THE UNIVERSITY
OF QUEENSLAND
AUSTRALIA

Sensor Management for Enhanced Catalogue Maintenance of Resident Space Objects

Tyler A. Hobson
Bachelor of Engineering (Aerospace Avionics) - First Class Honours

A thesis submitted for the degree of Doctor of Philosophy at
The University of Queensland in 2014
School of Information Technology and Electrical Engineering

Abstract

The research detailed in this thesis explores the deployment of steerable sensors as an efficient means of improving the capability, capacity and timeliness of existing space surveillance systems, to provide superior levels of Space Situational Awareness (SSA) through enhanced sensor management. These improvements are necessary as the world's increasing reliance on spacefaring has brought about the accumulation of an enormous quantity of man-made objects orbiting the Earth. For almost 60 years the number of objects has grown, causing a commensurate increase in the likelihood of destructive collisions involving manned missions and important space assets. To predict and ideally prevent collisions, a number of agencies endeavour to track as many Resident Space Objects (RSOs) as possible. Recent events involving unprecedented surges in the number of RSOs have made it clear that the ability to continue to operate safely in Earth orbit will require enhancements to existing levels of SSA. The act of maintaining SSA is reliant on many sources of information, of which a primary source is the direct observation of RSOs by space surveillance sensors. These observations are utilised to compile and maintain a catalogue of RSO's orbital state estimates that is analysed to determine the likelihood of collision. Surveillance of this environment is a challenging task that currently has a large dependence on legacy systems and techniques that can benefit from modernisation via the introduction of contemporary technologies and methodologies. Due to potential benefits such as low cost, high accuracy, scalability, flexibility and automation, the large scale deployment of steerable sensors is proposed as a means of improving existing catalogue maintenance systems.

Researching a judicious means of deploying and exploiting steerable sensors to improve the capability, capacity and timeliness of space surveillance networks requires consideration of the management of sensors at both a network and an individual level.

The exploration begins at the network level with an analysis of existing practices for maintaining RSO catalogues to understand how catalogue accuracy is affected when steerable sensors are employed. Through numerical simulation, the effectiveness of the current state of the art in steerable sensors, a class of electro-optical sensor, is contrasted with traditional radar surveillance. The findings indicate that if the current state of the art in steerable sensors were to be widely deployed as the primary contributing sensors, catalogue accuracy would increase significantly. The findings also show that greater catalogue accuracy could be expected if effects caused by passive optical sensing to observability of RSO range and sensor availability can be minimised.

Methods for improving observability and availability when using networks of optical sensors are investigated next. Measurement level sensor fusion and efficient analysis of the

network's visibility of the RSO catalogue are considered. The investigation's results show that measurement level sensor fusion is capable of reducing catalogue error caused by weak observability of range. However, the effectiveness of the result is highly dependent on the distribution of sensors and ensuring multiple sensors are only tasked to observe a single object when it benefits the catalogue as a whole. Parallel General Purpose computing on a Graphical Processing Unit (GPGPU) is employed to achieve efficient, full-scale simulation and visibility prediction of alternative network configurations. The findings indicate that using a high ratio of optical to radar sensors can achieve high levels of system availability when monitoring a realistic distribution of RSOs. The practicality of the network is further enhanced via visibility prediction by enabling the sensor manager to anticipate small lapses in coverage and schedule observations accordingly. These techniques may be used to implement a surveillance network using steerable electro-optical sensors that overcome the identified constraints to observability and availability.

The final investigation aims to improve upon system capacity and timeliness via enhanced management and control of individual steerable sensors using real-time, GPU-augmented decision making at the sensor. A novel method called dynamic steering is proposed to exploit this architecture and enable a sensor to autonomously switch between tracking and searching whilst reacquiring catalogued objects. Successful, automatic reacquisition was demonstrated in an experimental field trial of the system using targeting data as much as six months old, far surpassing the limits of existing surveillance systems. This result demonstrates an ability to achieve greater levels of system capacity as certain RSOs may be observed less regularly. The autonomy achieved increases the timeliness of the system, as dynamic steering permits the replacement of tasks currently conducted by human operators and reduces the level of tracking refinement necessary before information may be shared between sensors. Such capabilities have the potential to improve existing methods and inspire new techniques for obtaining SSA.

The proposed techniques for managing sensors at the network and individual sensor level may be used to improve the capability, capacity and timeliness of existing space surveillance systems to achieve improved SSA.

Declaration by Author

This thesis is composed of my original work, and contains no material previously published or written by another person except where due reference has been made in the text. I have clearly stated the contribution by others to jointly-authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my research higher degree candidature and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that an electronic copy of my thesis must be lodged with the University Library and, subject to the General Award Rules of The University of Queensland, immediately made available for research and study in accordance with the Copyright Act 1968.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material. Where appropriate I have obtained copyright permission from the copyright holder to reproduce material in this thesis.

Publications during candidature

- Tyler A. Hobson and I. Vaughan L. Clarkson *Sensor-scheduling simulation of disparate sensors for Space Situational Awareness.* presented at the 12th Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference, Maui, Hi, Sep. 2011

- Tyler A. Hobson and I. Vaughan L. Clarkson *Collaborative sensor scheduling for Space Situational Awareness.* presented at the 2012 IEEE Statistical Signal Processing (SSP 12) Workshop, 5-8th Aug. 2012

- Tyler A. Hobson and I. Vaughan L. Clarkson *GPU-based Space Situational Awareness simulation utilising parallelism for enhanced multi-sensor management.* presented at the 13th Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference, Maui, Hi, Sep. 2012

- Tyler A. Hobson and I. Vaughan L. Clarkson *A Particle-based search strategy for improved Space Situational Awareness.* presented at the Asilomar Conference on Signals, Systems & Computers, Asilomar, CA, 3-6th Nov. 2013

- Tyler A. Hobson and I. Vaughan L. Clarkson *An experimental implementation of a particle-based dynamic sensor steering method for tracking and searching for space objects.* accepted for presentation at the 39th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Florence, Italy. 25-30th May 2014

- Tyler A. Hobson and I. Vaughan L. Clarkson*Dynamic Steering, Tracking and Automated Reacquisition for Space Surveillance.* Submitted for publication in IEEE Transactions on Aerospace and Electronic Systems, 2014

Publications included in this thesis

No publications included.

Contributions by others to the thesis

No contributions by others.

Statement of parts of the thesis submitted to qualify for the award of another degree

None.

Acknowledgements

First and foremost, I would like to thank my principal research advisor A.Prof. Vaughan Clarkson. I am very grateful for his dependable guidance, endless enthusiasm and willingness to lend an ear throughout all stages of candidature. I am also indebted to him for his efforts forging new opportunities for me and his willingness to introduce and promote me to his high-calibre colleagues and friends.

I would also like to express my gratitude to Dr Kim Luu of the Air Force Research Laboratory (AFRL) for offering her wealth of knowledge regarding my topic of research, her steadfast support—particularly in making opportunities such as the field trials a reality—and for her efforts making visits to Maui welcoming, rewarding and enjoyable. I also owe her family my thanks for their hospitality and company during visits to Maui. On that note, I would like to acknowledge Dr Chris Sabol of AFRL, whose experience and knowledge regarding Raven sensors was highly valued.

I would like to thank a great many of the friendly and talented staff at Pacific Defence Solutions (PDS). Their collaboration and willingness to share their extensive knowledge regarding space surveillance was critical to achieve the results I have detailed in this thesis. I would particularly like to thank Randy Cortez, Keric Hill and Paul Sydney for their unabated support and enthusiasm at all stages throughout my research and for going above and beyond to ensure my visits to Maui were rewarding. In addition to these names, I would like to thank Kris Hamada, Laurie Hozaki, Kawai Kuluhiwa, Daron Nishimoto and Vicky Sensano for their perseverance, experience and talents making the field trials reported in this thesis possible.

I am very grateful to Prof. Neil Gordon of DSTO for sharing his vast knowledge and passion for surveillance and tracking. His willingness and enthusiasm to discuss his own work, promote others and take the time to answer the questions of a lowly research student are a testament to his character.

I would like to acknowledge the Air Force Office of Scientific Research (AFOSR) for funding the research reported by this thesis via FA9550-10-1-0493. Without this funding, I would not have been able to invest the time necessary to research this topic and collaboration with researchers in the middle of the Pacific would not have been possible.

My thanks also go out to my fellow students and staff at the University of Queenland (UQ) who were willing to lend an ear, assist me during my many endeavours and who have made UQ an enjoyable place to work throughout my candidature. I would like to extend a special thanks to Marnie Lamprecht and Kimberley Nunes who were particularly supportive and always willing to discuss the nontrivial task of piecing together a thesis.

Finally, I would like to thank my family, to whom I owe the most. I would like to thank my parents Alan and Barbara for their unconditional and boundless support; Paula for her strength and belief in me and all that I pursue as well as the remainder of my immediate and extended family, for their ceaseless interest and encouragement in my efforts.

Keywords

space situational awareness, sensor management, parallel computing, automation, tracking, search

Australian and New Zealand Standard Research Classification (ANZSRC)

ANZSRC code: 090609, Signal Processing, 60%

ANZSRC code: 090602, Control Systems, Robotics and Automation, 20%

ANZSRC code: 090199, Aerospace Engineering not elsewhere classified, 20%

Fields of Research (FoR) Classification

FoR code: 0906, Electrical and Electronic Engineering, 80%

FoR code: 0901, Aerospace Engineering, 20%

Contents

List of Figures

List of Tables

List of Acronyms

AFRL	Air Force Research Laboratory
AGI	Analytic Graphics Incorporated
API	Application Programming Interface
CKF	Conventional Kalman Filter
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture (NVIDIA)
ECEF	Earth Centred Earth Fixed
EKF	Extended Kalman Filter
EOP	Earth Orientation Parameter
FOV	Field of View
GEO	Geosynchronous Orbit
GLONASS	Globalnaya Navigatsionnaya Sputnikovaya Sistema
GPGPU	General-purpose computing on Graphics Processing Units
GPS	Global Positioning System
GPU	Graphics Processing Unit
HAMR	High Area to Mass Ratio
HEO	High Earth Orbit
HST	Hubble Space Telescope
i.i.d.	Independent and Identically Distributed
IEEE	Institute of Electrical and Electronics Engineers

IERS	International Earth Rotation Service
IPE	Inflated Process Error
ISS	International Space Station
ITRF	International Terrestrial Reference Frame
JPL	Jet Propulsion Laboratory
LEO	Low Earth Orbit
LKF	Linearised Kalman Filter
LOS	Line of Sight
MASSAS	MATLAB Space Situational Awareness Simulation
MCPE	Median Catalogue Position Error
MEO	Medium Earth Orbit
NANU	Notice Advisory to NAVSTAR Users
NASA	National Aeronautics and Space Administration
NOMAD	Naval Observatory Merged Astrometric Dataset
NORAD	North American Aerospace Defence
OAR	Optically Augmented Radar
ODTBX	Orbit Determination Toolbox
p.d.f.	Probability Density Function
PC	Personal Computer
PDS	Pacific Defence Solutions
PF	Particle Filter
RAAN	right ascension of the ascending node
RAM	Random Access Memory
RHS	Right Hand Side
RLS	Recursive Least Squares
RPF	Regularised Particle Filter
RSO	Resident Space Object
SDK	Software Development Kit

SGP4	Simplified General Perturbations model 4
SIS	Sequential Importance Sampling
SMC	Sequential Monte Carlo
SNSS	Sensor Network for Space Surveillance
SPARSE	Space Particle Search Evaluation program
SRP	Solar Radiation Pressure
SSA	Space Situational Awareness
SSN	Space Surveillance Network
TASMAN	Tasking Autonomous Sensors in a Multiple Application Network
TEME	True Equator Mean Equinox
TLE	Two Line Element
UKF	Unscented Kalman Filter
USNO	US Naval Observatory
USSTRATCOM	United States Strategic Command
UTC	Universal Time Coordinated
WCPE	Worst-case Catalogue Position Error
WGS-84	World Geodetic System - 1984

1 Introduction

On 4th October 1957 the Soviet Union launched the first artificial satellite into Earth orbit, signalling the beginning of the Space Age. Ever since this event, the world has become increasingly reliant on the advantages that spacefaring brings. Some of these advantages include global telecommunications, precision terrestrial navigation, Earth observation, exoatmospheric astronomy, exploration, space-based experimentation as well as a number of other applications concerning national security [1]. Furthermore, a number of private companies have recently begun or are on the verge of offering transportation [2,3], tourism [4,5] and resource gathering [6,7]. These emerging applications demonstrate that the list of advantages continues to expand and indicates that space technology will remain in a state of growth for the foreseeable future. The process of placing most of these space technologies in orbit involves propelling a comparatively large launch vehicle into space that produces a trail of rocket stages, separation devices, fairings and rocket exhaust products which fall behind the relatively small payload and are, in general, left to drift about the Earth [8]. Ordinarily, the asset itself has a finite lifetime such that once it is out of fuel and power, it too becomes a piece of debris passively drifting in Earth orbit. While a significant portion of this debris re-enters the Earth's atmosphere, commonly disintegrating due to the heat of re-entry, many objects persist in orbit [9]. Fig. 1.1 has been reproduced from a recent National Aeronautics and Space Administration (NASA) publication [10]. It displays the number of unclassified

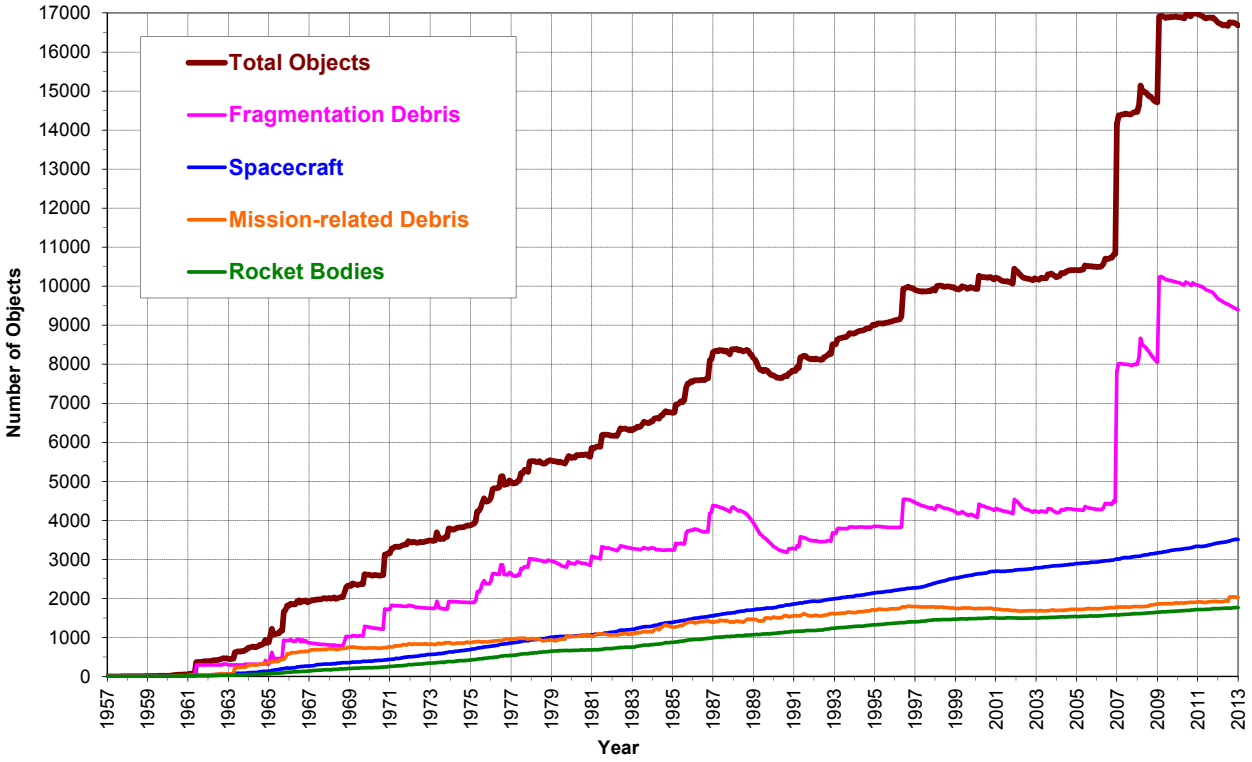

Figure 1.1: The number of unclassified man-made space objects tracked by the US each month from 1957 to 2013. "Fragmentation debris" includes debris associated with satellite and launch vehicle breakups, "spacecraft" includes active and defunct spacecraft, "mission-related debris" includes debris discarded during a planned mission and "rocket bodies" includes all primary rocket-stages separated during launch.

man-made objects tracked by the US since 1957. The trend portrayed by Fig. 1.1 demonstrates that for over half a century the number of man-made objects orbiting the Earth has continued to rise. The primary concern regarding this trend, is that as the number of objects in Earth orbit rises, there is a commensurate rise in the likelihood of collision between these objects [11].

Consisting of manned vehicles, active payloads and passive debris, these resident space objects (RSOs) all orbit the Earth at extreme velocities, resulting in relative velocities up to the order of tens-of-kilometres per second [8]. Even if one of the objects is but a few grams in mass and centimetres in length, collisions at these hyper-velocities may result in the destruction of both objects due to the enormous kinetic energies involved [11]. This is of great concern as the size of objects included in Fig. 1.1 are reportedly larger than 10 cm [12]. There are believed to be hundreds of thousands more smaller debris-objects currently beyond the sensing capabilities of any agency's surveillance network that may

still cause severe damage or destruction if they collided with other resident space object (RSO)s [11]. A recent example of a destructive collision occurred in 2009 between a defunct Russian Cosmos satellite and an active US Iridium satellite [13]. The collision resulted in the obliteration of both objects and produced thousands of pieces of smaller debris. Large amounts of fragmentation debris has also resulted from unintentional fragmentation events such as explosive hardware failures of rocket motors and batteries [11], as well as intentional fragmentation events such as those caused by the use of anti-satellite weapons. Two recent examples include a test by China in 2007 [14], which produced thousands of debris fragments that will remain in orbit for a number of decades, and the destruction of a defective satellite by the US in 2008 [15], whose comparatively small amount of debris took approximately 40 days to re-enter Earth's atmosphere. Both of these events, in addition to the 2009 Iridium/Cosmos collision, are clearly defined in Fig. 1.1 as sharp rises in the amount of fragmentation debris. Events such as these, that turn one or two pieces of debris into many, can dramatically increase the likelihood of subsequent collisions. This is because the fragmentation debris spreads throughout Earth orbit due not only to changes in velocity caused by the initial fragmentation event, but also due to long-term perturbations as a result of variances in material composition and area-to-mass ratio. In the worst case, a cascade of collisions is possible, as each subsequent collision may cause secondary collisions resulting in entire regions of Earth orbit becoming unserviceable and possibly impenetrable due to a cloud of debris [16]. For the continued use of space technology, the safety of manned missions and in the interests of national security, it is therefore essential to maintain situational awareness of these RSOs.

Whilst the prevention of further collisions in space will require the implementation of long-term solutions [17], the short-term planning, monitoring and prediction necessary to prevent losing space-vehicles and the advantages that spacefaring brings, requires space situational awareness (SSA). The term SSA is used somewhat inconsistently [18–22], but for the purposes of this thesis, will refer to a knowledge of the man-made space environment surrounding the Earth. A primary means of contributing to SSA is by performing space-surveillance and tracking of RSOs [23–25]. In the ideal scenario, each object's dynamics would be measured and modelled mathematically to the extent that all future instances of collision could be reliably predicted and prevented. Accurately measuring the state of an enormous number of unique objects, at great range is however a challenging task for even the most capable space surveillance sensors [9]. Furthermore, the best models for orbital dynamics are incomplete in the sense that they do not describe all perturbations to orbital motion, as it is intractable to measure all influences or predict their effects above a certain level of accuracy [26]. Unannounced manoeuvres and launches, equipment failures

and fragmentation events all add additional uncertainty.

For these reasons, the act of tracking RSOs to contribute to SSA is one of repeated measurement [12]. Whilst supplemented by secondary sources such as reporting by operators and public disclosures, a catalogue of RSOs is primarily maintained by routinely reacquiring each object with space surveillance sensors. Their state estimates are updated using the latest observations in order to estimate, characterise, detect changes and perform predictions on each object's orbit. The accuracy of the RSO catalogue thereby directly influences the level of SSA that may be attained. Its contents hold the most complete record of the state of Earth Orbit and is the ultimate source for predicting the likelihood of one RSO colliding with another [23].

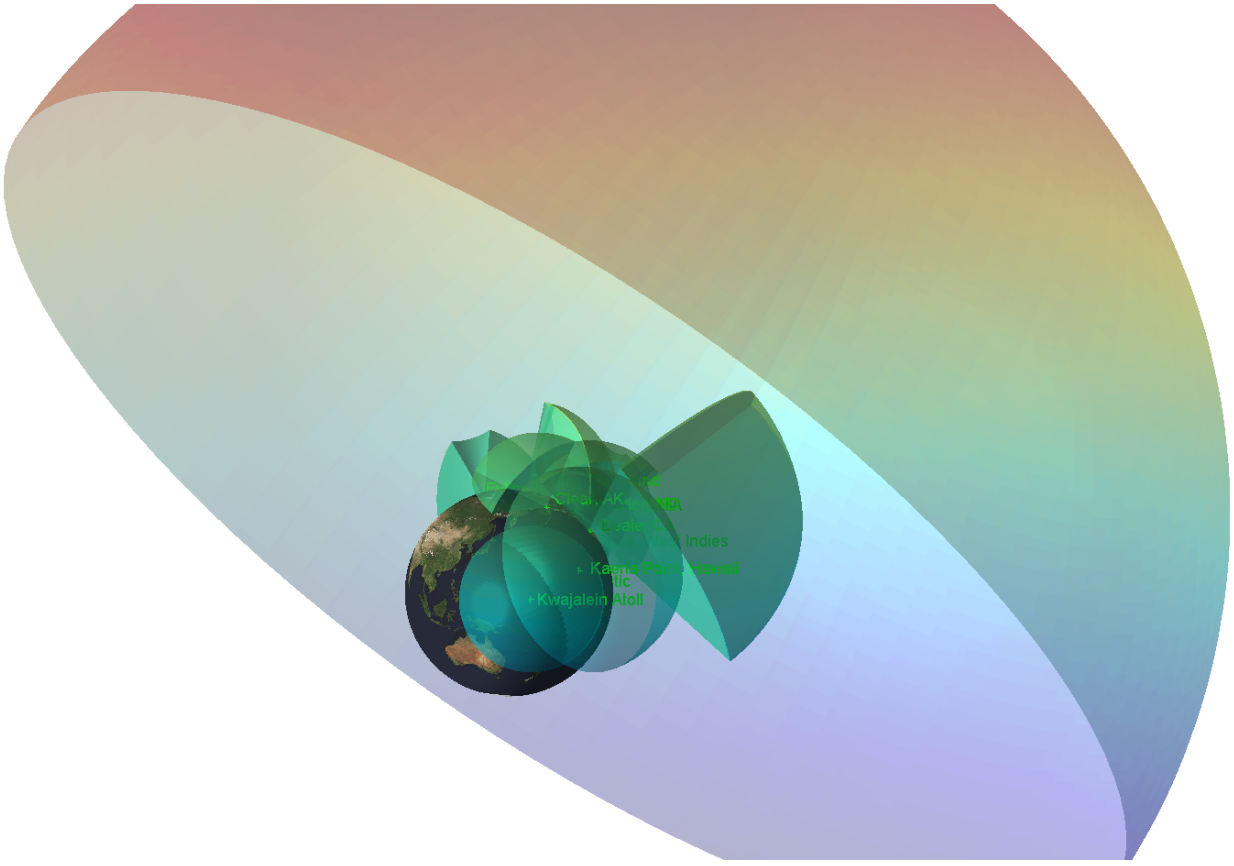

Figure 1.2: A visualisation of the surveillance volumes of a number of the SSN's radars.

Agencies from a number of countries—including America, Russia and China [18] in addition to an upcoming European system [21]—endeavour to maintain SSA by tracking as many RSOs as their respective surveillance systems permit. United States Strategic Command (USSTRATCOM) is one such agency. USSTRATCOM's Space Surveillance Network

(SSN) is purportedly world-leading in terms of coverage and capability [12, 18]. USSTRATCOM uses a global network of approximately 30 space sensors to routinely observe over 20 000 RSOs [22]. Around 95% of these objects are non-active debris. Fig. 1.2 displays a visualisation of the surveillance volumes of a number of the radars that contribute to the SSN's surveillance capability [9]. In light of the many difficulties faced when tracking a large number of unique objects in a challenging and dynamic environment and in spite of an agency's best efforts, existing techniques and procedures can lead to objects becoming lost [19, 27, 28]. In a recent report appraising the US Air Force's space surveillance practices it is noted that:

> "The space catalog lost list... is currently at an all-time high. ... recovering of lost satellites from uncorrelated tracks are manually intensive and require the talents of subject-matter experts..." [22].

Furthermore, projected growth, as a result of increasing commercial interest in space technology and the unveiling of advanced sensors capable of observing more debris than ever before [28, 29], has the space surveillance community expecting a significant increase in the number of RSOs requiring surveillance in the near future [19, 22, 30].

To overcome these challenges, the investigation of methods for increasing system capacity, capability and timeliness is crucial if the level of SSA is to keep pace with the evolving space environment.

1.1 Research Rationale and Methodology

Prominent sensor technologies for observing the state of an RSO include radar, electro-optics, laser range finding, radio transponders and the US Global Positioning System (GPS) [9, 31, 32]. The latter three options are not as commonly employed as the former two due to a number of limitations. The primary reason is that they, in general, require a 'cooperative target'. In the case of radio transponders and GPS, specialised electrical equipment must be installed and operational on the RSO for it to provide any tracking information. Laser range finding is typically applied to vehicles fitted with retro-reflectors permitting lower, and therefore safer, levels of light to be used to measure the RSO's state [9]. Whilst these sensing technologies provide many benefits to catalogue maintenance and research pertaining to SSA, they are not suited for regularly maintaining a large catalogue containing mostly uncooperative objects [32].

Space surveillance radars have been the traditional workhorse for RSO catalogue maintenance due, in large part, to a related heritage in ballistic missile warning systems [9, 18].

Their usage is well justified in the missile defence arena due to the need for a system that routinely scans a surveillance volume for targets with high reliability and availability. By means of actively projecting radio waves into a surveillance region and analysing any reflections off targets which pass through the volume, the radar provides exceptional situational awareness.

Use of active sensing nonetheless comes at a high cost and technical effort when observing RSOs in higher orbits, due to the power and level of sensitivity required at great distances [9,32]. This range limitation has prompted the use of optical sensing devices more traditionally employed in the field of astronomy, to assume surveillance roles for which radar is ill suited [33]. Whilst not without their own limitations, due in large to an optical sensor's reliance on passive illumination of the target, the cost of observing RSOs at range is significantly reduced. Exposure of optical surveillance strategies to space surveillance agencies and researchers of SSA has resulted in developments such as space-based space surveillance [28,34], wide-field optical surveillance [29] and low cost, high accuracy steerable sensors [35–37]. Whilst these latest innovations in space surveillance technologies promise a number of improvements, it is the attributes of the steerable sensors that have prompted investigation of the nonconventional practices for maintaining RSO catalogues that are detailed in this thesis.

The research reported in this thesis aims to investigate if novel sensor management techniques may be applied to large deployments of steerable sensors to achieve superior levels of SSA. The term 'steerable sensor' is used henceforth in reference to any space surveillance sensor whose finite field of view (FOV) is small enough to warrant non-uniform steering of the sensor about its surveillance volume in order to maximise its utility. These sensors necessitate an alternative but complementary approach to space surveillance as they intermittently interrogate confined regions of space rather than stare indefinitely or regularly sweep their entire surveillance volume. The current state of the art in steerable sensors is the Raven class electro-optical sensor [35,36]. This class of sensor offers features such as low cost, high accuracy, scalability, flexibility and the potential for higher than existing levels of automation [37]. The surveillance attributes of steerable sensors have the capacity to offer capabilities that warrant the SSA community to re-evaluate their heavy reliance on traditional high-cost active sensing strategies. The wider deployment and promotion of steerable sensors to a primary role in space surveillance is therefore proposed as an efficient means of improving the capability, capacity and timeliness of existing space surveillance systems.

Investigating the deployment and exploitation of steerable sensors requires consideration of the management of these sensors at both a network and an individual level. The investigation began at the network level involving an analysis of existing practices for maintaining

RSO catalogues to identify the benefits to catalogue accuracy when steerable sensors are employed. Exploration of this topic quickly warranted the development of an efficient means of simulating the process of catalogue maintenance via space surveillance. The simulation environment was used throughout the reported research to aid in conception, design, visualisation, analysis and verification of most outcomes. An investigation of alternative scheduling strategies was conducted next, to determine a means of maximising the effectiveness of a network of steerable sensors. Computational parallelism using graphics processing units (GPUs) was employed to provide the computational power necessary for simulation and management of sensor networks and RSO catalogues of an authentic scale and distribution. The final investigation considers the management of steerable sensors on an individual level. Aiming to increase system capacity and timeliness in particular, a method named dynamic steering is proposed for improving the autonomy of steerable sensors and enabling them to react quickly and intelligently when performing the reacquisition of RSOs. The GPU-based simulation developed throughout the previous topics provides a suitable platform on which to build such a sensor controller without compromising the benefits to cost, scalability and flexibility already offered through the use of steerable sensors.

This thesis presents a collection of techniques to enhance the capability, capacity and timeliness of catalogue maintenance systems. Via the exploitation of steerable sensors, the application of statistical signal processing techniques and use of modern computational hardware, sensor management strategies are proposed to enhance RSO catalogue maintenance for improved SSA.

1.2 Thesis Summary

The thesis is composed of six chapters. Chapters 2–6 document the research contributions.

Chapter 2 begins with an exploration of the current SSA environment via an analysis of publicly available RSO tracking data. The practicalities involved in maintaining surveillance of this environment is explored next resulting in a justification for maintenance of RSO catalogues. The differences between classical tracking scenarios and the surveillance of RSOs are highlighted. The chapter concludes with a discussion regarding contemporary statistical orbit determination techniques for updating RSO catalogues with recent measurements. The analysis of each topic reinforces the motives of the thesis and aids in justifying the approaches employed in subsequent chapters.

An original approach to simulating the SSA environment is developed in Chapter 3 named MATLAB Space Situational Awareness Simulation (MASSAS). Whilst many improvements and adaptations to MASSAS are detailed throughout the thesis, its core objectives are to

provide a flexible and rapid means of simulating and evaluating various strategies for catalogue maintenance. To ensure MASSAS is dependable for this role, the results of a prior investigation conducted on an established and proven third party simulation are replicated. Each simulation's results are compared to ensure consistent conclusions are drawn. A disparate sensor analysis is subsequently explored to contrast the relative performance between a traditional radar-based surveillance network and a network comprising of the state of the art in steerable sensors, a type of electro-optical sensor.

The results of Chapter 3's disparate sensor analysis identify two key aspects requiring consideration in a network reliant on electro-optical sensors. To maximise the benefits of employing electro-optical sensors, methods to improve the observability of RSO range and compensate for the constraints to optical-sensor availability are investigated throughout the following two chapters.

Chapter 4 details the development and implementation of a multi-sensor scheduling technique. The technique is used to better equip a sensor management system to observe the range of RSOs whilst employing a network of sensors incapable of directly measuring range. Through the use of measurement level sensor fusion, the technique aims to combine the simultaneous measurements of space sensors to achieve complementary measurement error characteristics to those achieved using measurements made by individual sensors. The technique is reliant on the appropriate distribution of sensors to ensure that the geometries are conducive for the simultaneous measurement of an RSO by multiple sensors. While the use of multi-sensor observations may provide superior orbital estimates for a specific RSO, the total number of RSOs that could have been observed per scheduling period will, in all likelihood, fall. Scheduling too many simultaneous measurements can consequently reduce the accuracy of the catalogue as a whole. A significant portion of the chapter therefore discusses the appropriate distribution of sensors about the globe and offers a method for controlling the ratio of simultaneous and individual measurements.

The surveillance network availability analysis performed in Chapter 5 requires a significant improvement in the computational efficiency of MASSAS to that originally proposed in Chapter 3. The chapter therefore begins with a detailed description of its augmentation with highly parallelised code for execution using a PC's GPU. Employing the parallel GPU architecture, MASSAS is used to determine sensor availability for a range of network configurations using the tracking data of approximately 15 000 RSOs. As a result of the analysis and to further demonstrate the suitability and capability of GPU computation for solving problems pertaining to SSA, the chapter concludes by detailing a scheduler that utilises long range visibility forecasting to improve catalogue accuracy. Using the proposed method, the scheduler is capable of prioritising the observation of RSOs that it predicts are about to

undergo a period of non-observability to the surveillance network.

Prior to Chapter 6, network level sensor management and the utilisation of MASSAS to investigate, simulate and verify various scenarios is discussed. Chapter 6 details the retasking of MASSAS as an intelligent sensor controller for sensor-level management. Incorporating the use of MASSAS's high performance parallel computation, a process named dynamic steering is proposed to accomplish real-time decision making at the sensor. Specifically designed for steerable sensors with a restrictive FOV, dynamic steering enables the autonomous reacquisition of targets, even when the probability of detecting the target is very low. This feature is intended to reduce the regularity at which RSOs require observation, thereby increasing system capacity. The process employs parallelism to permit the efficient evaluation and update of high fidelity probability density functions (p.d.f.s) with each observation. These high fidelity products may also be used to develop new or enhance existing capabilities for SSA. Furthermore, the level of sensor-autonomy that the dynamic steering process achieves increases the timeliness and versatility of the system. To demonstrate the practicality of employing dynamic steering to existing sensors, the chapter concludes with a detailed account of a successful field trial of an experimental implementation.

1.3 Original Contributions

The following details the areas in which the author believes this thesis makes a valuable and original contribution.

A previously unseen analysis and set of predictive models for determining the rate of reoccurrence of space objects transitioning a terrestrial sensor's surveillance volume is offered in Chapter 2. The outcomes are useful for understanding the challenges involved in maintaining surveillance of Earth orbit and are worthy of consideration whilst planning future space surveillance networks.

A disparate sensor analysis presented in Chapter 3 offers the results of the first study to compare the error characteristics of an RSO catalogue as the ratio of primary contributing sensors for catalogue maintenance is varied between favouring traditional radar sensors to the current state of the art in steerable sensors. The findings indicate that space surveillance networks can more capably maintain an RSO catalogue when using the measurements from steerable sensors in place of or to augment conventional radar measurements. These findings may lead to the wider adoption of steerable sensors in a primary surveillance role and the retirement or repurposing of current radar assets.

Chapter 4's multi-sensor covariance-based scheduling method is new and original. It may be used to improve the observability of RSO range whilst using a network of optical sensors

that are incapable of measuring this value directly. The proposed technique may be used to incorporate sensor-fusion in a surveillance system's scheduling solution to improve the observability of a target's state whilst maintaining regular surveillance of a larger catalogue of targets.

The findings of a previously unexplored area of analysis are offered in Chapter 5 involving the availability of sensing resources when optical sensors are employed as the primary sensors for observing the entire RSO catalogue. As a result of this analysis, an original scheduling method employing parallelised visibility forecasting is proposed. Via efficient long term prediction of catalogue-wide visibility, the system is capable of anticipating and compensating for lapses in visibility of catalogued objects. This process extends the scheduling solution's time horizon well beyond that currently employed by surveillance agencies. The technique maximises the timeliness of observations of elusive RSOs thereby minimising catalogue error and the likelihood of losing the object.

A unique approach to SSA research is demonstrated throughout the thesis that involves the utilisation of an SSA simulation system capable of quickly, flexibly and efficiently simulating an SSA environment on a realistic scale, on a single PC. Uniquely, this simulation capability addresses an important gap in SSA computational research methodologies, by providing a level of capability in between simplified analytic analysis and the utilisation of supercomputing in high performance computing centres. It achieves this level of computation via the use of General-purpose Computing on Graphics Processing Units (GPGPU). Researchers of SSA adopting this methodological approach may broaden the scope of their current research or it may provide new avenues of investigation.

Chapter 6 details the development of a steerable sensor controller that employs a novel technique named dynamic steering. The technique enhances sensor automation when reacquiring catalogued targets to improve the reaction time and increase the system capacity of space surveillance systems. Dynamic steering may be more generally applied to improve tracking performance in surveillance applications in which probability of detection by a steerable sensor is anticipated to be low. Adoption of the technique in existing space surveillance systems may lead to further improvements in SSA capabilities particularly concerning automatic refinement of tracking data and generation of high-fidelity RSO catalogues for improved SSA. These capabilities would provide a foundation on which a dynamic sensor scheduling system may be devised that autonomously adapts in real time as observations are conducted and assessed.

2

An Intuition for SSA

The intention of this Chapter is to rapidly instil in the reader an intuition for the broad and intricate area of research to which this thesis ascribes. Specifically, this chapter discusses three fundamental topics of concern to any party desiring, maintaining or researching SSA. Section 2.1 begins the Chapter with an overview of the man-made environment surrounding the Earth. This topic is followed by a discussion in Section 2.2 concerning the practicalities involved in maintaining surveillance of this environment. The chapter is concluded by Section 2.3 which provides a brief overview of prominent tracking techniques that have the capacity to ingest raw observations made by space surveillance sensors and produce management systems capable of providing its operators the means on which to build a level of SSA.

A reader already familiar with the near-Earth space environment, current methods for maintaining SSA and statistical tracking techniques might know or intuit some of this chapter's content. The intention is nonetheless to introduce and analyse these topics in such as way, as to set the appropriate contexts, provide prerequisite knowledge and supply justifications for the approach taken throughout the thesis.

2.1 The SSA Environment

Johannes Kepler's and Sir Isaac Newton's famous 17$^{\text{th}}$ century works on planetary motion describe the orbits of planets with an accuracy that was largely unsurpassed until the 20$^{\text{th}}$ century [9]. Kepler described the kinematics of planetary motion using the following three laws:

1. The orbit of each planet is an ellipse with the Sun at one focus.
2. The line joining the planet to the Sun sweeps out equal areas in equal times.
3. The square of a planet's orbital period is proportional to the cube of its orbit's semi-major axis.

Newton successfully described the dynamics of this motion using his law of gravitation

$$\mathbf{F}_g = \frac{-Gm_1m_2}{|\mathbf{r}|^2}\left(\frac{\mathbf{r}}{|\mathbf{r}|}\right) \tag{2.1}$$

where $\mathbf{F}_g$ is a vector describing the gravitational force attracting the second body toward the first, G is the universal gravitational constant equalling approximately $6.674\,28\times10^{-11}$ $\text{m}^3\text{kg}^{-1}\text{s}^{-2}$, m_1 and m_2 are the respective masses of the bodies and $\mathbf{r}$ is a vector pointing from the first body's centre of mass to the second body's centre of mass. Assigning an index to each body is arbitrary as the force between the bodies is equal in magnitude but opposite in direction.

More broadly, these early works of celestial mechanics describe the orbital motion of two bodies caused by a mutual attractive force as a result of each body exhibiting an isotropic gravitational field. As such, the ideal two-body orbit described by Kepler and Newton is useful for describing planetary motion, but it is also useful for describing the motion of RSOs about the Earth.

The two-body orbit cannot be solely relied upon to accurately describe an RSO's motion, as the two-body orbit assumptions would be an oversimplification of the many and often evolving forces that influence an RSO's orbital dynamics. The Earth's gravitational force will, in general, be the strongest force acting on an RSO, but it is not isotropic and is only one of many forces experienced by RSOs [9, 26]. Each RSO will be subject to a range of secular, periodic and impulsive forces due to influences such as the Earth's asphericity, multiple—albeit weaker—gravitational fields, solar radiation pressure (SRP), electro-magnetic fields and aerodynamic pressures [38]. These influences will cause an RSO's orbital trajectory to be perturbed from the fundamental two-body orbit with time. Describing and analysing an

RSO's two-body orbit nonetheless serves as a useful starting point to obtain a first order appreciation of an RSO's instantaneous Keplerian, or sometimes referred to as 'osculating', orbit profile.

Kepler's third law can be used to obtain the relation between an RSO's orbital period and orbital radius. The law can be rewritten mathematically as

$$\mathcal{T}^2 \propto a^3 \tag{2.2}$$

where $\mathcal{T}$ is the orbital period and a is the length of the semi-major axis of the orbit's elliptical locus or trajectory. This proportionality is appropriately scaled for RSOs by their orbital angular frequency ω_{RSO} and the Earth's gravitational parameter $\mu_\oplus$ such that

$$\omega_{\text{RSO}}^{-2} = \frac{a_{\text{RSO}}^3}{\mu_\oplus}. \tag{2.3}$$

Earth's gravitational parameter $\mu_\oplus$ is the product of the universal gravitational constant G and the Earth's mass $m_\oplus$ where $G.m_\oplus = \mu_\oplus$ and is approximately 398 600.441 8 km^3s^{-2} [9]. Equation (2.3) can subsequently be rewritten to conform to (2.2), resulting in

$$\mathcal{T}_{\text{RSO}} = 2\pi\sqrt{\frac{a_{\text{RSO}}^3}{\mu_\oplus}}. \tag{2.4}$$

A plot the relation described by (2.4) is presented in Fig. 2.1. Upon consideration of the Earth's mean radius of approximately 6378.1 km and the approximate 23.93 hour length of a sidereal day—the time the Earth takes to make a full rotation on its axis in inertial space—two important features of (2.4) are revealed. Even if an RSO has a perfectly circular orbit, without continuous thrust, its orbital period can not be smaller than approximately 1.4 hours as it would be prevented from completing an orbit due to severe atmospheric drag and/or collision with the Earth's surface. Additionally, because (2.4) increases monotonically, if an RSO's orbit is circular, prograde—orbiting in the same direction that the Earth rotates—and its semi-major axis length is approximately 42 164 km, the RSO will reside in a unique region of space that permits an orbit that is synchronised with the Earth's rotation. Such an orbit is called a geosynchronous orbit (GEO) [9, 11, 26]. These fundamental physical characteristics of Earth-orbit govern and bound the utility, and therefore popularity, of certain orbital regimes. This in turn has caused certain regions about the Earth to become more densely populated by RSOs than others.

Fig. 2.2 displays two histograms detailing the distribution of approximately 16 000 unclassified RSOs that were tracked by USSTRATCOM throughout December 2013 and made

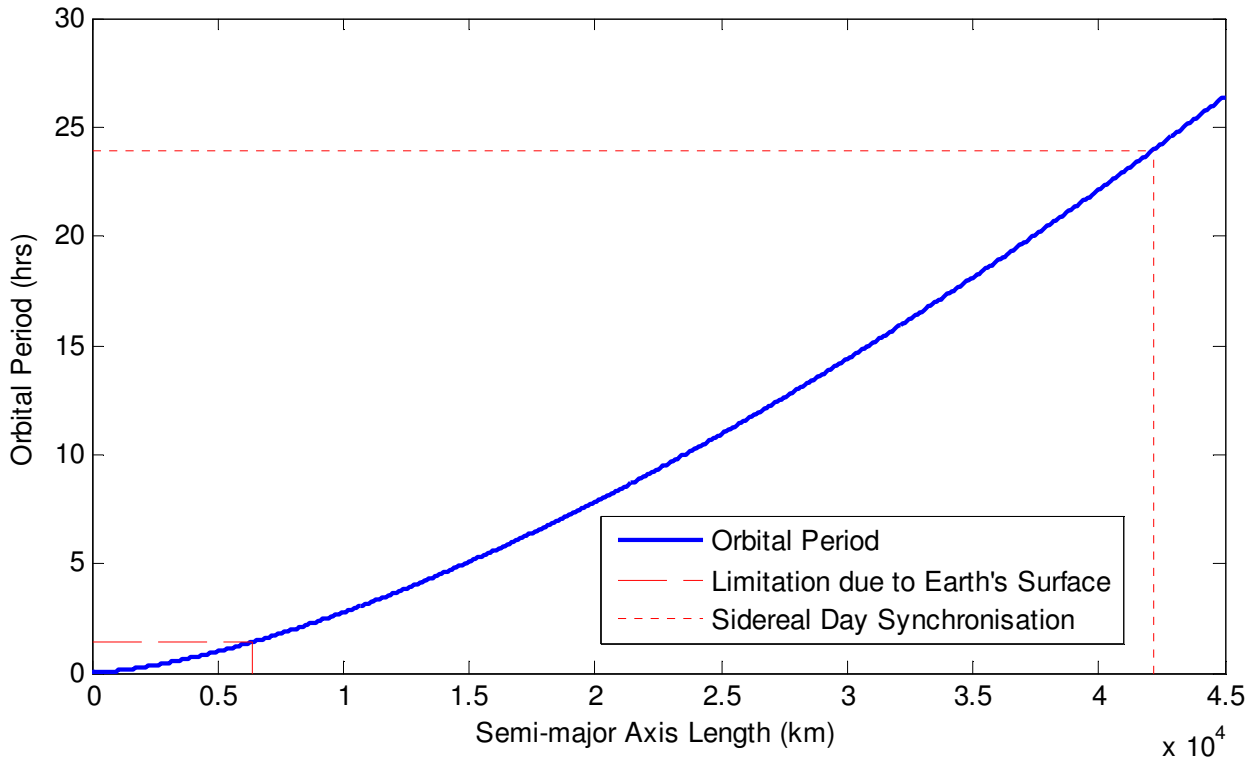

Figure 2.1: A plot of the relation between an RSO's semi-major axis length and orbital period.

publicly available on the Space-Track website [39]. The first histogram displays the frequency of objects with respect to their mean orbital period and the latter the mean length of the semi-major axis of their respective osculating orbits. Each histogram contains a similarly shaped but scaled version of the same distribution in agreement with Kepler's third law as described by (2.2).

Three distinct groups of RSOs are observed in each histogram of Fig. 2.2. The first group is by far the largest group and is located in low Earth orbit (LEO). It comprises approximately 80% of the unclassified catalogue. This group includes such objects as launch debris, fragmentation debris, Earth-observation satellites, some communication satellites such as the Iridium constellation as well as the well-known Hubble Space Telescope (HST) and International Space Station (ISS) [1, 11, 40]. The region is favoured for its proximity to the surface of the Earth resulting in lower launch costs, smaller power requirements for communication back to Earth and improved resolution for Earth observation [38]. LEO is also a useful location for space vehicles headed for higher orbits, or beyond, to discard excess rocket mass prior to and/or after a transfer-orbit burn which sends the space vehicle into an intermediary orbit that intersects with the desired orbit for its mission [8, 38]. The limits of

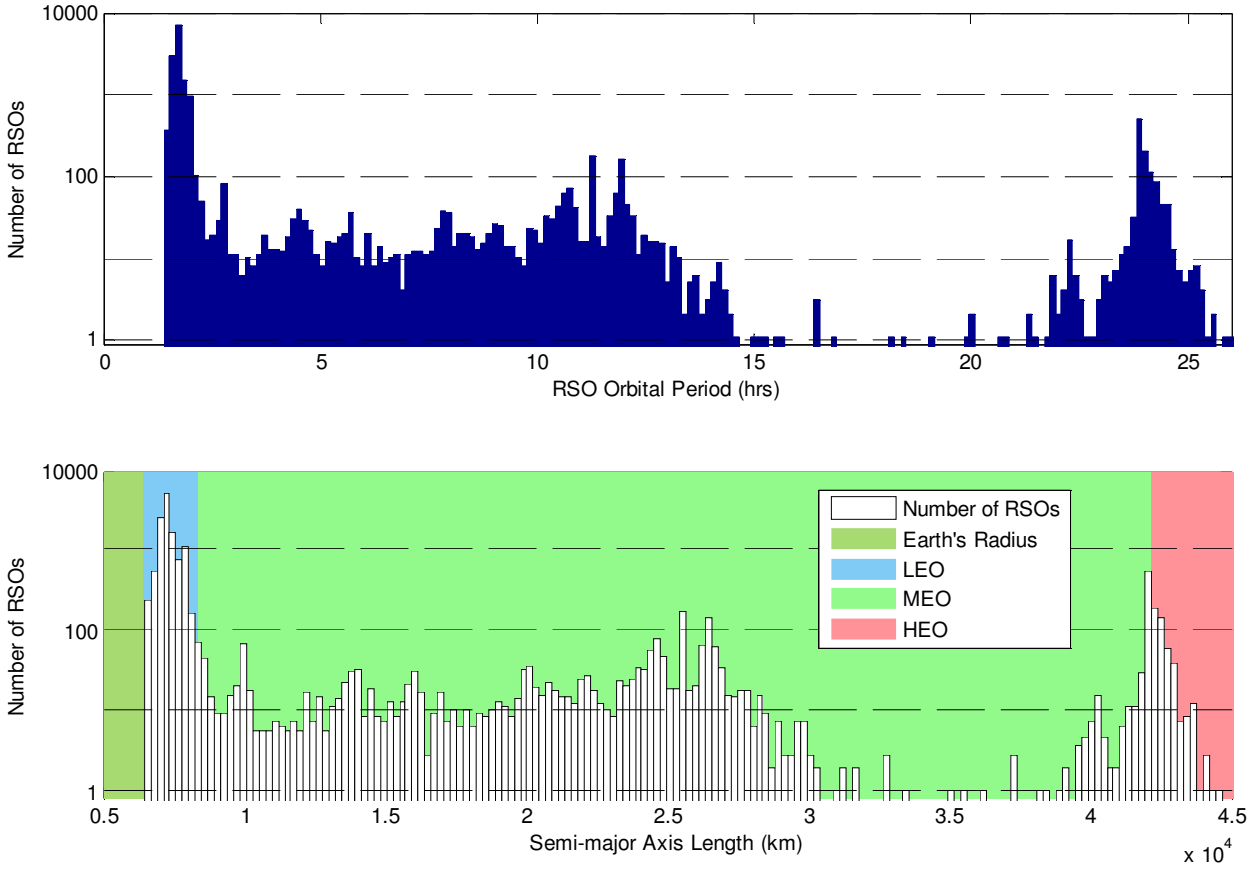

Figure 2.2: Histograms of USSTRATCOM's unclassified catalogue of objects showing the frequency of orbital periods and semi-major axis lengths.

LEO are defined differently in many different texts. However, it is generally accepted and will be defined henceforth to reside within the region between 100 km and 2 000 km in *altitude* above the Earth. Fig. 2.3 contains a plot of the positions, of the objects represented in Fig. 2.2. Their positions are plotted with respect to the Earth centred Earth fixed (ECEF) International Terrestrial Reference Frame (ITRF) at approximately 05:00 14th December 2013 UTC. The previously identified large LEO group can be observed in Fig. 2.3 as the dense ball of objects in the centre of the image. The ball-like shape of this group is a result of the diverse range of orbital inclinations adopted by LEO satellites and debris. Keplerian orbital elements, including orbital inclination, are discussed in detail in Appendix A.

The second group is found at a little over 25 000 km in the middle of the region known as medium Earth orbit (MEO). MEO extends from the upper bounds of LEO at around 8 000 km to GEO as indicated by Fig. 2.1 and shown in Fig. 2.4. This second group contains many navigation satellites such as the US NAVSTAR satellites, which constitute the space segment of GPS [40]. Currently the Russian GLONASS satellites — the Russian equivalent

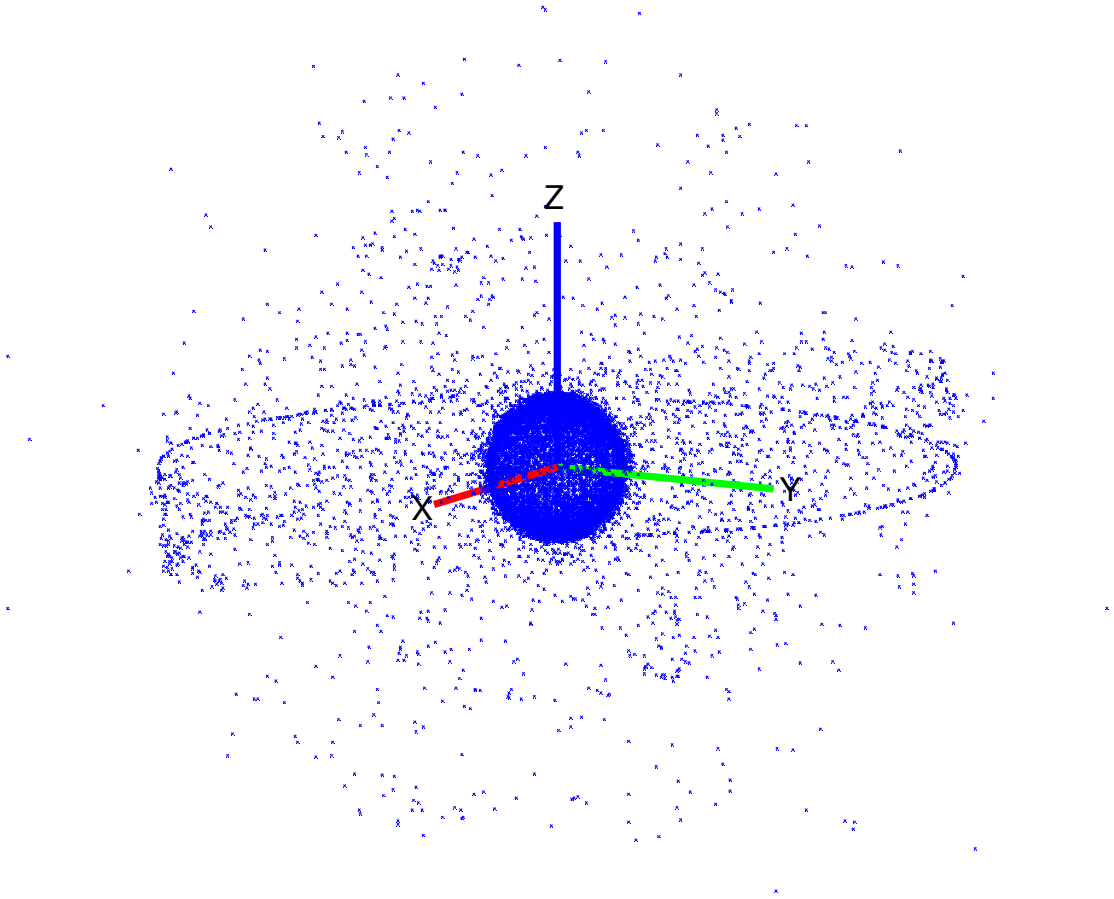

Figure 2.3: Positions of RSOs from USSTRATCOM's unclassified catalogue: as viewed from an oblique angle above the equatorial plane.

to GPS [40] — also reside in this region and are soon to be joined by the satellites of the fledgling European and Chinese versions named Galileo [41] and COMPASS [42] respectively. This region is favourable for navigation satellites as it strikes a balance between a number of factors, most notably the number of satellites and the availability of the system across the globe [1].

The final and highest of the previously identified groups of RSOs is located at GEO, which is also the lower bound of high Earth orbit (HEO). The upper bound of HEO is again loosely defined however a suitable definition for the purposes of this thesis would be any object that is orbiting within the Hill sphere of the Earth—the region in which the Earth's gravity is the dominant gravitational force—and not in lunar orbit. RSOs near the GEO boundary are mostly composed of telecommunication satellites and their debris [11, 40]. GEO orbits are favoured for telecommunications as satellites at GEO altitudes can maintain visibility of a specific region of the Earth throughout its entire orbit. Furthermore, if the satellite's

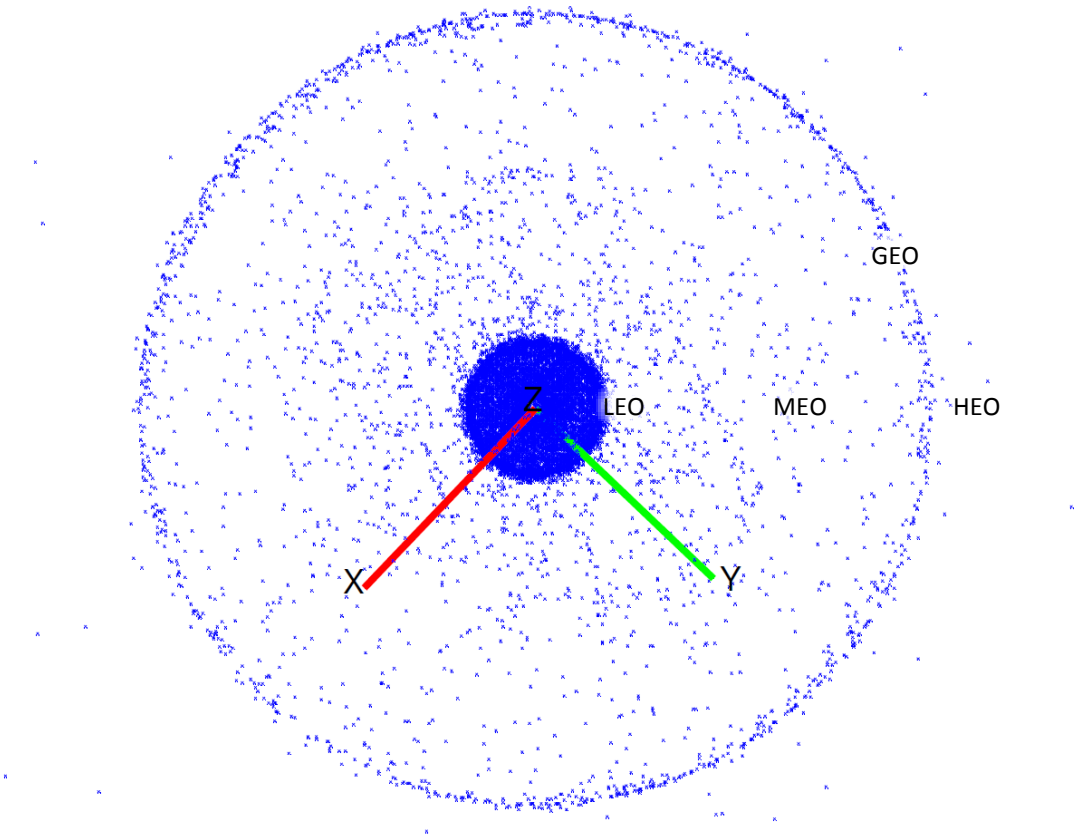

Figure 2.4: Positions of RSOs from USSTRATCOM's unclassified catalogue: as viewed from the north polar axis looking toward the Earth.

orbit is equatorial, it is said to be geostationary which enables non-steerable satellite dishes, used by low cost, high bandwidth terrestrial receivers, to be permanently aligned with the satellite. The broadcast of satellite television is a prevalent example of this concept. The belt of geostationary satellites has been labelled in Fig. 2.5.

Whilst most of the catalogue belongs to one of the three groups discussed so far, Fig. 2.1 and Figs. 2.3–2.5 show a number of objects inhabiting the space in between. These objects, in general, comprise of satellites undergoing orbital manoeuvring, debris whose orbits have been highly perturbed or have been released during orbital manoeuvres as well as missions requiring specialised orbital regimes [1, 11]. Some of these RSOs may undergo highly eccentric orbits that transition between a combination of LEO, MEO and HEO altitudes during a single orbit. An example of a highly eccentric orbital regime is the Russian 'Molniya orbit' which utilises the characteristics of Kepler's second law to achieve geosynchronous-like visibility for high latitudes [9]. Renderings of an example Molniya orbit may be viewed in Appendix B: Figs. B.1 and Figs. B.2. The large eccentricity and orbital inclination cause

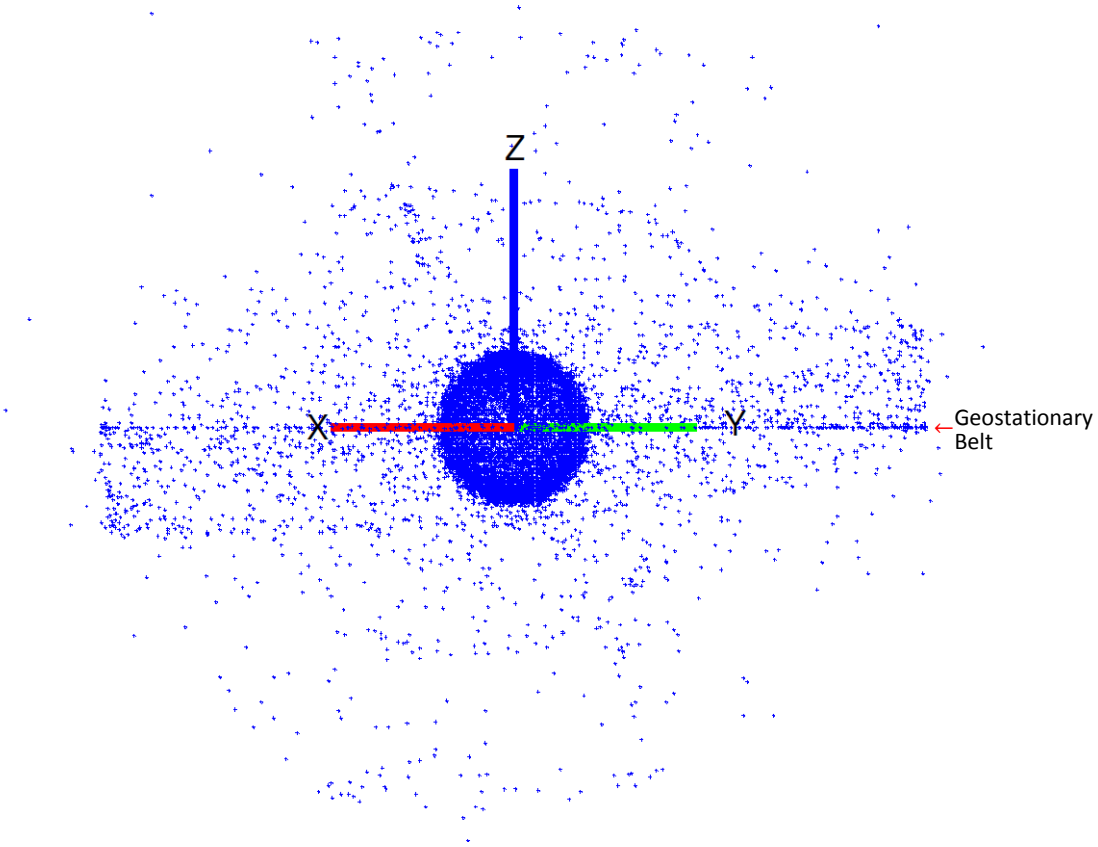

Figure 2.5: Positions of RSOs from USSTRATCOM's unclassified catalogue: as viewed from the equatorial plane.

a space-vehicle undergoing a Molniya orbit to 'linger' above high latitudes for long periods of time. Nonetheless this approach requires multiple satellites to achieve 100% system availability.

2.2 Surveillance of Earth Orbit

Performing surveillance for maintaining an RSO catalogue does not conform to the traditional tracking scenario [22,43] of one or more sensors regularly observing a finite surveillance volume, outside of which, targets are not of concern or are unknown. Instead, all targets inside and outside of the surveillance volume, which is constantly rotating through inertial space, are catalogued and of interest. Modern tracking techniques aim to achieve higher levels of awareness than traditional ad hoc tracking can provide, by retaining history, performing predictions and combining multiple sources of information [44, 45]. Predicting the arrival of catalogued targets within a sensor's surveillance volume is particularly useful for

space surveillance. Doing so enables sensor management systems to exploit opportunities for observing RSOs to maximise the accuracy of the orbital state estimates within the RSO catalogue [19].

Throughout the following sections, the rationale behind and implications of this form of surveillance will be investigated. The investigation will attempt to accomplish this objective by considering what would be necessary to implement an SNSS (sensor network for space surveillance) that has a surveillance volume that encompasses the entire region in which RSOs may reside, at all times. It is anticipated that the result will establish the motivation behind maintaining catalogues of RSOs and why it is important to improve catalogue refinement techniques. Trigonometric analysis of the geometry as well as first order approximations and assumptions will be utilised to obtain a low-fidelity, high-level appreciation for the intricacies involved in maintaining surveillance of Earth orbit. We begin by investigating how much of the Earth's surface is capable of observing an RSO at a predefined altitude.

2.2.1 Terrestrial Visibility of an RSO

Mathematically modelling a sensor's ability to observe an RSO is a complex task. For simplification appropriate to this chapter's level of analysis, it will be assumed that line of sight (LOS) is all that is necessary for a sensor to achieve visibility of an RSO. The practicalities and implications of this assumption will be discussed in more detail in Section 2.2.4's concluding remarks and is a topic of discussion in Chapter 3.

The LOS assumption means that the visibility of an RSO is dependant only on the relative geometry between sensor, Earth and RSO. To describe this geometry, we begin by defining a number of parameters as displayed in Fig. 2.6. A spherical Earth approximation is assumed thereby setting $r_{\oplus}$ to equal the mean radius of the Earth, r_{eff} is the distance from the centre of the Earth to the RSO, θ_{Vis} is the angle subtended by the region of the Earth's surface from which an RSO is visible and ϕ_{min} is the minimum elevation required for an observing sensor to 'see' the RSO. The variable ϕ_{min} is included to cater for the low elevation limitations that many terrestrial sensors require [9] to minimise atmospheric disturbance, avoid terrestrial occlusion and abide by physical pointing limitations. The trigonometry of Fig. 2.6 is solved using the equation

$$\theta_{\text{Vis}} = 2\left(\arccos\left(\frac{r_{\oplus}\cos(\phi_{\text{min}})}{r_{\text{eff}}}\right) - \phi_{\text{min}}\right). \tag{2.5}$$

Consequently, the solid angle Ω_{Vis} defining the portion of the Earth's surface permitting

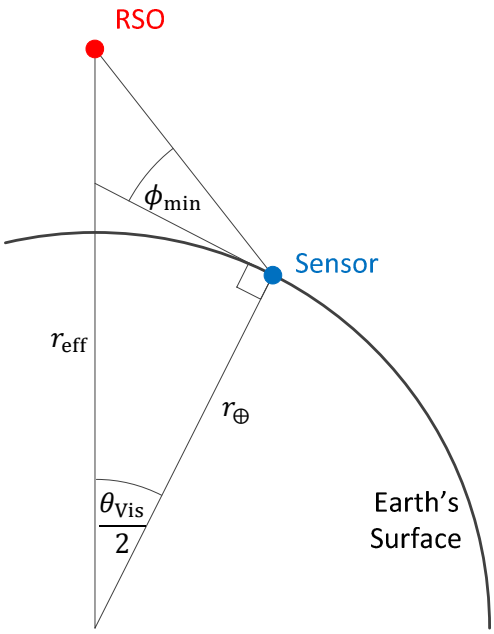

Figure 2.6: Geometrical constraints for RSO visibility by a terrestrial sensor.

visibility of the RSO can be found using,

$$\Omega_{\text{Vis}} = 4\pi \sin^2\left(\frac{\theta_{\text{Vis}}}{4}\right). \tag{2.6}$$

Fig. 2.7 presents a plot of (2.5) for common orbits, when $\phi_{\min}$ is set to a conservative 20°. As r_{eff} increases, the plot asymptotically approaches 140° as $140° = 180° - 2\phi_{\min}$. The plot also indicates that a 20° sensor limitation restricts geostationary—and therefore equatorial—satellites to be viewed by sensors at latitudes no higher than approximately 60°. In contrast, a LEO RSO whose r_{eff} is approximately 7 378 km, and has an equatorial orbit, would require a sensor to be located at a latitude no larger than 15° from the equator.

Whilst the Earth is not perfectly spherical, using an elliptical model such as the World Geodetic System - 1984 (WGS-84) to incorporate the Earth's largest component of asphericity, causes variations in θ_{Vis} of less than 0.4% depending on latitude. Therefore in spite of the spherical Earth approximation, Fig. 2.7 provides a highly accurate representation of the expected terrestrial visibility of an RSO located at a known altitude.

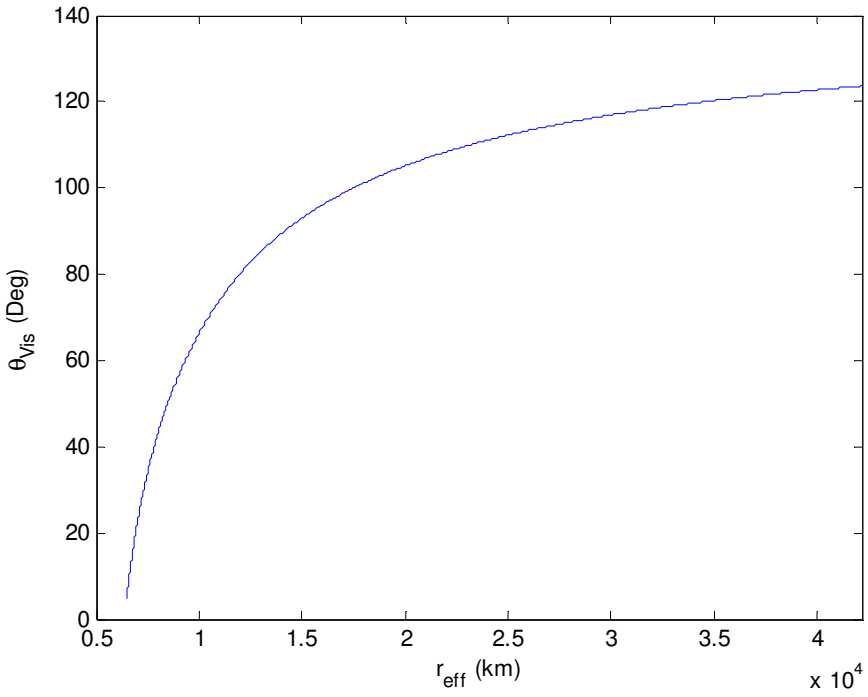

Figure 2.7: The angle subtended by the region of the Earth's surface that may observe an orbiting LEO → GEO RSO when limited by a minimum elevation $\phi_{\min} = 20°$.

2.2.2 Estimation of the Number of Sensors Required for Global Coverage

Agencies currently operating a sensor network for space surveillance (SNSS), such as USSTRATCOM's SSN [18], utilise approximately 20 to 30 primary sensors globally [31]. Whilst the exact number of sensors may vary due to deployable sensors and the ability to share observations, the distribution of sensors within these networks is highly dependent on geopolitical boundaries, international relations, utilisation of legacy systems, proximity to existing sensors, supportability of the site and funding for future development [18, 20, 31]. As will be demonstrated in Chapters 3 and 5, the type, number and distribution of sensors affects how regularly and effectively different regions of Earth-orbit can be observed [20]. This means that each SNSS is not equally capable of contributing to SSA.

Nonetheless, in the hopes of determining the practicalities involved, the objective of this section is to spend a brief time considering how one might distribute a hypothetical space surveillance network's sensors, if the aforementioned list of practical considerations were removed. Furthermore, we will begin by setting our hypothetical SNSS the task of observing all Earth-orbit—LEO, MEO and HEO—at all times. Theoretically, it would then be possible to track all objects in real time, or at least observe any object at any time. Such

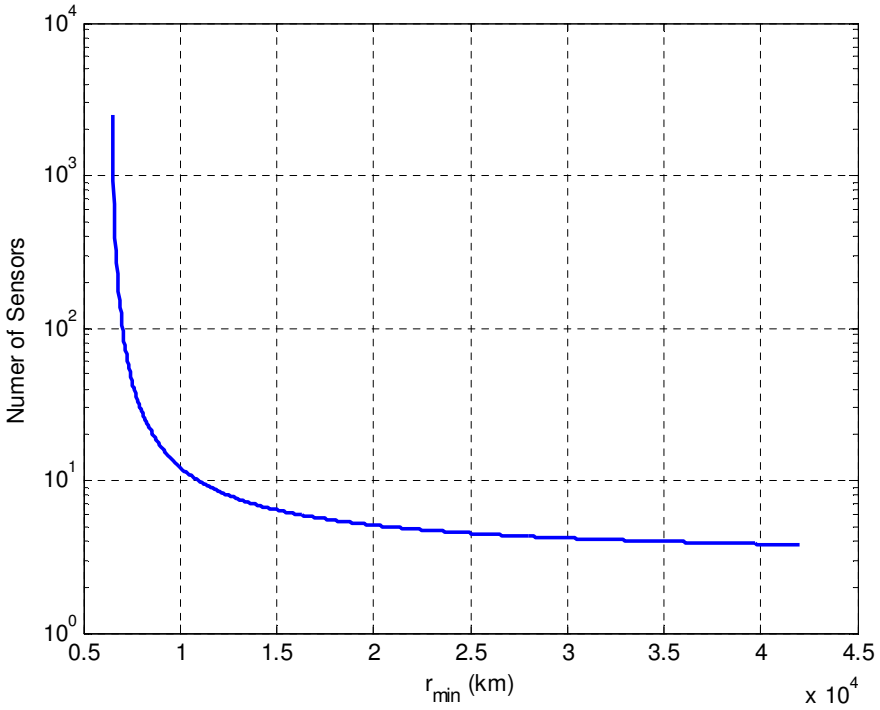

Figure 2.8: A plot of the relation between the number of sensors necessary for continuous coverage of Earth-orbit above a minimum distance of $r_{\min}$ from the centre of the Earth.

an implementation would describe a more conventional tracking scenario and the notion of making sporadic observations for catalogue maintenance loses its meaning.

Fig. 2.7 demonstrates that RSOs close to the surface of the Earth require sensors to be separated by smaller distances than RSOs at higher orbits, if a smooth hand-off from one sensor to the next is desired. As (2.6) increases monotonically, a minimum distance $r_{\min}$, can be chosen to obtain a value for the portion of the Earth's surface in which a sensor would need to be placed to observe all objects above that minimum distance. If this region is then divided by the surface area of the Earth, the theoretical minimum for the number of sensors necessary to observe all Earth orbit can be estimated. By adapting (2.6) accordingly, the minimum number of sensors n_s can be computed by means of

$$n_s = \sin^{-2}\left(\frac{\theta_{\text{Vis}}(r_{\min})}{4}\right) \tag{2.7}$$

The resulting values for an $r_{\min}$ between 100 km altitude and GEO is plotted in Fig. 2.8. Fig. 2.8 indicates that thousands of sensors are necessary to attain complete coverage of LEO, whilst only four sensors are necessary if gaps in coverage are acceptable below GEO.

Of course, in practicality, the surveillance volume of each sensor may not necessarily tessellate. Consequently overlapping regions, will warrant additional sensors. For example, if sensor surveillance volumes are modelled as spherical cones, a cross section of the spherical cone, parallel to local horizontal, produces a circle. As circles cannot tessellate with each other, we will borrow from the fundamental design principles of cellular networks [46] to achieve an efficient arrangement. As shown in Fig. 2.9, the sensors are arranged in a triangular formation. In this form, the limit of the surveillance volume of three adjacent sensor's would be designed to intersect at a radius of $r_{\min}$ from the centre of the Earth. The overlap would be small, but the consequent loss in coverage area means that additional sensors will be necessary above the theoretical minimum proposed in (2.7). The geometry described in

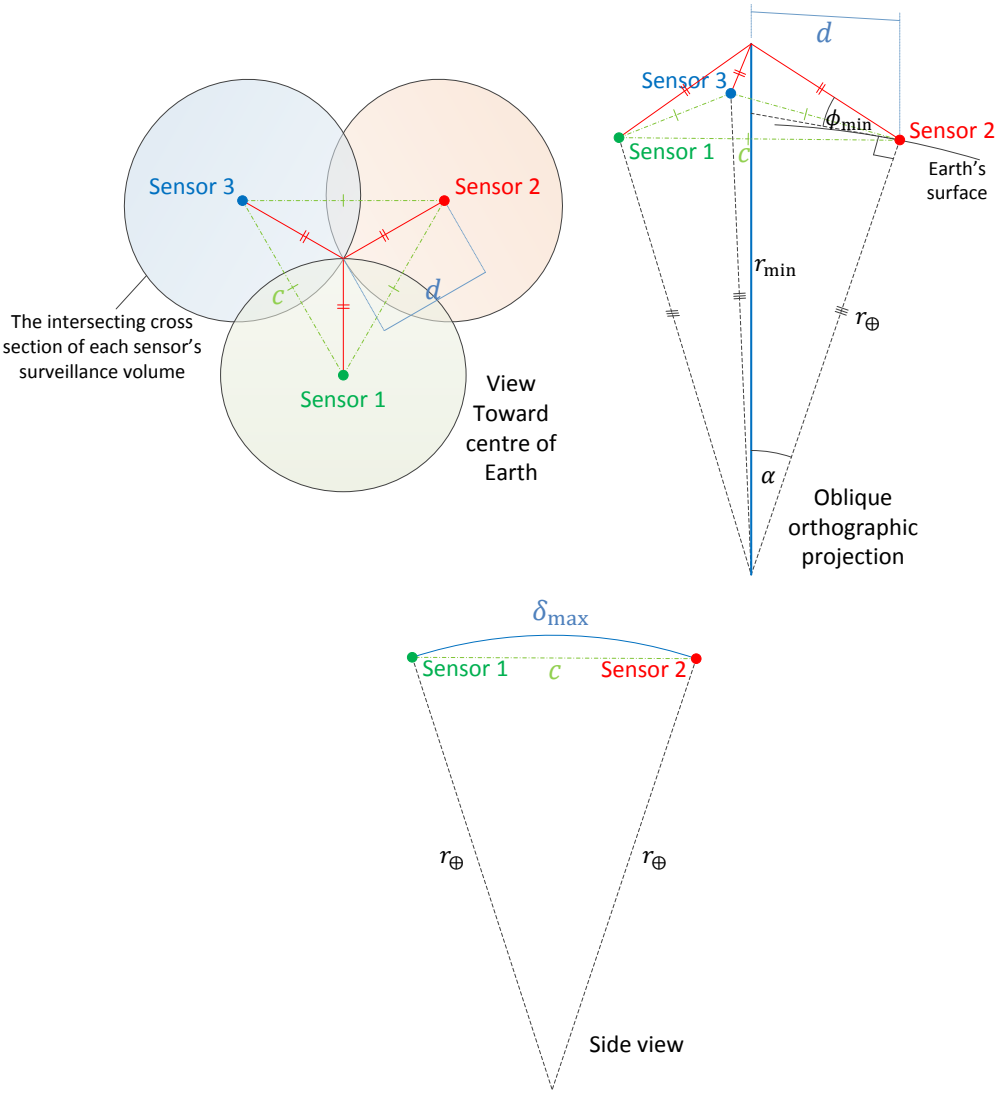

Figure 2.9: Cellular network-like configuration of space surveillance sensors for continuous global coverage.

Fig. 2.9 is solved using the equations

$$\delta_{\max} = 2r_{\oplus} \sin^{-1}\left(\frac{c}{2r_{\oplus}}\right) \tag{2.8}$$

$$c = \sqrt{3}d \tag{2.9}$$

$$d = r_{\oplus} \sin(\alpha) \tag{2.10}$$

$$\alpha = \frac{\pi}{2} - \phi_{\min} - \sin^{-1}\left(\frac{r_{\oplus} \cos(\phi_{\min})}{r_{\min}}\right) \tag{2.11}$$

where $\delta_{\max}$ is the maximum distance of separation between sensors in the triangular formation to achieve complete coverage above $r_{\min}$, α is the angle subtending the vector pointing from each sensor to the point of intersection between three adjacent surveillance volumes and c and d are intermediate linear dimensions as indicated in Fig. 2.9.

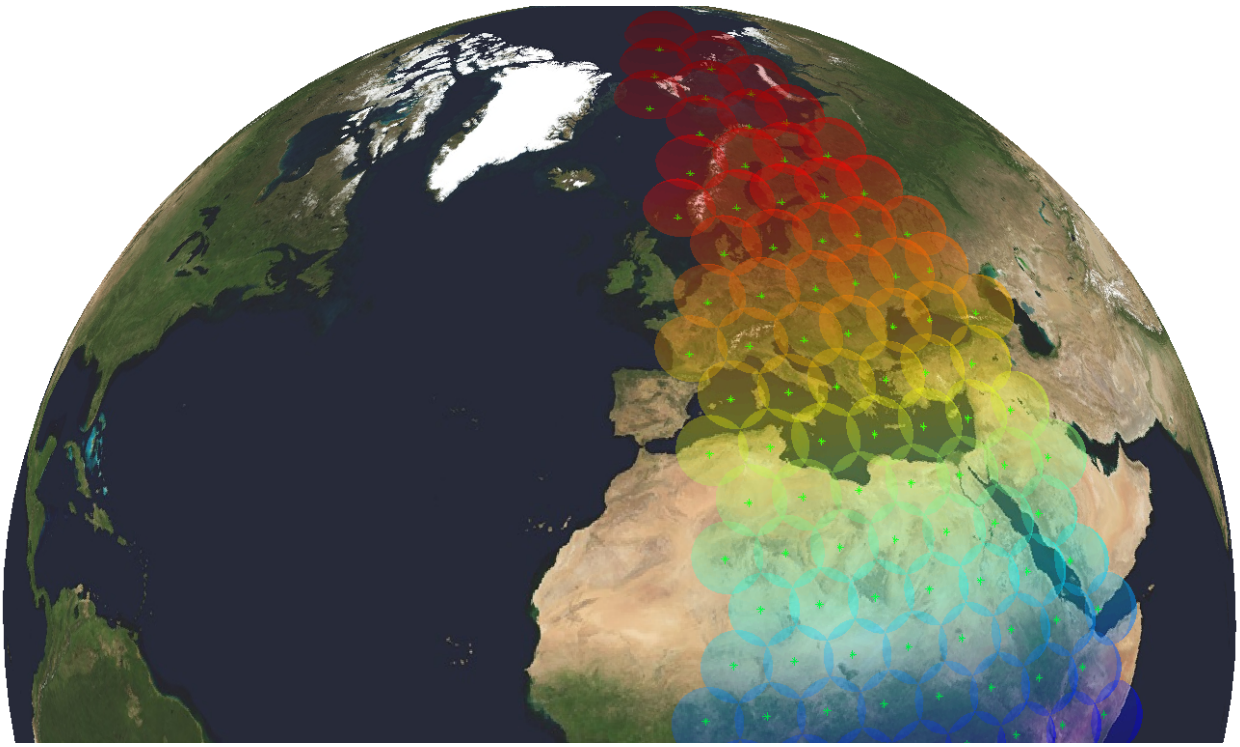

Figure 2.10: A visualisation of the sensor spacing necessary to achieve 100% surveillance coverage above 150 km altitude applied to $1/16^{th}$ of the Earth's surface.

Fig. 2.10* displays a visualisation of the sensor distribution strategy described in Fig. 2.9 when $r_{\min} = 150$ km altitude but is applied only to latitudes $0° - 90°$ N and longitudes $0° - 45°$ E. Note that whilst we assume the sensors have unlimited range, the surveillance

*The texture of the Earth's surface used in Fig. 2.10 and number of other images throughout the thesis was produced by NASA [47]

volumes have been represented by spherical cones with a radius that is equal to the range between a sensor and the point at which its surveillance volume intersects with the adjacent sensors' surveillance volumes. By visualising the network in this manner, the intersection of the surveillance volumes can be more clearly observed.

Fig. 2.10 demonstrates how impractical it would be to achieve 100% coverage for all objects orbiting higher than 150 km in altitude. In spite of the fact that only 6.25% of the proposed network is shown in Fig. 2.10, the image includes 96 sensors as $\delta_{\max}$ is approximately 646 km. Extrapolating this value indicates that over 1500 sensors are necessary to complete the SNSS. As anticipated, this value is slightly more than the 1171 sensors (2.7) predicts. To the best of the author's knowledge, the finances and level of international cooperation involved would make an SNSS of this magnitude the most ambitious global project ever undertaken. Most RSOs are passive debris incapable of manoeuvring and active satellites are manoeuvred as little as possible to converse their limited resources [10,38]. It is therefore possible to predict the location of most RSOs sufficiently for reacquisition many days after a previous observation [9]. The practicalities involved in implementing and operating the proposed SNSS are therefore extremely unlikely to outweigh the gains in SSA, achieved by 100% coverage. A practical space surveillance network's use of interpolation and prediction of RSO trajectories from intermittent observations will consequently result in a reduced level of catalogue accuracy and SSA. Investigating methods and techniques to minimise this loss is a core objective of this thesis and indeed a great deal of research pertaining to SSA.

Nevertheless, before we abandon the notion of devising a hypothetical high-availability SNSS, perhaps a compromise can be made to increase its practicality. If regular as opposed to continuous observation is an acceptable compromise, reconsidering Kepler's laws may offer a more practical solution. As Kepler's first law states that the Earth will be at one foci of an RSO's orbit, each RSO must either cross twice per orbit or traverse indefinitely, a plane that bisects the Earth. The two points of intersection between an RSO's orbit and such a plane are commonly referred to as the RSO's ascending and descending nodes, when the plane of reference is the equatorial plane [9]. Consequently, if sensors are placed along the equator, there are at least two opportunities for observations per orbit for each RSO. The geometry for this scenario is illustrated in Fig. 2.11. Provided a minimum range $r_{\min}$, Fig. 2.11 shows that the maximum spacing between sensors for obtaining 100% coverage above $r_{\min}$ along the equator, can be obtained by multiplying (2.5) with the Earth's radius to achieve

$$\delta_{\max} = r_{\oplus}\theta_{\text{Vis}}(r_{\min}), \tag{2.12}$$

the arc length or distance across the Earth's surface between sensors. The number of sensors

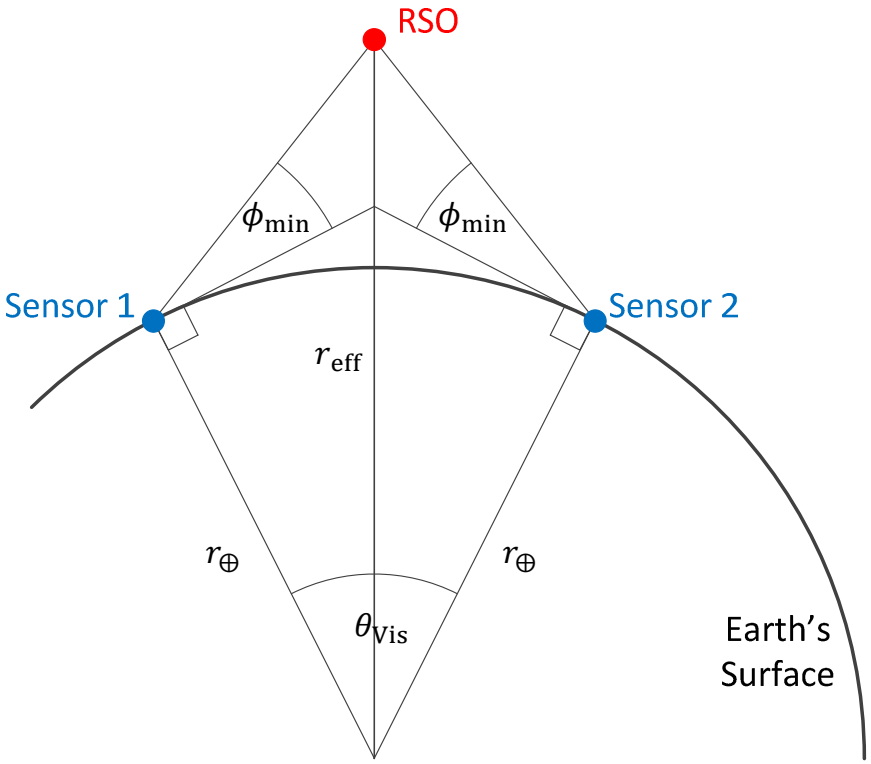

Figure 2.11: Maximum sensor spacing geometry for simultaneous observation of an RSO at a distance of r_{eff} from the centre of the Earth.

necessary to encircle the globe can subsequently be computed using the following equation

$$n_s = \frac{2\pi}{\theta_{\text{Vis}}(r_{\min})}, \tag{2.13}$$

assuming $\phi_{\min}$ is the same value for all sensors. According to (2.13) and using a WGS-84 equatorial radius of 6 378.1370 km, the resulting network of sensors would reduce the number of sensors from over 1500 to 54. Such an arrangement has been visualised in Fig. 2.12 by generating spherical cones, again range limited to observe the point of intersection, at regular intervals of $\delta_{\max} = 746$ km. The slight overlap, observed at the bottom of Fig. 2.12, is foretold by the non-integer result of $n_s = 53.731$ when $r_{\min} = 150$ km altitude is applied to (2.13).

Whilst this arrangement results in a reduction, by two orders of magnitude, in the number of sensors necessary to implement the SNSS, the network would still span a number of countries and large oceans. The practicalities involved are greatly improved however, this alternative SNSS would continue to be an enormous undertaking. The network, even with ideal sensors, continues to require almost double the number of sensors that are currently contributing to the SSN [9, 12]. Furthermore, such a network would only ever make observations when RSOs were near the equator, particularly for RSOs in LEO. Failing to observe

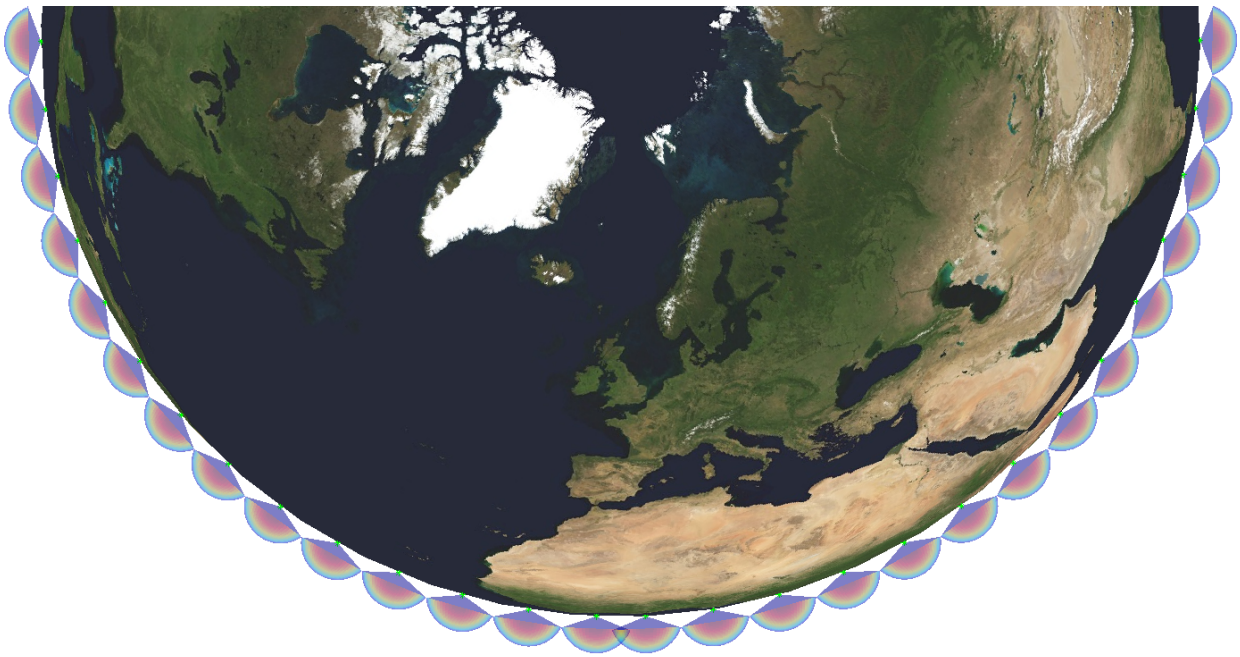

Figure 2.12: View from North pole looking South toward the proposed equatorial SNSS.

RSO's orbits with more diversity would likely lead to large uncertainties in the state of the RSO near the poles.

For this reason, discounting a breakthrough in space-sensor technology and a surge in international cooperation, terrestrial-based space surveillance will continue to be dependent on maintaining catalogues of objects and serendipitous opportunities for observation. As existing and future SNSS operators have not and may not ever have the freedom to locate their sensors wherever they choose, it is important to understand how often they can expect these serendipitous events to occur, and if the surveillance of certain types of RSOs is adversely affected when reliant on such events.

2.2.3 Longitudinal Revisit Rate of an RSO

This section's objective is to determine how often an RSO is expected to cross a line of longitude given that the longitudinal meridians rotate with the Earth, with respect to inertial space[†]. With this information, it is hoped that the regularity with which an RSO may visit a

[†]A coordinate system whose origin is located at the centre of the Earth, but remains oriented with respect to a celestial reference, experiences accelerations due to influences such as the Earth's orbit about the sun and the sun's orbit about the Milky-Way's galactic core. These coordinate frames are therefore non-inertial, however the effects are small. As such, Earth-fixed coordinates are commonly [9, 26, 48] referred to as quasi-inertial, pseudo-inertial or inertial frames. For simplicity, these frames will be named inertial throughout this thesis.

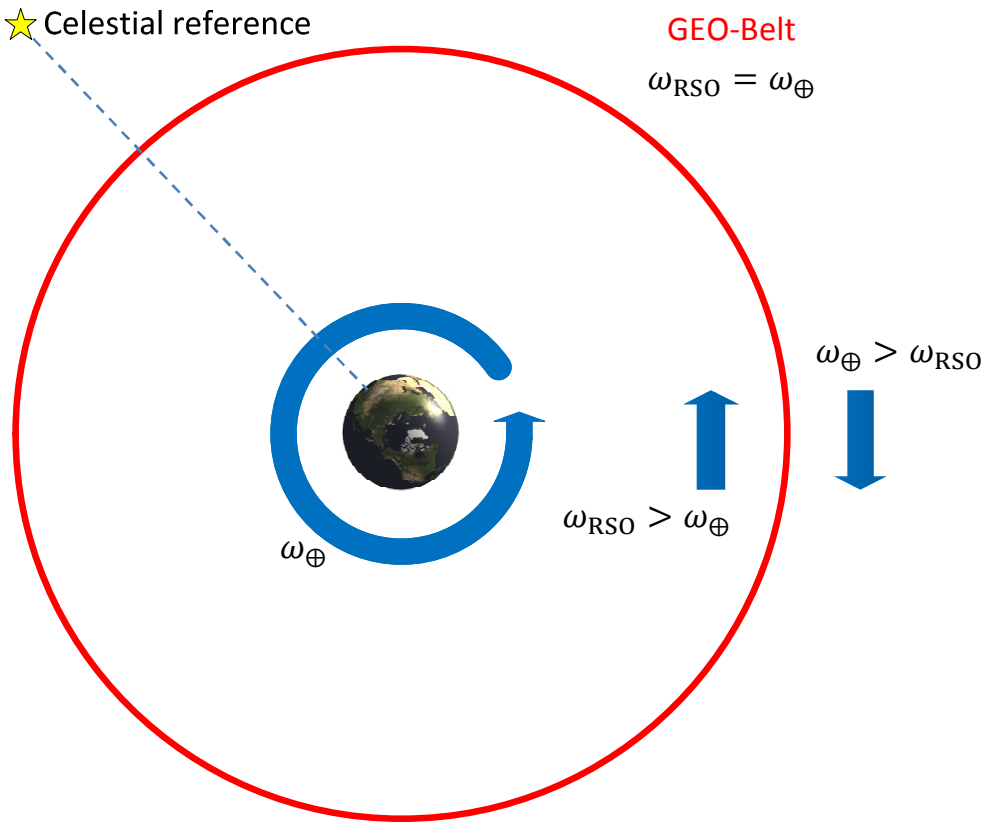

Figure 2.13: Prograde RSO orbital rates relative to Earth's surface, in an inertial reference frame.

terrestrial sensor can be established. Fig. 2.13 illustrates the concept of a longitudinal revisit rate, where $\omega_{\oplus}$ represents the angular frequency of Earth's sidereal rotation. Depending upon which angular frequency is greatest, $\omega_{\oplus}$ or ω_{RSO}, an RSO will either catch up to its original longitude, or its original longitude will catch up to it. The time between Earth-fixed longitudinal crossings by an RSO is therefore inversely proportional to the absolute difference between the Earth's and RSO's angular frequency of rotation and orbit respectively. The time between longitudinal crossings is obtained using

$$\mathcal{T}_{Long} = \begin{cases} 0 \text{ or } \infty & \omega_{\text{RSO}} = \omega_{\oplus} \\ \frac{2\pi}{|\omega_{\text{RSO}} - \omega_{\oplus}|} & \text{otherwise} \end{cases} \tag{2.14}$$

The implications of (2.14) become more apparent when the independent variable ω_{RSO} is substituted for an RSO's period of orbit resulting in

$$\mathcal{T}_{Long} = \begin{cases} 0 \text{ or } \infty & \mathcal{T}_{\text{RSO}} = \mathcal{T}_{\oplus} \\ \frac{2\pi}{|\frac{2\pi}{\mathcal{T}_{\text{RSO}}} - \omega_{\oplus}|} & \text{otherwise} \end{cases} \tag{2.15}$$

where $\mathcal{T}_\oplus$ is the length of a Sidereal day. A plot of this relation is shown in Fig. 2.14, where a

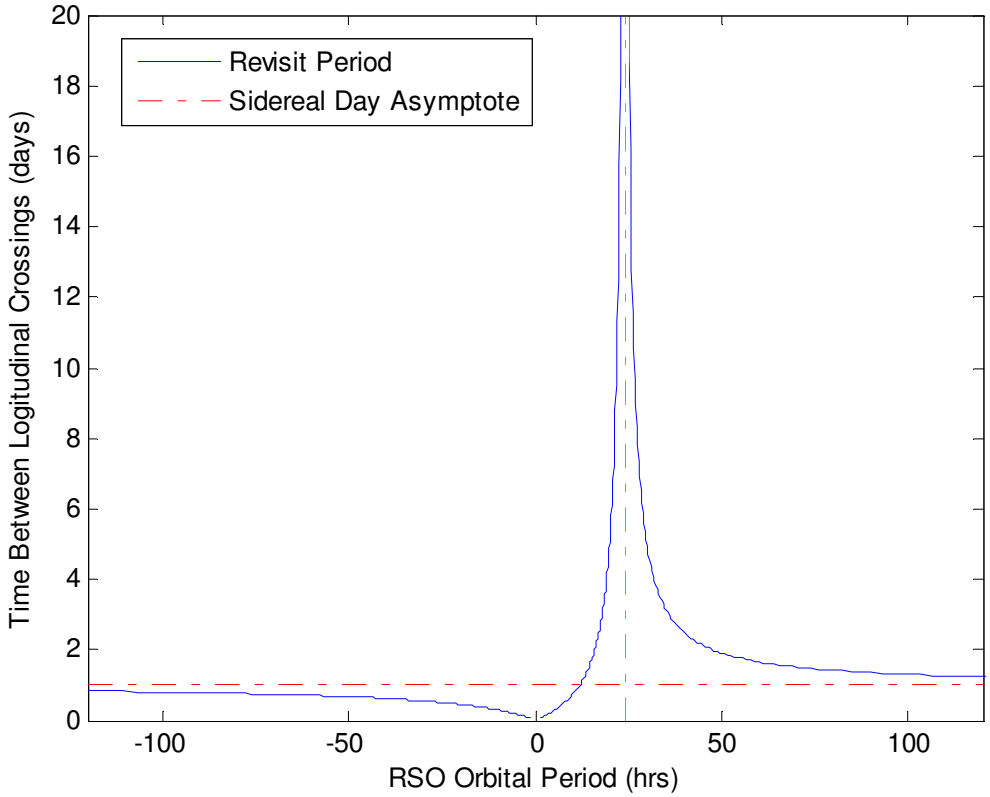

Figure 2.14: The time between longitudinal crossings verses RSO orbital periods.

negative orbital period represents a retrograde—against Earth's rotation—orbit. Horizontal and vertical asymptotes are observed.

The vertical asymptote is readily explained when $\frac{2\pi}{T_{\text{RSO}}} = \omega_{\text{RSO}}$ approaches $\omega_\oplus$, causing T_{Long} to approach infinity. This indicates that an object that orbits at precisely the speed the Earth spins on its axis, $\omega_{\text{RSO}} = \omega_\oplus$, will indefinitely reside at a line of longitude and never arrive at another. This result is of course synonymous with the previously discussed geosynchronous orbit. Nonetheless, as ω_{RSO} is the mean angular rate of the RSO, if the RSO's orbit is synchronised but eccentric, the object will instead oscillate about a specific line of longitude.

Fig. 2.14's horizontal asymptote is easier to interpret when (2.15) is inverted such that

$$T_{\text{RSO}} = \frac{1}{\frac{1}{T_{Long}} + \frac{\omega_\oplus}{2\pi}} \quad \text{where } T_{Long} > 0. \tag{2.16}$$

As $T_{Long} \to \infty$, the $\frac{1}{T_{Long}}$ term tends to zero, hence $T_{\text{RSO}} \to \frac{2\pi}{\omega_\oplus} = T_\oplus$, a sidereal day. This suggests that an object undergoing a retrograde or prograde orbit with a large orbital

period—and therefore large semi-major axis—will take approximately a sidereal day to appear at the same longitude. This result is unsurprising as the RSO's progress about the Earth would be small when compared to the Earth's own rotation.

Fig. 2.14 also indicates that an object orbiting the Earth with a high angular frequency, $|\omega_{\text{RSO}}| \gg \omega_{\oplus}$, will revisit a meridian of longitude many times in a single day. Consequently, LEO RSOs progress relatively quickly across the sky when observed by a terrestrial sensor.

2.2.4 Approximating the Expected Number of Passes per Surveillance Volume per Day

If RSOs were restricted to orbit along the equatorial plane in circular orbits, (2.14) and (2.5) could be used to accurately predict the frequency at which RSOs pass within a specific sensor's surveillance volume. Fig. 2.15, a histogram of the unclassified catalogue's eccentricities, indicates that the majority of RSOs do in fact have near-circular orbits. Using the equations,

$$r_p = a(1 - e) \tag{2.17}$$

$$r_a = a(1 + e), \tag{2.18}$$

Table 2.1 details the ratio between perigee, r_p, and apogee, r_a,—the smallest and largest distances between the Earth's and RSO's centre of mass respectively—for common eccentricities e, where a ratio of 100% would indicate a perfectly circular orbit. This data indicates that near-circular orbits are common, which supports the simplification that, in general, $r_{\text{eff}} \approx a$ throughout an RSO's orbit.

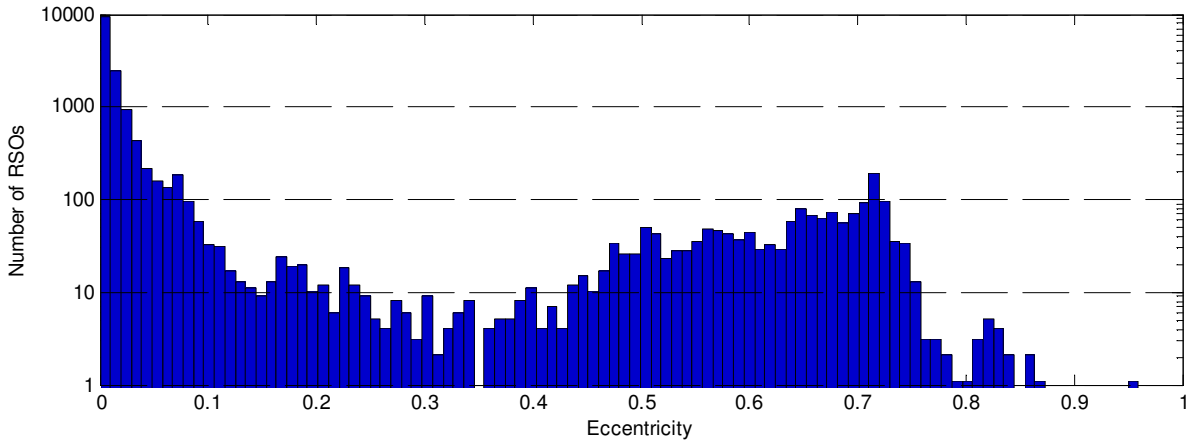

Figure 2.15: Frequency of RSO eccentricities within the unclassified catalogue.

Table 2.1: Unclassified catalogue eccentricity statistics as of 21 Jan 2014 06:30 UTC.

Eccentricity	Percentage of Catalogue	Ratio of r_p to r_a
< 0.1	88%	$> 81.8\%$
< 0.01	59%	$> 98.0\%$
< 0.001	12%	$> 99.8\%$

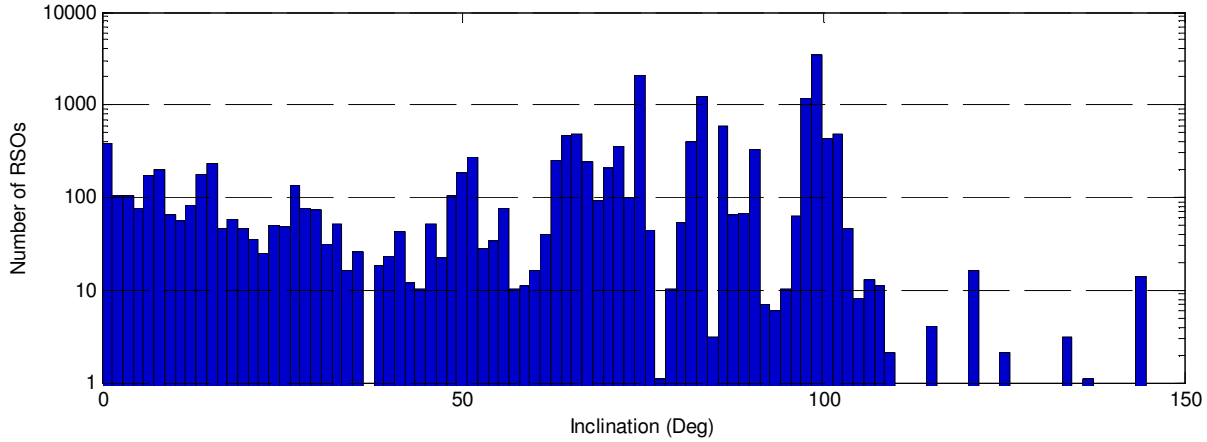

Figure 2.16: Frequency of RSO inclinations within the unclassified catalogue.

Nevertheless, Fig. 2.16 shows that most RSOs orbit at an inclination to the equatorial plane. Therefore making the assumption that RSOs have equatorial orbits would be detrimental to the analysis. Consideration of an RSO's inclination is therefore required.

As depicted in Fig. 2.17, an RSO's inclined orbit will cross the equatorial plane twice per orbit. Since it is assumed that most orbits are somewhat circular, the longitude at which this occurs will vary at a rate determined by T_{RSO} and $\omega_{\oplus}$. Therefore the amount of longitude that is lost or gained during equatorial crossings of an RSO's orbit can be approximated by

$$\Delta_{\text{Long}} \approx \frac{\mathcal{T}_{\text{RSO}}}{\mathcal{T}_{\oplus}} 2\pi. \tag{2.19}$$

By combining (2.19) with (2.5), the number of passes per day for a terrestrial sensor can also be approximated, if the following assumptions are made:

1. The sensor is located much closer to the equator than either pole.

2. Each RSO's orbit has sufficient inclination to be observable at the sensor's latitude.

3. The sensor's visibility limitations are completely described by a minimum elevation ϕ_{min} and LOS.

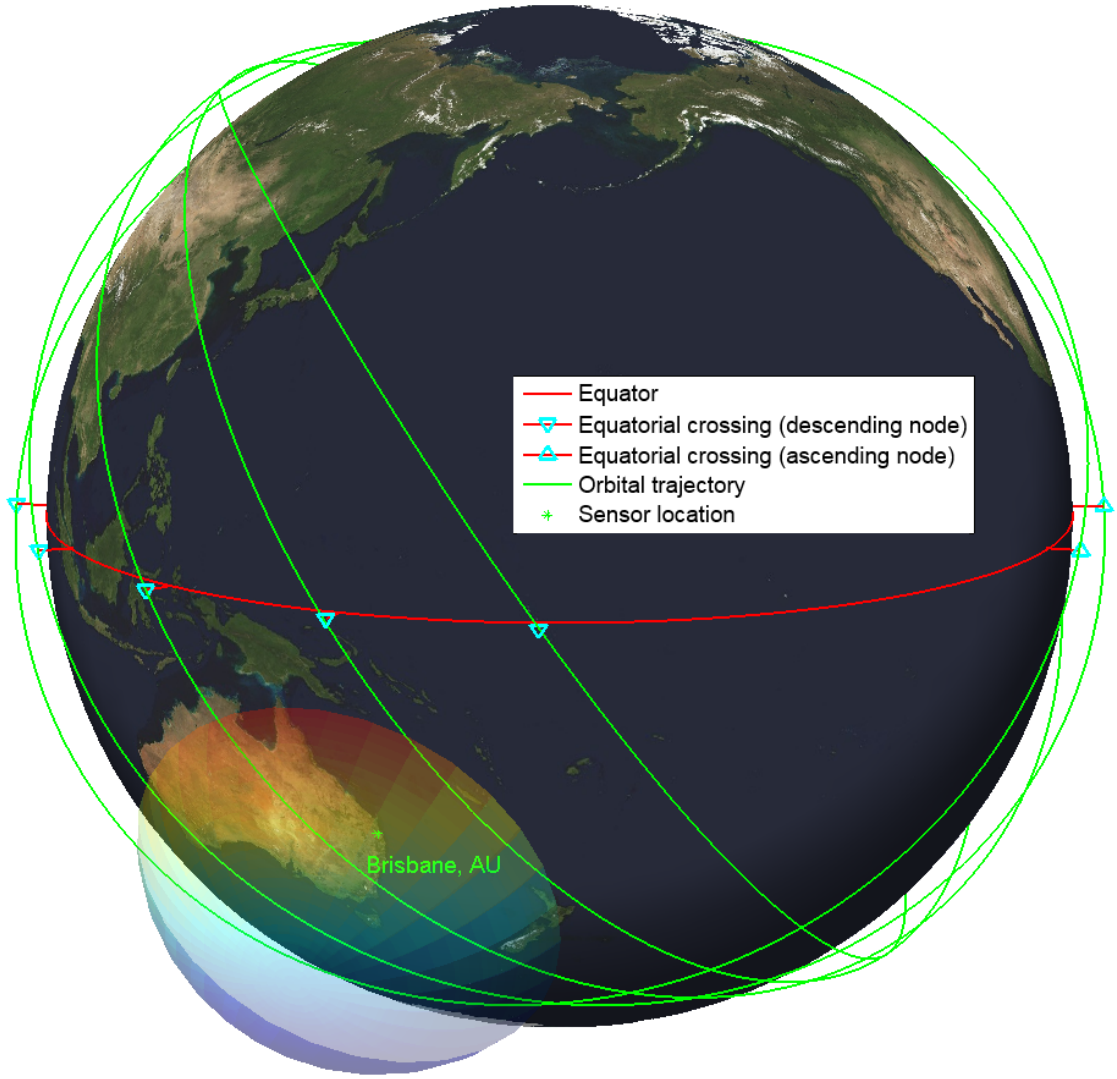

Figure 2.17: The eight hour progression of the ISS's ascending and descending node with respect to the Earth's surface and a fictional sensor's surveillance volume.

4. An RSO with an equatorial orbit is uncommon other than around the geosynchronous belt.

The reliability of assumption 4 is illustrated by Fig. 2.18 in which a histogram of the inclination of objects with a semi-major axis length greater than 27 000 km have been superimposed on top of the histogram, of all objects, previously shown in Fig. 2.16. By observing the deficit, it can be seen that the majority of low-inclination objects are indeed the near-GEO objects.

Using the aforementioned assumptions, the number of passes of an RSO through a surveillance volume per sidereal day can be approximated by obtaining the ratio between the range

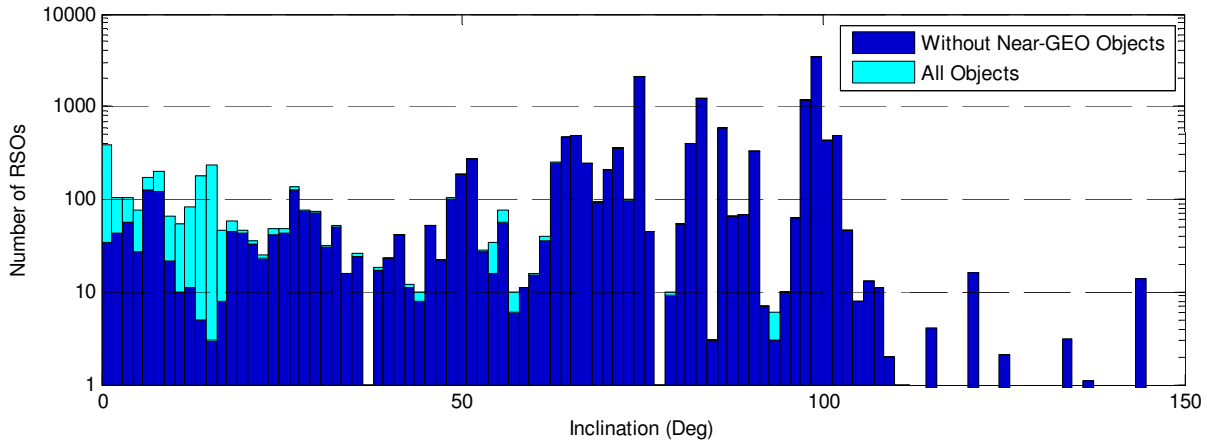

Figure 2.18: The superimposed histograms of RSO inclinations within the unclassified catalogue when including and excluding near-GEO objects.

of longitudes visible from a specific altitude divided by the longitudinal revisit rate, as estimated by (2.5) and (2.19) respectively, and doubling the result to account for RSO-visibility on both sides of the globe. The resulting equation is described by

$$n_{\mathrm{p}} \approx 2\frac{\theta_{\mathrm{Vis}}}{\Delta_{\mathrm{Long}}} = \frac{\mathcal{T}_{\oplus}\theta_{\mathrm{Vis}}}{\pi\mathcal{T}_{\mathrm{RSO}}}, \tag{2.20}$$

where, for a single sensor, n_{p} is the expected number of passes per day by an RSO with an orbital period of $\mathcal{T}_{\mathrm{RSO}}$. Furthermore, as we are only considering Earth orbiting bodies and $r_{\mathrm{eff}} \approx a$, for consistency (2.5)'s independent variable can be switched from $\theta_{\mathrm{Vis}}(r_{\mathrm{eff}})$ to $\theta_{\mathrm{Vis}}(\mathcal{T}_{\mathrm{RSO}})$ by rearranging Kepler's relation between semi-major axis and orbital period (2.4) such that

$$r_{\mathrm{eff}} \approx a_{\mathrm{RSO}} = \sqrt[3]{\mu_{\oplus}\left(\frac{\mathcal{T}_{\mathrm{RSO}}}{2\pi}\right)^2}. \tag{2.21}$$

Because it is assumed that only GEO-like objects will have low-inclination as well as equatorial orbits and because their orbital periods are approximately a sidereal day, it will be also be assumed that GEO-like objects will be seen either once or never during a day depending on their initial visibility from a terrestrial sensor. The resulting equation for obtaining an estimate of the number of passes of an RSO through a sensor's surveillance volume per day is consequently approximated using

$$n_{\mathrm{p}} \approx \begin{cases} 1 \text{ or } 0 & \mathcal{T}_{\mathrm{RSO}} \approx \mathcal{T}_{\oplus} \\ \frac{\mathcal{T}_{\oplus}\theta_{Vis}(\mathcal{T}_{\mathrm{RSO}})}{\pi\mathcal{T}_{RSO}} & \text{Otherwise} \end{cases}. \tag{2.22}$$

Fig. 2.19 displays a plot of the relation described by (2.22) for orbits up to and including

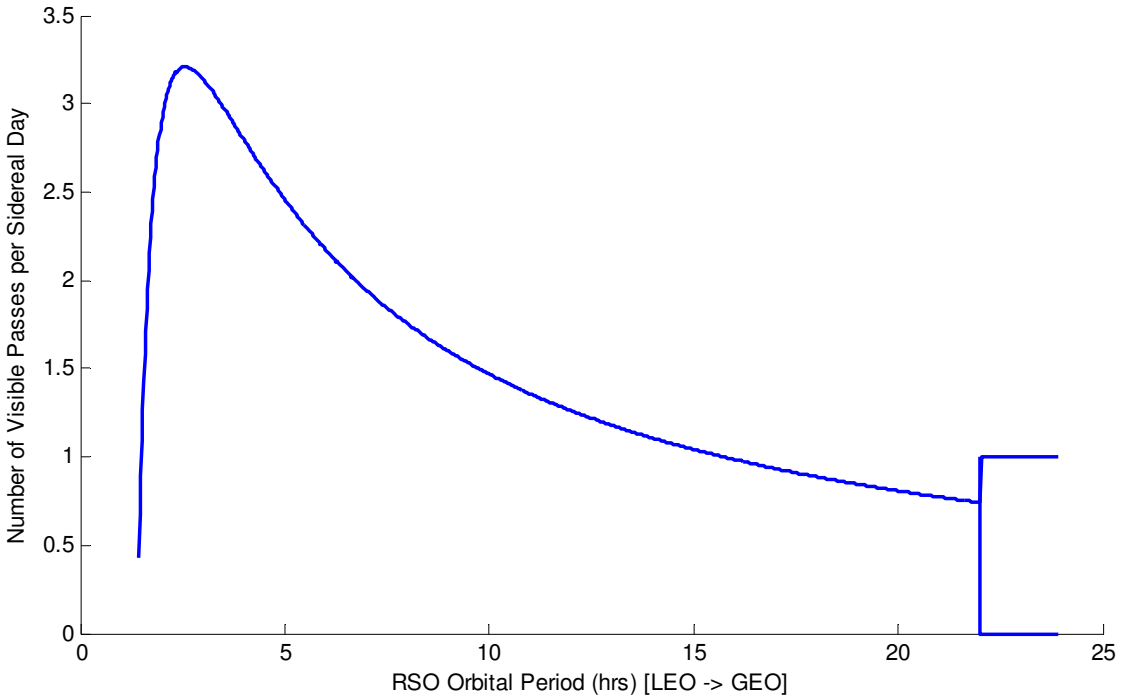

Figure 2.19: The expected number of passes of an RSO through a sensor's surveillance volume per day, for LEO → GEO orbits, with a minimum sensor elevation of 20°.

GEO. Fig. 2.19 exhibits a peak in pass regularity when $\mathcal{T}_{\text{RSO}}$ is approximately 2.6 hours, an equivalent average altitude of approximately 3200 km. This indicates that RSOs that exhibit this characteristic will be observable more regularly than all other objects. Furthermore, the number of passes per sidereal day drops much faster for objects with orbital periods shorter than 2.6 hours, than for longer orbital periods.

This information tells us that some LEO objects will be observed less often than objects in the low-to-middle MEO bound. Furthermore, objects in high-MEO and near-GEO orbits will be similarly challenging to regularly observe if sensors are not evenly distributed about the globe to watch these objects as they slowly progress across the sky. Nonetheless, these methods are only true of active sensing techniques with limitless range. So in reality (2.22) is somewhat optimistic and would in all likelihood require windowing to include a sensors' minimum and maximum range limitations. In addition, seasonal and diurnal scaling may be necessary for optical sensors as visibility is generally unattainable during daylight [49]. A sensor's non-zero latitude would also entail modification of Δ_{Long} to account for reduced motion of the sensor about the polar axis during an RSO's orbit.

2.2.5 Validation of Results

Whilst there are many assumptions in play, and (2.22) is far from a high fidelity predictor, it is hoped that it details a fundamental trend in RSO observation prediction. To achieve any level of validation of this result requires simulation and/or experimentation to test how closely the proposed relation fits simulated or real data. Common methods for performing such a validation in the field of SSA research resort to field trials, commercial software and when warranted, due to the number of RSOs, supercomputing [50–52]. Such methods require time and finances on a scale that would have inhibited the research detailed in this thesis, had an alternative not been sought. To produce meaningful outcomes in a timely and judicious manner, a highly adaptable simulation environment was created and developed for this research. Chapter 3 and Chapter 5 describe in detail the key elements of its implementation. Chapters 3 to 6 detail its configuration to meet the challenges of validating the research objectives of each chapter.

To gain an early appreciation for the value of such a simulation, the theoretical results proposed by (2.22) are compared to the results of such a simulation in Fig. 2.20. Theoretical and simulated results have been compared by utilising the catalogue of RSOs introduced earlier, in Section 2.1, as simulated objects. To produce the simulated data, a simulated sensor was positioned on the equator and the prime meridian. A counter maintained a record of the number of times each RSO entered the sensor's surveillance volume. These times were then averaged over the length of the simulation and plotted accordingly.

The simulated results presented in Fig. 2.20 exhibit a very strong visual correlation with the relation described by (2.22). The previously identified steep fall in passes in LEO, the gradual decline in passes in MEO and the bimodal result in GEO, are particularly noticeable. Nevertheless, there is also variation about the trend described by (2.22). The variation observed around the previously identified peak at $\mathcal{T}_{\text{RSO}} \approx 2.6$ hours is readily explained by the objects that do not quite conform to the assumption that only near-GEO satellites have small inclinations. These objects are likely to have relatively low inclination and are therefore visible to the equatorial sensor more often than objects with greater inclination. This hypothesis has been supported in Fig. 2.20 by adding a trend described by (2.19) which details the revisit rate of an RSO to a specific longitude. Further variation can be seen around mid-MEO with a slight increase in the rate of observation around $\mathcal{T}_{\text{RSO}} \approx 10$ hours and a significant drop at $\mathcal{T}_{\text{RSO}} \approx 12$ hours. The cause of this variation can be more easily deduced by colour coding the simulated data according to eccentricity and inclination – as has been shown in Figs. 2.21 & 2.22.

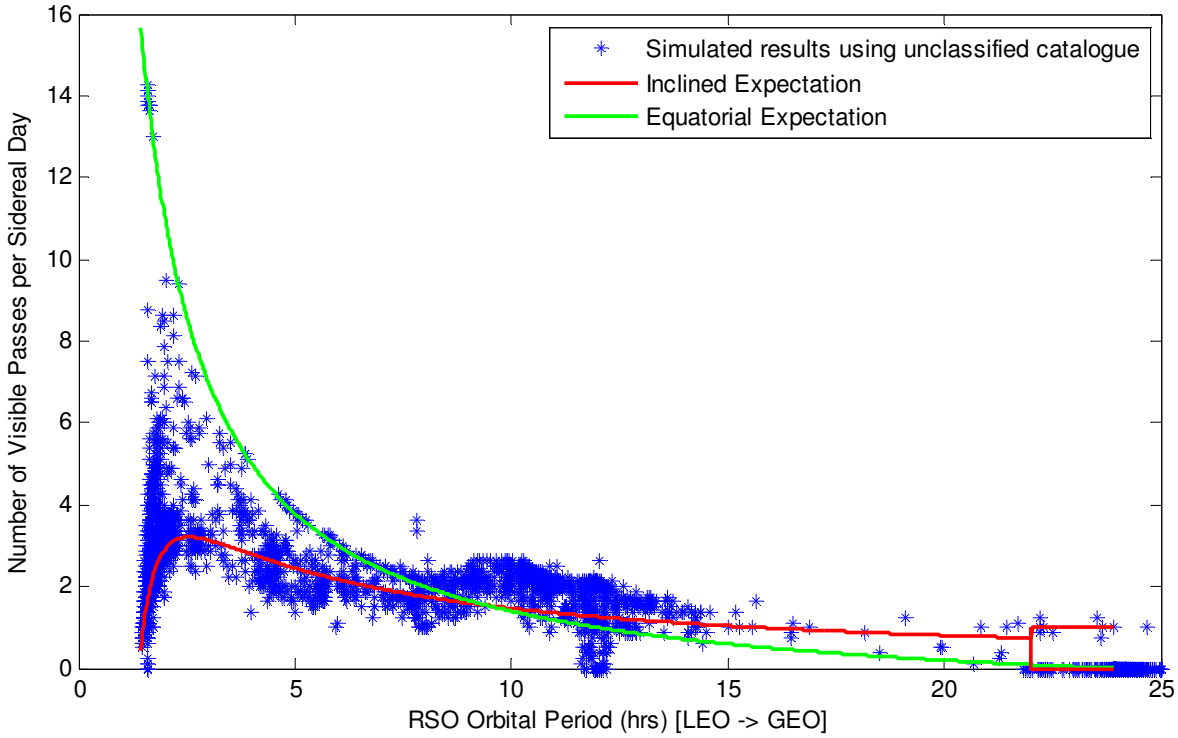

Figure 2.20: The number of expected and simulated passes of an RSO through a sensor's surveillance volume per day, for LEO → GEO orbits, with a minimum sensor elevation of 20°.

The assumptions made throughout Section 2.2.4 are generally well represented throughout Figs. 2.21 & 2.22. Nonetheless, the area in which the assumptions are least representative of the catalogue is in the mid-MEO area, which is seemingly populated by objects with highly eccentric orbits and inclinations varying from 0° to 65°. A particularly strong variation in the number of expected visible passes is noted at $\mathcal{T}_{\text{RSO}} \approx 12$ hours. Using the eccentricity e and inclination i information detailed in Figs. 2.21 & 2.22, these objects are likely candidates for Molniya objects—$e \approx 0.7$ & $i = 63.4°$—and their associated debris [40]. This deduction was confirmed by using the simulation's meta-information regarding these particular objects. The large drop in visibility is a direct result of their mission-specific orbital characteristic. As these objects are Russian equivalents to GEO-synchronous satellites [40], their orbits have been engineered to result in a loiter, at apogee, above Russia. Due to their characteristic 12 hour orbital period, they also spend an equal amount of time loitering 180° around the other side of the globe, over North America. During the brief period they spend at the simulated sensor's longitude at 0°, which is roughly the midpoint between Russian and North American longitudes, the Molniya object is near its perigee, at low latitudes and altitudes. This makes it difficult to observe anywhere other than at similar latitudes – as foretold by (2.5). As

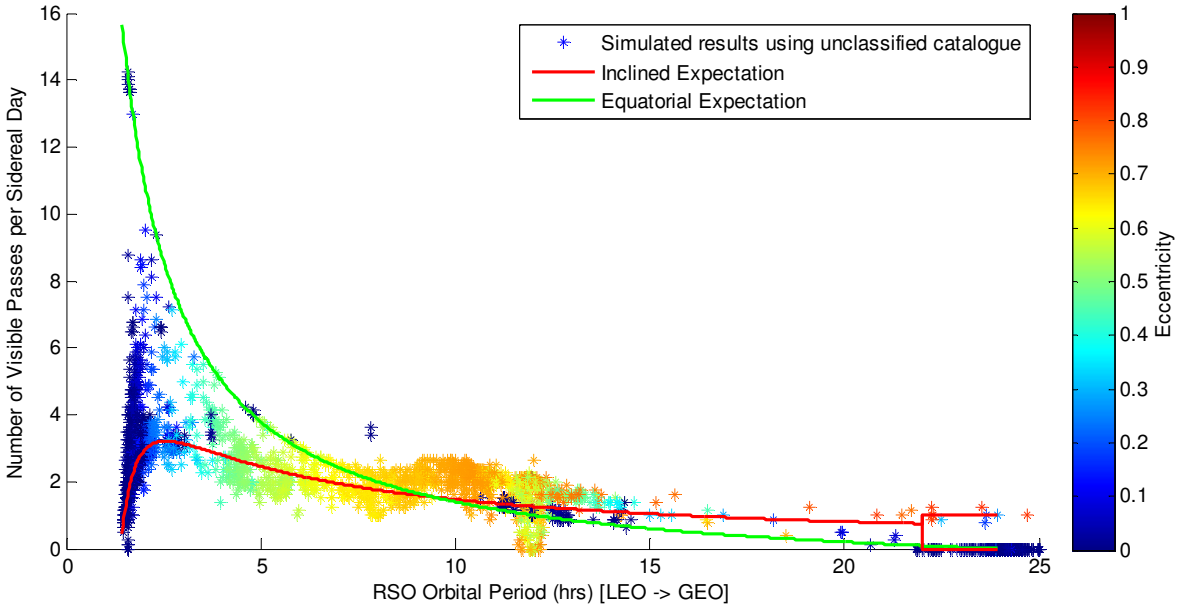

Figure 2.21: The number of expected and simulated passes of an RSO through a sensor's surveillance volume per day, for LEO → GEO orbits, with a minimum sensor elevation of 20°. The simulated results are colour coded by eccentricity.

the simulated sensor is located on the equator and in line with the prime meridian, objects undergoing Molniya orbits can not be observed by it. Appendix B: Figs. B.2's visualisation of an example Molniya orbit, offers some clarification of this concept.

By utilising the simulated data, it has been demonstrated that the relations described by (2.22) and (2.15) provides meaningful insight regarding the regularity at which sensors can expect certain RSOs to traverse their surveillance volumes. Furthermore, by utilising authentic RSO datasets, a greater level of understanding of the SSA environment has been gained. It is for this reason that the results of such a simulation have been pursued whilst completing the research detailed throughout this thesis.

2.3 Tracking Techniques for Maintaining SSA

Sections 2.1 & 2.2 have reinforced the motivation behind obtaining and refining an RSO catalogue. Having a catalogue is crucial for agencies concerned with SSA as observing all of Earth orbit in real time is currently intractable. Furthermore, as the number of objects orbiting the Earth continues to grow, the ability to routinely observe all catalogued RSOs will be increasingly challenging. Obtaining an accurate RSO catalogue to enhance the reliability of RSO trajectory prediction and analysis is therefore a primary objective of agencies attempting to achieve SSA. Nevertheless, even if a complete and error free catalogue

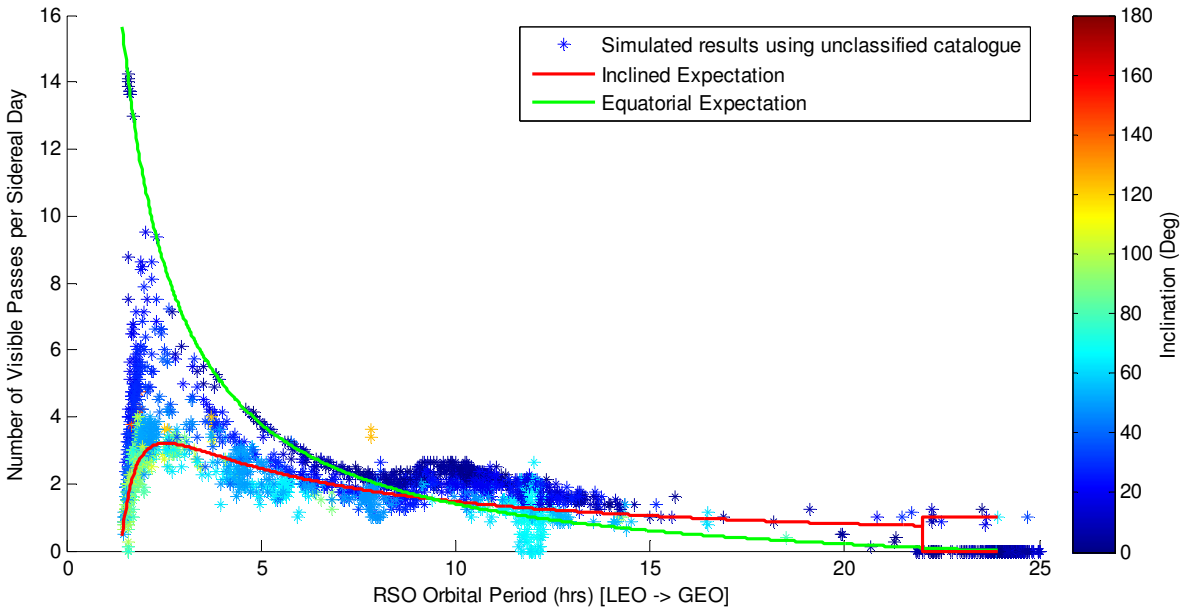

Figure 2.22: The number of expected and simulated passes of an RSO through a sensor's surveillance volume per day, for LEO → GEO orbits, with a minimum sensor elevation of 20°. The simulated results are colour coded by inclination.

could be compiled, utilisation of its contents would immediately begin diverging from reality due to new launches, unreported spacecraft manoeuvring and non-deterministic processes which include re-entry, catastrophic failures of spacecraft, RSO fragmentation events and a variety subtle forces [9, 38]. So whilst obtaining a complete and precise RSO catalogue is the ultimate goal, in reality it requires constant maintenance and will likely only ever be an estimate of the true state of the man-made objects orbiting the Earth.

A primary means of maintaining and ideally improving the accuracy of the catalogue is therefore achieved by utilising the sensors of an SNSS to routinely observe each catalogued RSO and update its state estimate. Measurement noise from the sensors nonetheless introduces additional sources of error to the catalogue. The remainder of the chapter will therefore discuss a number of techniques that are capable of filtering the information received via prediction and measurement to refine a catalogued RSO's state estimate, in spite of these sources of uncertainty. This brief overview of relevant state-tracking techniques will provide a foundation from which these techniques will be referenced, utilised and adapted throughout the remainder of the thesis in the pursuit of enhancing SSA by increasing the reliability of RSO catalogues.

2.3.1 Recursive Filtering for SSA

The Bayesian filtering framework is widely adopted for developing tracking techniques for many applications [43, 53–56], including the tracking of space objects [9, 26]. The framework anticipates indirect or direct measurements of a system's state and a system model, which it uses to recursively estimate an unknown state error probability density function (p.d.f.) as each measurement arrives. This probabilistic approach is designed to compute the most likely state of the system given the observations received so far, whilst in the presence of systematic process error and measurement noise. This framework is well suited for space surveillance, as sensor data is generally collected as a sporadic set or as individual observations. This contrasts with other filtering applications during which complete sets or periodic streams of observations are expected. Intermittent and sometimes infrequent measurements are anticipated due to operational constraints on sensors and orbital trajectories that may limit the visibility of an RSO to the surveillance network. For instance, an RSO's highly elliptical orbit may take it beyond the range of the sensor-network for large periods of time. It is therefore necessary to make use of state estimates between observations.

The Bayesian framework may be described by firstly defining $\mathbf{x}_k$ as the state of the target at time k, $\mathbf{z}_{1:k} = \{\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_k\}$ as the set of observations from the first observation up until the kth observation and $p(\mathbf{x}_k \mid \mathbf{z}_{1:k})$ as the conditional *a posteriori* or posterior state error p.d.f., conditioned by all observations so far. The objective is to transition the distribution $p(\mathbf{x}_{k-1} \mid \mathbf{z}_{1:k-1})$ to $p(\mathbf{x}_k \mid \mathbf{z}_{1:k})$ as each observation is made. In part, Bayes' theorem provides a mechanism to accomplish this and is defined by the equation $p(A \mid B) = p(A, B)/p(B)$ where A and B are random variables and $p(A, B)$ denotes the joint p.d.f. of A and B. From this theorem, the following three equations can be derived:

$$p(\mathbf{x}_k \mid \mathbf{z}_{1:k}) = \frac{p(\mathbf{x}_k, \mathbf{z}_{1:k})}{p(\mathbf{z}_{1:k})}, \tag{2.23}$$

$$p(\mathbf{x}_k, \mathbf{z}_{1:k}) = p(\mathbf{z}_k \mid \mathbf{x}_k, \mathbf{z}_{1:k-1}) p(\mathbf{x}_k, \mathbf{z}_{1:k-1}) \tag{2.24}$$

and similarly

$$p(\mathbf{x}_k, \mathbf{z}_{1:k-1}) = p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}) p(\mathbf{z}_{1:k-1}). \tag{2.25}$$

If conditional independence between all observations can be assumed such that

$$p(\mathbf{z}_k \mid \mathbf{x}_k, \mathbf{z}_{1:k-1}) = p(\mathbf{z}_k \mid \mathbf{x}_k), \tag{2.26}$$

then substituting (2.26) and (2.25) into (2.24) results in

$$p(\mathbf{x}_k, \mathbf{z}_{1:k}) = p(\mathbf{z}_k \mid \mathbf{x}_k)p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})p(\mathbf{z}_{1:k-1}). \tag{2.27}$$

Finally, substituting (2.27) into (2.23) produces the Bayesian update equation

$$\begin{aligned} p(\mathbf{x}_k \mid \mathbf{z}_{1:k}) &= \frac{p(\mathbf{z}_k \mid \mathbf{x}_k)p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})p(\mathbf{z}_{1:k-1})}{p(\mathbf{z}_{1:k})} \\ &= \frac{p(\mathbf{z}_k \mid \mathbf{x}_k)p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})}{p(\mathbf{z}_k \mid \mathbf{z}_{1:k-1})}. \end{aligned} \tag{2.28}$$

Notably, (2.28) only achieves the previously defined objective of transitioning the state p.d.f. from the last observation epoch to the next if each of the densities on the RHS can be accurately described. To analytically implement (2.28), we begin by defining the state update equation

$$\mathbf{x}_k = f(\mathbf{x}_{k-1}, \mathbf{v}_{k-1}), \tag{2.29}$$

where f is the system dynamical model and $\mathbf{v}$ is process noise, which is used to compensate for any deficiencies in the dynamical model. Equation (2.29) describes a Markov process of order one. This characteristic permits us to discard historical information preceding $k-1$ and evaluate the prior distribution $p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})$ as $p(\mathbf{x}_k \mid \mathbf{z}_{k-1})$. Evaluation of $p(\mathbf{x}_k \mid \mathbf{z}_{k-1})$ is achieved by utilising the Chapman-Kolmogorov equation to transition the posterior density at $k-1$ to the prior at time k, such that

$$p(\mathbf{x}_k \mid \mathbf{z}_{k-1}) = \int_{\mathbb{R}^{n_x}} p(\mathbf{x}_k \mid \mathbf{x}_{k-1})p(\mathbf{x}_{k-1} \mid \mathbf{z}_{k-1})\, d\mathbf{x}_{k-1}, \tag{2.30}$$

where n_x is the number of state elements and $p(\mathbf{x}_k \mid \mathbf{x}_{k-1})$ is defined by (2.29). Analytic implementation of (2.28) also requires the equation

$$\mathbf{z}_k = h(\mathbf{x}_k, \mathbf{u}_k), \tag{2.31}$$

where h is the measurement model and $\mathbf{u}$ is the measurement noise, as defined by the sensor's measurement-error characteristics, to define the p.d.f. $p(\mathbf{z}_k \mid \mathbf{x}_k)$. Computing the density $p(\mathbf{z}_k \mid \mathbf{z}_{k-1})$ is also reliant upon the Chapman-Kolmogorov equation and is obtained using

$$p(\mathbf{z}_k \mid \mathbf{z}_{k-1}) = \int_{\mathbb{R}^{n_x}} p(\mathbf{z}_k \mid \mathbf{x}_k)p(\mathbf{x}_k \mid \mathbf{z}_{k-1})\, d\mathbf{x}_k. \tag{2.32}$$

This framework describes a powerful process involving a predictive time-update (2.30)

and a corrective measurement-update (2.28), providing the foundation for filtering techniques used in many tracking applications. The extended Kalman filter (EKF) [53, 56], the unscented Kalman filter (UKF) [57] and the particle filter (PF) [54] are some common filtering methods resulting from the Bayesian filtering framework [9, 26, 43, 55]. Example applications of these filters for tracking, include inertial guidance systems, computer vision, air traffic surveillance and orbit determination.

2.3.2 Kalman Filtering

Whilst in theory the Bayesian filtering process would provide a means for tracking RSOs, the procedure cannot be realised in closed form except in some special cases. This is because the ability to analytically describe the p.d.f.s discussed so far, and thereafter apply the integral equations, are generally not tractable except when the system can be described by linear equations and is influenced only by Gaussian noise [53, 54]. In this special case, the Bayesian filtering process can be implemented using a filtering method commonly referred to as Kalman filtering. The Kalman filter has been shown to be the optimum method [53,58], for obtaining the maximum-likelihood estimate of the state based upon the set of all observations up until time k under the prescribed conditions. Nonetheless, sub-optimal variants such as the EKF and UKF have been developed to permit the application of Kalman filtering principles to nonlinear systems with non-Gaussian densities. This section will provide a brief description of each filter's implementation.

The Kalman Filter

When provided a time varying process that can be completely described by linear equations and is influenced only by i.i.d. Gaussian noise, the Kalman filter is capable of computing the optimal state estimate and its error covariance. It does this by applying statistical mathematics to linear system theory to obtain an estimate of the system state that has the maximum likelihood of producing the observations received so far.

To describe the Kalman filter, we begin by reformulating the Bayesian framework under linear and Gaussian assumptions. This in turn enables us to redefine all densities using only their first two moments. Under these assumptions the state error distribution is represented by is mean $\hat{\mathbf{x}}_k$ and its covariance

$$\mathbf{P}_k = E[(\mathbf{x} - \hat{\mathbf{x}}_k)(\mathbf{x} - \hat{\mathbf{x}}_k)^{\mathrm{T}}] \tag{2.33}$$

where $E[\cdot]$ is the expectation and $(\cdot)^{\mathrm{T}}$ represents the transpose operator. Furthermore,

the state update and measurement equations (2.29) and (2.31) can be redefined as linear equations, such that

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{v}_{k-1} \tag{2.34}$$

and

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{u}_k \tag{2.35}$$

respectively. The matrix $\mathbf{F}$ now applies a linear system model and the matrix $\mathbf{H}$ applies a linear measurement model. The random sequences $\mathbf{v}$ and $\mathbf{u}$ are assumed zero mean independent and identically distributed (i.i.d.) Gaussian noise, they are mutually independent and their covariances are represented by $\mathbf{Q}_{k-1}$ and $\mathbf{R}_k$ respectively.

The Bayesian framework was previously described as exhibiting two primary steps, a time update and a measurement update. The purpose of the time update is to determine what the state error distribution will be at the next time step. Under the linear and Gaussian assumptions, this distribution is now completely described by the mean and covariance [53]. The Kalman filter updates each independently. Whilst a time update equation of the state has already been defined by (2.34), we are lacking a covariance update. This can be achieved by considering the expectation of (2.34), and applying it to (2.33) to obtain a time-update equation for the covariance. This process is computed by means of

$$\begin{aligned}
\mathbf{P}_{k|k-1} &= E\left[(\mathbf{x}_k - \hat{\mathbf{x}}_k)(\mathbf{x}_k - \hat{\mathbf{x}}_k)^{\mathrm{T}}\right] \\
&= E\left[(\mathbf{F}_{k-1}(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1}) + \mathbf{v}_{k-1})(\mathbf{F}_{k-1}(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1}) + \mathbf{v}_{k-1})^{\mathrm{T}}\right] \\
&= E\left[\mathbf{F}_{k-1}(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1})(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1})^{\mathrm{T}}\mathbf{F}_{k-1}^{\mathrm{T}} + \mathbf{v}_{k-1}\mathbf{v}_{k-1}^{\mathrm{T}}\right. \\
&\quad \left. + \mathbf{F}_{k-1}(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1})\mathbf{v}_{k-1}^{\mathrm{T}} + \mathbf{v}_{k-1}(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1})^{\mathrm{T}}\mathbf{F}_{k-1}^{\mathrm{T}}\right]
\end{aligned} \tag{2.36}$$

where the covariance's prior status is indicated using the conditional subscript $[\cdot]_{k|k-1}$. Under the assumption that $(\mathbf{x}_k - \hat{\mathbf{x}}_k)$ is uncorrelated with $\mathbf{v}_{k-1}$, the final two terms of (2.36) reduce to zero resulting in

$$\begin{aligned}
\mathbf{P}_{k|k-1} &= E\left[\mathbf{F}_{k-1}(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1})(\mathbf{x}_k - \hat{\mathbf{x}}_{k-1})^{\mathrm{T}}\mathbf{F}_{k-1}^{\mathrm{T}} + \mathbf{v}_{k-1}\mathbf{v}_{k-1}^{\mathrm{T}}\right] \\
&= \mathbf{F}_{k-1}\mathbf{P}_{k-1|k-1}\mathbf{F}_{k-1}^{\mathrm{T}} + \mathbf{Q}_{k-1}
\end{aligned} \tag{2.37}$$

where $\mathbf{P}_{k-1|k-1}$ is the posterior covariance at time $k-1$.

To compute the measurement update, the Kalman filter utilises a recursive least squares estimation technique [9, 53, 55, 59]. The recursive element was built upon the least squares estimation technique first attributed to Gauss [56]. Its original formulation was devised to compute the most likely parameters of a system from a set of noisy observations, by

minimising the sum of the squares of the differences between the observed values, and the values predicted according to the measurement model $\mathbf{H}\hat{\mathbf{x}}$. The process was adapted to a recursive formulation to make it easier to apply to systems, such as tracking systems, that receive observations sequentially and indefinitely. Without a recursive technique, the degree of computation necessary to determine the current state quickly surpasses the computational capabilities of the tracking system as at each time step it must retain and utilise all observations made so far.

A recursive version is achieved by 'storing' the information gained by all observations prior to time k, $\mathbf{z}_1 \dots \mathbf{z}_{k-1}$, within the state error mean $\hat{\mathbf{x}}_{k|k-1}$ and covariance $\mathbf{P}_{k|k-1}$. The optimal recursive estimate [9, 53, 55, 56] is thereafter achieved by use of the equation

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}) \tag{2.38}$$

where $\mathbf{K}_k$ is the estimator or Kalman gain matrix and the innovation $(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k-1})$, is the difference between the current observation $\mathbf{z}_k$ and the predicted observation as computed by $\mathbf{H}_k\hat{\mathbf{x}}_{k-1}$. The role of $\mathbf{K}_k$ is to weight the innovation such that the sum of the variances of the estimation errors are minimised at time k. The Kalman gain matrix is consequently dependent on the covariance $\mathbf{P}_{k|k-1}$ to represent the confidence in the prior's estimate due to all previous observations whilst the influence on the covariance due to the current measurement is related to $\mathbf{H}_k$ and $\mathbf{R}_k$. The optimal Kalman gain matrix has been shown [9, 53, 55, 56] to be computed by means of

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\mathbf{H}_k^{\mathrm{T}}(\mathbf{H}_k\mathbf{P}_{k|k-1}\mathbf{H}_k^{\mathrm{T}} + \mathbf{R}_k)^{-1}. \tag{2.39}$$

When equations (2.38) and (2.39) are considered in combination, an intuition for how the optimal Kalman gain $\mathbf{K}_k$ influences the posterior estimate may be attained. If confidence in the target's position is high, the covariance $\mathbf{P}_{k|k-1}$ will be small resulting in a small Kalman gain. Consequently, the state update will be more reliant on the prediction $\hat{\mathbf{x}}_{k|k-1}$ and give little weight to the observation. Conversely, if confidence is low and the covariance large, the prediction will be heavily modified by the innovation vector $\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}$ and the relative size of $\mathbf{R}_k$. In so doing, the Kalman filter modulates its estimate based upon what it determines to be the most reliable of the two sources of state information.

The covariance also needs to reflect the influence of the Kalman gain on the recursive

update. This is achieved using (2.38), (2.39) and $\mathbf{P}_k \triangleq E[(\mathbf{x}-\hat{\mathbf{x}}_k)(\mathbf{x}-\hat{\mathbf{x}}_k)^{\mathrm{T}}]$ such that

$$\begin{aligned}\mathbf{P}_{k|k} &= (\mathbf{I}-\mathbf{K}_k\mathbf{H}_k)\mathbf{P}_{k|k-1}(\mathbf{I}-\mathbf{K}_k\mathbf{H}_k)^{\mathrm{T}} + \mathbf{K}_k\mathbf{R}_k\mathbf{K}_k^{\mathrm{T}} \\ &= (\mathbf{P}_{k|k-1}^{-1} + \mathbf{H}_k^{\mathrm{T}}\mathbf{R}_k^{-1}\mathbf{H}_k)^{-1} \\ &= (\mathbf{I}-\mathbf{K}_k\mathbf{H}_k)\mathbf{P}_{k|k-1}\end{aligned} \tag{2.40}$$

where $\mathbf{I}$ is the identity matrix.

The culmination of these equations is the Kalman filtering process which can be summarised using the following equations.

Time Update:

$$\begin{aligned}\hat{\mathbf{x}}_{k|k-1} &= \mathbf{F}\mathbf{x}_{k-1|k-1} + \mathbf{v}_{k-1} \\ \mathbf{P}_{k|k-1} &= \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^{\mathrm{T}} + \mathbf{Q}_{k-1}\end{aligned} \tag{2.41}$$

Measurement Update:

$$\begin{aligned}\mathbf{K}_k &= \mathbf{P}_{k|k-1}\mathbf{H}^{\mathrm{T}}(\mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^{\mathrm{T}} + \mathbf{R}_k)^{-1} \\ \hat{\mathbf{x}}_{k|k} &= \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}) \\ \mathbf{P}_{k|k} &= (\mathbf{I}-\mathbf{K}_k\mathbf{H})\mathbf{P}_{k|k-1}\end{aligned} \tag{2.42}$$

Using these equations, it is theoretically possible to obtain the best estimate of the system's state $\hat{\mathbf{x}}_{k|k}$ and accurately determine the error distribution as represented by the covariance $\mathbf{P}_{k|k}$ as each noisy measurement $\mathbf{z}_k$ arrives. Nevertheless, unanticipated discrepancies such as minor computational error introduced by a tracking computer or, more overtly, attempting to apply Kalman filtering to highly nonlinear dynamics, violate the assumptions made during the Kalman filter's derivation. For this reason, in practicality, non-optimal adaptations to the Kalman filtering framework are necessary to achieve a stable and implementable filtering method for most tracking applications.

The Extended Kalman Filter

As observed in Section 2.1, even the most fundamental formulae necessary to describe orbital motion are nonlinear. Orbital dynamics are well known [9, 26] for their nonlinearity and as such the system of equations necessary to describe an RSO's state propagation and its observation, according to a measurement model, are intrinsically nonlinear. For this reason the EKF is now introduced as the first tractable filtering process for space object tracking. It builds upon the conventional Kalman filter (CKF) to produce a related process that is

compatible with nonlinear system dynamics.

The extension to the CKF, from which the EKF derives its name, incorporates additional components that compromise its optimality and increase estimation error, but gain it the ability to filter error from nonlinear systems. The first extension is the linearisation of the system equations, (2.29) and (2.31), at a fixed set point $\mathbf{x}_{\mathcal{S}}$ or, for tracking, a nominal trajectory. The canonical method for linearisation is to utilise a first order Taylor series expansion of each equation, thereby assuming any higher order components are negligible. The result of this process [56], is the production of matrices $\hat{\mathbf{F}}_{\mathcal{S}}$ and $\hat{\mathbf{H}}_k$ by obtaining the Jacobian matrices of the nonlinear system equations, evaluated at the set point and prior estimate by means of

$$\hat{\mathbf{F}}_{\mathcal{S}} = \left[\nabla_0 f^T(\mathbf{x}_{\mathcal{S}})\right]^T \tag{2.43}$$

and

$$\hat{\mathbf{H}}_k = \left[\nabla_{k|k-1} h^T(\mathbf{x}_{k|k-1})\right]^T, \tag{2.44}$$

where ∇ is the vector differential operator resulting in, for example,

$$\hat{\mathbf{F}}_{\mathcal{S}} = \begin{bmatrix} \frac{\partial f(\mathbf{x}_{\mathcal{S},1})}{\partial \mathbf{x}_{\mathcal{S},1}} & \cdots & \frac{\partial f(\mathbf{x}_{\mathcal{S},1})}{\partial \mathbf{x}_{\mathcal{S},n_{\mathbf{x}}}} \\ \cdot & \cdot & \cdot \\ \frac{\partial f(\mathbf{x}_{\mathcal{S},n_{\mathbf{x}}})}{\partial \mathbf{x}_{\mathcal{S},1}} & \cdots & \frac{\partial f(\mathbf{x}_{\mathcal{S},n_{\mathbf{x}}})}{\partial \mathbf{x}_{\mathcal{S},n_{\mathbf{x}}}} \end{bmatrix} \tag{2.45}$$

where $\mathbf{x}_{\mathcal{S},i}$ denotes the ith component of the vector $\mathbf{x}_{\mathcal{S}}$, $f(\mathbf{x}_{\mathcal{S},i})$ denotes the ith component of the vector produced by $f(\mathbf{x}_{\mathcal{S}})$ and n_x is the dimension of the state vector.

Once computed $\hat{\mathbf{F}}_{\mathcal{S}}$ and $\hat{\mathbf{H}}_k$ may be used in place of $\mathbf{F}$ and $\mathbf{H}$ to apply the standard CKF equations. The effectiveness of this first extension, which may be referred to as a linearised Kalman filter (LKF) [9] or an EKF [56], is limited by the degree of nonlinearity of the system equations.

Unless the system is somewhat linear in the first place, this implementation may suffer in time from divergence between the observation and the nominal trajectory, unless the system is reinitialised using a new, better suited set point. A second useful extension is therefore to pre-empt this problem by re-linearising both system equations with each time step, as each observation arrives. Whilst re-linearising at each time step is more computationally intensive, this solution is particularly suited for applications such as orbit determination, due to highly nonlinear system equations with large time-scales between observations. This combined extension to the CKF may be referred to as an iterated EKF [56], but it will nonetheless be labelled an EKF for the remainder of this thesis.

As the EKF re-linearises at each time step and it is a straightforward process to apply a nonlinear equation to a single point, the most appropriate set point about which to linearise

may be found by applying the most recent state estimate to the nonlinear system equations. The set points for devising the local linear system equations are therefore obtained by finding the state and observation estimate using

$$\hat{\mathbf{x}}_{k|k-1} = f\left(\hat{\mathbf{x}}_{k-1|k-1}\right) \tag{2.46}$$

$$\hat{\mathbf{z}}_k = h\left(\hat{\mathbf{x}}_{k|k-1}\right). \tag{2.47}$$

Adaptation of the Kalman filtering process utilising these principles can consequently be described [56] using the following equations.

Time Update:

$$\begin{aligned} \hat{\mathbf{x}}_{k|k-1} &= f\left(\hat{\mathbf{x}}_{k-1|k-1}\right) \\ \mathbf{P}_{k|k-1} &= \hat{\mathbf{F}}_{k-1}\mathbf{P}_{k-1|k-1}\hat{\mathbf{F}}_{k-1}^{\mathrm{T}} + \mathbf{Q}_{k-1} \end{aligned} \tag{2.48}$$

Measurement Update:

$$\begin{aligned} \mathbf{K}_k &= \mathbf{P}_{k|k-1}\hat{\mathbf{H}}_k^{\mathrm{T}}(\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1}\hat{\mathbf{H}}_k^{\mathrm{T}} + \mathbf{R}_k)^{-1} \\ \hat{\mathbf{x}}_{k|k} &= \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k(\mathbf{z}_k - h\left(\hat{\mathbf{x}}_{k|k-1}\right)) \\ \mathbf{P}_{k|k} &= (\mathbf{I} - \mathbf{K}_k\hat{\mathbf{H}}_k)\mathbf{P}_{k|k-1} \end{aligned} \tag{2.49}$$

where the matrices $\hat{\mathbf{F}}_{k-1}$ and $\hat{\mathbf{H}}_k$ are obtained by computing the Jacobian matrices of $f(\cdot)$ and $h(\cdot)$ at the set points described by (2.46) and (2.47) respectively, such that

$$\hat{\mathbf{F}}_{k-1} = \left[\nabla_{k-1|k-1} f^T(\mathbf{x}_{k-1|k-1})\right]^T \tag{2.50}$$

and

$$\hat{\mathbf{H}}_k = \left[\nabla_{k|k-1} h^T(\mathbf{x}_{k|k-1})\right]^T. \tag{2.51}$$

The Unscented Kalman Filter

The EKF, be it in the form presented so far or as a similar linearised recursive least squares technique, is used widely for tracking [9, 26, 43, 56]. Its effective use in each application is dependent on the validity of the original assumptions used to define the Kalman filtering process. While the EKF is utilised for space-based applications [9] and even has a space heritage [60], an SSA management system's dependence on a tracking filter that relies on linearisation to propagate the state error p.d.f. has shortcomings and limitations [61]. Some of these pitfalls include difficulties implementing linearisation, maintaining filter stability and achieving a meaningful Gaussian estimation of an underlying density that is, in reality,

highly non-Gaussian. This is especially true in applications with highly nonlinear system equations such as those involving tracking of space objects [22, 62]. It is for these reasons that methods such as Unscented Kalman Filtering are being considered for replacing present EKF-based techniques for SSA applications [22, 50].

The UKF provides a relatively efficient middle ground between a highly-efficient low-fidelity EKF implementation and, as will be discussed in greater detail in Section 2.3.3, an arbitrarily accurate Monte Carlo or particle-based implementation. The UKF achieves a greater level of fidelity than the EKF by utilising an "unscented transformation". The unscented transformation entails estimating a density using a discrete distribution of points that capture its fundamental moments – typically but not restricted to its central moments. Each discrete point is then transformed using the well defined, albeit nonlinear system dynamics. Once transformed, the resulting distribution of points is used to estimate the continuous distribution that would have resulted from transforming the entire state space. In doing so, even if a Gaussian approximation is made, the resulting Gaussian density's mean and covariance are based on the nonlinear transformation. This produces a more accurate outcome when compared to the EKF's more challenging task, of obtaining an accurate estimate from a Gaussian estimate that has been transformed by a linearised approximation of the nonlinear transform. Consequently, the UKF's mean and covariance match the mean and covariance of the true, unknown p.d.f. more closely. The unscented transform was successfully integrated into the EKF's, and by extension the Bayesian, recursive filtering framework by Julier and Uhlmann [61] resulting in the UKF.

Time Update:

Representing the central moments of the posterior distribution as a set of discrete points $\boldsymbol{\chi}$, is practically achieved by firstly obtaining the matrix square root of its scaled covariance, to produce an nth-dimensional standard deviation contour. Consequently, these discrete points are referred to as "sigma points". Prior to taking the matrix square root, the covariance is appropriately scaled by the dimension of the state vector n_x. Nonetheless, numerous sources [56, 61, 62] suggest 'adjusting' n_x by adding it to an integer scalar ζ for minimising higher order errors in the prediction. The discrete standard deviation contour is thereby obtained using

$$\boldsymbol{\sigma}^i_{k-1} = \begin{cases} \left[\sqrt{(n_x+\zeta)\mathbf{P}_{k-1|k-1}}\right]^{\mathrm{T}}|_i & i = 1, 2, \ldots, n_x \\ \left[-\sqrt{(n_x+\zeta)\mathbf{P}_{k-1|k-1}}\right]^{\mathrm{T}}|_{i-n_x} & i = n_x+1, \ldots, 2n_x \end{cases} \tag{2.52}$$

where $\left[\sqrt{(n_x+\zeta)\mathbf{P}_{k-1|k-1}}\right]^{\mathrm{T}}|_j$ denotes the jth column of the transposed matrix square root

of $(n_x + \zeta)\mathbf{P}_{k-1|k-1}$. This contour is then added to the $k-1$ posterior mean to locate it appropriately within state-space. A central 'sigma point' is also defined to represent the mean, completing the posterior's discrete representation. The resulting set of sigma points $\boldsymbol{\chi}_{k-1}$ are therefore defined by

$$\boldsymbol{\chi}_{k-1}^{i} = \begin{cases} \hat{\mathbf{x}}_{k-1|k-1} & i = 0 \\ \boldsymbol{\sigma}_{k-1}^{i} + \hat{\mathbf{x}}_{k-1|k-1} & i = 1, \ldots, 2n \end{cases} . \tag{2.53}$$

Propagation of the sigma points from the current epoch $k-1$ to the observation epoch k is achieved by applying the nonlinear state propagation model f to each of the sigma points such that,

$$\boldsymbol{\chi}_{k}^{i} = f\left(\boldsymbol{\chi}_{k-1}^{i}\right). \tag{2.54}$$

The process noise $\mathbf{v}$, defined in the Bayesian system equation (2.29), is taken into consideration by adding the CKF's process error covariance $\mathbf{Q}_k$ to the resulting discrete representation of the prior density. Computation of the prior distribution's mean $\hat{\mathbf{x}}_{k|k-1}$ and covariance $\mathbf{P}_{k|k-1}$, as represented by the sigma points, are thereby obtained by evaluating,

$$\hat{\mathbf{x}}_{k|k-1} = \frac{1}{n_x + \zeta}\left(\zeta\boldsymbol{\chi}_{k}^{0} + 0.5\sum_{i=1}^{2n}\boldsymbol{\chi}_{k}^{i}\right) \tag{2.55}$$

$$\mathbf{P}_{k|k-1} = \frac{1}{n_x + \zeta}\left(0.5\sum_{i=1}^{2n}[\boldsymbol{\chi}_{k}^{i} - \hat{\mathbf{x}}_{k|k-1}][\boldsymbol{\chi}_{k}^{i} - \hat{\mathbf{x}}_{k|k-1}]^{\mathrm{T}}\right) + \mathbf{Q}_k \tag{2.56}$$

Measurement Update:

Observation sigma points $\mathcal{Z}$ are generated next in a very similar manner. These sigma points are generated to represent the Bayesian distribution $p(\mathbf{z}_k \mid \mathbf{x}_k)$, as was previously achieved via the EKF by means of the matrix $\mathbf{H}_k\mathbf{P}_{k|k-1}\mathbf{H}_k^{\mathrm{T}} + \mathbf{R}_k$. The set $\mathcal{Z}$ is generated by applying the state sigma points $\boldsymbol{\chi}$ to the nonlinear observation model (2.31) such that,

$$\mathcal{Z}_k(i) = h(\boldsymbol{\chi}_k(i)). \tag{2.57}$$

Computation of the resulting mean observation estimate $\hat{\mathbf{z}}_{k|k-1}$ and covariance $\mathbf{P}_k^{zz}$ are evaluated using,

$$\hat{\mathbf{z}}_{k|k-1} = \frac{1}{n_x + \zeta}\left(\zeta \mathcal{Z}_k^0 + 0.5\sum_{i=1}^{2n} \mathcal{Z}_k^i\right) \tag{2.58}$$

$$\mathbf{P}_k^{zz} = \frac{1}{n_x + \zeta}\left(0.5\sum_{i=1}^{2n}[\mathcal{Z}_k^i - \hat{\mathbf{z}}_{k|k-1}][\mathcal{Z}_k^i - \hat{\mathbf{z}}_{k|k-1}]^{\mathrm{T}}\right) + \mathbf{R}_k, \tag{2.59}$$

where in this case, observation noise $\mathbf{u}$ is applied by adding the CKF's measurement error covariance matrix $\mathbf{R}_k$ to the covariance. Computation of the cross-correlation matrix $\mathbf{P}_k^{xz}$ is performed next to achieve an equivalent form to the EKF's $\mathbf{P}_{k|k-1}\mathbf{H}_k^{\mathrm{T}}$ matrix. This permits computation of the estimator or Kalman gain $\mathbf{K}_k$ as devised previously in (2.39). This process can be described using the equations

$$\mathbf{P}_k^{xz} = \frac{1}{n_x + \zeta}\left(0.5\sum_{i=1}^{2n}[\boldsymbol{\chi}_k^i - \hat{\mathbf{x}}_{k|k-1}][\mathcal{Z}_k^i - \hat{\mathbf{z}}_{k|k-1}]^{\mathrm{T}}\right) \tag{2.60}$$

$$\mathbf{K}_k = \frac{\mathbf{P}_k^{xz}}{\mathbf{P}_k^{zz}}. \tag{2.61}$$

Finally, as at this point the unscented transform has been suitably applied to each of the nonlinear system equations and equivalent EKF matrices and vectors have been achieved, the standard EKF recursive update may be applied using

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k[\mathbf{z}_k - \hat{\mathbf{z}}_{k|k-1}] \tag{2.62}$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K}_k\mathbf{P}_k^{zz}\mathbf{K}_k^{\mathrm{T}} \tag{2.63}$$

to obtain the posterior state estimate and covariance at time k.

2.3.3 Particle Filtering

A PF (particle filter) is an alternative Bayesian filter that recursively estimates the posterior state error distribution to an arbitrary level of fidelity by using a Monte Carlo-based method called sequential importance sampling (SIS) [54, 63]. While the UKF attempted to improve upon the EKF by applying the nonlinear system equations to low fidelity discrete density models, the PF applies this concept on a greater scale. It utilises Monte Carlo sampling to, in a sense, use as many 'sigma points' as is necessary to represent each density to a desired level of fidelity. It then applies each particle to the nonlinear system equations to be subsequently used as a high fidelity representation of the transformation of each p.d.f.

This permits the Bayesian Filtering process to be implemented using discrete mathematics, thereby avoiding the major difficulties that drove us originally to consider Kalman filtering.

The PF uses a point-mass method whereby the points are a set of N random proposal states $\mathbf{x}^i; i = 1, \ldots, N$ and their masses are represented by weights w^i, where $\sum_{i=1}^{N} w^i = 1$, to approximate the prior and posterior distributions. Continuing under first order Markovian assumptions, SIS approximates the posterior density using

$$\hat{p}\left(\mathbf{x}_k \mid \mathbf{z}_k\right) = \sum_{i=1}^{N} w_k^i \delta\left(\mathbf{x}_k - \mathbf{x}_k^i\right) \tag{2.64}$$

where $\delta(\cdot)$ is the Dirac delta function. The SIS framework requires particles to be independently propagated when the target's location is changing or its uncertainty is increasing to predict the prior distribution. This propagation can be implemented directly using (2.29). SIS recognises that N must approach infinity to accurately represent any p.d.f. throughout all the dimensions of the system state space. As this is impractical to implement using machines with finite memory and processing power, an importance density $q(\mathbf{x})$, similar to the posterior, is used to generate a finite set of weighted samples that are concentrated within regions of state space that exhibit high probability [54]. To achieve this result, the importance density is sampled and weighted to imitate the true posterior density by means of the proportionality

$$w_k^i \propto \frac{p(\mathbf{x}_k^i \mid \mathbf{z}_k)}{q(\mathbf{x}_k^i \mid \mathbf{z}_k)}, \tag{2.65}$$

and appropriate weight scaling is subsequently applied via normalisation. In addition, the importance density can be chosen such that

$$q(\mathbf{x}_k \mid \mathbf{z}_k) = q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{z}_k) q(\mathbf{x}_{k-1} \mid \mathbf{z}_{k-1}), \tag{2.66}$$

enabling particles at time $k-1$ to be augmented using the distribution $q(\mathbf{x}_k|\mathbf{x}_{k-1}, \mathbf{z}_k)$ at time k. Finally, when the denominator in (2.28) is dropped in favour of a proportionality, it can be rewritten as

$$p(\mathbf{x}_k \mid \mathbf{z}_k) \propto p(\mathbf{z}_k \mid \mathbf{x}_k) p(\mathbf{x}_k \mid \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1} \mid \mathbf{z}_{k-1}), \tag{2.67}$$

permitting the substitution of (2.66) & (2.67) into (2.65) to produce

$$w_k^i \propto w_{k-1}^i \frac{p(\mathbf{z}_k \mid \mathbf{x}_k^i) p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)}{q(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i, \mathbf{z}_k)}, \tag{2.68}$$

the modified Bayesian measurement update equation for SIS-based particle filtering.

The product of the particle filtering process is a method capable of filtering non-Gaussian noise from systems involving highly nonlinear dynamical models. When using a PF, there are however well documented difficulties [54,63,64] regarding the use of optimal importance densities, problems with "particle degeneracy"—when all but a few particles have trivial weight—and "loss of particle diversity"—when many particles share the same state space—that must be addressed within each application specific implementation.

3

Simulated Catalogue Maintenance by Disparate Sensors

A number of the fundamental principles of SSA were discussed and explored in Chapter 2. The aim of the research detailed by this thesis is nonetheless to investigate the viability of incorporating the wider deployment of steerable sensors, controlled via judicious approaches to sensor management, to improve catalogue accuracy. To investigate this concept, a means of contrasting and comparing the influence of alternative sensors and surveillance strategies on the RSO catalogue is required.

Research and experimentation via the use of an active SNSS would undoubtedly produce accurate and reliable results. Nonetheless, taking an active SNSS off-line or building a new SNSS for ongoing research is clearly an expensive, and somewhat self-defeating, exercise. A common and logical alternative to field or lab experimentation in such circumstances is to leverage the power of modern computing and simulate the problem at a significantly reduced cost and with increased speed. Whilst confidence in the reliability of results may be reduced when adopting this approach, a prudent balance between efficiency and reliability can be achieved by rigorously and deliberately controlling the fidelity of the simulation. Furthermore, once the simulation has provided sufficient justification for a particular strategy, a scaled experiment may be considered to verify the results. It is for these reasons that

a capability to simulate an arbitrary SNSS observing an RSO population, with adaptable fidelity, is a desirable resource for this thesis' research objectives.

A simulation system is described in a recent paper by Hill et al. [50] who have named their simulation 'Tasking Autonomous Sensors in a Multiple Application Network' ('TASMAN'). TASMAN, like its European equivalent [51], is created and controlled by government agencies. Access to such simulations are restricted on the grounds of national security. They are operated under their respective agency's policies and guidelines for operation. So whilst collaborative use of these simulations was desirable and indeed sought after, sustained and direct control of such simulations was not a practical option for performing ongoing research for this thesis. Whilst there are commercial products such as Analytic Graphics Incorporated's (AGI's) space-mission analysis products [52], the software is priced for commercial and government operators and it is proprietary, causing its application-flexibility to be limited to roles envisaged by its creators and the needs of its consumers. AGI, for example, has focused on aerospace-mission analysis, design, planning and visualisation. Its goal is to assist operators and planners working within, rather than beyond, established techniques and procedures of the aerospace industry. Alternatively, there are projects under way [65] developing open source software suites pertaining to space surveillance, that aim to provide even greater capability to that already provided by NASA's Orbit Determination Toolbox (ODTBX) [66]. Nonetheless, these open-source alternatives, have similar objectives to AGI's products, in that they focus on supporting existing processes rather than developing new techniques. Whilst elements of ODTBX may be useful for certain thesis objectives, it does not provide a foundational structure that is readily adapted to SSA research. To achieve an appropriate level of autonomy and efficiency of research, Chapter 3 consequently explores the creation of a new low-cost simulation using a modular design to independently control the fidelity of its constituent elements.

To determine if such a simulation is capable of characterising and comparing alternative techniques for maintaining SSA, an initial aim of this chapter is to reproduce the results of a higher fidelity simulation, such as TASMAN, to ensure similar outcomes may be produced. As such, a large component of this chapter is devoted to detailing the reimplementation of techniques used during a recent investigation performed on TASMAN [50]. TASMAN simulates an SNSS that attempts to maintain an RSO catalogue using high-fidelity modelling software and hardware-in-the-loop to achieve reliable outcomes. The system is realised by coordinating stand-alone computers that simulate the roles of various system components in addition to supercomputing to permit high-throughput parallel computation. Whilst the cost and time required for researching and testing alternative catalogue maintenance strategies on TASMAN is smaller than utilising a full-scale SNSS, our aim is to achieve an

equivalent result, without the need for the supercomputing or mission-specific commercial and government software that is necessary to support such a sophisticated system.

Consequently, a new tool for SSA research is proposed named MASSAS (MATLAB Space Situational Awareness Simulation). The motivation behind MASSAS is to use a high-level and flexible programming language, such as MATLAB, to generate a software architecture for achieving timely and efficient characterisation of alternative SSA maintenance methodologies. For this reason MASSAS is designed to imitate the high-level functionality of TASMAN, without the focus on high-fidelity modelling. Ideally, MASSAS's fidelity will be minimised to reduce complexity and runtime, whilst ensuring it is sufficiently high to achieve the desired research objectives. MASSAS has consequently been designed for modularity, to enable efficient modification or addition to its constituent components. As a demonstration of this ability, this chapter also presents a characterisation study of optical and radar SSA sensors by means of interchanging the system sensor model. By doing so, the relative catalogue accuracy when using various ratios of steerable sensors and traditional radar sensors may be compared. It is proposed that with this tool, competing methods and alternative approaches can be researched and characterised with significantly less resources and in a reduced amount of time, than if an alternative method was employed.

Section 3.1 begins the chapter with a brief overview of TASMAN and how it achieves its high-fidelity. A general description of the methods used to implement the aforementioned study evaluated on TASMAN is provided in Section 3.2. The section includes the description of any challenges faced whilst attempting to reproduce these elements using MASSAS. Section 3.3 subsequently provides an overview of the development of MASSAS paying special interest to any differences between the simulations. The results of a number of comparative tests between TASMAN and MASSAS are provided in Section 3.4. In addition, results of comparative tests between MASSAS's optical and radar models are also presented. Section 3.5 concludes the chapter with a discussion of the viability of the proposed research methodology, the implications of MASSAS's own results and ideas for future applications for MASSAS.

3.1 High-fidelity Simulation via TASMAN

TASMAN is an SSA research tool tasked by the Air Force Research Laboratory (AFRL) and operated by Pacific Defence Solutions (PDS). Its current mission is to simulate the use of Raven class electro-optical sensors for performing catalogue maintenance. Raven class electro-optical sensors are the current state of the art in high accuracy, steerable space surveillance sensors. Raven class sensors have demonstrated a high-calibre capability for an

atypically low cost [32, 35, 36]. The Raven class is described as a design paradigm for an autonomous electro-optical system that achieves high accuracy whilst comprising of commercial off-the-shelf (COTS) hardware and software to enhance supportability and minimise cost. The class' ability to make high accuracy angular measurements is achieved by exploiting the quality of modern commercially available electro-optical sensors. The imagery produced by these sensors is processed via astrometry to precisely locate an RSO's position relative to the star field in an image's background, using astronomical star catalogues. The class typically involves the use of commercially available telescopes, electro-optical sensors, robotic mounts and associated control software which makes each system small, adaptive and even rapidly deployable [67].

TASMAN uses an asynchronous, reconfigurable and multinodal architecture to simulate the function and interaction of the many components that constitute an SNSS. To achieve high-fidelity, TASMAN relies on a number of commercial and government packages to provide high-quality modelling and data sets. This is particularly evident in its sensor nodes. An illustration of the sensor node architecture has been taken from the aforementioned TASMAN publication [50] and reproduced in Fig. 3.1.

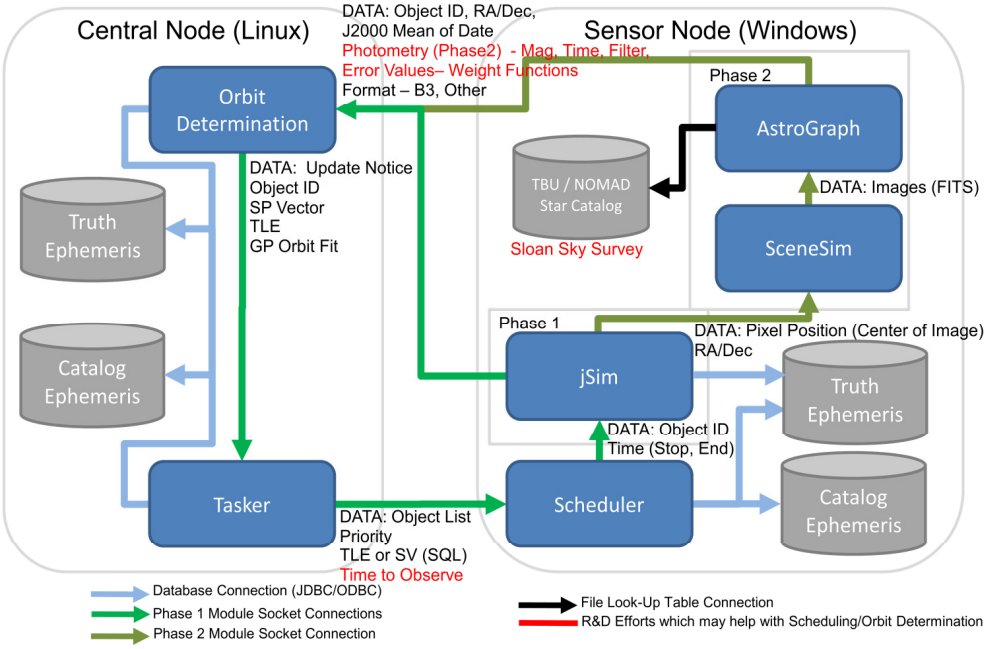

Figure 3.1: TASMAN's multi-phased nodal architecture.

Fig. 3.1 depicts how two levels of fidelity may be achieved by a sensor node by selecting one of two phases. Phase 1 sensor fidelity utilises a government-deployed package called jSim to compute sensor-target geometry according to sensor mount models. The precise sensor-target geometry is used to determine how an RSO would appear to the sensor. This process

includes determining if stars from the Naval Observatory Merged Astrometric Dataset (NOMAD) would be visible in the sensor's FOV during each measurement. Implementing Phase 2 fidelity sees this information passed on to another government-deployed package called SceneSim to generate realistic imagery including simulated distortions and noise. This imagery may then be passed on to software called AstroGraph to perform astrometry on the resulting imagery, as would be typically employed by a Raven class sensor to achieve the class-defining level of angular accuracy [50]. Using this architecture, simulated observations can be produced directly from the sensor model or, for extra fidelity, astrometry processing of imagery may be included in the process.

TASMAN's 'truth orbit' data, the error-free states of RSOs throughout the simulation, is generated using one of two packages. The first package, called TurboProp [50], can reportedly employ a fixed-step fourth-order Runge-Kutta solver, a variable-step fourth/fifth-order Runge-Kutta solver and a variable-step seventh/eighth-order Runge-Kutta solver to model an RSO's orbital dynamics, thereby implementing the system state propagation model. Alternatively a North American Aerospace Defence (NORAD) SGP4 [68] orbit propagator may be used as a less computationally intensive substitute. Propagation of RSO's orbits are implemented according to a variety of industry standard gravitational, atmospheric and solar radiation pressure force-models of which Hill et al. claim to include Jet Propulsion Laboratory (JPL) planetary ephemerides DE403, DE405, and DE421, lunar gravity models GLGM-2, LP100K, and LP150Q, Earth gravity models GGM02C, JGM-3, and WGS-84 and the NRLMSISE-00 atmospheric drag model.

This extensive set of high-fidelity models provides TASMAN's operators a high degree of confidence in the accuracy of its truth data and simulated sensor measurements. Consequently it may be possible to use TASMAN to predict the level of RSO catalogue accuracy to be expected if an SNSS configuration simulated on TASMAN were to be implemented in experimental field trials. Nevertheless, it is unclear if this level of accuracy is necessary to evaluate the relative performance of alternative sensor scheduling techniques. So long as the absolute accuracy of the result is not critical to the analysis, a smaller scale simulation may provide sufficient fidelity to reliably compare the characteristics of an SNSS sensor management strategy to another, using far fewer resources.

3.2 TASMAN Result Replication

The recently published results by Hill et al. [50] detail an investigation that aimed to compare the implementation of alternative sensor management strategies, to determine which would result in the most accurate RSO catalogue. Specifically, three alternative SNSS topologies

were under trial in addition to a scheduling method designed to maximise RSO state error reduction, during each observation. Each strategy was evaluated by numerical simulation on TASMAN.

As the focus of Hill et al.'s investigation concerns relative performance and system characterisation, as opposed to the anticipated accuracy of the catalogue, this investigation serves as an appropriate test case on which to test the hypothesis proposed in the concluding remarks of Section 3.1. To determine if a smaller scale simulation can similarly characterise competing SNSS configurations, MASSAS was configured to adopt as many of the methodologies and initialisation parameters as chosen by Hill et al., as was practical. The remainder of this section details the aspects of Hill et al.'s investigation that were adopted by MASSAS as well as any challenges faced while attempting to reproduce their method.

3.2.1 Sensor Management Strategies

Each sensor management strategy was presented as one of three alternative scenarios named Scenario 1, 2 & 3. Each scenario involves 8 days of simulated sensor management which is segmented into eight 24 hour periods. During these 24 hour periods, tasking and scheduling of the sensors is completed according to each scenario's constraints. Tasking is the process in which each sensor is assigned a subset of all catalogued RSOs to be observed during the upcoming tasking period. Scheduling involves generating a list of times for each of the RSOs that have been assigned to a specific sensor to be targeted and observed by that sensor. A primary source of information for each Scenario's tasking and scheduling system is the RSO catalogue and the orbit error covariance information presumed contained within. The principal differences between each scenario are defined by the network topology and the availability of covariance information throughout the simulated SNSS. Observations of RSOs are obtained exclusively by simulated electro-optical sensors which provide 120 second tracks, consisting of five angle pairs containing right ascension and declination measurements. Each scenario, detailed below, was implemented within MASSAS as the primary source of comparison between each simulation's result.

Scenario 1

The first scenario is designed to imitate the mission planning strategies currently implemented by USSTRATCOM's own SNSS, the SSN. Scenario 1's tasking of RSO observations is performed in the same geographical location in which the RSO catalogue is compiled. Crucially, although the tasker has access to the RSO catalogue and the orbit error covariance

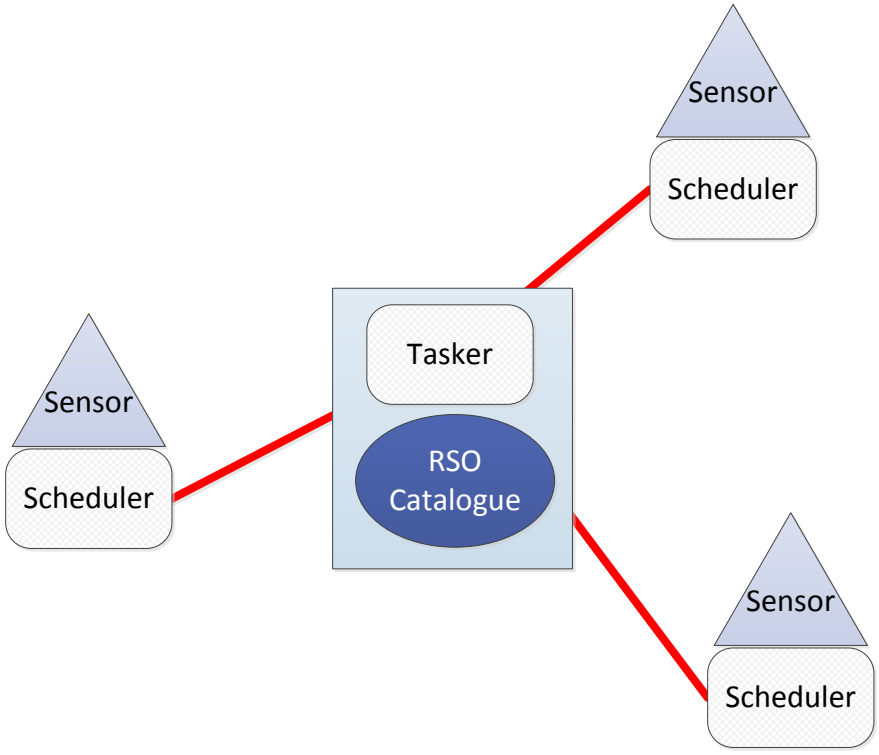

Figure 3.2: Scenario 1 network topology - current scheduling method.

information contained within, it has only rudimentary knowledge of each sensor's capabilities. The tasker creates a prioritised list of objects to be passed onto the scheduler by assigning each object to a category according to its orbit error covariance.

As depicted in Fig. 3.2, each scheduler is co-located with a sensor and has detailed knowledge of its sensor's capabilities. Each scheduler/sensor pair is placed at geographically disparate locations, providing adequate global coverage to observe all RSOs within the catalogue. Each scheduler progresses through the tasking list and produces an observation schedule tailored to its own sensor. Scheduling is accomplished through the application of sensor specific weighting criteria designed to maximise the information gained during an observation using the limited data it has at its disposal. Scheduling criteria include probability of detection, target visibility, concurrent visibility with other targets and remaining opportunities for diversifying observations of the target's orbit. This last scheduling criterion was achieved by ensuring that multiple observations of the same RSO are scheduled a sufficient time apart to ensure each observation is separated by at least 10° in true latitude about its orbit.

Scenario 2

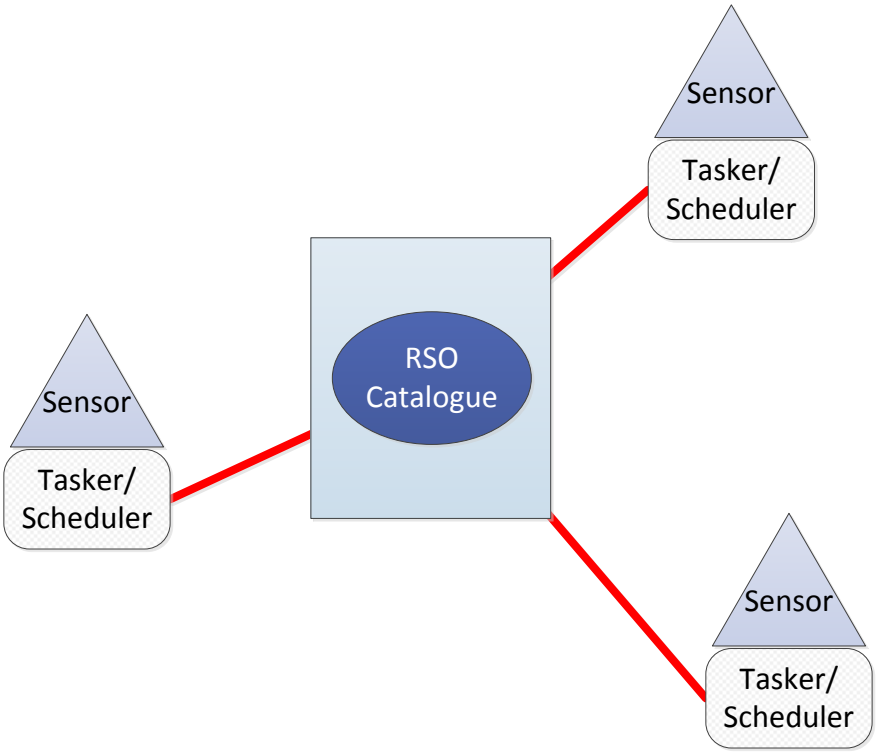

Figure 3.3: Scenario 2 network topology - decentralised scheduling.

Although the RSO catalogue remains centralized in Scenario 2, the distinction between tasker and scheduler is lost as the role of the tasker is effectively absorbed into each scheduler. In addition, each scheduler now has access to the orbit error covariance information contained in the RSO catalogue. Fig. 3.3 provides a depiction of the Scenario 2 topology. This topology enables the schedulers to take advantage of a covariance-based, observation effectiveness scheduling algorithm further explained in Section 3.2.2. Covariance-based scheduling enables the schedulers to not only select the RSO in most need of observation, but also decide when to make an observation to achieve the greatest reduction in orbit error covariance. Additionally, once each scheduler has chosen an object and an appropriate observation epoch, the scheduler can predict how the newly scheduled observation will affect the observation effectiveness of the same object at an alternate time. This feat is achieved by reverse time propagation of the predicted object covariance and enables the scheduler to effectively assign multiple observations to the same object within a single observation period. The weakness of this scenario however is the lack of coordination between schedulers. Each scheduler does not know how the schedules of other sensors will affect the catalogue. This results in redundant observations of some objects from multiple sites, which has the secondary effect of reducing the total number of objects that could have been observed for maximum benefit.

Scenario 3

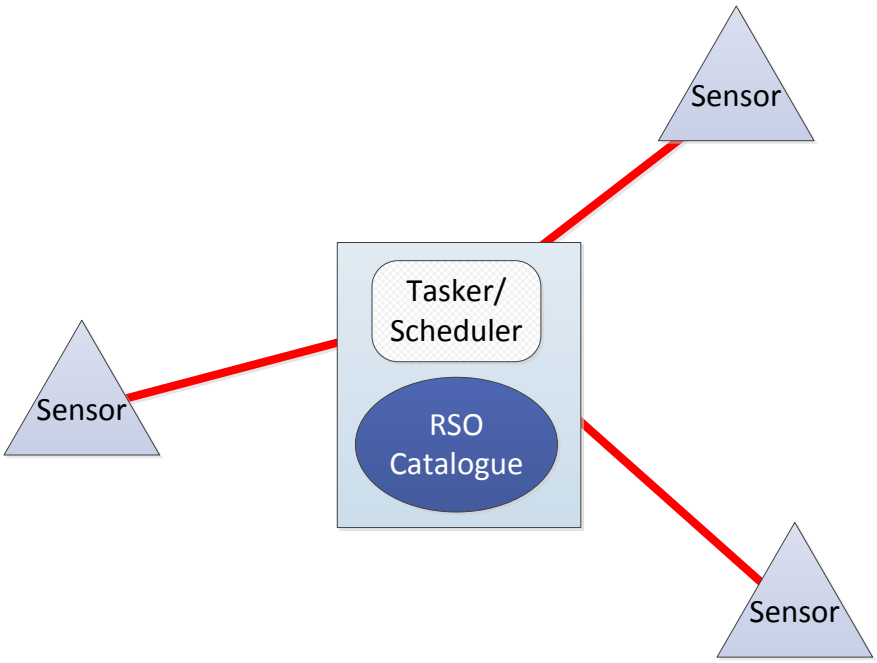

Figure 3.4: Scenario 3 network topology - centralised mission planner.

The final scenario, as depicted in Fig. 3.4, introduces the centralized mission planner which performs the role of tasker and scheduler for all sensors. The mission planner has access to the RSO catalogue as well as detailed knowledge of each sensor's capabilities. Similar to Scenario 2, the mission planner uses a covariance-based, observation effectiveness metric to perform tasking and scheduling. The vital difference however, is the mission planner's ability to overcome Scenario 2's weakness by comparing the observation effectiveness of all sensors for a single RSO. Scenario 3's mission planner can therefore predict how the scheduling of a sensor-object-time combination will affect the orbit error covariance of a catalogued RSO and use this prediction when performing subsequent scheduling.

3.2.2 Covariance-Based Scheduling

In addition to system topology, Hill et al. investigated a new method for scheduling observations of catalogued RSOs. This method uses a catalogued RSO's orbit error covariance to compute a metric for evaluating observation effectiveness. As described in Section 2.3.2, tracking RSOs with a recursive least squares estimator, such as the Kalman filter, generates error covariance. This covariance is an estimate of the target's underlying state p.d.f. Unless the p.d.f. is Gaussian, the covariance is unable to capture all information. Although the

highly non-linear dynamics involved in RSO propagation are unlikely to produce Gaussian distributions, the relative magnitude and orientation of the covariance is commonly employed to gauge a level of confidence in an RSO's current state estimate [9, 22, 26, 31, 50].

The covariance based scheduling method employed by Hill et al. was devised from a fusion of methods presented by Blackman [43] and Tapley [26]. The method presumes the availability of orbit error covariance data to predict an RSO's relative state error reduction when a particular sensor is used to observe it at a specific observation epoch. The chosen technique employs the measurement update of an EKF or UKF to determine a value representative of the reduction in three dimensional position error variance from a single observation. The applicable EKF measurement update equations are,

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\hat{\mathbf{H}}_k^{\mathrm{T}}(\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1}\hat{\mathbf{H}}_k^{\mathrm{T}} + \mathbf{R}_k)^{-1} \tag{3.1}$$

and

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K}_k\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1} \tag{3.2}$$

as previously specified in Section 2.3.2.

The ability to use these EKF equations to determine observation effectiveness is realised by considering the function of (3.2). The posterior covariance is updated by subtracting the matrix $\mathbf{K}_k\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1}$ from the prior covariance. The matrix $\mathbf{K}_k\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1}$ is therefore the predicted reduction in error covariance due to a measurement update. A scalar metric β_{red} is used to denote this reduction in position error covariance, thereby signifying the effectiveness of the proposed observation. β_{red} is computed by obtaining a scalar representation of the position component of the matrix $\mathbf{K}_k\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1}$. Hill et al. chose a Cartesian state vector represented by

$$\mathbf{x} = [x\ y\ z\ \dot{x}\ \dot{y}\ \dot{z}]^{\mathrm{T}}, \tag{3.3}$$

thereby incorporating three position and three velocity components in rectangular coordinates. Obtaining a scalar value for the reduction in position error, β_{red}, was consequently achieved by taking the 3×3 trace of the upper left quadrant of the 6×6 covariance reduction matrix $\mathbf{K}_k\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1}$ as detailed by

$$\beta_{\mathrm{red}} = \mathrm{trace}\left([\mathbf{K}_k\hat{\mathbf{H}}_k\mathbf{P}_{k|k-1}]_{p,3\times3}\right). \tag{3.4}$$

Alternatively, a UKF can similarly be used by employing its equivalent measurement update equation

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K}_k\mathbf{P}_k^{zz}\mathbf{K}_k^{\mathrm{T}}, \tag{3.5}$$

as previously described in Section 2.3.2. Repeating the method employed to produce the EKF-based metric, the equivalent UKF-based β_{red} can be obtained by using the matrix $\mathbf{K}_k\mathbf{P}_k^{zz}\mathbf{K}_k^{\text{T}}$, such that

$$\beta_{\text{red}} = \text{trace}\left([\mathbf{K}_k\mathbf{P}_{k|k-1}^{zz}\mathbf{K}_k^{\text{T}}]_{p,3\times3}\right). \tag{3.6}$$

The UKF-based method was chosen to compute observation effectiveness whilst producing the numerically-simulated results presented in Section 3.4.

A scheduling method that exploits the information provided by β_{red} is subsequently devised by computing β_{red} at all possible observation epochs during a tasking period for each sensor and RSO. Computing β_{red} at each epoch requires less computation than a full EKF or UKF implementation, as all that is necessary is the propagation of the prior covariance to the proposed observation epoch and knowledge of the sensor, such as its measurement noise characteristics and measurement model. After doing so, the scheduler can pick the most appropriate time to schedule an observation to have the greatest effect on the orbit error covariance, by selecting the configuration that resulted in the largest value of β_{red}. Once scheduled, the estimated prior covariance can be computed using (3.5) and reverse time propagated to the beginning of the tasking period. Subsequently, new values for β_{red} for the newly scheduled RSO can be computed and the scheduling process can begin again.

Hill et al. make no claim of optimality concerning their method of devising β_{red}. Indeed they have since published a paper [69] describing alternative β_{red} implementations that also target reduction in velocity error, semi-major axis length and an all-encompassing Frobenius norm. Their results nonetheless did not show any significant improvement over those produced by a position-based β_{red} metric.

The larger issue with this approach, is that the method focuses on maximising error reduction rather than minimising catalogue error. As acknowledged by Hill et al., employing β_{red} alone to implement covariance-based scheduling, causes the scheduler to favour observations late in the scheduling period, once the state error has had time to grow. Whilst this results in a larger error reduction, it is not necessarily the optimal time to schedule the observation. For instance, if an RSO is observed well before its covariance has had time to grow, its error covariance may be lower at the end of the scheduling period, than if the system had waited for large error before making an observation.

3.2.3 The Simulated RSO Catalogue

Truth orbit and state information for initialising TASMAN's RSO catalogue [50] was obtained by Hill et al. from the Space-Track website [39] in the form of Two Line Element (TLE) sets [9] on 15 December 2009. The list of RSOs to be catalogued for the simulation

was selected using the criteria found in Table 3.1.

Table 3.1: Criteria for selection of simulated RSOs.

	Semi-major axis (km)	Eccentricity	Inclination (°)	Radius of perigee (km)
Minimum	25 000	0	50	25 000
Maximum	28 000	0.05	70	28 000

The resulting list of MEO objects contained 214 entries. To ensure comparable results would be obtained when implemented using MASSAS, the same 214 object TLEs were again obtained from the Space-Track website [39]. Furthermore, historical archives were accessed to obtain the TLEs that would have been present on the day Hill et al. obtained their copy.

3.2.4 Sensor Model

In order to implement each of the scenarios, TASMAN was configured to simulate a small SNSS, consisting of three optical sensors. Sensor locations were selected at known space surveillance sites that also provide adequate global coverage. The sites selected by Hill et al. [50], are presented in Table 3.2.

Table 3.2: Sensor site locations.

Site	East Longitude (°)	North Latitude (°)	Height Above Ellipsoid* (m)
Kwajalein, Pacific Ocean	167.7333†	8.716667	50
Albuquerque, USA	253.502717	34.96305	1725
Moron, Spain	354.41194	37.1511	101

Sensor functionality was intentionally limited to ensure simulated observations were physically realisable using existing hardware. Limitations imposed on sensors by terrestrial occlusion and atmospheric distortion were implemented by restricting sensors from obtaining observations at elevations lower than 20° above horizontal. Observations made by the simulated optical sensors produced angle pairs, according to each RSO's simulated truth data. Measurement error covariance for each sensor was modelled using the matrix

$$\mathbf{R} = \begin{bmatrix} {\sigma_\alpha}^2 & 0 \\ 0 & {\sigma_\delta}^2 \end{bmatrix} \tag{3.7}$$

*The assumed elliptical Earth model was WGS-84.

†The East longitude for Kwajalein was erroneously published [50] as 192.2667°.

where the standard deviations of optical measurements of right ascension and declination are denoted as σ_α and σ_δ respectively and were each set to equal 1 arcsecond. This level of measurement error was chosen to replicate the levels of noise present in Raven Class electro-optical sensors [35,36]. Association of observation data with catalogued objects was assumed perfect. Each sensor was also limited to a maximum of 200 observations per tasking period. This limitation was applied to ensure the resulting catalogue accuracies were not related to the frequency with which each sensor had visibility of the catalogued RSOs but were instead a consequence of effective sensor scheduling.

Each of the site locations and sensor limitations were readily adopted by MASSAS's own sensor models, which are detailed further in Section 3.3.2.

3.2.5 Catalogue Initialisation and Evaluation

The process used by Hill et al. to initialise the state estimate and orbit error covariance information in the RSO catalogue was not detailed in [50]. A similar process was nonetheless determined via correspondence with the authors and a degree of deduction. It was learned via communication with PDS, that the RSO catalogue state estimate $\hat{\mathbf{x}}$ and covariance $\mathbf{P}$ was initialised by firstly adding Gaussian noise to the truth states with standard deviations of σ_p and σ_v for position and velocity respectively. The covariance was initialised by assuming an initial independence between the state components. Generation of the 6×6 covariance matrix $\mathbf{P}$ was consequently achieved by the equation,

$$\mathbf{P} = \text{diag}\left[{\sigma_p}^2 \ {\sigma_p}^2 \ {\sigma_p}^2 \ {\sigma_v}^2 \ {\sigma_v}^2 \ {\sigma_v}^2\right]. \tag{3.8}$$

This method was chosen, not only because covariance values for objects obtained from the TLE repository are not directly available to the public—although methods for obtaining error covariance information from TLE records have been developed [70,71]—but also because the intention was to begin the simulation with a catalogue in a 'state of disrepair'. Relatively large uncertainty was desired, to see how capably each scenario would improve the condition of the catalogue throughout the simulation.

The magnitude of the standard deviations σ_p and σ_v were devised empirically by comparing the levels of catalogue error reported by TASMAN with the levels reported by MASSAS. To quantify the catalogue accuracy throughout the simulations and to permit subsequent comparison of the resulting catalogues, two metrics were routinely computed. These metrics are defined the median catalogue position error (MCPE) and the worst-case catalogue position error (WCPE). To compute each metric, the determination of the largest three-dimensional position error between the truth and catalogue estimated states of each RSO,

when propagated over the subsequent 24 hour period, is initially required. Once these position errors are compiled into a list, MCPE and WCPE are obtained by computing the median and maximum of the compilation respectively.

When TASMAN's reported [50] initialisation values of MCPE and WCPE were compared to MASSAS's own values for initial MCPE and WCPE, an approximate, empirically determined match could be obtained. This empirical study was however, not without challenge. It was found that the level of MCPE and WCPE could be easily controlled by altering the value of σ_p. Nonetheless the system was very sensitive to changes in σ_v. In some cases, the resulting orbit error covariance matrix's eigenvalues were of a similar magnitude to computational error, sometimes causing loss of positive definiteness. Consequently the tracking filter and β_{red} scheduling routines experienced a large degree of instability whilst trialling alternative values for σ_v. After conferring with PDS and establishing that they had seen similar instability, suitable values of $\sigma_p = 3.5 \times 10^3$ m and $\sigma_v = 1 \times 10^{-4}$ ms^{-1} were empirically determined to initialise the catalogue to an appropriate level of MCPE and WCPE without causing filter instability.

3.3 MASSAS

Section 3.2 detailed a number of attributes and methodologies to be reimplemented by a smaller scale simulation named MASSAS, that has been specifically designed for our research objectives. As our research objectives call for the efficient characterisation of various methods for maintaining an RSO catalogue, rather than predicting the absolute expected error of the catalogue, MASSAS has been developed with a focus on rapid, flexible SSA research, rather than high-fidelity. The initial goal is to determine if it is capable of characterising competing sensor management strategies in a similar manner to that already achieved by TASMAN. Upon achieving this goal, MASSAS is intended to be used to investigate and evaluate the research objectives of this thesis beginning with the disparate sensor analysis, as detailed in Section 3.4.

MASSAS is a computer program written in MATLAB, that is designed to simulate an SNSS attempting to maintain an RSO catalogue, whilst executing on a single PC. It is flexible in functionality, due to high levels of modularity, as well as computation, due to advanced data structuring techniques afforded by the use of a high level language. The adaptability of MASSAS enables components such as its tasking/scheduling modules, sensor models, orbit propagators, orbit determination modules, physics modules and visualization features to be easily and/or dynamically interchanged and adapted to varying grades of fidelity. The flexibility afforded facilitates a large degree of control over the simulation complexity and

runtime. This section provides an overview of the implementation of a number of MASSAS's primary modules.

3.3.1 Truth Data

Depending on the available computational power and scenario constraints, MASSAS can be modified so that truth data is precompiled or dynamically evaluated during simulation. The benefit of precompiled truth data is a reduction of execution time when running alternate scenarios with the same initialisation parameters. Due to the specified scenario constraints, MASSAS has been configured to precompile with 30 second temporal resolution.

Similar to TASMAN [50], simulated truth orbit data is obtained from TLEs. Rather than perform a full coordinate conversion to internationally recognised reference frames [68], MASSAS was initially configured to perform a classical, unperturbed transformation to Cartesian elements, by assuming TLE elements are equivalent to standard Keplerian elements [9]. For further clarification, the differences between Keplerian and TLE elements are detailed in Appendix A. Whilst this simplification may result in a significant divergence between the true and simulated RSO positions, the simulation requires much less computation while retaining a representative catalogue of RSOs. Using the resulting truth data, propagation of truth state vectors and the application of sensor models enable low fidelity computation of observation angles and object visibility at desired epochs.

3.3.2 Sensor Models

MASSAS and TASMAN use sensor models to determine if it is possible to make an observation of an RSO using a particular sensor and if so, to what quality the sensor will be capable of measuring the RSO's state. MASSAS's emphasis on modular design permits the integration of alternative sensor models with relative ease. Therefore during this investigation, MASSAS has been supplemented with a radar sensor model in addition to an optical sensor model.

Due to the differences in their respective methods of observation, optical and radar observations provide disparate yet complementary orbit track update information [9, 72]. In the absence of observation diversity, optical sensors are likely to reduce in-track error while supplying poor radial information due to accurate angle measurements and the absence of range information. High angular accuracy, as produced by Raven class sensors [35], is achieved by utilising the star field behind the target to perform astrometric measurements to obtain higher accuracy than could otherwise be achieved by extrapolation from the sensor's mount orientation. While no range information is directly observed, the timing between

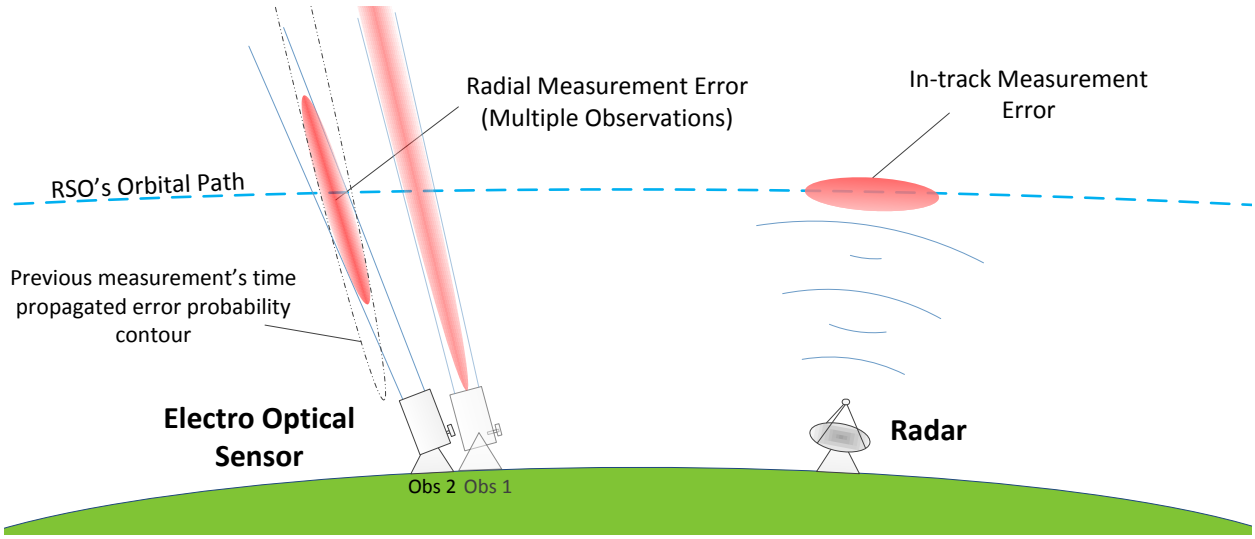

Figure 3.5: Measurement error characteristics of electro-optical and radar sensors.

multiple measurements, each with their own error tolerances, may be used to reduce the region in which the target is likely to reside from an infinite cone to finite volume. Conversely, radar is likely to provide superior radial error performance while providing inferior in-track error reduction. By emitting radio waves and analysing their reflections off of RSOs, a radar can accurately determine the range to an RSO, but a radar's ability to discern the direction to the RSO is significantly inferior to that supplied by astrometry [9, 32]. These complimentary error characteristics have been depicted in Fig. 3.5. MASSAS's respective sensor models are therefore configured to simulate these differences in capability, to compare the effects on the resulting catalogue accuracies.

Both optical and radar sensors are capable of observing more than one object whilst targeting a specific region inside their surveillance volumes. In spite of this, MASSAS is configured to process a single measurement per scheduled observation. This configuration has been chosen for this study for the following reasons:

1. Generality is enhanced using this method as the resulting catalogue accuracies will contrast the measurement error models rather than the idiosyncrasies of a specific sensor type or design.

2. Dismissing the possibility of multiple-object observations improves MASSAS's efficiency when scheduling and processing observations.

3. The proposed covariance based scheduling does not incorporate multi-object visibility in its analysis. Discounting such observations thereby ensures the accuracy of the catalogue is representative of the effectiveness of covariance based scheduling and will not be influenced by fortuitous occurrences of multi-object observations.

Optical

The optical sensor model is configured to simulate the highly accurate angular measurements of Raven sensors to replicate Hill et al.'s results and to model the current state of the art in high accuracy steerable sensors. As such, to imitate Raven class specifications [35, 36], its noise characteristics are modelled using i.i.d. Gaussian noise with a standard deviation equalling 1 arcsecond. Optical sensors measure the visible radiation from the sun, that has passively reflected off of RSOs. The optical sensor model therefore relies on knowledge of sensor capability as well as a solar illumination model. The solar illumination model, as depicted in Fig. 3.6, requires spatial awareness of the Sun $\odot$, Earth $\oplus$, sensor site and RSO. Low fidelity models to compute these geometries were obtained from [9], particularly making use of the algorithms `LSTime`, to perform conversions between Earth-fixed and inertial coordinates, and `Sun`, to obtain the sun's location with respect to the Earth. Using the resulting geometries, factors influencing illumination of the RSO such as solar eclipse, sensor-nightfall and solar phase angle—defined hence forth as the angle between the vectors $\mathbf{r}_{\text{site}\,\odot}$ and $\mathbf{r}_{\text{site}\,\text{RSO}}$—can be evaluated.

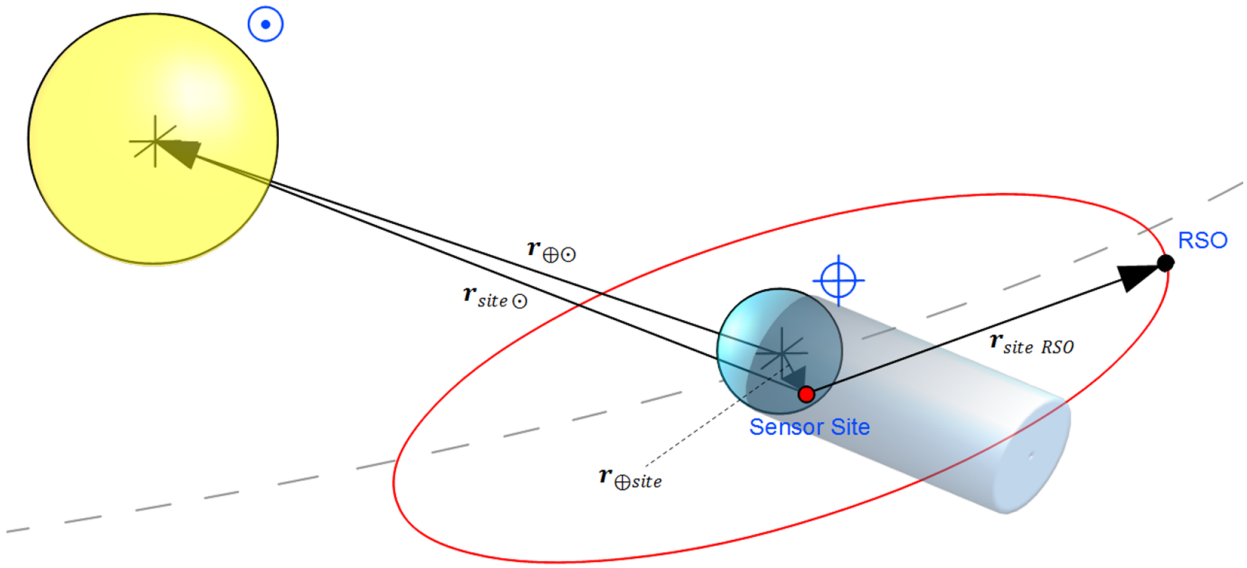

Figure 3.6: An illustration of MASSAS's solar illumination model.

The complexities of a high fidelity illumination model are simplified by a number of prudent assumptions and visibility criteria. Such assumptions include: the sun is the only light source, inclement weather does not occur, the light time involved is negligible, illumination of an RSO is independent of its orientation and the Earth casts a conical shadow, free of diffraction and refraction of light. In accordance with these assumptions, visibility criteria

are applied to determine if the RSO is visible to a specific sensor at time k. The electro-optical sensor model will claim an RSO is invisible to the sensor, unless all of the following criteria are met:

- The RSO must have line of sight with a minimum of 50% of the solar disk before it is adequately illuminated.
- The angle between the $\mathbf{r}_{\oplus\odot}$ and $\mathbf{r}_{\oplus site}$ vectors must be greater than 102° for adequate sensor-nightfall.
- The solar phase angle must be greater than 90°.
- The target must be at least 20° in elevation above the sensor's local horizon.

Whilst these criteria were applied to all electro-optical sensors for this investigation, MASSAS assesses the above criteria on a per-sensor basis enabling site specific configuration constraints. This also permits multiple types of sensors to be active in the simulated SNSS at the same time.

Radar

In contrast to the optical model, the radar model employed for this investigation requires knowledge only of each sensor's specific capabilities. Radars in general can return a number of alternative observation parameters [9]. However, for the purposes of this investigation, the model will provide three parameters: range, azimuth and elevation measurements. Each radar measurement has i.i.d. Gaussian noise added to produce an error standard deviation of 30 metres in range and 54 arcseconds when measuring angles as representative error values from genuine SSN radars [43, 72].

Because the radar model assumes active sensing, and therefore operates independently of passive radiation, the radar model achieves much greater availability when compared to the optical model. This difference in availability is demonstrated in Fig. 3.7 which has been generated using MASSAS's Optical and Radar models. Not only does the optical sensor have to be in adequate nightfall to 'see' the RSO, but it also suffers from gaps in visibility due to eclipse of the Sun by the Earth. In practice however, physical limitations such as range constraints and geopolitical limitations can heavily influence this theoretical advantage for observing certain orbits. Therefore, for a fair comparison, the number of observations made by either sensor were limited to 200 observations per sensor to ensure the scheduling and sensing capabilities are being compared independently of the sensor availability. Similar to the optical model, radar measurements must also occur at elevations exceeding 20°. Finally,

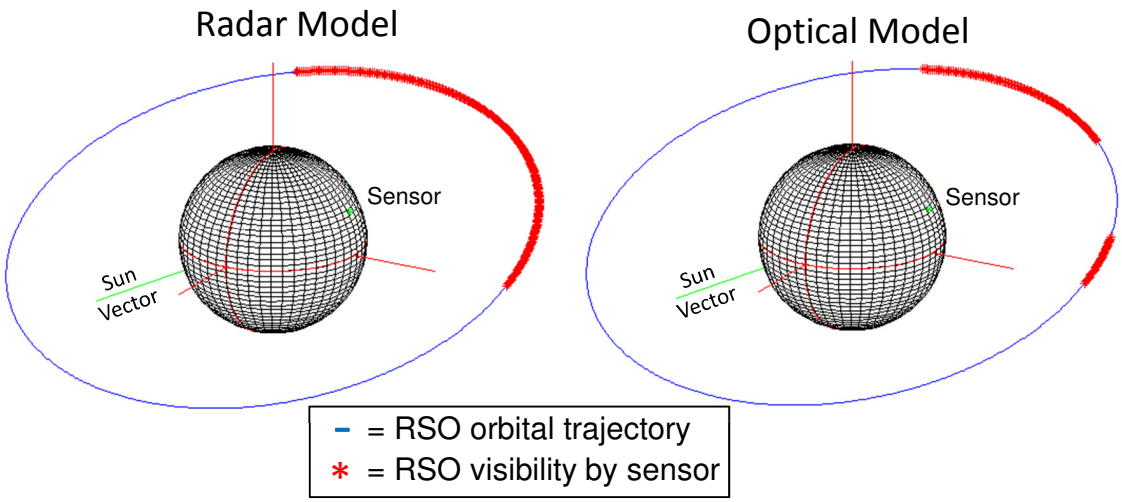

Figure 3.7: Demonstration of the difference in availability of radar and optical sensors using MASSAS's sensor models.

radar range limitations were extended well beyond the Semi-major axis criteria, as listed in Table 3.2, so that the elevation limitation was the only limiting factor for visibility by radar.

3.3.3 Orbit Propagation

Extremely high fidelity orbit propagation is a complex and costly exercise as it not only requires complex semi-analytic models, but also continuous measurement of space weather and the Earth's orientation [9]. While reasonably high fidelity can still be achieved over small time spans when using simplified techniques, the scale of the proposed simulation increases exponentially with resolution in time, observations per tasking period, the number of RSOs and the number of sensors. This is because TASMAN and MASSAS require computation of truth orbits, catalogue accuracy metrics and computation of observation effectiveness for all time-RSO-sensor combinations. Furthermore, observation effectiveness is re-evaluated for the remainder of the tasking period and for each sensor, each time an RSO is scheduled for observation. Therefore, because of the assumptions made during catalogue initialisation, the approximations made in the scenario methodologies, the proposed length of simulation and the fact that both truth orbit data and the models used for filtering observations employ the same propagator, any value gained from the accuracy of high fidelity simulation is likely to be lost in the error caused by these assumptions and approximations. The computational burden necessary to achieve high fidelity simulation is therefore not warranted.

For the reasons considered, MASSAS was configured to use a classical two-body orbit propagation model as defined in [9], using the Keplerian orbit propagation algorithm named

`KeplerCOE`. While MASSAS retains the ability to increase the fidelity of the orbit propagator as required for future investigations, the objective of this investigation is to run the simulation with the lowest fidelity possible, whilst retaining the ability to effectively characterise the performance of the system.

3.3.4 Visualisation Module

A crucial design feature of MASSAS is the incorporation of modules to produce computer graphics. By generating accurate and insightful graphical information, MASSAS's results may be communicated in an intuitive and informative manner. These visualisation capabilities have been used to produce the majority of graphics presented throughout this thesis. Furthermore, the modules were not only designed for generating imagery of post-processed results, but also to generate real-time visualisations. This adds a great deal of efficiency to the process of configuring and utilising MASSAS, as animated visual information is available during solution conception, program debugging and result analysis. The ability to view results as they are being generated is particularly useful as it provides real-time situational awareness of MASSAS's state during simulation and experimentation.

The primary elements MASSAS has been programmed to visualise are the Earth's orientation, RSO locations, RSO orbital trajectories, RSO state error p.d.f.s, the locations of sensors, sensor surveillance volumes and a sensor's FOV. Fig. 3.8 displays an example of the utilisation of multiple sensor models to accurately visualise the surveillance volumes of the SSN's radar assets according to the parameters listed in [9]. Fig. 3.9 displays a visualisation of a Raven class sensor's surveillance volume and its FOV—indicated by the blue cone—during a hypothetical observation. Both images also include the Earth and the positions of the unclassified RSOs as previously detailed in Chapter 2 for scale. Additional visualisations produced using MASSAS's visualisation modules are presented throughout the thesis, as necessary, to aid in the explanation of results. A number of additional demonstrative images, that were not vital for reproduction within the body of the thesis, are presented in Appendix B. Some of this imagery includes graphical displays for programs built to run in conjunction with MASSAS, that were used for demonstrative and exploratory purposes, as well as visualising MASSAS's own real-time SSA capability.

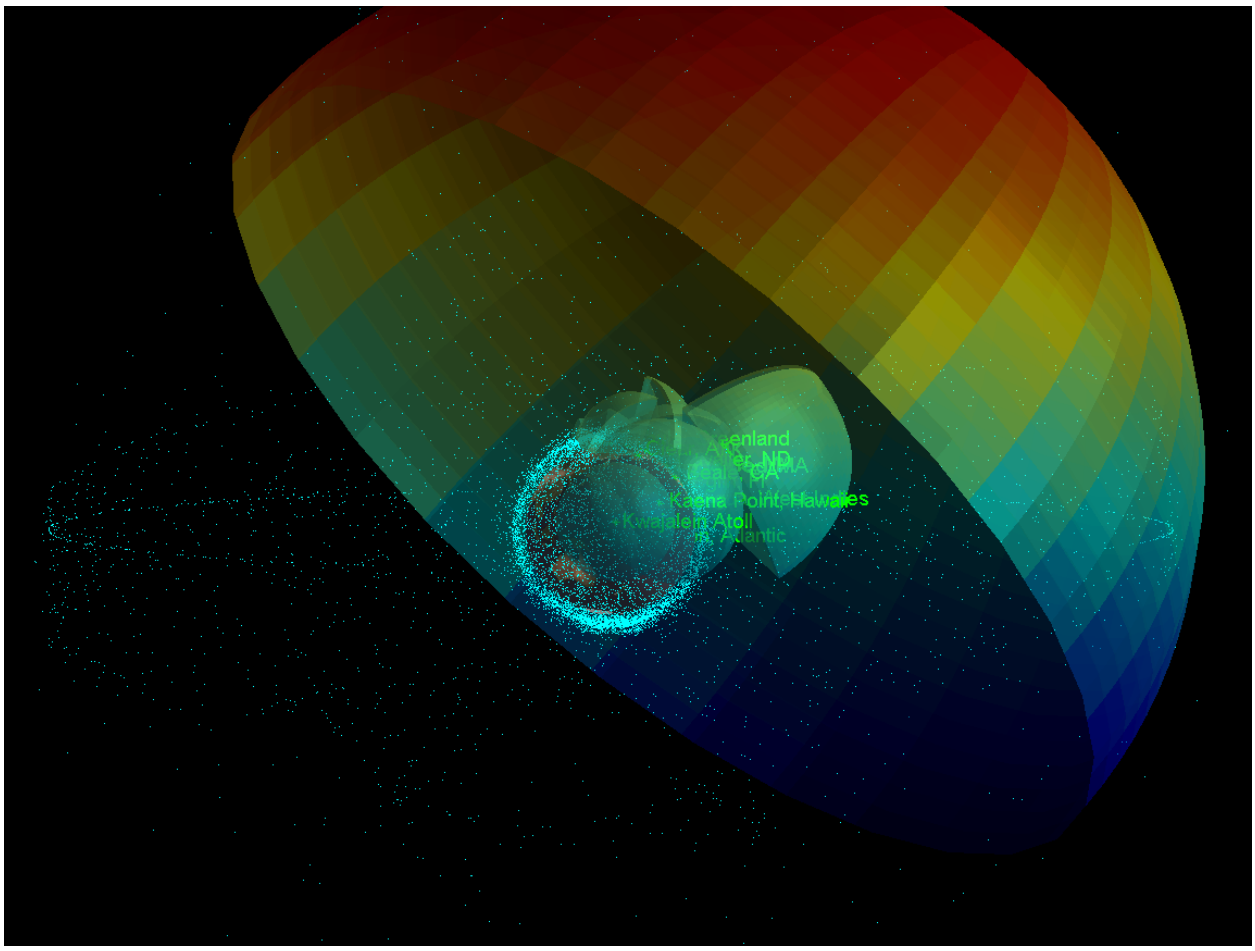

Figure 3.8: A visualisation generated by MASSAS to depict the surveillance volumes of SSN radars.

3.3.5 Orbit Determination

Section 2.3.2 detailed a number of benefits of using a UKF to perform orbit determination. The primary benefits include the ability to manage highly-nonlinear system dynamics whilst maintaining the efficiency gained by representing the state error p.d.f. using only its central moments. For these reasons, and for the sake of consistency with Hill et al.'s implementation [50], an orbit determination module has been implemented in MASSAS that utilises a UKF. The heuristic ζ was detailed in Section 2.3.2 as a means of anticipating higher-order error characteristics, to minimise prediction errors [57]. A method of implementation of the UKF appropriate to SSA has been proposed by Crassidis and Markley [62]. It is suggested that applying a scalar weighting parameter of $\zeta = -3$, in the presence of Gaussian noise and a 6-dimensional state vector, may reduce prediction errors. It is however acknowledged that after nonlinear transformation by the system equations, the state p.d.f. is no longer Gaussian and

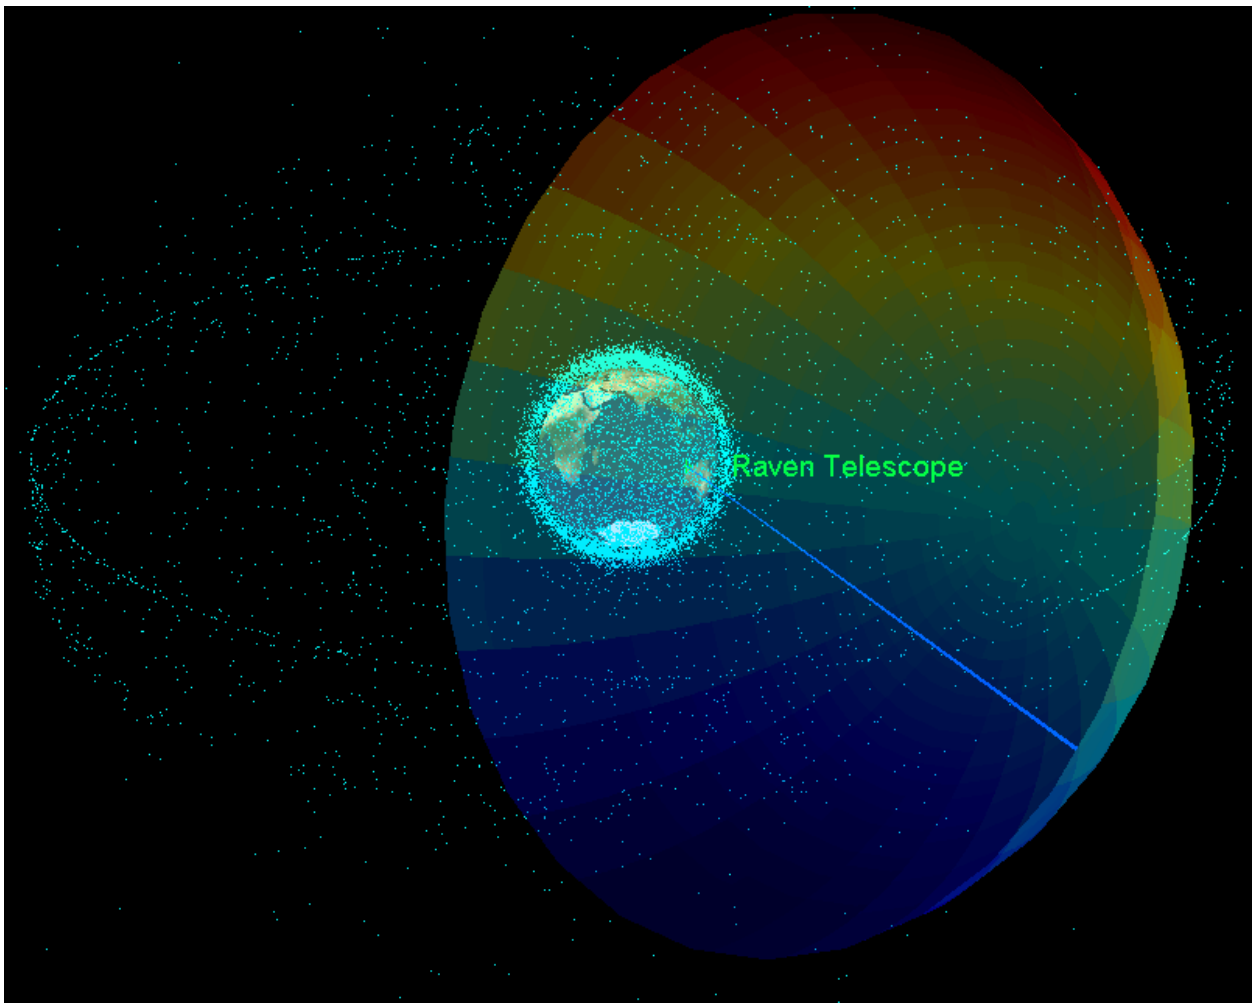

Figure 3.9: A visualisation generated by MASSAS that, as a demonstration, depicts a 20 arcminute angle of view of a Raven class elecro-optical sensor with respect to its surveillance volume.

sub-optimal solutions may result. It was empirically determined nonetheless that $\zeta = -3$ provided enhanced stability over the alternatives discussed by Crassidis and Markley. Filter stability was also enhanced by introducing a modified form of covariance evaluation [61] that prevents the matrix $\mathbf{P}^{zz}$ from losing positive definiteness due to a negative value for ζ. This modification is implemented by evaluating each covariance about the central sigma point, $\boldsymbol{\chi}_k^0$ and $\mathcal{Z}_k^0$ rather than about the estimated mean $\hat{\mathbf{x}}_{k|k-1}$ and $\hat{\mathbf{z}}_{k|k-1}$. These changes, to the implementation discussed in Section 2.3.2, alter the form of (2.56), (2.59) and (2.60). Application of the proposed changes result in the following respective equations:

$$\mathbf{P}_{k|k-1} = \frac{1}{n_x - 3}\left(0.5\sum_{i=1}^{2n}[\boldsymbol{\chi}_k^i - \boldsymbol{\chi}_k^0][\boldsymbol{\chi}_k^i - \boldsymbol{\chi}_k^0]^{\mathrm{T}}\right) + \mathbf{Q}_k \tag{3.9}$$

$$\mathbf{P}_k^{zz} = \frac{1}{n_x - 3}\left(0.5\sum_{i=1}^{2n}[\mathcal{Z}_k^i - \mathcal{Z}_k^0][\mathcal{Z}_k^i - \mathcal{Z}_k^0]^{\mathrm{T}}\right) + \mathbf{R}_k \tag{3.10}$$

$$\mathbf{P}_k^{xz} = \frac{1}{n_x - 3}\left(0.5\sum_{i=1}^{2n}[\boldsymbol{\chi}_k^i - \boldsymbol{\chi}_k^0][\mathcal{Z}_k^i - \mathcal{Z}_k^0]^{\mathrm{T}}\right) \tag{3.11}$$

Finally, it should be noted that the observation model $h(\cdot)$ and the measurement covariance $\mathbf{R}$ are dynamically updated to reflect the characteristics of the sensor that obtained the measurement $\mathbf{z}_k$. Consequently, in the case of radar, the matrix $\mathbf{R}$, as presented in (3.7), is expanded from a 2×2 to a 3×3 matrix to incorporate an additional independent parameter for range-variance. Furthermore, depending on the sensor type used to obtain $\mathbf{z}_k$, $h(\cdot)$ will produce range and/or angular information, requiring the dimension of the observation sigma points and corresponding matrices to vary accordingly.

3.4 Numerical Simulation

3.4.1 Comparison of TASMAN and MASSAS's Scenario Characterisations

Upon configuring MASSAS to imitate the test parameters detailed in [50] as closely as was practical, the prescribed 8 day simulation was performed for each scenario, using three, appropriately modelled, optical sensors.

The results of these simulations are presented in Figs. 3.10, 3.11 and 3.12. Each plot displays the evolution of the catalogue evaluation metrics, WCPE and MCPE, for TASMAN and MASSAS at the end of each 24 hour tasking period throughout each 8 day simulation.

Upon comparing the results, it is evident that both simulations' resulting catalogue accuracies exhibit greater similarity in Scenarios 2 and 3 than in Scenario 1. Nevertheless, when run multiple times, Scenario 1 produces the least consistent result as it relies on a pseudo-heuristic scheduling regime. While Fig. 3.10 displays a large difference between each simulation's WCPE, it should be noted that each simulation's MCPE show a similar trend in the reduction of the magnitude in error.

The most noticeable difference between the behaviours of each simulation's Scenario 2 and 3 results is the large difference observed in the WCPE metric at the end of the first and second days. MASSAS's delayed reduction in WCPE was verified to be caused by a complete lack of observability of two objects from any sensor on these two crucial days. This discrepancy points to either a discrepancy in sensor-visibility models or possibly a discrepancy between RSO positions due to MASSAS's simplified treatment of TLE position data. Whilst this discrepancy is undesirable, the differences between each simulation is sufficient to produce this outcome, without necessarily indicating a failure to meet the research objectives. Most importantly, both systems display similar trends in catalogue refinement and confirm that each scenario is an improvement over its predecessor. At the end of the 8 day simulation, MASSAS experienced a 43% improvement in MCPE between Scenarios 1 and 2. In contrast TASMAN indicated a 13% improvement. Nonetheless, MASSAS experienced a 12% improvement between Scenarios 2 and 3 which is similar to TASMAN's reported 10% improvement.

Whilst a more detailed discussion regarding the implications of these results is provided in Section 3.5, the level of agreement was sufficient to warrant performing the disparate sensor analysis detailed in the following sections.

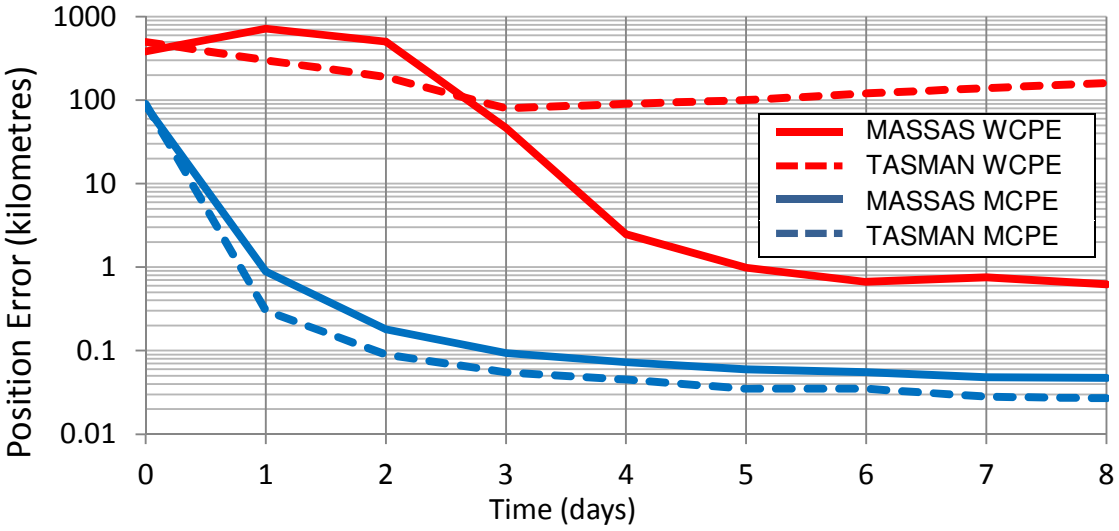

Figure 3.10: MASSAS-TASMAN comparison - Scenario 1.

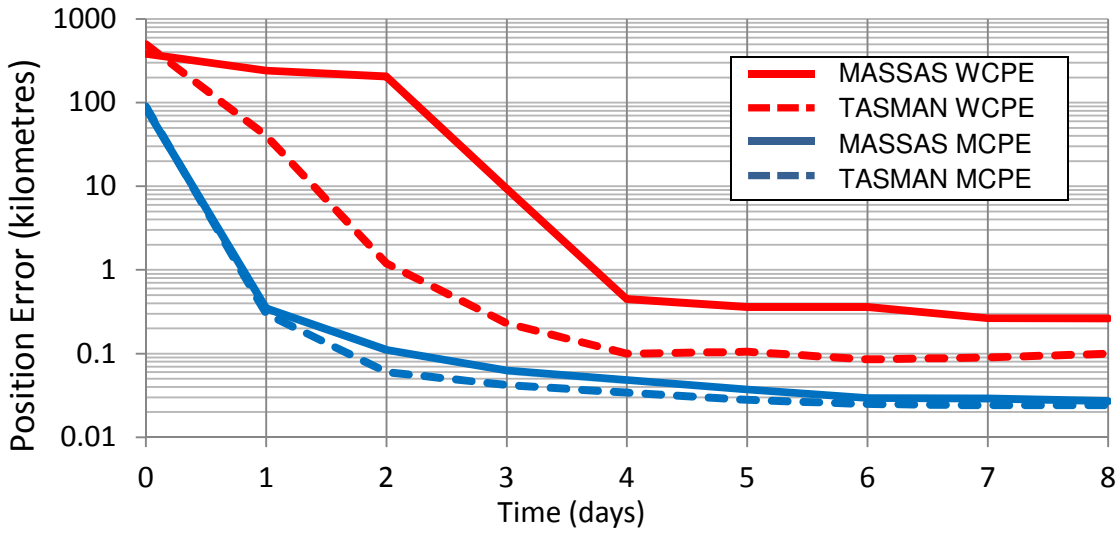

Figure 3.11: MASSAS-TASMAN comparison - Scenario 2.

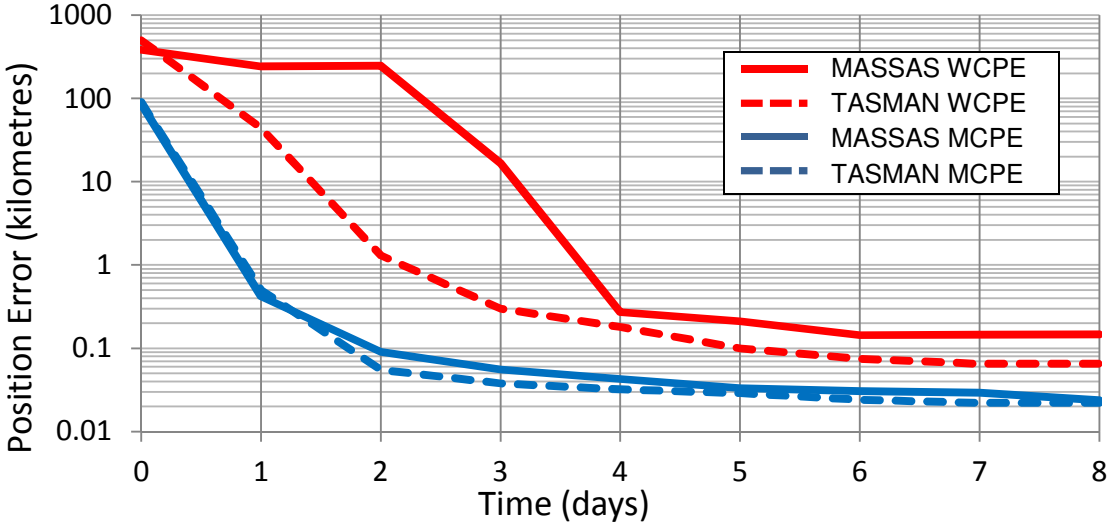

Figure 3.12: MASSAS-TASMAN comparison - Scenario 3.

3.4.2 Comparison Between Homogeneous Radar and Optical Networks

In order to compare the effectiveness and behaviour of the simulated SNSS when it is composed of only optical or radar sensors, further tests were performed, involving three simulated radar sensors. For the most part, the tests were performed under identical test constraints to Section 3.4.1 however, the catalogue was reinitialised using the slightly reduced error values of $\sigma_p = 1 \times 10^3$ m and $\sigma_v = 1 \times 10^{-4}\ ms^{-1}$ to maintain filter stability during the divergent behaviour seen in Fig. 3.13. For a fair comparison between sensor types, optical results were reproduced using the new initialisation settings for each scenario. The results are presented in Figs. 3.13, 3.14 and 3.15.

The results indicate that under the proposed experimental conditions, optical sensors achieve an MCPE with greater accuracy than radar with an approximate 60% increase in catalogue accuracy for both Scenarios 2 and 3. Scenario 1 produced a reduced yet still significant increase in accuracy of 38%. The results also highlight an unsurprising weakness in optical-surveillance due to its dependence on passive radiation. Radar's self-reliance on object illumination enabled it to avoid the stunted reduction in WCPE experienced by optical sensors during the previously identified period, in which two RSOs were not visible to any optical sensors throughout the first two days.

The homogeneous networks demonstrate increasing effectiveness when each scenario is compared to its predecessor. This is especially evident when Scenarios 2 & 3 are compared to Scenario 1. The homogeneous radar network's WCPE actually grows with time, as RSOs in urgent need of observation fail to be scheduled by Scenario 1's scheduling method in an effective manner for the radar's error characteristics. As optical measurements were successfully scheduled using the Scenario 1 scheduler, this result suggests, that the optical network was more resilient to Scenario 1's inferior scheduling.

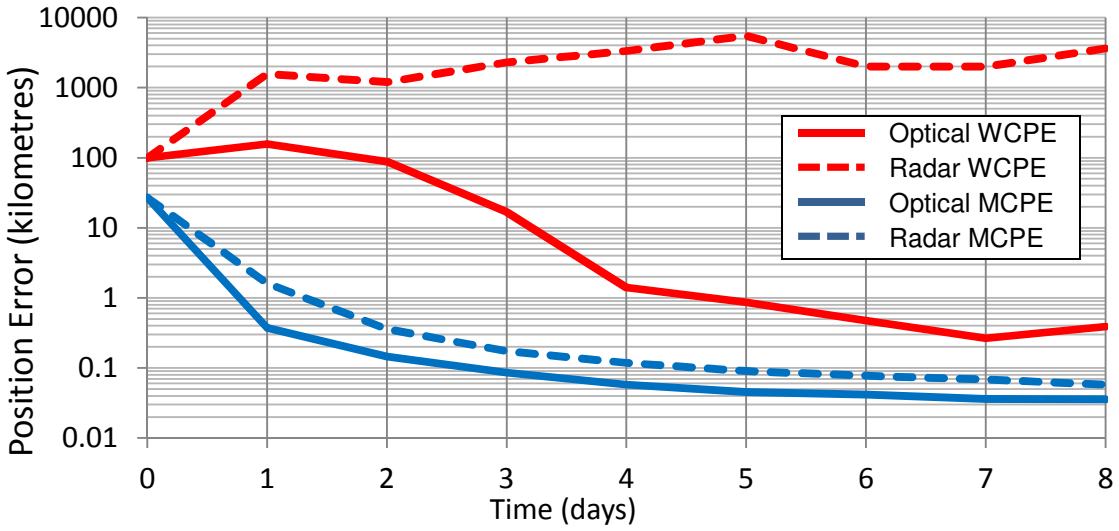

Figure 3.13: Optical-Radar comparison - Scenario 1.

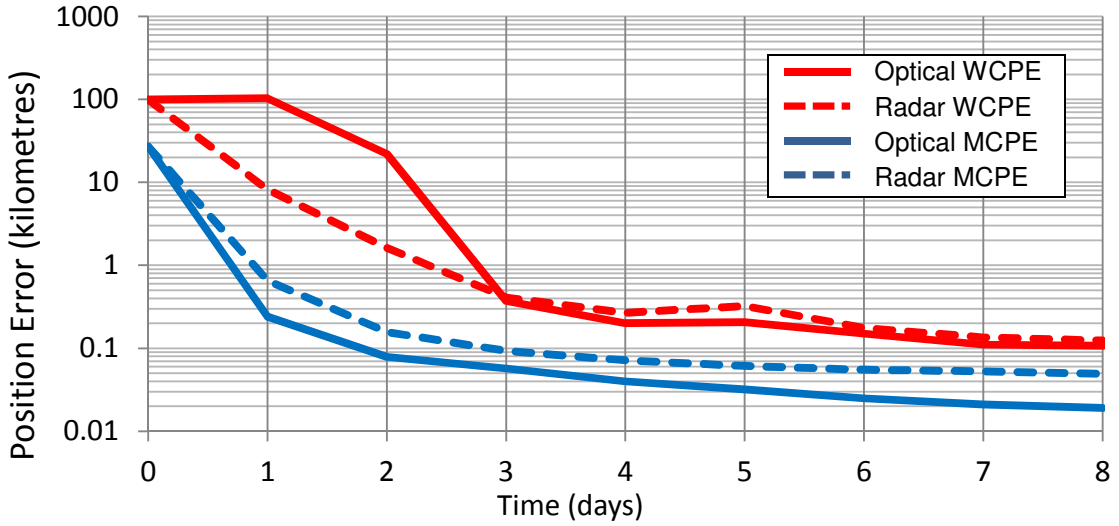

Figure 3.14: Optical-Radar comparison - Scenario 2.

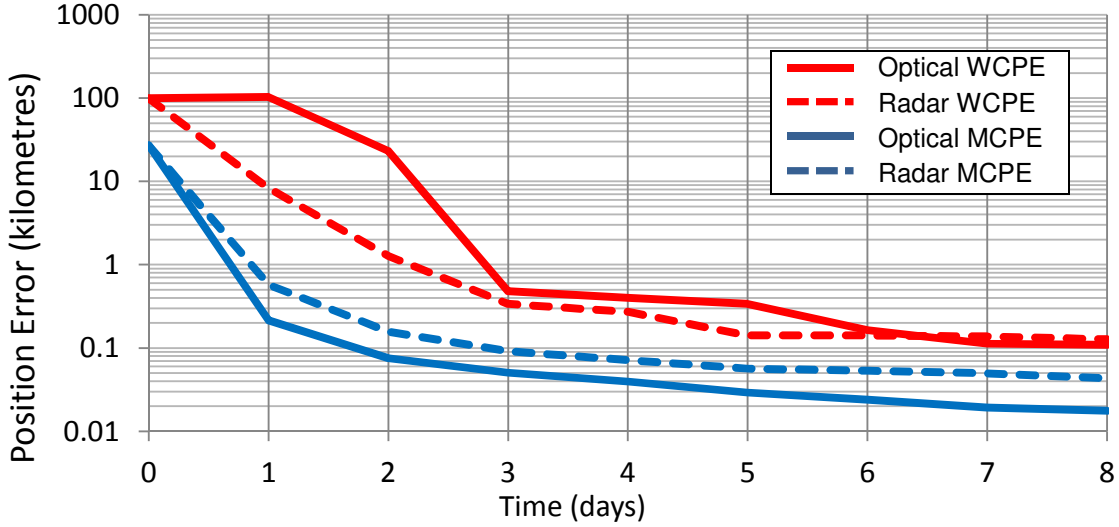

Figure 3.15: Optical-Radar comparison - Scenario 3.

3.4.3 Error Characterisation of Disparate Sensor Types

A third and final set of tests were performed to characterise and contrast the relative errors that result from using sensor networks comprising of homogeneous and non-homogeneous sensor types. The location and number of sensors were again maintained whilst producing these results, however only the results of Scenario 3 have been presented due to its consistently superior performance over the other two scenarios. A single combined-sensor-type, incorporating a co-situated optical and radar sensor, is also included for further contrast and to demonstrate the flexibility of MASSAS's sensor model. The additional combined-sensor-type labelled optically augmented radar (OAR) is a fused result combining the five optical and radar measurements within the UKF orbit determination computation. Evaluation of β_{red} was similarly augmented to represent the available measurement information. In order to better compare the effects of using disparate sensor types, as discussed previously in Section 3.3.2, the catalogued evaluation metrics have been broken down into three orthogonal error components:

1. In-track Error – along the object's velocity vector;

2. Cross-track Error – parallel to the orbital plane's angular momentum vector;

3. Radial Error – the error component orthogonal to the first two components, which is approximately radial for near-circular orbits.

The results of the final set of tests are presented in Fig. 3.16 and Fig. 3.17. Mixed sensor networks have been labelled 'Hybrid' with a numeric suffix indicating the ratio of optical to radar sensors.

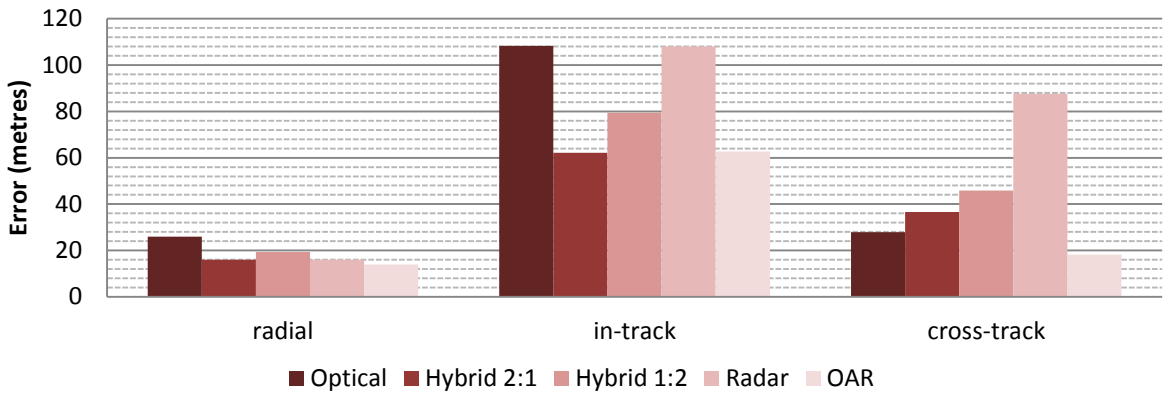

Figure 3.16: Sensor configuration comparison of WCPE on day 8, using Scenario 3.

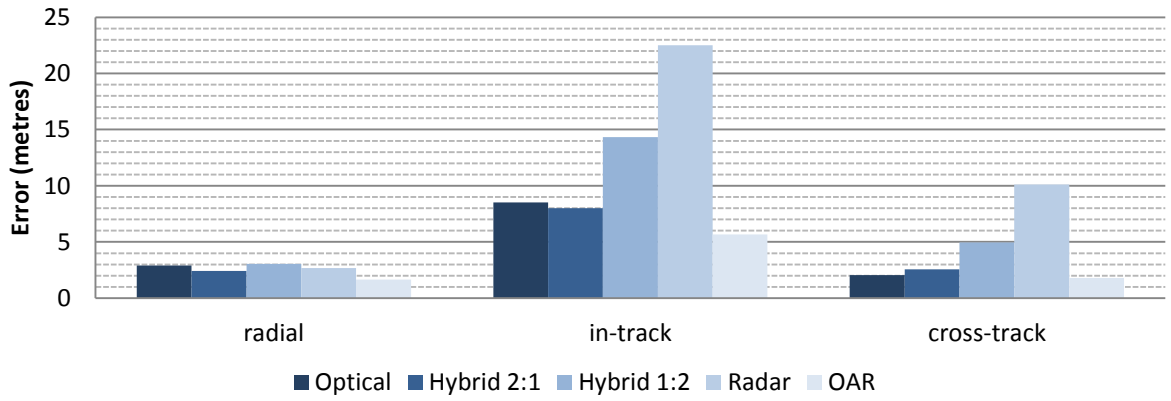

Figure 3.17: Sensor configuration comparison of MCPE on day 8, using Scenario 3.

The results confirm that radar measurements suffer from a large amount of in-track error and to a lesser extent also suffer from cross-track error. Interestingly, the difference between the optical and radar models' radial errors are disproportionate to the disparity in in-track error. This suggests that utilising a network full of low cost optical sensors can surpass the error characteristics, in magnitude, of a network filled with radars. Nonetheless, optical sensors' radial error component is slightly higher than in the all-radar result.

Momentarily disregarding the OAR, the Hybrid 2:1 configuration has proved to be the most effective balance between the attributes and weaknesses of an optical-only and radar-only SNSS, by obtaining the lowest error in radial and in-track components. Replacing an optical sensor with a radar has reduced the optical-only SNSS's visibility limitations previously identified in Section 3.4.2. Furthermore, the Hybrid 2:1 configuration does not suffer from such high levels of in-track and cross-track error as seen in the radar-only and Hybrid 1:2 configurations.

The fused information of co-situated optical and radar sensors has resulted in a very accurate result. Nonetheless, to compare its results to the others is somewhat unfair as it essentially performed double the observations performed by the other configurations. The OAR results thereby serve only to demonstrate the benefit of doubling the number of observations and MASSAS's flexibility to incorporate a variety of sensing capabilities.

3.5 Discussion and Conclusion

The strongest indications of disparity between TASMAN and MASSAS's characterisation of space sensor networks are attributable to minor differences in physical models, best-guess initialisation parameters and the variability of the stochastic process. As Hill et al. published a single realisation of each scenario, as opposed to multiple realisations or an averaged result,

random variation between simulated results cannot be examined in any conclusive manner. Nonetheless, in spite of MASSAS's reduced fidelity, it succeeded in producing commensurate scenario behaviour to that originally produced on TASMAN. The results therefore support the conclusion, that whilst the absolute catalogue errors may not be directly comparable, MASSAS is capable of providing a similar level of insight to that provided by TASMAN, regarding the predicted behaviour and characteristic of alternative surveillance techniques.

MASSAS's disparate sensor analysis has highlighted the limitations of utilising sensor networks with varying ratios of radars exhibiting traditional error characteristics and sensors capable of high-accuracy angular measurements. The results suggest that an SNSS comprising of sensors favouring high angular accuracy, as opposed to range accuracy, will result in less catalogue error. Furthermore, the catalogue accuracy achieved via radar measurements demonstrated higher susceptibility to the inferior scheduling employed by Scenario 1. This result indicates that catalogues exhibiting reduced in-track error may be more robust to scheduling methods that do not incorporate the level of observation evaluation afforded by covariance-based analysis.

The disparate sensor analysis also provides encouraging evidence that the current state of the art in steerable sensors, has the potential to provide a level of catalogue accuracy not yet achieved by existing space sensor networks. This investigation held many elements, such as the number of observations per day and the ability to observe the entire catalogue constant, to isolate the effects caused by disparate measurement error characteristics. An increase in the level of radial error was observed when using optical sensors due to the lack of range information and lapses in availability of the sensor to observe some of the catalogue were detected. Consequently, if the influence of these factors could be reduced, it may be possible to achieve even greater levels improvement.

Supplementary topics for future investigation using MASSAS and its capabilities should include characterising the benefits of alternative types of steerable sensors and the potential benefits of fusing their measurements [73]. MASSAS's speed could also be exploited to rapidly investigate steady state system behaviour and catalogue-building techniques. Additional investigations suitable to MASSAS's attributes also include the investigation of alternative scheduling processes and associated metrics, the involvement of data association strategies, investigating the effects of sensor outages and missed detections as well as the inclusion of maintenance and search modes of operation.

4

Scheduling for Collaborative Sensing

During Chapter 3, the use of disparate sensor types in combination with a range of sensor management strategies was investigated to determine their influence on catalogue accuracy. It was concluded that it is advantageous to utilise a centralised sensor manager to coordinate a sensor network composed of a diverse range of sensor types. This conclusion was made as a result of observing that a simulated homogeneous network would exhibit higher levels of catalogue error due to a lack of sensor diversity. The inherent weaknesses of any one particular sensor type, such as range limitations, reduced availability or unfavourable measurement-error characteristics, could limit an SNSS's ability to maintain an accurate catalogue, if disparate and complementary sensing techniques were not used elsewhere in the SNSS.

The findings indicated that the ratio of disparate sensor types influence the components of catalouge error—for instance, along or across RSOs' orbital trajectories—in a nonuniform manner. This outcome is particularly relevant for maintenance of RSO catalogues as error along an RSO's orbit is a prevalent issue for SNSS operators and researchers in the field of SSA [32, 74–78]. Namely, in-track error is routinely found to be the predominant component of error in RSO state estimates. Chapter 3's results indicated that an SNSS that employs sensors that are capable of obtaining high-accuracy angular measurements are able to minimise error of this type. The concluding recommendation was therefore to utilise a

network augmented with a number of these sensors, in addition to conventional radar assets, to achieve a highly effective result. Nonetheless, due to disproportionate costs and the relative capabilities between sensor types, populating an SNSS with a diverse range of sensors is not necessarily the most judicious solution to the problem.

Recently developed technologies, such as the Raven class electro-optical sensors, demonstrate that a highly capable space surveillance sensor can be developed for relatively low cost [36]. Implementing an SNSS comprising of mostly or only low cost sensors could be a cost-effective option for supplying the RSO catalogue with high accuracy measurements. A comprehensive implementation should nonetheless consider any restrictions, such as the aforementioned weaknesses in availability and measurement characteristics, when planning a heavy reliance on such sensors. For instance, optical sensors are reliant on passive radiation to sense their target and are unable to directly measure a target's range. As Raven class electro-optical sensors currently represent the state of the art in steerable sensors, in this chapter, we seek to improve the error characteristics of an SNSS comprising of electro-optical sensors. Investigating the availability of passive sensors for regularly observing the catalogue is a topic of investigation in Chapter 5.

A potential method for achieving this objective was indirectly proposed by Nash [79] in the late 1970s. Whilst Hill et al. [50] cite Tapley [26] and Blackman [43] as the impetus behind the scheduling technique recounted in Section 3.2.2, Blackman references Nash as the first to suggest that a surveillance system, comprising of multiple sensors and a centralised sensor management system, should utilise the predictive elements of Kalman equations to obtain an error covariance-based metric, to objectively allocate sensing resources. Crucially, Nash's original work not only implies that this technique could be used to allocate sensors to targets, but should be more generally applied to allocate a *combination* of the available sensing resources to observe a target. Application of this technique to an SNSS's centralised sensor management system, would enable it to predict the consequence of combining sensor measurements via measurement-level sensor fusion [73]. This process enables the sensor manager to exploit ancillary sensing capabilities with alternative measurement error characteristics, without necessarily introducing additional types of sensors to the surveillance network. This specifically addresses our goals as, if applied to RSO catalogue maintenance, the sensor management system gains the ability to predict the reduction in an RSO's state error covariance when the angular measurements of two or more optical sensors are fused to produce a combined, range-bounded measurement error distribution. Thus, a system reliant on optical measurements is thereby capable of evaluating and performing timely range bounded measurements, in spite of its sensor's inability to observe range directly when operating independently.

We therefore look to measurement-level sensor fusion [73] as a means of augmenting the proposed observation effectiveness technique, to implement a homogeneous network of electro-optical sensors that is capable of compensating for its weakness in range measurement. By giving the sensor manager the flexibility to assign a single sensor to observe an RSO or to coordinate sensors and simultaneously observe the same RSO, the effectiveness of the observation may be dramatically improved. Even so, simultaneous observations must be scheduled with care. The total number of objects observed must be balanced against the higher quality of simultaneous observations. A good balance will ensure that multiple sensing resources are only assigned to a single RSO's track update if it benefits the accuracy of the catalogue as a whole. Through numerical simulation, this chapter's objective is to determine if simultaneous observation can be used to increase the effectiveness of an SNSS comprising of steerable, electro-optical sensors.

The chapter is organised in the following manner. Section 4.1 proposes a scheduling method that integrates measurement-level sensor fusion with the observation effectiveness metric discussed in the previous chapter. The spacing of SNSS sensors is given consideration in Section 4.2 to determine the conditions necessary to achieve simultaneous observations. In Section 4.3, the reconfiguration and use of MASSAS to compare the previous and proposed scheduling methodologies is described. Section 4.4 presents and discusses the simulated results and Section 4.5 provides some concluding remarks.

4.1 Collaborative Observation Effectiveness

Whilst electro-optical systems are capable of excellent angular accuracy, they are unable to make measurements such as range and range-rate that many space-surveillance radars provide [9]. Nonetheless, combining two or more optical sensors' measurements via measurement-level sensor fusion, has the ability to provide both an accurate and timely position measurement of the target. Depending on the relative accuracies of the sensors involved, applying sensor fusion to optical measurements may minimise or overcome this limitation. This concept is illustrated in Fig. 4.1 where the standard deviation of two geographically separate optical sensor's angular measurements are shown to intersect at the target's approximate location. Whilst the cones defining the standard deviations are conceptually infinite in length, their intersection defines a finite volume, thus providing range information in a single simultaneous measurement.

The objective of this section is therefore to devise a means of predicting how effectively a simultaneous observation by two or more sensors will reduce the state error of a catalogued RSO. By including this prediction in the observation effectiveness update, we aim to enable

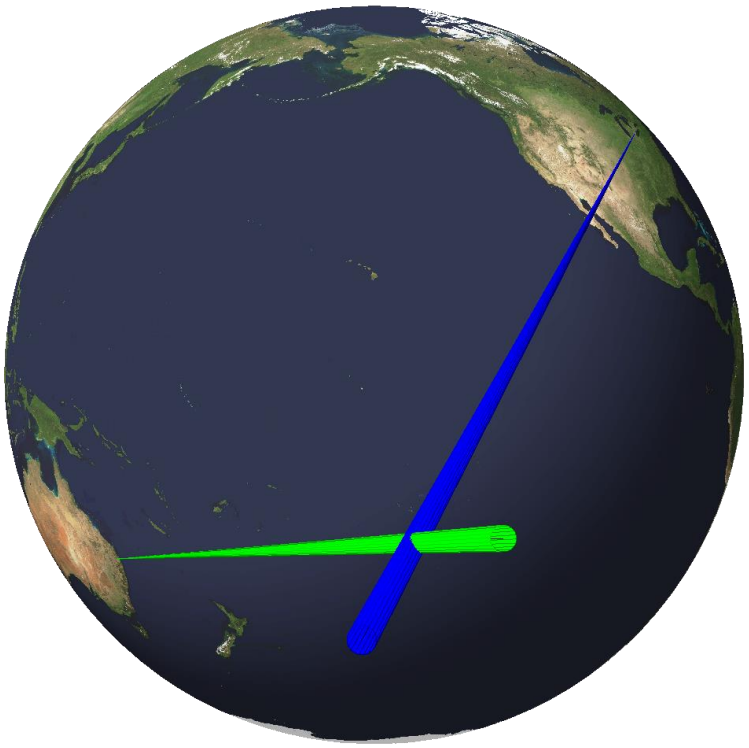

Figure 4.1: An illustration of two electro-optical angles-only measurement error probability contours, whose intersection describes a range bounded measurement.

the sensor manager to decide if and when it would be beneficial to utilise collaborative sensing to reduce catalogue error. We begin by revisiting and refining the covariance-based technique described in Chapter 3.

4.1.1 Review of Covariance-based Observation Effectiveness

Section 3.2.2 presented a covariance-based scheduling method that produces a scalar metric, β_{red}, that enables the sensor manager to schedule observations according to the predicted reduction in state error covariance. Computation of β_{red} allows the sensor manager to anticipate and compare the result of observing an RSO, using a specific sensor, at a specific epoch, to all other combinations of RSOs, sensors and epochs during a predefined tasking period. As we intend to build upon this result by including the capability to predict observation effectiveness for a combination of sensors, we begin by redefining the β_{red} observation effectiveness metric, to explicitly identify the intended sensor, time and RSO.

Let the ith target in the RSO catalogue have a state vector $\mathbf{x}_k^i$ and an error covariance matrix $\mathbf{P}_k^i$ at time k. The observation vector $\mathbf{z}_k$ produced by the jth sensor while observing

the ith target may be described using the EKF measurement equation

$$\mathbf{z}_k^{i,j} = \hat{\mathbf{H}}_k^j \mathbf{x}_k^i + \mathbf{u}_k^j \tag{4.1}$$

where $\hat{\mathbf{H}}^j$ is the linearised observation model and $\mathbf{u}^j$ is the noise vector, of the jth sensor. Each noise component is assumed zero mean and each sensor's noise auto-covariance matrix is defined

$$\mathbf{R}_k^j = \mathrm{E}\left[\mathbf{u}_k^j {\mathbf{u}_k^j}^{\mathrm{T}}\right]. \tag{4.2}$$

The applicable Kalman filter update equation is also redefined using the ith target and jth sensor notation such that

$$\mathbf{P}_{k|k}^i = \mathbf{P}_{k|k-1}^i - \mathbf{K}_k^{i,j} \hat{\mathbf{H}}_k^j \mathbf{P}_{k|k-1}^i \tag{4.3}$$

where

$$\mathbf{K}_k^{i,j} = \mathbf{P}_{k|k-1}^i {\hat{\mathbf{H}}_k^j}^{\mathrm{T}} \left(\hat{\mathbf{H}}_k^j \mathbf{P}_{k|k-1}^i {\hat{\mathbf{H}}_k^j}^{\mathrm{T}} + \mathbf{R}_k^j\right)^{-1}. \tag{4.4}$$

As detailed in Section 3.2.2, the crucial term of the Kalman measurement update-equation, newly described by (4.3), is the matrix $\mathbf{K}_k^{i,j} \hat{\mathbf{H}}_k^j \mathbf{P}_{k|k-1}^i$. It represents the predicted reduction in the prior covariance, $\mathbf{P}_{k|k-1}$, of the ith RSO due to an observation at time k by the jth sensor. Whilst the question of optimality remains open [69], the method employed in Chapter 3 will again be followed to turn $\mathbf{K}_k^{i,j} \hat{\mathbf{H}}_k^j \mathbf{P}_{k|k-1}^i$ into a scalar metric. Namely, the position component of the matrix $\mathbf{K}_k^{i,j} \hat{\mathbf{H}}_k^j \mathbf{P}_{k|k-1}^i$ will be used to define a metric that indicates the predicted level of reduction in position error variance that would result from taking the proposed measurement. The state vector is assumed to contain six elements, comprising of three-dimensional positions and velocities. The observation effectiveness metric β_{red} is thereby computed by taking the trace of the upper left position quadrant such that

$$\beta_{\mathrm{red}}^{i,j,k} = \mathrm{trace}\left(\left[\mathbf{K}_k^{i,j} \mathbf{H}_k^j \mathbf{P}_{k|k-1}^i\right]_{p,3\times 3}\right). \tag{4.5}$$

4.1.2 Sensor Scheduling Incorporating Measurement Level Sensor Fusion

Using the observation effectiveness metric β_{red} in its current form, the scheduler has the ability to choose when to use a particular sensor to observe a specific RSO for maximum reduction in error. To facilitate the assessment of a combined sensor measurement, the notion of the jth sensor is instead replaced by the mth sensor combination. Let s be the

number of sensors in the SNSS and c be the number of combinations, where

$$c = \sum_{j=1}^{s} \binom{s}{j} = 2^s - 1 \tag{4.6}$$

and $m = \{1, 2, \ldots, s, \ldots, c\}$.

This change necessitates the computation of multi-sensor noise auto-covariance matrices and observation models to determine β_{red} for sensor combinations containing more than one sensor. Let $\mathcal{L}^m = \{l_1, \ldots, l_N\}$ be the set that contains the indices which identify the sensors in the mth sensor combination. $\mathcal{L}^m$ contains N elements where $1 \leq N \leq s$. Assuming observation noise is independent of any other observation, multi-sensor updates are achieved by replacing $\hat{\mathbf{H}}_k^j$ and $\mathbf{R}_k^j$ with

$$\hat{\mathbf{H}}_k^m = \left[\hat{\mathbf{H}}_k^{l_1\,\mathrm{T}}, \hat{\mathbf{H}}_k^{l_2\,\mathrm{T}}, \ldots, \hat{\mathbf{H}}_k^{l_N\,\mathrm{T}} \right]^{\mathrm{T}} \tag{4.7}$$

and

$$\mathbf{R}_k^m = \text{blockdiag}\left(\mathbf{R}_k^{l_1}, \mathbf{R}_k^{l_2}, \ldots, \mathbf{R}_k^{l_N} \right), \tag{4.8}$$

to obtain

$$\beta_{\text{red}}^{i,m,k} = \text{trace}\left(\left[\mathbf{K}_k^{i,m} \hat{\mathbf{H}}_k^m \mathbf{P}_{k|k-1}^i \right]_{p,3\times 3} \right). \tag{4.9}$$

As the number of sensors grow, the $\hat{\mathbf{H}}_k^m$ and $\mathbf{R}_k^m$ matrices can become very large. Depending on the architecture of the system to be utilised for computation of $\beta_{\text{red}}^{i,m,k}$, it may be advantageous to use a sequential update procedure. This entails applying the appropriate $\hat{\mathbf{H}}_k^j$ and $\mathbf{R}_k^j$ matrix as an iterative update, for each of the N sensors, in the mth combination. To simplify the notation, let $\mathbf{P}_{k|k-1}^i = \mathbf{p}_0$, where the subscript indicates the number of sensor updates applied to the prior covariance matrix. Updates are thereafter applied iteratively using

$$\mathbf{p}_\alpha = \left(\mathbf{p}_{\alpha-1} - \mathbf{K}_k^{i,l_\alpha} \hat{\mathbf{H}}_k^{l_\alpha} \mathbf{p}_{\alpha-1} \right) \tag{4.10}$$

where $\alpha = 1, 2, \ldots, N$. Finally β_{red} is calculated by evaluating

$$\beta_{\text{red}}^{i,m,k} = \text{trace}\left(\left[\mathbf{p}_0 - \mathbf{p}_N \right]_{p,3\times 3} \right). \tag{4.11}$$

The augmented solution space, as depicted in Fig. 4.2, allows the scheduler to choose a sensor combination to observe a specific RSO at a favourable observation epoch. The process of scheduling is thereafter an iterative method involving the following steps:

1. The largest value in the solution space is chosen and a record made of the event.

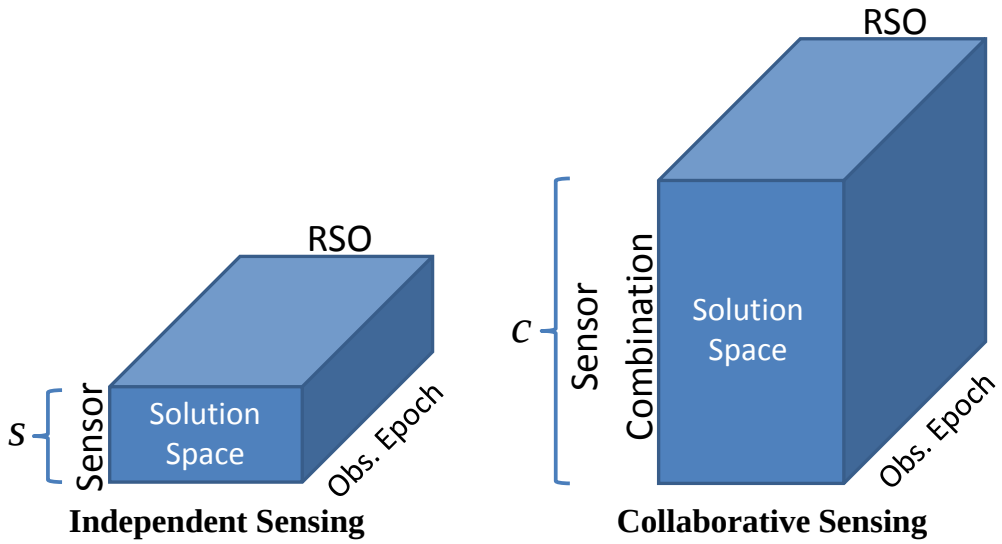

Figure 4.2: A visualisation of the solution spaces that are evaluated to obtain the highest value of observation effectiveness for independent and collaborative sensor scheduling.

2. The newly scheduled RSO's predicted posterior error covariance matrix $\mathbf{P}^i_{k|k}$ is computed using the Kalman measurement-update equation (4.3).

3. $\mathbf{P}^i_{k|k}$ is thereafter reverse time propagated, by applying a negative increment in time to the Kalman filter's time update, enabling re-evaluation of β_{red} throughout the entire scheduling period and for all sensor combinations, in light of the newly scheduled event.

4. Steps 1-3 are repeated until each sensor's schedule is full.

Whilst this process may be capable of anticipating which sensor combination will produce the greatest reduction in state-error covariance, consideration needs to be given to the global-cost of utilising multiple sensing resources for observing a single RSO, as opposed to many. As a naïve first attempt, a weighting method is applied to evaluate the consequence of using multiple sensors for observing a single RSO. β_{red} is therefore inversely weighted according to the number of sensors, N, involved in the prospective observation such that

$$\beta^{i,m,k}_{\text{w}} = \beta^{i,m,k}_{red} N^{-1}. \tag{4.12}$$

Nonetheless, the ability to evaluate simultaneous measurements will be of no benefit at all, unless the SNSS's sensors are suitably spaced to permit collaborative observation. More specifically, sensors must be sufficiently close to each other such that a portion of the RSO catalogue is capable of being observed by more than one sensor at a time.

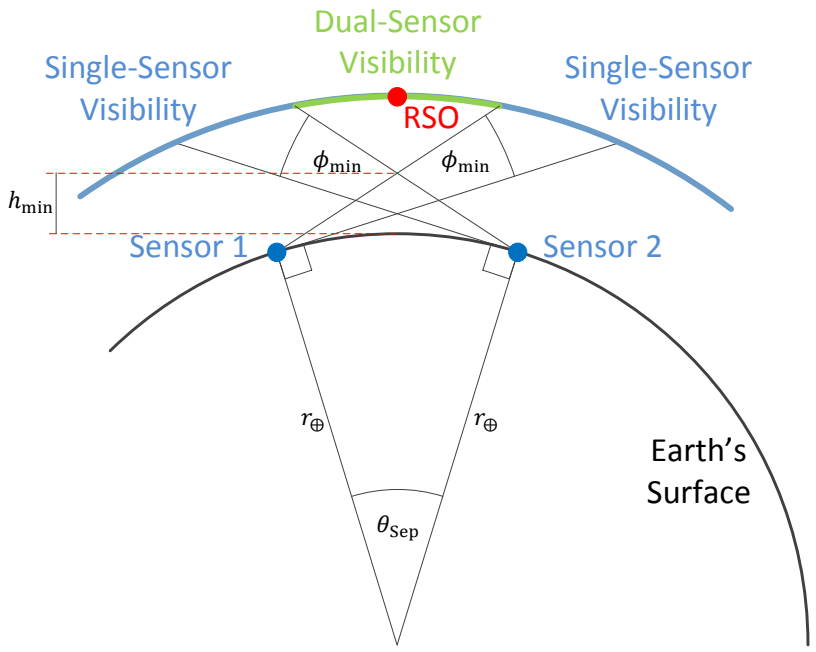

Figure 4.3: An illustration of the geometric requirements necessary for collaborative sensing.

4.2 Sensor Spacing Requirements for Collaborative Sensing

In order to utilise and evaluate the proposed collaborative scheduling technique, it is crucial to understand how the arrangement of SNSS sensors about the globe influences the SNSS's ability to make multi-sensor observations of the catalogue. We begin by building upon the visibility constraints of terrestrial sensors, discussed earlier in Section 2.2.1.

In order to permit collaborative observation, at least two sensors must have visibility of the target at the same time. Fig. 4.3 displays the applicable geometry where h_{min} is the minimum altitude from the surface of the Earth to the point of intersection between two sensors' surveillance volumes and θ_{Sep} is the angle subtending the arc, that traces the Earth's surface, between two sensors. Maintaining their definitions from Section 2.2.1, $\phi_{\min}$ represents each sensor's minimum elevation and $r_\oplus$ represents the mean radius of the Earth. Assuming a spherical Earth model, unlimited sensor range and uniform sensor limitations, the geometry of Fig. 4.3 is solved using the equation

$$h_{\min} = r_\oplus \left(\frac{\cos(\phi_{\min})}{\cos\left(\frac{1}{2}\theta_{\text{Sep}} + \phi_{\min}\right)} - 1 \right). \tag{4.13}$$

If a conservative value of $\phi_{\min} = 20°$ [9] is chosen for this investigation, the domain of (4.13)

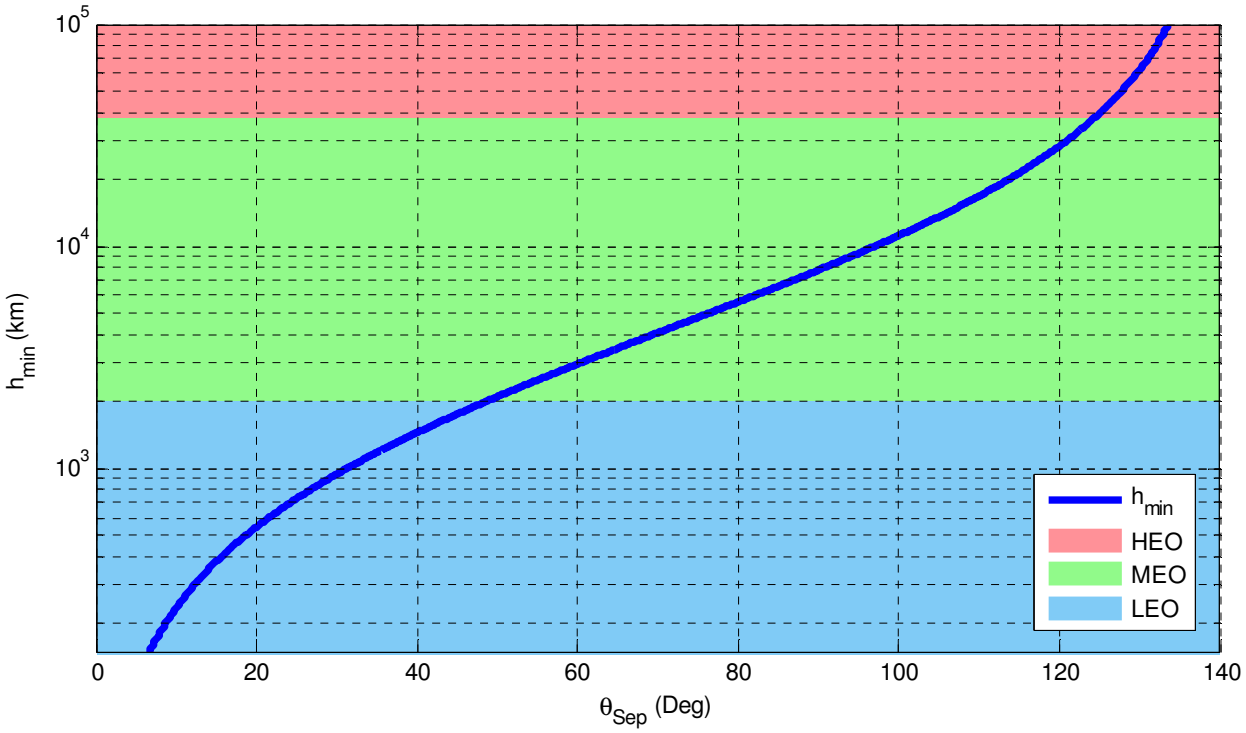

Figure 4.4: The minimum altitude, $h_{\min}$, at which collaborative scheduling can take place given a sensor separation of θ_{Sep} and a minimum sensor elevation of 20°.

may be confined to $0° \leq \theta_{\text{Sep}} \leq 140°$, as for any $\theta_{\text{Sep}} > 140°$ each sensors' surveillance volumes fail to converge*. This indicates that if all sensors have a minimum elevation limitation of 20°, we should not expect any collaborative scheduling to occur when sensors are placed greater than 140° apart, which equates to a separation of approximately 15 600 km across the surface of the Earth. This result can be used to reduce the size of the collaborative solution space described in Fig. 4.2.

Fig. 4.4 displays a plot of the relation described by (4.13) when it is assumed that $\phi_{\min} = 20°$. The curve presented in Fig. 4.4 indicates the lower bound of RSO altitudes capable of being observed via collaborative sensing by two sensors spaced apart by an angle of θ_{Sep}. For example, Fig. 4.4 indicates that it is possible to use simultaneous observations to observe objects orbiting above 150 km in altitude, when sensors are placed less than 6.7° apart, approximately 700 km across the Earth's surface.

Armed with this knowledge, and the RSO catalogue, it is possible to plan an SNSS capable of employing collaborative scheduling, for RSOs above a minimum altitude.

*Noting that θ_{Sep} is implicitly confined to $0° \leq \theta_{\text{Sep}} \leq 180°$ as θ_{Sep} represents the smallest non-zero angle between two sensors.

4.3 Numerical Simulation

An additional observation effectiveness evaluation module was created and added to MASSAS, to evaluate the effectiveness of incorporating the measurement level sensor fusion proposed in Section 4.1.2. To assess its effectiveness, MASSAS was configured to execute two alternative scenarios.

Scenario A employs an independent sensor scheduler by means of the original β_{red} metric described in Section 3.2.2. Scenario B incorporates the augmented method of evaluating β_{red} for collaborative sensor scheduling, as described in Section 4.1.2. Crucially, although the scheduler component of the sensor manager is dissimilar between scenarios, all other aspects are held equal. The intention is to compare the relative performance between the two scheduling techniques to determine if collaborative sensing can enhance the performance of a homogeneous SNSS.

4.3.1 RSO Catalogue

To demonstrate the effects of collaborative scheduling on a catalogue with a realistic distribution of RSOs, a larger and more diverse catalogue of simulated RSOs was compiled for this chapter's investigation. TLE data was once again obtained from NORAD's space-track website [39]. To maintain an acceptable run-time, the size of the simulation was limited to using 1000 TLEs to describe a realistic distribution of LEO, MEO and HEO RSOs. The distribution of semi-major axis lengths of these object's orbits are plotted in Fig. 4.5. Whilst the profile of the distribution is not a perfect match to the more comprehensive version produced in Fig. 2.1, it exhibits the same three characteristic groupings, as discussed in Section 2.1, in proportionate ratios. Similar to Chapter 3's methodology, the intention is to once again contrast alternative scenarios by observing how they minimise the level of error in a catalogue that begins with relatively large errors in RSO state estimates. To initialise the system, the TLEs are compiled into an RSO catalogue by firstly converting all objects' orbital elements into rectangular Earth-centred coordinates. This process results in a six element state vector containing position and velocity components, whereby the ith target in the RSO catalogue at time $k = 0$ has the truth-state vector

$$\mathbf{x}^i_{k=0} = [x \ y \ z \ \dot{x} \ \dot{y} \ \dot{z}]^{\mathrm{T}}. \tag{4.14}$$

An initial error covariance matrix $\mathbf{P}_{k=0}$ and state estimate $\hat{\mathbf{x}}_{k=0}$ are produced for each RSO by means of

$$\mathbf{P}^i_{k=0} = \text{diag}[{\sigma_p}^2 \ {\sigma_p}^2 \ {\sigma_p}^2 \ {\sigma_v}^2 \ {\sigma_v}^2 \ {\sigma_v}^2], \tag{4.15}$$

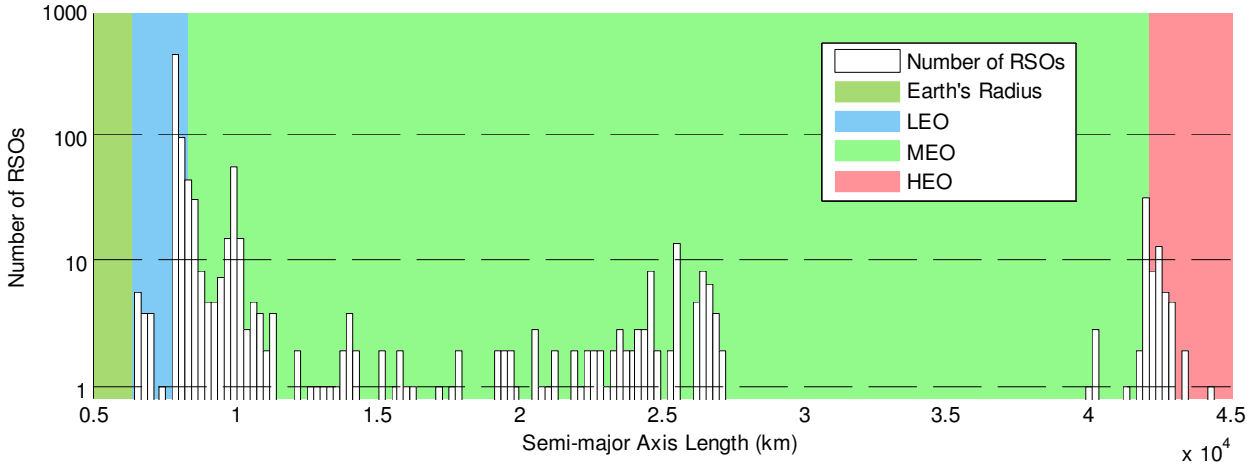

Figure 4.5: Histogram of the semi-major axis lengths of the orbits of the chosen 1000 RSOs.

and

$$\hat{\mathbf{x}}^i_{k=0} = x^i_{k=0} + \mathbf{L}\boldsymbol{r} \tag{4.16}$$

where $\mathbf{L}\mathbf{L}^{\mathrm{T}} = \mathbf{P}$ such that $\mathbf{L}$ is the Cholesky factorization of $\mathbf{P}$ and $\boldsymbol{r}$ is a vector of six i.i.d. $\mathcal{N}(0,1)$ random numbers. Standard deviations of $\sigma_p = 3.5 \times 10^3$ m and $\sigma_v = 1 \times 10^{-4}$ ms^{-1} were chosen due to their demonstrated stability in Chapter 3. Orbit propagation and determination of RSOs were conducted as detailed in Sections 3.3.3 and 3.3.5 respectively.

4.3.2 Sensor Network Configuration

MASSAS's sensor models were employed to generate simulated truth information about sensor-RSO visibility and for generating noisy measurements. For production of this investigation's results, hypothetical sensor sites were selected, involving six electro-optical sensor networks in two alternative sensor configurations. The USA sensor configuration consists of sensors positioned at Diego Garcia; Kwajalein Atoll; Haleakala, Hawaii, USA; Oregon, USA; Albuquerque, USA and West Virginia, USA. The global configuration consists of sensors located at Kwajalein Atoll; Albuquerque, USA; Moron, Spain; Bangalore, India; Betim, Brazil and Brisbane, Australia. The precise locations used to simulate each of the sensors may be obtained from Appendix D: Table D.1. The locations of these sensors have also been plotted on a world map in Fig. 4.6.

These two configurations were chosen to study the effects of forming small groups of sensors to permit collaborative sensing at low altitudes and, conversely, evenly distributing sensors about the globe to increase the separation of sensors and the minimum altitude at which collaborative sensing can occur. The relative performance, according to (4.13), of

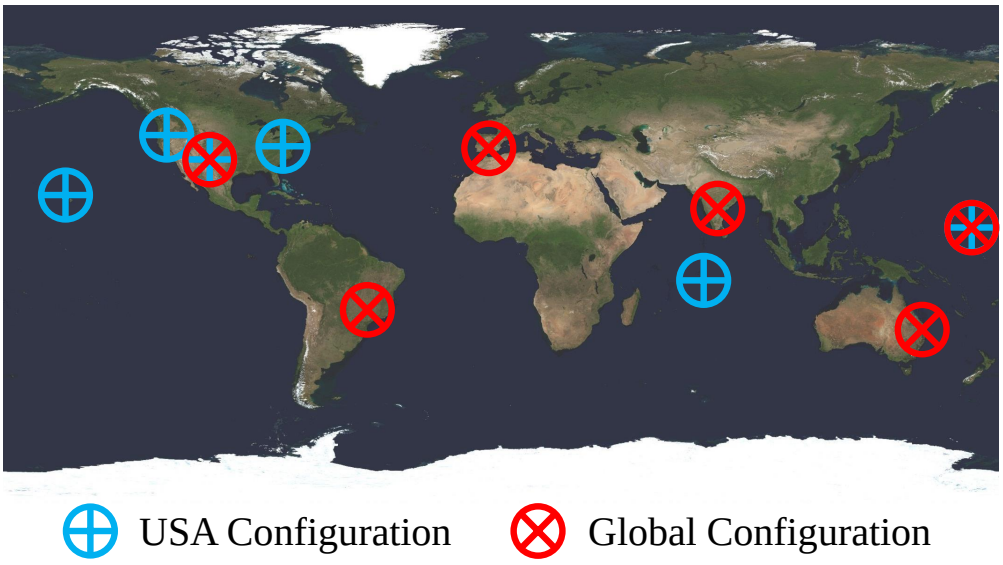

Figure 4.6: The locations of each configuration's sensors plotted on a world map.

each sensor configuration is displayed in Fig. 4.7, where the separations between each of the sensors in each configuration have been applied to (4.13) and overlaid on the relation previously illustrated in Fig. 4.4. Fig. 4.7 shows that whilst the USA configuration is capable of utilising collaborative sensing on most of the catalogue, the Global configuration has only a single sensor pair capable of collaborative scheduling at LEO altitudes. As approximately 80% of the catalogue is found within the mid-LEO to lower-MEO bound—as detailed in Section 2.1—the Global configuration is more restricted in the number and types of RSOs it is capable of observing via collaborative sensing. Whilst the use of small groups of sensors is favourable for simultaneous LEO observation, the results presented in Section 2.2.2, regarding global sensor coverage, motivated the placement of two of the USA configuration's sensors at Diego Garcia and Kwajalein Atoll to maintain adequate coverage of the higher orbiting RSOs.

During numerical simulation, sensors were limited to making 400 observations per scheduling period. If a sensor met its observation limit, MASSAS was instructed to ignore all further opportunities for independent and simultaneous observations for the applicable sensor. The use of collaborative sensor scheduling significantly increases the volume of the scheduling solution space, exponentially increasing the amount of computation. Fortunately, the computation of β_{red} as well as a number of other processes are well suited for parallel computation. For this reason, processes such as visibility analysis, observation effectiveness determination, orbit propagation, catalogue error analysis, observation simulation and orbit determination have all been parallelised on a GPU (graphics processing unit) for fast execution. The principles behind parallel computing for SSA are detailed in Chapter 5.

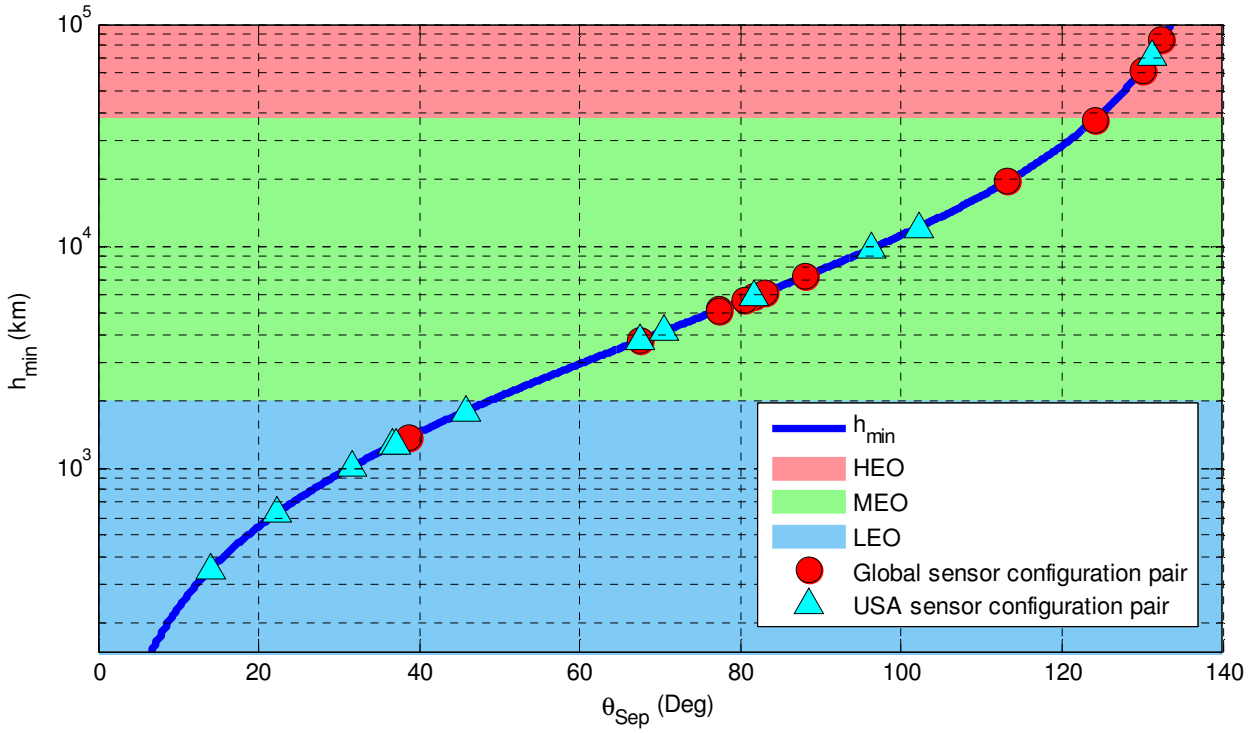

Figure 4.7: Minimum altitude $h_{\min}$ at which collaborative scheduling can take place for each sensor for each configuration.

4.3.3 Evaluation Metric

The metric chosen to represent the state of the RSO catalogue during the simulation is MCPE, as detailed in Section 3.2.5. Upon obtaining the daily MCPE for each scenario, the percentage *reduction* in MCPE due to the use of Scenario B—the collaborative scheduling method—in place of Scenario A—the independent scheduling method—is computed. Whilst this representation is convenient to monitor any improvement in accuracy due to the proposed method, each simulation's MCPE and WCPE were recorded individually. A plot of these results have been provided in Appendix C. To evaluate the effectiveness of the weighting applied via (4.12), Scenario B's collaborative sensing is also tested with and without the weighting applied via β_{w}. By comparing the performance of weighted and unweighted implementations, the anticipated effects of over-allocation of sensing resources may be directly observed.

4.4 Results & Discussion

Fig. 4.8 and Fig. 4.9 show the results of a 24-day simulation of each scenario, using each sensor configuration, with and without the weighting applied to β_{red}. Fig. 4.10 shows the

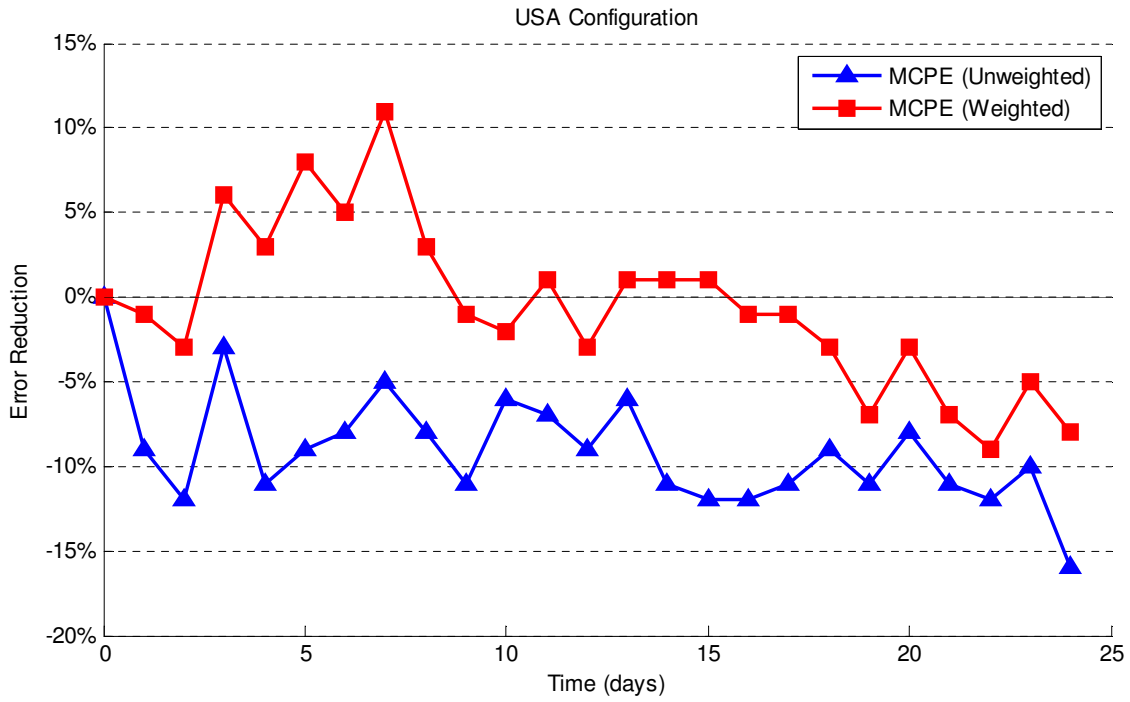

Figure 4.8: The percentage error reduction as a result of using collaborative sensor scheduling in place of independent scheduling, when using weighted and unweighted values for β_{red} to direct the USA configuration of sensors.

percentage of observations scheduled each day, for all configurations, that involved more than one sensor.

In almost all instances, Fig. 4.8 and Fig. 4.9 demonstrate that an unweighted β_{red} metric results in an increase in catalogue error when compared with the results of independent scheduling. Nonetheless, when the weighted metric β_{W} is employed, collaborative scheduling is observed to enhance catalogue accuracy. This result indicates that without a method for moderating the use of simultaneous measurements, application of a covariance-based observation effectiveness metric, such as β_{red}, to a collaborative scheduler, will likely produce an inferior level of catalogue accuracy than if sensors had been scheduled independently. This conclusion is further supported by Fig. 4.10 which demonstrates that multi-sensor observations were scheduled much more regularly for each SNSS configuration, when β_{red} was not weighted. This evidence supports our hypothesis that a balance must be struck between making too many or too few multi-sensor observations. Because multi-sensor observations involve two or more sensors that could have been observing multiple RSOs and each sensor has a finite schedule, overuse of information-rich multi-sensor observations will reduce the total number of RSOs observed per scheduling period, which may result in a negative effect on catalogue accuracy as a whole.

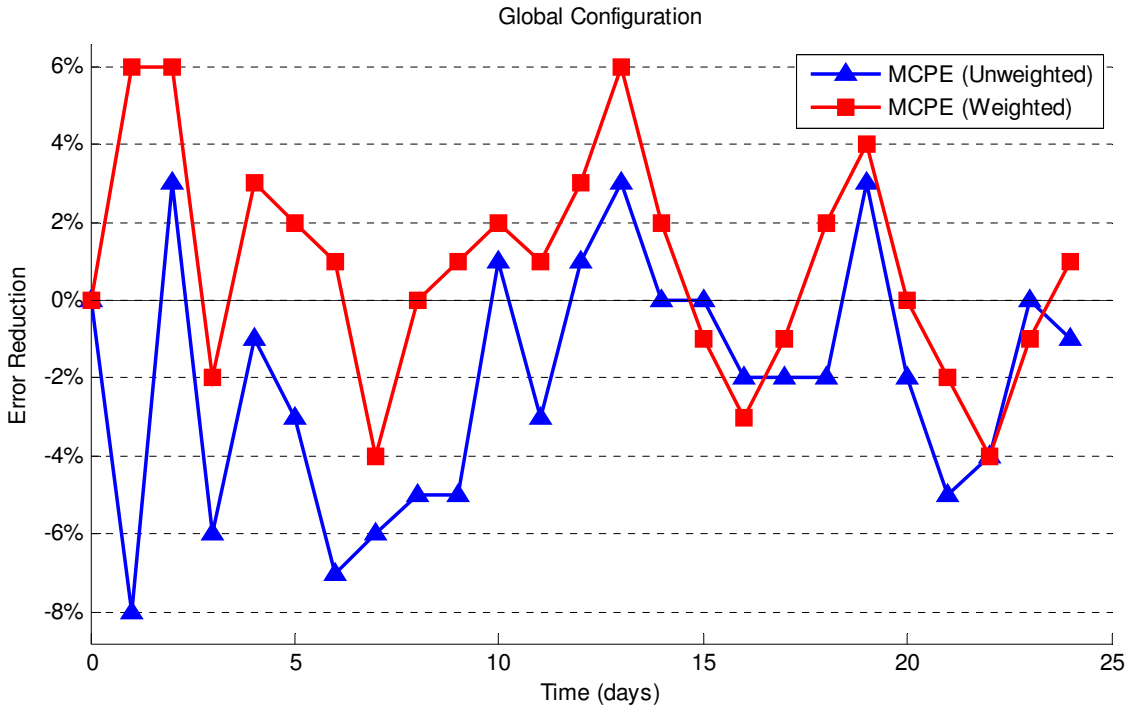

Figure 4.9: The percentage error reduction as a result of using collaborative sensor scheduling in place of independent scheduling, when using weighted and unweighted values for β_{red} to direct the Global configuration of sensors

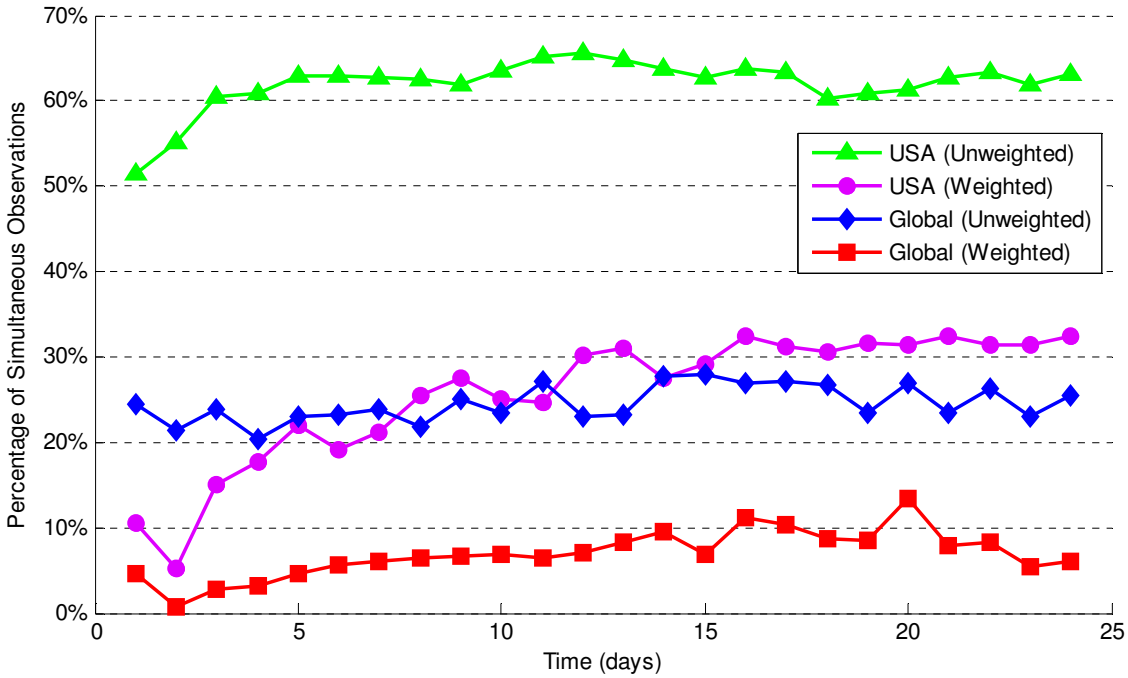

Figure 4.10: Percentage of multi-sensor observations scheduled per scheduling period for each simulation.

Whilst Fig. 4.8 and Fig. 4.9 indicate that the weighted collaborative-scheduler was capable of improving the level of accuracy of the catalogue, the improvement was inconsistent. Fig. 4.8, which presents the results of collaborative scheduling via the USA sensor configuration, displays the greatest peak reduction in error at approximately 11% on the 7th day. Nevertheless, the improved accuracy diminishes over the next 17 days ending in an increase in error of approximately 7% on the final day. The Global sensor configuration resulted in a smaller peak at a 6% reduction in error but the use of collaborative scheduling was more consistent at maintaining a positive influence on the catalogue. The effectiveness of the technique appeared to cycle between a positive and negative influence over a number of days. Encouragingly, the cycle was biased toward error reduction, averaging a 1% improvement over the 24 day period.

A decomposition of the recorded MCPE into radial, in-track and cross-track error components, for each weighted simulation, has been plotted in Appendix C: Figs. C.5 & C.6. Table 4.1 summarises this information. The average percentage reduction in error and the gradient of a least-squares linear regression of the data is presented for each component of error. These results indicate that over the 24 day simulation, weighted collaborative scheduling was, on average, detrimental to the catalogue's accuracy when using the USA configuration of sensors. Furthermore, as described by the gradients, collaborative scheduling became increasingly detrimental over the 24 day period. However, when using the Global configuration, all error components, but cross-track, showed improvement in catalogue accuracy. Whilst the average error reduction indicated limited improvement in in-track and radial error, the reduction in error was consistent for the duration of the simulation. This evidence suggests that collaborative scheduling may be of greater value to RSOs in higher orbits.

Table 4.1: Percentage reduction in MCPE error components throughout each weighted simulation.

Configuration	Statistic	Error Component			
		all	radial	in-track	cross-track
USA	Average	-0.6%	-4.3%	-3.4%	-1.0%
	Gradient	-0.5%/day	-0.5%/day	-0.4%/day	-0.5%/day
Global	Average	0.9%	1.0%	2.0%	-1.8%
	Gradient	-0.1%/day	0.0%/day	0.0%/day	-0.5%/day

Fig. 4.10 indicates a tendency of the collaborative-scheduler to increasingly schedule simultaneous observations as time progresses and the catalogue approaches steady-state. It is therefore proposed that controlling this increase could help sustain the improvements to

catalogue accuracy observed in Fig. 4.8 during days 3-8 and rectify the negative gradients observed in Table 4.1. Fig. 4.10 also highlights the influence the placement of sensors about the globe has on the scheduler's ability to take advantage of multi-sensor observations. Plots pertaining to the USA configuration, in which sensors are grouped more closely than in the global configuration, show significantly more multi-sensor observations.

4.5 Conclusion

A method has been proposed for scheduling and fusing simultaneous observations for improved catalogue accuracy when utilising a homogeneous SNSS comprising of electro-optical sensors. It was demonstrated through simulation that simultaneous observations can improve the accuracy of the RSO catalogue when using optical sensors capable of high accuracy angular measurement. Nonetheless, its use needs to be moderated to ensure multiple sensors are only used for observing a single object, when it is in the interest of the catalogue as a whole. The findings suggest that the inclusion of a weighted value for β_{red} is a positive step toward addressing β_{red}'s anticipated bias for employing as many sensors as possible to maximise error reduction. However the proposed method requires further refinement to obtain a consistent improvement to catalogue accuracy. A suggested avenue of improvement would be to pursue and utilise a metric other than β_{red}, that focuses on minimising catalogue error rather than maximising error reduction.

It was further proposed that effective use of collaborative scheduling requires the appropriate geographical placement of sensors. It was determined that using sensor networks with small groupings of sensors would increase the number of catalogued RSOs capable of being observed via simultaneous observation. Nonetheless, co-locating all sensors is clearly not optimal either as, for example, a number of geostationary objects may never be visible to a single small 'sensor cluster'. It is therefore necessary to strike a balance between clustering sensors with the intention of collaboration and distributing sensors for geographical diversity.

When sensors were evenly distributed about the globe, thereby restricting simultaneous observation to a subset of the catalogue at MEO and GEO altitudes, the proposed method appeared to more consistently improve catalogue accuracy. This may suggest that collaborative scheduling may be better suited to RSOs at higher orbits. A possible explanation is the fact that oblique optical measurements are more likely at LEO altitudes than GEO, thereby supplementing the lack of range information from optical measurements. Consequently, multi-sensor optical-measurements may be better suited to higher orbiting RSOs as they are likely to exhibit higher levels of radial error. Further testing is nonetheless required to substantiate this hypothesis.

5

GPU-Based Prediction of Non-Observability

Chapter 3 detailed how and why the SSA simulation named MASSAS was developed for performing the research presented throughout this thesis. It was acknowledged that while research topics of interest to the SSA community, both past [80–83] and present [22], require varying levels of fidelity and sophistication, MASSAS is intended to provide a middle-ground between simplified analytic-analysis of fundamental principles and a high-fidelity simulation environment requiring vast resources and supercomputing. By doing so, MASSAS can exploit the reduction in resources that reduced fidelity permits, whilst permitting the analysis of intricate concepts that are difficult to describe and predict via closed-form expressions.

Locating the appropriate middle ground between each approach is dependent on supplying sufficient fidelity to the primary topic or topics of research, to ensure meaningful results are obtained. A simulation's reliability and validity must be suitably balanced against runtime. Many physical processes may be modelled in progressively increasing levels of fidelity by introducing more, albeit less significant, physical phenomena. The expectation is therefore, that the reliability of the simulation is proportionate to the runtime. Consequently, a judicious selection of physical phenomena requiring simulation, may produce an acceptable runtime whilst maintaining an acceptable level of numerical error. This principle has been demonstrated by MASSAS as discussed in Chapter 3. Nonetheless, when the scale of the physical process is a crucial element under investigation, it is more challenging to minimise

the runtime whilst achieving meaningful results. This concept was demonstrated in the transition between Chapters 3 & 4. Five times as many RSOs were used in Chapter 4 as in Chapter 3 to achieve a realistic distribution of RSOs. Furthermore, the scheduling technique employed caused the solution space to scale exponentially by the number of sensors in the SNSS. As a single sensor was insufficient to observe the distribution of RSOs, the simulation was forced to scale to accommodate at least the minimum number of sensors necessary to achieve global coverage.

As demonstrated by such simulations as AFRL's supercomputer-based TASMAN [50], achieving an arbitrary level of fidelity and scale is solved by increasing the amount of processing power. This principle is believed to be the root of the large gap in computational complexity observed in SSA research techniques. Namely, unless the problem can be simplified in a useful manner or supercomputing is a viable option, certain areas of interest have received limited attention by researchers in the field [22]. Nonetheless, recent developments in computational science have a potential answer to this problem that enables a simulation such as MASSAS to maintain its low-resource, rapid research objectives whilst incorporating higher degrees of scale and fidelity.

Recent developments in GPGPU (General-purpose Computing on Graphics Processing Units) hardware and associated application programming interfaces (APIs) enable off-the-shelf consumer-grade graphics hardware to perform the parallel computing necessary for much more ambitious single-PC SSA simulations [84, 85]. In this chapter, the principles of GPGPU [86] are employed to increase the computational power of MASSAS via parallelism, whilst maintaining its independence and flexibility. Making use of GPGPU is, however, not without challenge. Mathematics libraries taken for granted in other languages are not yet routinely available for GPU processing. Parallel error trapping and the transportation of information to and from the GPU also require consideration. The use of highly parallel computation on a GPU requires a different approach to conventional algorithm design, requiring a trade-off between computation and communication with memory. To demonstrate, the augmentation of MASSAS with highly parallel GPU-executed code for enabling large-scale, single-PC simulation is detailed throughout this chapter.

In addition to detailing the application of GPGPU to MASSAS, an inaugural application of the parallelised system is presented. Sensor availability, throughout an SNSS, is investigated using a full-scale RSO catalogue. The investigations show that the sensor types used throughout the SNSS affects availability not only in terms of the amount of time a sensor is capable of observing RSOs per tasking period, but also by influencing the types of RSOs each sensor type is capable of regularly observing. It is found that when employing practical surveillance networks composed of mostly optical sensors, certain RSOs may not be visible to

any sensors for as long as days at a time. This outcome has motivated further enhancement of the scheduling techniques discussed in previous chapters via visibility forecasting. With the enhanced computational power afforded by the use of GPGPU, visibility forecasting is utilised to increase the time horizon of scheduling techniques. Incorporation of visibility forecasting enables the sensor manager to anticipate and prepare the RSO catalogue for periods when certain RSOs cannot be observed. This example is intended to demonstrate that GPGPU is not only a valuable tool for SSA related research, but that it also has great potential for use in existing sensor management computers and for improving the intelligence of space sensors. By enabling such systems to efficiently process large quantities of data via GPGPU, they are able to make more informed decisions, without a large overhead with regard to cost and ancillary resources. This potential is explored further in Chapter 6.

Section 5.1 begins by providing an overview of GPGPU and what it offers in an SSA context. Subsequently, Section 5.2 describes the parallelised simulation system and highlights a number of areas in which the GPGPU architecture influenced its development and functionality. Utilising the improved architecture, Section 5.3 provides an analysis of SNSS availability for observing the RSO catalogue as the number and type of sensors are varied. Section 5.4 builds upon these results to describe and simulate further enhancements to the observation effectiveness scheduling of previous chapters. The concluding Section 5.5 discusses this chapter's findings as well as the suitability of continued research and deployment of GPGPU-based techniques for improving SSA.

5.1 GPU Computation

GPGPU is a method for increasing the computational performance of a PC for scientific and engineering applications by utilising a GPU for large scale parallel processing. Although the foundations for GPGPU were laid in the 1970s [87], GPGPU has seen a slow adoption rate until around 2003 when CPU manufacturers began struggling to maintain the rate of increase that chip-speeds had experienced over the preceding decades [85]. In the short term, CPU manufacturers have circumvented this problem, to some extent, through the introduction of multicore processors which allow small scale computational parallelism. GPGPU nonetheless enables much greater levels of PC-parallelism by exploiting the parallel processing architecture found in conventional GPUs, for purposes other than producing computer graphics. In 2006, graphic computing company NVIDIA released a versatile application programming interface (API) and software development kit (SDK) for the express purpose of enhancing accessibility to GPGPU. Although alternative GPGPU architectures are available [85], MATLAB's compatibility with NVIDIA's well established Compute Unified Device

Architecture (CUDA) [88] made it the obvious choice to augment and enhance MASSAS's computational capabilities.

5.1.1 GPGPU Augmented SSA Simulation

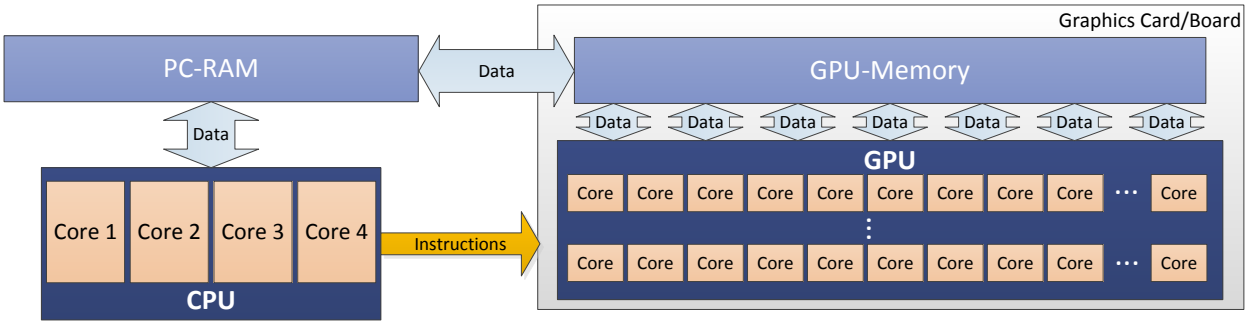

Figure 5.1: High-level GPGPU system architecture.

NVIDIA's CUDA, like many other GPGPU APIs, permits the transportation of data to and from a PC's conventional RAM and GPU memory for subsequent computation using specialised GPU-executable programs. Fig. 5.1 illustrates the high-level architecture of such a system. Applications such as MATLAB permit high-levels of abstraction of GPGPU programs. Nonetheless, if desired, MATLAB also permits the compilation and execution of relatively low-level computer code for greater versatility, control and efficiency. This process is achieved by compiling CUDA compliant C code via NVIDA's nvcc compiler. The nvcc compiler enables programs to be written for execution on both the 'host' PC and the GPU 'device' [84]. This enables sequential and parallel tasks to be executed on the appropriate processing device—Central Processing Unit (CPU) or GPU—for efficient execution, whilst enabling the developer to write a single computer program. Nonetheless, to maintain MASSAS's flexibility afforded by MATLAB-based programming and decrease the time necessary for porting MASSAS to a new language, only tasks that can benefit from parallel execution were ported to nvcc compliant code. The remaining aspects of MASSAS continue to run from their original MATLAB scripts. The only changes that were necessary to these scripts were the replacement of the calls to the superseded sequentially-executed functions with calls to and from GPU memory and execution of the replacement pre-compiled device code.

Device code, intended for parallel execution, is compiled as a 'kernel' to be run on the GPU's parallel processors as one of many 'threads' [84]. Each thread executes the same kernel however, each thread has a unique identity or index enabling the programmer to address different parts of GPU-memory relative to the thread index. In theory, this enables

a single task, which would typically reside in one or more nested for-loops for CPU execution, to be applied to each of the values in a large array or matrix, as a single parallel execution. In reality the GPU has a finite number of parallel cores and therefore has a limit to the amount of parallel execution. Nonetheless, as described more completely in the CUDA programming guide [84], CUDA's Scalable Programming Model separates groups of threads into 'blocks'. Each GPU has a finite number of blocks it can execute in parallel. If there are more blocks than it is capable of processing in parallel, the GPU will process groups of blocks sequentially. Each group will comprise of as many blocks as is possible to execute in parallel by the device's specific hardware. Therefore, regardless of the size of the GPU, the task is parallelised as much as the hardware permits.

Although utilising even a modest GPU can achieve very high numbers of operations per second, the advantage over sequential CPU computation is only gained if the problem is parallelisable on a large scale. For instance, if an application only needs to process four independent streams of data, it is likely to be a disadvantage to perform the calculation on the GPU. Whilst in theory, equivalent device clock speeds would permit the process to execute four times faster, time is lost while transferring the applicable data to and from the GPU's memory before and after execution. Consequently, the time saved via parallel execution must be sufficient to warrant the overhead of transferring the data to and from the GPU to obtain any kind of advantage. Depending on the specific implementation, the break-even point between CPU and GPU computation may require hundreds or perhaps hundreds of thousands of parallelised tasks before GPGPU becomes an advantage [85, 86]. This means GPGPU is not appropriate for all scientific and engineering tasks. Fortunately for SSA management and research, because of an RSO's comparably minor mass when compared to the Earth and each RSO's relatively large separation with respect to their small gravitational sphere of influence, the fundamental problem can be reduced to many individual sensors observing a large population of independently orbiting RSOs. This large-scale system-independence permits effective parallelisation and is therefore well-suited for GPGPU. Furthermore, even if the increased process error due to independence assumptions is deemed unacceptable, work by Nyland et al. [89] demonstrates the GPU's ability to also enhance the speed of n-body—each body influences every other body—orbit propagations.

In addition to independence of RSO motion and sensor tasking, under certain conditions SSA-related tasks can also be parallelised in time if closed-form orbit propagators are suitable. Large-scale parallelisation of tasks such as visibility analysis, observation effectiveness determination, orbit propagation, catalogue error analysis, observation simulation and orbit determination are therefore tractable, making SSA management and simulation an excellent match for GPGPU. For these reasons, all of the aforementioned SSA applications have been

implemented as GPU kernels and used to enhance the speed and scalability of MASSAS.

As the exact speed increase is intricately related to the precise balance between algorithm, memory usage, hardware, API version and scale, it is superfluous or misleading to quote or promise a specific speed increase for all SSA applications by using this method. Nonetheless, to provide at least a vague indication, the run-times of the simulations used to produce the results presented in Chapter 3 were compared with the run-times of MASSAS after the adoption of GPGPU. Using only a modest consumer-grade graphics card, an equivalently scaled simulation is observed to run at least 100 times faster. But as discussed earlier, this value improves with scale. Furthermore, as MASSAS is an experimental system, GPU utilisation is difficult to optimise for all possible research applications. It is therefore fair to assume that a mature system utilising high-end GPUs, perhaps implemented in an SSA management centre, could expect even higher levels of speed-increase.

5.1.2 Pitfalls of Adopting GPGPU

Although GPGPU holds great promise for SSA management, there are a number of limitations and hurdles that must be judiciously managed to achieve an effective system. As mentioned previously, parallelised tasks require large scale independence primarily due to memory transfer delays. This means that although many SSA-related tasks are suitable for parallelisation, ideally the runtime of CPU and GPU implementations of each element should be compared, to select the most appropriate hardware for the task and ensure a net gain is achieved. In general, tasks involving large amounts of data but only a small amount of parallelism and processing are best left to one or more CPU cores. Nonetheless the overhead for transferring large amounts of information to and from the GPU can be alleviated in some cases. By recomputing known values and leaving frequently accessed arrays and matrices of data stored within the GPU's memory, the transference of information may be minimised. In the case of SSA simulation, leaving sensor and RSO state information in GPU-memory has proved to be advantageous. This methodology however presumes the GPU has sufficient memory to store the information throughout the simulation.

Additional limitations to GPGPU are a consequence of the late-blooming of the technology. Although basic math libraries have been included in nvcc, higher level libraries taken for granted in well-established programming languages are yet to be transitioned into GPU-device executable code for single thread execution. For this reason a tailored matrix-mathematics library ranging from basic matrix operations to more sophisticated matrix inversion and Cholesky decomposition had to be developed to produce the simulated results presented later in Sections 5.3 and 5.4 of this chapter. Furthermore, to permit the transfer

of debugging and error information from the GPU, additional and often redundant data had to be transported with each parallel execution due to a lack of structured error handling. CUDA code is also hardware specific and needs to be compiled with a minimum architecture version in mind in order to take advantage of the benefits of the version's attributes. For example, double precision processing, which is crucial for many aspects of MASSAS, has only been possible in recent versions of CUDA and its compatible hardware. In addition, the double precision arithmetic does not fully comply with the Institute of Electrical and Electronics Engineers (IEEE) standards used by many CPU manufacturers [90]. This may lead to slight differences between the results generated by CPUs and GPUs. Nevertheless, the aforementioned pitfalls are manageable and it is anticipated that as the technology matures, its idiosyncrasies will diminish in prominence.

5.2 Simulating Sensor Availability

The RSO-visibility-prediction topics described in Section 2.2.4 shed some light on the regularity at which a terrestrial space-surveillance sensor can expect to see certain RSOs. Nonetheless, many assumptions were necessary to arrive at the result and they did not factor in limitations imposed by utilising passive sensing techniques, such as when making electro-optical measurements. Therefore to explore this topic in more detail and as a demonstration of MASSAS's scalability via GPGPU, the remainder of the chapter details the implementation of a full-scale simulation using as much RSO data as is available for public use. The objective is to better understand the practicalities of utilising sensor networks with a heavy dependence on optical sensing. In particular, we hope to learn how such a dependence might impact the accuracy of catalogued RSOs' state estimates. The primary method of investigation entails analysing the ability of a simulated SNSS to maintain regular observation of a full-scale RSO catalogue whilst the number and type of sensors are varied. Consequently, the results are not necessarily applicable to other roles for which an SNSS may be tasked, such as searching for previously uncatalogued RSOs.

This chapter's investigation of sensor-availability, continues to utilise many of the procedures established in the previous chapters. When appropriate, these procedures have been parallelised for fast GPU computation.

5.2.1 Full-Scale RSO Catalogue

To obtain a full-scale catalogue, simulated objects were obtained from genuine, publicly available TLE data on the Space Track.org website [39]. The file `all_sat_191.zip`, containing all of the latest unclassified TLE updates, recorded no earlier than 30 days prior to the

191st day of the year 2012 (9th July), was downloaded and processed to extract the mean orbital elements of each of the recorded objects. Via this method, 14 876 unique objects were obtained. The number of RSOs published as in-orbit TLEs fluctuate on small time scales due to re-entry, new launches and in-orbit events. On larger time-scales, as detailed in Chapter 1, the average number of objects continues to rise. Fig. 5.2 shows the distribution of the objects obtained from the TLEs with respect to the semi-major axis length of each RSO's respective orbit.

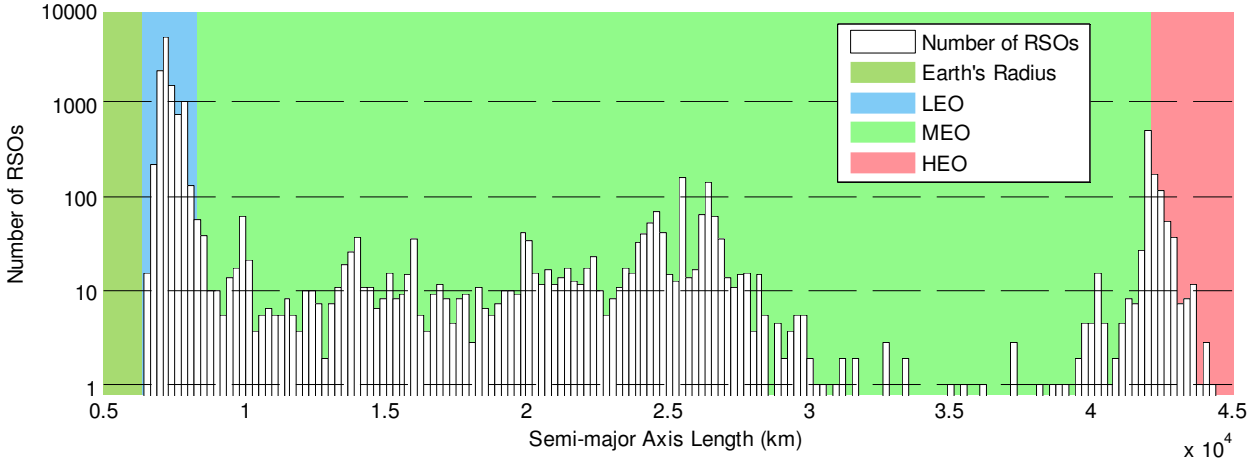

Figure 5.2: The frequency of all catalogued RSOs verses their orbit's respective semi-major axis length.

RSO catalogue initialisation, orbit propagation and orbit determination were implemented as detailed in Chapters 3 & 4.

5.2.2 Sensor Modelling

It is assumed a minimum elevation of 20°, line-of-sight and a maximum range is all that is necessary for determining visibility by an active sensor such as radar. In addition to these requirements, passive sensing via optical sensors will again employ MASSAS's solar illumination and optical sensor model.

Due to their prevalent use for space surveillance and the limitations on other sensor types discussed in Chapter 1, only radar and electro-optical terrestrial sensors will be considered for our current investigation. Sensor modelling parameters and scripts were transitioned to the GPU for efficient visibility determination, observation effectiveness determination and scheduling. MASSAS's optical sensor models are configured to measure elevation and azimuth while the radar models are configured to measure range, elevation and azimuth. In genuine space surveillance networks, the measurement characteristics and range limitations

of a particular type of space surveillance sensor often varies from sensor to sensor [9,31]. To attain a higher level of generality than if the specific sensor characteristics of any particular SNSS were employed, MASSAS's radar and optical sensor models will continue to be configured using homogeneous characteristics for each sensor type. The configuration parameters used by MASSAS's sensor models are summarised in Table 5.1.

Table 5.1: Sensor measurement characteristics.

Sensor Type	Range Error Std. Dev.	Angular Error Std. Dev.	Maximum Range
Optical	-	1 arcsecond	∞
Radar	30 m	50 arcseconds	8300 km

The maximum range of simulated radar measurements was selected by averaging the operational characteristics of genuine space-surveillance radars [9]. Whilst the sensing capabilities of electro-optical sensors depend on range, the optical sensor model has not been given a range limitation. Instead, as the Catalogue is composed of objects that have been successfully observed by genuine sensors and as most orbits are circular—see Section 2.2.4—for the purposes of this chapter's investigations, optical sensors are assumed to be capable of observing any object in the catalogue. For even greater generality, near-uniform positioning of sensors about the globe has been chosen to minimise any geographical dependencies in the resulting data. The precise locations and sensor types used during the simulations detailed throughout Sections 5.3 and 5.4 are presented in Appendix D.

5.2.3 Sensor Scheduling

The simulations detailed in Section 5.3 utilised the centralised sensor-management network-topology discussed in Chapter 3, employing the β_{red} observation effectiveness scheduler as described in Sections 3.2.1 & 3.2.2. The sensor manager holds the RSO catalogue, controls the global sensor network and collates sensor observations to update the catalogue. The sensor manager controls each sensor by creating and distributing a schedule for each sensor. The schedule instructs each sensor when to observe specific RSOs during each upcoming 24 hour scheduling period. Section 5.4 discusses the use of a similar sensor management process but explores the use of an alternative scheduling methodology to improve catalogue accuracy.

The evaluation of RSO-sensor visibility and β_{red} are both examples of independent processes. RSO visibility was therefore pre-computed on the GPU for the entire scheduling period. Subsequently, the visibility data was used in conjunction with the GPU to evaluate the initial $\boldsymbol{\beta}_{\text{red}}$ solution space and re-evaluate β_{red} as observations were scheduled. Because

re-evaluation of β_{red} modifies the solution space and it is assumed sensors can only observe one RSO at a time, scheduling of each observation is dependent on the previously scheduled events. The scheduling of each observation is thus an inherently sequential process, unable to exploit the GPU's architecture. The scheduling process is consequently the slowest process due to its lack of parallelism. This trait means the simulation runtime is strongly influenced by the number of sensors and the cumulative sum of the number of observations permitted per scheduling period. But, due to parallelisation of many other processes, it is very weakly influenced by the size of the catalogue. Consequently, the implementation is favourable for simulating a system involving a comparatively small number of sensors observing a large number of targets.

During Chapter 3's disparate sensor analysis, each sensor was limited to making 200 track updates per RSO per sensor per scheduling period. The simulations detailed in this chapter were operated without this restriction, causing the number of observations to be influenced only by the visibility characteristics of the network's sensors and the scheduling framework. Chapter 3's analysis also used a 120 second tracking period, during which each sensor was restricted to making a single track update of a single RSO. This scheduling framework has been retained, so it should be noted that the methodology employed does not attempt to accurately represent the number of track updates capable of being made per scheduling period. This is particularly significant in the case of conventional radars observing RSOs in LEO. They are generally capable of quickly sweeping their surveillance volumes and observing many RSOs in a relatively short time scale. The framework is nonetheless retained for two reasons. The first is that the 120 second window is used to efficiently discretise the scheduling period for parallel computations on the GPU. The visibility of each object by each sensor is computed and stored for each increment in time. The second motive is that the precise time necessary to make a measurement of a particular RSO is highly circumstantial. To accurately model the process, knowledge involving precise sensor characteristics, modes of operation, power settings, each RSO's physical characteristics as well as the state of steerable sensor's mounts and their respective dynamics would be required. As the time that radars and optical sensors take to observe an RSO can vary widely, the approach taken in this investigation is to maintain generality by assuming both sensor types take equally long to observe any particular RSO. The precise numbers of RSOs that could be observed per day will therefore be inconsistent with practical networks. Consequently, the precise number of sensors necessary to achieve regular observations of a finite number of RSOs is beyond the scope of this study. The relative number of track updates per scheduling period will instead indicate the cumulative period of time an SNSS's sensors have visibility of any particular catalogued RSO and are thereby able to contribute to catalogue maintenance.

In light of this approach, another aspect that is not modelled in this study is the reduction in availability of a sensor due to maintenance and weather. Including this level of detail in the current study increases the result's dependence on the precise type and location of the sensor. As detailed in Section 5.2.2, such a dependence works in opposition to our objectives. Rather than evaluate and discuss the merits of a specific SNSS, we aim to investigate a method for analysing various combinations of sensors. Modelling weather as well as scheduled and unscheduled maintenance would nonetheless be useful for operators to perform in a similar analysis concerning their specific SNSS.

5.2.4 Orbit Determination

Once observations have been scheduled and simulated, state estimates are augmented with new observation data by producing a best-fit estimated state via an appropriate nonlinear recursive filtering process. The UKF [54] was again adopted for this role for its proven performance in previous chapters and favourable tracking attributes as discussed in Section 2.3.2. All scheduled track updates comprise of five independent observations of the target RSO, taking a total of 120 seconds to complete. A non-trivial observation queuing process enables simulated data to be applied in a catalogue-wide parallel update. After scheduling, the queuing process obtains the data necessary for simulating the next observation of each RSO. This information can thereafter be applied in parallel, permitting a fast catalogue-wide update for the entire tasking period. This process further exploits the GPU's parallel architecture but ensures multiple UKF state-updates of single RSOs are applied in chronological order. Current state, covariance and update-epoch information remains on the GPU for the entirety of the simulation to enhance efficiency by minimising the transference of data.

5.3 Surveillance Network Availability Analysis

The availability analysis detailed in this section aims to determine how the availability of a space surveillance network is influenced when the number and type of sensors it employs are varied. In doing so, we aim to determine if it is possible to construct an SNSS using only or primarily optical sensors whilst retaining the level of availability of more traditional configurations.

The availability of a space surveillance network is appraised in two important areas:

1. the cumulative amount of time its sensors are capable of making observations; and

2. the regularity at which each catalogued RSO can be observed by the SNSS.

The analysis is conducted by analysing the simulated data of 8 alternative network configurations involving a combination of 21 sensor locations and 2 sensor types. The precise configuration details of Simulations 1–8 are presented in Appendix D.

5.3.1 Homogeneous Networks

To contrast the differences in availability of an SNSS for maintaining the RSO catalogue when either radars or optical sensors are employed, we begin the availability analysis by performing simulations involving the use of homogeneous sensor networks. Simulations 1 & 2 involve the scheduling of 8 sensors over an 8 day period. The catalogue's 8th day MCPE and WCPE of each simulation are displayed in Fig. 5.3 and Fig. 5.4.

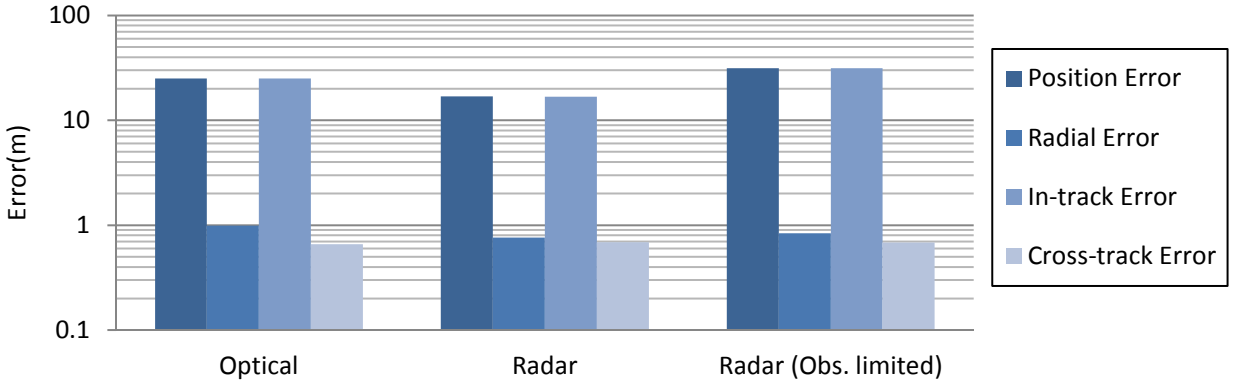

Figure 5.3: Simulations 1, 2 & 3: 8th day MCPE when using an 8-sensor homogeneous surveillance network.

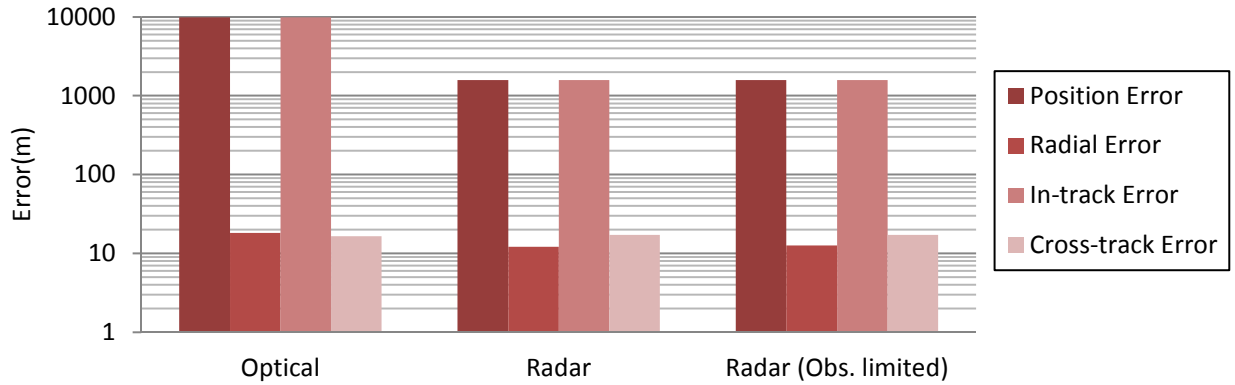

Figure 5.4: Simulations 1, 2 & 3: 8th day WCPE when using an 8-sensor homogeneous surveillance network.

Fig. 5.3 and Fig. 5.4 display the magnitude as well as radial, in-track and cross-track components of the MCPE and WCPE for each simulation. Contrary to the homogeneous

network results presented in Section 3.4.3 which favoured optical measurements, Fig. 5.3 and Fig. 5.4 indicate that the homogeneous radar network has achieved the superior result. The simulations detailed by this section differ from Chapter 3's simulations in two significant ways. A much more realistic distribution of RSOs is used and a limit for equalising the number of track-updates performed by each sensor type per day was not applied. We look to these aspects to ascertain the cause of the disparity.

Table 5.2 displays the average number of track updates made by the simulated SNSS each day, using each sensor type. These values indicate that the visibility constraints of the optical sensors have halved the number of track updates that were possible when using a network of radars. This is not unexpected as the optical sensor model has been configured to simulate the fact that, except in certain circumstances [49, 91], optical sensors require night-fall to operate. The season will therefore govern how much of the day the sensor will spend in daylight, inhibiting observations. Whilst the precise amount of time an SNSS's optical sensor will spend in day-light will vary throughout the year, as detailed in Appendix D, the sensors used in this simulation are roughly balanced between the Northern and Southern hemispheres. This in turn balances these seasonal effects, reducing the variation in availability of all optical sensors, on average, throughout the year. To verify that the reduced availability of the optical sensors was the cause of the disparity between the results presented in Section 3.4.3 and Figs. 5.3-Fig. 5.4 a third simulation was performed using the homogeneous radar network. However, during this simulation, each radar was restricted to make 360 observations per day, thereby matching the number of optical observations, albeit as closely as possible. The results of this third simulation are also presented in Fig. 5.3 and Fig. 5.4 for ease of comparison.

Table 5.2: Comparison of the average number of track updates made by eight homogeneous sensors per day.

Sensor Type	Avg. Updates/Day
Optical	2881
Radar	5760

Simulation 3's results present conflicting information. The level of MCPE aligns more closely with Chapter 3's conclusions. As observed previously, the optical network's in-track error is superior to that of radar and, also in agreement, its radial error is slightly worse. This correlation suggests that the relative number of updates, at least in part, explains the shift in superiority. The resulting WCPE however, appears entirely unaffected by limiting the number of radar observations. To investigate the cause of this inconsistency, we look to

an analysis of the regularity at which certain types of RSOs could be observed during each simulation.

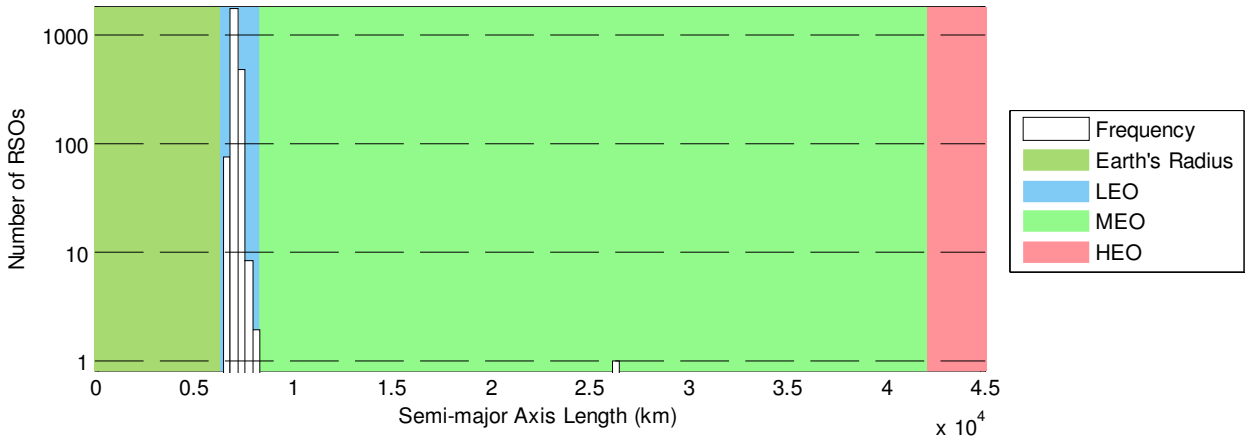

Figure 5.5: Simulation 1: The number and distribution of RSOs that were invisible to an all-optical SNSS throughout an 8 day simulation.

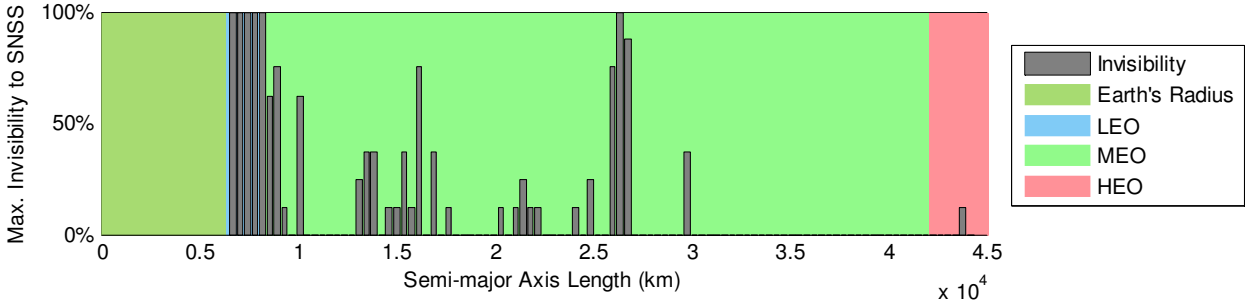

Figure 5.6: Simulation 1: The interval maximums of the average percentage of scheduling periods each RSO was invisible an all-optical SNSS throughout an 8 day simulation.

By collating sensor-RSO visibility information that was recorded during each of the simulations, objects that were difficult to observe using each sensor type can be identified. Fig. 5.5 displays a histogram of the RSOs that were invisible to the electro-optical sensor network throughout Simulation 1, with respect to their semi-major axis of orbit. In addition, Fig. 5.6 displays the results of a complementary analysis that indicates how often certain types of RSOs were capable of being observed. Each vertical bar indicates the maximum percentage of tasking periods in which RSOs belonging to each discrete interval of semi-major axis length were invisible to Simulation 1's SNSS. In total, 2489 RSOs remained invisible to the SNSS for the entire simulation. All but one was located in the lower half of LEO. The non-LEO RSO was verified via its TLE to be a single Molniya satellite. Our previous investigation regarding the regularity at which we can expect RSOs to transition a terrestrial sensor's surveillance volume—in Section 2.2.4—correlates well with these results.

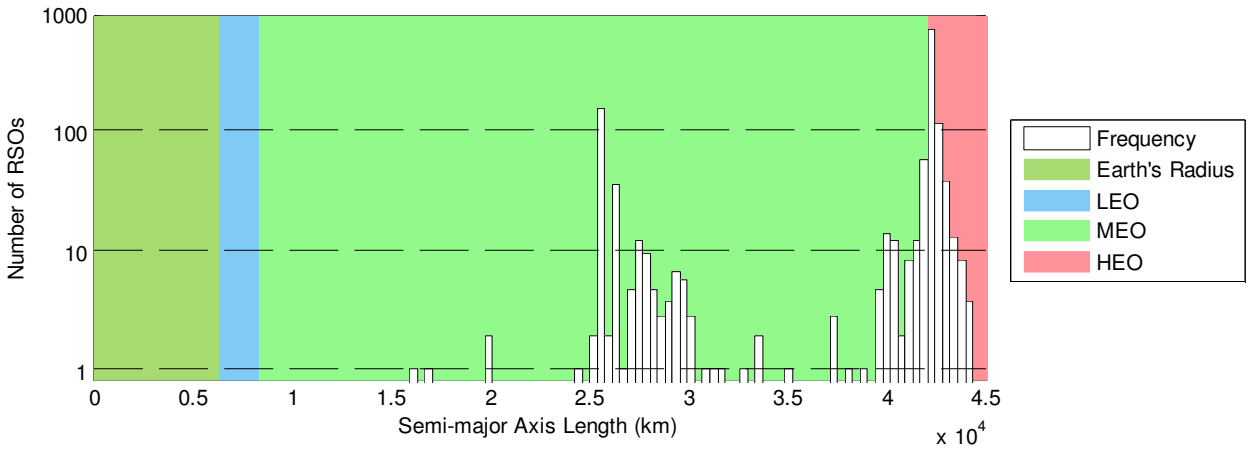

Figure 5.7: Simulations 2 & 3: The number and distribution of RSOs that were invisible to an all-radar SNSS throughout an 8 day simulation.

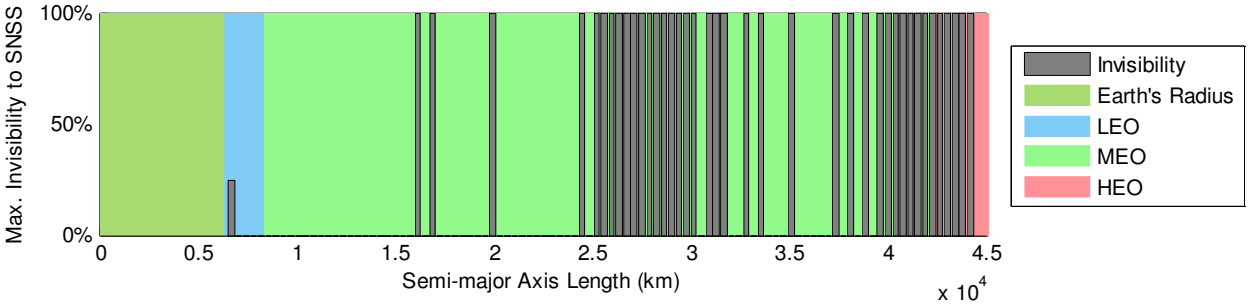

Figure 5.8: Simulations 2 & 3: The interval maximums of the average percentage of scheduling periods each RSO was invisible to an all-radar SNSS throughout an 8 day simulation.

The correlation allows us to infer that the LEO RSOs that were invisible to the SNSS were difficult to observe, because these RSOs are rarely within the surveillance volume of any sensor. Due to the added restrictions to visibility as a result of passive sensing, the chances of these LEO objects residing within an optical sensor's surveillance volume, while each sensor is in nightfall and while RSOs are not in eclipse, are very low. As a consequence, many LEO objects are not observed during the 8-day period. A peak of limited-visibility is also observed in Fig. 5.6 at mid-MEO altitudes. In agreement with the findings of Section 2.2.4, the reduced visibility in this region correlates with, and was verified to be, the notoriously difficult to observe Molniya objects.

Whilst a dip in observability near HEO was also predicted in Section 2.2.4, the objective was to predict the visibility constraints of a single sensor. Such constraints are absent from our current analysis as regular observation of near-HEO objects has been achieved by distributing a network of sensors about the globe.

Fig. 5.7 and Fig. 5.8 show the equivalent plots for Simulation 2 & 3's network of radars.

These figures indicate that the radars were much more effective at observing LEO RSOs than the optical sensors. This is not unexpected as the moment an RSO enters a radar's surveillance volume, in spite of daylight or eclipse, it is capable of observing it via active sensing. Nevertheless, each radar's range limitations have prevented the observation of objects whose perigees were not below 8300 km in altitude. Consequently Fig. 5.8 displays a somewhat binary result with respect to RSO visibility. In total, 1390 RSOs remained invisible to the radar-dependant SNSS throughout the simulation. This data provides the final evidence we require to provide an explanation as to why Simulation 2's radar network achieved superior catalogue accuracy over Simulation 1's optical network, in spite of our contradictory conclusions in Chapter 3.

The culmination of these analyses indicate that, on average, a sensor network comprising of electro-optical sensors will have half of the availability of an equivalent radar network to update an equally visible RSO catalogue. The physical characteristics of the surveillance environment and sensor types make optical sensors better suited to observing higher orbiting objects whilst radar is well suited for observing lower orbiting objects. Reconsidering the aberrant WCPE values observed in Fig. 5.4, it is now apparent that the worst-case results are likely attributed to the inability to perform track-updates for certain RSOs rather than due to sensor measurement characteristics, as large numbers of RSOs could not be observed. This outcome, along with our prior conclusions, suggest two possible explanations for the higher levels of WCPE when reliant on optical observations, that are not necessarily mutually exclusive. The first is that LEO objects experience a higher rate of growth in orbital-error than HEO objects over the same period of time, thus causing Simulation 1's WCPE to be worse than Simulation 2's. Alternatively, as we are recording the worst-case value of a stochastic process, we are more likely to obtain a larger outlier for a large population when compared to a small population. As there are many more RSOs in LEO than in mid-MEO to HEO altitudes, it was more likely that larger outliers were obtained in Simulation 1 than in Simulation 2. Both of these aspects were absent in Chapter 3's analysis. The RSO catalogue contained objects with almost identical orbits and every RSO was observable by each sensor type. The objective is now to exploit this knowledge to improve the effectiveness of an SNSS that is highly dependent on electro-optical sensors.

5.3.2 Optically Biased Networks

So far we have found that the types of sensors employed by an SNSS predominantly affect the system's availability by influencing the number of track-updates per day and the types of RSOs visible to the network. Due to the relative costs of implementing and operating each

sensor type, the reduced number of track-updates per day when utilising optical sensors, could naïvely be overcome by simply employing more sensors. It is nonetheless unclear how utilising large numbers of optical sensors affects the types of RSOs that can be regularly observed. We therefore aim to determine if an SNSS with a heavy dependence on optical sensors can be devised to achieve similar levels of availability to that of hybrid optical/radar implementations. A hybrid network is chosen as a benchmark in performance for its performance in Chapter 3 and because such a network was used to compile the catalogue under test [9, 31].

To objectively compare the availability achieved by each SNSS configuration, a metric to indicate how comprehensively an SNSS configuration can observe the catalogue is introduced. An SNSS's 'visibility factor' α_{vis} is computed using the following formula:

$$\alpha_{\text{vis}} = \frac{1}{N_{\text{RSO}}} \sum_{i=1}^{N_{\text{RSO}}} \left(\frac{1}{N_{\text{TP}}} \sum_{j=1}^{N_{\text{TP}}} \mathcal{I}(i,j) \right) \tag{5.1}$$

where N_{TP} is the number of tasking periods completed in the simulation, N_{RSO} is the number of RSOs in the catalogue and $\mathcal{I}(i,j)$ is an indicator function defined by

$$\mathcal{I}(i,j) = \begin{cases} 1 & \text{when the } i\text{th RSO is visible in the } j\text{th tasking period} \\ 0 & \text{otherwise} \end{cases} . \tag{5.2}$$

Equation (5.1) computes the average of the number of tasking periods in which each RSO is visible, then averages the result over the number of RSOs to find the expected percentage of RSOs that were observable by the SNSS during each tasking period. The visibility factor α_{vis} of an SNSS configuration therefore varies from $0 \rightarrow 1$, where a value of 1 would indicate that the entire catalogue is visible to the SNSS during every scheduling period. For improved readability, α_{vis} will henceforth be written as a percentage. As (5.1) is dependant on the length of the simulation to achieve an accurate value for α_{vis}, all further simulations are extended from 8 to 24 days to improve the reliability of the comparison.

We continue our availability analysis with a forth simulation that implements a hybrid sensor network. The original optical sensor network is supplemented with 2 additional radar sensors, achieving an expected 4320 track-updates per day. The resulting catalogue invisibility analysis is presented in Fig. 5.9.

Fig. 5.9 indicates that consistent visibility is achieved throughout the catalogue except for two peaks found in lower-LEO and mid-MEO. Further analysis revealed that the lower peak contained a single RSO that could not be observed by the surveillance network. The secondary peak belongs to the Molniya-like objects. These features earn Simulation 4 an

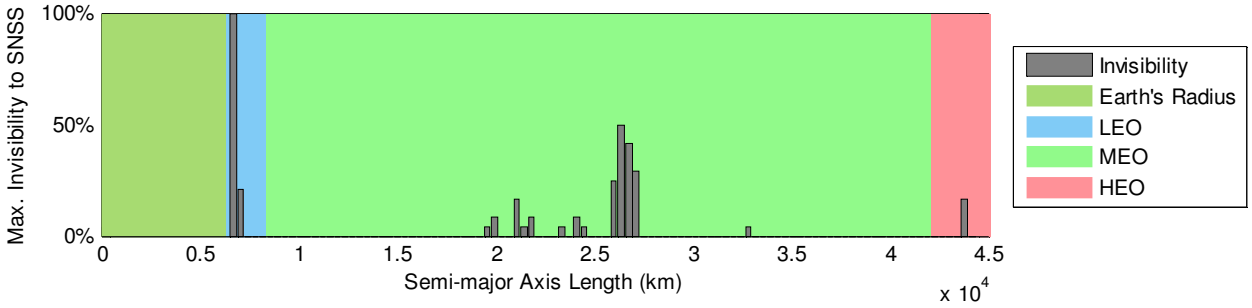

Figure 5.9: Simulation 4: The interval maximums of the average percentage of scheduling periods each RSO was invisible to an SNSS composed of 8 optical and 2 radar sensors throughout a 24 day simulation.

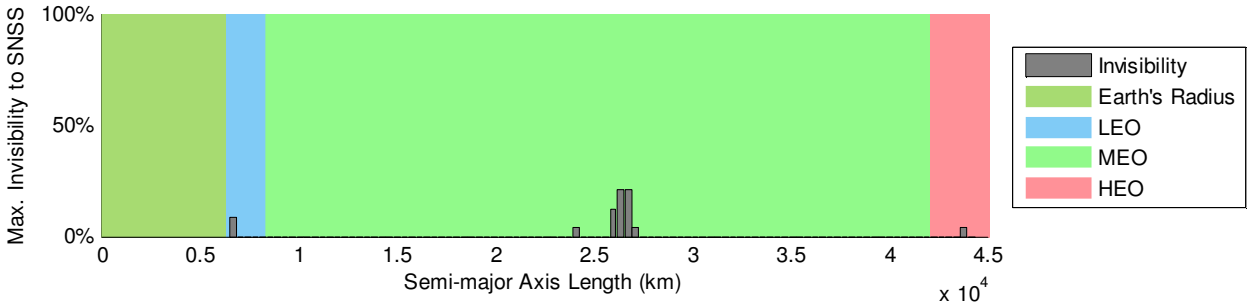

Figure 5.10: Simulation 5: The interval maximums of the average percentage of scheduling periods each RSO was invisible to an SNSS composed of 8 optical and 4 radar sensors throughout a 24 day simulation.

α_{vis} value of 99.92%.

In comparison, Fig. 5.10 displays the results of a second hybrid simulation, Simulation 5, involving the introduction of two additional radars resulting in a total of 12 sensors. The results show that the peaks observed in Fig. 5.9 are significantly reduced by the additional radars. Simulation 5's 5760 expected track-updates per day are now commensurate with the 8-radar, Simulation 2. Fig. 5.10 indicates that the entire catalogue has been visible to the SNSS for at least 75% of the simulation and the additional two radars earn Simulation 5 an improved α_{vis} of 99.95%.

With this new benchmark in performance, we can now attempt to learn what it would take to match this level of availability using relatively large numbers of optical sensors. The results of two additional homogeneous electro-optical simulations—Simulations 6 and 7—are presented to demonstrate the effects of dramatically changing the number of optical sensors. Simulation 6, a network comprising 15 optical sensors, achieves an expected track-update rate of 5 400 track-updates per day and Simulation 7, a network comprising 20 optical sensors, achieves 7 200 track-updates per day. These values have been chosen to provide sufficient

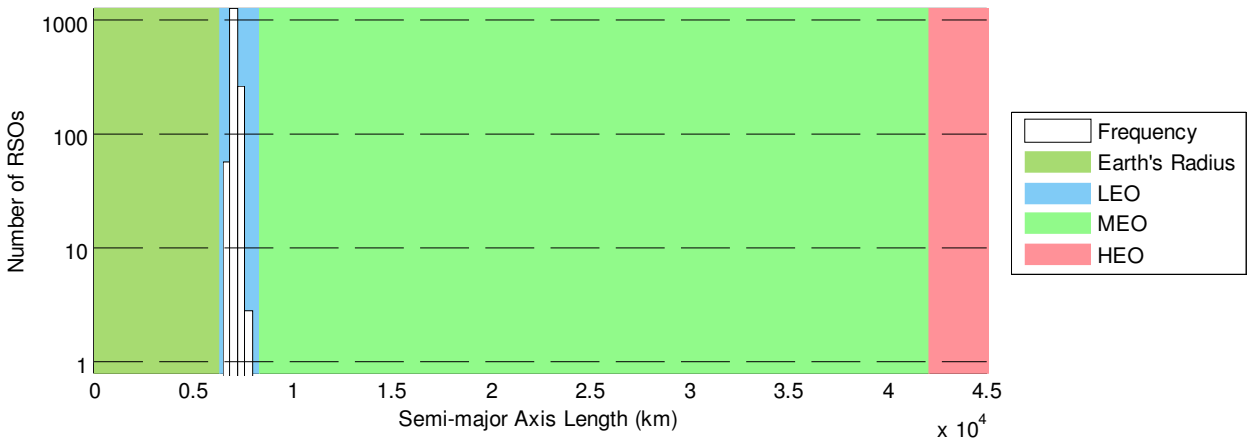

Figure 5.11: Simulation 6: The number and distribution of RSOs that were invisible to an SNSS composed of 15 optical sensors throughout a 24 day simulation.

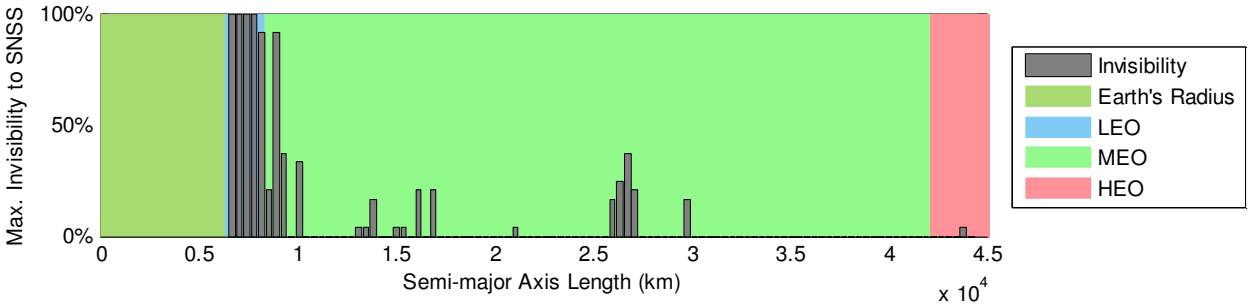

Figure 5.12: Simulation 6: The interval maximums of the average percentage of scheduling periods each RSO was invisible to an SNSS composed of 15 optical sensors throughout a 24 day simulation.

contrast for viewing the effects to SNSS availability. The respective invisibility analyses are presented in Figs. 5.11–5.14.

Figs. 5.11–5.14 display a group of non-observable RSOs in LEO, in addition to other sporadic lapses in visibility throughout MEO and HEO. Upon comparison of these figures, two important features are identified. By increasing the number of optical sensors from 15 to 20, the number of RSOs invisible to the SNSS almost halves and the consistency of visibility of all other RSOs improves. This observation is supported by calculating each scenario's visibility factor. The 15-optical sensor simulation has an α_{vis} of 76.83% and the 20-optical sensor simulation has an α_{vis} of 81.64%. The results demonstrate that whilst the expected number of track-updates per day and catalogue visibility reduce with increasing numbers of optical sensors, the percentage increase in sensors is much larger than the increase in α_{vis}. Therefore, if we chose to ensure the last of the low-altitude LEO objects are observable every scheduling period, we would likely require very large numbers of optical sensors. Since Simulation 4 and 5, we have almost doubled the number of sensors but α_{vis} is falling well

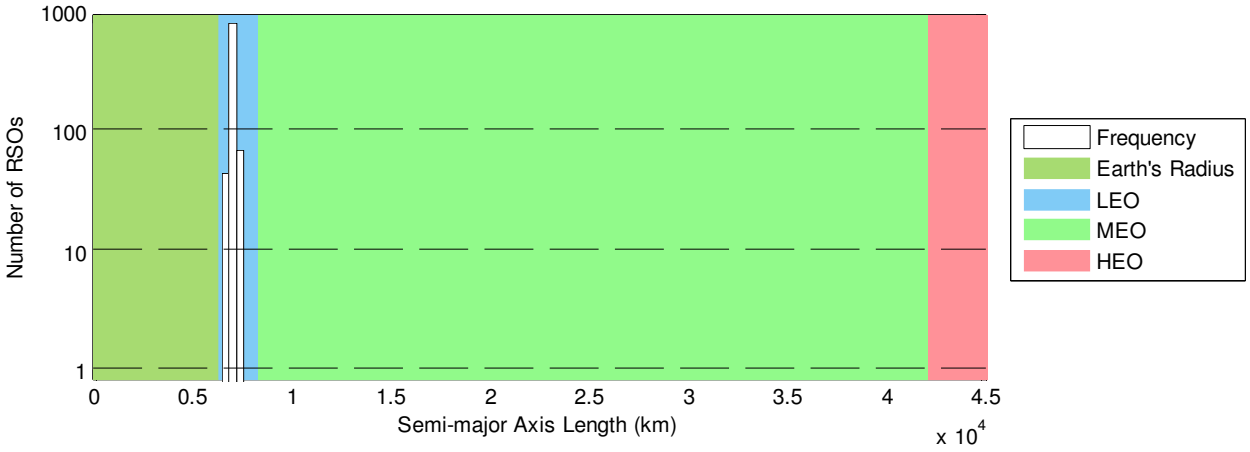

Figure 5.13: Simulation 7: The number and distribution of RSOs that were invisible to an SNSS composed of 20 optical sensors throughout a 24 day simulation.

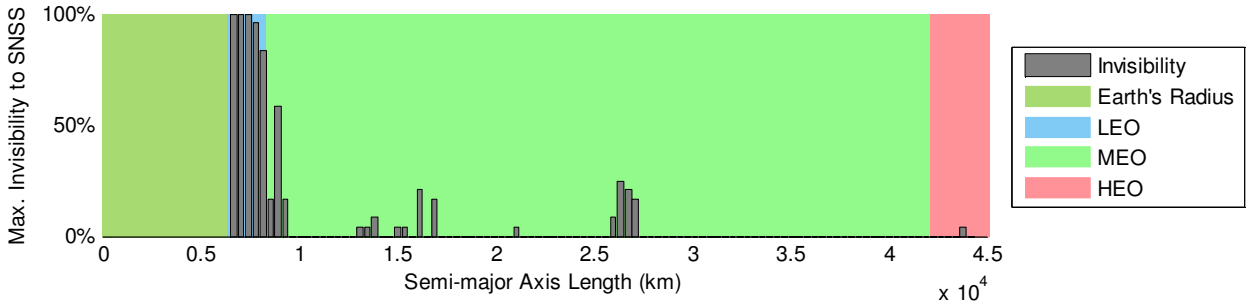

Figure 5.14: Simulation 7: The interval maximums of the average percentage of scheduling periods each RSO was invisible to an SNSS composed of 20 optical sensors throughout a 24 day simulation.

short of the hybrid network's greater than 99.9% coverage.

A compromise is therefore proposed. As range is a limiting factor for radar [32], in terms of cost and capability, perhaps it is more prudent to leave observation of the lowest regions of LEO to radar to guarantee adequate coverage, whilst utilising optical measurements for all other regions. To test this hypothesis, a final SNSS configuration was implemented in Simulation 8, utilising a single radar and 17 optical sensors. Its invisibility analysis is presented in Fig. 5.15.

Fig. 5.15 indicates that a very effective SNSS results from such a configuration. All but a single RSO, in LEO, are observable for the majority of the simulated period. Further analysis indicates that the single non-observable RSO was about to undergo re-entry shortly after its TLE was produced. This element is therefore considered to be an outlier rather than an expectation. To efficiently compare the results of Simulation 8 to all simulations presented so far, the resulting visibility factors, the number of RSOs invisible to each SNSS configuration and estimated number of track updates per day, are summarised in Table 5.3.

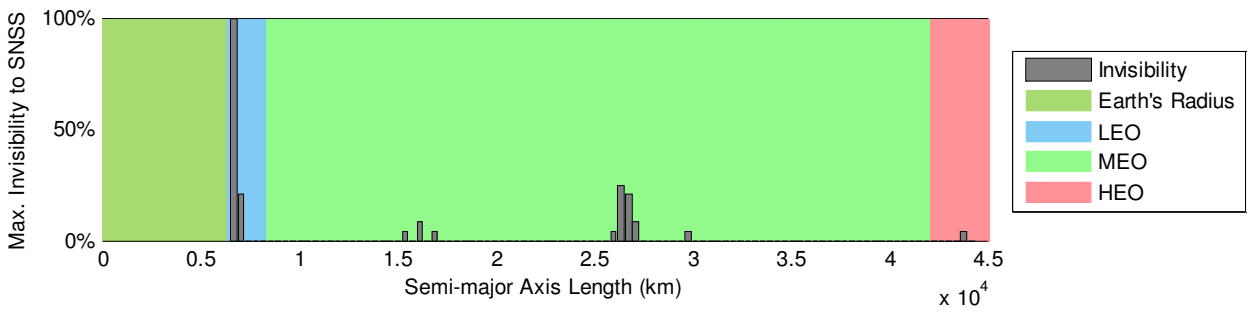

Figure 5.15: Simulation 8: The interval maximums of the average percentage of scheduling periods each RSO was invisible to an SNSS composed of 17 optical sensors and 1 radar throughout a 24 day simulation.

Table 5.3 shows that the final simulation involving 17 optical sensors and only a single radar is able to achieve a level of availability similar to that of the former hybrid simulations.

Table 5.3: Availability characteristics of simulated SNSS configurations

Sim.	Sim. Length (days)	N^o Optical Sensors	N^o Radar Sensors	α_{vis}	N^o Invis. RSOs	Exp. Updates per Day
1	8	8	0	75.37%	2489	2880
2 & 3	8	0	8	90.55%	1390	5760
4	24	8	2	99.92%	1	4320
5	24	8	4	99.95%	0	5760
6	24	15	0	76.83%	1712	5400
7	24	20	0	81.64%	960	7200
8	24	17	1	99.86%	1	6840

This outcome demonstrates that a high ratio of steerable optical sensors to conventional radar sensors is necessary to deliver similar levels of availability to that of more typical hybrid configurations. More importantly, it has been shown that commensurate levels of availability can be achieved when using large numbers of optical sensors, whilst simultaneously ensuring the majority of the catalogue can be routinely updated using the superior optical-based updates detailed in previous chapters.

This investigation does not include sufficient scope to surmise the precise number or ratio of sensors necessary to routinely observe a finite number of RSOs. Introducing influences such as weather outages to the model would inevitably increase the number of optical sensors necessary to achieve commensurate availability. This detail is of limited significance however, when consideration is given to the large disparity in cost concerning the construction and operation of electro-optical sensors when compared to radar [36]. While a network that is

heavily reliant on optical sensors is likely to require many more sensors than a conventional SNSS, the disparity in price suggests that a ratio of optical to radar sensors many times larger than 17:1 can be implemented, without increasing the cost of the system.

5.4 Enhanced Scheduling via Visibility Prediction

Section 5.3's results suggest the large scale deployment of steerable sensors is necessary to attain the levels of availability achieved by existing hybrid networks. But it is a brute force approach that quickly suffers from diminishing returns as extra sensors are added. A balance must therefore be struck between the number of sensors and the regularity at which certain RSOs may be observed, to achieve a certain level of catalogue accuracy. To aid in minimising the number of sensors necessary to achieve a specific level of accuracy, a scheduling method that exploits Section 5.3's efficient visibility prediction is proposed.

Via the use of visibility forecasting, a sensor management system can incorporate a scheduler that prepares the catalogue for periods when specific RSOs are not visible to the network by prioritising the observation of those objects. Incorporating visibility prediction into the scheduling process is however a computationally intensive task that will scale with the size of the RSO catalogue, the number of sensors and the length of the forecast period. Nevertheless, this type of task is well suited to MASSAS's newly described, parallel GPU architecture, as many elements may be computed independently in parallel.

5.4.1 Scaling of Observation Effectiveness by Predicted Visibility

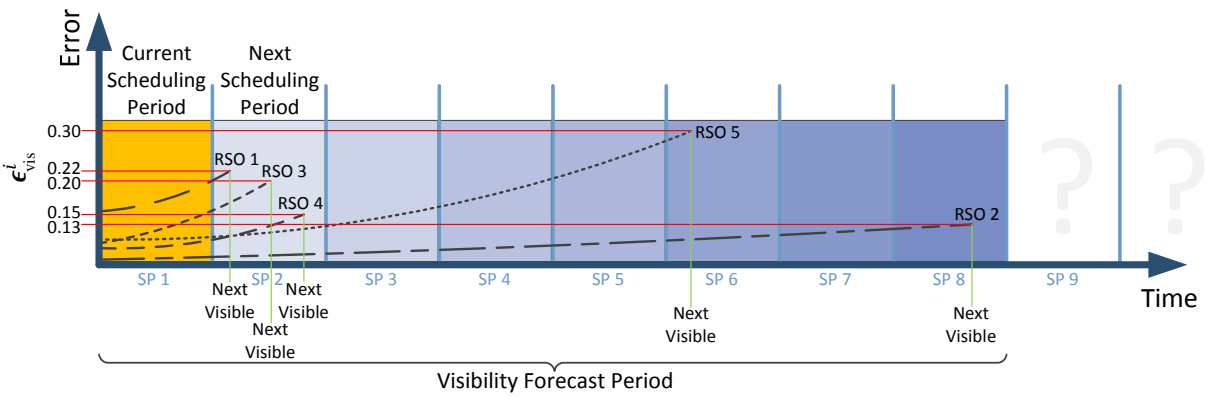

Figure 5.16: Observation effectiveness scaling via visibility prediction.

As depicted in Fig. 5.16, after an RSO is observed, any observation error or process error will cause the variance of its state estimate to grow with time. If an RSO is not visible for an

extended period of time, its variance can grow so large that its estimated orbital parameters become unreliable for the purposes of conjunction analysis and, worse still, the RSO may be difficult to reacquire by an SNSS's sensors. Therefore, to compensate for periods of non-observability, the augmentation of an observation effectiveness metric—such as Chapter 3's β_{red}—is proposed, to prioritise the observation of RSOs about to enter such a period. Ideally the process would ensure the error in an RSO's orbital estimate would be sufficiently small to ensure the reliability of its estimate up until it can next be observed. But the process should also weigh the consequence of making these observations against compromising the accuracy of the larger catalogue.

Algorithm 1 Enhanced observation effectiveness scaling.

Parallelisable Variables	Pseudo Code	Comment
i, j & k	$\mathbf{v}_{\text{bin}}^{i,j,k} = \texttt{CheckVisibility}(i,j,k)$	Determine when each object is next visible by the sensor network during the visibility forecast period: $\mathbf{v}_{\text{bin}}$ is a three-dimensional binary matrix indicating visibility. $\mathbf{v}_{\text{fs}}$ is a two-dimensional matrix recording the time t_k of the first instance of RSO-Sensor visibility. $\mathbf{v}_{\text{next}}$ is a vector storing each RSO's earliest time of visibility for the entire network.

i & j	$k \equiv +1$ day
	LOOP
	IF ($\mathbf{v}_{\text{bin}}^{i,j,k}$ = TRUE) OR ($k = t_{\text{MAX}}$)
	EXIT LOOP
	ELSE
	$k = k + 1$
	END IF
	END LOOP
	$\mathbf{v}_{\text{fs}}^{i,j} = t_k$	

i	$\mathbf{v}_{\text{next}}^{i} = \text{MIN}_j \left(\mathbf{v}_{\text{fs}}^{i,j} \right)$	MIN_j returns the earliest instance of visibility amongst all of the sensors.
i	$\mathbf{P}_{k-\Delta k\|k-\Delta k}^{i} \rightarrow \mathbf{P}_{k\|k-\Delta k}^{i}$; where $\Delta k = \mathbf{v}_{\text{next}}^{i}$ $\boldsymbol{\epsilon}_{\text{vis}}^{i} = \text{trace}\left(\left[\mathbf{P}_{k\|k=0}^{i} \right]_{p,3\times 3} \right)$	Propagate $\mathbf{P}^i$ from the start of the scheduling period to $\mathbf{v}_{\text{next}}^{i}$ using standard Kalman filtering propagation and record.
None.	$\tilde{\boldsymbol{\epsilon}}_{\text{vis}}^{i} = \frac{\boldsymbol{\epsilon}_{\text{vis}}^{i}}{\sum_{i=1}^{N_{RSO}} \boldsymbol{\epsilon}_{\text{vis}}^{i}}$	Normalise $\boldsymbol{\epsilon}_{\text{vis}}$, where N_{RSO} is the total number of RSOs.
i	$\tilde{\boldsymbol{\beta}}_{\text{red}}^{i} = \boldsymbol{\beta}_{\text{red}}^{i} \cdot \tilde{\boldsymbol{\epsilon}}_{\text{vis}}^{i}$	Apply scaling to the ith layer of $\boldsymbol{\beta}_{\text{red}}$.

A method for achieving such a process is proposed in Algorithm 1. At the beginning of a scheduling period, the time until an object is next visible to the sensor network but after the

current scheduling period, is determined. This is achieved by exploiting the computational power of the GPU to efficiently compute visibility for the ith RSO by the jth sensor at time k, well into the future $k = t_{\text{MAX}}$. The orbit error covariance for each object is then propagated forward to the first instance of visibility via a Kalman time update and recorded. Similar to the β_{red} metric, the trace of the position component is used to obtain a new metric named ϵ_{vis} for comparison. Finally, all ϵ_{vis} values are normalised and used to scale the observation effectiveness of each respective RSO for the current scheduling period. This in turn causes the observation effectiveness of an RSO about to enter a period of non-observability, to be scaled according to its estimated variance when it is next visible to the network. This process retains the information concerning when an RSO should be observed to obtain the maximum reduction in state estimate error, but prioritises the observation of objects according to how long it will be before they can be observed again. Whilst the β_{red} observation effectiveness metric has been displayed in Algorithm 1, this process could be similarly applied to alternative compatible methods for evaluating observation effectiveness.

The implementation of the proposed method provides a number of opportunities for parallelisation. Algorithm 1 summarises the dimensions of each element that may be implemented in parallel on the GPU.

5.4.2 Numerical Simulation

To evaluate the proposed scheduling method, Simulation 8's 17-optical 1-radar network was utilised for testing. A visibility forecast period ending at $t_{\text{MAX}} = k_0 + 8$ days, where k_0 is the end of the current scheduling period, was chosen. As Fig. 5.15 indicates that all but a single object experienced periods of non-observability for less than 8 days, this forecast period is likely to encapsulate most lapses in RSO visibility. If an object is not visible in this time, it is assumed to be visible at the end of this period.

Two 24-day scenarios were simulated for comparison. Scenario A used the β_{red} observation effectiveness method to schedule observations for its sensors, as defined previously in Section 3.2.2. Scenario B utilised the proposed GPU-augmented scheduling method.

Fig. 5.17 displays the catalogue error, once again denoted by WCPE and MCPE, recorded at the end of each day for each simulation. The most obvious difference between Scenario A and B's results can be seen upon comparison of their respective values for WCPE. After approximately six days, the Scenario B WCPE begins to diverge from Scenario A's and continues diverging until the end of the simulation when Scenario B's WCPE is an order of magnitude more accurate than Scenario A's.

A small difference is also observable between each scenario's MCPE. Scenario A's MCPE

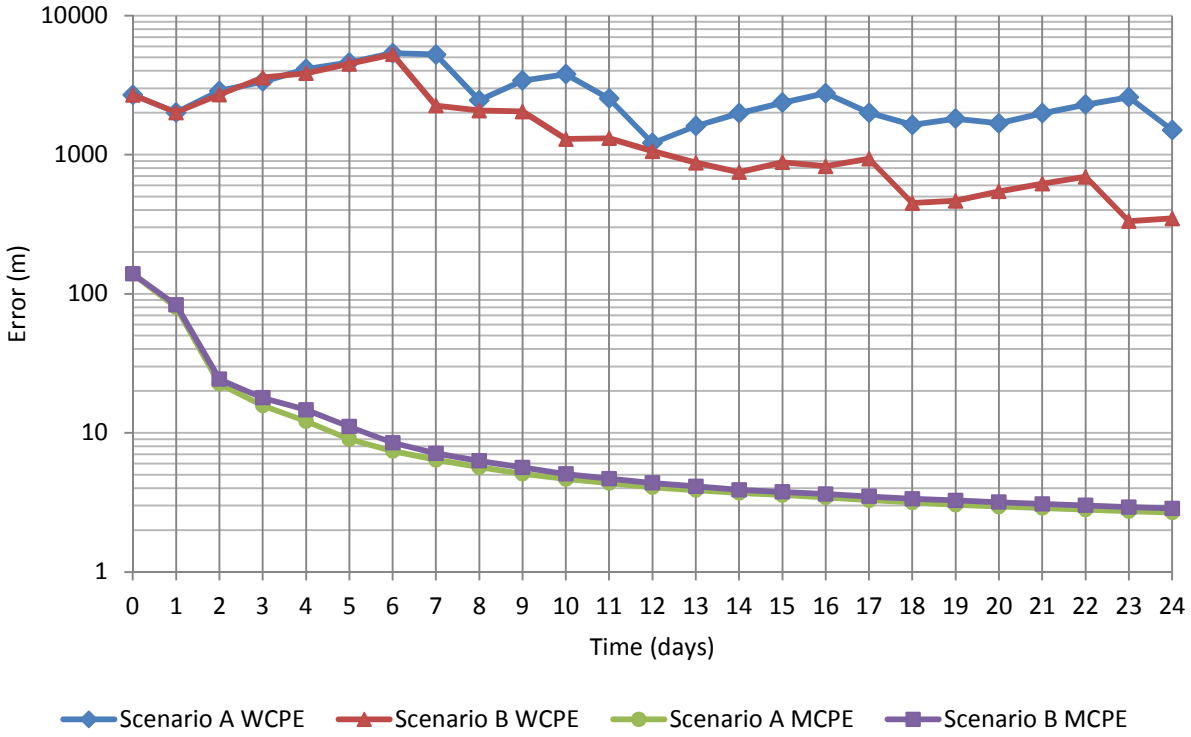

Figure 5.17: Scenario comparison of a 24-Day simulation with and without enhanced scheduling via visibility prediction.

appears slightly more accurate than Scenario B's however, the discrepancy is minor and the shape of the curve, describing the catalogue's improving accuracy, is shared.

5.5 Discussion and Conclusion

The results of the availability analysis suggest that replacing a network of radars with a network of optical sensors, that achieves the same level of availability, will entail the coordination of a much larger number of sensors than before. However, the relative price and capability of these sensors means that this outcome is not necessarily discouraging. It was also found that the number of optical sensors necessary to attain a relatively high level of availability in lower-LEO could be significantly reduced by retaining even a single radar. A highly effective network would therefore consist of many widely distributed optical sensors targeting all altitudes, but for a significant reduction in the number of optical sensors, a very small number of radars should also remain active. These radars can use relatively low power settings to monitor only those RSOs in the lower regions of LEO.

The results of the enhanced-scheduling via visibility forecasting display some growth in MCPE due to an increased focus on any elusive RSOs by the sensor manager. Nonetheless

this reduction in accuracy is very minor when compared to the significant improvement in WCPE as state estimate errors are no longer left to accumulate during periods of non-observability. It is therefore suggested that the proposed scheduling could be used to further reduce the number of sensors necessary to achieve a viable surveillance network composed of a large number of optical sensors. Not only does the catalogue accuracy improve as a result of this method but it also decreases the likelihood of loosing RSOs due to periods of non-observability. It is therefore recommended that use of small, low cost optical sensors can be used to achieve accurate catalogues and enable the network to scale and adapt as the catalogue continues to grow. This recommendation is made not only due to their demonstrated performance and the reduced cost of building new sensors, but also due to their mobility [67].

Future enhancements to proposed scheduling technique could include cueing the deployment of mobile sensors. In the event very large periods of non-observability are anticipated, the system could prompt the deployment of mobile sensors to reacquire and refine the track of specific RSOs in urgent need of observation. Although network visibility prediction has been presented as a benefit to passive-sensing networks, the method could also be applied when using steerable sensors that utilise active-sensing. The capabilities of the proposed method may be further exploited by including other predictive information. The inclusion of information regarding routine maintenance, inclement weather predictions or any other foreseeable outage event could be used to prepare the catalogue for periods of reduced visibility.

When utilising the enhanced observation effectiveness scheduling method, the computationally intensive visibility analysis is performed prior to the highly sequential scheduling process. Because the visibility analysis is performed in parallel, the process is very fast and will very weakly affect the run time of the simulation regardless of the number of catalogued RSOs. Comparison of the test scenarios has demonstrated how valuable the increased computational power of the sensor manager can be, while only lightly influencing the run time through parallelism.

Due to the success of this preliminary application of GPGPU to SSA research, it is hoped that this work encourages further research regarding methods for leveraging the benefits of parallelised algorithms. In spite of the growing number of RSOs orbiting the Earth, emerging technologies such as GPGPU could enable the production of smarter and faster sensor management systems to deliver a commensurate growth in tracking capability.

6

Dynamic Steering

As previously discussed in Chapters 1 & 2, a primary source of information for achieving SSA is the direct observation of RSOs using global networks of space surveillance sensors. Catalogued RSOs are regularly observed in order to minimise their state estimate error [19, 92]. Common methods for updating state estimates with new measurements employ recursive Bayesian filtering techniques requiring Gaussian p.d.f. approximations. Such techniques include the EKF and—a very close relative of the EKF—the recursive least squares (RLS) filter [9, 26, 93]. The US SSN (Space Surveillance Network), for example, uses an automated batch least squares technique [59] to update the state estimates for tens of thousands of objects, utilising any observations made over the preceding 8 hour period [22]. The reliance on Gaussian-based p.d.f. techniques has resulted in a space-tracking paradigm of monitoring the state error covariance of all catalogued objects for tracking, conjunction analysis and scheduling observations. For instance, current scheduling procedures implemented by the SSN [50] utilise the relative size of an RSO's error covariance to weight its priority for observation in the next scheduling period. After observations are made, all subsequent SSA analyses [9] are performed under the assumption a Gaussian distribution is an adequate approximation for the state error p.d.f.

While covariance-based analysis may be adequate for some aspects of SSA, ensuring the state error distribution is reliably approximated by a Gaussian error ellipsoid requires

frequent observation of the target [22,26,31]. In addition, such techniques fail to account for instances when an object is scheduled for observation, but it is not observed in the FOV. In this situation, the 'negative information' [94,95] contained in the *failed* observation, about where the target is not, is discarded as the error covariance cannot easily capture this detail. For these reasons, factors such as sensor availability, operational limitations, limited sensor FOV, data misassociation, manoeuvring targets, unmodelled dynamics and non-Gaussian distributions limit the effectiveness of the covariance-based approach. Therefore in spite of an agency's best efforts, from time to time objects can become lost [9, 19, 27, 28]. As detailed in Chapter 1, the number of lost RSOs has never been larger and reacquisition is a manually intensive process reliant on specialist sensor operators [22]. The continuing reliance on Gaussian p.d.f.s and rather standard Kalman filtering is not a reflection of any consensus that such an approach is optimal. Rather, the shortcomings have been noted before [9,22]. Due to increasing pressure to overcome current limitations and provide new capabilities for the reasons outlined in Chapter 1, techniques capable of modelling the state error to higher levels of fidelity are receiving increased interest [22].

The culture of covariance-based analysis that surrounds the discipline of, and research pertaining to, maintaining SSA [22], has been a strong influence on the research described throughout this thesis. This is evident by its regular use in the previous chapters. Until this chapter, factors such as inclement weather and a sensor's FOV have been ignored, idly assuming that when an observation is scheduled, it will occur with certainty. Whilst this approach is not uncommon for many tracking techniques [43], the validity and usability of every measurement attempt, is of great significance to the effective use of a steerable sensor. Inappropriate steering of the sensor will have a cumulative negative influence on catalogue accuracy if attempted measurements are routinely ineffective. If a high degree of observation reliability cannot be guaranteed, it would be advantageous to have the system appraise its own performance, so that the system can detect problems and respond appropriately in a timely manner.

During this chapter, a modernisation of the space-tracking paradigm is proposed. By incorporating high-fidelity real-time state error p.d.f. estimation at the sensor using a dedicated parallel architecture, a more robust and capable system is achieved than a traditional covariance-based system can accomplish. Within the Bayesian filtering framework, we seek sufficient fidelity to permit augmentation of the error estimate with information regarding attempted observations that fail to observe the target within the sensor's FOV. We incorporate a limited search capability enabling tracking and, if necessary, searching to be autonomously performed by the sensors to better utilise sensing resources and increase system capacity. Whilst a number of Bayesian filters may exhibit sufficient fidelity to implement the proposed

strategy, we utilise a PF, also known as a sequential Monte Carlo (SMC) filter [54, 64], for its ability to represent a p.d.f. to an arbitrary level of fidelity and its ability to be parallelised for efficient computation on the GPU. We exploit the PF's particle-based estimate to judiciously decide in real time where to point the sensor next to improve the probability of gaining desirable information with each observation.

Use of a PF for exploiting the information obtained when a target is not observed in the FOV has previously been implemented by the US Navy and Coast Guard as a means of dynamically revising maritime search strategies for finding lost fishermen and submarines [96]. Their strategy involves devising a particle representation of the prior p.d.f. for one or more targets within a finite search region which is subsequently updated according to oceanographic and observation data. More recently, research has been performed regarding the application of this principle to guiding robotic aircraft to perform a related marine rescue role [97,98]. These related works are posed as methods for coordinated search in a finite two-dimensional search region, in some cases restricted to a grid, and are not directly applicable to Space Surveillance.

Autonomous control of surveillance sensors has recently been employed to improve video-based security systems [99–101]. This research describes a system that coordinates the steering and zooming of cameras to maximise coverage of a surveillance zone while minimising the FOV encompassing a target to improve feature detection. Whilst related to the proposed strategy, this work focuses on selecting an appropriate camera and steering it or an adjacent camera to recapture a target when it momentarily leaves the FOV. This process, in addition to the use of Gaussian estimators, means that a relatively small time scale is required between observations. Our goal, by comparison, is to maximise the time required between observations by maintaining a high-fidelity representation of the target's state p.d.f. over long periods of time.

We develop a modified Bayesian filtering method that incorporates real-time observation evaluation to enable a space-surveillance system to seamlessly transition between tracking and searching as required. By dynamically refining the system's state p.d.f. as each observation arrives, in spite of the possibility of failing to observe the target in the FOV, the observing sensor is directed toward regions of high probability or value. This process permits tracking and automatic reacquisition of the target even if the probability of detection is low. The dynamically steered surveillance sensor also has the capacity to reacquire targets whose dynamics are not known in great detail or are not easily defined. This includes tracking targets during manoeuvres and observing targets such as high area to mass ratio (HAMR) objects, known for their challenging orbital dynamics [27].

This chapter further details the practical implementation of the proposed dynamically

steered system. The effectiveness of various methods of implementation have been assessed via numerical simulation and the most promising methods tested in an experimental field trial. It is shown that replacing or supplementing covariance-based filtering methods with the proposed dynamic steering technique, enables space-surveillance sensors to judiciously and autonomously reacquire targets that an existing SNSS would deem lost. Targeting data as much as 6 months old is utilised to reacquire RSOs. Consequently the system can be operated with relaxed thresholds for probability of detection such that certain RSOs can be reliably reacquired even months after their last observation. The current state of the art, by contrast, requires re-observation every few days [19, 22]. The method proposed during this chapter therefore has the potential to enable greater numbers of targets to be tracked with fewer sensing resources. Furthermore, the system can be implemented using existing technology. This was demonstrated during the field trial when a system capable of dynamic steering was practically implemented using a Raven class electro-optical sensor [35, 36].

The chapter is divided into the following sections. Section 6.1 begins with a detailed description of the proposed method in which the Bayesian filtering framework may be adapted to incorporate dynamic steering. Section 6.2 further adapts the method to make it practical and realisable using current computational techniques. A computer program incorporating these techniques is detailed in Section 6.3 and the results of numerical simulations are provided in Section 6.4. Section 6.5 describes a recent field trial of the software and presents the results. Some concluding remarks and suggestions for future research are discussed in Section 6.6.

6.1 Augmentation of Bayesian Filtering for Dynamic Steering

During Chapter 2, the Bayesian filtering process was detailed and its suitability for tracking RSOs was discussed. It was noted in Section 2.3 that when an observation occurs, the Bayesian filtering process involves two primary elements to update the state estimate; a time update and a measurement update. This process, in a tracking context, ascribes to the following paradigm. When the target is observed, the state error p.d.f. is to be propagated to the observation epoch, it is subsequently refined according to the observation and a new 'best estimate' of the target's state is produced. Whilst this paradigm is appropriate for implementations such as radar-based air traffic control, in which a specified surveillance area is routinely and robustly interrogated for targets, this paradigm is not well suited for space surveillance via steerable sensors. Crucially, the time required by a steerable sensor to

observe its entire surveillance volume may be significantly larger than the amount of time necessary for a potential target to traverse the volume. Judiciously steering the sensor to the most appropriate region within the surveillance volume is therefore of great significance and deserves attention inside the filtering process itself. The decision to point a sensor in any particular direction comes with an associated risk that the target will not be captured within the sensor's FOV. Furthermore, limitations associated with passive detection may limit the sensor's ability to 'see' the target whether the target is in the FOV or not. Consequently, the direction the sensor is steered and the reliability of the observation weighs heavily on the effectiveness of its observations.

To produce a filtering process that recognises the significance of effective steering and measurement reliability when using a steerable sensor, a modification to the Bayesian filtering process is proposed. This modification incorporates a sensor-steering step and an observation evaluation step between the standard time-update and measurement-update steps. Furthermore, the act of making the observation itself is positioned between these new steps so that, in general, the system will perform a sequence comprising of a prediction, steering of the sensor, observation, evaluation of its effectiveness and the application of an appropriate measurement update, in a dynamic loop. This concept, as depicted in Fig. 6.1, forms the iterative process that occurs at each observation to enable the desired real-time, autonomous system behaviour.

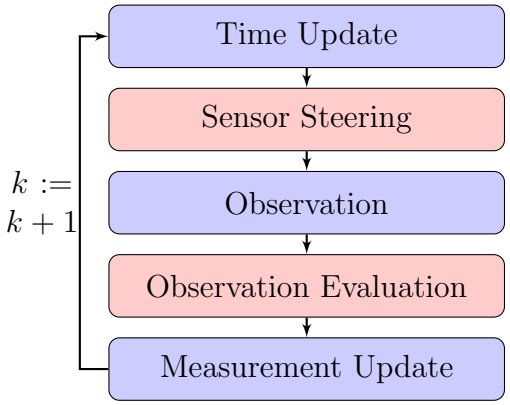

Figure 6.1: The proposed modified Bayesian sequential update process.

6.1.1 Sensor Steering

The first step in the dynamic loop is the time-update, which involves the implementation of the system equation (2.29) to produce the prior distribution $p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})$. If the prior is unimodal and Gaussian, aligning the sensor's bore-sight with the distribution's mean will

result in the maximum probability of detecting the target. As discussed in this chapter's introduction, an EKF or RLS filter models the prior as a Gaussian density. Implicitly, the standard procedure for each observation is to target the mean for directing steerable space-surveillance sensors. However, as the intention is now to anticipate and capture information regarding the absence of the target from the FOV and allow the nonlinear dynamics to distort the p.d.f. with time, the expectation of a unimodal and Gaussian-like prior distribution is no longer valid. It is therefore the role of the sensor steering step to exploit the information contained in the prior, to judiciously select the sensor pointing angles that are most likely to result in the greatest benefit.

Depending upon the application, steering objectives such as quickly locating the target or minimising p.d.f. variance, may be desirable. The latter may be useful if a high-fidelity uni-modal p.d.f. is desired, for example, when interfacing with legacy systems. Physical limitations of the sensor may restrict where the sensor can be pointed and the cost of slewing the sensor from one set of pointing angles to the next may be of concern. Therefore, obtaining the pointing angles with, for example, maximum probability of detection may not be the only consideration. We therefore generalise the steering process by defining the sensor aiming vector $\mathbf{\Psi}$ which will contain at least the sensor pointing angles, such as azimuth and elevation, but may also contain information such as the cost of slewing to the proposed angles and metrics to enforce pointing limitations such as unobservable low elevations or gimbal lock. Obtaining the most advantageous aiming parameters is therefore achieved by locating the global maxima or minima of the objective function $\mathbf{J}(\mathbf{\Psi})$.

The form of the objective function governs the search strategy that the process will execute during each iteration. Whilst it is possible to use the prior to plan observations that collectively describe a search path, this investigation will remain focused on dynamically updating the prior with each observation to guide the next. The following are example objective functions that could be applied to locate the target as quickly as possible or to bound the error as quickly as the sensor's FOV permits. As the examples are general, their formulation assumes that $\mathbf{\Psi}$ contains only the pointing angles and the precise FOV is known but undisclosed.

- To find the target as quickly as possible, a greedy pointing strategy is proposed that points the sensor in the direction that achieves the maximum probability of locating the target. Pointing angles are selected by choosing the aiming parameters that result in the global maximum of the function

$$\mathbf{J}(\mathbf{\Psi}) = \int_{\chi} \mathcal{I}(\mathrm{FOV}(\mathbf{\Psi}))p(\mathbf{x}_k \mid \mathbf{z}_{k-1})\, d\mathbf{x}, \tag{6.1}$$

where $\mathcal{I}$ is an indicator function defined by

$$\mathcal{I}(\mathrm{FOV}(\boldsymbol{\Psi})) = \begin{cases} 1 & \mathbf{x} \in \mathrm{FOV}(\boldsymbol{\Psi}) \\ 0 & \text{otherwise} \end{cases}, \tag{6.2}$$

χ is the state space and $\mathrm{FOV}(\boldsymbol{\Psi})$ is the volume, within state space, that is directly or indirectly observable by the sensor's FOV when steered according to $\boldsymbol{\Psi}$.

- To bound the error, we propose a 'pessimistic' strategy that targets the tails of the distribution in anticipation of an unsuccessful observation. This method prioritises reduction in uncertainty over locating the target. We begin by defining the unfavourable outcome $\mathbf{z}_k = \boldsymbol{\varepsilon}$, in which an observation was attempted but the target is not found in the FOV. A first order approximation for bounding the error can be achieved by minimising the expected variance at time k. This is implemented by selecting the parameters that result in the global minima of the objective function

$$\mathbf{J}(\boldsymbol{\Psi}) = \|\mathrm{cov}\left(\mathbf{x}_k \mid \mathbf{z}_k = \boldsymbol{\varepsilon}\right)\|_{\mathrm{F}} \tag{6.3}$$

where $\|\cdot\|_{\mathrm{F}}$ is the Frobenius norm,

$$\mathrm{cov}\left(\mathbf{x}_k \mid \mathbf{z}_k = \boldsymbol{\varepsilon}\right) = \int_{\chi} (\mathbf{x}_k - \boldsymbol{\mu}_k)(\mathbf{x}_k - \boldsymbol{\mu}_k)^{\mathrm{T}} p\left(\mathbf{x}_k \mid \mathbf{z}_k = \boldsymbol{\varepsilon}, \boldsymbol{\Psi}\right)\, d\mathbf{x}_k, \tag{6.4}$$

and the vector $\boldsymbol{\mu}_k$ denotes the mean of $p\left(\mathbf{x}_k \mid \mathbf{z}_k = \boldsymbol{\varepsilon}, \boldsymbol{\Psi}\right)$.

6.1.2 Observation Evaluation

The observation evaluation step introduces a deviation from the standard Bayesian measurement update by accounting for the uncertainty in the outcome of an attempted observation. Once the sensor has been steered and the observation is made, it is the role of the observation evaluation step to determine how the measurement update should be applied depending on the outcome and quality of the data returned. This is achieved by assessing performance metrics about the observation in addition to the measurement data itself. Performance metrics may include signal-to-noise ratio, data association confidence and environmental considerations such as inclement weather. The ultimate aim is to calculate the probability of detecting the target when it is in the FOV, p_{D}, and the probability of falsely detecting the target when the true target is not in the FOV, p_{FA}, for the newly acquired observation. The probabilities can then be used to appropriately weight the state error distribution to record the information contained in the measurement, even if the target was not observed

in the sensor's FOV.

As first detailed in Section 6.1.1, the set of acceptable observations has expanded to include an unsuccessful observation, denoted by $\boldsymbol{\varepsilon}$. To incorporate unsuccessful observations into the Bayesian framework, the state $\boldsymbol{\varepsilon}$ will be included as an additional dimension to the measurement vector. Consequently, $\mathbf{z}_k \in \mathbb{R}^{n_\mathbf{z}+1}$ where $n_\mathbf{z}$ is the number of measurements returned by a sensor and the $(n_\mathbf{z}+1)$th element, $\mathbf{z}_{k,n_\mathbf{z}+1}$, contains a boolean value such that

$$\mathbf{z}_{k,n_\mathbf{z}+1} = \begin{cases} 1 & \text{when target detected in FOV, } [\mathbf{z}_{k,1}, \ldots, \mathbf{z}_{k,n_\mathbf{z}}]^{\mathrm{T}} \in \mathbb{R}^{n_\mathbf{z}} \\ 0 & \text{when target } \textit{not} \text{ detected in FOV, } \mathbf{z}_k = \boldsymbol{\varepsilon} = \mathbf{0}. \end{cases} \tag{6.5}$$

In spite of adding the error state to the measurement vector, the standard Bayesian measurement-update described by (2.28) remains unchanged. Adaptation of the method is instead applied via the constituent distribution $p(\mathbf{z}_k \mid \mathbf{x}_k)$, by incorporating the observation evaluation information in addition to the sensor's measurement error characteristics. The equation

$$p(\mathbf{z}_k \mid \mathbf{x}_k) = \begin{cases} \delta^{n_\mathbf{z}+1}(\mathbf{z}_k)(1-p_\mathrm{D}) + \delta(\mathbf{z}_{k,n_\mathbf{z}+1}-1)p_\mathrm{D}p_\mathrm{Z}(\mathbf{z}_k \mid \mathbf{x}_k, \boldsymbol{\Psi}_k) & \mathbf{x} \in \mathrm{FOV}(\boldsymbol{\Psi}_k) \\ \delta^{n_\mathbf{z}+1}(\mathbf{z}_k)(1-p_\mathrm{FA}) + \delta(\mathbf{z}_{k,n_\mathbf{z}+1}-1)p_\mathrm{FA}p_\mathrm{U}(\mathbf{z}_k \mid \boldsymbol{\Psi}_k) & \mathbf{x} \notin \mathrm{FOV}(\boldsymbol{\Psi}_k) \end{cases} \tag{6.6}$$

is used to describe it and is thus a mixed discrete/continuous distribution that depends on the FOV, the presence or absence of a measurement and the reliability of the sensor. $\delta^n(\cdot)$ denotes an nth-dimensional Dirac delta function. The distribution $p_\mathrm{Z}(\mathbf{z}_k \mid \mathbf{x}_k, \boldsymbol{\Psi}_k)$ describes the measurement error characteristics of the sensor and may incorporate the probability of data misassociation. $p_\mathrm{U}(\mathbf{z}_k \mid \boldsymbol{\Psi}_k)$ describes the assumed uniform distribution across the FOV in the event of a false alarm.

6.2 A Particle-based Implementation

As discussed in greater detail in Section 2.3, use of a Bayesian filter is generally not tractable except when the system can be described by linear equations and is influenced only by Gaussian noise [53, 54]. Popular solutions to this problem include approximating the p.d.f. with simplistic models, discretising the state space to turn integrals into sums and summing simpler models together to create models with increased fidelity. Filtering of nonlinear non-Gaussian systems therefore requires a balance to be struck between the degree of computation involved in estimating the posterior and the desired level of fidelity. This principle has made filters such as the EKF and RLS filter a popular choice for many tracking applications. When using these filters, computational efficiency is achieved at the expense of relatively

low fidelity through linearisation of system dynamics and modelling each p.d.f. using only its first and second moments.

Nevertheless, in certain applications, high rates of mis-detections, long time spans between updates, nonlinear dynamics and non-Gaussian densities can result in a severe disparity between a covariance-based representation of a prior or posterior p.d.f. and their true form. This is demonstrated in Fig. 6.2 by illustrating the difference between a low fidelity Gaussian and higher fidelity Monte-Carlo-particle p.d.f. approximation when applied to orbital dynamics. It is in applications such as SSA, that increased fidelity may be worth the added computational load. As discussed previously in Chapter 5, many elements of SSA can be implemented as independent processes enabling parallel architectures to better cope with higher levels of computation. For these reasons, we look to a filtering method capable of implementing the Bayesian framework using high-fidelity p.d.f. approximations to enable implementation of dynamic steering for improved SSA.

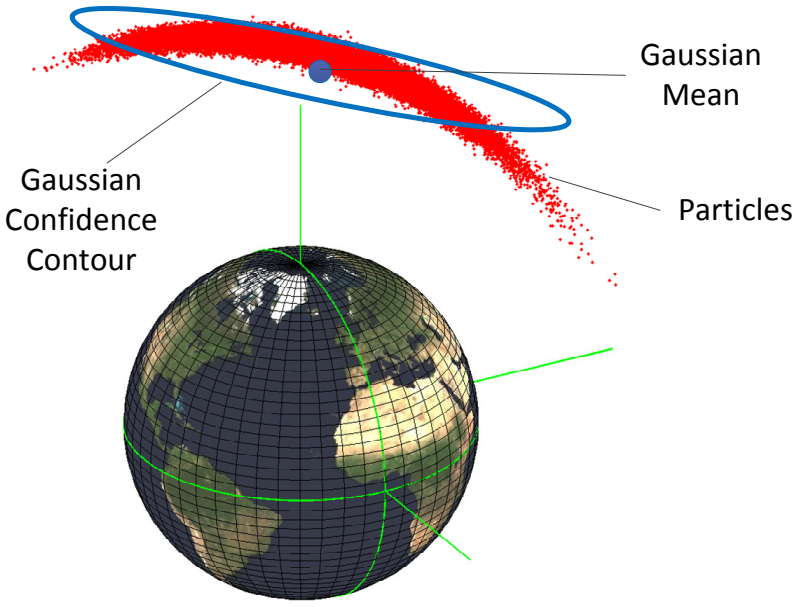

Figure 6.2: An illustration of Gaussian and particle representations of an RSO's state p.d.f.

The particle filtering process detailed in Section 2.3.3 is one such filtering process. By utilising a PF, the system is capable of maintaining a high-fidelity state p.d.f. and it may be parallised in an efficient GPGPU implementation. Whilst a PF is suitable for SSA related applications and is compatible with MASSAS's architecture, its use introduces some application specific complications that must be overcome to be effective. As noted in Section 2.3.3, utilisation of the optimal importance density may be challenging depending on the application. Furthermore, well documented difficulties [54, 63, 64] regarding "particle degeneracy"—when all but a few particles have trivial weight—and "loss of particle diversity"—when many

particles share the same state space—must be addressed. Popular methods for mitigating these concerns have been published as a range of SIS-based PFs [54, 63, 64]. One such method named a regularised particle filter (RPF) [54, 64, 102] was chosen for implementation of dynamic steering for SSA.

6.2.1 Regularised Particle Filtering

The RPF was chosen above alternative methods due to its relatively straightforward implementation, whilst incorporating just enough complexity to overcome some of the filtering challenges simpler methods could not [54, 102]. The RPF combats the two major PF challenges, particle degeneracy and loss of diversity, by incorporating resampling and regularisation respectively. Particle degeneracy is caused by an increase in the variance of the importance weights, as an unavoidable consequence of utilising the SIS framework to filter successive observations [103]. Because there are a finite number of particles, as the variance increases so too does the number of particles representing trivial weights on the periphery of the p.d.f. This in turn has a detrimental effect on the accuracy of the particle-based approximation. Without intervention, eventually a single particle will achieve near-unity weight. It is this feature that is monitored to measure the degree of degeneracy. This is achieved by estimating the number of effective particles using the equation

$$\hat{N}_{\text{eff}} = \frac{1}{\sum_{i=1}^{N} \left(w_k^i\right)^2}. \tag{6.7}$$

If the number of effective particles drops below a threshold, $\hat{N}_{\text{eff}} < N_{\text{Thr}}$ a resampling process is executed. Resampling is employed to better represent regions of high probability within state space and to remove or relocate particles with trivial weight. There are many methodologies suggested for resampling [104], however systematic resampling [105] is chosen for its computational efficiency and empirically demonstrated performance. Systematic resampling redistributes particles to regions of high weight by relocating particles with low weight to collocate them with particles with high weight, such that the weights of all particles become equal. Nonetheless, this process is sensitive to the magnitude of process error applied in (2.29) as loss of particle diversity may result when process error is low. This is because particles that have been intentionally collocated will continue to share the same or a very similar state if not separated over time by process noise [54]. In the worst case, all particles inhabit the same location and an outcome similar to worst-case particle degeneracy is all that has been achieved. For this reason, the RPF augments systematic resampling with a regularisation step that involves adding noise to particles from a continuous distribution, the

kernel density $K(\cdot)$. The intention is to jitter the particles, discouraging them from sharing the same location in state space, without changing the stochastic properties of the particle-based representation of the posterior density [54, 64]. The cumulative result is a process that effectively resamples the particles from a continuous approximation of the posterior distribution. This process is described by the equation

$$\hat{\mathrm{p}}\left(\mathbf{x}_k \mid \mathbf{z}_{1:k}\right) = \sum_{i=1}^{N} w_k^i K_h(\mathbf{x}_k - \mathbf{x}_k^i) \tag{6.8}$$

where

$$K_h(\mathbf{x}) = \frac{1}{h^{n_\mathbf{x}}} K\left(\frac{\mathbf{x}}{h}\right) \tag{6.9}$$

is the kernel density scaled by the kernel bandwidth h and $n_\mathbf{x}$ is the dimension of the state vector. Kernel scaling is necessary to ensure the amount of jitter is proportionate to the larger approximate posterior density. Appropriate kernel densities and bandwidth scaling are discussed in detail in [54] and [64]. Their suggested implementation strategy was followed during implementation.

As discussed in Section 2.3.3's introduction to particle filtering, selecting an appropriate importance density $q(\cdot)$ is required by the SIS framework to update particle weights. The optimal importance density has been shown [103] to be $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}^i, \mathbf{z}_k)$ resulting in a weight update equation of

$$w_k^i \propto w_{k-1}^i p(\mathbf{z}_k \mid \mathbf{x}_{k-1}^i). \tag{6.10}$$

Nonetheless, [54] notes that this is not straightforward unless the proposal particles are part of a finite set or if $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}^i, \mathbf{z}_k)$ is Gaussian. For this reason it is common to adopt the transitional prior instead, setting

$$q(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i, \mathbf{z}_k) = p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i). \tag{6.11}$$

The substitution of (6.11) into (2.68) results in the sub-optimal but effective weight update proportionality

$$w_k^i \propto w_{k-1}^i p(\mathbf{z}_k \mid \mathbf{x}_k^i). \tag{6.12}$$

An unfortunate by-product of this process is that in applications, such as the proposed application, the uncertainty of the transitional prior $p(\mathbf{x}_k \mid \mathbf{x}_{k-1})$ may be significantly larger than the uncertainty of $p(\mathbf{z}_k \mid \mathbf{x}_k)$. In this case, the density of the particles representing $p(\mathbf{x}_k \mid \mathbf{x}_{k-1})$ may be low enough, with respect to $p(\mathbf{z}_k \mid \mathbf{x}_k)$, that the application of (6.12) may result in particle degeneracy. Whilst significantly increasing the number of particles is a

method of combating this problem, solutions involving a multi-stepped, gradual application of (6.12) via resampling have been proposed [102, 106] and are implemented, to achieve a robust RPF implementation without changing the number of particles, N.

6.2.2 Adaptation of the Sensor Steering Step for the PF

Irrespective of implementation method, the sensor steering step assumes an approximation of the prior is available to inform the system where it should next point the sensor. Elements of the PF time-update can be utilised to achieve this goal. Initially, the particles are independently propagated from the previous epoch $k-1$ to the current epoch k using (2.29). When N is sufficiently large, time propagated particles take on a form representative of the prior, albeit with some Monte Carlo and process error. We utilise this approximation, now of the form

$$\hat{p}\left(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}\right) = \sum_{i=1}^{N} w_{k-1}^{i} \delta\left(\mathbf{x}_k - F\left(\mathbf{x}_{k-1}^{i}, \mathbf{v}_{k-1}^{i}\right)\right), \tag{6.13}$$

to predict how the sensor should be pointed to obtain the observation with the most value after scheduling an observation to occur at time k.

In conjunction with (6.13), the search strategies previously discussed in Section 6.1.1 can be implemented using a PF adaptation. Equation (6.1) is adapted to locate the target as quickly as possible by obtaining looking angles with maximum weight. The second strategy, adapted from (6.3), looks to minimise the variance of the state vector's position component to reduce the position-uncertainty in $\hat{p}\left(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}\right)$.

Maximum Weight

The maximum weight objective function evaluates the ith element of the finite set of azimuth and elevation angles $\boldsymbol{\mathcal{P}} = \left\{\boldsymbol{\Psi}^1, \ldots, \boldsymbol{\Psi}^N\right\}$ that align the sensor's boresight with one of the N particles at time k, where $\boldsymbol{\Psi}^i = h(\mathbf{x}^i)$ is the angle pair resulting from the measurement model h—see (2.31)—when applied to the ith particle. The value assigned to each set of steering angles is governed by the FOV of the sensor, when aimed using $\boldsymbol{\Psi}^i$, and the cumulative sum of the particle-weights within the FOV. The resulting objective function is described by

$$J\left(\boldsymbol{\Psi}^i\right) = \sum_{\mathbf{x}_{k|k-1}^{s} \in \mathrm{FOV}(\boldsymbol{\Psi}^i)} w_{k-1}^{s}, \tag{6.14}$$

where s is an element of the set of indices for particles that fall within the sensor's FOV when the sensor is aimed at the ith particle. The objective is therefore to find the pointing

angles with the maximum probability of detection, as estimated by (6.13), when restricted to the finite set of angles $\boldsymbol{\mathcal{P}}$. This is computed by finding the maxima of (6.14).

Minimum Variance

The elements of $\boldsymbol{\mathcal{P}}$ are also evaluated by the minimum variance objective function. The variance of the expected posterior distribution is assessed assuming the target will not be observed and consequently the weights of the particles within the FOV are to be reduced. The resulting objective function is described by

$$J\left(\boldsymbol{\Psi}^i\right) = \left\|\mathbf{X}\mathbf{W}\mathbf{X}^{\mathrm{T}}\right\|_{\mathrm{F}} \tag{6.15}$$

where,

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}^1_{\mathrm{pos},k|k-1} - \mu_{\mathbf{x}_{\mathrm{pos}}} & \dots & \mathbf{x}^N_{\mathrm{pos},k|k-1} - \mu_{\mathbf{x}_{\mathrm{pos}}} \end{bmatrix}, \tag{6.16}$$

$$\mathbf{W} = \mathrm{diag}\begin{bmatrix} w^1_{k|\boldsymbol{\varepsilon},\boldsymbol{\Psi}^i} & \dots & w^N_{k|\boldsymbol{\varepsilon},\boldsymbol{\Psi}^i} \end{bmatrix}, \tag{6.17}$$

$$w^s_{k|\boldsymbol{\varepsilon},\boldsymbol{\Psi}^i} = \begin{cases} w^s_{k-1}(1-p_{\mathrm{D}}) & \mathbf{x}^s_{k|k-1} \in \mathrm{FOV}(\boldsymbol{\Psi}^i) \\ w^s_{k-1}(1-p_{\mathrm{FA}}) & \mathbf{x}^s_{k|k-1} \notin \mathrm{FOV}(\boldsymbol{\Psi}^i) \end{cases}, \tag{6.18}$$

and $\mathbf{x}_{\mathrm{pos}}$ are the position components of the state vector. The best choice of $\boldsymbol{\Psi}$, of the set $\boldsymbol{\mathcal{P}}$, is therefore found by selecting the value resulting in a global minimum according to (6.15).

Alternative strategies

Whilst the previous strategies are theoretically attractive generalised solutions, practical use of the dynamic steering methodology to a specific application can motivate the use of alternative techniques. It is for this reason the following two techniques are proposed below. The motivation behind and relative performance of these techniques will be discussed further in Section 6.4.2.

Raster Search

Unlike the previously proposed objective functions, this process is not general, and is designed to exploit the knowledge that the primary component of an RSO's state uncertainty most commonly lies along the RSO's estimated trajectory [107]. This characteristic means that the search problem, when applied to SSA and depending on the relative FOV of the sensor, is expected to reduce to a linear search problem along the RSO's estimated trajectory. A

methodical scan along the trajectory can be devised by means of the objective function

$$J\left(\boldsymbol{\Psi}^i\right) = \mu^{\mathrm{T}}_{\mathbf{x}_{\mathrm{vel},k}}\left(\mathbf{x}^i_k - \mu_{\mathbf{x}_{\mathrm{pos},k}}\right) \tag{6.19}$$

where $\mu_{\mathbf{x}_{vel}}$ is the arithmetic mean of all particles' velocity components, $\mathbf{x}^i_k = h^{-1}(\boldsymbol{\Psi}^i)$ at time k and h^{-1} is the inverse function of the measurement model. The raster motion is achieved by initially using (6.19) to locate the particle with the greatest positive distance from the mean, with respect to the velocity vector. This particle is then positioned on the edge of the FOV whilst centering the boresight on the average track by use of $\boldsymbol{\Psi} = H(\mathbf{x}^{\mathrm{boresight}}_k)$ where

$$\begin{aligned} \mathbf{x}^{\mathrm{boresight}}_k &= \mathbf{x}^{\max}_k - \frac{\mu_{\mathbf{x}_{\mathrm{vel},}}}{|\mu_{\mathbf{x}_{\mathrm{vel},k}}|}d, \\ d &= R_{\mu_p}\tan\left(\frac{\theta_{\mathrm{FOV}}}{2}\right), \end{aligned} \tag{6.20}$$

$\mathbf{x}^{\max}_k$ is the particle that maximises $J\left(\boldsymbol{\Psi}^i\right)$ of the set $\{\mathbf{x}^i_k \mid w^i > w_{\mathrm{Thr}}\}$, to ensure non-trivial weights are excluded, R_{μ_p} is the average slant range from the site to the position component of the particle's arithmetic mean and θ_{FOV} is the angle the sensor's FOV subtends. The iteration of this process results in a 'nose to tail' raster along the p.d.f., however a 'tail to nose' raster is equally achievable by setting $\mu_{\mathbf{x}_{\mathrm{vel},k}} := -\mu_{\mathbf{x}_{\mathrm{vel},k}}$.

Maximum Particles

The maximum particles objective function is used to maximise the number of particles within the FOV whose weights are greater than a minimum threshold w_{Thr}. The maximum particles objective function is similarly applied to each element of the set $\boldsymbol{\mathcal{P}}$. The objective function is described as

$$J\left(\boldsymbol{\Psi}^i\right) = \#\{\mathbf{x}^s_{k|k-1} \in \mathrm{FOV}(\boldsymbol{\Psi}^i) \mid w^s_{k-1} > w_{\mathrm{Thr}}\} \tag{6.21}$$

where # denotes the cardinality of the set and w_{Thr} is the minimum particle weight necessary to be included in the set. The objective is to find the angle pair that maximises (6.21).

6.2.3 Adaptation of the Observation Evaluation Step for the PF

Only minor adaptation of the process proposed in Section 6.1.2 is required to implement the observation evaluation step using a PF. The combination of (6.6) with (6.12) results in the

weight update equation

$$\tilde{w}_k^i = \begin{cases} w_{k-1}^i \delta^{n_{\mathbf{z}}+1}(\mathbf{z}_k)(1-p_{\mathrm{D}}) + w_{k-1}^i \delta(\mathbf{z}_{k,n_{\mathbf{z}}+1}-1) p_{\mathrm{D}} p_{\mathrm{Z}}(\mathbf{z}_k \mid \mathbf{x}_k, \boldsymbol{\Psi}_k) & \mathbf{x} \in \mathrm{FOV}(\boldsymbol{\Psi}_k) \\ w_{k-1}^i \delta^{n_{\mathbf{z}}+1}(\mathbf{z}_k)(1-p_{\mathrm{FA}}) + w_{k-1}^i \delta(\mathbf{z}_{k,n_{\mathbf{z}}+1}-1) p_{\mathrm{FA}} p_{\mathrm{U}}(\mathbf{z}_k \mid \boldsymbol{\Psi}_k) & \mathbf{x} \notin \mathrm{FOV}(\boldsymbol{\Psi}_k) \end{cases} \tag{6.22}$$

As (6.12) is a proportionality, subsequent weight normalisation is required and is achieved by means of

$$w_k^i = \frac{\tilde{w}_k^i}{\sum_{j=1}^{N} \tilde{w}_k^j}. \tag{6.23}$$

In the intended scenario, where p_{D} is greater than p_{FA}, the implementation of this process will cause the weights of particles surrounding a successfully observed object to increase. If however the target cannot be located within the FOV, the weight of all particles within the FOV will decrease to deter further observation.

6.2.4 The Iterative Result

By incorporating the sensor steering and observation evaluation components within the PF, the system is capable of assessing how best to aim the sensor and how to apply the information obtained from an observation in a dynamic loop. As an example, Fig. 6.3 illustrates the use of the maximum weight objective function to find an elusive RSO. The iteration of the modified PF process results in the convergence of the bore-sight with the target. Equally, if the target is observed, the system will continue steering the sensor toward the target to maintain a track. Furthermore, if the target is lost again, the system will seamlessly readopt its search behaviour and autonomously reacquire the target during future observations.

6.3 SPARSE

The methodology described in Section 6.2 has been practically implemented in a computer program named Space Particle Search Evaluation (SPARSE). SPARSE has been produced as a high-level module for MASSAS. Consequently, SPARSE utilises many of MASSAS's modules in addition to contributing a number of task specific modules of its own. Making SPARSE field-capable has also necessitated enhancing the fidelity of some of MASSAS's modules including coordinate transformations, orbit propagation and visibility prediction. These enhancements enable the system to utilise industry/military standard coordinate frames, include many of the primary sources of orbit perturbation [9, 68] and prevent saturation of optical sensors when directed near the moon.

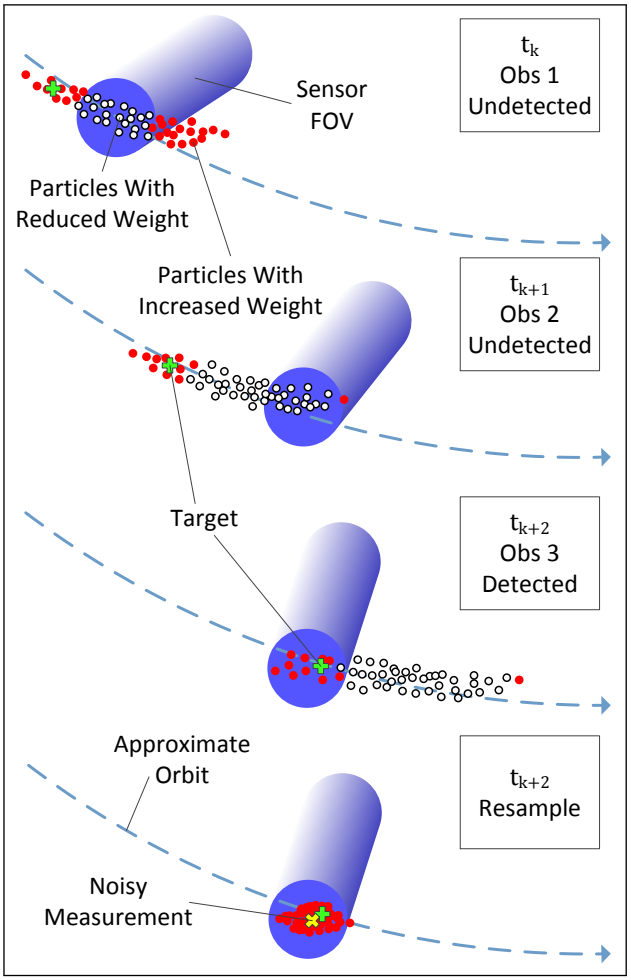

Figure 6.3: Illustration of a reacquisition achieved using the maximum weight steering strategy.

To accomplish the level of computation required for particle filtering, when SPARSE's various models required many—typically a multiple of N—independent executions, they were implemented as parallel code to be run on a GPU using the GPGPU framework as detailed in Chapter 5. Examples of some of the models required by SPARSE to monitor and control the modified PF as well as interface in real time with external systems include planetary motion, planetary orientation, lunar motion, solar illumination, orbit propagation, coordinate transformation and sensor operation.

SPARSE's high-level algorithmic architecture is displayed in Algorithm 2. The process begins with the download of the RSO catalogue and Earth Orientation Parameter (EOP) data from the internet. TLE sets were chosen as a source of authentic RSO data, as they can be readily obtained from the internet [39], are routinely updated and are compatible with a propagation model named Simplified General Perturbations model 4 (SGP4) [68]. The

Algorithm 2 SPARSE high-level pseudo algorithm.

Initialisation:
Obtain EOP data and RSO catalogue
Initialise particle distribution

Main Loop:
loop
  while *No Objects Visible* **do**
    $k := k + 1$
    Compute visibility for all RSOs at time k
  end while
  Choose a visible RSO
  Propagate particles to time k
  Run sensor steering step
  Create & send targeting data to sensor
  Observation is made
  Obtain:
    -Measurement angles
    -Observation evaluation parameters
  Run observation evaluation step
  Apply appropriate measurement update
  if *Required* **then**
    Apply progressive correction
    and/or resampling
  end if
end loop

SGP4 model was chosen for SPARSE as it was found to be parallelisable, with respect to time and each RSO, for fast execution. It is also capable of supplying sufficient accuracy to reliably locate RSOs when using recently published TLEs and an optical sensor with a 20 arcminute angle of view. Nonetheless, TLEs are not published with error covariance data. For this reason, once SPARSE has obtained TLEs for each of the objects in its catalogue, it subsequently creates a hypothetical particle distribution via the equation

$$\mathbf{x}^i = \mathbf{x}_o + \boldsymbol{\lambda}^i \tag{6.24}$$

where $\mathbf{x}_o$ is the original set of six SGP4-compatible mean-elements obtained from the TLE and $\boldsymbol{\lambda}$ is a vector containing a random sample of the initialisation noise. The TLE elements contained in the state vector are represented by

$$\mathbf{x} = [i\ \Omega\ e\ \omega\ M\ n]^{\mathrm{T}} \tag{6.25}$$

where i is the inclination, Ω is the right ascension of the ascending node (RAAN), e is the eccentricity, ω is the argument of perigee, M is the mean anomaly and n is the mean motion. Appendix A offers a description of how these TLE elements relate to the Keplerian elements discussed in earlier chapters. In conjunction with the TLE's epoch and ballistic coefficient B^*, this state data is sufficient to initialise the catalogue, propagate state estimates and target an RSO.

SPARSE uses EOP data to accurately relate inertial orbit frames to the Earth's surface. EOP data is measured, filtered and published by a number of sources. To obtain a comprehensive set of parameters, SPARSE obtains its EOP information—which includes precise timing information—from the International Earth Rotation Service (IERS), National Geospatial-Intelligence Agency (NGA) and the US Naval Observatory (USNO). Using this data, it is then possible to convert between the TLE compatible Cartesian frame True Equator Mean Equinox (TEME), inertial J2000 coordinates, ITRF and a sensor-relative South-East-Z (SEZ) frame [9] as necessary.

Following system initialisation, SPARSE enters the main observation loop in which a single track update is performed each iteration. Initially, the particle-estimated mean of each RSO is propagated to the present observation epoch and is used to determine if the RSO is currently visible to the observing sensor. If any objects are visible, SPARSE selects an RSO according to programmable selection criteria which may involve such features as the update history and the predicted observation effectiveness of each RSO.

Once an object is scheduled for observation, its particle distribution is propagated to the observation epoch. Thereafter, the particles are evaluated by the pre-programmed sensor steering step to obtain the desired pointing angles $\mathbf{\Psi}$ to aim the sensor.

After sending the steering commands to the sensor, SPARSE waits to receive back the measurements made by the sensor and the observation evaluation parameters. The parameters are used to decide how SPARSE should apply the appropriate measurement update step to the particle distribution. If SPARSE detects that the application of the weight update described by (6.12) will cause particle degeneracy, a progressive correction [102] technique is applied to concentrate particles around the observation. Similarly, if particle degeneracy is detected after the weight update (6.7), systematic resampling is applied.

This process can thereafter repeat indefinitely, enabling the system to maintain a track of its catalogue of RSOs. It should also be noted that although SPARSE's algorithm was described and implemented as an asynchronous process, there is no reason why a set of observations cannot be scheduled for each RSO over a finite horizon to enable the use of observation effectiveness optimisation procedures [50] to choose when to observe a specific target or to interleave search-observations inside a pre-existing schedule.

6.4 Numerical Simulation

Prior to field testing, SPARSE has undergone a dual-phased evaluation via numerical simulation. The first phase was designed to compare the outcome when dynamic steering is used in place of a traditional Gaussian based tracking system. The second phased was designed to verify the system is capable of practical implementation on a genuine steerable sensor and to determine which steering strategies are most effective for our application.

6.4.1 Numerical Simulation - via MASSAS (Phase 1)

Throughout Phase 1, SPARSE utilised MASSAS's in-built SSA-environment-simulation capabilities to evaluate the effectiveness of dynamic steering when compared to the use of a more conventional Gaussian estimator for maintaining a small catalogue of near-geostationary objects. As all aspects were simulated, all elements could be perfectly replicated—including time, sensor location and target truth data—to ensure a fair comparison between each tracking approach. During each simulation, the reliability of the sensor was assumed ideal such that, $p_{\mathrm{D}} = 1$ and $p_{\mathrm{FA}} = 0$.

A UKF [54, 56] was chosen over an EKF or RLS filter for running the conventional Gausian tracking scenario because—as detailed in earlier chapters—it can more accurately approximate the p.d.f. covariance for nonlinear systems and is consequently being considered as a possible replacement for current methodologies [22, 50]. The dynamically steered implementation was programmed to utilise the maximum weight steering strategy. This strategy was chosen as it was most comparable to targeting the mean of a UKF's Gaussian p.d.f. and, as explained by the following section, the results when utilising the remaining strategies were either irrelevant and/or uninformative for this exercise.

The test procedure began by obtaining 40 near-geosynchronous objects from the Space Track TLE repository [39]. Each object's state error was initialised with enough uncertainty such that there was approximately a 75% chance that the target would reside outside of the FOV if the state estimate was targeted by the sensor. This was achieved by utilising (6.24) in combination with

$$\boldsymbol{\lambda}^i \sim \mathcal{N}(0, \boldsymbol{\gamma}_\lambda^2);\ \boldsymbol{\gamma}_\lambda = \mathrm{diag}\begin{bmatrix} 0.01^\circ \\ 0.01^\circ \\ 0.00001 \\ 0.01^\circ \\ 0.5^\circ \\ 0.0001 \text{ orbits/day} \end{bmatrix} \tag{6.26}$$

to initialise the particle distribution. Subsequently, the 'truth' state of the object was drawn from the resulting distribution. The simulation was run twice, once using an RPF to perform dynamic-steering and once using the UKF. When using the UKF, its covariance and mean were initialised using the covariance and mean of the RPF's particle distribution and its state estimate—the UKF sigma point mean—was used to steer the sensor during each observation, to imitate existing steering procedures. If the target was not observed in the FOV, the observation was disregarded as though the observation was never scheduled. As the number of targets was relatively small and due to the computational power afforded by the GPU, the RPF was configured to use 50 000 particles per object and N_{Thr} was set to 100 particles.

To test the applicability of the dynamic steering method to HAMR objects, in contrast to 'cannonball-like' objects whose dynamics can be more easily modelled, 20 of the 40 objects were designated HAMR. This special designation informed MASSAS that during orbit propagation, a random perturbation was to be applied in order to simulate unpredicted perturbation of the RSO. The random perturbation was scaled to imitate the observed level of perturbation experienced by known near-GEO HAMR objects [27]. Consequently, after an empirical analysis on MASSAS, the error vector $\mathbf{v}_{\mathcal{H}}$ as defined by

$$\mathbf{v}_{\mathcal{H}} \sim \mathcal{N}(0, \gamma_{\mathcal{H}}^2);\ \gamma_{\mathcal{H}} = \text{diag} \begin{bmatrix} 0.0005^\circ \\ 0.0005^\circ \\ 0.00001 \\ 0.0005^\circ \\ 0.00001^\circ \\ 0.00001 \text{ orbits/day} \end{bmatrix} \tag{6.27}$$

was added to the state vector of each HAMR object prior to SGP4 propagation to achieve appropriate levels of unmodelled dynamics.

The simulated duration of the experiment was eight days. Each day, the simulated electro-optical sensor was permitted to make one set of five contiguous observations per RSO, to force SPARSE to find each object quickly and efficiently. The results of this comparative simulation are presented in Tables 6.1 & 6.2.

The results show a very clear difference in performance between the two tracking methods. Table 6.1 displays the comparison between UKF and the dynamically steered RPF for the 20 cannonball objects. It shows that 4 out of the 20 objects resided close enough to the UKF's mean that they were captured in the FOV on the first observation. In seven other cases, passive reacquisition is assumed to have occurred due to simulated perturbation forces on the RSOs and changes to the sensor-RSO geometry causing just enough relative movement

Table 6.1: Comparison between a dynamically steered RPF and a generic UKF implementation when observing 20 cannonball objects.

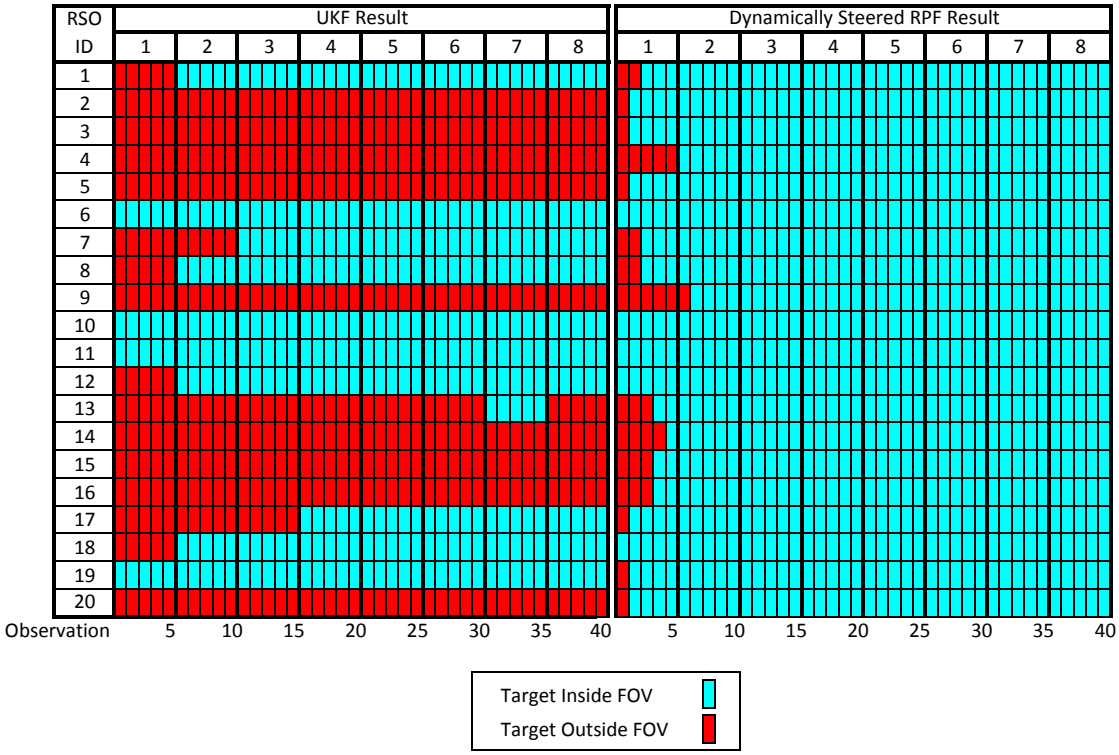

to enter the FOV. Nonetheless, slightly less than half of the RSOs were never observed, as failed observations could not improve the probability of targeting the RSO during subsequent observations. Conversely, the RPF augmented with dynamic steering, was able to reacquire all RSOs and track them for the remainder of the simulation. If a target was not observed in the first observation, the augmented RPF was capable of utilising the information to search and in the worst case, took six observations to reacquire an elusive target.

The results for the HAMR objects, as displayed in Table 6.2, report a similar outcome for the UKF implementation however, the number of instances when the target is outside of the sensor's FOV are noticeably elevated. This outcome is expected as the unmodelled motion of the simulated HAMR objects is likely to make them more challenging to track. In agreement with this hypothesis, an equivalent increase—albeit more sporadic—in the number of instances when targets are not observed in the FOV, whilst using the dynamically-steered implementation, is also present. Table 6.2's RPF results demonstrate an important attribute of the dynamically steered implementation. Although it was challenging to maintain a track of HAMR targets, each time a target was lost, the system autonomously and immediately began looking for the target once again. Notably, the dynamically steered RPF switches

Table 6.2: Comparison between a dynamically steered RPF and a generic UKF implementation when observing 20 HAMR objects.

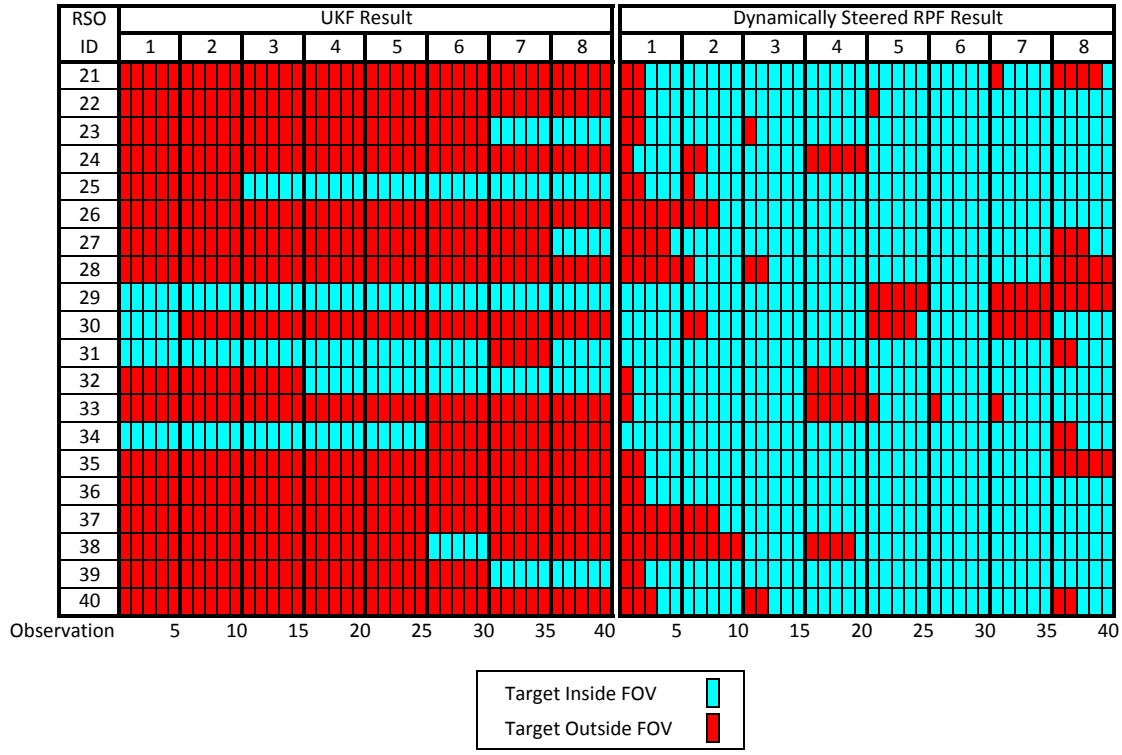

between tracking and searching roles as required. The resulting rate of reacquisition is encouraging. On average, a lost target is reacquired after three observations spent searching for it.

6.4.2 Numerical Simulation - via TASMAN (Phase 2)

The objective of the Phase 2 simulations was to verify that SPARSE is capable of practical implementation by replacing MASSAS's simulated 'truth' data with a mock-sensor-interface provided by a high-fidelity, third-party SSA simulation. Thanks to AFRL and PDS, time utilising TASMAN—as previously introduced in Chapter 3—was very kindly offered to satisfy this role. Unlike Phase 1, this phase requires practicalities such as causal tasking of the sensor, anticipation of system latencies as well as implementing functions to format data to send to, and parse information returning from, the sensor.

The RPF was again configured as described in Section 6.2.2. However, unlike Phase 1, each of the proposed sensor steering methods were trialled in four separate tests and the observation evaluation element was implemented using the operational characteristics of a Raven class electro-optical sensor. Observation evaluation metrics $\boldsymbol{\Omega}_k$ such as the number

of objects within the FOV and the number of stars observed were therefore obtained and evaluated during the tests. Respectively, these parameters inform the observation evaluation method of obscuration of the target or inclement weather and, to work within the constraints of the experimental system, reduce the likelihood of misassociation of observation data. As the experimental system did not have the capacity to determine a quantitative assessment of p_{D} and p_{FA}, the observation effectiveness element assessed the outcome of each observation qualitatively. The outcome was defined as one of the three following states:

1. WITHIN FOV - the target was observed in the FOV,
2. OUTSIDE FOV - the observation is valid but the target was not observed in the FOV,
3. INVALID - the target can not be confidently declared in or outside of the FOV.

In the event an image was returned by the sensor and the data association was poor, contained more than one object or not enough stars to be sure the sensor had a clear view of the sky, SPARSE considered $\mathbf{z}_k = \boldsymbol{\varepsilon}$, $p_{\mathrm{D}} = 0$ and $p_{\mathrm{FA}} = 0$. The observation was therefore INVALID and had no impact on the state error distribution. If the observation was valid, SPARSE considered $p_{\mathrm{D}} = 1$, $p_{\mathrm{FA}} = 0$ and (6.22) was employed to implement the measurement update. If the target was observed in the FOV, the WITHIN FOV state was accepted and a measurement consisting of right ascension and declination, such that $[\mathbf{z}_{k,1}, \mathbf{z}_{k,2}]^{\mathrm{T}} \in \mathbb{R}^2$, was anticipated. Conversely, if the object was not observed but a valid observation was made, an OUTSIDE FOV state resulted in $\mathbf{z}_k = \boldsymbol{\varepsilon}$. Whilst the system was tested to see if it would respond correctly to obscuration of the stars due to cloud, the results of these simulations have not been included in this section as they add little to the current analysis and there are many examples included in the field trial results shown in the following section.

The previous phase indicated that SPARSE is capable of, but not necessarily limited to, finding objects that have not been observed by a space sensor network for some time and objects whose orbital dynamics are not easily modelled, such as HAMR objects. To test this capability, the Phase 2 numerical simulation methodology involved the acquisition of a range of historical TLEs from the Space Track website for 2 generic near-GEO objects and 2 HAMR near-GEO objects, as displayed in Table 6.3. The age of the historical TLEs were selected to provide a contrasting range of targeting data as a TLE's error is presumed to increase with age. TLEs of ages 0, 60, 120, 150, 180 and 210 days old were obtained from the Space Track repository [39]. Thereafter, SPARSE was instructed to locate each of the objects using an independent particle distribution for each TLE. This meant that although in reality, only four objects were required for simulation by TASMAN, SPARSE was instructed to observe up to 24 unique targets. As HAMR objects can be difficult to

Table 6.3: Simulated targets for Phase 2's numerical simulation.

US Cat. ID	Description	HAMR
29106	MSG 2 DEB (COOLER COVER)	Yes
29676	MSG 2 DEB (BAFFLE COVER)	Yes
858	SYNCOM 3	No
2639	ATS 3	No

track and observe, TLEs for HAMR objects are not published as regularly as others objects. Consequently, only five historical TLEs for object 29676 were obtained, resulting in a total of 23 unique TLEs for reacquisition and tracking by SPARSE. Evaluation of the resulting test data therefore provides a comparison between SPARSE's ability to reacquire and track conventional and HAMR targets with varying degrees of initial state estimate accuracy. The truth information, required by TASMAN to simulate the four 'true' objects, utilised the most recent of each object's historical TLEs.

As considered earlier in Section 6.3, initialisation of the particle distribution from TLEs relies upon a distribution generation method (6.24) and appropriate values for $\boldsymbol{\lambda}$. Whilst an existing space surveillance system would have an active catalogue of RSO's orbital state uncertainties, a single set of noise values were chosen to be added to each RSO's TLEs. This meant that the relative state uncertainty between RSOs was dependant on TLE age alone. The magnitude of noise was chosen empirically such that any unmodelled motion by the SGP4 propagator was captured by the RPF when using the latest, and therefore most accurate, TLE. Consequently, random values for $\boldsymbol{\lambda}^i$ were obtained by means of the equation

$$\boldsymbol{\lambda}^i \sim \mathcal{N}(0, \boldsymbol{\gamma}_\lambda^2);\ \boldsymbol{\gamma}_\lambda = \text{diag}\begin{bmatrix} 0.01^\circ \\ 0.005^\circ \\ 0.00001 \\ 0.1^\circ \\ 0.1^\circ \\ 0.00005 \text{ orbits/day} \end{bmatrix}. \tag{6.28}$$

Similarly, appropriate levels of process error were necessary to account for perturbations that are modelled by TASMAN's high-fidelity propagator that are not modelled by SPARSE's SGP4 propagator. An empirical investigation was conducted to learn how changes in process error, as modelled by zero-mean additive Gaussian noise, to each of the orbital elements influenced the state uncertainty with time. Of particular interest were the effects to in-track, cross-track and exo-planar uncertainty. The conclusion of this study suggested that

appropriate levels of process noise could be generated for the conventional objects using,

$$\mathbf{v}^i \sim \mathcal{N}(0, \boldsymbol{\gamma}_{\mathrm{v}}^2);\ \boldsymbol{\gamma}_{\mathrm{v}} = \mathrm{diag} \begin{bmatrix} 0.893 \times 10^{-6^\circ} \\ 18.380 \times 10^{-6^\circ} \\ 0.067 \times 10^{-6} \\ 1553.731 \times 10^{-6^\circ} \\ 1557.268 \times 10^{-6^\circ} \\ 0.103 \times 10^{-6}\ \text{orbits/day} \end{bmatrix} \tag{6.29}$$

and for HAMR objects

$$\mathbf{v}^i \sim \mathcal{N}(0, \boldsymbol{\gamma}_{\mathrm{v}}^2);\ \boldsymbol{\gamma}_{\mathrm{v}} = \mathrm{diag} \begin{bmatrix} 30 \times 10^{-3^\circ} \\ 0.2 \times 10^{-3^\circ} \\ 0.0005 \times 10^{-3} \\ 50 \times 10^{-3^\circ} \\ 50 \times 10^{-3^\circ} \\ 0.001 \times 10^{-3}\ \text{orbits/day} \end{bmatrix}. \tag{6.30}$$

To simulate appropriate levels of process error, the resulting process noise was incorporated in truth-state propagation by adding the appropriate noise vector to each object's state vector prior to SGP4 propagation.

To generate high-fidelity truth data, TASMAN requires more than just a TLE to initialise and propagate the truth-object's orbital motion. In addition to state information, TASMAN's propagator requires information about the shape, mass, orientation, rotation and reflectivity of the object's surfaces. This information enables TASMAN to simulate the unique orbital trajectories of HAMR objects by including influences such as transient SRP (solar radiation pressure) [27]. While the parameters for the non-HAMR objects were chosen to ensure visibility, the parameters for the HAMR objects were obtained from [108] to ensure their orbital-dynamics were as realistic as possible. The noteworthy parameters supplied for each of the objects are detailed in Table 6.4.

Table 6.4: Truth object initialisation parameters.

US Catalogue ID	29106	29676	858	2639
Shape	Flat Plate	Flat Plate	Sphere	Sphere
Surface Area (m^2)	1.9	0.7	42.3	36.7
Mass (kg)	8.4	4	822.2	691.6
Area to Mass Ratio (m^2/kg)	0.23	0.18	0.05	0.05

Table 6.5: A record of the ability of a MASSAS-based simulated Raven sensor to acquire and track four near-GEO MASSAS simulated objects over 8 days using an RPF to implement dynamic steering.

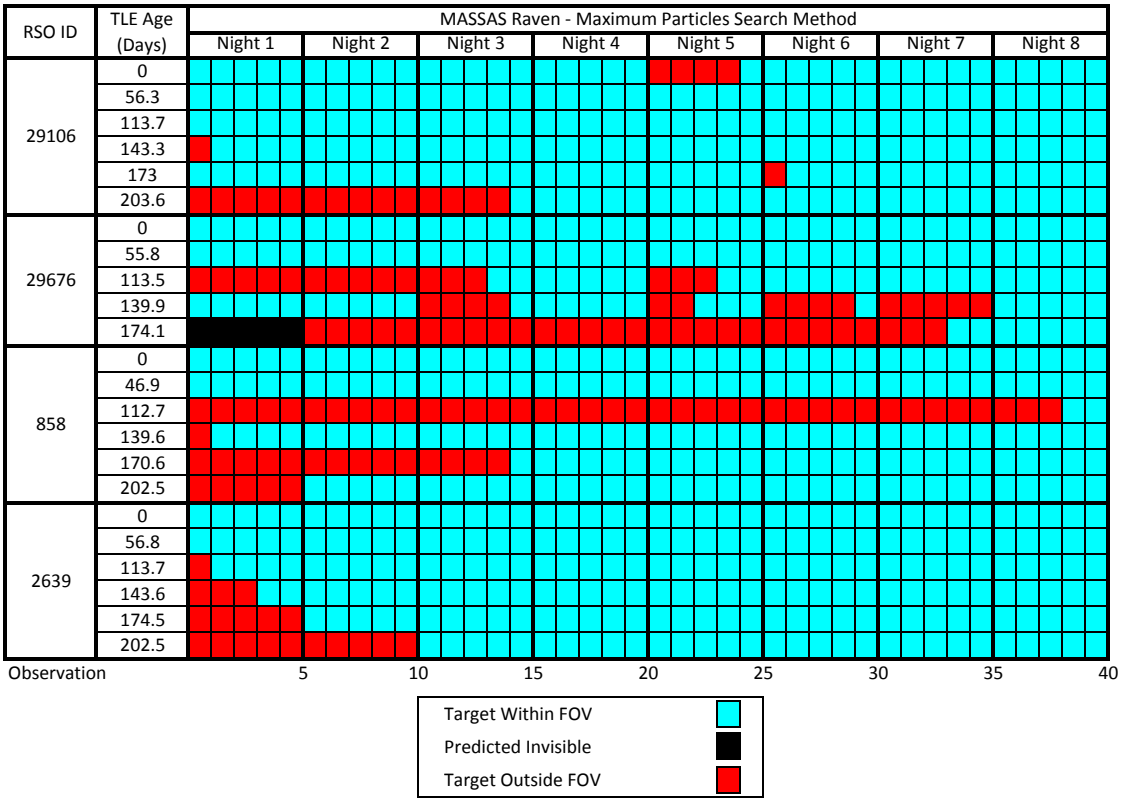

Transitional Testing

For verification of the Phase 1 results and to compare the difference upon transitioning from the lower fidelity MASSAS environment to the higher-fidelity TASMAN environment, the Phase 2 configuration was simulated on MASSAS first, utilising an internal simulated raven sensor mode. Each simulation was once again run for a length of 8 days, and each historical TLE's particle distribution was supplied a single track update of five observations per day. Utilisation of alternative search methodologies will be discussed and evaluated later in this section, however for an initial comparison Table 6.5 and Table 6.6 show the result of utilising an RPF configured using the maximum particles search method using both systems.

Table 6.5 displays the result of performing the Phase 2 simulation on MASSAS. Three important features are shown. All four objects are reacquired regardless of the age of the TLEs used to reacquire their respective targets. The results also show an increasing, albeit non-monotonic, rate of detection as the age of a target's TLE increases. Similar to Phase 1's

Table 6.6: A record of the ability of a TASMAN-based simulated Raven sensor to acquire and track four near-GEO TASMAN simulated objects over eight days using an RPF to implement dynamic steering.

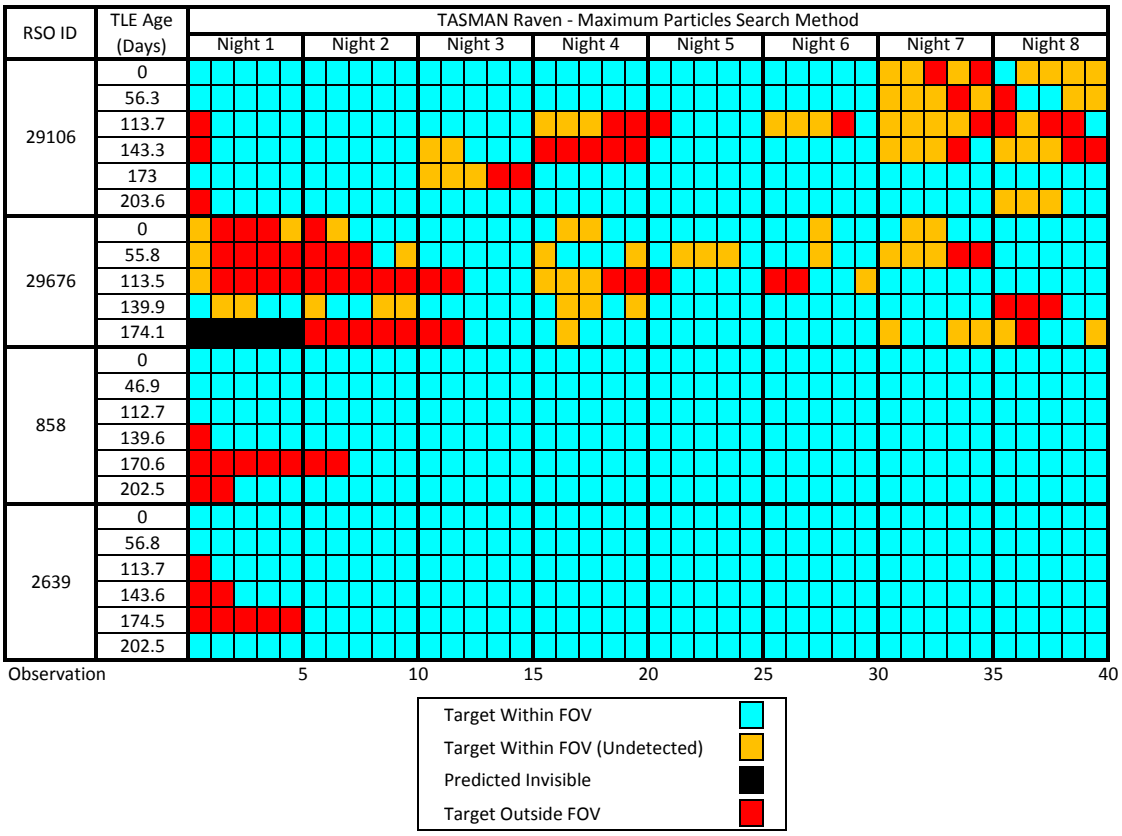

results, the HAMR objects occasionally require a brief period of reacquisition due to their unmodelled motion. An additional noteworthy feature is object 29676's predicted invisibility to the sensor on Night 1 when using its oldest TLE's distribution. As object 29676 was observed using other TLE distributions, this indicates that although it was visible, the oldest TLE's distribution mean was not. This feature highlights an area for future consideration regarding the reliability of utilising the state estimate to predict visibility. Particularly for objects whose state error p.d.f.s are bisected by the sensor's local horizon or other visibility limitations.

Table 6.6 displays the result of running the same test procedure on the combined SPARSE-TASMAN system. Observation requests formatted for Raven sensors were produced by SPARSE and sent to TASMAN. Thereafter, TASMAN utilised its internal high-fidelity models to simulate the SSA environment during each observation and return an appropriately formatted response for a Raven sensor. For verification, TASMAN also returned a truth-file including the precise state of each target during each observation. Table 6.6 displays the

same three features identified in MASSAS's results in addition to some aberrant behaviour demonstrated by a much reduced rate of reacquisition of the HAMR objects. At first glance, this data appeared to indicate that the HAMR objects are much more challenging to acquire than MASSAS had predicted. This hypothesis was nonetheless dispelled when the aforementioned verification files were scrutinised. Upon comparing TASMAN's truth state vectors with SPARSE's steering information, it was found that both HAMR objects were routinely within the sensor's FOV, but on many occasions TASMAN's sensor model concluded that they were not visible. As acknowledged by TASMAN's operators, the lapses in visibility may have been caused by an overly sensitive illumination model and/or the HAMR object's evolving orientation, which consequently caused a time evolving cross sectional area and luminosity of reflected light. Consequently, it was concluded that non-simulated observations of these particular objects were necessary to calibrate and resolve the illumination issues and obtain a true value for p_{D}. In spite of this, SPARSE demonstrated an ability to robustly reacquire these challenging targets, even in the presence of unexpected visibility constraints. Furthermore, the results indicate that MASSAS's propagation model made RSOs slightly more challenging for SPARSE to track than would necessarily be expected in a field trial. For all of these reasons, testing of field trial readiness was continued.

High-Fidelity Evaluation of the Proposed Steering Methods

As the validity of TASMAN's sensor model could not be confirmed prior to field trials, TASMAN's Raven sensor model was modified to increase its sensitivity such that the HAMR objects would be reliably detected each time they entered the FOV. This permitted subsequent testing of the four proposed steering methodologies with independence from the unknown constraints to visibility.

After the modification of TASMAN's sensor model, new TLEs were obtained for trialling each of the search methods proposed in Section 6.2.2. To aid in discerning the relative effects of utilising TLEs of various ages, the spacing between the ages of TLEs was reselected to achieve a consistent 50 day increment in age. Consequently TLEs of ages 0, 50, 100, 150 and 200 days old were sought for each object. A full set of five historical TLEs were obtained for objects 29106, 858 and 2639, but only three TLEs were available and appropriate for object 29676.

Prior to detailing and discussing the effectiveness of each steering method, the presence of a small number of invalid observations in some of the results should be noted. At the time of these simulations, early September 2013, object 29106—the MSG 2 cooler cover debris—and object 29676—MSG 2 baffle cover debris—were passing each other so closely that on the

3rd night of simulation, invalid observations were recorded as TASMAN's simulated sensor could occasionally observe both objects in the FOV. In accordance with the observation evaluation measures described previously, these observations were labelled invalid and discarded.

Maximum Weight Steering Method:

The first of the four steering methods to be tested, utilising the updated system configuration, was the maximum weight steering method. Table 6.7 displays the results of utilising the maximum weight steering method to search for the four targets utilising the 18 historical-TLE distributions. The results produced by TASMAN's high-fidelity simulation are encouragingly similar to the Phase 1 results produced by MASSAS. Any targets that were not immediately observed in the FOV were, in all but a single case, quickly reacquired. As the dynamics of objects 858 and 2639 could be more accurately modelled by an SGP4 propagator than objects 29106 and 29676, once acquired 858 and 2639 were flawlessly tracked for the remainder of the simulation. Whilst 29106 and 29676 were occasionally lost by the sensor, they were routinely reacquired.

Table 6.7: The recorded states of eight nights of high-fidelity simulated observations of four objects whilst using increasingly aged targeting data and employing the maximum weight steering method to perform dynamic steering.

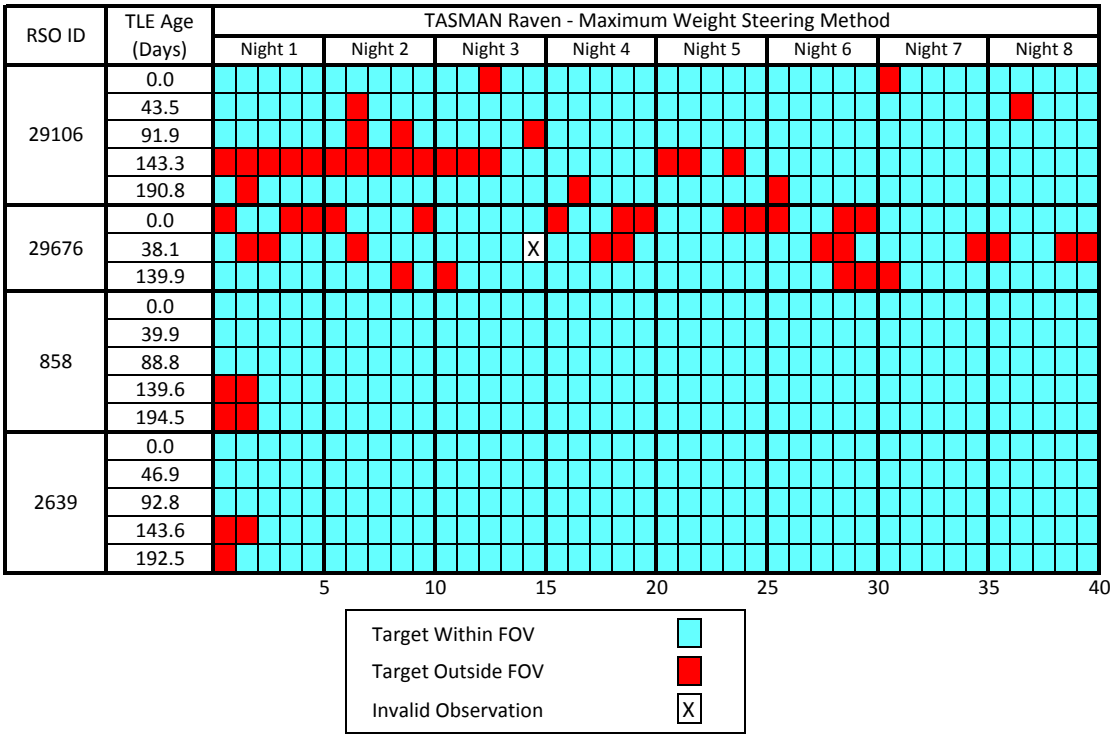

Minimum Variance Steering Method:

When the minimum variance steering method is applied to the same Phase 2 simulation configuration, a very different outcome is observed. Table 6.8 displays the results of the second test, which employed the minimum variance method for steering the sensor. Whilst the distributions of the most recent TLEs could be used for reacquisition, distributions resulting from TLEs approximately 100 days in age or older could not be used to reacquire objects within the simulation time frame. The only exceptions are object 29106's 100 and 150 day old TLE distributions which were used successfully on Night 2—presumably due to unmodelled motion of the target. Table 6.8 & 6.7's results indicate that it is more beneficial to be optimistic and steer toward the region of highest probability than to be pessimistic and aim to reduce the dispersion of the p.d.f. as quickly as the FOV permits.

Table 6.8: The recorded states of eight nights of high-fidelity simulated observations of four objects whilst using increasingly aged targeting data and employing the minimum variance steering method to perform dynamic steering.

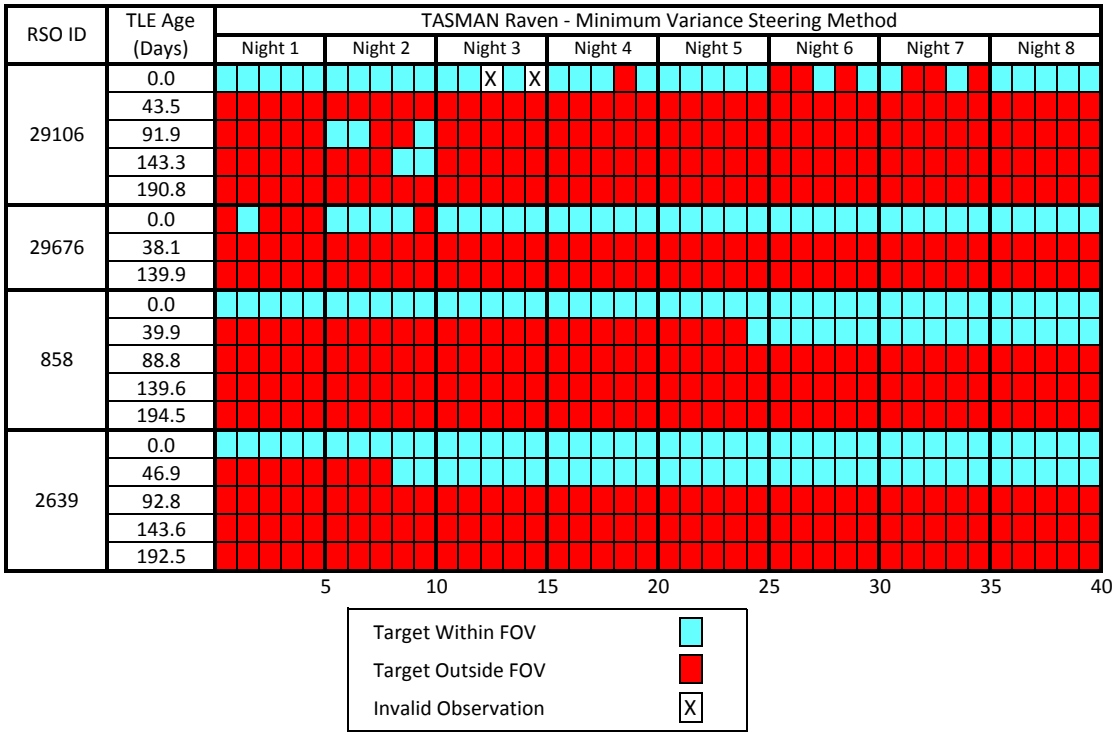

This conclusion is supported by Fig. 6.4 which shows the standard deviations of the positions of particles within each distribution, throughout each of the simulations. The values of standard deviation have been grouped and averaged over the set of distributions belonging to each of the TLE ages used to initialise each distribution. As intended, Fig. 6.4

indicates that the minimum variance strategy consistently reduces the standard deviation—and therefore the variance—of the position of each distribution's particles throughout the simulation. Nonetheless, as $p_{\text{FA}} = 0$ is assumed and due to the accuracy of the sensor, the variance drops much more significantly when the object is observed in the FOV than when the tails of the distribution are observed. Although the minimum variance strategy causes the variance to consistently fall, the maximum weight strategy is locating the RSOs much more quickly and consequently reducing the variance much more effectively over the length of the simulation.

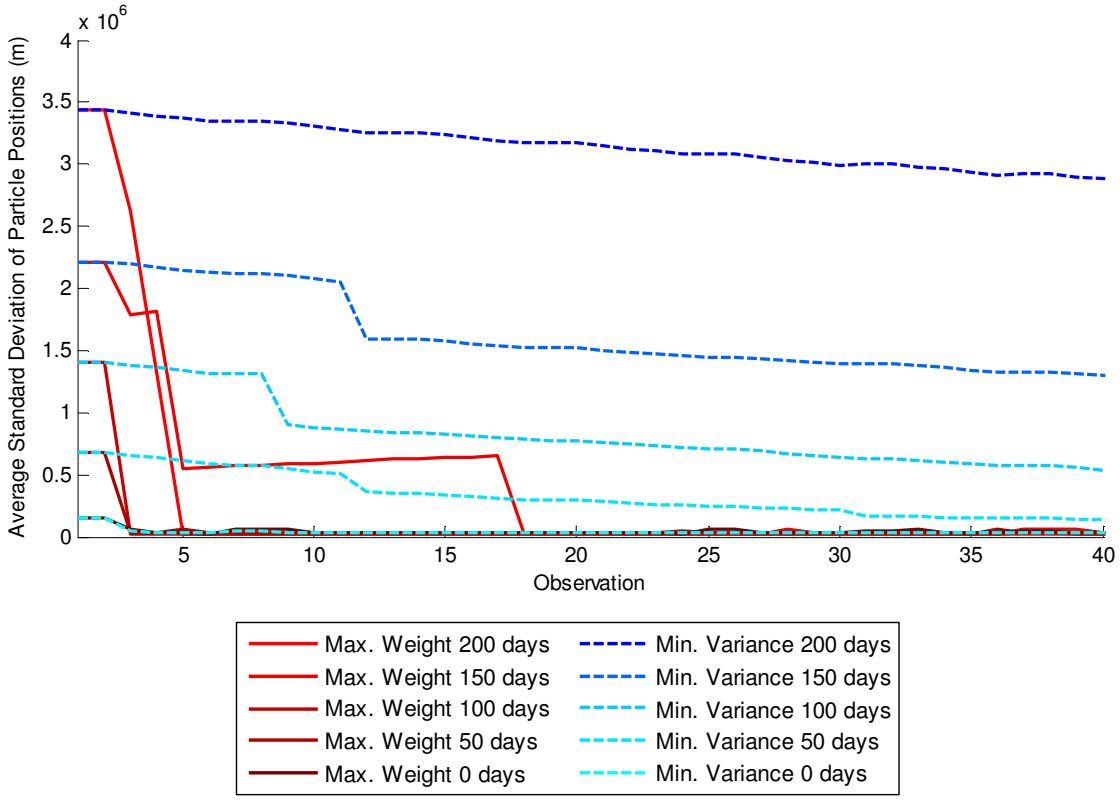

Figure 6.4: The average—grouped by TLE age—standard deviation of the RPF's particle positions throughout the simulation whilst utilising the maximum weight and minimum variance steering methods.

Whilst the maximum weight strategy is clearly a much superior choice for the sensors and targets in these simulations, the effectiveness of either method is related to the rate of error growth or diffusion in the p.d.f., the physical characteristics of the sensor and the objects involved in the simulation. The minimum variance strategy may therefore be of greater benefit for different configurations and applications.

Raster Steering Method:

The raster steering method, as detailed in Section 6.2.2, is proposed as a method that might be adopted by a human operator, to scan along the estimated path of a lost RSO. Whilst alternative methods may be preferred by operators to locate lost RSOs, a 'nose to tail' raster implementation has been chosen as a naïve brute-force approach that moves methodically along the distribution until the target is detected. The intention is to demonstrate the amount of time it would take to work methodically along the high probability region as defined by the particles to locate the target. The raster method uses the prior distribution at each observation to find the 'front' of the high probability region defined by the RPF's particles. Therefore, if the sensor's FOV is wider than the distribution's growth in error between each observation, the sensor will progress along the distribution faster than it can grow, eventually converging on the target. Nonetheless, if the opposite is true, the particle distribution can never be fully observed and the target may remain lost.

The results of a third simulation, which employed the raster steering method, are presented in Table 6.9. Whilst the sensor is shown to have succeeded in reacquiring and tracking the target in some instances, the results are generally poorer than demonstrated by the minimum variance steering method. The non-HAMR objects 858 and 2639 were observed using their most recent TLE distributions in a single night. Their 50 day old TLE distributions permitted reacquisition after four nights spent searching. Other than some sporadic success finding the HAMR objects, most other distributions were unsuccessfully used.

If nothing else, the results generated by this steering method indicate that the relative growth in the p.d.f. with respect to the sensor's FOV makes a methodical search a poor choice when compared to maximum weight steering. Nevertheless, in spite of failing to reacquire most targets, information was gained from each and every observation. The chances of reacquisition occurring if the resulting data was used to cue another sensor, was thereby improved. If human operators were to employ similar or more successful strategies, but failed to reacquire the target, they are currently unable to make similar use of the information obtained by their failed observations, thereby receiving no gain for their resource expenditure. It is therefore suggested, that if a dynamically steered system were to permit manual steering of the sensor, observation evaluation should remain active to record the findings, irrespective of the outcome.

Maximum Particles Steering Method:

When employing dynamic steering, the ability of the sensor to converge on the target is dependent on the size of the FOV relative to any diffusion of the p.d.f. with time. Whilst there is limited control over these factors, there are operational circumstances that may

Table 6.9: The recorded states of eight nights of high-fidelity simulated observations of four objects whilst using increasingly aged targeting data and employing the raster steering method to perform dynamic steering.

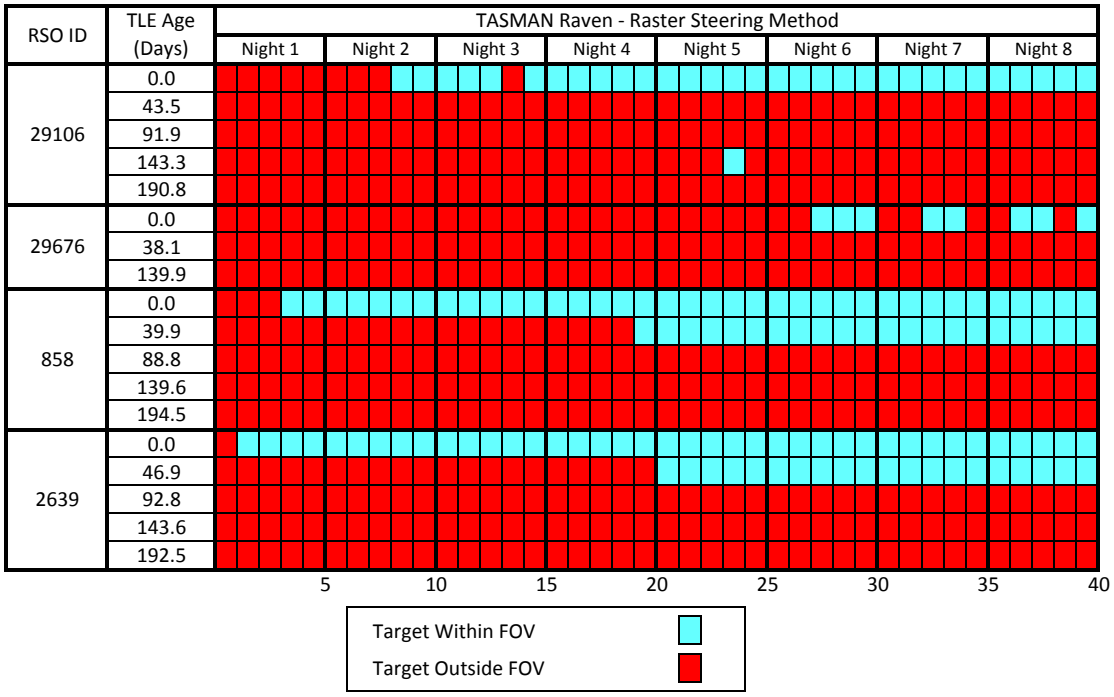

increase diffusion, causing the rate of convergence to reduce. An example would include a situation when the filter has been prepared for the manoeuvring of a target or unmodelled motion, by inflating the system process error with additional noise to capture the unmodelled but bounded motion. It is in this situation that the motivation behind the maximum particles objective function is revealed.

A noticeable reduction in convergence, between the FOV and the target, was observed when the maximum weight steering method was employed to reacquire HAMR objects. Convergence was observed using an augmented version of MASSAS's visualisation module, designed to visually monitor the dynamic steering process. Appendix B: Figs. B.9-B.13 display example visualisations produced by MASSAS whilst performing dynamic steering. Fig. 6.5 displays a simplified illustration of the observed behaviour. When utilising maximum weight search and employing inflated process error to capture the unmodelled HAMR dynamics, the sensor was noted to continually 'look' in the same region of the p.d.f. over multiple observations. High-weight particles were observed to bleed back into the newly nullified region due to the inflated process noise, even for relatively small time steps. Whilst the sensor *should* continue to look in this region to maximise the probability of detection with respect to the particle estimate, the sensor is less likely to look about the distribution

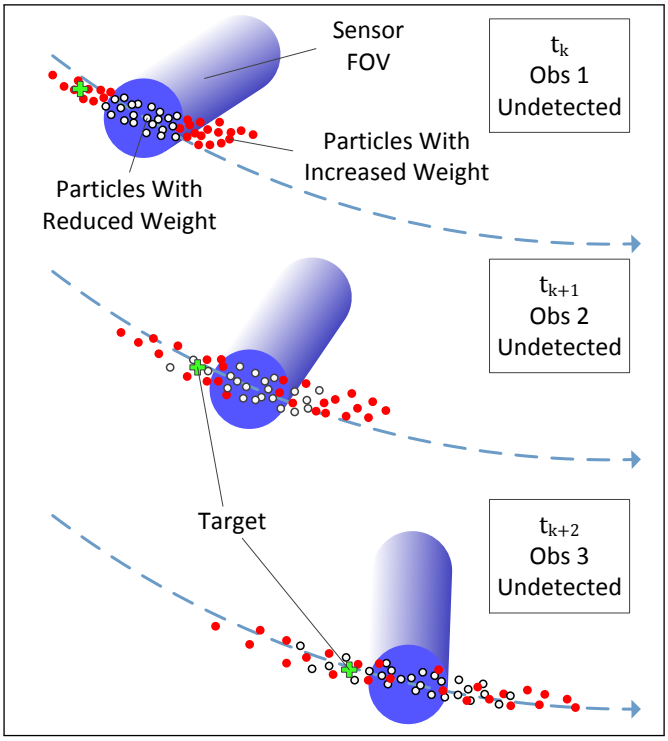

Figure 6.5: Attempted reacquisition via maximum weight steering in the presence of inflated process error.

and locate the target quickly. It is therefore proposed that in this situation, it may be more prudent to pick a steering solution whose FOV includes the most non-trivial particles, rather than the highest weight.

If the p.d.f. is unimodal and the particles have recently been resampled resulting in even weights, maximising the number of particles in the FOV results in the same solution as the maximum weight steering strategy. Nonetheless, the inevitable multimodal form of the distribution will force the maximum particles search method to either side of the nullified region as it is deterred from centring over the nullified region. Consequently the hypothesis is: by including as many non-trivial particles in the sensor's FOV as possible, the sensor is more likely to look about the distribution and locate the target more quickly.

This hypothesis was tested on TASMAN using the Phase 2 simulation configuration. The results are presented in Table 6.10. As $p_{\mathrm{D}} = 1$ and $p_{\mathrm{FA}} = 0$ were utilised when an object was observed, the threshold for defining a trivial weight was zeroed, such that $w_{\mathrm{Thr}} = 0$ when evaluating the maximum particle objective function (6.21). In general, Table 6.10 indicates that steering toward the most non-trivial particles is similarly effective to steering toward the maximum weight. Both methods appear to achieve somewhat comparable results. Nevertheless, a difference in the speed at which RSOs are reacquired is more easily observed

Table 6.10: The recorded states of eight nights of high-fidelity simulated observations of four objects whilst using increasingly aged targeting data and employing the maximum particles steering method to perform dynamic steering.

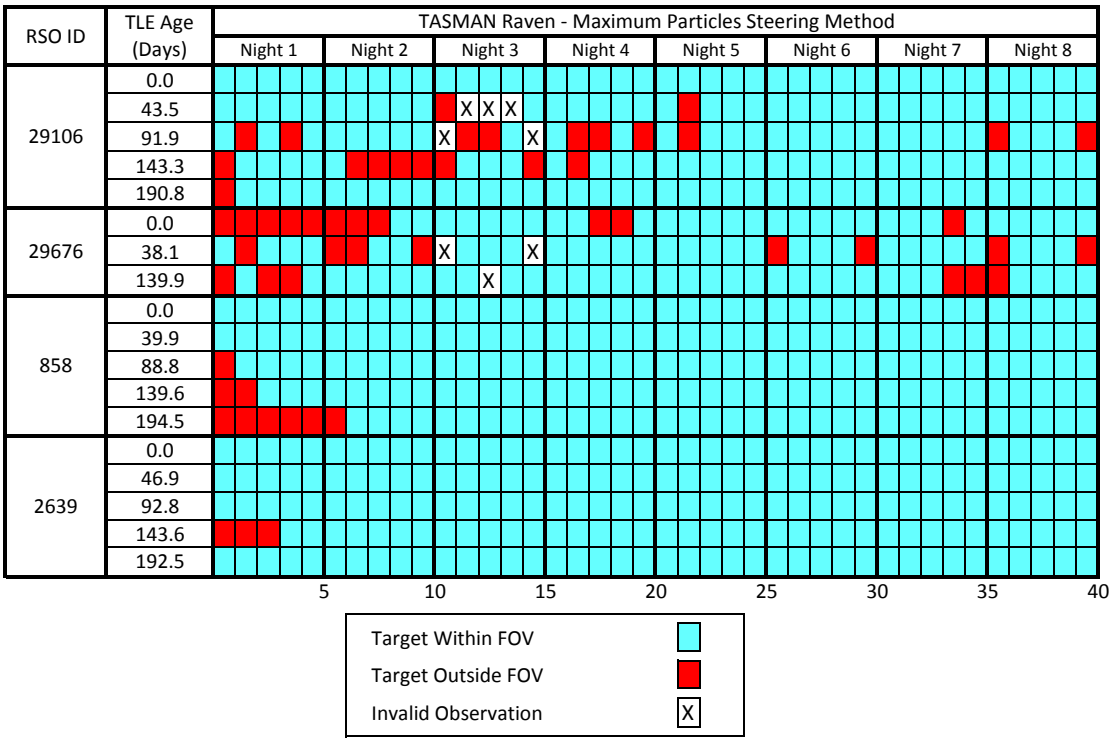

in Table 6.11, which displays the average number of contiguous observations in which a target was not observed in the FOV for each steering method and RSO type. Table 6.11 indicates that there was some improvement in the speed at which HAMR targets were reacquired when using the maximum particle method, as the average drops from 1.9 to 1.6 observations. Nonetheless, there is also a noticeable detrimental effect to the speed at which the cannonball targets could be reacquired. The average number of contiguous observations not containing the target is shown to increase from 1.8 to 3.0.

Table 6.11: The average number of contiguous observations when RSOs were outside of the FOV when employing maximum weight and maximum particles steering methods.

RSO Type	Maximum Weight	Maximum Particles
HAMR	1.9	1.6
Cannonball	1.8	3.0
Both	1.9	1.8

This evidence suggests that there is some merit to the hypothesised gain in speed of reacquisition when using maximum particle steering to find an object with inflated process error. But it appears to be wiser still to select a steering regime appropriate to the level of process error necessary to capture the target's unmodelled dynamics, to achieve the greatest benefit from either strategy.

6.5 Field Trial of Dynamic Steering

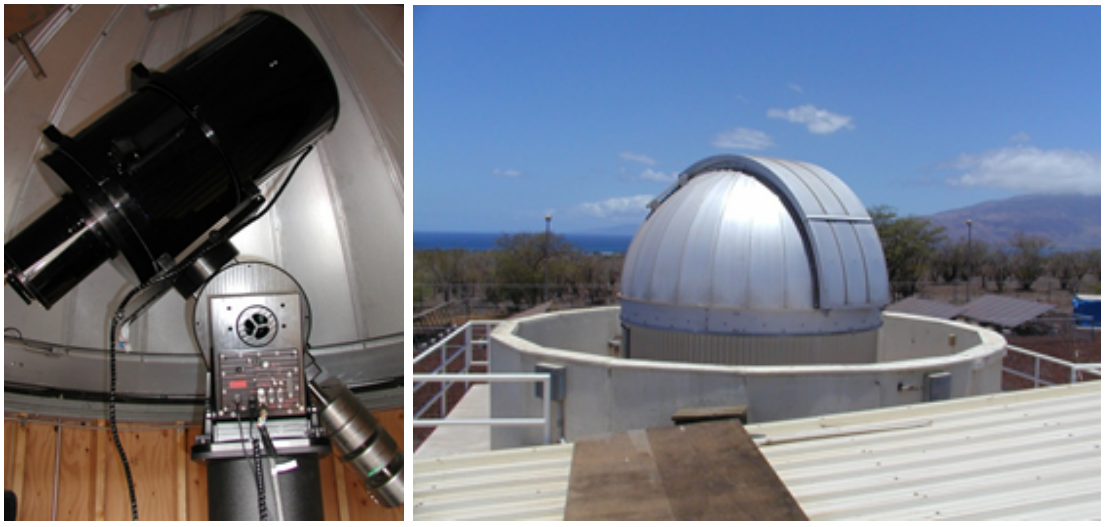

Figure 6.6: Photos of an example Raven class electro-optical sensor in Kihei, Hawaii. Photographs graciously provided by Nelson [91].

In preparation of a field trial, but prior to running the Phase 2 numerical simulations on TASMAN, AFRL and PDS offered time utilising an operational Raven class sensor for precursory tests. Therefore, to verify SPARSE's sensor interface was operationally compliant, an open-loop test was performed throughout February and March of 2013. Many of MASSAS's internal models were utilised to generate a schedule to regularly observe any objects that were orbiting between the altitudes of approximately 20 000 km 36 000 km, within the unclassified Space Track repository [39] and visible at the sensor's location in Learmonth, Australia. At the time, these requirements resulted in a list of up to 1175 objects to be scheduled for observation. Throughout the testing period, MASSAS performed very capably as most objects were regularly observed. Of the 4-9 objects that failed to be observed each evening, some of these objects were already noted by PDS staff to be notoriously difficult to acquire using a Raven sensor. Other objects failed to be observed due to their visual-proximity to the moon. It was during this time that a lunar model was integrated

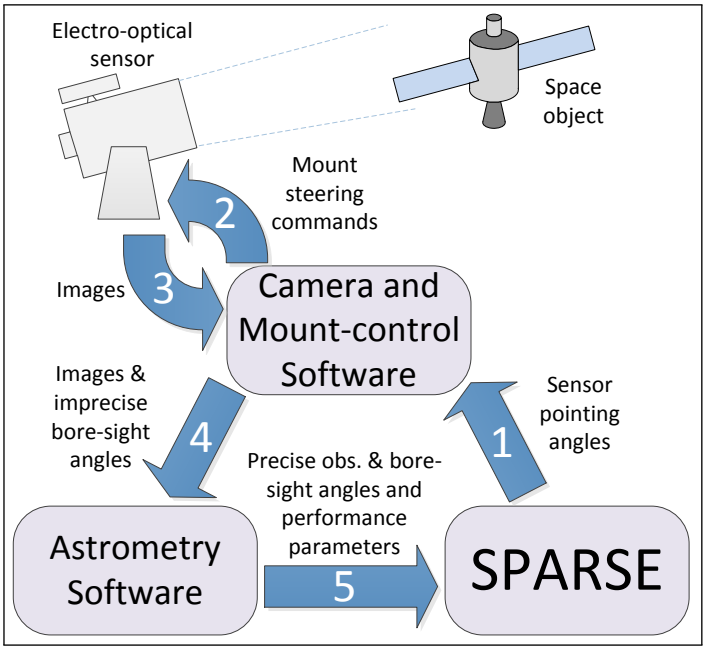

Figure 6.7: Architecture of the experimental dynamically steered system.

into MASSAS to avoid further issues due to saturation of the sensor by moon-light. Other than these minor cases of failed observations, these initial trials demonstrated that SPARSE would be capable of successfully communicating with an AFRL Raven class sensor and that MASSAS's level of fidelity was sufficient to consistently steer a sensor and observe the desired target, when using accurate targeting data.

Due to the success of these precursory field trials and the Phase 2 numerical simulations, a field trial of an experimental dynamically steered system, comprising of SPARSE and a single electro-optical sensor, was performed. The field trials occurred throughout October of 2013 with the cooperation of AFRL at Maui, under the leadership of Dr. Kim Luu, and with the support of PDS. Controlled access to a Raven class electro-optical sensor in Kihei was graciously offered by AFRL for the extent of the field trial. An image of a Raven sensor, similar in design and location to that used during the field trial, is displayed in Fig. 6.6. The high-level system architecture of the experimental system is shown in Fig. 6.7. The arrows indicate the direction and order in which information was passed throughout the system in a dynamic loop.

Aside from incorporating some additional site specific steering metrics into SPARSE to prevent steering the sensor to undesirable elevations, SPARSE was left largely unchanged from Phase 2 numerical simulations. As the primary control software, SPARSE initiated observations by sending commands to the sensor's camera and mount-control software and received the results from the astrometry software named Astrograph, previously introduced

in Chapter 3. The camera and mount-control software were used to command the sensor's mount to move to the pointing angles provided by SPARSE to subsequently take images with the camera. Fig. 6.8 displays an example image which demonstrates the preferred rate-tracking method used to follow the target and integrate its light, ideally, on a single pixel. This process causes the stars to streak and the target to appear as a bright dot. The astrometry software was used for processing the resulting images to obtain measurements with 1 arcsecond standard deviation of error. In addition, the astrometry software returned information about data association confidence, the number of objects in the FOV, a qualitative assessment of occlusion by weather and the precise borsesight angles of the sensor to enable SPARSE to accurately model the sensor's FOV after each observation. Combined, these metrics comprised the system's observation performance parameters used during the observation evaluation step. SPARSE's observation evaluation was configured as previously described in the Phase 2 numerical simulations.

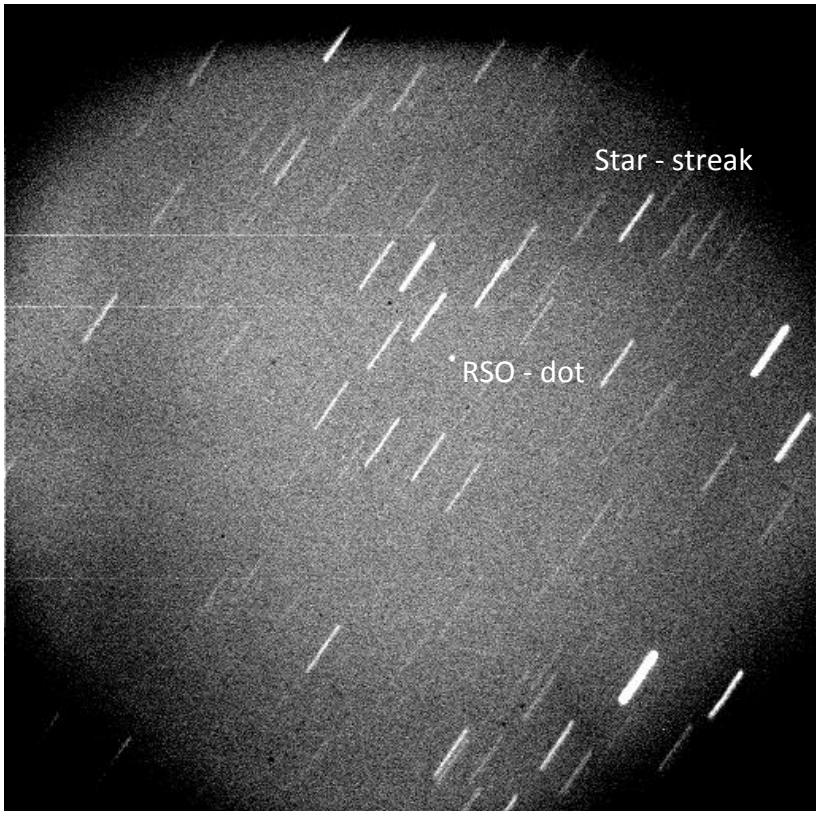

Figure 6.8: An image captured by a Raven sensor showing the rate-tracking technique used for robust astrometry.

The primary objective of the field trial was to test the ability of the proposed method to reacquire objects that would be deemed lost to a system using the prevailing methods for maintaining SSA. An objective definition of a lost RSO is not easily obtained. Nonetheless,

TLEs are known to exhibit relatively low precision due to arbitrary parameter rounding that causes error that worsens with orbital altitude and time [9, 39]. Furthermore, it is generally accepted in the US surveillance community [19, 22] that once an object has not been observed for greater than 30 days, the object is deemed lost as the orbital estimate is now likely to be unreliable. The continued use of TLEs of various ages greater than 30 days old was therefore accepted as a source of targeting information that was both imprecise and old enough to claim that it would not be considered for use in an existing space surveillance system. Therefore, TLEs that were approximately 0, 50, 100, 150 and 200 days old were again obtained for each test to observe if a correlation existed between the TLE's age and SPARSE's ability to track and reacquire each object.

Rather than utilise the near-GEO objects that were used during Phase 2 numerical simulations, GPS satellites were chosen as preferred test RSOs due to the operational requirements of the experimental system, because they are routinely manoeuvred for station-keeping and their element sets are updated often. Consequently there are always relatively current TLEs available and there is a rich history of past TLEs to choose from for aged initialisation data. Whilst targeting the MSG-2 debris objects was desirable, their inclusion within testing was dependent on serendipitous visibility at the site and the availability of current TLE targeting data. Unfortunately, these requirements were not met during the test period and consequently only GPS satellites were targeted by SPARSE during field trials.

Initialisation of the particle distribution was implemented as previously described for Phase 2 numerical simulations. The levels of process noise, to account for deficiencies in the SGP4 propagator, were also maintained for field trials, with one exception. To observe how the experimental system behaves when the process error is inflated, a subset of objects were propagated with inflated process error (IPE) defined by

$$\mathbf{v}^i_{\text{IPE}} \sim \mathcal{N}(0, \boldsymbol{\gamma}^2_{\text{IPE}});\ \boldsymbol{\gamma}_{\text{IPE}} = \text{diag}\begin{bmatrix} 30 \times 10^{-3^\circ} \\ 0.2 \times 10^{-3^\circ} \\ 0.0005 \times 10^{-3} \\ 50 \times 10^{-3^\circ} \\ 50 \times 10^{-3^\circ} \\ 0.001 \times 10^{-3} \text{ orbits/day} \end{bmatrix}. \tag{6.31}$$

The intention behind using the IPE values was to learn about the benefits or complications that may arise when utilising dynamic steering, when the particle distribution has been artificially inflated in anticipation of unmodelled dynamics.

6.5.1 Results

After some initial system verification tests run in the beginning of October 2013, a 5 day field trial involving six GPS satellites—as listed in Table 6.12—was held on 22nd–27th October 2013. During the 5 day trial, the maximum weight and maximum particles steering methods were tested simultaneously. SPARSE was instructed to update the track of each of, what is perceived to be, 60 objects via a set of five contiguous observations per evening. Every second GPS satellite's set of five historical-TLE distributions were propagated using the IPE values. The results of the first field trial are shown in Tables 6.13 and 6.14. In spite of some light cloud cover and the occasional inability to schedule observations due to time constraints, Table 6.13 and Table 6.14 show many examples of tracking and reacquisition of RSOs using the full range of historical TLEs. A clear correlation exists between the age of the TLE and the time SPARSE takes to reacquire each object.

Upon comparing Table 6.13 to Table 6.14, the maximum particles method appears to be, in general, slightly less effective than the maximum weight method. Nonetheless, when the results for object 32711 and 32260 are compared between each table, there is some evidence, while inconclusive, that the maximum particles objective function achieved superior performance for these IPE objects. Otherwise, in accordance with previous numerical tests, utilising the maximum particles steering method on non-IPE objects has appeared to generate inferior results. This evidence provides further weight to the proposed strategy of utilising the steering method that best suits the level of process noise applied during orbit propagation.

Table 6.12: Test RSOs used for Field Trial 1.

US Cat. ID	Description
26360	NAVSTAR 47 (USA 150)
32711	NAVSTAR 62 (USA 201)
25030	NAVSTAR 44 (USA 135)
35752	NAVSTAR 64 (USA 206)
22014	NAVSTAR 26 (USA 83)
32260	NAVSTAR 60 (USA 196)

Table 6.13: Observation states recorded during Field Test 1 - maximum weight result.

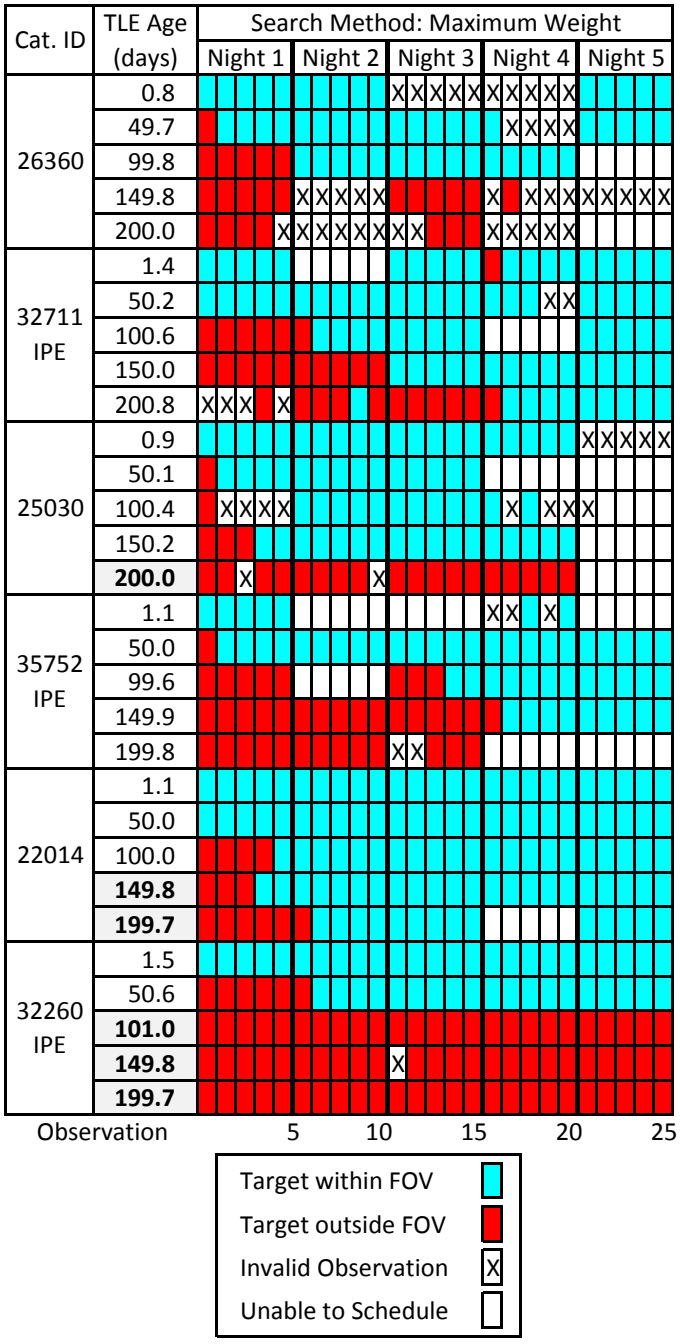

Whilst many examples of positive outcomes are displayed in each table, both tables also indicate that the experimental system appeared to have significant trouble reacquiring object 32260 and 25030 when their respective distributions were based on their oldest TLEs. In an attempt to explain why this occurred to these specific distributions and not others, data regarding manoeuvres of GPS satellites, that occurred prior to the trial, were compiled from

Notice Advisory to NAVSTAR Users (NANUs). NANUs are published by the U.S. Coast Guard Navigation Center [109] for reporting changes to the GPS satellite constellation to the public. Fig. 6.9 summaries the relevant data on a timeline. The timeline shows that three of the six GPS satellites were manoeuvred within 200 days of the field trial and may have consequently influenced the result.

Table 6.14: Observation states recorded during Field Test 1 - maximum particles result.

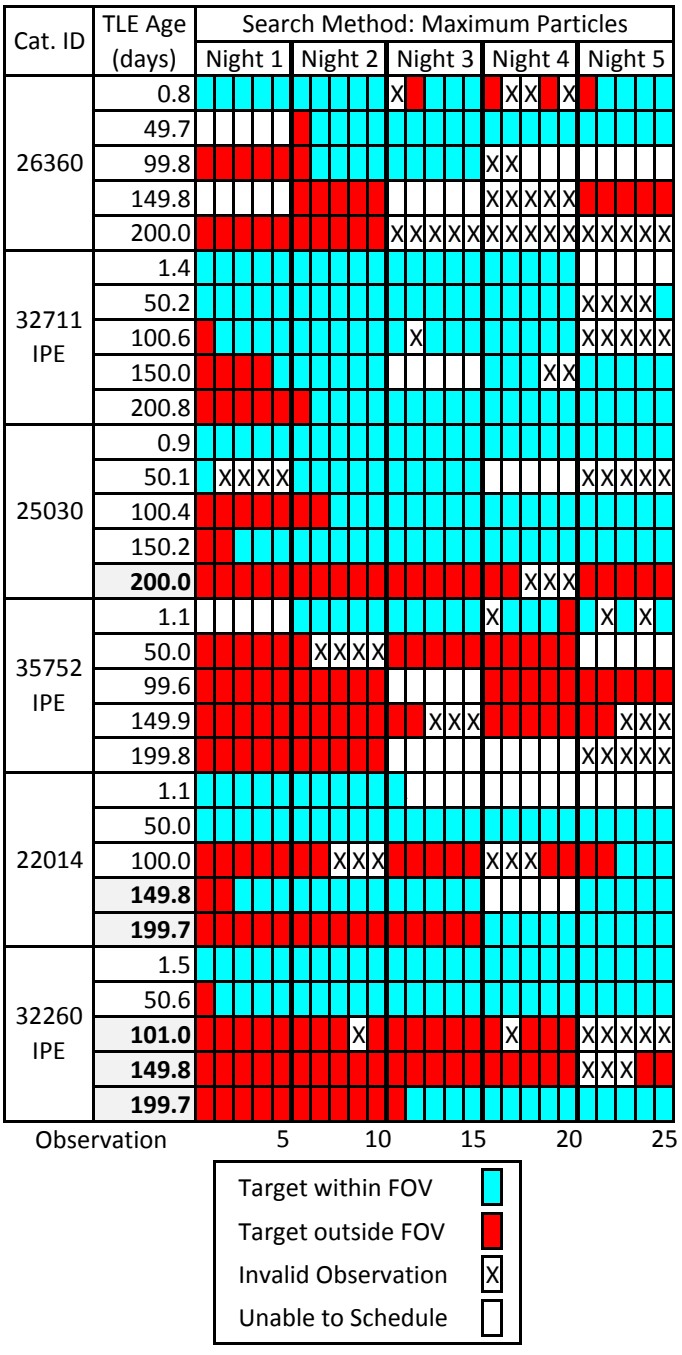

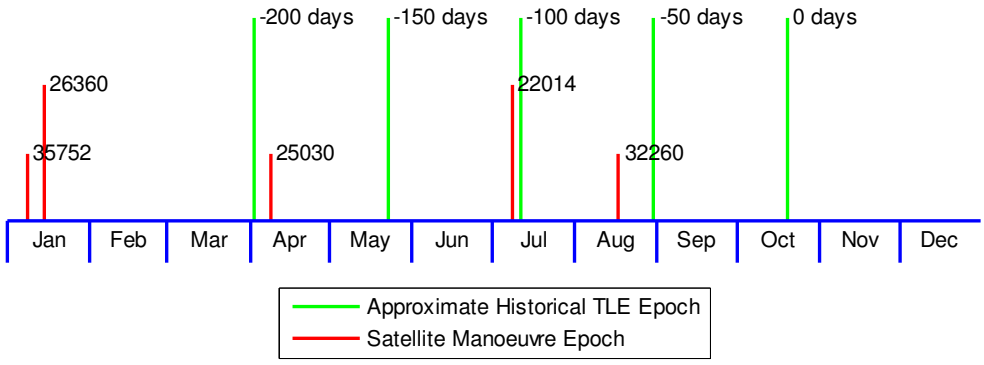

Figure 6.9: Manoeuvre timeline inferred from NANUs published during 2013 and plotted relative to Test 1's approximate historical TLE dates.

Using the dates indicated in Fig. 6.9, the affected rows in Table 6.13 and Table 6.14 can be deduced. This information has been represented in each of the "TLE Age" columns. Specifically, if the historical TLE was obtained prior to the last manoeuvre of the satellite, the TLE's age has been made 'bold' and its background grey. With this additional insight, a strong correlation can be observed between the pre-manoeuvre TLEs and difficulty by SPARSE to reacquire RSOs such as objects 25030 and 32260. Nonetheless, the correlation is not as strongly shared by other objects. In spite of utilising two pre-manoeuvre TLEs, object 22014 was reliably reacquired—especially using maximum weight steering—using all TLEs. To aid in explaining why 22014 appeared to be unaffected by the manoeuvres, the relative size of the manoeuvres were inferred by obtaining TLEs before and after of the recorded manoeuvre dates. These TLEs were then propagated over an identical 24 hour period, centred on the date of manoeuvre, and the resulting maximum position and velocity error was compared. To gauge the validity of this procedure, additional TLEs for GPS satellites that were and were not manoeuvred in this time were compiled and compared to each other. The results of this analysis are displayed in Fig. 6.10, noting that the non-manoeuvred control examples have been clearly labelled.

Fig. 6.10 indicates that even the controls—TLEs without evidence of manoeuvring in NANUs—show some change in position and velocity. In comparison, the manoeuvred GPS satellites show significantly more change between TLEs, with one exception. Whilst object 22014 was listed as a manoeuvred satellite within 2013's NANU archive, a detailed account of the manoeuvre is not provided. According to the relative values presented in Fig. 6.10, object 22014 shows even less evidence of a manoeuvre than one of the controls. For this reason it is assumed either the manoeuvre was very minor or a record of its cancellation could not be found. This evidence therefore adds significant weight to the hypothesis that manoeuvring was the cause of the most significant lapses in reacquisition observed during

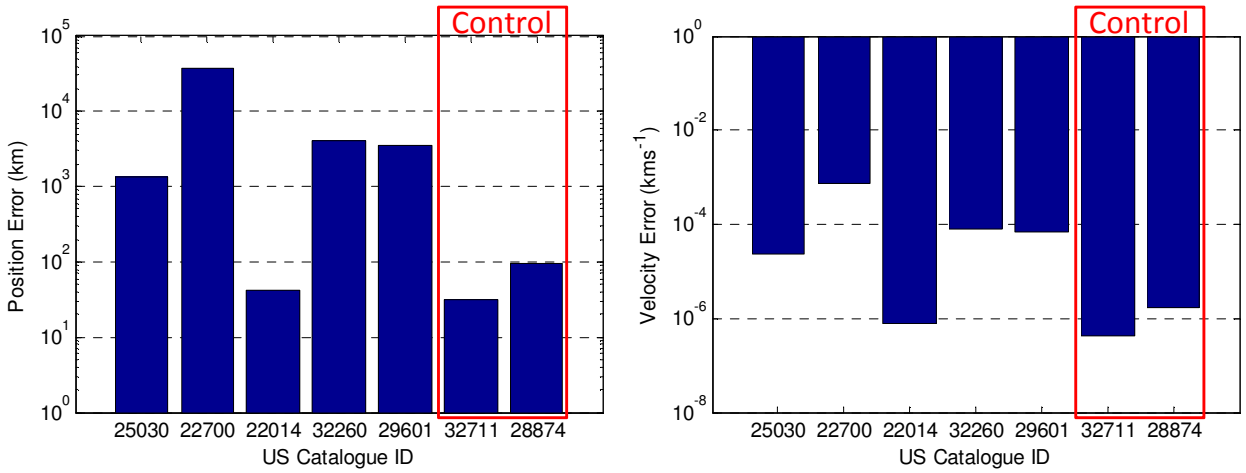

Figure 6.10: Relative TLE error due to manoeuvring of objects 25030, 22700, 22014, 32260 and 29601, in addition to two non-manoeuvred control examples for objects 32711 and 28874.

Field Trial 1.

Whilst objects 26360 and 35752 were sometimes challenging to reacquire throughout the test, many instances of invalid and unscheduled observations make it plausible that they simply didn't get enough observations. In addition to this, 35732 was an IPE object, which would likely have reduced the speed at which it could be reacquired. Nevertheless, to provide additional insight and validate these hypotheses, a second field trial was conducted.

To provide greater consistency over a wider set of targets, a second field trial involving the 15 RSOs listed in Table 6.15 was run on 27$^{\text{th}}$ October – 1$^{\text{st}}$ November. During this test, only the maximum weight steering method was utilised to perform dynamic steering. Otherwise, all other elements remained the same for this second trial. The results of the second trial are presented in Table 6.16. It should be noted, that a hardware failure occurred during the evening of 29$^{\text{th}}$ October resulting in many invalid results. A secondary hardware fault halted operations altogether on 30$^{\text{th}}$ October. Nonetheless, the system was operational again to conclude on 31$^{\text{st}}$ October.

Table 6.16 once again shows many examples of RSOs being tracked, automatic reacquisition and a correlation between the speed of reacquisition and TLE age. Manoeuvring information was compiled once more and the adjusted timeline is presented in Fig. 6.11. The relative changes in TLEs due to manoeuvring are also shown in Fig. 6.10 and Table 6.16's affected rows have been similarly identified. The summary of this evidence shows a very strong correlation between poor observations and manoeuvring of the spacecraft, with the justifiable exception of object 22014. Objects 26360 and 35752 were again challenging to

Table 6.15: Test RSOs used for Field Trial 2.

US Cat. ID	Description
37753	NAVSTAR 66 (USA 232)
26360	NAVSTAR 47 (USA 150)
28361	NAVSTAR 55 (USA 178)
32711	NAVSTAR 62 (USA 201)
24876	NAVSTAR 43 (USA 132)
25030	NAVSTAR 44 (USA 135)
22700	NAVSTAR 33 (USA 92)
23953	NAVSTAR 38 (USA 126)
26407	NAVSTAR 48 (USA 151)
35752	NAVSTAR 64 (USA 206)
22014	NAVSTAR 26 (USA 83)
28874	NAVSTAR 57 (USA 183)
32260	NAVSTAR 60 (USA 196)
38833	NAVSTAR 67 (USA 239)
29601	NAVSTAR 59 (USA 192)

acquire using their oldest TLEs suggesting that they may have undergone significant perturbation since the TLEs were produced or, possibly, the TLEs were somewhat inaccurate at the time. In almost all other instances, even with some interference from cloud and interruptions to scheduling, in a relatively short period of time the system very capably reacquired each object using targeting data that would not currently be considered serviceable.

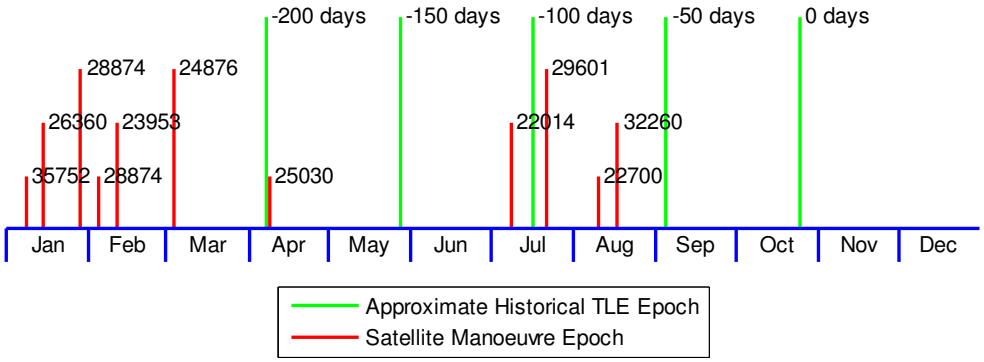

Figure 6.11: Manoeuvre timeline inferred from NANUs published during 2013 and plotted relative to Test 2's approximate historical TLE dates.

Table 6.16: Observation states recorded during Field Test 2.

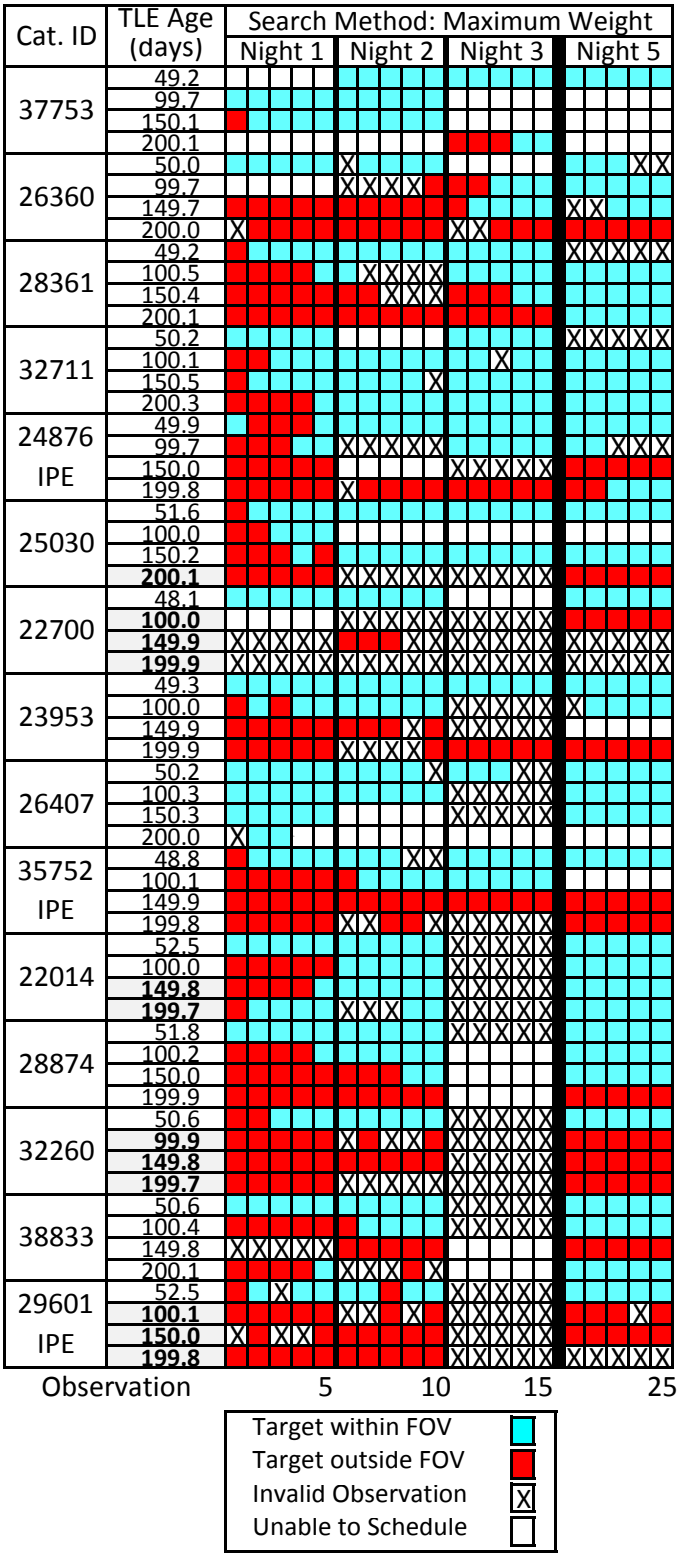

6.5.2 Discussion

Obtaining data whilst trialling the maximum weight and maximum particles steering methods was prioritised during field trials as they were found to be of greater value during simulations. Other field-trial objectives, such as reacquiring objects in contrasting orbital regimes, during manoeuvres and whilst using each of the proposed objective functions were deemed secondary and were conditional on serendipitous events and operational limitations. These constraints prevented their inclusion in this thesis however little to no changes were necessary to perform these tests, had the conditions been favourable.

The act of tracking a target using targeting data from well before a manoeuvre is a challenging task and was not an objective of this investigation. While the results indicate that using TLEs that were created prior to a GPS satellite manoeuvre often resulted in an inability to reacquire the satellite, the methodology described so far is likely capable of accounting for this unmodelled motion if it can be bounded. It is anticipated that the system would be more effective at tracking an object during a manoeuvre. Albeit preliminary, this is supported by the IPE results, as the system continued to achieve reacquisition in spite of the excessively large process error.

Ultimately, in all cases 50-day old TLEs were reacquired. In most other cases, 100, 150 and even 200 day old TLEs were used successfully for reacquisition. In the remaining cases, almost all failed attempts at reacquisition correlated with manoeuvring of the space vehicles or were hampered by inclement weather, time constraints and hardware faults. The results therefore show that dynamic steering, using the maximum weight steering method was very successful at using targeting data that existing space surveillance networks would discard due to its old age and poor accuracy. Furthermore, the use of the maximum particle steering method has shown potential as an effective alternative when large process errors are required.

6.6 Conclusions & Future Work

Dynamic steering has demonstrated an ability to autonomously reacquire RSOs using a real-time high-fidelity p.d.f. prediction and observation evaluation method when implemented at the sensor. By improving the intelligence of the sensor, it has been shown that p.d.f. processing can be decentralised to good effect, the act of maintaining SSA can be further automated and the capabilities of the system can be increased. A primary enhancement which dynamic steering can provide to an existing SNSS, is the potential to increase system capacity by reducing the frequency of observation of certain RSOs. Whilst this approach would provide extra time for the orbital state error to grow, the higher fidelity representation

of the p.d.f. reduces the uncertainty, enhancing the reliability of conjunction analysis.

Improvements to Monte Carlo error minimisation during the prediction of the prior p.d.f. and the enhancement of visibility prediction by incorporating the particle-approximated state error distribution, were identified as a means of improving the reliability of dynamic steering strategies. The reliability and objectivity of the process is also likely to increase with the development of purpose-built hardware and software, to minimise the reaction time of the system and maximise the throughput of information. The utility of dynamic steering may also be expanded via the investigation of alternative steering regimes, for use under a broader range of surveillance scenarios.

The ability to capture and exploit information pertaining to false alarms and probability of detection were limited by the experimental apparatus. These elements ought to receive greater attention in future implementations. Initial attempts could incorporate multi-object detections by utilising over-lapping densities to compute relative probabilities of association and anticipate the detection of new objects that are yet to be catalogued. More advanced implementations might include auxiliary information, such as object characterisation, to enhance target discrimination. The mathematical formulation would also benefit from the incorporation of a non-uniform false alarm distribution and an arbitrary probability of detection profile.

The results of the field trial have demonstrated that dynamic steering enables sensors to automatically reacquire targets whose orbital state data is too poorly conditioned to be reliably targeted using existing SSA tracking techniques. Dynamic steering directly addresses an identified [22] need by the SSA community to devise a means of automating the manually intensive process of reacquiring lost RSOs, in a timely manner. Dynamic steering also has the potential to improve, and in some cases enable for the first time, tracking of objects whose orbital dynamics are difficult to model, expeditious acquisition of targets that have been detected but their orbits are yet to be refined as well as detection and tracking of unanticipated manoeuvres. This investigation's progress regarding real-time decision making at the sensor also opens the way in future to develop a network-wide dynamic scheduling system, capable of interacting with its sensors in real time to dynamically adapt a global schedule as observations are made.

7 Conclusion

The objective of this thesis has been to explore the deployment of steerable sensors as an efficient means of improving the capability, capacity and timeliness of existing space surveillance systems, to provide superior levels of SSA through enhanced sensor management. The research methodology has focused on utilisation of statistical signal processing techniques to improve sensor management at both a network and an individual sensor level. The research began at the network level by studying the benefits to catalogue accuracy when steerable sensors are employed as the primary sensing resource. It was found that the utilisation of a high ratio of steerable sensors improved the network's capability to maintain an RSO catalogue, thereby improving the level of SSA attained. This research identified two key areas pertaining to range observability and sensor availability that if improved through effective network management, would see an additional improvement to catalogue accuracy when using optical steerable sensors. These areas were investigated next resulting in the development of two scheduling techniques designed to address each topic. Thereafter the focus turned to research at the sensor level to improve the capacity and timeliness of the surveillance system. A technique for achieving real-time decision making at the sensor, named dynamic steering, was found to achieve these aims.

In parallel, this research has necessitated the development of a practical means to efficiently perform research on a topic that is typically costly to investigate and challenging to

validate. The result was an adaptive and scalable SSA simulation capability that evolved to incorporate GPGPU computation. The involvement of GPU hardware provided the simulation system the computational power necessary to achieve full scale simulations and real time high-fidelity state p.d.f. analysis and update without compromising the flexibility of the system. This new capability permitted the development of a sensor controller with intelligence surpassing that of existing space sensors. By providing a steerable sensor with the means to simulate aspects of the surveillance environment, the sensor gains a higher degree of awareness and the ability to perform high fidelity prediction. These attributes permit the development and utilisation of capabilities beyond the scope of existing space surveillance systems.

The Raven class electro-optical sensor—state-of-the-art in steerable sensors at the time of writing—features frequently in discussion. The intention is nonetheless to apply the presented research to any class of sensor or surveillance application with which these techniques are compatible.

7.1 Research Summary

7.1.1 Network Level Sensor Management for Improved System Capability

In Chapter 2 the current SSA environment was explored by analysing the distribution of man-made objects throughout Earth orbit using publicly available tracking information. While it was acknowledged that this information is by no means a complete picture of the state of all man-made items currently in space, it provides insight as to where RSOs are likely to reside in Earth orbit and why they are likely to be sent to these regions. By demonstrating how impractical it would be to apply a traditional surveillance and tracking regime for maintaining uninterrupted coverage of all Earth orbit, the necessity of obtaining and maintaining an RSO catalogue for achieving SSA was emphasised. This demonstration highlighted the value of maintaining SSA in between serendipitous opportunities for observation of RSOs. A follow-on investigation culminated in an equation for predicting how often these serendipitous opportunities might occur for common RSO orbital trajectories. The desire to appraise the validity of this equation using simulated data, foreshadows the reliance on SSA simulation software for investigation and validation throughout the remainder of the thesis. Lastly, a review of the prominent tracking techniques of relevance to SSA was conducted. The review details the attributes of each techniques' ability to combine raw measurements made by space surveillance sensors and produce a catalogue capable of predicting

the dynamic motion of all RSOs for conjunction analysis. Through the presentation of these introductory analyses, it is hoped that the reader gained sufficient intuition for maintaining SSA to appreciate the motives and justifications for performing the research presented in the following chapters.

Chapter 3 detailed the development of the simulation environment named MASSAS for comparing the performance of alternative SNSS configurations. MASSAS was designed to achieve this goal whilst utilising minimal resources to enhance the flexibility of the system and speed of execution. Consequently, MASSAS was designed with a focus on system characterisation rather than precisely predicting the magnitude of catalogue error. To demonstrate that MASSAS can effectively characterise the ability of an SNSS to minimise catalogue error, MASSAS's results were compared to a third party high fidelity simulation developed for AFRL named TASMAN. Whilst some understandable differences were observed, the results of each simulation led to the same conclusions regarding each configuration's ability to minimise catalogue error. These results indicate that while MASSAS's reduced fidelity may have implications for predicting the absolute error expected by a particular SNSS configuration, it can capably characterise and compare their ability to minimise catalogue error. Additional tests, run on MASSAS alone, produced the results of the first analysis to compare the relative error characteristics of an RSO catalogue when a space surveillance network's radars are gradually replaced by steerable high-accuracy optical sensors. The results indicate that a network comprising of a high ratio of optical sensors could produce catalogues with lower levels of error than a network dependent on radars, with typical measurement characteristics. The study also demonstrates that, when using optical sensors, superior results would be achieved if the system is capable of minimising constraints to radial measurements and reduced sensor availability. These topics were investigated in the following chapters.

Enhancing the Effectiveness of Optical Sensor Networks

To provide a means to overcome an optical sensor network's inability to directly observe RSO range, a collaborative space sensor scheduling technique was developed in Chapter 4. The technique aims to provide a centralised sensor management system the ability to utilise the simultaneous measurements of multiple sensors, to achieve complementary measurement characteristics that could not otherwise be achieved using sensors independently. The technique predicts error covariance reduction when utilising all combinations of sensors. This information is then used to identify and exploit opportunities to leverage measurement level sensor fusion to obtain measurements of inferred quantities, such as range. Two critical areas were identified in which balance must be maintained to achieve a favourable result.

The first is to ensure the distribution of the network's sensors are appropriate to permit collaborative measurements at desirable altitudes, whilst remaining mindful of maintaining global coverage. A method to appraise a network's ability to perform collaborative measurements at various altitudes is proposed that relates the relative spacing between sensors to the geometry necessary to achieve multi-sensor measurements of a single RSO. The second area identified to be critical to the success of the technique is to ensure that multi-sensor measurements are only employed for observing a single object if it is beneficial for the catalogue as a whole. Namely, the method is of limited value if a small subset of RSOs attain very accurate orbital estimates while the remainder of the catalogue is neglected. The proposed solution is to inversely weight the priority of collaborative measurements according to the number of sensors involved in the measurement. Via this weighting, the technique demonstrated improvement in catalogue error. Whilst consistent improvement was not achieved, the technique is likely to benefit from the development and utilisation of an observation effectiveness metric that aims to minimise catalogue error, rather than maximise a single RSO's position covariance. Overall, it was found that the process is likely better suited to RSOs in higher orbit such as MEO and GEO. This result is attributed to the fact that range information can be especially challenging for optical sensors to achieve when targeting RSOs at high altitudes.

Chapter 5 details an original investigation regarding the availability of optical sensors when supporting routine maintenance of a large catalogue of genuine RSOs. To perform this investigation, all tracking data belonging to a publicly available catalogue was obtained and utilised. Enhancements to MASSAS's computational performance were necessary to achieve a practical runtime whilst simulating an SNSS comprising of up to 20 sensors and approximately 15 000 objects. MASSAS's computational efficiency was significantly enhanced by converting a number of crucial modules into parallelised algorithms to be executed efficiently using a GPU. In doing so, the number of objects has a much weaker influence on the runtime of the system. Utilising these enhancements, an availability analysis was performed to compare the regularity at which each RSO can be observed when employing various proportions of radar and optical sensors in an SNSS. Whilst radars are traditionally used to observe RSOs in LEO to lower MEO altitudes, the analysis determined that, in sufficient numbers, steerable optical sensors could be effectively utilised to observe objects as low as mid-LEO altitudes. A particularly effective arrangement was identified when a very high ratio of optical to radar sensors were used. In this scenario, most RSOs could be regularly observed using high-accuracy optical measurements, while the use of expensive, lower accuracy radars was minimised. When employing such an arrangement, it was demonstrated that lapses in visibility of some RSOs should be anticipated. This was particularly true for RSO's

with challenging trajectories such as highly eccentric orbits, but may also be anticipated in practical networks due to sensor maintenance or commandeering of sensors for specialised tasks. Consequently, the chapter concluded with the development of a scheduling method that utilises visibility forecasting and observation effectiveness evaluation, to predict the consequence of failing to observe an RSO during an upcoming period of non-observability and compensate accordingly. By predicting the level of state estimate error when an RSO is visible after the current scheduling period, the method is capable of prioritising the observation of RSOs about to enter a period of non-observability. This ensures that their state estimate error will not be unusually large when they are next visible to the SNSS. To further leverage the utility of deployable steerable sensors, such as Raven class electro-optical sensors, the system could foreseeably be further developed to anticipate extended periods when it may be necessary to deploy mobile sensors to take measurements of RSOs at risk of being lost.

7.1.2 Enhanced Sensor Level Management for Improved System Capacity and Timeliness

The computational power and flexibility afforded by the newly parallelised MASSAS simulation environment, permitted a retasking of the system into an adaptive sensor controller. Chapter 6 details the development and implementation of a novel method named dynamic steering for controlling steerable space sensors. The chapter also presents the results of a field trial of an experimental implementation of the method. Dynamic steering utilises real-time decision-making at the sensor to permit it to autonomously steer itself for enhanced tracking and, when necessary, to search for a target. Dynamic steering utilises an adaptation of the conventional Bayesian filtering method to incorporate consideration of steering costs and the reliability of observations into the filtering strategy. Crucially, the method relies upon the update of a high fidelity state error p.d.f. with each observation. In doing so, the system can distinguish in which direction the sensor should be pointed next, to maximise the likelihood of a favourable observation. This is particularly important in situations when an RSO is scheduled for observation, but it is not observed in the sensor's FOV. In such a case, the system is capable of retaining knowledge of where the target was not found, such that it may subsequently converge on the target's location. This demonstrates a capability beyond the current operational protocol that necessitates a comparatively slow manual intervention by suitably qualified operators. Furthermore, the adoption of dynamic steering by many SNSS sensors, would permit catalogue maintenance to be performed with a reduced threshold on the probability of detecting an RSO with the next observation. The catalogue maintenance system would thus be capable of increased capacity as RSOs can be observed less often.

For instance, experimental field trials demonstrated the successful reacquisition of RSOs using TLEs that were more than 6 months old. Whilst reduced regularity of observations permits error to grow, the dynamic steering method requires the use of high fidelity state error p.d.f.s. Uncertainty in the p.d.f. is therefore reduced permitting higher confidence in collision predictions via conjunction analysis. In spite of the fact that the fidelity of p.d.f.s required for dynamic steering require much higher levels of computation and memory than conventional tracking, the use of GPGPU hardware proved to be a practical solution. Using consumer level componentry, the system was capable of performing the required levels of computation in real time. Additional development of purpose-built hardware and greater consideration for data association would further enhance the reliability of the system. The dynamic steering method also has the potential to permit a number of enhancements to existing catalogue maintenance methods. Such advances could include routine automatic reacquisition of objects whose dynamics are difficult to model, rapid refinement of newly discovered or perturbed RSOs in addition to improved manoeuvre detection and tracking.

7.2 Recommendations for Future Research

The application of the dynamic steering method to a number of space surveillance sensors, would result in a stream of high fidelity state estimate products being generated as observations are being made. If these products can be returned and processed by a centralised sensor management system, a coordinated system that responds in real time to the effectiveness of its sensors is achievable. The system would be capable of adapting to unforeseen circumstances such as sensor outages, unexpected changes to RSO orbits and ad-hoc tasking during events that require prioritised attention from the surveillance network. Ad-hoc events need not only be initiated by human operators, but should also include automatic cueing from other sensors which have detected anomalous activity. For instance, a radar or wide-field optical sensor may detect the fragmentation of an RSO and automatically call on multiple dynamically steered sensors to expediently refine the orbits of debris before the debris has time to disperse. Such a response would enhance the system's ability to compute the likelihood of secondary collisions and provide a timely warning.

Furthermore, it is anticipated that the insights gained throughout Chapters 2–5 may assist with the planning necessary for expansion and adaptation of a space surveillance network, to scale with the increasing numbers and evolving distributions of RSOs. Via collaborative scheduling, prioritised targets may be identified and observed by multiple sensing resources for timely reduction of orbital uncertainty, whilst maintaining sufficient coverage

of lower priority targets. Through visibility forecasting, the system could identify RSOs requiring observation before they become invisible to the network and update its orbital track in preparation. Both techniques can also be augmented to exploit the high fidelity p.d.f.s generated by dynamic steering. Using this additional information, the system will have a greater understanding of the effectiveness of multi-sensor measurements and can anticipate how difficult an object would be to reacquire when it is next visible to the surveillance network and schedule accordingly. Depending on the immediate results of these observations, additional sensing resources may be allocated for improved evaluation and, if necessary, advise the deployment of mobile sensors. The result would be a space sensor management system with unprecedented adaptability, capable of tolerating and exploiting ad hoc addition and deletion of heterogeneous sensors and sensor sub-networks.

Ultimately, it is hoped that the results of the investigations, the proposed techniques and implementation strategies detailed in this thesis contribute to a foundational framework on which an adaptive and robust space sensor management system may be devised. The proposed techniques will ideally be used to produce a system for SSA that is capable of keeping pace with the expanding use of space technologies.

A

Orbital Elements

As detailed in Chapter 2, it is often useful to examine the Keplerian 'two-body' orbit of an RSO to gain a first order appreciation for the RSO's current or instantaneous orbital characteristics. This appendix provides an overview of a prevalent method for describing the Keplerian orbit's elliptical trajectory, sufficient for understanding its usage throughout the thesis. Further detail, particularly regarding fundamental orbital dynamics and its relation to conic sections, may be found in Chobotov [8] and Vallado [9].

Table A.1: Keplerian elements.

Symbol	Description	Dimension
a	semi-major axis length	length
e	eccentricity	dimensionless or length/length
i	inclination	angle
Ω	right ascension of the ascending node (RAAN)	angle
ω	argument of perigee	angle
ν	true anomaly	angle

The six Keplerian elements listed in Table A.1 may be used to describe the elliptical osculating orbit of an RSO.

The scale and precise shape of the ellipse are described by the elements a and e. Whilst a is the length of the semi-major axis, e is a dimensionless quantity, that may be used to

relate a to the length of the ellipse's semi-minor axis b via the equation

$$e = \sqrt{1 - \frac{b^2}{a^2}}. \tag{A.1}$$

Elements a and e also fully define the length of perigee (2.17) and apogee (2.18); the closest and farthest distance between the centre of the Earth and an RSO, throughout its orbit.

The application of Kepler's first law—as detailed in Section 2.1—to an RSO's orbit, implies that the RSO's orbit is an ellipse and the Earth's centre is located at one of the ellipse's foci*. The Earth's North celestial pole and a celestial reference vector, that lies along the Earth's equatorial plane, are used to locate and orient the ellipse in inertial space. The celestial reference vector, by definition, remains stationary in inertial space as the Earth rotates. Most commonly, the celestial reference vector points from the centre of the Earth toward the Sun at the Northern Hemisphere's vernal equinox [9]. Namely, the celestial reference vector will point toward the Sun, as the Sun's declination—an angle North or South of the Earth's equator—transitions from a negative Southern value to a positive Northern value, due to the Earth's progress around the sun and its axial tilt with respect to its orbital plane. The precise moment of each equinox varies from year to year. At the time of writing this these, it is occurring on about 20$^{\text{th}}$ March each year [110]. Due to this variability, an internationally recognised inertial coordinate frame will use a specific year's equinox as a common reference [9].

As depicted in Fig. A.1, the orientation of the ellipse with respect to the aforementioned reference vectors and the equatorial plane, is described using the elements i, Ω and ω. The inclination i defines the angle between the equator and the RSO's orbital plane. Inclination is restricted to the domain $0 \leq i \leq \pi$, where $0 \leq i < \frac{\pi}{2}$ indicates a prograde orbit, with the Earth's rotation; $\frac{\pi}{2} < i \leq \pi$ indicates a retrograde orbit, against the Earth's rotation and $i = \frac{\pi}{2}$ indicates a polar orbit. The right ascension of the ascending node Ω specifies the angular distance between the celestial reference vector and the ascending node; the point at which the RSO's orbit passes from the Southern hemisphere into the Northern hemisphere. A positive value for Ω describes an anti-clockwise rotation about the Earth's polar axis when viewed from the North, looking South. The argument of perigee ω is the angular distance between the ascending node and the point of perigee. A positive angle is measured in the direction of motion of the RSO.

The location of the RSO along the elliptical trajectory is denoted by the true anomaly

*In actuality, one of the two foci would coincide with centre of mass of the two masses, not the Earth's centre of mass. But due to the large difference in mass between Earth and RSO, the centre of mass of the hypothetical system would be located extremely close to the Earth's centre of mass.

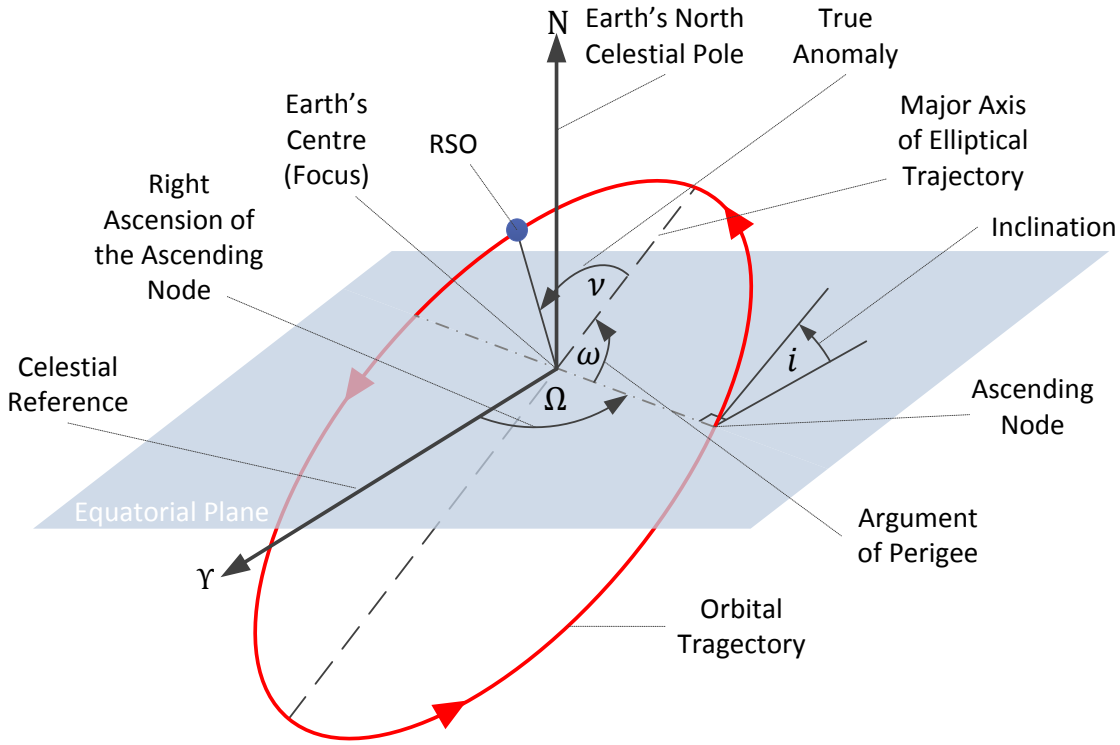

Figure A.1: An illustration of the Keplerian elements necessary to describe the orientation of an RSO's osculating orbit.

ν. It is a measure of the angular distance from the point of perigee to the RSO's current location along the elliptical trajectory. Similar to ω, a positive angle is measured in the direction of motion of the RSO.

The Relation Between Keplerian and TLE Elements

The state vector defined in Section 6.3 utilises six 'TLE elements' as reproduced in Table A.2. These elements may be combined with a TLE's ballistic coefficient, mean motion derivatives and TLE epoch information to define an orbit according to the Simplified General Perturbations (SGP) model [39, 68].

The TLE elements are similar to the Keplerian elements except that they are specifically defined within the NORAD TEME reference frame [9]. Furthermore, the semi-major axis length a and true anomaly ν have been replaced by the mean motion n and mean anomaly M respectively. The mean motion n may be related to the semi major axis length via the equation

$$n = \frac{86400}{2\pi}\sqrt{\frac{\mu_\oplus}{a^3}}. \tag{A.2}$$

But it is noted in Vallado [9] that n is a Kozai-mean value, which implies that the resulting

Table A.2: TLE elements utilised in Chapter 6's state vector.

Symbol	Description	Dimension
i	inclination	angle
Ω	RAAN	angle
e	eccentricity	dimensionless or length/length
ω	argument of perigee	angle
M	mean anomaly	angle
n	mean motion	orbits/day

semi-major axis will be a mean value rather than an osculating value. The mean anomaly M is the angular distance that the RSO would have travelled in time t seconds since it last past the point of perigee, if the RSO's orbit was circular with the same orbital period. This relation can be readily described using the equation

$$M = \frac{2\pi}{86400} nt. \tag{A.3}$$

Conversion of M to ν while not straightforward, has a number of efficient solutions, which include using the algorithms `KepEqtnE` and `Anomaly2v`, as detailed in Vallado [9]. Nevertheless, the intended purpose of the TLE mean elements is to apply them to an appropriate SGP model to approximate an RSO's *perturbed* orbit. The SGP models thereby utilise the TLE elements to approximate the influence of many primary sources of perturbation [68] to, ideally, achieve a more accurate result than could otherwise be achieved using a two-body approximation.

B

Supplementary MASSAS Visualisation Examples

Displayed throughout Appendix B are a number of image captures from programs and visualisations that were created in conjunction with the research detailed in this thesis. Whilst they were not necessary for explanation in the main text, some items may aid in appreciation of a number of concepts that were discussed. The remainder are provided to demonstrate the quality of visual information that could be produced by MASSAS whilst conducting the reported research.

B.1 Real Time SSA Visualisation

Utilising the PC's GPU, MASSAS is currently capable of visualising the position of all RSOs, who's tracking data has been published publicly, in real time. Figs. B.1 & B.2 are example frames from a continuously updating visualisation. In each image, a spacecraft undergoing a Molniya orbit has been selected and its predicted orbital trajectory for the next 24 hours has been propagated and plotted. Whilst each image appears to display different trajectories, they are in fact the same orbit, but in different reference frames. Fig. B.1 displays the Molniya orbit in an inertial reference frame. This means that as time progresses, the visualisation shows the Earth rotating with respect to the orbital plane. Conversely, Fig. B.2 displays the same orbit in an Earth-fixed reference frame. The elliptical orbit is distorted by the rotation of the reference frame within inertial space. In this reference frame, the visualisation displays the Earth remaining stationary with respect to the spacecraft's non-elliptical trajectory.

It is in the Earth-fixed coordinate frame that the value of the Molniya orbit is revealed. Fig. B.2 demonstrates how a spacecraft undergoing a Molniya orbit, oscillates between two prominent locations at high altitude and latitude, above Russia and USA. At these locations, the spacecraft is nearing apogee in its highly elliptical orbit, causing it to loiter in a confined region of the Russian and American sky for a significant portion of its 12 hour orbit. In addition to an eccentricity of approximately 0.7 and an orbital inclination of 63.4°, its 12 hour orbital period accounts for the observed oscillatory behaviour in Earth fixed coordinates. In the time the spacecraft takes to complete a single orbit, the Earth will rotate 180° on its axis. This causes the spacecraft to consistently reach apogee at one of two longitudes, 180° apart, every orbit. Twice per day, the spacecraft will also pass 'underneath' the Earth at low altitude and at high velocity with respect to the surface of the Earth.

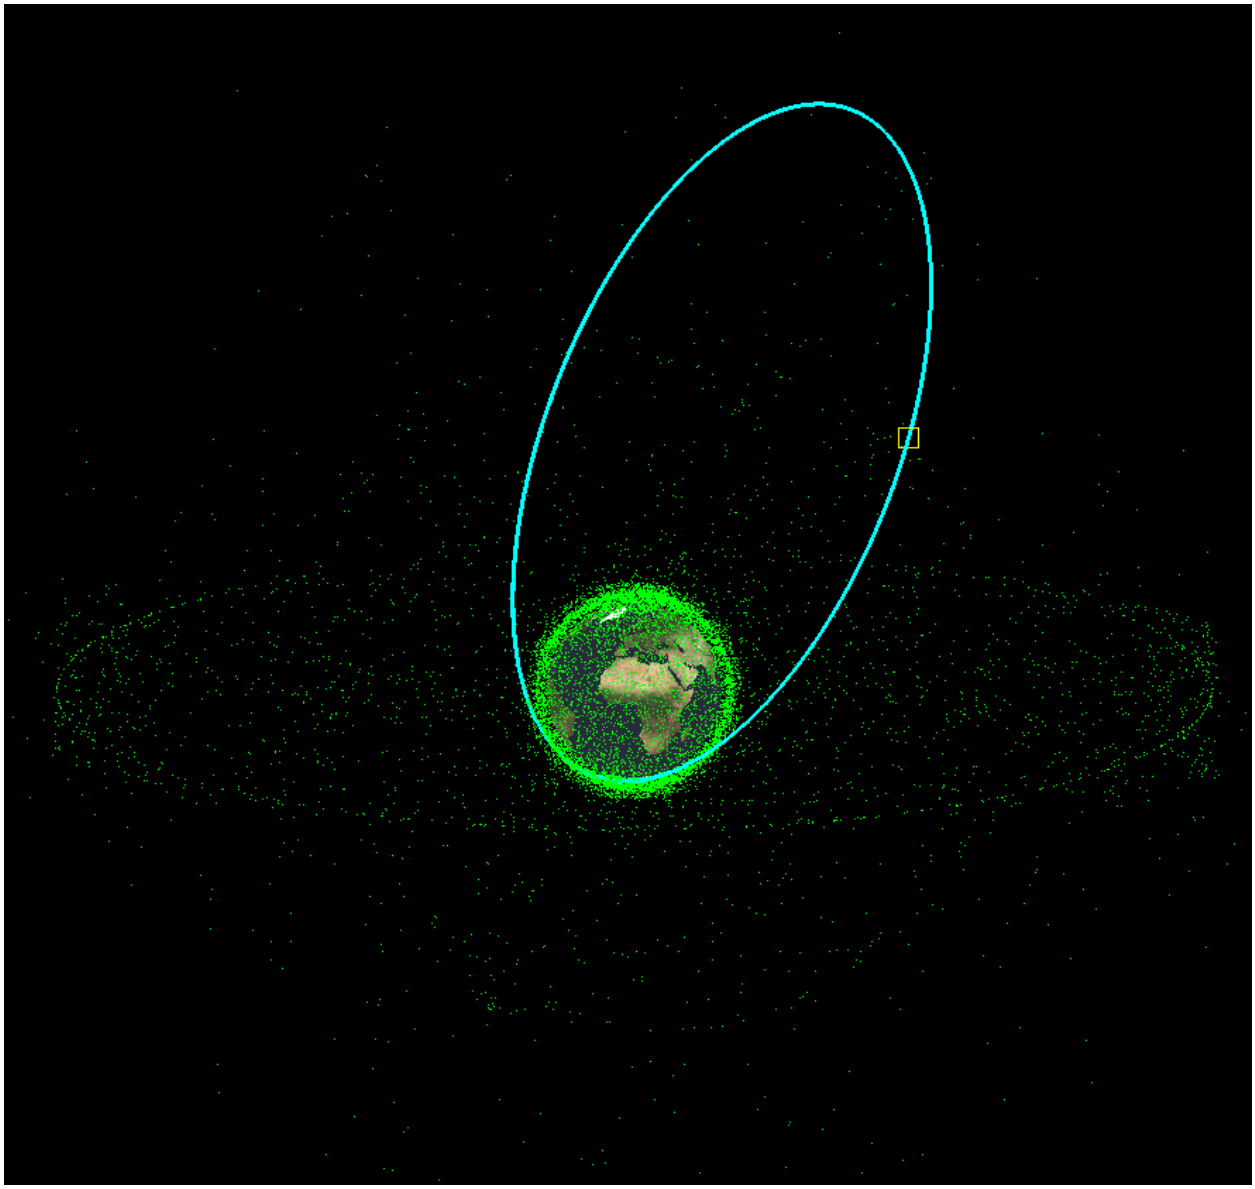

Figure B.1: A screen capture of a real-time visualisation of an RSO undergoing a Molniya orbit in an inertial reference frame.

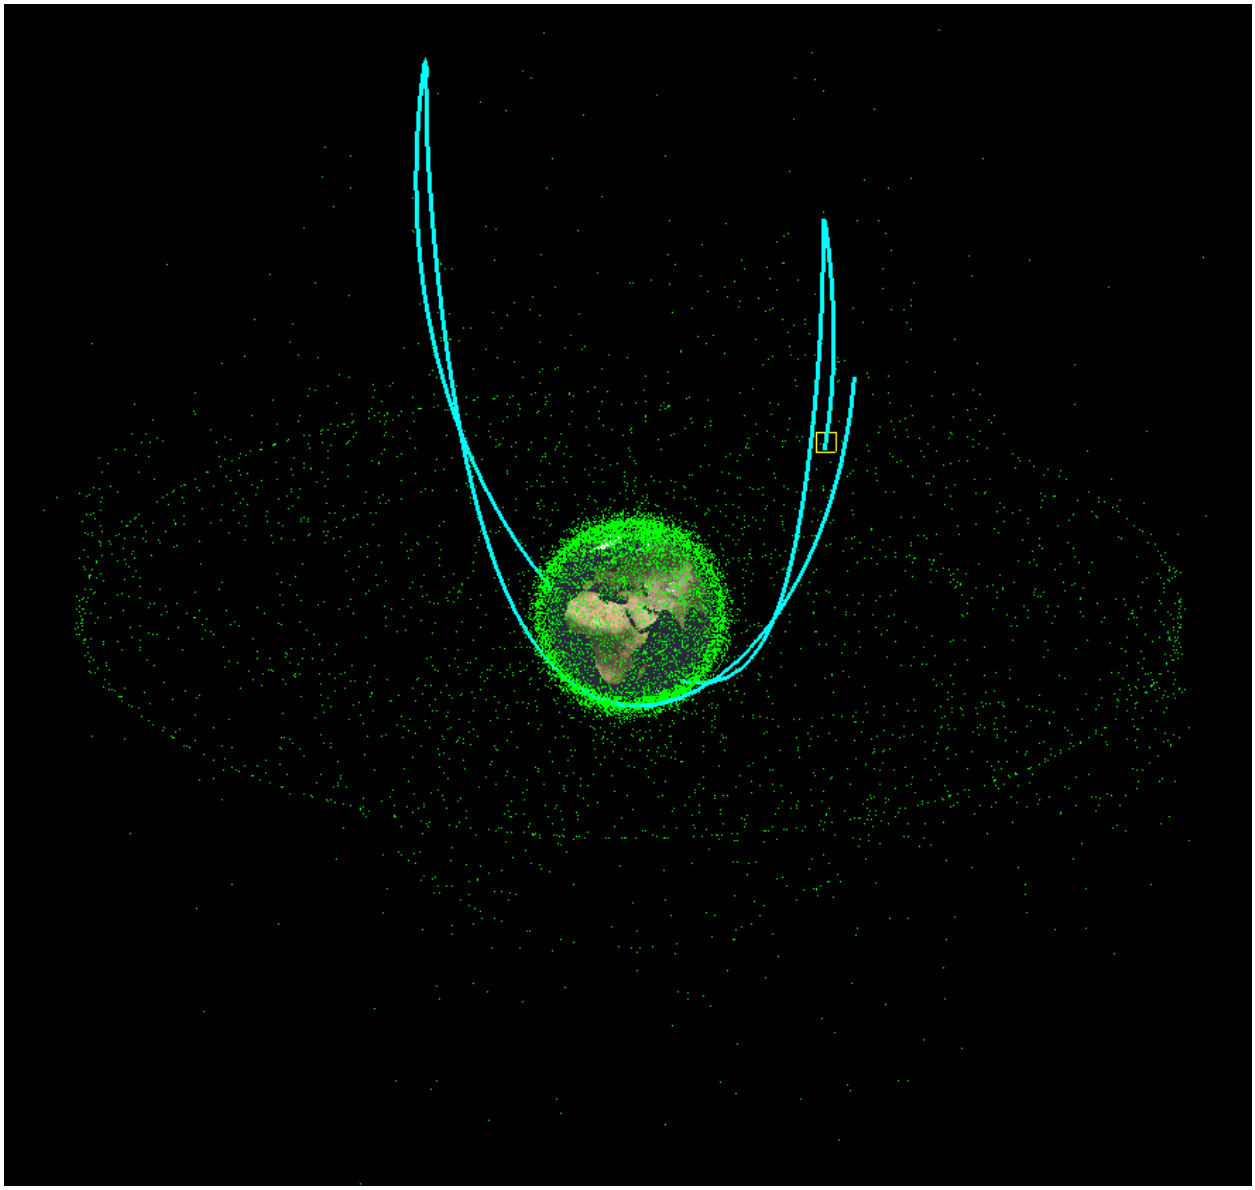

Figure B.2: A screen capture of a real-time visualisation of an RSO undergoing a Molniya orbit in an Earth-fixed reference frame.

B.2 Orbital Trajectory and Visibility Prediction Visualisations

Fig. B.3 displays a visualisation produced by MASSAS that was generated by a module that may be used to display the orbital trajectory of an RSO. As an example, the orbit of the ISS in an Earth-fixed coordinate system has been demonstrated. This particular visualisation also incorporates real time lighting of the globe using a solar and lunar model to position

light sources and apply appropriate lighting to the scene.

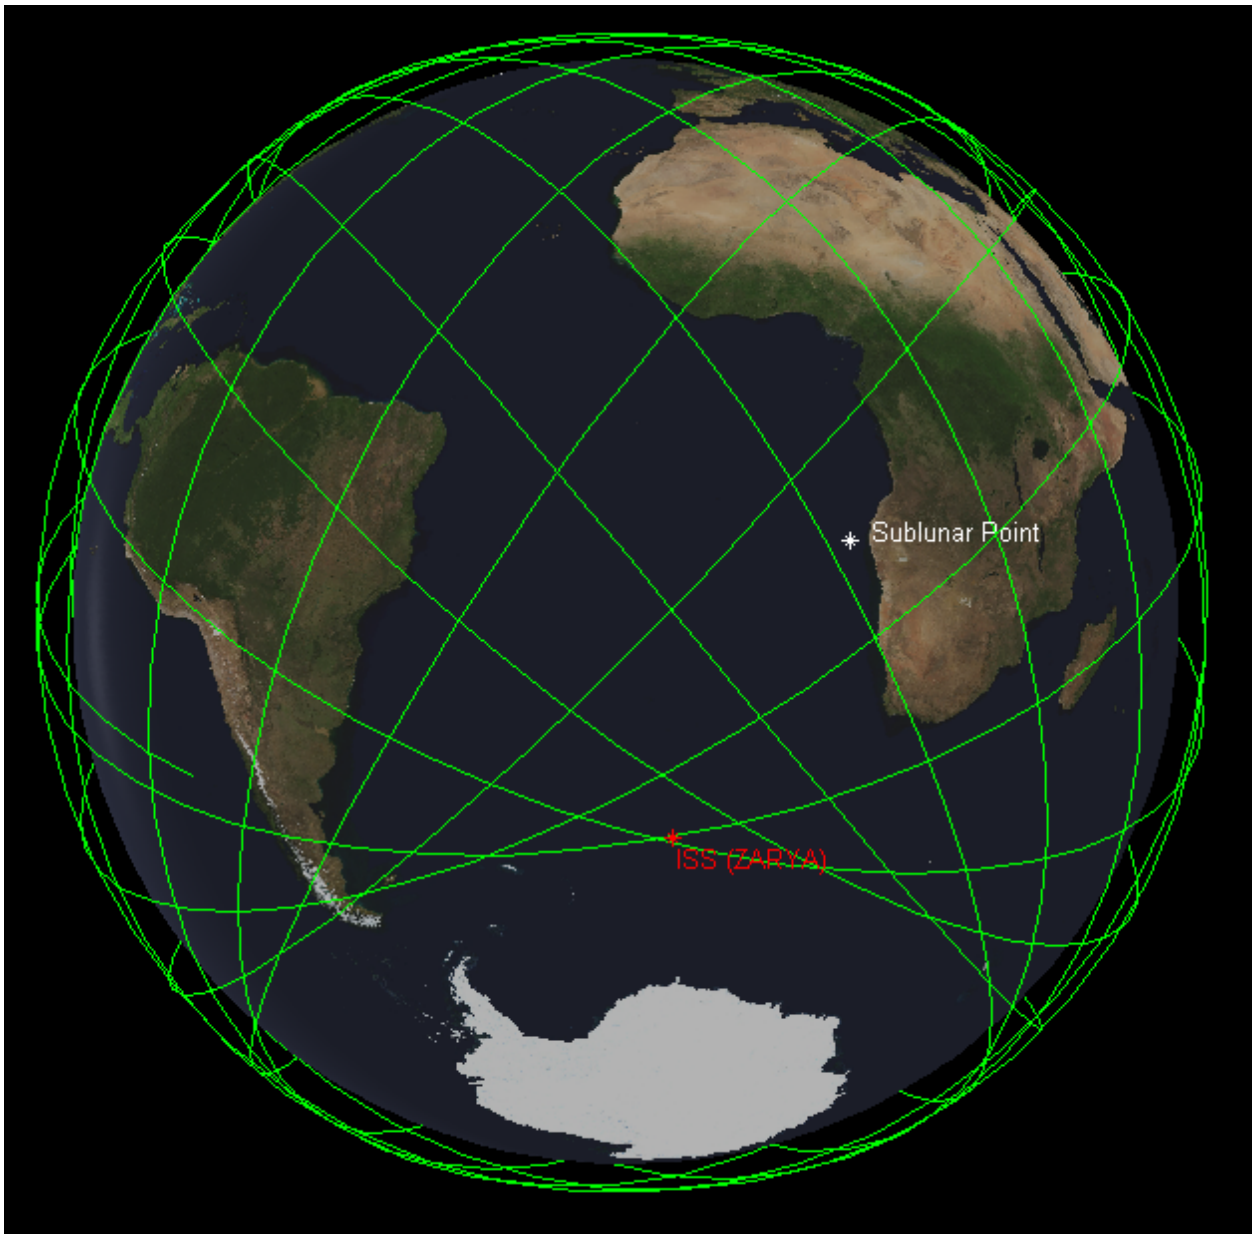

Figure B.3: ISS ECEF orbital trajectory.

MASSAS's visibility prediction module generated the visualisation displayed in Fig. B.4, when instructed to predict the 4-day visibility forecast for the ISS on 28th Jul 2014. Fig. B.4 displays the trajectory of the ISS over the 4 day period and highlights instances of optical visibility. The visibility prediction module also generates a report detailing information about any instances of visibility. The report produced by this example is displayed below:

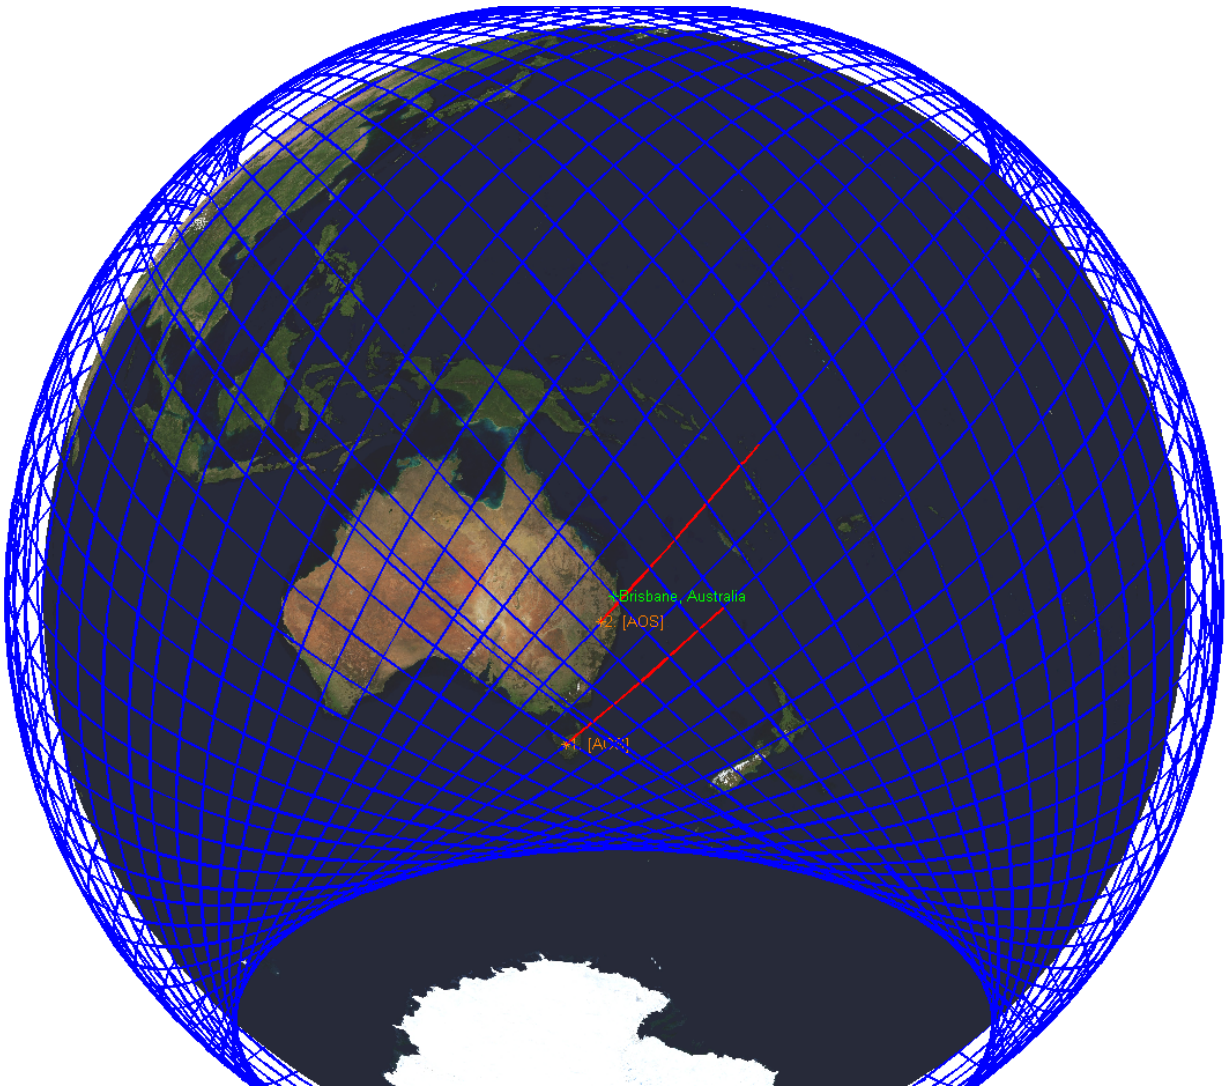

Figure B.4: An example visualisation of a 4-day visibility forecast for the ISS, as produced by MASSAS's visibility prediction module.

```
Site:  Brisbane, Australia
Object:  ISS (ZARYA) [USCATID: 25544]
Observation Period:  4.0 days
10:00:00.00 28/7/2014 to
10:00:00.00 1/8/2014 Local Time.

Observation 1.
Duration:  0 hrs :  5 min :  17 s
AOS: 17:57:07.90 28/7/2014
```

```
Azi 197.81 Deg, Ele 6.60 Deg, Range 1743.26 km, SPA 89.98 Deg

MID: 17:59:46.42 28/7/2014
Azi 158.63 Deg, Ele 19.91 Deg, Range 1035.37 km, SPA 51.29 Deg

LOS: 18:02:24.93 28/7/2014
Azi 95.92 Deg, Ele 14.16 Deg, Range 1263.60 km, SPA 10.93 Deg

Observation 2.
Duration:  0 hrs :  6 min :  6 s
AOS: 17:59:01.41 30/7/2014
Azi 207.29 Deg, Ele 51.80 Deg, Range 525.34 km, SPA 89.67 Deg

MID: 18:02:04.43 30/7/2014
Azi 46.59 Deg, Ele 18.17 Deg, Range 1086.28 km, SPA 57.91 Deg

LOS: 18:05:07.45 30/7/2014
Azi 43.64 Deg, Ele 0.03 Deg, Range 2332.38 km, SPA 62.38 Deg

Nomenclature:
Azi - Azimuth
Ele - Elevation
SPA - Solar Phase Angle
AOS - Acquisition of Signal
MID - Middle of Visibility Period
LOS - Loss of Signal
```

B.3 Visualisation of Particle Representations

Utilising MASSAS's parallelised SGP4 orbit propagator, MASSAS is capable of processing and visualising the orbital dynamics of many more objects than currently reside in the largest RSO catalogues. Fig. B.5 displays a scene from a visualisation that displays the position of hundreds of thousands of randomly generated RSOs as they orbit the Earth. All object's positions can be computed quickly enough to permit real time visualisation of their motion.

This process can also be adapted to visualise the fragmentation of an object into many

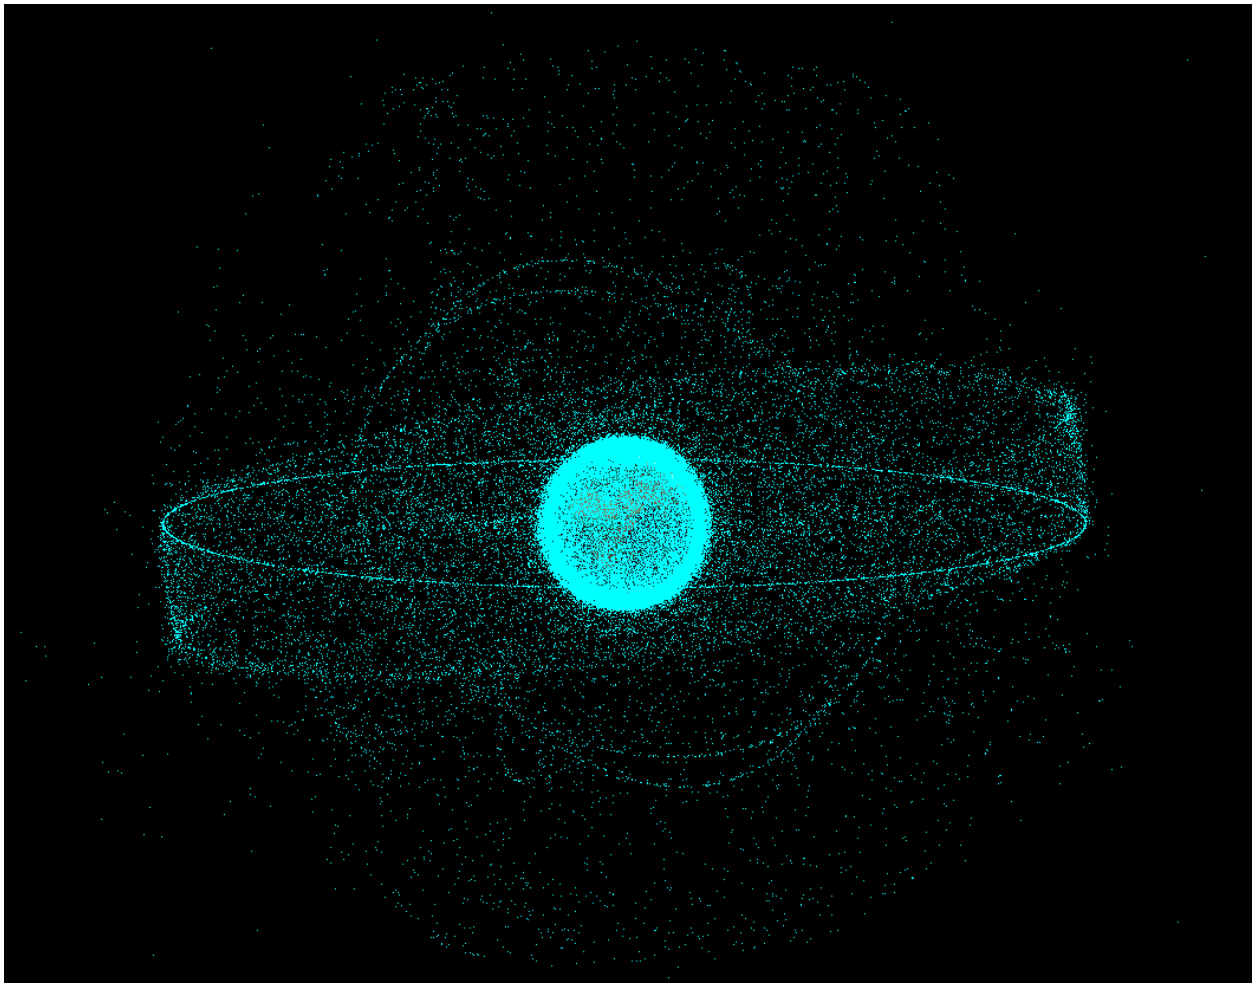

Figure B.5: A visualisation of hundreds of thousands of randomly generated RSOs orbiting the Earth.

thousands of pieces. Fig. B.6 displays a screen-shot of such a visualisation from two viewpoints. The hypothetical RSO's trajectory at the time of fragmentation is shown via an ellipse. The fragments are generated by applying random perturbations to many copies of the RSO's state vector at the time of fragmentation. Fig. B.6 shows the resulting cloud of debris distributed along the original trajectory of the RSO. Additionally, collision detection permits the simulation to recognise and visualise a line of debris that has impacted the Earth in the Indian Ocean and Western Australia.

The visualisation of hypothetical catalogues and fragmentation events are examples of applications that MASSAS could be used to investigate and visualise in future. Nonetheless, visualisations of this kind were also used to aid in the development and implementation of the research presented in this thesis. Processing and visualising particle representations was particularly useful for visualising Monte Carlo samples whilst performing the particle

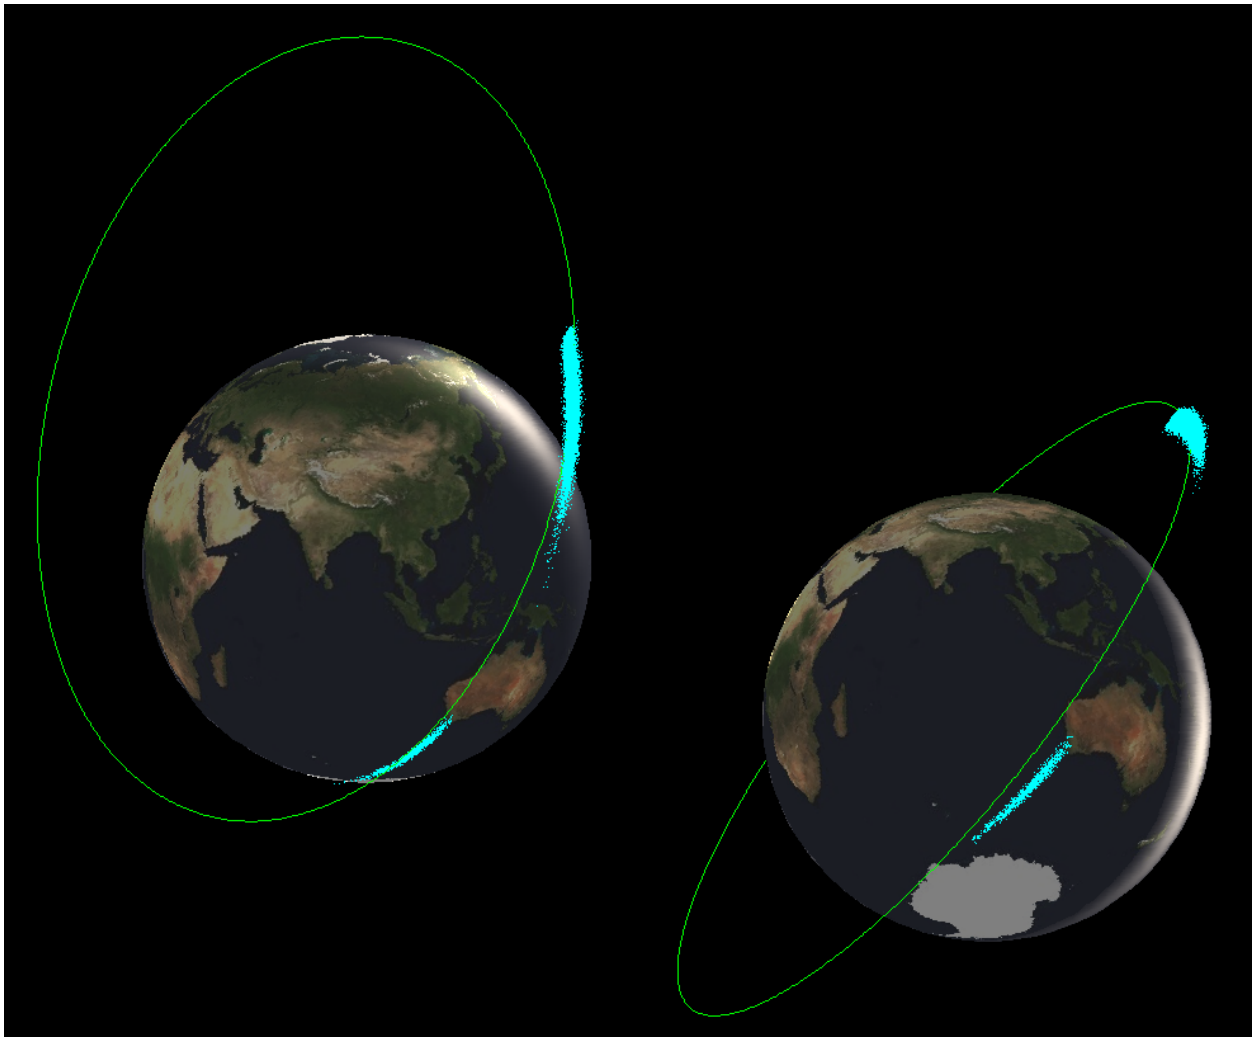

Figure B.6: Two camera angles of a scene generated by a module in MASSAS that visualises the fragmentation of an RSO. Collisions of debris with the Earth are visible in the Indian Ocean and Western Australia.

filtering detailed in Chapter 6. Fig. B.7 displays an exaggerated example of the propagation of particles representing the position component of an RSO's state p.d.f. Displayed chronologically from left to right, top to bottom, each image indicates the evolving shape and dispersion of an RSO's initially Gaussian p.d.f. Images such as Fig. B.7 demonstrate the non-Gaussian form that an RSO's p.d.f. can assume in a relatively short period of time.

A more practical demonstration of a particle visualisation is displayed in Fig. B.8. The visualisation displays the geometric relationship between the Earth, a sensor's steerable FOV and a particle representation of an RSO's p.d.f. The inset image on the right of Fig. B.8 is a zoomed version of the intersection between FOV and p.d.f. on the left. The white background was added for improved contrast when viewing the particles. Using such a visualisation, it is

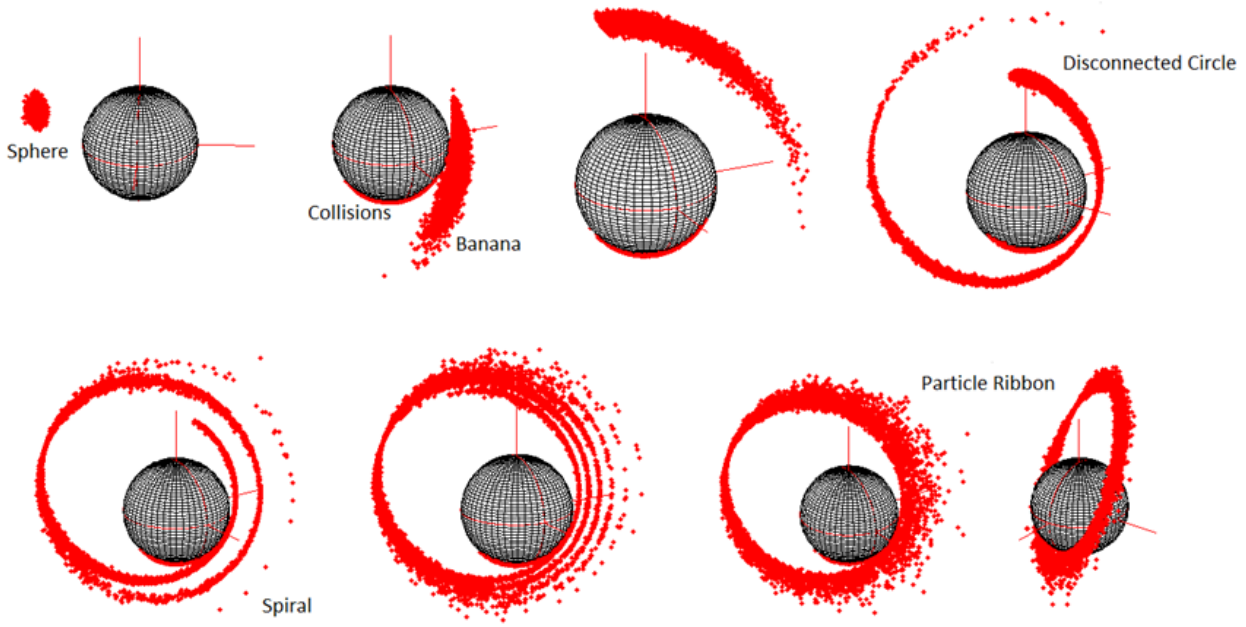

Figure B.7: Evolving particle representation of an RSO's p.d.f.

clear to the user—be they potential researcher or SNSS operator—where the system predicts the RSO is likely to be and how effectively the observing sensor is targeting the p.d.f.

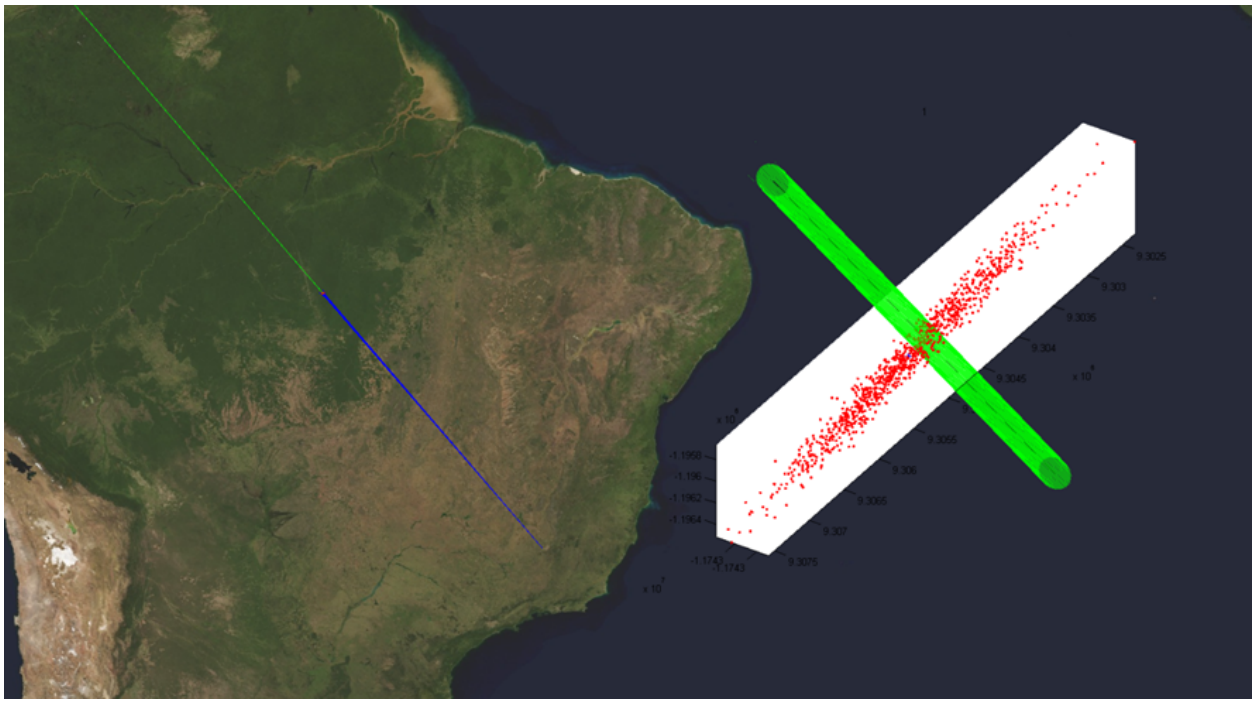

Figure B.8: A visualisation generated by MASSAS for providing situational awareness about an observation that is currently underway.

B.4 Visualisation of Dynamic Steering

Utilising a particle visualisation, as described in Section B.3, played a significant role in the development and implementation of the dynamic steering process proposed in Chapter 6. The visualisation was invaluable during development of the dynamic steering method as it ensured the software was behaving as intended and aided in conceptual understanding and planning. Furthermore, the visualisation can be generated during field trials to provide a means of monitoring the state of the system as experiments are conducted.

Figs. B.9-B.13 display example imagery, produced by MASSAS whilst implementing a simulation of a dynamically steered optical surveillance sensor. Each image shows consecutive attempts by the sensor to reacquire an elusive RSO, in spite of a relatively low probability of detection. Similar to Fig. B.8, each image contains an inset on the right that displays a zoomed view of the process occurring on the left. The observing sensor's FOV is represented by the blue cone. Whilst absent during practical experiments, the red line is a vector pointing from the sensor to the RSO's true location. It thereby indicates the direction—yet unknown to the sensor—that it needs to point to observe the target. The position component of the RSO's state error p.d.f. is represented by red and black particles. Red particles indicate particles with non-trivial weight while black particles represent particles whose weights have been reduced, as they were within the FOV when an observation was attempted but the target was not observed.

Figs. B.9-B.12 demonstrate how the dynamically steered system searches the p.d.f. for the target, recording failed observations as it progresses. Fig. B.13 displays a final successful attempt to reacquire the target. Due to the successful reacquisition, the p.d.f.'s particle representation has been resampled to better represent the posterior p.d.f.

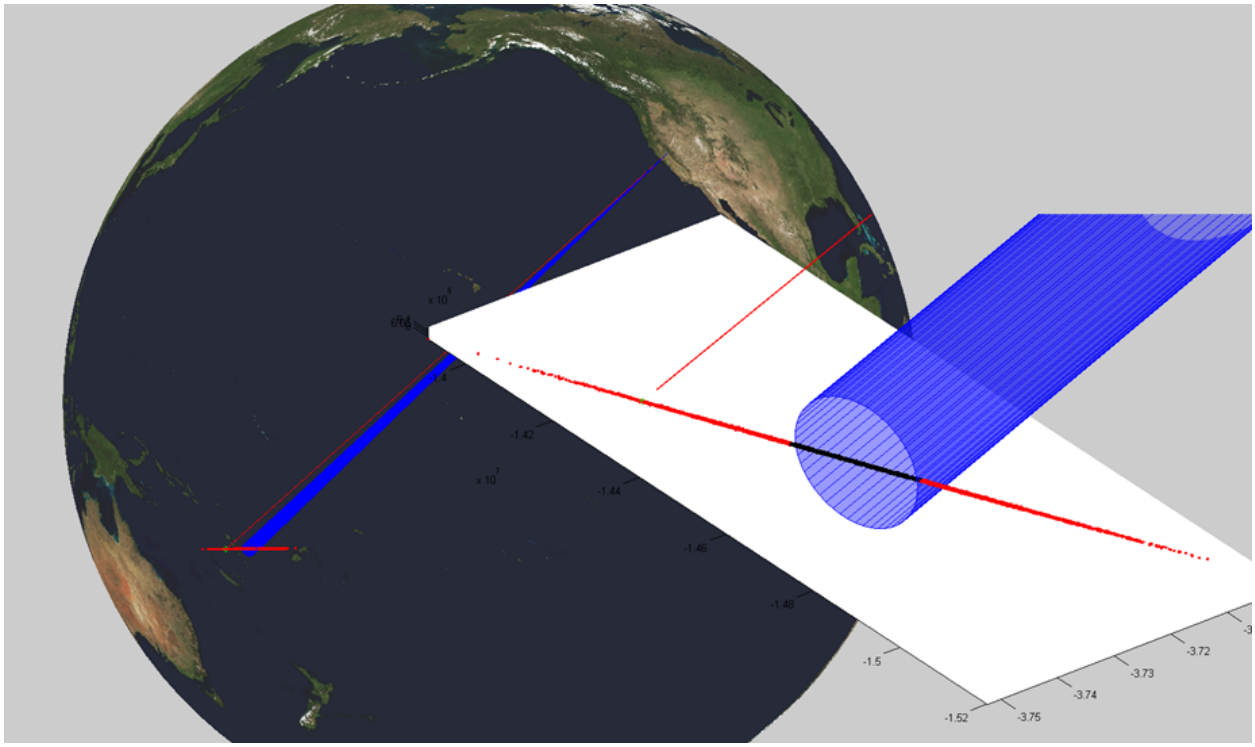

Figure B.9: Dynamic steering example: 1st observation, target not in FOV.

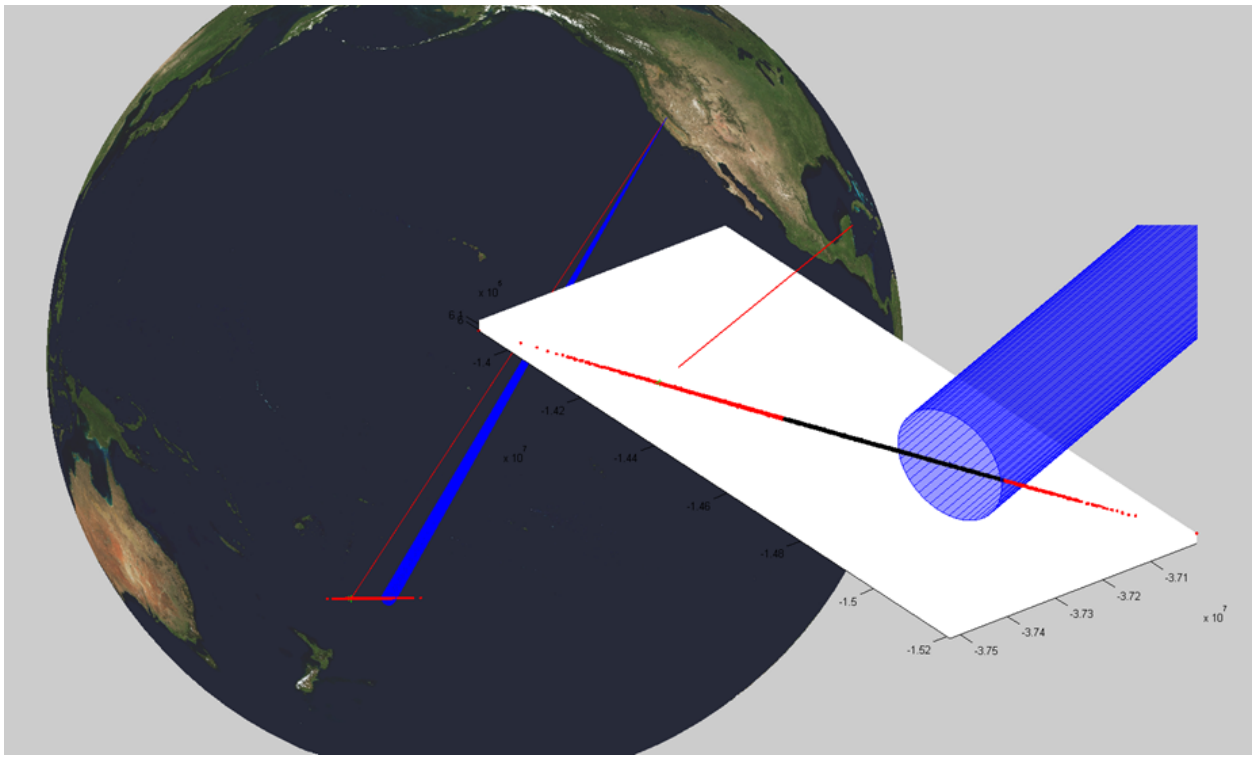

Figure B.10: Dynamic steering example: 2nd observation, target not in FOV.

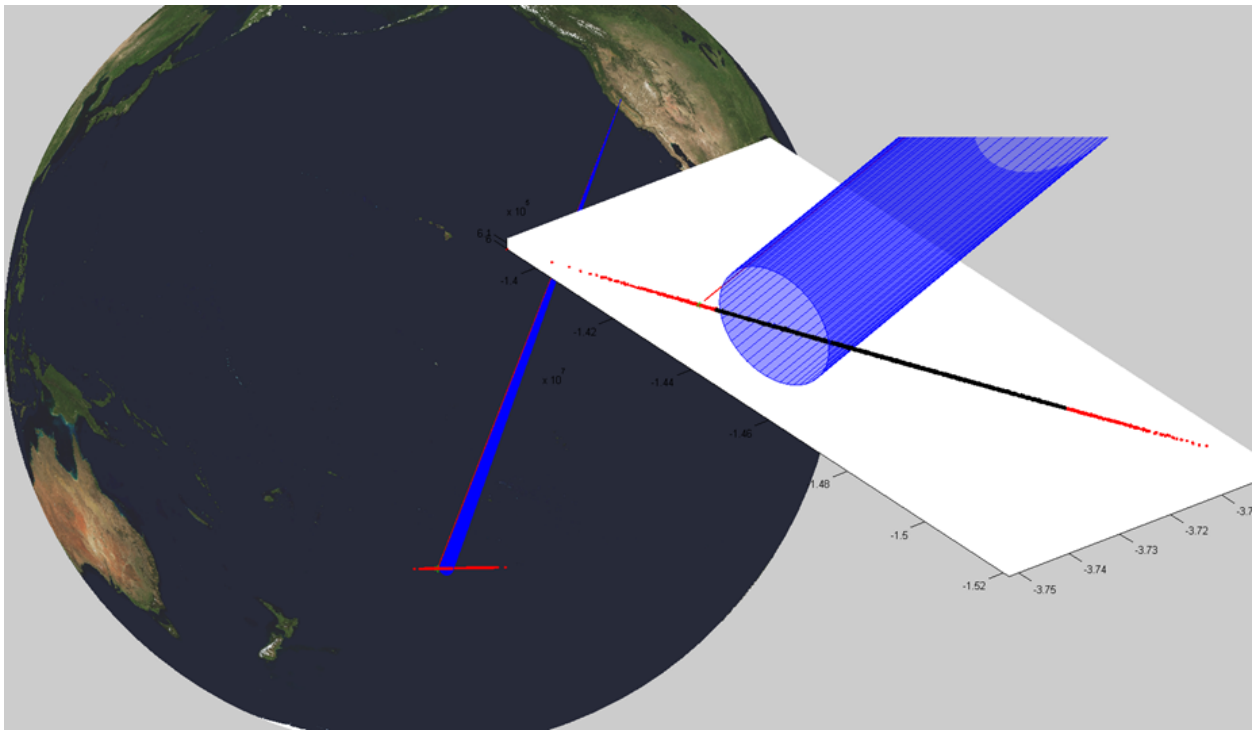

Figure B.11: Dynamic steering example: 3rd observation, target not in FOV.

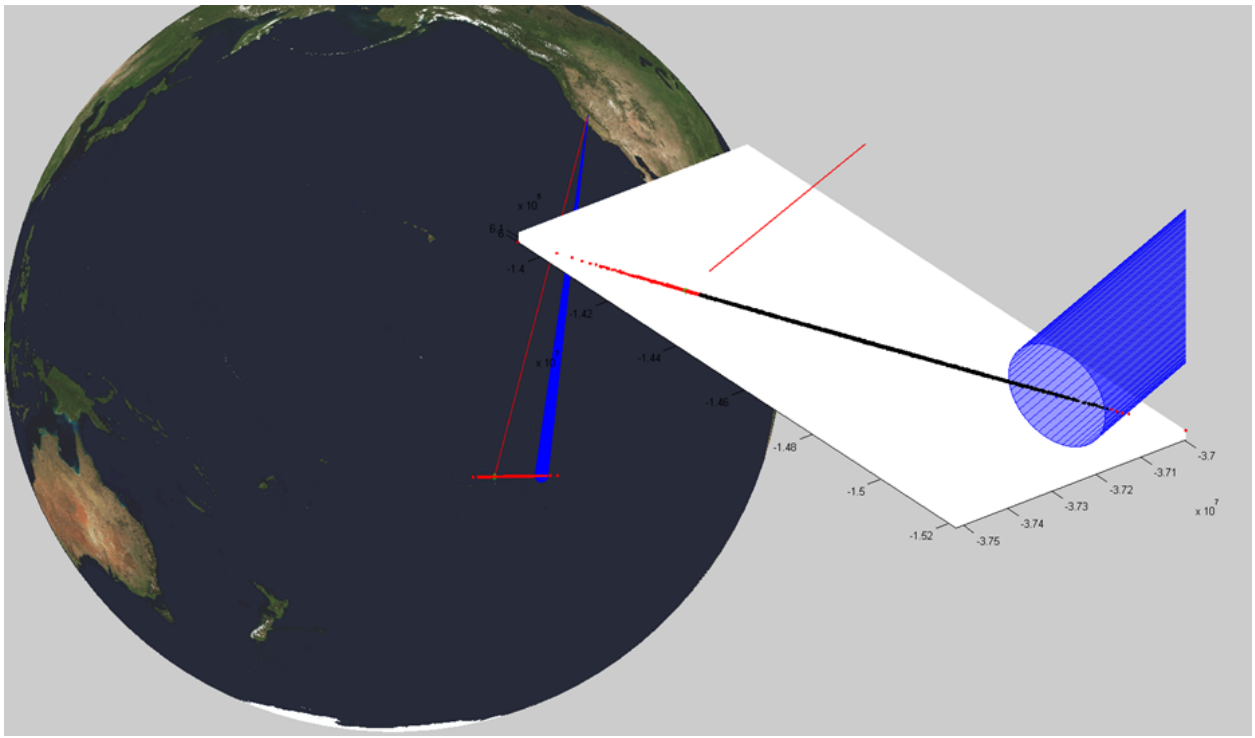

Figure B.12: Dynamic steering example: 4th observation, target not in FOV.

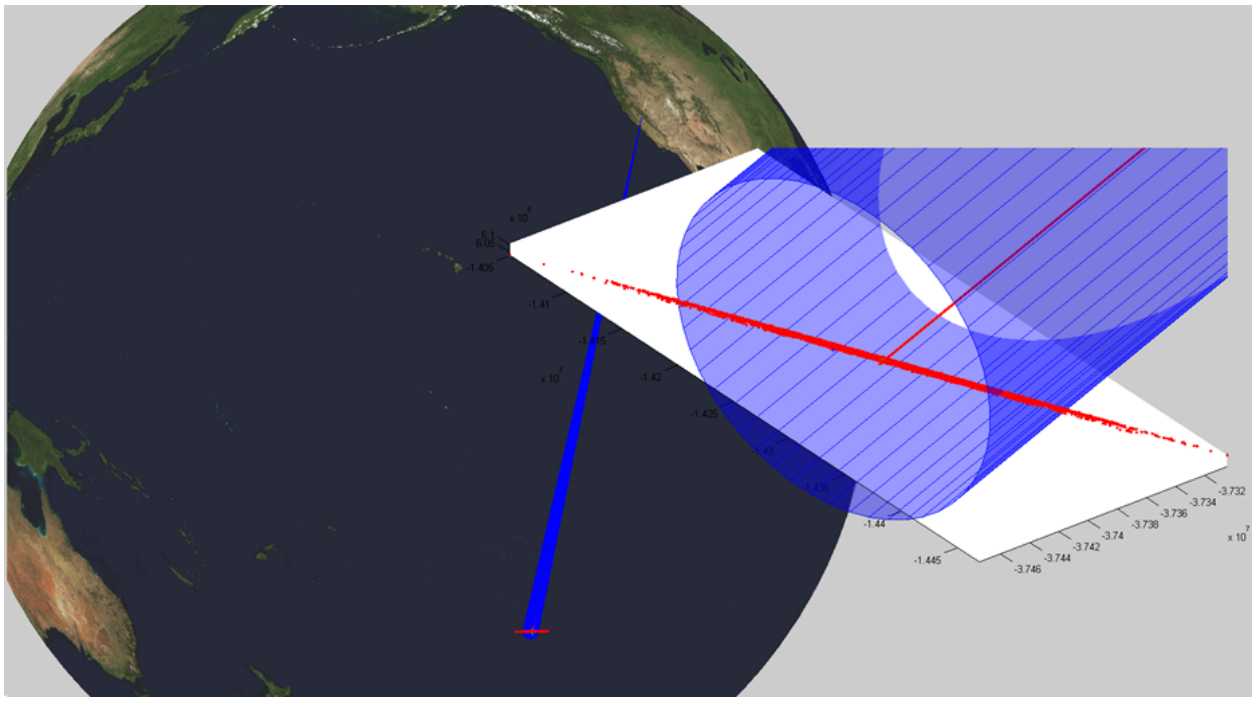

Figure B.13: Dynamic steering example: 5^{th} observation, target reacquired.

C

Collaborative Sensing - Ancillary Figures

Appendix C provides ancillary data for Chapter 4. The following items were detailed during the chapter:

- Scenario A - independent scheduling of sensors
- Scenario B - collaborative scheduling of sensors

Figs. C.1-C.4 present the WCPE and MCPE as recorded at the end of each scheduling period throughout each of the simulations presented in Chapter 4. Figs. C.1 & C.2 detail the catalogue's accuracy whilst utilising the USA sensor configuration. Figs. C.3 & C.4 detail the catalogue's accuracy whilst utilising the Global sensor configuration. Each pair of figures, display the results of simulations utilising unweighted observation effectiveness and weighted observation effectiveness respectively.

Figs. C.5 & C.6 display the percentage reduction in MCPE when utilising Scenrio B in place of Scenario A, for each catalogue configuration. The MCPE has been decomposed into radial, in-track and cross-track error components, prior to computing the reduction in error.

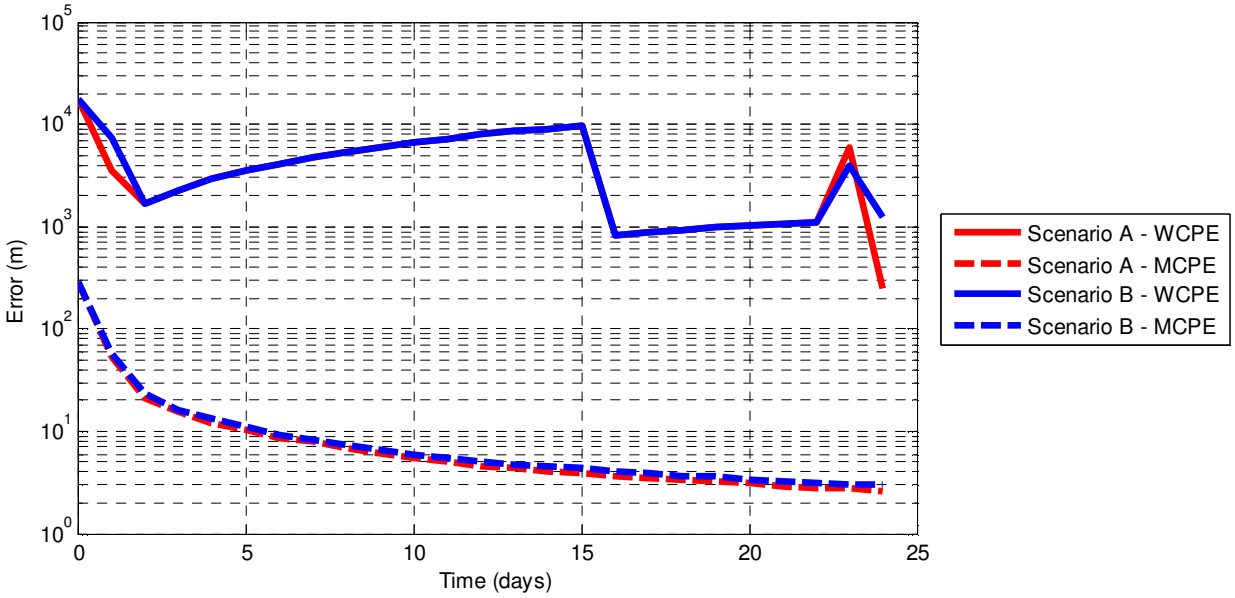

Figure C.1: USA sensor configuration (unweighted).

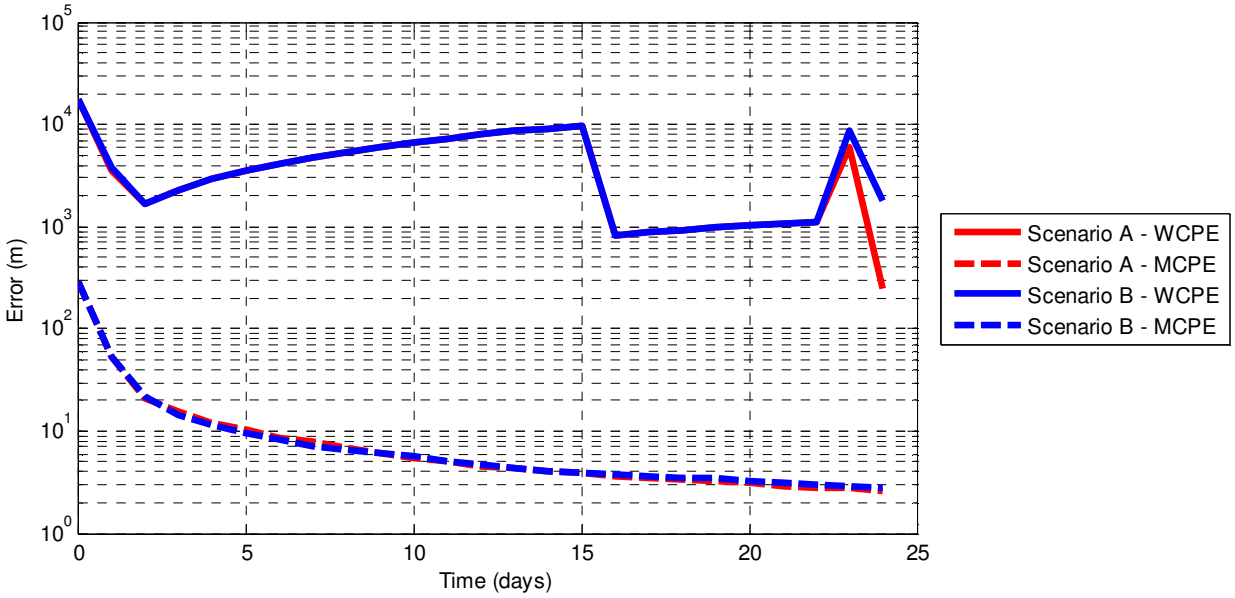

Figure C.2: USA sensor configuration (weighted).

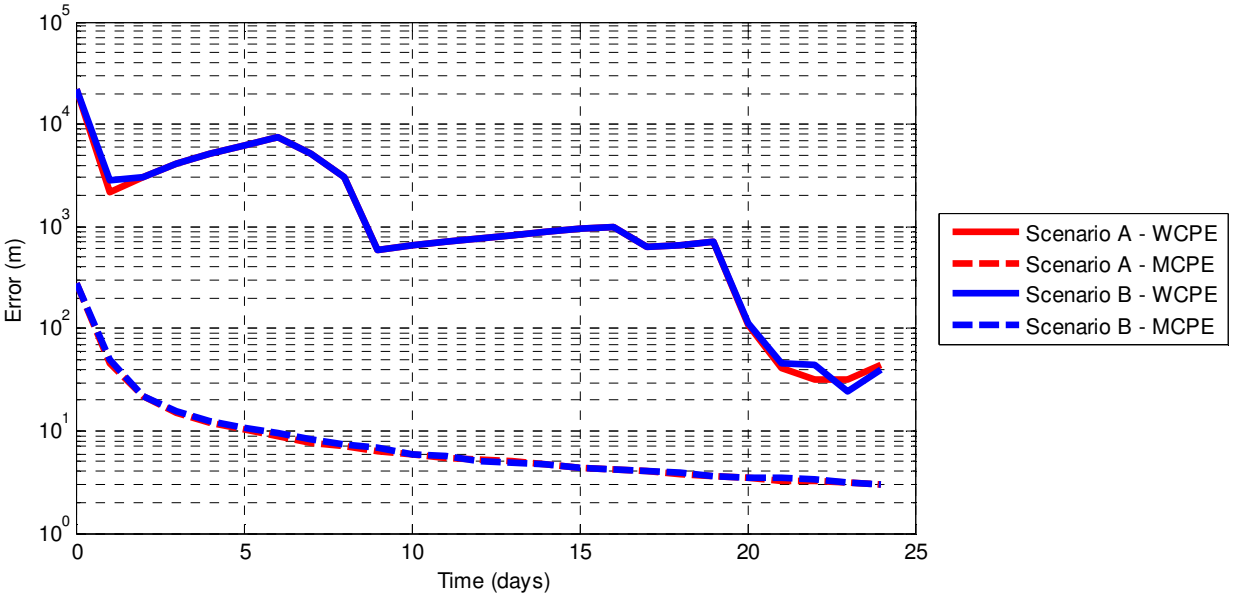

Figure C.3: Global sensor configuration (unweighted).

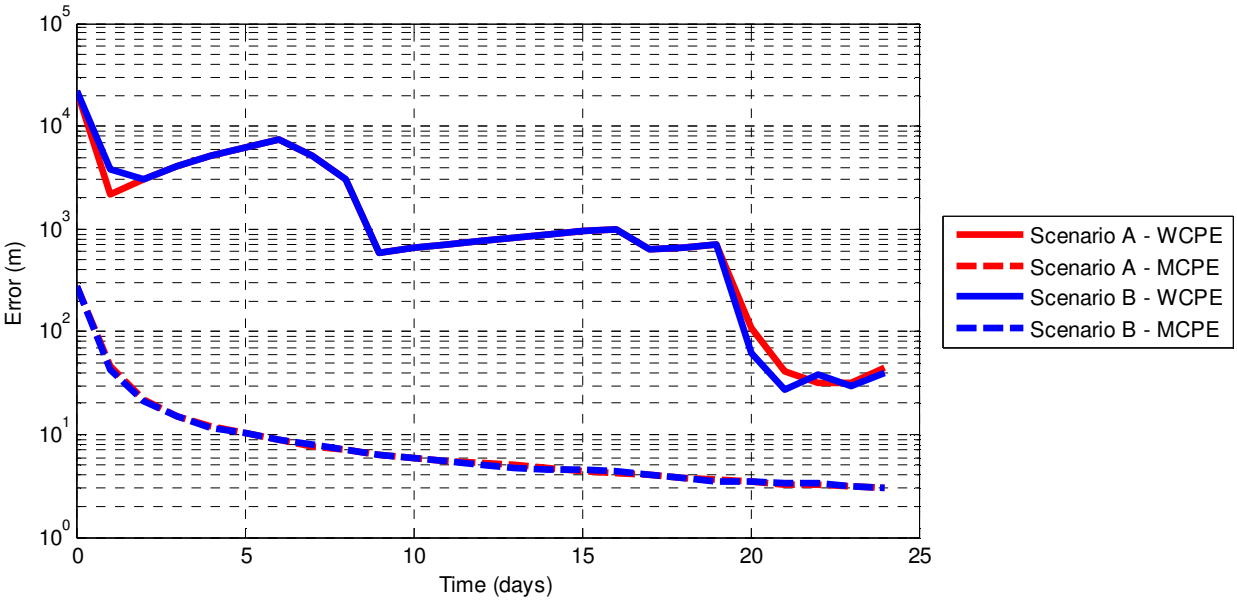

Figure C.4: Global sensor configuration (weighted).

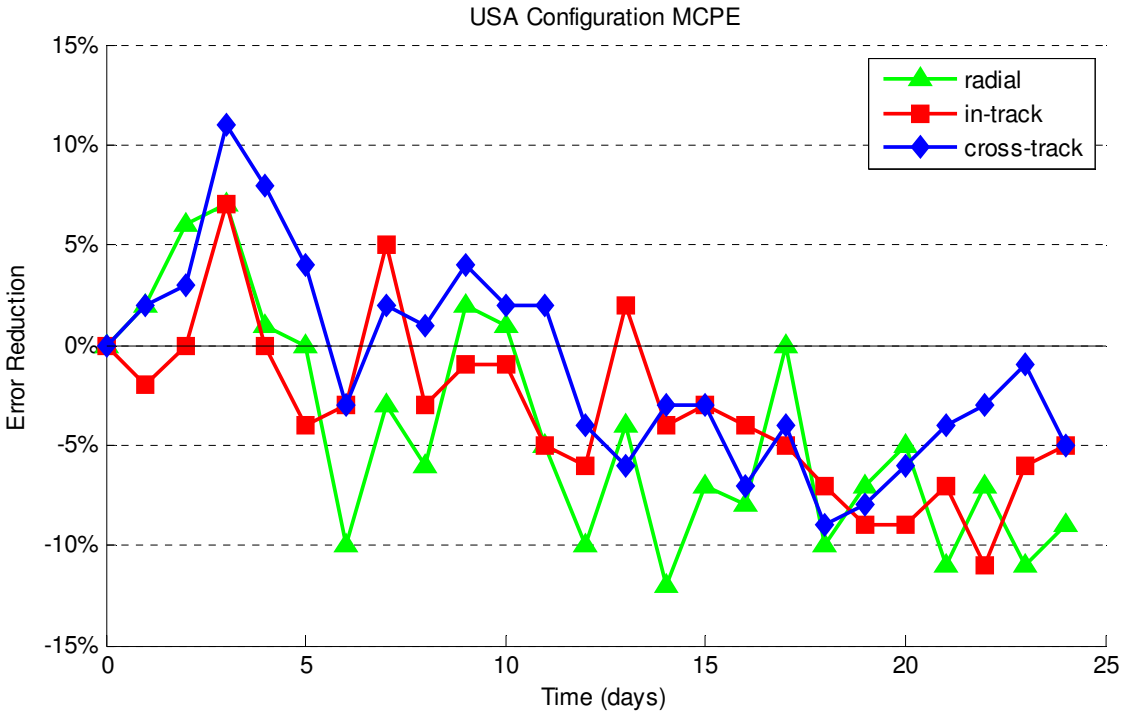

Figure C.5: MCPE error components during a weighted simulation utilising the USA sensor configuration.

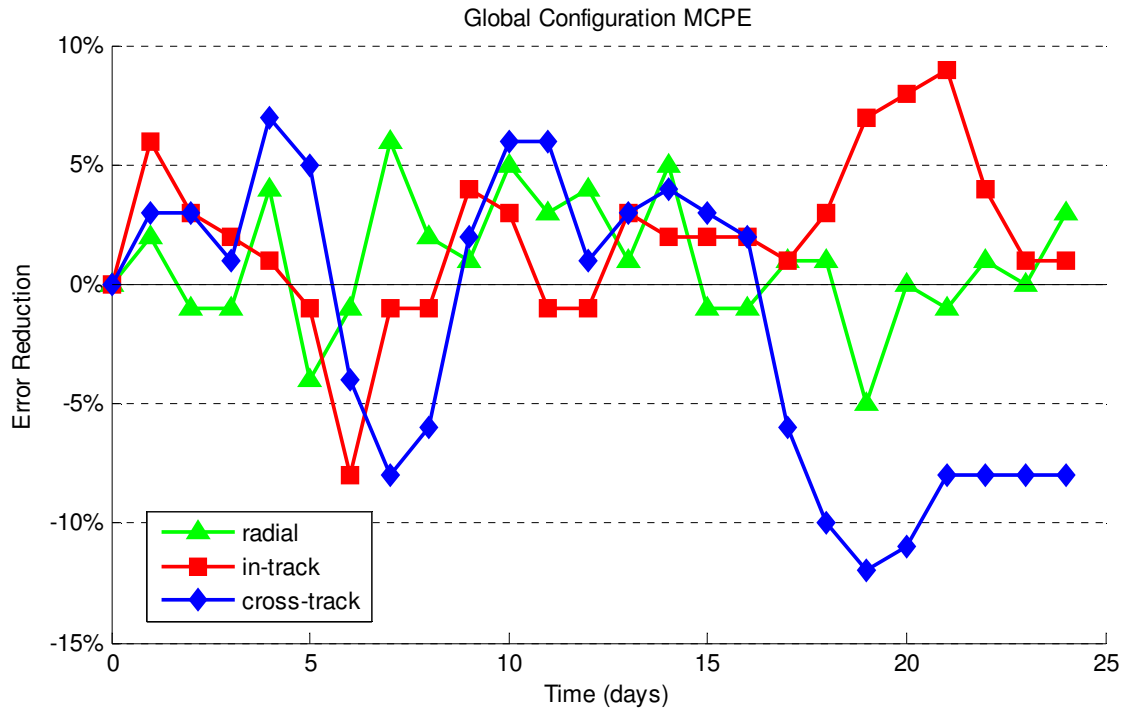

Figure C.6: MCPE error components during a weighted simulation utilising the Global sensor configuration.

D

Supplementary Sensor Site Specifications

Appendix D contains the precise locations of the sensors used for all simulations conducted in Chapters 4 & 5.

Table D.1 lists the names and locations of all simulated sensors. Whilst a number entries are the approximate locations of genuine space surveillance sensors, some sites were chosen solely to achieve a uniform global distribution.

Table D.1: The precise locations of simulated sensor sites used during Chapters 4 & 5.

Site Name	East Longitude (°)	North Latitude (°)	Height Above WGS-84 Ellipsoid (m)
Kwajalein, Pacific	-192.26670	8.71667	50
Albuquerque, USA	253.50272	34.96305	1725
Moron, Spain	354.41195	37.15110	101
Bangalore, India	77.80518	12.67178	920
Betim, Brazil	315.80612	-19.97404	840
Haleakala, USA	203.74337	20.70821	3000
Oregon	-121.02539	43.37205	400
West Virginia	-79.01367	38.87279	101
Brisbane, AU	153.02839	-27.47257	300
Diego Garcia, Indian Ocean	72.43080	-7.44307	10
Beale AFB	238.65000	39.14000	116
Lyndon, AU	115.26750	-23.62794	250
Katherine, AU	133.66242	-13.91612	180
Atacama desert, Chile	-70.40417	-24.62722	2625
Cerro Armazones, Chile	-70.19222	-24.58917	3040
Cederberg, South Africa	19.02701	-32.14047	1025
Yeste, Spain	-2.47858	38.38860	1605
Mauna Kea, USA	-155.47599	19.82627	4109
Delareyville, South Africa	25.78697	-26.87002	140
Cook Islands	-159.75818	-21.24404	465
Millstone, USA	288.51000	42.62000	123

Table D.2 details the seven SNSS configurations used to conduct the eight simulations detailed in Chapter 5. Each column indicates the number of optical 'O' and radar 'R' sensors used in each configuration. The locations of each of the sensors may be obtained from Table D.1.

Table D.2: Sensor site configurations used for each simulation detailed in Chapter 5.

Simulation Number	1		2 & 3		4		5		6		7		8	
Sensor Type	O	R	O	R	O	R	O	R	O	R	O	R	O	R
Total Sensors	8	0	0	8	8	2	8	4	15	0	20	0	17	1
Kwajalein, Pacific	•			•	•		•		•		•		•	
Albuquerque, USA	•			•	•		•		•		•		•	
Moron, Spain	•			•	•		•		•		•		•	
Bangalore, India	•			•	•		•		•		•		•	
Betim, Brazil	•			•	•		•		•		•		•	
Haleakala, USA	•			•	•		•		•		•		•	
Oregon								•	•		•		•	
West Virginia									•		•		•	
Brisbane, AU	•			•	•		•		•		•		•	
Diego Garcia, Indian Ocean	•			•	•		•	•	•		•		•	
Beale AFB						•		•	•		•		•	
Lyndon, AU									•		•		•	
Katherine, AU									•		•		•	
Atacama desert, Chile									•		•		•	
Cerro Armazones, Chile									•		•		•	
Cederberg, South Africa											•		•	
Yeste, Spain											•		•	
Mauna Kea, USA								•			•			•
Delareyville, South Africa											•			
Cook Islands											•			
Millstone, USA						•								

Glossary

active sensing A method of probing or interrogating a region of interest by generating and projecting energy into the region and analysing the reflected energy.
angle of view The angle that subtends a sensor's field of view.
apogee The largest distance between the centre of mass of the Earth and an orbiting object's centre of mass, as predicted by its osculating orbit.
availability The degree to which or ratio of time a device or system of devices are capable of achieving their intended mission objectives.

bore-sight The central axis of a sensor's field of view. For example, if an RSO is aligned with an optical sensor's bore-sight, the RSO will appear in the centre of an image taken by the sensor.

celestial reference Typically a star or similarly bright celestial body that is located an extremely long distance away from the Earth, such that it appears stationary relative to the Earth's own motion. The term may also refer to a geometric reference that remains stationary in inertial space, such as the direction of a vector point from the Earth to the Sun at a specific equinox in a specific year.
conventional radar Radars that sweep their surveillance volumes with radio waves in regular intervals — as opposed to using an agile beam or a fixed fan of radio waves.
cooperative target An RSO that has been specifically engineered or augmented with devices to aid in the determination of its orbit. The term may be further refined to indicate if such devices and augmentations may be used by a specific agency as certain RSOs may not 'cooperate' with all surveillance agencies.

data association The process of assigning a single or multiple measurements to an identified target while reducing the likelihood of attributing measurements to the wrong target due to the use of an imperfect sensor.

Earth orbit The orbit of a celestial body or man-made space object whose dominant

gravitational influence is the Earth.

epoch A specific moment in time—used commonly to signify an event such as the time at which an observation occurred or the moment from which a TLE's nominal state may be used to extrapolate an RSO's state at any other time.

field of view The volume in which a sensor may detect the presence of a target at a given moment in time.

fragmentation debris The shards or fragments of an RSO that has been torn apart as a result of an event such as an intentional or unintentional explosion or the collision of two or more objects.

geostationary orbit A type of geosynchronous orbit that causes an RSO to appear stationary with respect to the Earth's surface, as the RSO is constrained to the equatorial plane.

geosynchronous orbit A type of orbit that results in the synchronisation of an RSO's motion with the rotation of the Earth.

greedy A class of algorithm that aims to solve a problem by making a choice that optimises the outcome at a particular stage as opposed to seeking optimality over all stages.

high Earth orbit An object is in high Earth orbit if the perigee of its orbit is greater than 42 164 km from the centre of the Earth and it is orbiting the Earth.

Hill Sphere The region within which an astronomical body's gravitational sphere of influence is dominant.

inertial space A frame of reference that does not undergo acceleration.

instantaneous orbit Also known as an unperturbed two-body or osculating orbit, an instantaneous orbit is the theoretical orbital motion of the smaller of two bodies caused by a mutual attractive force as a result of each body exhibiting an isotropic gravitational field.

jitter A process of intentionally adding a random deviation to a parameter from a nominal value.

Keplerian orbital elements A six element parameter set that may be used to uniquely define the instantaneous or osculating orbit of an object.

light time The time light would take to propagate from its source or a reflective surface to a receiver such as an observer's eye or an optical sensor.

loss of particle diversity In the context of particle filtering: when many particles share the same state space.

low Earth orbit An object is in low Earth orbit if the apogee and perigee of its orbit fall within, approximately, 100 km and 2000 km of the surface of the Earth.

Markov process A stochastic process whose conditional probability distribution of future states can be predicted using the current state of the process; thereby not requiring knowledge of any previous states.

medium Earth orbit An object is in medium Earth orbit if the apogee and perigee of its orbit fall within 8000 km and 42 164 km from the centre of the Earth.

Molniya orbit A type of highly elliptical orbit used by Russian space agencies to achieve non-geostationary orbits that may be used to approximate geostationary mission profiles at high latitudes when using multiple spacecraft.

osculating orbit Also known as an unperturbed two-body or instantaneous orbit, an osculating orbit is the theoretical orbital motion of the smaller of two bodies caused by a mutual attractive force as a result of each body exhibiting an isotropic gravitational field.

particle degeneracy In the context of particle filtering: when all but a few particles have trivial weight.

passive debris RSOs that are not themselves, or in anyway part of, an operational spacevehicle. Examples include fragmentation debris, discarded components or satellites that are no longer operational.

passive sensing A method for interrogating a region of interest by analysing the energy emitted or reflected by objects within the region.

perigee The smallest distance between the centre of mass of the Earth and an orbiting object's centre of mass, as predicted by its osculating orbit.

prograde orbit To orbit about the Earth's polar axis in the same direction as the Earth's rotation.

resident space object A man-made object that is currently orbiting the Earth.

retro-reflector A device that is designed to reflect light back to the source, irrespective of the angle of incidence with which light strikes the device.

retrograde orbit To orbit about the Earth's polar axis in the opposite direction to the Earth's rotation.

sensor management The process in which sensors are tasked to observe a region within their surveillance volume via an ad-hoc command or a prepared schedule.

sidereal day The period of time—approximately 86164.090517s—the Earth takes to rotate once on its axis in inertial space.

space situational awareness Refers to a knowledge of the man-made space environment surrounding the Earth. This knowledge is useful for conducting, planning and monitoring space activities.

steerable sensor A space surveillance sensor whose finite FOV is small enough to warrant non-uniform steering about its surveillance volume in order to maximise its utility.

Supercomputing A method of combining a large number of computer processors together to achieve a parallel computing capability far surpassing that of conventional computers.

surveillance volume The volume in which a sensor or network of sensors have the ability to detect and measure information about targets currently residing within the volume.

References

[1] American Institute of Aeronautics and Astronautics, *Utilization of Space: Today and Tomorrow*, B. Feuerbacher and H. Stoewer, Eds. Springer Berlin Heidelberg, 2006.

[2] E. Seedhouse, *SpaceX: Making Commercial Spaceflight a Reality*, ser. Springer Praxis Books. Springer New York, 2013.

[3] J. Aron, "Private space race comes of age." *New Scientist*, vol. 219, no. 2935, pp. 8–9, 21 Sep 2013.

[4] L. Billings, "Exploration for the masses? Or joyrides for the ultra-rich? Prospects for space tourism." *Space Policy*, vol. 22, no. 3, pp. 162–164, 2006.

[5] J. P. Penn and C. A. Lindley, "Requirements and approach for a space tourism launch system," *Acta Astronautica*, vol. 52, no. 1, pp. 49–75, 2003.

[6] M. Slezak, "Space miners hope to build first off-earth economy," *New Scientist*, vol. 217, no. 2906, pp. 8–10, 2013.

[7] C. Lewicki, P. Diamandis, E. Anderson, C. Voorhees, and F. Mycroft, "Planetary resourcesthe asteroid mining company," *New Space*, vol. 1, no. 2, 2013.

[8] V. A. Chobotov, *Orbital mechanics*, ser. AIAA education series. Washington, DC: American Institute of Aeronautics and Astronautics, 1991.

[9] D. A. Vallado and W. D. McClain, *Fundamentals of Astrodynamics and Applications*, 3rd ed., ser. Space Technology Library. New York: Springer, 2007.

[10] NASA Orbital Debris Program Office, "Image: Monthly number of objects in earth orbit by object type," *Orbital Debris Quarterly News, National Aeronautics and Space Administration*, vol. 17, no. 1, 2013, j.C. Liou (Tech. Ed.) D. Shoots (Mng. Ed.) [Online]. Available: http://orbitaldebris.jsc.nasa.gov/newsletter/pdfs/odqnv17i1.pdf, Accessed: 11 Jul 2014.

[11] N. N. Smirnov, Ed., *Space Debris: Hazard Evaluation and Mitigation*, ser. Earth Space Institute. London, New York: CRC Press, 2002, vol. 6.

[12] USSTRATCOM, "USSTRATCOM space control and space surveillance fact sheet," 2013, [Online]. Available: http://www.stratcom.mil/factsheets/11/Space_Control_and_Space_Surveillance/, Accessed: 7 Feb 2014.

[13] T. S. Kelso, "Analysis of the Iridium 33-Cosmos 2251 collision," 2009.

[14] ——, "Analysis of the 2007 Chinese ASAT test and the impact of its debris on the space environment," 2007.

[15] J. Mackey, "Recent US and Chinese antisatellite activities," *Air University Maxwell Air Force Base Air Force Research Institute*, 2009.

[16] D. J. Kessler and B. G. Cour-Palais, "Collision frequency of artificial satellites: The creation of a debris belt," *Journal of Geophysical Research*, vol. 83 (A6), p. 26372646, 1978.

[17] U. N. G. Assembly, "A/res/62/217 62nd session - international cooperation in the peaceful uses of outer space," 2008.

[18] B. Weeden, P. Cefola, and J. Sankaran, "Global space situational awareness sensors," in *The Advanced Maui Optical and Space Surveillance Conference*, 2010.

[19] M. Wasson, "Space situational awareness in the Joint Space Operations Center," 2011.

[20] K. Becker, E. Detsis, C. Nwosa, V. Palevska, M. Rathnasabapathy, and M. Taheran, *Space Situational Awareness*, ser. SSS Educational Series. Vienna, Austria: Space Generation Advisory Council, 2012.

[21] N. Bobrinsky and L. Del Monte, "The space situational awareness program of the European Space Agency," *Cosmic Research*, vol. 48, no. 5, pp. 392–398, 2010.

[22] N. R. C. of the National Academies, *Continuing Kepler's Quest: Assessing Air Force Space Command's Astrodynamics Standards*. The National Academies Press, 2012.

[23] K. T. Alfriend, M. R. Akella, J. Frisbee, J. L. Foster, D.-J. Lee, and M. Wilkins, "Probability of collision error analysis," *Space Debris*, vol. 1, no. 1, pp. 21–35, 1999.

[24] H. Klinkrad, J. R. Alarcon, and N. Sanchez, "Collision avoidance for operational ESA satellites," in *Proceedings of the 4th European Conference on Space Debris (ESA SP-587)*. ESA/ESOC, 2005, p. 509.

[25] J. Stupl, J. Mason, C. Levit, W. Marshall, A. G. Salas, C. Smith, S. Olivier, A. Pertica, W. D. Vries, and W. Ting, "Lightforce: An update on orbital collision avoidance using photon pressure," in *The Advanced Maui Optical and Space Surveillance Conference*, 2012.

[26] B. D. Tapley, B. E. Schutz, and G. H. Born, *Statistical Orbit Determination*. Amsterdam; Boston: Elsevier Academic Press, 2004.

[27] T. Kelecy and M. Jah, "Analysis of orbit prediction sensitivity to thermal emissions acceleration modeling for high area-to-mass ratio (HAMR) objects," in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf., (AMOS)*, 2009.

[28] J. Sharma, G. H. Stokes, C. v. Braun, G. Zollinger, and A. J. Wiseman, "Toward operational space-based space surveillance," *Lincoln Laboratory Journal*, vol. 13, no. 2, pp. 309–334, 2002.

[29] J. T. McGraw, M. R. Ackermann, J. B. Martin, and P. C. Zimmer, "The Air Force Space Surveillance Telescope," in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf., (AMOS)*, 2003.

[30] H. Thatcher, "AFSPC A5CS SSA activities," 2011.

[31] D. A. Vallado and J. D. Griesbach, "Simulating space surveillance networks," in *Proc. AAS/AIAA Astrodynamics Specialists Conf.*, 2012.

[32] C. Sabol, "A role for improved angular observations in geosynchronous orbit determination," PhD, University of Colorado at Boulder, 1998.

[33] K. G. Henize, J. F. Stanley, C. A. O'Neill, and B. S. Nowakowski, "Detection of orbital debris with GEODSS telescopes," *Optical Engineering and Photonics in Aerospace Sensing*, pp. 76–84, 1993.

[34] R. Leitch and I. Hemphill, "Sapphire: A small satellite system for the surveillance of space," in *24th Annual AIAA/USU Conference on Small Satellites*, 2010.

[35] P. F. Sydney, J. L. Africano, A. Fredericks, K. M. Hamada, V. Soo Hoo, D. L. Nishimoto, P. W. Kervin, S. Bisque, and M. Bisque, "Raven automated small telescope systems," *SPIE , Imaging Technology and Telescopes*, vol. 4091, no. 237, 2000.

[36] C. Sabol, K. Luu, P. Kervin, D. Nishimoto, K. Hamada, and P. Sydney, "Recent developments of the raven small telescope program," *Advances in the Astronautical Sciences*, vol. 112, no. AAS 02-131, pp. 397–416, 2002.

[37] R. D. Coder and M. J. Holzinger, "Autonomy architecture for a raven-class telescope with space situational awareness applications," in *2013 AAS/AIAA*, 2013.

[38] J. Wertz and W. J. Larson, *Space Mission Analysis and Design*, ser. Volume 8 of Space Technology Library. Springer Netherlands, 1999.

[39] USSTRATCOM, "Space Track," 2004, [Online]. Available: www.space-track.org, Accessed: 27 Oct 2013.

[40] O. Montenbruck and G. Eberhard, *Satellite orbits : models, methods, and applications*. Berlin New York: Springer, 2000.

[41] K. OKeefe, "Availability and reliability advantages of gps/galileo integration," in *14th International Technical Meeting of the Satellite Division of the Institute of Navigation*, 2001.

[42] Y. Yang, J. Li, J. Xu, J. Tang, H. Guo, and H. He, "Contribution of the compass satellite navigation system to global PNT users," *Chinese Science Bulletin*, vol. 56, no. 26, pp. 2813–2819, 2011.

[43] S. S. Blackman and R. Popoli, *Design and analysis of modern tracking systems*, ser. Artech House radar library. Boston, MA: Artech House, 1999.

[44] R. P. Mahler, *Statistical multisource-multitarget information fusion*. Artech House ˆ eBoston Boston, 2007, vol. 685.

[45] ——, *Advances in Statistical Multisource-Multitarget Information Fusion*. Artech House, 2014.

[46] A. R. Mishra, *Fundamentals of Cellular Network Planning and Optimisation: 2G/2.5G/3G... Evolution to 4G*. John Wiley & Sons, 2004.

[47] R. Stöckli, E. Vermote, N. Saleous, R. Simmon, and D. Herring, "The blue marble next generation - a true color earth dataset including seasonal dynamics from MODIS," 2005, published by the NASA Earth Observatory, Corresponding author: rstockli@climate.gsfc.nasa.gov.

[48] C. Ma, E. F. Arias, T. M. Eubanks, A. L. Fey, A. M. Gontier, C. S. Jacobs, O. J. Sovers, B. A. Archinal, and P. Charlot, "The international celestial reference frame as realized by very long baseline interferometry," *The Astronomical Journal*, vol. 116, no. 1, pp. 516–546, 1998.

[49] J. Nelson, M. Goda, P. Sydney, C. Sabol, D. Talent, D. OConnell, and M. Murai, "Daylight astrometry and design studies for the LEO Raven," in *14th AAS/AIAA Space Flight Mechanics Conference*, 2004.

[50] K. Hill, P. Sydney, K. Hamada, R. Cortez, K. Luu, M. Jah, P. W. S. Jr., M. Coulman, J. Houchard, and D. Naho'olewa, "Covariance-based network tasking of optical sensors," in *Proc. AAS/AIAA Space Flight Mechanics Meeting*, 2010.

[51] N. Sánchez-Ortiz, E. Olmedo-Casal, N. Guijarro-Lpez, and M. Bell-Mora, "AS4, a simulator supporting the definition of the European space surveillance segment of ssa," 7-9 Jun 2011.

[52] Analytical Graphics, Inc., "AGI: Software to model, analyze and visualize space, defence and intelligence systems," 2014, [Online]. Available: www.agi.com.

[53] S. Haykin, *Adaptive Filter Theory*, 3rd ed. Upper Saddle River, New Jersey: Prentice-Hall, 1996.

[54] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter*. Artech House, 2004.

[55] S. Challa, M. R. Morelande, D. Musicki, and R. J. Evans, *Fundamentals of object tracking*. Cambridge University Press, 2011.

[56] D. Simon, *Optimal state estimation : Kalman, H and nonlinear approaches*. Hoboken, N.J.: Wiley-Interscience, 2006.

[57] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte, "A new approach for filtering nonlinear systems," in *Proceedings of the 1995 American Control Conference*, vol. 3, Jun 1995, pp. 1628–1632.

[58] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[59] D. A. Danielson, D. Canright, D. N. Perini, and J. Schumacher, Paul W., "The Naval Space Command automatic differential correction process," 1999, american Astronautical Society (AAS) / American Institute of Aeronautics and Astronautics (AIAA) Astrodynamics Conference.

[60] M. S. Grewal and A. P. Andrews, "Applications of kalman filtering in aerospace 1960 to the present [historical perspectives]." *Control Systems, IEEE*, vol. 30, no. 3, pp. 69–78, 2010.

[61] S. J. Julier and J. K. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," in *AeroSense: 11th Int. Symp. Aerospace/Defence Sensing, Simulation and Controls*, 1997, pp. 182–193.

[62] J. L. Crassidis and F. L. Markley, "Unscented filtering for spacecraft attitude estimation," in *Proc. AIAA 2003-5484, AIAA Guidance, Navigation and Control Conf.*, 2003.

[63] O. Cappe, S. J. Godsill, and E. Moulines, "An overview of existing methods and recent advances in sequential Monte Carlo," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 899–924, 2007.

[64] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.

[65] P. J. Cefola, B. Weeden, and C. Levit, "Open source software suite for space situational awareness and space object catalog work," in *4th International Conference on Astrodynamics Tools Techniques*, 2010.

[66] NASA's Goddard Space Flight Center, "Orbit Determination Toolbox (ODTBX)," 2014, [Online]. Available: http://opensource.gsfc.nasa.gov/projects/ODTBX.

[67] P. Kervin, V. Soo Hoo, D. Nishimoto, and D. Liang, "Rapidly deployable raven-class systems SSA support in the field," 1-4 Sep 2009.

[68] D. A. Vallado, P. Crawford, R. Hujsak, and T. S. Kelso, "Revisiting Spacetrack Report NO.3," in *Proc. AIAA/AAS Astrodynamics Specialist Conf.*, no. AIAA 2006-6573, Aug 2006.

[69] K. Hill, P. Sydney, R. Cortez, K. Hamada, D. Nishimoto, K. Luu, and J. Schumacher, Paul W, "Dynamic tasking of networked sensors using covariance information," in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf., (AMOS)*, Sep 14-17 2010.

[70] T. Flohrer, H. Krag, and H. Klinkrad, "Assessment and categorization of tle orbit errors for the us ssn catalogue," in *The Advanced Maui Optical and Space Surveillance Conference*, 16-19 Sep 2008.

[71] V. P. Osweiler, "Covariance estimation and autocorrelation of norad two-line element sets," Master's thesis, Air Force Institute of Technology, 2006.

[72] K. T. Alfriend, C. Sabol, and K. K. Luu, "Orbit update comparison using optical and radar systems," in *Proc. the Advanced Maui Optical and Space Surveillance Technologies Conf., (AMOS)*, September 13-17 2004.

[73] D. L. Hall and J. Llinas, "An introduction to multisensor data fusion," *Proceedings of the IEEE*, vol. 85, no. 1, pp. 6–23, 1997.

[74] C. Sabol, S. Carter, and D. A. Vallado, "A fresh look at angles-only orbit determination," in *AAS/AIAA Astrodynamics Specialist Conference*. AAS Publications Office, 16-19 Aug 1999.

[75] M. L. Thrall, "Orbit determination of highly eccentric orbits using a raven telescope," Master's thesis, Naval Postgraduate School, Monterey, 2005.

[76] H. Yim and D. Chung, "Validation on residual variation and covariance matrix of USSTRATCOM Two Line Element," *Journal of Astronomy and Space Sciences*, vol. 29, no. 3, pp. 287–293, 2012.

[77] S. Aida, M. Kirschner, M. Wermuth, and R. Kiehling, "Collision avoidance operations for LEO satellites controlled by GSOC," in *SpaceOps 2010*, 2010.

[78] S. Aida, T. Patzelt, L. Leushacke, M. Kirschner, and R. Kiehling, "Monitoring and mitigation of close proximities in Low Earth Orbit," in *Proceedings of the 21st International Symposium on Space Flight Dynamics*, 2009.

[79] J. M. Nash, "Optimal allocation of tracking resources," *Decision and Control including the 16th IEEE Symposium on Adaptive Processes and A Special Symposium on Fuzzy Set Theory and Applications*, pp. 1177–1180, 1977.

[80] J. Ender, L. Leushacke, A. Brenner, and H. Wilden, "Radar techniques for space situational awareness," pp. 21–26, 7-9 Sep 2011.

[81] R. S. Erwin, P. Albuquerque, S. K. Jayaweera, and I. Hussein, "Dynamic sensor tasking for space situational awareness," in *Proc. American Control Conf.*, 30th June - 2nd July 2010, pp. 1153 – 1158.

[82] V. Fateev, S. Sukhanov, Z. Khutorovsky, S. Kamensky, A. Samotokhin, V. Stepaniants, and Y. V. Burtzev, "Analysis of collision prediction characteristics," in *Eighth US/Russian Space Surveillance Workshop Space Surveillance Detecting and Tracking Innovation*, 2010.

[83] C. Mingyan, X. Jianmin, Z. Wei, and W. Xuegang, "Study on simulation of signal processing system of space surveillance radar," in *2008 International Conference on Radar*, 2-5 Sep 2008, pp. 474–477.

[84] NVIDIA Corporation, "NVIDIA CUDA: Programming guide," 31st Oct 2011 2011.

[85] GPGPU.org, "General-purpose computation on graphics hardware," 2012.

[86] J. D. Owens, M. Houston, D. Luebke, S. Green, J. E. Stone, and J. C. Phillips, "GPU computing," *Proceedings of the IEEE*, vol. 96, no. 5, pp. 879 – 899, 2008.

[87] R. Vuduc and J. Choi, *A Brief History and Introduction to GPGPU*, X. Shi, V. Kindratenko, and C. Yang, Eds. Springer US, 2013.

[88] I. The MathWorks, "MATLAB GPU computing support for NVIDIA CUDA-enabled GPUs," 2012.

[89] L. Nyland, M. Harris, and J. Prins, "Fast n-body simulation with CUDA," *GPU Gems*, vol. 3, no. 1, pp. 677–696, 2007.

[90] N. Corporation, "Cuda c best practices guide," 31st Oct 2011 2011.

[91] J. E. Nelson, "Infrared methods for daylight acquisition of LEO satellites," Master's thesis, Air Force Institute of Tech Wright-Patterson AFB Ohio School of Engineering and Management, 2004.

[92] R. Sridharan and A. F. Pensa, "US space surveillance network capabilities," *Characteristics and Consequences of Space Debris and Near-Earth Objects*, pp. 88–100, 1998.

[93] J. Seago, J. Griesbach, J. Woodburn, and D. A. Vallado, "Sequential orbit-estimation with sparse tracking," in *Proc. AAS/AIAA Space Flight Mechanics Conf.*, 2011.

[94] W. Koch, "On negativeinformation in tracking and sensor data fusion: Discussion of selected examples," in *Proceedings of the Seventh International Conference on Information Fusion*, vol. 1. IEEE Publ. Piscataway, NJ, 2004, pp. 91–98.

[95] J. Hoffman, M. Spranger, D. Gohring, and M. Jungel, "Making use of what you don't see: Negative information in markov localization," in *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on.* IEEE, 2005, pp. 2947–2952.

[96] H. R. Richardson, L. D. Stone, W. R. Monach, and J. H. Discenza, "Early maritime applications of particle filtering," in *Proc. SPIE 5204, Signal and Data Processing of Small Targets*, 2003.

[97] T. Furukawa, F. Bourgault, B. Lavis, and H. F. Durrant-Whyte, "Recursive Bayesian Search-and-Tracking Using Coordinated UAVs for Lost Targets," in *Proc. IEEE International Conf. on Robotics and Automation*, 15-19 May 2006, pp. 2521 – 2526.

[98] F. Bourgault, T. Furukawa, and H. F. Durrant-Whyte, "Coordinated Decentralized Search for a Lost Target in a Bayesian World," in *Proc. 2003 IEEE/RSJ International Conf. on Intelligent Robots and Systems*, vol. 1, 27-31 Oct 2003, pp. 48 – 53.

[99] A. T. Kamal, C. Ding, A. A. Morye, J. A. Farrell, and A. K. Roy-Chowdhury, *An Overview of Distributed Tracking and Control in Camera Networks.* Springer Berlin Heidelberg, 2014, pp. 207–234.

[100] N. T. Pham, K. Leman, R. Chang, J. Zhang, and H. L. Wang, *Fusing Appearance and Spatio-temporal Features for Multiple Camera Tracking.* Springer International Publishing, 2014.

[101] J. Wan and L. Liu, "Distributed data association in smart camera networks using belief propagation," *ACM Transactions on Sensor Networks (TOSN)*, vol. 10, no. 2, 2014.

[102] N. Oudjane and C. Musso, "Progressive correction for regularized particle filters," *Proceedings of the Third International Conference on Information Fusion. IEEE.*, vol. 2, 2000.

[103] A. Doucet, S. Godsil, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.

[104] R. Douc and O. Capp, "Comparison of resampling schemes for particle filtering," in *4th International Symposium on Image and Signal Processing and Analysis. IEEE*, 2005, pp. 64–69.

[105] G. Kitagawa, "Monte Carlo filter and smoother for non-Gaussian nonlinear state space models," *Journal of Computational and Graphical Statistics*, vol. 5, no. 1, pp. 1–25, 1996.

[106] S. Godsil and T. Clapp, *Improvement Strategies for Monte Carlo Particle Filters*. A. Doucet, J. F. G. de Freitas and N. J. Gordon, Eds. Springer-Verlag: New York, 2001.

[107] J. T. Horwood and A. B. Poore, "Orbital state uncertainty realism," in *Advanced Maui Optical and Space Surveillance Technologies Conference*, Sep 2012.

[108] J. Ortega, F. Quintana, and G. Ybarra, "Mechanisms and pyros subsystem for the meteosat second generation," in *8th European Symposium on Space Mechanisms and Tribology*, vol. 438, 1999, p. 169.

[109] U.S. Coast Guard Navigation Center, "GPS NANUs, Almanacs, & OPS Advisories," 2014, [Online]. Available: http://www.navcen.uscg.gov/?pageName=gpsAlmanacs, Accessed: 22 Jan 2014.

[110] U.S. Naval Observatory, "Earth's seasons: Equinoxes, solstices, perihelion, and aphelion, 2000-2020," 2014, [Online]. Available: http://www.usno.navy.mil/USNO/astronomical-applications/data-services/earth-seasons, Accessed: 22 Jul 2014.