

Master of Science Thesis

Supervisor: Prof. Dr. Ir. P.N.A.M. Visser

Numerical optimisation of constellation and orbit design for the OLFAR radio interferometry swarm in orbit of the Earth-Moon L4 point.

Jurriaan Alexander van 't Hoff



Front cover image is a composite using images by:

- Moon: Muhammad Hamed / Reuters
- Earth: National Geographic
- Satellites: UNSW and S. Engelen
- Milky way: Felix Mittermeier

Abstract

In the field of radio astronomy the Ultra-Long Wavelength ($f < 30\text{MHz}$) band is unique, it has been studied for nearly a century yet it is still almost completely uncharted. Due to the reflective properties of Earth's ionosphere this band is nearly completely inaccessible for ground-based observatories, which can only receive frequencies down to 8 MHz in ideal conditions. This band may only be explored from space, but the creation of space based observatories is problematic due to the large telescope sizes associated with these frequencies. The only viable solution to space-based radio observation is the use of radio interferometry, a process which combines measurements of multiple radio receivers to act as a single instrument with a larger aperture. Plans to establish these constellations have existed for some time, but have always been too expensive. With the rise of micro-satellites the concept of a radio interferometry constellation has become economically viable. The OLFAR mission is one among the many recent radio interferometry concepts, which stands out among its peers through the application of a swarm design philosophy.

The most important instrumental property of a radio interferometer is the number of available instrument pairs (baselines), and the distribution of their relative orientation in uvw space. To facilitate radio interferometry at 10 MHz and lower frequencies the OLFAR swarm requires orbits which offer relative velocities below 1 m/s, while also being sufficiently stable to keep the baseline between its members below a maximum of 100 km. A minimal separation of 500 meters between satellites is used for collision safety. Early concepts for the OLFAR mission made use of Lunar orbits, where the Moon would act as a radiation shield against Earth's interference. Previous studies have shown that such orbits expose the swarm to unacceptably large relative velocities, which is why alternative deployment locations are still being studied. This thesis studies the applicability of swarm orbit designs around the fourth Lagrangian point, denoted as L4. This point offers very promising orbital properties, but it has not been studied in detail due to the swarm's exposure to interference. As a basis for this work it is assumed that the 9dB of interference might be worked around through the dynamic range of hardware and longer integration times, making L4-centric orbits a viable deployment location.

Particular attention is paid to developing an accurate numerical model for a perturbed orbital environment, which is necessary to provide long-term orbits of good quality. The inclusion of perturbations in this model is based on their maximum demonstrable effect on baselines and overall satellite positions over a year in orbit. After establishing the numerical simulation environment the satellite swarm design problem around L4 is posed as an optimisation problem for heuristic algorithms. Based on small-scale experiments the efficiency of different algorithms and architectures is evaluated, and the best-suited solution is used to optimise swarm designs for the OLFAR missions. The resulting method of choice maximises the potential of multi-threading, using 32 connected differential evolution algorithms with 48 population members to perform as a single algorithm with a population of 1536 individuals.

Using this method satellite swarm designs and orbits of up to 35 elements are found, which demonstrably meet all interferometry-related mission requirements for over a year in orbit while only relying on passive formation flight. These swarm designs rely on a process which is described as swarm folding to achieve this result. By distributing the swarm as a column mirrored in the barycentric z direction with uniform velocities, the swarm initially folds over itself. This folding motion is periodically repeated throughout the designed orbit after the initial fold as a result of the swarms natural orbits. The folding motion greatly enhanced long-term cohesion of the swarm, and it creates a very dynamic baseline distribution pattern for radio interferometry. It is demonstrated that these orbit designs have near-ideal baseline distributions, wherein they are only limited by the natural limitations of the Lunar orbital plane.

Though the application of swarm folding these designs remain compact for long time periods, even in a perturbed environment. The demonstrated designs have feasible mission lifetimes up to 3 years in-orbit, requiring only a few manoeuvres during this time to enforce the 500 meter separation for collision safety. Through the application of active formation control it is expected that this can be extended to 5 years in-orbit with the right swarm orbit design. It is also expected that using the methods in this thesis larger swarm designs might also be found, potentially ranging into 50 satellites. The largest hindrance to further extending the size of the swarm is the risk of near-collision events, which are inherent to the folding motion.

Preface

The work which lies before you represents the cumulation of close to 10 months of work, condensed into as few pages as I physically could. Starting with a literature study into a mission and topic I was wholly unfamiliar with at that time, it has been a very rewarding experience to carve my own path through a research topic of my interest. After such a long and arduous process as this thesis has been, writing the preface is proving to be one of the more difficult tasks. Prefaces I find to be the most daunting part of writing any reports, the content I envision for them changes by the hour and there is the ever-lasting fear of forgetting to properly allocate thanks to any associated party. The latter is particularly challenging since there are too many people to thank altogether.

First and foremost I'd like to thank my supervisor, Pieter Visser, for all his good advice and for reigning in my ambitions where necessary. Without the latter the scope of this work would surely span triple the pages, while only having half the content. During the difficult times of the pandemic I always enjoyed our weekly (distanced) meetings, which have helped tremendously with keeping my focus and maintaining a steady rate of progress. I'd also like to thank Raj Thilak Rajan, whom spent a great deal of time and energy reading through my theories in order to provide relevant feedback and insights into optimisation strategies. I'd also like to acknowledge that without technical support from Marc Naeije I would not have been able to achieve half the results, as the capabilities of the faculty's Eudoxos server have been vital to my work¹.

Outside of the direct academic environment I would also like to extend my thanks to all others in the field whom have been very generous and enthusiastic in sharing about their work and exchanging ideas: Mark J. Bentum (University of Twente/ASTRON), Steven Engelen (Hyperion Technologies), Alisa Nevinskaia (Hyperion Technologies), Sung-Hoon Mok (Technical University of Delft), and Chris Verhoeven (Technical University of Delft). Their receptiveness to questions and openness is most appreciated, and in many forms directly improved the work which is presented now.

Finally, a word of recognition to the efforts made to keep me sane throughout this process by my direct family, friends, and in particular my girlfriend, Charlotte. Without their unwavering support (and the necessary loads of caffeine and snacks) there is no doubt in my mind that this work would not be in the state in which it is today. I am grateful for all the support I've received throughout this process, and the patience that was given when I rambled on about these subjects.

Jurriaan van 't Hoff,
August 2020

Contact details;

Should the reader wish to contact me regarding the work in this thesis, for example to request a copy of the literature study or data files, feel free to email me at javanthoff@gmail.com

¹Should "yuxin" read this, I am sorry if me hogging resources have hindered you in any capacity.

Contents

1	Introduction	5
1.1	The history of radio astronomy	6
1.2	The rise of micro-satellites	10
1.3	Scientific interests within the ULW domain	11
1.4	Research Questions	12
1.5	Thesis Structure	12
2	The OLFAR mission concept	13
2.1	The swarm design philosophy	13
2.2	Technological challenges	14
2.3	Deployment of the OLFAR swarm	15
2.4	Alternative deployment solutions	17
2.5	Satellite design for the OLFAR mission	18
3	Synthetic Aperture Radio Interferometry	19
3.1	Basic principles of radio interferometry	19
3.2	Visibility measurements	21
3.2.1	The three-dimensional Cittert-Zernike equation	21
3.3	Imaging using radio interferometers	22
3.4	Influence of the PSF on imaging capabilities	23
3.5	The ideal point-spread function	24
3.6	Sample function resolution determination and scaling	25
3.6.1	Resolving the PSF at very low resolutions	25
3.6.2	A PSF-based cost function	26
3.7	Orbit requirements of the OLFAR mission	27
3.7.1	Collision avoidance	27
3.7.2	Diffusion-limited resolution	27
3.7.3	Maximum baseline rate	27
3.7.4	Confusion-limited sensitivity	27
4	Constellation design around the 4th libration point	28
4.1	The barycentric three-body problem	28
4.2	Simplified three-body problem	29
4.3	The Pseudo-potential	31
4.4	Surfaces of Hill	33
4.5	Motion around the L4 point	34
4.6	Designing around orbit decay	35
5	Numerical simulation environment	36
5.1	Used third-party libraries and tools	36
5.2	Important remarks	36
5.3	Acceptable model uncertainties	37
5.4	Environmental setup	37
5.4.1	Planetary modelling and positioning	37
5.4.2	Perturbation sources	37
5.4.3	Modelling solar radiation pressure	38
5.5	The necessity of modelling perturbations	39
5.6	Analysis of perturbing forces	40
5.6.1	Standard reference model	40
5.6.2	Influence of perturbation sources	41
5.6.3	Resonance of perturbations	44
5.6.4	Degree and order of harmonic gravitational models	45

5.7	Analysis of integration settings	46
5.7.1	Standard reference model	46
5.8	Final environment model	48
6	Optimisation strategy	49
6.1	Defining the optimisation problem	50
6.1.1	Cost function	50
6.1.2	Problem boundaries	52
6.1.3	Second-stage optimisation	53
6.2	Comparison of optimisation algorithms	53
6.2.1	Reference problem description	53
6.2.2	Single-algorithm performance	54
6.3	The archipelago method	55
6.4	Topology design	56
6.5	Scaling of the optimisation problem	57
7	Results	59
7.1	Search pattern for satellite swarm design	59
7.2	The effect of satellite body-pointing	60
7.3	Small scale swarm optimisation ($N \geq 15$)	62
7.3.1	160 day period swarm design:	62
7.3.2	Full year period swarm design:	63
7.4	20 element swarm design for one year in orbit	66
7.5	25 element swarm design for one year in orbit	69
7.6	35 element swarm design for one year in orbit	72
7.7	Comparison of swarm designs	75
7.8	Second-stage optimisation results	77
7.9	Collision detection and interpolation frequency	81
8	Conclusions	83
8.1	Filling the uvw space for radio interferometry	83
8.2	Longevity of passive formation flying swarms around the L4 point	84
8.3	Swarm orbit design around the L4 point	84
8.4	Swarm folding	85
8.5	The flexibility of swarm orbits	85
8.6	Conclusion	86
9	Recommendations	87
9.1	The application of body pointing for formation control	87
9.2	The use of the Point-Spread Function as cost function	87
9.3	Application of this optimisation method for low-lunar orbit	88
9.4	Investigating the use of multiple smaller swarms	88
9.5	Smarter methods of constellation design	88
9.5.1	Focussed solution spaces for optimisation	88
9.5.2	Solution-seeding	88
9.5.3	Staggered optimisation	89
9.5.4	Swarm design adjustments through linear approximation	89
9.5.5	Changes to the cost function	90
9.6	The relation between swarm size and orbit design	90
9.7	The applicability of the 5th Lagrangian point	90
9.8	Orbit insertion and de-orbiting	90
Appendices		95
A	Verification of developed functions	96
A.1	Verification of the Point-Spread Function computation	96
A.1.1	Comparison with the theoretical result	96
A.1.2	Simple sample function reconstruction	97
A.1.3	Complex sample function reconstruction	97
A.2	Verification of frame transformations	98
A.2.1	Relative frame transformation	98
A.2.2	Barycentric frame transformation	98
A.3	Verification of baseline extraction	99
A.4	Verification of two-dimensional projection	99

B J2000 - Earth-Centered Inertial frame	100
C Swarm design tables	101
C.1 15 and 20-satellite swarm design	101
C.2 25-satellite swarm design	102
C.3 35-satellite swarm design	102
C.4 35-satellite designs, second stage	103

1 Introduction

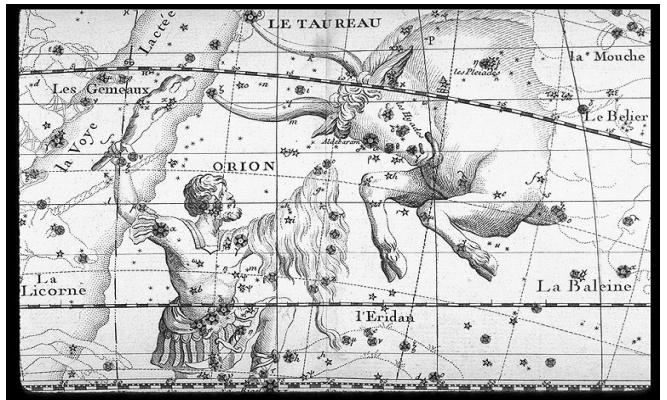
"Man must rise above Earth - to the top of the atmosphere and beyond - for only thus will he fully understand the world in which he lives"

-Socrates, Philosopher

For what we know of human history, the stars have always held the interest of humankind. It is impossible to tell exactly when humanity started looking up, but the influence of astronomy is present for as far as historical records go back. Many ancient religions and cultures incorporated astronomical elements into their mythology, such as the Egyptian god Ra whom was believed to be physically transporting the Sun across the sky during the day (Figure 1.1a)[31]. Gods in the Greek and Roman pantheons had the habit of creating stars to commemorate heroes or important developments in their respective mythologies ¹(Figure 1.1b). Even modern astronomy cannot escape its ancient association with religion and mythology, the planets in our solar system are named after the Roman pantheon, and the names of characters from ancient Greek mythology sire the moons of Jupiter and Saturn.



(a) Depiction of god Ra traversing the celestial sea in his boat, whom had to battle the snake Apep every night to ensure the sun could rise the next day [31].



(b) In Greek mythology Orion was a hunter, whom was immortalised among the stars after his death by Artemis, God of the hunt. Image: *Atlas Céleste de Flamberg* [26]

Figure 1.1: Depictions of mythological figures in Egyptian and Greek cultures associated with astrology.

Apart from being a cultural influence, astronomy has captivated scientists throughout history as well. The practice of scientific astronomy can be traced back to early pictographs from Mesopotamia in 3200-2100 BC, but the true extend of their knowledge is disputed to a lack of written records [52]. The first real written application of astronomy is commonly recognized as Homer's *Iliad* (8th century BC). Not only does the *Iliad* contain detailed descriptions of constellations, it also describes how the circumpolar rotation of the Bear may be used for nautical navigation [24]. Even without access to proper instruments or tools ancient cultures show to have impressive knowledge of astronomy. One of the earliest arguments for a Heliocentric solar system model was presented by Aristarchus of Samos in the third century BC, whom had arrived on this conclusion using nothing more than observations with measuring sticks[24].

Despite the impressive knowledge early astronomers were able to obtain with primitive tools, the field of astronomy stagnated until the 17th century. Aside from historical or religious aversion to astronomy, the capabilities of astronomers were primarily limited by a lacklustre observation platform. There is only so much the human eye can see, its optics are very limited, and its spectral coverage is relatively narrow. With the invention of the telescope in the 17th century astronomers could now see beyond the reach of the naked eye, which led to an explosive growth of the field. A similar breakthrough happened with the invention of electronic systems, which allowed astronomers to observe the skies beyond frequencies visible to the human eye. The latest advancement was the applications of satellites for astronomy, allowing astronomers to observe from places where no human could ordinarily go. Modern astronomy is now fully independent of human limitations, instead the platform capabilities are now entirely hindered by technological or economical boundaries. The latter is particularly true for the field of radio astronomy, which is in dire need of a new breakthrough to continue moving forward.

¹For those interested in this subject, [3] presents a concise overview of myths associated with constellations in Greek mythology.

1.1 The history of radio astronomy

In 1931 a radio engineer named Karl Jansky discovered a continuous background noise around 20.5 MHz of unknown origin in his measurements, which could not be explained with conventional knowledge, nor by a fault in his antenna setup. Through further observations he discovered that the source of the noise was directional, and that its relative position on the horizon changed over time in accordance with the rotation of the Earth. Jansky eventually concluded that the source of the noise had to lie outside of our solar system[32], a moment which is commonly seen as the birth of radio astronomy. It was not until 1961 however that the next significant advance in the field was made, when Rougoor and Oort created the first map of hydrogen distribution in the celestial sphere [54].

After the creation of this map the field of radio astronomy developed rapidly, alongside astronomy in other regions of the frequency spectrum². Not much later in 1967 the field had expanded to the point where observatories covered the entire frequency range accessible from Earth[27]. Further expansion was not possible using Earth-based observatories, due to the opaqueness of our Atmosphere at lower frequencies. Figure 1.2 shows the transmission values of Earth's atmosphere for a wide range of wavelengths. Note the radio frequency window is entirely opaque for wavelengths larger than 30 meters, due to radiative deflection by our ionosphere[33].

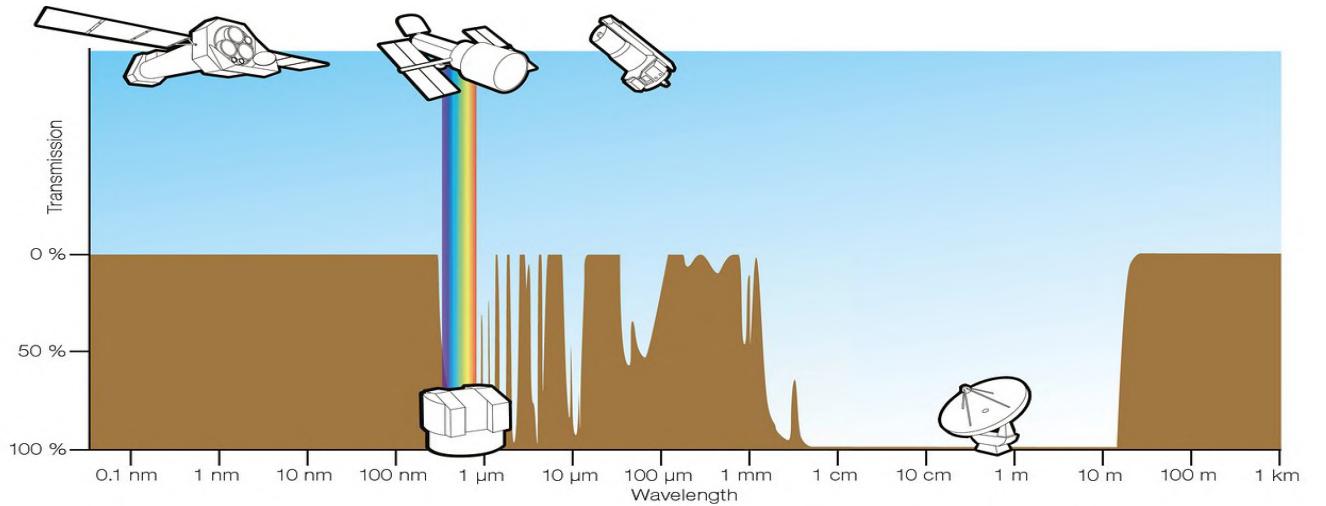


Figure 1.2: The transparancy of the atmosphere over a wide array of wavelengths. Figure: European Southern Observatory[47]

The pictographs in Figure 1.2 showcase observation platforms used to observe these wavelengths. For a large swath of the radio frequencies ($\lambda > 1\text{mm}$) observations are performed using ground stations, the size of which can span up to several hundred meters in diameter. The primary challenge of instrument design for radio astronomy is related to the wavelength which is to be observed. Consider the equation for the Rayleigh criterion, which is used to determine the angular resolution θ of instruments observing electromagnetic frequencies. The resolution of an instrument depends on the size of its aperture D compared to the wavelength λ :

$$\theta = \arcsin \left(1.22 \frac{\lambda}{D} \right) \quad (1.1)$$

In order to achieve a decent imaging resolution the antenna aperture is required to be at least an order of magnitude larger than the wavelength, which starts to become very challenging for lower parts of the radio spectrum.. The largest full-aperture telescope is currently the Chinese FAST³ radio telescope(Figure 1.3), which has an aperture of 500 meters in diameter[63]. Telescopes of this size are rare since they are very expensive to build and maintain, while also being very restricted in view angles. Due to their massive weight they are often "fixed" to the ground with very small movement capabilities.

Instead of using full-aperture telescopes the preferred approach to long-wave radio observation is the use of interferometric arrays. In an interferometric array the measurements of multiple smaller telescopes are combined to simulate a much larger radio telescope. The virtual instrument has a synthetic aperture equal to the distance between the used telescopes, a measure which is often called the baseline[21]. Through this technique two imaginary telescopes placed 1km apart with apertures of 100m can be used to simulate measurements with a 1km aperture telescope, although the observable wavelength is still limited by the 100m apertures. Interferometric arrays are often cheaper to build than full-aperture observatories, while also being more versatile since small telescopes are easier to aim.

²The radio frequency range is commonly defined as 20kHz to 300 GHz

³Five-hundred meter Aperture Spherical radio Telescope



Figure 1.3: Aerial photograph of the FAST radio telescope in Guizhou Province, China[63].

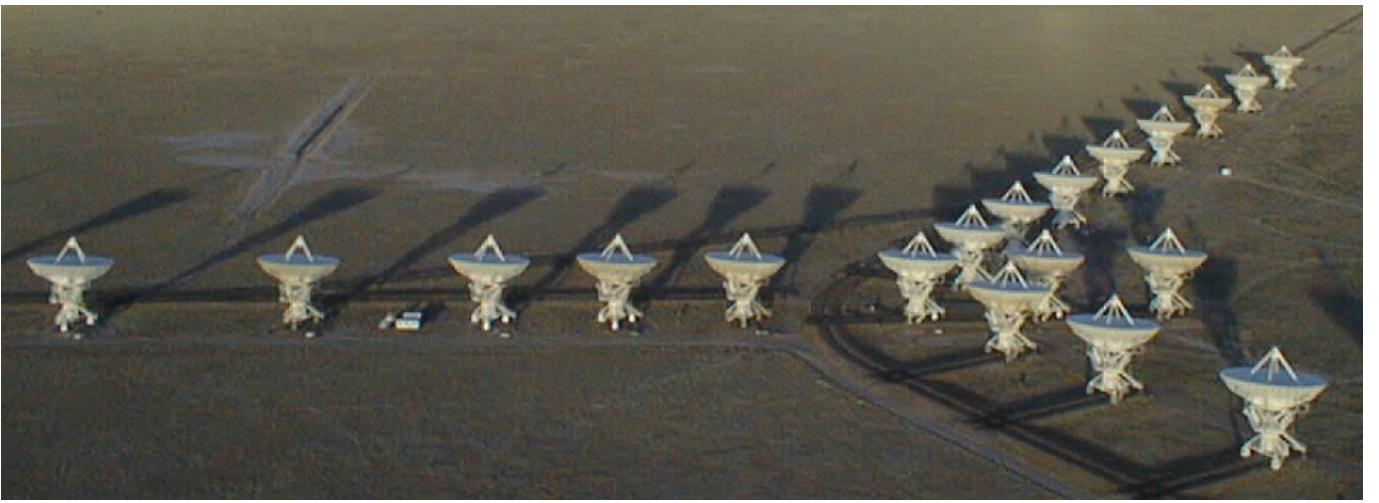


Figure 1.4: The Very Large Array photographed by Dave Finley, NROA.

The Very Large Array (VLA, Figure 1.4) in New Mexico is a well-known example of an interferometric radio array telescope. It uses a system of 27 different 25-meter radio telescopes placed on rails to create synthetic apertures of up to 36 kilometres with equivalent sensitivities of 130m large radar dishes[2]. The LOFAR radio telescope is the largest interferometric array currently in use in the world, it combines observatories spread across the Netherlands and Europe to create synthetic apertures up to 1500 kilometres[7]. Despite the impressive capabilities of the resulting imaging system, the LOFAR array is still limited by the transmissive properties of our ionosphere..

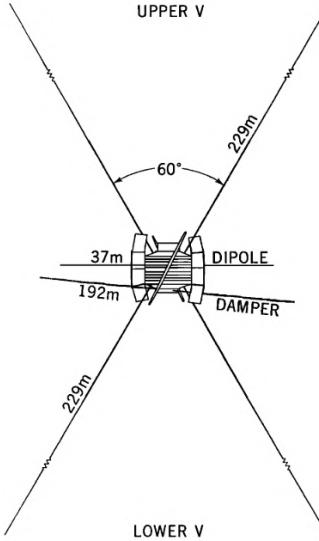
The 10m and larger wavelength band belongs to what will be described as the Ultra-Long Wavelength (ULW) domain ⁴. Below 10 MHz (30m) the ionosphere is entirely opaque, which makes this domain nearly⁵ inaccessible from Earth based observatories. Despite having lived centuries ago the words of Socrates are still very applicable in the modern day, in order to fully explore the ULW domain astronomers have to reach beyond the atmosphere. There is a definite need for space-based observations in the ULW band, which is easier said than done.

With wavelengths of 30 meters and larger, it is near-impossible to launch satellites with full aperture antennae. In order to observe these wavelengths scientists will have to make use of radio interferometry, combining multiple satellites to use as a singular interferometric instrument. The concept for a space-based radio interferometer constellation was proposed by French et al. as early as 1967 [27]. But this concept never left the drawing board, as it was recognized that the technology was not ready at that time to build such a constellation at any reasonable cost. French et al. concluded that their concept was ahead of their time, and the idea has never seen fruition until the 21st century. Only a few small-scale missions have been launched to explore the ULW band in history, and all of them only used singular satellites. The Radio Astronomy Explorer missions (RAE) by NASA were the first and the largest missions to explore the ULW band[61] using two⁶ independent satellites with large foldable antennae (See Figure 1.5).

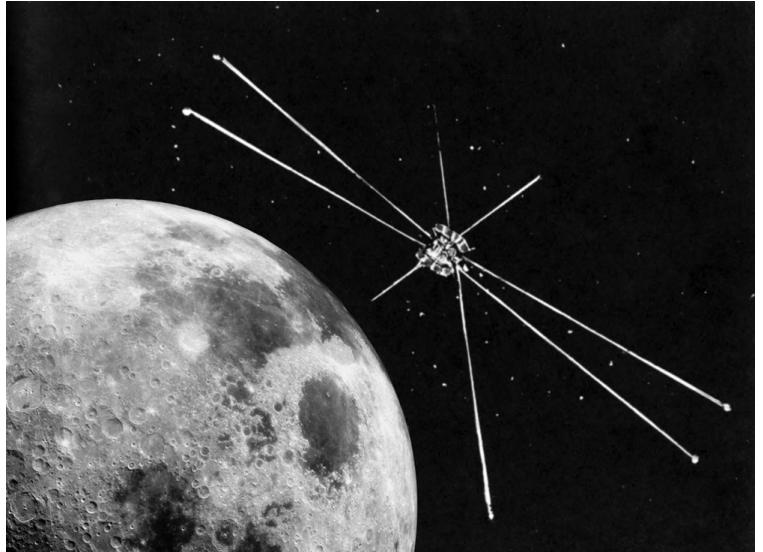
⁴A naming convention which is adopted from Heino and Falcke[33]

⁵Depending on weather conditions, altitude, solar exposure, and ionospheric density signals down to 5MHz might fall through.

⁶Yet still singular, as they were launched over a year apart.



(a) Antenna setup of RAE-1[61]



(b) Artist rendition of the RAE-2 satellite in Lunar orbit[42]

Figure 1.5: The Radio Astronomy Explorer missions were the first step in space-borne radio astronomy

Despite the limited scope the RAE missions led to some very important scientific discoveries, such as Auroral Kilometric Radiation(AKR) which originates from cyclotron⁷ mechanisms in the magnetosphere. [30]. AKR forced mission designers to change the designed orbit for RAE-2 to a Lunar orbit, as it made the instruments on RAE-1 virtually worthless from over-saturation. While this was a large setback at that time this discovery has allowed scientists to study magnetospheres of planets remotely, and it is expected that AKR can be used to search for exoplanets[8]. The switch to the Lunar orbit also lead to the discovery of the shielded region behind the Moon, which is a leading concept in many interferometric mission designs today. In 1977 Novaco and Brown[45] produced multiple maps of the celestial sphere using the ULW measurements from the RAE-2 satellites, such as the 9.18 MHz map shown in Figure 1.6. At this time these maps still represent the best maps astronomers have of this frequency domain to date, despite the resolution being inadequate to differentiate between individual point-sources [21].

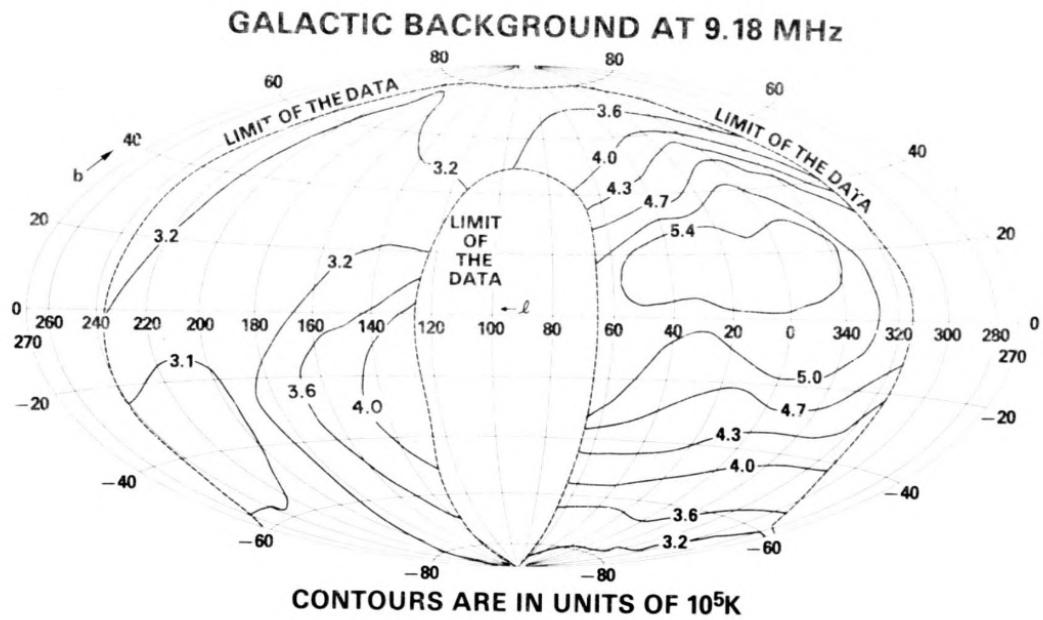


Figure 1.6: Map of galactic background noise constructed by Novaco and Brown[45].

Figure 1.7 shows an excellent visual comparison of the different spectral regions by E. Dekens[21], which highlights the stark contrast in resolution between maps of the ULW band and other spectral regions in 2012. As long as single satellites are used for observation the resolution of such maps cannot be improved for the ULW domain, creating a dire need for a radio interferometry constellation. Despite this need being identified as early as 1967 by French et al., the concept only recently started to gain real traction with the rise of micro-satellites.

⁷Radiation which is emitted when a charged relativistic particle is deflected by a magnetic field [18].

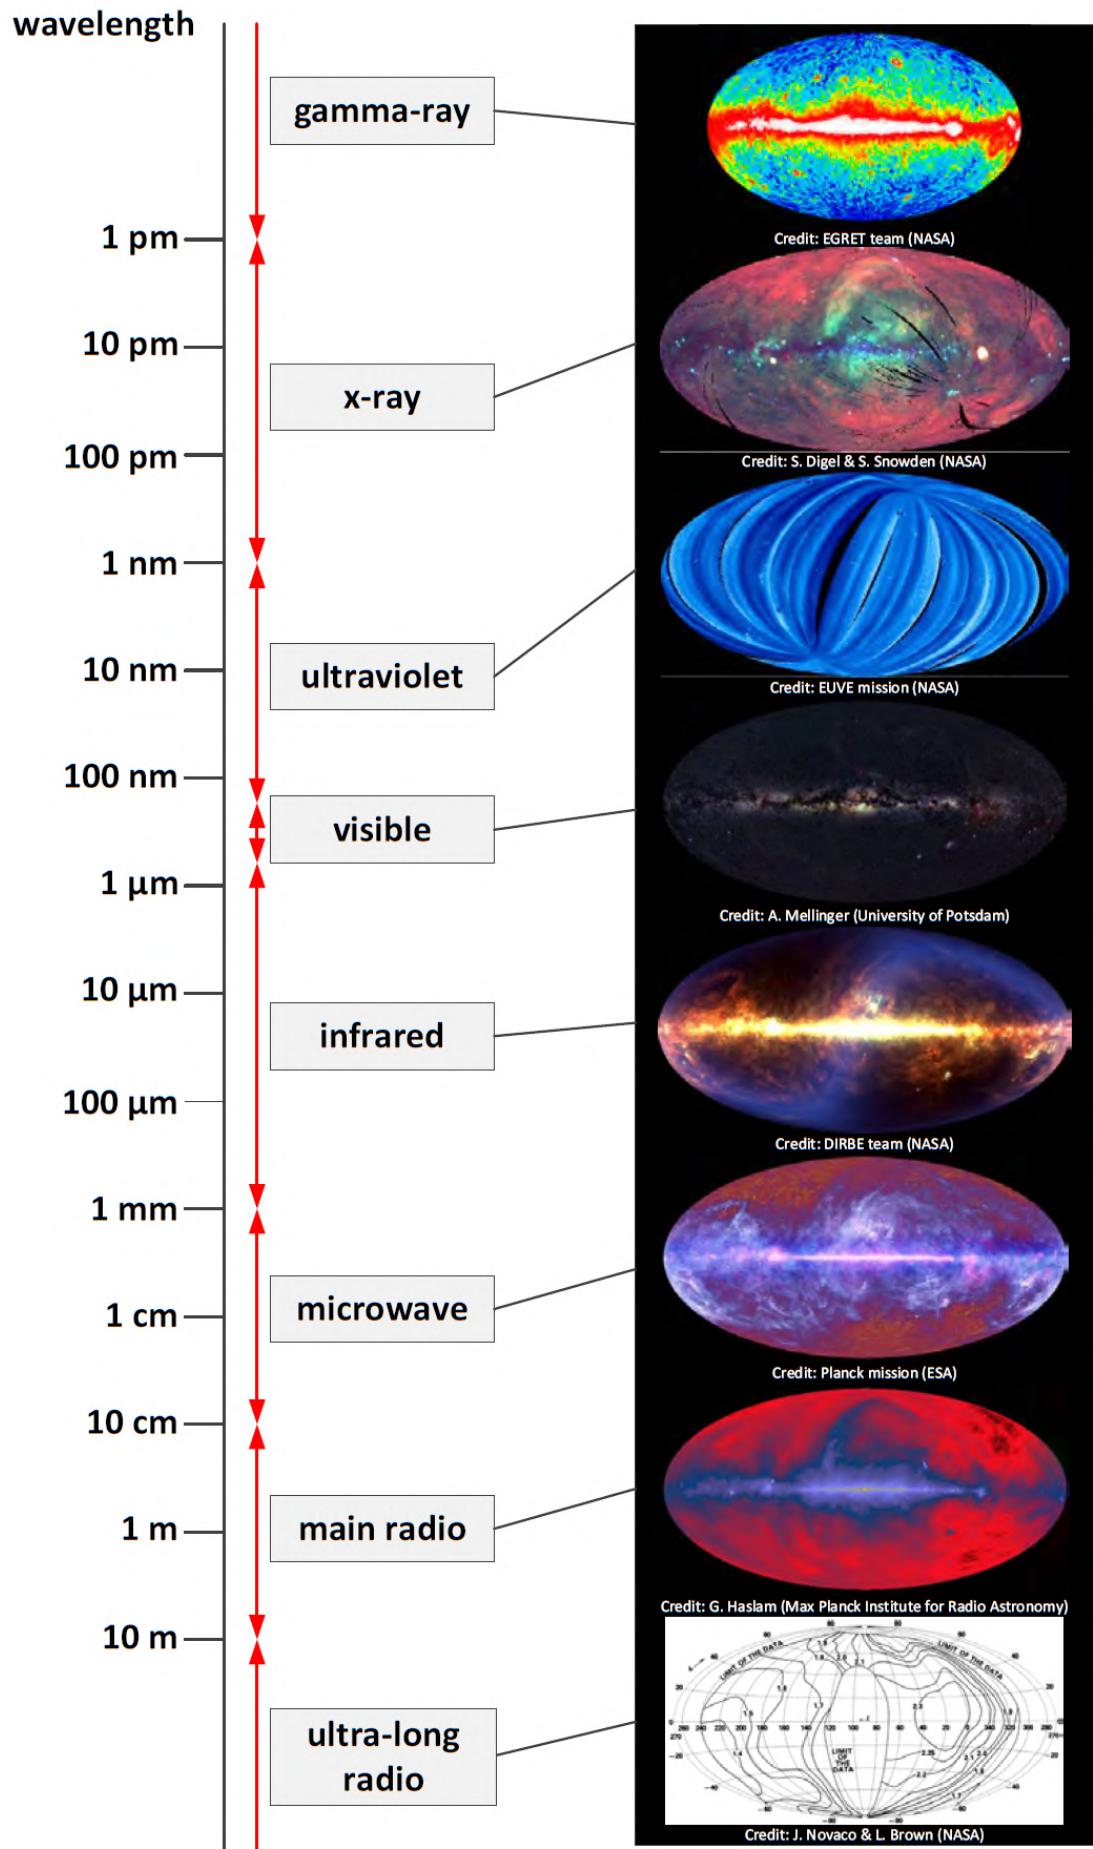


Figure 1.7: Comparison of the most current sky maps of different spectral regions in 2012, produced by E. Dekens[21]

1.2 The rise of micro-satellites

Despite having the interest of multiple fields of science, no large-scale radio interferometry constellation for the ULW region has been launched so far. The technological challenge of producing a large constellation of cooperating satellites at an affordable cost had continually proven to be a bridge too far. It is only in the past decade that astronomers have started to see a way to break through this economic barrier by using micro-satellite technology. The rapid development and widespread adoption of micro-satellites (see Figure 1.8), alongside the rapid growth of Commercial-Of-The-Shelf (COTS) components have led to a very economical option to deploy large numbers of reliable satellites.

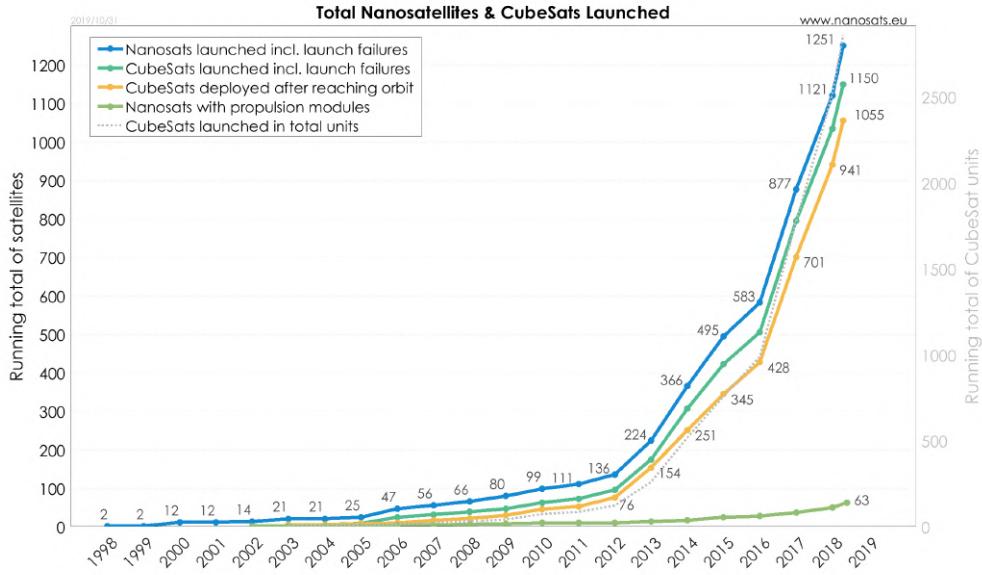


Figure 1.8: A chart documenting the number of cumulative nano-satellite missions launched over the past two decades, showing rapid growth in adaption of the technology. Chart by nanosats.eu

The adoption of these technologies has revitalised the concept originally envisioned by French et al., leading to a vast number of new mission concept studies for radio interferometer constellations such as OLFAR. Table 1.1 summarizes important details of most of the major concept studies which are closely related to the OLFAR concept.

Table 1.1: Overview of radio interferometer constellations concepts started in the last two decades, adapted from what was presented for the literature study [58].

Name	Origin	Year	Deployment	Satellites	Hierarchy
ALFA[34]	USA	2000	Earth retrograde orbit	16	Decentralized ¹
FIRST[57]	UK/ESA	2009	Earth-Moon L2 point	7	Centralized ²
DARIS [56]	The Netherlands	2010	-	9	Centralized
DARE[35]	NASA	2015	-	1	-
SURO[13]	Europe	2013	-	9	Centralized
PARIS[46]	USA/France	2005	Sun-tracing orbit	16	Centralized
NOIRE[18]	France/The Netherlands	2018	-	-	Decentralized
DSL[14]	Europe/China	2016	Low lunar orbit	9	Centralized
OLFAR[10]	The Netherlands	2009	-	>50	Decentralized

¹ All satellites in the constellation are identical and distribute workloads or tasks.

² Constellation has specialised members (e.g. "motherships") for different functions such as data correlation or communications.

Fields marked with - are either not discussed in the reference material, or still open for debate. A more detailed synopsis of these mission concepts is presented in the literature study[58]. It stands out that most mission concepts rely on a centralized hierarchy, using a few specialised "motherships" to handle data downlink and correlation. OLFAR, and other mission concepts with Dutch roots instead prefer to rely on a decentralized hierarchy, which uses large amounts of identical satellites and distributed workloads to achieve the same goals. The similarity between the OLFAR, DARIS and NOIRE concepts is not surprising considering the large overlap between these teams[58]. The work in this thesis will be based around orbit design for the OLFAR mission, which will be introduced in detail in Chapter 2.

1.3 Scientific interests within the ULW domain

Being nearly unobservable from Earth, the ULW domain garners significant interest simply because it is nearly completely unexplored. The domain covers a very broad part of the radio spectrum of which very little has been studied so far, this makes it a prime target for research even without pre-defined goals. Apart from the allure of the unknown, a more practical interest is the lack of existing spectral maps of the celestial sphere in this domain [50]. Mapping the celestial sphere is commonly listed as the primary goal of most proposals listed in Table 1.1.

Apart from mapping the celestial sphere, there are also more targeted scientific interests within the ULW domain. One of the other common goals which was identified for most of the mission concepts during the literature study[58] was the desire to study galactic history. It is expected that the ULW domain contains red-shifted signals emitted from Hydrogen during the period of re-ionization⁸ in galactic history (see Figure 1.9). These emissions are estimated to have red-shifted to approximately 10 MHz by Heino and Falcke [33]. Mapping these emissions would allow astronomers to look back at this crucial period in galactic history, making this a very valuable target for radio interferometers.

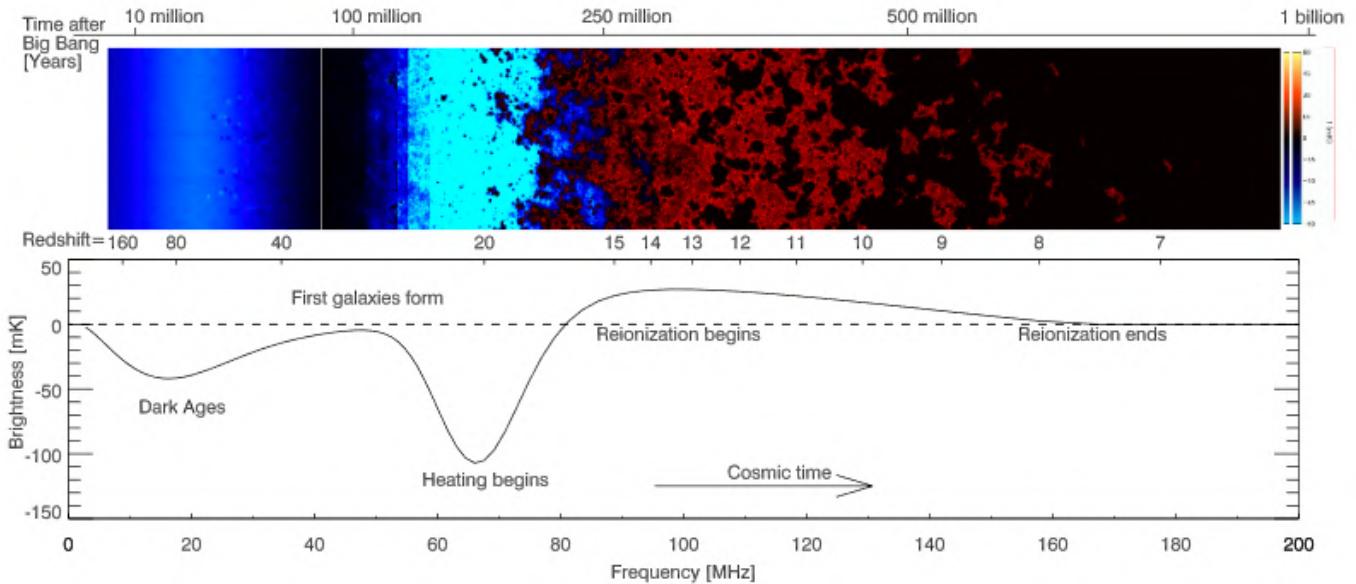


Figure 1.9: The simulated evolution of the 21 cm Hydrogen line over galactic history by Pritchard and Loeb[48]. Colors represent phases of radio absorption (blue), radio emission(red), and balanced absorption/emission ("neutral" / black). It may be observed that the period of re-ionization is a strong emitter of radiation.

In addition to mapping galactic history a promising prospect of ULW band astronomy is planetary exploration by studying ULW emission sources. For the NOIRE mission concept it is proposed that AKR emissions (see Figure 1.10) will prove useful to study the Earth's magnetosphere, as well as that of other bodies in the solar system. Lightning is another source of ULW radiation bursts, making an interferometer suitable to study weather patterns on Earth or other planets[18]. Apart from granting scientists new insights into bodies we know, these emissions can also be used to search for new exoplanets since they have very distinct profiles compared to background radiation [8],[33],[9].



Figure 1.10: A well known sibling of AKR is the Aurora Borealis, which is fundamentally created from the same source, albeit at a different wavelength. Photo courtesy of Holland America Line.

⁸Period when the first stars started re-ionizing hydrogen particles.

Finally, it is estimated by Heino and Falcke that a full map of the ULW band might yield the discovery of several million new radio galaxies [33]. Radio galaxies, such as the Hercules A galaxy pictured in Figure 1.11, are large structures of low-density plasma which can often be observed surrounding common galaxies. Radio galaxies emit strong radio signals due to cyclotron and synchotron⁹ interactions with charged particles travelling through space[21]. Most of the radiation emitted by radio galaxies falls within the ULW domain, making them prime observation targets.

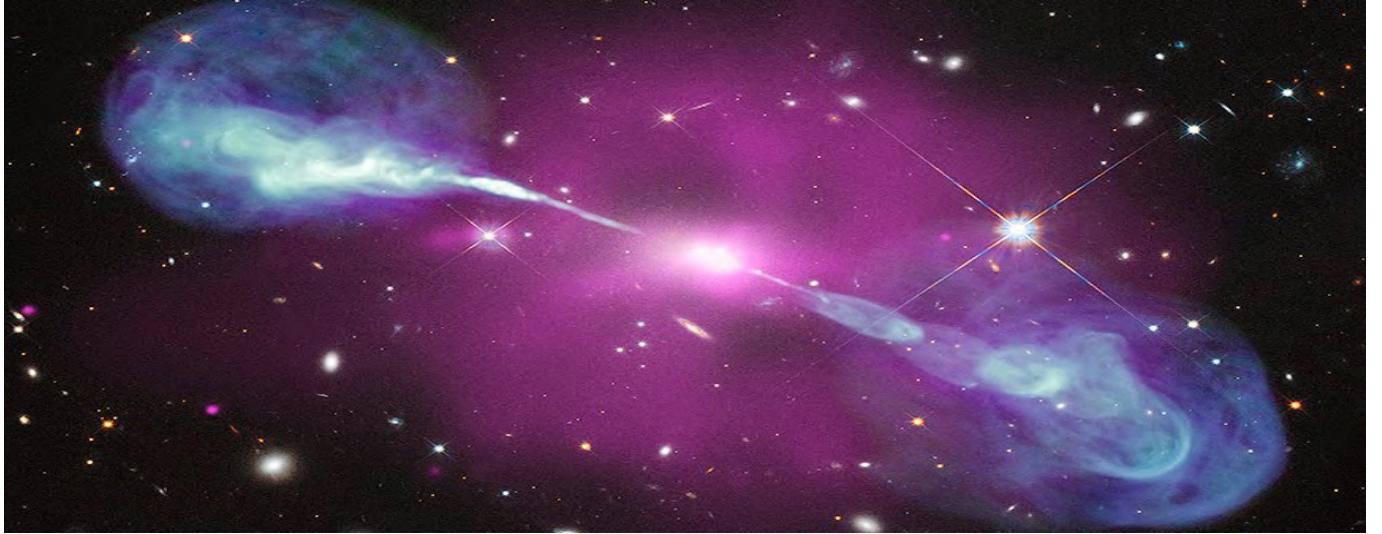


Figure 1.11: False-color image of the Hercules A radio galaxy, with blue showing radio emissions from large plasma jets. Pink shows X-ray emissions from a surrounding heated gas cloud [1].

In addition to all the goals mentioned in this section, there is still a plethora of smaller scientific goals to which the ULW domain might provide the answer. The list of goals presented in this section does not cover the full extend of possibilities, it is primarily meant to establish the primary goals for the OLFAR mission concept and to demonstrate the versatility of radio astronomy. In truth there is no telling how much discoveries proper exploration of this domain will yield, since most of this domain is entirely uncharted.

1.4 Research Questions

The work presented within this thesis, and the precursor literature study set out to answer a set of three primary research questions. These will be summarized briefly, without nested sub-questions, for the convenience of the reader:

1. *What are the relevant requirements for the orbital and constellation design of the OLFAR radio interferometry swarm?*
2. *How should a system be built for the simulation of L4-centric swarm orbits and the optimisation towards the application of radio interferometry of these swarms?*
3. *Can a radio interferometry swarm be designed which is passively stable for at least a year in orbit around the fourth Lagrangian point?*

The first research question was entirely answered within the literature study[58], the results of this study will be summarized throughout the initial chapters of this report. Chapter 5 and 6 treat the second research question, and the third¹⁰ research question is discussed in the final chapters.

1.5 Thesis Structure

This chapter served as a introductory chapter to the concept and history of radio interferometry. Chapter 2 will introduce the OLFAR mission concept, the underlying design philosophy, and its limitations. Chapter 3 introduces basic and advanced concepts of radio interferometry necessary to understand the work presented in this thesis. Chapter 4 treats general theories regarding motion of objects around the fourth Lagrangian point in order to provide insight in the expected results. Chapter 5 establishes the numerical environment and techniques to be used for satellite orbit propagation. Chapter 6 introduces the definition of the optimisation problem for swarm orbit design for radio interferometry, and studies different algorithms and architectures for their efficiency. The results of the optimisation process are presented and discussed in Chapter 7. Chapter 8 evaluates the suitability of L4-centric orbits for radio interferometry based on these results, and summarizes important conclusions. It is followed by a set of recommendations for future study in Chapter 9.

⁹Also known as Magnetobremsstrahlung; Radiation emmited when charged particles are accelerated radially.

¹⁰And arguably most important.

2 The OLFAR mission concept

OLFAR (Orbiting Low Frequency ARray) is a mission concept which aims to make use of a large amount of identical nano-satellites to build a large-scale radio interferometer swarm for the ULW band. The OLFAR concept was initially proposed in 2009 by a multi-disciplinary team stationed at ASTRON¹ and the Technical University of Delft [10]. The concept originally aimed to make use of a swarm of nano-satellites in orbit of the moon as a radio interferometer, utilising the moon as a radiation shield against Radio Frequency Interference (RFI) from Earth (see Figure 2.1). This deployment option would allow the swarm to make relatively noise-free observations, while remaining in close enough proximity to Earth to transmit its data. The OLFAR concept aims to create a scalable radio interferometer swarm, with a proposed size of up to thousands of satellites, wherein it surely is not limited in ambition.

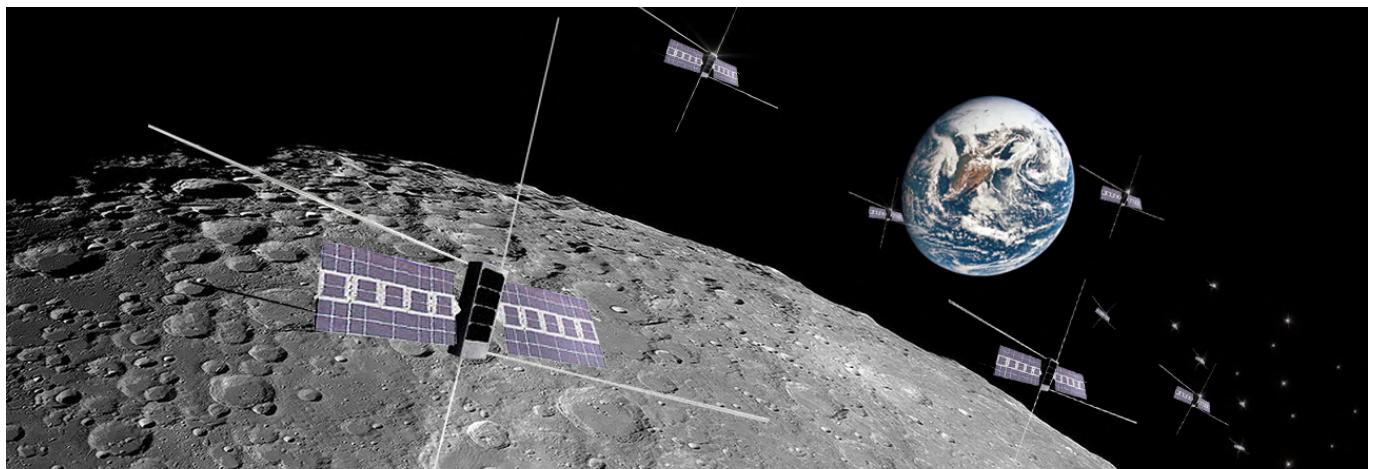


Figure 2.1: Artistic impression of the OLFAR swarm elements in orbit of the Moon[58].

The OLFAR swarm aims to cover the spectral region of 30 MHz and below, having some deliberate overlap with the LOFAR telescope system for mutual calibration and verification [23]. This would allow astronomers to achieve all of the scientific goals mentioned in section 1.3. In order to create a full-sky map with confusion-limited sensitivity it is estimated that the swarm needs a cumulative measurement time of 52,000 days[21]. By itself this requirement shows why the OLFAR swarm needs to be ambitious in scale. The root of these requirements, and how they affect the orbital design, will be discussed in section 3.7 after the necessary background theory is established.

The original mission proposal specified a launch date set in 2019[10], but this could not be achieved due to the significant technological challenges associated with the concept (see section 2.2). The OLFAR swarm aims to make use of the swarm design philosophy, in which the constellation (called a swarm) is made up of a large number of identical satellites. This philosophy has some very distinct strengths which make it well suited for radio interferometry, but it is not without its own challenges.

2.1 The swarm design philosophy

The OLFAR concept is set apart from contemporary radio interferometer proposals by the application of the swarm design philosophy. This philosophy is derived from observing the evolutionary success of swarm-like behaviour observed in nature [23]. Beings such as sardines, locusts, ants, or bees are not particularly intelligent, strong, or hardy when compared to predatory species. Yet they have thrived in nature due to the strength these species garner in numbers. Take for example the Japanese honey bee, utilises strength in numbers as a defence strategy against Asian Giant Hornets, a predator specialised in the eradication of entire bee hives. When Japanese honey bees are invaded by a hornet hundreds of them cluster around the invader, vibrating their bodies in unison in order to cook the attacker to death [26]. These bees demonstrate that even simple organisms can achieve great deeds through cooperation, which might also be observed from their intricate hives and the world-wide spread of their kin.

¹Netherlands institute for radio astronomy

The application of large numbers and cooperation to achieve a larger goal is the driving force behind the swarm design theory, a concept which is particularly well-suited to exploit the recent advances in COTS satellite technology. The idea behind the swarm design philosophy is not necessarily new, but it has only recently become economically viable after advances in COTS technology[23]. Compared to traditional² mission design philosophies the application of swarm design has some distinct advantages. A swarm is very resistant to single-point failures, very flexible in terms of scaling, and capable of producing much more scientific data. The latter is particularly important for the field of radio interferometry, in which the scientific output is primarily dependent on the number of instruments and their relative orientation.

This advantage is lost for scientific purposes in which qualitative measurements are easily achievable using single satellites, such as with direct optical imaging. In this case the imaging capabilities of a large satellite cannot be matched with a swarm approach. Likewise it must be assumed that swarm elements are less reliable than custom-made counterparts, making swarm constellations less suited for the observation of transient³ subjects[23]. Finally the lack of specialised units within a swarm makes it difficult to deal with technological extremes. One of the primary challenges with the OLFAR mission is the distribution of the data handling and downlink capabilities across the swarm, which might otherwise be easily handled through a single dedicated satellite. The inclusion of such an element however would make the swarm susceptible to single-point failures, negating the strengths of the concept.

2.2 Technological challenges

The technological challenges which are related to the delay of OLFAR's envisioned launch date are directly related to the weakness of swarm design, being rather extreme technological requirements for micro-satellite platforms. One of the primary challenges facing the swarm concept is processing power. With a swarm of the envisioned size a very large amount of data can be gathered simultaneously, a prospect which is problematic considering the limited storage capabilities of micro-satellite platforms. Equipping all swarm elements with dedicated large storage is very expensive, whereas the raw data is too big to send down at the same rate in which it is measured. The necessary alternative is to perform most of the data processing in-orbit, which significantly reduces the size of the data which needs to be stored and transmitted[50]. Designing system architectures which can distribute data processing tasks across a swarm is challenging, especially when this system also needs to be scalable to a wide range of swarm sizes[50],[11].

In addition to processing power the concept faces several other challenges, such as the communications systems. The small satellites cannot carry, or power, large-scale communications equipment, which presents challenges for both inter-satellite communications and ground station telemetry. The OLFAR swarm either needs significant advances in the capabilities of micro-satellites, or smarter communication methods in order to meet this challenge[11]. With such a large number of satellites flying in relatively close proximity one of the other major challenges is the design of autonomous formation flying systems. Having access to autonomous systems is a necessity, since the scope of the swarm would easily overwhelm the capabilities of ground stations. The size of the swarm also introduces problems regarding instrument calibration, which is challenging for a combined instrument of this proportion[11].



Figure 2.2: Roadmap for the OLFAR mission concept adapted from Bentum et al. (2019) [11].

In lieu of these challenges the new timeline proposal places the launch of the first OLFAR elements in 2030[11], which is shown in Figure 2.2. This new timeline gives ample time to find solutions for the technological challenges, and to seek an orbit design which would support a scalable swarm.

²The application of a few highly specialised satellites

³Sources which can only briefly be observed, such as lightning.

2.3 Deployment of the OLFAR swarm

During the first RAE mission it was discovered that the Earth itself is a strong emitter of ULW radiation, a part of this was accredited to AKR after its discovery [30]. AKR is not the exclusive source of RFI however, its contribution is relatively minor compared to spectral pollution from human activity[9]. Since ULW band signals are reflected by the ionosphere they are commonly used for long-distance communications or over-the-horizon radar systems. With every reflection approximately 4% of this energy leaks out into space. Nowadays these technologies are used all over the world, making mankind the largest producer of RFI in the ULW band [9].

In order to avoid Earth's spectral pollution radio interferometer proposals such as PARIS[46], FIRST[57], SURO[13], SULFRO[64], and DARIS[15] considered using heliocentric orbits, placing a constellation close to the Sun-Earth L4 or L5 point. This is not possible for the OLFAR mission concept, due to the limitations of communications for nano-satellite platforms sufficient telemetry cannot be established using the swarm philosophy. The OLFAR swarm needs to be placed in close proximity of the Earth, hence it aims to make use of the Moon as a radiation shield. Figure 2.3 shows data from the RAE-2 satellite in which it can be observed that the background noise drops between 10 and 30 dB when the satellite crossed behind the Lunar body[6].

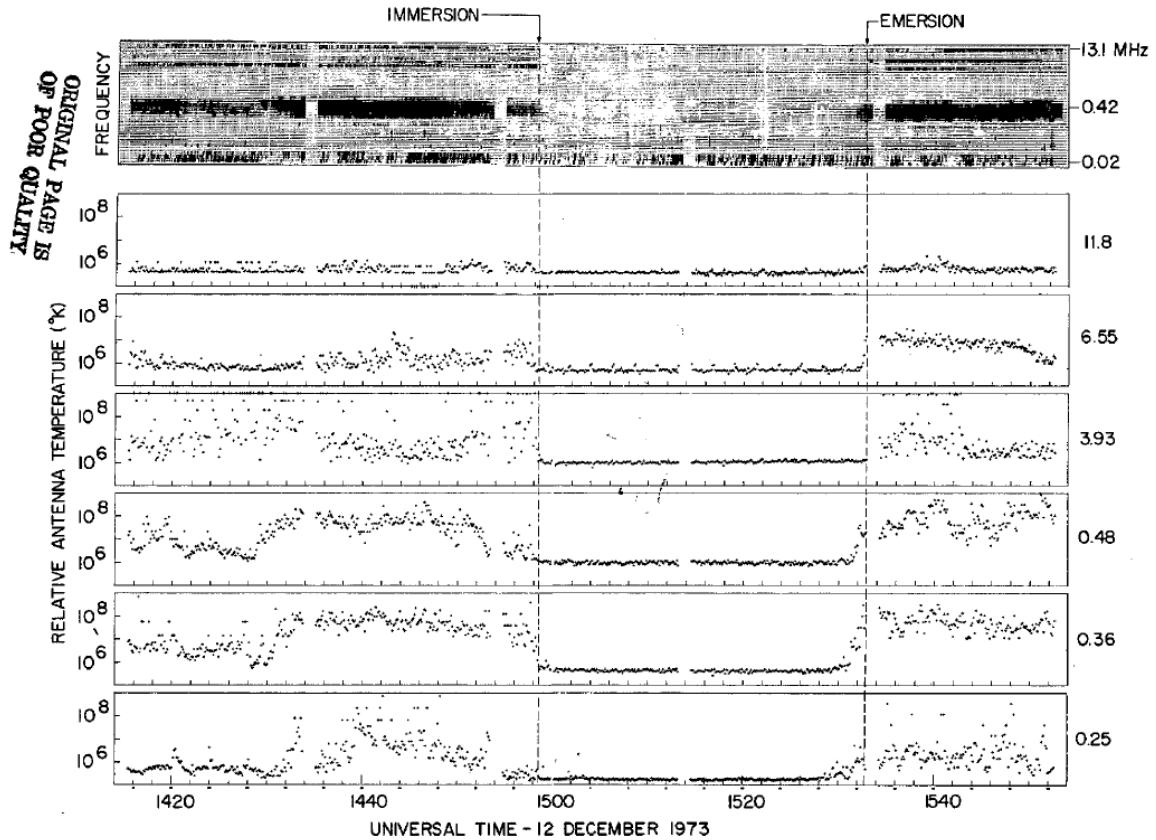


Figure 2.3: Instrument data from the RAE-2 satellite which orbited the Moon, a noticeable reduction in background noise can be seen during a period of Lunar occlusion. Figure by Alexander et al. [6].

The shielded area behind the moon is commonly referred to as the Lunar Quiet Zone (LQZ) , a concept first discussed by C. Maccone[37]. Figure 2.4 shows the geometric definition used by C. Maccone for the LQZ, derived from geostationary satellite orbits to shape the LQZ (which he names the Quiet Cone).

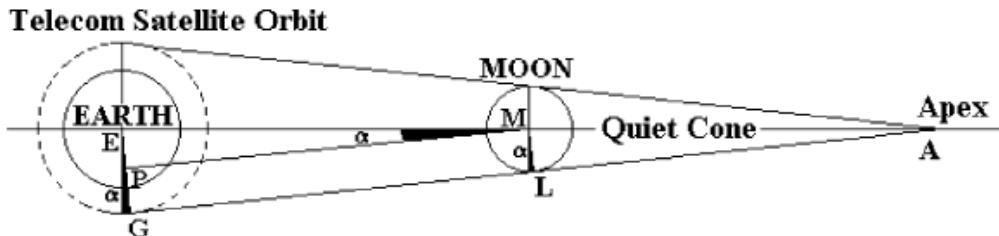


Figure 2.4: Mathematical definition of the LQZ by C. Maccone [37]

The shape definition for the LQZ by Maccone is built on the assumption that diffraction is negligible, which is valid when the wavelength is much smaller than the occluding body. For the purpose of ULW radio interferometry however, this assumption cannot be used in good faith. The exact extend of ULW diffraction and how it affects the shape of the LQZ is currently unknown, although it can be predicted using diffraction models. Figure 2.5 shows the threshold frequencies at different lunar radii behind the lunar body computed using a Fresnel diffraction model[58]. The threshold frequencies indicate the lower frequency boundaries beyond which background noise is larger than a given percentage of the original intensity on the far side of the moon.

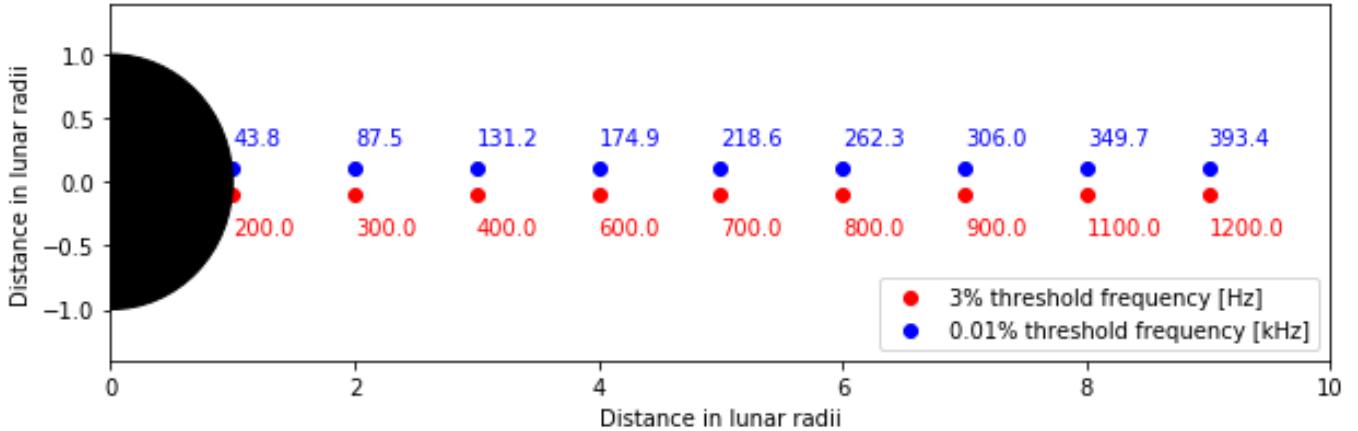


Figure 2.5: Threshold frequencies for 0.01% intensity fallthrough (blue) and 3% fallthrough (red) [58].

In Figure 2.5 it can be seen that at an orbital radius of 2 lunar radii around the moon frequencies below 87.5 kHz can be expected to have background noise stronger than 0.01% of their original intensity, and that on the Lunar surface itself some noise might be expected for frequencies below 43.8 kHz. For the higher tolerances it can be seen that even at the lunar surface 200 Hz and lower frequency signals diffract around the Moon. 200 Hz represents a boundary where the Lunar body is too small to successfully occlude lower frequency signals [23].

Regardless of the ambiguity surrounding the exact shape of the LQZ, it is known that the OLFAR swarm needs be as close to the Lunar surface as possible in order to maximise the measurement time in the occluded area [50],[23], [21]. At these low altitudes the irregular gravity field of the Moon poses a severe problem for OLFAR's swarm design. One of the main performance limitations of the OLFAR swarm is the capability to process data during measurements[11],[50]. This affects both the rate at which measurements can be taken and the required integration time, and as a result the swarm needs to have very low relative velocities between elements during measurements. Preferably these velocities are as low as possible, for practical reasons 1 m/s is commonly used as a mission design requirement[23]. In his thesis E. Dekens studied the possibilities of low-lunar orbits, during which he found baselines rates up to 117 m/s between swarm elements in the occlusion zone[21].

Despite the disheartening results of this study there is still confidence in the possibility of using a low-lunar orbit for the OLFAR swarm. Low-Lunar orbit designs are again being studied by Mok et al. [41]⁴ who have developed a different systematic approach to swarm orbit design. By introducing strict confines for the observation points and along-track offsets their approach minimizes baseline rates during observations, revitalising the option of a Lunar orbit. Mok et al. derived a system of three equations which tie the important orbit requirements into relative orbit design for swarm members, and these will be used to find optimal constellation designs for *uvw* coverage. Mok et al. are currently working on delta-v analysis and validation of this theory against high detail perturbation models, of which the results are expected to be published in October 2020⁵

A final consideration regarding the use of the LQZ for the OLFAR swarm comes from a more ethical perspective⁶. In 2008 C. Maccone proposed a worldwide treaty between astronomers to protect the LQZ from spectral pollution[38]. The LQZ is the last environment in close proximity of Earth which is not polluted by man-made emissions, and this area represents a unique opportunity for the establishment of shielded observatories on the dark side of the moon by exploiting existing craters. The treaty aims to preserve this opportunity by preventing any radio emissions originating from within the LQZ. This problem is also recognized by the International Telecommunications Union, which has recommended the designation of the LQZ as a radio quiet zone[19]. The OLFAR swarm concept is inherently reliant on inter-satellite communications for measurement cross-correlation, calibration, and relative navigation. This makes the OLFAR swarm a potential recurring pollutant to the LQZ, the impact of which should at least be considered.

⁴This paper should be published in the near future

⁵From personal communications with Sung-Hoon Mok regarding his work on Lunar orbits for OLFAR, August 2020.

⁶The author is in no position to make a statement on the ethics of spectral pollution and preservation, but feels that this topic deserves to be acknowledged.

2.4 Alternative deployment solutions

Despite the associated challenges the Lunar orbit is still being considered as the primary candidate for OLFAR's deployment, but research is also actively being done into alternative solutions. An example of such an alternative is the use of an Earth-Trailing heliocentric orbit using an adjusted OLFAR concept, which is currently being studied by Hyperion Technologies⁷. This option will require some changes to the core OLFAR concept to support data telemetry over long distances, but it shows promising orbital properties. Other near-Earth orbits are also being considered, which do not require changing the OLFAR concept. Most of these candidate locations have been studied in some capacity for OLFAR or similar mission concepts[58]⁸.

Out of these deployment candidates the fourth (L4) and fifth (L5) Lagrangian points have been studied the least, despite being the only stable Lagrangian points [60]. These deployment locations have often been dismissed as they fully expose the swarm to Earth's RFI, offering no option for shielding. Using interpolation between data from the WIND/WAVES instrument (Figure 2.6) it can be interpolated that the level of background noise at the altitude of the L4 point (385,000 km) is approximately 9 dB higher than the cosmic background at 10 MHz. While not optimal, this level of noise can be dealt with using instrumentation with high dynamic ranges and longer integration times.

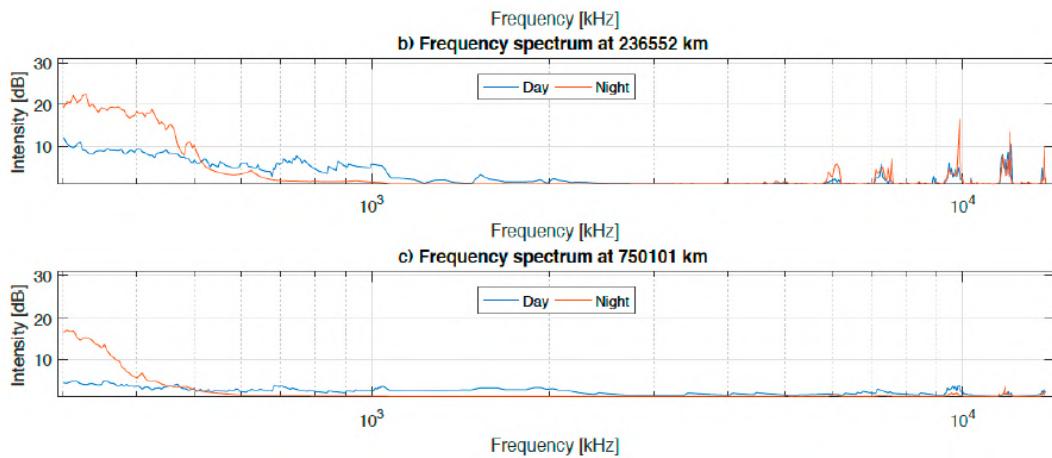


Figure 2.6: Intensity of RFI compared to cosmic background radiation at different altitudes, measured using the WIND/WAVES instrument. Figure by Bentum et al. [9].

If such a level of noise could be circumvented through systems design, the triangular Lagrangian points offer very promising orbital features for a radio interferometer. The stable dynamics of these points could allow for the creation of cohesive swarms with long lifetimes, small relative velocities, and a wide range of relative motion[58]. The stability of these orbits would be immensely beneficial to the design of autonomous formation flight algorithms, as the relative motion is both slower and easier to predict. It also means that the swarm would need to spend much less fuel to maintain its orbit for longer mission lifetimes than it would need in low Lunar orbits.

The second advantage of triangular Lagrangian orbits over Lunar orbits is the option of continuous observation. The Lunar orbit concept can only take measurements during a small fraction of its orbit, since it relies on occlusion by the Lunar body[21]. By adapting the system to deal with RFI, it might be used nearly full-time while in orbit, depending on how fast data can be processed and downlinked. This also allows for a different swarm observation mode which has different satellite groups observing or downlinking data simultaneously, which can allow for continuous coverage for the observation of transient events.

Recognizing the potential of orbits around the triangular points, as well as the limited knowledge from previous studies into these locations, it is chosen to focus this thesis on orbit design around the fourth Lagrangian point. The choice for the fourth Lagrangian point was made primarily out of convenience, since this point is better known than its sister. In the end the fifth Lagrangian point is an equally interesting topic of study, and many of the conclusions from this work might equally apply to L5-centric orbit design.

⁷From personal communications with Steven Engelen and Alisa Nevinskaia of Hyperion Technologies, January 2020.

⁸An extensive overview of previous studies, and studies for related missions, is given in the literature study [58]

2.5 Satellite design for the OLFAR mission

At this time very little is set in stone regarding the satellite design for the OLFAR mission, a conclusion needs to be made using what is presented during conceptual designs. For the purpose of this work the exact details of satellite designs are not necessary, but some general knowledge is required to accurately model the orbit of the satellites. The OLFAR mission aims to make use of a swarm of micro-satellites based of small CubeSats, equipped with at least 3 sets of orthogonal monopole antennas with lengths of at least 5 meters. This creates a very distinct satellite profile which is visualised by M.J. Bentum in the render shown in Figure 2.7.

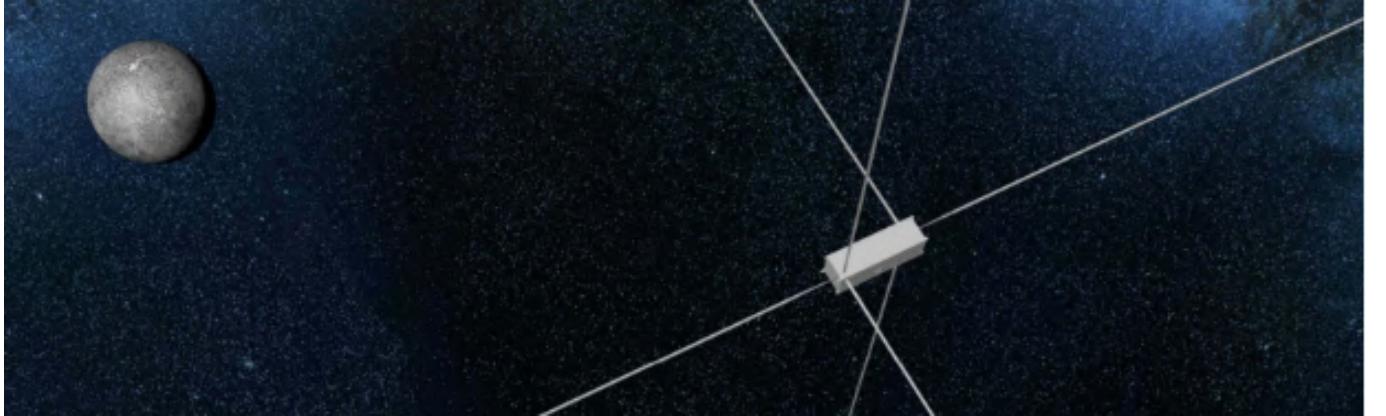
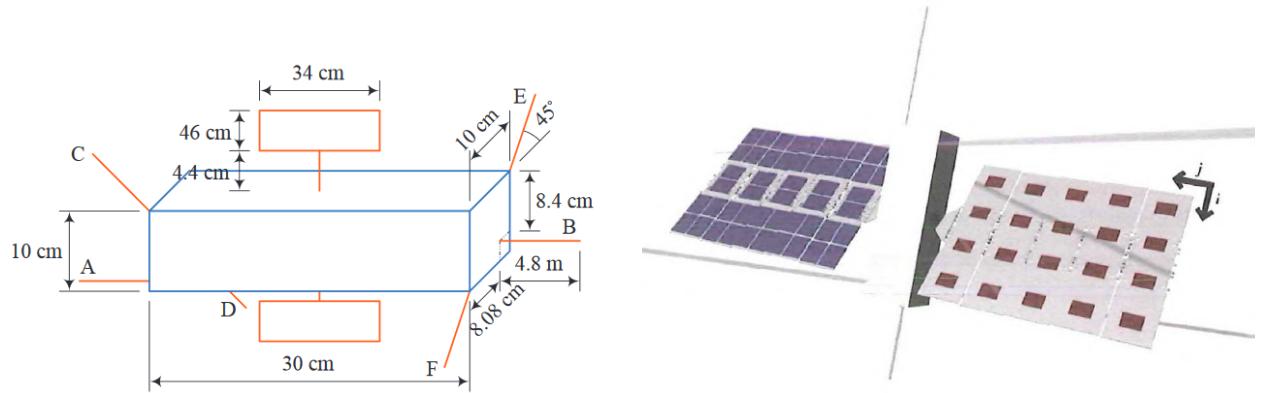


Figure 2.7: A render of a single OLFAR element in orbit of the moon, pictured without its solar panels. Original by M.J. Bentum, retrieved from [59].

The base body shape of such a satellite is assumed to be equivalent to at least a 3U CubeSat, which is tied to the minimal space requirement to support the electronics associated with the monopole antennas [49]. Another distinct feature of the satellites are two sets of very large solar panels. In order to provide sufficient power for inter-satellite communications, processing, and data downlink to Earth it is expected that each element needs at least 0.31 m^2 of solar panels. The schematics (Figure 2.8a) by Quillen et al. propose the use of two collapsible solar panels of 46 by 32 centimeters alongside the cubesat body[59]. Likewise very little is known about the mass of these elements, estimates range between 5 and 10[21] kilograms. This paper will use the lower end mass of 5kg to model the satellites, since these represent the worst-case scenario for the influence of perturbations.



(a) Schematic for an OLFAR satellite presented by Quillen et al.[49]. AB, CD and EF represent the antenna pairs.

(b) Render of an OLFAR satellite by S. Engelen [23] with its solar panels.

Figure 2.8: Different representations of the current OLFAR satellite concept.

The schematics shown in Figure 2.8a will be used as a basis to model OLFAR satellites for orbit propagation, using 5kg for the estimated mass. subsection 5.4.3 will discuss how surface properties are modelled for the purpose of Solar radiation pressure modelling.

In order to adequately discuss the OLFAR mission in more detail some general knowledge about radio interferometry is required. The next chapter will go through the relevant theory and terminology, before the mission requirements for OLFAR's orbital design are discussed.

3 Synthetic Aperture Radio Interferometry

This chapter will introduce the concept of synthetic aperture radio interferometry and discuss the related theoretical background for the work presented in this thesis. The content of this chapter is written under the assumption that the reader already has some understanding of the field of radio astronomy, and/or radio theory.

3.1 Basic principles of radio interferometry

As the name suggests radio interferometry makes use of the interference between two measurements from different points in space or time to extract additional information about the subject. This section will predominantly treat spatial interferometry, since this is the technique used by the OLFAR mission concept. Consider the situation sketched in Figure 3.1 to demonstrate the core idea of spatial interferometry:

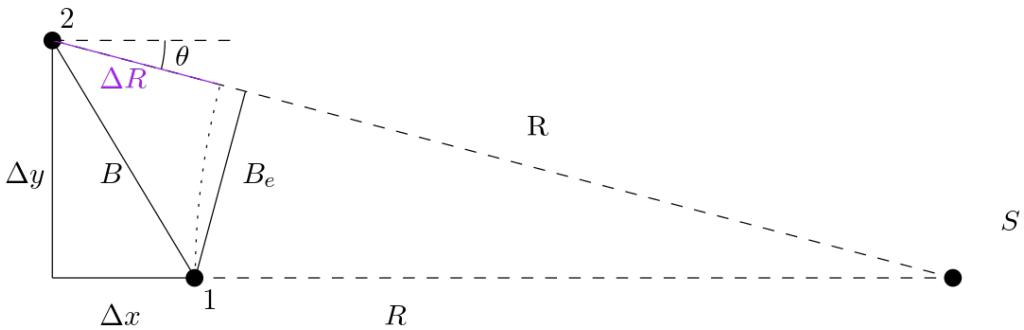


Figure 3.1: A geometric example of synthetic aperture radio interferometry, adapted from the literature study[58].

For these two satellites 1 and 2 with known relative position vector B , a subject S is observed at unknown distance R . Considering the distance associated with astronomical sources, it can be assumed that the received radiation is approximately coherent and planar. Using this assumption, the vector change ΔR can be estimated as a function of the phase difference between observations. If the phase difference is known the unknown distance R may then be solved using simple trigonometry. The relative position vector B is commonly referred to as the *baseline* for radio interferometry, and the projected vector B_e is the *effective baseline* towards the observation of a subject.

The imaging capabilities of an interferometer are predominantly determined by its collection of baselines, which span between every satellite within the constellation. It is important to realise that a baseline can be used for an interferometry measurement in two ways, and a single baseline should always be considered along with its negative pair. This yields the following relation between the number of baselines in a constellation and the number of satellites N [50]:

$$N_{\text{baseline}} = 0.5N(N - 1) \quad (3.1)$$

To represent the distribution of baselines the so-called *uvw space* is most commonly used, in this space baselines are represented as vectors originating from a common origin point. When mapping out *uvw* sample points it should be known that a baseline can be interpreted in two ways (S_1S_2, S_2S_1), which results in a mirrored copy of the *uvw* coordinate of the baseline. Such a reflection will be associated as a "negative" pairing of a sample point. The number of total sample points will thus equal twice the number of baselines[28]. The axes of the *uvw* space are commonly normalized to units of wavelengths, which makes it easy to correlate baselines with imaging capabilities. This standard will not be used in this report however, because OLFAR aims to measure a very wide frequency range¹. Figure 3.2 shows an example of the *uvw* baseline distribution for a set of 5 arbitrarily distributed satellites in a three-dimensional space.

¹Besides, the use of kilometres is more convenient for discussion later on in the results.

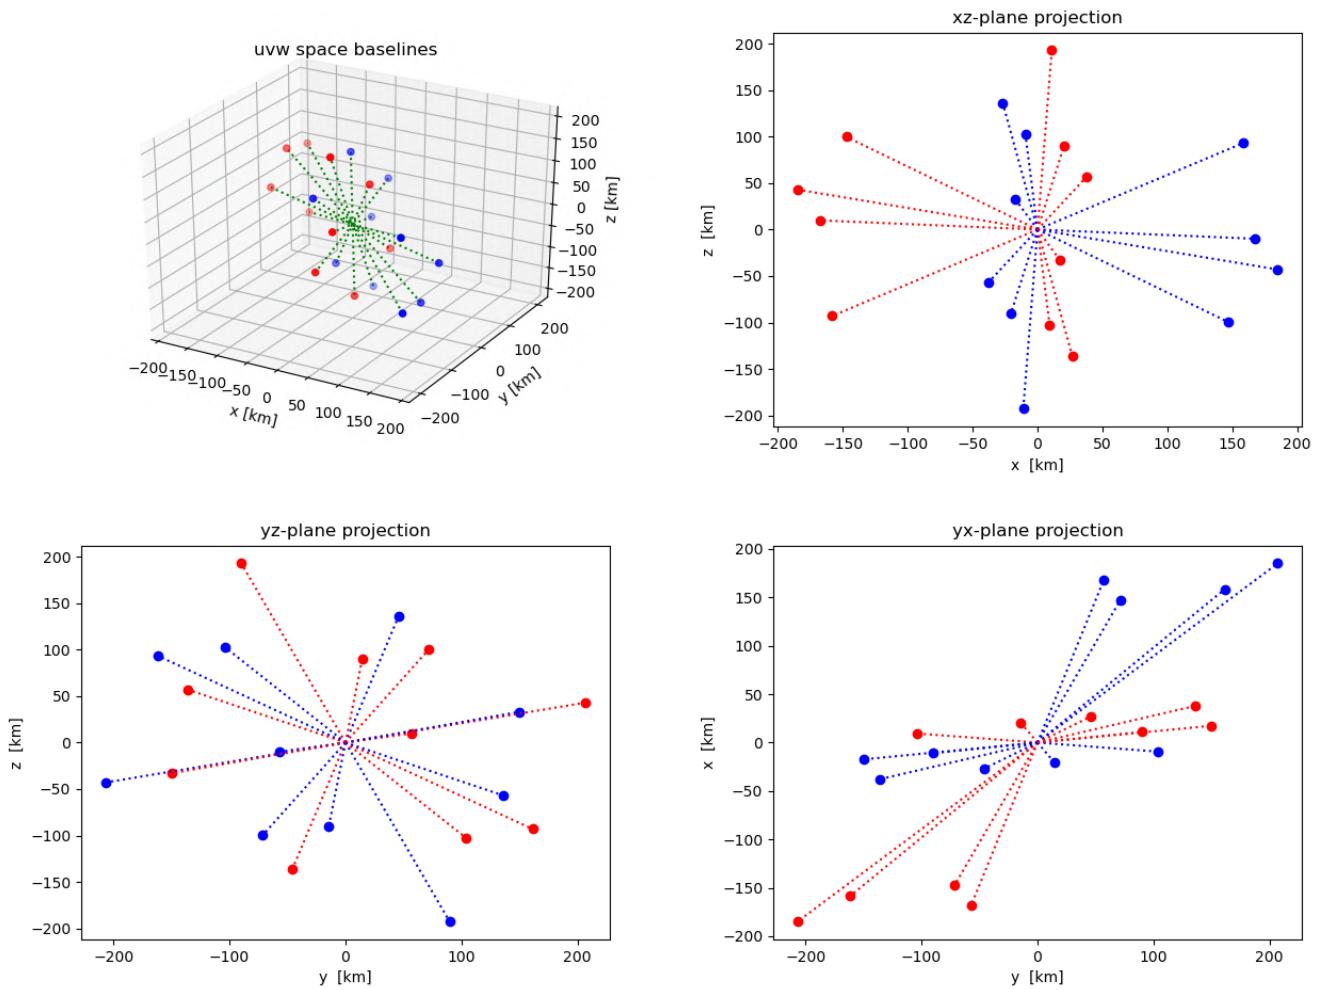
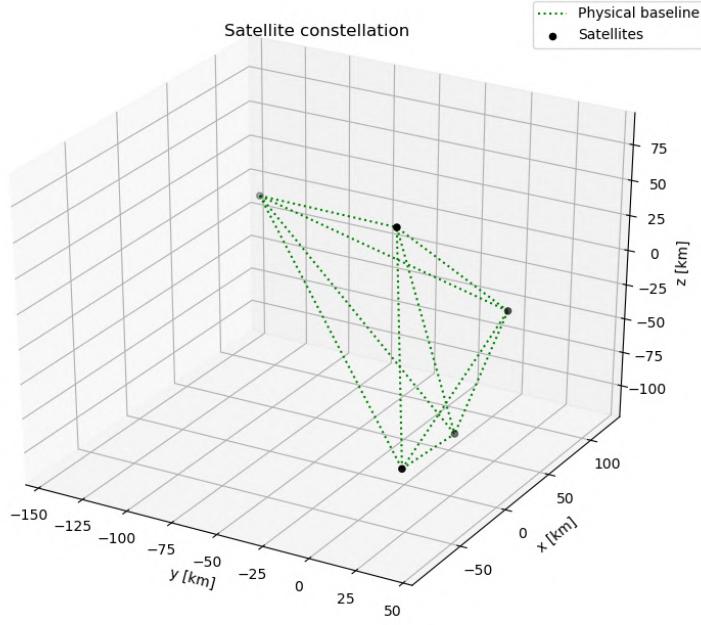


Figure 3.2: A visualisation of a set of 5 random satellites, their respective baselines, and the resulting baseline distribution in uvw space [58]. Blue marking represent the baselines while red presents the associated negative pairing.

3.2 Visibility measurements

In order to understand the importance of a proper baseline distribution for a radio interferometer one must first understand the nature of an interferometer measurement. The OLFAR mission aims to make use of an interferometry technique called signal cross-correlation, a technique which multiplies two incoming signals before dividing by a time-average [21]. This results in either constructive or destructive interference between the two signals, and this output is commonly referred to as the visibility measurement. This measurement is described by the Cittert-Zernike equation, which describes the cross-correlation measurement between two observers on a plane perpendicular to the source vector[20]:

$$V(u, v) = \int \int I_{sky}(l, m) e^{-2\pi i(u l + v m)} dldm \quad (3.2)$$

Variable	Description	Unit
$I_{sky}(l, m)$	Intensity image of the sky as a function of angles l and m	K
l, m	Direction cosines, 2 orthogonal angles from the target vector	$mJy/beam$

In Equation 3.2 it can be seen that the visibility measurement V is actually a sample of the Fourier transform of the observed sky image I_{sky} , in which the sample coordinates u, v are determined by the effective baseline between both observers. With relation to satellite interferometry these coordinates directly represent the baselines between two satellites, projected on a plane perpendicular to the image direction. Figure 3.3 shows a three-dimensional example of two satellites sampling an arbitrary location in the sky.

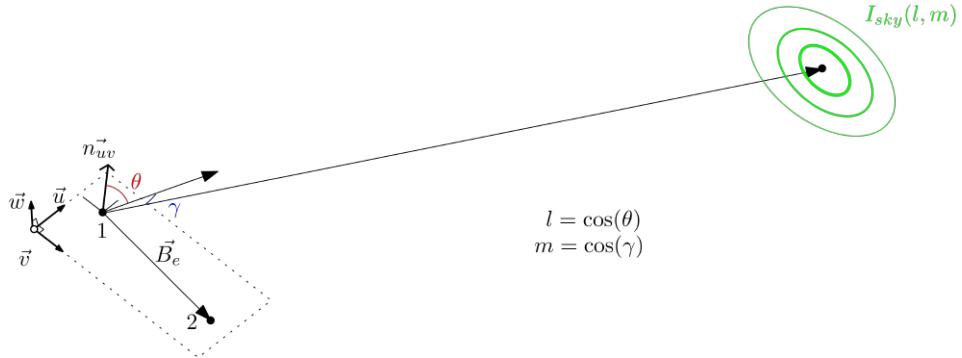


Figure 3.3: Geometry of two coplanar observers 1,2 on the uv plane, observing a target with local intensity distribution I_{sky} as a function of cosine coordinates l and m [58].

Due to the nature of this technique both the baseline magnitude and orientation relative to the target affect the measuring capabilities of the system. Since the visibility measurements by a satellite interferometer pair represent a spatial sample of the Fourier transform of the observation target, it becomes imperative that a wide array of baselines is used to sample a single target. As seen in Equation 3.2 the spatial coordinates of the baseline also directly influence the frequency range of the sample, meaning that large baselines are required to measure elements with small spatial frequencies. Likewise small baselines are required to measure elements with large spatial frequencies, and a mixture of the two is required to create a clear image of the subject. Since the sample is also directional in a two-dimensional plane, there is also the requirement to ensure the samples cover the entire angular range[21].

3.2.1 The three-dimensional Cittert-Zernike equation

For a spatially distributed interferometer array in orbit it cannot be safely assumed that the receivers are all coplanar, nor will the displacements be negligible compared to the wavelengths. In practice the three-dimensional variant of the Cittert-Zernike equation would be applied in-orbit[20]:

$$V(u, v, w) = \int \int \frac{I_{sky}(l, m)}{\sqrt{1 - l^2 - m^2}} e^{i2\pi(u l + v m + w m)} dldm \quad (3.3)$$

Solving the three-dimensional variant requires significant computational power, which is one of the primary bottlenecks in the design of any interferometry constellation. This is one of the primary reasons why most of the mission concepts presented in Table 1.1 rely on centralized hierarchies to handle the computational requirements. For the purpose of the work in this thesis it will be opted to rely plane projections combined with the two-dimensional Zittert-Zernike equation to visualise point spread functions and interferometer performance. This variant is both much faster to use and produces results which are easier to interpret when presented in a report.

3.3 Imaging using radio interferometers

By itself, a visibility measurement yields only a sample of a small region of the Fourier transform of the sky image. In order to construct an image, the set of visibility measurements needs to undergo a reverse Fourier transform (FT^{-1}) after projecting them onto the orthogonal observation plane. Equation 3.4 describes this process:

$$I_{\text{dirty}}(l, m) = F^{-1}(V(u, v)) = \int \int V(u, v) e^{i2\pi(ul+vm)} du dv \quad (3.4)$$

The image constructed during this process is commonly referred to as being "dirty". When only a part of the entire Fourier spectrum has been sampled, the reconstructed image will appear smudged or distorted as a result (see Figure 3.6). Mathematically this process can be described as a convolution between the original image and the Point-Spread Function (PSF) of the interferometer.

$$I_{\text{dirty}} = PSF(l, m) \times I_{\text{sky}}(l, m) \quad (3.5)$$

Simply put the PSF describes the response of an imaging system to impulses, comparable with point-sources in the case of astronomy. To properly use an imaging system it is thus vital to have knowledge of its PSF. The PSF of a radio interferometer entirely relies on its sample function. The sample function is defined as a binary function which equals to 1 where visibility measurements were taken[17]. The PSF can be acquired by taking the inverse 2-dimensional Fourier transform of this sample function on a projected uv plane $S(u, v)$, as shown in Equation 3.6.

$$PSF(x, y) = FT^{-1}[S(u, v)] \quad (3.6)$$

The sample function of a space-borne interferometer array consists of the collection of baselines used for the measurement, projected onto a uv plane. For the purpose of cross-correlation observations temporally different visibility measurements may be combined, allowing for the use of the collection of baselines apparent during orbit [50]². It is important to realise that the sample function is dependent on a projection of the baselines onto a plane perpendicular to the observation direction. This means that the shape of the PSF is directional, and that the observational capabilities of a collection of baselines will differ directionally. Figure 3.4 demonstrates this principle, showing three entirely different sample functions from a single set of baselines.

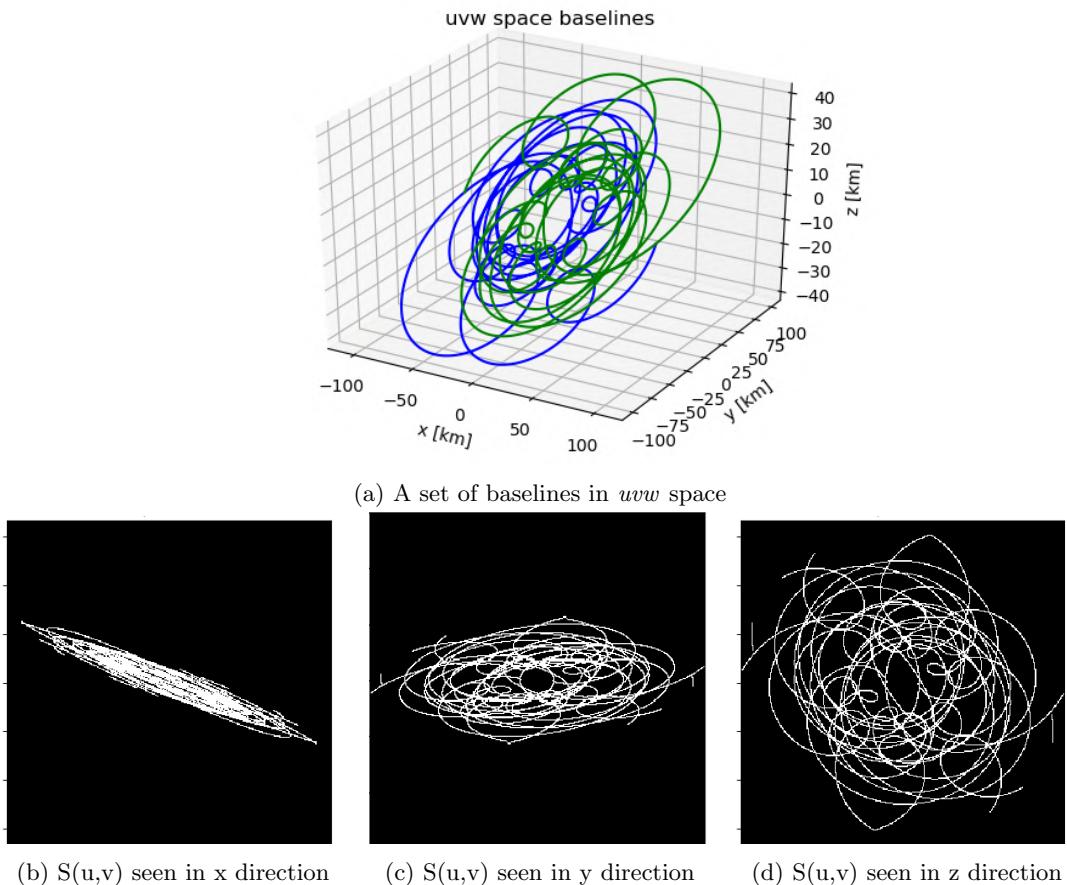


Figure 3.4: Projection of the uvw baselines of a single pair of satellites along the three major axes. $S(u,v)$ appears slightly different since baselines greater than 100 km are not included, as well as different z-axis limits.

²Although in practice transient phenomena must be observed in very short timeframes, in which case this is not possible.

3.4 Influence of the PSF on imaging capabilities

The effect of the sample function and the PSF is best explained when accompanied by visual examples. The examples shown within this chapter are generated from a simulated interferometry satellite pair in orbit around the L4 point, utilising the model described in section 5.8. The resulting baseline collection (shown in Figure 3.4a) of these satellites is projected onto different planes to simulate observations of the Andromeda galaxy using Figure 3.5 as I_{sky} ³.



Figure 3.5: Image of the Andromeda galaxy, after processing to grayscale. Original image by NASA.

In order to properly image a subject the shape of the PSF is extremely important. To avoid directional smudging or repetition full angular coverage of the sample function is required, whereas a good range of different baselines is necessary to capture the full range of spatial frequencies that might be present in the subject. If either of these two are lacking, the resulting measurement will lose significant value. Figure 3.6 shows an example of a sample function with insufficient angular coverage, resulting in a smudge-like PSF. The resulting observation of the Andromeda galaxy is clearly smudged, but it is still recognizable as a galaxy.

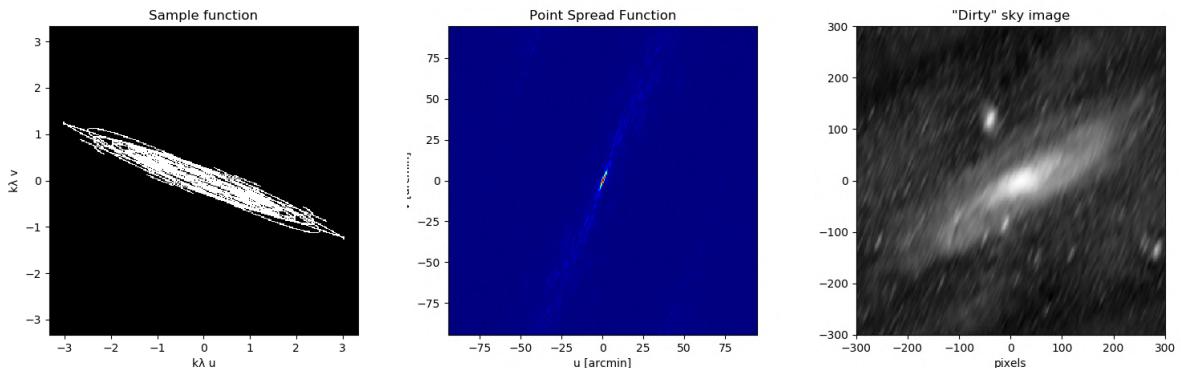


Figure 3.6: Dirty image of Andromeda observed with insufficient angular baseline coverage.

A lack in angular coverage results in smudging, or even repetitions of the PSF outside of its central feature. A lack of baseline magnitude coverage on the other hand means certain spatial frequencies cannot be captured. Figure 3.7 demonstrates the danger of this shortcoming by using the z-projection from Figure 3.4 as sample function.

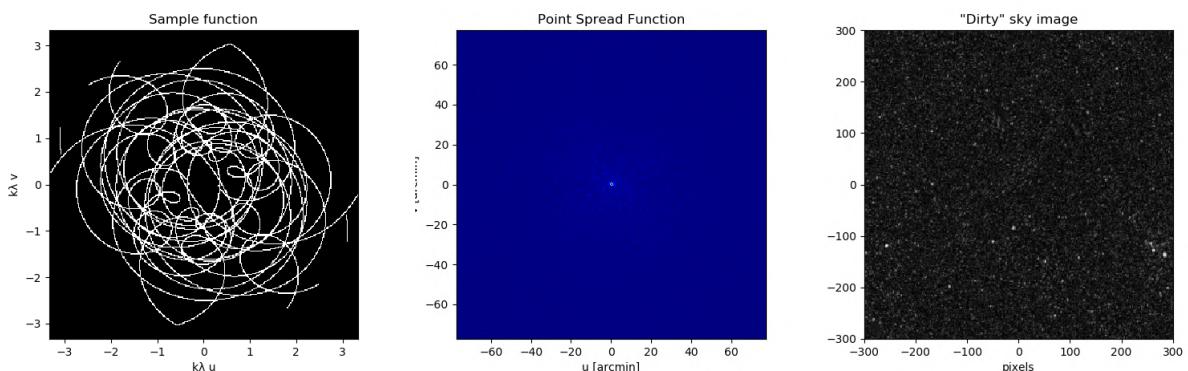


Figure 3.7: Dirty image of Andromeda observed with insufficient baseline range coverage.

³While reviewing this work it has become apparent that the image display by Python is mirrored compared from the original, due to improper definition of the origin. An unfortunate detail, but it does not detract from the value of the demonstrations.

Due to the lack of small baselines in Figure 3.7 low-frequency spatial components of Andromeda cannot be observed in the z-direction, resulting in an observation in which the Andromeda galaxy is invisible. It can be observed that the full-angular coverage of this sample function yields a very focussed PSF, but the lack of small-baseline measurements results in a large amount of low-magnitude outlier noise. When convolved with the original image, this noise smudges out all original image features with low spatial frequencies. What remains is a starry background, made up from the well-sampled higher frequency components of the original image.

In order to successfully observe any subject, it is thus imperative that the set of baselines provide both angular and magnitude coverage. Obtaining a perfect PSF is unfortunately impossible to achieve in practice, which is why extensive "cleaning" algorithms exist which deconvolve the dirty image with the PSF. Cleaning dirty images requires models of the PSF of the instrument, as well as computer algorithms which need to tell real image features from features created by the PSF convolution[51]. The latter part in particular is extremely challenging, especially when the observed sky image is unknown. Despite being very interesting, the subject of image cleaning will not be discussed in detail in this thesis.

3.5 The ideal point-spread function

Image cleaning will always be a necessity, since no system will yield a perfect PSF. At the very best case the PSF will be limited by the diffraction of the radiation it measures, resulting in the Airy Disk pattern. The Airy disk resembles the image of a point source that a theoretically perfect imaging instrument could make, and in practice it represents the limit of an instrument's capability to distinct point sources from one another. The size of the Airy disk is directly related to the Rayleigh criterion⁴, which is also commonly used as the upper limit for the imaging capabilities of a system. The Rayleigh criterion establishes a diffraction limit θ which depends on the system's baseline B , and the wavelength λ :

$$\theta = \arcsin \left(1.22 \frac{\lambda}{B} \right) \quad (3.7)$$

For the OLFAR constellation this limit approximates the diffusion limit at $\theta = 75.49''$ for an observation frequency of 10 MHz at a baseline of 100 kilometres. An airy disk of this size thus represents the ideal PSF for the OLFAR swarm. The Airy disk is mathematically described by the Fraunhofer diffraction pattern of a circular aperture with radius B , resulting in Equation 3.8 [16].

$$I(\theta) = I_0 \left(\frac{2J_1(kB \sin \theta)}{kB \sin \theta} \right)^2 \quad (3.8)$$

In which k is the wave number, described by $\frac{2\pi}{\lambda}$, B is half the effective baseline in meters, and $J_1(x)$ is the first kind Bessel Function of order one. I_0 is the original source intensity. Figure 3.8 shows the Airy disk that can be generated using this equation for the OLFAR system settings, it can be observed that the primary lobe has a radius of approximately $63''$, which is slightly smaller than the Rayleigh criterion. The latter is defined as the distance between point sources where the first minimum overlaps with the maximum. It can be observed that the first minima ring has an outerradius of approximately 1.22 arc-minutes, yielding $73.2''$. The pattern in Figure 3.8 resembles the best possible PSF that a radio interferometer at 10MHz achieve, making it a suitable benchmark to evaluate the quality of a set of baselines.

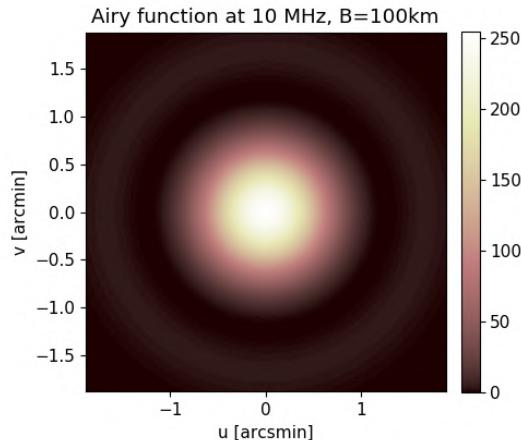


Figure 3.8: Airy disk generated using Equation 3.8 with baseline of 100 km (R=50 km) and a frequency of 10 MHz. Colorbar values scale between 0 and 255 due to data formatting for image generation.

⁴The minimum angle for which a system can distinguish two point sources.

3.6 Sample function resolution determination and scaling

In order to determine a PSF from a set of uvw sample points spread in three-dimensional space, the sample points need to be projected onto a two-dimensional orthogonal plane to a view direction. Projecting three-dimensional data onto an arbitrary two-dimensional plane is straight forward, but sampling this to a discrete map requires a bit of thought. The sample function is a binary map with values of 1 at points of measurements. To create the sample function the projected uvw data has to be discretized, and sampled into pixels of a given spatial size. It is instinctive to want to preserve as much accuracy as possible in this "imaging" process by using a very small pixel resolution, but in practice this approach is far from beneficial. Two-dimensional Fourier transforms are very memory intensive operations which should preferably not be done with extremely large matrices⁵. In order to frequently use the PSF, for example as part of a cost function, the resolution of the sample function needs to be limited. Reducing this resolution may be done without considerable loss in PSF quality, due to the unique nature of the Fourier transform.

3.6.1 Resolving the PSF at very low resolutions

The PSF is computed using the inverse two-dimensional transform of the sample function $S(u, v)$, which ties its pixel resolution directly to that of the sample function. Due to the nature of the transform however, the angular resolution of the pixels of the PSF is not affected by the spatial resolution of the sample function. Instead, the angular resolution scales with the size of the uv space and the observed wavelength. When the PSF is directly computed from discretized data sets, it is the equivalent of directly imaging it using an idealized system. In this case it means that every single pixel representing the PSF must have an angular size equivalent to the Airy disk. While the angular resolution might be decoupled from the spatial resolution of the sample function, it is still connected to the quality of the captured PSF. The spatial resolution of the sample function determines the frequency range that can be displayed in the PSF, following the Nyquist criterion dependent on the number of pixels N and baseline size B :

$$f_{PSF,max} = \frac{1}{2f_{pixel}} = \frac{1}{2B/N} = \frac{N}{2B} \quad (3.9)$$

Thus the level of detail that can be observed in the PSF is directly tied to the spatial resolution of the sample function. The major features of any given PSF will be within the lower end of the frequency domain, meaning that there should be some room to work with lower sampling resolutions while losing minimal information regarding the shape of the PSF. This knowledge is very promising for the use of the PSF as a cost function, since much time can be saved by processing smaller images. If the theory holds true the PSF generated from the same set of orbit data should retain its general shape regardless of spatial sampling resolution. This will be put to the test by evaluating the PSF of the same set of baselines multiple times, using different spatial sampling resolutions. The similarity of the images will be judged based on the grayscale values V of individual pixels P :

$$S = \left(1 - \frac{\sum_{j=0}^{N_{pixel}} \|V_P - V_{ref}\|}{N_{pixel} * V_{max}} \right) * 100\% \quad (3.10)$$

A reference image will first be generated using a fine spatial sampling size of 1x1 wavelength. At 10 MHz this yields a spatial sampling size of 29.6x29.6m, requiring a total of 45.653 Megapixels to picture the uv plane. Only the central $N \times N$ pixels of the reference will be used to compute the similarity, thus the smaller resolution images will be compared to sub-selections of the reference image. Figure 3.9 shows the progression of image similarity for different sample function resolutions.

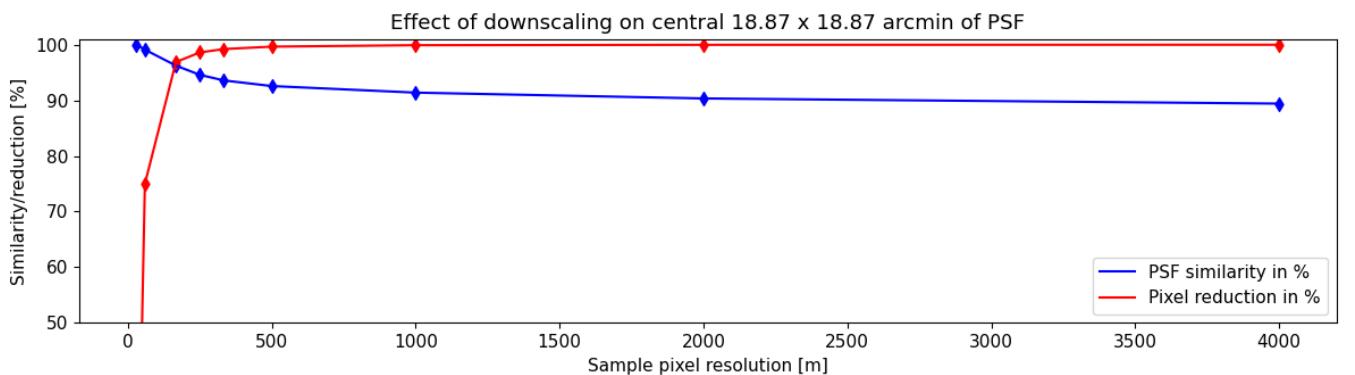


Figure 3.9: Effect of different sampling resolutions on PSF similarity, and the total percentage reduction in pixels from the smaller resolution

⁵At this point in time with hardware available to a master student. Point might be moot in the considerable future.

Figure 3.9 shows that there is a definitive change in the PSF constructed from larger pixel sizes, but the level of change is hardly proportional to the reduction of the total amount of pixels. Over 90% of image similarity may be retained with a larger resolution that reduces the number of pixels by 97%. Figure 3.10 shows the PSF generated using the reference settings, and two PSFs generated using much larger pixels. It can be observed that using larger sampling pixels primarily reduces the details captured in the PSF, while the main features from lower spatial frequencies of the PSF remain visible.

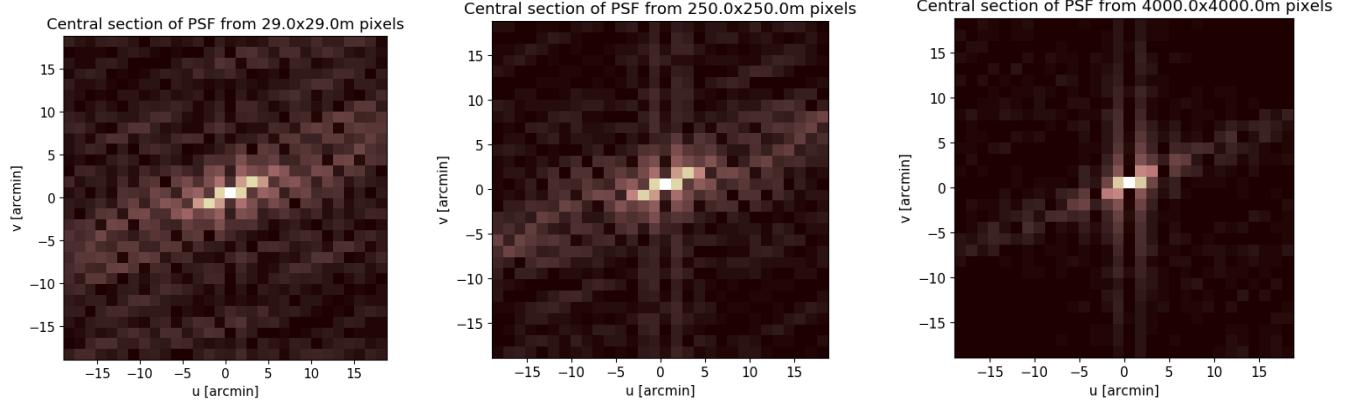


Figure 3.10: Central 31.5 arcminutes of the PSF reconstructed from the dataset shown in Figure 3.4 with different resolutions for the sample function. Sampled with approximately $1\lambda \times 1\lambda$ (29.5m) pixels, 250x250 m pixels, and 4x4 km pixels.

Figure 3.10 shows the reference PSF made after sampling with 1λ pixel size, and the PSF reconstructed using significantly larger pixels sizes. It can be observed that the higher-frequency components of the PSF are lost with lower resolutions, this is particularly evident in the high-frequency fringe patterns. Despite this loss the main body of the PSF remains similar in shape since it is made of low-frequency components primarily. The step from sampling at 29x29 m pixels to 250x250m represents a reduction of 98.6% in the total amount of pixels to be processed, while the primary shape of the PSF is retained quite well. The PSF constructed using 4x4km pixels shows that there is a definite limit to how far resolution reduction can be taken. It's elongated shape is still visible, but all details have been lost.

Even with this loss of detail the low-resolution PSF might still be usable in a cost function. A computer algorithm could still tell that this PSF is elongated, which evidently is an error compared to the reference Airy disk. The loss of detail might hinder it from fine-tuning a constellation, but this coarse PSF should allow for fast global optimization. In that regard all three of these examples could be suitable for a cost function, where the resolution used for the sample function would need to be chosen according to the finesse required.

3.6.2 A PSF-based cost function

The idea of using the PSF as part of the cost function for the optimisation of a radio interferometer was based on the recommendation by E. Dekens in his thesis work [21], who suggested this use as an improvement to the methods he developed. Knowing that a reference PSF generated using Equation 3.8 represents the ideal case, a simple cost function is created by taking the mean of the summed residual value for all pixels. Each pixel is represented by an integer value representing the grayscale brightness, ranging from 0 (black) to 255 (white).

$$C = \frac{\sum \left(\sqrt{(PSF - PSF_{airy})^2} \right)}{N} \quad (3.11)$$

This cost function allows for a quick comparison of the PSF with the ideal case in any arbitrary view direction, ultimately yielding a cost of 0 for perfect solutions. This cost function will need to be evaluated for a fixed set of view directions, which are preferably well-distributed along a unit-sphere. In order to generate a reproducible and well-distributed set of view directions a spherical Fibonacci lattice [29] will be used. This choice is made over a grid-lattice, since the Fibonacci lattice offers a more efficient point distribution [58]. The Fibonacci lattice will yield a set of M direction vectors for which the PSF-cost of the uvw motion will be evaluated. This allows for the evaluation of a cumulative orbit cost accounting for all observation directions. An additional benefit of this method is the capability to optimise for directional measurements as well, by adjusting the weight of view samples based on their direction.

3.7 Orbit requirements of the OLFAR mission

In the observation of astronomical phenomena a distinction is generally made between transient and static radiation sources. As the name suggests transient sources (e.g. a super-nova) emit short-term signals, which can only be observed for a short (in astronomical terms) period of time. Static sources on the other hand may be observed over very large time frames, such as radio galaxies. Following the logic of E. Dekens[21] the observation of static sources will be used as a guideline for the requirements of the orbit design. Static observations might require as much time as the entire mission lifetime, meaning that the general orbit requirements for observations should be adhered to at all times. Table 3.1 shows a selection of the mission requirements for the OLFAR mission, which are relevant to the orbit and constellation design[58]. The origin of these requirements is detailed below.

Table 3.1: Selection of OLFAR mission requirements relevant for orbit and constellation design

Description	Requirement	Source
Minimum baseline	500m ¹	Collision avoidance / safety
Maximum baseline	100 km[11]	Diffusion limited resolution
Maximum baseline rate	1 m/s [11]	Data processing rate
Cumulative measurement time	52000 days[21]	Confusion limited sensitivity

¹ A more conservative 1km is also used for lunar-centric orbits, which inhibit much higher relative velocities.

3.7.1 Collision avoidance

The reason for the preservation of a minimum distance between satellites is entirely related to orbital safety. When such a minimum distance can be guaranteed the swarm should be safe to fly even in case of a total systems failure, at least for a known period of time. The minimum distance which will be used in this thesis is lower than figures commonly used for Lunar orbit studies, since the relative velocities in the L4 domain will be at least a full order of magnitude lower than velocities in the Lunar orbital domain.⁶

3.7.2 Diffusion-limited resolution

The diffusion limited resolution is an absolute lower-end to the achievable resolution of any astronomical telescope system. It originates from electromagnetic signals scattering as they pass through the inter-stellar medium: a collection of ionized particles, dust, and solar winds. The interstellar medium is found in what is commonly seen as "empty" space between galaxies and stars. The density of these particles is often not more than a few atoms per cubic centimetre, but over the vast distances between astronomical sources the presence of this medium introduces a lower limit to any observable resolution. Jester and Falcke estimate that this resolution limit will be around 1 arc-minute at 10 MHz [33], making it futile to use larger baselines than 100 km for radio interferometry at this wavelength. Measurements taken with larger baseline would still provide usable results, but designing around a higher maximum puts more strain on the inter-satellite communication systems.

3.7.3 Maximum baseline rate

In order to reach their desired sensitivity radio interferometers need to observe continually for a prolonged period of time, which is necessary to integrate data and correlate results between the interferometer elements. It is estimated that at least 1 full second will be necessary for measurements[11], during which it is necessary that the baseline does not change much more than a fraction of the wavelength. The exact acceptable limit to the baseline rate is dependent on many system parameters, which makes it difficult to predict. A general upper limit is 1 m/s, while in practice baseline rates are preferably as low as possible to allow for longer integration times.

3.7.4 Confusion-limited sensitivity

The confusion limit is a term which represents the maximum effective sensitivity of any observation system, beyond which lower-magnitude noise sources will be impossible to distinguish from background noise. Jester and Falcke estimate that the maximum effective sensitivity of a interferometer with a baseline of 100km at 10MHz is 65 mJy[33]. Using this figure E. Dekens estimated that achieving this system sensitivity requires a cumulative observation time of 52,000 days[21]. This cumulative figure is an indication of the collective observation time of all baselines, which highlights why a radio interferometer constellation benefits from the use of a large amount of satellites.

⁶Demonstrated in Chapter 7.

4 Constellation design around the 4th libration point

Having selected the 4th libration point as deployment location, this chapter is dedicated to studying the dynamics of motion around this location. If the dynamics of the location are well understood, this might lead to more efficient methods of designing swarm geometries and orbits. This chapter starts out with analysing the simplified description of the three-body problem, before branching in more complex descriptions such as the surfaces of Hill.

4.1 The barycentric three-body problem

Mapping or describing the motion of satellites is complex, in space everything is continually moving relative to each other and rotating around its own axes. A well-defined reference frame is necessary to display this complex system in an orderly fashion. One such frame is the barycentric reference frame, which is particularly well suited to describe orbits around the L4 point. A barycentric frame is commonly defined using a three-body problem, with a central body (Earth / E), a secondary orbiting mass (Moon / M), and a third orbiting body (S) of negligible mass. The barycentre¹ between the Earth and the moon is used as origin, with the local x axis cast along the barycentre-moon vector. The local y axis is perpendicular to this axis along the orbital plane, and the Z-vector completes the frame following the right hand rule. Figure 4.1 shows the barycentric frame with a third body S .

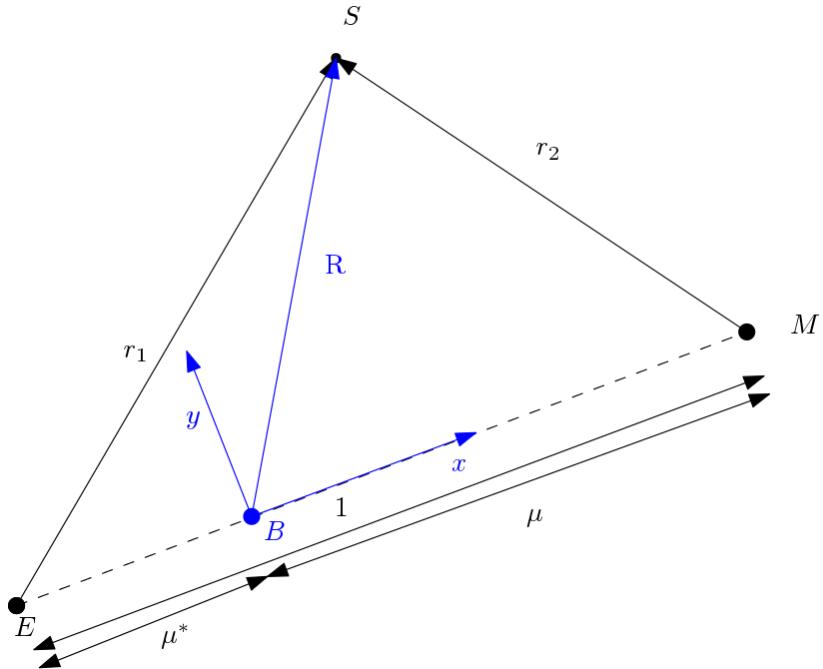


Figure 4.1: Definition of the barycentric frame in the Earth-Moon system. Satellite S orbits around the L4 point trailing the Moon (M) around the Earth (E)

Within the barycentric reference frame distances are normalized to the magnitude of the Earth-Moon vector. The mass ratio between the primary and secondary bodies is used to express the position of the barycentre along the main vector:

$$\begin{aligned}\mu^* &= \frac{m_M}{m_E + m_M} = 0.0122 \\ \mu &= 1 - \mu^* = 0.9878\end{aligned}\tag{4.1}$$

¹Centre of mass

Using the definition of μ^* , the range vectors may be expressed as:

$$\begin{aligned} r_1 &= \sqrt{(x + \mu^*)^2 + y^2 + z^2} \\ r_2 &= \sqrt{(x - 1 + \mu^*)^2 + y^2 + z^2} \end{aligned} \quad (4.2)$$

It is important to realise that this frame of reference is co-rotating with the Earth-Moon system. Hence, the velocities of state vectors in this frame are relative to the rotation of the system as well. A satellite orbiting the Earth with the same orbit as the L4 point would thus show no velocity within this reference frame. When the barycentric frame is used there may also be mention of its characteristic time, this is a normalized unit of time which will be used as a simplification. The characteristic time is defined as a time divided by the rotation period of the reference frame:

$$t_{char} = \sqrt{(EM^3/(G * (m_1 + m_2)))} \quad (4.3)$$

4.2 Simplified three-body problem

Solving the three-body problem in the defined barycentric system yields the following equations of motion (EoM) for the satellite S [40]:

$$\begin{aligned} \ddot{x} - 2\dot{y} - x &= -\frac{(1 - \mu^*)(x + \mu^*)}{r_1^3} - \frac{\mu^*(x - 1 + \mu^*)}{r_2^3} \\ \ddot{y} + 2\dot{x} - y &= -\frac{(1 - \mu^*)y}{r_1^3} - \frac{\mu^*y}{r_2^3} \\ \ddot{z} &= -\frac{(1 - \mu^*)z}{r_1^3} - \frac{\mu^*z}{r_2^3} \end{aligned} \quad (4.4)$$

The simplified three-body problem can be solved for the locations of the Lagrangian points where all accelerations are zero. In the barycentric system this yields the following coordinates for the L4 point, normalized to the length of the Earth-Moon vector:

$$L_4 = \begin{bmatrix} \mu - 1/2 \\ \sqrt{3}/2 \\ 0 \end{bmatrix}$$

Assuming that the satellites start in the vicinity of the L4 point, the initial position of the satellite is expressed as the position of the L4 point modified by an arbitrary displacement vector \vec{d} :

$$S_0 = \begin{bmatrix} \mu - 1/2 + d_x \\ \sqrt{3}/2 + d_y \\ d_z \\ \dot{x} \\ \dot{y} \\ \dot{z} \\ \ddot{x} \\ \ddot{y} \\ \ddot{z} \end{bmatrix} = \begin{bmatrix} 1/2 - \mu^* + d_x \\ \sqrt{3}/2 + d_y \\ d_z \\ \dot{x} \\ \dot{y} \\ \dot{z} \\ \ddot{x} \\ \ddot{y} \\ \ddot{z} \end{bmatrix} \quad (4.5)$$

Within this state vector the values $d_x, d_y, d_z, \dot{x}, \dot{y}$, and \dot{z} may be seen as the design variables of the swarm geometry. Assuming no initial thrust, the acceleration components are a result of the natural motion of the satellite in orbit. The EoM of the satellite may be used to find an expression for these values, which might prove useful in designing a cohesive swarm. Starting by filling in the values of this state vector for the EoM in the x-direction:

$$\begin{aligned} \ddot{x} - 2\dot{y} - 1/2 + \mu^* - d_x &= -\frac{(1 - \mu^*)(1/2 - \mu^* + d_x + \mu^*)}{r_1^3} - \frac{\mu^*(1/2 - \mu^* + d_x - 1 + \mu^*)}{r_2^3} \\ \ddot{x} &= -\frac{(1 - \mu^*)(1/2 - \mu^* + d_x + \mu^*)}{r_1^3} - \frac{\mu^*(1/2 - \mu^* + d_x - 1 + \mu^*)}{r_2^3} + 2\dot{y} + 1/2 - \mu^* + d_x \\ \ddot{x} &= -\frac{(1 - \mu^*)(1/2 + d_x)}{r_1^3} - \frac{\mu^*(-1/2 + d_x)}{r_2^3} + 2\dot{y} + 1/2 - \mu^* + d_x \end{aligned}$$

Doing the same for the y and z components:

$$\begin{aligned}\ddot{y} &= -\frac{(1-\mu^*)(\sqrt{3}/2 + d_y)}{r_1^3} - \frac{\mu^*\sqrt{3}/2 + \mu^*d_y}{r_2^3} - 2\dot{x} + \sqrt{3}/2 + d_y \\ \ddot{z} &= -\frac{(1-\mu^*)d_z}{r_1^3} - \frac{\mu^*d_z}{r_2^3}\end{aligned}$$

The complexity of these equations of motion stems from the presence of the two range vectors, which make the accelerations dependent of all 3 position variables. With the new definition of the initial state vector the range vectors yield:

$$\begin{aligned}r_1 &= \sqrt{(1/2 - \mu^* + d_x + \mu^*)^2 + (\sqrt{3}/2 + d_y)^2 + d_z^2} = \sqrt{(1/2 + d_x)^2 + (\sqrt{3}/2 + d_y)^2 + d_z^2} \\ r_2 &= \sqrt{(1/2 - \mu^* + d_x - 1 + \mu^*)^2 + (\sqrt{3}/2 + d_y)^2 + d_z^2} = \sqrt{(-1/2 + d_x)^2 + (\sqrt{3}/2 + d_y)^2 + d_z^2}\end{aligned}\quad (4.6)$$

The problem may be simplified significantly by realising that the displacement vector is also scaled relative to the Earth-Moon distance. Assuming that the swarm is initially placed within 1000 kilometres of the L4 point the maximum value of any displacement would be $d = \frac{10^3}{384 \cdot 10^3} = 2.6 \cdot 10^{-3}$. Realising this, the range vectors can be approximated as:

$$\begin{aligned}r_1 &\approx \sqrt{(1/2)^2 + (\sqrt{3}/2)^2} = 1 \\ r_2 &\approx \sqrt{(1/2)^2 + (\sqrt{3}/2)^2} = 1\end{aligned}\quad (4.7)$$

This simplification allows for the resolution of simplified initial acceleration terms:

$$\begin{aligned}\ddot{x} &\approx -(1-\mu^*)(1/2 + d_x) - \mu^*(-1/2 + d_x) + 2\dot{y} + 1/2 - \mu^* + d_x \\ \ddot{x} &\approx -1/2 - d_x + 1/2\mu^* + \mu^*d_x + 1/2\mu^* - \mu^*d_x + 2\dot{y} + 1/2 - \mu^* + d_x \\ \ddot{x} &\approx d_x(-1 + \mu^* - \mu^* + 1) + \mu^*(1/2 + 1/2 - 1) + 2\dot{y} = 2\dot{y} \\ \\ \ddot{y} &\approx -(1-\mu^*)(\sqrt{3}/2 + d_y) - \mu^*\sqrt{3}/2 + \mu^*d_y - 2\dot{x} + \sqrt{3}/2 + d_y \\ \ddot{y} &\approx -\sqrt{3}/2 - d_y + \mu^*\sqrt{3}/2\mu^*d_y - \mu^*\sqrt{3}/2 - \mu^*d_y - 2\dot{x} + \sqrt{3}/2 + d_y \\ \ddot{y} &\approx -2\dot{x} \\ \\ \ddot{z} &\approx -(1-\mu^*)d_z - \mu^*d_z = -d_z\end{aligned}$$

Using these approximations the initial state vectors of the satellites near the L4 point may be expressed as S_0 . Applying this to a second satellite as well this allows for an estimation of the initial baseline state vector B_0 between these two satellites, where Δ denotes differences between satellites:

$$S_0 \approx \begin{bmatrix} \mu - 1/2 + d_x \\ \sqrt{3}/2 + d_y \\ d_z \\ \dot{x} \\ \dot{y} \\ \dot{z} \\ 2\dot{y} \\ -2\dot{x} \\ -d_z \end{bmatrix} \quad B_0 = \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \\ \Delta \dot{x} \\ \Delta \dot{y} \\ \Delta \dot{z} \\ 2\Delta \dot{y} \\ -2\Delta \dot{x} \\ -\Delta z \end{bmatrix} \quad (4.8)$$

This assumption does not hold up well in the long term however, so its use will be limited to providing insight in the factors which are of importance to the stability of the swarm. To keep the swarm cohesive, having similar acceleration properties will be very important in the long term. In this formulation it is easy to see that variation in the out-of-orbit plane of the placement of satellites is not desirable for swarm cohesion, since it creates an immediate relative acceleration. An exception to this rule would be a symmetrical deployment around the $z = 0$ plane, which creates a mirrored relative acceleration in the swarm. Positioning in the orbital plane itself is not as important as making sure the initial velocities are identical.

4.3 The Pseudo-potential

Based on the three body problem, a psuedo-potential can be defined as a product of the gravitational pull of both bodies, compensated for the rotating reference frame. This potential may be used to gain more insight into the stable regions within the three-body problem. Using the definition by A. Rubinsztejn [55] the pseudo-potential is defined as the product of gravitational potential and the force from the co-rotating frame. The pseudo-potential field should not be confused with a surface of Hill, of which it is a part off. Hill surfaces will be discussed in section 4.4.

$$\Omega^* = \frac{1 - \mu^*}{r_1} + \frac{\mu}{r_2} + \frac{x^2 + y^2}{2} \quad (4.9)$$

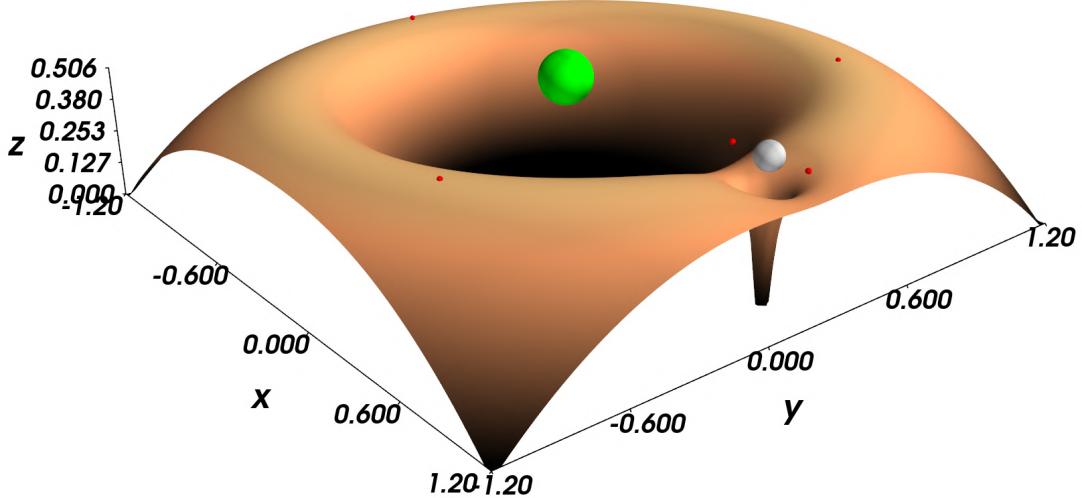


Figure 4.2: Shape of $2 - \Omega^*$ in a normalized Earth-Moon system, with the position of the lagrangian points indicated by red dots.

Figure 4.2 depicts the inverted pseudo-potential field and its relation to the two major bodies and the libration points. It can be seen that the libration points lie on saddle points (L_1, L_2), or local valleys (L_3, L_4, L_5), taking into account that these are hills in the inverted display. The motion of objects around these libration points can better be explained through further study of this potential. A potential gradient can be found by seeking the derivative of this potential field in the cardinal directions of the barycentric frame:

$$\begin{aligned} \frac{\Omega^*}{dx} &= x - \frac{(1 - \mu^*)(x + \mu)}{((x + \mu^*)^2 + y^2 + z^2)^{3/2}} - \frac{\mu^*(2x - 2 - 2\mu^*)}{2((x - 1 - \mu^*)^2 + y^2 + z^2)^{3/2}} \\ \frac{\Omega^*}{dy} &= y - \frac{(1 - \mu^*)y}{((x + \mu^*)^2 + y^2 + z^2)^{3/2}} - \frac{\mu^*y}{((x - 1 - \mu^*)^2 + y^2 + z^2)^{3/2}} \\ \frac{\Omega^*}{dz} &= - \frac{(1 - \mu^*)z}{((x + \mu^*)^2 + y^2 + z^2)^{3/2}} - \frac{\mu^*z}{((x - 1 - \mu^*)^2 + y^2 + z^2)^{3/2}} \end{aligned} \quad (4.10)$$

Using these derivatives the equations of motion can be shortened to the following expressions:

$$\begin{aligned} \ddot{x} &= 2\dot{y} + \frac{\Omega^*}{dx} \\ \ddot{y} &= -2\dot{x} + \frac{\Omega^*}{dy} \\ \ddot{z} &= \frac{\Omega^*}{dz} \end{aligned}$$

The derivative form can be used to further explore the properties of a satellite in motion within a 3 body problem. By seeking a stable solution with no inherent acceleration ($\ddot{x}, \ddot{y}, \ddot{z} = 0$) it can be found that truly stable velocities only exist on the $z = 0$ plane, meeting the conditions that:

$$\begin{aligned} \dot{x} &= 1/2 \frac{\Omega^*}{dy} \\ \dot{y} &= -1/2 \frac{\Omega^*}{dx} \end{aligned}$$

A second property of interest is what will be called the auto-velocity acceleration (AVA) field. This field describes the most natural initial acceleration for a motionless object in the barycentric frame. It is based on the assumption that for a coarse grid the satellite accelerates to the point where its velocity equals the initial acceleration. This assumption corrects the acceleration field for the velocity a satellite will receive during its time in the grid section. For a grid with grid size S the time it would take to move to the next grid field from the initial acceleration is approximated by:

$$t_S = \sqrt{\frac{2S}{\ddot{a}_0^2}} \quad (4.11)$$

$$\ddot{a}_0 = \sqrt{\ddot{x}^2 + \ddot{y}^2 + \ddot{z}^2} = \sqrt{\frac{\Omega^*}{dx}^2 + \frac{\Omega^*}{dy}^2 + \frac{\Omega^*}{dz}^2}$$

Accounting for the velocity gained by the time the satellite will leave the grid sector, the acceleration at that point is recomputed using this velocity. This yields the following description for the AVA field for objects with $z=0$:

$$t_S = \sqrt{\frac{2S}{\ddot{a}_0^2}}$$

$$\ddot{x} = 2t_S \frac{\Omega^*}{dy} + \frac{\Omega^*}{dx} \quad (4.12)$$

$$\ddot{y} = -2t_S \frac{\Omega^*}{dx} + \frac{\Omega^*}{dy}$$

$$\ddot{z} = \frac{\Omega^*}{dz}$$

The resulting acceleration field is a better approximation of the motion of an object with no initial relative velocity, although it is no substitute for numerical propagation of the motion. Figure 4.3 shows the stable velocity field and the AVA field of the 3-body problem in the barycentric system made unitless using the size of the Earth-Moon vector:

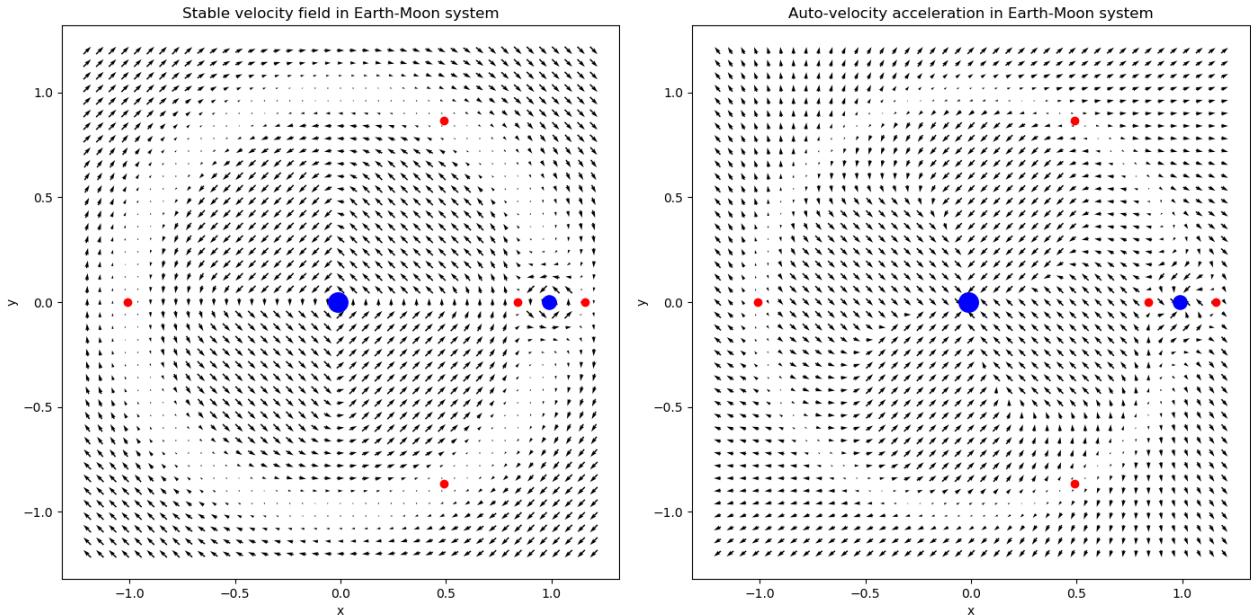


Figure 4.3: Stable velocity field(L) and auto-velocity acceleration(R) field in the Earth-Moon system, with red markers depicting the position of the libration points and blue markers the Earth and Moon.

These fields explain the tendency for satellites to enter a clockwise motion around the L4 point, which lies in a vortex between the inward and outward accelerations. The stable velocity field shows a large ring which seems to be inherently stable, but the AVA field shows that the dynamics of this ring become unstable around the -0.5 x coordinate. The AVA field shows a large difference in size between what could be described as "stable" fields surrounding the L4 and L5 points, compared to the other libration points. It also shows how fickle the L1 and L2 points are in comparison, as these lie trapped between outward vortexes. The outward acceleration in the AVA-field outside of the stable ring is due to the assumption that the satellite has no initial velocity to the co-rotating frame. At a higher radius than the moon's orbit the initial velocity will be faster than the orbiting velocity, creating outward motion compared to the barycentric frame.

Figure 4.3 shows that when left to its own devices, an object orbiting the L4 point will tend to orbit in a clockwise motion around the L4 point, when its motion is considered in the barycentric frame. The AVA field shows that the natural accelerations in this orbit are outwards, along with the natural direction of motion. Based on this it can be expected that any swarm utilising passive formation flight around the L4 point will follow suit, showing clockwise motion around L4 with a natural expansion of the orbit. Eventually expansion of this orbit will draw the constellation away from the stable region around the L4 point, leading to a natural degeneration into an Earth-centralised orbit. In order to maximise the orbit time around the L4 point, it will thus be important to deploy the swarm into an orbit which initially lies close to the L4 point.

4.4 Surfaces of Hill

A more commonly used modification of the three-body problem is Jacobi's integral along with the so-called surfaces of Hill. Jacobi's integral is built by first multiplying the derivatives of Ω^* shown in Equation 4.3 with the velocities in their derivative directions, which yields[60]:

$$\dot{x} \frac{\Omega^*}{dx} + \dot{y} \frac{\Omega^*}{dy} + \dot{z} \frac{\Omega^*}{dz} = \dot{x}\ddot{x} + \dot{y}\ddot{y} + \dot{z}\ddot{z} \quad (4.13)$$

By realising that the potential field is independent of time, the leftmost terms can be described as the derivative of Ω^* in regards of time. This allows for solving this equation through integration:

$$2\Omega^* - C = \dot{x}^2 + \dot{y}^2 + \dot{z}^2 = V^2 \quad (4.14)$$

The integration constant C is called Jacobi's Constant, which is determined by the initial position and velocity of a satellite in the three-body system from Equation 4.15:

$$C = 2\Omega^* = x^2 + y^2 + \frac{2(\mu^* - 1)}{r_1} + \frac{2\mu^*}{r_2} - V^2 \quad (4.15)$$

Surfaces of Hill are defined as surfaces for which the velocity V of the third body is equal to zero, leading to surfaces in the XYZ plane where $2\Omega^* = C$. Upon realizing that for any real body $V^2 \geq 0$ this equation can be used to find regions in space where objects can move for a set of initial conditions[60]:

$$2\Omega^* = x^2 + y^2 + \frac{2(\mu^* - 1)}{r_1} + \frac{2\mu^*}{r_2} \geq C \quad (4.16)$$

The triangular Lagrangian points represent minima of the potential field Ω^* , and resulting initial Jacobi constants for satellites orbiting these locations will be low. This condition also shows that a satellite starting in orbit around a triangular Lagrangian point has sufficient energy to reach all other locations in the barycentric system. Whereas reaching the L4 point is impossible for any satellite which has not started its orbit in its vicinity[60]. Figure 4.4 shows how a satellite starting at $(0.5,0,0)$ with an initial velocity of 50 m/s relative to the barycentric frame is stuck in a near-Earth orbit, while with 1020 m/s it could escape the barycentric frame through the "gate" near the moon. Yet despite being able to escape the barycentric system, the latter satellite lacks the energy to move to the L4 point.

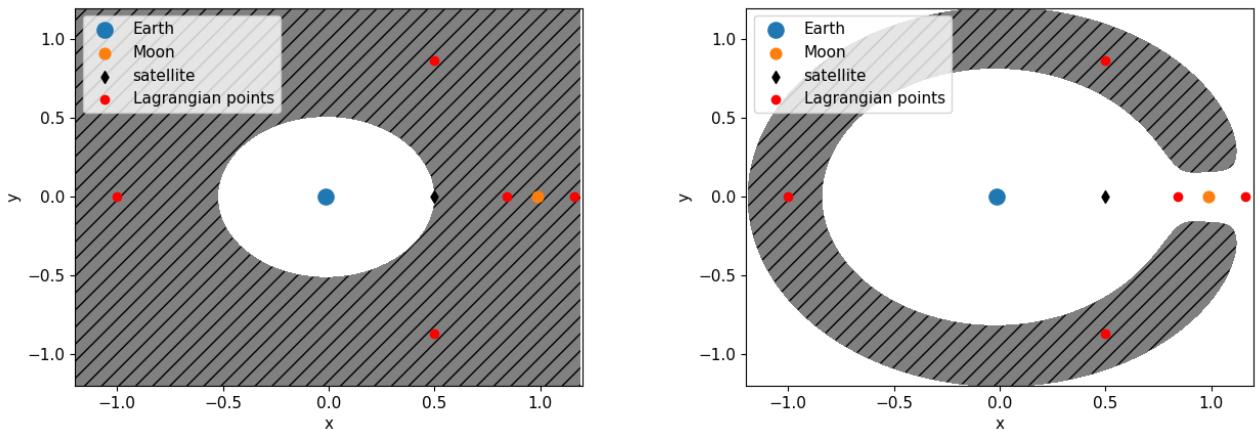


Figure 4.4: Inaccessible areas through the condition $2\Omega^* \geq C$ for a satellite starting at $(0.5,0,0)$ with relative velocities of 50 and 1020 m/s.

While this knowledge is not necessarily useful for designing a swarm to stay in an L4-centric orbit, it does garner some insight into the paths the swarm might take once it naturally evolves out of it. Not being constrained by energy limitations, it can be expected that a swarm enters a very unpredictable wide orbit after de-orbiting the L4 point. Keeping this in mind the de-orbit process of the swarm of individual satellites ought to be controlled, rather than leaving it up to chance or natural evolution. The surface of Hill shows that any satellite starting in vicinity of the L4 point has the potential to endanger all near-Earth missions once it de-orbits.

4.5 Motion around the L4 point

While the AVA field shown in Figure 4.3 already gives a general idea of the clockwise evolution and growth of a swarm around the L4 point, the model used to describe Hill surfaces may be used to garner more insight. K. Wakker used a set of relative coordinates (x', y', z') to an initial location (x_0, y_0, z_0) to derive an adjusted set of equations of motion which can describe the movement of an object around the L4 point[60]:

$$\begin{aligned} x' &= x - x_0 \\ y' &= y - y_0 \\ z' &= z - z_0 \\ \ddot{x}' - 2\dot{y}' - x'\Omega_{xx} - y'\Omega_{xy} &= 0 \\ \ddot{y}' + 2\dot{x}' - x'\Omega_{xy} - y'\Omega_{yy} &= 0 \\ \ddot{z}' - z'\Omega_{zz} &= 0 \end{aligned} \tag{4.17}$$

Within these equations of motion it can be seen that motion in the z axis is entirely decoupled from motion in the xy plane, resulting in a harmonic periodic motion in the z direction[60]. Barring perturbation sources outside of the three-body model, it should thus be expected that the motion in the z plane experiences no natural growth or decay. The motion in the xy plane is much more difficult to interpret however, these equations require more manipulation² to become comprehensible. For orbits in proximity of the triangular Lagrangian points K. Wakker rewrites the first two equations of motion to:

$$\begin{aligned} \ddot{x}' - 2\dot{y}' - \frac{3}{4}x' \mp \frac{3}{4}\sqrt{3}(1-2\mu)y' &= 0 \\ \ddot{y}' + 2\dot{x}' \mp \frac{3}{4}\sqrt{3}(1-2\mu)x' - \frac{9}{4}y' &= 0 \end{aligned} \tag{4.18}$$

Through a long and arduous derivation which will not be shown in this thesis, K. Wakker alters these equations of motion to describe them as a sum of harmonical motions around the triangular Lagrangian points:

$$\begin{aligned} x' &= C_1 \cos s_1 t + S_1 \sin s_1 t + C_2 \cos s_2 t + S_2 \sin s_2 t \\ y' &= \vec{C}_1 \cos s_1 t + \vec{S}_1 \sin s_1 t + \vec{C}_2 \cos s_2 t + \vec{S}_2 \sin s_2 t \end{aligned} \tag{4.19}$$

In which $C_1, C_2, S_1, S_2, \vec{C}_1, \vec{C}_2, \vec{S}_1$, and \vec{S}_2 are harmonic constants which will not be further detailed in this thesis. Their exact definitions and values are not important enough to warrant including three pages of derivations necessary to shed light on them³. What is important is the realisation that there are two distinct frequency terms in these equations of motion: s_1 and s_2 . Motion in the xy plane around the L4 point is made up from two distinct frequencies, both a long- and short-term harmonic which are related to s_1 and s_2 . K. Wakker describes the periods of these motions as a function of the orbital period of the Moon (T_2) and the harmonic components:

$$T_{xy,l} = \frac{T_2}{s_1} \quad T_{xy,s} = \frac{T_2}{s_2} \tag{4.20}$$

The harmonic motion in the z plane also has a distinct period, which is expressed as:

$$T_z = \frac{T_2}{\frac{1-\mu}{r_1} + \frac{\mu}{r_2}} = \frac{T_2}{\sqrt{K}} \tag{4.21}$$

In which the value of K equals 1 for points in close proximity to the triangular Lagrangian points. K. Wakker calculated the values for the harmonic terms s_1 and s_2 for the Earth-Moon system to be 0.2982 and 0.9545 for the L4 point[60].

²To prevent padding this chapter with the long and arduous derivations most steps are left out, leaving only a set of condensed equations necessary to explain the subject matter.

³Interested readers are referred to Fundamentals of Astrodynamics by Karel F. Wakker [60], section 3.10

Using the values provided by Wakker the orbital components of a L4-centric orbit can be computed using an orbital period of 27.32 days for the Moon:

$$\begin{aligned}
 T_{xy,l} &= \frac{27.32}{0.2982} = 91.62 \\
 T_{xy,s} &= \frac{27.32}{0.9545} = 28.62 \\
 T_z &= \frac{27.32}{1} = 27.32
 \end{aligned} \tag{4.22}$$

These orbital periods around the L4 point are constant regardless of initial state and shape of the L4-centric orbit, as long as the initial state is close to the L4 point. This property is very beneficial for formation flight, as it means that the entire swarm will experience the same periods in their relative motion. Over prolonged periods of time however this property will be lost, the orbits are not entirely periodic since the fraction of periods $91.62/28.62$ yields an irrational number[60]. Even if some components of the swarm would be put into very different orbits this property can be used to link up again near the origin point, allowing for more complex swarm designs consisting of distributed smaller constellations.

The derivations of these relations by Wakker are very useful to understanding the nature of motion around the L4 point, as well as doing some basic predictions of the relative motions a satellite swarm will undergo. A very important question however is how accurate these methods are compared to real-world environments, where satellites experience perturbations from other third bodies and solar radiation pressure. The next chapter is dedicated to investigating the effect of orbital perturbations, and studying the applicability of a three-body model for longer-term applications.

4.6 Designing around orbit decay

The models based on the three-body model show that the natural tendency of motion for the swarm is clockwise growth around the L4 point on the xy , xz , and xz planes. The evolution of this motion is inevitable, and it can at best only be delayed by starting as close to the L4 point as possible. The closer to the L4 point a satellite is, the smaller the gradient of the potential field and induced accelerations are. Thus, it will be important to keep the swarm close to the L4 point for as long as possible. An important consideration is that the swarm should not start on-top of the L4 point, and it needs to have some form of initial motion. Since the L4 point is placed at a top of the potential field satellites starting at this location have the potential of rapid chaotic collapse if given no initial velocity, which can be disastrous for maintaining swarm cohesion. To avoid this swarms should instead start at slight sub-optimal potentials with a defined initial velocity relative to the frame.

This motion is primarily driven by the differences in the potential field, and it can be expected that the satellite swarm itself will also drift apart over time because of the spatial differences between these satellites. Satellites closer to the L4 point within the swarm will experience smaller potential gradients, and thus naturally experience less accelerations. The potential gradient is strongest in the xy plane, meaning that an ideal swarm design would show little spread in this plane. The initial swarm design should thus be compact, and be allowed to grow to larger baselines using natural orbit decay.

These conclusions are made using simplified models, and a question of interest is how well these conclusions hold up in a perturbed environment. In particular the conclusions regarding small-scale swarm movement might be based on too simple a model to use as guideline, which is why broader search areas will be used for optimisation.

5 Numerical simulation environment

Considering that physical experiments with satellite constellations are not a reasonable request for a thesis, the evolution of constellation designs and orbits will need to be studied in a digital environment. This chapter will discuss the formation of the digital environment in detail, as well as the chosen propagation and integration settings that will be used to evaluate the orbit of satellites. Since most of the work in this section is built upon pre-existing toolboxes and libraries the first section in the chapter will introduce those, to establish future terminology and give credit where it is due.

5.1 Used third-party libraries and tools

The work presented in this thesis is built upon some pre-existing software libraries, without which the scope of work presented could not be possible. The majority of the work presented in this thesis is built upon the astronomical toolbox developed by the Technical University of Delft, which in turn integrates third party tool sets and libraries.

Technical University of Delft Astronomical Toolbox (TUDAT)

The entire numerical environment used for orbit propagation and modelling in this thesis is built upon the TUDAT toolset developed at the Technical University of Delft [22]¹. This toolkit combines a wide variety of existing models and integrated tools to allow its users to create detailed simulations easily.

PAGMO2

A large part of the presented work relies on the Pagmo2 optimisation library for C++ developed by F. Biscani and D. Izzo [12]². This library has proven to be vital to this thesis, allowing for the easy implementation of various algorithms and a unique architecture which is used to speed up the optimisation process.

SPICE

The CSPICE toolkit³ developed by NASA [4], [5] is integrated within the TUDAT environment and it is relied on to provide position data and trajectory models for all planetary bodies used in the simulation environment. Internally SPICE combines data kernels from observatories from all over the world, which makes it difficult to further appropriate credit.

5.2 Important remarks

Unless explicitly mentioned otherwise, the reference date for scenario start positions used is the first of January 2030. This means that in any future case studies and propagations the planetary start positions will be equivalent to the best estimates for that date, and full-year propagations will likewise propagate the planetary positions from 1 January 2030 to 1 January 2031. This is done to ensure the research is most representative of the future plans for OLFAR.

x , y , and z coordinates denoted in this, and future, chapters will refer to the coordinate systems of the J2000 Earth-Centered Inertial (ECI) reference frame (see Appendix B). Coordinates in the barycentric frame will be denoted using x_B , y_B , and z_B and their use will be explicitly mentioned.

¹For more information about TUDAT see tudat.tudelft.nl

²See <https://esa.github.io/pagmo2/> for more information

³See <https://naif.jpl.nasa.gov/naif/index.html> for more information

5.3 Acceptable model uncertainties

When determining which perturbations to include in the environment model it is important to consider the acceptable uncertainty for the model as a whole. After all, deciding to remove a perturbation source from the model introduces an error compared to the full model. Whether this can be done will have to depend on what uncertainty limits are deemed acceptable beforehand. Since the perturbation sources considered in this section influence the orbit of an object over time, the time-frame in which this consideration is made will also need to be kept in mind. For this work the considered timeframe will be a full year in orbit, aligning with the sought result.

For the purpose of the OLFAR mission modelling the most important figures of merit are the evolution of the baselines, and the baseline rate. It is more important to be confident about these values, than the relative placement of the swarm compared to the L4 point. It is important to consider that the proximity of the satellites is so small compared to the distance to the perturbation sources, that the resulting error in displacement is approximately equal for all satellites in the swarm. Due to this property the error in relative motion modelling is smaller than the error in global positioning. The latter is also much less important for the consideration of the suitability of an orbit for interferometry, which is why a less strict accuracy requirement will be used for the global position of the satellites compared to relative motion.

The acceptable uncertainty of the baseline rate will be based on $1/10^th$ the baseline rate limit of 1 m/s, resulting in 10 cm/s. The overall position accuracy of the swarm will use a less restrictive 10 m requirement. The combination of these requirements means that the environment will be designed to propagate the relative motion of satellites with good certainty, while retaining a decent estimate of the overall position of the swarm.

5.4 Environmental setup

This section will briefly address the general modelling of the numerical environment used throughout this chapter, and later chapters in the thesis. Not all the planetary bodies and perturbations described in this section will be present in the final model, these deliberations will be discussed in section 5.6, and the components kept in the final model will be discussed in section 5.8.

5.4.1 Planetary modelling and positioning

For the environment model all planetary bodies in the solar system will be modelled, as for the minor bodies only the Moon itself will be included in the environment. The influence of Pluto and other minor bodies and moons will be assumed negligible for near-Earth orbits, and their modelling will thus not be considered. Finally the Sun itself completes the solar system, being the source of radiation pressure and a strong gravitational force. A custom SPICE kernel is used to compute starting positions of these bodies on the 1st of January 2030, and to propagate the motion of these planets up to 2040 where necessary. This kernel combines data from the following ephemeris files released by JPL's Navigational and Ancillary Information Facility (NAIF) [44], [4]: *de438*, *mar097*, *jup310*, *sat427*, *ura115*, and *nep090*.

Both the Earth and the Moon are close enough in proximity to warrant investigating the influence of their harmonic gravitational fields, in which case it also becomes important to consider their rotation. For both these bodies the Spice library will again be used to propagate their rotation using their respective IAU rotation models. Apart from positioning and rotation, there is no need for more extensive modelling of the planets. The primary purpose of this part is to correctly place the perturbation sources affecting the orbit of the satellites.

5.4.2 Perturbation sources

The primary perturbations that will influence the orbit of satellites around the Earth-Moon L4 point will be gravitational in nature. The magnitude of the force exerted on the satellites is easily predicted using Newton's gravitational law, but the long-term effects of forces on an orbit are much more difficult to predict, as even small forces might be influential over time. Gravitational pull might be modelled using the simple equation of Newton's law (point-mass gravity), or by using the more complicated spherical harmonic models. Apart from gravitational perturbations Solar radiation pressure might be a considerable influence on the orbital evolution, this will be modelled as described in subsection 5.4.3. Throughout the next chapters the following notations will be used to distinguish these perturbation source models:

- PM: Source is modelled as a Point Mass, with gravitational pull from Newton's equation.
- SH: Gravitational pull modelled as a Spherical Harmonic model originating from the source. Degree and order will be denoted by (D/O).
- RP: Radiation Pressure modelled with this body as origin.

Using these notations, Table 5.1 shows the modelled bodies in the system, along with associated perturbation models which will be investigated later in section 5.6:

Table 5.1: Solar system bodies and associated perturbation models.

Model	Earth	Moon	Sun	Mercury	Venus	Mars	Jupiter	Saturn	Neptune	Uranus
Satellites	PM, SH ¹	PM, SH ²	PM, RP	PM	PM	PM	PM	PM	PM	PM

¹ EGM96 model [36]

² LPE200 model [43]

This table does not represent the final environment model used for swarm geometry optimisation, this will be determined after an extensive analysis of the effect of individual perturbation sources. While a full perturbation model might be the most accurate, it will also be very computationally expensive. Hence, is why it is preferred to leave out the least-influential perturbation sources. The next section will first introduce the methods used to model solar radiation pressure on the OLFAR elements, before the effect of perturbation sources is investigated.

5.4.3 Modelling solar radiation pressure

The satellite designs proposed for the OLFAR mission are lightweight at 5 kilogrammes, but equipped with very large solar panels with a surface area of 0.31 m^2 . This combination makes the satellites particularly susceptible to orbit perturbations from solar radiation. The method used to model solar radiation is based on the methods used by Marshall et al. [39] for the TOPEX/Poseidon satellite. Equation 5.1 is adjusted from the work presented by Marshall et al., dividing by the speed of light to convert it to an expression of force [39].

$$F = C \frac{G_{sc}}{4\pi R^2 c} A_i = CP_R A_i \quad (5.1)$$

Variable	Description	Unit
C	Radiation pressure coefficient	-
G_{sc}	Solar radiation constant	W/m^2
R	Distance from the sun	-
c	Speed of light	m/s
A_i	Incident surface area	m^2
P_R	Radiation pressure	Pa

This is modelled internally through the cannon-ball radiation interface in TUDAT, which automatically takes the distance to the radiation source into account, as well as occultation from the Earth. All that remains is a model for the exposed surface of these conceptual satellites, which will be based on the schematics (Figure 2.8a) proposed by Quillen et al. [49]. Throughout the propagation it will be assumed that the solar panels will be pointed at the Sun. This introduces a base surface area of 0.31 m^2 . The central body adds a small additional surface area which can range between 0.01 m^2 and 0.0316 m^2 depending on the relative angle of the body. The propagations will by default assume the worst-case surface area to investigate the effect of solar radiation pressure on the swarm element orbit. Figure 5.1 shows the two extremes in angling of the body relative to a light source, which will further be referred to as full and minimal body pointing.

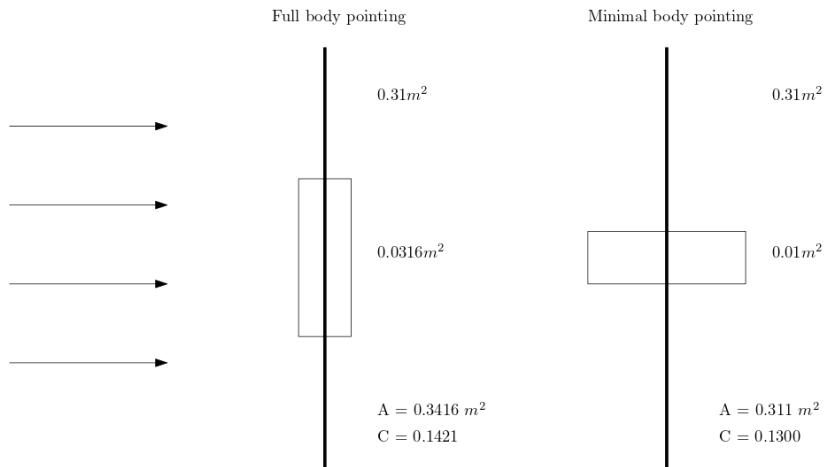


Figure 5.1: Definition of full and minimal body pointing, with respective surface area and reflectivity

In addition to the satellite geometry, the reflectivity of the satellite is a key parameter which needs to be determined to model the radiation pressure. The reflectivity is represented by the radiation pressure coefficient C . A fully reflective surface ($C = 2$) will experience stronger radiative accelerations than a matte surface ($1 < C < 2$). Surface reflectivity is a combination of δ specular (mirror-like) and ρ diffuse light reflection, and the overall reflectivity can be approximated for incident angle α using [39]:

$$r = \frac{\delta}{3} + \rho \cos \alpha$$

The surface area of the satellite predominantly stems from the large solar arrays, which are by design meant to absorb as much light as possible. Since no concrete concept for the OLFAR elements has been developed yet, its surface properties will be estimated using the values for the TOPEX/Poseidon satellite in the work by Marshall et al. as reference [39]. It will be assumed that the surface properties of the OLFAR element will have similar reflective properties. Using diffuse and specular values used by Marshall et al. it can be computed that solar panels have an approximate reflectivity of 0.123 when viewed head-on. The body itself will be assumed to be similar to Poseidon's X+ body plate model, yielding a head-on reflectivity of 0.326. To calculate the combined reflectivity of the OLFAR element surfaces an average surface reflectivity is used:

$$r_a = \frac{r_{SA} * A_{SA} + r_B * A_B}{A_{SA} + A_B} \quad (5.2)$$

This yields an average reflectivity which will range between 0.1421 at full body facing, and 0.1300 at minimal body facing. This yields radiation pressure coefficients of 1.1421 and 1.1300 respectively. These results show an interesting relation between the pointing of the central satellite body and the solar radiation pressure. When the body is pointed for maximum surface area (full body pointing) the surface area and the average surface reflectivity increase, strengthening the effect of solar radiation perturbations. Due to the low mass of the satellites this effect could be useful for passive course corrections between constellation members. The obtainable results from body-pointing methods will further be discussed in section 7.2.

5.5 The necessity of modelling perturbations

In section 4.2 an analytical solution to the three body problem in the barycentric frame was presented, as well as a derivation of this method to propagate baselines using the three-body model. The three-body model and derivative equations of motion are very useful for understanding the problem topology and getting an estimate of the relative motion of objects in orbit. The model is built upon a very bold set of assumptions however, which make it hard to reliably apply to long term real-world applications. This becomes evident when the propagated orbit of a three-body model is compared to an orbit propagated with a full perturbation model. The motion of two satellites with identical starting conditions will be propagated using the equations of motions presented in Equation 4.4, as well as a model with external propagations as described in section 5.8. Both orbits are integrated over a year using RK4 integration with 30 minute time steps. Figure 5.2 shows a comparison between the year-long orbits, in which a significant difference may be observed due to the influence of perturbations over time.

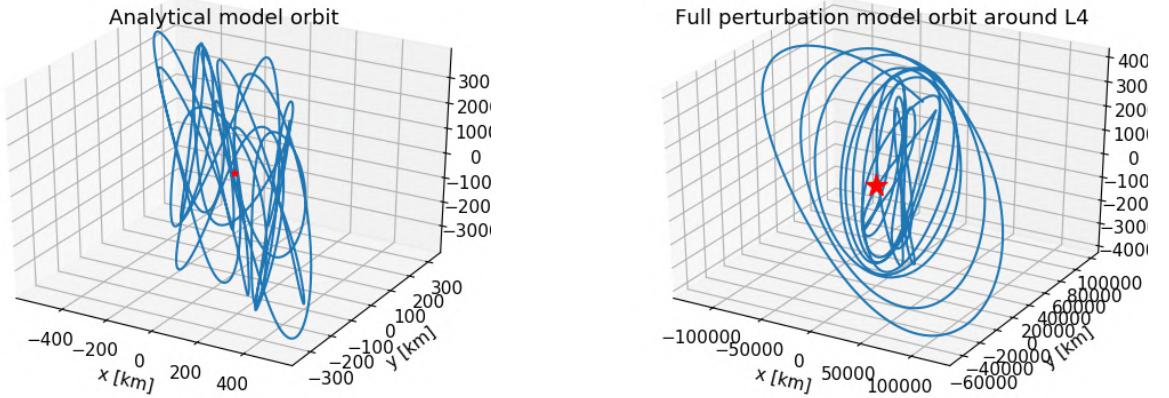


Figure 5.2: Comparison of the barycentric orbit profiles relative to the L4 point for a 3 body model and an orbit model with external perturbations.

It can be seen that the primary influence of perturbations affects the development of the orbit on the x_{BYB} plane in the barycentric frame, where the full orbit model shows a movement range which is two degrees of magnitude larger than the analytical model. The effect in the z_B axis on the other hand is minor, showing very little relative growth.

Figure 5.3 shows the magnitude of the position error of the 3-body model orbit compared to the full perturbation model over the first few days of propagation. It can be seen that there is a significant difference in the global position propagation of both satellites, and the relative position propagation. The baseline magnitude error is several orders of magnitude smaller, showing an error smaller than a meter after an hour.

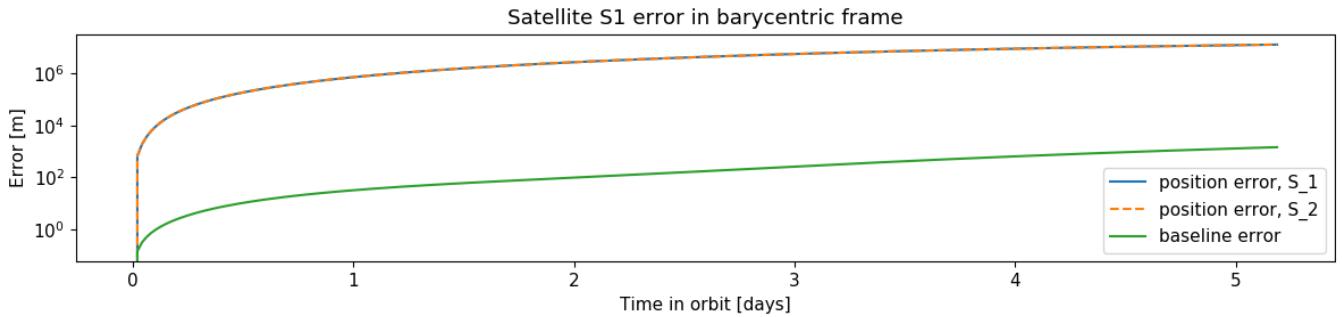


Figure 5.3: Position error of satellites S_1 and S_2 in the 3-body propagation compared to the full-perturbation model.

For propagation of overall positions the accuracy of the three body model is noticeably worse, with errors ranging in several hundred meters after the first half-hour timestep. This error grows over time, resulting in entirely different orbits after a few days. It needs to be concluded that for long-term propagation it is necessary to use a perturbed model, and the three-body model is only useful for very short-term propagation (e.g. collision avoidance). One important observation is the increased accuracy for baseline magnitude prediction for the three-body model however, which might prove very useful for predicting baselines during observation with a computationally cheap model. The rest of this chapter will be dedicated to investigating and showing the effect of individual perturbations, and selecting which to include in the final perturbation model.

5.6 Analysis of perturbing forces

During the OLFAR mission the orbit of the satellite swarm will be influenced by varying perturbing forces, such as the irregularity of the Earth's gravitational pull or the gravitational pull of other celestial bodies. In an ideal numerical model all of these perturbations would be fully modelled, but in practice this is not feasible due to the large computational cost that is associated with a full model. In the end the model is supposed to be used in an iterative optimisation algorithm, which will require small computation times to run efficiently. Thus, the effect of the different perturbations will be investigated for their actual effect on the evolution of the swarm's orbit. By doing so a conscious choice may be made about which perturbations to include within the model, allowing for a balanced choice between model accuracy and computational cost.

To investigate the effect of individual perturbation sources on the evolution of a satellite orbit a full-perturbation model will be used to generate a reference orbit. Individual perturbation models will then be removed from the environment model and the resulting change in orbit compared to the reference will be used to study the influence of separate perturbations.

5.6.1 Standard reference model

To investigate the effect of different perturbing forces a full environment model with perturbations as described in subsection 5.4.2 will be used to propagate the orbit of a satellite pair with a known initial spatial distribution. The reference model will consist of a pair of satellites placed into a orbit close to the L4 point. Instead of opting for a small-scale halo orbit the initial conditions will be chosen to ensure a wide, sweeping orbit around the L4 point over the course of a full year (see Figure 5.4). This exposes the satellite pair to a wide variety of dynamical conditions. During propagation of the reference orbit the satellites will be exposed to a highly perturbed environment, Table 5.2 contains all perturbation sources modelled to generate the reference orbit.

Table 5.2: Perturbation sources per body included in the reference model

Model	Earth	Moon	Sun	Mercury	Venus	Mars	Jupiter	Saturn	Neptune	Uranus
Satellites	SH(5,5)	PM	PM,RP	PM	PM	PM	PM	PM	PM	PM

The position of the planetary bodies will be propagated over the course of a year using the model described in subsection 5.4.1. All of the perturbation sources presented in Table 5.2 will be studied for their effect on the orbit, as well as higher-order degrees and orders for SH models of the Earth and Moon.

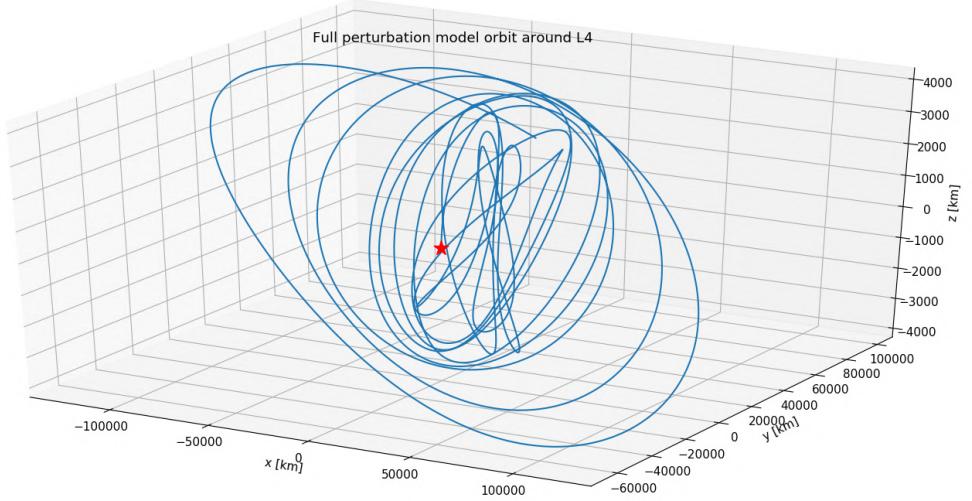


Figure 5.4: Orbit of the reference model around the L4 point, denoted by a red star, during the first full year. Shown in the barycentric frame.

Propagation and integration settings

Every iteration of the model will be propagated using a Cowell propagator and a Runge-Kutta 4 (RK4) integrator with a fixed step-size of 30 minutes. The motion of the two satellites will be simulated for the duration of a year under influence of the perturbations listed in Table 5.2.

Initial conditions of satellites

The investigation model will be using a pair of satellites (S_1, S_2), which allows studying changes in both relative motion and global positioning. Satellite S_1 is placed on top of the L4 point at ephemeris 1 January 2030, while the second satellite S_2 is placed at controlled distance R from the first satellite into the x direction (Following the J2000 frame). Both satellites start with the initial orbital velocity of the L4 point, in addition with a small offset velocity of 11 m/s in the z -direction of the J2000 frame. Table 5.3 summarizes these initial conditions, these insert the satellites into a broad orbit surrounding the L4 point.

Table 5.3: Initial conditions of satellites relative to the L4 point, expressed in the J2000 frame.

Satellite	x	y	z	V_x	V_y	V_z
S_1	0	0	0	0	0	11
S_2	R	0	0	0	0	11

To investigate the effect of different initial baselines a range of relative distances R will be studied as well. Initial separations of 1, 10, 30, 50, 70 and 100 kilometres will be used for the analysis. The largest initial separation will be used to determine which perturbations to retain, since this represents the worst-case scenario for perturbations affecting the evolution of baselines. The selection of which perturbations to retain will be based on the thresholds discussed in section 5.3

5.6.2 Influence of perturbation sources

In order to investigate the influence of perturbation sources on satellite positions and baselines over the duration of a year in orbit the full reference model as described in subsection 5.6.1 will be used. Single perturbation sources will individually be removed, after which the orbit will be propagated again with the same initial conditions and settings. In case of spherical harmonic gravitational models the effect of changes in modelled degree and order (D/O) will also be investigated using this method. Figure 5.5⁴ shows the effects that changes in the perturbing environment had on the evolution of the baseline between the reference satellites over the duration of one year.

⁴Generated using the adjustText library for Python [25]

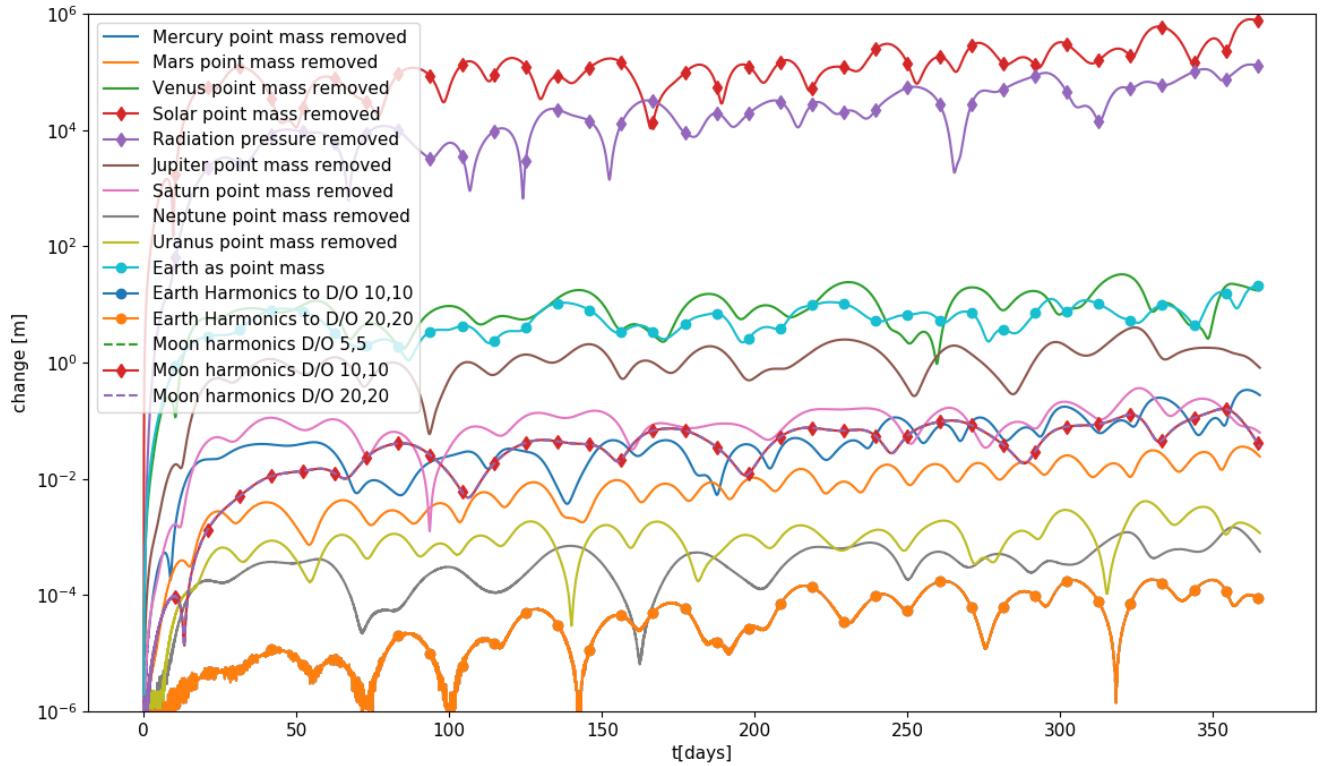


Figure 5.5: Changes in baseline originating from perturbation sources over time, for a pair of satellites in orbit of the L4 point

Within this data a few perturbation sources immediately stand out. It was already expected that Solar radiation pressure would be very influential, but it is surprising to see that it (at times) rivals the effect of Solar gravitation in magnitude. By themselves both Solar perturbations eclipse the other perturbations sources in their effect, showing that both are absolutely necessary to model L4-centric orbits. Figure 5.6 shows the maximum observed changes in both positional data and relative baseline of the satellites for all perturbation sources.

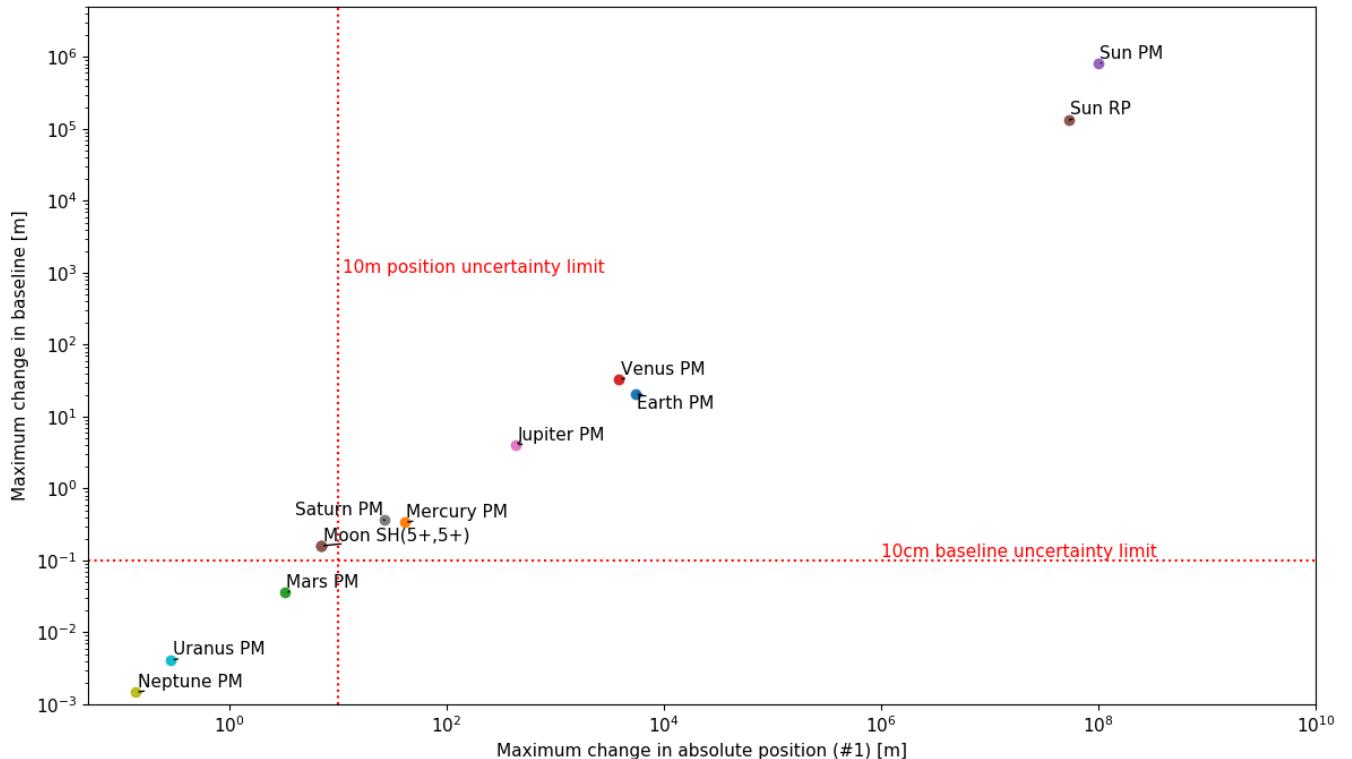


Figure 5.6: Maximum baseline changes over maximum positional change of satellite S_1 .

An interesting detail within this data is the near-linear trend shown by the perturbation sources when the effect on relative movement (baselines) is plotted against absolute position change. In all cases the magnitude of the effect on relative motion (baselines) is generally two orders of magnitude smaller than the effect on global positioning. This shows that the relative motion is slightly independent of the accuracy of the satellite positioning, the satellites show very similar baseline motion even if displaced by a few hundred meters. This result reinforces the notion in section 5.3 that the uncertainty in both baseline and global position may be evaluated using different standards. Table 5.4 shows the data displayed in this figure. At a first glance it can be seen that the effects from the gravitational pull of Mars, Neptune, and Uranus are well below the determined uncertainty requirements, making them candidates for exclusion from the model. This topic will be discussed in section 5.8.

Table 5.4: Perturbation sources and their maximum magnitude of effect in meters on orbital position and baseline during a year in orbit. "Effect" measured as maximum deviation from the reference orbit in the full-year period.

Source Type	Earth SH ¹	Moon SH(5,5) ²	Sun PM	Sun RP	Mercury PM	Venus PM
Effect on baseline	20.84	0.1593	8.0157×10^5	1.336×10^5	0.338	32.90
Effect on position	5.454×10^3	6.962	9.920×10^7	5.416×10^7	41.20	3.865×10^3

Source Type	Mars PM	Saturn PM	Jupiter PM	Uranus PM	Neptune PM
Effect on baseline	0.03578	0.3629	4.008	1.471×10^{-3}	4.121×10^{-3}
Effect on position	3.255	26.98	435.3	0.2913	0.1378

¹ The change to a PM model shows the influence of the low D/O SH effects, not the entire exclusion of Earth gravitation (Which is vital for an Earth-centric orbit).

² And higher order models, which show no discernible additional effect in this data.

Figure 5.6 does not show the full range of investigated perturbations, since some inclusions cannot be visualised as clearly. This is the case with different orders of SH models, which instead will be discussed in subsection 5.6.4. Note that the reference model includes the Earth's gravitational pull as a 5th order spherical harmonics model, and the Moon as a point mass. Unlike with the other bodies markers pertaining the Earth or Moon indicate changes to a different model, rather than removal. The marker "Earth PM" represents a downgrade from a (5,5) SH model to a point-mass model, making it best suited to demonstrate the effect of the lower degrees and orders.

Venus has a surprising amount of influence on the orbit, surpassing even the influence of the J_2 effect and other low-order harmonics of the Earth. This is not in line with what one would instinctively expect based off its mass and distance. The source of the strength of Venus's perturbation is better explained once the initial conditions of the solar system are considered. Figure 5.7 shows the initial placement of the inner solar system bodies on the first of January 2030. It can be observed that the inner planets are aligned, resulting in a strong initial perturbation from Venus.

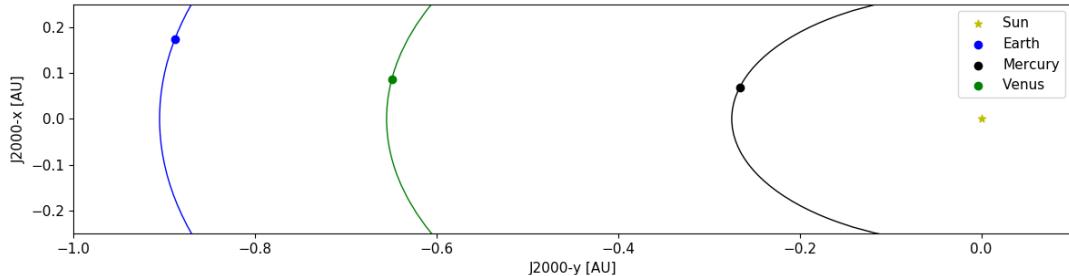


Figure 5.7: Position of the interior planets on the first of January 2030, showing the close proximity of Venus. Orbits appear smaller due to xy-projection.

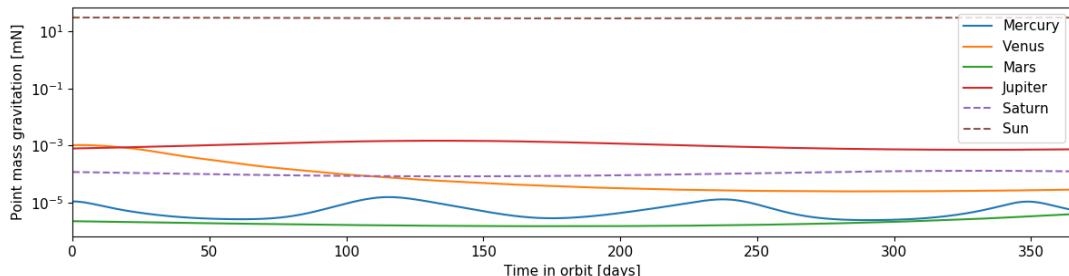


Figure 5.8: Magnitude of gravitational pull of solar system bodies on a 5 kg satellite placed at the L4 point.

Figure 5.8 shows the magnitude of the gravitational pull of the solar system bodies on a 5kg satellite, fixed to the L4 point. Within these results it can be observed that the initial gravitational pull of Venus is slightly larger than that of Jupiter, which occurs during a critical part for the long-term development of the orbit. The effect of these initial perturbations particularly develop over the full duration of the orbit, which provides some explanation regarding the apparent severity of Venus's gravitational pull. In combination with the early resonance this makes Venus a considerable perturbation source for long-term orbit development.

5.6.3 Resonance of perturbations

One reason which may amplify the effect of one perturbation over the other could be the presence of resonance. With resonant perturbations the variation in magnitude or direction coincides with the natural orbital frequencies of bodies, greatly increasing the effect of these perturbations on orbits. An example of resonance can be found in the satellite moons of Jupiter, where Io, Ganymede and Europa are known to be in resonance-locked orbits [65]. Figure 5.9 shows the normalized coordinates of satellite S_1 's orbit relative to L4 from the reference case, as well as normalized components of relative direction vectors from L4 to Venus and Jupiter. Easily visible in these plots are the long-term spatial movement frequencies of 28.62 days computed by Wakker[60], which were discussed in section 4.5.

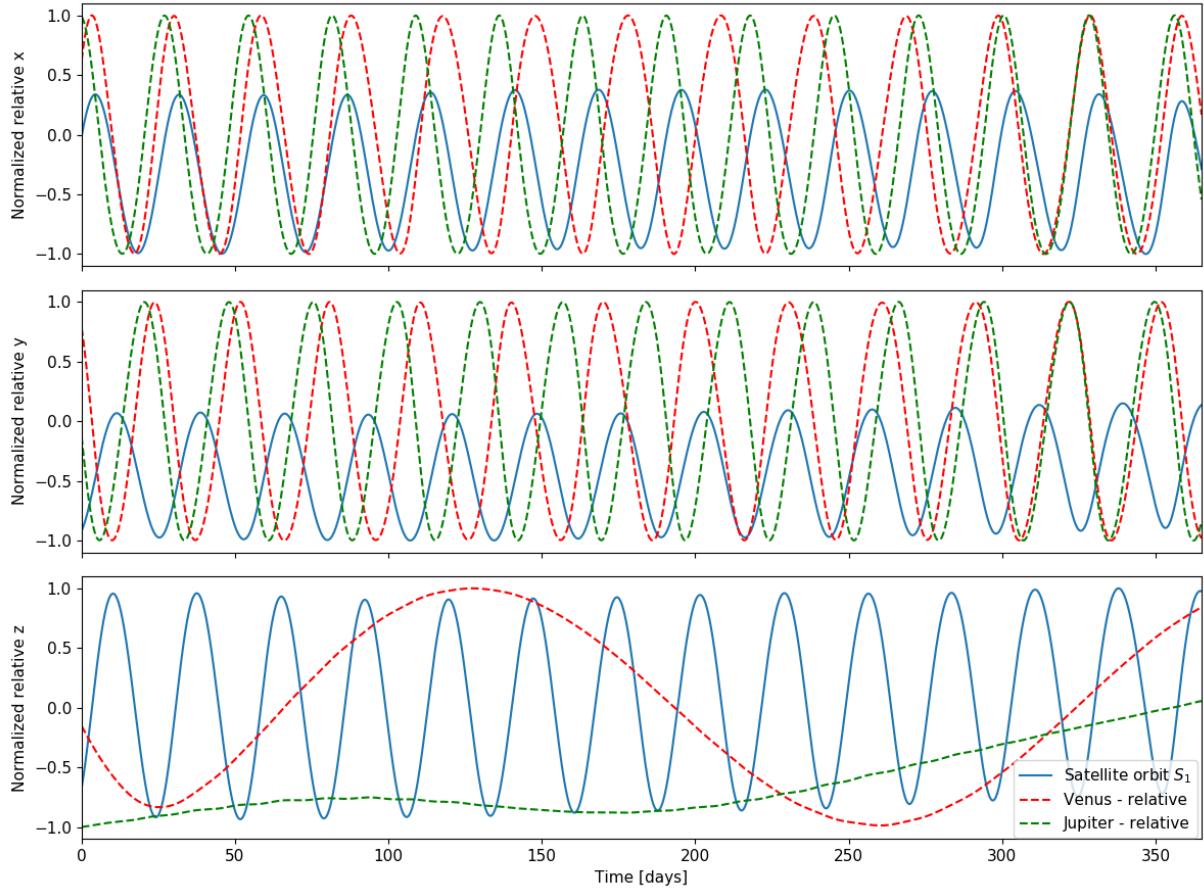


Figure 5.9: Relative positions of the satellite S_1 , Venus, and Jupiter to the L4 point. Coordinates normalized for visibility, J2000 coordinate frame.

In this figure it can be seen that Venus is particularly resonant with the x-motion of this satellite orbit during the first three months of the mission. In this case the angular velocities of Venus, the Earth and the L4's rotation around the Earth line up well to create short periods of resonance with the chosen reference orbit. Considering that this also happens very early in the modelled period, it explains why the absence of this perturbation yields such significant changes in baselines and positions down the line. The cause and scale of this effect are naturally very reliant on the properties of the satellite orbit around L4, but the possibility of this effect shows a clear need to model Venus for these types of orbits. Resonance might also occur with other planetary bodies, but the effect with Venus is strongest due to its proximity.

Typically orbital resonance is not considered for these short time frames since the effects are negligible over short times for most bodies. For long term orbit design and evolution, third body resonance might still be influential for constellation design around the L4 point. Using statistical modelling, it is predicted that the average lifetime of a single OLFAR element can range from 2.8 to 4 years, yielding mission lifetime of up to 48 years with 100 elements [23]. With this vision, it will become important to look at other resonance relations as well for long-term mission planning.

5.6.4 Degree and order of harmonic gravitational models

The selection of the degree and order for the harmonic gravitational models has up to this point been rather arbitrary, three sample points have been taken at degree and order (5,5), (10,10), and (20,20) but this still leaves a lot of ambiguity regarding to what is minimally required. To investigate this the convergence of the effect of degree and order for both the Earth's and the Moon's models will be studied. Table 5.5 describes the adjusted reference environment that will be used to map the effects of the degrees and order of SH models. The effect of the lowest D/O will be measured against an orbit propagated exclusively PM models, and higher order SH models will be compared against the lower order SH orbits. The induced changes from the different SH models are shown in Figure 5.10.

Table 5.5: Perturbation sources per body included in the reference model for spherical harmonics

Model	Earth	Moon	Sun	Mercury	Venus	Mars	Jupiter	Saturn	Neptune	Uranus
Satellites	PM	PM	PM,RP	PM	PM	PM	PM	PM	PM	PM

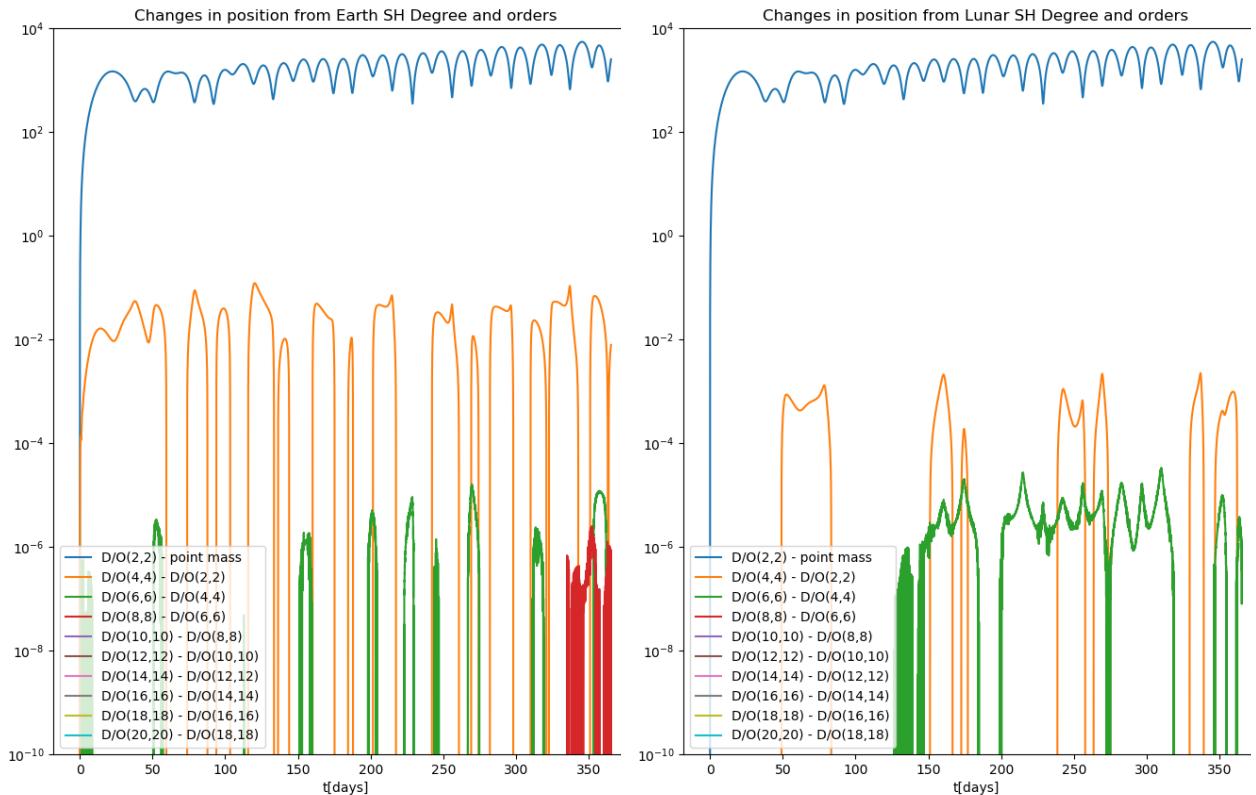


Figure 5.10: Changes in the propagation of the position of satellite S_1 over the duration of one year, caused by increases to the degree and order of the Earth's and Lunar harmonic modelling. Y-axis shows difference between modelled orbits in meters.

Figure 5.10 shows that for both modelling the Earth and the Moon's spherical harmonic gravitational field the use of D/O (6,6) is sufficient. At these D/O values the resulting orbits converge below millimetre levels over the course of one year, which is well below the used uncertainty thresholds which were established previously. Including higher order models is, bluntly put, a waste of computational resources considering that the difference will go unnoticed.

This conclusion, like any other conclusion made in this chapter, needs to be reconsidered when this method is applied for other deployment locations. In particular for Lunar or Earth orbits there might be a considerable difference in between the higher order D/O models, but this is not the case for L4-centric orbit modelling.

5.7 Analysis of integration settings

Besides figuring out *what* to include in the orbital model, a secondary question is *how* the model should be propagated. The choice of integration techniques affects both the overall speed of the process, as well as its accuracy. Figure 5.11 demonstrates that a more complex integrator yields very different orbits for a simple three-body model, and it can be expected that this is even more relevant in perturbed environments. Up to this point a Runge-Kutta 4 propagator was used along with a Cowell integrator. This combination is reliable, but hardly cost-efficient and it requires a few seconds to simulate the orbit of the pair of satellites. If this system is to be upscaled to include tens of satellites in an optimisation loop there are significant benefits to finding a more efficient solution for the computations.

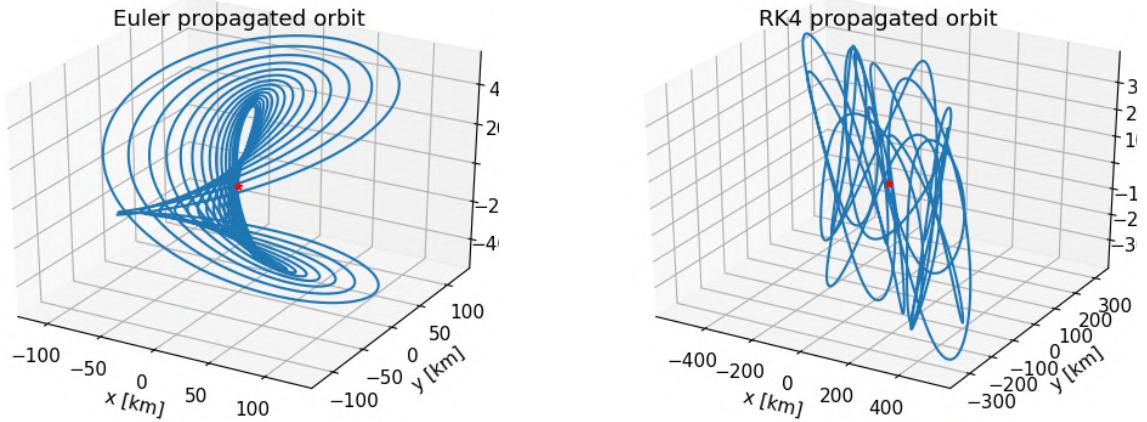


Figure 5.11: Difference in orbit propagations of the 3-body model using Euler(L) and RK4(R) integration.

The largest improvement in computational efficiency will stem from the use of variable step-size integrator methods. These methods seek to estimate the proper step-size to be taken based on the dynamics at the current time, and they are well-suited to environments with stable dynamics. The use of different propagators would change the way the problems are framed, which could serve to further reduce the complexity of the required computations. This change will be small compared to the choice of integrator however, which is why the integrator will be investigated firstly.

5.7.1 Standard reference model

To keep the analysis of integrator settings as relevant to the optimisation loop as possible, the analysis will make use of the environment model as described in Table 5.2. Like in the previous section, the orbit of a pair of OLFAR elements with an initial spread of 100 kilometres will be propagated for the duration of a year using different integration settings. The reference orbit will be generated using a Runge-Kutta 4 (RK4) integrator with Cowell propagation. The output data will be interpolated to 4 hour intervals using 8th order Lagrangian interpolation, which allows for easier comparison of variable step size orbits. The resulting orbits from the variable time-step methods will be interpolated to the same timeframe as the reference using 8th order Lagrangian interpolation.

Integration settings

The full spectrum of available integrator techniques to the TUDAT toolset will be investigated in this section. Table 5.6 shows this range of integrators along with the settings that will be used for the analysis.

Table 5.6: Investigated integrator techniques along with used settings. NA: not applicable, x-y: investigated range

Integrator	Step size	Min. step size	Max. step size	Tolerance	Order(min,max)	Max. steps
RK4	1 - 15 minutes	NA	NA	NA	NA	NA
RKF45-78	NA	1 s	24 hr	$10^{-6} - 10^{-12}$	NA	NA
DOPRI87	NA	1 s	24 hr	$10^{-6} - 10^{-12}$	NA	NA
ABM	NA	1 s	24 hr	$10^{-6} - 10^{-12}$	(6,11)	NA
BS	NA	1 s	24 hr	$10^{-6} - 10^{-12}$	NA	8

In order to evaluate the performance of these integrators they will be used to integrate the motion of the pair of satellites as described in the standard reference model for a range of step sizes or tolerance settings. By comparing the evolution of the difference relative to the reference orbit, as well as the gains in computation time the suitability of these integrators will be evaluated. Figure 5.12 shows the performance of the different integrator setting combinations compared to the reference orbit.

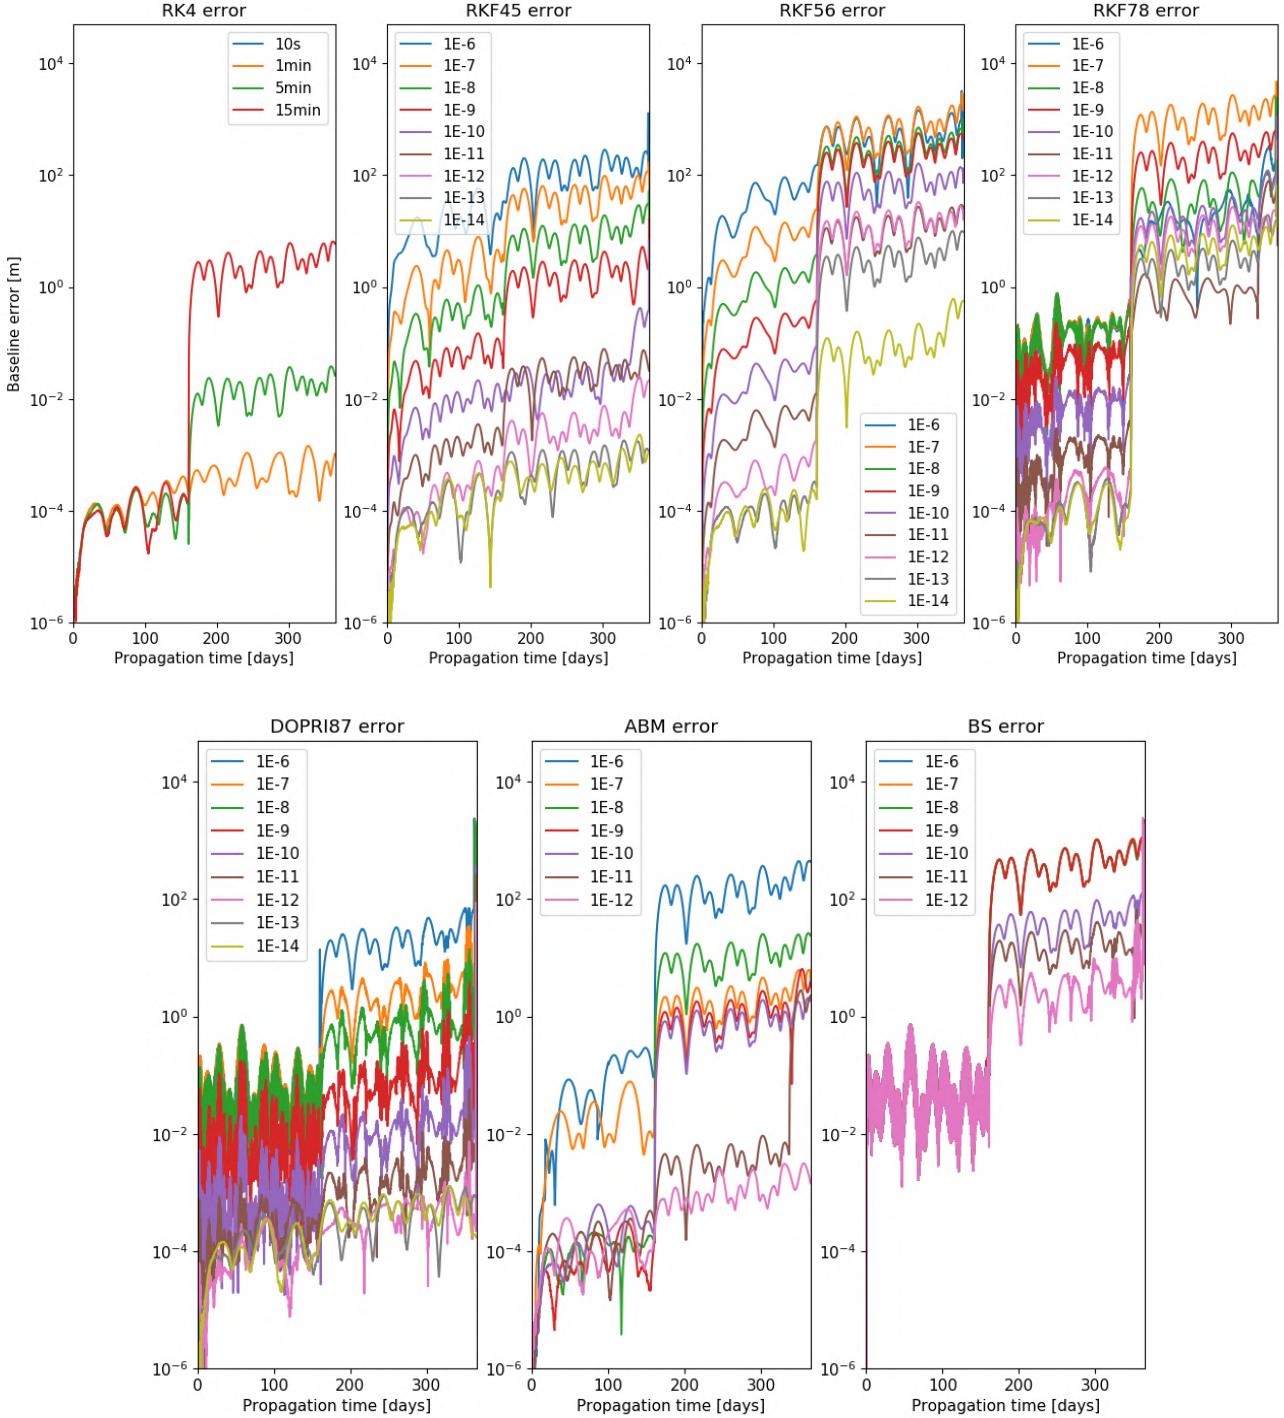


Figure 5.12: Propagation errors in the baseline magnitude compared to the reference orbits from different integrator techniques and settings over time.

It can be observed that many of the algorithms encounter a major error around 170 days into the propagation. The cause of which is likely to be the allowance of too large a timestep during a critical part of the orbit, misplacing the satellites into a state belonging to a very different orbit. This is evident from the fixed step size RK4 data, where the same error occurs only for step sizes of 5 and 15 minutes. Only the RK45 algorithms and strict tolerance ABM integrators manage to avoid this misstep, from which the RKF45 integrator shows the most promise due to the wider array of usable settings.

In addition to the position error compared to the reference orbit a second parameter of importance is the time required to achieve this result. After all the main purpose of the variable time-step methods is to estimate the true orbit more efficiently, both computationally and time-wise. Figure 5.13 shows the final error in baseline magnitude after a year in orbit compared to the time required for the computations on an Intel i7-7700 CPU.

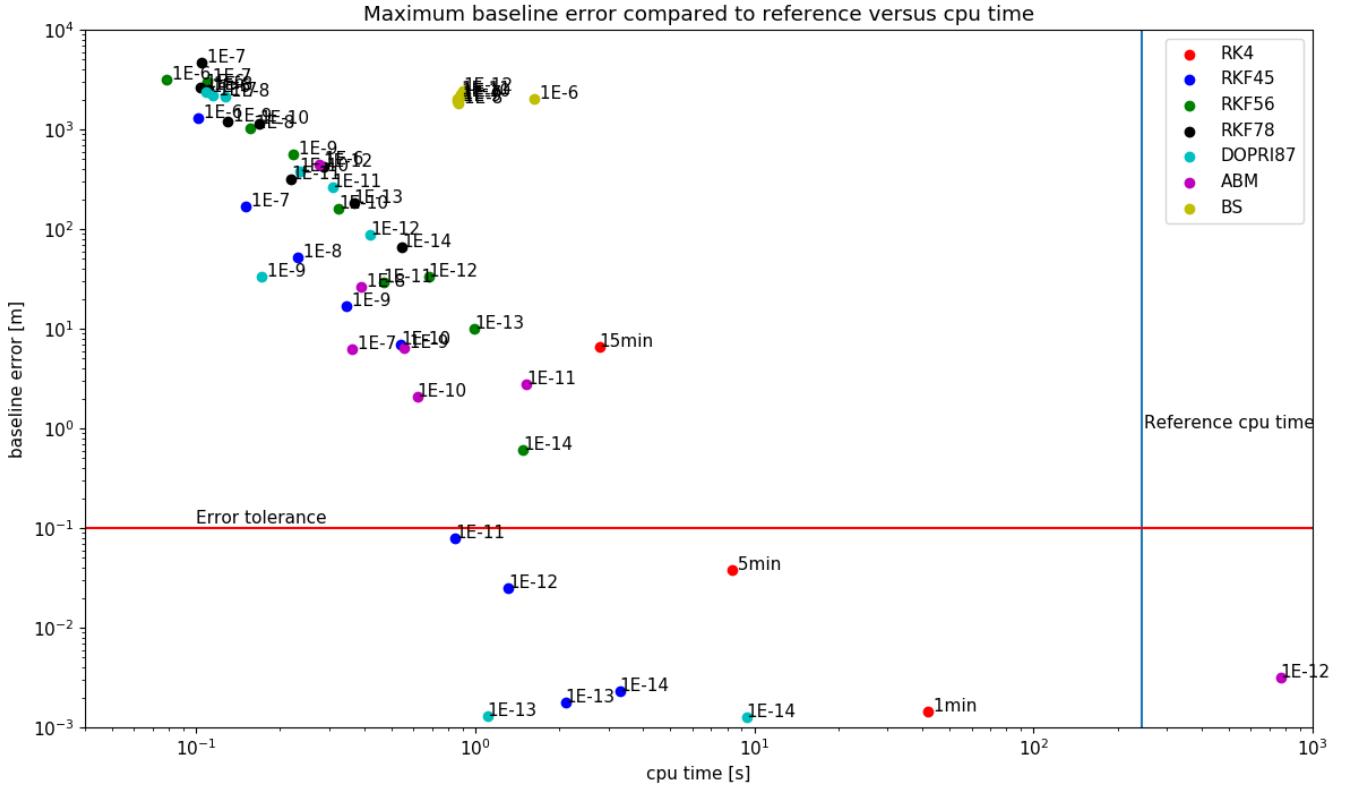


Figure 5.13: Maximum baseline error during propagation and required time to run a 1 year propagation. Reference cpu time was 243.7 seconds using RKF4 integration with 10 second time steps.

It can be observed that the general trend is what one would expect, lower tolerance settings improve the propagation time at a cost of accuracy. A notable exception is the subset of Bulirsch-Stoer integrators, which are clustered tightly regardless of tolerance settings. Uncaring for the settings, this integrator yielded near-identical results both for the propagated error and computational time. This highlights either a general unsuitability of this algorithm for this type of orbit integration, a faulty implementation, or set of instructions. Since no fault could be located in the code, and there are plenty of suitable alternatives, it is decided not to investigate the BS integrator further.

5.8 Final environment model

Based on the baseline error relative to the reference orbit and the cpu time the RKF45 propagator clearly offers the best performance at low tolerance settings, closely followed by the DOPRI87 integrator. Following these results it is chosen to use the RKF45 integrator with a tolerance of 10^{-11} , this setting offers the best computation time below the 10cm error threshold. All that remains is a conclusion on the perturbation models which will be included. In section 5.3 a set of acceptable inaccuracies was defined, when these are applied to the results shown in subsection 5.6.2 the perturbations originating from Mars, Uranus and Neptune may be excluded from the model on an individual basis. Even though the option of mutual exclusion is obvious from the magnitudes of these errors, the combined effect will be calculated to be thorough. Since these perturbations come from very different points in space the cumulative effect will be estimated using a propagation of uncertainty:

$$RMS_{baseline} = \sqrt{(0.03578)^2 + (1.471 * 10^{-3})^2 + (4.121 * 10^{-3})^2} = 0.0360m \quad (5.3)$$

$$RMS_{position} = \sqrt{(3.255)^2 + (0.2913)^2 + (0.1378)^2} = 3.271m$$

Considering that both these values are below the set requirements, it is concluded that these perturbations can be excluded from the model environment without breaching the set desired of 10cm baseline uncertainty, and the 10m position uncertainty. After omitting these sources, and applying the results of the subsection 5.6.4 a selection for the final environment is made. Table 5.7 contains the perturbation models used in the final environment.

Table 5.7: Selection of perturbation sources to be included in the final model for satellite propagation

Model	Earth	Moon	Sun	Mercury	Venus	Mars	Jupiter	Saturn	Neptune	Uranus
Satellites	SH(6,6)	SH(6,6)	PM,RP	PM	PM	-	PM	PM	-	-

6 Optimisation strategy

After establishing the problem and the numerical methods and environment, all that remains to discuss is the optimisation strategy which will be used for the constellation design. This chapter will introduce the optimisation problem derived from the constellation design, the scope of the problem, and discuss the most efficient method of resolving such a problem based on experiments with a smaller scale reference problem. This chapter is dedicated to finding an efficient and scalable optimisation strategy for the OLFAR swarm problem.

Terminology: Optimisation

Within this chapter it is assumed that the reader has some basic understanding of genetic algorithms and their operation. This chapter will rely on terminology which is used by the Pagmo2 library to describe optimisation algorithms and architectures, which has proven to be very convenient to use for the discussion of optimisation architectures. For the convenience of the reader this terminology will briefly be addressed in this section beforehand.

Within the Pagmo2 package a single generational optimisation algorithm is usually referred to as an *island*, which contains its own set of candidate solutions named the *population*. Every island will have its own individual population groups of a set size, which it evolves through generational algorithms. A group of islands might work together to solve a single problem by distributing the workload, such a group is referred to as an *archipelago* of islands. Islands may exchange population members between each-other through the archipelago's *topology*. The topology describes the network within an archipelago through which the islands can communicate or migrate population members. An easy visualisation of this concept is to imagine a set of bridges connecting different islands, allowing the population to migrate over the archipelago between generations.

A term which will also be used often is the *solution space*, this refers to the M -dimensional space in which the optimisation algorithm is seeking a solution, wherein M is the number of variables in the optimisation problem. The solution space represents the space with all viable solutions, which is bound by a set of constraints. Due to the large amount of variables with the problems introduced in this chapter, the solution space will not directly be pictured in this chapter.

Terminology: Swarm orbits

The orbit of interferometry swarms will be defined on two levels, global and local. The *global* orbit refers to the orbit of the swarm as a single entity, it is used to describe periods of motion at scales where the movement of individual swarm members cannot be discerned. The *local* orbit refers to the movement of swarm elements relative to the core of the swarm.

6.1 Defining the optimisation problem

The primary goal of the optimisation algorithm will be to search a swarm geometry for N satellites and an accompanying orbit around the Earth-Moon L4 point that will be suitable for the OLFAR mission concept. In this case this means relying on passive formation flight whilst being stable for at least a year in orbit. As per the mission requirements for the OLFAR swarm, this means that satellites may not drift apart further than 100 kilometres, while also not approach each-other within 500 meters for collision avoidance. Accounting for the position of the satellite relative to the L4 point, its velocity, and acceleration this yields an optimisation problem with $9 * N$ variables and a very broad solution space. Especially for large swarms this problem scope is promising to be a problem, which is why attention is given to methods that allow for reducing the problem scope while retaining swarm size.

The primary method to increase the speed at which this problem can be resolved is to reduce the scope of the problem directly through simplifications. By using some simplifications in the description of the problem the number of variables can nearly be halved, as well as have their solution space reduced. A major simplification of the constellation design problem is to assume that at the initial time all swarm satellites share an identical velocity, and have no inherent acceleration. This reduces the number of variables from 9 (position, velocity, acceleration) to 3 per swarm element. This simplification does exclude some possible configurations with better long term stability, which will need to be sought after an initial configuration is found using a secondary optimisation. The position of the satellites will then be described relative to a central "core" position for the swarm P_c , this allows for the use of a much smaller solution space for a majority of the variables.

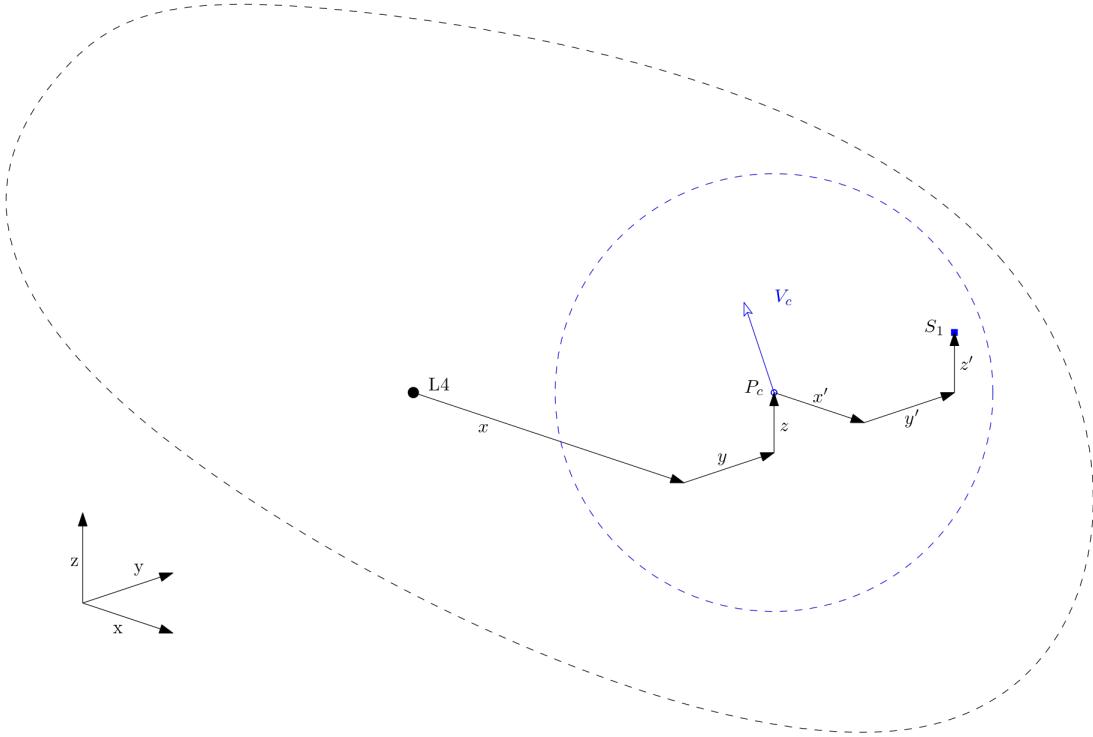


Figure 6.1: Definition of the swarm geometry design for optimization: The swarm core has position vector P_c and it is given an initial velocity V_c . Satellite elements inherit this velocity, and only require a position vector relative to the core. Dashed black and blue lines represent placement boundaries for the core and satellite respectively, they are not representative in shape or size.

Using this problem description the number of variables is scaled back to $M = 6 + 3 * N$, which will result in scope reductions of 48 % and higher for swarm problems involving more than 20 satellites compared to full-state vector optimisation with $M = 6 * N$ variables. This reduction of scope is expected to drastically increase the convergence time for large-scale optimisation problems, which might be too complex to solve within a decent time-frame otherwise.

6.1.1 Cost function

In subsection 3.6.2 a cost function based on the point-spread function of a radio interferometer was proposed to be used in combination with boundary functions enforcing the baseline constraints. Unfortunately, this idea could not be implemented in the final optimisation setup due to technical challenges integrating new libraries into the TUDAT environment. Foregoing the direct implementation of the PSF into a cost-function, the cost-function for optimisation will be based entirely on the set of baseline requirements. The loss of the PSF integration into the cost function is regrettable, but it is not a vital parameter for the design of a passively stable swarm geometry.

The set of requirements for the baseline formulated in Table 3.1 will be enforced as soft boundary constraints by penalizing infractions. Infractions on these conditions can be categorized to three different scenarios: a near-collision (NC) event when the baseline is below the minimum requirement of 500m, a wandering (W) satellite when a baseline is larger than the allowed 100 km, and a breach of the allowed relative velocity (V). The cost function will evaluate the magnitude of all baselines during the orbit, as well as their relative velocities. The cost of an orbit is based on the sum of the number of times when an infraction is measured, where baseline measurements are taken at 4 hour intervals¹. Equation 6.1 describes this function mathematically as the sum of the evaluated number of infractions.

$$C = N_{NC} + N_W + N_V \quad (6.1)$$

For an orbit spanning between times t_0 and t_e , with N_R individual baselines, the number of infractions is measured as sums of every constraint breach during the entire orbit:

$$\begin{aligned} N_{NC} &= \sum_{t=0}^{t=t_e} \sum_{R=0}^{R=N_R} (R < 500m) \\ N_W &= \sum_{t=0}^{t=t_e} \sum_{R=0}^{R=N_R} (R > 100km) \\ N_V &= \sum_{t=0}^{t=t_e} \sum_{R=0}^{R=N_R} (\dot{R} > 1m/s) \end{aligned}$$

Using this cost-function the optimisation algorithm is guided towards finding a swarm orbit and geometry which is inherently stable for the targeted mission time. The suitability of the orbit for interferometry will then have to be evaluated on a separate basis using the PSF as metric. This cost function results in an M-dimensional solution space for a small scale problem, consider an existing constellation of 4 satellites with a design problem of adding a fifth element. Figure 6.2 shows the resulting solution space in the X-Y plane for the addition of the satellite, with a clear single local optimum for the fixed configuration.

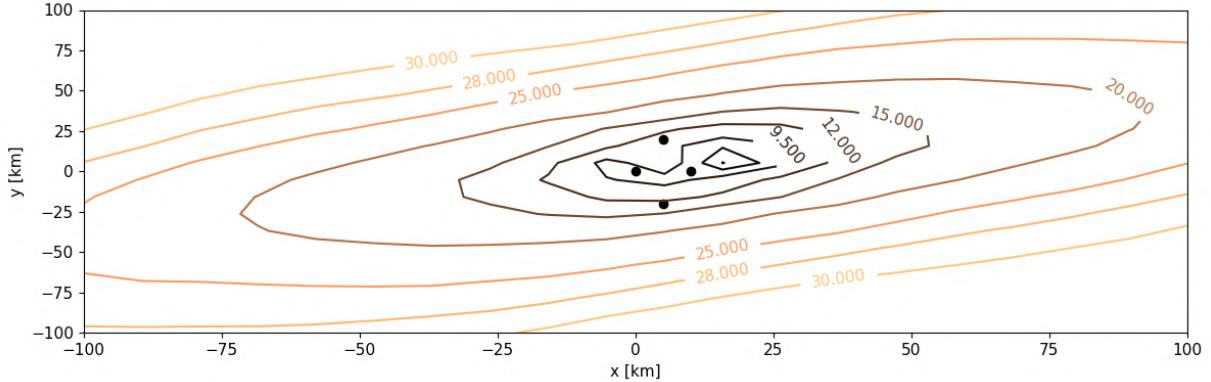


Figure 6.2: Solution space in the X-Y plane for the addition of a fifth satellite to an existing constellation for a year in orbit. Origin is the L4 point. Unit of contours is in thousands. Black dots represent the original satellite configuration, contour values show cost computed with Equation 6.1.

When posed as the addition of a single satellite to an existing configuration the problem seems very solvable. However, this application is far from what will be possible in practice. The given example is a staggered resolution of the variables, based on the solutions for previous variables. The orbit of the swarm is pre-determined, and positions for pre-existing satellites are also resolved. Due to the nature of the problem this approach is impossible in practice, it requires all variables to be resolved simultaneously.

The primary challenge stems from how tightly knit together the global swarm orbit and the initial swarm geometry are. A swarm geometry for passive formation flight cannot be designed independently from the global orbit, since the local motion is entirely dependent on the global orbit. Likewise it is difficult to find a global orbit with advantageous dynamics without a local swarm design. Deciding to design either in advance would severely restrict the potential outcomes, or even make it impossible to find a stable solution at all. Even if a decoupled approach should be considered, deciding which would need to exist first is reminiscent of trying to solve the question pertaining the chicken or the egg. Neither an orbit or swarm geometry can exist without the other, and the easiest solution is to just assume both of them start existing at the same moment in time.

¹Ideally smaller intervals are preferable, but this is a practical consideration with regards to the size of the resulting data files.

Thus, it is concluded that the best approach is solving both the orbit design and swarm geometry at once, since these aspects can not be separated. The downside of this approach is that the optimisation algorithm becomes very complex due to the large amount of variables, some of which have large possible ranges of solutions. The resulting solution-space becomes very ill-suited for first or second-order optimisation algorithms, leaving only heuristic² algorithms as viable candidates. Section 6.2 will review different candidate algorithms, and measure their performance on this type of problem.

6.1.2 Problem boundaries

It has been established that this optimisation problem involves a total number of $M = 6 + 3N$ variables for a swarm of N satellites, likewise an equivalent large amount of constraints will need to be defined in order to keep the problem bounded. This process is simplified by sharing constraints across the three different types of variables, which are the core position (3), the core velocity (3), and relative satellite positions ($3 \times N$). The positioning of all N satellites will be done using the same set of shared constraints, in order to keep the problem scalable. A differentiation will be made between "hard" and "soft" constraints, where hard constraints are enforced through the allowed range of variables, whilst soft constraints are enforced through the cost function.

Due to the way in which boundaries are handled with the Pagmo2 package certain compromises need to be made regarding the definition of constraints. Pagmo requires the use of box-bounds, meaning that fixed ranges are defined for the values of individual variables. As of such, using relative constraints that depend on the assigned values of other variables is not possible. Boundaries which would ordinarily be described as spherical instead need to be approximated as cuboid spaces. The result is a slightly larger search area than what would be ideal. The core position determines the general starting point for the constellation. In chapter 4 it was concluded that the swarm would ideally start close to the L4 point to delay the natural expansion of the L4-centric orbit. This close-proximity region is converted into boundaries by modelling it as a box surrounding the L4 point with a length and width of 5000 km, and a height of 1500 km. These values are deliberately chosen to be large to allow the optimisation algorithm freedom to seek solutions outside of what was expected, while retaining a bounded problem.

The initial velocity of the swarm relative to the L4 point determines the orbit, and in turn the long-term evolution of the swarm profile. Since the primary goal is to orbit the L4 point the relative velocity can not be allowed to be too large. Based on small-scale experiments with optimisation it is decided that the components of the core velocity are constrained to a maximum of 20m/s in any cardinal direction, this range allows for sufficient freedom of orbit design without including unnecessarily large values. For the relative positioning of satellites the maximum allowed baseline is considered as a guide. In order to ensure the initial baseline from placement is kept within 100 km, the placement boundary is described as 100 km sphere around the core position, yielding ranges of 50 km from the core. Similar to the core position and velocity, this constraint needs to be expressed as a box boundary. This results in the following individual hard boundaries for variable assignment:

Relative core position:

$$\begin{aligned} -5000\text{km} &\leq x_c \leq 5000\text{km} \\ -5000\text{km} &\leq y_c \leq 5000\text{km} \\ -1500\text{km} &\leq z_c \leq 1500\text{km} \end{aligned}$$

Relative core velocity:

$$\begin{aligned} -20\text{m/s} &\leq v_x \leq 20\text{m/s} \\ -20\text{m/s} &\leq v_y \leq 20\text{m/s} \\ -20\text{m/s} &\leq v_z \leq 20\text{m/s} \end{aligned} \tag{6.2}$$

Relative satellite position:

$$\begin{aligned} -50\text{km} &\leq x_s \leq 50\text{km} \\ -50\text{km} &\leq y_s \leq 50\text{km} \\ -50\text{km} &\leq z_s \leq 50\text{km} \end{aligned}$$

The unfortunate result of the limitations of the existing PagMo2 framework is that the minimum distance between satellites will need to be enforced as a soft boundary, as they cannot be expressed using a box boundary. Combined with the larger-than ideal search space this makes the established optimisation process less efficient than ideal. All things considered this loss in efficiency is negligible compared to the time saved by the use of the PagMo2 environment over developing these tools from scratch.

²Methods which rely on chance rather than gradient information.

6.1.3 Second-stage optimisation

For the primary optimisation some simplifications were made which can be quite limiting for the possible solutions. The choice for a unified velocity between satellites in particular can be quite limiting for long-term constellation design. To alleviate these concerns a second-stage optimisation will be introduced, which optimises a more detailed set of state vectors for the relative initial positioning and velocities of satellites. In order to allow doing so the second-stage optimisation will make use of the swarm core position and velocities from a first-stage optimisation.

Apart from the swarm core data the second-stage optimisation will be decoupled from the results of the first-stage optimisation. Decoupling the relative satellite positions allows for the investigation of larger swarm sizes utilising the same orbit, as well as allowing for more freedom in swarm geometry designs for longer periods of time. The resulting optimisation problem has a total of $M = 6N$ variables for a swarm of N satellites.

Optimisation constraints

The constraints for the relative positioning of satellites to the swarm core are entirely inherited from the primary optimisation problem. A set of new constraints are introduced for relative satellite velocities to the core swarm velocity, the magnitude of these velocities will be limited to 5 m/s. This velocity boundary allows for some creative freedom in swarm design, while being restrictive enough to retain the cohesion of the swarm.

Relative satellite position:

$$\begin{aligned} -50\text{km} &\leq x_s \leq 50\text{km} \\ -50\text{km} &\leq y_s \leq 50\text{km} \\ -50\text{km} &\leq z_s \leq 50\text{km} \end{aligned} \tag{6.3}$$

Relative satellite velocity:

$$\begin{aligned} -5\text{m/s} &\leq v_{sx} \leq 5\text{m/s} \\ -5\text{m/s} &\leq v_{sy} \leq 5\text{m/s} \\ -5\text{m/s} &\leq v_{sz} \leq 5\text{m/s} \end{aligned}$$

6.2 Comparison of optimisation algorithms

In order to best evaluate the performance of different algorithms a reduced version of the optimisation problem will be used as reference. At a smaller scale this problem faces the same challenges as the planned full scale optimisation, making it a suitable method for pre-evaluation of algorithm performance. The optimisation algorithms will be evaluated on their speed of convergence, the quality of the found optima, and the time required to evolve the algorithm for a set number of generations. A selection of the 5 available global optimisation algorithms offered by the Pagmo library will be applied to a reference problem, this selection includes global heuristic algorithms and particle-based algorithms.

- **[de1220]** Self-adaptive Differential Evolution 1220 ³
- **[sade]** Self-Adaptive Differential Evolution
- **[pso]** Particle Swarm Optimisation
- **[pso_gen]** Generational Particle Swarm Optimisation
- **[gaco]** Generational Ant Colony Optimisation

6.2.1 Reference problem description

The reference problem is identical to what the full-scale optimisation will entail, albeit at a smaller scale. The environment model and integration settings used will thus be identical to those presented in section 5.8. In the reduced problem a swarm of 5 satellites will be optimized for orbital stability using passive formation flight over the course of 160 days, relying on the same cost-, and penalty-functions as described in section 6.1. Finally, the boundaries of the initial placement of the core are stricter, spanning 50 km from the L4 point in all cardinal directions. The resulting problem is much faster to resolve, whilst retaining the same challenging solution space and mission requirements. This makes it well suited to compare the performance of algorithms without requiring a large time investment.

Within the reference algorithm there will be no fixed initial values for the population members, as this feature is not supported by Pagmo⁴. Instead the population is initialized to be entirely random on every run. To guarantee that the processes are compared on equal footing their performance will be compared using results from the same random seed. This forces the algorithms to share initial population members and random mutations throughout the process.

³A variation upon the Self-adaptive method which also utilises adaptive mutations. It's performance is praised by the developers, which is why it is investigated in addition to the normal algorithm.

⁴To the authors knowledge...

6.2.2 Single-algorithm performance

First the performance of the different algorithms will be compared on an individual basis by comparing the results of single-threaded optimisations using a population of 128. The performance is evaluated over 75 generations, with no early interruption in case of a found optimum. Figure 6.3 shows how the performance of different algorithms compares based on the mean, minimum, and maximum cost values of the population per generation.

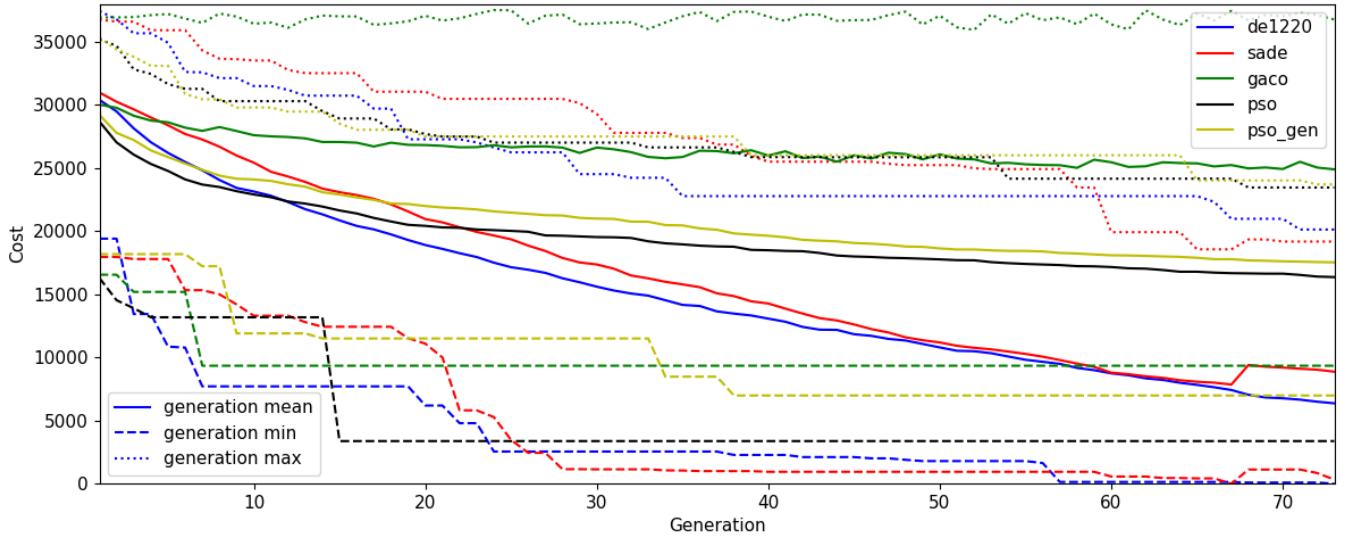


Figure 6.3: Performance comparison of single-algorithm application to the reference problem, with lines for the mean, maximum, and minimum cost per generation.

It can be observed that the particle-based methods are not well suited to tackling this optimisation problem. The generational ant-colony in particular struggles to converge with its population. The cause for this struggle becomes evident from the maximum population cost line alone, which instead of improving over generations oscillates around the original maximum. This shows that one of the "ants" is trapped in a local optimum in a global sub-optimal of the solution space. The particle-swarm based methods show similar difficulties with converging towards a solution, and the plateauing of the curve shows that these methods also get trapped around sub-optima. This can be seen by looking at how the similarity between population members increases over time. Figure 6.4 shows the average cosine similarity between population members, looking at the core position and velocity.

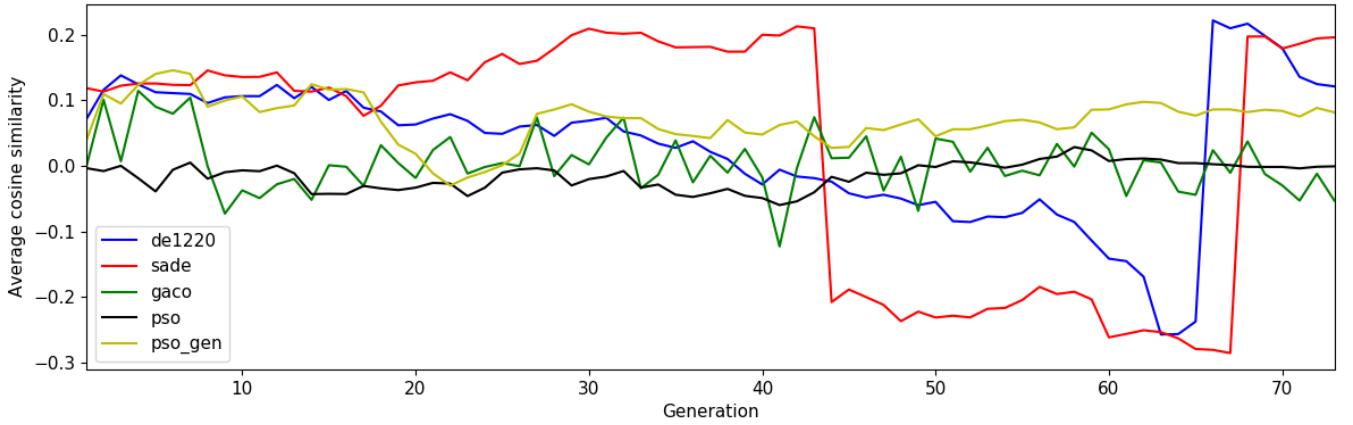


Figure 6.4: Average cosine similarity of the core positioning and velocity components between population members, per generation.

Looking at Figure 6.4 a difference in operation between the individual-, and particle-based methods can be observed. Both differential evolution methods show a wide initial variance in (dis)similarity as they search the solution space, before starting to converge to a more homogeneous population as they get close to the found optimum. The particle-based methods on the other hand are not converging at all. Instead, the individual particles get trapped in sub-optimal solutions as the allowed step size becomes smaller. Only the generational particle swarm optimisation shows some progress towards convergence, allowing particles to escape through generational evolution.

Based on these observations, it is obvious that particle-based methods are not well suited for the swarm design problem. These methods could still prove to be suitable for local optimisation, where the problem of entrapment will not occur. But they are shown to be inadequate for the global solution space search, and thus will not be considered for further use. Between the evolutionary algorithms both de1220 and sade succeeded in finding an optima. The preference for algorithm will be given to de1220, which proved to be faster overall for finding the optima in multiple scenario tests when compared to the sade algorithm.

6.3 The archipelago method

Having chosen the preferred algorithm the next step is to study the efficiency of different problem architectures. Even with the reduced scope of the reference problem the optimisation process requires a significant time investment. The complexity of both the solution space, as well as the orbit propagation, yields an optimisation process that is unsuitable to work with if no solution for the processing time can be found. The results presented thus far have been relying on only a single thread, and the first obvious choice to save time would thus be the implementation of multi-threading in the optimisation scheme. The Pagmo library allows for an easy implementation of this functionality through the archipelago method. This method allows for the parallelisation of multiple islands by distributing these over different cpu cores. A distinct feature of this method is that it allows the islands to exchange information between evolutions through a pre-defined topology, which allows islands to communicate and exchange population to speed up the search.

Figure 6.5 shows the optimisation performance over generations of 3 different architectures for the de1220 algorithm. Note that the distribution across 4 cores by itself yields little difference in the convergence profile, while it yields a substantial reduction in processing time. It can be observed that the use of internal iterations (locally evolving each island between generations) yields a substantial improvement to the rate of convergence, at a small average time cost increase per generation. Figure 6.5 shows a comparison of optimisation performance of different archipelago architectures. Note that for this comparison the algorithms do not yet take advantage of the topology, meaning that the islands are disconnected.

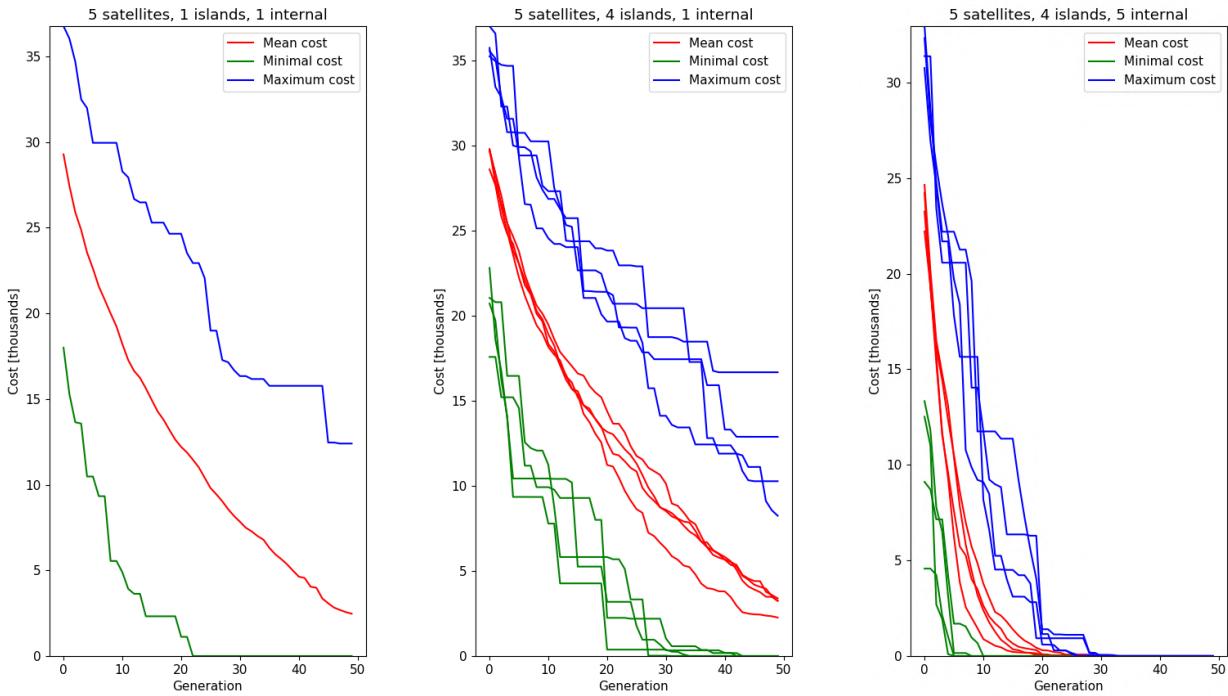


Figure 6.5: Optimisation performance for a 5 satellite problem over generations for different de1220 optimisation architectures: a single-threaded 128 population island, 4 islands of 32 population, and 4 islands of 32 population with 5 internal iterations.

It can be observed that the inclusion of internal iterations per island offers a massive improvement to the generational convergence speed. This is exactly what would be expected, considering the number of function evaluations per generation is multiplied fivefold. Taking this into consideration this method is actually more time-efficient, as shown by the average evaluation time. Single-internal islands required on average 1.9 seconds per cost evaluation, the islands with 5 internal evaluations required only 0.6 seconds on average. Table 6.1 shows the properties and results achieved by the different methods for both a 5 and 10 satellite optimisation.

Table 6.1: Optimisation performance figures of 50 generations of evolution for different architectures for de1220

5 satellites:

Islands	Internal	Evaluations	cpu time [s]	Avg. evaluation time [s]	Speed increase ¹	TTO ² [s]	GTO ³
1	1	6400	12548	1.92	-	5521	28
4	1	6400	3772	0.5893	3.258	2338	31
4	5	32000	13330	0.4166	4.609	1333	5

10 satellites:

Islands	Internal	Evaluations	cpu time [s]	Avg. evaluation time [s]	Speed increase	TTO[s]	GTO
1	1	6400	25043	3.9130	-	-	-
4	1	6400	7393	1.1551	3.128	-	-
4	5	32000	35415	1.1067	3.536	12041	17

¹ Compared to a single-threaded solution, computed as $\frac{T_{avg0}}{T_{avg}}$

² Time Till Optimum: Cpu-time required until the optimum was found, - if it was found.

³ Generations Till Optimum: Generation after which the optimum was found, if it was found.

The data presented in Table 6.1 shows that the use of internal iterations yields the overall best speed increase for evaluations, but also that its efficiency diminishes when the scope of the problem increases. From these two data points it can be theorized that at a certain constellation size single internal iterations become faster overall due to this efficiency loss, but it is hard to draw a definitive conclusion from this limited data without extensive knowledge of the source code. Due to time constraints it was chosen not to further investigate this behaviour by generating more data, since this topic is only tangential to the research question of the thesis.

Four of these 6 small-scale experiments already were successful in finding an optimal solution with cost 0, meaning a swarm geometry and orbit which follows the OLFAR mission requirements for the baselines as described in Table 3.1 during a duration of 160 days. The swarm designs found by the different architectures differ from each-other as well, showing that there is no single optimum to this simple reference problem. The found optima particularly differ in baseline evolution profile, which is most important to the imaging capabilities of an interferometer. The presence of multiple optima with differing properties is beneficial to the notion of including the PSF for interferometry swarm optimisation, although it will remain to be seen if multiple solutions for larger scale swarms could be found. Figure 6.6 shows two *uvw* baseline profiles for the optima found using a single island, and using 4 islands for 5 satellites.

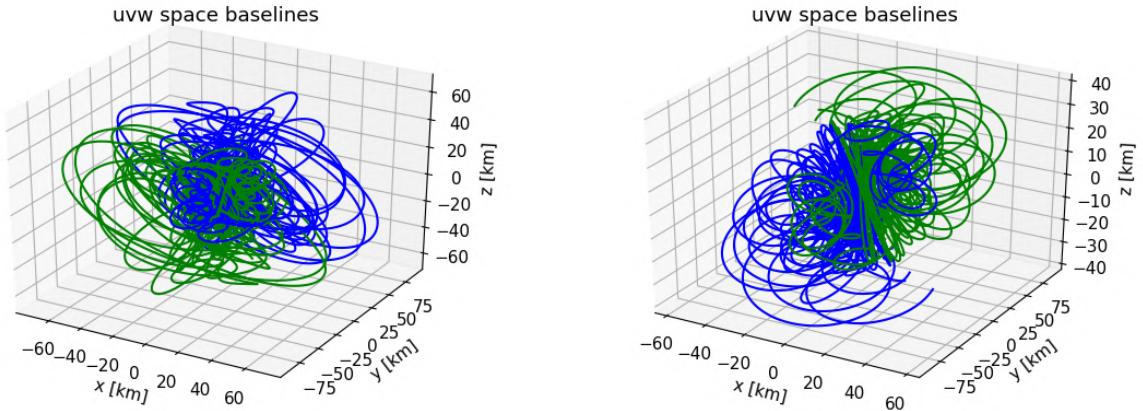


Figure 6.6: *uvw* baseline profile for 5 satellites over 160 days from two optima, found by the 1 and 4 island architectures. Positive baselines in blue, with negative pairs in green

6.4 Topology design

As could be expected the results show that the use of multi-threaded solutions offer significant time-savings over a single thread solution, which will be particularly important for the large-scale optimisations. A remaining question is related to the effect of topology design. It may be expected that the inclusion of topology further increases the overall effectiveness of architectures. This is evident in the comparison between a disconnected and full topology shown in Figure 6.7, where the fully connected topology shows massive improvements in convergence.

A secondary advantage of the use of the topology framework is that it allows combining different algorithms to tackle a single problem. It was previously concluded that particle-based algorithms are not well suited to this problem on their own, yet they might be valuable working in combination with the de1220 algorithm. By adopting population data from the differential algorithm island the particle-based methods might be able to escape entrapment, allowing particle based methods to be used to speed up the final steps of optimisation. To investigate this potential a group of 8 fully connected de1220 islands will be used to optimise the reference optimisation problem, and the performance will be compared with an optimisation done using 7 de1220 islands along with 1 pso island. The different optimisation curves of these topologies are shown in Figure 6.7.

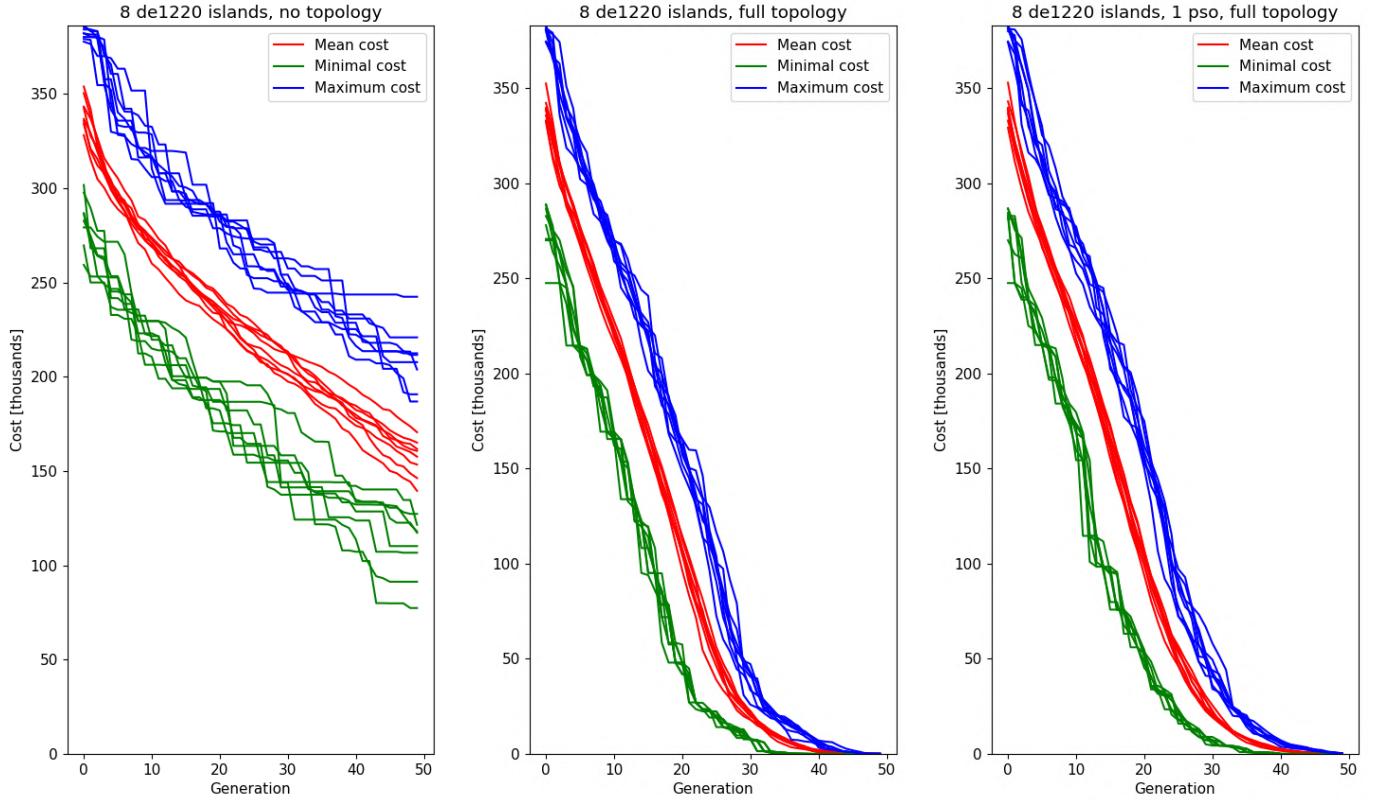


Figure 6.7: A comparison of the optimisation performance by different 8-island topologies on a reference problem of 15 satellites.

It can be seen that the inclusion of a single pso-island only has a very slight effect on the convergence profile. The overall progression towards the optimum is slightly slower, which is what was expected considering the unsuitability of this algorithm for global optimisation in this problem. The final phase of the optimisation by both algorithms slows down, as resolving the near-collision events is the most challenging part of the process. At these final generations it can be seen that the progress is slightly smoother, but overall slower than a full de1220 topology. Overall the results appear to be very comparable, and there is no apparent reason to favour the used of a mixed topology for this problem.

6.5 Scaling of the optimisation problem

Whilst studying the suitability of different algorithms for this problem a secondary research goal was to investigate how well the optimisation problem scaled. This was done by comparing optimisation times towards optima for various problem scales, yielding data such as shown in Table 6.1. It was initially assumed that the large number of variables could pose a problem for large-scale swarm optimisation, and it appears as if this fear was well placed. Due to the way this swarm problem is defined, it is not well suited for scaling to large quantities.

Using the requirement that *all* swarm elements need to be within 100 km from each-other at *all* times, the three-dimensional space in which all satellites can move is rather confined. Besides increasing the overall complexity of the problem the use of more satellites further contests this movement space, making it more difficult to find a solution which does not infringe on the 500 m minimum separation rule. As the number of satellites grows, it will continually become harder to find a configuration which poses no risk of collision. On the other hand the introduction of more variables results in a larger solution space to be searched for the shrinking number of viable combinations.

A remedy to the increasing complexity is the application of brute force, in the form of dedicated servers full of raw processing power. This does not provide a sophisticated answer to the problem, but it can be effective to a certain degree. At a certain point, further increases into the scale of the swarm simply can no longer be mitigated through the application of more processing power. After all, a viable solution might no longer exist beyond a certain swarm scale. Estimating the initial solution space as a sphere with a diameter of a hundred kilometres, this may be filled with several tens of thousands smaller 500m diameter spheres without any of them touching. The complexity of the problem lies in the fact that these also need to move through this spherical solution space without collision for at least a year with paths determined by their relative orbits, which drastically reduces the viable scope of the swarm.

There is no reliable method to estimate this cutoff point beforehand, not without some idea about the ideal swarm design and its orbit. The cutoff point will need to be estimated through studying increasingly larger swarm designs, seeking to find out at which point it is no longer viable to further extend the swarm.

7 Results

Hopefully the title of this chapter is adequately descriptive of the contents which may be expected within. This chapter will firstly present the results of a small study related to the effect of body pointing for orbit manoeuvres using solar radiation pressure, before tackling the results of the main body of work related to swarm orbit design. The latter have all been obtained using the methods described in chapter 6 with the numerical model described in section 5.8. Following the conclusion of Chapter 5 these results were all generated using a variable time-step integration method, and these results were later interpolated to 4 hour timesteps using 8th order Lagrangian interpolation. While this time interval is larger than ideal, it does not pose a threat of missing near-collision events, this is shown in section 7.9.

Chapters 4,5, and 6 of this thesis also present a few significant results which will not directly be repeated in this chapter, instead it was opted to integrate these into their relevant chapters to construct a clear narrative. By doing so the thought process behind the major choices in this work is easier to understand than if the results were all condensed into a single chapter. This chapter predominantly revolves around the results regarding the primary topic of this thesis, optimised orbit designs for radio interferometers utilising passive formation flight.

After presenting the results of body pointing alternation this chapter will first focus on the results of smaller-scale swarm optimisation to point out relevant trends and orbit features, before presenting the largest optimised swarm design using 35 individual satellites. The next section will first briefly discuss the overall strategy used to search this optimum, the associated difficulties with this method, and present some thoughts on the subject of scaling up the problem further.

In this chapter, as well as the Conclusion, some orbits may be referred to as "optimal". This will indicate that the orbits found achieved a cost of zero from the cost function described in subsection 6.1.1. This is to say that said swarm designs fully adhere to the baseline mission requirements described in Table 3.1 during their designed lifetime.

7.1 Search pattern for satellite swarm design

In order to make the structure of this chapter as clear as possible, some light should be shed on the search pattern which was used to produce these results. From observing the progress of increasingly larger small scale optimisations it was known that the expected runtime for larger-scale swarm optimisation could be quite long, even with access to some very powerful servers to run the optimisation algorithms¹. Deciding on how to approach the envisioned swarm scale for this thesis was a consideration of risk versus reward.

A 50-element passive formation flight swarm design presented the ultimate goal for this work, and it would demonstrate substantial capabilities of passive constellation design around the L4 point if such an orbit could be found. After all, such a result would demonstrate that smaller constellations could be designed following these requirements as well². On the other hand, it was not known if such a result could be found for a 1 year orbital period, and if it could be found it was impossible to predict how long this would take. Essentially this option would present a very high reward, at risk of potentially wasting very much time looking for an optimum which might not exist at all. Instead it was chosen to gradually work up the swarm size for the optimisation rather than start with the targeted size of 50 elements. This approach would ensure an increasingly attractive result would always be found, though this safety comes at a cost of requiring more time than immediately scaling up.

The initial intent of this thesis was to seek swarm designs for 50 satellites at least, but both time and resource constraints placed the achievable limit at 35 satellites. The latter optimisation already strained the resources available on the server, as well as the practical limit for computation time³. Larger scale swarms might still be found using these methods, but this was not possible within the constraints of this thesis. section 9.5 contains several recommendations for the further search of even larger optima using the methods developed in this thesis, with particular attention to recommendations which make the process more efficient.

¹Eudoxos server of the Technical University of Delft, of which 32 cores and 32 Gb of RAM were available for this work.

²Removing satellites from a design can be done without risk of breaching any of the set requirements

³The 35 swarm optimisation took 8 days with 32 cpu-cores running simultaneously, occupying just over 32 Gb of working memory.

7.2 The effect of satellite body-pointing

In subsection 5.4.3 it was speculated that pointing the central body of the OLFAR element might be used for relative course corrections, which was put to the test using the full perturbation model environment as described in section 5.8. To investigate the effect of changes in body pointing the propagated orbit of a minimal-body pointing satellite will be compared to a propagation using the full-pointing configuration, which has a larger overall surface area and a higher surface-averaged reflectivity. The satellites were propagated using the settings described in subsection 5.6.1, meaning a Cowell propagator with RK4 integration using 30-minute fixed time steps.

Due to the small mass of individual OLFAR members and their large surface area the acceleration from solar radiation pressure is considerable enough that even a detail as small as the pointing of the central body has large effects on the evolution of the orbit. Figure 7.1 shows the positional change of the satellite over the duration of the orbit, as well as the change in total system energy (kinetic + potential). Within 4 days the total position change compared to a minimal-body pointing orbit is over a kilometre, which is substantial for a course correction method which consumes no fuel.

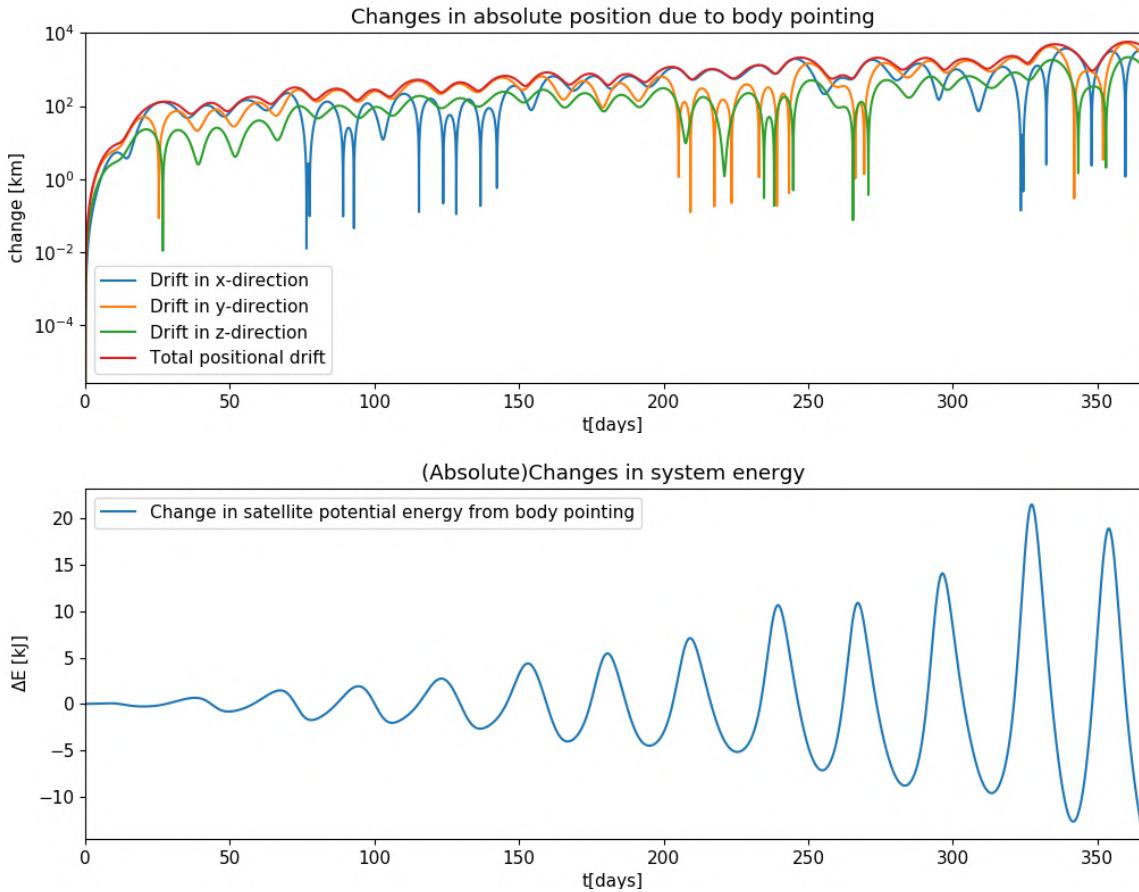


Figure 7.1: Absolute positional change of the satellite orbit from the change in body pointing, and the change in total system energy from the attitude adjustment.

In the changes of the total system energy it may be observed that the reduced effect of solar radiation pressure works to both increase and reduce the energy of the system periodically. This cycle is aligned with the monthly cycle of the L4 point's motion towards, or from, the Sun. During which solar radiation pressure assists to accelerate or decelerate the satellite depending on the direction of motion. With full body pointing the satellite receives more radiation pressure, causing larger accelerations and decelerations. This creates the fluctuating pattern of the total system energy change. The continual growth of the fluctuations show that the effects of additional acceleration and deceleration do not sufficiently cancel out each-other. The continual growth of the fluctuations shows that energy is added to the satellite over time from the use of full body pointing.

These forces do not cancel out due to the relative position of the Sun to the equatorial plane along which the L4 point orbits. The Sun does not lie on the Lunar orbital xy plane of the barycentric system, but instead is offset by a small angle in the z direction. This causes solar radiation pressure and gravitation to be a consistent perturbation in the z direction in addition to the xy plane, meaning that changes in body pointing causes the orbit of satellites to evolve differently in all cardinal directions. Figure 7.2 demonstrates this by plotting the change vector between the two satellite orbits for the duration of the orbit in the barycentric frame.

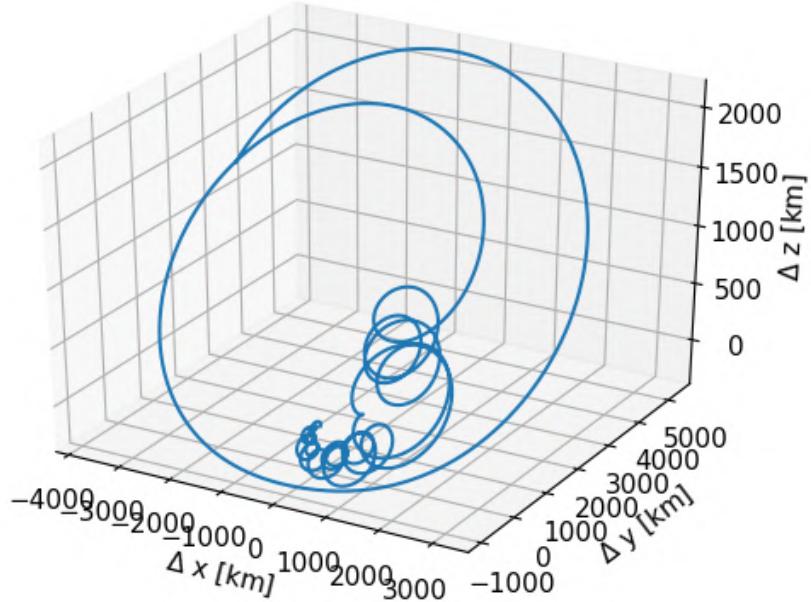


Figure 7.2: Plot of the evolution of the positional difference vector from body pointing over a year in orbit

It can be seen that the changes in the satellite orbit from body pointing are constricted to a single plane, which is equivalent to the relative orbital plane on which the Sun lies to the barycentric system. The change vector demonstrates a sweeping, gradual growth along the Solar plane. Individual revolutions of the L4 point around the Earth can be distinguished in the "curly" pattern in the growth and decay of the difference vector, which grows in magnitude with every evolution. Figure 7.3 shows a detailed view of the growth of the change vector during the first 10 days.



Figure 7.3: Detail view of the first 10 days of the positional difference plot from body pointing alteration

Even in much shorter timespans the influence of body pointing is considerable, within 6 hours corrections of up to 1 meter can be made using body pointing, which might be valuable for collision avoidance. These results show that the OLFAR swarm will be very susceptible to solar radiation pressure, and especially movements induced from pressure differences between elements. This is a double edged sword which might be used to the advantage of maintaining swarm cohesion or orbit safety, but it also represents a considerable danger. If the presented surface area and reflective properties are not cohesive throughout the swarm its cohesion can be affected adversely within a few days. For a swarm orbit around the L4 point it will thus be necessary to control the attitude and the presented surface area of every swarm element to maintain cohesion of the swarm, especially with long mission lifetimes in mind. If the presented surface area is not controlled the swarm will surely disintegrate over time due to the differences in experienced radiation pressure. Having satellites present different surface areas from the rest of the swarm should be a deliberate choice, for example to adjust the course of satellites or to remove defective elements from the vicinity of the swarm.

To prevent this the attitude of every swarm element needs to be controlled and synchronized to maintain swarm cohesion, especially over long mission lifetimes. The potential severity of solar radiation wander from different surface areas is demonstrated for L4-centric orbits, which continually expose the swarm to solar radiation. A secondary question is how this concept translates to a Lunar orbit which offers less and less directional solar exposure. Several avenues for future research are recommended based on these results, which will be discussed in section 9.1

7.3 Small scale swarm optimisation ($N \geq 15$)

This section will be focused on the analysis of a large amount of optimal swarm design solutions for satellite swarms of 15 satellites which were found during experiments with optimisation strategies. The work presented in section 6.4 produced a total of 213 distinct swarm designs for short term (160-days) orbits, and 214 designs for year-long orbits. These results were obtained using a variant of the optimisation algorithm discussed in chapter 6, with smaller solution space constraints. Despite originating from a more constrained problem the large number of distinct solutions allow for some conclusions on general trends, and a comparison showing how swarm designs change for longer-term orbits.

Due to the nature of the design problems and its requirements, a 15-element swarm design can easily be downscaled to smaller swarms by removing arbitrary satellites without risk to said requirements. Keeping this property in mind, it is chosen not to discuss even smaller-scale solutions. It will be trusted that the large variety of solutions for 15 satellite orbits adequately demonstrates the boundless options for smaller swarm designs. Due to the large number of results they will not all be attached as full design tables, instead Appendix C.1 contains the design table for the closest-to-median full-year 15-element swarm designs.

7.3.1 160 day period swarm design:

Figure 7.4 shows the superimposed results of the core positions and initial velocities from 213 individual optima relative to L4 in the barycentric frame. These results were obtained with a reduced search space of 50km in all cardinal directions from L4, this setting was used to speed up the search for smaller scale optimisations which were primarily used to garner insight into the swarm design problem.

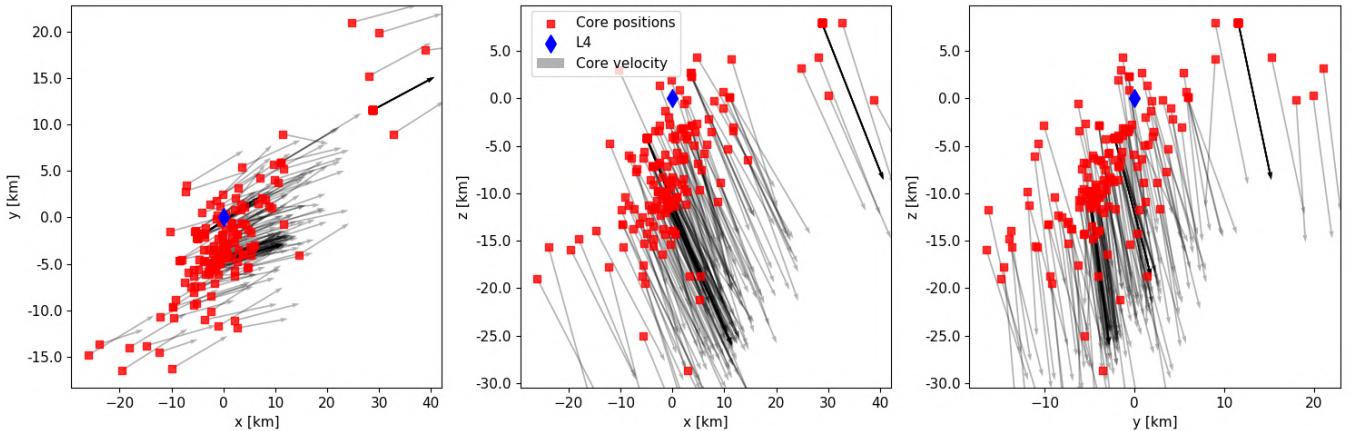


Figure 7.4: Superimposed initial relative core positions and velocities of 213 individual zero-cost solutions for a swarm with 15 elements, relative to L4 in the barycentric frame. Arrows enlarged for visual clarity, size is related to initial velocity magnitude.

Despite the small search space a variety of initial core positions can be observed, focused around a single hotspot with a few outlier solutions. A common trend which may be observed within these results is that all optima introduce an initial velocity in the negative z_B direction of the barycentric frame, and that this velocity trends to be in the positive x_B and y_B directions. This places the swarm into a wide swooping orbit around the L4 point close to the $x_B = y_B$ plane, which slowly expands as perturbations grow the motion along the $x_B y_B$ plane. Figure 7.5 shows the propagated motion of 1 of these swarms relative to L4 in the barycentric frame.

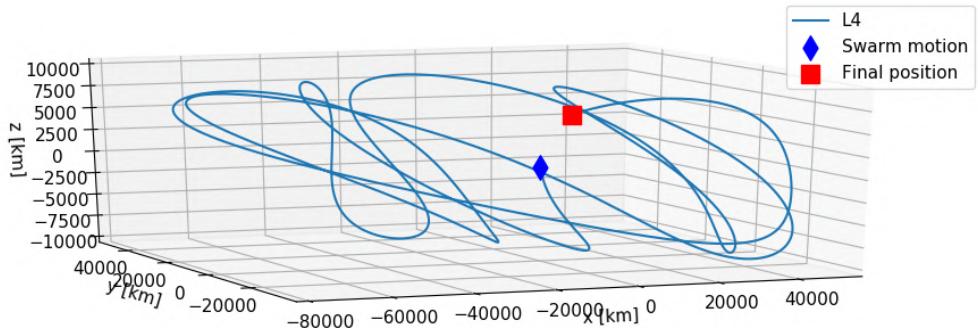


Figure 7.5: Motion of a found optimum swarm relative to L4 in the barycentric frame during 160 days.

Figure 7.6 shows the superimposed initial satellite placements relative to the core position in the barycentric frame, along with the initial velocity vectors. It can be observed that all satellites are placed at a positive x_B displacement from the core position, and that there is an apparent symmetry in the $x_B = y_B$ plane when observed on the $x_B y_B$ plane. The swarm designs are all very compact when viewed on the $x_B y_B$ plane, but the other views show a much larger distribution in the z_B axis.

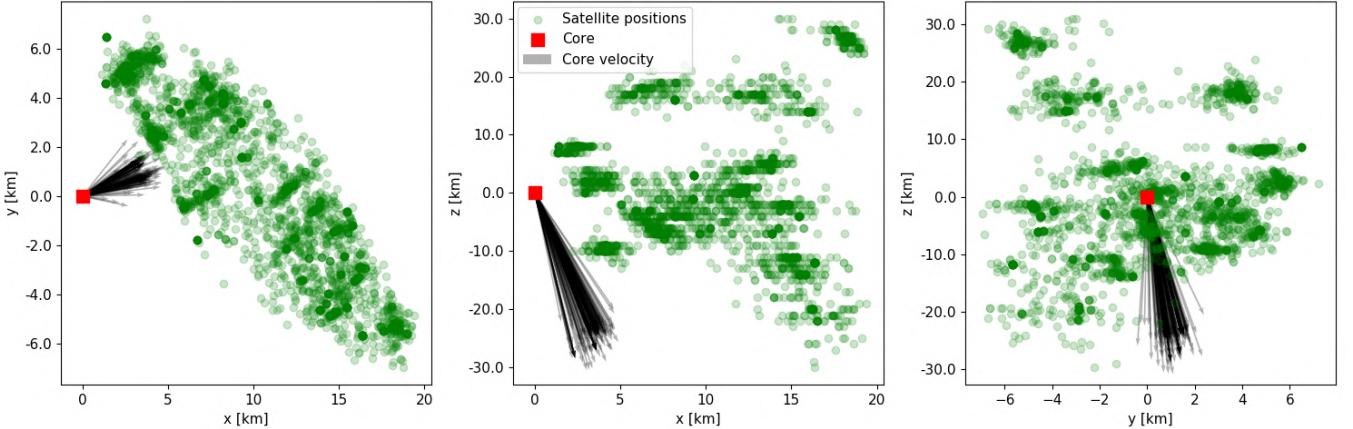


Figure 7.6: Superimposed satellite locations of 213 individual swarm solutions for 15 satellites, placed relative to the swarm core. Reference frame converted to the barycentric frame. Vectors denote initial velocity of the swarm within the barycentric frame

Within this superimposed image several "hotspots" for satellite placements can be observed, the frequency of which is disproportional with the distribution of core positions shown in Figure 7.4. This once again reinforces the notion that for the surroundings of the L4 point, the relative movement between satellites is more-or-less consistent regardless of global swarm position. Besides these observed "hotspots" these results show some relative freedom in placement of remaining satellites, considering that all outliers also belong to an "optimal"⁴ solution.

7.3.2 Full year period swarm design:

Designing a swarm to stay cohesive for a full year is significantly more difficult than a half year, which is evident from the optimisation process itself. For 160 days a solution was found within 35 generations, for 365 days the algorithm required 52 generations to find the first solution. Due to the higher difficulty of this optimisation problem the presented solutions are also much more uniform. Figure 7.7 displays the initial core placements and velocities for the 365 day optimisation, in which a uniform front regarding relative placement and in particular initial velocity may be observed.

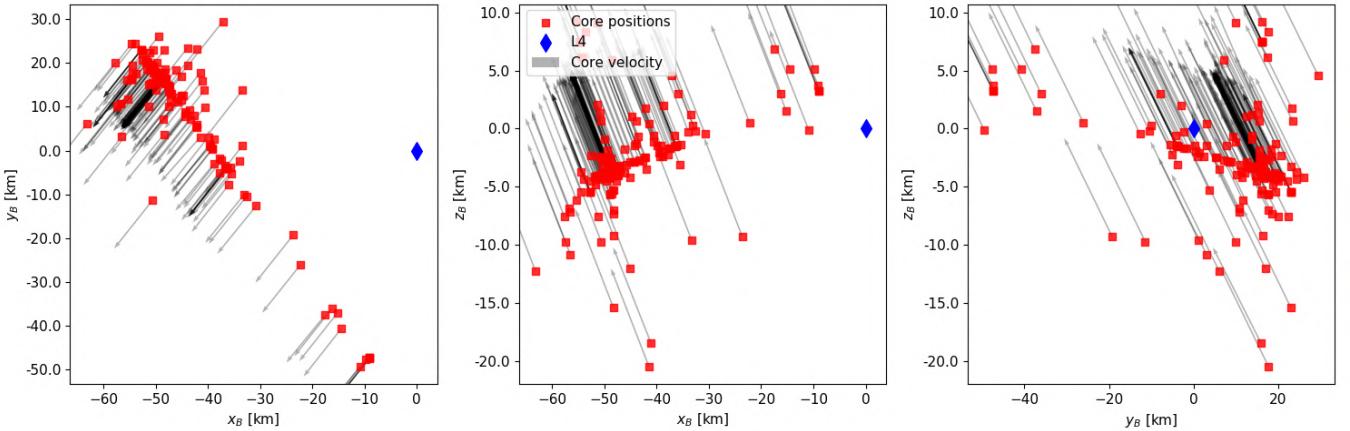


Figure 7.7: Superimposed initial relative core positions and velocities of 213 individual zero-cost solutions for a swarm with 15 elements, relative to L4 in the barycentric frame

A very interesting feature is that the resulting optimum seems to be a contra-design to the short term solution, the initial placement with regards to the L4 point is very similar but the velocity vectors are approximately complete opposites. Figure 7.8 shows the resulting orbit for one of these solutions, which is much more constrained than the short-term equivalent. The major difference between these two is that when observed on the $x_B y_B$ plane the long-term solution counteracts the natural clockwise growth, where the short term orbit embraces it.

⁴With regards to the optimisation algorithm cost function.

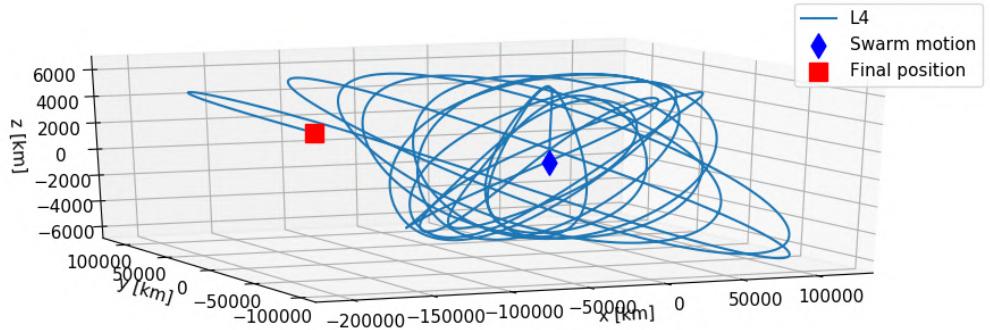


Figure 7.8: Motion of a found optimum swarm relative to L4 in the barycentric plane during a year.

By countering this natural growth the orbit can retain its proximity to the L4 point longer without intervention, resulting in a more stable environment for passive formation flight. Figure 7.9 shows the initial placement of the satellites compared to the core position and the initial velocity in the barycentric plane. The increased restrictiveness of the long-term problem is evident when this figure is compared to Figure 7.6, which shows a much larger spread in satellite locations and velocity vectors.

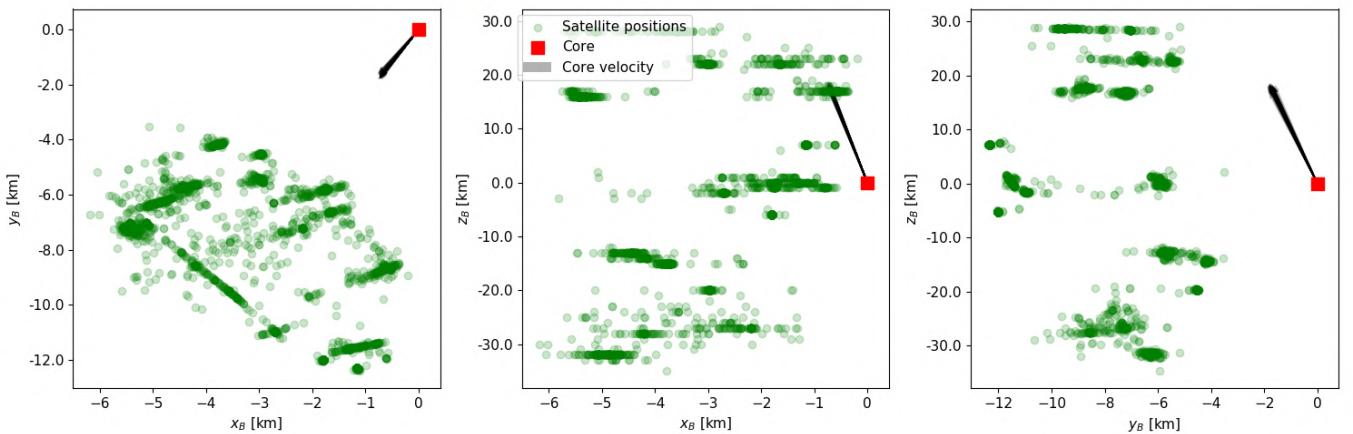


Figure 7.9: Superimposed satellite locations of 213 individual swarm solutions for 15 satellites, placed relative to the swarm core. Reference frame converted to the barycentric frame. Vectors denote initial velocities.

Regardless of the term or direction of the velocity vector a common trend is that for all optima, the satellites are placed in front of the velocity vector in the $x_B y_B$ plane, with the largest spread being in the z_B direction. This is a direct response to the choice to distribute identical initial velocities across the entire swarm, following the theory presented in chapter 4 the swarm is built around the z_B axis with elements both above and below it. The elements closer to $z_B = 0$ have an overshoot velocity compared to the outer satellites, causing them to enter slightly larger and faster relative orbits. As a result the swarm will "fold" over itself around the core position, which is pictured in Figure 7.10 by plotting the initial motion of all elements relative to the core position for the first 3 months.

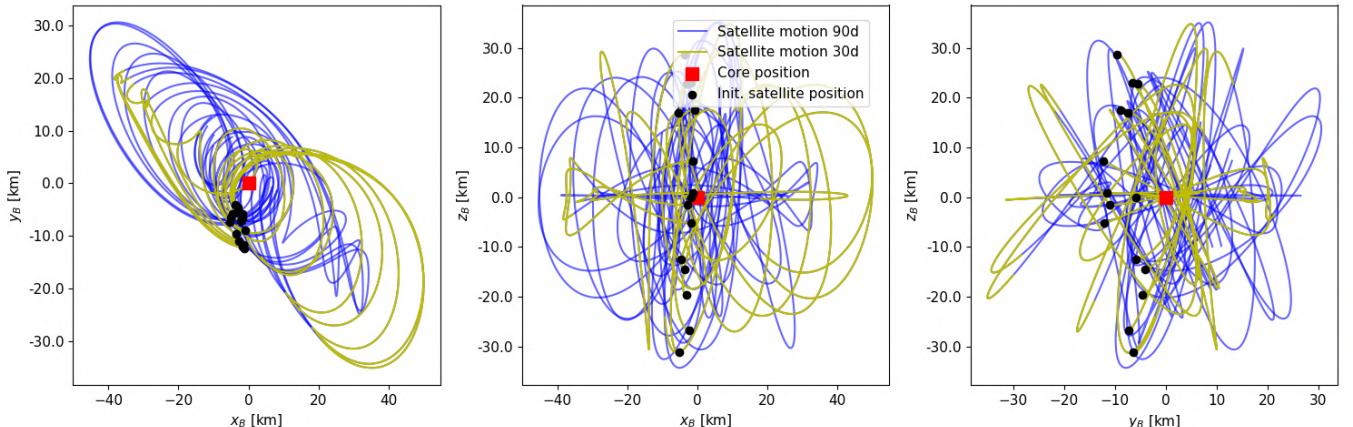


Figure 7.10: Motion of the swarm elements relative to the propagated core over the first 90 days in orbit.

In the $y_B z_B$ view the "folding" of the swarm is best seen by the golden crossover which occurs during the first 30 days. Such crossovers keep occurring over the lifetime of the swarm, in a pattern which initially seems harmonic due to the symmetry in the initial 3-month's motion. Over time however the harmonic symmetry collapses when third-body perturbations alter the course of these elements, resulting in a much more chaotic motion. Figure 7.11 shows the same figure with an extended period of a full year, showing a much more chaotic motion with little visible symmetry.

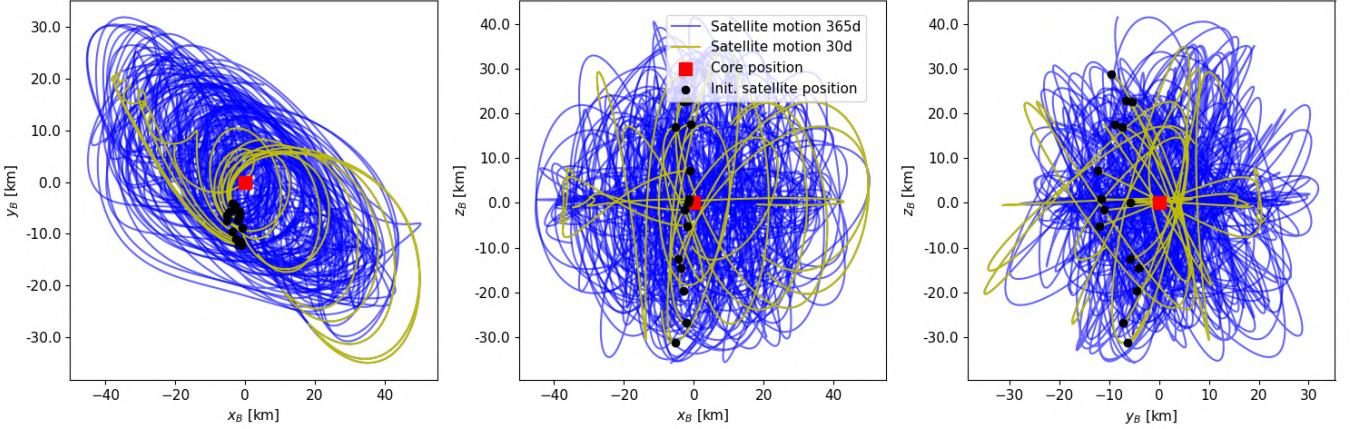


Figure 7.11: Motion of the swarm elements relative to the propagated core over the first year in orbit, with golden lines highlighting the first 30 days.

Figure 7.12 shows the evolution of individual coordinate elements for the 15 satellites relative to the core over time. Within this figure the decay from an initially semi-harmonic motion to a more chaotic motion is clearly visible in the x_B and y_B directions. The motion in the z_B direction on the other hand remains harmonic over time, as this direction is barely affected by perturbations. Clearly visible in this figure are the folding points, which occur as part of 91-day period sinusoidal motion along the $x_B y_B$ plane. It can be seen that the motion in the z_B direction is regulated by a much shorter sinusoidal with a period of approximately 27 days. This motion is very close to the analytical predictions by K. Wakker, which were discussed in section 4.5.

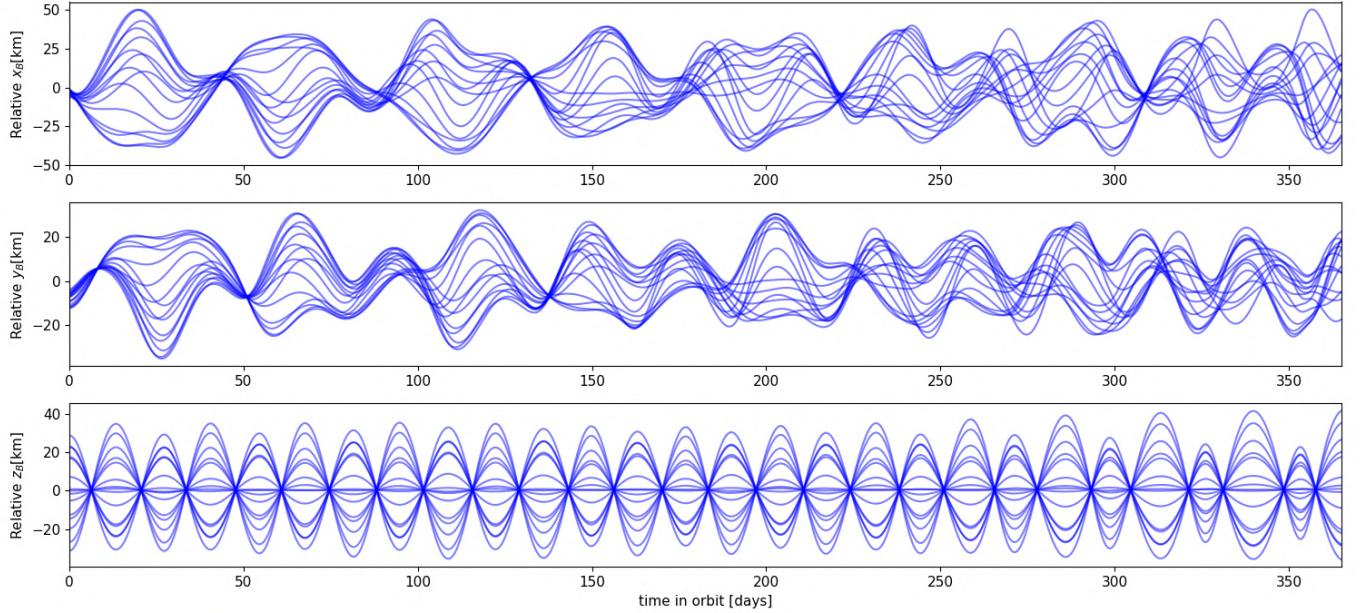


Figure 7.12: Motion of individual swarm elements relative to the core position in x_B , y_B and z_B over time.

The act of folding the swarm over itself in-orbit is beneficial both to its long-term cohesive stability and the distribution of its baselines across uvw space. It may be a unintended side-effect from the use of a identical initial velocity among the entire swarm, but it is very useful for the baseline distribution of radio interferometry. The fold loses its symmetry over time from perturbations, and a key to succesful swarm design would be to find methods that delay the collapse of the symmetrical folding process. As long as the global swarm orbit is suitable, this could be done using minor corrections and active formation control to greatly extend the stable lifetime of the orbit design. Sustaining this motion will continually become harder as the swarm drifts farther from L4. A downside to the folding motion of the swarm is the risk of near-collision events during the fold. This will be particularly challenging with the design of larger swarms, where near-collision events are inherently a greater risk.

7.4 20 element swarm design for one year in orbit

Figure 7.13 shows the initial placement of the swarm core, its relative velocity, and the placement of individual satellites in the barycentric frame relative to L4. Appendix C.1 contains the design table for this solution. Alike the smaller scale solutions the initial velocity of the swarm is downward, with a slight counter-clockwise initial velocity away from the L4 point.

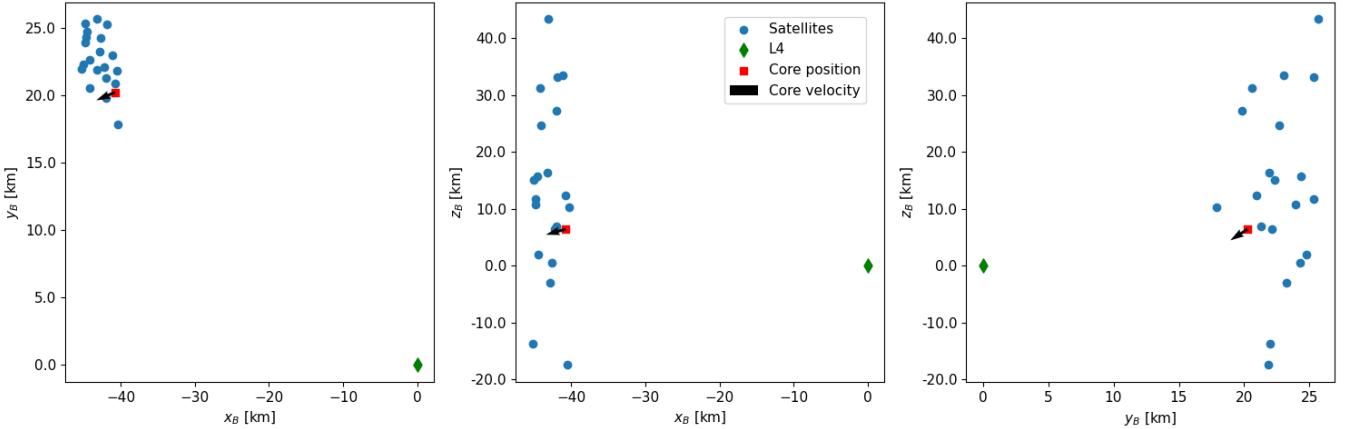


Figure 7.13: Views of initial swarm design for the 20 element satellite swarm in the barycentric frame, relative to L4.

In chapter 4 it was concluded that any swarm orbit around the L4 point will have a tendency of clockwise gradual growth along the xy plane of the barycentric frame. A tendency which this swarm design counteracts by introducing a small counter-clockwise velocity. Figure 7.14 shows the resulting orbit of the satellite swarm. In the $x_B y_B$ plane projection it can be seen that the orbit initially rotates from a nearly y_B -aligned relative orbit, before it starts growing in a clockwise pattern. It can also be seen that the swarm experiences very stable motion in the z_B direction.

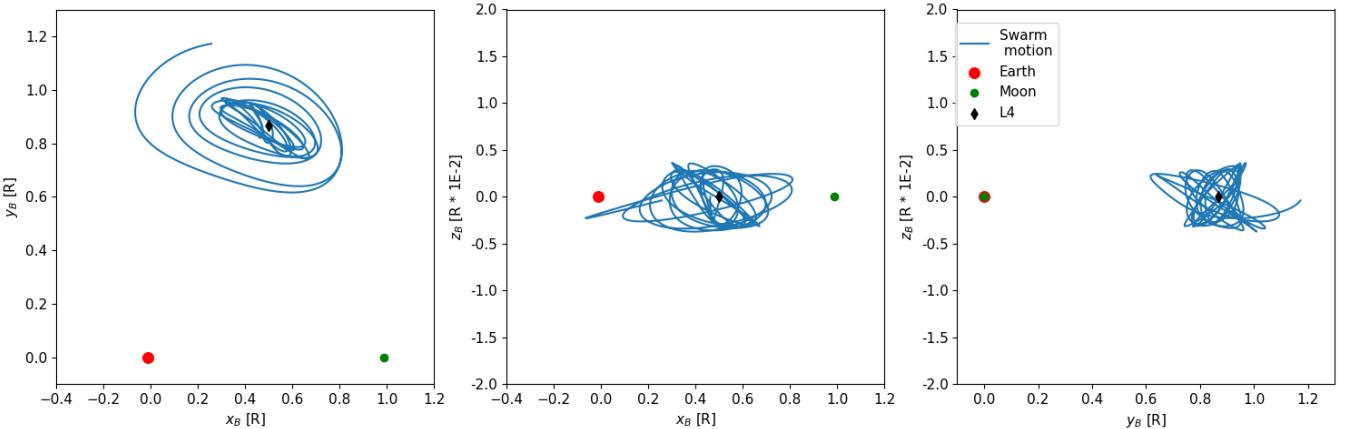


Figure 7.14: Orbit of the 20-element swarm in the barycentric system over a year

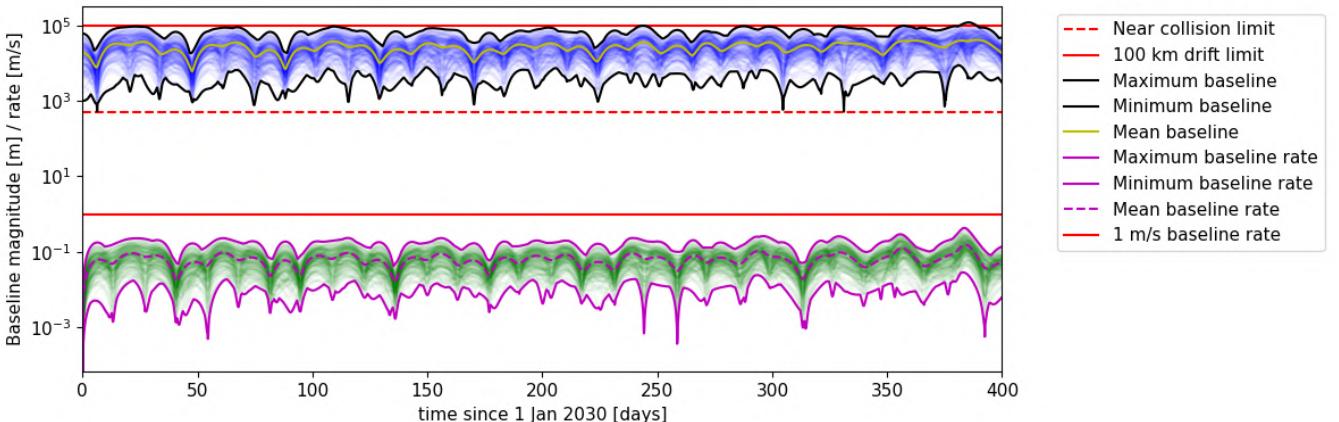
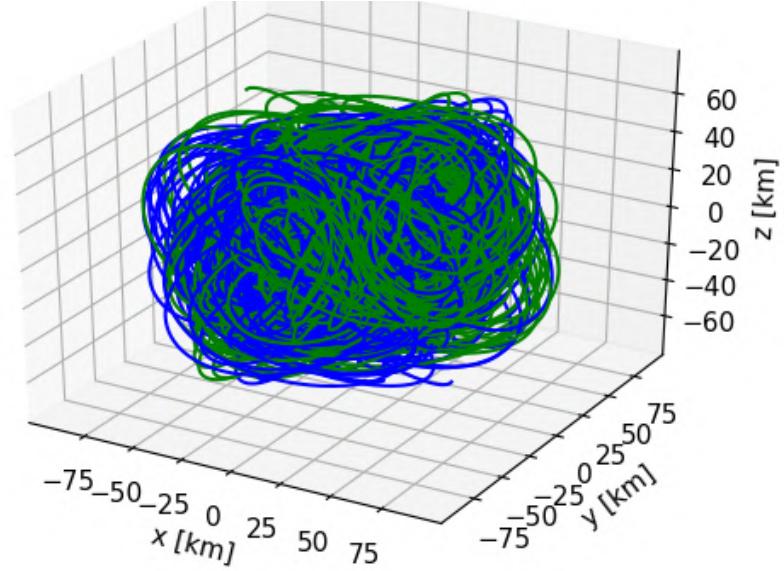
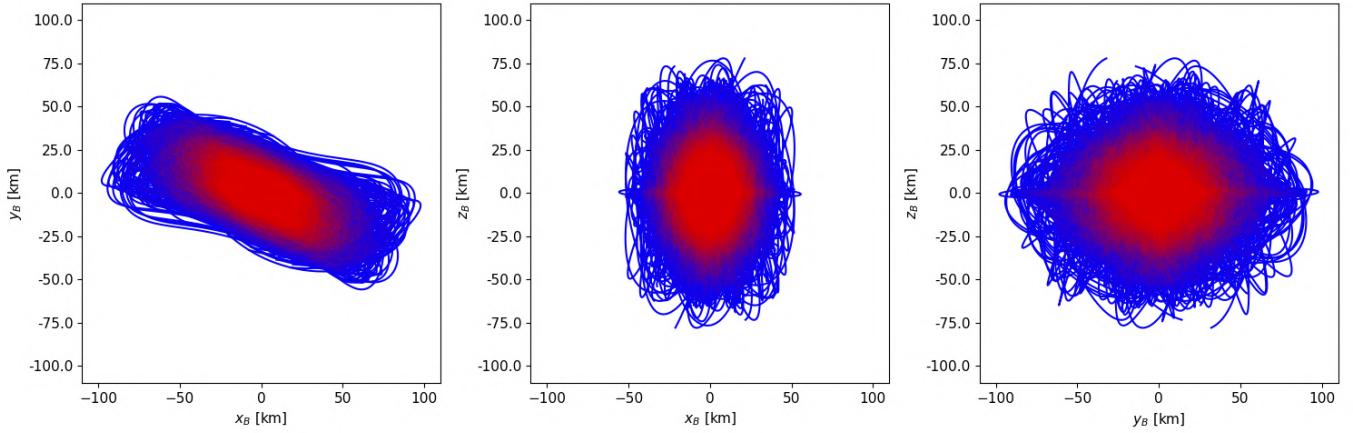


Figure 7.15: History of baseline magnitudes and baseline rates over the first year in orbit.

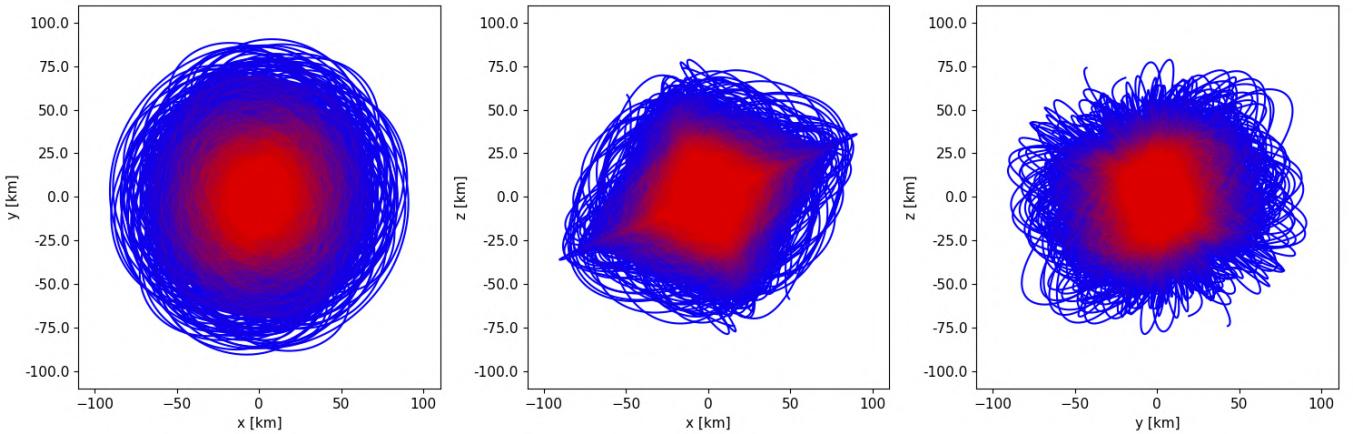
Figure 7.15 shows how all 380 baselines within this swarm evolve over time in both magnitude, and relative baseline rate. It can be seen that at no point either of these values breaches any of the set mission requirements during the designed timeframe. Figure 7.16 shows the resulting baseline distribution projected to the J2000 frame of the entire swarm after this first year in orbit.



(a) 3-dimensional baseline distribution in the J2000 frame after a year in orbit, green lines denote negative baseline pairs.



(b) Projected views of baselines in the Barycentric frame after a year in orbit, red accents indicate areas of higher concentration.



(c) Projected views of baselines in the J2000 frame after a year in orbit, red accents indicate areas of higher concentration.

Figure 7.16: Baseline profile of the 20 satellite optimum orbit after a year in orbit.

Figure 7.16 shows that the resulting baseline profile is particularly well rounded in the xy plane of the J2000 frame. In this frame the relative rotation of the L4 point around Earth helps tremendously with angular baseline distribution, which is not particularly promising when seen in the barycentric frame. Likewise the effect of the obliquity of the Lunar orbital plane can be observed in the projected xz and yz views, as axes of higher concentration of baselines.

The resulting baseline profiles are somewhat lacking in the z direction of the J2000 frame, which is particularly hard to obtain due to the oblique fixed nature of the orbital plane the swarm is on. Filling the uvw space in this direction will be hard for any constellation which relies on a orbit around the Lunar orbital plane, which is why it won't be seen as a problem for L4-centric constellation designs.

Long-term evolution

A secondary question of interest is how well this swarm design holds up for longer periods of time than the design period of a year. Using the same environment model the swarm design is propagated for a period of 5 years to investigate long-term behaviour. Figure 7.17 shows the resulting orbital profile and baseline properties over this timeframe.

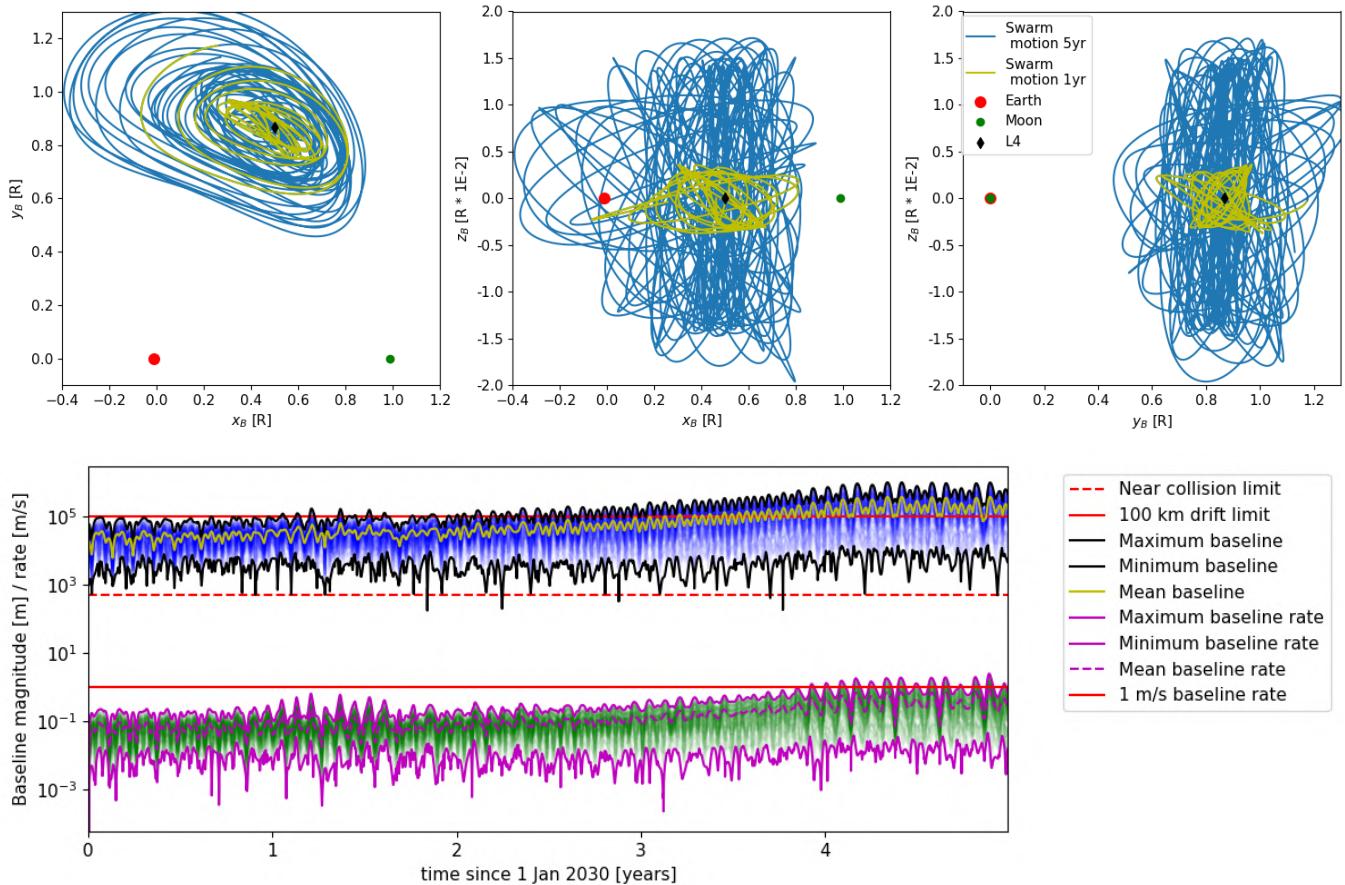


Figure 7.17: Orbit of the 20 satellite swarm in the barycentric frame over 5 years, with the first year orbit overlaid (top). Baseline magnitude and rate history over 5 years in orbit (bottom).

The global orbit found for the swarm is inherently quite stable, and after 5 years the swarm can be expected to still orbit the L4 point. As a result of the folding behaviour the swarm remains cohesive for almost three years before the average baseline exceeds 100 km. Over time the swarm continues to grow beyond this point, which is not fully surprising considering the modelled swarm relies entirely on passive formation flight.

The swarm design was optimised for a mission period of 365 days, during which all requirements from Table 3.1 are met. The first infraction on the mission design requirements is a baseline which overshoots the 100km requirement after 381 days. This temporarily renders a single baseline unusable, but it is not of large influence to the overall capabilities of the swarm. Using passive formation flight only the swarm remains moderately usable for up to 3 years before the average baseline magnitude exceeds the 100km threshold. The first near-collision event occurs after 681 days of passive formation flight with a separation of 181 m between two satellites. Such situations are not desirable, but they also pose little actual danger as relative velocities are in the order of centimetres per second. Over the course of 5 years only 4 of these events are registered, none of them with closer proximities than 181m. Given the small relative velocities these events could be avoided at little fuel cost, making them a small problem compared to the global expansion of the swarm.

If the infractions upon the mission requirements could be accepted this swarm design remains usable for at least 2.5 years utilising only passive formation flight. With 380 individual baselines every element only needs 142.1 days of observation to reach the cumulative 54,000 day sensitivity limit. This design already has ample lifetime in theory, and it could easily be extended by implementing active formation flight to occasionally contract the swarm, which is viable for as long as the swarm retains a L4-centric orbit.

7.5 25 element swarm design for one year in orbit

Despite being optimised with the same random seed as the 20 satellite swarm, the 25 element swarm design is based on an entirely different global orbit. Figure 7.18 shows that this swarm design starts much further from the L4 point, with a downward velocity along the $x_B = y_B$ plane. The general swarm design and its initial z_B velocity are consistent with the previous design. The swarm is distributed as a column in the z_B direction, displaced in the positive y_B direction from the core. Appendix C.2 contains the design table for this swarm.

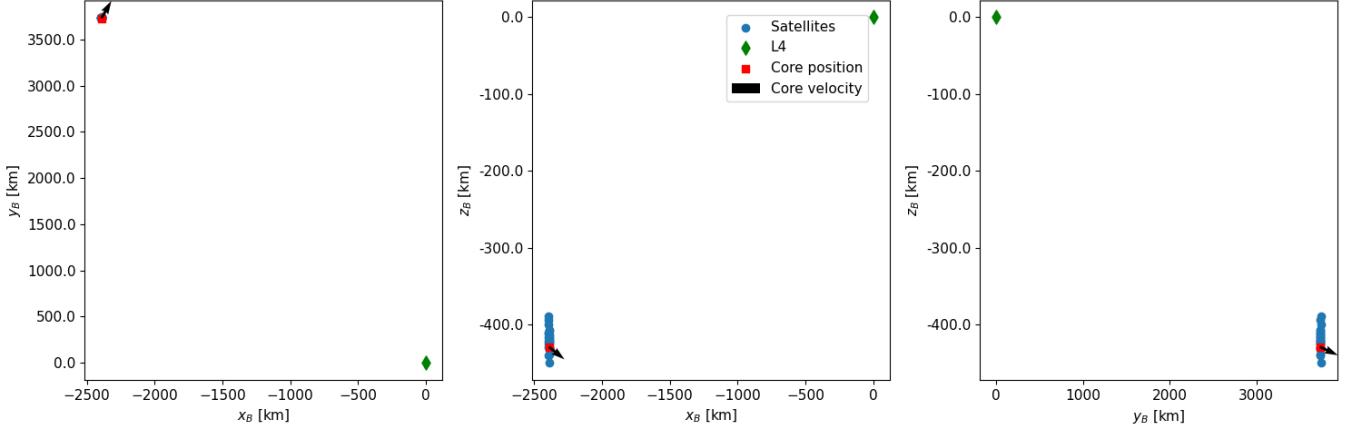


Figure 7.18: Views of initial swarm design for the 25 element satellite swarm in the barycentric frame, relative to L4.

The resulting global swarm orbit for this design shown in Figure 7.19 is much more concise than the 20-element variant, when compared on the x_B, y_B plane expansion over a single year. This solution once again relies on the natural evolution in clockwise direction along the $x_B y_B$ plane to expand the swarm from an initially focused orbit, in which it is more successful at resisting expansion than the 20-satellite solution.

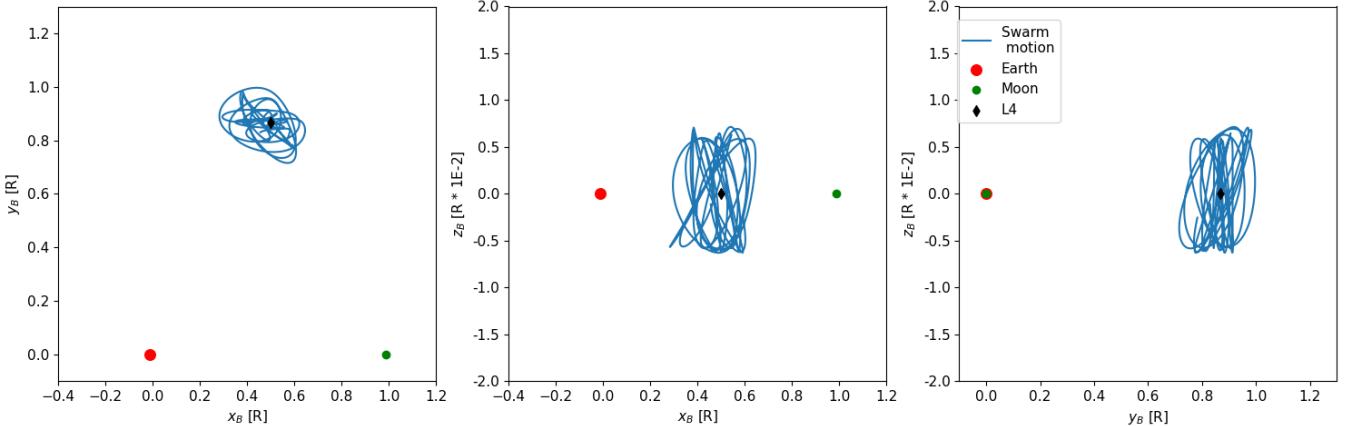


Figure 7.19: Orbit of the 25-element swarm in the barycentric system over a year

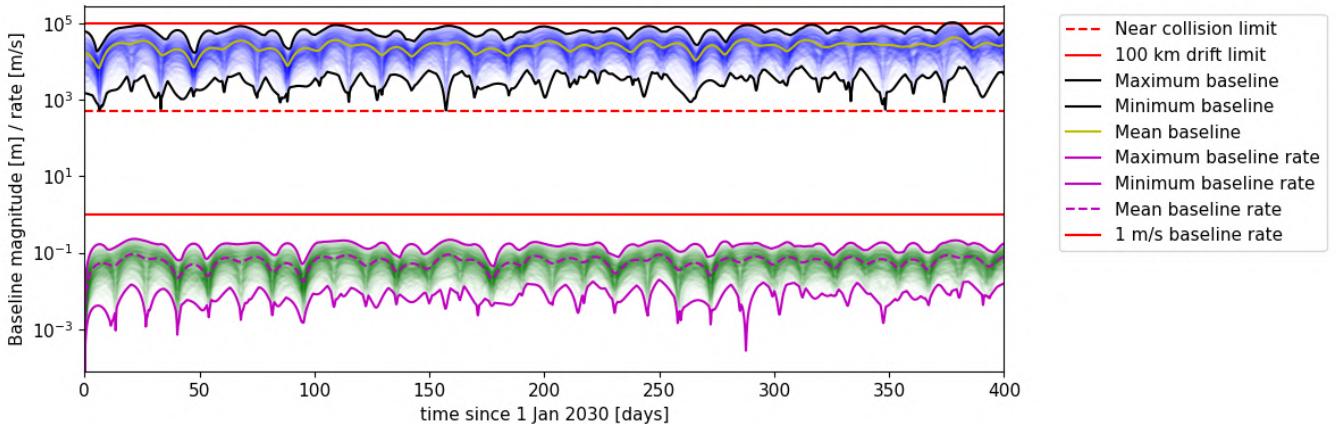


Figure 7.20: History of baseline magnitudes and baseline rates over the first year in orbit for 25 satellites.

Figure 7.20 shows the state history of all 600 baselines within this swarm over the first year in orbit, in which it can be seen that the mission requirements might be skirted but they are never breached. Figure 7.21 shows that the global uvw space baseline profile is very similar, albeit better rounded out through the addition of 220 baselines.

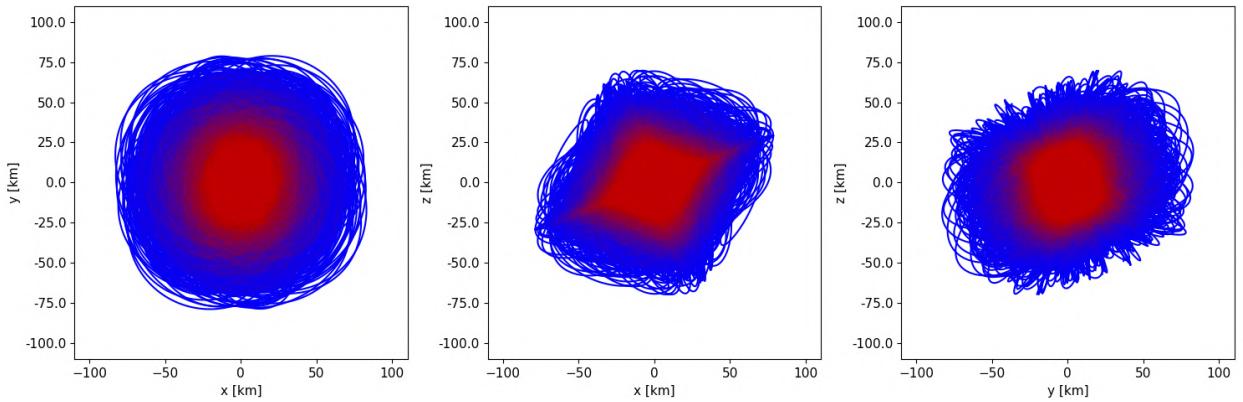


Figure 7.21: Projected views of baselines after a year in orbit with 25 satellites. Projected along the xy (L), xz (M), and yz (R) planes of the J2000 frame, red accents indicate areas of higher concentration.

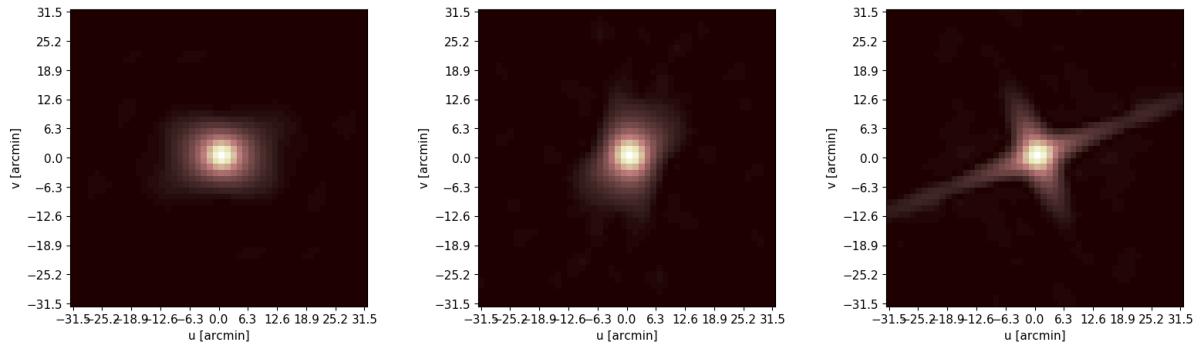


Figure 7.22: Point-spread functions of the 25-satellite swarm for the xy (L), yz (M), and yz (R) projections in the J2000 frame.

The limits of being confined to the Lunar orbital plane start to become more apparent in the xz and yz views, as the fill shape is noticeably more flat in some directions. The effects of these shortcomings are apparent in the PSF profiles of these views shown in Figure 7.22. The idealized PSF cannot be achieved due to the lack of 100 km baselines in this projected view. The baseline history in Figure 7.20 shows that baselines this size rarely appear in the first year of this swarm design's orbit. The average baseline instead is lower as a result of the design around avoiding baselines larger than 100 km during the design period. Due to the natural expansion of the swarm over time it is nearly impossible to reach baselines of this magnitude without eventually skirting over this boundary, hence the optimised design is overtly conservative with its baseline magnitudes. That said, the second and third year offer ample opportunity to observe with 100 km baselines through the natural expansion of the swarm over time. Figure 7.23 shows the same baseline projections, accounting for all valid baselines for 5 years in this orbit.

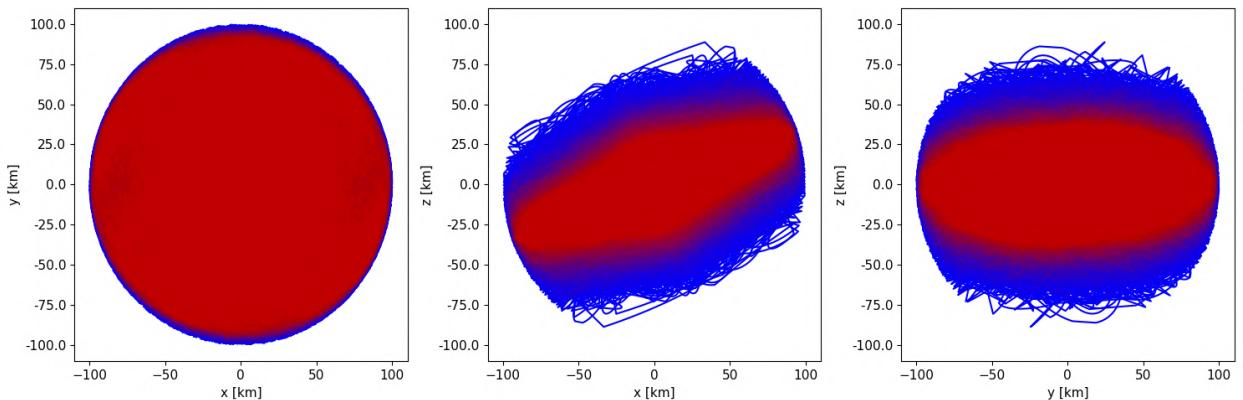


Figure 7.23: Projected views of valid baselines after five years in orbit with 25 satellites. Projected along the xy (L), xz (M), and yz (R) planes of the J2000 frame, red accents indicate areas of higher concentration.

Taking into account all the valid baselines over 5 years for this swarm design the coverage profiles become significantly more appealing. The natural rotation of the L4 point around Earth is a tremendous help in filling in the xy plane with ample coverage to spare. In the other two planes the influence of the obliquity of the orbital plane with regards to the J2000 system can clearly be seen. Without sacrificing some swarm cohesion it will be extremely difficult for any swarm on the Lunar plane to achieve similar coverage as the xy plane for these views, as it requires large out-of plane relative motions which do not naturally occur between two similar orbits. The baseline profiles for the 35 element design presented in the next section do show that having more baselines is beneficial towards this coverage.

Long-term evolution

Like the 20-satellite swarm the same numerical environment is used to propagate this swarm design for 5 years to study long term behaviour. Figure 7.24 shows the resulting long-term orbit and baseline properties. The global orbit of this swarm design is very similar to that of the 20 satellite swarm design, both in short and long-term evolution. However upon comparison it can be seen that the long-term evolution of this orbit appears to be less stable, as it is distributed among a much larger space.

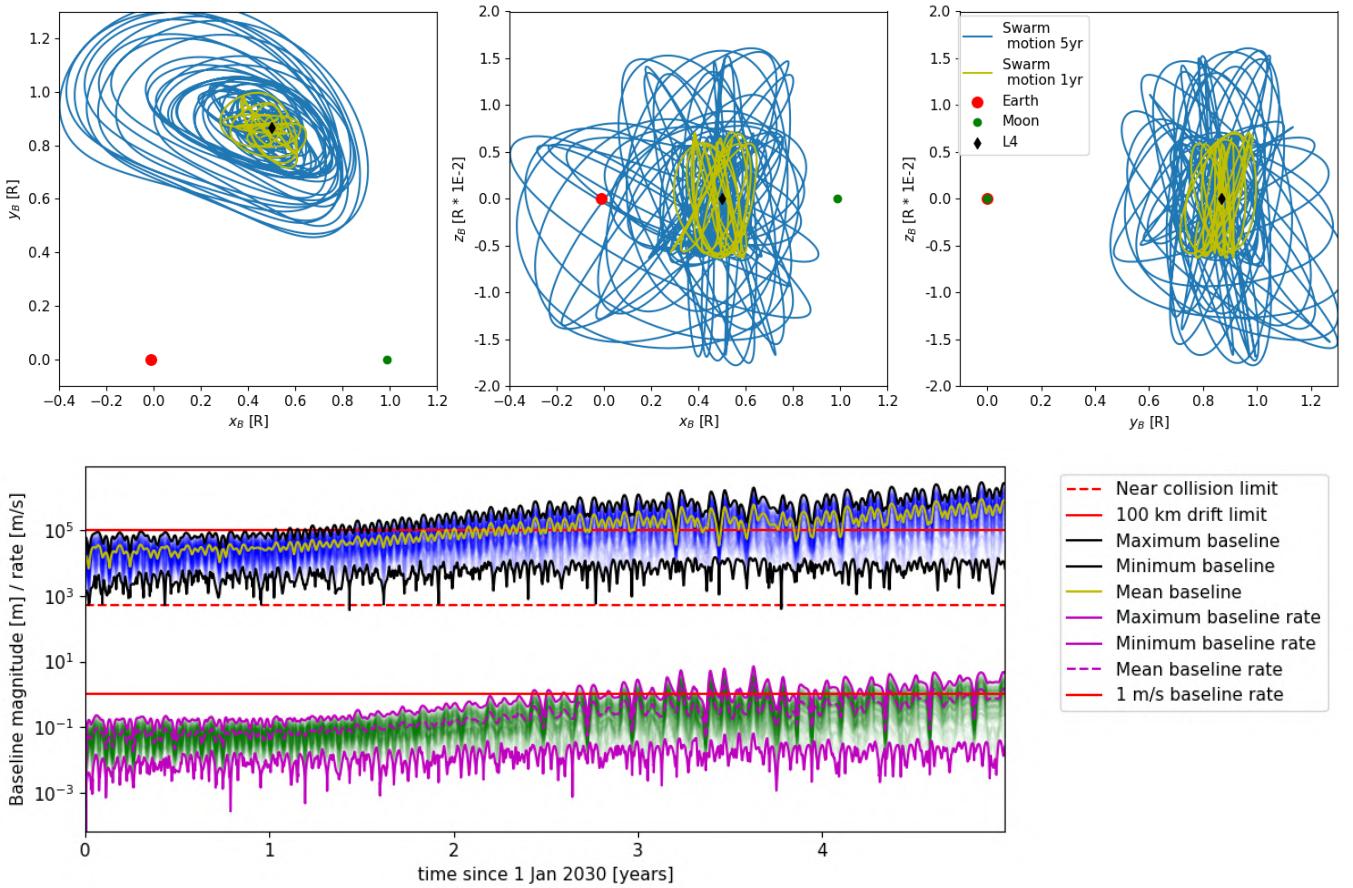


Figure 7.24: Orbit of the 25 satellite swarm in the barycentric frame over 5 years, with the first year orbit overlaid (top). Baseline magnitude and rate history over 5 years in orbit (bottom).

The lesser stability is also visible when the swarm design's longevity is observed. This design drifts apart much faster than the smaller scale design, resulting in a lower overall lifetime than the 20-satellite configuration. The baseline history shows that the overall magnitude and baseline rate increase much faster than the 20-satellite design, rendering the swarm unsuitable for interferometry after 2.3 years when the average baseline is larger than 100 km. This difference is allocated to the less stable global orbit which was found for this design. Despite the similar initial swarm orbits the "late stage" orbit of the swarm is very different from the 20-satellite design, spanning a much larger swath of space in the z_B direction. This design might sport 5 more satellites, but its global orbit makes it unsuitable for very long-term use.

With a total of 600 baselines this design needs only 90 days of observation time per baseline to reach the sensitivity limited 54,000 days cumulative observation time. Assuming that the swarm can handle measurements across all baselines simultaneously this gives ample time to reach the targeted observations during the first year, and for a repetition of these measurements during the second year in orbit to make use of a generally longer baseline profile. The orbit which was found for this design might offer a smaller overall mission lifetime but it is still more than sufficient to meet the targeted mission criteria multiple times.

7.6 35 element swarm design for one year in orbit

After 8 days of optimisation with 32 cpu cores a swarm design for 35 satellites was found which also met all of the conditions for at least a year in orbit, for which the design table can be found in Appendix C.3. Figure 7.25 shows the initial swarm geometry of this configuration, which is not entirely dissimilar from the previous solutions.

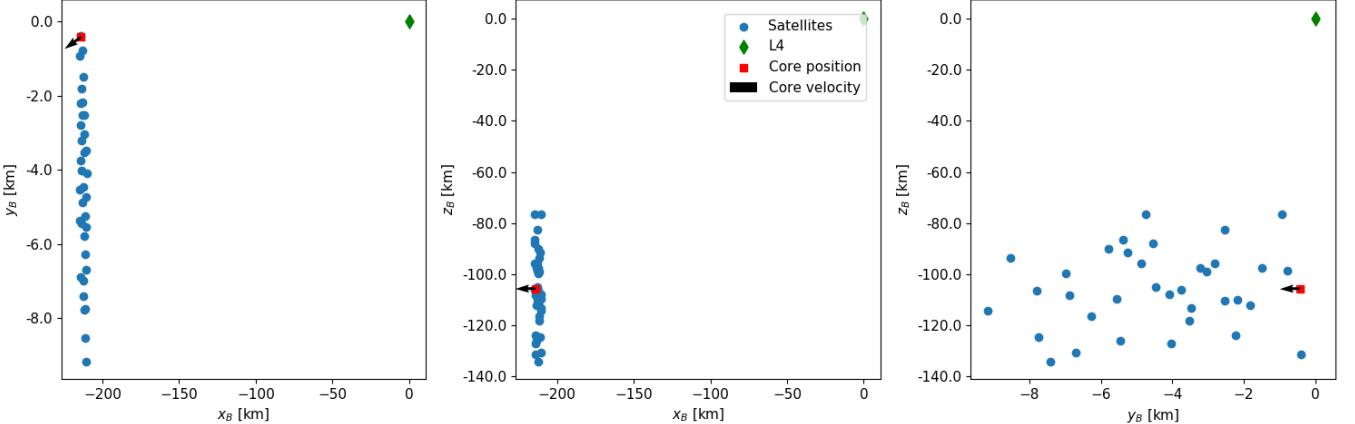


Figure 7.25: Views of initial swarm design for the 35 element satellite swarm in the barycentric frame, relative to L4.

Like the 20 and 25 swarm optima this design starts at smaller x_B , y_B and z_B coordinates than the L4 point. A difference with those solutions is that this design has very little initial velocity in the z_B direction, and that this velocity is aimed away from the L4 point⁵. The resulting orbit is very similar on the x_B - y_B plane, but offers minimal motion in the z_B direction. Figure 7.26 shows the orbit and baseline properties for the first year in orbit, in which the first near-collision event can be observed just after the design period at 371 days.

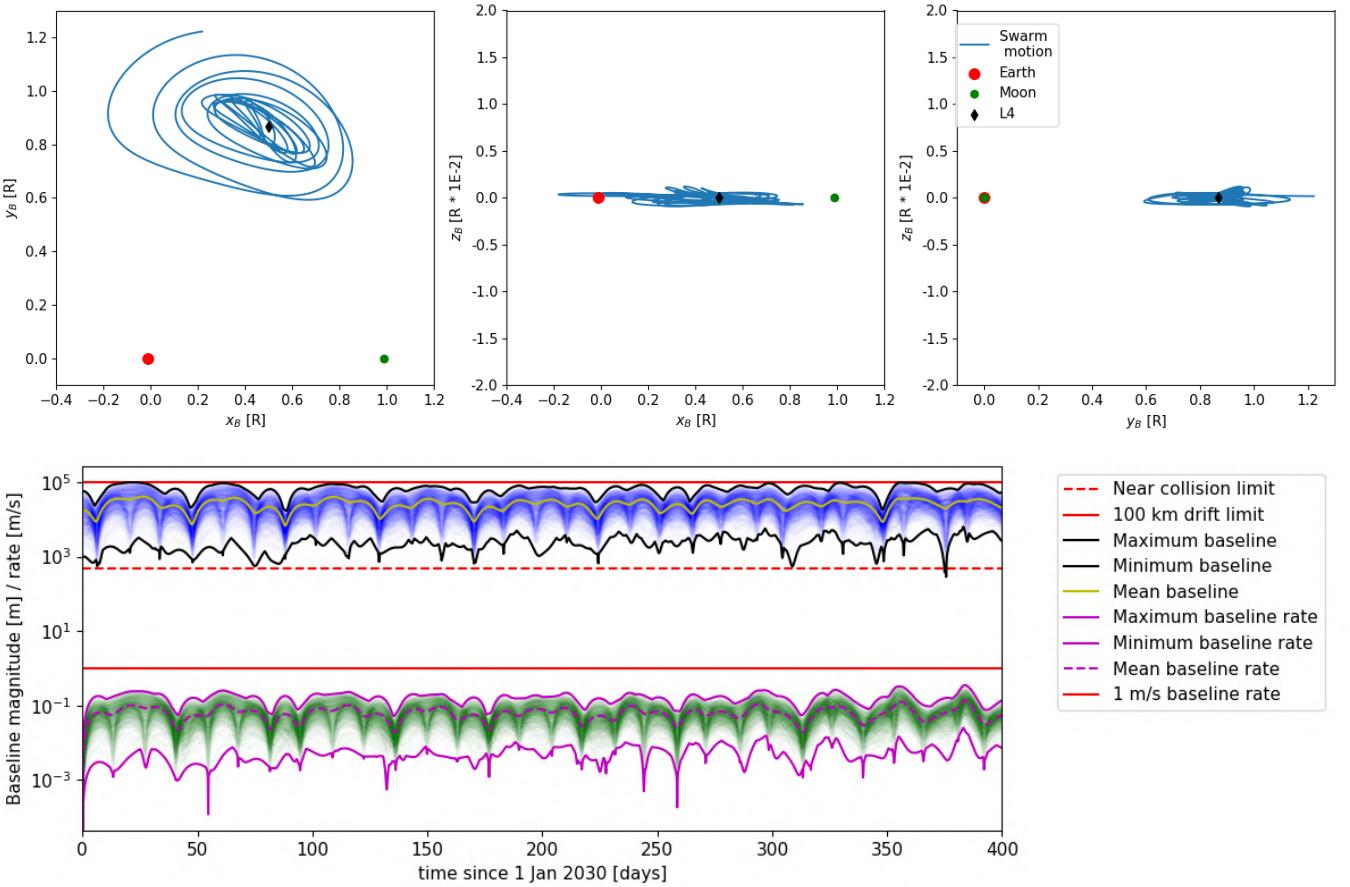


Figure 7.26: Orbit of the 35 satellite swarm in the barycentric frame over one year (top). Baseline magnitude and rate history over one year in orbit (bottom).

⁵The axis on the x_B - y_B view are deceptive.

Figure 7.27 shows the projected baseline profiles in the J2000 system, in which it can be seen that the profiles are slightly better rounded and wider through the addition of new baselines. This results in slightly more compact PSFs, which are shown in Figure 7.28.

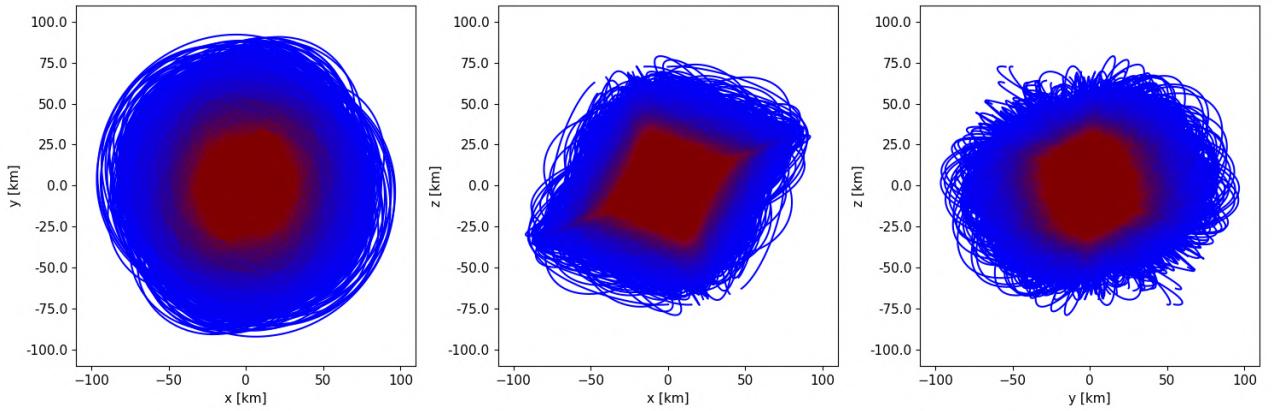


Figure 7.27: Projected views of baselines after a year in orbit with 35 satellites. Projected along the xy (L), xz (M), and yz (R) planes of the J2000 frame, red accents indicate areas of higher concentration.

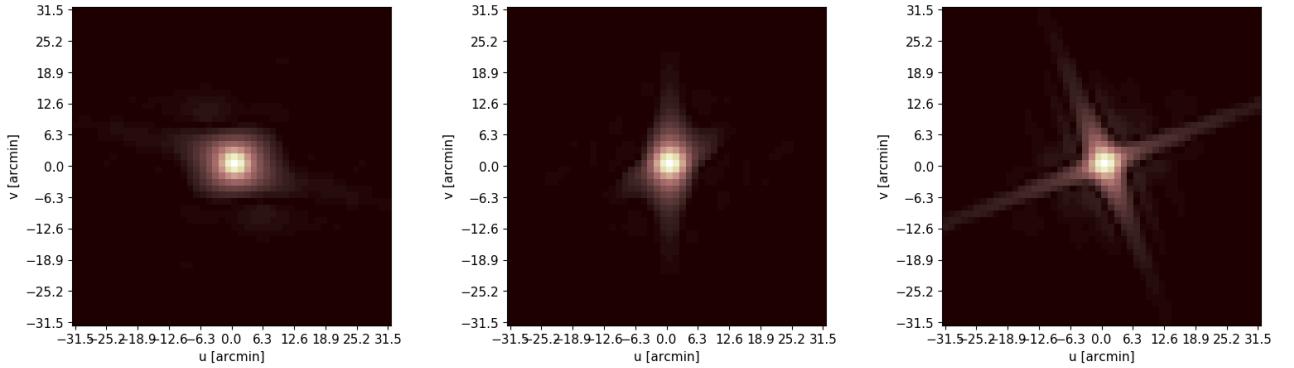


Figure 7.28: Point-spread functions of the 35-satellite swarm for the xy (L), yz (M), and xz (R) projections in the J2000 frame.

Due to the presence of slightly larger baselines in the xy projection the shape of the resulting PSF shrunk compared with the 25 satellite design, but it is still far from the ideal Airy disk. The yz plane in particular yields an imperfect PSF due to the flattened shape of the baseline profile along this plane. These profiles only account for the first year in orbit however, and the interferometry results may be improved significantly when the full set of valid baselines is considered. Figure 7.29 shows the baseline profiles when all valid baselines during the 5-year period are considered.

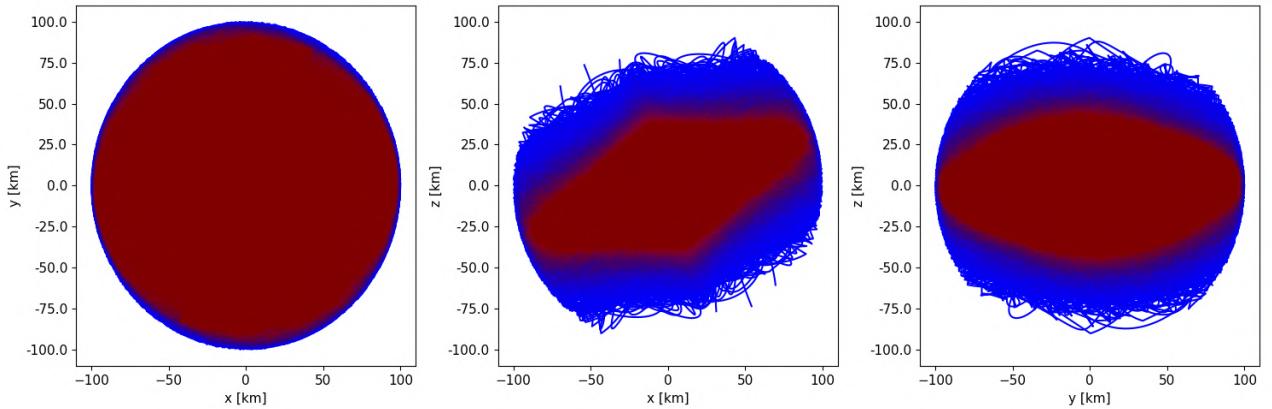


Figure 7.29: Projected views of baselines after five years in orbit with 25 satellites. Projected along the xy (L), xz (M), and yz (R) planes of the J2000 frame, red accents indicate areas of higher concentration.

Aided by the natural revolution around Earth, the xy plane coverage is fully realised with significant surplus coverage in the 4 years of stable orbit around Earth. On the other planes the limitations of being attached to the Lunar orbital plane can be seen once more, although the coverage has improved from the addition of 590 baselines. In the end these views will be challenging for any swarm design which is attached to the Lunar orbital plane. A method which might remedy this limitation is the use of multiple swarms, which will be discussed in section 9.4.

Long term evolution

In the long term the designed global swarm orbit is noticeably more unstable than the smaller scale counterparts. Despite to the lack in initial motion within the z_B axis the swarm starts to oscillate in this direction over time from third body perturbations, and eventually this resonance causes the swarm to depart the stable envelope around L4 in the fourth year of the orbit. After leaving the stable orbit envelope the swarm enters a chaotic Earth centric orbit, which eventually ends up in a slingshot manoeuvre around the moon into a global escape trajectory of the swarm. Figure 7.30 shows the result of the 5-year propagation.

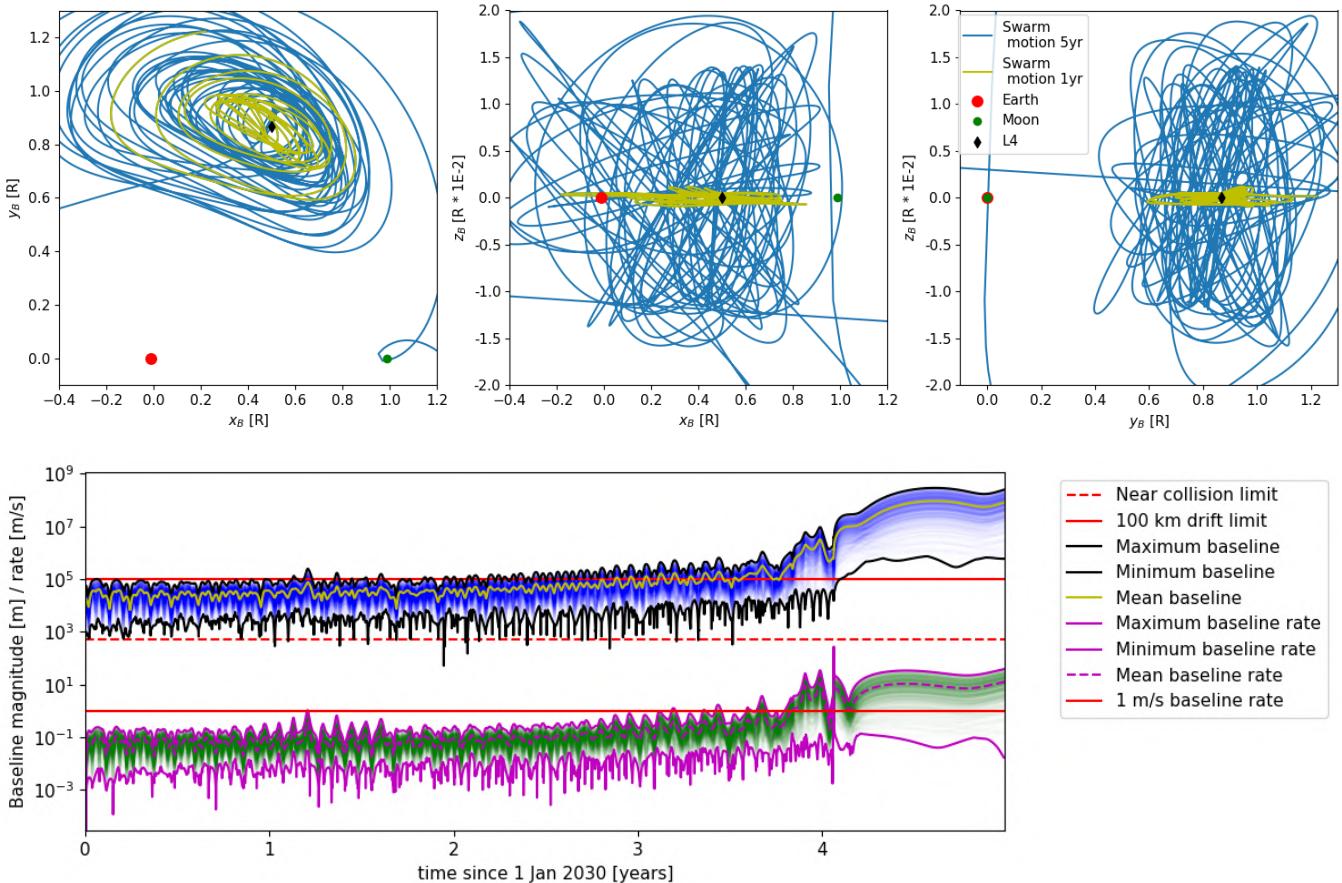


Figure 7.30: Orbit of the 35 satellite swarm in the barycentric frame over 5 years, with the first year orbit overlaid (top). Baseline magnitude and rate history over 5 years in orbit (bottom).

In the baseline history the slingshot around the Moon can be distinguished by the sharp spike in relative velocities just after the 4th year in orbit. From this point on the swarm is naturally irrecoverably lost, meaning that this orbit has a well-defined end-of-life. Apart from the disastrous end, the first four years of this orbit are remarkably stable. Much like the 20 element swarm design the average baseline of the solution remains below 100 km for almost 3 years, with only 2 significant near-collision events occurring in this period. The biggest improvement of this swarm design over the earlier solutions is the step up to a total of 1190^6 individual baselines, requiring only 45.38 days of swarm observation time to reach the sensitivity limit.

With minor course adjustments along the way this design should serve for at least 3 years in orbit, before it might become too expensive to further maintain the swarm's cohesion. The slingshot manoeuvre of this found solution might not represent the ideal result, but on the other hand it does represent an opportunity to dispose of the entire swarm at little fuel cost. The topic of de-orbiting a L4-centric swarm design, as well as placing the swarm in the necessary initial state, is left as a recommendation for future study.

⁶Counting negative pairs.

7.7 Comparison of swarm designs

The optimal designs for different swarm sizes show some interesting similarities, but also some differences which might be even more worthwhile to investigate. The differences between these designs raise the question whether they are imposed by differences in the optimisation process, or if they offer different features necessary to support larger swarms.

Swarm geometry relative to core

One aspect which was not discussed in detail so far was the overall swarm geometry design relative to the core position and its initial velocity. Figure 7.31 shows how the barycentric swarm designs relative to the core positions compare for 20, 25 and 35 satellite optima. The common trend which can be seen is the design toward "columns" in the z_B direction around the $z_B = 0$ axis. This design facilitates the process of swarm folding throughout the swarm's flight. Unlike the smaller swarm designs these columns are placed adjacent to the velocity vector in the $x_B y_B$ plane, instead of directly behind it. The direction in which this column is offset from the main velocity depends on the relative initial placement of the swarm, where it tends to be on the inside of the resulting global orbit.

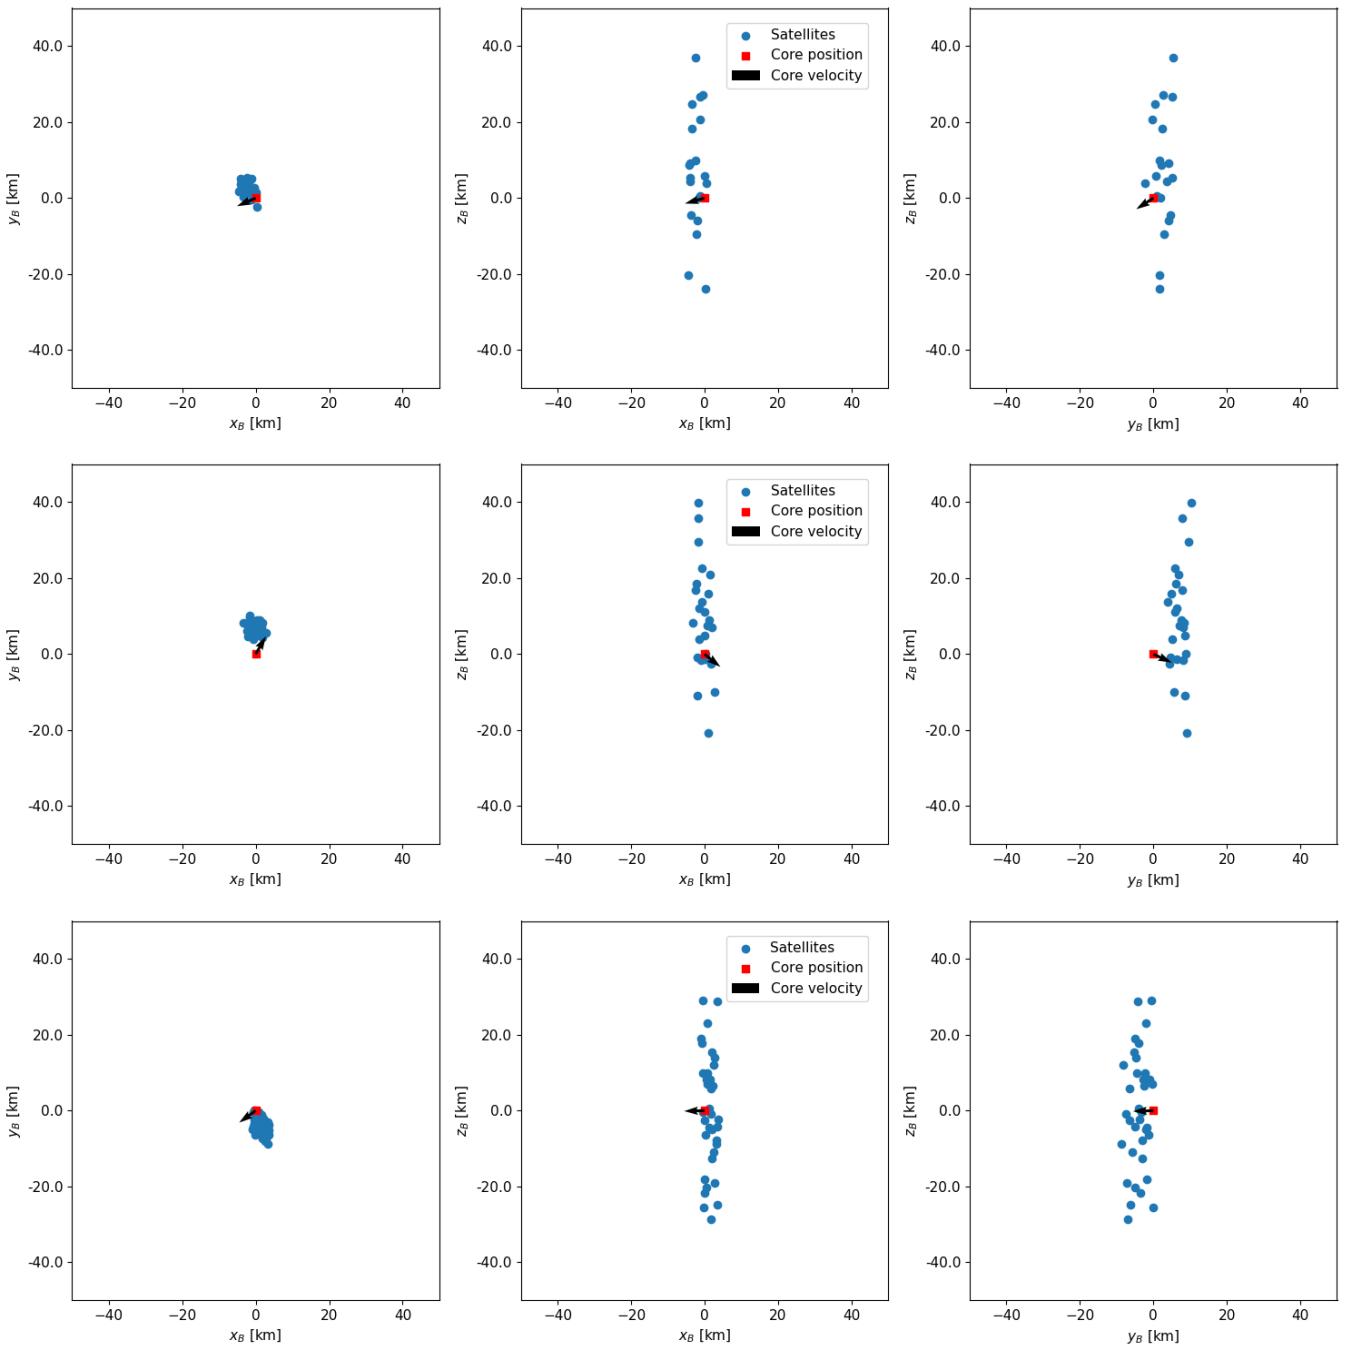


Figure 7.31: Initial swarm placement relative to the core in the barycentric frame for 20, 25 and 35 satellite swarm designs.

Swarm orbit design

For ease of comparison Table 7.1 contains the initial placement and velocity vectors for the three presented swarm designs in the J2000 and barycentric reference frames.

Table 7.1: Comparison of swarm core position and velocity relative to L4 on the first of January 2030, from optimisation results for different swarm sizes. The right-hand side of the table contains these coordinates and velocities converted to the barycentric system.

N Satellites	20	25	35	20_B	25_B	35_B
Core position						
C_x [km]	38.0670	4451.5958	122.6701	-40.7376	-2391.1283	-214.0109
C_y [km]	17.9474	45.6624	204.0254	20.2554	3730.5595	-0.4225
C_z [km]	18.442	-3.7657	-16.3404	6.4395	-429.2273	-105.5483
Core velocity						
\dot{C}_x [m/s]	-3.539	2.182	-6.995	2.2812	8.0018	-1.1748
\dot{C}_y [m/s]	10.415	-13.298	14.720	13.0504	-10.6626	17.3558
\dot{C}_z [m/s]	0.430	-12.380	5.159	2.1295	-10.6803	6.8591

As expected the largest spread in initial swarm locations is on the $x_B y_B$ plane, and the results show that the allowed solution space in the z_B directions was much larger than necessary. Every solution introduces its own motion in the z_B direction in addition to a velocity along the $x_B = y_B$ plane, and all of them follow a column design around the $z_B = 0$ plane to facilitate the folding swarm motion. The directions of the initial relative swarm velocities all differ from what would be the natural direction of orbital motion at the core positions. This initially places the swarm into a relative orbit which counteracts what would be expected from it naturally. The placement of the swarms into these counter-orbits seems to aid the swarm in counteracting the natural decay of the L4-centric orbit, extending the longevity of the orbit around the Lagrangian point. In turn this motion is very beneficial to the cohesion of the swarm, as it is exposed to much smaller potential gradients in close proximity of L4. A secondary benefit of this design choice is that it results in a lot of relative motion between elements, making it a welcome addition towards filling the uvw sphere.

A question which is raised by these results is why the different swarm designs follow different global orbits, despite having very similar initial swarm column structures. The differences might indicate that multiple global orbits are feasible for a single folding swarm design, or it could indicate that certain orbital requirements change as the number of satellites in the swarm increases. The second-stage optimisation algorithm was used to garner some insight into this factor, by optimising a 35-satellite swarm onto the global orbit determined by the 25-satellite swarm design. The second-stage algorithm successfully designed a 35-element swarm on this orbit, showing that there is definite leniency in the global swarm orbit design. The next section will discuss the results of the second-stage optimisation algorithm.

7.8 Second-stage optimisation results

While the results obtained using the first-stage optimisation are already convincing regarding the suitability of the fourth Lagrangian point for interferometer swarm orbits, they do not sufficiently explore the limits of this deployment location. The second-stage optimisation (subsection 6.1.3) was designed to push found solutions further. It was used to study whether the design period of a swarm design could be extended to two years, as well as whether a found global orbit of a smaller swarm could facilitate a large swarm design. Appendix C.4 contains the design tables of the second-stage optimisation results.

Lifetime extension of the 35-element swarm design

The first question which will be treated is the possibility of extending the design lifetime of a swarm solution. The first-stage results were all optimised towards passive formation flight for a full year, and the main question is whether the existing designs can be extended towards two years. To do this the global orbit and swarm velocity of the 35-element design were re-used, but the second-stage algorithm was allowed to re-arrange the 35 satellites and assign relative velocities to search a configuration which is passively stable for two full years.

Being given nearly a full week of computation time, this optimisation was unable to find a design which would be fully compliant with the requirements over a period of two years in orbit employing passive formation flight. Figure 7.32 shows the baseline history of the original and adjusted swarm designs. It can be seen from the evolution of these baselines that the adjusted swarm design is more compact, leading to much less occasions where baselines are too large and more near-collision events compared to the original solution.

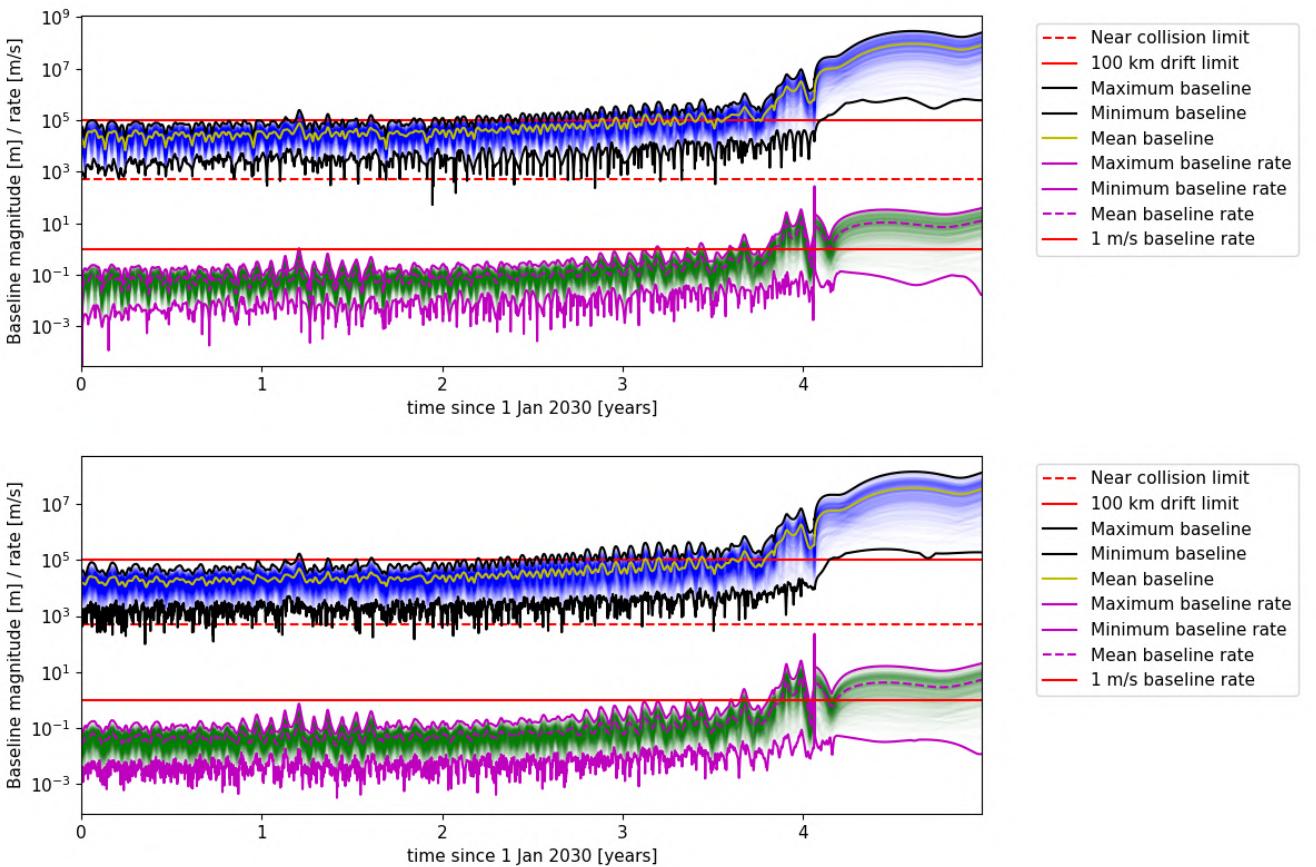


Figure 7.32: Baseline magnitude and rate history over 5 years in orbit for the original 35-element swarm design (Upper) and the second-stage optimised design (lower).

Being unable to find a perfect optimum, the algorithm optimised towards a solution which minimizes baseline overshoot at the cost of much more frequent near-collision events. Reflecting on these results this is a result of a bias in the cost function which equally penalizes baseline overshoots and near-collision events on a per-baseline basis. When a single satellite wanders from the swarm multiple of its baselines will overshoot the 100 km limit, whereas with near-collision events only singular baselines breach the 500m requirement. As a result the cost function presented in subsection 6.1.1 is biased towards designs which favour near-collisions over temporary wandering when faced with problems which it cannot fully resolve. Considering the severity of a potential collision compared to the temporary loss of a few baselines, this behaviour is far from ideal. Based off these realisations small changes to the cost function are recommended in subsection 9.5.5 which improve this behaviour for future studies.

Within the resulting baseline history it could be seen that the design found after this optimisation is much more compact during flight than the initial design, but it is also based on some fundamental differences compared to previous designs. Figure 7.33 shows the initial swarm geometry compared to the core position for the original design and the adjusted design. In this comparison it can be seen that this results starts to deviate from the previous trend of columns in the barycentric z_B direction.

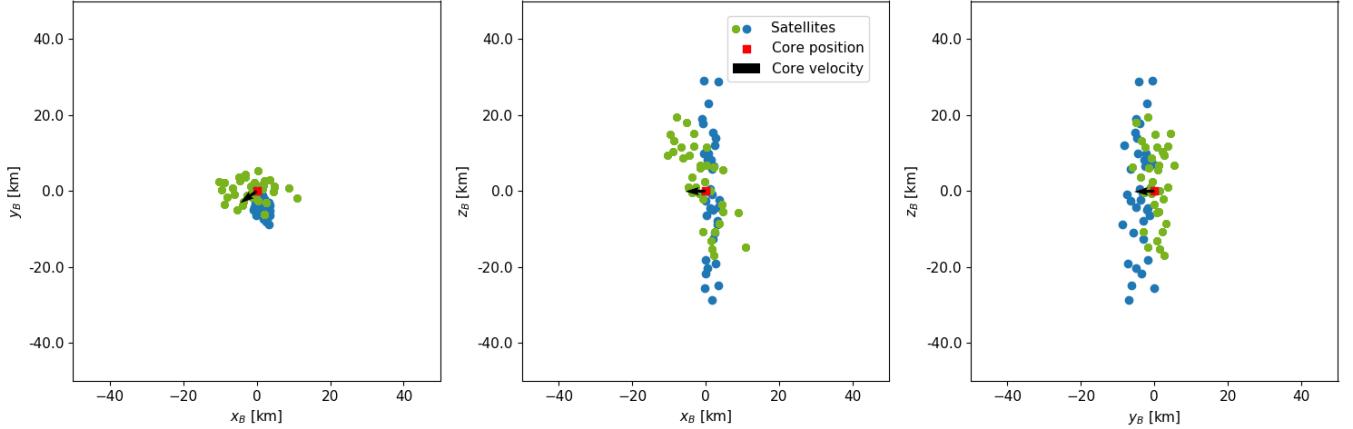


Figure 7.33: Initial swarm designs in the barycentric frame relative to the core position, for the original design (Blue) and the adjusted design (Green).

While the relative motion of the swarm is much more compact, it is clear from this comparison that this is not the case for the initial swarm configuration. The inclusion of relative motion between satellites allows for a larger spatial distribution at the start position, granting much more freedom compared to the fixed velocity designs. The second-stage optimisation allowed for a large range of up to 5 m/s in any of the cardinal directions for initial relative velocities, but this design has initial relative velocities of 5.89 cm/s at maximum. In addition to showing that the given problem boundaries were too large, this also shows that near-uniform initial velocities are preferable for long-term swarm cohesion. The largest difference in the overall motion of this swarm is the use of the folding process, or to be precise the lack thereof. Figure 7.34 shows the motion of the satellites relative to the core during its first year in orbit. It can be seen that the periodic folding motion shown in Figure 7.12 is now nearly non-existent.

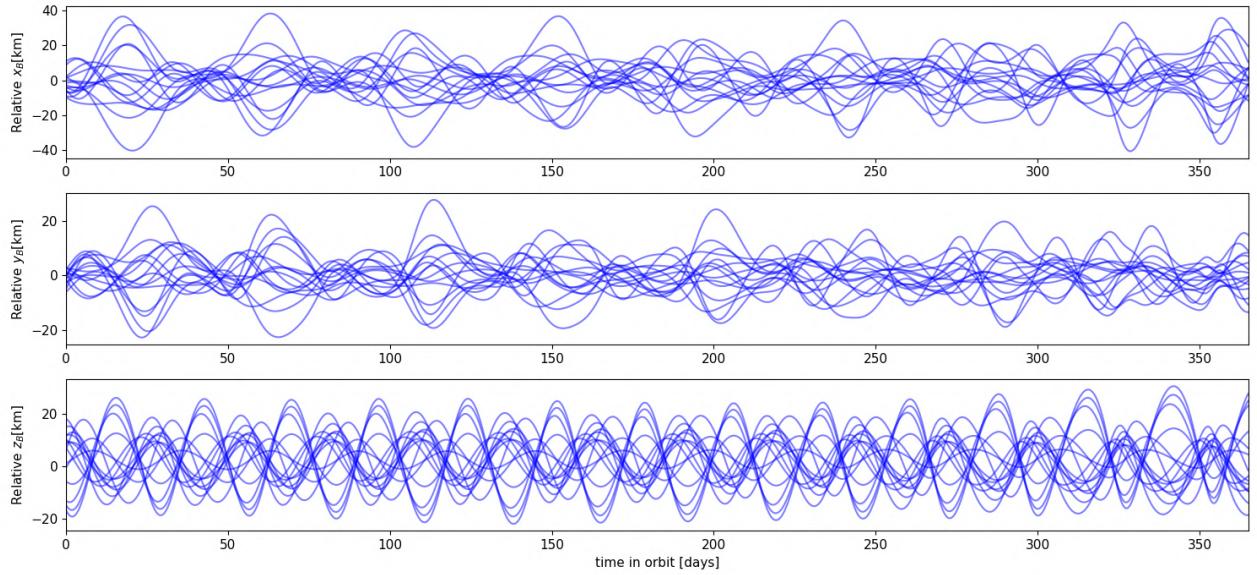


Figure 7.34: Motion of individual swarm elements relative to the core position in x_B , y_B and z_B over time.

The break from the harmonic folding motion may simply be explained by the inability to sustain this motion for periods as long as two years using passive formation flight. In Figure 7.12 it can be seen that the harmonic folding pattern of a 15 satellite swarm starts to fall apart after approximately 300 days in orbit. Being unable to find a folding design which is stable for two years, it makes sense that the optimisation algorithm strays from this solution. The fact that this design breaks with the established trend, in combination with the lack of success in finding an optima, suggests that a perfect solution may not exist for a two-year period. After all this would be a very long time to rely only on passive formation flight, even in an environment as stable as this. It may be concluded that while a period of a full year is feasible, achieving a perfect design for two years is likely to be impossible with a swarm of this size.

35-satellite swarm on the 25-satellite design's orbit.

The second-stage optimisation algorithm was capable of finding a near-ideal swarm design before being cut off after reaching the maximum amount of 75 generations after 6 and a half days. This cutoff was implemented since there was no guarantee that a design could be found at all, but in practice it has proven to be a few generations too short. In the end the choice for 75 generations was unfortunate, as the algorithm only needed to resolve 6 near-collision events to find an optimum. Figure 7.35 shows the swarm design relative to the core position in the barycentric frame and the resulting baseline history over 5 years in orbit. In the latter it can be seen that the design performs according to the requirements, barring 6 slight infractions on the 500 meter lower baseline limit. Had the algorithm been given a few more generations, there is little doubt that it would have found an optimum.

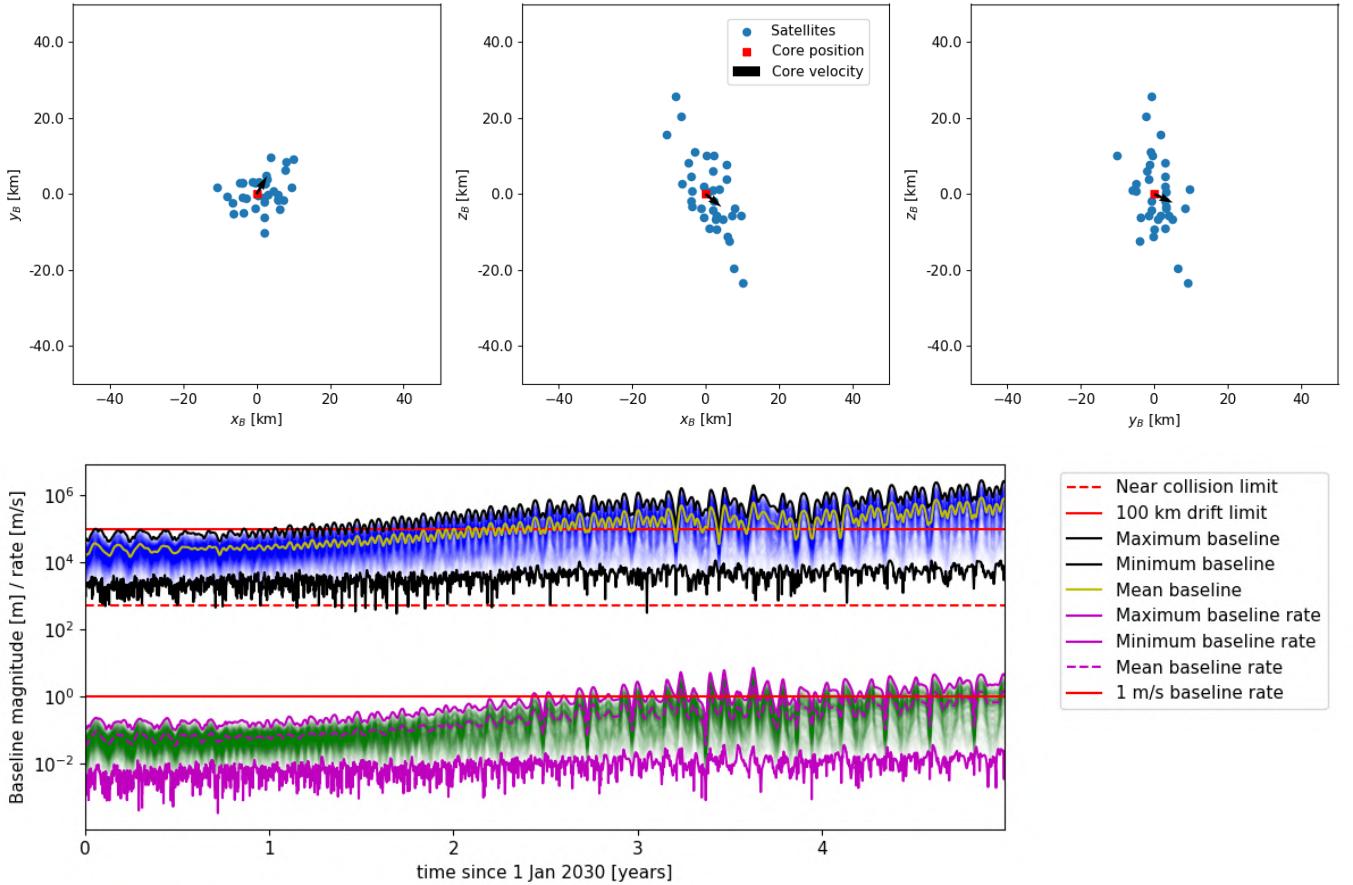


Figure 7.35: Absolute positional change of the satellite orbit from the change in body pointing, and the change in total system energy from the attitude adjustment.

It was already well covered that resolving near-collision events is much more difficult for the algorithm with swarms of this size, and the nature of the folding motion. It stood out however that this seemed much more difficult for this optimisation than ordinarily. It also stands out that the swarm design breaks away from the norm of vertical column design. Compared to the original designs this pillar design is much more spread out and slightly skewed towards the direction of the initial velocity. In the previous section this break from the norm was due to the extension to a longer 2-year lifetime, but in this case this break from the norm is entirely related to the change of orbit. This deviation from the norm confirms the suspicion that swarm designs are inherently bound to global orbits, and that they ideally ought to be resolved together. Despite this, this result shows that it is possible to scale up the number of satellites along an existing orbit, although the result will stray from the idealized column design.

The baseline history of this orbit design shows a very similar profile to the 25-satellite design, with a similar lifetime of roughly 2 years before the average baseline exceeds 100 km. This indicates that the global lifetime properties of passively flying swarms are primarily dictated by the orbit of the swarm, instead of the number of satellites. Compared to the design presented in section 7.6 this orbit offers a longer potential lifetime, since it remains in orbit of L4 for at least 5 years compared to the original's 4. The potential of a longer lifetime comes at a significant cost however, the cohesion of the swarm is significantly worse than the original orbit of the 35-element design. When compared to the data shown in Figure 7.30 the baseline history presented in Figure 7.35 is more divergent, meaning that the swarm would need to consume much more fuel to retain cohesion using this orbit.

Even when given the same freedom as the previous result, this design again tends towards the use of swarm folding to retain long-term cohesion. Despite the swarm geometry differing from the column design, the relative velocities of the satellite are assigned in a way in which the swarm still facilitates the folding process. Creating a clean fold is much more difficult however, which is evident when Figure 7.36 is observed. This figure shows the position of individual swarm elements relative to the swarm core over time:

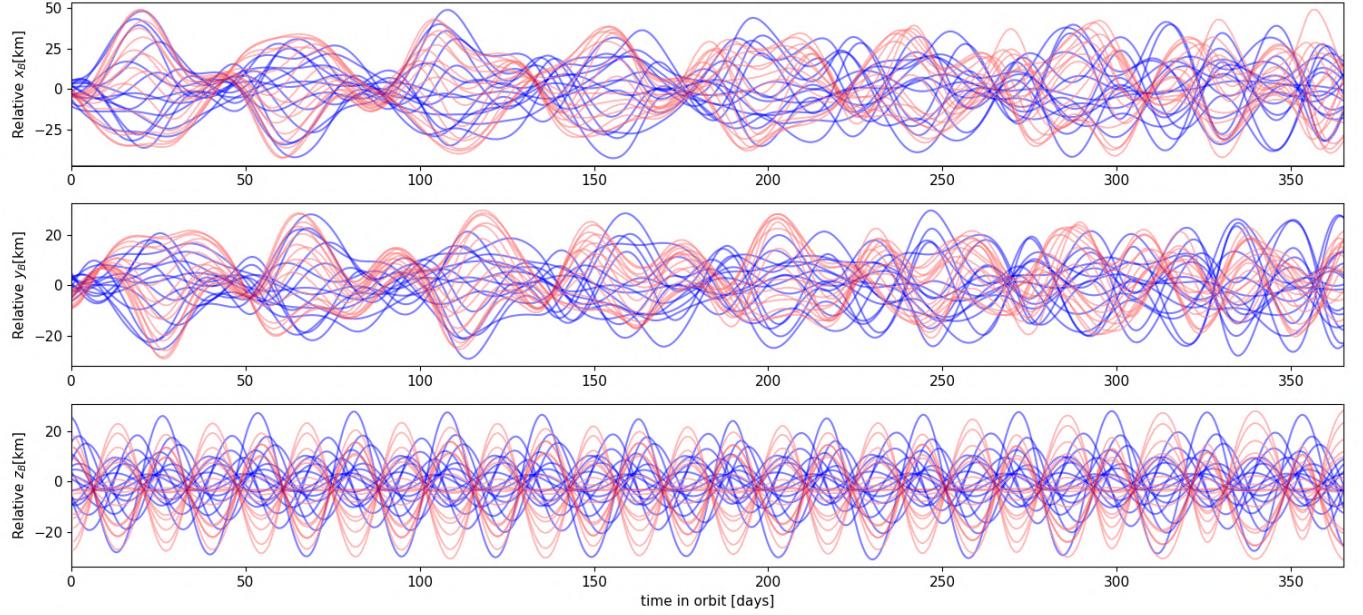


Figure 7.36: Motion of 35 individual swarm elements relative to the core position in x_B , y_B and z_B over time (Blue). Motion data from Figure 7.36 is overlaid in Red for ease of comparison

When compared to Figure 7.12, which showed the same data for a 15-element swarm design, this resulting relative motion is much more chaotic. It still consists of a combination of sinusoidal motions, but due to the introduction of initial relative velocities the folding motion is far less cohesive. The addition of velocities between swarm members is hardly beneficial for 1-year design periods, as the designs with fixed velocities achieved better results. Ideally velocity variables should be removed from short-term design algorithms entirely, which also aids in reducing the required computation time.

Changing the number of satellites from 25 satellites to a 35-satellite swarm while retaining the same orbit has created some interesting changes in the overall swarm design. Only 7 out of the 35 satellites were granted initial velocities in the order of centimetres per second compared to other swarm members, other relative velocities were minimized. This indicates that these changes are the result of a forced adaptation of a 35-element swarm to an orbit which was designed for 25 elements, rather than a byproduct of the freedom in initial velocities. It is not fully understood exactly which mechanisms drive the adaptation of the swarm design to a different orbit, as this single set does not provide sufficient data. Understanding which mechanisms force the swarm design to adapt on non-ideal global orbits might allow for the estimation of the maximum swarm scale compatible with a given orbit, and the topic of design adaptation to different orbits is thus highly recommended for future research.

7.9 Collision detection and interpolation frequency

It was previously noted that the data from orbit optimisation had to be interpolated to 4-hour intervals in order to keep data sizes manageable. This time interval is bigger than ideally would be desired, and one concern is the risk of undetected collision between interpolation points. This section addresses this concern by determining just how likely an undetected collision is for the given orbit design problem. Attention is paid only to the risk of actual (near) collisions, "grazes" of the 500 meter bubbles are not as much a concern due to the low relative velocities.

In order for a collision to go undetected a pair of satellites need to move in, and out, of the collision detection/warning bubble of 500 meters between two interpolation points. For a head-on collision this means satellites need to traverse a distance of twice the detection bubble, 1km, in the 4 hours between interpolation points. This requires both satellites to travel at speeds of 6.944 cm/s, for a relative velocity of 13.888 cm/s. The absolute worst-case scenario for missing a collision is imaged in Figure 7.37, in this case the collision velocities are orthogonal, resulting in a minimal move distance R_D between two interpolation points within triggering a detection, as long as R_0 and $R_1 \geq 500$ meters.

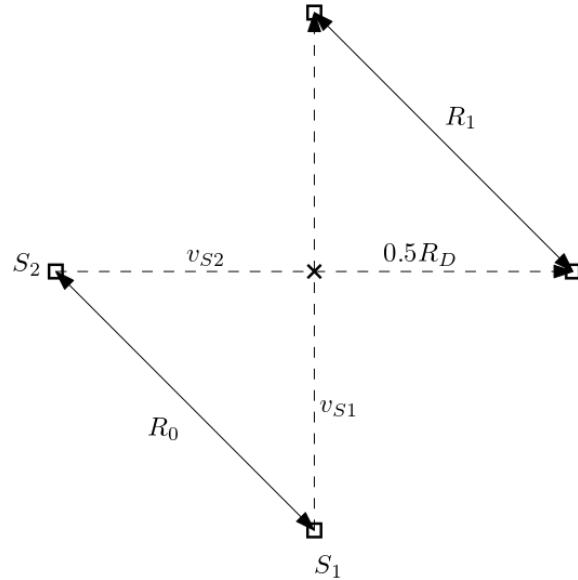


Figure 7.37: The worst-case collision geometry with the smallest geometrical move distance is an orthogonal velocity collision.

Consider a bare minimal detection escape with a distance barely larger than 500m, which for the sake of convenience will be rounded down to 500m. Knowing that the velocity vectors are orthogonal, it can be computed that the distance satellites need to traverse between interpolation points is $2 * 0.5R_D = 707.11m$. Using an interpolation time of 4 hours, both satellites need relative orthogonal velocities of at least 4.910 cm/s to pass by undetected. Considering that the relative velocities are orthogonal the relative velocity between these satellites can be found using a square triangle, yielding a velocity of $\sqrt{2} * 4.910^2 = 6.944$ cm/s. Using this worst case scenario it can be concluded that the system can only fully guarantee collision detection for relative velocities up to 6.944 cm per second. Figure 7.38 shows all observed baselines and associated relative velocities of the 35 satellite swarm solution during its design period of a full year, as well as the 6.944 cm/s observation threshold.

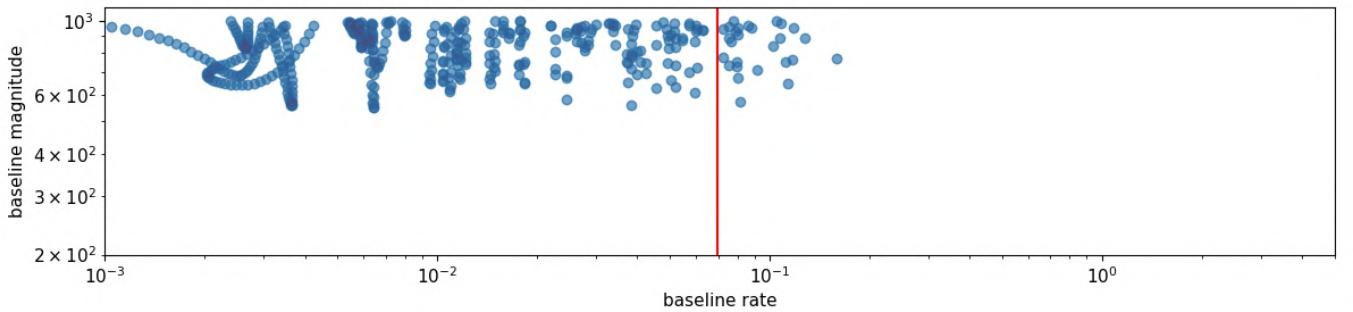


Figure 7.38: Selection of baselines smaller than 1 km and their associated baseline rates observed during the first year in orbit for the 35 element swarm design.

Off all 828 occurrences 54 baselines (6.522 %) had relative velocities larger than the worst-case observation threshold, none had a velocity larger than the best case threshold. This means that 6.522 % of the observed small baselines pose risk of a potential undetected near-collision, if these satellites have near-orthogonal velocities. Figure 7.39 shows the superimposed angles between all satellite velocities for the 35-element swarm during its year in orbit. It can be seen that during normal flight the angle between velocity angles is very small, making it virtually impossible to have a (near) collision event go unnoticed.

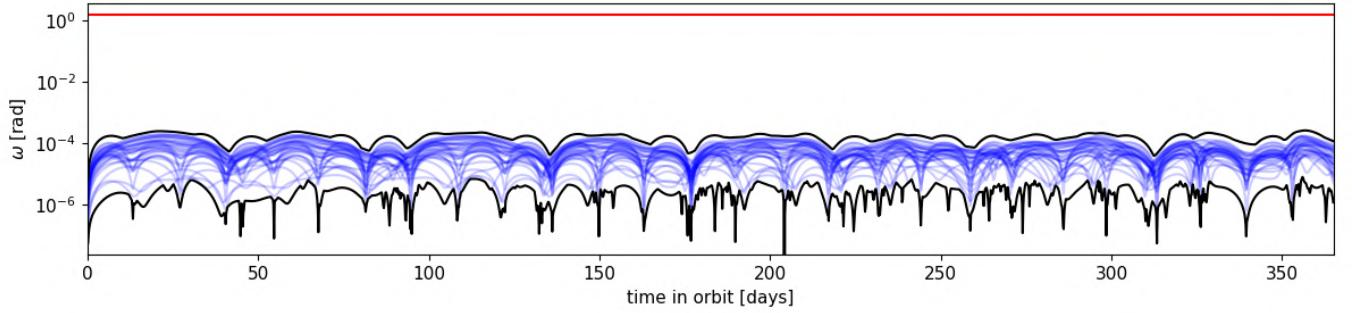


Figure 7.39: Angle between satellite velocity vectors ω in flight for the 35-element swarm design, with a red line denoting the angle for orthogonal velocities ($\pi/2$)

While the system is not perfect with this interpolation time, the chance that a collision goes undetected is extremely small given the observed trajectories and detection limits for relative velocities. Nevertheless a higher interpolation frequency is recommended for future studies, in order to exclude the possibility of freak accidents.

8 Conclusions

The suitability of L4-centric orbits for radio interferometry

Historically L4-centric orbits have not been studied in detail for the use of radio interferometry, since they continually expose the instrument to Earth's radio frequency interference. Based on previous measurements with the WIND/WAVES instrument it is expected that compensating for this interference around the L4 point requires 9dB of dynamic range, which is not ideal but it could be bridged with modern systems. This work used this assumption to study the potential of L4-centric orbit design for interferometry swarms, orbits which offer very promising features.

The orbit design for the OLFAR mission should primarily serve to facilitate the purpose of radio interferometry, which relies on the adequate distribution of baselines over the mission lifetime and sufficiently low relative velocities to allow for data processing. Assuming a diffusion-limited resolution for the interferometer the baseline between any satellite pair should not be larger than 100 km at any point in the orbit, and with regards to orbital safety it should at all times be larger than 500 m. Due to data processing limitations the relative velocities between all satellites in the swarm need to be smaller than 1 m/s for the OLFAR radio interferometer. As a result radio interferometry requires swarm orbits which keep the swarm concise and cohesive with low relative velocities, yet which offer large enough swaths of relative motion to cover the entire uvw sphere.

In this work it is demonstrated that swarm orbit designs up to at least 35 satellites may be found which meet all of these requirements for at least a full year in orbit, while only utilising passive formation flight. This period could easily be extended to 3 or 4 years with only minor course corrections for collision avoidance. L4-centric swarm orbits make for very promising candidates which can deliver on very broad uvw space coverage, long mission lifetime with minimal need for manoeuvring, and relative velocities in the order of centimetres per second. Regarding the posed research question it has to be concluded that L4-centric orbits are very suitable for radio interferometry constellations, if the exposure to Earth's RFI can be compensated for.

8.1 Filling the uvw space for radio interferometry

It is demonstrated that L4-centric radio interferometry swarms can achieve very thorough uvw space coverage over mission times due to their wide range of natural motion and the use of swarm folding for increased cohesion. L4-centric swarms can be designed to be cohesive with minor corrections for up to 3 years, during which near-perfect coverage in the uvw space can be obtained. Figure 8.1 shows the achievable baseline profile of the 35 element swarm design found during this thesis, which highlights very good saturation on the xy plane in particular.

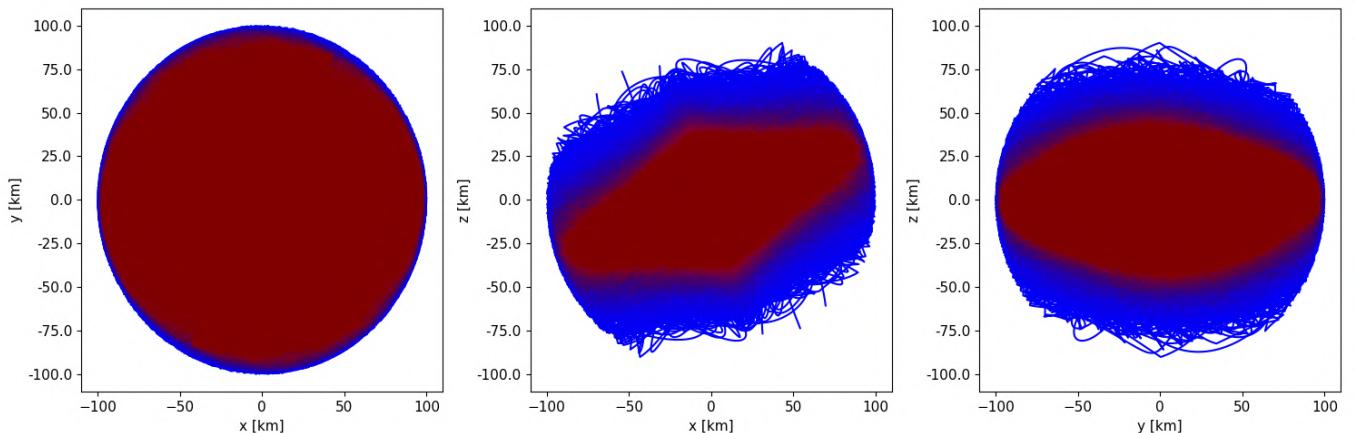


Figure 8.1: Projected views of valid baselines in the J2000 frame of the 35-satellite swarm design, after four years in orbit. Red accents indicate areas of higher baseline saturation.

By exploiting the rotation of the L4 point around Earth radio-interferometric systems can achieve near-optimal baseline coverage on the xy plane of the J2000 reference frame after a single year, with longer mission lifetimes this coverage can be perfected. Coverage of the xz and yz planes on the other hand is limited, as these view projections require significant motion outside of the Lunar orbital plane which is not achievable while maintaining a cohesive swarm. This is a limitation that any swarm design in a orbit along the Lunar plane will experience, if anything the large freedom of movement of L4-centric orbits allow the swarms to better cope with this limitation.

This shortcoming might be compensated for by using multiple smaller swarms. Orbits around the L4 point theoretically share orbital periods in the barycentric xy , xz , and yz planes regardless of their exact shape. This could be exploited to design multiple smaller swarms which periodically meet, using cross-swarm measurements to achieve better baseline distributions in the z direction. This is a very complicated topic of study however, which will be left as a recommendation for future work.

8.2 Longevity of passive formation flying swarms around the L4 point

In this work it was demonstrated that passive formation flight is possible up to periods for at least a year in a perturbed environment with up to 35 satellites, while fully meeting all mission requirements. After this initial year the optimised swarm design remains very serviceable for interferometry up to 3 years while relying only on passive formation flight, before the average baseline becomes too large for measurement. During this time some minor corrections might be necessary to maintain the 500 meter separation between satellites. When such a manoeuvre is required average relative velocities range below 8 cm/s, meaning that such corrections require the use of little fuel.

The 1-year design period for the swarm geometries may easily be extended using more sophisticated optimisation techniques, after all the methods used to design these swarms are relatively crude. By fine-tuning the initial relative velocities of the swarm members it might be possible to extent the envisioned swarm lifetime significantly. Passive formation flight will eventually reach a limit to which it can be relied upon. Regardless of its orbit or design the swarm should have some method of active course correction, the application of which might significantly extend the envisioned lifetimes of these designs. Extending the viable lifetime and size of these swarms through the application of active formation control is highly recommended for future research.

The major limitation to swarm longevity is the design of the general swarm orbit, as long as the swarm remains in the stable region around the L4 point the relative movement between elements is predictable and easily corrected. At a certain point in time however, it is inevitable that the orbit decays and the swarm starts to orbit the Earth in a chaotic fashion. At this point maintaining cohesion of the swarm becomes extremely challenging, and the chaotic orbit of the swarm becomes a threat to any other missions. The subject of finding the most-suitable global swarm orbit, which allows both for very large swarm designs and long-term global swarm stability is recommended for future work, since it is expected that this work only scratches the surface of the possibilities.

8.3 Swarm orbit design around the L4 point

Throughout all found solutions a common trend is the initialisation of the orbit in close proximity to L4 with a slight disposition from the L4 point in direction of the barycentre, located with a slight offset from the orbital plane of the Moon. The initial velocity of all swarms relative to the L4 point aligns with the $x_B = y_B$ plane of the barycentric system. The resulting orbit is located on the approximate $x_B = y_B$ plane of the barycentric system for the initial weeks, before the slow clockwise growth expands the orbit to cover the x_BY_B plane.

By starting the initial orbit along the $x_B = y_B$ plane the swarm design has maximum resistance to the natural clockwise growth of this orbit along the x_BY_B plane of the barycentric system. By introducing an initial velocity in the z_B direction as well this resistance is increased further, as the clockwise growth is strongest around the Lunar orbital plane. Inducing an initial orbit around L4 on the $x_B = y_B$ plane is beneficial both for longevity of the swarm's global orbit, as well as its uvw space coverage. The inclusion of a large movement in the z_B direction allows for the swarm to achieve much more baseline coverage than a coplanar orbit could achieve.

An unavoidable conclusion to the natural decay of the global swarm orbit is that instead it must be opted to de-orbit the swarm from the L4 point in a very controlled fashion. From the L4 point it is possible to de-orbit the swarm in a fashion where it would escape from the Earth-Moon system, which might be a more economical solution than attempting an atmospheric burn or finding a suitable graveyard orbit. A second recommended topic is studying methods to insert swarms into these orbits, which has not been touched upon during this work.

8.4 Swarm folding

The swarm designs found in this thesis all made use of "column" swarm designs, which initially distributed the satellites as a vertical column in the barycentric frame, which is mirrored in the Lunar orbital plane. In combination with a shared initial velocity this results in a swarm which continually folds over itself relative to its core position. The application of this folding was not planned, but it clearly greatly improves the duration for which the swarm remains cohesive during passive formation flight. In addition to being beneficial for swarm cohesiveness, the folding motion also greatly enhances the filling of the uvw sphere by introducing very dynamic relative motion of the swarm. As a result of this motion the initial velocity vectors of the swarm ought to be as cohesive as possible, preferentially the velocity vectors of all swarm elements is uniform. This property is very beneficial towards optimisation processes, as it allows for a large reduction in the number of variables.

The folding motion is very helpful for retaining swarm cohesion while relying on passive formation flight, but its harmonic nature decays over time due to external perturbations affecting the swarm elements. The designs presented in this thesis show that 300-day long folding motions are feasible without need for correction, but that periods beyond a year require occasional manoeuvres to sustain the folding process. While the lifetime of this motion may be increased, the decay of the harmonic folding motion is inevitable once the swarms orbit around L4 grows larger. As the orbit grows the swarm is exposed to stronger potential gradients, making it continually more difficult to correct for its natural expansion. Breaking away from the folding pattern is not disastrous however, as the small-scale nature of the relative folding orbits still serves to keep the swarm concise for much longer after the folding motion is lost. This allows for much longer feasible mission lifetimes, even though the interferometer will be less efficient from the loss of some baselines.

The downside of the folding motion is the inherent risk of near-collision events at folding points. This becomes increasingly more difficult to avoid when designing for larger swarm sizes, which was already noticeable between the process of designing 25 and 35 satellite swarms. In the latter it got very difficult to resolve swarm configurations without near-collision events, and for a larger swarm size this might even be impossible for prolonged periods of time. This may prove to become a upper-boundary to the maximum number of satellites which can be included in a folding swarm design for a given mission lifetime. It is expected that this boundary will exist as long as the minimum distance between elements needs to be strictly enforced, but the data within this thesis is not conclusive to provide an estimate for this boundary at a full year of passive formation flight. The designs show that the use of 35 satellites is feasible, at least.

Near-collision events are by far the biggest challenge to resolve for passive formation flying swarms, and in practice autonomous course correction will be a necessity for any swarm. The latter will be true regardless of deployment location for any large swarm. If anything handling near-collision events is slightly more forgiving for L4-centric orbits, since the relative velocities between satellites hardly exceed ten centimetres per second during near-collision events.

8.5 The flexibility of swarm orbits

The results of the second-stage optimisation show that the global orbit of the swarm around the L4 point is tied directly to the scale of the swarm and its column geometry. Despite the orbit being inherently tied to the swarm scale there is shown to be room for flexibility in its design. A 25-satellite swarm orbit design was successfully up-scaled to facilitate a swarm of 35 satellites for the same full year lifetime, but this required some deviation from the previously established trend of folding column swarm designs. Despite having more satellites and a different initial geometry, the upscaled swarm design inhibited similar feasible lifetimes and expansion patterns to the small-scale design. This shows that the global orbit of the swarm primarily determines the feasible lifetime and natural expansion rate of swarm designs, rather than the number of swarm elements or their initial configuration.

With how important the design of the global orbit is now understood to be, it is promising to know that there is some inherent flexibility regarding the swarm scale and the orbit. It means that larger swarm designs could theoretically be fitted to work with non-ideal orbits, which may provide better long-term stability or easier insertion options. Adapting a design to a different orbit requires some changes from the folding column design, the studied case still represented a column but it was spread out and skewed towards the initial velocity vector relative to L4. The adaptation of this design worked to increase the swarm size by 10 elements, but the resulting folding motion was much less cohesive. Designs adjusted to sub-optimal orbits will thus always be less efficient than designs using orbits optimal for their scale, requiring more fuel investment to maintain the folding motion over prolonged periods of time. Despite this the fuel consumption of such designs can still be very low, making it an attractive option when the optimal global orbit is problematic.

The exact mechanisms behind this adaptation are still not fully understood, as there is insufficient data to study at this point. Understanding what drives certain changes in the swarm designs may be vital to estimate the maximum swarm scale a certain orbit might support, making it a prime topic for future research.

8.6 Conclusion

All things considered it is concluded that L4-centric orbits offer substantial potential for the design of radio interferometry swarms. Within this deployment location orbits can be found which offer cohesive swarm designs with broad uvw space coverage whilst requiring very little fuel to maintain swarm orbits suitable for interferometry up to three years. Orbit designs up to 35 swarm elements have been demonstrated thus far, but it is expected that these results only scratch the surface of the potential of L4-centric orbits. Based of these results several recommendations are made to improve the method developed for this work, which will allow for more efficient searches for even larger swarm designs.

The different swarm designs show a clear trend towards a column swarm designs which facilitate the folding motion. Knowing this design tendency the devised methods can be improved by the application of a more targeted search. An additional area of potential improvement is the implementation of active formation flight to the swarm designs. With the occasional course correction it should be possible to push the viable lifetime of these swarm designs up to 4 years or longer, depending on the global orbit of the swarm. Research into improved optimisation schemes and the application of active formation control is thus highly recommended for future work.

9 Recommendations

A common saying is that science often raises more questions than it answers, and this is certainly the case for the work presented in this thesis. Although it was attempted to cover as much subject matter as feasible under this work, there are still many remaining questions and topics which require further study. As final remarks, this chapter contains identified topics of interest for future work and recommendations to improvements of the presented methods. The sectioning of this chapter defines broad categories, in which more specific recommendations are given.

9.1 The application of body pointing for formation control

In section 7.2 it was demonstrated that the OLFAR swarm will be very susceptible to the influence of solar radiation pressure, which extends to the attitude and body pointing of its individual satellites. Solar radiation pressure is identified as both a danger and a potential benefit to maintaining swarm cohesion, both of which are recommended as avenues for further study. This topic is particularly interesting considering the potential improvements that may be made to the presented designs through the introduction of active formation control.

The application of body pointing and attitude control for manoeuvring

A promising subject of study is the use of body pointing and solar radiation pressure as a means for relative navigation and course correction in orbit of the triangular Lagrangian points. It was shown that 6 hours of exposure with a different attitude was sufficient for a course correction of a meter, which is not insignificant for a method which requires no fuel. Future research is recommended to study the potential of applying body pointing for relative satellite manoeuvring within the OLFAR satellite swarm. This application may provide a very cheap, albeit directionally limited, method of relative course correction or formation control.

The influence of solar radiation pressure on low-lunar orbits

With Solar radiation pressure being as influential as demonstrated it is recommended that it is studied for long-term effects on low-Lunar orbit designs as well, even though the effects are likely to be less severe in these conditions. The low mass of OLFAR elements and the large surface area make even short exposures influential if there is a difference in attitude. The potential occurrence of these differences certainly applies to Lunar orbit concepts as well, and it is recommended that these are studied with proper models. For this deployment the periodic occlusion from the Moon might also be a considerable effect, which ought to be included.

9.2 The use of the Point-Spread Function as cost function

The original intention for the optimisation strategy presented in this report was to include the PSF of the baseline profiles of a swarm design in the cost function of the optimisation algorithm. This idea had to be discontinued due to technical difficulties with the implementation in the TUDAT environment. subsection 3.6.2 presented a method of implementing the PSF in the cost function by quantifying the quality of the PSF through comparison with the Airy disk, and further study into the effects of optimisation through this metric is recommended.

The main body of work in this thesis has shown that optimisation can be used to find suitable geometries for passive formation flight, which also show some promising baseline profiles. Yet these swarm designs were not optimised for radio interferometry, and there is room to improve their baseline profile. It is shown that resulting uvw baseline distributions are not very good in achieving proper distribution along the J2000 z axis due to the nature of the Lunar orbital plane. By integrating the PSF into the optimisation function different stable orbits might be found which are better at rounding out the PSF. After all the only condition used thus far is the search for a stable swarm configuration.

The remaining question is whether the implementation of the PSF in the cost function can yield stable constellation designs which do a better job of covering the uvw space outside of the primary orbital plane. The use of this different metric might yield very different swarm designs which meet the same criteria, but do a better job for radio interferometry. Continued research in this application is highly recommended for future work.

9.3 Application of this optimisation method for low-lunar orbit

The core idea of the OLFAR mission still revolves around the use of low-lunar orbits, utilising the moon as a radiation shield. This environment offers much more challenging dynamics, making it difficult to find an orbital solution which meets all set mission requirements. The optimisation strategies developed for this thesis could equally well be applied to low-lunar orbit problems, after a short study to establish some changes to the environment model. The search of Lunar orbit swarm designs using this methodology is recommended for future studies, if anything it would be interesting to observe how optimised swarm designs compare to the designs which are currently found for Lunar orbits. Though the brute-force methods are not as sophisticated as analytical approaches, it can offer very surprising results¹ to design problems.

9.4 Investigating the use of multiple smaller swarms

With the swarm folding method there will be a definitive upper limit to the amount of elements which can safely be placed within a single swarm. Based on the results observed so far, it is not likely that this design method will safely support hundreds of satellite flying in a single close-proximity swarm. The envisioned scale of the OLFAR concept is bold, and it is not likely that this will be achievable in a single swarm if a L4-centric orbit is used, unless it is chosen to not strictly enforce the maximum baseline limits.

An alternative, and potentially more attractive, solution is the use of multiple smaller swarms in different orbits around the L4 points. When placed in orbit around the L4 point the orbital periods in the x_B , y_B and z_B directions are similar regardless of the shape of the orbit, according to theoretical motion models. This property might be used to have multiple smaller swarms synchronized their orbits, occasionally meeting in close proximity to allow for inter-swarm measurements. Such use of multiple smaller swarms could produce very promising results in regards to global baseline distribution and potential science yield. This topic is very challenging to take on however, as it requires first exploring the limits of single-swarm designs to properly investigate the use of multiple swarms. This topic promises to be very challenging, but it might yield some considerable improvements for L4-centric radio interferometry concepts.

9.5 Smarter methods of constellation design

The optimisation approach presented in this thesis was effective in achieving results, but doing so required a considerable time investment. This inevitably led to having too little time and resources to achieve a 50-satellite global swarm design. Based on this experience several recommendations are made for improvements to the presented optimisation approaches. The primary goal is to make the developed methods more efficient at finding results through a more targeted application and design.

9.5.1 Focussed solution spaces for optimisation

The resulting swarm design shows that the preferred initial geometry closely resembles a column in the z direction of the barycentric frame. In comparison the initial distribution in x and y directions is very small. A simple efficiency improvement would be to adjust the search space boundaries accordingly. For the purpose of the second-stage optimisation this is mandatory as well, in order for these designs to work the initial velocity needs to be approximately constant throughout the swarm. The boundaries which were used were simply much too loose, making the process extremely inefficient at resolving a design.

9.5.2 Solution-seeding

Seeding a genetic algorithm is a term used to describe modifying some individuals of the first generation to resemble what is expected to be the optimal result. This allows the designer to transfer some knowledge from previous results to new optimisation processes, potentially greatly enhancing the efficiency of the application. For the swarm design problem an example of seeding would be to use the 35-element swarm design as a base for a 40-element optimisation, which would put the entire algorithm at a much better starting state than if it is entirely left up to chance.

¹Such as the use of swarm geometries which fold over themselves.

Although the option of solution-seeding was desired, this could not be implemented in the limited timeframe of this thesis. Doing so would require alterations to the PagMo2 package which would be very time-consuming to make (and verify). For future studies the application of this concept is highly recommended.

9.5.3 Staggered optimisation

A secondary method to efficiently grow the size of the optimised swarm is to apply a staggered optimisation. In this approach an existing swarm design is taken, and only a single satellite is fitted to be added to the existing design. The resulting problem is of much smaller scale, and it should be suitable to expand the satellite count of existing swarm designs by a handful of satellites or more.

The downside of this method is that the results will be very dependent on the initial configuration, and the order in which satellite placement is resolved. For example it might not be possible to grow a 25 satellite swarm to a swarm of 35 satellites, while a design for the latter can be found through a large-scale approach. This optimisation approach should not be used to build swarm designs from the ground up, but it could be valuable to expand existing designs.

9.5.4 Swarm design adjustments through linear approximation

In the results of the perturbation analysis for the different initial baselines there is an apparent linearity to the magnitude of displacement caused by the perturbations over time. This prompted for a short additional investigation into the linear behaviour of the effects from the perturbation sources. Figure 9.1 displays the change of the final baseline between the two satellites after a year in orbit due to solar gravitational perturbations. The initial range between this satellite pair is varied for 6 simulations, wherein the second satellite is displaced from the first in the x direction of the J2000 frame. The figures also show a fitted linear relation between these points.

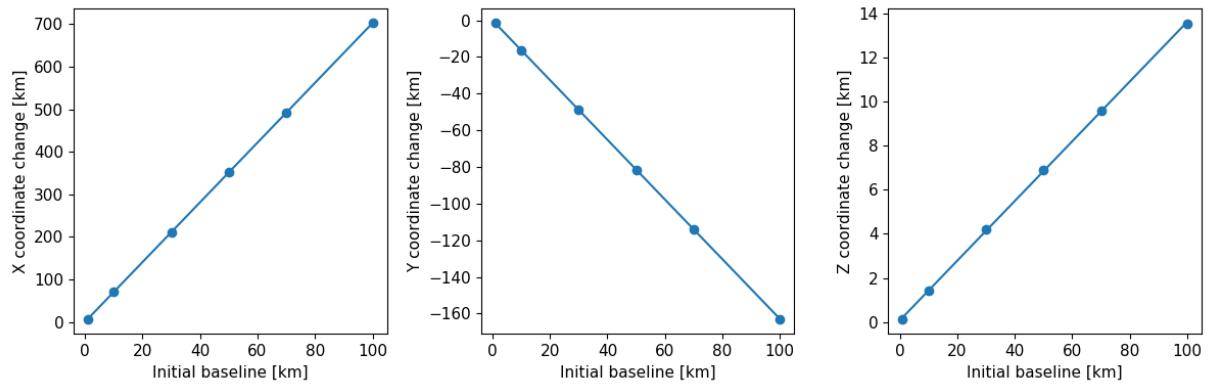


Figure 9.1: Magnitude of baseline coordinate changes in x,y, and z directions of the J2000 frame from Solar gravitational pull, modelled as point mass. Lines are fitted linear functions utilising linear regression, yielding $R=1.000$ for all plots.

While the sample size is arguably small, this dataset shows a clear linear relation between the initial satellite displacement and the magnitude of effect of solar gravitation over time. Applying a linear regression fit yields an R -squared value of 1.000 when rounded to three decimals, further confirming the notion of the linear relation. Considering the size of the initial baselines compared to the distance of the perturbing source, which in this case is the sun, it is not all that surprising to find a near-linear relation.

These near-linear relations for the effect of perturbation sources might prove to be useful in fine-tuning swarm designs in a more controlled manner. If² all perturbation sources could be approached using a linear model a change in the initial placement of a satellite could easily be propagated to a linear change of the propagated orbit. This might allow for a much more time-efficient method of swarm design that needs not rely on propagation in a fully perturbed environment.

²In practice this is probably only applicable to distant, consistent perturbation sources such as Solar gravitation.

9.5.5 Changes to the cost function

It was found that the current form of the cost function presented in subsection 6.1.1 was biased towards designs which favour near-collision events over baselines temporarily extending above a hundred kilometres. In practice the opposite would be favourable, as the latter has much less severe implications towards the health of the swarm. Based on this bias in the original equation an adjusted cost function is proposed for future research:

$$C = N_{NC} + N_W + N_V \quad (9.1)$$

For an orbit spanning between times t_0 and t_e , with N_R individual baselines, the number of infractions is measured as sums of every constraint breach during the entire orbit:

$$\begin{aligned} N_{NC} &= \sum_{t=0}^{t=t_e} \sum_{R=0}^{R=N_R} \alpha N (R < 500m) \\ N_W &= \sum_{t=0}^{t=t_e} \sum_{R=0}^{R=N_R} (R > 100km) \\ N_V &= \sum_{t=0}^{t=t_e} \sum_{R=0}^{R=N_R} (\dot{R} > 1m/s) \end{aligned}$$

The change is simply adjusting the weight of near-collision events N_c from 1 to a factor α the number of satellites in the swarm (N). As long as $\alpha \geq 1$ the avoidance of near-collision events will always be punished harder than the wandering of a single satellite, which prevents the algorithm from optimising towards near-collision events.

9.6 The relation between swarm size and orbit design

The results in Figure 7.8 showed that a 35-element swarm could feasibly be designed to work with the global swarm orbit which was found for the 25-satellite design. A curious property is that the resulting design broke with the established trend of a vertical column in the barycentric z_B direction, instead the column was skewed towards the initial velocity vector. This result shows that it is possible to upscale existing swarm designs to facilitate more satellites, but it requires some deviation from the vertical column design trend. The mechanism which directs this deviation is not understood yet, and it is highly recommended that this is studied. Understanding how scaling up the swarm affects its idealized design in non-ideal orbits may provide means to estimate the feasible swarm sizes of a given orbit, which will allow for alternative means of swarm design where the global swarm orbit is chosen beforehand.

9.7 The applicability of the 5th Lagrangian point

The work in this thesis has entirely been focused around the fourth triangular Lagrangian point, which has shown itself to be a very promising candidate. An open question is how well the concepts and designs presented in this thesis translate to the fifth Lagrangian point. In theory these points are very similar, but a major difference is the direction of the Lunar gravitational perturbation with respects to the orbital velocity of these points. It is unknown to the author how much this affects swarm orbit design, which makes it an interesting topic for future research.

9.8 Orbit insertion and de-orbiting

The orbits which were found for the swarm designs are all far from what would normally be considered convenient for insertion. The initial velocities all counteract what would be normal for satellites placed at that point in the Earth-Moon system. The topic of how to insert the swarm into these orbits has not been touched upon during this thesis, as it was deemed too big a problem to tackle in addition to the presented work. It is highly recommended for future work however, as the concept might be entirely impossible if satellites cannot be placed in these orbits.

Likewise, figuring out how to de-orbit these satellite swarms is an equally interesting problem. Due to the high energy that these satellites have on initial placement they can revert to very chaotic orbits across the Earth-Moon system after falling out of a L4-centric orbit. De-orbiting these swarms ought to be a well-planned part of the process, and research into this topic is recommended.

”Long days, and pleasant nights”

-Steven King, *The Gunslinger*

Bibliography

- [1] Chandra archive collection, Aug 2020. Accessed 9. Aug. 2020.
- [2] NRAO Very Large Array. <http://www.vla.nrao.edu>, Aug 2020. Accessed 13. Aug. 2020.
- [3] Star Myths - Theoi Greek Mythology. <https://www.theoi.com/greek-mythology/star-myths.html>, Aug 2020. accessed 10. Aug. 2020.
- [4] C. Acton. Ancillary Data Services of NASA's Navigation and Ancillary Information Facility. *Planetary and Space Science*, 44(1):65–70, 1996.
- [5] C. Acton, N. Bachman, B. Semenov, and A. Wright. A look toward the future in the handling of space science mission geometry. *Planetary and Space Science*, 2017.
- [6] Alexander, J K and Kaiser, M L and Novaco, JC and Grena, FR and Weber, RR. Scientific instrumentation of the Radio-Astronomy-Explorer-2 satellite. 1974.
- [7] ASTRON. About LOFAR. <http://www.lofar.org/about-lofar/about-lofar>, 2019. Accessed: 19/12.
- [8] M. J. Bentum. The search for Exoplanets using Ultra-long wavelength radio astronomy. In *2017 IEEE Aerospace Conference*, pages 1–7. IEEE.
- [9] M. J. Bentum, A. J. Boonstra, W. Horlings, and P. v. Vugt. The radio environment for a space-based low-frequency radio astronomy instrument. In *2019 IEEE Aerospace Conference*, pages 1–7.
- [10] M. J. Bentum, C. J. M. Verhoeven, A. J. Boonstra, A. J. Van Der Veen, and E. K. A. Gill. A novel astronomical application for formation flying small satellites. In *60th International Astronautical Congress 2009, IAC 2009*, volume 2, pages 1254–1261.
- [11] M. J. Bentum, M. K. Verma, R. T. Rajan, A. J. Boonstra, C. J. M. Verhoeven, E. K. A. Gill, A. J. van der Veen, H. Falcke, M. K. Wolt, B. Monna, S. Engelen, J. Rotteveel, and L. I. Gurvits. A roadmap towards a space-based radio telescope for ultra-low frequency radio astronomy. *Advances in Space Research*, 2019.
- [12] F. Biscani and D. Izzo. esa/pagmo2: Pagmo 2.15.0, Apr. 2020.
- [13] R. Blott, W. Baan, A. Boonstra, J. Bergman, D. Robinson, D. Liddle, N. Navarathinam, S. Eves, C. Bridges, and S. Gao. Space-based ultra-long wavelength radio observatory (low cost)-SURO-LC. In *European Planetary Science Congress 2013*, pages 1–2. EPSC.
- [14] A. J. Boonstra, M. Garrett, G. Kruithof, M. Wise, A. Van Ardenne, J. Yan, J. Wu, J. Zheng, E. K. A. Gill, J. Guo, M. Bentum, J. N. Girard, X. Hong, T. An, H. Falcke, M. Klein-Wolt, S. Wu, W. Chen, L. Koopmans, H. Rothkaehl, X. Chen, M. Huang, L. Chen, L. Gurvits, P. Zarka, B. Cecconi, and H. De Haan. Discovering the sky at the Longest Wavelengths (DSL). In *IEEE Aerospace Conference Proceedings*, volume 2016-June.
- [15] A. J. Boonstra, N. Saks, M. Bentum, K. Van't Klooster, and H. Falcke. DARIS, a low-frequency distributed aperture array for radio astronomy in space. In *61st International Astronautical Congress 2010, IAC 2010*, volume 10, pages 8036–8042.
- [16] M. Born and E. Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.
- [17] J. Braatz. Lecture on Imaging. https://science.nrao.edu/facilities/alma/naasc-workshops/alma_dr/Braatz_Imaging2.pdf, 2020. Accessed: 12/05/2020.
- [18] B. Cecconi, M. Dekkali, C. Briand, B. Segret, J. N. Girard, A. Laurens, A. Lamy, D. Valat, M. Delpech, M. Bruno, P. Glard, M. Bucher, Q. Nenon, J. M. Griemeier, A. J. Boonstra, and M. Bentum. NOIRE study report: Towards a low frequency radio interferometer in space. In *IEEE Aerospace Conference Proceedings*, volume 2018-March, pages 1–19.

[19] R. Cohen, G. Delgado, E. Hardy, T. Hasegawa, and L.-. Nyman. *Radio-Quiet Zones*, pages 225–259. Springer, 2003.

[20] J. J. Condon and S. M. Ransom. *Essential Radio Astronomy*, volume 2. Princeton University Press, 2016.

[21] E. Dekens. Orbit Analysis of a Low Frequency Array for Radio Astronomy. 2012.

[22] D. Dirkx, E. Mooij, and B. Root. Propagation and estimation of the dynamical behaviour of gravitationally interacting rigid bodies. *Astrophysics and Space Science*, 364(2):37, 2019.

[23] S. Engelen. Proefschrift S.Engelen Swarm Satellites: Design, Characteristics and Applications. 2016.

[24] J. Evans. *The history and practice of ancient astronomy*. Oxford University Press, 1998.

[25] I. Flyamer, Colin, Z. Xue, A. Li, V. Vazquez, N. Morshed, C. V. Neste, "scaine1", and "mski iksm". Phlya/adjustText. Nov 2018.

[26] J. N. Fortin. *Atlas Céleste de Flamstéed*. Paris: Lamarche, 1795.

[27] S. French, F. W., S. Rodman, A. K., and G. R. Huguenin. A synthetic aperture approach to space-based radio telescopes. *Journal of Spacecraft and Rockets*, 4(12):1649–1656, 1967.

[28] D. E. Gary. Radio astronomy lecture number 6. <https://web.njit.edu/~gary/728/Lecture6.html>, Mar 2019. Accessed 15. Aug. 2020.

[29] I. Gonzlez. Measurement of areas on a sphere using Fibonacci and latitudelongitude lattices. *Mathematical Geosciences*, 42(1):49, 2010.

[30] C. L. Grabbe. Auroral kilometric radiation: A theoretical review. *Reviews of Geophysics*, 19(4):627–633, 1981.

[31] IsGeschiedenis. Ra in Egyptian mythology. <https://isgeschiedenis.nl/nieuws/ra-in-de-egyptische-mythologie>, 2020. Accessed: 23/01/2020.

[32] K. G. Jansky. Electrical disturbances apparently of extraterrestrial origin. *Proceedings of the Institute of Radio Engineers*, 21(10):1387–1398, 1933.

[33] S. Jester and H. Falcke. Science with a lunar low-frequency array: From the dark ages of the universe to nearby exoplanets. *New Astronomy Reviews*, 53(1-2):1–26, 2009.

[34] D. L. Jones, R. J. Allen, J. P. Basart, T. Bastian, W. H. Blume, J. L. Bougeret, B. K. Dennison, M. D. Desch, K. S. Dwarakanath, W. C. Erickson, W. Farrell, D. G. Finley, N. Gopalswamy, R. E. Howard, M. L. Kaiser, N. E. Kassim, T. B. H. Kuiper, R. J. MacDowall, M. J. Mahoney, R. A. Perley, R. A. Preston, M. J. Reiner, P. Rodriguez, R. G. Stone, S. C. Unwin, K. W. Weiler, G. Woan, and R. Woo. The ALFA medium explorer mission. *Results and the Future of Space*, 26(4):743–746, 2000.

[35] D. L. Jones, T. J. W. Lazio, and J. O. Burns. Dark Ages Radio Explorer mission: Probing the cosmic dawn. In *IEEE Aerospace Conference Proceedings*, volume 2015-June.

[36] F. Lemoine, S. Kenyon, J. Factor, R. Trimmer, N. Pavlis, D. Chinn, C. Cox, S. Klosko, S. Luthcke, and M. Torrence. The development of the joint NASA GSFC and the National Imagery and Mapping Agency (NIMA) geopotential model EGM96. 1998.

[37] C. Maccone. The Quiet Cone above the Farside of the Moon. *Acta Astronautica*, 53(1):65–70, 2003.

[38] C. Maccone. Protected antipode circle on the Farside of the Moon. *Acta Astronautica*, 63(1-4):110–118, 2008.

[39] J. A. Marshall and S. B. Luthcke. Modeling radiation forces acting on TOPEX/Poseidon for precision orbit determination. *Journal of Spacecraft and Rockets*, 31(1):99–105, 1994.

[40] G. McCaine. Halo orbit design and optimization. Report, NAVAL POSTGRADUATE SCHOOL MONTEREY CA, 2004.

[41] S. Mok, J. Guo, E. Gill, and R. Rajan. Lunar orbit design of a satellite swarm for radio astronomy. *IEEE Aerospace Conference Proceedings*, 2020.

[42] NASA. RAE-B. <https://nssdc.gsfc.nasa.gov/nmc/spacecraft/display.action?id=1973-039A>, 2019. Accessed on 29/12.

[43] NASA Planetary Data System. PDS: Data Set Information. <https://pds.nasa.gov/ds-view/pds-viewProfile.jsp?dsid=LP-L-RSS-5-GLGM3/GRAVITY-V1.0>, Jul 2020. Accessed 05-09-2020.

[44] T. Navigation and A. I. Facility. Index of ephemeris kernel releases. https://naif.jpl.nasa.gov/pub/naif/generic_kernels/spk/satellites, Jul 2020. Accessed: 20/07/2020.

[45] J. C. Novaco and L. W. Brown. Nonthermal galactic emission below 10 mhz. 1977.

[46] D. Oberoi and J.-L. Pinon. A new design for a very low frequency spaceborne radio interferometer. *Radio science*, 40(04):1–19, 2005.

[47] E. S. Observatory. Opacity of the atmosphere. https://www.eso.org/public/images/atm_opacity/, 2019. Accessed: 29/12.

[48] J. R. Pritchard and A. Loeb. 21 cm cosmology in the 21st century. *Reports on Progress in Physics*, 75(8):086901, 2012.

[49] K. Quillien, S. Engelen, E. Gill, D. Smith, M. Arts, and A.-J. Boonstra. Astronomical antenna for a space based low frequency radio telescope. 2013.

[50] R. T. Rajan, A. J. Boonstra, M. Bentum, M. Klein-Wolt, F. Belien, M. Arts, N. Saks, and A. J. van der Veen. Space-based aperture array for ultra-long wavelength radio astronomy. *Experimental Astronomy*, 41(1-2):271–306, 2016.

[51] M. J. Rioja. Radio Astronomy - Interferometer Mapping. International Centre for Radio Astronomy Research.

[52] J. H. Rogers. Origins of the ancient constellations: I. The Mesopotamian traditions. *Journal of the British Astronomical Association*, 108:9–28, 1998.

[53] D. Rose. International space station familiarization. *User Manual for International Space Station ISS FAM C*, 21109, 1998.

[54] G. Rougoor and J. Oort. Distribution and motion of interstellar hydrogen in the galactic system with particular reference to the region within 3 kiloparsecs of the center. *Proceedings of the National Academy of Sciences of the United States of America*, 46(1):1, 1960.

[55] A. Rubinsztein. Dynamics of the 3 body problem. <https://gereshes.com/2018/11/12/dynamics-of-the-3-body-problem/>, 2018. Accessed: 05/05.

[56] N. Saks, A. Boonstra, R. T. Rajan, M. Bentum, F. Belin, and K. van t Klooster. DARIS, A FLEET OF PASSIVE FORMATION FLYING SMALL SATELLITES FOR LOW FREQUENCY RADIO ASTRONOMY. 2010.

[57] J. E. S. B. Stavrinidis, R. J. Blott, A. B. Forbes, D. A. Humphreys, D. W. Robinson, and Constantinos. FIRST Explorer - An innovative low-cost passive formation-flying system. 2009.

[58] J. van 't Hoff. A literature study regarding the orbital design of the OLFAR interferometry swarm. 2020.

[59] H. Vermeiden. Literature Study: Optimal Translunar Lagrange Point Orbits for OLFAR. 2013.

[60] K. F. Wakker. Fundamentals of Astrodynamics. 2015.

[61] R. Weber, J. Alexander, and R. Stone. The radio astronomy explorer satellite, a low-frequency observatory. *Radio Science*, 6(12):1085–1097, 1971.

[62] J. R. Wertz. *Orbit & Constellation Design & Management: Spacecraft Orbit and Attitude Systems*. Microcosm Press, 2001.

[63] A. Wood. Construction complete on the world's largest radio telescope in China. <https://newatlas.com/china-fast-telescope-complete/44163>, July 2016. Accessed 13. Aug. 2020.

[64] S. Wu, W. Chen, Y. Zhang, W. Baan, and T. An. SULFRO: A Swarm of Nano-/Micro-Satellite at SE L2 for Space Ultra-Low Frequency Radio Observatory. 2014.

[65] C. F. J. N. Yoder. How tidal heating in Io drives the Galilean orbital resonance locks. 279(5716):767–770, 1979.

Appendices

Appendix A : Verification of developed functions

This appendix will treat the verification methods which were used to validate the proper functioning of crucial functions developed for parts of this thesis. This appendix will be limited to only the functions which alter data, and which are used regularly.

A.1 Verification of the Point-Spread Function computation

A key parameter of the point spread function is that it is a fourier transform of the sample function which is projected into a two-dimensional plane. As such, as long as the underlying functions such as the plane projection are verified, validating the PSF computation is relatively straightforward by reconstructing the sample function using an inverse Fourier transform. If the inverse Fourier transform can be used to rebuild the original sample function it can be assured that the functions work adequately. Three different scenarios will be used to ensure proper functioning.

A.1.1 Comparison with the theoretical result

The two-dimensional discrete inverse Fourier transform in a periodic two-dimensional system spanning N by M sample points is described by:

$$f[x, y] = \frac{1}{\sqrt{NM}} \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} F(m, n) e^{j2\pi(\frac{nx}{N} + \frac{my}{M})} \quad (\text{A.1})$$

In which f is the sampled function within the two-dimensional plane. Assuming that the sample function is a single baseline with magnitude R along the local x-axis the inverse Fourier transform of the system would simplify to:

$$f[x, y] = \frac{1}{\sqrt{NM}} \left(f(R, 0) e^{j2\pi(\frac{xR}{N})} \right) \quad (\text{A.2})$$

Since the sample function is represented as a binary matrix with values of either one or zero this yields:

$$f[x, y] = \frac{1}{\sqrt{NM}} \left(e^{j2\pi(\frac{xR}{N})} \right) \quad (\text{A.3})$$

$$f[x, y] = \frac{1}{\sqrt{NM}} \left(\sin(2\pi(\frac{Rn}{N})) + i \cos(2\pi(\frac{xR}{N})) \right) \quad (\text{A.4})$$

The resulting two-dimensional inverse transform will thus be a sinusoid in the x-direction with a frequency of baseline magnitude R divided by the number of samples in x-direction N . This property will be used as the first verification test, taking a Fourier transform of an arbitrary slice of the PSF in the x-direction should in theory yield a peak at this frequency. If the largest frequency peak does not coincide with the R/N derived frequency, the verification fails. Figure A.1 visualises this validation step with results from the PSF generation computation.

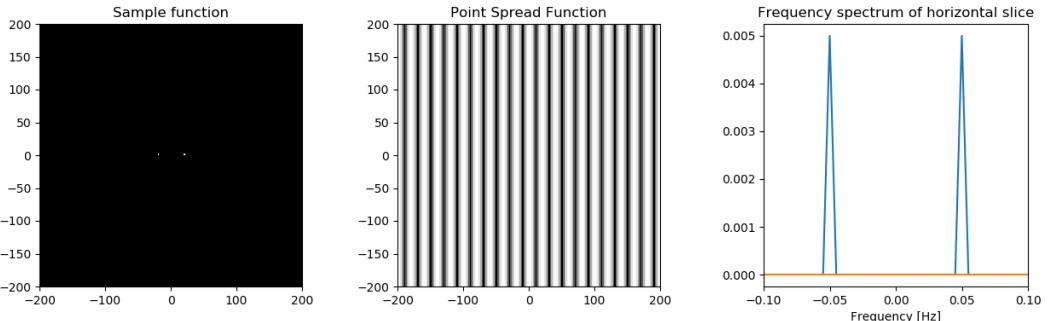


Figure A.1: A one-dimensional sample function and its resulting sinusoidal PSF with frequency R/N . The frequency spectrum is made from a one-dimensional horizontal slice of the PSF. $R=1.0$, $N=200$

The proper placements of the frequency spikes on ± 0.05 Hz confirm that the system is working properly, validating this first test case.

A.1.2 Simple sample function reconstruction

Using the same test scenario as the theoretical test, the second step is attempting to rebuild the sample function by taking the Fourier transform of the PSF. If the result is not identical to the x-axis baseline input the verification fails. The reverse Fourier transform of the PSF shown in Figure A.1 yields two points along the x axis, identical to the sample function provided. This validates that the PSF algorithm can work in reverse.

A.1.3 Complex sample function reconstruction

The final element which needs testing is reconstructing a more complex sample function that spans two dimensions. To do this an array of sample points is randomly generated in three dimensions, and projected onto an arbitrary plane described by a random view direction vector. The projected baselines are evaluated for their resulting PSF, which is then exposed to a Fourier transform. If the resulting set of points is not identical to the random input sample function, the validation fails. If all three scenarios are passed, the function is shown to be compatible with the established theoretical background, and suitable for reverse-engineering of the results.

This step is also passed without problems, verifying that the algorithm performs according to its design even in arbitrary three-dimensional cases and projections.

A.2 Verification of frame transformations

The work throughout this thesis regularly relied on frame transformations to convert J2000-centric orbits into barycentric views. This section will discuss the verification and validation of the functions developed for these conversions.

A.2.1 Relative frame transformation

A frame transformation to a frame relative to a moving object is easily tested by defining a very well known 3-dimensional test case. The "global" reference frame with origin O . In this plane a point P is defined. The coordinates of target point P are converted to a frame relative to object F which moves in time relative to the origin, yielding position vectors F_1 and F_2 .

$$O = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad P = \begin{bmatrix} 3 \\ 3 \\ 3 \end{bmatrix} \quad F_1 = \begin{bmatrix} 0 \\ 2 \\ 2 \end{bmatrix} \quad F_2 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \quad P_{F1} = \begin{bmatrix} 3 \\ 1 \\ 1 \end{bmatrix} \quad P_{F2} = \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix} \quad (\text{A.5})$$

The coordinates of P relative to point F are easy to compute by hand, and serve to validate the proper working of the relative frame conversion.

A.2.2 Barycentric frame transformation

A similarly simple test case will be used for the barycentric frame conversion, two masses M_1, M_2 are moving through three-dimensional Cartesian space with a third point of interest P located nearby. The masses of M_1 and M_2 are 9kg and 1 kg respectively, meaning that the barycentre B should be located at a tenth of the vector between both masses. Figure A.2 shows the geometry created by this problem description.

$$M_1 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad M_2 = \begin{bmatrix} 2 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad B = \begin{bmatrix} 0.2 \\ 0 \\ 0 \end{bmatrix} \quad P = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

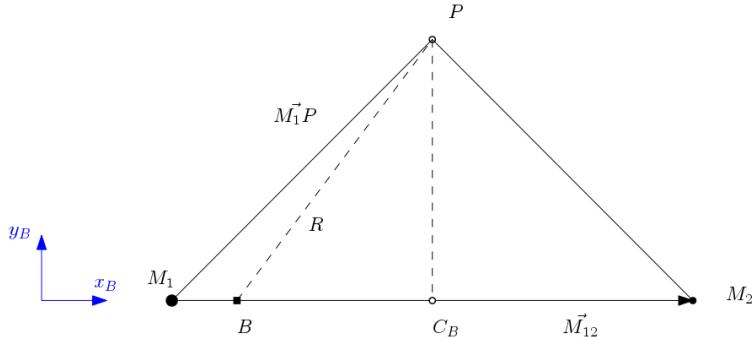


Figure A.2: Definition of the barycentric frame

To determine the coordinates of point P in the barycentric frame (vector P_B) a construction point C_B is used, which is placed alongside the projection of vector \vec{M}_1P on \vec{M}_{12} . The coordinates of P in the barycentric system can then be expressed as the magnitude of vector C_B corrected for the location of the barycentric frame. It is known that this construction point is placed at an orthogonal corner of the M_1PC triangle due to this projection, allowing for the resolution of the barycentric y coordinate of P through simple geometry. This yields coordinates in the barycentric frame for both the construction point and point P :

$$C_B = \begin{bmatrix} \vec{P} * \frac{\vec{M}_{12}}{|\vec{M}_{12}|} \\ 0 \\ 0 \end{bmatrix} \quad P_B = \begin{bmatrix} \vec{P} * \frac{\vec{M}_{12}}{|\vec{M}_{12}|} - 0.1|\vec{M}_{12}| \\ \sqrt{|\vec{M}_1P|^2 - |C_B|^2} \\ P_z \end{bmatrix} = \begin{bmatrix} 0.8 \\ 1 \\ 1 \end{bmatrix}$$

By carefully selecting the coordinates of the reference problem, the solution for the position vector P_B can be computed by hand. The conversion algorithm on the other hand needs to reconstruct this simple test case by computing the position of the barycentre, determining the direction of the barycentric axes from the momentum of the rotating system, and converting the third body position through vector projection along these axes. This process can successfully reproduce the answer through this approach, validating that it works as intended. This is also the case for a more complex system with displaced initial coordinates and relative velocities.

A.3 Verification of baseline extraction

To verify the baseline extraction algorithm two points are distributed across a three-dimensional Cartesian space with initial velocities which are used to compute the location at the next timeframe.

$$A_0 = \begin{bmatrix} 1 \\ 6 \\ 0 \\ -6 \\ 5 \\ -4 \end{bmatrix} \quad B_0 = \begin{bmatrix} 3 \\ 1 \\ 0 \\ 0 \\ -6 \\ 3 \end{bmatrix}$$

Computing the baselines between these two dummy satellites should yield two opposing vectors with the same magnitude as the direct vector between these points. If the function yields two opposing vectors with the same magnitude and relative velocities, the function is validated.

A.4 Verification of two-dimensional projection

To evaluate the point-spread function of a constellation the satellite baselines need to be projected onto a two-dimensional plane in any arbitrary view direction. To verify this projection function a simple test scenario uses three points that are spread randomly in a Cartesian three-dimensional space:

$$A = \begin{bmatrix} a_x \\ a_y \\ a_z \end{bmatrix} \quad B = \begin{bmatrix} b_x \\ b_y \\ b_z \end{bmatrix} \quad C = \begin{bmatrix} c_x \\ c_y \\ c_z \end{bmatrix}$$

By projecting these points onto the xy , xz , and yz plane the resulting projected coordinates are easily predictable. For example, consider the 2-dimensional projection of these points on the xy plane:

$$A = \begin{bmatrix} a_x \\ a_y \end{bmatrix} \quad B = \begin{bmatrix} b_x \\ b_y \end{bmatrix} \quad C = \begin{bmatrix} c_x \\ c_y \end{bmatrix}$$

This test case allows for very straight-forward verification of proper functioning of the projection function, which first needs to establish a 2-dimensional plane based off the given view vector ($[0,0,1]$). After determining the vectors which act as basis for the two-dimensional frame the coordinates off all these points are converted through vector projection.

This algorithm passed all the projection tests for the cardinal axes of the system, verifying that this approach works.

Appendix B : J2000 - Earth-Centered Inertial frame

The J2000 ECI frame defines the x-axis as being aligned with the vernal Equinox at 12:00 terrestrial time on the first of January 2000. The local z-axis is aligned with the celestial pole, and the y-axis completes a right-hand coordinate system. The frame is inertial, meaning that it does not rotate relative to the celestial sphere over time, yet its centre remains fixed to the center of mass of the Earth [62]. Figure B.1 visualises the J2000 reference frame coordinate system:

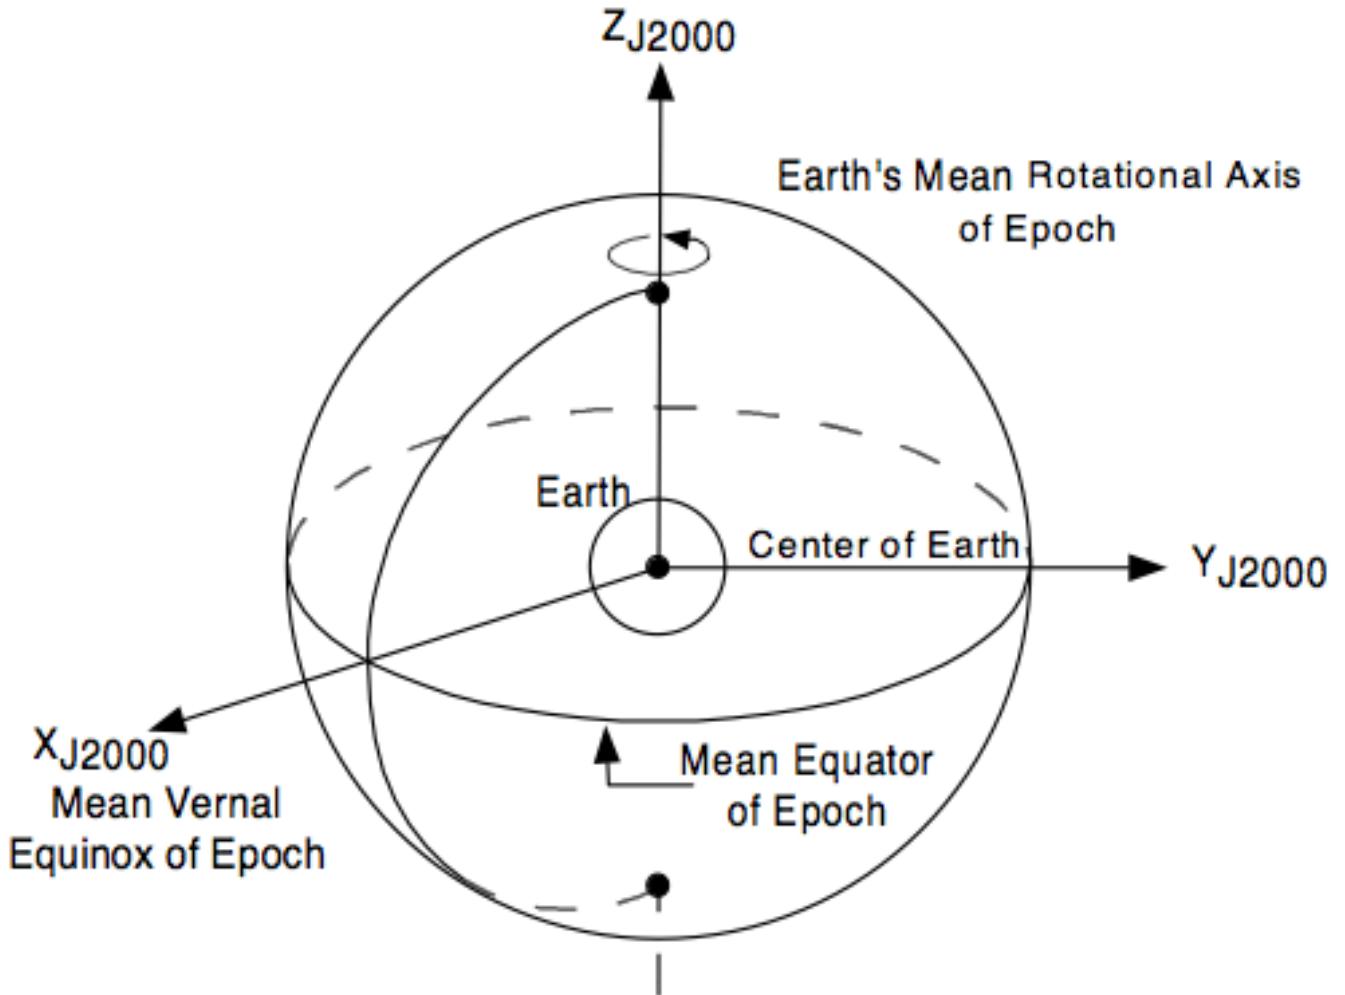


Figure B.1: Definition of the J2000 reference frame, taken from [53].

Appendix C : Swarm design tables

This appendix includes the swarm design tables of the optimal solutions found for 20, 25 and 35 satellite configurations presented in Chapter 7. The coordinates within these design tables are presented in the J2000 frame, on the 1st of January 2030. The core positions are expressed relative to the coordinates of L4 at this time, and the core velocities relative to the velocity vector of L4. The L4 coordinates are computed using the distance between the Earth and Moon at that time, instead of using the average.

$$\vec{L}_4 = 10^8 * \begin{bmatrix} 1.69665159 \\ -3.02772120 \\ -1.11443546 \end{bmatrix} [m] \quad \vec{L}_4 = \begin{bmatrix} 943.16877462 \\ 427.13901237 \\ 275.44974209 \end{bmatrix} [m/s]$$

Finally, individual satellite coordinates within the swarm are expressed relative to the core position. Should the reader have interest in re-using these results, feel free to contact¹ the author for pre-processed data files with individual satellite state vectors in the J2000 frame.

C.1 15 and 20-satellite swarm design

Table C.1: Design table for a 15-element radio interferometer swarm

	x	y	z			x	y	z
Core position [km]	40.2719	30.2410	12.8253					
Core velocity [m/s]	-4.8000	2.7000	18.1000					
	x[km]	y[km]	z[km]			x[km]	y[km]	z[km]
S_1	-1.3155	1.1334	9.2450		S_9	-6.8109	-4.2125	22.5365
S_2	-3.9823	4.3531	1.4430		S_{10}	-10.4271	4.4846	8.7071
S_3	-2.4768	15.797	-22.688		S_{11}	-4.6041	1.2297	18.4724
S_4	-7.6780	8.3866	1.0497		S_{12}	-9.1663	6.6735	2.8295
S_5	-8.7145	9.6218	-2.4076		S_{13}	-8.7651	-3.51030	28.9842
S_6	-0.4394	12.2561	-16.2912		S_{14}	-0.1921	10.5986	-11.2593
S_7	-8.6424	-1.9455	17.6387		S_{15}	-5.0668	-3.7208	22.6989
S_8	0.2090	19.4074	-26.1245					

Table C.2: Design table for a 20-element radio interferometer swarm

	x	y	z			x	y	z
Core position [km]	38.067	17.9474	18.4417					
Core velocity [m/s]	-3.5386	10.4148	0.4299					
	x[km]	y[km]	z[km]			x[km]	y[km]	z[km]
S_1	2.4538	-15.3296	33.9844		S_{11}	1.7696	-2.8532	9.7615
S_2	2.3462	0.0827	0.3297		S_{12}	4.8073	-0.5439	5.0184
S_3	0.0223	-11.6574	24.5556		S_{13}	-2.5862	-0.5895	3.7357
S_4	4.5183	3.7513	-8.2912		S_{14}	6.2059	2.1995	-3.2218
S_5	-0.2156	-7.2397	23.9963		S_{15}	4.6488	-2.7911	9.3928
S_6	1.4764	0.2019	0.829		S_{16}	5.6609	10.3634	-17.0213
S_7	4.9582	1.6645	-5.2941		S_{17}	0.0491	-2.6104	5.3792
S_8	5.9233	-1.6568	5.6732		S_{18}	-1.5997	-6.9147	19.5749
S_9	2.1472	-5.7992	17.7005		S_{19}	3.2284	-1.1768	9.2965
S_{10}	2.4452	-12.1129	24.1904		S_{20}	3.401	8.3156	-22.141

¹Contact details can be found in the Preface.

C.2 25-satellite swarm design

Table C.3: Design table for a 25-element radio interferometer swarm

	x	y	z		x	y	z
Core position [km]	4451.5958	45.6624	-3.7657				
Core velocity [m/s]	2.1817	-13.2982	-12.3799				
	x[km]	y[km]	z[km]		x[km]	y[km]	z[km]
S_1	7.9148	-5.0381	7.7531	S_{14}	7.3776	-2.8681	-2.2423
S_2	4.2439	-1.1072	-10.9251	S_{15}	3.1692	-11.3139	20.207
S_3	5.4649	-2.8878	-2.2065	S_{16}	4.0172	-7.3477	9.3651
S_4	4.9978	-8.3128	6.6866	S_{17}	6.7412	-6.2152	3.2533
S_5	2.9839	-2.5992	-3.6861	S_{18}	2.1556	-9.516	13.5329
S_6	5.0419	-6.7938	10.8208	S_{19}	9.1507	1.2394	-10.3591
S_7	4.6937	-8.6412	17.0026	S_{20}	6.2712	-8.7395	15.4174
S_8	4.7746	-3.0253	3.4359	S_{21}	2.3212	-6.8594	12.3567
S_9	6.2728	-15.1159	26.3456	S_{22}	5.0915	-0.4912	-0.7077
S_{10}	3.1111	-12.9079	17.6447	S_{23}	8.9528	2.698	-20.6953
S_{11}	4.2893	-16.6108	32.2752	S_{24}	7.3194	-4.8344	-1.1631
S_{12}	4.9265	-7.2616	5.6501	S_{25}	5.2222	-8.4202	4.6766
S_{13}	5.8812	-19.5097	35.8188				

C.3 35-satellite swarm design

Table C.4: Design table for a 35-element radio interferometer swarm

	x	y	z				
Core position [km]	122.6701	204.0254	-16.3404				
Core velocity [m/s]	-6.9946	14.72019	5.1595				
	x[km]	y[km]	z[km]		x[km]	y[km]	z[km]
S_1	-0.7898	2.9492	-5.8446	S_{19}	-7.0249	-0.2409	5.637
S_2	-3.5704	2.1535	-7.8662	S_{20}	-5.4702	-4.1846	18.515
S_3	-4.0612	-2.9946	5.5964	S_{21}	-2.6336	-2.1855	9.5343
S_4	-8.3186	5.4306	-7.9659	S_{22}	-2.8118	-10.6665	27.0123
S_5	-4.88	10.2686	-23.2885	S_{23}	-3.3931	-2.1363	7.6022
S_6	-1.3888	-3.144	6.0971	S_{24}	-5.0859	-2.1465	9.3782
S_7	-6.9756	-4.711	14.1913	S_{25}	-2.5016	2.3106	0.1841
S_8	-2.6704	10.1106	-18.1997	S_{26}	-2.5538	4.9374	-11.9088
S_9	-7.0997	2.7977	-0.5809	S_{27}	-4.337	-8.5064	21.1079
S_{10}	0.1744	8.0451	-16.4527	S_{28}	-5.7651	1.6565	-4.395
S_{11}	2.5281	10.1751	-23.386	S_{29}	-4.1722	13.4849	-26.006
S_{12}	-8.0768	-11.5495	25.8034	S_{30}	-5.9107	9.0856	-17.5228
S_{13}	-2.4774	-3.8073	7.0142	S_{31}	-5.2753	5.3882	-10.2077
S_{14}	-4.8207	0.027	-2.9804	S_{32}	-4.1365	0.9024	0.6012
S_{15}	-5.1737	4.405	-1.3716	S_{33}	-6.7878	-4.9636	12.4311
S_{16}	-2.4178	1.4308	-4.9605	S_{34}	-9.2494	-2.3588	11.2631
S_{17}	-1.714	1.7148	-4.2682	S_{35}	-4.7284	-4.2445	17.0872
S_{18}	-1.1126	10.2152	-19.3974				

C.4 35-satellite designs, second stage

Original orbit

Table C.5: Design table for a 35-element radio interferometer swarm

	x	y	z				
Core position [km]	122.6701	204.0254	-16.3404				
Core velocity [m/s]	-6.9946	14.72019	5.1595				
	x[km]	y[km]	z[km]		x[km]	y[km]	z[km]
S_1	-0.3093	4.3613	1.2793	S_{19}	-2.102	-2.0173	-4.8753
S_2	3.8774	-1.6758	9.9979	S_{20}	1.0226	-0.6269	21.009
S_3	3.9163	1.3848	17.2566	S_{21}	-2.2299	-6.2798	3.1081
S_4	2.6444	3.5376	-16.5599	S_{22}	-2.9461	3.846	-10.2628
S_5	3.0167	0.0895	13.0209	S_{23}	-0.985	-2.1335	-7.0037
S_6	-3.377	-3.4474	10.8887	S_{24}	-1.5607	3.4925	5.2403
S_7	-6.8261	-0.8935	5.7839	S_{25}	5.6103	1.2917	2.2488
S_8	-6.0103	-1.5433	-17.3925	S_{26}	3.9237	-5.7633	14.5931
S_9	5.5816	1.5828	12.4813	S_{27}	1.7031	3.7392	-14.8504
S_{10}	-2.9802	-0.4351	19.1837	S_{28}	-0.6427	-0.6015	2.4565
S_{11}	1.7007	1.7725	10.3154	S_{29}	0.3566	1.9996	0.1972
S_{12}	-3.5079	-4.8776	-8.6708	S_{30}	3.2328	3.4924	-9.8385
S_{13}	0.3366	2.2357	2.1377	S_{31}	3.7803	-4.0376	11.4726
S_{14}	0.9618	3.5797	-12.7224	S_{32}	6.8062	2.8818	12.2992
S_{15}	1.2208	-2.1339	6.7488	S_{33}	-1.1272	-0.4791	6.2356
S_{16}	1.4484	-1.0258	-9.5997	S_{34}	3.7476	-5.6413	5.4214
S_{17}	0.5427	-1.7293	-0.6683	S_{35}	2.5811	0.179	-1.9411
S_{18}	0.54	3.2602	15.954				

25-satellite orbit

Table C.6: Design table for a 35-element radio interferometer swarm

	x	y	z				
Core position [km]	4451.5958	45.6624	-3.7657				
Core velocity [m/s]	2.1817	-13.2982	-12.3799				
	x[km]	y[km]	z[km]		x[km]	y[km]	z[km]
S_1	4.1434	-0.375	5.2698	S_{19}	2.2734	-2.4488	-0.1021
S_2	-3.4691	-2.8281	5.0399	S_{20}	-2.3178	5.0559	2.8617
S_3	-0.9262	1.3667	-9.6136	S_{21}	4.5147	-3.3457	-26.4236
S_4	-1.1168	-1.3857	-7.9303	S_{22}	-2.358	4.6633	-5.0389
S_5	-0.4469	-1.4643	11.3906	S_{23}	1.5547	4.0657	-0.2056
S_6	2.6106	1.2731	-9.1268	S_{24}	5.9415	-8.3946	-1.5434
S_7	4.2056	-1.1619	8.8468	S_{25}	-4.6475	-2.389	-7.7323
S_8	3.3497	-1.9936	-7.6759	S_{26}	4.8512	2.5821	-2.0668
S_9	1.3664	-3.3486	26.6286	S_{27}	2.5429	-2.0245	1.5591
S_{10}	-6.3366	1.1797	1.0026	S_{28}	-3.1338	-5.8967	-9.026
S_{11}	-4.761	-4.959	1.5838	S_{29}	3.2008	-9.061	-7.4832
S_{12}	3.579	0.642	-3.3769	S_{30}	3.1361	-1.5322	-21.6733
S_{13}	5.6495	1.1398	18.1932	S_{31}	-1.4273	0.4837	-4.4254
S_{14}	-1.1992	6.4432	5.461	S_{32}	-1.5143	-3.9491	9.1499
S_{15}	2.3092	-2.0692	-6.7773	S_{33}	-0.2131	-1.6861	21.4903
S_{16}	-4.8796	-6.5643	5.1261	S_{34}	-5.5267	1.984	-13.1572
S_{17}	-2.2511	-0.0137	-12.3243	S_{35}	-10.5882	-0.3558	9.8715
S_{18}	1.3435	-2.4135	-2.9891				