MULTICOLOUR OPTICAL PHOTOMETRY OF ACTIVE GEOSTATIONARY SATELLITES USING A SMALL APERTURE TELESCOPE

PHOTOMÉTRIE OPTIQUE MULTI-COULEUR DE SATELLITES GÉOSTATIONNAIRES ACTIFS AVEC UN TÉLESCOPE À PETITE APERTURE

A Thesis Submitted

to the Division of Graduate Studies of the Royal Military College of Canada

by

Andrew Robert Jolley, B.Tech(Aero Eng)
Flight Lieutenant

In Partial Fulfillment of the Requirements for the Degree of

Masters of Science in Physics

May, 2014

To my wife Jodie and daughter Tallulah.

Acknowledgements

I would like to thank my supervisor, Dr. Gregg Wade for his enthusiastic and proactive support of my work. Without his thorough attention to detail and constructive criticisms the quality of the work presented in this thesis and the conclusions that were reached as a result would be significantly poorer than they are.

A great deal of gratitude is owed to Major Don Bedard for introducing me to the Royal Military College of Canada (RMCC) Space Surveillance Research and Analysis Laboratory (SSRAL), helping to troubleshoot the abundance of problems encountered during the experiment, and offering his expertise and opinion along the way. Without Don's encouragement this thesis would never have been started, and the innumerable discussions had with him were instrumental in keeping it on track. His patience, friendship and generous assistance are greatly appreciated.

I would also like to thank Dave Watts and Bryce Bennett for their practical support in repairing and maintaining the equipment that was used. Without their support, obtaining data would have been significantly more challenging.

Finally, it would be remiss of me not to acknowledge the strong support provided by my wife, Jodie, over the last two years. In many ways completing this thesis was incompatible with engaging in family life whilst living across the world from our home country. Jodie's unwavering encouragement and support, despite the challenges presented, have made this project possible, and for that I am deeply grateful.

Abstract

Jolley, Andrew Robert. M.Sc. Royal Military College of Canada, May, 2014. *Multi-colour Optical Photometry of Active Geostationary Satellites Using a Small Aperture Telescope.* Supervised by Dr. Gregg A. Wade.

Although broadband photometry has been used to infer information about artificial satellites since soon after the launch of Sputnik 1, to date there has not been a coherent progression of understanding regarding the information that can be obtained, and the most effective methods of doing so. Many individual studies, and some long ongoing experiments, have used costly metre-class telescopes to obtain data despite other experiments demonstrating that much more flexible and affordable small aperture telescopes may be suitable for the task. In addition, due to the highly time consuming and weather dependent nature of obtaining photometric observations, many studies have suffered from limited data sets or relied upon simulations to support their claims. Therefore, an experiment was conducted with the aim of determining the utility of small aperture telescopes for conducting broadband photometry of satellites. A small aperture telescope (Celestron CG-14, 35.6 cm aperture) was used to gain multiple night-long, high temporal resolution data sets of a limited number of geostationary satellites (Anik F1, Anik F1R, Anik G1, Echostar 17, Galaxy 11, and Intelsat 805). The results of the experiment demonstrate that small aperture telescopes can be highly effective for obtaining accurate photometric measurements of geostationary satellites. In the process some novel discoveries were made that cast doubt on the efficacy of some of the previous approaches to obtaining and analysing photometric data. For example, it was discovered that satellite lightcurves vary to a significantly greater degree than has been generally recognised, and significant features such as oscillations were detected in some lightcurves with frequencies higher than would be detectable at temporal resolutions that have typically been employed

in the past. Furthermore, these results suggest that high temporal resolution, simultaneous, multicolour observations may have significant utility in characterising or identifying satellites in some cases.

Résumé

Jolley, Andrew Robert. M.Sc. Collège militaire royal du Canada, May, 2014. *Photométrie optique multi-couleur de satellites géostationnaires actifs avec un télescope à petite aperture.* Thèse dirigée par M. Gregg A. Wade, Ph.D.

Bien que la photométrie à large bande ait été utilisée pour obtenir de l'information sur les satellites artificiels depuis peu de temps après le lancement de Sputnik 1, il n'y a pas eu jusqu'à maintenant une progression cohérente dans la compréhension de l'information pouvant en être déduite, ainsi que sur les méthodologies les plus efficaces. Plusieurs études individuelles et quelques longues expériences toujours en cours ont fait l'usage de télescopes dispendieux de la catégorie des mètres pour obtenir des données malgré que d'autres expériences aient démontré que des télescopes à petite aperture beaucoup plus flexibles et abordables peuvent être convenables pour la tâche. De plus, étant donné la variabilité de l'état du ciel et le coût élevé en temps pour obtenir des observations photométriques, plusieurs études reposaient sur des ensembles de données limités ou bien sur des simulations afin d'appuyer leurs allégations. Par conséquent, une expérience a été conduite ayant comme objectif d'évaluer l'utilité des télescopes à petite aperture effectuant de la photométrie à large bande de satellites. Un télescope à petite aperture (Celestron CG-14, aperture de 35.6 cm) a été utilisé pour obtenir plusieurs ensembles de données à haute résolution temporelle pour la durée des nuits, à partir d'un nombre limité de satellites géostationnaires (Anik F1, Anik F1R, Anik G1, Echostar 17, Galaxy 11 et Intelsat 805). Les résultats de cette expérience démontrent que les télescopes à petite aperture peuvent être hautement efficaces pour obtenir des relevés photométriques exactes de satellites géostationnaires. Tout au long du processus, quelques nouvelles découvertes ont été faites mettant en doute l'efficacité de certaines approches antérieures pour obtenir et analyser des données photométriques. Par exemple, il a été trouvé que le degré de variation

des courbes de lumière de satellites est beaucoup plus important que ce qui avait généralement été reconnu jusqu'à maintenant. De plus, des caractéristiques significatives telles que des oscillations ont été détectées dans certaines courbes de lumière avec des fréquences plus grandes que ce qui serait détectable avec les résolutions temporelles typiquement utilisées dans le passé. En outre, ces résultats indiquent que des observations à haute résolution temporelle, simultanées et multi-couleur peuvent avoir une utilité significative pour la caractérisation ou l'identification de satellites dans certains cas.

Mots clés : Photométrie, satellites géostationnaires, connaissance de la situation spatiale, télescopes à petite aperture

Table of Contents

List of Tables

List of Figures

xiii

Acronyms

ADC   Analogue to Digital Converter

ADU   Analogue to Digital Unit

AFRL   Air Force Research Laboratory

ANDE   Atmospheric Neutral Density Experiment

BRDF   Bi-directional Reflectance Distribution Function

CCD   Charge Coupled Device

DRDC   Defence Research and Development Canada

FWHM   Full Width at Half Maximum

GEO   Geostationary Orbit

GPS   Global Positioning System

HEO   Highly Elliptical Orbit

IDCSP   Initial Defense Communication Satellite Program

JSpOC   Joint Space Operations Center

LEO   Low Earth Orbit

MEO   Medium Earth Orbit

MLI   Multi Layer Insulation

PAGEOS   Passive Geodetic Satellite

PSF   Point Spread Function

QE	Quantum Efficiency
RMCC	Royal Military College of Canada
RSO	Resident Space Object
SED	Spectral Energy Distribution
SILC	Space object identification In Living Colour
SNR	Signal to Noise Ratio
SOI	Space Object Identification
SSA	Space Situational Awareness
SSRAL	Space Surveillance Research and Analysis Laboratory
TJPV	Triple Junction Photovoltaic Cell
USNO	United States Naval Observatory
UTC	Universal Time Coordinated

Chapter 1

Introduction

1.1 Background

Prior to the launch of Sputnik 1 in 1957 the near-Earth space environment was considered pristine: a satellite could be placed into any desired orbit with virtually no risk of it colliding with another object. Today, only six decades later, over 15 000 objects in low Earth orbit greater than ten centimeters in diameter are being continuously tracked, with greater than 500 000 objects smaller than ten centimeters estimated to be in orbit [1]. With such a large population of objects in orbit, and evidence that the debris population in some regions is either unstable or in the early phase of collisional cascading [2], the task of enhancing Space Situational Awareness (SSA) is becoming ever more urgent. This urgency is amplified by the increasing reliance on space assets by governments, militaries, companies and civilians throughout the world. Characterising Resident Space Objects (RSOs) in terms of their size, mass, surface material composition, or area to mass ratio can assist with efforts to propagate the locations of tracked objects into the future, in turn increasing the accuracy of conjunction analyses and estimates of the time until objects deorbit. Potential has been shown for the ability of photometric analysis to provide such information about RSOs [3–5]. But very useful information can be obtained even without determining such specific satellite characteristics. Substantial effort has been invested into satellite discrimination and health status determination based solely on measurements of the intensity of reflected light [6–8]. In terms of operational satellites, SSA can refer to information regarding the status and motion of the satellite. If an event occurs that causes an operational satellite to behave in an unexpected way, then determining the nature and cause of the event can be of critical importance. From a commercial perspective this information can potentially be used to help solve problems with the operation of a satellite; this can be worth many millions of dollars to a company's

bottom line. From a military intelligence perspective, attributing causes to problems with satellites can help to identify possible interference by an adversary. There are very few ways of obtaining health information from a non-communicative satellite, but photometry has shown some promise in doing so [9]. Also, characterising an unknown object can give insight into its possible mission, power consumption, or other valuable information. Much work has been done that demonstrates the potential for such information to be derived using photometric analysis.

1.1.1 Photometry

Photometry is the science of the measurement of photon flux. Typically, in astronomical photometry, broad regions of the incident spectrum of light are isolated using transmission filters. The typical bandpasses of such filters, for example in the widely used Johnson-Cousins UBVRI system illustrated in figure 1.1, are approximately 50 to 200 nm Full Width at Half Maximum (FWHM). This is known as broadband photometry.[1]
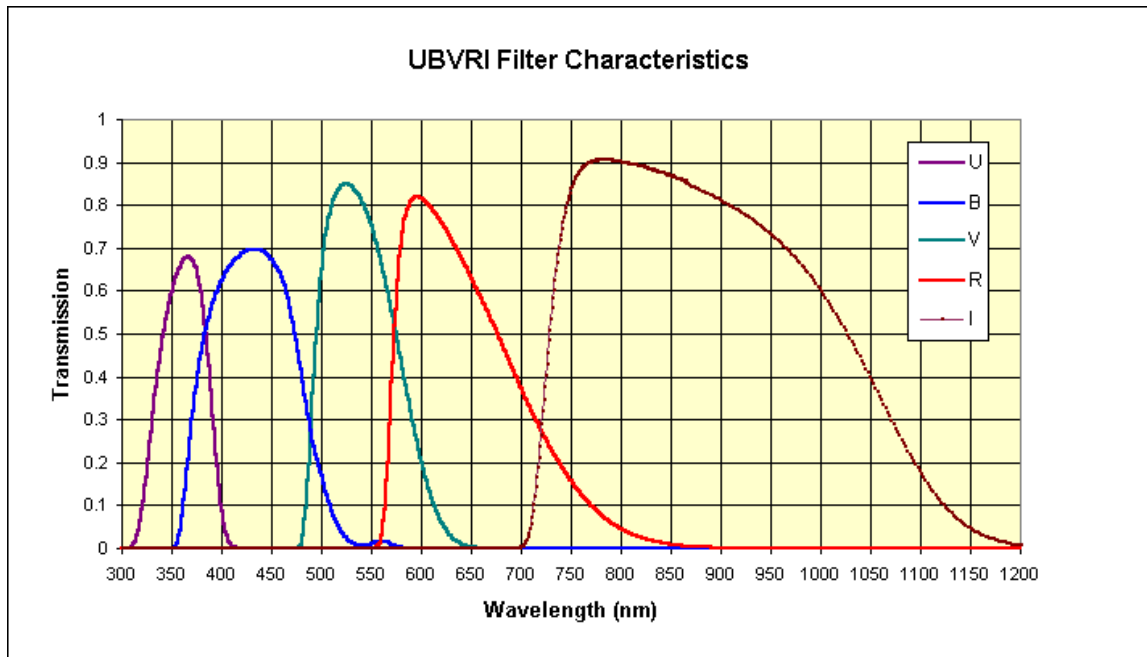
Figure 1.1: Johnson-Cousins Photometric System Passbands [10]

---

[1]When the regions of the incident spectrum are categorised according to extremely narrow wavelength bands, their measurement is known as spectroscopy.

1.1.2   Specular and Diffuse Reflection

When incident light interacts with an object's surface materials, a proportion of it is reflected, whilst the remainder is absorbed or transmitted. The terms *specular* and *diffuse* are commonly used to qualitatively describe the way in which the reflected light exits the material's surface. A specular reflection is one where the incident light vector, the material's surface normal vector, and the exitant light vector exist on the same plane; and the angle between the incident light vector and the surface normal vector is equal in magnitude to the angle between the exitant light vector and the surface normal vector. A perfectly specular reflection is one where 100 percent of the reflected light is specular, although not all of the incident light needs to be reflected; some may be absorbed or transmitted. In contrast, a diffuse reflection is characterised by exitant light that is distributed across a broad range of directions. A perfectly diffuse reflection is one characterised by a lambertian distribution; where the reflected light intensity is isotropic, or equal in all directions in the hemisphere above the reflecting surface [11]. In practice it is almost impossible for a given surface to be a perfectly specular or perfectly diffuse reflector; some of the incident light will be reflected specularly and some will be reflected in a somewhat diffuse (but probably not perfectly lambertian) manner.

1.1.3   The Sun-Satellite-Sensor Geometry

Photometry, as it applies to SSA, involves measuring the photon flux due to sunlight reflected off an RSO towards a sensor. The ratio of measured reflected light to incident light for a given surface depends upon the azimuth angle and zenith angle of both the incident light source and the observer. This geometry is depicted in figure 1.2. The orange cube represents the satellite, θ_{sun} and θ_{obs} represent the azimuth angles to the sun and the observer, respectively, and ϕ_{sun} and ϕ_{obs} represent the zenith angles to the sun and observer, respectively. The function that describes the reflected intensity as a function of the incident angles θ_{sun} and ϕ_{sun} and exitant angles θ_{obs} and ϕ_{obs} is known as the Bi-directional Reflectance Distribution Function (BRDF) [11].
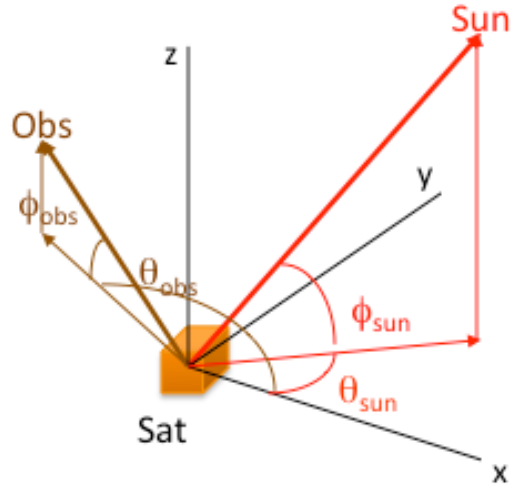
Figure 1.2: Geometry of incident and reflected light in the satellite body reference coordinate system [12].

When a material reflects light, the proportion of the incident light that is reflected is dependent upon its wavelength and the material in question. How diffuse or specular the reflection is (and therefore how much of it is reflected towards a detector) also depends upon the characteristics of the reflecting material. As a result, if one can measure the intensity of the reflected light and its variation with wavelength and illumination/observation geometry, information may be gleaned about the reflecting surface. Such measurements are often acquired repeatedly as a satellite moves across the sky, with the brightness plotted as a function of the sun-satellite-observer angle, known as the *phase angle*. There are difficulties associated with the use of phase angle in analysing satellite brightness variations because there is an inherent loss of information when combining the four angles that describe the illumination/observation geometry into a single angle. As a result, several different definitions of phase angle exist, each with their own advantages and disadvantages. A discussion on the various different definitions of phase angle that have been used to date can be found at Appendix B.

1.1.4 Magnitude Scales

The basic form of the intensity scale most commonly used in visible light photometry has remained essentially unchanged since it was first proposed by William Pogson in the mid nineteenth century when he noticed that the rather arbitrary system of assigning stars to one of six classes, where the first class was the brightest and the sixth class was the faintest, roughly followed a logarithmic scale. Pogson's scale notes the brightness of objects by assigning them *magnitude* values, defined as follows: [13]

$$m = -2.5 \log(f) + k \tag{1.1}$$

In the preceeding equation, m is the magnitude, f is the measured photon flux in photons per second, and k is a constant added to standardise measurements such that different systems will measure the same magnitude for the same object. k is known as the *zero point*, and it varies from one equipment setup to another, and with changing atmospheric conditions. The minus sign is required to maintain the original convention that was devised such that brighter objects are assigned a lower value and fainter objects are assigned a higher value. Pogson's logarithmic magnitude scale has survived quite successfully for passbands in the visible light range because it allows a very wide range of light intensities to be recorded in a simple fashion.

1.1.5 The Application of Photometry to SSA

A significant amount of research has been conducted to determine the utility of photometry in deriving information about artificial Earth orbiting satellites. In general, most efforts have focussed on either satellite discrimination, which is the process of differentiating between two or more (usually closely spaced) satellites [6, 14, 15]; or satellite characterisation, which is the process of determining specific characteristics of a given satellite such as its surface material charateristics and composition, solar array size, probable mission, health status, and bus type [4, 7, 16–18]. A satellite's bus type refers to its structural design. It is analogous to a specific satellite model. In practice, using photometry to obtain useful and accurate information about satellites

is not necessarily a straightforward process because the spectral energy recorded by a detector is highly sensitive to a range of variables that are difficult to measure. For example, the reflected light intensity measured from a highly specularly reflecting surface can be more than one hundred times brighter when the specular reflection is directed towards the sensor than when it is not [19]. However, knowing the precise orientation of all of the significant specularly reflecting surfaces on a satellite, and thus whether the specular reflection is directed towards the sensor, would be extremely challenging. Developing a procedure to conduct photometry on artificial Earth orbiting satellites that is both inexpensive and flexible, but capable of delivering substantial information, thus requires a detailed understanding of the variables that affect the way light is reflected off materials, how it is affected by the atmosphere, how it is measured when it arrives at a sensor, and how it is then interpreted. Importantly, because various simplifying assumptions will almost certainly need to be made when conducting photometry of an unknown satellite, it is also necessary to understand *the degree* to which each variable affects photometric measurements to provide confidence that the assumptions that are made are reasonable.

The space environment is unique in that it is host to a multitude of highly advanced capabilities that provide critically important contributions to the wellbeing of modern global society. Space capabilities are now fundamental to a wide range of fields including precision navigation and timing, Earth observation for land management, communications, search and rescue, weather forecasting, and intelligence gathering, among others. But the space environment is also unique in the difficulty of obtaining and maintaining information regarding the objects within it and what they are doing. As part of the effort to do so, both small and large aperture telescopes have been used for many years to measure the brightness of Earth-orbiting satellites. The analysis of the observed variation of brightness versus phase angle, known as a lightcurve, has long formed the foundation of attempts to use photometry to enhance SSA. Any reliable method of substantially increasing SSA would certainly be highly valued, particularly if that method is inexpensive and rapidly deployable. In practice, it is likely that a combination of different methods will be used to enhance a

nation's SSA, as is the case currently, including optical imaging, radar surveillance, and spectroscopy. However, ground-based photometry may have the potential to be a significant contributor to the overall intelligence picture: the extent of its possible contribution is still not entirely clear.

1.2 Aim of the Masters Thesis

The primary aim of this thesis was to develop a procedure that uses a small aperture telescope to conduct multicolour photometric analysis of geostationary satellites; and to determine the utility of the procedure with regard to the ability to consistently and accurately record apparent magnitudes. The secondary aim was to analyse the data that were obtained, in the context of previous work conducted in this field, to determine whether further contributions towards photometric analysis of geostationary satellites could be made by this work. This experiment specifically required that a small aperture telescope be used because their relative portability and low cost imply that the method used in this thesis could potentially represent a flexible and highly affordable means of enhancing SSA.

1.3 Thesis Outline

The following chapter provides a summary of the publicly available literature relating to the photometric analysis of operational satellites. Particular attention is paid to geostationary satellites and advances relating to the information that photometry can provide to characterise a satellite. Chapter 3 describes the hardware and software used to conduct the experiment. The data collection procedure is detailed, including limitations of the experiment and challenges that were faced in obtaining the data. The chapter concludes with the image reduction, data extraction and data manipulation procedures. Chapter 4 introduces typical results that were achieved in the analysis of the data.

Understanding the potential accuracy of the data obtained by this method is critical to achieving the aim of determining the feasibility of using a small aperture

telescope for enhancing SSA. Therefore, Chapter 5 details the methods used to determine the accuracy of the data that were obtained, and to identify the presence or absence of systematic errors. The data are examined to determine the degree to which systematic errors are present, followed by a discussion on the degree of accuracy obtainable with the equipment used for this experiment. The results of the data obtained on satellites are then discussed in chapter 6, in the context of previous work conducted in this field, to determine whether using a small aperture telescope for photometry can significantly contribute to enhancing SSA. In the process, some novel contributions to the field are identified.

Chapter 7 provides a summary of the conclusions identified throughout the thesis and recommendations for future work. Finally, Appendix A contains the complete set of lightcurves and colour variations that were derived from the data.

Chapter 2

Literature survey

The following chapter presents a summary of the significant work that has been conducted to advance the understanding of the contribution that photometry can make to SSA, as it relates to artificial Earth orbiting satellites. Spectroscopy is arguably a subset of photometry, albeit with a much higher spectral resolution. Therefore, aspects of the spectroscopic analysis of satellite reflectance are presented, when they relate directly to photometry. The chapter begins with a summary of the laboratory experiments that have been conducted to advance the understanding of photometry of artificial satellites. The next section then discusses techniques that have been developed to discriminate between two or more satellites without specifically identifying their physical characteristics; and the final section presents the work that has been done to this point to characterise satellites from their photometric signatures.

2.1 Laboratory Experiments

Whilst a significant amount of work has been conducted using photometric observations of satellites, or computer simulation studies, only two purely laboratory based photometric studies have been conducted to assist with the interpretation of photometric observations of satellites. The first study, by Deiotte *et al*, was reported in 2006 [19]. Deiotte's team created a laboratory setup designed to enable photometric measurements to be taken of a small object illuminated by a collimated light source. The setup consisted of a five axis articulated arm robot, an Apogee 9E CCD camera, and an Oriel 13 watt white light source with a custom made collimator, all mounted on an optics bench. The robot arm was computer controlled using a program that allowed the object being investigated to be precisely oriented to achieve a desired attitude relative to the light source and camera. The camera and light source were mounted on rails and positioned by hand. Limitations of the system include that

the light source, object, and camera must all remain in the same plane, parallel to the optics bench; the maximum object size was limited by the 6.5 cm diameter light source; and the minimum possible phase angle was 15 degrees.

Five objects were selected for the experiment, which was conducted only at phase angles of 15 degrees and 90 degrees for each object. In addition, since the light source, object and camera were required to remain in the plane parallel to the optics bench, no latitude dependent brightness variations could be investigated. Each object was rotated through between two and four full rotations, with unfiltered images taken at one degree intervals. Lightcurves were produced representing the measured brightness of each object as a function of rotation angle at 15 degrees and 90 degrees phase angle. The five objects were a ping pong ball; a white cube; a sphere with alternating white and black painted, vertically oriented, quarters; a cube with alternating faces painted white and black; and a miniature replica of an Atmospheric Neutral Density Experiment (ANDE) sphere, which is similar to the black and white sphere with a white band around the equator.

The results of the experiment demonstrate that the laboratory technique is potentially useful in determining the characteristics of lightcurves of simple objects. The spherical objects tended to have more rounded peaks in the lightcurve, whereas the cubes exhibited pointed peaks in their lightcurves. In addition, the purely white objects showed less variation in amplitude than the black and white objects. Figure 2.1 depicts the lightcurve resulting from two full rotations of the cube with alternating black and white faces at 15 degrees phase angle. The top and bottom faces of the cube were oriented parallel to the optics bench, and the axis of rotation was perpendicular to the bench.

It was noted that relatively strong specular reflections were sometimes observed off the black paint, but they were only visible through very small angles of rotation. Given that the experiment was conducted at only two fixed phase angles by rotating the illuminated object, it is more representative of a tumbling satellite at a constant phase angle rather than a stabilised satellite at changing phase angles. No attempt was made to create lightcurves representative of a stabilised satellite by maintaining
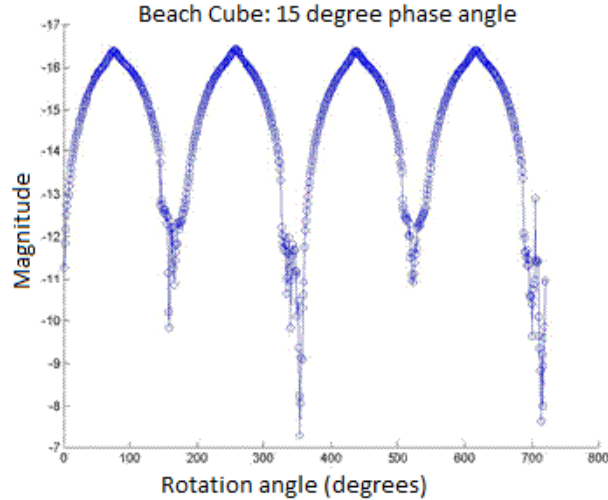
Figure 2.1: Lightcurve of a cube with alternating black and white faces [19].

a constant object attitude relative to the camera and moving the light source; or by maintaining a constant illumination angle and moving the camera. Also, the fact that the lamp-object-camera plane remained parallel to the optics bench means that the effect of changes in declination could not be investigated. Although each object was measured over more than one full rotation, the experiment was not repeated for each object for a given phase angle to determine the repeatability of the results. Areas where the data appears more noisy seem particularly unpredictable. For example, in Figure 2.1 the measured intensity near 160 degrees and 520 degrees corresponds to the same part of the object being illuminated, yet those parts of the lightcurve are more different to each other than the peaks of the lightcurve. No reference was made to the magnitude of the uncertainties in each of the observations.

The authors note that the aim of the experiment was only to validate the setup and observation method, and that it was only a small step towards the ultimate goal of characterising the lightcurves of complex objects following complex trajectories; however the experiment does not appear to have been followed up.

2.1.1 Photometric and spectroscopic measurements of debris-like materials

The second laboratory experiment aiming to model photometric lightcurves of satellites was conducted by Cowardin as part of her doctoral thesis in 2010 [20]. The lab

setup was functionally similar to that used by Deiotte *et al.*, with the exception that the light source, camera, and target object were mounted on separate benches, and the light source was not collimated. The same five axis articulated-arm robot that Deiotte employed was used by Cowardin to manipulate the position and orientation of the target object. A Xenon arc lamp was used to simulate solar illumination, and the detector was a Charge Coupled Device (CCD) camera fitted with a five position filter wheel with Johnson/Bessel BVRI filters. A white spectralon panel was used as a standard intensity reference to enable relative intensity measurements and colour ratios to be produced.

In addition to the photometric measurements, Cowardin repeated her experiment using a spectrometer, integrating the measured spectra over the photometric filter bandpasses such that the photometric and spectroscopic data could be compared. Finally, the laboratory results were compared to published colour photometric data for a group of orbital debris targets.

Cowardin obtained 14 material samples believed to be representative of the space debris population, including flight ready parts, impact experiment fragments, and explosion fragments from pressure tested objects. Photometric measurements were taken of each object at 6 degrees, 36 degrees, and 68 degrees phase angle, averaged over a 360 degree rotation about one axis, with 10 degree rotation intervals between measurements. Using an average brightness value over a full rotation of each object implies the assumption that the observed magnitudes are largely independent of illumination and observation angle, instead depending only upon phase angle. Cowardin noted that at first there did appear to be a phase angle dependence of the magnitude, but it was subsequently concluded that the illumination equipment was causing a bias in the data. No further explanation was given regarding the nature or cause of the bias, however only the data obtained at the smallest phase angle (six degrees) were used. From the average magnitude measurements a single value was determined for the B-V, B-R and B-I colour indices of each material sample, with associated variance values for each colour indice. The colour indices were therefore calculated based upon magnitude measurements taken throughout an entire 360 degree rota-

tion, again implicitly implying that the magnitude is independent of the illumination and observation angles, or that if the magnitude does vary it does so to the same degree for each of the filter bandpasses (i.e. the Spectral Energy Distribution (SED) is independent of object attitude).

When the experiment was repeated with the spectrometer, only measurements of diffuse reflections were able to be included in the data set due to limitations of the ability of the equipment to measure specular reflections, whereas the photometric data included all of the observed specular and diffuse reflection components. One of the key observations made by Cowardin was that many of the material samples exhibited far stronger reflections at the longer wavelengths when measured using photometry than when measured using spectroscopy. In addition, it was observed that many of the materials showed far greater variance in the photometric observations than the spectroscopic observations. It was hypothesised that these effects could be due to differences between specular and diffuse reflectance properties, or dilution of the light reflected off the surface of interest by light reflected off other surfaces. Some of the material samples, such as the Aluminised Kapton, showed much better agreement between the photometric and spectral data than others.

Cowardin stated that the photometric laboratory data are a better source to compare to telescopic data than spectroscopic laboratory data because the spectroscopic data measure reflections from a single, small area of material, whereas the photometric data measure light reflected off the whole sample. She did not elaborate on why the spectroscopic data only measure reflections from small areas of each material. Cowardin attempted to compare the laboratory data to photometric measurements of two debris objects in geostationary orbits. It was observed that correlating laboratory measurements with unstable objects that regularly glint is extremely difficult due to the large scatter in magnitude and colour data that results. The only conclusion that could confidently be drawn from observations of the debris is that they consist of highly reflective material. Cowardin concluded that photometric characterisation of satellites is only feasible for stabilised objects, and that knowledge of their area to mass ratios could be of benefit.

2.2   Satellite Discrimination

Research using telescope observations of satellites can be divided into two general categories. The first category includes experiments solely aiming to develop a means to discriminate between multiple satellites; and the second category consists of investigations aiming to classify satellites based on certain characteristics of their observed lightcurves and colour variations. The following is a summary of advances made to date in the understanding of how to discriminate between multiple satellites.

2.2.1   Discrimination based upon magnitude versus colour analysis

The United States Air Force Research Laboratory (AFRL) has conducted extensive and ongoing research into using photometry for satellite discrimination and characterisation, with data collection efforts going back at least as far as 1996 [21]. By far the majority of the AFRL's published work was conducted by teams led by Tamara Payne or subgroups of her teams.

In 1999, Payne *et al.* undertook a project with the goal of creating a new set of filters that is optimised for discriminating between multiple satellites [14, 22]. The newly developed filters were to be known as Space object identification In Living Colour (SILC) filters. Payne *et al.* gathered spectroscopic data in two broad wavelength bands, which they called blue and red, on nine active GEO satellites and one rocket body. The blue band is from 329.9 nm to 668.9 nm, and the red band is from 610.2 nm to 952.9 nm. A total of 717 blue spectra and 580 red spectra were taken. To simulate real filters, Gaussian functions were generated with FWHMs of 10, 20, and 80 nm spaced 5 nm apart across the entire wavelength band. Magnitudes measured through each simulated filter were calculated by multiplying the simulated filter functions by each observed spectra. To account for brightness variations with phase angle and season, the data were divided into two groups: one group consisted of data acquired from May 1999 through to July 1999, and the second group consisted of data acquired in September 1999, which contained significantly brighter specular glints. The data from the first group were then divided into ten degree solar phase

angle bins; and the initial analysis was conducted using only these data. The analysis consisted of comparing the median magnitude of each satellite, in each phase angle bin, as calculated by each simulated candidate SILC filter. The best filters for satellite discrimination were judged to be those that yield magnitudes that on average differ by the greatest amount for different satellites. The selected SILC filters, referred to by their median nanometer wavelength values, were called 415, 475, 590, 630, 820, and 850.

Once the six most optimum SILC filters had been selected, data from both the first and second group were used to verify their effectiveness. The effectiveness of the SILC filters was measured by determining how frequently a satellite could correctly be identified, out of the group of ten that were observed, from a single observation, using a specially developed software program that compared colour ratio versus magnitude data. The same verification process was repeated using Johnson B and V filters simulated by Gaussian functions. From the data collected during the experiment it was determined that the correct satellite was identified on average 85.3 percent of the time using the Johnson filters, and 90.3 percent of the time using the SILC filters. This result was obtained from the same data that were used to select the SILC filters, and only included data from the blue spectra measurements, also excluding all data collected in September. When the brighter September data was included the correct satellite was only identified on average 85 percent of the time using the simulated SILC filters. Results were not presented for the simulated Johnson filters with the September data included. The effectiveness of the SILC filters was not tested on satellites that were not used to select them in the first place. From the results of the study it was concluded that the SILC filters perform better than the Johnson filters at discriminating between satellites.

Payne *et al.* continued their research into the utility of the SILC filters in 2001 [15, 23]. Photometric observations were taken of a cluster of five GEO satellites using the SILC and Johnson filters in one week in May 2000 and May 2001. Plots were made of the magnitude through one filter versus the colour ratio of two other filters, in both the Johnson and SILC systems, of all five satellites in the cluster. Figure 2.2

is an example of one of the plots using the SILC filters. Of note, measurements in different colour bands were obtained sequentially, not simultaneously.
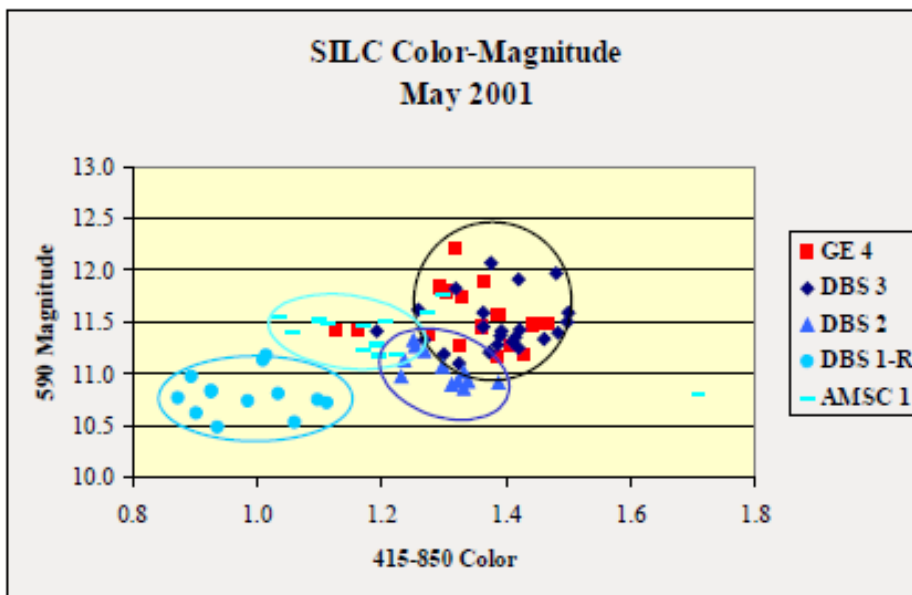
Figure 2.2: Example of a SILC Colour-Magnitude Plot [15]. The legend labels refer to the respective satellites' names.

The aim of the experiment was to confirm the improved ability of the SILC filters to discriminate between satellites compared to the Johnson filters. It was pointed out that the colour-magnitude data points for each satellite tend to cluster together in different areas on the plot, although with significant overlap for some of the satellites. Comparing the degree of separation between these clusters, the authors concluded that the SILC filters do provide an improved ability to differentiate between satellites. Details were not provided regarding the length of time between each observation, or the length of time between the first and last observations, however there appear to be only approximately 12 to 22 data points per satellite.

The data gathered in 2000 were compared to the data from 2001 to see how well they match, and to determine the likelihood of detecting cross-tagged satellites. A Johnson B magnitude plot created from observations of the same satellite cluster in 1998 was used to assist in interpreting the results of the colour-magnitude plots, as depicted in Figure 2.3. Although only approximately 20 data points were plotted per satellite over a seven hour lightcurve, often with over one hour between points and sig-

nificant variability over short timeframes, it was concluded that satellites' lightcurves are smoothly varying and predictable. No comment was made regarding the possibility of rapid or irregular magnitude changes ocurring over time scales shorter than the typical time between measurements, other than that transient specular reflections may be responsible for some of the unusually bright individual measurements observed. The date that the 1998 data were obtained was not provided.
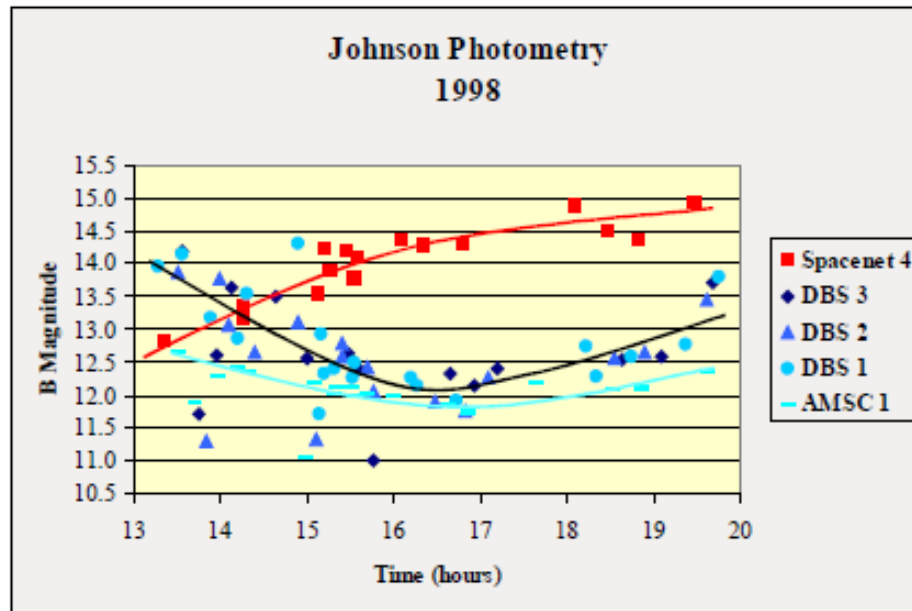
Figure 2.3: Johnson B Magnitude Plot of 1998 data [15].

The general magnitude differences exhibited by the satellites, as well as their relative positions on the colour-magnitude plots, were used to hypothesise about the identities of each of the satellites in the cluster. It was also assumed that satellites of the same bus type should have similar lightcurves. The colour ratio values for all of the year 2000 SILC data on the 590 magnitude versus 415-850 colour plot range from approximately 1.2 to 2.0. By contrast, the range of values for the year 2001 data is approximately 0.85 to 1.5, with one outlying data point at greater than 1.7. Data from 1998, 2000, and 2001 were used in Payne's analysis, assuming that the lightcurves of satellites are stable over long timeframes. The fact that the 2000 colour data spans a different range of values from the 2001 data was not discussed. Comparing the data from 2000 to the data from 2001, in conjunction with the 1998

B magnitude data, it was concluded that there had been substantial misidentification of satellites within the cluster, and that the SILC filters provide an improved ability to differentiate between them.

In 2003 Payne published an update of the progress made by her team towards using photometry for satellite discrimination [21]. Payne used data gathered on twelve three-axis stabilised GEO satellites in the Johnson B and R filters from 1998 and 2002 to investigate the relationship between magnitude, colour, and phase angle. The phase angle used throughout the analysis is the equator projected phase angle. The data collection process was not described in detail, however observations by Payne's team have always been reported as being conducted sequentially through the different filters rather than simultaneously. The timing of the data collection was reported by season and year, rather than by specific dates.

Payne discovered that a least squares best fit plane could be applied to data plotted in three dimensions using B magnitude, B-V colour ratio, and phase angle; and that the plane remained stable over time. It was also observed that different satellites did not necessarily share the same plane, possibly representing a means to discriminate between satellites. Standard deviation values from the best fit plane for each satellite were not provided, nor was a statistical analysis of the degree of fit of a given plane to data from other satellites. For some satellites as few as six data points were used in the analysis. Payne concluded that colour index is tightly correlated to brightness and phase angle, stating that the physical processes responsible for brightness changes as a function of phase angle are the same as those responsible for colour ratio changes. Finally, it was hypothesised that a diagnostic could be developed based upon the presented results whereby satellites that do not fit the plane must have an abnormal ratio of solar panel size to bus size or solar panel size to solar panel brightness.

2.2.2 Discrimination based upon magnitude versus phase angle analysis

Beginning in 1999, a team of researchers from the United States Naval Observatory (USNO) and the Aerospace Corporation, led by Frederick Vrba, begun obtaining optical and near infrared photometry measurements of Global Positioning System

(GPS) satellites using the one metre telescope at the USNO Flagstaff Station [24, 25]. Observations of GPS Block II, IIA, and IIR satellites obtained between 1999 and June 2003 were analysed, initially with the aim of calibrating their brightness as a function of Sun-observer-satellite phase angle for later USNO astrometric measurements, and to determine reflectivity values for input into solar force models. During the analysis of the data it was observed that photometric observations might be able to provide some level of discrimination between satellites.

Vrba reported several unexpected results from the study, including: the occasional presence of glints with brightnesses of up to one hundred times greater than at phase angles either side of the glint; an apparent decrease in reflectivity of approximately ten percent per year for the Block II and IIA satellites; and that all of the satellites displayed a scatter in the data that was much greater than the uncertainty of individual measurements. Throughout the analysis Vrba often reported results only from measurements obtained within a few degrees of 140 degrees Sun-observer-GPS phase angle rather than scaling all results to a single phase angle to avoid low Signal to Noise Ratio (SNR)s when the satellites are low in the sky and GPS reorientations when the satellites are coming out of eclipse.

Figure 2.4 depicts the essentially linear relationship that Vrba discovered in all bands between magnitude and phase angle. Vrba noted that the slopes of the Block II light curves were significantly different from the Block IIA light curves, indicating a possible means of discrimination despite the fact that the two satellite types are nearly identical. Vrba also observed that the lightcurves and colour ratio curves were only repetitive if taken within a few days of each other, however the effect of this lack of repeatability on satellite discrimination methods was not examined in detail.

Figure 2.5 depicts I-band magnitude normalised to 20 000 km at 140 degrees phase angle as a function of time on orbit. Although it was noted that the scatter in the data is much greater than the errors of individual measurements, it was concluded that diminishing reflectivity with age is clearly a complicating factor for satellite discrimination. In each of the colour bands it was noted that brightness seemed to diminish equally, by about ten percent per year, implying that the satellites do
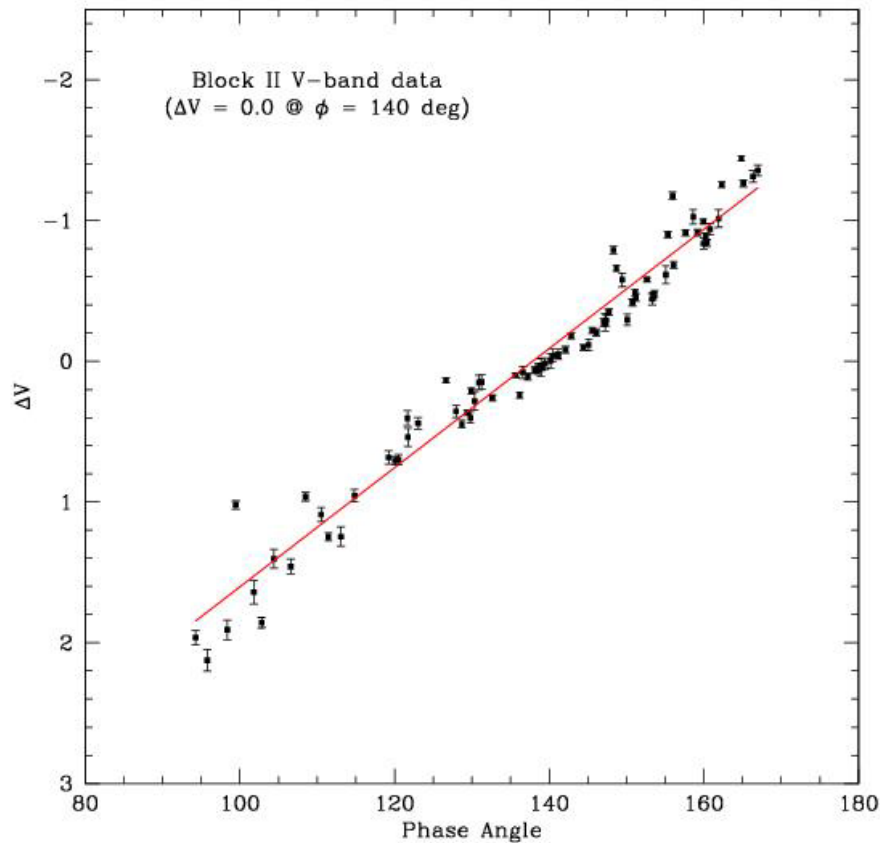
Figure 2.4: GPS Block II V-band data normalised to a magnitude of 0.0 at 140 degrees phase angle [25].

not change colour as they age, due to space weather effects or any other factor. The possibility that the apparent reduction in brightness with age might be due to differences in the sun-satellite-observer geometry or subtle differences between different satellites of identical bus type was not investigated.

In addition, all of the data indicate that for all satellites the observed colours are redder than expected from known laboratory measured material properties [26]. It was therefore hypothesised that rapid surface changes may occur very soon after launch that cause the apparent redenning.

Vrba concluded that with care discrimination between even very similar satellites might be possible, by analysing the slope of magnitude versus phase angle plots, with repeated observations and broad wavelength coverage.
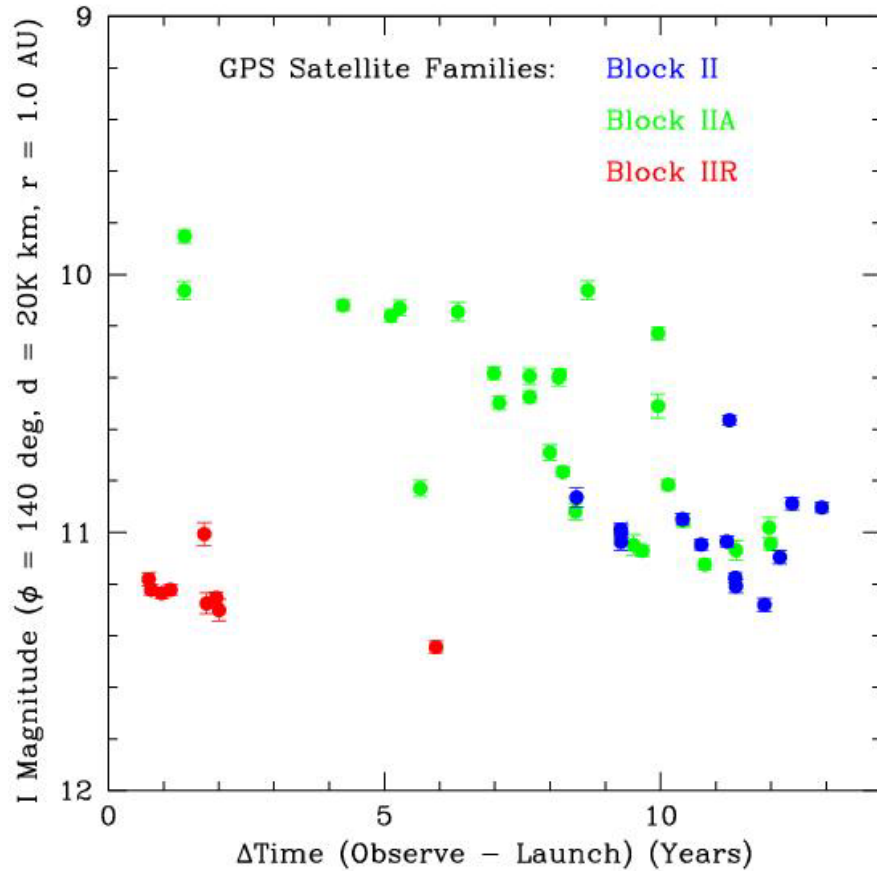
Figure 2.5: I-Band magnitude at 140 degrees phase angle normalised to 20 000 km as a function of time on orbit [25].

In 2008 Scott and Wallace from Defence Research and Development Canada (DRDC) used a 35.6 cm small aperture telescope to obtain unfiltered photometric data on 14 three-axis stabilised GEO satellites on 24 February, 25 February and 17 March [27]. The objective of the experiment was to investigate the use of small aperture telescope photometry to differentiate between satellites based upon brightness. Each observing session began with a measurement of the easternmost of the 14 satellites, then progressing westwards until all of the satellites had been imaged. This process was repeated every five minutes from sunset until sunrise. The resultant lightcurves of each of the satellites were compared subjectively in terms of their general shape. There were six different bus types among the 14 satellites that were investigated. It was noted that satellites of the same bus type exhibited the same

general lightcurve shape. It was hypothesised that Anik F1 and Anik F1R had been cross-tagged based upon the shape of their lightcurves relative to the other observed satellites of the same bus type. It was concluded that satellites of the same bus type exhibit similar lightcurves and that they are recurring, thereby offering a means of discrimination.

In 2010 LeVan, also of the United States AFRL, presented a method for distinguishing between closely spaced GEO satellites in clusters to enable their reidentification after they pass too close to each other to visually resolve them as separate objects [6]. Adapting a technique that had previously been applied to characterising an object's spectra, LeVan proposed that just three measurements are required to correctly identify all of the satellites in a cluster, dramatically reducing the effort required compared to techniques that rely on analysing a satellite's full lightcurve. Choosing three arbitrarily selected equator-projected phase angles, which for the purposes of this discussion will be called ϕ_1, ϕ_2, and ϕ_3, where $\phi_1 < \phi_2 < \phi_3$, a brightness measurement is made at each phase angle, $S_{\phi 1}$, $S_{\phi 2}$, and $S_{\phi 3}$. A quantity that LeVan called the *curvature* of the lightcurve is then calculated for each satellite according to the following equation:

$$\text{Curvature} = \frac{S_{\phi 1} \times S_{\phi 3}}{S_{\phi 2}^2} \tag{2.1}$$

LeVan asserts that each satellite's lightcurve curvature is not necessarily expected to remain constant with every measurement over time, however the ratio of one satellite's curvature to another's is expected to be stable over long time periods. Therefore, given a cluster of any number of satellites, if the ratio of every permutation of two satellite curvatures is calculated then the identity of every satellite can be determined from previous curvature calculations using simultaneous equations. An example was provided to demonstrate how the process might work, however a detailed analysis was not conducted into the sensitivity of the process to changes in variables such as the phase angles that are selected or day of the year. No information was provided into how effective the process is expected to be when differentiating between satellites of the same bus type.

2.3 Satellite Characterisation

2.3.1 Early photometry studies

The effort to understand how photometry can be used to characterise artificial satellites has been ongoing for many years. The development of CCD cameras has made conducting accurate photometry easier, and the effort has accelerated since the turn of the century, however some studies using photomultiplier tubes were published as early as the 1970s [28].

Beginning in 1966 a four year study consisting of six observation periods of the Passive Geodetic Satellite (PAGEOS) and Echo 1 satellites was conducted [28]. PAGEOS and Echo 1 are identical inflatable, spherical satellites with highly specular vapour deposited aluminium coatings, and diameters of 100 ft. PAGEOS was initially placed into a nearly circular orbit at 4248 km average altitude, and Echo 1 was placed into a nearly circular orbit, averaging 1609 km. The purposes of the study were to obtain photoelectric measurements of the magnitude of the satellites in the Johnson U, B, V, R, and I bands; to measure their specular and diffuse reflectances; and to measure their mean radii of curvature and variations in radius of curvature.

The specularity of both of the satellites, from soon after launch until greater than seven years on orbit, was found to be very stable at close to 96 percent of the laboratory determined value in nine out of the ten measurements that were reported. Only the final measurement of Echo 1, at 7.2 years on orbit, was found to have dropped by approximately five percent. It was hypothesised that the reduced specularity in the final measurement could be due to surface degradation caused by atmospheric effects since there was increased solar activity in that year and the satellite's mean orbital altitude fell to below 1207 km.

The reflectance values in the five colour bands also appear to be essentially stable within the variance shown in the data, although as with the specularity only ten measurements were made in total per filter. The data also show a drop in the reflectance of Echo 1 after 6.5 years on orbit, corresponding to the increased solar activity of 1967. The drop in reflectance was greatest for the shortest wavelengths,

and there appeared to be no drop in reflectance for the V filter; although there was insufficient data in the R and I filters to determine if a drop in reflectance had occurred in those wavebands. The degradation in reflectance after 6.5 years on orbit for Echo 1 correlates with its reduction in specularity at about the same time. Aside from the final two measurements of Echo 1 the data indicate that specularity and reflectance are essentially unchanged even after at least six years on orbit, implying that space weathering effects may be negligible at the altitudes at which the satellites were operating.

2.3.2 Categorisation Based On Collections of Individual Photometry Measurements

In 2010 Hejduk proposed a method that uses collections of individual photometric measurements obtained over many years, rather than long continuous datasets, to categorise the relationship between brightness and phase angle of satellites [29]. The motivation for such an approach lies in the fact that individual measurements represent the most common photometric collection technique, thus representing an ideal opportunity to investigate the relationship between brightness and the illumination-observation geometry. The Joint Space Operations Center (JSpOC) in California currently maintains the largest dataset of single point photometry of satellites, with approximately 3.5 million individual measurements [29]. Hejduk stated that the dataset was obtained using three, one meter class telescopes, located in Diego Garcia, New Mexico, and Hawaii, however the passbands used to obtain these data were not specified. Instead of focussing only on one particular type of satellite or orbit regime, Hejduk analysed all payloads, rocket bodies, and other debris that were in Highly Elliptical Orbit (HEO), Medium Earth Orbit (MEO), or GEO orbits, with orbital periods greater than 225 minutes and measurement histories exceeding 500 measurements. These categories were maintained throughout the analysis to determine whether the photometric response of each type of satellite and orbit regime could be categorised equally effectively.

Hejduk's analysis began by placing measurements for each satellite into signed solar phase angle bins, each with a size of three degrees. The median magnitude

was selected as the representative intensity for each bin. A standard lightcurve was produced for each satellite by plotting the median measurement in each phase angle bin against the central phase angle of each respective bin. A linear function was then fit to the plotted data, which was observed to be an effective yet simple representation of spacecraft lightcurves in most cases. Given that brightness generally decreases with increasing absolute values of phase angle, two lines were fit to the data: one with a negative slope, and one with a positive slope, either side of zero phase angle. Figure 2.6 is an example of such a plot.
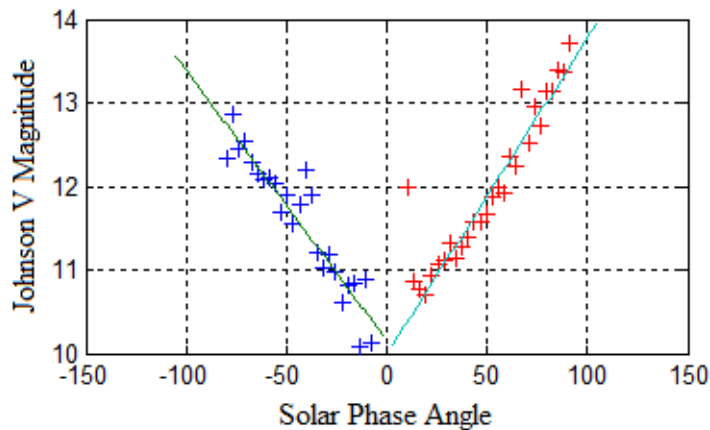
Figure 2.6: Example of linear fit of magnitude to phase angle [29].

Hejduk proposed four basic tests, regarding the linear fit, to categorise each satellite. The first test would simply be to assess whether the sign of the slope of each half of the linear fit was such that the brightness increased with decreasing absolute value of phase angle. A slope in one or both halves of the plot that did not satisfy this criterion, or the inability to establish a linear fit, would indicate a satellite whose lightcurve does not conform to the usual lightcurve behaviour. Through analysing the JSpOC dataset, Hejduk discovered that all types of satellite, with the exception of HEO debris, conformed to the expected behaviour at least 95 percent of the time. The slopes of HEO debris lightcurves were found to satisfy this criterion in nearly 80 percent of cases.

The second test separates out satellites with linear fit slopes of magnitude greater than a given threshold. The choice of threshold to apply is somewhat arbitrary,

however Hejduk noted that in almost all cases a significant majority of lightcurves had slopes with an absolute value of less than 0.022 magnitude per phase angle degree, which corresponds to a change of two magnitudes over the observation range of 0 to 105 degrees. Hence 0.022 magnitude per degree was suggested as a suitable threshold beyond which photometric responses could be flagged as unusual. Payloads, both stabilised and unstabilised, were noted to be the only satellite category that frequently did not exhibit the same distinctive slope distribution.

The third test examines the difference in the magnitude of the slope on either side of zero phase angle. Hejduk observed that at least 80 percent of satellites in all categories and orbit regimes exhibit slope asymmetry less than 0.005 magnitudes per phase angle degree, except for stabilised payloads in GEO. In the case of stabilised payloads in GEO only approximately 40 percent of satellites have slope asymmetry less than 0.005 magnitudes per degree.

The final test uses a stochastic modelling approach to analyse the residuals in the linear fit model to determine how well it fits the data. It was noted, as an aside, that some satellites' data differed quite considerably from the linear fit. Further, stabilised payload data tended to vary with phase angle in a continuous and predictable way, whereas data from all other satellites varied randomly with phase angle.

Hejduk concluded that the four tests could be used to separate all satellites into five different groups that categorise the relationship between their brightness and phase angle, although the fifth group, that passes all of the four tests, would constitute more than 50 percent of all satellites. There was no investigation into the typical variation in magnitude that might be expected in a given three degree phase angle bin, or whether that variation is consistent over the entire range of phase angle bins. There was also no analysis of the effect of season dependent solar declination angle on brightness.

2.3.3 Characterisation of lightcurves

Much of the work conducted into using broadband photometry in SSA since the late 1990s has been conducted by Payne and colleagues from the AFRL and their

collaborating organisations. Following efforts relating to satellite discrimination, as previously discussed, Payne *et al.* begun focussing on more ambitious techniques for satellite characterisation. Most of the work focussed on GEO satellites because their stabilised orientation and location relative to Earth's surface reduces the number of variables involved and simplifies the characterisation problem. In 2005 Payne described three main observing strategies that each make different compromises between the number of satellites that can be observed in a given night and the quantity of data that can be obtained on each satellite. The potential uses for a database comprising all of the measurements collected using such strategies were also discussed [30]. Such a database was initiated at the AFRL in 1996 and has continuously been expanded ever since, however it does not appear to be publicly available [30]. The three general observing strategies can be described as follows:

- Taking relatively few measurements (up to about 30) of as many satellites as possible in a single night in order to obtain their basic lightcurve shape or to enable the creation of magnitude versus colour plots or colour versus phase angle plots. Payne stated that colour-phase angle plots are not as effective as other techniques for Space Object Identification (SOI), however the claim was not elaborated.

- Intensively observing a satellite for one or two hours around zero phase angle to determine solar panel orientation information. By only observing a given satellite for one or two hours, several satellites can be observed in a single night.

- Intensely observing a single satellite for an entire night to obtain the most detailed information possible.

Any of these observing techniques may be used to gain information on more than one satellite at a time whenever more than one satellite is visible in each CCD frame, such as when observing a satellite cluster. The third observing strategy provides more information about a satellite than either of the first two strategies, however the first two strategies allow for information to be gathered on more satellites per night.

Payne describes the first observing strategy as being useful for discriminating between closely spaced satellites, and for identifying them by bus type through examining the general lightcurve shape and magnitude-colour plots. It was claimed that GEO satellites could be identified by bus type with 97 to 98 percent confidence using this technique.

The second observing strategy is described as being of use for deriving power consumption and health information. Since it is usually desirable for a satellite's solar panels to point directly at the sun for maximum power generation a strong specular glint is commonly found to occur near zero phase angle. A simple geometric analysis of the asymmetry around zero phase angle can reveal the angular difference, in the longitudinal direction, between the satellite-sun vector and the normal to the satellite's solar panels. A highly asymmetric signature might also be indicative of problems with the health of the satellite's power subsystem.

Payne described the third observing strategy as being of use for producing three dimensional surface plots where day of the year and equator-projected phase angle are the independent variables, and magnitude is the dependent variable. It was noticed, from the data on some of the frequently observed satellites, that the shape of a given satellite's lightcurve changes according to the differing sun-satellite-observer geometry throughout the year. Three dimensional plots of these data, built up over time, can subsequently be used for change detection at a later time. Figure 2.7 is an example of such a plot using interpolated lightcurves of Telstar 7 obtained on five nights throughout a six month period beginning in December 2004.

Payne concluded that a detailed database of GEO broadband colour photometric observations would have significant utility for SSA.

Also in 2005, a subgroup of Payne's team, led by Gregory, sought to improve the understanding of the physics of satellite lightcurves by comparing simulated lightcurves with actual observations [17]. Using artists conceptions and photographs eight satellites were modelled as basic structures consisting of three components: a solar array, satellite body, and large antennae. Using these models and a materials reflectance database created by Jorgensen as part of her doctoral thesis [5], Gregory
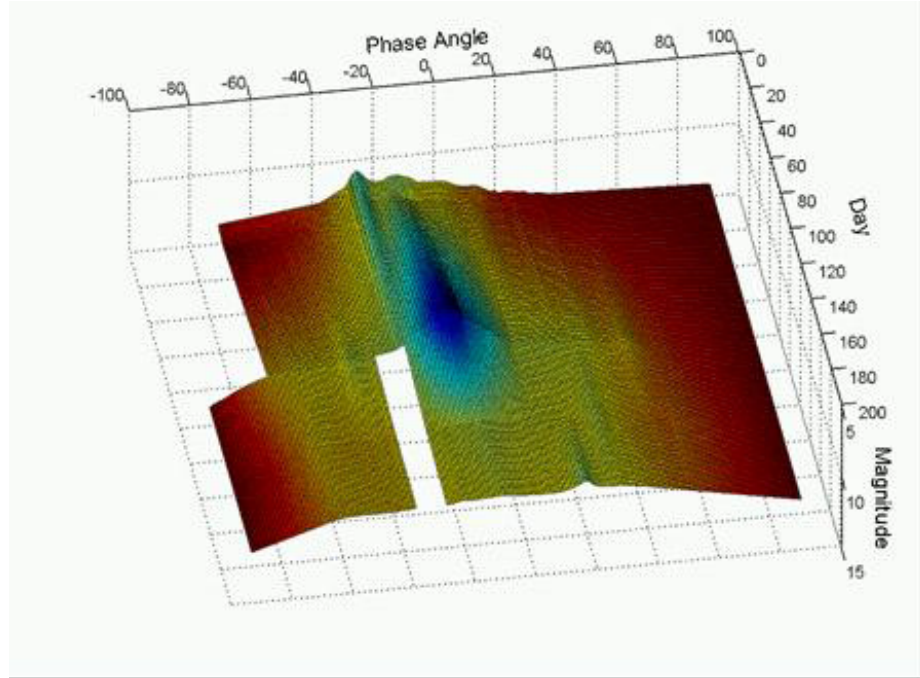
Figure 2.7: An interpolated plot showing how Telstar 7's lightcurve changes throughout the year [30].

calculated simulated lightcurves for the eight satellites. The simulated lightcurves have Johnson B,V,R, or I magnitude, or SILC filter magnitudes as the dependent variable, and equator-projected phase angle as the independent variable. The day of the year and observing location that the lightcurves were simulated for was not specified. Jorgensen's material reflectance database is not available in the public domain, however some examples of reflectance spectra for various materials are provided in her doctoral work, and it is noted that each material's normalised reflectance spectrum is considered to be invariant with changing illumination or observation angle. Importantly, all of Jorgensen's measurements were taken using the same illumination and observation geometry, which avoided specular reflections. The effects of changing the geometry were not reported. Further, only minimal details are provided about each material for which a reflectance spectra were measured. For example, the surface roughness and the exact material type are not always specified. Figure 2.8 depicts the measured reflectance spectra for three different materials commonly used in the construction of satellites [5]. Note that a reflectance function is shown for steel, however

the particular alloy that was measured, and its surface roughness are not provided. Similar observations can be made about the other two materials shown.
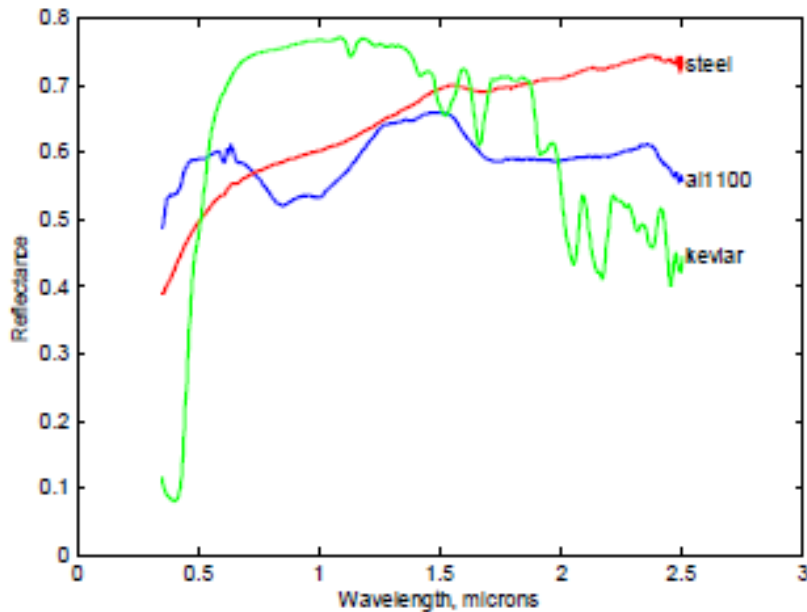
Figure 2.8: Example reflectance spectra as measured by Jorgensen [5].

Gregory reported limited success at replicating observed lightcurves using his simulation; three examples were shown. The simulations were generally unable to reproduce glints and small surface features, and it was noted that the limited laboratory measured reflectance data pertaining to the specular regime was one of the primary problems with the lightcurve simulation model. When an observation-based specular reflectance function for solar panel material was added to the model, the simulated lightcurve of AMSC1 was far more successful at reproducing the specular reflection that was observed on 08 October 2004. It is not known whether the specular reflectance function is as successful at simulating the lightcurve on different days of the year or with different types of solar panels.

In 2006, as a result of the ongoing analysis of the AFRL's observation database, Payne defined five lightcurve classes into which all satellite lightcurves could be categorised [18]. Each lightcurve class is defined by its general shape characteristics, and represents one of the most fundamental pieces of identifying information for a satellite. The five different lightcurve classes are defined as follows:

- Canonical: A brightness peak exists near zero phase angle, with the brightness decreasing either side of the peak in a smooth and symmetric manner.

- A2100. A local minimum exists near zero phase angle. Brightness increases smoothly to either side of the local minimum to reach two local maxima at ±40 degrees. Beyond ±40 degrees phase angle brightness smoothly decreases again. Most A2100 lightcurves are generally, but imperfectly, symmetric.

- Telstar. Similar to the canonical class, however secondary minor brightness peaks exist near ±40 degrees, seen most strongly in the shorter wavebands.

- BSS702C. Similar in shape to the Telstar class, however the secondary peaks exist near ±60 degrees. In addition, the secondary peaks are brighter than the zero phase angle peak, and broader than the Telstar secondary peaks.

- Peculiar. These lightcurves exhibit either a lack of symmetry or seemingly random, irregular brightness variations, or both.

Sufficient data was considered to be available for 36 GEO satellites to have their lightcurves classified, although the number of nights on which data were gathered per satellite was not specified. Payne provided possible explanations as to which aspects of various satellite's structures may be responsible for their lightcurve's features. As a result, she deduced that satellite lightcurves are tightly correlated with their basic bus structure. Of note, however, Payne categorised some satellites of identical bus structure into different lightcurve classes. For example, Anik F1 and Galaxy 11 both have the BSS-702C bus structure, however Anik F1 is classified as having a BSS702C lightcurve, whereas Galaxy 11 is classified as having a peculiar lightcurve.

Some limitations with the data were emphasised, including the fact that some were restricted to the central 20 degrees of phase angle, and that different filters were used during different observing periods. It was also noted that the data set was at that point still relatively small, and that more satellite lightcurve classes may be identified in the future. To date, to the best of the author's knowledge, there have been no further lightcurve classes identified.

In 2009 Payne *et al* published a method for extracting the spectral response of the solar panels, and of the satellite body, from a GEO satellite's photometric lightcurve [31]. The method relies upon several hypotheses that Payne proposed as a result of the analysis of over a decade of photometric data on GEOs. The hypotheses most significant to the method are as follows:

1. The undulations in lightcurves are the result of specular glints off small features on the satellite.

2. Since the solar panels track the sun, their contribution to the lightcurve is symmetric about the phase angle at which the peak of the specular glint occurs (usually close to zero). Given that lightcurves have usually been observed to not be symmetric, the asymmetry is therefore a result of asymmetries in the body.

3. The solar panel signature is a linear function of the cosine of the difference between the phase angle and the specular peak phase angle; and is thus an even function.

4. The rate of change of the solar panel brightness is positive prior to the specular peak, and negative after the peak; and is thus an odd function.

5. Assuming the specular peak is narrow, the rate of change of the satellite body's contribution to the brightness can be considered to be constant, and is thus an even function over the period of the specular glint.

Using basic arguments relating to the symmetry of the solar panel and satellite body contributions to the lightcurve, Payne argued that by using a small number of measurements taken at key phase angles the solar panel contribution to the brightness can be removed, leaving only the body contribution. Alternatively, the body contribution can be removed, leaving only the solar panel contribution. If this process is completed in more than one waveband then the spectra of the solar panels and the satellite body can be reproduced, where the spectral resolution is defined

by the number and spacing of the passbands of the filters used. In the interests of brevity the process will not be described here in detail: the reader is referred instead to Payne's work [31] for a full description. There are, however, three assumptions relevant to the process described by Payne that are worth considering.

First, the assumption that the rate of change of the satellite body's contribution to the brightness is constant over the period of a specular peak is critical to the success of the method. To date, convincing evidence has not been provided to show that a satellite body's brightness always varies smoothly and consistently with phase angle over times similar in length to solar panel specular glints. On the contrary, it is conceivable that a glint off even a small component on the satellite body could significantly alter the calculated body spectrum.

Second, Payne's method relies upon being able to accurately characterise a satellite's lightcurve as being a combination of a solar panel component and a body component. Even if it is assumed that the solar panel contribution could be removed from the lightcurve, the remainder of the signature is usually made up of contributions from many different satellite parts and materials including antennas, insulation, structural components, and various instruments. Although Payne acknowledges that any calculation of the body spectra will be an average of the body components, no assessment is made as to the variation that could be expected as a result of different parts of the body being illuminated or shaded as the sun-object-sensor geometry changes.

Finally, Payne argues that the calculated solar panel and body spectra should be invariant throughout the year because, "...although the phase angle conditions under which the signature data is collected vary, the solar panel and body material spectra do not" [31].

At this point it is instructive to review a recent study that examined the effect of changing illumination and observation geometry on a material's reflected SED. It has long been assumed, both explicitly and implicitly, that a given material's reflected SED remains constant under changing illumination and observation geometries. It is well known that the total energy reflected in a given direction can decrease or in-

crease with changing illumination geometry, however the assumption has always been that the percentage change in energy is equal across the spectrum. This assumption has been implicitly and explicitly evident in many studies that have attempted to compare laboratory measured spectra of various materials with observed satellite spectra, without considering the effect of the illumination and observation geometry [5, 17, 32–35]. The validity of the assumption is of significant importance because it allows spectral measurements of satellites made at different times to be compared, enabling results to be independently verified. Further, if the body material's SED does change with changing illumination and observation geometry throughout a given night, then it may render Payne's spectral extraction technique entirely invalid.

In 2011, Bedard published the results of a laboratory experiment studying the effects of the illumination-object-sensor geometry on the reflected SED of various materials used on spacecraft [36]. In the decade prior to Bedard's study almost all ground based spectrometric measurements of artificial Earth orbiting satellites exhibited an unexpected increase in relative brightness of the longer wavelengths of the visible part of the spectrum when compared with laboratory measurements. This phenomenon has become known as the *reddening effect* [36]. Until the results of Bedard's study had been published, the leading contender for the cause of the reddening effect had been space weathering [37]. In fact, several studies point to the reddening effect as evidence that space weathering needs to be incorporated into predictive models and simulations [4, 38]. Although numerous attempts have been made to model the reddening effect to match predicted lightcurves with observed measurements; and although several theories have been proposed to explain the mechanism that causes it, to date a satisfactory explanation has not been widely accepted [36]. In addition, the reddening effect has not been observed in spectral measurements of satellites that have been returned to Earth, leading to the conclusion that the effect is reversible [37].

Bedard's experiment, elaborated upon in his 2013 doctoral thesis [39], set out to investigate the effect of varying incident light angles on the spectral content of reflections off typical spacecraft materials. Bedard's results demonstrated that the relative

reflectance at longer wavelengths increased as the illumination-object-sensor angle decreased. The degree of variation with phase angle depended strongly upon the type of material, as well as its surface roughness. Bedard stopped short of concluding that the answer to the reddening effect had been found, although the demonstrated variation of SED with phase angle and material surface roughness was offered as a promising possible explanation. Regardless of whether this is the sole cause of the reddening effect or not, the fact that the SED has been shown to vary significantly with changing geometry has implications for much of the work that has been conducted into the analysis of satellite photometric signatures. For example, a satellite, or even an individual satellite component, cannot be considered to have a single value for each colour ratio: instead, the measured colour ratios would be expected to change throughout an observation period. Therefore, it would be difficult to positively identify a satellite or discriminate between satellites based solely upon the comparison of a small number of colour ratio measurements with previously measured colour ratios. Also, changing colour ratios with phase angle cannot necessarily be said to be the result of different materials being illuminated. For example, a relative increase in the brightness of longer wavelengths might be the result of increased reflection off gold-coloured Multi Layer Insulation (MLI), or it might be simply the result of a specular reflection altering the SED of other, already illuminated, materials.

In 2011, the same year that Bedard published the initial findings of his investigation into the effects of changing illumination angle on a material's SED, Payne and Chaudhary published a work tying their previous efforts into a single procedure for characterising lightcurves of three-axis stabilised satellites [40]. Likening the procedure to biometric fingerprinting, Payne and Chaudhary suggest that satellite fingerprinting could be considered to be a process that broadly categorises lightcurves according to their general characteristics, followed by three levels of increasingly detailed, and uniquely identifying, features. A collection of features at each of these different levels would then represent a satellite's fingerprint.

The broad categorising features might include overall brightness; or shape and position of specular glints. Level one features describe the shape of the lightcurve

and its colour indices. The level two features involve the separation of the solar panel and spacecraft body contributions to the lightcurve, as described by Payne in [31], allowing the solar panel and body spectra to be extracted. Level three features describe the temporal evolution of the relative abundance of the solar panel and body contributions. In order to reduce the number of unknowns when separating a signature into its solar panel and body component parts, the satellite body is represented as a single planar facet. In addition, the body is assumed to be a perfectly diffuse reflector. Making these simplifications leads to an important aspect of the fingerprinting process, known as the two point method, whereby a satellite's solar panel and body spectra can be separated using only two lightcurve data points, provided that they satisfy certain geometric conditions. These simplifications are acknowledged to be limitations with the model, however a representative magnitude of the uncertainty introduced by making the simplifications was not determined. Chaudhary and Payne note that the two facet model and two point method may produce different results for different pairs of selected points, however they state that it has been observed that by repeating their calculations using many different eligible pairs of data points, a clear trend emerges. Payne and Chaudhary also explain that one of the key principles that allows the satellite fingerprinting procedure to work is their knowledge that each material's SED is invariant.

2.3.4 Simultaneous Single-Site Photometry

In 2008 Frith *et al.* [41] conducted an experiment to gain high temporal resolution, simultaneous, single-site colour photometry on several Low Earth Orbit (LEO) satellites. Three essentially co-located small aperture telescopes were used to obtain high temporal rate photometry data, at 2.35 seconds between frames, in the Johnson B, V, and R bands, for various LEO satellites. Only three satellites were actually included in the analysis, and only for one night each: Latinsat-A, Formosat 3D, and Lincoln Calibration Sphere-1 (LCS-1).

LCS-1 is a specularly reflecting, polished aluminium sphere used for radar calibration measurements. Only B and V magnitude data were able to be collected on

LCS-1 due to the unavailability of one of the three telescopes on that particular night. Frith noted that the B and V magnitude lightcurves were essentially flat, and that the B-V colour ratio was stable and very close to that of the illumination source, the sun, as expected of a specular reflector. The phase angles at which the measurements were taken were not shown.

Latinsat-A is a cubesat with faces mostly covered by solar panels and antennae. Frith noted strong glints spaced approximately 40 seconds apart that occurred simultaneously in each of the three colour bands; presumably the result of specular reflections off the faces of the satellite as it rotated. Of particular note, Frith also observed that the colour ratios changed very rapidly during the glints; becoming significantly more red. It was stated that the implication of the changing B-V values is that specular and diffuse reflections have different colour ratios, however Frith did not explicitly extend the argument to suggest that a given material's SED can vary with changing illumination and observation geometry.

Formosat 3D is a much more complex satellite than LCS-1 and Latinsat-A, and the lightcurves and colour plots were correspondingly much more complex. No attempt was made to analyse the results from Formosat 3D. The authors conclude that colour ratios might be able to be used to differentiate between specular and diffuse reflections, but that more work is required to better predict an object's colour index change as a function of BRDF. This work does not appear to have been followed up to date.

2.3.5   Surface Material Characterisation

One of the desirable pieces of information about a satellite is the type and relative abundance of materials that are present on its surface. In 2010, Hall described a method for characterising the surface materials of a satellite [4], which he expanded upon with the addition of a second, slightly different, method in 2012 [38]. In the context of obtaining information about an unknown satellite, both of the methods are particularly suited to LEOs because they require a detailed wireframe model and attitude information as inputs; which cannot be obtained using remote sensing methods at ranges beyond LEO.

The first method, termed the Material Abundance Estimation (MAE) method, requires a library of BRDFs of candidate materials that are commonly used on satellites. Using a wireframe model of the spacecraft, which may have been derived from optical or radar observations, it is represented as a collection of flat facets arranged in such a way as to reconstruct the known shape of the satellite. If accurate attitude information can also be obtained at any given time, then the illumination and observation geometry for each facet can be determined. A guess is then made as to the relative abundance of each material that each facet is made of. Using the known BRDF of each of these materials, and their areas, their contribution to the total reflection can be calculated. Summing the contribution of each facet gives the total calculated reflected intensity. The process is then repeated using a different set of guesses for the relative abundance of materials for each facet. A correct guess for each facet should yield a calculated total magnitude that, in the absence of any errors, is equal to the measured magnitude. The entire process is repeated for each wavelength band that a measurement was made in. A least squares method is then used to find the best fit fractional areas of each satellite component covered by each material.

Hall used two computer simulated satellites, a cubesat and a more complex satellite with solar panels and a protruding cylindrical component, using scenarios both with and without random noise, to test whether the correct surface materials can potentially be extracted. Each face of the cubesat, and each major component of the complex satellite, was comprised entirely of one material. Hall used a variety of scenarios in his simulation, including: a rotating satellite where all sides are seen; a stabilised satellite where some faces are never seen; a BRDF library containing more candidate materials than the number of materials used on the satellite; and a satellite partially comprised of one material that is not one of the candidate materials in the BRDF library. Hall concluded that in the perfectly noise free simulations the component materials were successfully identified, but that the model became less accurate with the introduction of noise, and when one of the materials was not included in the BRDF library.

No analysis was conducted into the effect of errors in attitude estimation or the

difference between the wireframe model and the actual satellite structure. Also, the satellite is modelled as a set of flat facets, which imperfectly simulates rounded surfaces or crinkly insulation. In addition, the success of the model relies upon candidate BRDFs being significantly different from each other such that materials can easily be distinguished. If each material's BRDF is very similar, then a small error in measured magnitude may lead to large errors in estimated material abundance. Hall also noted that the effects of space weathering are not able to be taken into account since the BRDF library is constructed from laboratory measurements of non-weathered materials.

The second method described by Hall is termed the Reflectance Distribution Estimation (RDE) method [38]. The RDE method differs from the MAE method in that a library of candidate BRDFs is not required. Instead of using known BRDFs to estimate the relative abundance of materials, the RDE method attempts to estimate the BRDF of each satellite component. An expansion function models the BRDF of each satellite component as the weighted sum of multiple reflection modes ranging from perfectly diffuse to perfectly specular. RDE relies upon the availability of a very large data set of measurements, taken in multiple bands, and over a wide range of different geometries. Using the available dataset, RDE attempts to solve for the best fit albedos of constituent components. Hall used a computer simulated satellite with solar panels and protruding component again to test the performance of the RDE method. It was discovered that RDE is heavily reliant upon a large quantity of measurements, taken under a diverse range of geometries, to converge to a sufficiently accurate result. Even using eight days worth of simulated data gained from two geographically displaced observation sites, comprising eighteen seperate passes of a rotating satellite, the method still produced somewhat imperfect results. In addition, as with the RDE method, the effect of errors in attitude estimation or the difference between the wireframe model and the actual satellite structure were not investigated.

Hall concluded that the MAE method suffers from limitations of the BRDF database, making it inappropriate for unknown or ageing satellites. The RDE method

overcomes this problem but requires a prohibitively large number of observing sites and/or observation passes, and ideally space based sensors, to converge to sufficiently accurate results.

In 2012 a team led by Lederer conducted a proof of concept study to compare the relative merits of photometry and spectroscopy as tools for identifying spacecraft and debris surface materials [35]. The study consisted of four distinct data collection efforts: photometric measurements of satellites; spectrometric measurements of satellites; laboratory based photometric measurements of solar cells; and laboratory based spectrometric measurements of solar cells. For the study Lederer chose a set of 18 satellites that were launched into slightly sub-synchronous orbits in the 1960s, called Initial Defense Communication Satellite Program (IDCSP) satellites. The IDCSP satellites were chosen for their structural simplicity: they are 26 sided polygons with all sides covered in solar cells. The satellites are no longer operational and thus are not stabilised.

Photometric data were collected on the IDCSP satellites over four consecutive nights in the Johnson-Cousins BVRI bands. Between 10 and 20 measurements were taken of each satellite in each band, and an average of these measurements, weighted by their respective errors, was taken as a representative magnitude for each satellite in each band. Importantly, the measurements were not taken simultaneously through each of the colour filters. Instead, all of the measurements through a particular filter were taken consecutively before cycling to the next filter. Thus the sun-object-sensor geometry may have changed substantially between measurements in different filters; particularly because the rotation rate of each of the satellites is unknown. Lederer noted that the variation in the measured magnitudes for each satellite was often far greater than the calculated uncertainty of each individual measurement, and therefore the standard deviations of the measurements in each band were also calculated. Despite the IDCSP satellites being almost identical to each other, the average brightness in the V band varied between 13.689 and 17.645 magnitude, and the standard deviation in B-V colour indices was as much as 0.364 magnitudes. Lederer noted that variability in colour could be expected if a satellite's lightcurve varied with rotation

angle, however a thorough investigation into the magnitude of the variability for each satellite was not conducted. Nonetheless, the average B-R colour ratio of each satellite was plotted against the average B-V colour ratio, with error bars representing the standard deviation of each satellite's colour ratios. On the same plot, laboratory data of five different types of solar cell were shown for comparison to the telescopic data. Three of the solar cells had their colour ratios derived from spectroscopic data, one used photometric data, and one solar cell had both photometric and spectroscopic data measured. The laboratory data were not presented with associated error bars, and the number of laboratory measurements taken and the associated illumination and observation geometry were not stated. Lederer noted that the laboratory spectral data only measured diffuse reflections, whereas the laboratory photometric data measured both specular and diffuse reflections, and therefore there may be some differences in measurements using both techniques. Three of the five solar cells appeared in generally the same area on the plot as the telescopic data.

An example plot of spectral reflectance over the wavelength range 4500 Å to 8000 Å of one of the IDCSP satellites was presented for comparison to the laboratory measured spectral reflectance of the solar cells over the same range. It was noted that there was an increase in reflectance at the longer wavelengths for the satellite data. This was suggested to be the result of either reddening due to space weathering or increased noise in the longer wavelength data. None of the four solar cells' measured spectra closely matched the IDCSP spectra, however it was noted that the solar cells covering the IDCSP satellites were manufactured in the 1960s, whereas the sample solar cells were manufactured in the 1990s. Therefore, a perfect match was not expected.

Lederer concluded that it is difficult to tell which solar cell is most similar to the satellite's solar cells using either photometry or spectroscopy, and that the effort is hampered by changes due to space weathering and noise. It was suggested that photometric data must be gathered simultaneously through each colour filter to obtain more accurate results.

2.3.6 Seasonal Effects on Lightcurves

In November and December of 2010 Hall obtained multiple, detailed, long duration lightcurves of the GEO satellites Galaxy 12 and Galaxy 15 in the I band, using a small aperture telescope [42]. Hall studied these lightcurves to determine how sequences of faint glints that repeat from one night to the next can contribute to satellite characterisation. Galaxy 12 was observed on five out of eight consecutive nights, and Galaxy 15 was observed on nine out of fourteen consecutive nights. Upon initial inspection of the lightcurves Hall noted that in many cases small glints recurred on several consecutive nights at almost exactly the same time, but that the nature of the repetition was not the same for all glints. For example, some glints appeared on one night and disappeared several nights later, whereas other glints recurred on all of the nights that the satellites were observed. Also, sometimes there were subtle changes in the brightness and timing of the glints, and sometimes the glints were of equal magnitude on each night. Hall hypothesised that changes in the faint glints were the result of season driven illumination changes due to the changing position of the sun with respect to the satellite and observer. The fact that they recur at nearly the same time and magnitude indicates that they originate from the same satellite component each night.

One particular glint that was observed on eight nights was examined in detail. The time at which this glint occurred was observed to change incrementally each day, by approximately five minutes in total over the course of the fourteen day period. To determine where on the satellite the glint originated from, Hall plotted the glint sequence as a function of solar phase angle, equator projected phase angle, the azimuth angle offset from the phase angle bisector in the solar array reference frame, and the azimuth angle offset from the phase angle bisector in the satellite body reference frame. In all cases except for the body reference frame case the glints did not line up on each night. Figure 2.9 depicts an expanded view of the lightcurve from just prior to the glint to just after the glint, as a function of solar phase angle on the left, and as a function of the azimuth angle offset from the phase angle bisector in the

satellite body reference frame on the right. Each subsequent night's lightcurve has been shifted up in brightness by 0.25 magnitudes over the preceding one to allow for unobstructed inspection of each one.
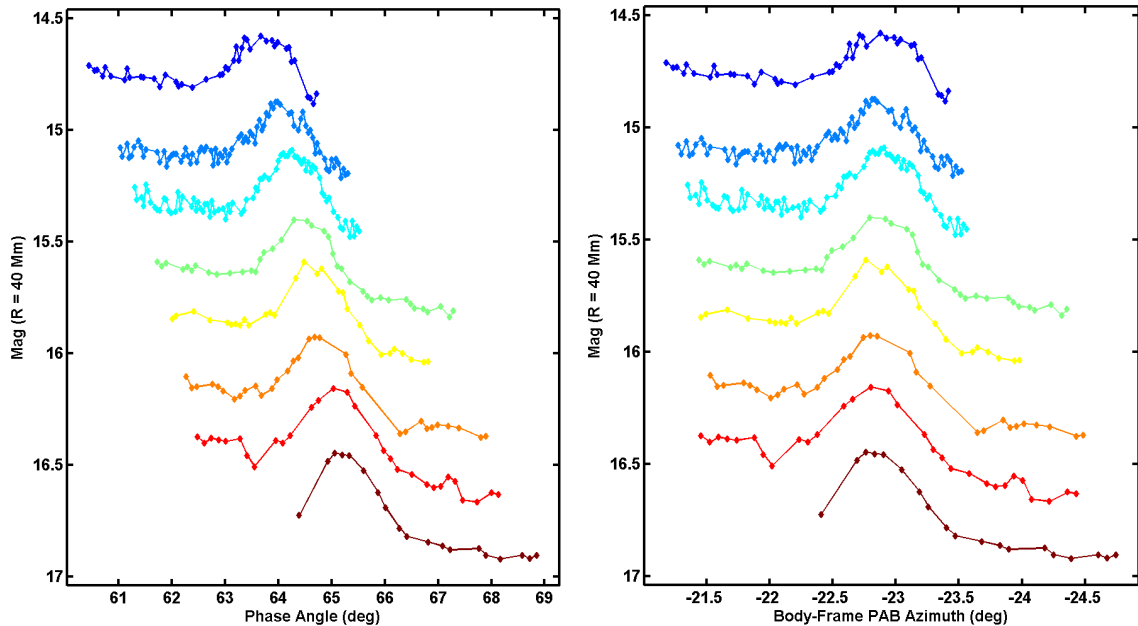
Figure 2.9: A portion of a series of lightcurves taken on eight nights over a two week period showing the same small glint off Galaxy 15 as a function of solar phase angle (a), and body frame phase angle bisector azimuth angle (b) [42].

As can be seen, when plotted as a function of the azimuth angle offset from the phase angle bisector in the satellite body reference frame, the glints line up at approximately 23 degrees. Hall suggests that this indicates that the glints originated from the same satellite component each night, which must be part of the satellite body. Further, the angle at which the glint occurs constrains the geometry of the component that is glinting.

2.4   Unresolved Issues

This literature survey has highlighted that there has not been a consistent progression in the understanding of how photometry can contribute to SSA. Researchers have used multiple different methods, techniques and assumptions to investigate a range

of disparate ideas; however a strong foundation of basic principles has not been established. Therefore, it was considered necessary to take a step back and re-examine some of the basic understandings that currently exist regarding photometry of geostationary satellites. As a result, this experiment set out to answer the following questions:

- Can geostationary satellites be categorised according to a set of standard lightcurve classes? (Section 6.1).

- Is it possible to detect evidence that a SED varies with illumination and observation geometry in lightcurves or colour plots? (Section 6.2).

- Can colour ratios be used to quickly and easily characterise or discriminate between different geostationary satellites? (Section 6.3).

- Is it possible to use basic models to simulate the structure of satellites during the conduct of an experiment without substantially affecting the conclusions of the experiment? (Section 6.4).

- To first order, what is the typical minimum temporal resolution required to enable all significant information to be obtained from a satellite's lightcurves? (Section 6.5).

Chapter 3

Experimental Setup and Procedure

This chapter describes in detail the equipment that was used, and the procedures that were followed, to gain photometric data on six GEO satellites. Hardware and software requirements are discussed in terms of both collecting data and processing it. Several challenges were encountered in the conduct of the experiment, including partial equipment failures, light pollution, and technical limitations of the equipment. In the end, despite these difficulties, satisfactory results were obtained, successfully demonstrating the viability of using inexpensive small aperture telescopes to conduct photometric analyses of Earth orbiting objects, for which valuable time on larger, meter-class, telescopes has often been used to this point [29, 43].

3.1 Scope of the Experiment

Significant progress has been reported over the last decade regarding the potential for photometry to reveal details about artificial Earth orbiting objects [7, 38, 40, 44]. Of note, however, is that much of the photometric analysis conducted to date has suffered from limitations in the data sets. Examples of these limitations include: data only being obtained on very few different nights for given satellites, low temporal resolution in lightcurves, and large gaps in lightcurve data. As a result, questions remain as to how well all possible variables have been considered in the analysis of reflected intensities. For example, Bedard (2013) has demonstrated that the sun/object/sensor geometry involved in analysing the photometric or spectrometric signature of a satellite is probably far more important than has previously been assumed [39]. Whilst the longitude dependent phase angle has been used extensively to compare photometric signatures, the specific orientation of the satellite has not always been considered. Therefore, it was decided that it was necessary to return to first principles and collect high cadence, long period observations (i.e. over the course of an entire night) on

a limited number of satellites, over as many nights as possible. This approach was taken to ensure that undersampling of data was less likely to lead to inconclusive results or erroneous conclusions. In addition, in order to isolate as many variables as possible, it was decided that only geostationary satellites would be observed. Advantages of geostationary satellites, in relation to this experiment, include the fact that they can be observed for an entire night through an unchanging airmass; and they usually employ three-axis stabilisation, which constrains their orientation relative to the observer.

3.2 Experimental Setup

Throughout the experiment, observations were taken using a 35.6 cm Celestron CG-14 telescope fitted with an Apogee Alta U42 camera. The telescope is mounted on a Paramount ME mount, and includes a temperature compensating focusser and a five slot filter wheel. Standard Bessel BVRI filters were used in the filter wheel [45]. The most widely used photometric system to date, of the hundreds that have been created, is the Johnson-Cousins UBVRI system [13]. Bessel filters are manufactured to very closely match the Johnson-Cousins UBVRI passbands such that measurements made using Bessel filters can be placed on the Johnson-Cousins scale.

The Apogee Alta U42 has a 2048 x 2048 pixel CCD. Electrons are captured by CCD pixels when impinging photons knock them from the silicon from which the CCD is constructed. However, not every photon that arrives at a pixel is converted to an electron. The fraction of photons that are converted to electrons is wavelength dependent, and is known as the CCD's Quantum Efficiency (QE). Figure 3.1 illustrates the QE as a function of wavelength for the Apogee Alta U42 CCD chip . The midband CCD, represented by the red line, was used to conduct photometry during this thesis.

When a CCD image is read out, the charge in each pixel is converted into a number by an Analogue to Digital Converter (ADC). This number has units known as Analogue to Digital Unit (ADU)s, and is stored in computer memory as the value of its associated pixel for future use. There are two points to note regarding ADU
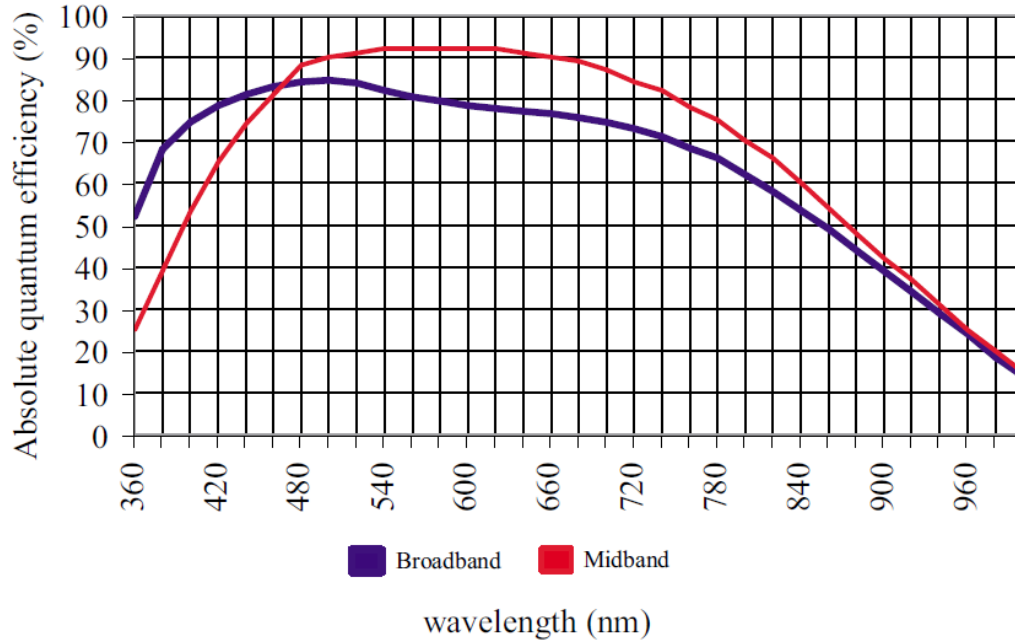
Figure 3.1: Apogee Alta U42 Quantum Efficiency [46].

values. First, the ADU values generated by the ADC are binary numbers. The range of possible values that can be generated by the ADC is therefore limited by its bit depth. For example, a 16 bit ADC can generate 2^{16} (65536) distinct values. The actual range of ADU values in that case would be zero to 65535. Second, a property of the ADC known as its *gain*, which is a manufacturing setting that can sometimes be altered by the user, determines the number of electrons required to increase a pixel's value by one ADU. For example, a gain of 1.5 would mean that 98303 electrons are required to generate a pixel value of 65535 ADU. The Apogee Alta U42 used in this experiment uses 16 bit digitisation.

The entire telescope and camera setup is housed in a dome on the roof of one of RMCC's five-floor buildings. Figure 3.2 is a photograph taken from within the dome showing the telescope on its mount, with the U42 camera and filter wheel attached.

The camera, telescope and dome are controlled remotely from a computer in the SSRAL, operating The Sky X Pro, CCDSoft 5, Automadome, and Orchestrate for automation, all by Software Bisque. Figure 3.3 depicts the workstation and control room.

Figure 3.2: Image of the telescope used in the experiment.

3.3   Data Acquisition

Twelve nights of observations were conducted, excluding nights where the data obtained were not of sufficiently high quality, on six GEO satellites from 27 July 2013 through to 20 November 2013. Since all of the imaged satellites were part of GEO clusters, each image contained more than one satellite on all except for two nights. As a result, every satellite was imaged on between three and eight nights. Table 3.1 details the satellites that were observed and how many nights of imagery were obtained for each. The satellites in GEO cluster one are all located at 55.5° West [47], and the satellites in GEO cluster two are all located at 107.3° West [48].

In selecting the nights on which to conduct observations, the primary criterion was that the sky should be forecast to be completely clear of cloud until after sunrise. This decision was made because, since the experimental setup did not include an infrared camera capable of showing the presence of clouds, doubt could otherwise arise as to whether variations in brightness were the result of clouds or true variations in the signal. Consideration was also given to phases of the Moon and the proximity of the

Figure 3.3: Photograph of the SSRAL.

Satellite	Bus Type [49]	Nights Imaged	GEO Cluster
Galaxy 11	BSS702C	8	One
Intelsat 805	AS-7000	6	One
Anik F1	BSS702C	4	Two
Anik F1R	Eurostar-3000S	4	Two
Anik G1	LS-1300	4	Two
Echostar 17	LS-1300	3	Two

Table 3.1: Number of nights that each satellite was observed.

Moon's orbit to the satellites being imaged. One night's data were discarded entirely because the nearly full Moon passed sufficiently close to the satellite being imaged to significantly affect the measurements.

3.3.1 Detailed data acquisition procedure

The first task that was performed each night, after powering up and initialising all of the equipment, was to set the CCD temperature and wait for it to cool down. CCDSoft 5 provides a continuously updated indication of the current effort of the

camera's thermoelectric cooling system as a percentage of its maximum effort. On each night, the CCD temperature was selected such that the camera cooling system was required to operate at approximately 80 percent of its maximum effort to maintain the set temperature. The temperatures that were set throughout the experiment ranged from -17 ℃ to -30 ℃.

Before imaging the satellites of interest, bias, dark and flat field frames were taken to allow for image reduction. Because the CCD temperature affects the CCD response to bias, dark and flat field images, and since new dust particles on the telescope optics can change the flat field frames from night to night, bias, dark and flat images were obtained on each night that data were obtained. All of the images on a given night were taken at CCD temperatures within one degree celsius of each other, and usually within 0.1 degree of each other. The following paragraphs describe the process followed for obtaining bias, dark, flat, standard star, and data images, in chronological order.

On each night of observations 10 twilight sky flat frames were taken per filter. As the sun sets, longer exposure times are required to maintain a given ADU count per pixel. However, at exposure times of less than 1 second effects related to the finite time taken for the shutter to open and close begin to become noticeable, and at exposure times greater than about 2 seconds many stars start to become visible in the images. Therefore, there is a very narrow time window during which twilight flats can be taken; on the order of about 5 minutes. As a result, and due to the different CCD response to the different filter bands, flats were always taken in the order B, V, R, I, and then unfiltered. All of the flats taken throughout the experiment used either 1 second or 1.5 second exposure times, resulting in an exposure level of approximately 20 000 ADU to 30 000 ADU per pixel. To take the twilight flats the telescope was pointed near zenith, approximately 20 degrees in the opposite direction to the setting sun, where the sky is most evenly illuminated. Figure 3.4 is an example of an unfiltered flat frame taken on the 28[th] of October 2013. Some stars are visible in the image, however they will be eliminated by median combining with the remaining nine unfiltered flat frames.
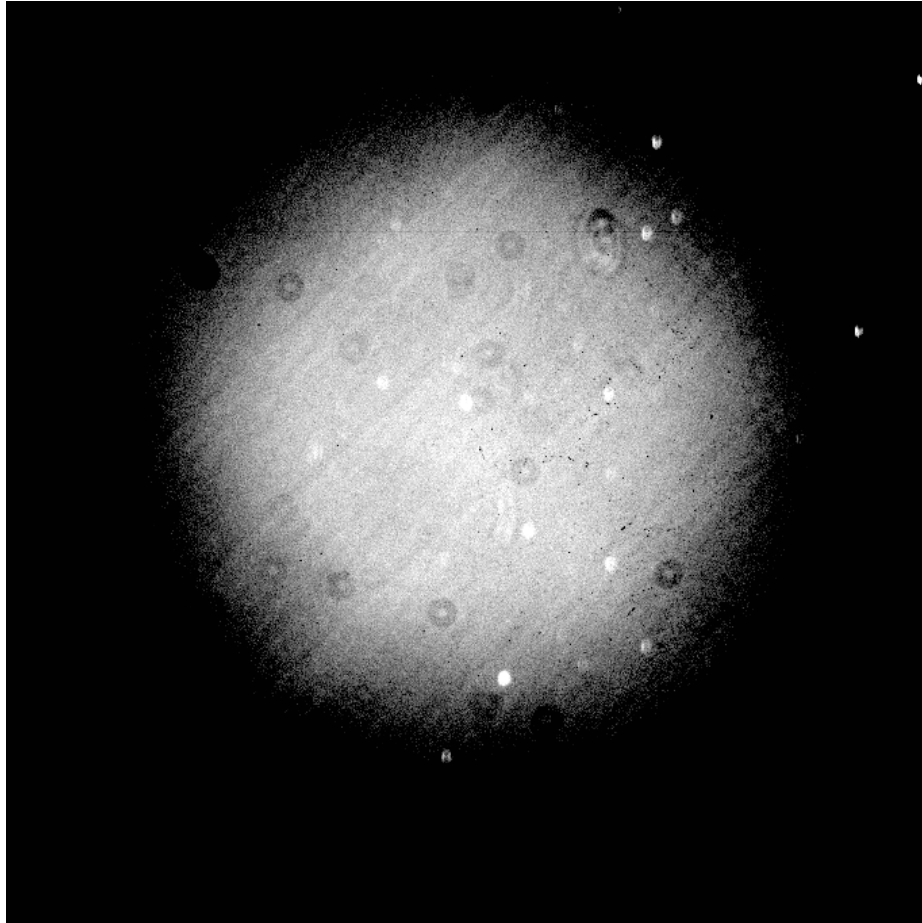
Figure 3.4: Example flat field image taken on 28 Oct 2013. The bright spots are stars.

After obtaining the flats, 25 bias frames were taken, followed by 10 dark frames of 1 minute duration, which is at least as long as each of the data frames. Multiple bias, dark and flat images were taken, rather than only one of each, so that median master frames could be produced later. Producing median combined bias, dark, and flat frames allows extreme ADU values, such as those resulting from cosmic ray hits, to be filtered out. Generally, the more images used to produce master frames the better, however only 10 sky flats were taken due to the limited time available at twilight, and only 10 dark images were taken to maximise the time spent collecting data on the satellites. Bias and dark images were taken well after sunset to guarantee that they would not be contaminated by daylight. Even though the shutter remained closed for these frames, stray photons could still find their way onto the CCD if they were taken

during the day. Figure 3.5 is an example dark frame taken on the 20^{th} of November 2013. A small bright dot is visible near the top centre of the image, which is likely due to a cosmic ray hit. Such events are not uncommon and they will be removed by median combining with other dark frames.
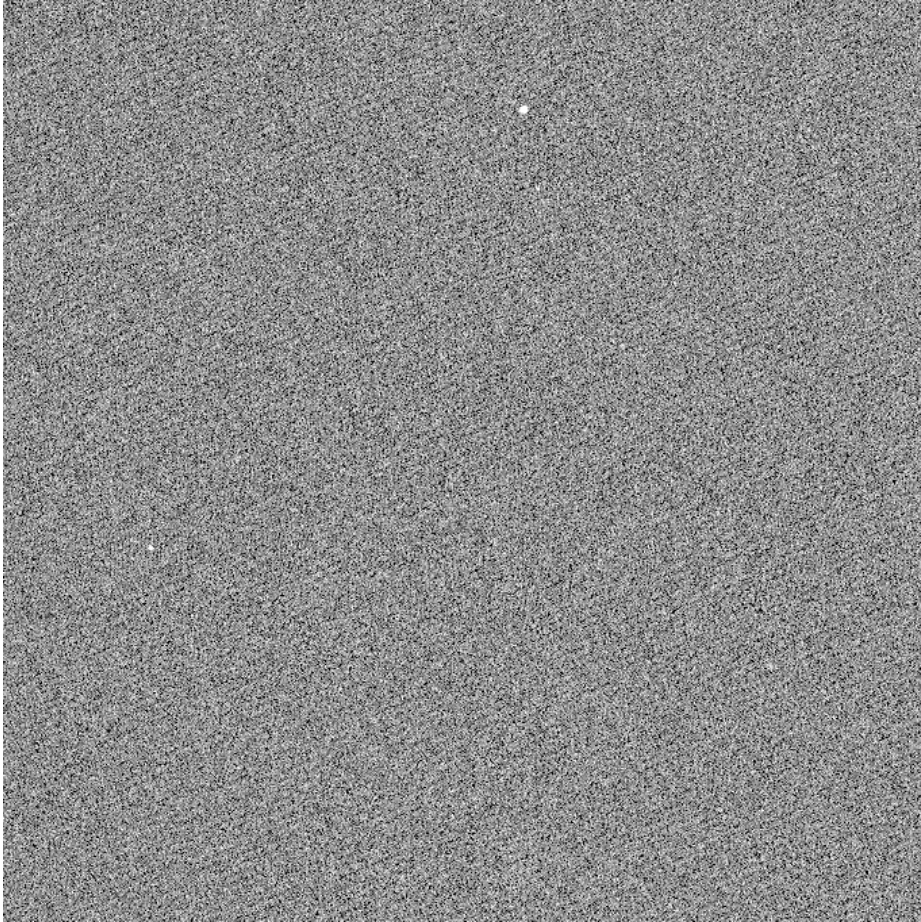
Figure 3.5: An example dark frame taken on 20 Nov 2013.

In addition to the bias, dark, and flat frames, images of Landolt standard star fields were taken either at the start or end of the night to enable the calibration of instrumental magnitudes to the standard apparent magnitude scale. Since the satellites that were observed were all GEOs, they were continuously observed through an unchanging airmass each night. Landolt stars, however, continuously move across the sky due to sidereal motion. Therefore, the Landolt stars were observed for a period each night, usually at the beginning of the night, over a range of airmasses that includes the airmass through which the satellites were observed. The full image

reduction and calibration process, using the Landolt images, is described in detail in Section 3.4.1. Once a Landolt star field had been selected, the telescope was slewed to a region of relatively high star density within that Landolt field and sidereal tracking was turned on. A test image was then taken to assess the number of stars that would be captured in the frame. The telescope pointing was then adjusted, followed by further test images and adjustment, until optimum telescope pointing was achieved. The main criterion used to determine optimum telescope pointing was that the number of Landolt stars captured in each image should be maximised. Also, brighter stars were preferred such that the SNR would be maximised; however, a single bright star was not preferred over multiple fainter stars. In addition, care was taken to ensure that very bright stars that might contaminate other nearby stars (through saturation or diffraction effects for example) were not captured in the frames. Once the desired telescope pointing had been achieved, one test image was taken through one of the filters using an exposure time estimated to achieve sufficient exposure depth. The brightest pixel in each of the Landolt stars was then measured, and the exposure time was adjusted if required. The exposure time was selected to achieve the maximum number of Landolt stars in the image with between 10 000 and 50 000 ADU for their brightest pixels. The process was then repeated for the rest of the filters and for the unfiltered case. Once the desired exposure times had been determined they were programmed into Orchestrate. Orchestrate, which was programmed to cycle through each filter taking one image per filter per cycle, was then set to run. As images were being acquired, the airmass was regularly monitored by examining the image headers in CCDSoft 5 to determine when the airmass through which the satellites would be observed fell within the range of airmasses through which the stars were observed. Figure 3.6 is an example of a Landolt star field image taken through the I filter on the 20[th] of November 2013.

Once the Landolt star images were obtained the system was set up to take the data images. This involved slewing the telescope to the GEO cluster of interest and turning off sidereal tracking. A test image was then taken to ensure that all of the desired satellites would be captured in the frame. The imaging sequence was programmed
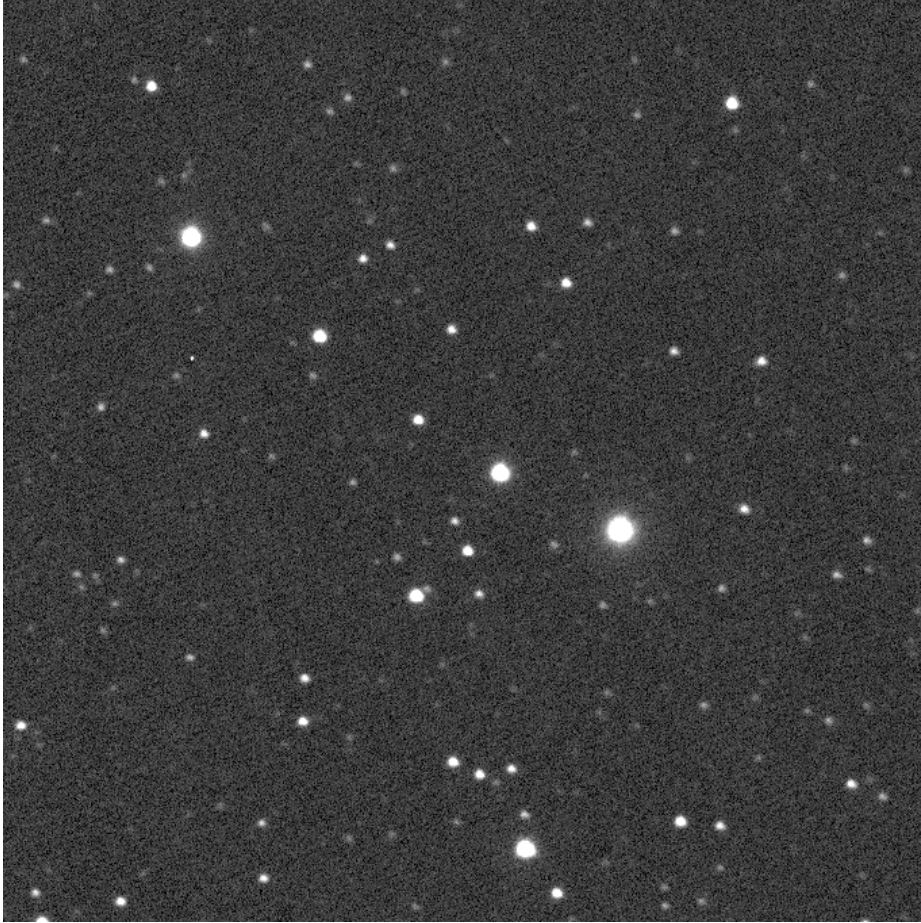
Figure 3.6: Example Landolt star field image taken on 20 Nov 2013.

into Orchestrate, which was again used to automate the process of cycling through the filters and taking images at the desired exposure time. Only one image was taken through each filter before changing to the next filter. Although Orchestrate significantly reduced the effort required to obtain the data, it remained necessary to closely monitor the images as each one was obtained, due to the rapid changes in brightness that were commonly encountered during satellite observations. Figure 3.7 shows how quickly satellites can increase in brightness during specular glints.

Throughout each night a signal level of between 10000 ADU and 50000 ADU was attempted to be maintained for the brightest pixel illuminated by the satellites being observed. When the brightness was seen to increase markedly it was necessary to reduce the exposure time to avoid saturating any of the pixels; and when the brightness decreased the exposure times were increased to maintain sufficiently high

(a) Image showing all four satellites not saturating the CCD.

(b) Image showing Echostar 17 (bottom of image) saturating the CCD.

Figure 3.7: Two consecutive images taken through the V filter on 12 Oct 2013 that demonstrate how rapidly satellites can increase in brightness. Both images show satellite cluster two, with Anik F1, F1R, and G1 at the top of the image, and Echostar 17 at the bottom of the image. The image on the left used a 20 second exposure, and the image on the right was taken using a 3 second exposure, after first cycling through the other filters. The image on the right shows Echostar 17 now significantly saturating the CCD less than 2 minutes after the image on the left was taken, despite the reduction in exposure time.

SNRs. This effort was made significantly more complicated whenever more than one satellite was being observed in each frame, which occurred on all except for two nights. In these cases the satellites in the frame often showed significant differences in brightness. Whilst great effort was made to maintain each satellite's measured signal between the limits stated above, sometimes it was not possible to do so. When this occurred, one satellite was selected to be the primary satellite of interest. The primary satellite was always maintained within the signal limits, which occasionally resulted in either saturated pixels or low SNR for the other satellites. Any satellite observations containing saturated pixels or excessively low SNR were later discarded.

Typically, exposure times at the beginning of a night were as follows: 45 to 60 seconds for the B filter, 20 to 30 seconds for the V filter, 5 to 15 seconds for the R and I filters, and 2 to 5 seconds for the unfiltered case. The exposure times usually

only decreased minimally until just prior to the main specular glint, when they would quickly drop to less than one second in each filter. After the specular glint the exposure times were returned to values similar to those used at the start of the night. On nights when the main specular glint was not particularly bright the exposure times for all filters never dropped below approximately 1 to 5 seconds for all filters.

3.3.2 Data Collection Challenges

A number of challenges were encountered in obtaining accurate data, with minimal systematic errors, throughout the experiment. These challenges included large, rapid changes in the brightness of each satellite, light pollution from RMCC buildings and the city of Kingston, frequent inclement weather, and technical difficulties with the equipment. The considerations taken into account, and methods used to deal with these challenges, are described in the following subsections.

3.3.2.1 Maximising Temporal Resolution Without Saturating Pixels

Competing priorities existed throughout the data collection process between attaining the highest possible temporal resolution and maintaining a sufficiently high signal level in each image, without saturating any pixels. To obtain the highest possible temporal resolution, each image should be well focussed such that the signal received from the satellite is concentrated on as few pixels as possible. This will maximise the signal being recorded by each pixel receiving light from the satellite per unit of time. Thus, the exposure time for each image can be minimised, in turn maximising the temporal resolution of the data. During glints, however, the brightness of some of the satellites increased so rapidly that it was difficult to predict what would be an appropriate exposure time for each successive image. In addition, at maximum brightness, to avoid pixel saturation the exposure time was occasionally required to be reduced to below 20 milliseconds, which is the minimum recommended according to the camera manufacturer's datasheet [46]. The brightest specular reflection observed from Galaxy 11 saturated some pixels in unfiltered images even at exposure times

as low as one millisecond. It was also noticed during the data analysis procedure that images with exposure times less than 20 milliseconds yielded unusually high apparent magnitudes. At the point on the lightcurve where the exposure time dropped below 20 milliseconds there was a noticeable, large step change increase in magnitude. There was subsequently a corresponding decrease in magnitude at the point on the lightcurve where the exposure time increased beyond 20 milliseconds. This was a clear indication that exposure times shorter than 20 milliseconds should be avoided. To achieve exposure times longer than 20 milliseconds even during bright specular reflections, without saturating any pixels, the telescope was occasionally defocussed to distribute the light over more pixels. The result is a donut shaped Point Spread Function (PSF) with a broader, shallower distribution of pixel values than would be seen in a focussed image. All images with exposure times less than 20 milliseconds were discarded.

On-chip pixel binning can also be used to increase temporal resolution. Binning reduces the spatial resolution of the images by combining groups of pixels, as they are read out, to form larger 'super pixels'. The readout time, however, is reduced by a factor equal to the number of pixels contained in each binned pixel. For example, 3x3 binning combines each group of nine pixels (in three pixel by three pixel squares) to form larger, binned pixels. As a result, the readout time is reduced by a factor of nine. Unfortunately, binning reduces the effective dynamic range of each pixel. This is because the electrons in each pixel contribute to their binned pixel's ADU value. For example, using 3x3 binning and 16 bit digitisation, the maximum value of 65535 ADU would be reached for a binned pixel when its individual constituent pixels each average only 7282 ADU (one ninth of 65535). The implication is that binning can increase the temporal resolution, however it potentially comes at the expense of exacerbating the previously described issues regarding short exposure times because the pixels saturate more quickly.

All of the images taken throughout the experiment, with the exception of the final night's data, were taken using 3x3 binning in order to minimise the readout times. Using 3x3 binning the readout times were on the order of five seconds, whereas

unbinned images would take approximately 45 seconds to be read out. The three satellites imaged on the final night of data collection were all located within close proximity of each other on each CCD frame. Therefore, instead of using 3x3 binning, CCDSoft's subframe function was used to crop each image before being read out, without binning, maintaining the temporal resolution attained on previous nights but with nine times the dynamic range. Thus, difficulties associated with short exposure times were avoided whilst also maximising the temporal resolution.

3.3.2.2   Light Pollution

Light pollution is always a consideration when conducting photometry. Even at dark sites, well away from city lights, light pollution levels can vary from one night to another. The Moon is probably the most obvious source of unwanted light in such situations; and even when it is located well away from the object of interest atmospheric scattering can increase the observed background noise level.

The location of the telescope used to conduct this experiment, on the roof of a five-floor building at RMCC, is far from ideal. Kingston's brightly lit Central Business District is clearly visible across the causeway from the telescope's location, and is situated in the general direction of the GEO satellites that were observed. In addition, lights from many buildings within RMCC were permanently operating. The general location of RMCC, next to Lake Ontario, was also a concern given the potential for lake effects to increase the density of particles in the atmosphere, resulting in a greater degree of light scattering and absorption.

Mitigating the light pollution concerns is the fact that the telescope is located in a dome with a retractable slit. The dome therefore blocks much of the unwanted light that does not originate from the general direction in which the telescope is pointing. Aside from background subtraction, which is described in section 3.4.1, no other attempt was made to reduce the effects of light pollution, which remains a source of systematic error in the data.

3.3.2.3 Technical Difficulties

Several technical difficulties were encountered that contributed to the loss of some data. The first problem that was discovered was that the telescope often failed to align properly with the slit in the dome. Upon startup the telescope would be aligned well with the slit, however the alignment was seen to progressively deteriorate with successive slews of the telescope such that the edge of the slit began to obscure the telescope's line of sight. With the assistance of RMCC staff some minor modifications were made to the system used by the dome to keep track of its alignment. The modifications proved to be effective and no further misalignment problems were experienced.

The next problem involved the partial failure of the camera cooling mechanism. Close inspection of the data obtained on several nights, early in the experiment, revealed that the camera was not maintaining a consistent CCD temperature. Initially the cooling system appeared to work well, with the temperature gradually dropping until the set temperature was reached and initially maintained. After approximately one hour, however, the temperature began to rise again, and was generally unstable beyond that point, commonly reaching well above 10 ℃. Troubleshooting revealed that although the camera's thermoelectric cooling system was working, the four small fans located on the back of the camera always ceased to operate after approximately one hour. The problem was rectified by placing a portable fan in the dome, located such that it directed airflow towards the rear of the camera. This proved to be a very effective solution, and the camera was thereafter easily able to maintain the set temperature. However several nights' data obtained prior to the problem being identified needed to be discarded due to the inconsistent and extremely high dark current.

Unfortunately, another temperature control problem continued to manifest itself throughout the experiment. For unknown reasons Orchestrate would routinely cease functioning once or twice per night. Whenever this occurred, CCDSoft 5 ceased to recognise which filters were in the filter wheel, and therefore Orchestrate could not be

immediately made to restart collecting data. The problem could only be resolved by disconnecting CCDSoft 5 from the camera, then reconnecting, re-entering the filters into the filter wheel slots, and re-setting the temperature controller. Unfortunately as soon as the camera was disconnected, the temperature immediately and rapidly began to rise. Lowering the temperature back to the selected value was a much slower process, however, and up to half an hour of data were lost each time whilst waiting for the temperature fall again. With practice, the recovery procedure was able to be completed quickly enough such that only about five minutes of data were lost on each occurrance.

3.3.2.4 Other Difficulties Encountered

Some challenges that were encountered during the data collection phase were completely unavoidable. The most apparent example was the weather. Obtaining high temporal resolution data throughout an entire night generally required completely clear conditions all night. Such conditions proved to be a rarity during this experiment, necessitating that all data collection opportunities be taken advantage of. Data were successfully obtained on twelve nights between July 2013 and December 2013, and there were very few opportunities, apart from those twelve nights, that were missed.

Section 3.4.1.3 described how short exposure times, particularly below half of a second, introduce an uneven exposure across each frame, caused by the finite time required for the shutter to open and close. Exposure times of less than half of a second were necessary over portions of most nights' observations during this experiment to avoid pixel saturation when bright specular reflections were observed. Therefore, this shutter effect is an unavoidable source of systematic errors.

3.4 Image Processing

Image processing was conducted using Mirametrics' Windows based Mira Pro Ultimate Edition software program (Mira Pro UE), hereafter referred to as Mira. The

image reduction procedure involved subtracting bias and dark frames and applying flat field corrections to the data images. Instrumental magnitudes were calculated by using Mira's aperture tool that measures the ADU count within a circular synthetic aperture. Landolt star fields were used to calculate the nightly zero point, which was added to each satellite's measured instrumental magnitude to arrive at the apparent magnitude. All magnitude data were exported from Mira to Microsoft Excel, and then input to Matlab for plotting. Each process is described in detail in the following subsections.

3.4.1   Image Reduction Principles

CCD photometry would be considerably more simple to conduct if the light signal from the object of interest was the only contributor to the number of electrons read off the CCD, and if there was no variation across the CCD in terms of response to incident light. Unfortunately this is far from true. In reality the total number of ADU recorded on a pixel depends upon manufacturing settings, manufacturing irregularities, temperature, the optics used, background light, cosmic rays, and a variety of systematic inexactitudes. To avoid large errors from reducing the quality of the data these effects need to be compensated for by adjusting the values recorded on each pixel in an appropriate way. Compensating for unwanted effects is known as image reduction. The main image reduction processes are bias subtraction, dark current subtraction, and flat fielding, which together usually eliminate the majority of the systematic errors in the data.

3.4.1.1   Bias Subtraction

If an image is taken with zero exposure time (by clearing built up charge on the CCD and then immediately reading it out) the ADU values recorded for each pixel will not be zero. Contrary to what might be expected, each pixel will record some positive value, which is usually fairly consistent across the CCD. An image taken in this way is known as a *bias* image. The reason that bias images record positive pixel values

rather than zero ADU for each pixel is that CCDs are manufactured to apply a small electronic offset that is added to each image. Without a positive bias offset, a zero exposure image will result in a Gaussian distribution of pixel values centred on zero. The shape of the Gaussian distribution will depend on the readout noise [50]. As a result, some pixels will be recorded as having negative values. The bias offset avoids these negative values.

Another reason that manufacturers may want to apply a bias offset is to avoid the non-linear response to incident light that CCDs commonly exhibit at low exposure levels. One of the key advantages of using CCDs for photometry is that the number of electrons captured by a pixel has an almost exactly linear relationship to the number of incident photons. This makes it very easy to measure light intensities and to compare different light sources. However, the linear relationship breaks down when pixels have captured only a very small charge, or when they are near their full well depth, known as saturation. Applying a bias offset allows the CCD to avoid the non-linear regime at low exposure levels [13].

Despite the very good reasons for introducing a bias offset, it must be subtracted from the data images before they can be measured. A common way of doing this is to take several bias images from which a master bias frame is created by selecting the median value of each pixel from the range of images. Using the median value eliminates cosmic ray effects and averages out random fluctuations; the more images used to create the master frame, the better. This master bias frame is then subtracted from each of the data images to produce bias corrected images.

3.4.1.2   Dark Current Subtraction

Thermal noise within any material can cause electrons within it to change energy levels. For a CCD this may result in electrons gaining enough energy to jump the band gap separating the valence band and the conduction band. These electrons can then be captured by a nearby pixel's potential well. When the CCD is read out these additional electrons generated by thermal noise are indistinguishable from electrons produced by photons, and are known as *dark current*. The dark current is a function

of the temperature of the CCD, and although the dark current can be subtracted from an image, the noise associated with it cannot be; which is why astronomical CCDs are almost always cooled. The noise introduced to an image is proportional to the square root of the dark current, which increases with increasing temperature. At room temperature a typical CCD can exhibit a dark current of approximately 25 000 electrons per pixel per second, whereas a CCD cooled to below -10 ℃ may have a dark current of well under one electron per pixel per second [50].

Removing the dark current from an image requires a similar process to that used for removing the bias. A set of images, called dark frames, is taken with the shutter closed such that no light is exposed to the CCD. It is preferable to do so at night because even with the shutter closed it is possible for some stray light to find its way onto the CCD during the day. Each image is taken with the same exposure time, that should also match the exposure time of the data images that will be taken separately. Importantly, because of the strong relationship between the CCD temperature and the dark current, each dark frame must be taken at the same temperature as the data images. The dark frames are then median combined to produce a master dark frame.

If at this point the master dark frame is subtracted from a data image, it is actually the dark current *and* the bias offset that is being subtracted because each dark frame also contains the bias offset. The master bias frame should therefore not be subtracted from the data frames in this instance. If it is known in advance that all of the data images to be taken on a given night will have identical exposure times then a master bias frame need not be produced at all. Instead, a master dark frame of the same exposure time and without the bias subtracted can be used to subtract both the bias and dark current simultaneously.

In many instances the exposure times for data frames will not be known in advance, or the exposure times will vary throughout an observation period. In such instances the procedure for removing the bias offset and dark current is slightly more complicated. The total dark current signal increases linearly with exposure time, so a master dark frame can be scaled up or down to match the exposure time of any data image. For example, the individual pixel values of a ten second master dark frame

can all be halved before being subtracted from a five second data frame in order to remove the dark signal from that image. The problem with this approach, however, is that the bias offset that is present in the master dark frame is also halved. Therefore, it is important when using variable exposure times to subtract the master bias from the master dark frame before scaling and subtracting the master dark frame from the data images. The master bias frame must also be subtracted from the data images in such cases.

3.4.1.3 Flat Fielding

The third main image reduction task is known as *flat fielding*. Vignetting and pixel to pixel variations in sensitivity result in different pixels recording different intensities from a uniform light source. Dust particles also produce characteristic dark donut shaped features in images where they cast shadows on the CCD. All of these effects need to be removed to be able to compare pixels with each other. There are a few different ways to eliminate variation across a CCD but they each involve imaging an evenly illuminated source. If it is known in advance that an evenly illuminated source is being imaged, then any pixel to pixel variation across the CCD, excluding cosmic rays and other interference, must be a systematic effect. Because the degree of sensitivity variation between pixels depends upon the wavelength of the incident light, flat fielding must be conducted separately for all filters that are intended to be used.

The primary difficulty with flat fielding is finding a source that is sufficiently uniformly illuminated across the field of view. Common methods include imaging an evenly illuminated white screen affixed to the inside of the observatory dome, imaging the twilight sky near zenith, and imaging the dark night sky. Each of these methods has its own advantages and disadvantages that should be taken into account when deciding which one to use. The twilight sky is inherently non-uniform, particularly in the direction of the setting or rising sun. Pointing the telescope about ten degrees away from zenith, in the opposite direction from the sun, finds a section of the sky that is uniform to within about one percent over a one degree field of view [13].

Unfortunately, the twilight sky does not remain appropriately luminous for very long. An observer usually has on the order of ten minutes to gather the required twilight flats, which can be challenging if many different filters will be used. Obtaining flats when it is too dark results in too many stars becoming visible in the images. If the sky is too bright then the exposure time must be reduced so that the CCD does not saturate. This in turn introduces shutter effects whereby the finite time taken by the shutter to open and close results in the centre of the image receiving more light than the periphery. Figure 3.8 is an example of a very short exposure image exhibiting very apparent shutter effects.
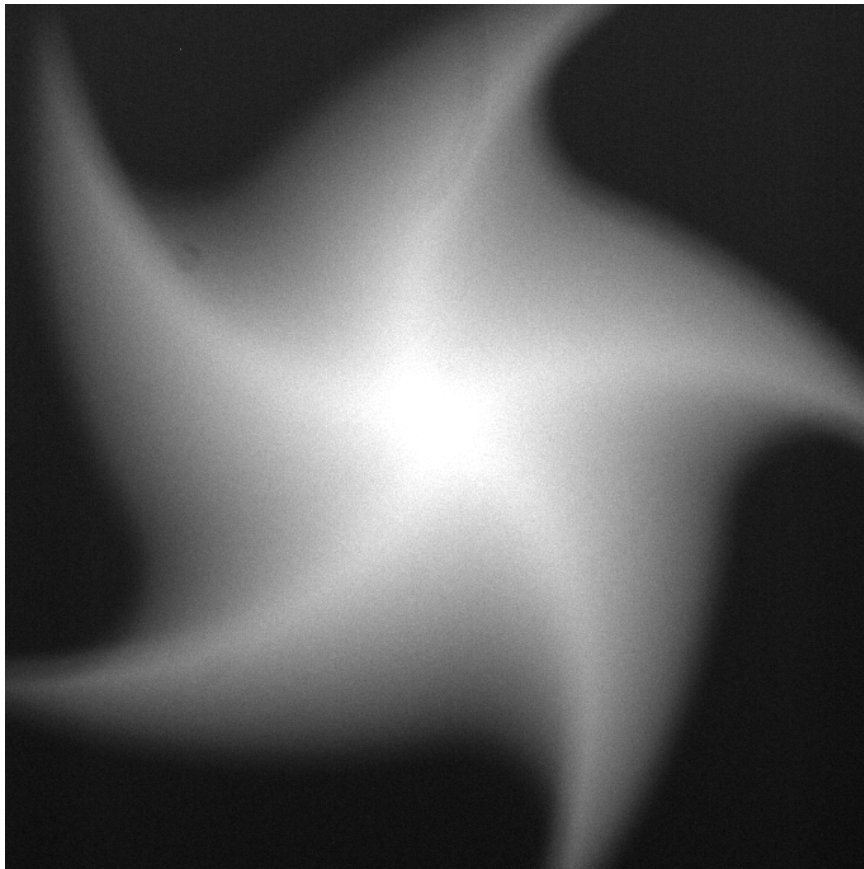

Figure 3.8: A 0.02 second flat field image showing the effect of the Apogee Alta U42 five bladed shutter on the relative exposure of different parts of the frame at short exposure times.

Dome flats give an astronomer control over the illumination intensity and time available to obtain the required set of flats. However, guaranteeing uniform illumination can be very difficult, with shadowing from secondary objects being hard to

eliminate. Unwanted light from reflections off other objects is also possible, as well as specular reflection components off the white screen used as the target. A light diffuser can be placed on the end of the telescope tube to reduce such effects.

A disadvantage with both twilight flats and dome flats is that the spectrum of the target is quite different from the spectrum of most sources of interest. Dark sky flats, taken well after sunset, have the advantage that the spectrum of the night sky is identical to the background that needs to be subtracted from most data images. Dark sky flats also do not suffer from the time limitations imposed upon twilight flats, and they are perfectly uniform at zenith. Unfortunately dark sky flats have several significant disadvantages [13]:

- Many stars will inevitably be present in the images. Halos created by scattered light from stars can be greater than one percent of the background intensity up to several tens of arc seconds in radius.

- The intensity of the night sky is very low, which means that very long exposures are required. Each pixel should be sufficiently exposed to accumulate at least 10000 electrons, which can take between 100 and one million times as long as twilight flat exposures for the visible light bands. This significantly interferes with the time available to obtain data images, especially if many flats are required.

- Dark sky flats must be taken only on moonless nights, to avoid scattered moonlight introducing non-uniformities across the sky.

As with bias and dark subtraction, flat fielding requires that multiple images be taken and median combined to produce a master flat for each filter that will be used, as well as an unfiltered master flat if unfiltered images will be taken. It is especially important to combine multiple images when using twilight or dark sky flats in order to remove the stars that will be present in the images. The telescope tracking should be turned off when taking the flats so that the stars do not appear on the same pixels in more than one image. As a result, the stars will be rejected when the images are median combined.

Each master flat frame needs to have the bias and dark current subtracted from it before it can be used to calibrate data images. In addition, the data images also need to have the bias and dark current subtracted from them before they can be flat fielded. This means that the flats must be taken at the same temperature as the darks and data images. The order is important when reducing the images because the bias and dark signals, which are relatively uniform across the CCD, will alter the degree of scaling required to flatten images if they are not subtracted first.

The three main image reduction tasks are relatively simple in concept, but difficult to achieve in practice. There are many variables that can affect the success of image reduction and it is important to pay close attention to them because the accuracy of the final data depends upon accurate image reduction.

3.4.2 Image Reduction Procedure

The first steps in image reduction for this experiment were to create nightly master bias, dark, and flat frames from each night's data. A master frame is a form of average image created from several individual frames in order to minimise the effects of extreme pixel values caused by cosmic rays and random fluctuations.

First, a master bias frame was created in Mira using the median combine method, which sets each pixel's value to the median value of the corresponding pixel intensities in the 25 bias images used to create it. The master dark frame was created next by first subtracting the master bias frame from each of the individual dark frames, and then median combining the dark frames. Finally, five master flat frames were created: one for each of the Bessel filters, and one for the unfiltered data. To create the master flats, each individual flat frame first had the master bias and master dark frames subtracted from it. When subtracting the master dark frame, a critical step was to ensure that the master dark was first scaled according to the relative exposure times of the flat frames compared to the master dark frame. Since the dark frame exposure times were one minute long, compared to one second for the flat frames, sixty times too much dark current would be subtracted from the flats if the master dark was not first scaled. This is also the reason why separate master bias and master dark frames were

produced. If all of the dark, flat and data images were taken with identical exposure times, then the master dark frame could be produced without first subtracting the bias. Subsequently subtracting the master dark from other images would thus have the effect of simultaneously subtracting both the bias and dark current. If the master dark in this situation was to be scaled, however, it would have the unwanted effect of scaling the included bias as well. As with the bias and dark frames, the master flats were created by median combining the set of constituent flats. Unlike the bias and dark frames, however, the flats were first required to be normalised according to the median value in a central region of each frame. The central region that was used was 20 percent of the width and 20 percent of the height of the overall frame. Normalising the flats was required to correct for the fact that the twilight sky continuously darkens whilst the flat field images are being obtained. Therefore, without scaling up each subsequent image to match the level of exposure of the first image, they could not meaningfully be median combined.

Once the master bias, dark, and flat frames had been produced, they were applied to all of the data images, including both the satellite images and the standard star images. The order in which this was done was bias subtraction first, followed by exposure-scaled dark subtraction, and finally the flats were applied.

3.4.3 Data Extraction

Before data could be extracted from the images, the gain and read noise was required to be added to each image header. The gain is important because it determines how many electrons, and thus the photon flux, correspond to a given ADU value. Without knowing the correct value for the gain, all magnitude measurements could potentially be very significantly in error. The read noise is necessary because it is used in the calculation of error values in magnitude measurements. The gain and read noise are both calculated in Mira using the 'Calculate Gain' command in the 'Measure' - 'Diagnostics' menu. The 'Calculate Gain' command requires the use of Mira's rectangular cursor to select an area within two flat images and two bias images. Alternatively, each image in its entirety can be used for the calculation. During this

experiment, the cursor was always used to select an area in each flat image that was free of stars. Exactly the same area was selected on both of the flats and also on the bias images. For the best accuracy, the gain calculation requires that the flats are taken sequentially, using the same illumination. Because twilight flats were used in this experiment, the sky illumination changed slightly between each image. Although the flats were taken sequentially, within approximately six seconds of each other, this change in illumination is a source of error in all of the gain measurements.

Mira was used to extract data from the reduced images in large batches. All of the images taken through a particular filter on a given night were opened in one window as an image set. Each satellite was marked in the first image with the aperture tool, available from the Aperture Photometry Toolbar. Mira automatically calculates the centroid of each satellite's PSF and centres the aperture tool on the centroid. Mira is then commanded to track the satellite through the entire image set, centering the aperture tool on the PSF centroid in each image and simultaneously calculating the instrumental magnitude of each satellite in each image. The instrumental magnitude of each satellite, along with a variety of other data, including the Universal Time Coordinated (UTC) time of each image, is presented in a photometry report window that facilitates the export of the data to Microsoft Excel.

To measure magnitudes of sources within images, Mira uses a three ring aperture tool. The inner ring defines the area within which the source signal is measured, and the area between the two outer rings is used to measure the sky background level. The modal pixel value between the two outer rings was taken to be the background level in each measurement. Figure 3.9 is a screenshot showing the aperture tool (coloured red in the image) centred on Galaxy 11. The data recorded by Mira is visible in the spreadsheet at the bottom of the image.

Throughout the experiment, great care was taken to ensure that the dimensions of the apertures were set such that as much of the signal from each satellite was being captured as possible without including excessive amounts of background noise. Two-dimensional radial intensity plots that plot pixel values as a function of their radial distance from the satellite's centroid, and three-dimensional surface plots, were both
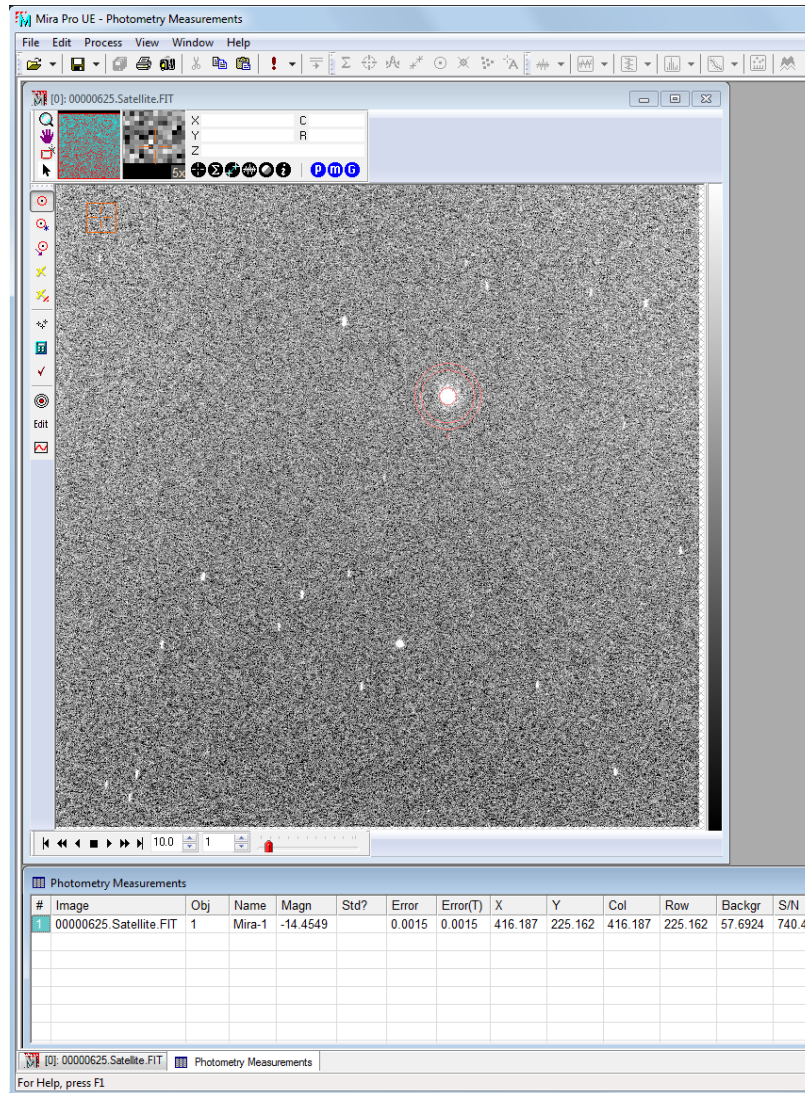
Figure 3.9: Screenshot of a satellite's magnitude being measured in Mira.

used in the effort to ensure that a consistent approach to selecting aperture sizes was utilised. Figure 3.10 shows a screenshot of a radial intensity plot and a 3D surface plot for an image containing Galaxy 11.

Once the satellites had been marked and measured throughout an image set, each image was carefully examined to detect and discard those in which the satellite's PSF had been contaminated by light from stars. Figure 3.11, for example, shows a star passing behind Anik F1, contaminating its signal, and thus making an accurate magnitude measurement impossible to achieve.
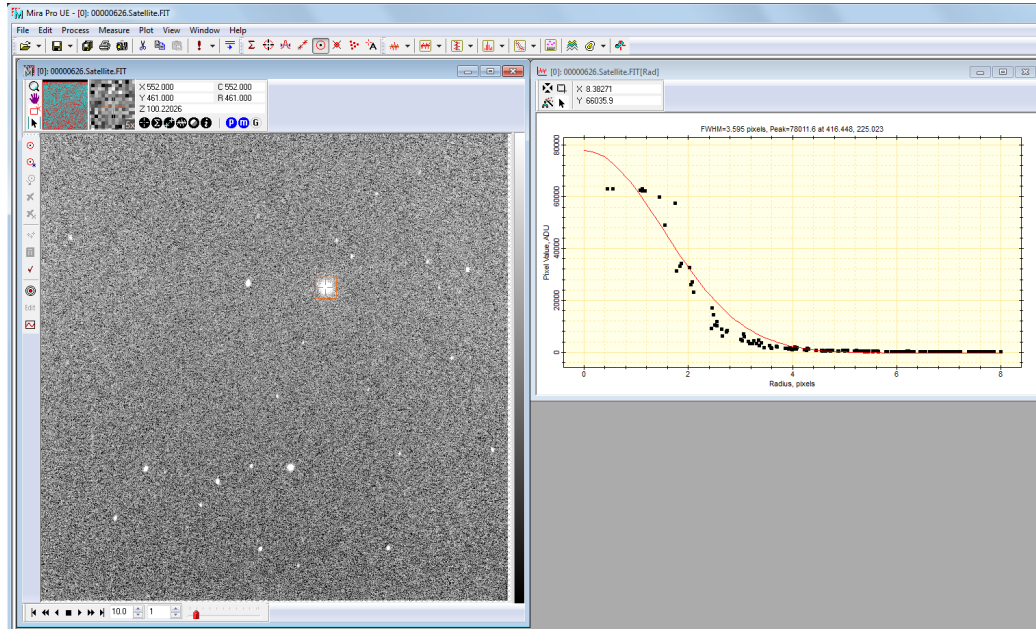
(a) Radial intensity plot.
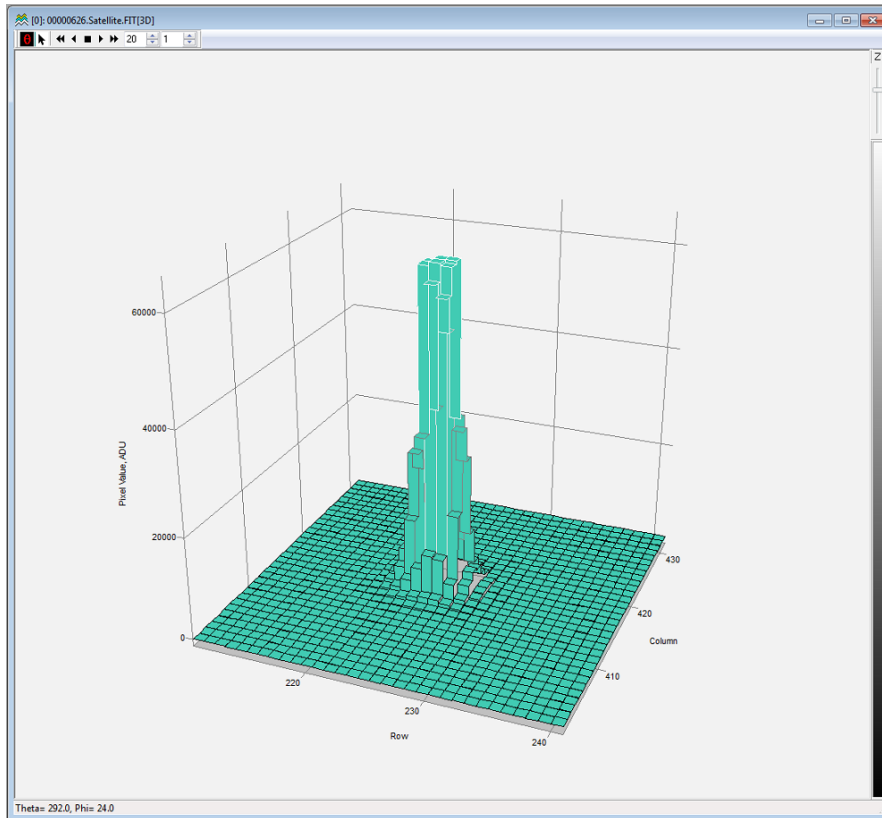
(b) Partial image surface plot.

Figure 3.10: Examples of a radial intensity plot (a) and a 3D surface plot (b) for a small portion of an image containing Galaxy 11. In these images it can be seen that Galaxy 11 has saturated a number of pixels at the centre of its PSF, as indicated by their ADU values being close to 65535.
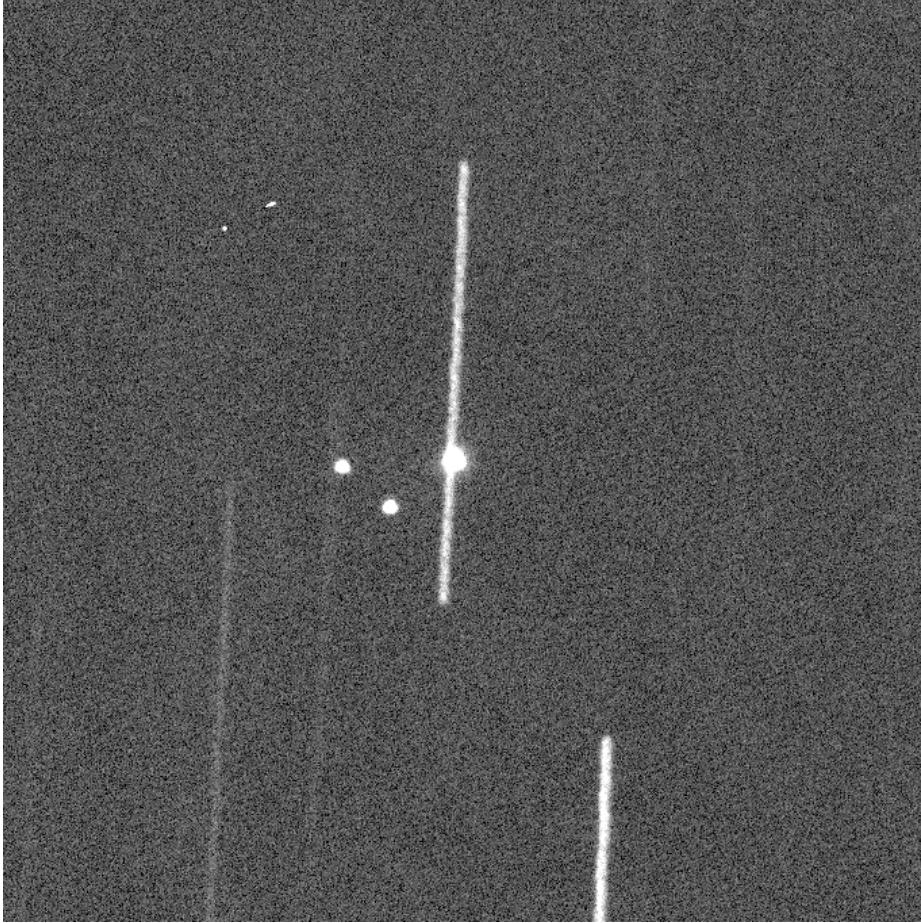
Figure 3.11: Image showing a star passing behind Anik F1, contaminating its signal. Anik F1R and Anik G1 are visible to the left of Anik F1, which has the long star streak passing through it.

There were two challenges resulting from the fact that multiple satellites were captured in each image. First, in images where the primary satellite was very bright, the SNRs of the other, fainter, satellites were often very low, because of the correspondingly short exposure times, resulting in noisier data. Conversely, when the primary satellite was very faint, sometimes the other satellites saturated some pixels, resulting in data needing to be discarded. The second challenge was somewhat unexpected. On two separate nights, Anik G1 and Anik F1R passed so close to each other that they contaminated each other's light signal. On 29 October 2013, their PSF's merged such that only one satellite could be distinguished; their centroids passed within two pixels of each other, which corresponds to approximately 200 metres lateral separation at GEO altitudes. As a result, accurate magnitude measurements could not be made

for these two satellites for approximately one hour on 12 October, and approximately two hours on 29 October.

Once all of the contaminated images had been discarded from the set, the photometry report was exported to Excel for further analysis.

3.4.4   Calibration to Standard Stars

The magnitude values measured by Mira are instrumental magnitudes. To convert them to apparent magnitudes on the standard scale required the addition of an offset called a zero point. The zero point is equal to the absolute value of the instrumental magnitude that would result from the measurement of a source with an apparent magnitude of zero according to the photometric system being used. The zero point varies with airmass and changing atmospheric conditions such that it is usually fairly stable over the course of a night, but may vary from one night to the next. The four main atmospheric effects that contribute to zero point variability are Rayleigh scattering, scattering by aerosols, absorption by ozone, and absorption by other molecules. Aerosols scatter photons through diffraction, whereas Rayleigh scattering is the result of a more complex interaction between electromagnetic waves and air molecules. The fact that the efficiency of Rayleigh scattering decreases by the fourth power with increasing wavelength explains why the sky appears blue. The degree of Rayleigh scattering varies directly with atmospheric pressure, but is otherwise stable over time. Aerosol scattering and absorption by molecules can vary over relatively short timescales as airmasses with different densities of aerosols and other particles move around the sky. Absorption by ozone is relatively constant in the short term in a given location [13].

Due to the possibility of zero points varying with atmospheric effects, zero points were calculated for each night on which Landolt stars were successfully imaged; and for each filter too, since the zero point is dependent on the waveband.

The precisely known apparent magnitudes of the Landolt stars that were imaged facilitated the calculation of each night's zero point. All of the images of the Landolt star field taken on a given night through a particular filter were opened in Mira as

an image set. Between two and five stars of sufficient SNR were then marked as standard stars with the aperture tool, and their known magnitudes were entered into Mira. The stars were then tracked through the image set. By comparing the measured instrumental magnitude of each standard star with its known apparent magnitude, Mira calculates the zero point of each image such that it minimises the sum of the squared differences between the measured and known apparent magnitude of each star in the image. A different zero point is therefore calculated for each image. Because the airmass through which each image was taken changed as the stars tracked sidereally across the sky, there would be one image taken through an airmass most similar to the airmass through which the satellites were imaged. The images subsequent and previous to that image were taken through increasingly different airmasses. The nightly zero point to be added to the instrumental data was taken to be the mean of the zero points calculated for the five images taken through airmasses most similar to the airmass through which the satellites were imaged. The airmass through which each image was taken was automatically calculated by CCDSoft 5 and added to the image header.

3.4.5 Manipulation of the Data for Presentation

After adding the zero points to all of the respective instrumental magnitude measurements in Excel, the resultant apparent magnitudes and their associated observation times were copied into MATLAB. MATLAB was used to plot lightcurves, represented as magnitude versus minutes since midnight (UTC), and colour ratio curves, represented as colour ratio versus minutes since midnight (UTC). The magnitude plots were relatively straightforward, however the colour plots required some manipulation of the data.

Because photometric measurements were not taken simultaneously in each of the colour bands, colour ratios could not accurately be calculated simply by subtracting the magnitude measured through one filter from the magnitude measured through another. Instead, the magnitude data were interpolated by fitting a quadratic curve to rolling sets of three data points and determining the interpolated magnitude in

one minute intervals. All of the data through each of the filters were interpolated to identical times, thus allowing the calculation of colour ratios. Small dots were plotted on the lightcurves to indicate the actual measured values, however smooth lines without data points were used to plot the colour curves since they were calculated from interpolated values.

Chapter 4

Overview of Observational Results

This chapter provides a description of the typical results that were obtained during the experiment. The full set of all lightcurves and colour variations constructed from the data obtained throughout the experiment can be found at Appendix A, and will therefore not be provided here, however general characteristics of the results will be described, and some examples will be shown.

4.1 Nightly Zero Points

Of the twelve nights during which observations were taken, zero points were calculated for nine. On all of the nights for which zero points were not calculated, GEO cluster 1 was observed. On the 16^{th} and 25^{th} of August, and the 28^{th} of September, cloud cover or fog interfered with the Landolt star field images, rendering zero point calculations impossible. Therefore, the zero point calculated on the 17^{th} of August was applied to the instrumental magnitudes measured on the 16^{th} and 25^{th} of August, and the zero point calculated on the 29^{th} of September was applied to the instrumental magnitudes measured on the 28^{th} of September. Although this remains a source of error for these three nights, the variation in the zero point data indicates that the error would likely be relatively small. The V band zero points showed the greatest variation during the remaining five nights of observations of GEO cluster 1, with a standard deviation of 0.155 magnitudes. The R band zero points showed the smallest variation, with a standard deviation of 0.016 magnitudes. Also, any error would result in the lightcurves and colour variations being shifted up or down slightly, without altering their overall shape. Table 4.1 lists the zero points that were calculated for each colour band on each night for the first GEO cluster that was observed, and Table 4.2 lists the zero points for the second GEO cluster.

Due to the different locations of the two GEO clusters in the sky, the first GEO

Date	B Band Zero Point	V Band Zero Point	R Band Zero Point	I Band Zero Point
27 Jul 2013	20.661	21.080	21.451	20.709
05 Aug 2013	20.813	21.059	21.452	20.708
16 Aug 2013	Landolt stars obscured by cloud. The 17 Aug 2013 zero points were used for this night instead.			
17 Aug 2013	20.567	20.960	21.420	20.720
25 Aug 2013	Landolt stars obscured by cloud. The 17 Aug 2013 zero points were used for this night instead.			
28 Sep 2013	Landolt stars obscured by fog. The 29 Sep 2013 zero points were used for this night instead.			
29 Sep 2013	20.585	20.699	21.414	20.643
03 Oct 2013	20.629	20.760	21.438	20.637
Mean	20.651	20.912	21.435	20.683
Standard Deviation	0.087	0.155	0.016	0.036

Table 4.1: Nightly Zero Points (in magnitudes) for the First GEO Cluster.

Date	B Band Zero Point	V Band Zero Point	R Band Zero Point	I Band Zero Point
12 Oct 2013	20.709	20.603	21.418	20.632
29 Oct 2013	20.550	20.709	21.419	20.656
04 Nov 2013	20.608	20.737	21.440	20.673
20 Nov 2013	20.850	20.685	21.414	20.668
Mean	20.679	20.684	21.423	20.657
Standard Deviation	0.114	0.050	0.010	0.016

Table 4.2: Nightly Zero Points (in magnitudes) for the Second GEO Cluster.

cluster was observed through approximately 1.7 airmasses, and the second cluster was observed through approximately 2.0 airmasses. Despite the different airmasses, the zero points for the two satellite clusters were very similar to each other, indicating that the magnitudes of the zero points are only weakly dependent upon airmass. Also, it can be seen that there was less variability in the R and I band zero points compared to the B and V band zero points for both GEO clusters, implying that the B and V bands are more significantly affected by the atmospheric changes that were encountered than the R and I bands. This result is consistent with night to night changes in wavelength dependent scattering processes such as Rayleigh scattering, which is more significant at shorter wavelengths, or aerosol scattering, which depends upon both the wavelength and the aerosol particle size.

4.2   Lightcurves

For each of the satellites that were observed, lightcurves were constructed for each of the filter bands and for the unfiltered data, on each night that observations were taken. The lightcurves for the data obtained using filters depict apparent magnitudes, whereas the lightcurves produced from unfiltered observations depict instrumental magnitudes since there exists no standard photometric system specific to unfiltered data obtained using the particular equipment used in this experiment. Each lightcurve consists of approximately 200 data points. Throughout this thesis, time (represented as minutes since midnight UTC) is placed on the horizontal axis for all plots, rather than phase angle. The primary reason for doing so is to simplify the process of reconstructing the sun-satellite-sensor geometry should a reader want to do so.

Each satellite was observed to be brightest in the I band, followed by the R band, then V band, and faintest in the B band. This was true for all of the satellites throughout every night of observation. Typically, each of the lightcurves varied smoothly, with only a few occasions where the brightness varied unpredictably by greater than a few tenths of a magnitude between data points. Further, instances of unpredictable lightcurve behaviour always lasted for a relatively short time: usually less than one hour. Figure 4.1 depicts the B, V, R, and I band lightcurves for Intelsat
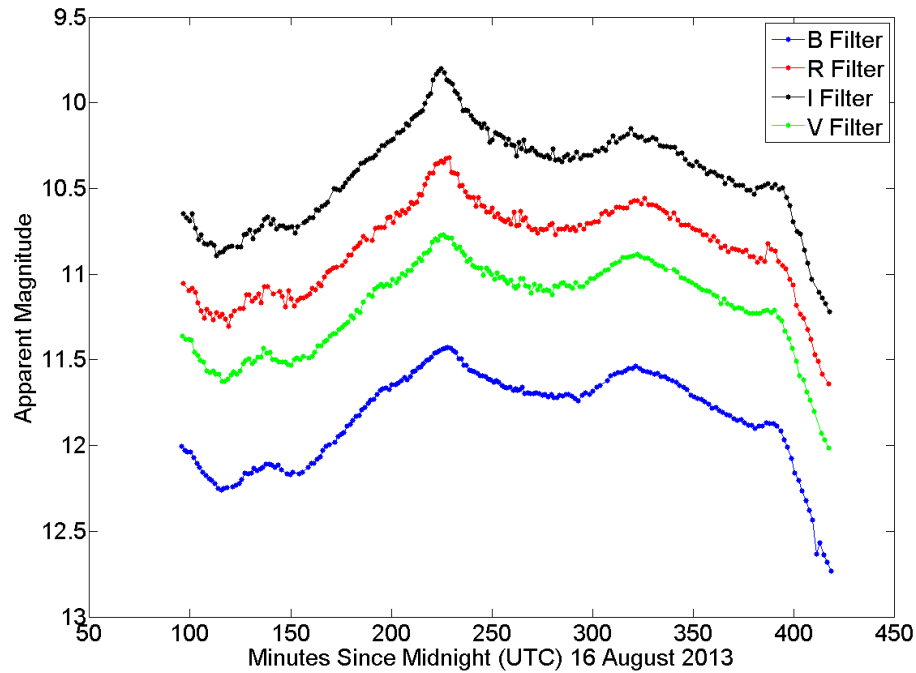
805 on the 16$^{\text{th}}$ of August.

Figure 4.1: Intelsat 805 lightcurves.

In general, the most prominent feature of each lightcurve was a short period of brightening lasting for approximately one hour, consistent with the description of glints that have frequently been reported as due to specular reflections off satellite solar arrays [18, 27, 30, 51]. The precise shape of the curve during these periods of brightening varied between the different satellites, and sometimes there was more than one peak visible. In addition, the degree of brightening changed for each satellite from one night to another, varying between about one magnitude and ten magnitudes above the level outside of the specular glint. Figure 4.2 presents example lightcurves for Anik G1 and Anik F1.
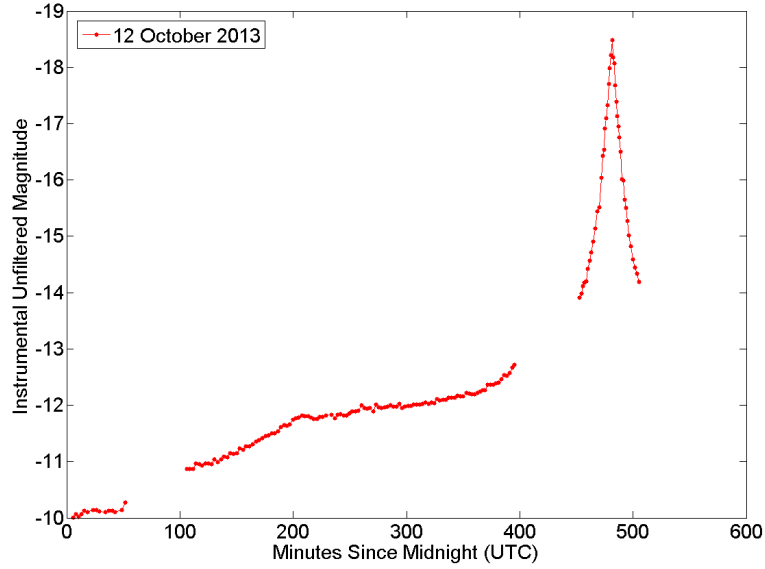
(a) Anik G1 Unfiltered Lightcurves.
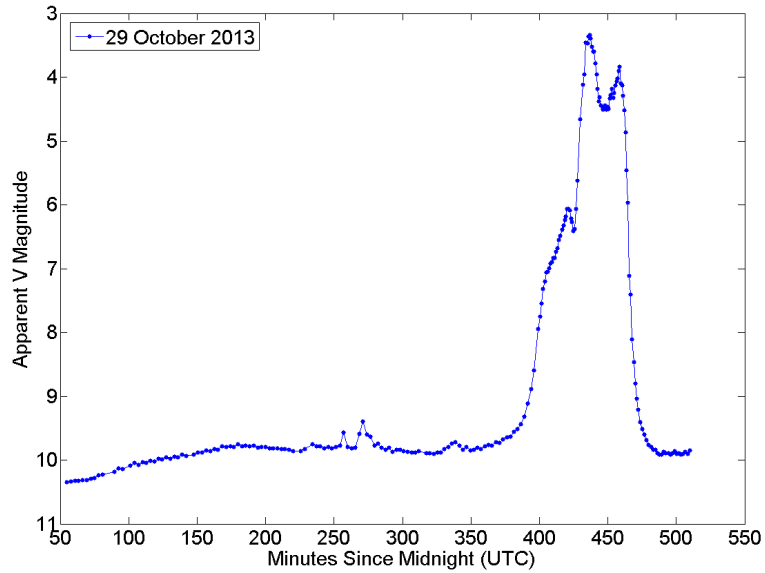
(b) AnikF1 V Band Lightcurves.

Figure 4.2: Unfiltered lightcurve for Anik G1 (a), and V band lightcurve for Anik F1 (b).

Note that Anik G1 has a single, sharp peak, whereas Anik F1 has a more complex peak with up to three separate local maxima. Although very bright peaks such as those depicted in Figure 4.2 were observed for all of the satellites, the degree of brightening at the peak, relative to the brightness just prior to and after the peak, changed considerably over several weeks. In Figure 4.3, for example, on the 29th of October Anik F1's B band brightness increased by approximately 6.5 magnitudes from 350 minutes since midnight to 430 minutes since midnight, but three weeks later, on the 20th of November, the brightness only increased by approximately 1.5 magnitudes over the same time. Such variation in the brightness of the lightcurve peaks was observed for all of the satellites. On the other hand, lightcurve segments where such peaks in brightness do not occur usually appear very similar to each other from one night to another, as can be seen in Figure 4.3.
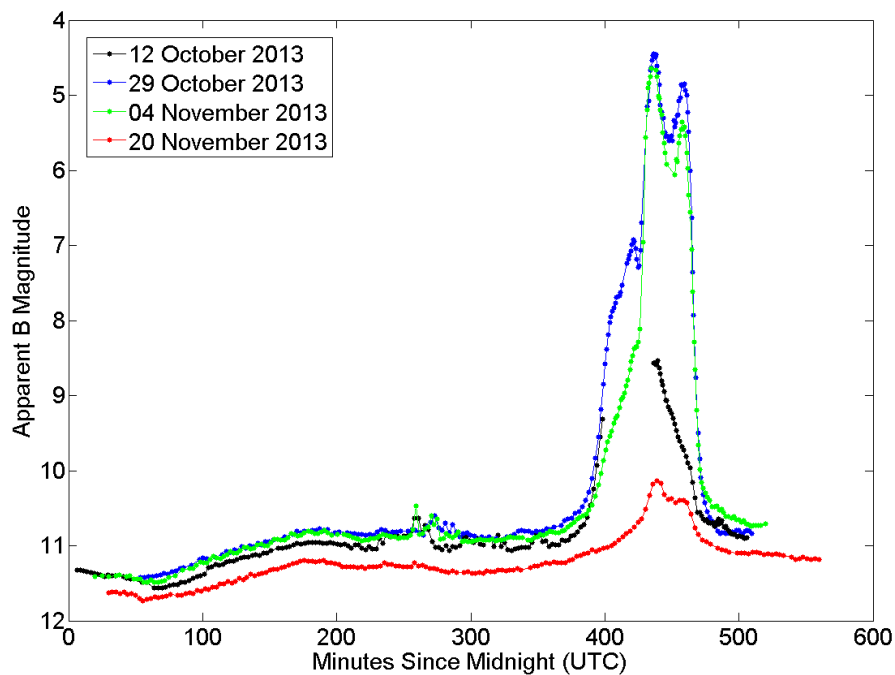

Figure 4.3: Anik F1 B band lightcurves.

The lightcurves in Figure 4.3 are generally smoothly varying throughout the night. On the other hand, in Figure 4.4, between 300 and 400 minutes since midnight on the 3rd of October Galaxy 11's I band lightcurve varies in a far less predictable manner.

Similar to the main brightness peaks, lightcurve segments that display unpredictable variations change over time, as can be seen in Figure 4.3 near 280 minutes since midnight. It is believed that the smooth, high amplitude, peaks are caused by specular reflections off satellites' large, flat solar arrays, whereas the smaller, more irregular peaks are caused by specular reflections off relatively small parts of the satellite body.
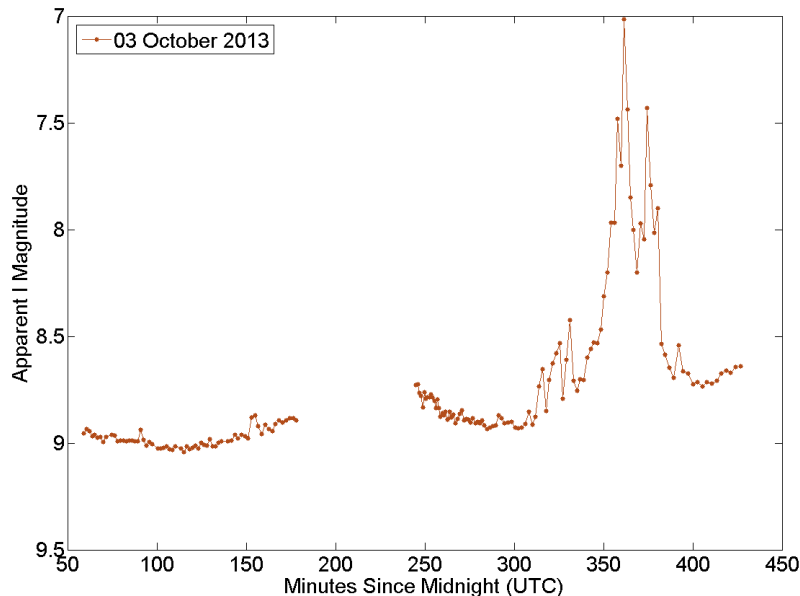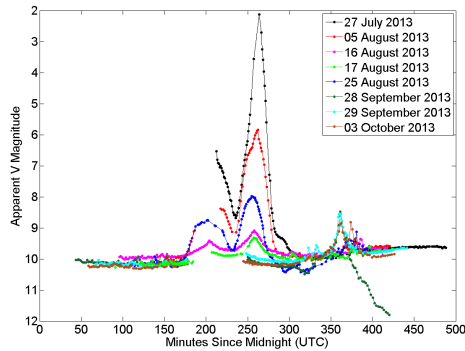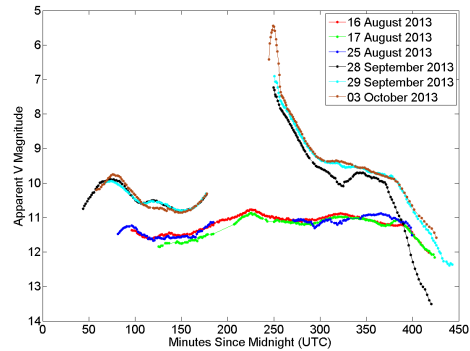

Figure 4.4: Galaxy 11 I band lightcurve.

Notwithstanding the care that was taken to ensure that observations were only conducted on clear nights, occasionally it was difficult to avoid interference by cloud or fog. Of note, the Galaxy 11 and Intelsat 805 data acquired on the 28[th] of September appear to be affected by fog from 375 minutes since midnight onwards. At the end of the observing period on the 28[th] of September, shortly before dawn, it was noticed that fog had begun forming at ground level. Although to the naked eye fog was not visibly present in the vicinity of the telescope dome on the RMCC roof, there is a steady reduction in brightness with time of both Galaxy 11 and Intelsat 805 apparent from 375 minutes since midnight onwards. Figure 4.5 depicts the Galaxy 11 and Intelsat 805 V band lightcurves. The reduction in brightness after 375 minutes since midnight is similarly apparent in all of the other bands. The Landolt star observations

on the 28$^{\text{th}}$ of September were taken at the end of the night, and they too displayed a steady decrease in brightness with time characteristic of obscuration by fog.

(a) Galaxy 11 V Band Lightcurves.

(b) Intelsat 805 V Band Lightcurves.

Figure 4.5: Galaxy 11 and Intelsat 805 lightcurves showing the effects of fog on the 28$^{\text{th}}$ of September.

Throughout the experiment uncertainties in the magnitude measurements were almost always smaller than the size of the plotted data points. As a result, errorbars are not depicted on any of the plots. Typical uncertainties are annotated on some plots in portions where the uncertainties are slightly greater than the size of the data points.

4.3   Colour Variations

In addition to the lightcurves, six different colour plots were constructed for each satellite. The colour indicies that were considered are B-V, B-R, B-I, V-R, V-I, and R-I. Rather than each satellite's colours remaining constant throughout a given night, there was often significant variability. In addition, the shape of a given satellite's colour variation did not necessarily remain unchanged from one night to another. For example, Figure 4.6 depicts Intelsat 805's B-I colour variation on the 29[th] of September and the 17[th] of August. It can be seen that not only do the two night's colour plots have distinctly different shapes, but the 29[th] of September curve is relatively more blue throughout the night than the 17[th] of August curve.
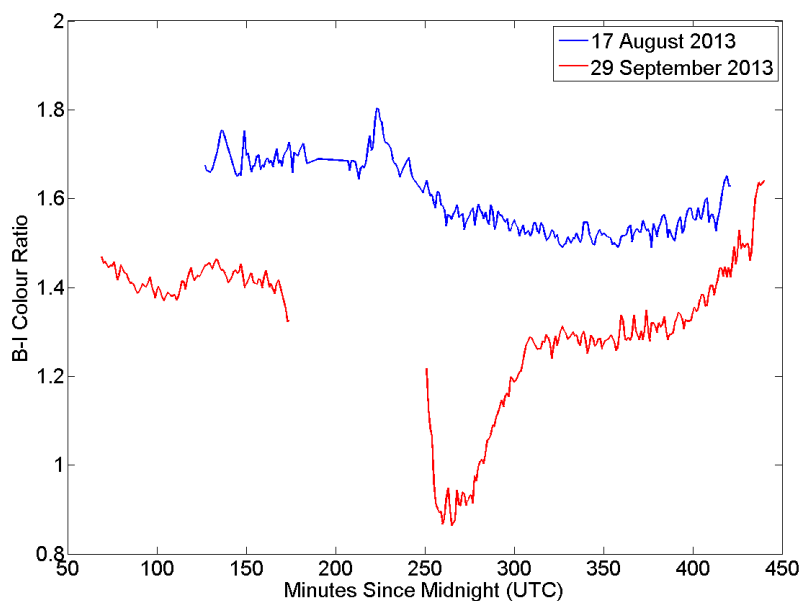

Figure 4.6: Intelsat 805 B-I colour variation.

Not all of the colour variations for different nights were vertically separated on the axes as in Figure 4.6, however. Galaxy 11's colour variations from different nights significantly obscured each other when plotted on the same axes. Therefore, to improve the ability for the figures to be interpreted, Galaxy 11's colours were plotted such that successive nights' data were progressively shifted vertically, as depicted in Figure 4.7.
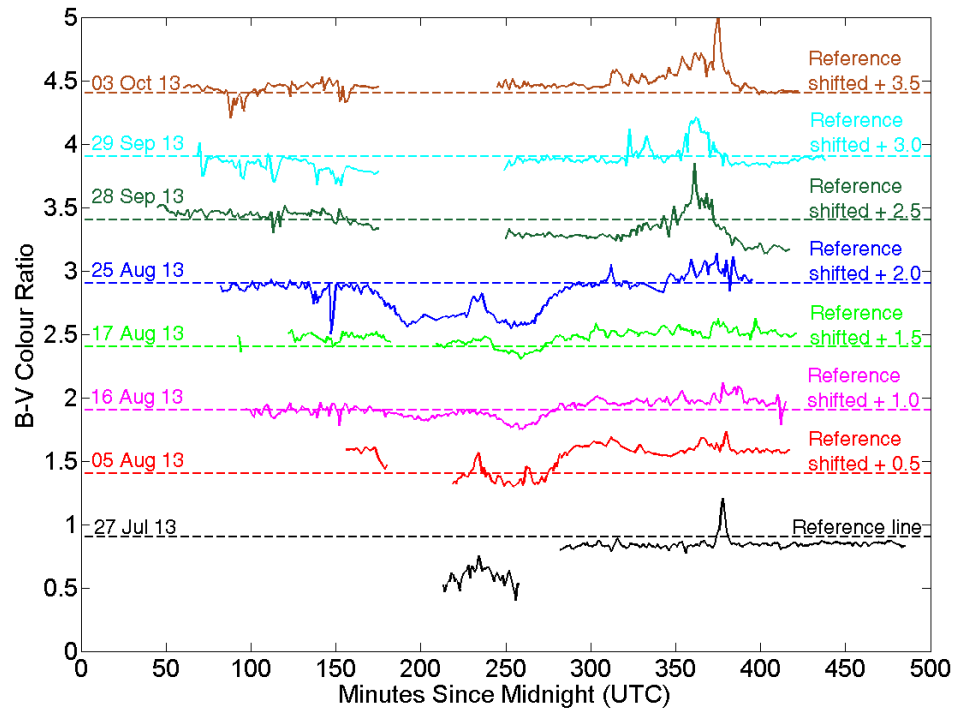
Figure 4.7: Galaxy 11 B-V colour variations.

Throughout the experiment uncertainties in the colour ratios were almost always smaller than the thickness of the lines. Therefore, errorbars are also not depicted on any of the colour plots. Typical uncertainties are annotated on some plots in portions where the uncertainties are slightly greater than the thickness of the lines. Figure 4.8 is an example of one such plot.
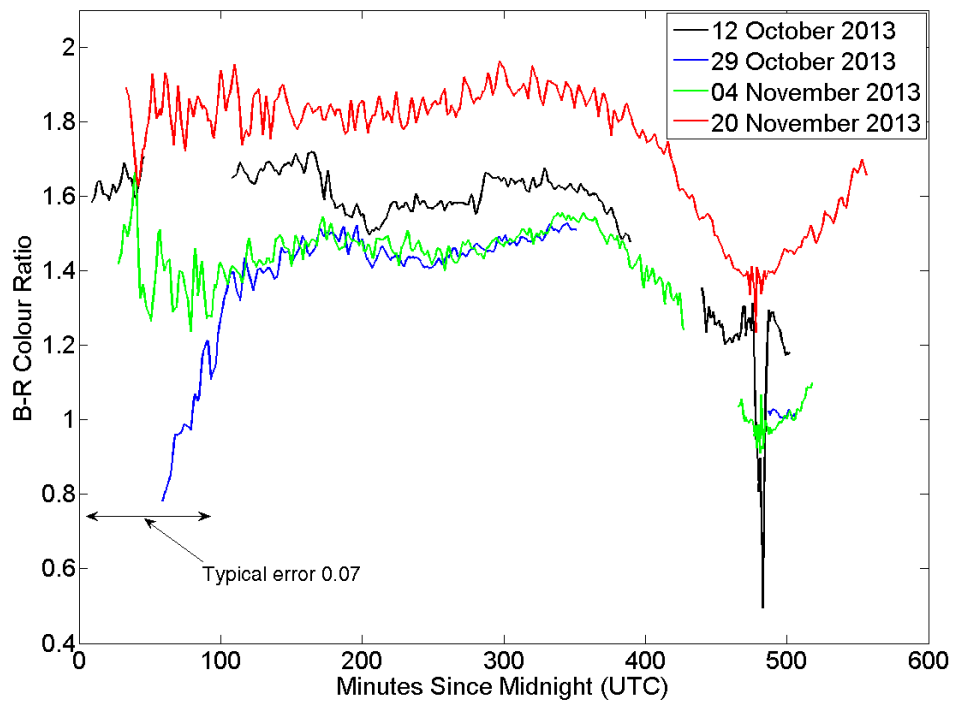
Figure 4.8: Anik G1 B-R colour variations.

Chapter 5

Tests for Systematic Errors

This chapter details the efforts made to determine the degree to which systematic errors are present in the data, including both the zero point calculations and the individual magnitude measurements. Three independent tests were undertaken for this purpose. First, on one night the zero points were calculated three separate times using three different Landolt star fields to confirm the accuracy of the zero point measurements and their stability over the period of the night. Second, individual stars of known magnitude that were incidentally present in the background of various images were measured. The measured magnitudes were compared to the stars' known magnitudes to determine the precision of the magnitude measurements made using the Mira software. Finally, colour ratio variations were compared for satellites that were imaged together in the same frames. More specifically, the colour plots were examined for significant features that were replicated for multiple satellites that might be indicative of systematic errors.

5.1   Multiple Zero Point Calculations During a Single Night

Comparing the general shape of different satellite's colour ratio curves is useful for identifying systematic errors that vary throughout the course of a particular night, however it will not be helpful in identifying errors in each night's zero point calculation because such errors would uniformly shift the entire curve without changing its shape. As demonstrated in figure 5.1, the general position of each satellite's lightcurve was seen to be shifted up or down on some nights relative to others. In figure 5.1, the lightcurve appears to become increasingly fainter overall at later dates. Similar effects can be seen in other satellites' lighcurves, and through other filter bands. What is not immediately clear is whether the shift in the lightcurves is as a result of changes in the illumination-observation geometry as the seasons change, or if it is the result

87

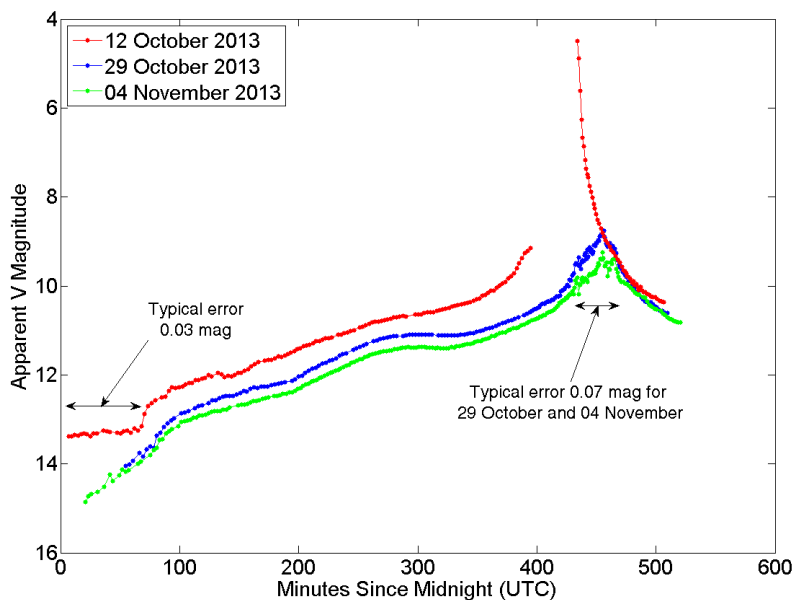of errors in the calculation of the zero points.

Figure 5.1: Echostar 17 V Band Lightcurves.

It is well known that atmospheric changes will cause the photometric zero point to vary from one night to another, however it is generally assumed that the atmosphere will be stable enough throughout any one particular night to allow the use of a single zero point to be applied to all measurements made that night [52]. Although the zero point would be expected to change if a different airmass moved into place over the observer, such as when a cold front brings a colder airmass with it, such events are usually associated with increased cloud cover. Therefore, the clear conditions necessary for photometry of satellites are generally associated with stable photometric properties over the course of a given night.

To test the repeatability of the zero point calculations, on 03 October 2013 seperate Landolt standard star fields were imaged three times during the night: at the beginning, middle, and end of the night. The nightly zero points were subsequently calculated three times for each of the colour filters; once each using the three Landolt star fields. Landolt star fields were only imaged once per night on each of the other nights because imaging the Landolt stars interrupts the collection of data on

Landolt Star Field	B Band Zero Point	V Band Zero Point	R Band Zero Point	I Band Zero Point
SA114	20.648	20.761	21.450	20.665
SA93	20.650	20.766	21.450	20.658
SA97	20.627	20.753	21.414	20.589
Mean Zero Point	20.642	20.760	21.438	20.637
Max Deviation from the Mean	0.015	0.007	0.024	0.048

Table 5.1: Comparison of three independent zero point measurements on 03 October 2013.

the satellites. The third of October, however, presented a good opportunity for the collection of multiple sets of images of Landolt stars because for approximately one hour in the middle of the night the satellites being imaged were in eclipse.

Table 5.1 displays the calculated zero points for 03 October 2013. The SA114 Landolt field was imaged at the start of the night, followed by the SA93 field in the middle of the night, and the SA97 field at the end of the night. The mean zero points for each colour band, listed at the bottom of the table, were taken to be the 03 October zero points for the purpose of analysing the satellite data. Zero point data for the remaining nights were detailed in Chapter 4.

As can be seen, only 0.023 magnitudes separates the highest and lowest zero points for the B band, 0.013 of a magnitude for the V band, 0.036 of a magnitude for the R band, and 0.076 of a magnitude for the I band. These variations are far smaller than the vertical separation between lightcurves depicted in figure 5.1. Interestingly, when only the first two sets of zero point measurements are compared, the difference in the B,V,R, and I bands drops considerably to 0.002, 0.005, less than 0.001, and 0.007 magnitudes, respectively. With such a small dataset it is impossible to know whether this is indicative of a slight change in the photometric conditions during the latter half of the night, or if it is simply coincidental that the results from the first two sets of measurements are much closer to each other than the third set of measurements. Irrespective of the reason, the results demonstrate that the zero point calculations are repeatable to within 0.1 magnitude, and that there is no significant variation due

to systematic errors, at least in the 03 October data.

5.2   Measurement of Incidentally Imaged Known Stars

To further test the accuracy of the zero point measurements, and to assess the stability of the photometric zero point throughout each night, various stars' known magnitudes were compared to their measured magnitudes where those stars were incidentally captured in the background of an image.

On 05 Aug 2013, 28 Sep 2013, 29 Sep 2013, 03 Oct 2013, and 12 Oct 2013, B, V and R band apparent magnitudes were calculated for various stars that happened to be in the field of view of some of the images, using the respective nightly zero points. To do this, images with exposure times of two seconds or less were identified. Only images with such short exposure times were used because in images with longer exposure times the star streaks were too long to obtain sufficiently accurate measurements. Stars appeared as streaks in the images because data were gathered with the telescope tracking the satellites rather than the stars. Although the eccentricity of Mira's aperture tool can be adjusted to measure elliptical light sources, accurately measuring the magnitudes of the stars becomes more difficult the longer the streaks appear on the images. Therefore, it was decided that only images with exposure times of two seconds or less would be used for the test. To calculate the apparent magnitude of a star, first the instrumental magnitude was measured using Mira, and then the nightly zero point was added to it to convert it into an apparent magnitude. Figure 5.2 is a screenshot showing a measurement being taken of a star in the background of an image taken with a two second exposure through the R filter. The red three-ring aperture tool is visible surrounding the star at the centre left of the image.

In order to identify the stars present in the selected images, an astrometric calibration was done on each image using the PinPoint software program and MaxIm DL 5. Once this step was achieved for a particular image, the name and location of each star in that image could easily be obtained using MaxIM DL 5's Astrometry tool. From there the SIMBAD astronomical database, operated by the Strasbourg Astronomical Data Center (CDS), was used to find published values for the magnitudes of these
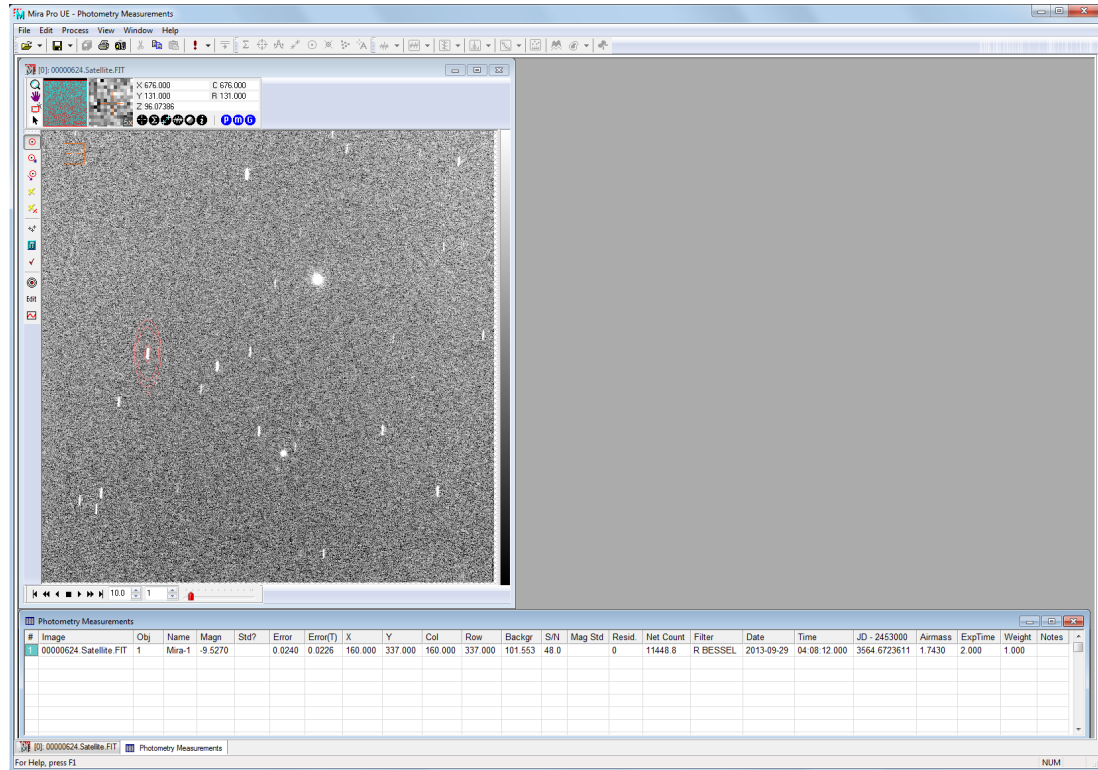
Figure 5.2: Taking a measurement of the magnitude of a background star in Mira.

stars. SIMBAD contains measurements of over seven million different astronomical objects obtained from a variety of sources. Some objects have been measured many times, using highly accurate methods, and other objects have been measured very few times, or with lower accuracy. As such the quality of the measurements is highly variable. Standard (I.e. 1 sigma) errors are given for most values, and the quality of the data is indicated subjectively by assigning them a value from A to E, where A is the highest quality and E is the lowest quality. Most of the stars that were able to be identified were given a quality indicator of either D or E. Where a published value of the apparent magnitude of a star was available, it was then compared to the calculated apparent magnitude. All stars that were able to be identified in the images that were examined were measured, and no measurements were excluded from the data set to improve the results. There were no published I band magnitudes for any of the stars identified in the images, and far fewer R band measurements than B and V band measurements. As a result, only B, V, and R band comparisons could be made. In total, 67 measurements were made.

Figure 5.3 plots the difference between the calculated and published apparent magnitudes for all of the stars that were measured. The one sigma error values for the known magnitudes reported in the SIMBAD database are indicated in the figure by the black errorbars spanning the line Y = 0. Where no errorbar is present, no error value was published for that particular colour magnitude for that star. The coloured errorbars centred on the datapoints represent the one sigma error values in the measurement of the instrumental magnitudes, as calculated by Mira.
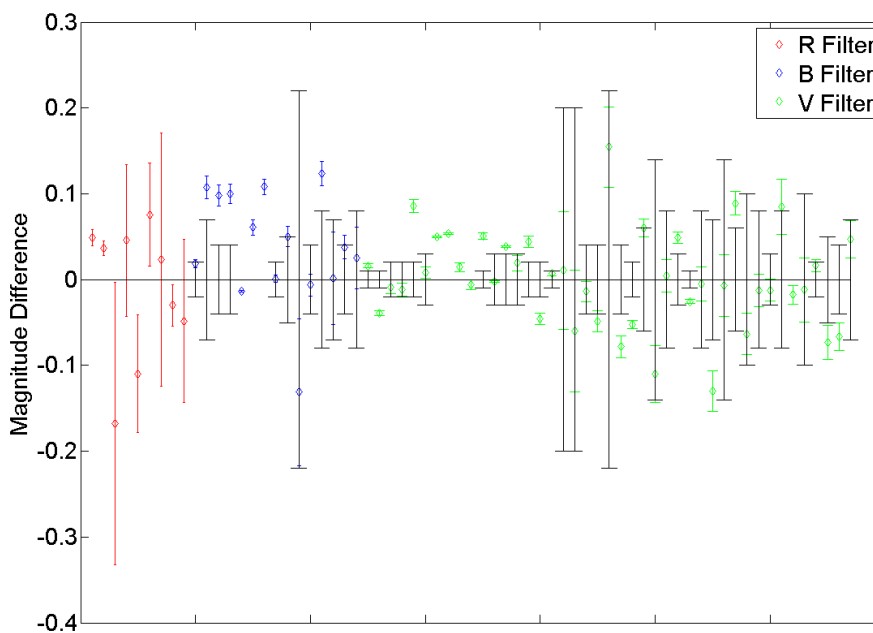

Figure 5.3: Difference between measured and published magnitudes of background stars. The horizontal axis is arbitrary.

The measured magnitudes are all very close to the published values. The errorbars of 44 of the 67 data points (66%) at least partially overlap the errorbars associated with the respective star's published magnitude; including all points for which the error is unknown. This is a positive result because one standard deviation from the mean, by definition, contains 68% of the data in a Gaussian distribution. Also, the data points appear to be distributed symmetrically about the zero error line. There is possibly a slight positive bias in the B filter data, however this could also be a

result of the limited size of the data set.

The typical error in the satellite magnitude measurements is likely to be even less than the error displayed in figure 5.3 for two reasons. First, the star magnitude measurements involved fitting an elliptical aperture to the cigar shaped star streaks, which introduces additional measurement errors. Second, background star measurements were only made in images with exposure times less than or equal to two seconds. The shorter the exposure time, the greater the error due to shutter effects. In addition, the stars for which R band magnitudes are known were usually relatively faint compared to other known stars. Therefore, the SNR of these stars was relatively low, resulting in larger than usual errors.

5.3 Comparison of Different Satellites' Lightcurves and Colour Variations

As discussed in chapter 2, the assumption that a material's SED is invariant with changing illumination and observation geometry has been instrumental in the development of techniques to discriminate between and characterise different satellites. If this assumption is correct, then the SED of the light reflected off a satellite, as measured by an observer, could only change as a result of different materials on the satellite becoming illuminated, or the relative flux off one material increasing relative to another, such as when a specular reflection off a satellite's solar panels dramatically increases the proportion of the measured light intensity originating from the solar panels. Given that the orientation of three-axis stabilised GEOs remains approximately constant with respect to an observer on the surface of the Earth, except for the solar panels, which constantly rotate to face the sun, each colour ratio would be expected to remain fairly constant throughout each night. In demonstrating the importance of the sun-satellite-observer geometry to the spectroscopic or photometric characterisation of satellites, Bedard even cited the repeatability of his satellite colour ratio curves on different nights as evidence for the high quality of his measurements [53]. One could surmise that this is the reason why Payne has declared that colour plots are not as effective for SOI as other techniques [30].

The discovery during this thesis that colour ratio curves often showed significant

variation in shape from one night to another was therefore viewed initially as a possible indicator of significant systematic errors. Fortunately, given that on almost all nights more than one satellite was visible in each image, the respective satellites' colour curves could be compared, allowing for the identification of possible systematic errors where significant features in a given plot are replicated in the other satellite's corresponding plot.

(a) Intelsat 805 B-I Colour Variation. (b) Galaxy 11 B-I Colour Variation.
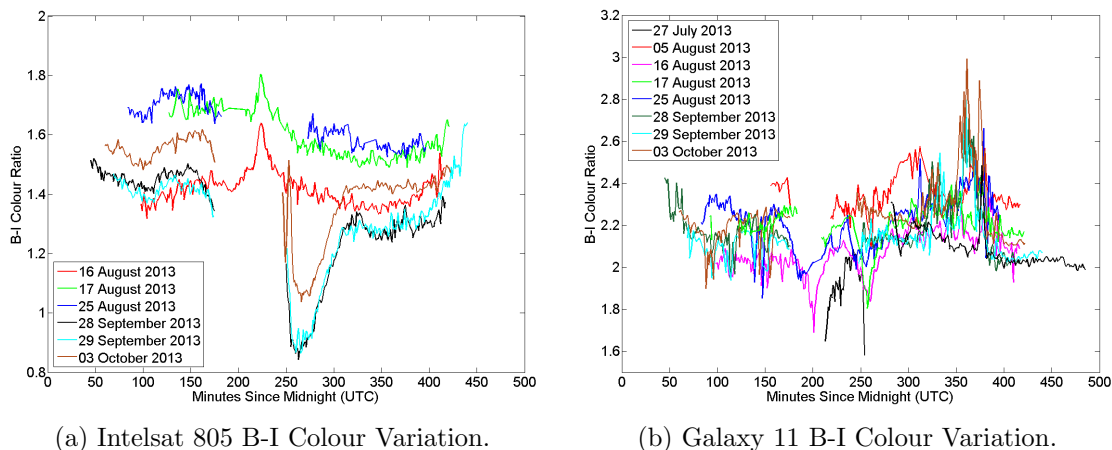
Figure 5.4: B-I Colour Ratio Curves for Intelsat 805 (a), and Galaxy 11 (b).

Figure 5.4 compares Galaxy 11's B-I colour ratio plot with that of Intelsat 805. Galaxy 11 and Intelsat 805 were captured together in the same images on all nights except for 27 July and 05 August, when only Galaxy 11 was imaged. Therefore any systematic errors in data capture or processing that uniformly affect the entire image should manifest themselves in similar ways in both satellite's colour plots and lightcurves. Examination of figure 5.4a reveals that aside from the overall colour plots shifting up or down between subsequent nights, there is a striking difference in shape between Intelsat 805's first three nights' plots and the last three night's plots between about 250 and 300 minutes since midnight. Specifically, the colour curve changes abruptly to become considerably more blue during this period on the last three nights compared to the first three nights. If this was the result of a systematic error, then a similar effect should be noticeable in the Galaxy 11 plots. Examination of figure 5.4b, however, reveals that the Galaxy 11 colour ratios are largely stable during this period on the last three nights. There is certainly no abrupt change in

colour ratio at all on those nights.

In addition, figure 5.4b shows that a temporary increase in the relative intensity of the longer wavelengths from Galaxy 11 occurs between 300 and 400 minutes since midnight. This hump in the curve, although having a noisy appearance, has a consistent shape over several nights, steadily growing in amplitude. Conversely, the Intelsat 805 colour ratio remains relatively stable during this period on all nights. Similar observations can be made about the colour curves derived from the other filters, which all share similar shapes to those depicted in figure 5.4.

Anik F1R and Anik G1, whose data were extracted from the same images as each other, also provide an effective opportunity to test for evidence of systematic errors. Close inspection of figure 5.5a reveals that the 29 October and 04 November curves for Anik F1R overlap almost perfectly, and that the 20 November curve is of similar shape, but offset consistently by a factor of about 0.4. The 12 October curve, however, closely matches the 20 November curve until about 220 minutes since midnight, at which point it abruptly changes to more closely follow the 29 October and 04 November curves. There is also a feature in the 12 October curve at 440 minutes since midnight, indicating a much more blue reflection, that is not present on the other nights. At this point one may be suspicious that the 12 October curve shows evidence of systematic errors.

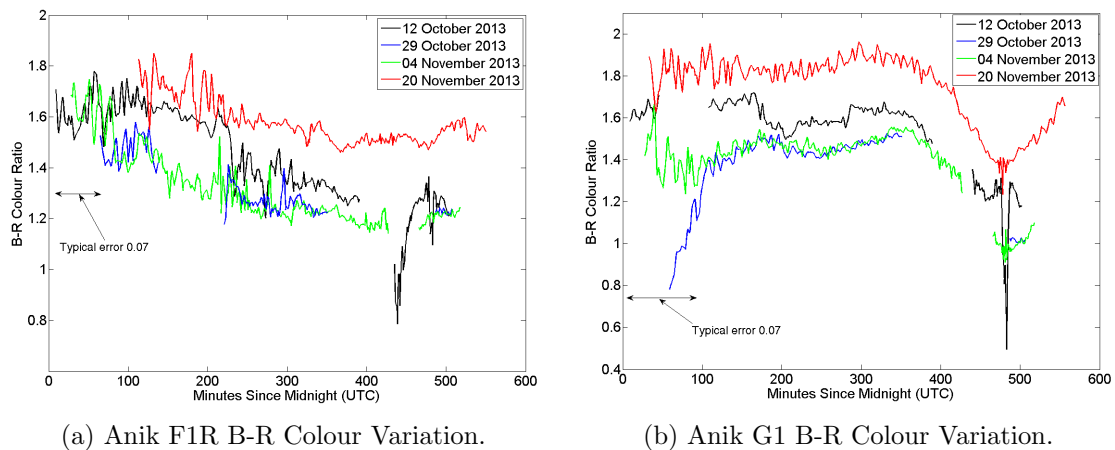(a) Anik F1R B-R Colour Variation.   (b) Anik G1 B-R Colour Variation.

Figure 5.5: B-R Colour Ratio Curves for Anik F1R (a), and Anik G1 (b).

Comparison with the Anik G1 data (figure 5.5b) is helpful. Anik G1's 12 October

curve does also show an abrupt blueward change, however it occurs prior to 220 minutes since midnight. Near 220 minutes since midnight the 12 October curve is essentially flat. Again, towards the end of the night there is a significant feature indicating much more reflection in the blue band, however it occurs at about 485 minutes since midnight, rather than 440 minutes since midnight. In addition, Anik G1's 29 October and 04 November curves diverge smoothly prior to 100 minutes since midnight, whereas the respective curves for Anik F1R overlap during the same period. Similar observations can be made about the colour plots derived from the other filters, and also for Anik F1 and Echostar 17, which were both also present in the same image sets. The reader is referred to Appendix A for the complete set of colour plots for all satellites.

Comparing the respective colour ratio curves of satellites that are derived from the same sets of images it is clear that the major features that differ from one night to another are unlikely to be the result of systematic errors that have a uniform effect across each image. The presence of such unique and differing features in the colour plots of satellites that are derived from the same images could only be the result of systematic errors if those errors had a non-uniform effect across each image. This seems particularly unlikely given that Anik F1, Anik F1R and Anik G1 are extremely close to each other in each of the images.

Comparison of the colour curves to the satellite's lightcurves reveals a possible explanation for the changing features present in the colour curves, which will be discussed in detail in Chapter 6.

Chapter 6

Discussion

This chapter addresses the questions that were raised at the conclusion of the litera-ture survey in chapter 2. It aims to establish some basic principles that might help direct the methodology of future experiments, including whether satellites maintain consistent lightcurve shapes over time, and the potential significance of changing SED on photometric analyses. The use of colours in satellite discrimination and character-isation, the feasibility of modelling satellites as simple structures, and the minimum temporal resolution for satellite characterisation are also investigated.

6.1   Examination of the Concept of Lightcurve Classes

One of the more fundamental, and widely accepted, proposals regarding satellite characterisation is that each satellite can be classified into one of five groups according to the shape of their lightcurves [17, 18, 40, 51, 54]. These lightcurve categories were described in section 2.3, and for brevity they will not be repeated here. Although it is understood that a particular satellite's glints will vary in magnitude at different times of the year, it is commonly expected that the overall lightcurve class will remain the same because it has been concluded that satellite lightcurves are tightly correlated with their basic bus structure [18]. There have occasionally been results published, however, that cast doubt on this conclusion. It has been noted that to avoid lightcurve mischaracterisation seasonal lightcurve variations should be considered [54]. Also, some measured lightcurves have differed greatly from their expected shapes based upon the satellite bus structure [51]. A lack of data seems to have prevented a thorough analysis of the consistency of satellite lightcurve shapes throughout an entire year.

This thesis, in focussing on obtaining detailed, long period observations on many different nights, presents a good opportunity to assess how well the data supports the

idea that satellites have consistent, characteristic lightcurve shapes. Figure 6.1 depicts the B band lightcurves of Anik G1 and Echostar 17, which were imaged on four and three separate nights, respectively. Although the peak magnitude during the specular glints varies considerably, the lightcurves, which span a period of approximately six weeks, appear to have smoothly varying, consistent shapes, supporting the idea of unique lightcurve classes.
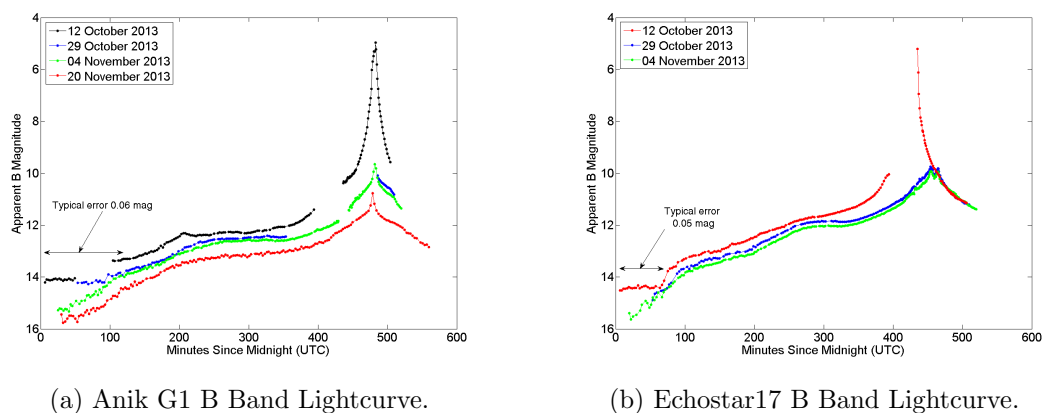

(a) Anik G1 B Band Lightcurve.

(b) Echostar17 B Band Lightcurve.

Figure 6.1: B Band Lightcurves for Anik G1 (a), and Echostar 17 (b).

Although these two sets of lightcurves do not contradict the idea of unique lightcurve classes, they only represent a small number of observations conducted over six weeks and four weeks respectively. Galaxy 11, on the other hand, was imaged on eight nights over a period of time exceeding two months. Although it still only represents less than one fifth of a year, the lightcurve seems to evolve over time in an interesting, and not so predictable way. There are three main features of the Galaxy 11 lightcurves, depicted in figure 6.2, that change considerably over the two months.

There are two relatively smooth glint features centred at approximately 200 and 260 minutes since midnight, and one glint feature centred near 370 minutes since midnight that has a much more noisy appearance. The sizes of the two smooth glint features, above the base magnitude level, appear to maintain the same ratio to each other as they vary from one night to another. They also seem to have an inverse relationship to the size of the noisy glint feature. Payne describes Galaxy 11's lightcurve as belonging to the Peculiar class, rather than the BSS702C class despite
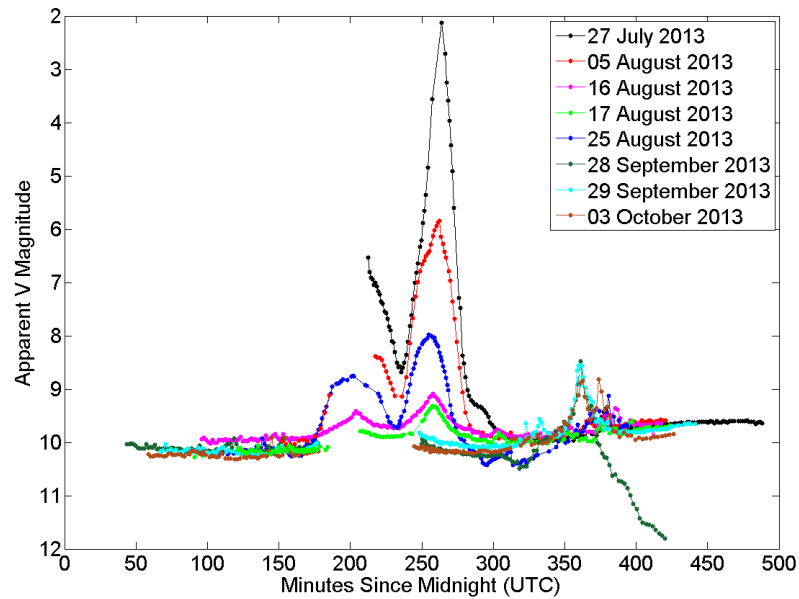
Figure 6.2: Galaxy 11 V Band Lightcurve.

it having a BSS702C bus, and also despite her assertion that lightcurve classifications are strongly correlated to basic bus type [18]. Figure 6.2 would seem to support classifying Galaxy 11's lightcurve as Peculiar, casting doubt on the degree of correlation between bus type and lightcurve class. Furthermore, when Galaxy 11's lightcurves are separated into individual plots for nights that are within one week of each other, as illustrated in Figure 6.3, the appropriate classification becomes less clear.

Close inspection of figure 6.3b reveals three smooth specular peaks, roughly evenly spaced in time but with greatly different amplitudes, centred near zero phase angle, which occurs at 214 minutes since midnight. These lightcurves would seem to be most similar to the Telstar category, although the secondary peaks are offset from the primary peak by much less than 40 degrees phase angle, and beyond about 320 minutes since midnight variations in the lightcurve take on a much more erratic appearance. Thus, the lightcurve might also be described as belonging to the Peculiar category. Examination of figures 6.3a and 6.3c reveals smoothly varying lightcurves with two specular peaks. There are no data for the 27[th] of July prior to 215 minutes since midnight because of the late sunset experienced during summer so it cannot be
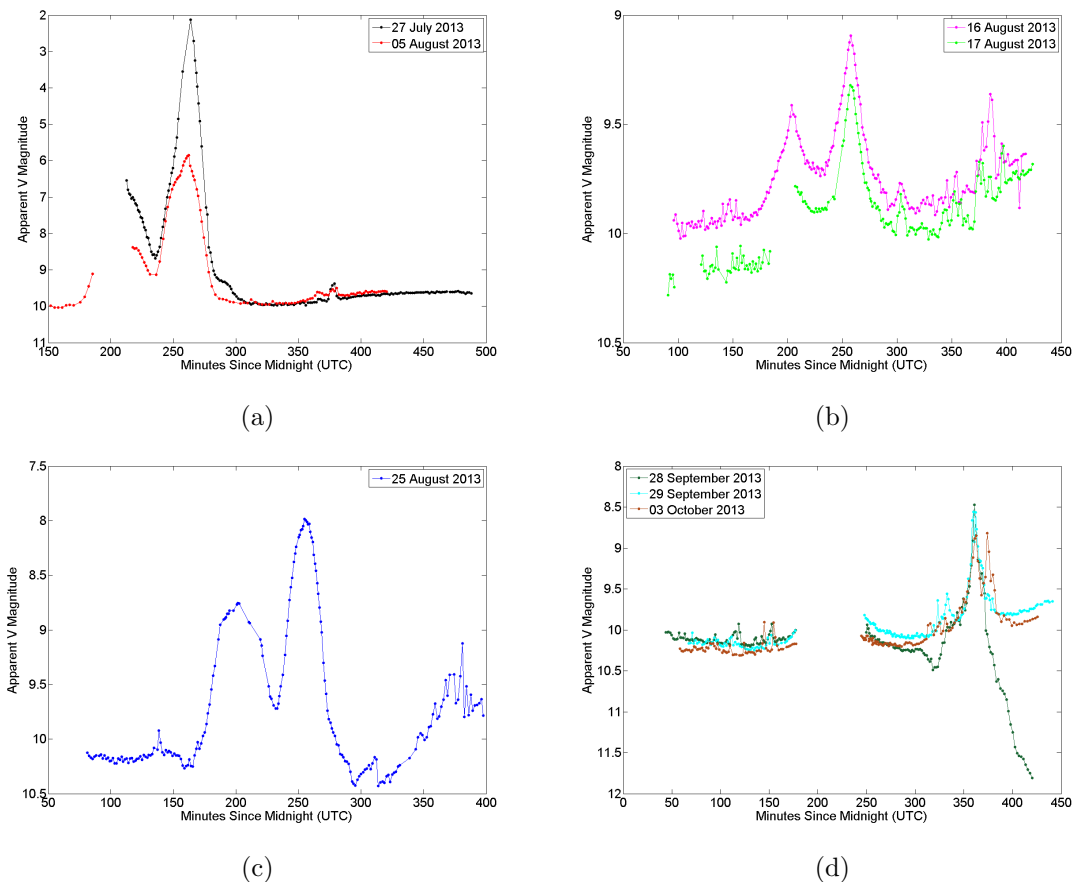
Figure 6.3: Galaxy 11 V Band lightcurves plotted on separate axes. Lightcurves with similar shapes are plotted together. The gaps in the data evident on the 05[th] and 17[th] of August are due to technical difficulties, and the gaps in the data in figure 6.3d are due to the satellite passing through Earth's shadow.

guaranteed that there are no more peaks prior to that time. There are no further peaks prior to that time apparent in the 05[th] and 25[th] of August data, however. None of the defined lightcurve classes perfectly describes these plots, however the most closely matching description is that of the A2100 class, which is characterised by two smooth peaks showing general but imperfect symmetry around a local minima near zero degrees phase angle [18]. Finally, figure 6.3d displays a completely different shape again. Where the primary specular peak is present in the other figures, near 260 minutes since midnight, there is no evidence of a glint at all. It does appear as though the eclipse period may have masked a broad peak that would have otherwise appeared near the one centred at 200 minutes since midnight in the other plots;

however it is difficult to know exactly what shape the peak would have. Also, the most prominent feature, present near 370 minutes since midnight, is located where there is relatively little or no discernable feature in the other nights' plots. Figure 6.3d is not symmetrical or smoothly varying, and could therefore best be classified as Peculiar.

Noting that these lightcurves only span a period of approximately two months, the fact that they might be classified in at least three different ways during this period indicates that a particular lightcurve class cannot be assigned uniquely to Galaxy 11, and that the appropriate lightcurve class is not tightly correlated to the satellite bus type.

Intelsat 805, based upon the AS-7000 bus [49], also arguably produces more than one specific lightcurve class. Figure 6.4 depicts the V band Intelsat 805 lightcurves. Figure 6.4a shows the two lightcurves for mid-August, and figure 6.4c the curve for late August. The gap in the 25 August curve is due to obstruction by clouds. The most appropriate classification based upon these curves is not entirely clear, however there is a peak near zero phase angle, which occurs at 214 minutes since midnight, and two slightly asymmetric smaller peaks either side of the primary peak; although the peak that occurs at the start of the night is not entirely covered by the 16th and 17th of August curves. It would therefore appear that the Telstar class would be the most appropriate based on figure 6.4a.

In contrast, figure 6.4b, depicting lightcurves for late September and early October, is much more asymmetric. The gaps in the data are due to the satellite passing into eclipse. Given the lack of symmetry, these curves would likely be classified as Peculiar.

Observations of both Galaxy 11 and Intelsat 805 occurred over a period of less than ten weeks. Despite this, each satellite produced more than one type of lightcurve class. In addition, the satellites' lightcurve classes do not necessarily seem to be very well correllated to their bus type. Although more data is required, ideally gathered over many nights throughout an entire year, it is clear that lightcurve shapes of individual satellites can change fundamentally with time. It seems likely that lightcurve classes

(a) Intelsat 805 V Band Lightcurve.

(b) Intelsat 805 V Band Lightcurve.

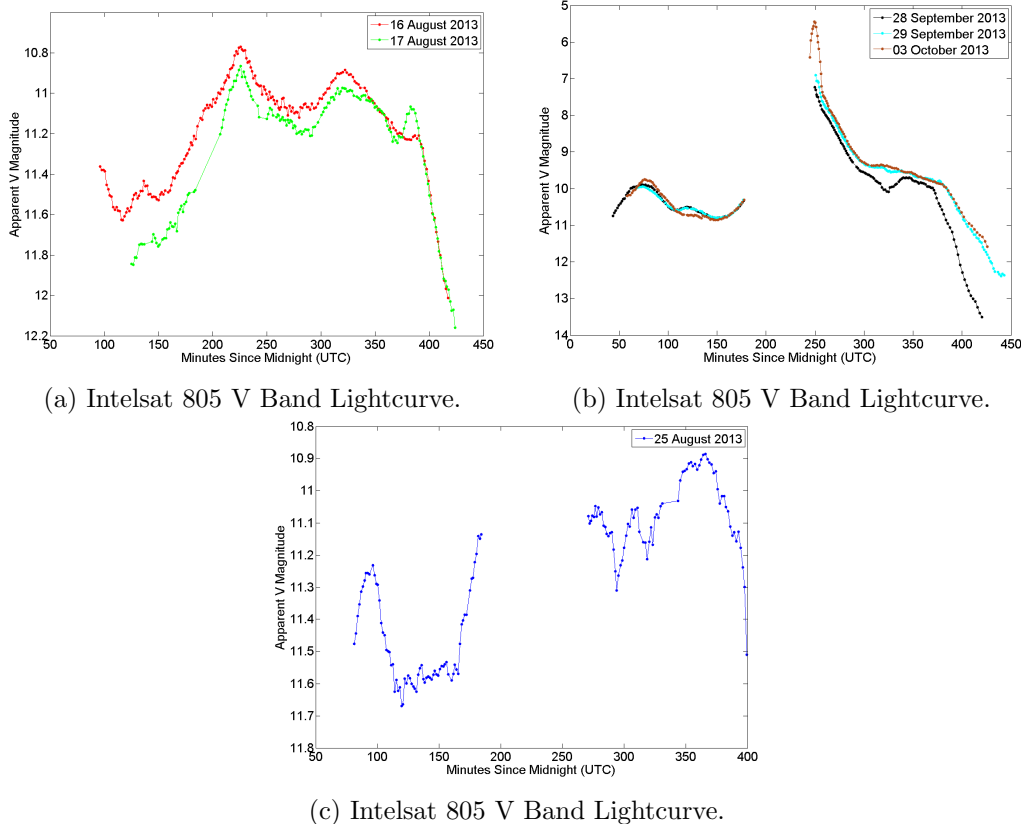(c) Intelsat 805 V Band Lightcurve.

Figure 6.4: Intelsat 805 V Band lightcurves on separate axes.

are not a very effective fundamental tool for grouping or classifying satellites. At the very least lightcurve classes would be specific to a particular period of time in at least some cases.

## 6.2   Variation of Spectral Energy Distribution

The invariance of a given material's SED is, as described in chapter 2, a key assumption that is fundamental to many proposed advances in the use of multicolour photometry for SSA. Bedard has presented strong evidence to suggest that this assumption is actually incorrect [36, 39]. The data collected during this experiment were analysed for evidence that supports or contradicts the assumption of invariable SED.

Colour ratio plots are useful when analysing a satellite's reflected SED because they give an effective visual indication of the evolution of the relative intensity through

one waveband compared to another waveband. If a particular satellite was made entirely of a single material, then identifying whether its SED changes with illumination-observation geometry would be simple: an unchanging SED would yield a flat colour curve. Regardless of changes in the overall light intensity, the colour ratio would remain the same. For a complex GEO satellite that is comprised of many different materials, however, the problem is not so easily resolved. Specular reflections can greatly increase the relative contribution of even a small satellite component's reflected light to the overall intensity. Therefore, a colour ratio that varies with time may simply be the result of specular reflections changing the relative contribution of certain materials to the overall reflected SED. Alternatively, self shadowing or changes in the attitude of either the entire satellite or one of its components, such as its solar array, can change the relative abundance of each material that is illuminated. Therefore, a more detailed analysis of the data is required to identify whether the SED of individual materials varies with illumination-observation geometry.

Examination of each satellite's colour plots near the brightest specular peaks reveals one interesting detail in particular. Each specular peak coincides with a change in colour ratio for every colour plot, for all of the satellites. The brightest specular reflections, however, exhibit a reversal in the colour ratio change, over a short time-frame, at the peak. Intelsat 805 provides a clear illustration of this point. Figure 6.5 depicts Intelsat 805's B band lightcurves above the B-R colour variations.

The most striking feature of the colour plot is the dramatic increase in the relative intensity of the shorter wavelengths between 250 and 300 minutes since midnight on the 28[th] and 29[th] of September, and the 3[rd] of October. Unfortunately, a significant gap in the data exists on these nights due to the satellite being eclipsed by Earth's shadow. The peak in the lightcurve was captured on the 3[rd] of October, and very narrowly missed during the September observations. By comparison to the B band lightcurve it is clear that the change in colour ratio begins to occur at the base of the specular peak, and that the colour becomes increasingly blue as the lightcurve magnitude increases. Within approximately 15 minutes of the peak, however, the colour reddens sharply to reach close to the same colour ratio as existed prior to the
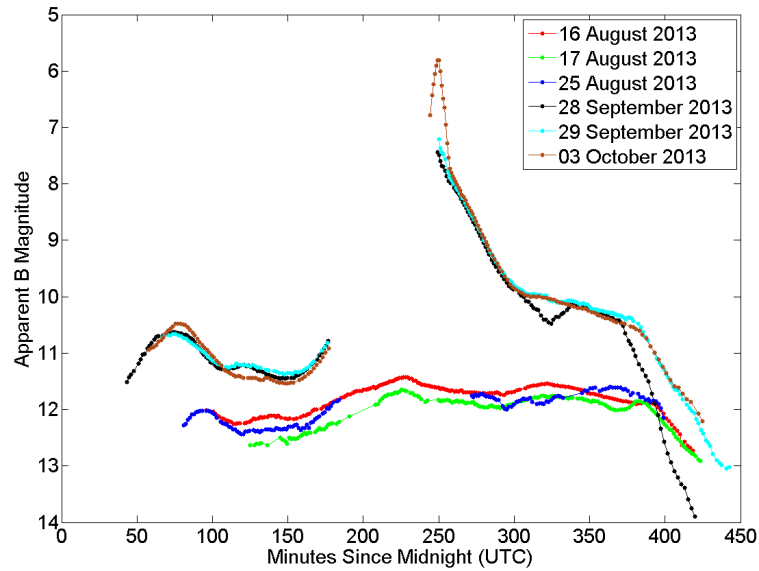
(a) Intelsat 805 B Band Lightcurve.
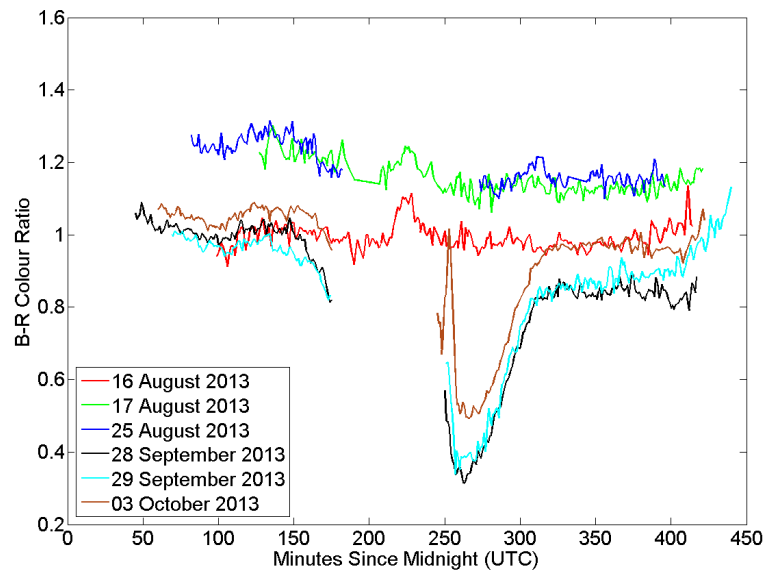
(b) Intelsat 805 B-R Colour Variation.

Figure 6.5: B Band Lightcurve for Intelsat 805 (a), compared with its B-R colour variation (b).

specular glint.

It is widely believed that the bright peaks common to most GEO satellite's lightcurves are caused by specular reflections off their solar panels, and there is much evidence to suggest this is correct. The true cause of the peaks is not important to this discussion, but for now it will be assumed that they are indeed caused by solar panel reflections. It seems intuitive that a dramatic increase in the solar panel's contribution to the overall intensity would lead to a more blue colour ratio because solar panels usually appear blue to the naked eye. The sudden reddening near the peak, however, is not consistent with this reasoning. There are three possible causes for the colour ratio reversal:

1. A second specular reflection, off a facet other than the solar array, is being observed;

2. The relative increase in shorter wavelengths is due to the SED of light reflected off the solar array changing with the illumination-observation geometry; or

3. A combination of both of the first two explanations.

It is important to note that these possible explanations are valid regardless of the source of the primary specular peak. If the first explanation is the correct one, then the plane of the facet causing the reddening must be very closely aligned with the plane of the solar array because the local maxima near 250 minutes since midnight on the colour plot coincides with the specular peak on the lightcurve. Also, the cone of specular reflection off the second facet must be narrower (more mirror-like) than off the solar array because the reddening occurs over a shorter time interval than the increase in blue due to the solar panels. In addition, very close to the specular peak, the increase in intensity at the red end of the spectrum occurs more steeply with time than the increase in intensity at the blue end of the spectrum. Therefore, the facet causing the reddening must be either larger than the solar array or exhibit much greater reflected flux at longer wavelengths than shorter wavelengths. As depicted in figure 6.6 there are no facets larger than the solar array on Intelsat 805, as is common

amongst GEO satellites [49]. Also, the satellite is stabilised such that the antennas always point towards Earth so that communication can continuously be maintained. Therefore, figure 6.6 would seem to indicate that the gold coloured MLI, which is known to be relatively red [39], does not present a significant flat surface towards Earth that might be capable of producing a bright specular reflection to rival that of the solar array.
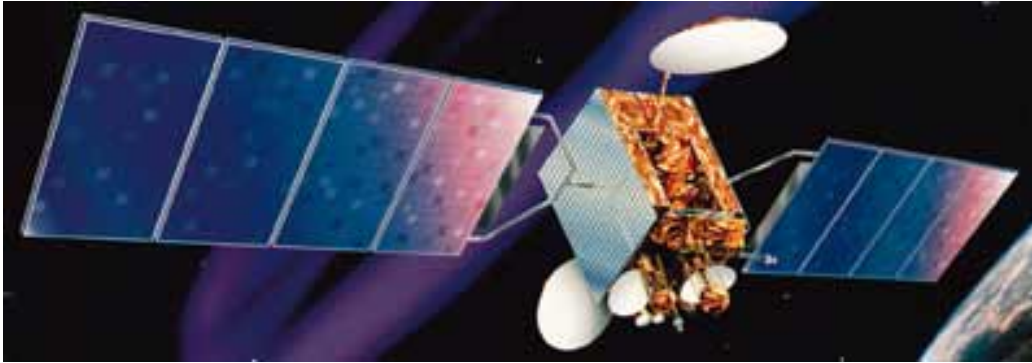

Figure 6.6: Artist's depiction of Intelsat 805 [49].

It seems unlikely that the colour ratio variation across the primary specular peak could be due to satellite components other than the solar panels. The alternative explanation, that it is due to the SED varying with changing illumination-observation geometry, fits more easily with observations made during this thesis, and matches the conclusions drawn from Bedard's doctoral work [39]. Figure 6.7 shows how colour ratios of one particular Triple Junction Photovoltaic Cell (TJPV) solar cell change with observation angle, at a constant illumination angle of 10° [39]. Although the plots only display a range of observation angles of less than two degrees, it is clear that outside of the very narrow cone of specular reflectance the colour ratios become sharply more blue. At the specular peak, the B-R colour ratio value is close to one.

Although the Intelsat 805 solar array is likely to use different solar panels to the one corresponding to figure 6.7, all of the four solar cells tested by Bedard yielded similar results. If the cause of the reversal in Intelsat 805's colour ratio change is varying SED, and if the satellite's solar array has similar reflectance properties to that measured by Bedard in figure 6.7, then the evolution of the B-R colour ratio with observation angle might be explained as follows:
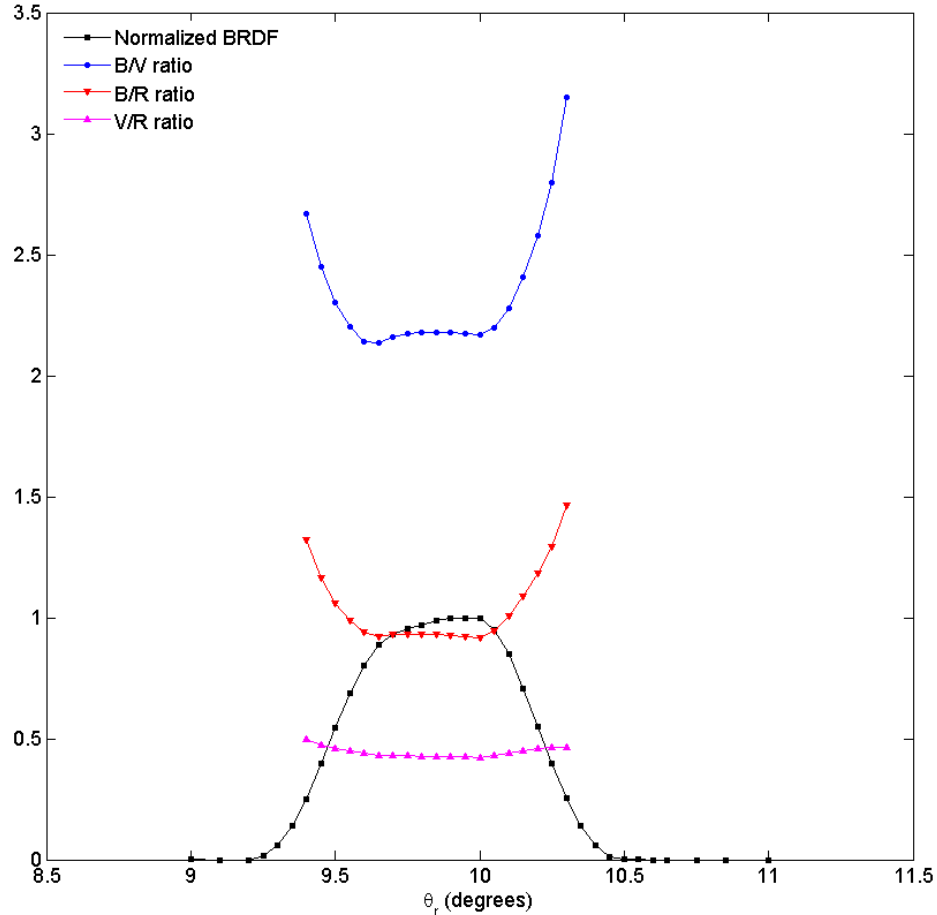
Figure 6.7: Normalised broadband BRDF measurement and colour ratios of a TJPV cell for varying observation angle (θ_r) at constant illumination angle of 10° [39].

- At all times the solar panels track the sun, thus maintaining a constant illumination angle of close to zero;

- At the edge of the primary glint feature the observation angle is such that the specular reflection off the solar array is just beginning to become visible.

- As the observation angle decreases, the contribution from the solar array to the total measured flux increases markedly. Because the SED of the light reflected off the solar array is relatively more blue at these observation angles, the overall colour ratio of the satellite becomes steadily more blue.

- When the observation angle decreases further, to within approximately two degrees of the centre of the cone of specular reflection, the SED abruptly becomes

more even across the visible spectrum. With a B-R colour ratio close to 1, as in figure 6.7, the satellite's total B-R colour ratio returns to a value similar to those measured prior to the specular reflection, as seen with Intelsat 805 in figure 6.5b.

- After the observation angle passes beyond the centre of the cone of specular reflection, the colour ratio evolves in the reverse of the manner just described.

The colour ratio reversal feature has also been observed in Anik F1, Anik G1, and Echostar 17 to varying degrees. In figure 6.8 the feature is most obvious in the 29 October and 04 November plots near 450 minutes since midnight. The reason that the feature is broader than observed in Intelsat 805 may be that Anik F1 displays three separate, closely space specular peaks.

Anik G1 also displays a very sharp reversal in colour ratio, similar to Intelsat 805, centred at approximately 490 minutes since midnight, as can be seen in figure 6.9.

The peak of the specular reflection was not observed for Echostar 17 due to shadowing by Earth, however a sharp reversal in the colour ratio change is apparent in the 12 October data, as depicted in figure 6.10.

This work is consistent with Bedard's carefully conducted experiment, strongly indicating that the SED of light reflected off a particular material, in this case solar panels, varies with changing illumination-observation geometry, and that the changing SED affects observed colour ratio curves. This conclusion has significant implications for previous work that has been conducted in the analysis of multicolour photometry of satellites, particularly where the assumption of invariant SED has been expressly fundamental to that work. On the other hand, knowledge of how different materials' SEDs change with observation-illumination geometry may provide new means of enhancing SSA. For example, it has been well demonstrated that examination of a satellite's lightcurve as a function of longitudinal phase angle can reveal the angular offset between the solar panel normal vector and the satellite-sun vector [18]. Such information constrains the solar panel orientation about one of its axes. When observed, the extremely narrow cone of specular reflection within which a material's
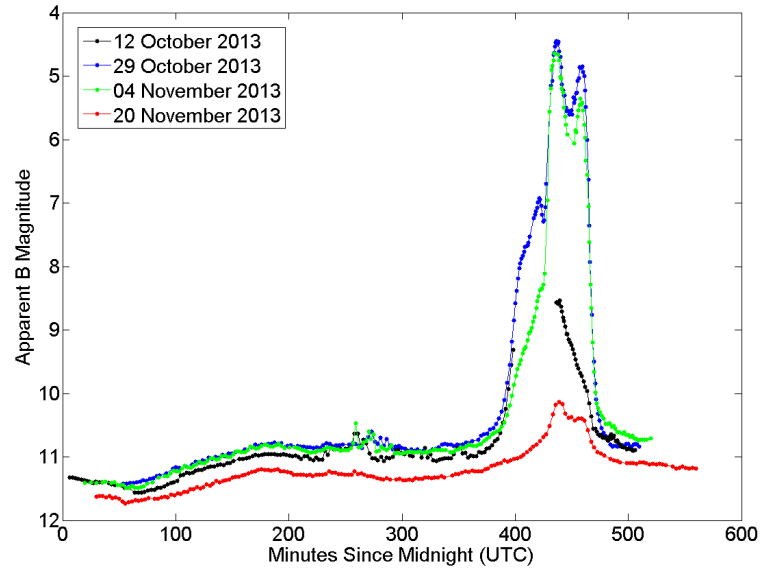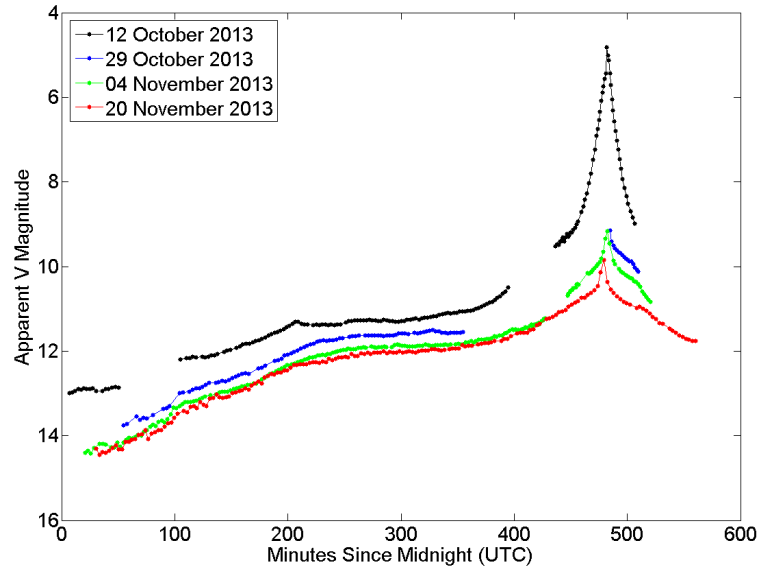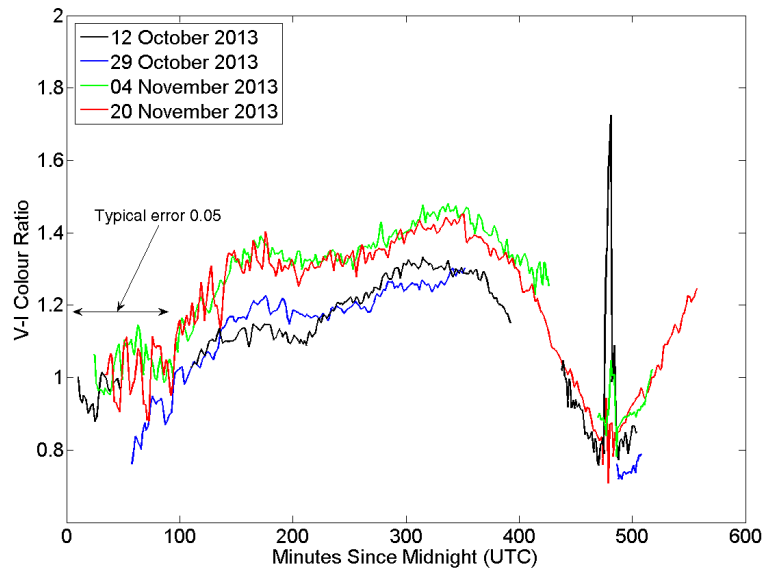
(a) Anik F1 B Band Lightcurve.
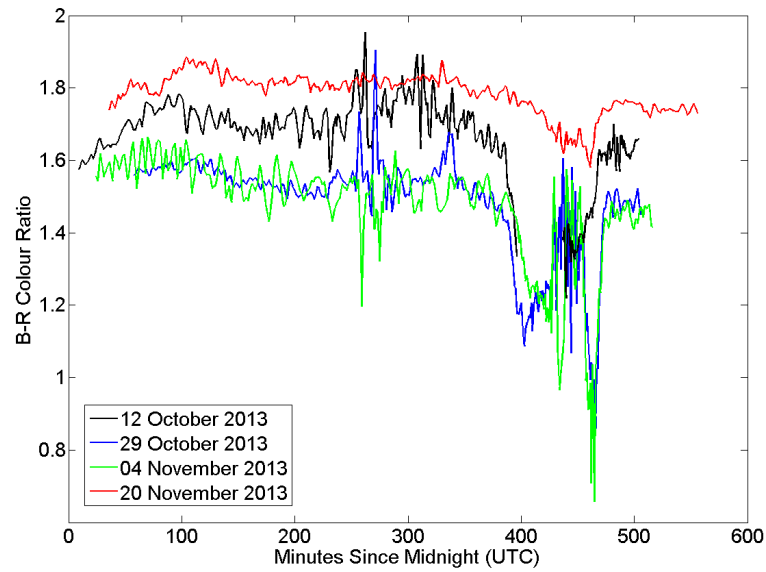
(b) Anik F1 B-R Colour Plot.

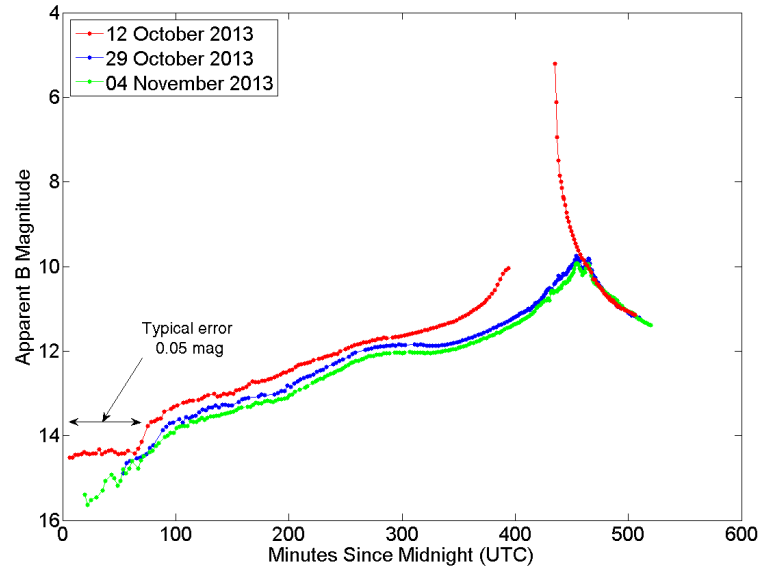Figure 6.8: B Band Lightcurve for Anik F1 (a), compared with its B-R colour variation (b).

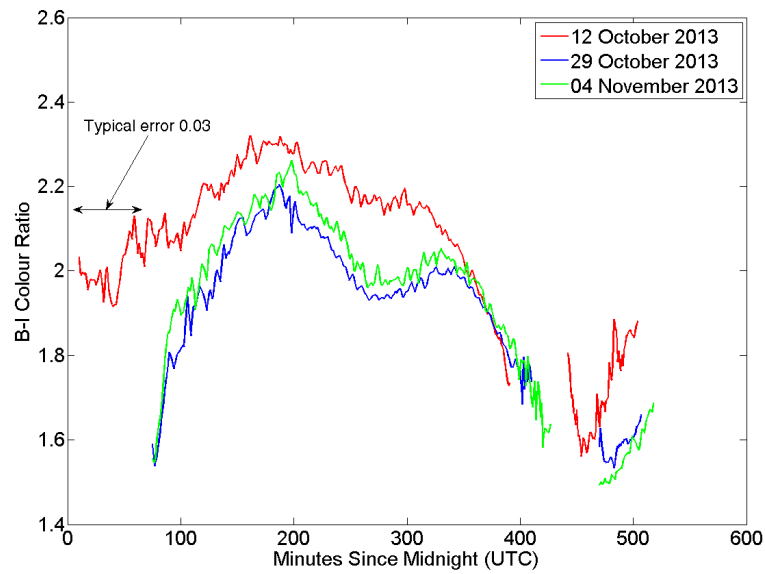(a) Anik G1 V Band Lightcurve.

(b) Anik G1 V-I Colour Variation.

Figure 6.9: V Band Lightcurve for Anik G1 (a), compared with its V-I colour variation (b).

(a) Echostar 17 B Band Lightcurve.

(b) Echostar 17 B-I Colour Variation.

Figure 6.10: B Band Lightcurve for Echostar 17 (a), compared with its B-I colour variation (b).

colour ratio suddenly changes can also place tight constraints on the attitude of a GEO satellite about its East-West axis. The solar array plane can thus be determined in three dimensions rather than just two.

Further examination of figure 6.5, in light of the discussion on varying SED, reveals another interesting feature. On the 16[th] and 17[th] of August, the primary specular peak is relatively faint. The associated colour plots do not display the colour reversal that was noted on the other nights, and the colour ratio change occurs in the opposite manner. Instead of the colour steadily becoming more blue as the specular peak is approached, a slight reddening is observed. The same colour ratio behaviour was noted by Payne *et al.* in 2007 [7] when the GEO satellite DTV 1R was observed in March and June. In March the colour ratio was observed to become bluer, whereas in June it became redder. The temporal resolution of Payne's data was not sufficiently high, however, to determine whether there was a reversal in the colour ratio change similar to that observed with Intelsat 805. Payne did not offer an explanation for the observed phenomena.

At first glance the Intelsat 805 colour plots for the 16[th] and 17[th] of August seem to contradict the discussion about SED variation: increased contribution from the solar array to the total flux reddens the observed reflection, rather than making it more blue. There is, however, one variable of particular importance that changes continuously with the seasons: the illumination angle. Although Intelsat 805's solar panels continuously rotate to track the sun, their axis of rotation remains fixed relative to Earth's equatorial plane; usually close to perpendicular to it. As a result, seasonal changes in solar declination result in variable solar panel illumination angles. On nights when there is a relatively small specular peak observed, usually at or near the solstices, the illumination angle is the greatest. At or near the equinoxes the illumination angle is smallest, and the specular peak is the brightest. The solar panel illumination angle on the 16[th] and 17[th] of August was therefore significantly greater than on the 28[th] and 29[th] of September, and the 3[rd] of October. At this point, further examination of Bedard's solar panel colour ratio plots is instructive. Figure 6.11 depicts various colour ratios of the same solar cell referenced in figure 6.7, as a

function of observation angle, for incidence angles of 30° and 60°.
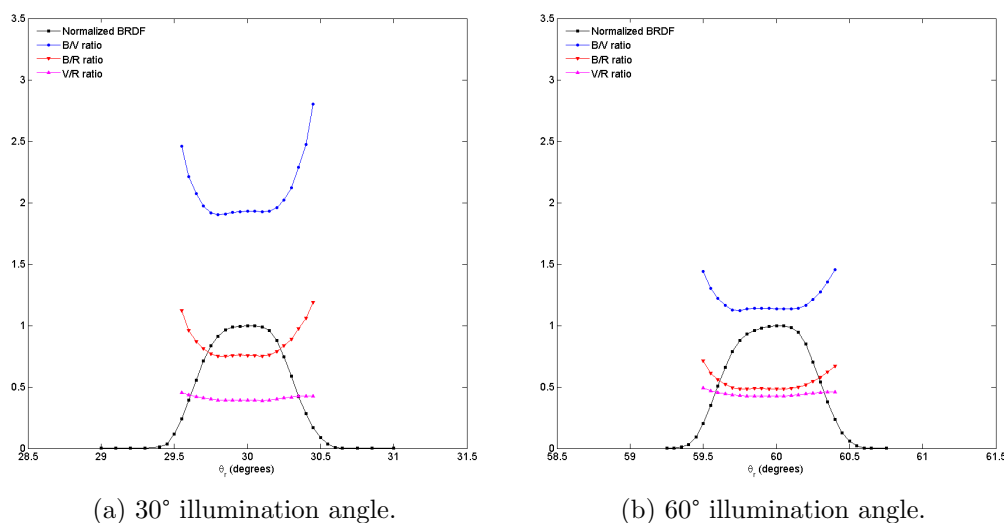

(a) 30° illumination angle.

(b) 60° illumination angle.

Figure 6.11: Normalised broadband BRDF measurements and colour ratios of a TJPV cell for varying observation angle (θ_r) at constant illumination angle of 30° (a) and 60° (b) [39].

Of note, the B-R colour ratio in figure 6.11b is lower across the entire curve than in figure 6.11a, indicating a redder colour. Both of these curves depict redder colour ratios than the B-R curve in figure 6.7. What these figures show is that higher illumination angles produce redder colour ratios than lower illumination angles for any given observation angle. At an illumination angle of 60° the observed colour ratio was relatively red across the entire range of observation angles measured by Bedard.

The variable nature of Intelsat 805's B-R colour ratio curves can be explained if the satellite's solar cells have similar reflectance properties to the TJPV cell measured by Bedard. At times of the year near the summer and winter solstices the angle at which the sun's light illuminates a GEO's solar array is the greatest. The specular reflection is thus directed away from Earth's equator; to the north during the northern hemisphere winter, and to the south during the northern hemisphere summer. On the 16[th] and 17[th] of August the specular reflection off Intelsat 805's solar array was therefore directed away from the observing location in Canada, causing relatively faint glint features. Since the illumination angle was large on those nights, the extra contribution made by the solar panel glint to the measured flux, above the level measured without the specular reflection, consisted of more longer wavelength light

than shorter wavelength light. Thus, the overall B-R colour ratio curve exhibits only a slight, brief reddening. These results indicate that a material's SED may vary not only with observation angle, but with illumination angle as well.

6.3 The Use of Colour Ratios in Satellite Discrimination and Characterisation

A clear implication of these results is that any SSA effort that involves the use of colour ratios can only be applied to space objects whose attitude is accurately known, along with the illumination-observation geometry at the time of each observation. This point is illustrated by a study conducted by Payne *et al.* in 2001. As was described in chapter 2, these authors compared magnitude versus colour variations of a cluster of GEOs, using Johnson filters and specially designed SILC filters, to conclude that the SILC filters developed by the AFRL were superior to the Johnson filters for satellite discrimination [15]. Figure 6.12 is an example of one such plot.
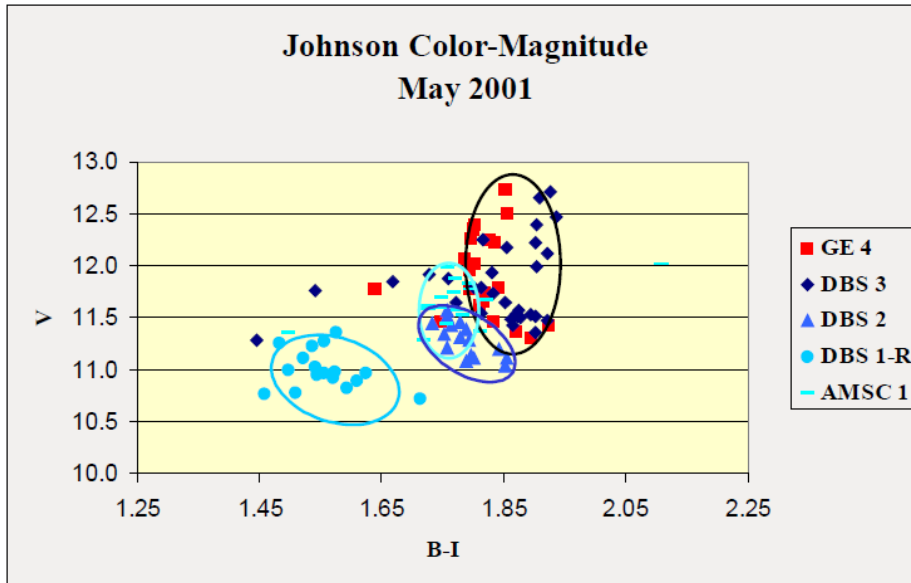

Figure 6.12: Magnitude versus colour plot used for satellite discrimination [15].

Specific details about the timing of the measurements used to construct the magnitude versus colour plots were not provided, however it can be seen that only approximately 12 to 24 measurements were made for each satellite. When data are plotted

throughout the night, for multiple nights, it becomes far more difficult to distinguish between different satellites. Figure 6.13 depicts Anik F1R and Echostar 17 measurements from this thesis, from three nights each, plotted on the same magnitude versus colour plot used by Payne et al.
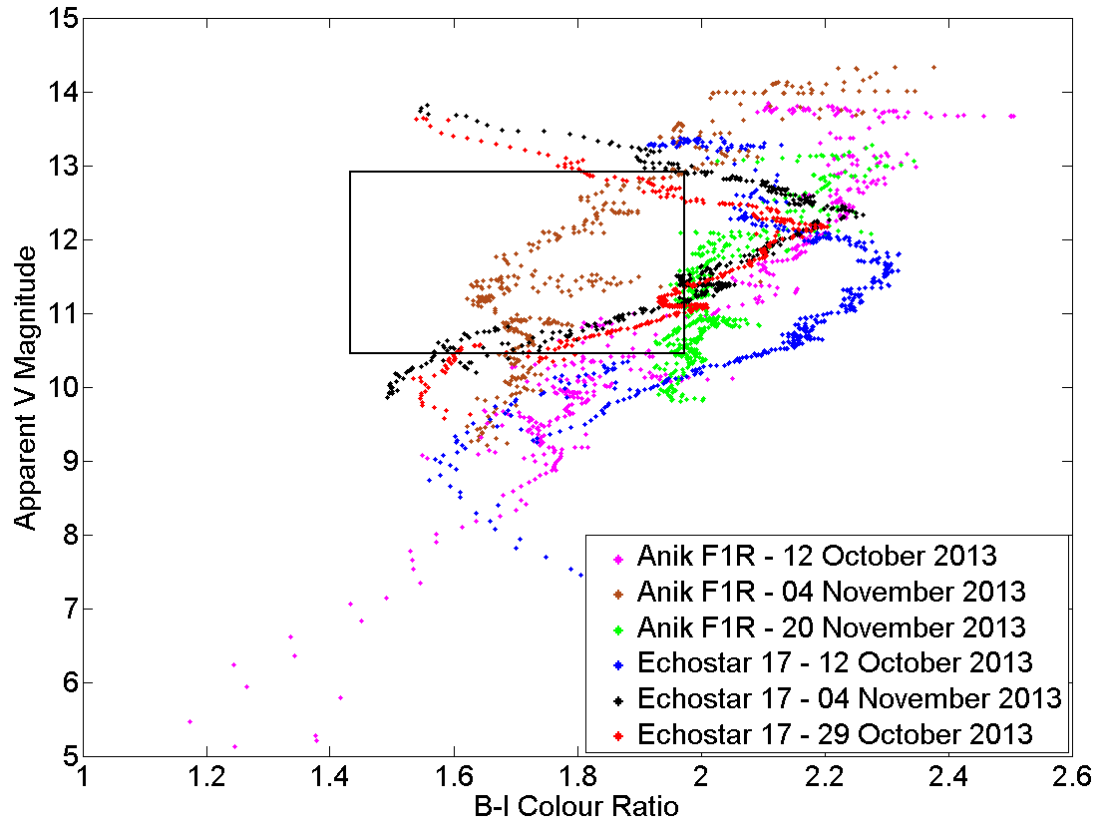

Figure 6.13: Plot of magnitude vs colour for Anik F1R and Echostar 17.

The data represent measurements taken throughout each night for each satellite. On the 12[th] of October and the 4[th] of November the data for each satellite were extracted from the same set of images, and thus represent simultaneous measurements. The black rectangle in the centre of the image represents the boundaries within which all of Payne's data from figure 6.12 reside. Also, Anik F1R and Echostar 17 have significantly different structures. Anik F1R is based upon Astrium's Eurostar 3000S bus, whereas Echostar 17 is based upon the Space Systems Loral SSL-1300 bus [49].

Three things are immediately apparent upon inspecting Figure 6.13. First, each satellite's data points are not tightly grouped in one small area as might be expected

from Payne et al.'s analysis of their data. The range of values for each satellite on any of the nights is far greater than the range of values presented for all five of the satellites analysed by Payne together. Second, the location on the plot occupied by each satellite's data points varies significantly from one night to another. Third, there is considerable overlap between the areas on the plot occupied by each satellite. It is clear that it would be impossible to differentiate between Anik F1R and Echostar 17 based upon a small number of measurements plotted on a magnitude versus colour plot, despite the significant differences between their physical structures. In analysing the data obtained throughout this work, the wide variation of magnitude versus colour values exhibited by Anik F1R and Echostar 17 were found to be characteristic of all of the satellites observed.

The observations made regarding the variable nature of a given material's SED are important to any photometric study of satellites that relies upon colour ratio measurements. Although it is widely understood that a satellite's observed colour ratios may change throughout a given night, the factors that influence the colour ratios have not been thoroughly investigated. The current, longstanding, assumption that any material's SED is invariant has led to the conclusion that a satellite's observed colour ratios are purely functions of the materials that are visible and the relative intensity of the light observed to be reflected off each material. For example, a relative increase in red wavelengths is commonly associated with either specular reflections off gold coloured MLI or a greater surface area of MLI becoming visible; whereas a relative increase in blue wavelengths is associated with increased reflection off solar panels. This work has shown that illumination angle and observation angle are both very important factors in determining observed colour ratios: reflections off solar panels may result in either redder or bluer colours. In addition, Bedard's doctoral work identified that almost all materials commonly used on spacecraft, not just solar panels, exhibit significant colour changes with varying illumination and observation angle. Each material's surface roughness was also found to be a significant factor in determining colour ratios.

These results imply that there is very limited scope to use colour ratios alone

for satellite discrimination or characterisation. Cowardin's 2010 experiment to try to identify two pieces of GEO debris by correlating their colour ratios to laboratory-measured colours of spacecraft material fragments was largely unsuccessful [55]. Only one single value was used for each different colour ratio, for each laboratory measured material. In addition, a very limited number of telescopic measurements were made of the GEO debris objects. Cowardin noted that a debris fragment's material composition could only possibly be identified if its attitude was relatively stable such that there were no large magnitude variations due to specular glints. Unfortunately, such a constraint does not limit the reflectance properties of the debris object, it merely limits what is observed. If the variation in colour ratio of each material is great enough such that their ranges overlap, then positively identifying a particular material without accurate knowledge of its shape and attitude, as well as the illumination and observation angles, would be impossible.

6.4   Modelling Satellites as Simple Structures

Most attempts to model satellites or simulate satellite lightcurves employ simplifications of either the satellite structure or the reflectance characteristics of its component materials. Given that solar panels are widely understood to produce the bright specular glints that are common amongst GEO satellites, one approach has been to model GEO satellites as consisting of only two components: a spacecraft body and a solar array. The spacecraft body is assumed to be a perfectly diffuse flat facet permanently facing Earth, and the solar array is assumed to be a specular reflector that tracks the sun. Sometimes additional facets are included in the models that represent antennas [17, 40]. Such simplifications are valuable for providing a means by which simulated satellite lightcurves can be generated without requiring unreasonable time or effort to do so. In many cases assumptions are required to be made to reduce the number of unknown variables to a manageable quantity. There is a trade off, however, between accuracy and simplicity when employing simplifications; the value in doing so depends upon the aim of the simulation.

Chaudhary and Payne used a two facet model to describe any three axis stabilised

GEO as a fundamental part of their method for extracting the spectra of the solar array and the spacecraft body from the non-resolved photometric data [40]. The relevant simplifying assumptions that were employed to enable the method to work are as follows:

- The satellite can be represented by two planar facets: one facet represents the solar array and the other facet represents the remainder of the satellite.

- The solar array has specular reflectance properties, and the body has perfectly diffuse reflectance properties.

- The body facet faces nadir, and the solar panels articulate to face the sun. It is not essential to the method for the solar panels to perfectly face the sun provided that any offset angle is consistently maintained.

- The solar panel contribution to the lightcurve is symmetrical about the specular peak; and thus the rate of change of the solar panel contribution is an odd function with respect to the specular peak.

- The rate of change of the body contribution is constant throughout the duration of the specular peak; and thus it is an even function with respect to the specular peak.

- Any material's SED is invariant with changing illumination-observation geometry.

Chaudhary and Payne acknowledge that their assumptions' imperfect fit with reality is a weakness of their technique, however the simplifications enable a great quantity of very valuable information to be obtained with relatively little effort. For example, by using the simplifications listed above, as well as some simple geometric arguments, Chaudhary and Payne describe how it is possible to extract the spectra of the satellite body and solar panels using only two measurements in each of the spectral wavebands. Unfortunately it is very difficult to know whether a procedure such as Chaudhary and Payne's is capable of producing sufficiently accurate results.

The authors observe that when the calculation is repeated using multiple different suitable pairs of data points a clear trend emerges. However, even if the procedure is carried out multiple times with consistent results, it is possible that systematic errors may be manifesting themselves in similar ways every time.

Although not explicitly stated, there are other implicit assumptions made by Chaudhary and Payne. First, it is assumed that both of the satellite's solar panels are perfectly aligned with each other. Given the very steep magnitude versus phase angle relationship usually observed near specular peaks, even a very small misalignment between the two panels would likely have a significant effect on the measured magnitude, and thus on the overall result.

Also, the satellite body is usually comprised of many different types of materials. By claiming that it is possible to extract the satellite body spectra, Chaudhary and Payne are implying that the combined spectra from all of the different materials that the body is composed of remains constant throughout the range of possible illumination-observation geometries. Specular components of the reflection off antennas, MLI, and other materials are likely to significantly change the actual total non-solar panel spectra that are directed towards the observer from one moment to another. This effect is apparent in all of the colour plots for all of the satellites that were observed throughout this experiment, particularly at the times when glints off the body were observed to occur. For the assumption that there exists a unique body spectrum to hold true, each colour ratio would need to remain constant across the respective colour ratio curves, outside the region where specular reflections off the solar panels are observed. For brevity, examples of colour plots will not be provided here. The reader is directed to Appendix A where the complete set of colour variations can be found. In addition to the colour ratios often changing from one night to another, it can be observed that throughout a given night the colour ratios change considerably. Sometimes the ratios change at a relatively constant rate throughout the night and other times they vary in a much more irregular fashion. The overall manner in which the ratios vary, however, does not appear to be linked to the general shape of the lightcurve. For example, Anik G1 and Echostar 17 both have relatively

120

smoothly varying lightcurves, that might be described as canonical. But Anik G1's colours mostly appear to vary with time in a more consistent manner than Echostar 17's colours, outside the region dominated by the solar array's specular reflection. Additionally, it is important to note that data were collected on these satellites on only four and three nights respectively, representing only a very small data set. Therefore it is uncertain whether it is accurate to assert that a particular satellite's colour plots generally vary in a consistent manner or not without specifying the particular nights that the claim applies to.

6.5   Temporal Resolution and Satellite Characterisation

In addition to the many simplifying assumptions that have already been discussed, it has often been assumed that the lightcurves of three-axis stabilised satellites are predominantly smoothly varying, or that additional useful information cannot be obtained by increasing the temporal resolution of the observations beyond what will reveal the large scale shape of the lightcurve. Payne stated so explicitly in 2001 [15], although she subsequently noted some exceptions [30]. Other researchers have implied the assumption by deliberately arranging observation cycles that result in low temporal resolution. Lederer conducted observations of the IDCSP satellites in 2012 that resulted in between 10 and 20 samples taken per satellite over a period of four separate nights [35]; LeVan proposed a satellite discrimination method in 2010 that requires only three observations per satellite over an entire night [6]; Vrba conducted observations in 2003 at a rate of approximately five observations per hour, per filter [24]; and Hejduk grouped single point photometric observations from different nights throughout the year into three-degree phase angle bins [29]. Some experiments have been conducted at the highest temporal resolution that the researcher's equipment would allow [41, 42, 56], however the utility of the higher sampling rate has generally been limited to studying the timing of small glints. Although it has been noted several times that the scatter in lightcurve and colour ratio data is greater than the errors of the individual measurements, to date no one has attempted to define the lowest temporal resolution that can yield the maximum useful characterising information.

One example from this experiment that illustrates information that can be obtained from high temporal resolution data is the rapid colour changes described in section 6.2. In figure 6.5b the slope of the 03 October colour plot changes sign three times in approximately ten minutes near 260 minutes since midnight. As described in that section, sharp peaks in colour ratio plots can be very narrow, but can yield information about the orientation of the plane of the solar array. The possibility of also inferring specific solar cell types from these plots remains to be investigated. If the temporal resolution of these data was less than approximately one observation every two minutes, then that information would be lost.

Careful examination of the Galaxy 11 lightcurves reveals indications that the temporal resolution of the data in this experiment is not sufficiently high to reveal all of the information that may potentially be of use. The lightcurves depicted in figure 6.3 appear relatively very noisy between about 300 to 400 minutes since midnight. Importantly, it is not suggested that the magnitude fluctuations in this region are the result of actual noise, rather that the appearance of the curve is consistent with what would be expected of noisy data. If the data were oversampled, then one would expect to see that even if there were rapid, frequent changes in magnitude, the curve would evolve smoothly. In undersampled data one would expect to see irregular, random looking fluctuations, which are apparent in some of the Galaxy 11 lightcurves. Higher resolution data would reveal the true evolution of the satellite's lightcurves, including the small glints from minor satellite components. Smoothly varying colour plots would then allow slight colour changes to be attributed to faint glints, the evolution of which could be monitored over several nights. It is possible that such information could contribute to satellite characterisation, including material identification.

Plots of Anik F1R also demonstrate the utility of high temporal resolution data. Figure 6.14 depicts the Anik F1R I band lightcurve and the B-I colour plot.
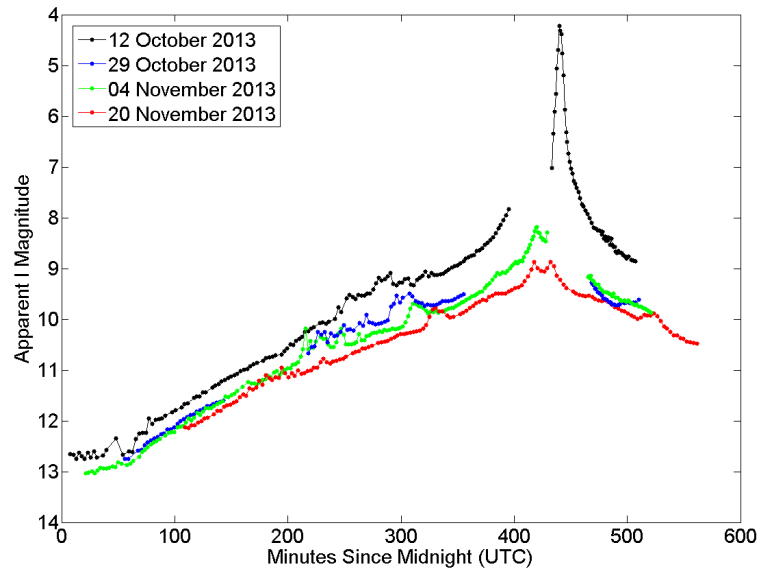
(a) Anik F1R I Band Lightcurve.
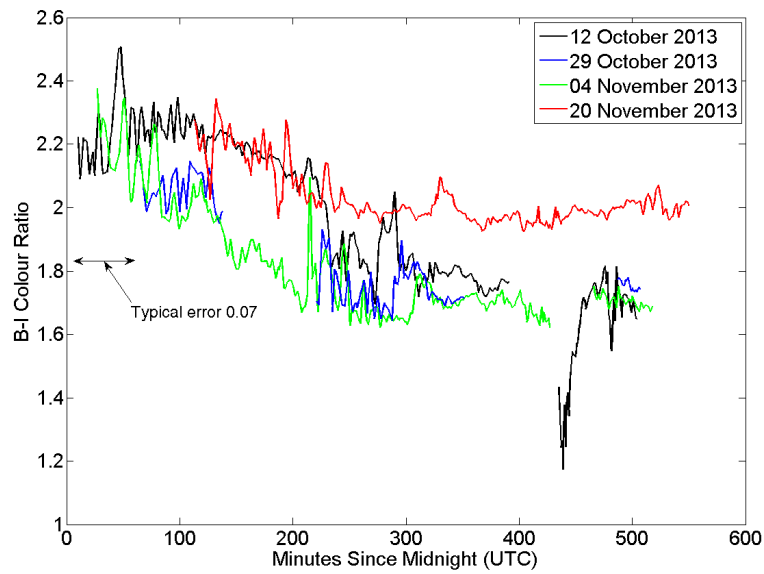
(b) Anik F1R B-I Colour Variation.

Figure 6.14: I Band Lightcurve for Anik F1R (a), compared with its B-I colour variation (b).

In figure 6.14a the solar panel specular glints centred at 430 minutes since midnight are present at the same time on each of the different nights. A broad glint feature, secondary in prominence to the solar panel glint and lasting approximately 30 minutes, is also visible near 300 minutes since midnight; however the feature does not occur at the same time each night. Instead, the glint is centred at progressively later times on later nights. Hall observed the same phenomenon in 2010 when observing the GEO satellite, Galaxy 15 [42]. Hall discovered that a faint glint observed on eight nights over a two week period occurred slightly later on each subsequent night. The glints could only be made to align on each night when they were plotted as a function of the satellite body frame phase angle bisector, indicating that the source of the glint was located on the satellite body, rather than the solar array. Such an analysis could be performed on the Anik F1R glint just described. There are, however, other more fleeting glints that appear on the Anik F1R lightcurve between 200 and 300 minutes since midnight. These glints lack clearly identifiable shapes because their duration is comparable to the data sampling rate, and it is not clear whether they are repeated on subsequent nights. A higher temporal resolution might enable these glints to be compared on different nights, including their effect on the observed colour ratios. The broad glint feature visible near 300 minutes since midnight corresponds to an increase in the intensity of longer wavelength light, which is clearly identifiable on each night in figure 6.14b. Correlating other short duration glints with their effect on colour ratios, and whether they are likely to originate from the same satellite components on different nights, is much more difficult.

Another observation regarding the Anik F1R data, which is noticable also in some of the other satellites' lightcurves, is that the B band lightcurves have a tendency to appear smoother than the I band lightcurves. In general it appears as though shorter wavelength band lightcurves are smoother than longer wavelength band lightcurves. It is proposed that this is a result of the fact that throughout this experiment images taken through shorter wavelength filters always had longer exposure times than images taken through longer wavelength filters. Figure 6.15 depicts Anik F1R's B band lightcurves (top), which have a smoother appearance than Anik F1R's I band

lightcurves (bottom).

Low SNR is not a decisive factor that could be responsible for the noisier looking plots, so it is thought that the duration of the B band images, which were 30 to 60 seconds in most cases, smoothed out short duration fluctuations that were captured by the other filters. For comparison, R and I band filter images typically had exposure times of between 0.02 seconds and five seconds. What is not apparent from the data is whether the short duration fluctuations are due to random noise or genuine variation in the signal.

Fortunately, Anik F1's lightcurve and colour ratio plots contain some features that can help to identify the extent to which variability in the data is due to signal variations, and to what degree it is the result of random noise. In figure 6.16 it can be seen that Anik F1's I band lightcurve generally appears to have a slightly more random appearance than the B band lightcurve, especially at the specular peak.
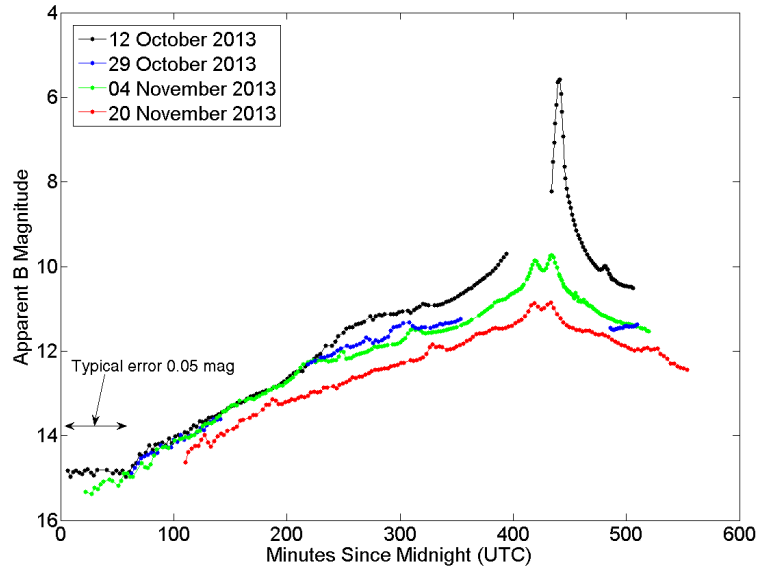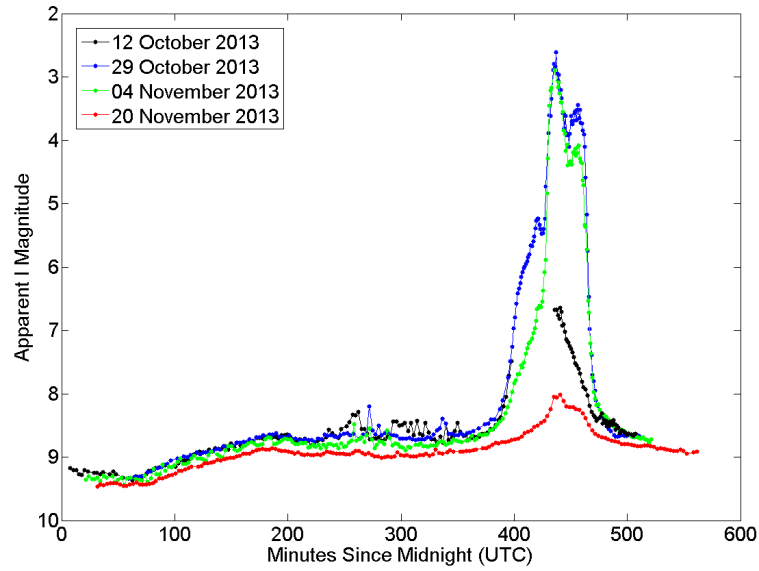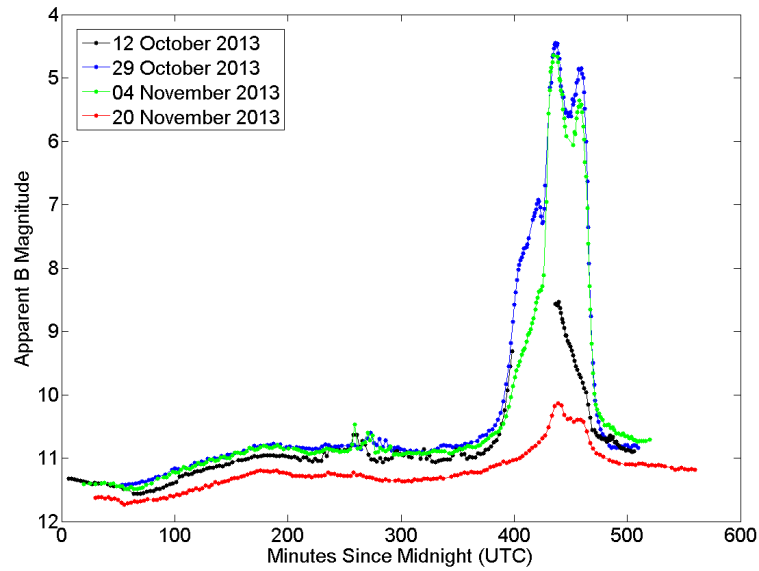
(a) Anik F1R B Band Lightcurve.
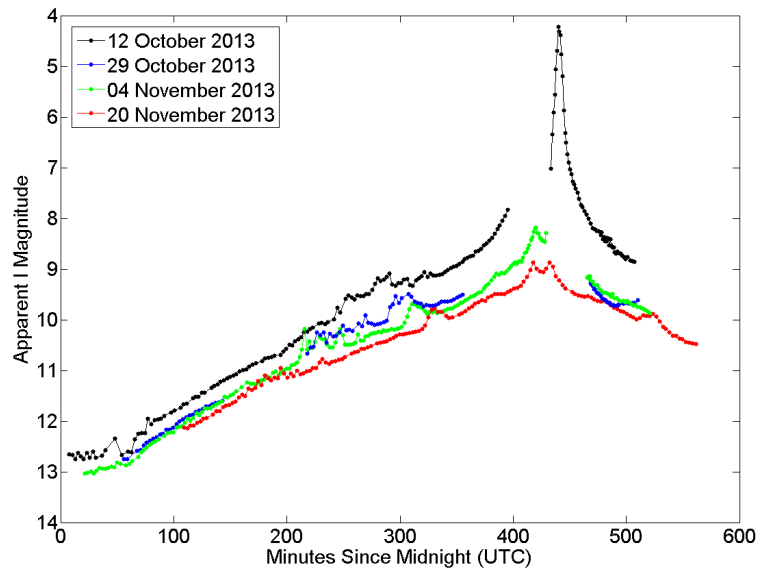
(b) Anik F1R I Band Lightcurve.

Figure 6.15: Anik F1R's B band lightcurve (a) has a smoother appearance than its I band lightcurve (b).

(a) Anik F1 I Band Lightcurve.

(b) Anik F1 B Band Lightcurve.

Figure 6.16: Anik F1 I Band Lightcurve (a), and B Band Lightcurve (b).

When the region around the specular peak is examined closely it can be seen that the shape of the 29 October and 04 November lightcurves is very similar, including the small, seemingly random, fluctuations. The similarities are noticeable in all of the four colour bands, as shown in figure 6.17. Due to the brightness of the satellite during the specular peaks, the exposure times through all of the filters were much shorter than outside of the specular regime. The longest exposure times were through the B filter, between 0.5 and 2 seconds; and the shortest exposure times were through the I filter, between 0.03 and 0.3 seconds. As a result, any potential smoothing effect caused by the averaging of rapid variations over the duration of an exposure is significantly reduced.

(a) I Band

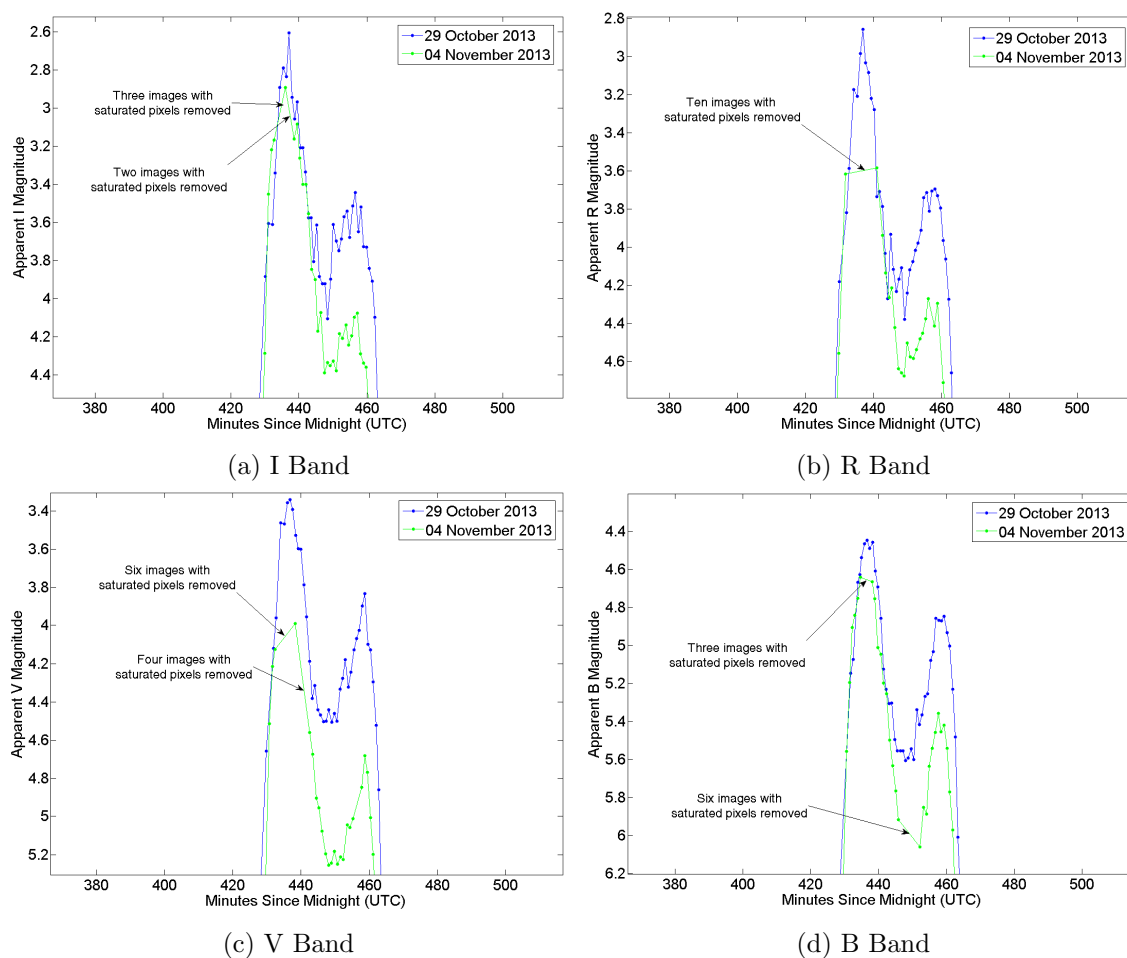(b) R Band

(c) V Band

(d) B Band

Figure 6.17: Enlarged Views of Anik F1 Lightcurve Peaks by Filter.

The unfiltered images of Anik F1 contained saturated pixels throughout the peaks of the brightest specular reflections, and therefore unfortunately cannot be compared. In addition, some images that were taken through the colour filters contained saturated pixels. Markers on each of the images indicate where these individual data points were removed from the data sets. The fact that similarities are apparent in all four sets of data indicate that it is unlikely to be a coincidental result: the random looking fluctuations are likely to be caused by actual variations in the signal.

In addition, a feature was identified in the Anik F1 lightcurves, prior to the specular peak, that is not apparent in any of the other satellites' lightcurves. An oscillation in the lightcurves, that is visible in figure 6.16 was noticed that varies in amplitude between the different nights, and between the colour filters. Figure 6.18 presents details of portions of Anik F1 lightcurves.
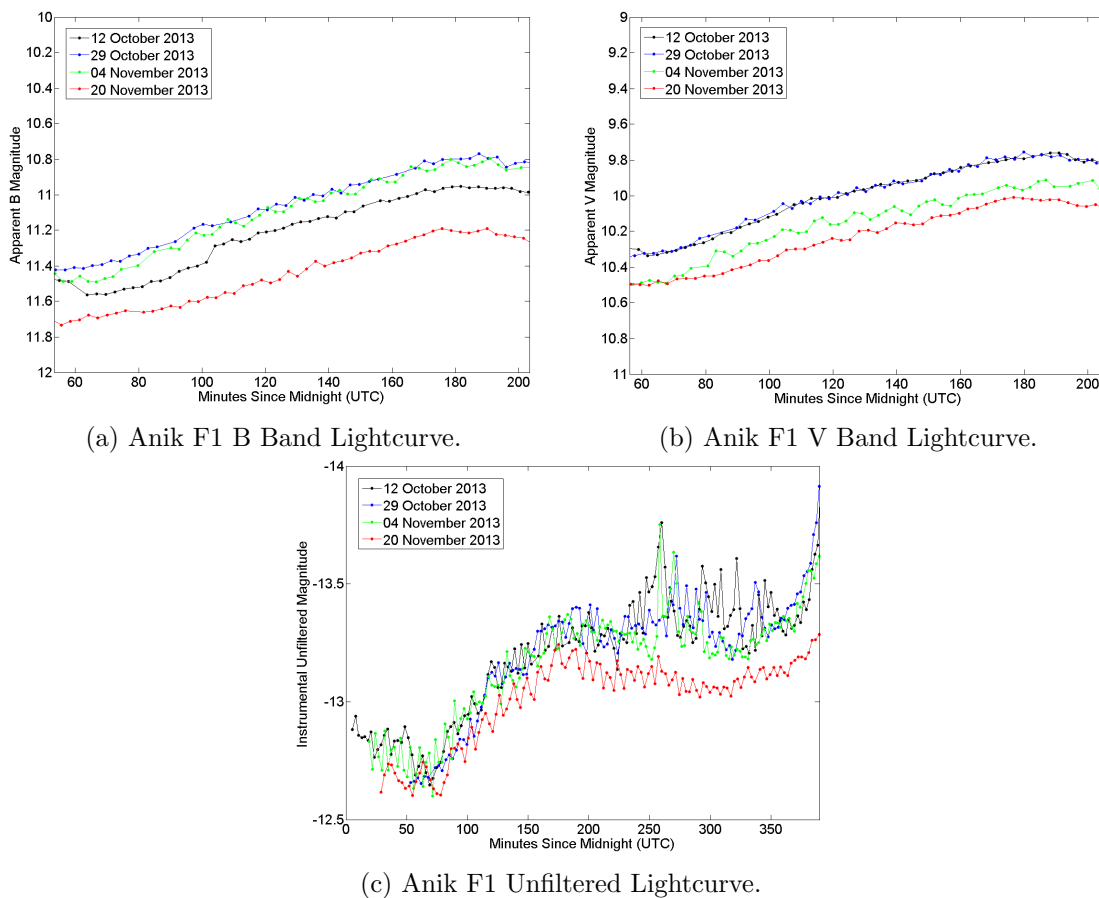

(a) Anik F1 B Band Lightcurve.

(b) Anik F1 V Band Lightcurve.

(c) Anik F1 Unfiltered Lightcurve.

Figure 6.18: Anik F1 B and I Band and Unfiltered Lightcurves.

The oscillation is most noticeable in the 4th of November data, however it is possible that this is simply a coincidence resulting from the timing of the data. It can be seen that on the 4th of November there are approximately four to five data points per oscillation, which is very close to the minimum sampling rate required to unambiguously determine the true oscillation frequency. On the other nights, the 29th of October in particular, there appears to be only two data points per oscillation: one at the peak and one at the trough. The data rate in those cases is too low to clearly ascertain the true nature of the signal variability. Frequency aliasing is possible, whereby the true oscillation frequency is actually a multiple of the apparent frequency; the amplitude of the oscillations is also ambiguous. For example, the data points on the 29th of October lightcurve in figure 6.18b may reside at the peaks and troughs of the oscillations, or they may reside close to the middle, reducing the apparent amplitude. Of note, the 20th of November unfiltered lightcurve, figure 6.18c, displays a clear oscillation with an amplitude of approximately 0.1 magnitude, whereas through all of the filters the oscillation is barely detectable. In figures 6.18a and 6.18b the apparent frequency of the oscillations is also visibly lower than in the unfiltered data. The oscillation period, according to these data, is approximately 10 minutes.

The relatively undersampled data sets resulted in the oscillations predominantly being undetectable in the colour ratio curves. The B-V curve in figure 6.19, however, displays a clear oscillation; particularly on the 4th of November and to a lesser extent on the 29th of October, between 100 and 200 minutes since midnight.

The precise cause of the oscillation is unknown (a rotating satellite component, flapping solar panels, or interference effects caused by reflections off the parabolic antennas are possible explanations), however an important conclusion can be made. The very smooth nature of the oscillation in the B-V colour ratio plot on the 4th of November (figure 6.19) indicates that there is very little error in the individual measurements on the plots. The probability of random errors arranging themselves in such a way as to create a smooth, regular oscillation is exceptionally low. Therefore, it can reasonably be assumed that the true noise level is indicated subjectively by
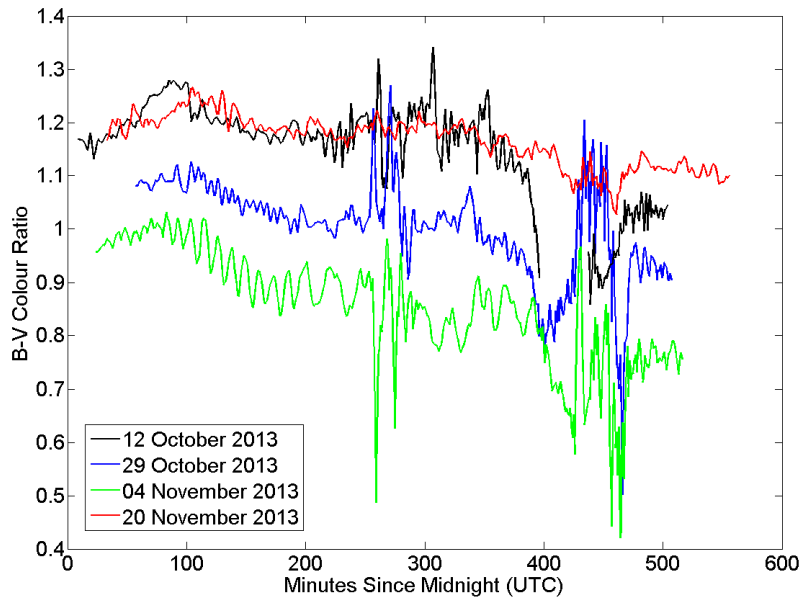
Figure 6.19: Anik F1 B-V Colour Variation.

the 'roughness' of the oscillation: and deviations from a smooth curve are barely perceptible on the 4[th] of November between 100 and 200 minutes since midnight. Given that there are no substantial differences in SNR between different parts of the plots, it can also reasonably be assumed that the noisy appearance of the lightcurves is actually indicative of the true variability of the signal. However, given that images were taken sequentially through the different filters, rather than simultaneously, the colour ratio plots are not likely to reproduce the rapid small amplitude variations as accurately as is done in the lightcurves.

Finally, through good fortune, a further test of the accuracy of data from this experiment is available. By coincidence a fellow postgraduate researcher at the RMCC SSRAL, Mr Mike Earl, who has extensive experience observing satellites, also happened to observe Anik F1 on the 4[th] of November using a small aperture telescope. Earl observed Anik F1 for just under four hours capturing high rate, unfiltered imagery at a frequency of approximately one image every three seconds. Figure 6.20a is a plot of Earl's data, with the author's lightcurve overlaid, and figure 6.20b is a short, enlarged segment of the lightcurve.
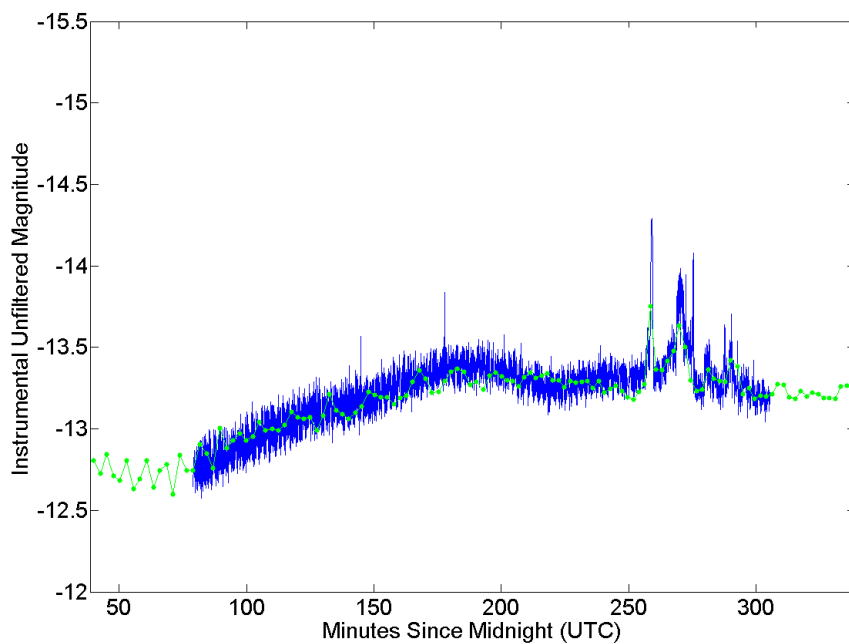
(a) Unfiltered lightcurve.
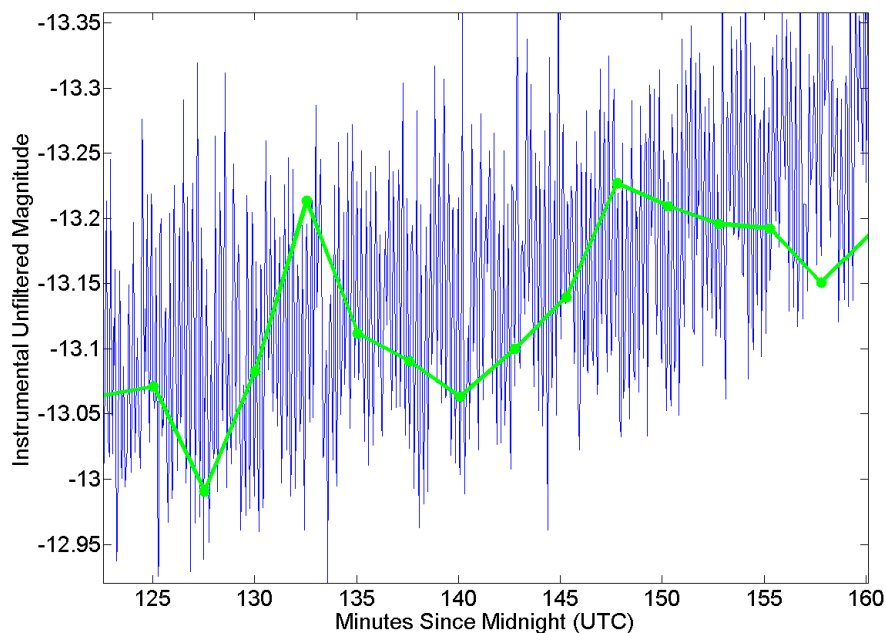
(b) Short segment of unfiltered lightcurve.

Figure 6.20: High temporal resolution unfiltered lightcurves of Anik F1. Earl's entire observation period is depicted in figure 6.20a, with the author's lightcurve overlaid, and an enlarged, short segment clearly showing the high frequency oscillations is depicted in figure 6.20b. Data courtesy of Mr Mike Earl.

It is immediately apparent that oscillations are present in Earl's data, but at a much higher frequency than what was observed by the author. A Fourier transform was performed on Earl's data, which revealed the presence of two significant oscillations that have periods of approximately two minutes, and 13 seconds respectively. At a data acquisition rate of one measurement every three seconds, even Earl's data just barely has a high enough temporal resolution to detect the high frequency oscillation. A Fourier transform of the author's data, however, revealed a period of approximately 10 minutes. Although the lack of evidence in Earl's data of a ten minute period oscillation seems to contradict the author's data, it can possibly be explained by either frequency aliasing or the different exposure times that were used. Earl employed 0.5 second exposure times, whereas the author used exposure times ranging from at least 1.5 seconds for the unfiltered measurements to 60 seconds for the B filter measurements. Longer exposure times have the effect of integrating over larger portions of the high frequency oscillation, resulting in magnitudes that are an average over portions of an oscillation or several oscillations.

In light of Earl's data, another conclusion that can be made is that the temporal resolution of the data in this experiment is too low to be able to accurately reconstruct the satellite's lightcurves and colour variations, including small features with amplitudes less than approximately 0.1 magnitude. As discussed, the frequency and amplitude of the lightcurve and colour ratio oscillations is ambiguous; it is very difficult to distinguish noise from true signal variations; and it is difficult to identify when short duration glints that appear on different nights come from the same satellite component. In addition, Earl's data show that there are interesting oscillations in the lightcurves that have periods too short to be detected using this experimental method.

This experiment demonstrates that satellite lightcurves are mostly not at all smoothly varying on a small scale. What looks like noise in many lightcurves may in fact be rapidly changing, small scale signal fluctuations. The precise aim of a particular experiment will determine whether it is critical to accurately measure such fluctuations, and how important it is to minimise image exposure times to prevent

smoothing. In addition, oscillations such as those measured from Anik F1 may potentially be present in other satellites. Determining which satellites exhibit such oscillations, at what frequencies, and whether they are ongoing or temporary may help to uncover their cause. Once the cause and long term nature of the oscillations is known, their precise frequency or amplitude may be very useful information for characterising or differentiating between satellites. If small scale fluctuations do prove to be of more interest than has been shown in the past, then a very high data sampling rate would be required to capture the necessary information.

A small aperture telescope such as the one used in this experiment is capable of measuring satellite lightcurves and colour ratios to a very high accuracy, however the required exposure times are longer than what would be necessary using a large aperture telescope, limiting the maximum obtainable temporal resolution.

Chapter 7

Conclusion

This thesis focussed on the utility of small aperture telescopes for the photometric analysis of GEO satellites. Observations of six GEO satellites in two clusters were taken both unfiltered and through Bessel B, V, R, and I filters on twelve separate nights. The first GEO cluster, consisting of two satellites, was observed on eight different nights, and the second GEO cluster, consisting of four satellites, was observed on four different nights. In order to determine the level of systematic error present in the data three separate tests were performed. The first test consisted of comparing the measured magnitude of incidentally imaged stars that were present in the background of various images to their known magnitudes. This test gives an indication of the precision of the magnitude measurements, as well as the accuracy of the zero point calculations. The second test involved using observations of three different Landolt star fields at different times throughout one night to produce three independent zero point calculations for each filter. This test gives an indication of the accuracy of the zero point calculations and the stability of the zero point throughout the night. The third test involved visually comparing the colour plots of satellites from the same cluster. Similar features occurring at the same time for different satellites could be indicative of systematic errors in the data.

This chapter summarises the main conclusions of the thesis, both in regard to the utility of small aperture telescopes for photometric analysis of artificial satellites, and the contribution of this work to understanding how photometry can be used to characterise artificial satellites. Recommendations for future work that could build upon these conclusions are also presented.

7.1 Summary of Conclusions

7.1.1 Utility of Small Aperture Telescopes for Photometric Analysis of Artificial Satellites

Much effort has gone into conducting photometric analysis of artificial Earth orbiting satellites, particularly GEO satellites, for the purposes of enhancing SSA. Often, metre class telescopes have been used for this purpose despite the lack of a thorough understanding of the precision and temporal resolution required to be able to obtain information useful for characterising satellites. The relative portability, affordability, and abundance of small aperture telescopes are significant advantages over metre class telescopes that could be realised should they be found to be suitable.

This experiment has demonstrated that an equipment setup that includes a small aperture telescope can obtain relatively accurate multicolour broadband photometric measurements, with uncertainties on the order of 0.1 magnitude or less, at a temporal cadence of between one and three minutes per measurement, despite the experiment being conducted in a location with a relatively high degree of light pollution. Two important observations were made during the experiment that are relevant to the choice of equipment that a researcher might use. First, although lightcurves often have a generally smooth appearance overall, it is common for low amplitude fluctuations to be present at various times. On two separate nights seemingly random, erratic fluctuations with amplitudes of approximately 0.1 to 0.2 magnitude were observed near the peak of Anik F1's lightcurves. The fact that there were many similarities between the fluctuations observed on both nights indicates that they may well be the result of real signal variations rather than noise. In addition, periodic oscillations with amplitudes of approximately 0.1 magnitude were discovered in the lightcurves of at least one satellite on multiple nights. Such unexplained features may turn out to be of critical importance in characterising satellites. If so, small aperture telescopes, which are capable of detecting such small amplitude fluctuations, would be well suited to the task due to their cost effectiveness and flexibility. Converseley, large aperture telescopes are more expensive to operate and often require their usage to be booked well in advance.

The second observation that was made is that despite the outwardly smooth appearance of many lightcurves, signal fluctuations frequently occur on very short timescales. The erratic appearance of various parts of many lightcurves produced during this experiment indicates that the temporal resolution of the data was too low to enable the detection of finer features. In those cases the sampling rate was usually between one and three minutes per data point. A clear example was made when a colleague at the SSRAL, Mr Mike Earl, observed Anik F1 on the same night as the author, employing a much higher sampling rate. Whilst the author detected oscillations in all of Anik F1's lightcurves on that night, with a period of 10 minutes per oscillation, Earl detected two further underlying oscillations with periods of 2 minutes and 13 seconds, respectively, but did not detect the 10 minute period oscillation. Small aperture telescopes are more limited than large aperture telescopes in the maximum data rate they are capable of achieving because they require longer exposure times to capture an equivalent amount of signal. In addition, longer exposure times may have the effect of smoothing out portions of a lightcurve that vary on timescales comparable to the exposure time. Large aperture telescopes are therefore, in principle, more effective than small aperture telescopes at detecting high frequency signal fluctuations. Because of their affordability, however, it would be possible to use multiple small aperture telescopes to simultaneously acquire data through only one filter each, thus greatly increase the data rate and mitigating their weaknesses.

Ultimately, whether or not a small aperture telescope can be used effectively for enhancing SSA depends upon the precise aim of the researcher. Large aperture telescopes are generally capable of defining finer features of a lightcurve than small aperture telescopes because they can capture an equivalent signal with shorter exposure times. However, this experiment has shown that highly accurate measurements can be made using small aperture telescopes, which have the potential to contribute significantly to SSA efforts.

7.1.2   Characterisation of Active Geostationary Satellites

This experiment has demonstrated that GEO satellite lightcurves and colour variations do not necessarily behave in a straightforwardly predictable manner. Spikes in brightness due to specular reflections off solar arrays or parts of the satellite body vary significantly in magnitude throughout the year. The precise sun-satellite-sensor geometry seems to be a critical factor in determining the flux that will be measured by the sensor at a given point in time. As a result, it is often difficult to uniquely classify a satellite by the general shape of its lightcurve. The illumination and observation geometry also affects the colour ratios that will be measured. Far from exhibiting constant colour ratios at all times, as has been assumed in the past [35], each satellite's colours vary considerably throughout a given night. The shape of the resultant colour plot also changes from night to night, making it difficult to use colour for satellite discrimination.

Because of the number of variables that affect measurements of flux and colour, and the degree to which measurements are affected by small changes in those variables, one must be very careful when making simplifying assumptions about a satellite. Studies that model satellites as simple structures consisting of only two or three basic components cannot capture the effects of small glints off minor satellite surfaces. Similarly, models that do not account for the effect of changing illumination and observation geometry on the SED and intensity of reflected light off each satellite surface cannot produce accurate results.

7.2   Recommendations and Future Work

The research conducted during the course of this study has successfully demonstrated the capacity of small aperture telescopes to measure the apparent brightness of a satellite to within 0.1 magnitude. The main limiting factor in reconstructing the finer details of a satellite's lightcurves is thus the maximum temporal resolution that can be obtained. Increasing the temporal resolution without making use of a large aperture telescope will require multiple telescopes to be used simultaneously: one for

each filter. Having a dedicated telescope for each colour filter will eliminate time wasted due to changing filters and waiting for images to be taken and downloaded through the other filters.

Another weakness that this experiment has identified in previous studies of photometric lightcurves is a lack of suitable data spanning a period of at least one entire year. General observations have been made, such as that specular glint features vary in amplitude over time, however to the best of the author's knowledge no satellites have been observed using a high sampling rate more than about five times over the course of a year. It has thus been impossible to investigate phenomena that occur over relatively short time frames, such the evolution of colour variations near specular peaks.

A key recommendation for future work, therefore, is to use multiple telescopes to simultaneously collect data on a single GEO cluster on as many nights as possible throughout the course of at least a year. Ideally, four telescopes would be used such that data could be acquired in each of the B, V, R, and I bands. However, if only two or three telescopes were available, using the R, I, and V filters, rather than the B filter, would allow a higher temporal resolution to be achieved because the satellites always appear fainter in the B band, thus requiring longer exposure times to achieve a suitable SNR.

References

[1] NASA. Orbital debris quarterly news. Technical report, National Aeronautics and Space Administration, 2013.

[2] D. Kessler, N. Johnson, J. Liou, and M. Matney. The kessler syndrome: Implications to future space operations. In *33rd Annual Guidance and Control Conference*, 2010.

[3] D. Hall, B. Calef, K. Knox, M. Bolden, and P. Kervin. Separating attitude and shape effects for non-resolved objects. In *Proceedings of the 2007 AMOS Technical Conference*, Kihei, Maui, HI, 2007.

[4] D. Hall. Surface material characterisation from multi-band optical observations. In *Proceedings of the 2010 AMOS Technical Conference*, Kihei, Maui, HI, 2010.

[5] K. Jorgensen. *Using Reflectance Spectroscopy to Determine Material Type of Orbital Debris.* PhD thesis, Colorado Centre for Astrodynamics Research, University of Colorado, 2000.

[6] P. LeVan. Discrimination of closely space geosynchronous satellites - phase curve analysis & new small business innovative research (sbir) efforts. In *Proceedings of the 2010 AMOS Technical Conference*, Kihei, Maui, HI, 2010.

[7] T. E. Payne, S. A. Gregory, J. Tombasco, K. Luu, and L. Durr. Satellite monitoring, change detection and characterization using non-resolved, electro-optical data from a small aperture telescope. In *Proceedings of the 2007 AMOS Technical Conference*, Kihei, Maui, HI, 2007.

[8] D. Hall. Amos galaxy 15 satellite observations and analysis. In *Proceedings of the 2011 AMOS Technical Conference*, Kihei, Maui, HI, 2011.

[9] P. Kervin. Anomaly resolution using optical signatures. In *Proceedings of the 2011 AMOS Technical Conference*, Kihei, Maui, HI, 2011.

[10] SBIG. Research grade ubvri filters, March 2002.

[11] F. Nicodemus, J. Richmond, I. Ginsberg, and T. Limperis. Geometrical considerations and nomenclature for reflectance. Technical report, National Bureau of Standards, 1977.

[12] P. Kervin, D. Hall, M. Bolden, and J. Toth. Phase angle: What is it good for? In *Proceedings of the 2010 AMOS Technical Conference*, Kihei, Maui, HI, 2010.

[13] F. Chromey. *To Measure the Sky.* Cambridge University Press, 2010.

[14] T. E. Payne, S. A. Gregory, D. Payne, L. Kann, D. Sanchez, D. Werling, C. Davis, and L. Finkner. Determination of photometric filters for satellite observations using silc data. In S E Andrews, editor, *Proceedings of the 2000 Space Control Conference*, pages 125–134. Massachusetts Institute of Technology Lincoln Laboratory, 2000.

[15] T. E. W. Payne, S. A. Gregory, D. J. Sanchez, T. W. Burdillis, and S. A. Storm. Color photometry of geosynchronous satellites using the silc filters. In *Proc. SPIE 4490*, volume 4490, pages 194–199, San Diego, CA, 2001.

[16] A. Chaudhary, N. Nosek, J. Africano, and D. Hall. A technique for space aging incorporation in satellite photometry models. In *Proceedings of the 2005 AMOS Technical Conference*, Kihei, Maui, HI, 2005.

[17] S. A. Gregory, T. E. Payne, and K. Luu. Comparisons between simulated and observed color photometric signatures of geosynchronous satellites. In *Proceedings of the 2005 AMOS Technical Conference*, Kihei, Maui, HI, 2005.

[18] T. E. Payne, S. A. Gregory, and K. Luu. Ssa analysis of geos photometric signature classifications and solar panel offsets. In *Proceedings of the 2006 AMOS Technical Conference*, Kihei, Maui, HI, 2006.

[19] R. Deiotte, M. Guyote, T. Kelecy, D. Hall, J. Africano, and P. Kervin. Maui space surveillance system (msss) satellite categorization laboratory. In *Proceedings of the 2006 AMOS Technical Conference*, Kihei, Maui, HI, 2006.

[20] H. Cowardin. *Characterization of orbital debris objects over optical wavelengths via laboratory measurements*. PhD thesis, Department of Natural Sciences and Mathematics, University of Houston, 2010.

[21] T. E. Payne, S. Gregory, and N. Houtkooper. Long-term analysis of geo photometric signatures. In *Proceedings of the 2003 AMOS Technical Conference*, Kihei, Maui, HI, 2003.

[22] J. Trimble. Space object identification in living color (silc). Technical report, United States Air Force Space Battlelab, Schriever AFB, CO, 2000.

[23] T.E. Payne, S.A. Gregory, N.M. Houtkooper, and T.D. Burdullis. Analysis of multispectral radiometric signatures from geosynchronous satellites. In *Proc. SPIE 4847*, volume 4847, pages 332–336, Albuquerque, NM, 2002.

[24] F. J. Vrba, H. F. Fliegel, and L. F. Warner. Optical brightness measurements of gps block ii, iia, and iir satellites on orbit. In *Proceedings of the 2003 AMOS Technical Conference*, Kihei, Maui, HI, 2003.

[25] F. J. Vrba, H. F. Fliegel, and L. F. Warner. The discrimination of gps block ii, iia, and iir satellites: Spectral energy distributions between 450-2220 nm. In *Proceedings of the 2004 AMOS Technical Conference*, Kihei, Maui, HI, 2004.

[26] H. F. Fliegel, L. F. Warner, and F. J. Vrba. Photometry of global positioning system block ii and iia satellites on orbit. *Journal of Spacecraft and Rockets*, 38(4):609–616, 2001.

[27] R. Scott and B. Wallace. Small-aperture optical photometry of canadian geostationary satellites. *Canadian Aeronautics and Space Journal*, 55(2):41–53, 2009.

[28] D. S. McDougal, R. B. III Lee, and D. C. Romick. Measured physical and optical properties of the passive geodetic satellite (pageos) and echo 1. In *Proceedings of the Third International Symposium on the Use of Artificial Satellites for Geodesy*, pages 253–259, Washington D.C., 1972. American Geophysical Union.

[29] M. D. Hejduk. Catalogue-wide satellite photometric behavior paradigms. In *Proceedings of the 2010 AMOS Technical Conference*, Kihei, Maui, HI, 2010.

[30] T. E. Payne, S. A. Gregory, F. J. Vrba, and K. Luu. Utility of a multi-color photometric database. In *Proceedings of the 2005 AMOS Technical Conference*, Kihei, Maui, HI, 2005.

[31] T. E. Payne, A. Chaudhary, S. Gregory, J. Brown, and M. Nosek. Signature intensity derivative and its application to resident space object typing. In *Proceedings of the 2009 AMOS Technical Conference*, Kihei, Maui, HI, 2009.

[32] K. Abercromby, J. Okada, M. Guyote, K. Hamada, and E. Barker. Comparisons of ground truth and remote spectral measurements of the formosat and ande spacecraft. In *Proceedings of the 2006 AMOS Technical Conference*, 2006.

[33] K. Abercromby, K. Hamada, M. Guyote, J. Okada, and E. Barker. Remote and ground truth spectral measurement comparisons of formosat iii. In *Proceedings of the 2007 AMOS Technical Conference*, 2007.

[34] H. Cowardin, K. Abercromby, E. Barker, P. Seitzer, M. Mulrooney, and T. Schildknecht. An assessment of geo orbital debris photometric properties derived from laboratory-based measurements. In *Proceedings of the 2009 AMOS Technical Conference*, Kihei, Maui, HI, 2009.

[35] S.M. Lederer, P. Seitzer, H.M. Cowardin, E.S. Barker, K.J. Abercromby, and A. Burkhardt. Preliminary characterisation of idcsp spacecrafts through a multi-analytical approach. In *Proceedings of the 2012 AMOS Technical Conference*, Kihei, Maui, HI, 2012.

[36] D. Bedard. Using a physics-based reflectance model to study the reddening effect observed in spectrometric measumeasure of artificial space objects. In *Proceedings of the 2011 AMOS Technical Conference*, Kihei, Maui, HI, 2011.

[37] M. Guyote, J. Okada, and K. Abercromby. Using space weathering models to match observed spectra to predicted spectra. In *Proceedings of the 2006 AMOS Technical Conference*, 2006.

[38] D. Hall, K. Hamada, T. Kelecy, and P. Kervin. Surface material characterization from non-resolved multi-band optical observations. In *Proceedings of the 2012 AMOS Technical Conference*, Kihei, Maui, HI, 2012.

[39] D. Bedard. *Spectrometric Characterisation of Artificial Earth-Orbiting Objects.* PhD thesis, Royal Military College of Canada, 2013.

[40] A. Chaudhary, T. Payne, S. Gregory, and P. Dao. Fingerprinting of non-resolved three-axis stabilized space objects using a two-facet analytical model. In *Proceedings of the 2011 AMOS Technical Conference*, Kihei, Maui, HI, 2011.

[41] J. Frith, B. Gibson, R. Knox, and K. Kuluhiwa. Simultaneous single site color photometry of leo satellites. In *Proceedings of the 2008 AMOS Technical Conference*, Kihei, Maui, HI, 2008.

[42] D. Hall and P. Kervin. Analysis of faint glints from stabilized geo satellites. In *Proceedings of the 2013 AMOS Technical Conference*, Kihei, Maui, HI, 2013.

[43] E. V. Ryan and W. H. Ryan. Space object characterization studies and the magdalena ridge observatory's 2.4 meter telescope. In *Proceedings of the 2007 AMOS Technical Conference*, Kihei, Maui, HI, 2007.

[44] A. Chaudhary, M. Nosek, and B. Klem. Applications of space-based infrared sensor data collections to spacespace health monitoring, space object identification and sensor calibration. In *Proceedings of the 2002 AMOS Technical Conference*, Kihei, Maui, HI, 2002.

[45] M. S. Bessel. Ubvri passbands. *Publications of the Astronomical Society of the Pacific*, 102:1181–1199, October 1990.

[46] Apogee Instruments Inc. Alta u42 high performance cooled ccd camera system, 2003.

[47] Intelsat. Satellites and coverage maps, February 2014.

[48] Telesat. Telesat - our fleet, February 2014.

[49] G. Krebs. Gunter's space page, February 2014.

[50] S.B. Howell. *Handbook of CCD Astronomy, Second edition.* Cambridge University Press, New York, 2006.

[51] F. J. Vrba, M. E. DiVittorio, R. B. Hindsley, H. R. Schmitt, J. T. Armstrong, P. D. Shankland, D. J. Hutter, and J. A. Benson. A survey of geosynchronous satellite glints. In *Proceedings of the 2009 AMOS Technical Conference*, Kihei, Maui, HI, 2009.

[52] B. D. Warner. *A Practical Guide to Lightcurve Photometry and Analysis.* Springer Science and Business Media Inc., 2006.

[53] D. Bedard, G. Wade, D. Monin, and R. Scott. Spectrometric characterisation of geostationary satellites. In *Proceedings of the 2012 AMOS Technical Conference*, Kihei, Maui, HI, 2012.

[54] J. Murray-Krezan, W. Inbody, P. Dao, A. Dentamaro, D. Fulcoly, and S. Gregory. Algorithms for automated characterization of three-axis stabilized geos using non-resolved optical observations. In *Proceedings of the 2012 AMOS Technical Conference*, Kihei, Maui, HI, 2012.

[55] H. Cowardin, P. Seitzer, K. Abercromby, E. Barker, and T. Schildknecht. Characterisation of orbital debris photometric properties derived from laboratory based measurements. In *Proceedings of the 2010 AMOS Technical Conference*, Kihei, Maui, HI, 2010.

[56] A. Chaudhary, M. Nosek, B. Klem, C. Birkemeier, and T. Kim. Calibration of four color photometry data gathered using a raven-like telescope system. In *Proceedings of the 2004 AMOS Technical Conference*, Kihei, Maui, HI, 2004.

Complete Set of Results

The following are the complete set of lightcurve and colour plots that were produced throughout the conduct of this experiment. Magnitude errors are all smaller than the size of the individual data points except where annotated on the plot, and similarly, colour ratio errors are all smaller than the thickness of the lines except where annotated. Galaxy 11's colour variations are progressively shifted up each subsequent night such that each night's plots are separated for ease of inspection.

A.1   Intelsat 805

(a) B Mag

(b) V Mag

(c) B-V Colour Variations

(d) B-R Colour Variations

(e) B-I Colour Variations

(f) V-R Colour Variations
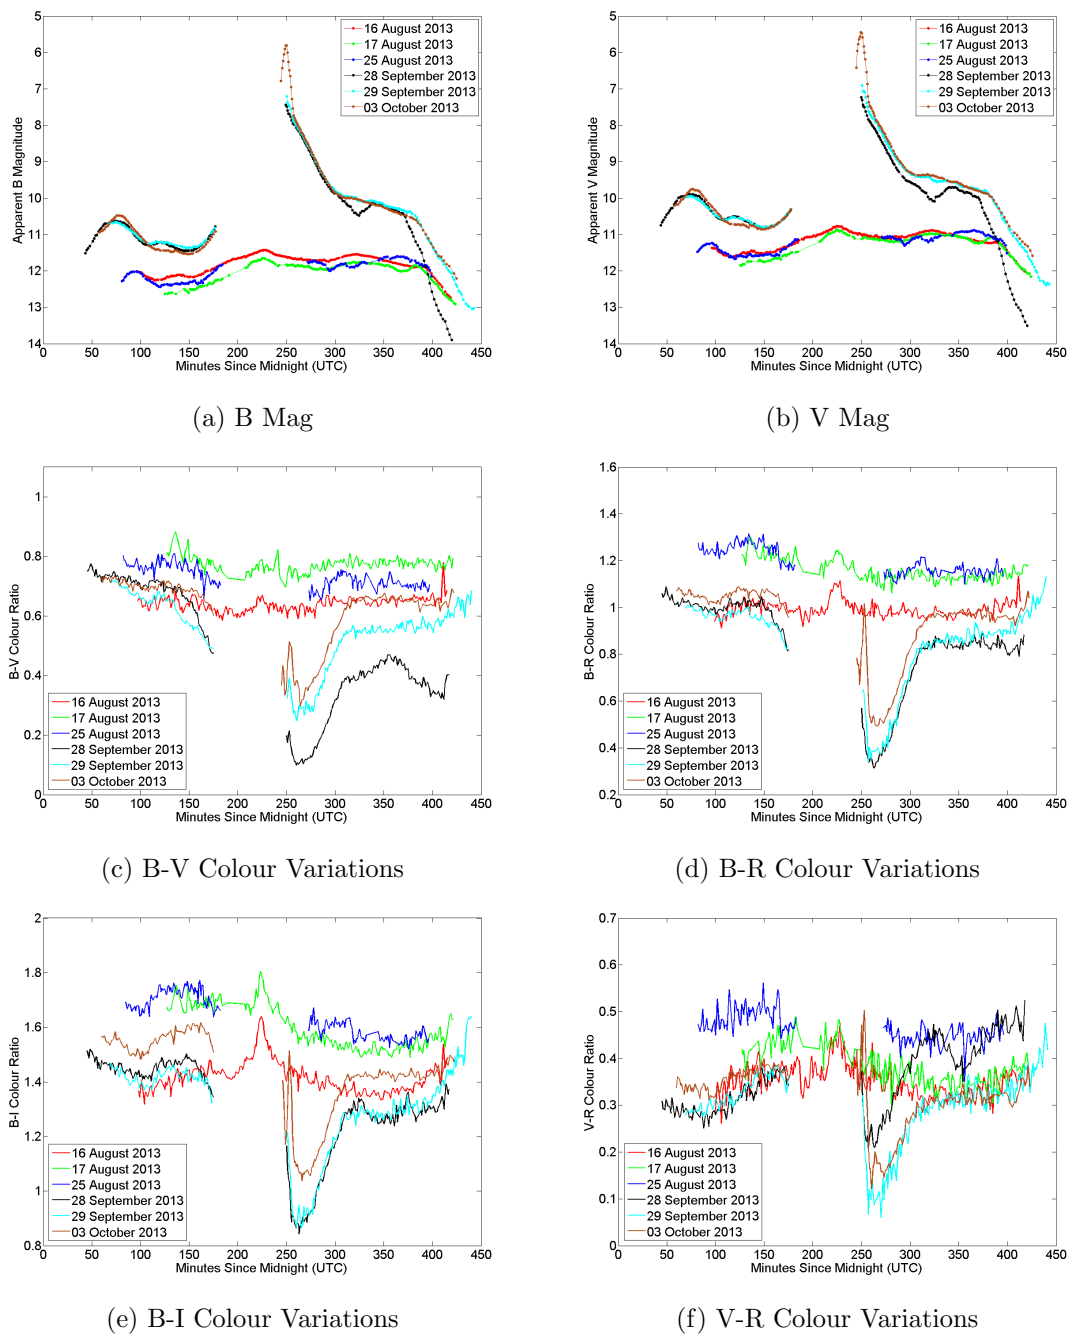
Figure A.1: Intelsat 805 Lightcurves and Colour Variations.

(a) R Mag

(b) I Mag

(c) R-I Colour Variations

(d) V-I Colour Variations

(e) Unfiltered Instrumental Magnitude
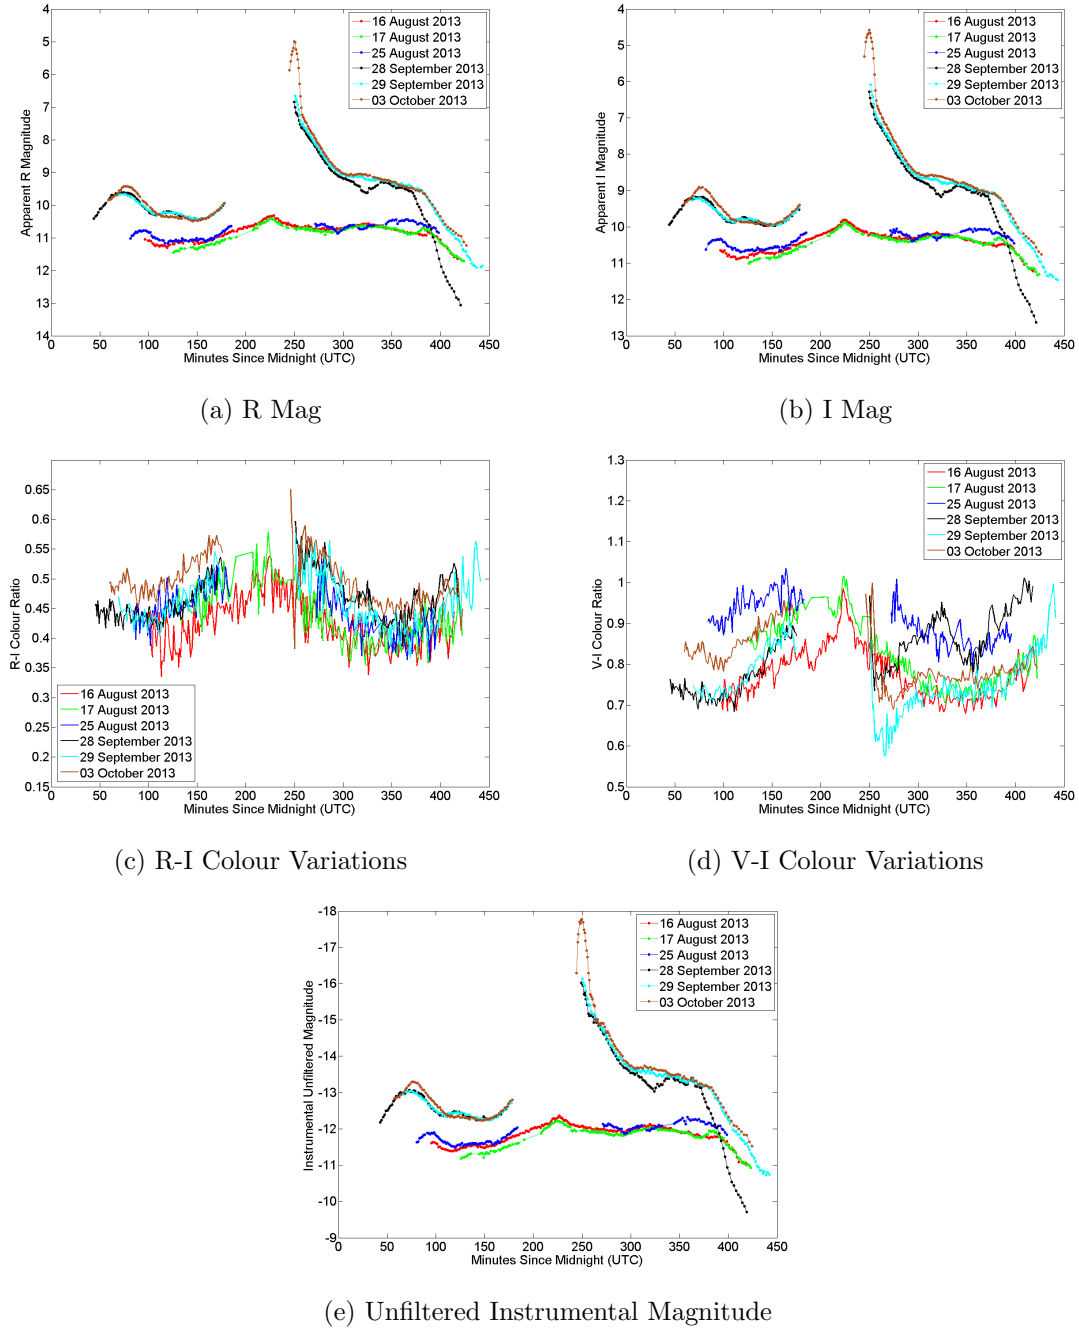
Figure A.2: Intelsat 805 Lightcurves and Colour Variations.

A.2 Galaxy 11

(a) B Mag

(b) V Mag

(c) B-V Colour Variations

(d) B-R Colour Variations

(e) B-I Colour Variations

(f) V-R Colour Variations
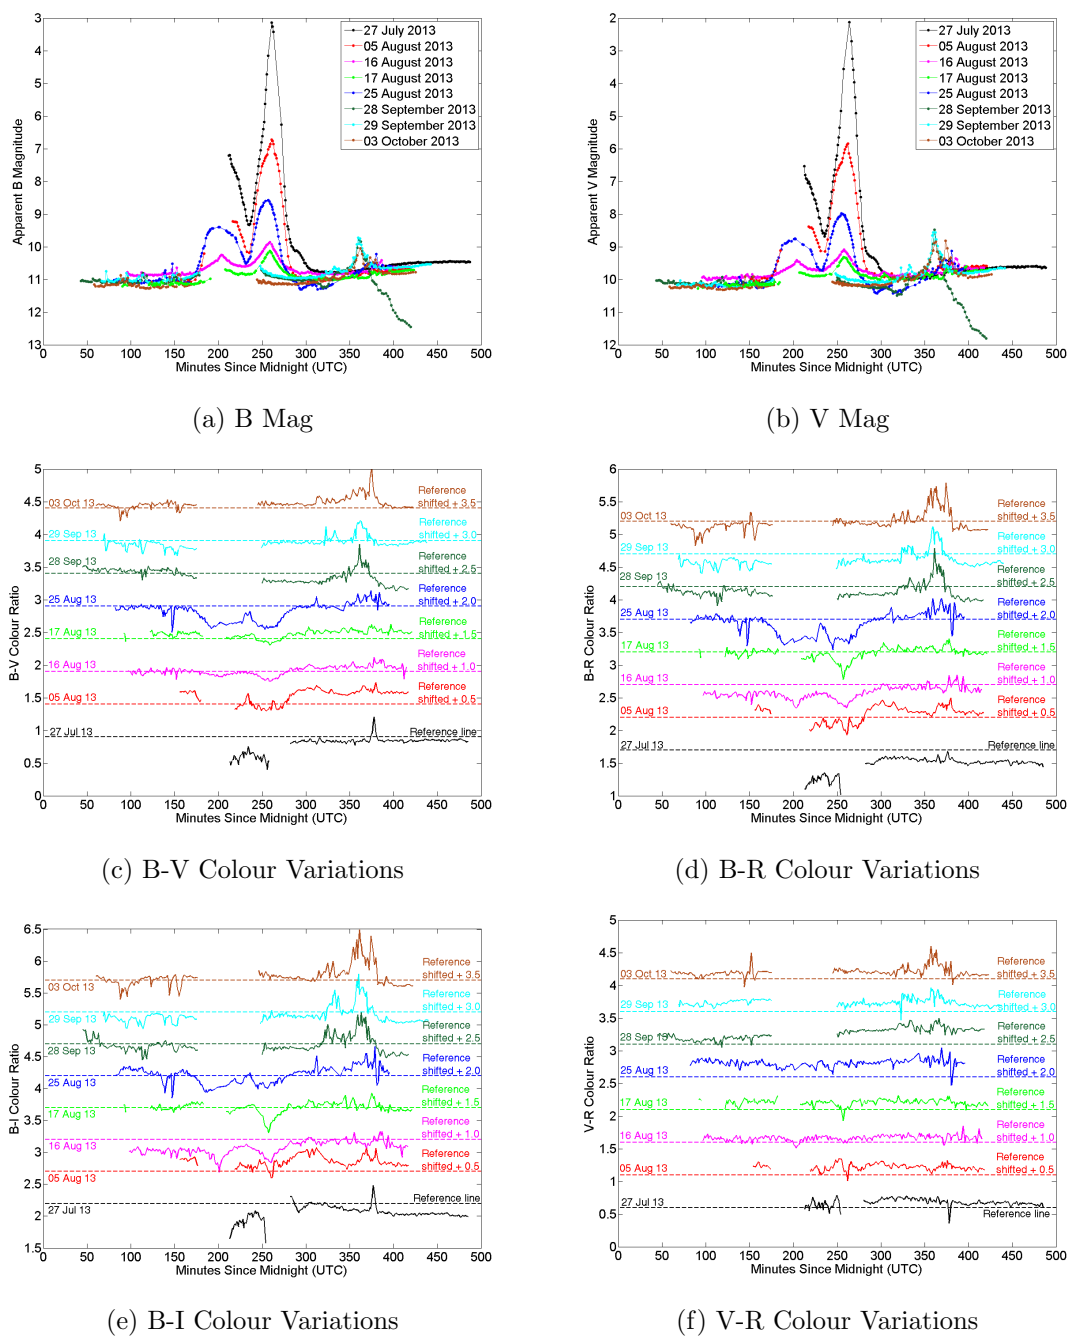
Figure A.3: Galaxy 11 Lightcurves and Colour Variations.
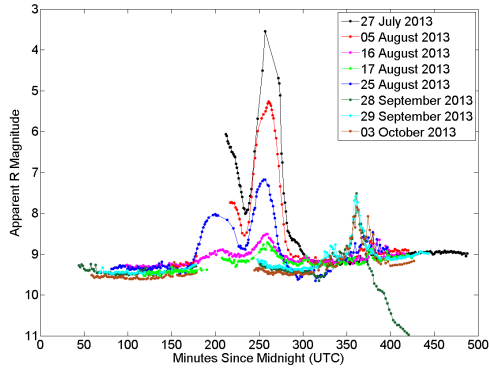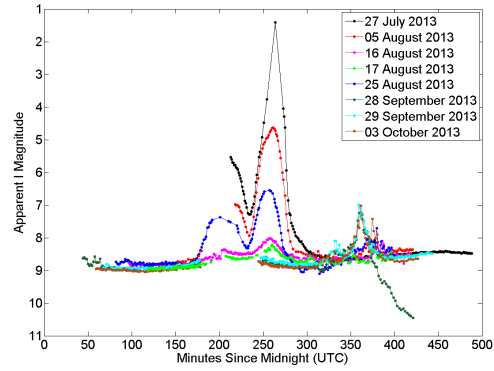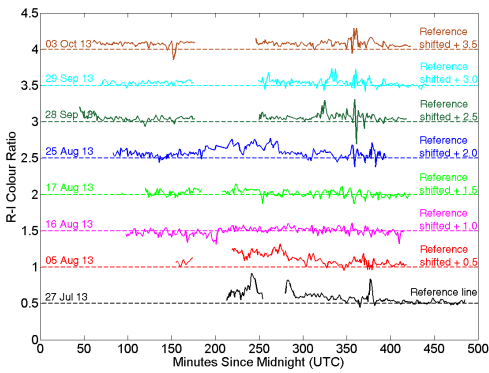
(a) R Mag

(b) I Mag

(c) R-I Colour Variations
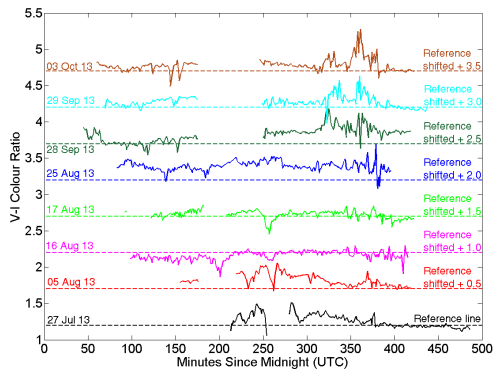
(d) V-I Colour Variations
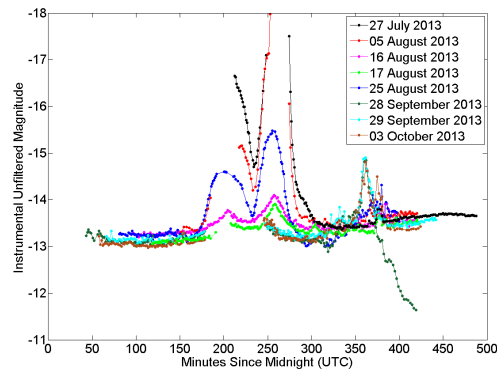
(e) Unfiltered Instrumental Magnitude

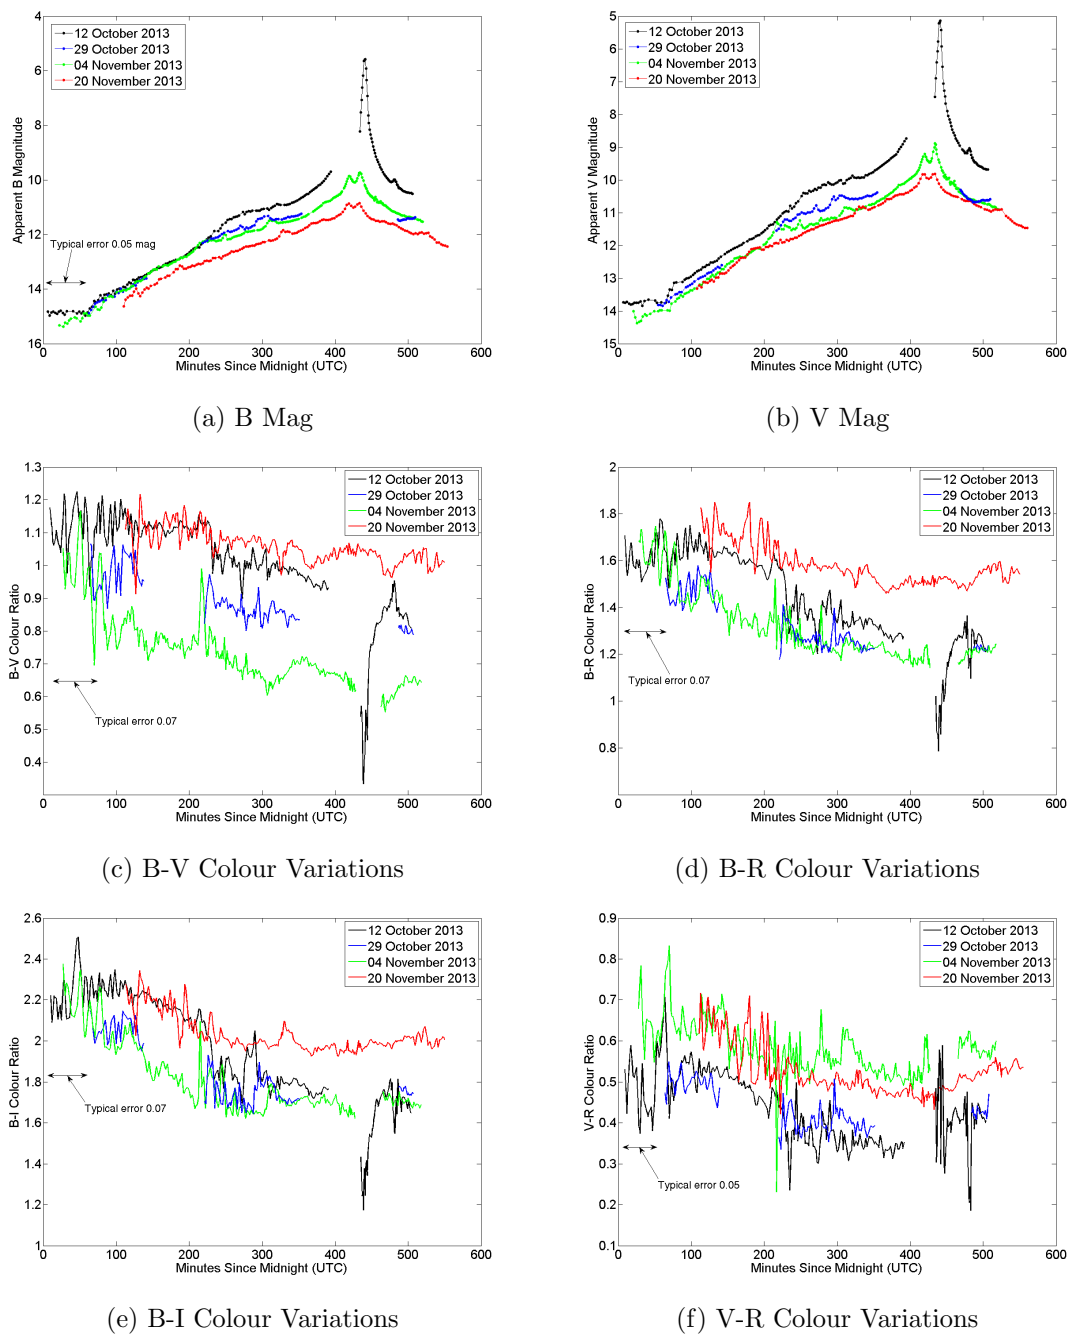Figure A.4: Galaxy 11 Lightcurves and Colour Variations.

A.3 Anik F1

(a) B Mag

(b) V Mag

(c) B-V Colour Variations

(d) B-R Colour Variations

(e) B-I Colour Variations

(f) V-R Colour Variations
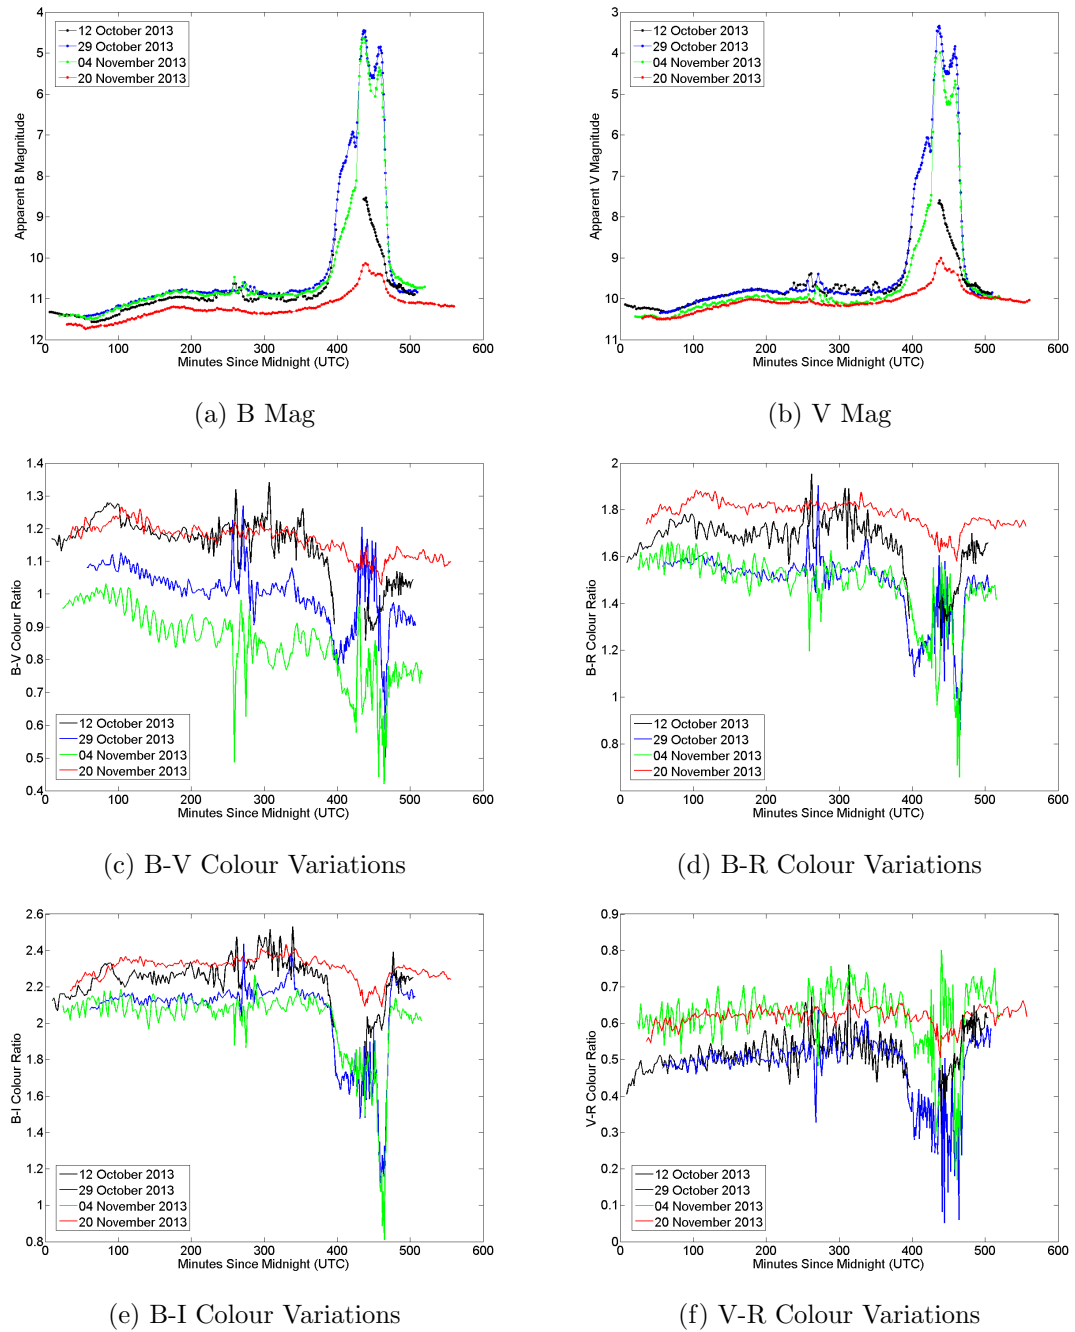
Figure A.5: Anik F1 Lightcurves and Colour Variations.
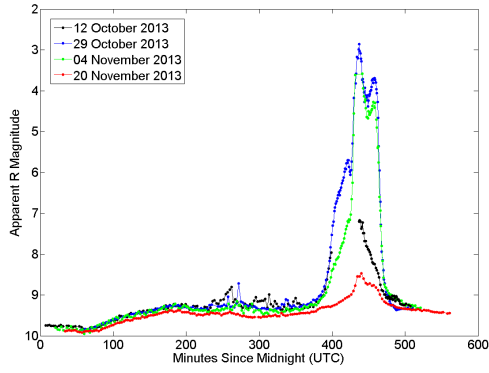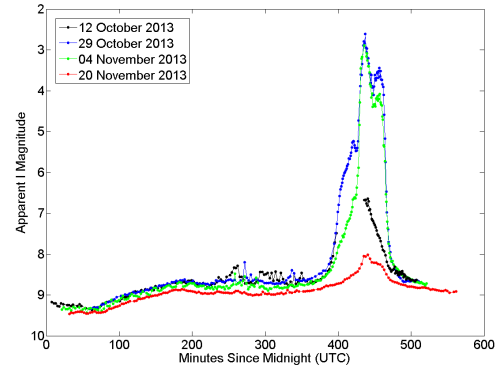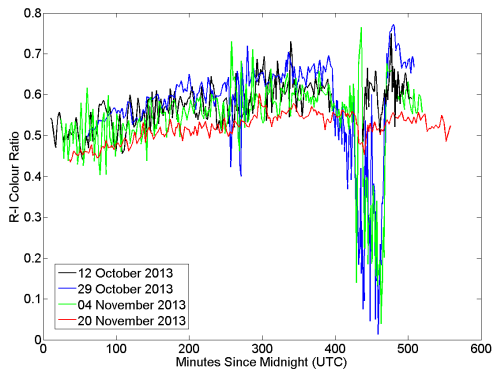
(a) R Mag

(b) I Mag

(c) R-I Colour Variations
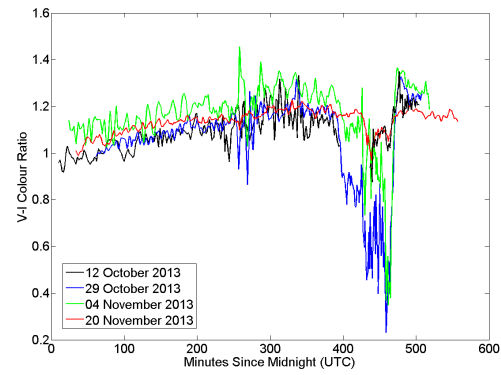
(d) V-I Colour Variations
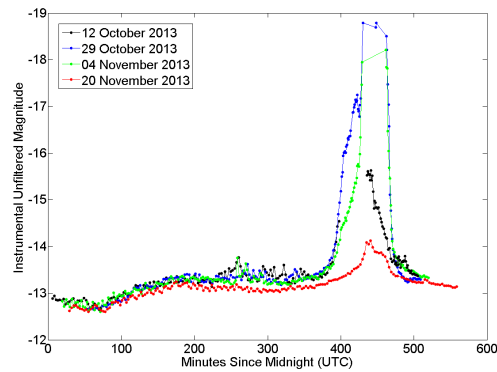
(e) Unfiltered Instrumental Magnitude

Figure A.6: Anik F1 Lightcurves and Colour Variations.
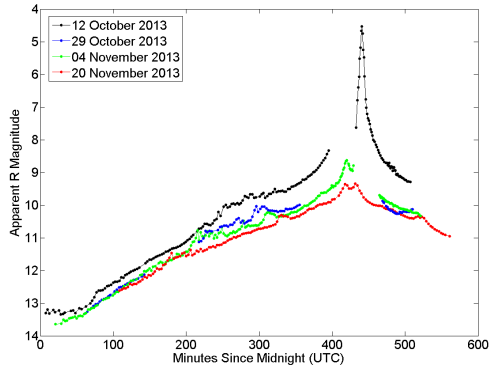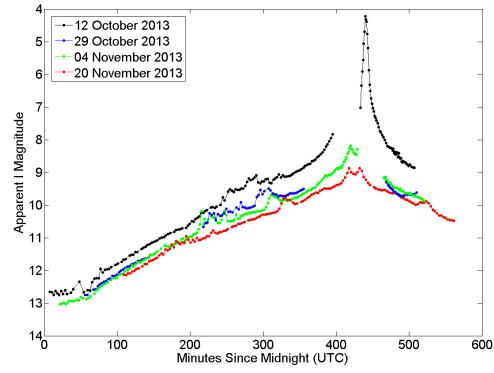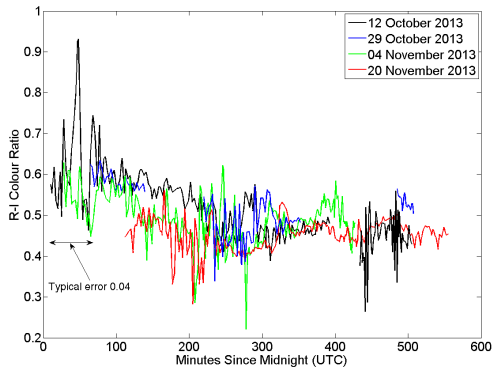
A.4   Anik F1R

(a) B Mag

(b) V Mag

(c) B-V Colour Variations

(d) B-R Colour Variations

(e) B-I Colour Variations

(f) V-R Colour Variations

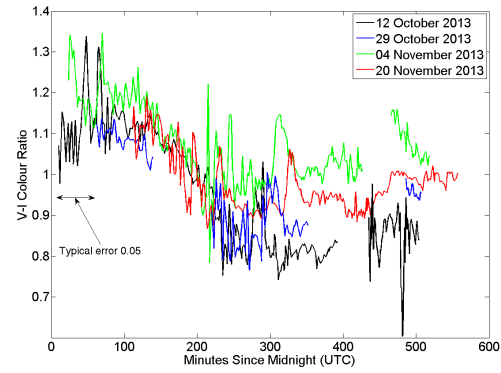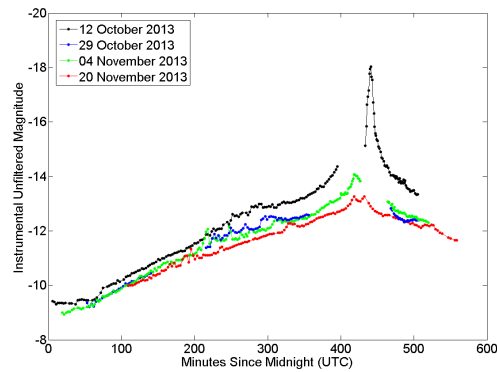Figure A.7: Anik F1R Lightcurves and Colour Variations.

(a) R Mag

(b) I Mag

(c) R-I Colour Variations

(d) V-I Colour Variations

(e) Unfiltered Instrumental Magnitude

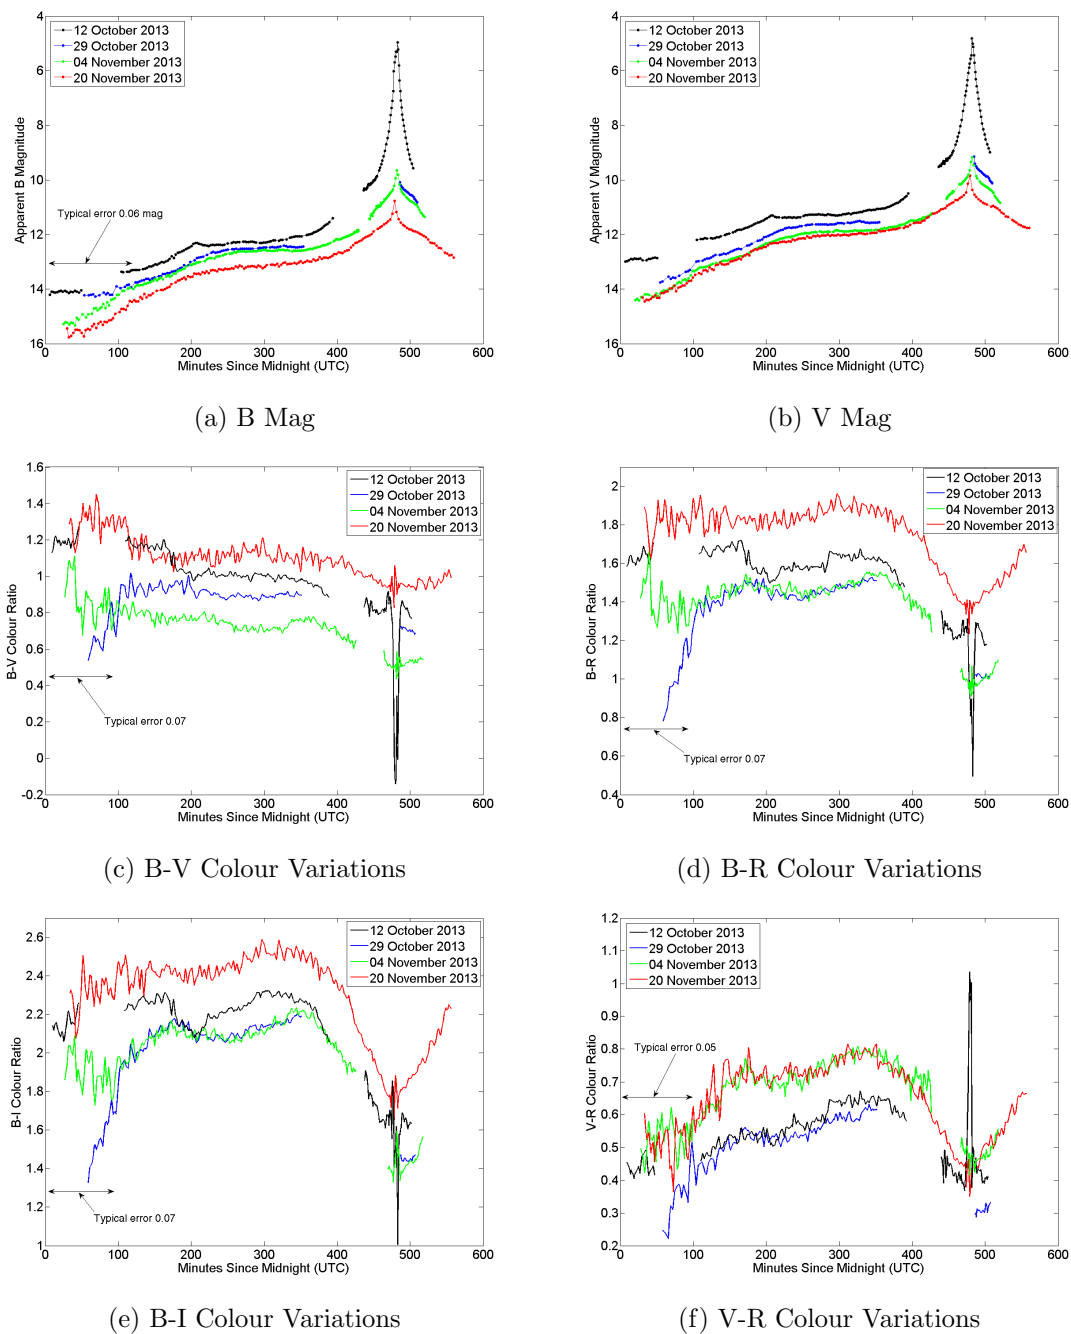Figure A.8: Anik F1R Lightcurves and Colour Variations.

A.5 Anik G1

(a) B Mag

(b) V Mag

(c) B-V Colour Variations

(d) B-R Colour Variations

(e) B-I Colour Variations

(f) V-R Colour Variations

Figure A.9: Anik G1 Lightcurves and Colour Variations.

(a) R Mag

(b) I Mag

(c) R-I Colour Variations

(d) V-I Colour Variations

(e) Unfiltered Instrumental Magnitude
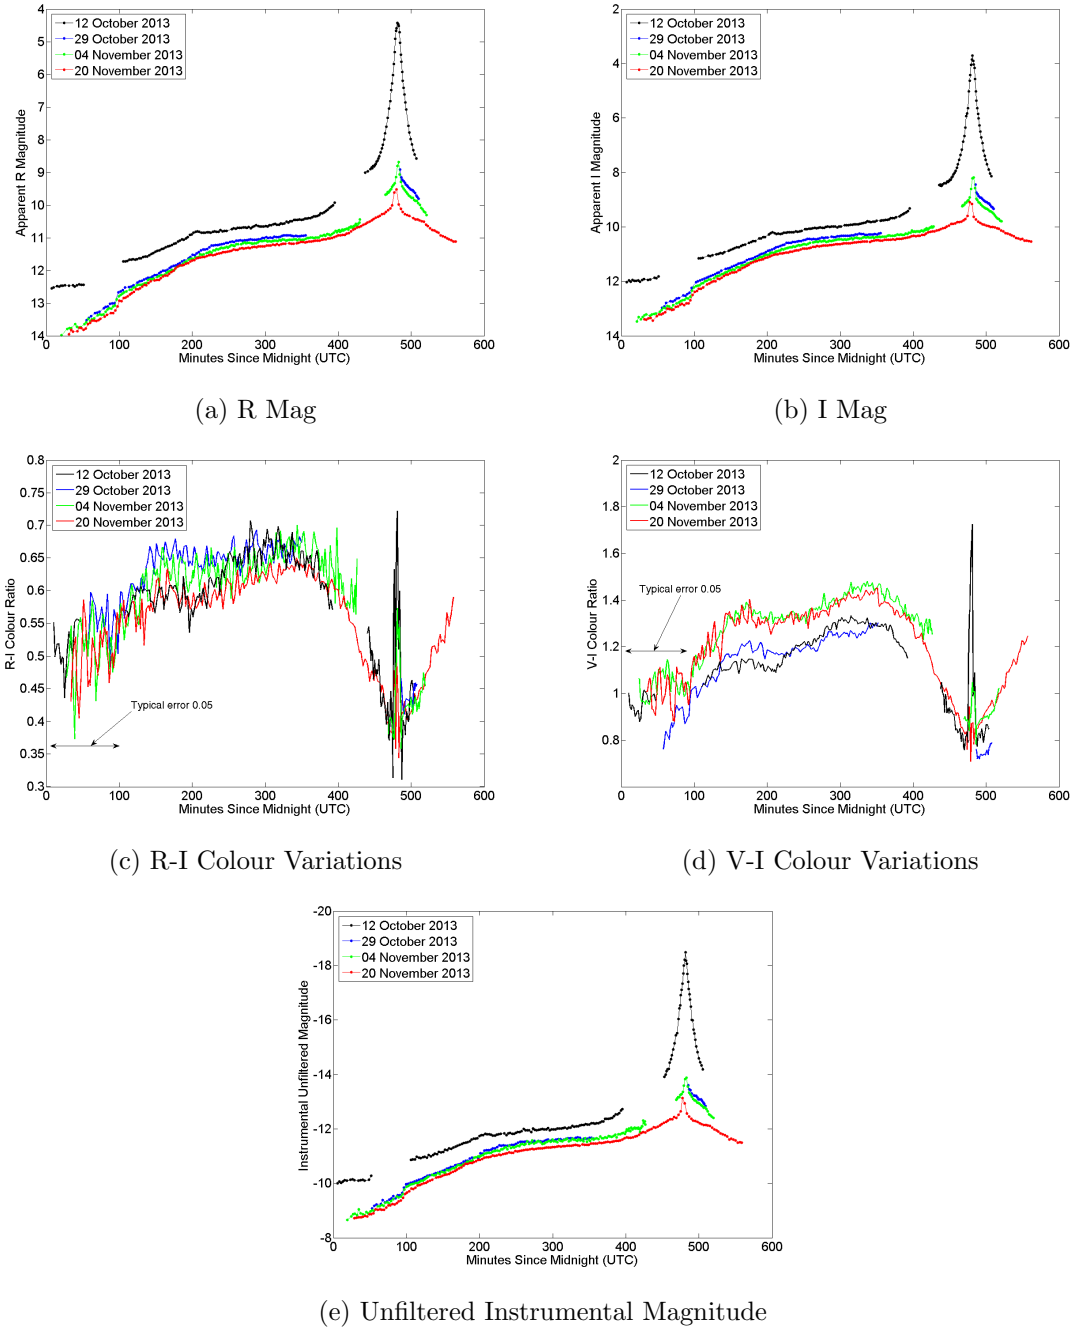
Figure A.10: Anik G1 Lightcurves and Colour Variations.

A.6   Echostar 17

(a) B Mag

(b) V Mag

(c) B-V Colour Variations

(d) B-R Colour Variations

(e) B-I Colour Variations

(f) V-R Colour Variations
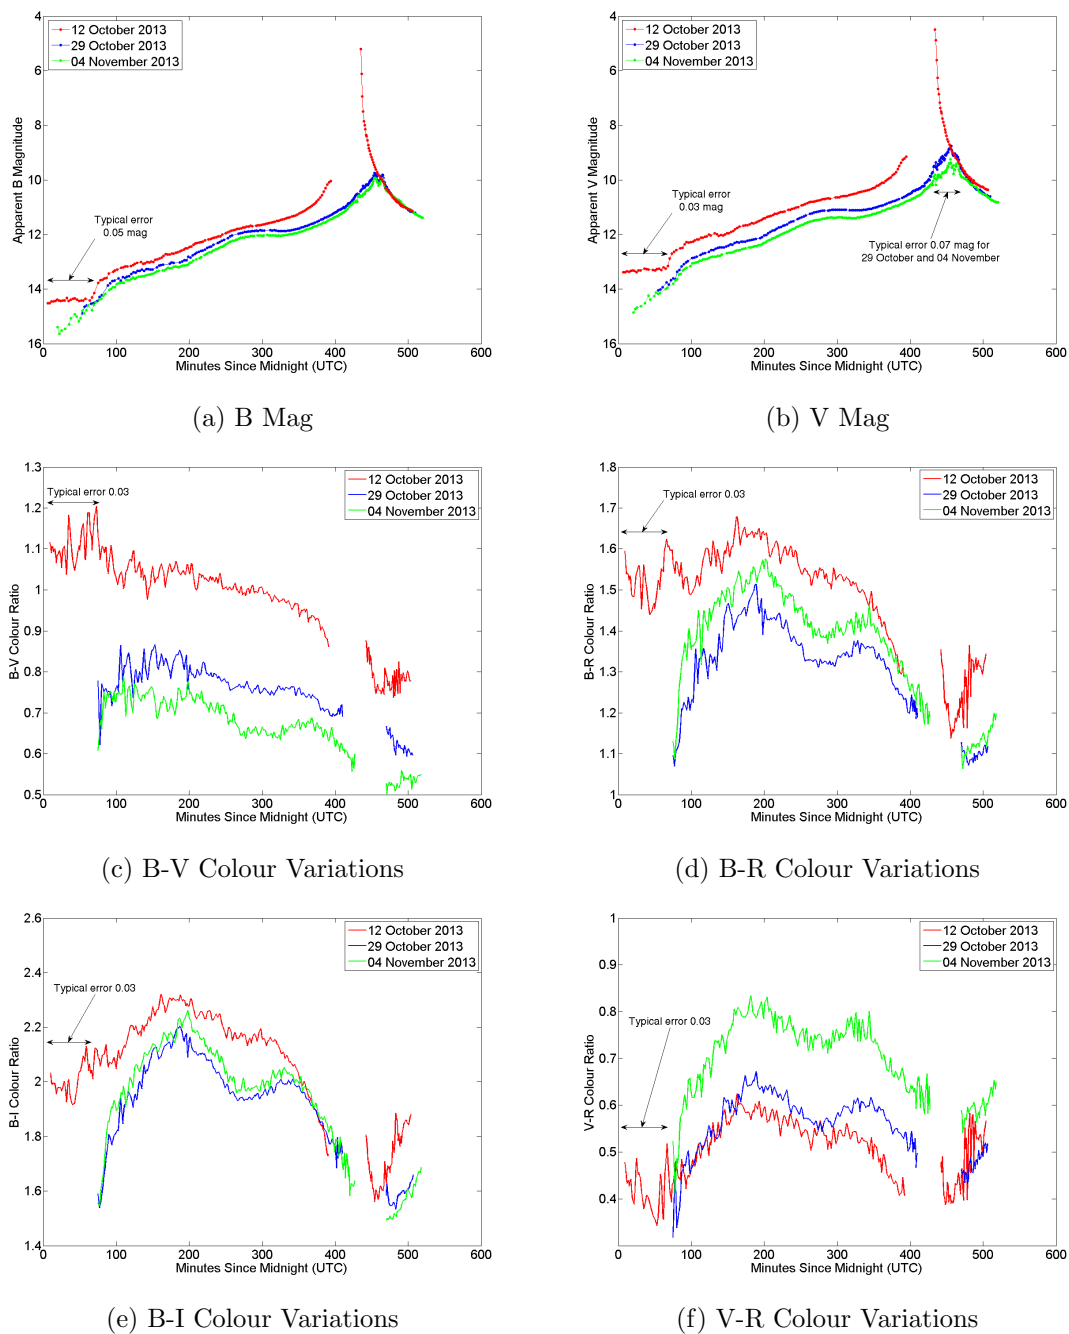
Figure A.11: Echostar 17 Lightcurves and Colour Variations.

(a) R Mag

(b) I Mag

(c) R-I Colour Variations

(d) V-I Colour Variations

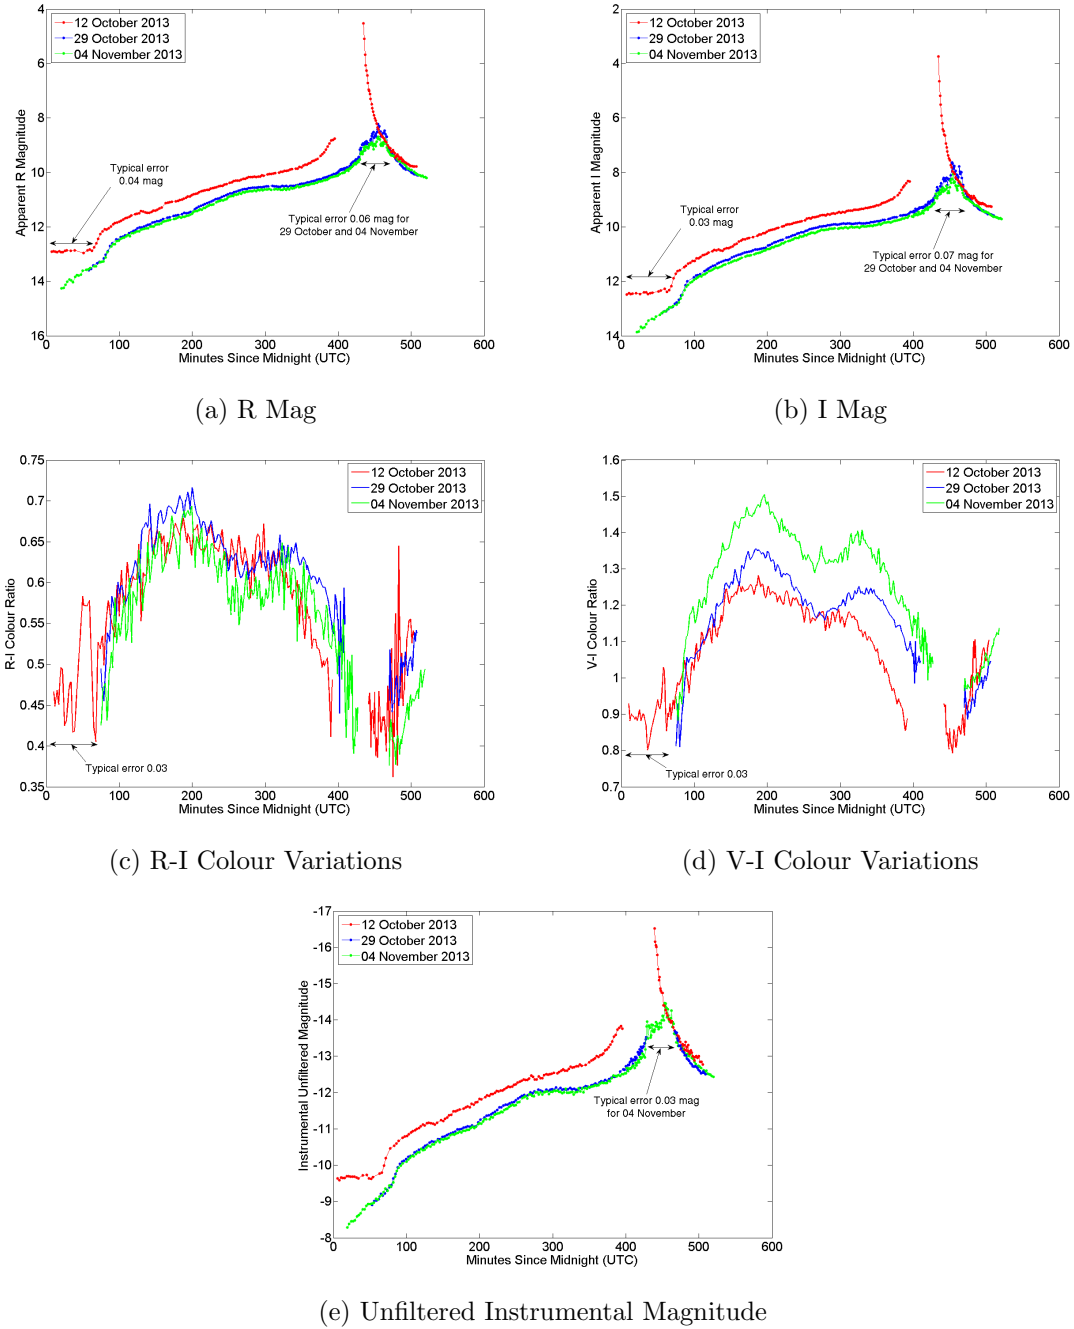(e) Unfiltered Instrumental Magnitude

Figure A.12: Echostar 17 Lightcurves and Colour Variations.

A Discussion on Phase Angle

When analyzing photometric magnitude or colour ratio measurements of a satellite, several options are available for what to use as the independent variable, and it may not be clear which is the most effective one to use. A time scale may seem appropriate, however local time is not a convenient measure for someone who could be examining the plot from anywhere in the world, and UTC time may also seem inappropriate for many observation sites. In addition, most time scales are not able to take into account geometrical changes that occur from one night to the next, and thus the time of day is not likely to uniquely determine photometric magnitude. A more appropriate measure for the horizontal axis might therefore be one that is based on a geometric consideration, since a material's BRDF dictates that the sun-object-sensor geometry ultimately determines the measured magnitude. By far the most common type of scale used to define the horizontal axis uses an angle between vectors relating the positions of the sun, satellite, and sensor to each other; called a *phase angle*. Unfortunately, at least four different definitions of phase angle have been used in the literature, which, to a newcomer to the field, could easily result in confusion. In some cases [6, 29] the type of phase angle that was used in a study is not precisely defined, and in others the phase angle appears to be incorrectly defined. For example, Gregory [17] states that he uses the satellite-sun-observer angle projected onto the equator as the phase angle definition, but the data seem to indicate that the sun-satellite-observer angle projected onto the equator was actually used. The variously defined phase angles encountered in the literature will now be defined, along with the major implications of using each one.

B.1 Equator-Projected, Signed, Satellite Centred Phase Angle

The most commonly used phase angle is measured at the satellite, between the projection of the satellite-sun and satellite-observer vectors onto the equatorial plane. This phase angle, often referred to as the *signed in-plane phase angle* or *equatorial phase angle* is depicted as ϕ in figure B.1. The orange circle represents the sun, the large shaded circle represents Earth, and the small green circle represents the observer location. The red arrows are the satellite-sun and satellite-observer vectors; and the red dotted lines join the satellite, sun, and observer to the projection of the satellite-sun and satellite-observer vectors onto the equatorial plane. The projected vectors are represented by the green arrows.
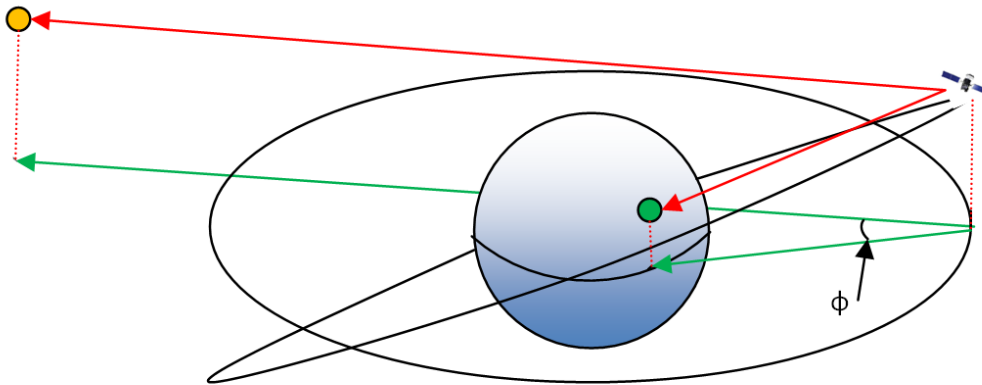

Figure B.1: Equator projected, signed, satellite centred phase angle (ϕ).

Negative phase angles occur earlier in the night than positive phase angles, when the sun is west of the extended satellite-observer vector. Because this phase angle is projected onto the equatorial plane, a satellite will always pass through zero degrees each orbit. Since the solar panels of three axis stabilised satellites are usually designed to track the sun, the maximum brightness resulting from specular reflection off the solar panels should occur when the sun, observer and satellite are aligned; assuming the solar panels perfectly track the sun. In practice Earth's shadow will often prevent the satellite from being illuminated in such a situation. However, when the satellite's orbit is such that it avoids Earth's shadow entirely, the maximum specular reflection intensity will always occur when the equator projected phase angle equals zero

(again assuming that the solar panels track the sun perfectly). Potential advantages of using this phase angle definition include the fact that there will be no discontinuities in lightcurves as they transition from negative to positive phase angles, and that differences between positive and negative phase angles will be readily apparent. A significant disadvantage is that it depicts all intensity variations as purely longitude dependent. There are two other variables that partially determine the angle between the satellite-sun vector and the satellite-observer vector (and thus the measured intensity). The latitude dependent component of the total sun-satellite-observer phase angle can vary on relatively short time scales with changes in the satellite's declination as it orbits Earth. This is on the order of a couple of hours per cycle for LEO satellites. The rate of change is more rapid, and greater in magnitude, for higher inclination orbits, and is minimal for equatorial orbits. Independent of the satellite's orbital inclination, the total phase angle varies cyclicly due to the tilt of Earth's spin axis relative to Earth's orbital plane around the sun. Since Earth's spin axis direction remains fixed in space, the solar declination changes as Earth orbits the sun. Each solar declination cycle therefore takes one year to complete. Solar declination changes contribute relatively more significantly to latitude dependent phase angle changes for low inclination orbits, such as geostationary orbits. This is because lower inclinations produce smaller changes in the latitude component of the phase angle. Figure B.2 depicts the sun-satellite-observer geometry during an equinox and solstice.

Any information regarding the effect of satellite declination or seasonal changes on the measured intensity will be lost if it is plotted against the equator-projected phase angle. In practice, seasonal effects can usually be recovered provided that the date of each observation was recorded, or by constructing a three dimensional plot of the data with day of the year on one of the axes [54].

B.2   Equator-Projected, Unsigned, Satellite Centred Phase Angle

The difference between the signed and unsigned equator projected phase angles is that negative values are not possible for the unsigned phase angle: only the magnitude of the angle is of interest. In the unsigned case the horizontal axis starts at zero
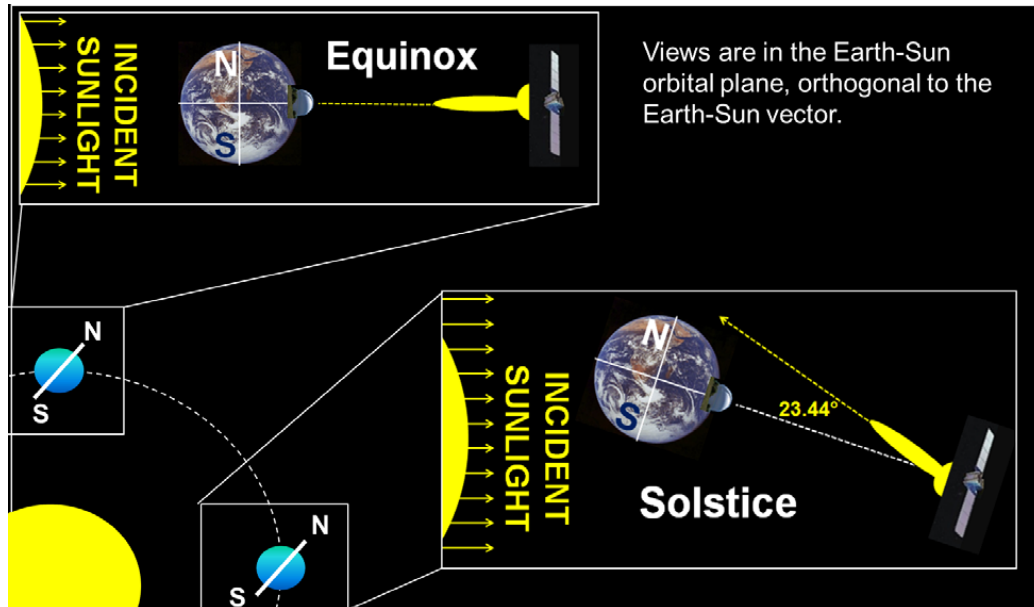
Figure B.2: Schematic showing the seasonal effect of Earth's tilt on the sun-satellite-observer angle for a GEO satellite [54].

degrees phase angle, and since it is the equator projected angle being measured, any lightcurve or colour plot that spans an entire pass across the sky will intercept the vertical axis. This phase angle definition also loses any information resulting from declination dependent variability; as well as information resulting from asymmetries about zero phase angle. In cases where there is very high temporal resolution in a lightcurve it can become obvious which measurements were taken prior to the satellite reaching zero phase angle and which measurements were taken after reaching zero phase angle. Comparing the two sets of measurements can then actually enhance the investigation of asymmetries about zero phase angle.

B.3 Solar Phase Angle

Sometimes it is desirable to plot the measured intensity as a function of the entire angle between the satellite-sun vector and the satellite-observer vector. This satellite centred phase angle, commonly known as the *solar phase angle*, but also sometimes referred to as *total phase angle* is depicted in figure B.3. As with the equator-projected phase angle, negative values occur when the sun is west of the extended satellite-observer vector, unless the unsigned phase angle is used, in which case all angles are positive. Adding to the confusion about phase angle, solar phase angle has also been used to refer to the equator-projected phase angle [54].
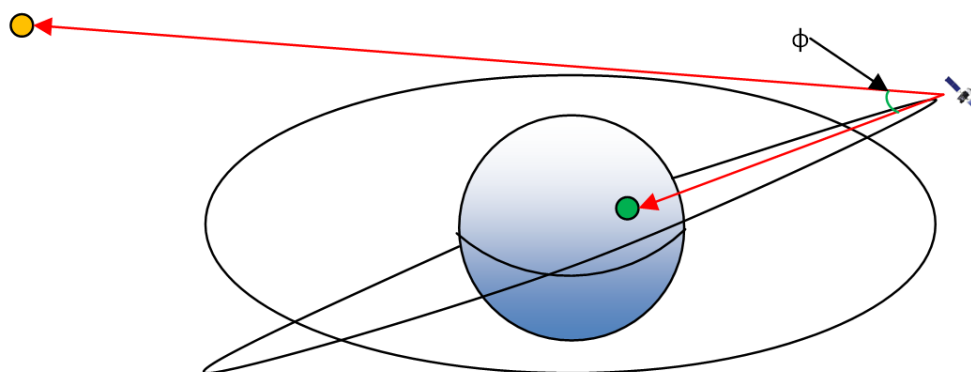

Figure B.3: Solar phase angle (ϕ).

Figure B.4 depicts the lightcurves of two GEO satellites, Anik F2 and Wildblue-1, observed over three nights, with solar phase angle as the independent variable on the horizontal axis. Note that there are no data points located within approximately ten degrees of zero phase angle. This is because at such low phase angles the satellites pass into Earth's shadow. A potential disadvantage of this phase angle definition is that there is no means of determining what proportion of the phase angle is latitude dependent and what proportion is longitude dependent. Two different measurements, taken at identical phase angles, may therefore have occurred under different geometrical conditions. As with the equator-projected phase angle, either the signed or unsigned solar phase angle can be used.

Other, less commonly used, phase angles include the signed and unsigned sun-
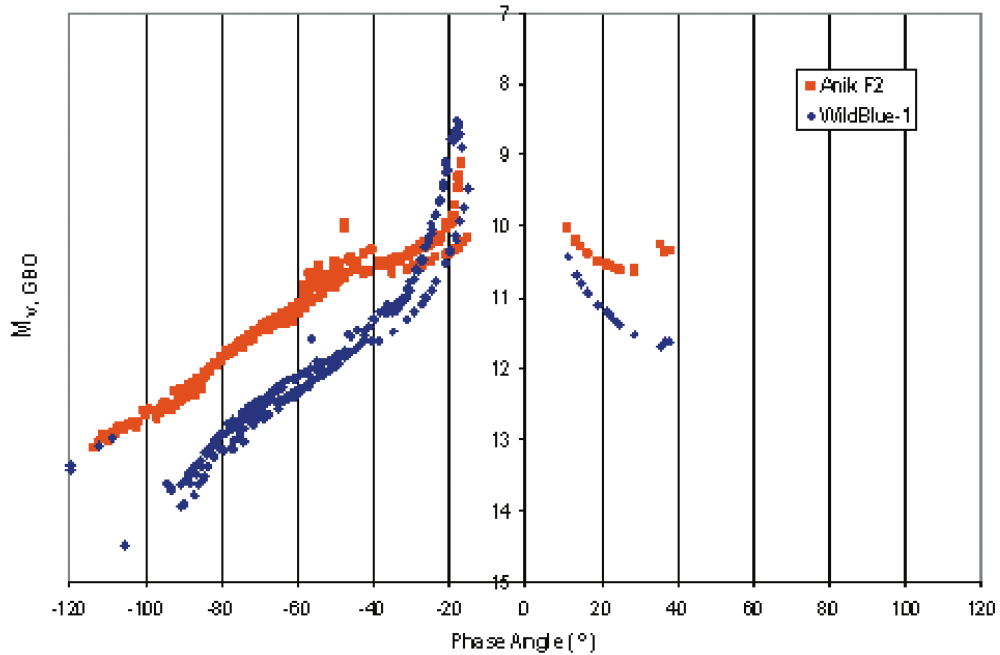
Figure B.4: Two satellite lightcurves plotted against total phase angle [27].

observer-satellite phase angle [24]; which is the angle, measured at the observer, between the observer-sun vector and the observer-satellite vector.

Given the multitude of different definitions for phase angle in the literature, and the frequency with which it is not specifically defined, it is important for the reader to understand which type of phase angle is being referred to, and its associated implications.

Curriculum Vitae

Curriculum Vitae

Name: Andrew Jolley
Place of birth: Swan Hill, Victoria, Australia
Year of birth: 1981

Post-secondary education and degrees:	**M.Sc. Candidate** Royal Military College of Canada Kingston, Ontario	**2012-2014**
	M.A. Candidate Australian Defence Force Academy Canberra, A.C.T.	**2010-2014**
	Diploma of Project Management Gippsland TAFE Sale, Victoria	**2012**
	Certificate IV in Training and Education Gippsland TAFE Sale, Victoria	**2010**
	Bachelor of Technology in Aeronautical Engineering Australian Defence Force Academy Canberra, A.C.T.	**2001-2003**

Employment history:	**Air Warfare Instructor Standardisation Officer** School of Air Warfare Sale, Victoria	**2012**
	Air Combat Officer Instructor School of Air Warfare Sale, Victoria	**2010-2012**
	Tactical Coordinator 10 Squadron Adelaide, South Australia	**2008-2009**
	Navigator/Communicator 10 Squadron Adelaide, South Australia	**2006-2008**
Honours and awards:	**Australian Active Service Medal**	**2006**
	Iraq Campaign Medal Operations Slipper and Catalyst	**2006**
	Afghanistan Campaign Medal Operations Slipper and Catalyst	**2009**