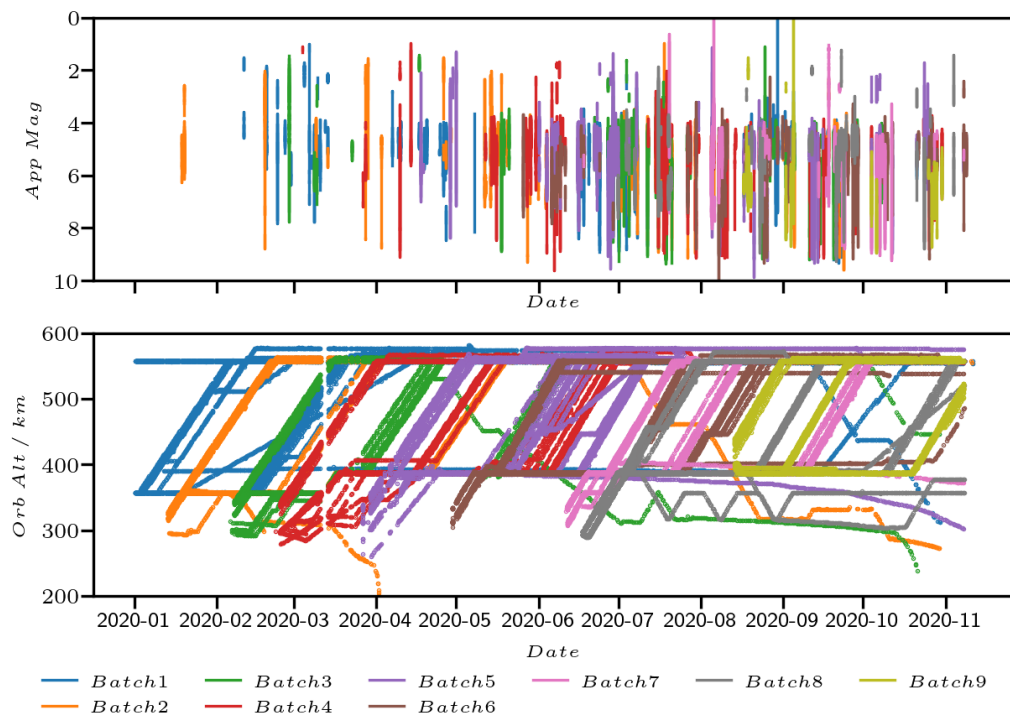


FIGURE 8.5: Apparent magnitude and orbital altitude, against date, for the first 9 batches of Starlink spacecraft.




FIGURE 8.6: Apparent magnitude and orbital altitude, against date, for the first 9 batches of Starlink spacecraft. Data has been filtered to show only the data for spacecraft whilst on-station.

8.2.1.1 Synthetic Light Curve Inputs

Synthetic light curves were generated for each real light curve in the subset shown in figure 8.6. To do this, a geometry model, BRDF parameters and attitude state were required. To test the hypothesis of this paper, one geometry model was required that represented only the spacecraft bus and a second geometry model was required that represented the full Starlink geometry including the solar array in shark-fin configuration. Figure 8.7 shows these two geometry models. Using Starlink images containing objects of known dimensions, such as the Falcon-9 payload fairing, the spacecraft bus was determined to be approximately $1.5 \text{ m} \times 3.0 \text{ m}$. By then examining the size of the solar array relative to the spacecraft bus in the imagery released by SpaceX, the solar panel was determined to be approximately $3.0 \text{ m} \times 12.0 \text{ m}$. The angle between the spacecraft bus and solar panel was maintained at 90° , as is shown in the Starlink renders released by SpaceX.

For the three BRDF parameters, diffuse albedo, diffuse-to-specular ratio and the roughness parameter, values of 0.45, 0.4 and 0.1 respectively were used for all facets. Best fitting values for these parameters were not determined. This was because the measurement important to this work was the difference between light curves generated with and without the solar panel, or the difference between light curves generated with or without self-shadowing. As a result, it is mostly unaffected by the accuracy of the BRDF modelling.

The final required input was the attitude state. No precise information is available on the attitude state adopted by Starlink satellites. However, a best-guess can be made based on the pointing requirements of the satellite components. Assuming a shark-fin configuration, the spacecraft was orientated such that the antenna-containing face of the spacecraft bus was Earth-facing. The spacecraft was then rotated about the remaining axis of freedom so that the angle between the normal vector of the solar panel and the Sun vector was minimised. This attitude state maintains the mission-critical pointing requirements of the payload while maximising the power-raising capability of the solar panel.

8.3 Results

8.3.1 Synthetic Light Curves

The first set of results examines the brightnesses of the synthetic light curves and their relationship to solar panel involvement and solar panel shadowing. Figure 8.8 plots the synthetic magnitudes of light curves generated with the full geometry in a shark-fin configuration, against solar panel involvement. Solar panel involvement was




FIGURE 8.7: The two geometries used to represent a Starlink system. On the left is the full spacecraft, including solar panel, and on the right is the payload-only geometry.




FIGURE 8.8: Minimum, maximum and average synthetic light curve brightness against solar panel involvement.

calculated as the RMSE between a synthetic light curve generated using only the spacecraft bus geometry and one generated with the full geometry. Therefore, the difference between these synthetic light curves is due to reflections from the solar panel. The left graph plots the maximum magnitude, hence the dimmest data-points, on the x-axis; in the central graph average magnitudes are plotted; and in the right graph the minimum magnitudes, hence brightest data-points, are plotted. Because magnitude is a reverse scale, the directions of the x-axes are reversed so that a position further to the right of the graph indicates a brighter light curve.

All three graphs in Figure 8.8 show a correlation between solar panel involvement and brightness: as brightness increases, so does solar panel involvement. A linear regression trend line has been added to each plot showing this correlation and the equation is provided above each plot. Note that the gradients of the lines are negative because of the reversal of the x-axis scale. In the central and right graphs, there are a small number of data points that are brighter than magnitude 6, but have zero solar panel involvement. These are caused by particular viewing conditions that result in specular reflections from the spacecraft bus, i.e., the angle of incidence was equal to the angle of reflection. Apart from this small number of exceptions, all light curves with brightnesses greater than magnitude 6.0 included reflections from the solar panel.

In Figure 8.9, solar panel shadowing is plotted in place of solar panel involvement. Solar panel shadowing is measured as the RMSE between two light curves, generated using the geometry that included the solar panel, with and without self-shadowing. This measures the reduction in light curve brightness that was caused by shadowing of the solar panel by the spacecraft bus. Once again, the left graph shows maximum magnitude, the middle graph shows average magnitude and the right graph shows the minimum magnitude. These three plots each have the same three features, which are indicated on the central plot using a dotted, dashed or solid ellipse.

The dotted ellipse highlights a trend where synthetic light curve brightness increases as shadowing increases. This can be explained by referring to Figure 8.8 which showed that, in this region, as brightness increases solar panel involvement also increases. Due to the assumed attitude state, where the antenna-containing spacecraft bus facets remain Earth-pointing, it is not possible both for the solar panel to be involved in the reflection and for the shadowing to be zero. There will always be at least a small fraction of the solar panel obscured by the spacecraft bus. Hence, as brightness increases because of solar panel reflections becoming more dominant, there will be a higher capacity for the brightness to be reduced by solar panel shadowing.

Peak shadowing occurs at approximately magnitude 6.0. This shows that there were a number of light curves with magnitudes of approximately 6.0, that would have been brighter if not for shadowing of the solar panel by the spacecraft bus.

The second feature, which has been highlighted using the dashed ellipse in the central plot of Figure 8.9, shows a trend where, beyond magnitudes of 6.0, shadowing decreases with increasing brightness. As Figure 8.8 showed, the majority of the light curves with brightnesses above magnitude 6.0 were bright as a result of higher solar panel involvement. Hence, these light curves increase in brightness precisely because the shadowing is decreasing. If shadowing could be increased, such as by modifying the attitude state, then both solar panel involvement and brightness would decrease in this region.

The final feature in these plots is the 'shoulder' inside the solid ellipse. These light curves increase in brightness for an approximately constant level of solar panel shadowing. There are few or no light curves that appear below this feature - i.e., magnitudes between 6.0 and 4.0 and shadowing below 0.2. This is showing a result that has already been discussed - the brightest light curves are caused by solar panel involvement, which, due to the selected shape and attitude state, must have some degree of solar panel shadowing. Figure 8.8 has shown that light curve brightness increases from magnitude 6.0 to 4.0 as a result of increasing solar panel involvement. Therefore, the group of light curves inside the solid ellipse and the group of light curves inside the dashed ellipse show two reasons why the solar panel causes increases in brightness. The first reason is that the illumination conditions can change




FIGURE 8.9: Minimum, maximum and average synthetic light curve brightness against solar panel shadowing.

in a way that results in a brighter reflection from the solar panel for a constant level of shadowing - the solid ellipse. The second reason is that the illumination conditions can change such that a lesser proportion of the reflection from the solar panel is shadowed by the spacecraft bus - the dashed ellipse.

In the right-most graph of Figure 8.9, which shows the brightest data-points, four light curves were brighter than magnitude 3.0 (coloured orange). These are the same four light curves that the right-most graph of Figure 8.8 showed to have zero solar panel involvement, where the high brightness was caused by specular reflection from the spacecraft bus.

Together, Figures 8.8 and 8.9 establish the expected behaviour of Starlink light curve brightness, derived from a synthetic light curve model using best-guess assumptions of configuration and attitude state. If the modelling has been performed accurately, then the same behaviour should be present in the real light curve data.

8.3.2 Real Light Curves

In Figures 8.10 and 8.11, solar panel involvement and solar panel shadowing from the synthetic results is plotted against the real light curve brightnesses from the MMT database.

The trend lines in Figure 8.10 suggest that there is a correlation between light curve brightness and solar panel involvement in the real light curve data. However, the strong correlation between brightness and solar panel involvement that was expected from the synthetic data is not present.

Figure 8.11 shows the MMT brightness data plotted against solar panel shadowing predicted from the synthetic data. The trend lines in these plots appear to indicate that solar panel shadowing increases with increasing brightness, similar to the trend within the dotted ellipse of Figure 8.9. However, the two other features that were




FIGURE 8.10: Minimum, maximum and average real light curve brightness against solar panel involvement predicted in the synthetic data.




FIGURE 8.11: Minimum, maximum and average real light curve brightness against solar panel shadowing predicted in the synthetic data.

present in the synthetic data (highlighted by a dashed and a solid ellipse in Figure 8.9) are not present in the real data.

An additional result is presented in Figure 8.12. In this plot, each of the light curves that were shown in Figure 8.6 have been replaced with a single data point of the average magnitude. These data points were then averaged to generate horizontal dashed lines for each batch. This plot shows that batches 1 through to 8 all exhibited approximately the same brightness on average. A small decrease in brightness can be observed in batches 6 and 7. This may be a result of procedural changes implemented by SpaceX, although this may also be the result of other mechanisms or variance in the data. For Batch 9 however, a clear decrease in average brightness can be observed. On average, these light curves were approximately 1 apparent magnitude dimmer, which corresponds to a brightness reduction factor of 2.5. Batch 9 was the first batch where every satellite had a sun-visor. This suggests that the spacecraft bus, specifically the antenna surfaces, were a major contributor to the brightness of Starlink satellites and that the visor is playing a role in mitigating the brightness problem.




FIGURE 8.12: Average light curve brightness of all available data in the MMT database for the first 9 batches of operational Starlink satellites.

8.4 Discussion

Using information for the Starlink configuration released by SpaceX with assumptions for attitude state, synthetic light curves were generated for Starlink satellites. In these synthetic light curves, the various positions of the satellite with respect to the observer created a range of viewing conditions. The calculated brightness of reflections from the solar panel, and the reduction in brightness caused by shadowing of the solar panel by the spacecraft bus, were determined by the combination of attitude state, configuration and these viewing conditions.

It was found that, in the synthetic data, brightness increased with solar panel involvement. Additionally, shadowing was shown to play a role in decreasing the brightness of many of the synthetic light curves. With the exception a very small number of light curves where the viewing conditions lead to extremely bright reflections from the spacecraft bus, the brightest light curves were caused when solar panel involvement was high and solar panel shadowing was low.

This synthetic light curve data was used to define expectations of the real data. If the modelling and the assumptions were accurate, then the brightnesses of the real light curves should correlate with the predicted solar panel involvement and shadowing in a similar way to the synthetic brightnesses. However, although real light curve brightness increased with increasing predicted solar panel involvement, the trend was diminished compared to the synthetic data. In the shadowing data, two of the features

that were present in the synthetic data were not present in the real data. The reason why the real data did not match the synthetic data more closely is most likely to be due to some breakdown in the modelling or assumptions.

One possibility is the BRDF modelling. It may be that, although visible, the reflections from the solar panel are very dim compared to those from the spacecraft bus. This could be caused by the reflection from the solar panel being mostly specular in nature and the viewing geometry being such that the specular reflections do not reach an observer on Earth, instead being reflected into space.

The next possibility is the inaccuracies of the geometry model. Some imagery released by SpaceX show the solar panel to be flat, such that each segment of the panel has the same normal vector (see Figure 8.2). This is how the panel was modelled for this work. However, other images show the solar panel as having more of a 'saw tooth' profile (see Figure 8.1). This 'saw tooth' shape would therefore exhibit different reflective characteristics.

Another possibility is spacecraft configuration. Imagery released by SpaceX only ever shows the spacecraft in a configuration where the solar panel is in the same plane as the spacecraft bus ('open book'), or angled from the spacecraft bus by 90° ('shark-fin'). However, the fact that the spacecraft can transition between these two configurations shows that it has the capability to actuate the solar panel. Hence, it may be that the solar panel is rarely in a position where the angle from the spacecraft bus is 90° , rather, the solar panel actuates in order to increase the incident flux from the Sun. Furthermore, SpaceX may actuate the solar panel to maximise shadowing and decrease potential brightness. This possibility is supported by the large reductions observed in synthetic light curve brightness that were caused by solar panel shadowing. It may be that solar panel actuation can be used to increase shadowing to a point that the solar panel is rarely the cause of any bright reflections.

The next possibility is the attitude state. The attitude state used in this work was derived from two key assumptions. The first was that the faces containing the antennas are maintained as Earth-pointing. A high level of confidence can be placed on this assumption, as performing the mission objective is the primary goal for any spacecraft. However, the margin of error on this pointing requirement is unknown. The Starlink satellites therefore may not always be required to be perfectly Earth-pointing. If this constraint is loosened, then the attitude can be further modified to improve power raising. This may be combined with solar panel actuation in order to adopt an attitude state and configuration combination that maximises power generation while maintaining the antenna-pointing vector within some margin of Earth-pointing.

The final result presented in this work, Figure 8.12, showed that the 9th batch of Starlink satellites were dimmer on average than all preceding batches by

approximately 1.0 apparent magnitude. All satellites in this batch were equipped with a sun visor for shading the antennas. The fact that such a large brightness decrease could be achieved by modifying the reflection characteristics of the spacecraft bus indicates that it is the bus, rather than the solar panel, that was the primary cause of high Starlink brightness.

8.5 Conclusions

The objective of this chapter was to examine the hypothesis that solar panel involvement is the cause of the bright reflections detected from Starlink satellites. This was shown to be true in synthetic light curve data - light curve brightness increased with solar panel involvement and the brightest light curves were caused by the highest solar panel involvements. In real light curve data, this trend was observed but substantially reduced. Instead, the results suggest that the spacecraft bus, specifically the antennas, are a larger cause of reflected brightness.

Some reasons for why this was the case have been discussed, which present good candidates for future work. By examining a less constrained range of BRDF, geometry and attitude state inputs, it may be possible to identify why the real light curve data suggests that the solar panel plays a smaller role in detected brightness than was predicted in the synthetic data. Performing this future work would allow more definitive conclusions to be made concerning solar panel involvement in bright Starlink light curves.

Chapter 9

Conclusions

9.1 Introduction

In this conclusion chapter, the key findings of each of the content chapters are summarised. This is followed by a discussion of the overall thesis conclusions, how they address the thesis objectives, and how this work can be used to support future research and development in the field of spacecraft attitude determination from optical light curve data.

9.2 Summary of Results Chapters

9.2.1 Chapter 3: Investigating Simple Light Curve Properties

In this first content chapter, the results of a pilot study were presented. The objective of this study was to test the hypothesis that simple light curve properties could be used to characterise the operational state of spacecraft. The specific property that was examined was standard deviation.

In general, an inactive spacecraft no longer maintains a particular attitude state, such as pointing antennas towards Earth. As a result, the spacecraft's attitude state will evolve over time due to on-orbit torques that are no longer being corrected by the onboard attitude control subsystem. Therefore, a spacecraft would tend to be observed in a variety of different attitude states and consequently would exhibit a wider range of different light curve brightnesses. Conversely, an Earth-pointing active spacecraft would be observed in a consistent attitude state and so would exhibit a narrower range of brightnesses.

Hence, the novel hypothesis of this work was that an inactive spacecraft would exhibit larger standard deviations in light curve brightness than an active spacecraft. This

hypothesis was tested by examining real light curve data to determine if a standard deviation threshold could be applied that differentiated active and inactive spacecraft. These light curves were also simulated using the synthetic light curve model. This was done with either a consistent attitude state with zero angular velocity in the orbital reference frame, or a random initial orientation and tumbling motion. The purpose of generating these synthetic light curves was to show that uncontrolled attitude motion can cause increases in standard deviation.

The results showed that a threshold value could be applied that separated real light curves of active and inactive Globalstar satellites, and the synthetic light curves showed that the differences in standard deviation could be caused by controlled or uncontrolled attitude states. Although additional objects need to be considered to improve the confidence in this conclusion, the results indicated that a standard deviation threshold could be used to identify the operational status of these box-wing spacecraft. This result was used to identify possible failure dates of three Globalstar satellites, for which this information is not publicly available.

The technique demonstrated in this chapter may be very beneficial for application in an operational setting. Assuming that the light curve data is available, this technique can be used operationally to perform computationally inexpensive analysis of the statuses of large numbers of spacecraft. Through continual monitoring of the light curve data, detection of anomalies in standard deviation can be used to identify changes to attitude state or configuration, which can then be followed up with further observations and more complex analysis.

9.2.2 Chapter 4: Investigation of Variable Rocket Body Light Curve Periods

This chapter performed the first detailed examination of the unusual variability in rocket body light curve period data that was one of the original motivations of this PhD. The objective of this chapter was to investigate possible causes of the light curve period variations.

The initial data set used in this chapter was a list of objects identified by the IADC as exhibiting unusual light curve period variability. Examination of the light curves of these objects showed that they could be categorised into one of two categories. In the first category, light curve period increased in an orderly fashion over time. In the second category, light curve period was observed to increase and decrease in a complex manner.

To ensure that the changes in light curve period were true reflections of changes in the signal periodicity, and not caused by the technique used to extract them, the light curve period extraction results were examined. It was found that the period-folded light curves exhibited a consistent shape, and that the variability of the extracted

periods was the result of a ‘stretching’ of this shape. This indicated that the period extraction results were representative of actual changes in signal periodicity.

From the initial set of objects that were examined, it was found that objects in the orderly category were in low-eccentricity orbits and objects in the complex category were in high-eccentricity orbits. This result suggested that it was the orbital characteristics that influenced the differences between the two categories. However, when the analysis was extended to include a much larger range of objects, separability by eccentricity was less reliable.

Although the results do not categorically prove that orbital characteristics are the cause of the differences between the categories, further work is required to better understand the categorisation process and its application to a larger data set.

9.2.3 Chapter 5: Synthetic Light Curve Model Sensitivity Study

If a synthetic light curve model is to be used in the analysis of spacecraft brightness data, it is important to understand the sensitivity of the model to the various modelling inputs. The objective of this chapter was to quantify the change in a synthetic light curve that resulted from a change in each of the inputs. Changes were quantified by measuring the RMSE between two light curves.

The inputs that were tested were as follows: object geometry model; phase angle; reflection model parameters (BRDF); initial attitude state; and rotation rate. Where possible, the change to the synthetic light curve caused by each of these inputs was tested under both diffuse and specular reflection conditions. A range of geometry models was tested that were used to represent a sphere with varying levels of complexity, which was achieved by changing the number of facets.

The results showed that the diffuse component of reflection was most sensitive to phase angle and diffuse albedo modelling. The specular component of reflection was more sensitive to all model inputs and was most sensitive to attitude state. When the diffuse and specular components were combined, the model was most sensitive to the modelling of the ratio between diffuse and specular reflection.

This chapter showed that modelling of the specular component of reflection is important for attitude determination applications, as it is a substantial cause of differences in light curve brightness resulting from attitude state. Modelling of specular reflections may greatly increase the complexity and computational expense of a synthetic light curve model. The conclusions of this chapter, therefore, form an important justification for making this addition.

9.2.4 Chapter 6: Investigation of the Reflection Characteristics of Primitive Geometries

The objective of Chapter 6 was to examine the relationship between illumination conditions and attitude state to identify the conditions under which different attitude states were most distinguishable. This was performed by quantifying the difference between light curves of rotating primitive geometries under a range of illumination conditions. The geometries considered were a cube, a 20-sided icosahedron and a 500-sided polyhedron. The differences between no-rotation and rotations about each of the three body-fixed axes were quantified under both diffuse and specular reflection conditions.

The results showed that the directions of the illumination and observation vectors, with respect to the angular velocity vector, determined the differences measured between light curves with different angular velocity vectors.

To demonstrate how these results may affect attitude determination from light curve data in practice, an example light curve generation scenario was constructed. This scenario showed that, for a cube geometry, when the angular velocity vector was in a plane defined by the observation and illumination vectors, the resultant light curve was exactly the same as a light curve with no rotation.

This chapter highlighted that the change in a light curve due to attitude motion is highly dependent on the illumination conditions, and that it is possible for a light curve of a tumbling object to appear aperiodic. This was an important conclusion, as it provided evidence that the attitude motion of an object may not always be captured within its light curve. Using this conclusion, observation and data selection strategies should be devised such that attitude determination analysis is performed using a data set that is most likely to have captured the effects of attitude motion.

9.2.5 Chapter 7: Angular Velocity Vector Determination of Rocket Bodies

This chapter presented a technique for angular velocity vector determination of rocket bodies. This was achieved by performing a brute-force fitting process whereby all input parameters of the synthetic light curve model were varied to find the set of inputs that produced the best-fitting synthetic light curve.

Because angular velocity vector truth data is not available for real rocket body light curves, this technique was first tested on purely synthetic data where the correct angular velocity vector solution was known. The results showed that the technique could identify the correct angular velocity vector direction as a unique solution. The technique was then applied to a real light curve of a tumbling rocket body, which also provided a unique angular velocity vector solution. The strong similarities between

the real and synthetic results that suggested that the correct angular velocity vector solution had been identified. This conclusion was further supported through comparing the real and best-fitting synthetic light curves, which were very similar and shared many of the same features.

The work of this chapter provides a foundation on which a robust process for attitude determination of rocket bodies can be developed. This could be applied in a wide range of areas in which rocket body attitude motion is of interest, such as high-accuracy orbital propagation for collision avoidance/reentry analysis, and ADR.

9.2.6 Chapter 8: Analysis of the Brightness of Starlink Satellites

This final content chapter applied the synthetic light curve model to Starlink satellites to determine reasons why they appeared as particularly bright objects in the night sky. The hypothesis of this work was that large-area solar arrays played an important role in the brightness of these reflections.

The synthetic light curve model was used to generate expected levels of solar panel involvement and solar panel shadowing in the Starlink light curves. These properties were then plotted against synthetic light curve brightness to uncover the expected relationship between solar panel involvement/shadowing and light curve brightness. The results showed that synthetic brightness increased with increasing solar panel involvement. The results also showed a clear relationship between synthetic brightness and shadowing.

The expected solar panel involvement and shadowing were then plotted against real light curve brightness. It was found that the real brightnesses did not correlate with solar panel involvement and shadowing in the same way, or as clearly, as in the synthetic data.

This led to one of two conclusions: either the hypothesis was incorrect and the solar panel does not contribute to light curve brightness to the extent that was expected; or the attitude and configuration modelling of the synthetic light curves was inaccurate. Some reasons why this may be the case were presented and discussed. An additional result was also presented that showed a 1.0 visual magnitude decrease in average brightness of Starlink satellites with sun-visors. The fact that modifying the reflection characteristics of the spacecraft bus had such a large effect on average brightness suggested that the solar panel was not the primary cause of the brightest reflections.

This chapter showed how light curve analysis and a synthetic light curve model can have a wider range of applications than attitude characterisation. In this example, the synthetic light curve model was used to analyse the brightness of an active object with the aim of discovering the causes of the brightest reflections and how they could

potentially be minimised. A synthetic brightness model would therefore be an ideal tool to companies such as SpaceX, who want to understand and mitigate the brightness of their satellites.

9.3 Thesis Conclusions

The research aim of this PhD thesis was to demonstrate how optical light curve data can be used to derive information on space objects, and potentially augment present day SSA/SST capabilities. This aim was broken down into five key objectives.

The first objective was to develop a synthetic light curve model capable of reproducing real light curve data. This objective was achieved and the development and operation of the model is discussed in Chapter 2.

The second research objective was to use the synthetic light curve model to improve understanding of real light curve data. This was achieved in Chapter 3, where the model was used to show that the differences in light curve standard deviation data observed between active and inactive Globalstar spacecraft could be explained by tumbling and non-tumbling attitude states. Deriving useful information from simple properties is a powerful tool, as it can be applied across large numbers of objects to persistently perform anomaly detection. Furthermore, a technique of this kind does not require complex modelling or analysis, encouraging its potential widespread use. Another simple light curve property that can provide useful insight is light curve period, which was examined in Chapter 4. Analysis of light curve periodicity has a long heritage in determining spin-rates of both natural space objects (such as asteroids) and artificial space objects (such as rocket bodies and satellites). In this work, changes in light curve periodicity were used as indicators of changes in the spin-rate of tumbling rocket bodies, which is important to understand for ADR applications. These results, therefore, form an important contribution towards understanding the dynamics of rocket body tumble rate.

The third research objective was to evaluate the performance of the synthetic light curve model characteristics that are most important to attitude determination. An example of a model characteristic important to attitude determination is the sensitivity of the synthetic light curve to attitude state. This objective was achieved in Chapter 5, which presented a study of the sensitivity to each of the model inputs. In addition, Chapter 6 further evaluated the performance of the model for attitude determination by identifying the illumination conditions under which the largest differences are observed between light curves generated with different attitude states. If synthetic light curve models are to see widespread use for analysis of space objects, then performing this work is important for ensuring the models are well understood and are used effectively.

The fourth objective was to evaluate how the synthetic light curve model could be used to perform attitude determination. This was achieved for defunct rocket bodies in 'flat spin' attitude states in Chapter 7. The future work of this chapter would be to examine the feasibility of applying this technique more generally to a wider range of object types.

The fifth and final research objective was to demonstrate how a synthetic light curve model could exploit the relationship between attitude and brightness for purposes other than attitude determination. This was achieved in Chapter 8, which was focused on how attitude state could be used to control spacecraft brightness. With the increase in the number of objects in LEO that is expected in the near future, ensuring that the brightnesses of these objects do not pollute the night sky will become a larger and larger concern, in particular for protecting astronomical research. A synthetic brightness model is ideally suited to this purpose.

This thesis has therefore achieved its research aim and the results have shown how optical light curve data can form a valuable source of information on space objects, and that a synthetic light curve model can be a valuable tool for the analysis of this data.

With the low cost of deploying an optical system for observing satellites, and the increased interest in low-cost satellite observation data, the availability and utilisation of optical light curves will increase. The results of this thesis will, therefore, be a valuable resource for the continued development of space object brightness analysis and modelling capabilities.

Appendix A

Quantifying the Difference Between Light Curves using Hilbert Space

A.1 Introduction

In Sections 2.3.6.1 and 2.3.6.2 of Chapter 2 (Synthetic Light Curve Model Development), two methods for quantifying the difference between light curves were presented. The first method was to calculate RMSE. The second method was a more complex approach using Hilbert space. Subsequently, it was found that quantifying differences between light curves using Hilbert space did not produce substantially different results than using RMSE.

In this appendix, use of Hilbert space distance calculations to quantify the difference between light curves is examined, and it is shown that the results are extremely similar to RMSE results. This is performed by presenting two key example results from this thesis using both RMSE and Hilbert space distance.

A.2 Recap: Hilbert Space Distance Calculation

To calculate the difference between two light curves, their position in Hilbert space with respect to a selected set of orthonormal basis functions must be determined. The set of basis functions selected for this work was the Fourier series due to their suitability for application to periodic signals. Hence, the coordinates of a function in Hilbert space can be described using a set of Fourier coefficients, and the distance between two functions in Hilbert space can be calculated using standard vector mathematics. The distance between two points in Hilbert space is zero if and only if the functions are exactly the same.

Calculation of a Hilbert space distance is performed as follows (key equations from Chapter 2 are repeated here for convenience). Firstly, the light curve is mirrored in the negative x domain using Eqn. A.1. This is done so that the Fourier series of the light curve is reduced to only cosine terms. Next, the Fourier coefficients of the light curve are calculated using Eqn. A.2, where n defines the total number of terms used. A value of n must be selected that is sufficiently large that the approximated signal closely matches the original light curve. The same process is applied to a second light curve from which the distance needs to be measured. Finally, the Fourier coefficients of both light curves are entered into Eqn. A.3 to calculate the distance between the two light curves in Hilbert space.

$$f(x) = f(-x) \quad (\text{A.1})$$

$$a_n = \frac{2}{L} \int_{-L}^L f(x) \cos nx dx. \quad (\text{A.2})$$

$$d^2 = \sum_{n=1}^{\infty} (a_n - b_n)^2. \quad (\text{A.3})$$

Additionally, using the Fourier coefficients generated from Eqn. A.2, the approximated signal can be generated by the equation

$$f(x) \sim \frac{1}{2}a_0 + \sum_{n=1}^{\infty} (a_n \cos nx + b_n \sin nx), \quad (\text{A.4})$$

where b_n remains zero for all values of n , as a result of the application of Eqn. A.1.

A.3 Example 1

This first example assesses the use of Hilbert space distance to quantify differences between synthetic light curves generated for Chapter 5 (Synthetic Light Curve Model Sensitivity Study). The particular result selected for this demonstration was Figure 5.11, which quantified the change to a synthetic light curve that resulted from changes in rotational velocity about the body-fixed x , y and z -axes. This result was selected because changes in light curve periodicity caused by changes in rotation rate was the result most likely to be well suited to a Hilbert space calculation using Fourier basis functions.

Figure A.1 shows the result from Chapter 2 using RMSE. This figure plots the RMSE between a baseline light curve generated with a non-rotating geometry, and light

curves with an increasing rotation rates for five different geometry models. Further detailed explanation of this figure is available in Chapter 2.

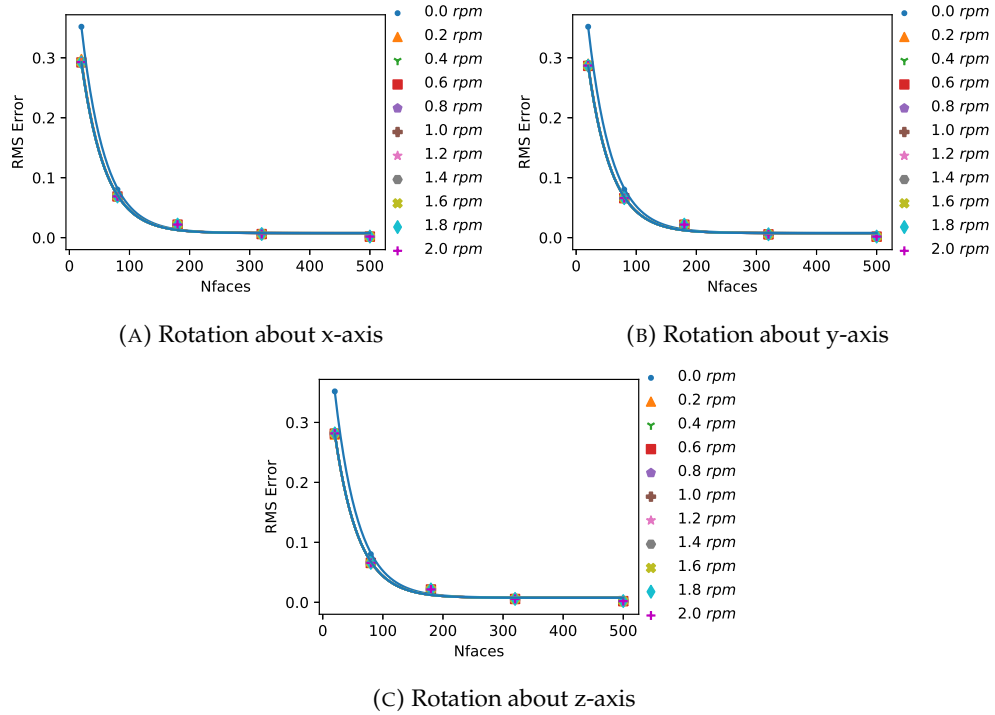


FIGURE A.1: The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.

The first step in reproducing this figure using Hilbert space distance was to find the Fourier coefficients of the baseline light curve. Figure A.2 shows the baseline signal and the approximated baseline signal produced using the determined Fourier coefficients. The plot shows that, when using a sufficient number of terms, the Fourier coefficients are a good approximation of the baseline signal. The RMSE between the baseline and the approximated baseline signal is 0.000248.

Figure A.3 shows an example light curve generated with a rotation about the body-fixed x -axis of 0.2 RPM, and its Fourier approximation. This plot shows that the approximated signal is a good match to the original light curve. The RMSE between these two curves is 0.000191.

Using the Fourier coefficients generated for each light curve, Hilbert space distances from the baseline can be calculated using Eqn. A.3. Figure A.4 shows the new result using Hilbert distance in place of RMSE. Comparing this figure to the original in A.1, the Hilbert distance values are different from the RMSE values, but the relative changes between data points are identical. Hence, changing to a Hilbert distance measurement does not change the interpretation of the results.




FIGURE A.2: The baseline light curve from Figure A.1 (black) and the approximated signal generated from the Fourier coefficients determined for the baseline light curve (red).

A.4 Example 2

In the next example, a result from Chapter 6 is reproduced using Hilbert space distance calculations. Figure A.5 shows the RMSE between a baseline light curve, where the observation vector direction is equal to the illumination vector, and a light curve with an observation vector direction indicated by the position of the pixel on the polar plot. This is performed with a non-rotating geometry, and three cases for rotations about each of the body-fixed axes. Figure A.6 shows the subtractions between the polar plots of the rotating cases. The geometry used in these light curves was a 5-frequency subdivided polyhedron. Again, more detailed descriptions of these figures are available in Chapter 6.

Figures A.7 and A.8 show the same results but using Hilbert distance in place of RMSE. There are some very subtle differences between these results and the results that used RMSE. However, the interpretation of the results is identical.




FIGURE A.3: A light curve from Figure A.1 with a rotation about the body-fixed x -axis of 0.2 RPM (black), and the approximated signal generated from the Fourier coefficients determined for this light curve (red).

A.5 Conclusion

In this appendix two representative examples were shown comparing the uses of RMSE or Hilbert space distance to quantify the differences between light curves. It has been shown that using a Hilbert space distance calculation does not change the interpretation of the results over using RMSE. It may be possible that Hilbert space distance calculations provide additional information over RMSE when considering the changes in the individual Fourier coefficients, but this was not the case when this information is reduced to a 1-dimensional distance calculation.

It was, therefore, concluded that RMSE was more suitable for this application. This was for two reasons. Firstly, RMSE is theoretically less complex than Hilbert space distance calculation and so improves the understandability of the results. Secondly, RMSE is much less computationally complex, which greatly reduces computation time when large numbers of calculations need to be performed. Hence, RMSE is used exclusively throughout this thesis.

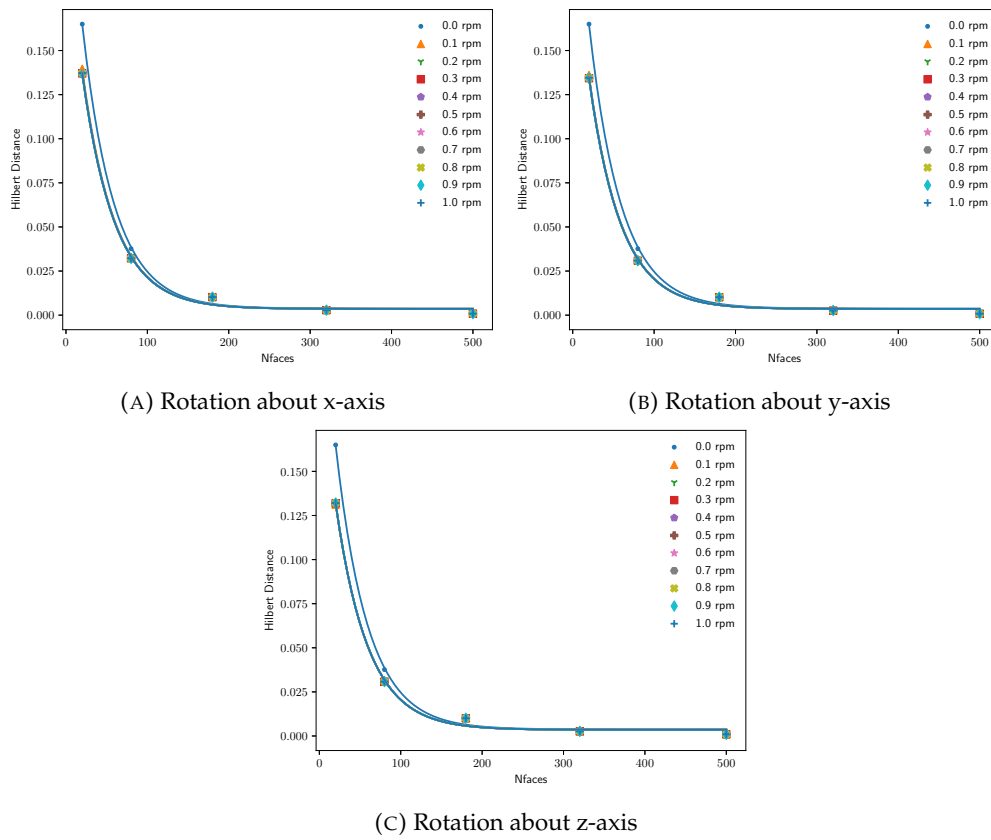


FIGURE A.4: The change in a synthetic light curve resulting from changes in attitude under entirely specular reflection.




FIGURE A.5: Results for a specular reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.




FIGURE A.6: The result of performing a subtraction between two of the plots from Figure 6.14. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a 5-frequency subdivided polyhedron

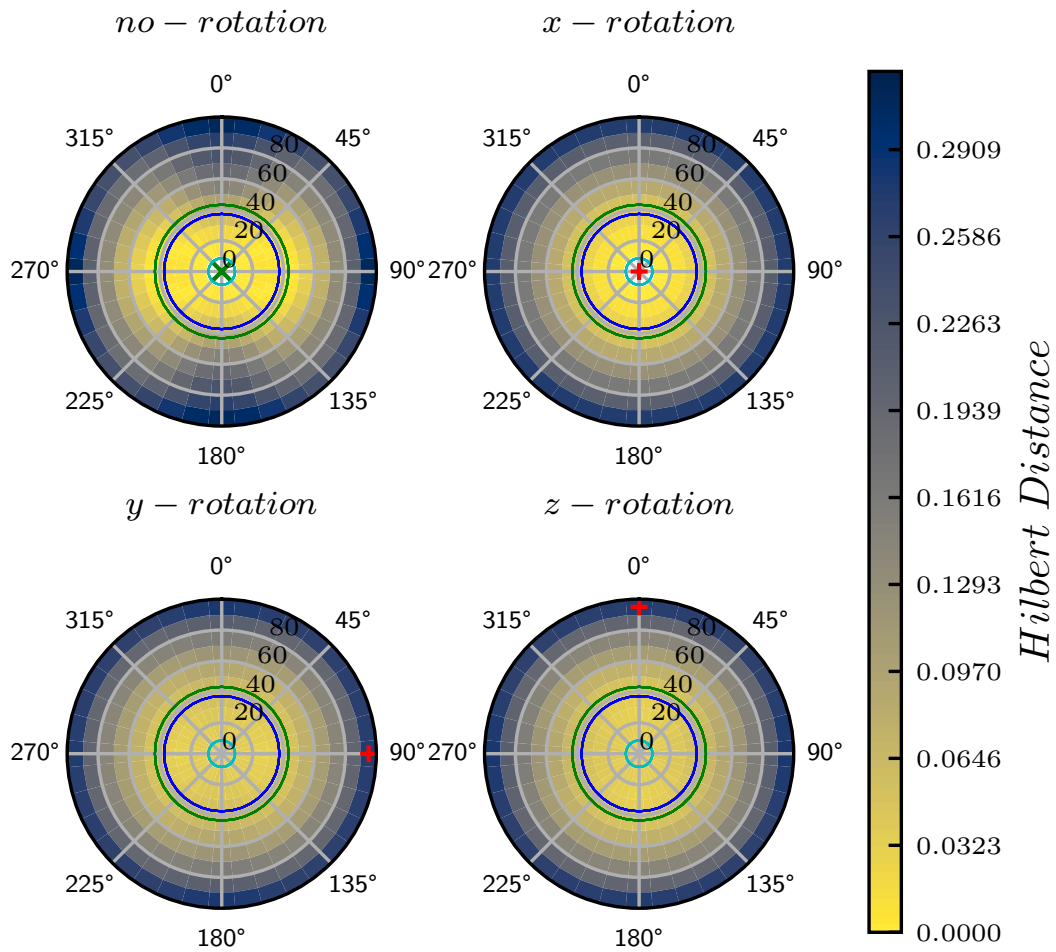


FIGURE A.7: Results for a specular reflection model and a 5-frequency subdivided polyhedron object geometry. RMSE is measured from a no-rotation no illumination offset baseline.

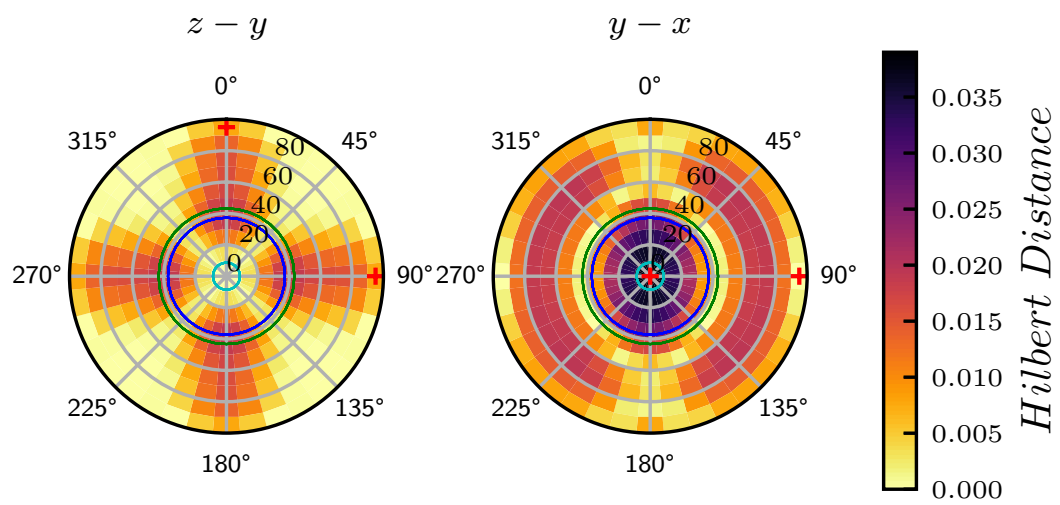


FIGURE A.8: The result of performing a subtraction between two of the plots from Figure 6.14. The two rotation cases in question are labelled above each polar plot. The reflection model is specular and the object geometry is a 5-frequency subdivided polyhedron

References

- Guglielmo S. Aglietti, Ben Taylor, Simon Fellowes, Thierry Salmon, Ingo Retat, Alexander Hall, Thomas Chabot, Aurélien Pisseloup, C. Cox, A. Zarkesh, A. Mafficini, N. Vinkoff, K. Bashford, Cesar Bernal, François Chaumette, Alexandre Pollini, and Willem H. Steyn. The active space debris removal mission RemoveDebris. Part 2: In orbit operations. *Acta Astronautica*, 168(September 2019): 310–322, 2020. ISSN 00945765. .
- P Anz-Meador, M Ward, A Manis, K Nornoo, B Dolan, C Claunch, and J Rivera. The Space Debris Sensor Experiment. *First International Orbital Debris Conference*, 3(6026): 10, 2019.
- R. B. Bent. Attitude Determination of the Ariel III Satellite. Technical report, 1967.
- A. Biryukov, G. Beskin, S. Karpov, S. Bondar, E. Ivanov, E. Katkova, A. Perkov, and V. Sasyuk. The first light of Mini-MegaTORTORA wide-field monitoring system. 24: 100–108, 2014. URL <http://arxiv.org/abs/1411.2552>.
- Christian Bizouard, Olivier Becker, and Jean Yves Richard. Combined solution C04 for Earth Rotation Parameters consistent with International Terrestrial Reference Frame 2014. Technical Report January, Observatoire de Paris, 2017.
- R.L. L Cook and K.E. E Torrance. A Reflectance Model for Computer Graphics. *ACM SIGGRAPH Computer Graphics*, 15(3):307–316, jan 1982. ISSN 00978930. . URL <https://doi.org/10.1145/357290.357293>.
- John Crassidis and F. Landis Markley. Unscented Filtering for Spacecraft Attitude Estimation. *Journal of Guidance, Control, and Dynamics*, 26(4):536–542, 2003. ISSN 0731-5090. .
- J.C. Dolado-Perez, Romain Di Costanzo, and Bruno Revelin. Introducing Medee – a New Orbital Debris. *Proceedings of the 6th European Conference on Space Debris*, 2012 (April), 2012.
- Michael A. Earl and Gregg A. Wade. Observations of the Spin-Period Variations of Inactive Box-Wing Geosynchronous Satellites. *Journal of Spacecraft and Rockets*, 52(3):

- 968–977, apr 2015. ISSN 0022-4650. . URL <http://arc.aiaa.org/doi/10.2514/1.A33077>.
- J Emmert, J Byers, H Warren, and A Segerman. Propagation of Forecast Errors from the Sun to LEO Trajectories: How Does Drag Uncertainty Affect Conjunction Frequency? In S Ryan, editor, *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E48, sep 2014.
- Executive Office of the President. Presidential directive on national space policy, 1988. URL <http://www.hq.nasa.gov/office/pao/History/policy88.html>.
- J. L. Farrell. Attitude determination by kalman filtering. *Automatica*, 6(3):419–430, 1970. ISSN 00051098. .
- V L Foreman, O Ladislav de Weck, and A Siddiqi. Large Satellite Constellation Orbital Debris Impacts: Case Studies of OneWeb and SpaceX Proposals. In *Proc. of the American Institute of Aeronautics and Astronautics Space Forum*, number 203999, Orlando, FL, 2017. ISBN 9781624104831. .
- Jason L. Forshaw, Guglielmo S. Aglietti, Nimal Navarathinam, Haval Kadhém, Thierry Salmon, Aurélien Pisseloup, Eric Joffre, Thomas Chabot, Ingo Retat, Robert Axthelm, Simon Barraclough, Andrew Ratcliffe, Cesar Bernal, François Chaumette, Alexandre Pollini, and Willem H. Steyn. RemoveDEBRIS: An in-orbit active debris removal demonstration mission. *Acta Astronautica*, 127(2016):448–463, 2016. ISSN 00945765. . URL <http://dx.doi.org/10.1016/j.actaastro.2016.06.018>.
- Alex Friedman, Siwei Fan, Carolin Frueh, and Thomas Schildknecht. Observability of Light Curve Shape Inversion Based on Optical Data. 2019.
- R H Giese. Attitude Determination from Specular and Diffuse Reflection by Cylindrical Artificial Satellites. Technical report, Smithsonian Institution Astrophysical Observatory, 1963.
- R. H. Giese. Investigations into a Method for the Optical Estimation of the Rotation Axis of Artificial Satellites. Technical report, 1967.
- J. Goguen, J. Veverka, J. L. Elliot, and C. Church. The lightcurve and rotation period of asteroid 139 Juewa. *Icarus*, 29(1):137–142, 1976. ISSN 10902643. . URL <http://www.sciencedirect.com/science/article/pii/0019103576901081>.
- Joseph A. Haimerl and G. P. Fonder. Space Fence System Overview. *Proc. of the Advanced Maui Optical and Space Surveillance Technology Conference*, 2015.
- J Hamilton. Development of the Space Debris Sensor. In *Proc. 7th European Conference on Space Debris*, Darmstadt, Germany, 2017. ESA Space Debris Office.

- Joanna C Hinks, Richard Linares, and John L Crassidis. Attitude Observability from Light Curve Measurements. In *AIAA Guidance, Navigation, and Control (GNC) Conference*, number August, pages 1–10, Boston, MA, 2013. ISBN 9781624102240. .
- Felix R. Hoots and Ronald L. Roehrich. Spacetrack Report No. 3–Models for Propagation of NORAD Elements Sets. *Spacetrack Report*, 3(3):1–91, 1980.
- IADC. IADC-02-01 Rev1 - IADC Space Debris Mitigation Guidelines. (22):10, 2007.
- International Telecommunications Union. ITU releases 2018 global and regional ICT estimates, 2018. URL <https://www.itu.int/en/mediacentre/Pages/2018-PR40.aspx>.
- Nicholas L Johnson, E. Stansbery, J. C. Liou, M. Horstman, C L Stokely, and D. Whitlock. The Characteristics and Consequences of the Break-up of the Fengyun-1C Spacecraft. pages 1–9, 2008.
- M Kaasalainen and J Torppa. Optimization Methods for Asteroid Lightcurve Inversion: I. Shape Determination. *Icarus*, 153(1):24–36, 2001. ISSN 0019-1035. . URL <http://www.sciencedirect.com/science/article/pii/S0019103501966734>.
- M. Kaasalainen, L. Lamberg, and K. Lumme. Interpretation of lightcurve of atmosphereless bodies - Part2. *Astronomy and Astrophysics Review*, 259:333–340, 1992a.
- M Kaasalainen, L Lamberg, K Lumme, and E Bowell. Interpretation of lightcurves of atmosphereless bodies. I - General theory and new inversion schemes. *Astronomy and Astrophysics*, 259:318–332, jun 1992b. ISSN 0004-6361. URL <https://ui.adsabs.harvard.edu/abs/1992A&A...259..318K>.
- M Kaasalainen, J Torppa, and K Muinonen. Optimization Methods for Asteroid Lightcurve Inversion: II. The Complete Inverse Problem. *Icarus*, 153(1):37–51, 2001. ISSN 0019-1035. . URL <http://www.sciencedirect.com/science/article/pii/S0019103501966746>.
- R. Kanzler, J. Silha, T. Schildknecht, B. Fritsche, T. Lips, and H. Krag. Space Debris Attitude Simulation - IOTA (In-Orbit Tumbling Analysis). *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, held in Wailea, Maui, Hawaii, September 15-18, 2015*, 2015.
- S. Karpov, G. Beskin, A. Biryukov, S. Bondar, E. Ivanov, E. Katkova, A. Perkov, and V. Sasyuk. Mini-MegaTORTORA Wide-Field Monitoring System with Sub-Second Temporal Resolution: First Year of Operation. *Revista Mexicana de Astronomía y Astrofísica*, (48), 2016a.
- S Karpov, E Katkova, G Beskin, A. Biryukov, S Bondar, E Davydov, E Ivanov, and A Perkov. MASSIVE PHOTOMETRY OF LOW-ALTITUDE ARTIFICIAL

- SATELLITES ON MINI-MEGA-TORTORA. *Revista Mexicana de Astronomía y Astrofísica*, 48:112–113, 2016b. ISSN 0185-1101.
- Patrik Käräräng, Ronny Kanzler, Bent Fritsche, and Stijn Lemmens. First Results of IOTA (in-Orbit Tumbling Analysis). In *6th International Conference on Astrodynamics Tools and Techniques*, Darmstadt, Germany, 2016. URL <https://indico.esa.int/indico/event/111/session/31/contribution/127/material/paper/0.pdf>.
- P.H. Krisko. The New NASA Orbital Debris Engineering Model ORDEM 3.0 P. H. Krisko. In *AIAA/AAS Astrodynamics Specialist Conference*, 2014. .
- Marc L. Kutner. *Astronomy: A Physical Perspective*. Cambridge University Press, Cambridge, 2 edition, jul 2003. ISBN 9780511802195. . URL <https://www.cambridge.org/core/product/identifier/9780511802195/type/book><https://www.cambridge.org/core/books/astronomy-a-physical-perspective/9F769B29B99992A4BC90356484394726>.
- S M Lederer, E G Stansbery, H M Cowardin, P Hickson, L F Pace, K J Abercromby, P W Kervin, and R J Alliss. The NASA Meter Class Autonomous Telescope: Ascension Island. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2013. URL <https://ui.adsabs.harvard.edu/abs/2013amos.confE..82L>.
- Ej Lefferts, Fl Markley, and Md Shuster. Kalman filtering for spacecraft attitude estimation, 1982. ISSN 0731-5090. URL <http://arc.aiaa.org/doi/pdf/10.2514/3.56190>.
- Hugh G. Lewis, Graham G. Swinerd, Rebecca J. Newland, and Arrun Saunders. Active removal study for on-orbit debris using damage. *European Space Agency, (Special Publication) ESA SP*, 672 SP(April), sep 2009. ISSN 03796566. URL <https://eprints.soton.ac.uk/68969/>.
- Hugh G. Lewis, Arrun Saunders, Graham Swinerd, and Rebecca J. Newland. Effect of thermospheric contraction on remediation of the near-Earth space debris environment. *Journal of Geophysical Research: Space Physics*, 116(8):1–10, 2011. ISSN 21699402. .
- Hugh G. Lewis, Adam E. White, Richard Crowther, and Hedley Stokes. Synergy of debris mitigation and removal. *Acta Astronautica*, 81(1):62–68, 2012. ISSN 00945765. . URL <http://dx.doi.org/10.1016/j.actaastro.2012.06.012>.
- Aleksander A. Lidtke, Hugh G. Lewis, Roberto Armellin, and Hodei Urrutxua. Considering the collision probability of Active Debris Removal missions. *Acta Astronautica*, 131(November 2016):10–17, 2017. ISSN 00945765. . URL <http://dx.doi.org/10.1016/j.actaastro.2016.11.012>.

- Richard Linares, Moriba K. Jah, John L. Crassidis, and Christopher K. Nebelecky. Space Object Shape Characterization and Tracking Using Light Curve and Angles Data. *Journal of Guidance, Control, and Dynamics*, 37(1):13–25, 2014. ISSN 0731-5090. . URL <http://arc.aiaa.org/doi/abs/10.2514/1.62986>.
- J. C. Liou. An active debris removal parametric study for LEO environment remediation. *Advances in Space Research*, 47(11):1865–1876, 2011. ISSN 02731177. . URL <http://dx.doi.org/10.1016/j.asr.2011.02.003>.
- J. C. Liou, D. T. Hall, P. H. Krisko, and J. N. Opiela. LEGEND - A three-dimensional LEO-to-GEO debris evolutionary model. *Advances in Space Research*, 34(5):981–986, 2004. ISSN 02731177. .
- J. C. Liou, N. L. Johnson, and N. M. Hill. Controlling the growth of future LEO debris populations with active debris removal. *Acta Astronautica*, 66(5-6):648–653, 2010. ISSN 00945765. . URL <http://dx.doi.org/10.1016/j.actaastro.2009.08.005>.
- J.-C. Liou, M. Kieffer, A. Drew, and A. Sweet. Project Review. *Orbital Debris Quarterly News*, 24(1):4–8, feb 2020.
- J.-C. C. Liou and N. L. Johnson. Collision activities in the future orbital debris environment. *Advances in Space Research*, 38(9):2102–2106, 2006. ISSN 02731177. .
- J.-C. C. Liou and N. L. Johnson. Instability of the present LEO satellite populations. *Advances in Space Research*, 41(7):1046–1053, 2008. ISSN 02731177. . URL <http://www.sciencedirect.com/science/article/pii/S0273117707004097>.
- J.-C. C. Liou and Nicholas L. Johnson. A sensitivity study of the effectiveness of active debris removal in LEO. *Acta Astronautica*, 64(2):236–243, 2009. ISSN 0094-5765. . URL <http://www.sciencedirect.com/science/article/pii/S0094576508002634>.
- JC Liou. USA Space Debris Environment, Operations, and Research Updates. Number February 2017, 2017.
- Charlotte Lücking, Camilla Colombo, and Colin R. McInnes. A passive satellite deorbiting strategy for medium earth orbit using solar radiation pressure and the J2 effect. *Acta Astronautica*, 77:197–206, 2012. ISSN 0094-5765. . URL <http://www.sciencedirect.com/science/article/pii/S0094576512000987>.
- Anthony Mallama. The Brightness of VisorSat-Design Starlink Satellites. 2021.
- Rosana Montes and Carlos Ureña. An Overview of BRDF Models. *Technical Report*, University of Granada, pages 1–26, 2012.
- NASA. Goldstone Solar System Radar Unique Capabilities of GSSR, 1996. URL <https://gssr.jpl.nasa.gov/gssrrep.pdf>.

- NASA. Handbook for Limiting Orbital Debris. Technical report, Washington, DC, 2008. URL https://explorers.larc.nasa.gov/APMIDEX2016/M0/pdf_files/NHBK871914.pdf.
- P. Papushev, Yu Karavaev, and M. Mishina. Investigations of the evolution of optical characteristics and dynamics of proper rotation of uncontrolled geostationary artificial satellites. *Advances in Space Research*, 43(9):1416–1422, 2009. ISSN 02731177. . URL <http://dx.doi.org/10.1016/j.asr.2009.02.007>.
- G rard Petit and Brian Luzum. IERS Conventions (2010). Technical report, International Earth Rotation and Reference Systems Service, 2010. URL <http://www.iers.org/TN36/>.
- C.R. R Phipps, G. Albrecht, H. Friedman, D. Gavel, E.V. V George, J. Murray, C. Ho, W. Priedhorsky, M.M. M Michaelis, and J.P. P Reilly. ORION: Clearing near-Earth space debris using a 20-kW, 530-nm, Earth-based, repetitively pulsed laser. *Laser and Particle Beams*, 14(01):1, 1996. ISSN 0263-0346. .
- Fabrizio Piergentili, Fabio Santoni, and Patrick Seitzer. Attitude Determination of Orbiting Objects from Lightcurve Measurements. *IEEE Transactions on Aerospace and Electronic Systems*, 53(1):81–90, 2017. ISSN 00189251. .
- N Praly, M Hillion, C Bonnal, J Laurent-varin, and N Petit. Acta Astronautica Study on the eddy current damping of the spin dynamics of space debris from the Ariane launcher upper stages. *Acta Astronautica*, 76:145–153, 2012. ISSN 0094-5765. . URL <http://dx.doi.org/10.1016/j.actaastro.2012.03.004>.
- Jonas Radtke, Sven Mueller, Volker Schaus, and Enrico Stoll. Luca2 - an Enhanced Long-Term Utility for Collision Analysis. In *Proc. 7th European Conference on Space Debris*, number April, 2017.
- Tyler G.R. Reid, Andrew M. Neish, Todd F. Walter, and Per K. Enge. Leveraging commercial broadband LEO constellations for navigation. *29th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2016*, 4(October):2300–2314, 2016. .
- A Rossi, E M Alessi, G B Valsecchi, H G Lewis, C Colombo, L Anselmo, C Pardini, F Deleflie, and K Merz. MEO Dynamics and GNSS Disposal Strategies. In *Proc. 7th European Conference on Space Debris*, number April, Darmstadt, Germany, 2017. ESA Space Debris Office.
- G. Sansone. *Orthogonal Functions*. Courier Corporation, 1959. ISBN 0486667308, 9780486667300.
- Fabio Santoni, Emiliano Cordelli, and Fabrizio Piergentili. Determination of Disposed-Upper-Stage Attitude Motion by Ground-Based Optical Observations.

- Journal of Spacecraft and Rockets*, 50(3):701–708, 2013. ISSN 0022-4650. . URL <http://arc.aiaa.org/doi/abs/10.2514/1.A32372>.
- Thomas Schildknecht. Optical surveys for space debris. *The Astronomy and Astrophysics Review*, 14(1):41–111, 2007. ISSN 1432-0754. . URL <https://doi.org/10.1007/s00159-006-0003-9>.
- P. Seitzer, R. Smith, J. Africano, K. Jorgensen, E. Stansbery, and D. Monet. MODEST observations of space debris at geosynchronous orbit. *Advances in Space Research*, 34(5):1139–1142, 2004. ISSN 02731177. .
- Minghe Shan, Jian Guo, and Eberhard Gill. Review and comparison of active space debris capturing and removal methods. *Progress in Aerospace Sciences*, 80:18–32, 2016. ISSN 03760421. . URL <http://dx.doi.org/10.1016/j.paerosci.2015.11.001>.
- D. E. Smith. An Approximate Method of Estimating the Reflection Characteristics of a Spherical Satellite. In *Proceedings of University of Miami Symposium on Optical Properties of Orbiting Satellites*, pages 139 – 156, 1967.
- R. Sridharan and Antonio F. Pensa. U.S. Space Surveillance Network Capabilities. *Image Intensifiers and Applications; and Characteristics and Consequences of Space Debris and Near-Earth Objects*, 3434(November 1998):88–100, 1998. .
- E. M. Standish and J. G. Williams. Orbital Ephemerides of the Sun, Moon, and Planets. In *Explanatory supplement to the astronomical almanac*, pages 1–33. 2003. ISBN 1891389459.
- R.~F. Stellingwerf. Period determination using phase dispersion minimization. *The Astrophysical Journal*, 224:953–960, sep 1978. .
- C L Stokely, J L Foster, E G Stansbery, J R Benbrook, and Q Juarez. Haystack and HAX radar measurements of the orbital debris environment; 2003. *Jsc-62815*, (November), 2006.
- Johanna Torppa, Mikko Kaasalainen, Tadeusz Michałowski, Tomasz Kwiatkowski, Agnieszka Kryszczyńska, Peter Denchev, and Richard Kowalski. Shapes and rotational properties of thirty asteroids from photometric data. *Icarus*, 164(2): 346–383, 2003. ISSN 0019-1035. . URL <http://www.sciencedirect.com/science/article/pii/S0019103503001465>.
- J. Tregloan-Reed, A. Otarola, E. Ortiz, V. Molina, J. Anais, R. González, J. P. Colque, and E. Unda-Sanzana. First observations and magnitude measurement of Starlink’s Darksat. *A&A*, 637:L1, 2020. . URL <https://doi.org/10.1051/0004-6361/202037958>.
- United Nations Office For Outer Space Affairs. Space Debris Mitigation Guidelines of the Committee on the Peaceful Uses of Outer Space. *United Nations Publication*,

- page 1, 2010. ISSN 02784017. . URL
http://www.unoosa.org/pdf/publications/st_space_49E.pdf.
- David Vallado, Paul Crawford, Ricahrd Hujsak, and T.S. Kelso. *Revisiting Spacetrack Report 3*. 2006. ISBN 978-1-62410-048-2. . URL
<http://arc.aiaa.org/doi/10.2514/6.2006-6753>.
- R. Walker, C. E. Martin, P. H. Stokes, J. E. Wilkinson, and H. Klinkrad. Analysis of the effectiveness of space debris mitigation measures using the DELTA model. *Advances in Space Research*, 28(9):1437–1445, 2001. ISSN 0273-1177. . URL
<http://www.sciencedirect.com/science/article/pii/S0273117701004458>.
- Ting Wang. Analysis of debris from the collision of the Cosmos 2251 and the Iridium 33 satellites. *Science and Global Security*, 18(2):87–118, 2010. ISSN 08929882. .
- Charles J Wetterer and Moriba K Jah. Attitude Determination from Light Curves. *Journal of Guidance, Control, and Dynamics*, 32(5):1648–1651, sep 2009. ISSN 0731-5090. . URL <http://arc.aiaa.org/doi/abs/10.2514/1.44254>.
- Adam E. White and Hugh G. Lewis. An adaptive strategy for active debris removal. *Advances in Space Research*, 53(8):1195–1206, 2014a. ISSN 02731177. . URL
<http://dx.doi.org/10.1016/j.asr.2014.01.021>.
- Adam E. White and Hugh G. Lewis. The many futures of active debris removal. *Acta Astronautica*, 95(1):189–197, 2014b. ISSN 00945765. . URL
<http://dx.doi.org/10.1016/j.actaastro.2013.11.009>.
- J G Williams. The Determination of the Orientation of a Tumbling Cylinder from the Shape of the Light Curve. Technical report, 1967.
- V. Williams. Location of the rotation axis of a tumbling cylindrical earth satellite by using visual observations. Part I: Theory. *Planetary and Space Science*, 27(6):885–890, jun 1979a. ISSN 00320633. . URL
<http://www.sciencedirect.com/science/article/pii/003206337990014X>
<http://www.sciencedirect.com/science/article/pii/0032063379900151>
<https://linkinghub.elsevier.com/retrieve/pii/003206337990014X>.
- V. Williams. Location of the Rotation Axis of a Tumbling Cylindrical Earth Satellites by Using Visual Observations Part II: Results. *Planetary and Space Science*, 27(6): 891–894, jun 1979b. ISSN 00320633. . URL
<http://www.sciencedirect.com/science/article/pii/0032063379900151>
<http://www.sciencedirect.com/science/article/pii/003206337990014X>
<https://linkinghub.elsevier.com/retrieve/pii/003206337990014X>.