POLITECNICO DI TORINO

Collegio di Ingegneria Meccanica, Aerospaziale,
dell'Autoveicolo e della Produzione

Corso di Laurea Magistrale In Ingegneria Aerospaziale

Tesi di Laurea Magistrale

Optimizing Satellite Constellation Orbits for Low-Latency Internet Connections

Relatore:
Prof. Lorenzo Casalino

Candidato:
Carlo De Micheli

Marzo 2021

Acknowledgements

First and foremost, I would like to thank my professors at the Polytechnic University of Turin for sharing their knowledge throughout the years.

I wish to especially thank the professors of the Aerospace department for having supported me in this effort.

I would also like to thank my family and friends for helping me throughout this journey. They have all been understanding and have always given me great inputs and ideas regarding my work.

- Thank You!

Abstract

This thesis analyzes different approaches used to provide internet connectivity via satellite constellations and uses this analysis to find the most effective coverage-to-cost solution through an optimal configuration of orbital planes. The analysis starts with a discussion on the options for the coverage of specific geographic regions and the extensibility of such solutions to global coverage. The following steps focus on global coverage and latency issues to provide solutions that allow a latency comparable to, or better than, traditional connections. The last section addresses cost implications and the estimation of the optimal number of satellites for each configuration to provide the necessary coverage to the end users. The final results report the launch cost associated with each solution and the optimal orbit choice is compared to the industry standards.

Abstract - Italian

Questa tesi analizza i diversi approcci utilizzati per fornire connettività Internet tramite costellazioni satellitari ed esamina la soluzione più efficace in termini di copertura e costo attraverso una configurazione ottimale dei piani orbitali. L'analisi inizia con una discussione delle opzioni che garantiscono la copertura di regioni geografiche specifiche e considera la possibilità di estendere tali soluzioni per fornire una copertura globale. I capitoli seguenti si concentrano sulla copertura e sui problemi di latenza per fornire soluzioni che consentano connessioni paragonabili a quelle tradizionali. La tesi si conclude con un capitolo che compara sia i costi delle diverse soluzioni, sia la stima del numero ottimale di satelliti per ciascuna configurazione per fornire connettività a banda larga agli utenti finali. I risultati comprendono i costi di lancio della costellazione per le diverse soluzioni e le orbite scelte vengono comparate alle soluzioni adottate dall'industria.

Table of Contents

List of Figures

List of Tables

1. Introduction

Internet connectivity is becoming increasingly-more essential to the world economy, as particularly evidenced by the current travel restrictions and increase in remote work caused by the COVID-19 pandemic The increasing dependence on internet connectivity and current conditions have highlighted once again the appeal of achieving a global coverage, particularly in emerging countries or in remote areas where fiber connectivity is not yet available. Providing digital connectivity via satellite has been in the works for decades. Back in 1962 [3] the first Telstar satellite was launched, and with its one megabit of bandwidth, it enabled the first communications between continents through Low Earth Orbit.

The developments of this technology have been relatively slow in the four decades following Telstar 1, but in the first decade of the 21st century, things were about to change. Undersea cables already provided high-speed connectivity between countries and continents, but the need for connectivity in every remote corner of the earth spawned new ventures and funding to study how to achieve this goal. Low bandwidth global communications were available via satellite through small constellations in geostationary orbits.

These connections enabled people in rural areas, or on moving objects such as ships and aircraft, to access the internet via expensive data plans, with extremely high latencies. While for some simple functions, such as sending emails, a high latency did not impact their performance, the average quality of audio and video calls was extremely degraded. The latency issue will be one of the main points of discussion throughout this paper.

The focus of this paper will be on creating a LEO constellation and optimizing the parameters to lower the number of necessary satellites to achieve global coverage and low latency connections.

The economic implications of the chosen orbits will be analyzed as well, with particular attention to variations in the launch costs. The analysis will compare the economic cost of polar orbits to the mixed setup of inclined and polar orbits. It will then estimate the range for the cost of launch, regardless of orbit type as well as the total cost once satellite unit costs are included in each cost estimate.

Finally, the analysis will discuss which orbit type appears more advantageous from the economic standpoint and how real-life applications of orbit choice are closest to the optimal solution found by this thesis.

2. Options for coverage

The first consideration to make when studying any orbits is to analyze which regions should be covered by the satellites. The coverage area of a satellite at a given position and time depends on multiple factors including the altitude of the satellite and the desired elevation angle of the ground station that transmits or receives information.

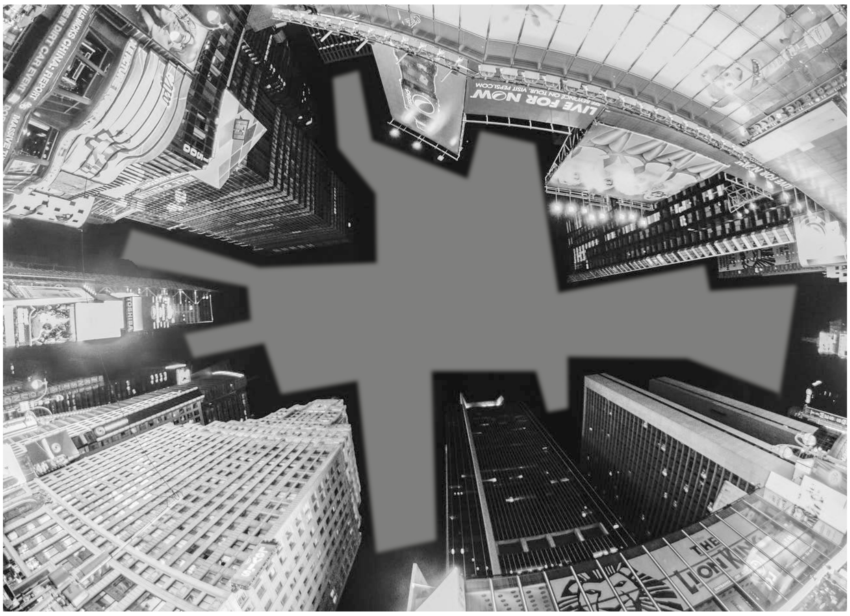

Figure 1 - Sky visibility, New York

The elevation angle is limited by natural and man-made obstacles. Satellite visibility at low, 0-20 degree, elevation angles may be possible in rural areas, but the portion of visible sky may be greatly reduced in urban environments. In Image 1.1 the sky represents only 18 percent of the picture in a highly-urbanized area, even though the photo was taken at 90 degrees of elevation with a lens that did not reach zero degrees of elevation at the edges. Numerous papers have been published about this topic using original images or pre-existing ones from software such as Google’s Street View to estimate the visibility of satellite constellations in different locations [1].

The implementation of a minimum elevation angles to grant satellite visibility has varied greatly, from Telesat using 40 degrees as minimum, OneWeb enabling users to always connect to a satellite with an elevation angle higher than or equal to 55 degrees, and SpaceX requesting a reduction from 40 to 25 degrees to the FCC [2]. The regulatory concerns are also important to take into account, as competitors in the space are very sensitive to the requests on behalf of other companies to access similar orbits and lower elevation angles. Potential interference concerns deferred the decision recently released on January 8th, 2021, and did not authorize SpaceX to reduce the elevation angle of its ground stations.

There is a trade-off between a more-optimal elevation angle and the cost associated with building the network to achieve it. The main reason is that the number of satellites required to achieve the same coverage increases approximately with the square of the elevation angle. Extensive research has analyzed the effect of parameters necessary to obtain a successful link at low elevation angles, such as the ground station's transmission power or the interference with Earth's magnetic field[4]. In LEO orbits, the characteristics of a satellite's footprint, the area from which the satellite is visible at a given altitude and elevation angle, can be approximated using trigonometrical formulas [5][6].

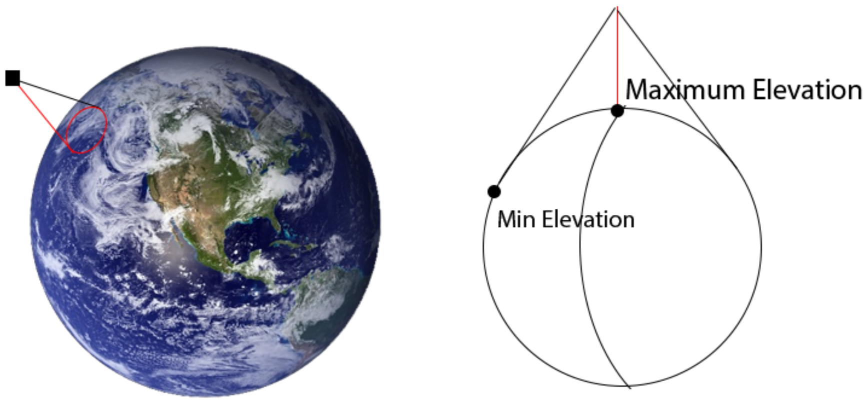

Figure 2 - Minimum and maximum elevation angles

For the purpose of this paper, calculating the footprint has been simplified to the area of intersection of a cone with a spherical object. It could be further simplified to the intersection of a cone with a flat surface, but given that the altitude of the satellite and the radius of the Earth are of the same order of magnitude, the calculation would be skewed toward significantly-smaller areas. We therefore obtain a formula based solely on the altitude and on the minimum elevation angle of the ground stations [7].

The area can be calculated as:

$Footprint = 2\pi r^2(1 - \cos\alpha)$ with α being the angle between the center of the Earth and the edge of the spherical cap, which marks the visibility area of a satellite at a given elevation angle.

The assumption of a reasonable error due to a perfectly-spherical Earth in the model was verified with the data found for the Iridium satellites footprint.

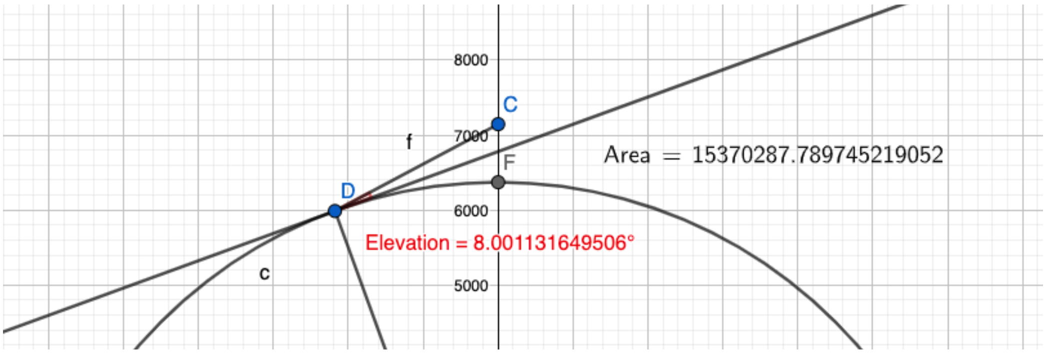

Figure 3 - Iridium footprint calculation

The Iridium constellation has a series of polar orbits at an altitude of 780 km with a minimum elevation angle of 8 degrees. The stated footprint is equal to 15.3 million square kilometers and the results with the simplifications introduced earlier were comparable.

The calculation was run for different altitudes and elevation angles, the results are presented in the table below. The unit is millions of square kilometers.

[10^6 km^2]	20 degrees	40 degrees	60 degrees
500 km	3.41	0.87	0.22
600 km	**4.51**	1.19	0.31
1500 km	16.1	**5.25**	1.45

Table 1 - Footprint (coverage area) based on elevation angle and altitude

It is clear from these results why SpaceX would have had to request a lower elevation angle of 20 degrees, from the previous 40 degrees, when lowering the orbits from an altitude of approximately 1500 km to just under 600 km. The values in bold in the table represent similar coverage areas, provided that, at a lower altitude, a lower transmission power would be necessary to maintain the link. The downside is clearly that lower elevation angles would provide less coverage in densely-populated areas with tall man-made structures. From a business perspective, the loss would be negligible as densely-populated areas are mostly served by high-speed internet connections.

Although the focus of this paper will be on LEO orbits, due to latency requirements, a brief analysis of alternative orbits has to be taken into consideration. Specific higher orbits have extraordinary characteristics which enable satellites to virtually hover over specific geographical regions. The prime example are geosynchronous orbits, which enable a satellite to oscillate at different latitudes while always covering a specific area in terms of longitude.

Figure 3 - IGSO Orbit Ground Track Example [8]

The properties of geosynchronous orbits more generally, and in the specific case of Figure 3 representing the ground track of an Inclined Geosynchronous Orbit (IGSO), are that the time required for a complete orbit around the Earth is equal to the time that the Earth requires to complete one revolution upon its axis, namely 23 hours 56 minutes 4 seconds. This allows the satellites in geosynchronous orbits to always be above the same location in the case of geostationary orbits or to appear at a fixed average longitude in the case of inclined geosynchronous orbits. Geosynchronous satellites have been used for decades for telecommunications, as they provide a wide footprint at an altitude of 35,786 km.

The maximum footprint of a GEO satellite, if the ground stations had a zero elevation, would cover approximately 42 percent of the Earth's surface. The approximation using a perfect sphere yields a maximum spherical cap area of 209 million square kilometers, which is 41 percent of the Earth's surface. Only three GEO satellites would be necessary to cover the entire surface of the Earth. The elevation could be increased up to 20 degrees, while still maintaining the coverage at the equator, as seen in the diagram below.

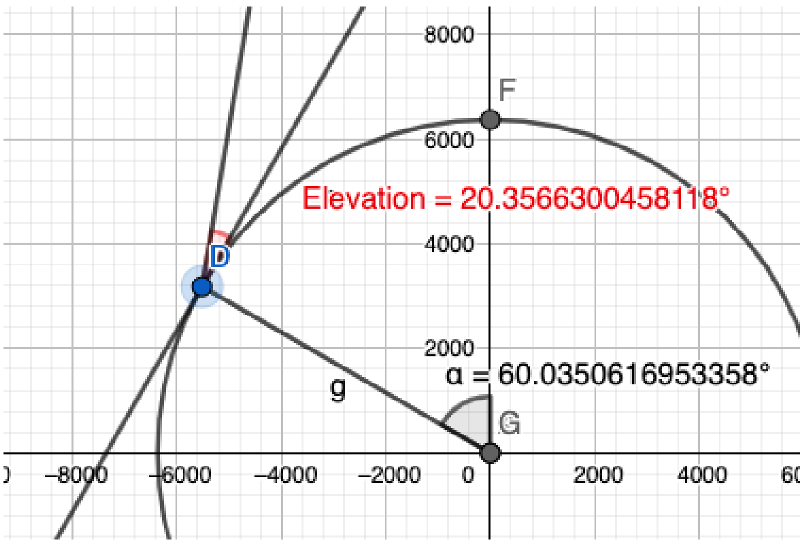

Figure 6 - GEO coverage at the equator with 20 degrees of elevation

Three GEO satellites in the above configuration would enable the coverage of the entire circumference of the Earth at the equator, but would have blind spots in northern or southern latitudes, if a lower degree of elevation was not feasible due to natural or man-made obstacles. GEO satellites are also usually region-specific and positioned to serve densely-populated areas rather than to aim for global coverage. Broadcasting satellites are usually in GEO orbits due to the large footprint and ease of use for ground users as the satellite dish can be pointed once and never moved even while being directed physically or electronically toward a new position.

This thesis will consider the possibility of utilizing GEO as backups, relays, or higher-latency redundancies for a LEO constellation. By contrast, using satellites in these orbits as primary nodes will not be a focus of the discussion because the mere altitude of 35,786 km would introduce a round-trip latency of 238.7 milliseconds at the speed of light which would be unacceptable in most modern-day real-time communications.

While a complete LEO constellation would be self-sufficient in the communication with all the satellites from a single ground station through mesh networking, the mesh may lack enough links to operate in the first phases of deployment. During the launch process of a constellation, which may take several years, GEO satellites may be used to provide stable links between groups of satellites in different orbital planes. In case of failures of the mesh network, they can also be used to relay information to otherwise-unreachable satellites. Another potential use case for GEO satellites in the study of a constellation is to cover polar areas despite the higher latency. These areas have often been neglected when building LEO constellations due to the small number of potential users.

The cost of launching a satellite in geostationary orbits is estimated to be six times more than in LEO. This thesis will take into account existing GEO satellites and the possibilities of relaying information through them, rather than analyzing the launch of new satellites in higher orbits.

3. Achieving global coverage

In describing how to achieve global coverage, a brief consideration of three common ways that have been used throughout the years in satellite constellations should be made:

- Hundreds of satellites in LEO.
- A smaller number, between 10 and 100, of satellites in Medium Earth Orbit (MEO).
- Fewer than 10 satellites in GEO.

MEO orbits have the advantage of allowing global coverage with a substantially-smaller number of satellites than in LEO, but do not possess the advantage of being stationary over specific geographical areas and suffer from higher latency than LEO orbits. They are popular for navigation-related applications such as GPS and GLONASS, but will not be considered in this thesis as they do not present sufficient advantages compared to LEO orbits.

The number of satellites required to achieve global coverage using GEO has been addressed in the previous chapter. LEO constellations present further challenges. The main issue is how many satellites are required to obtain and maintain global coverage. Companies have even changed the number of satellites for their constellation after the first launches.

From a geometrical perspective, in the tessellation of a plane using circles with minimum overlap, a hexagonal arrangement similar to a beehive would be used. The Iridium constellation pictured in Figure 7 portrays exactly this type of arrangement. Each satellite covers a specific area, with the edges of the coverage area overlapping with six adjacent satellites.

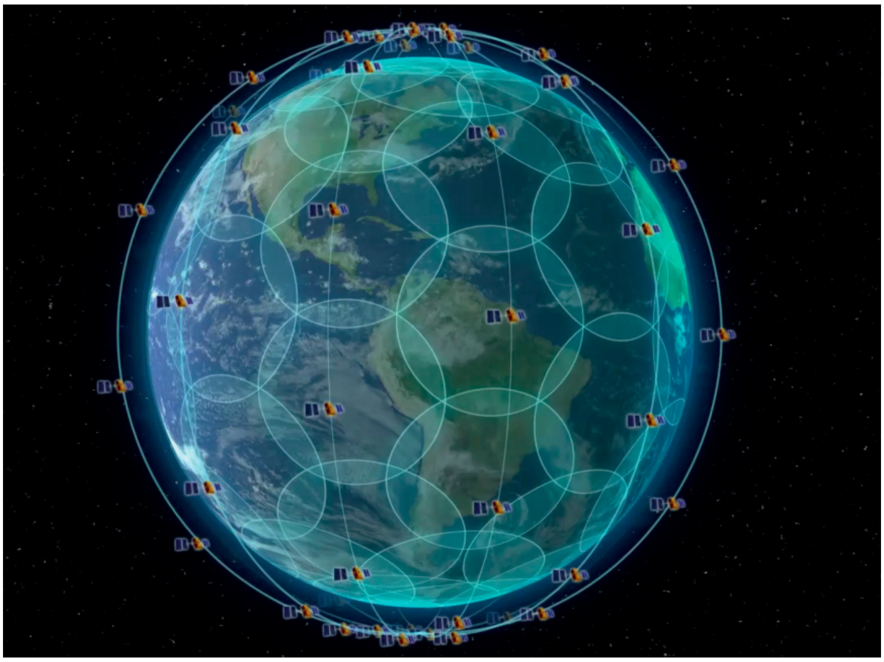

Figure 7 - The Iridium Constellation [9]

In a plane, the area which is covered twice by two overlapping circles is equal to the difference between the area of the inscribed hexagon and the area of the circle itself. These calculations yield that 17.3 percent of the area is covered more than once at any given time, with the caveat that they do not consider crowded spaces with greater overlapping in polar areas. In this framework, at a tentative altitude of 600 km, a global coverage would be obtained with numbers of satellites greater than:

Elevation angle	Coverage [10^6 km^2]	Unique Coverage [10^6 km^2]	Number of satellites required would be larger than
40 degrees	1.19	0.984	**519**
20 degrees	4.51	3.73	**137**

Table 2 - Number of satellites required at 600 km altitude

For comparison, the Iridium constellation pictured above uses 66 satellites in polar orbital planes at an altitude of 780 km, covering 100 percent of the Earth's surface.

Although a larger number of satellites increases the probability of having satellites directly overhead the users, improving coverage in areas with obstacles, the time between hand-offs in connecting from one visible satellite to the next decreases dramatically. If global coverage can be achieved with such a low number of satellites even in LEO, the reasons why companies such as OneWeb and SpaceX established networks with much larger numbers of satellites are clearly not related to coverage issues. The main difference between Iridium and newer constellations is the bandwidth available for each user. Iridium was conceived for end-user bandwidths in the order of kilobytes/second, while SpaceX considers modern requirements in the order of gigabytes/second [10].

This consideration leads to a key point: if coverage is not an issue, above a certain threshold of number of satellites, then the parameters to consider in the determination of the number of satellites to use will be related to the usage of the network. The distribution of the network usage across the globe varies greatly, hence there are different possibilities in terms of coverage:

- A) Building a constellation that has more satellites in regions where the usage is expected to be greater.
- B) Building a constellation that covers most of the Earth equally, but excludes polar areas.
- C) Building a constellation that covers 100 percent of the Earth at all times.

There are large differences between A and the other options, and trivial modifications between B and C. Starlink, for example, began with a plan resembling option B and recently switched to C to enable the connection of national-security-related clients operating in the polar areas. Adding very few polar orbital planes to numerous inclined planes switched Starlink from option B to C. Option A would be desirable from a usage optimization perspective. Satellites would have to follow orbits known as RGT (Repeat Ground Track), in order to visit the space above specific geographical locations at regular intervals.

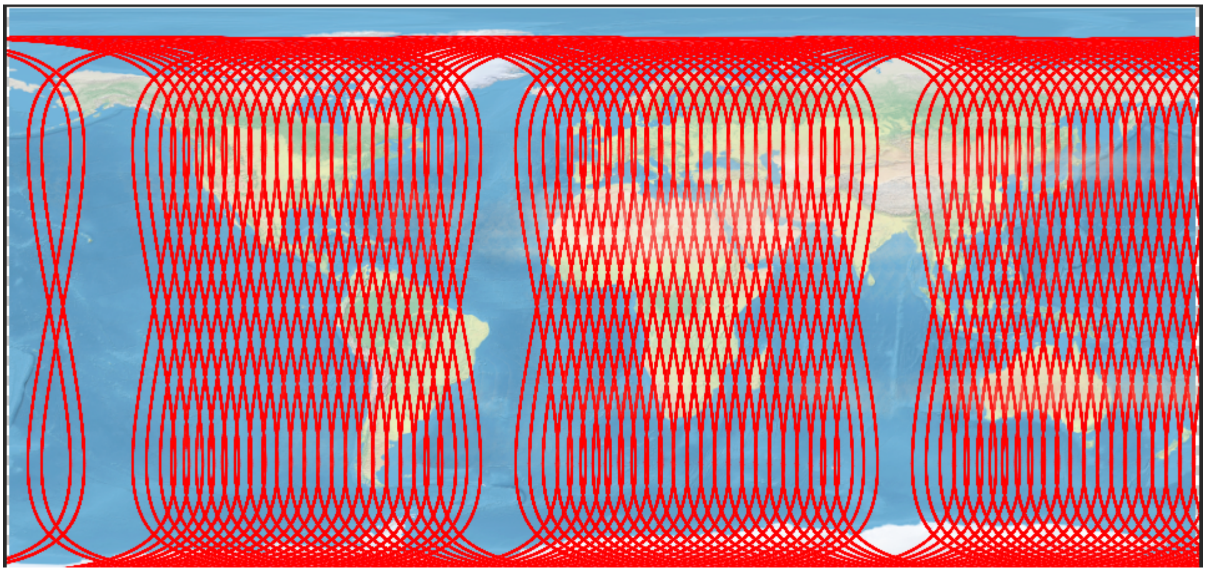

Figure 5 - Polar orbit RGT Simulation [11]

Simulating LEO, RGT polar orbits with multiple satellites, it is clear that the benefits from avoiding specific routes are minimal, as adding 10 percent more satellites would fill the gaps and enable global coverage, which would be definitely appealing from a business perspective for a relatively-small variation in cost. Figure 8 portrays a sample of the gaps that could be left uncovered. The uncovered sections are less populated areas, so most of the end-users would not be affected, but the gaps would deter clients related to shipping, air transportation, and defense from joining. The other way to obtain coverage over specific regions, as highlighted in the previous chapters, is using GEO orbits and their variations, but we excluded this possibility due to the inherently-high latency. The alternatives that remain are B and C.

In both options, the main factors taken into account when estimating the number of satellites will be the size of the satellites and the bandwidth necessities.

SpaceX opted for a large amount of smaller satellites, while companies like OneWeb chose a smaller amount of high-bandwidth satellites. Ironically, SpaceX's satellites, at 250 kg in weight for each spacecraft, are 100 kg heavier than the competitors'. It is hard to establish which configuration will work best in the long run. Opting for a larger number of satellites carries a higher cost, increases the number of potential failures, and increases the necessary deorbit efforts and potential for space debris creation. It also has advantages in terms of redundancy, lower number of simultaneously-connected users, and potentially higher bandwidths. With users connecting to different satellites and modern beam-forming technologies, the interference caused by many users communicating on the same frequencies toward the same satellites is reduced significantly, enabling higher data-transfer rates and a better connection quality overall. A simple parallel is given by modern cell phone networks, with numerous base stations serving users and then communicating with the backbones to relay the traffic.

If more satellites enable the end users to have a wider bandwidth, the optimum number of satellites would only be limited by budget concerns. Routing mechanisms and automatic checks of the network status become essential as the numbers grow, because human intervention would be too slow or simply unfeasible. An optimum number of satellites enables high bandwidth, low launch and maintenance costs, and future scalability. The potential for scalability and improvement of the network in the future plays a fundamental role in decision-making during the initial configuration.

To address scalability, an important assumption has to be made first. In all the configurations that were analyzed, inter satellite communication will be feasible, as every satellite will always have at least two other satellites in its field of view. The Iridium network, with its modest number of satellites and six orbital planes, only spans 30 degrees longitudinally between its polar orbital planes, meaning that at any given altitude, both the satellites contained in an eastern and a western orbital plane are visible from any given satellite. This includes visibility at equatorial latitudes, where the planes are the furthest apart. This characteristic enables the creation of mesh networks, in which satellites can route data through different paths to reach a destination in the shortest or less congested way.

Having established that mesh networking was, at least geometrically, available for all the configurations that were analyzed, let us turn to the issue of scalability. There are two ways to improve the bandwidth of a constellation: adding more satellites or offloading traffic from the mesh network using more ground stations.

Adding more satellites in the future can be achieved by increasing the number of satellites in an existing orbital plane or creating new orbital planes. In the first case, the existing satellites would most likely have to be re-arranged. The maneuvers would have to make use of the on-board thrusters and propellant, which would have to be taken into account before launch. Adding new satellites in new, intermediate orbital planes would also be feasible, but would require the launch of numerous satellites to balance out asymmetries in the capacity of the network portions depending on where the additional satellites would be at any given moment.

Offloading traffic to avoid congested paths through the mesh network is another way of scaling the system, although it does not address local hotspots in the numbers of users connecting from the same region. A LEO mesh network can carry traffic through multiple satellites from an origin, for example in Europe, to a destination, for example in South America, following a potentially more efficient route than a path on the ground through various countries and undersea cables. As the number of users increases, the most common paths would be congested and the inter-satellite links would be saturated, even if new satellites were added to solve regional hotspot issues. To solve the path congestion issue, one way would be routing through different paths. This is an NP-Hard to NP-Complete problem that each satellite would have to solve in real time to direct the packets to the correct next hop. An alternative solution could be to offload data streams via ground stations, which could then carry the data for the “last mile,” or in case of congestions in the mesh, also across longer distances on traditional fiber networks.

The mesh network can be seen as a parallel highway above a traditional road network. In case the highway is congested, the packets can flow back on the traditional network, despite the longer path, to reach the destination. The problem of routing and optimization of the routes will be addressed in the next section, along with the latency considerations and advantages of mesh satellite networks over land routes.

4. Latency and bandwidth study

Having established that mesh networking will always be a key factor, the parameters to fix to start the study of the optimal constellation for low-latency internet connections, without entering the realm of economic considerations yet, are:

Client-side requirements:

- The maximum acceptable round-trip latency.
- The minimum elevation angle.
- The minimum number of satellites visible at any given time.
- Geographical coverage needs.

Constellation-side parameters:
- The altitude of each satellite.
- The number of satellites.
- The number of orbital planes and the number of satellites in each plane.
- The inclination of each orbital plane.

Tackling these points individually introduces a number of upper and lower bounds to each parameter describing the orbits and greatly restricts the possibilities of the potential configurations.

The first requirement is the maximum acceptable one-way latency between any two points on Earth to be acceptable for real-time applications. The advantage of inter-satellite links is that most of the communications happen in a near-vacuum environment. Using lasers has been accepted as equivalent to traditional fiber connections in terms of speed, and a rough estimate of the path can be established by calculating the down-link and up-link paths from the ground to the endpoints plus the semicircle that connects the two endpoint satellites. This simplification does not take into account the oblate characteristics of the Earth, but most importantly it approximates the connection between the two endpoints with an arch, which will not be true in the case of multiple segments connecting the satellites in a mesh network. It also underestimates, in half of the situations, the distance of the satellite from the ground stations, as it does not take into account the gateway and repeater latencies and it assumes that the satellites are directly above the observers. However, this difference is negligible at approximately 300,000 km per second.. This approximation will give us an upper ceiling for the orbit altitude based on the latency requirement.

$$MinLatency = (2 * MaxAltitude + \pi * (MaxAltitude + R))/c$$

With *c* being the speed of light and *R* being the average radius of the Earth, in a vacuum, we obtain:

$$MaxAltitude = \frac{MinLatency * c - (\pi * R))}{(2 + \pi)}$$

We set the minimum latency equal to 80ms to guarantee near-to-real-time communications that are acceptable both for video-conference applications [11] and more critical applications such as remote piloting of UAVs. The resulting maximum acceptable altitude would have a ceiling of 775 km to guarantee a 80ms one-way latency for the longest possible distance traveling on an ideal path.

The next constraint we have to take into account is the minimum acceptable elevation angle for the users. Telesat makes use of a minimum elevation angle for users of 20 degrees, OneWeb of 55 degrees, and SpaceX recently asked to reduce theirs from

40 degrees to 20 degrees. Given that the user base is expected not to be in areas with tall man-made obstacles, and that in highly-populated areas the ground stations can be placed on rooftops, we will consider a minimum elevation angle of 20 degrees as a limit for a satellite to be considered in the field of view of the user. As a practical example, this elevation would grant connectivity throughout the majority of the Grand Canyon.

Another constraint will be the number of satellites visible at any given time by a user in an arbitrary location. To perform successful hand-offs, at least one satellite would have to be in the field of view at any given time, with a minimal overlap with the next satellite appearing while the first one disappears. For redundancy purposes, given that the current constellations have extremely-high failure rates, with approximately 3 percent of SpaceX satellites being unreachable or uncontrollable [12], we will require at least two satellites to always be in the field of view of each user, with a third one coming into view during a hand-off, at least in non-polar areas.

The geographical coverage is also a fundamental point in the user-based constraints. SpaceX has recently decided to cover the polar areas with Starlink satellites to enable specific customers to connect. Given the leverage and budget of the customers connections in those locations, the need for coverage in the polar regions will be taken into account. The bandwidth needs though will be less onerous than in other areas. This translates into two possible scenarios. The first option is a satellite constellation that by definition takes polar areas into account, similar to the Iridium network, in which only polar orbits are used and actually provide a denser satellite area near the poles. The second option is a satellite constellation which covers the polar areas with specific orbits dedicated to those areas, while the majority of satellites in the constellation do not reach over certain latitudes. Both options will be investigated, and in the case of dedicated orbital planes for polar areas, the low-latency requirements could be made less strict for those specific orbits as a higher altitude would allow coverage with fewer satellites.

In conclusion, the client-side parameters impose global coverage with an altitude less than 780 km in non-polar areas and at least two satellites visible, from any location and at any given time, with an elevation angle greater than or equal to 20 degrees.

Next, we will turn to the analysis of the constellation-side parameters, namely:

- The altitude of each satellite,
- The number of satellites,
- The number of orbital planes and the number of satellites in each plane,
- The inclination of each orbital plane,

to optimize the performance of the constellation and minimize the number of required satellites.

5. Optimization of the satellite numbers and creation of the paths

Satellite position estimation models such as Simplified General Perturbations 4 (SGP4) will be compared to a purely geometrical model with an average Earth radius of 6,371 km used to plot the orbital paths. After having established that the error was negligible in the study of the constellation as a whole, a Matlab simulation will be run to establish the visibility of different configurations from a series of points on the Earth's surface, establishing the optimal number of orbits and inclination angles, the numbers of satellites, and the altitudes, to obtain an accurate estimate of the constellation-side parameters.

Two Line Element (TLE) notation is commonly used to identify spacecraft orbiting the Earth at any given moment in time. Although this format was initially designed for punch cards, it is still the standard today and used by NORAD to publish the observed satellites list. Starting from the TLE of a satellite, different models can be used to predict the evolution of the position of the objects.

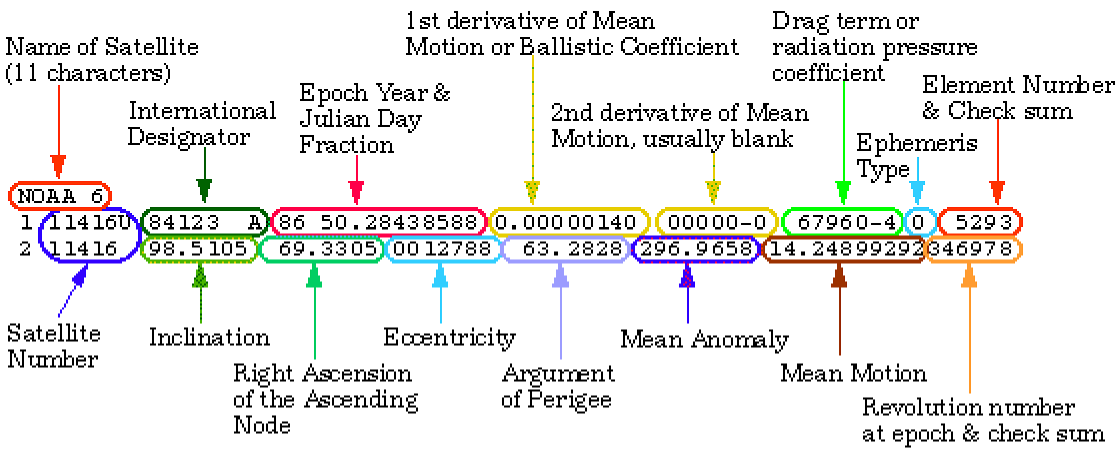

Figure 9 - TLE notation explanation [14]

The common model is known as SGP4 and is used as a propagator that takes into account a starting TLE and perturbations caused by the Earth's shape, the drag coefficient, and the effects of gravity [15]. The parameters that were fundamental to establish in the TLE were the inclination, the right ascension of the ascending node, the mean anomaly, and the mean motion.

The simplified model to translate TLE to coordinates [16] does not accurately predict the position of a satellite, including drag and perturbations, which may affect the satellites over the course of time. Thrusters, usually ion-based, on the satellites make sure that an adequate orbit can be maintained throughout the lifespan of the satellite. The instant positions of the satellites in the constellation were used for this study, so the evolution of the orbits over time was considered negligible due to the corrections which would take place.

TLE makes use of a "mean motion" instead of an altitude to represent the satellites. The mean motion refers to the number of revolutions the satellite performs around the Earth, in 1 Earth day. It is trivial to translate the mean motion in an average altitude.

Setting μ as the standard gravitational parameter, defined as the mass of the Earth times the constant of gravitation:

$$\mu = G * M$$

$$\mu = 398618; [km^3/s^2]$$

The average altitude of the satellite is equal to

$$averageAltitude = (\mu/(meanMotion) * 2\pi/(60 * 60 * 24))^{2/3}$$

The rounding error caused by the use of an average altitude instead of the use of the altitude as a function of the position in the orbit, was calculated to be less than 0.19 percent, and in particular, the approximated value calculated using a known altitude of geostationary satellites was equal to 42,242 km, while the true value is equal to 42,164 km.

This enables the altitude and the mean motion to be used in an interoperable manner throughout the calculations.

TLE was used as the format to represent the satellites in each constellation, due to its wide adoption in the industry and the interoperability with different computational models and software platforms to visualize the constellations.

To initially test the models, the Iridium constellation was reproduced. It is composed of 66 satellites, 11 on each of its 6 orbital planes. The planes are heavily inclined polar orbits, and alternate planes are rotated on their axis by a number of degrees equal to half of the angular distance between two satellites in the same orbit, resulting in a "`zig-zag" configuration between adjacent planes. This allows the entire constellation to assume the shape of a hexagonal tessellation.

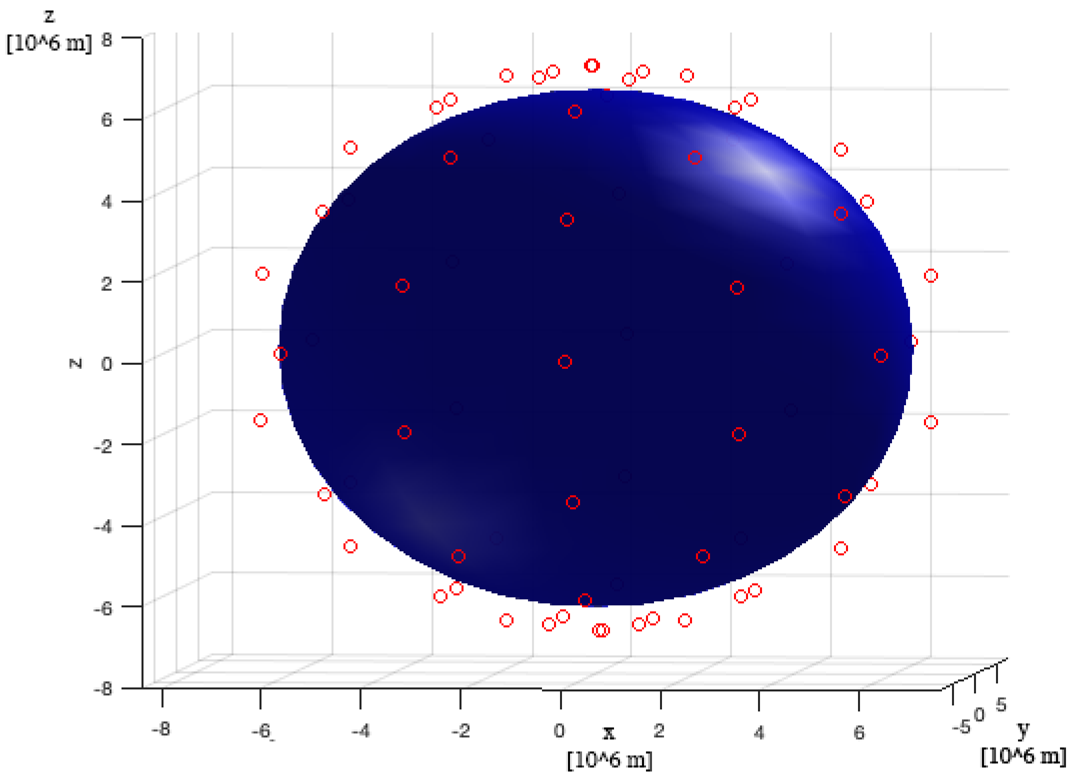

Figure 10 - Simulation of the Iridium constellation

One of the main requirements from a user's perspective was to always have at least two satellites in the field of view. To verify this requirement for each constellation, tentative points were placed on the surface of the Earth, and the sky from each point had to be observed, at the preset elevation angle, to verify how many satellites were indeed visible. The implementation of this test was performed through solving a set of equations, as follows.

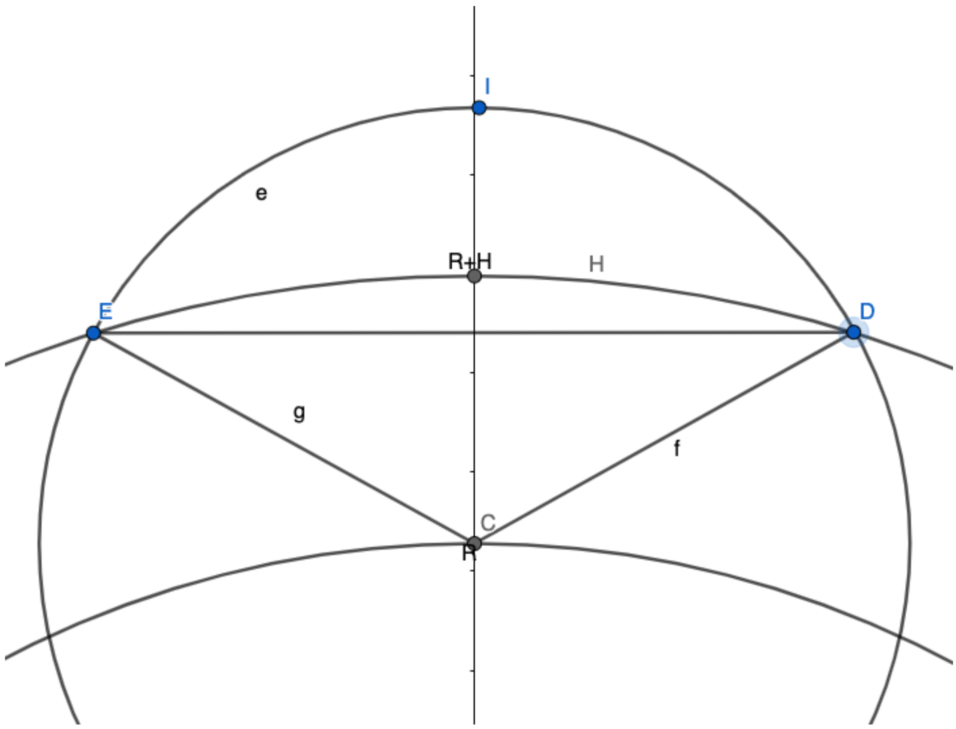

Figure 11 - Side view of a conical FOV from the ground

In Figure 11, Point C represents the observer on the surface of the Earth. The point denoted by the altitude R+H represents the satellite directly over the user, following a trajectory from E to D. The triangle in the center is a section of the cone generated by the field of view of the user with a specified angle of elevation. The interest is in obtaining the radius of the circle that indicates the entry points of the satellite in the field of view of the observer. The problem can be simplified to a geometrical problem in a plane, as the diameter of the disc under the spherical cap will correspond to the distance between E and D. To solve the problem numerically, in order to parametrize it in Matlab, the equation of the outer circle was placed in a system of two equations along with the line passing through E and C.

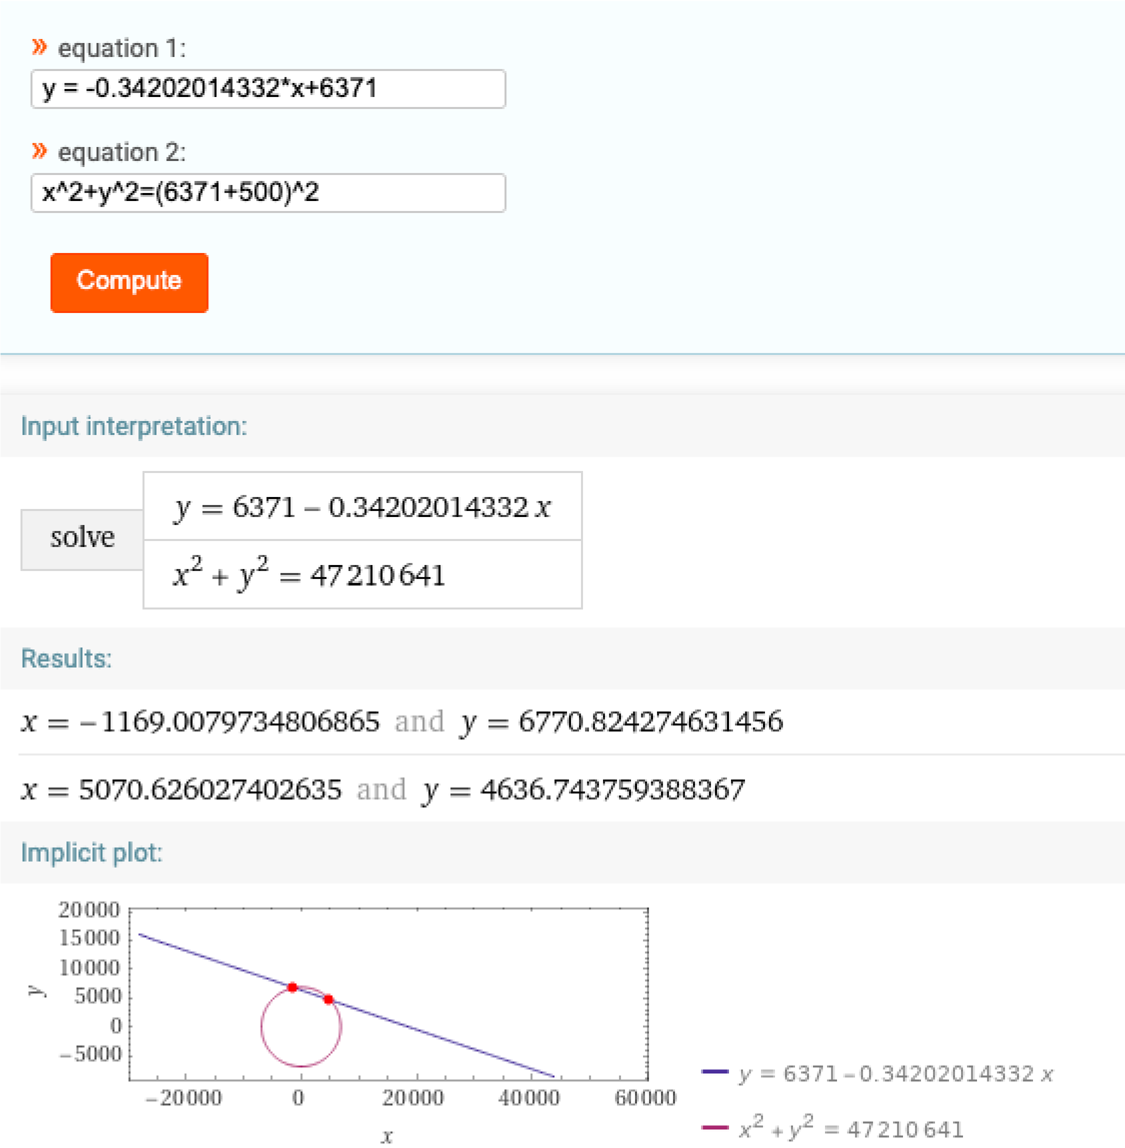

Figure 12 - Radius of the visibility cone at the satellite entry points

The line has a slope equal to the negative of the sine of the elevation angle, and a y-intercept at the radius of the Earth. The outer circle is concentric with the Earth and has a radius equal to the radius of the Earth plus the altitude of the satellite. The absolute value of the x-axis coordinate of the left-most intersection will be the radius of the visibility circle, the section of the visibility cone at which the satellite enters and exits the cone.

Implementing this simple system in Matlab enables the generation of this visibility circle for any altitude, elevation angle, and position on the ground. The circle can be therefore plotted on the map, with particular attention to the transformations that occur when plotting spherical coordinates through a Mercator projection.

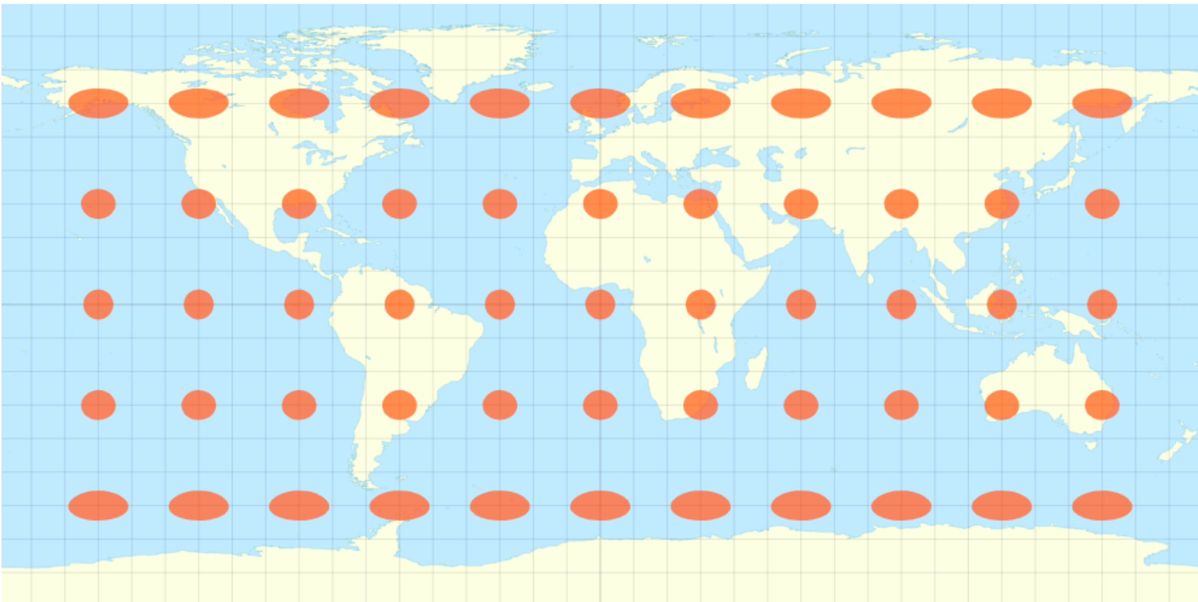

Figure 13 - Tissoy's Indicatrix on Mercator Projection

Figure 13 portrays Tissoy's Indicatrix on the Mercator projection. A visibility disc, which is close to a perfect circle near the equator, is distorted near the poles. This will be important when visually verifying the results, to check if indeed the satellites would be present inside a visibility circle in each region.

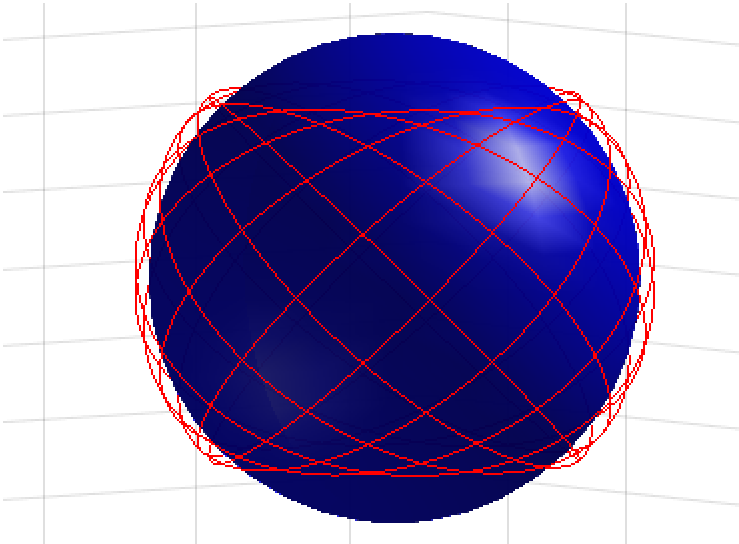

Figure 14 - Series of Inclined Circular LEO Orbits at 45 degrees

Coincidentally, constellations of circular LEO orbits have a lower concentration of satellites near the equator, as portrayed in Figure 14. The observation points will therefore be taken along the equator, in order to account for the worst-case scenario of user terminals connecting from zones with lower concentrations of satellites.

Through Matlab, different orbit combinations have been studied. The parameters that were taken into account were: the inclination of each orbital plane, the altitude of the satellites, the number of satellites in each orbital plane, and the number of observable satellites in different visibility cones at the equator.

After filtering the viable combinations of the three parameters, the initial results were plotted on a graph, Figure 15, in increasing order of number of total satellites.

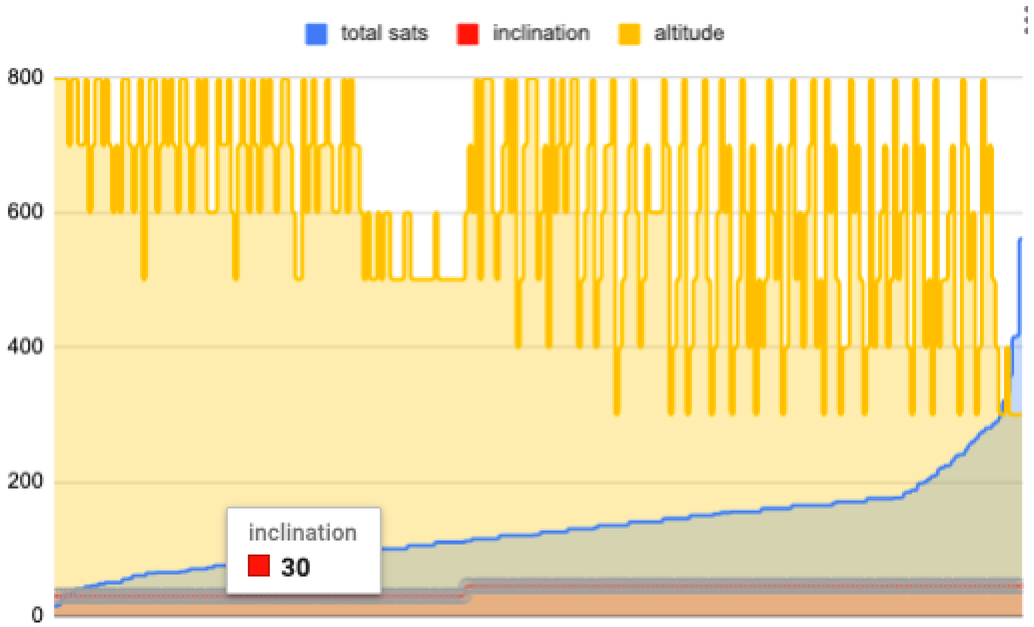

Figure 15 - Preliminary analysis of orbital parameters

The preliminary results are very intuitive. The number of satellites to provide global coverage is least when the orbital plane inclinations are lower and the altitude is higher. This is due to the observation points being placed on the equator, but it does not take into account a fundamental parameter, the minimum desirable inclination of the orbits. The maximum is clearly 90 degrees, with polar orbits, but the minimum if the poles are not covered is dictated by the user density.

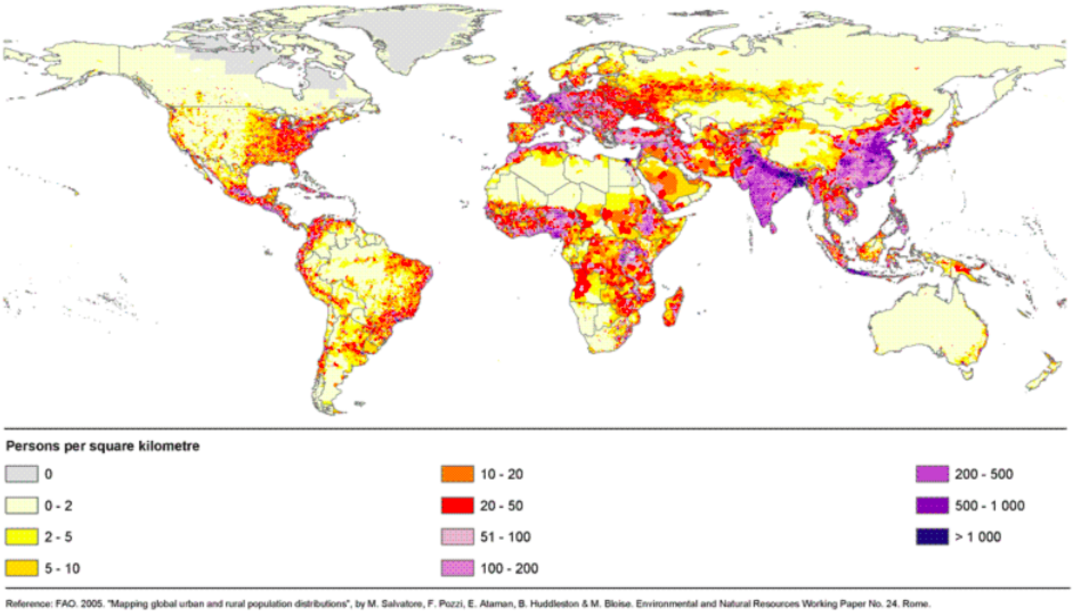

Figure 16 - Population Density Estimate 2015 [17]

A population density map gives us an estimate of the minimum latitudes that should be covered by the main portion of the constellation, in case of non-polar orbits. Aside from small portions of relatively low density in South America, Scandinavia is the critical norther-most area to be covered. Analyzing the map in more detail in Figure 17, a cutoff point has to be established between 55 and 65 degrees. The satellites orbiting at a maximum of 55 degrees will be seen at a variable radius towards the North depending on the altitude, and considering orbital altitudes of at least 350 km, 55 degrees can be chosen as a preliminary minimum latitude.

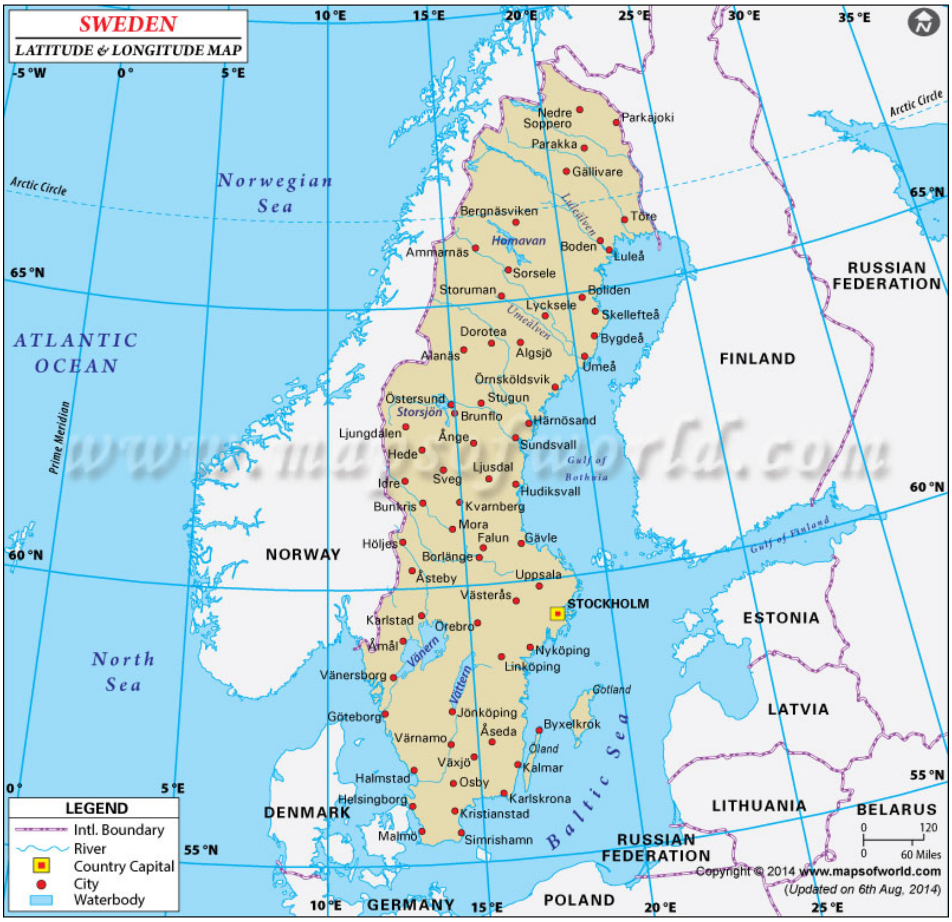

Figure 17 - Map of Sweden [18]

Running the calculations on Matlab with this more restricted set of latitudes, the following results were obtained

Satellites per plane	Orbital Planes	Total Satellite	Inclination [°]	Altitude [km]
5	23	**115**	55	700
5	24	120	55	700
5	25	125	55	700
5	27	135	55	700
5	28	140	55	700
5	28	140	60	700
9	16	144	**90**	700
5	29	145	55	700
7	21	147	75	700
5	30	150	55	700
7	22	154	55	**525**
7	22	154	70	700
5	31	155	55	700
5	32	160	55	700
7	23	161	70	700
7	23	161	80	700
5	33	165	55	700
5	33	165	60	700
7	24	168	65	700
5	34	170	55	700
7	25	175	55	525
5	35	175	55	700
5	35	175	65	700
7	25	175	70	700

7	26	182	65	525
7	26	182	75	700
7	27	189	55	350
7	27	189	75	700
7	27	189	80	700
12	16	192	**55**	**350**

Table 3 - Global coverage configurations at different inclinations

Table 13 represents a filtered list of results in which the requirement of having at least two satellites in the observer's field of view was taken into account. For each combination of altitudes and inclinations, the calculation was run for up to 500 cycles with an increasing number of planes and satellites per plane. The resulting successful configurations were recorded. The altitudes vary from 350 to 700 km and the inclinations from 55 to 90 degrees. The orbital planes were equally spaced between each other, as the number of planes varied. The highlighted lines represent the most relevant results.

The first highlighted line represents a configuration with 23 orbital planes, an inclination of 55 degrees, a total of 115 satellites, and an altitude of 700 km. This configuration could potentially be one of the least costly to implement as it contains the minimum number of satellites compared to the other configurations.

The second highlighted line represents a configuration using polar orbits, with 16 orbital planes and a total of 144 satellites. This configuration is particularly useful as it covers the polar areas without the need for secondary dedicated orbits.

The third and fourth lines in green represent configurations at lower altitudes, namely 525 and 300 km. With an increase to 154 and approximately 190 satellites, it is possible to lower the orbital altitude significantly, while maintaining an inclination of 55 degrees. It has to be noted that the line displaying the configuration at 350 km with the lowest number of satellites was excluded in favor of a configuration with three more satellites, but a significant reduction in the orbital planes, due to the complexity of adding eleven additional orbital planes instead of three extra satellites.

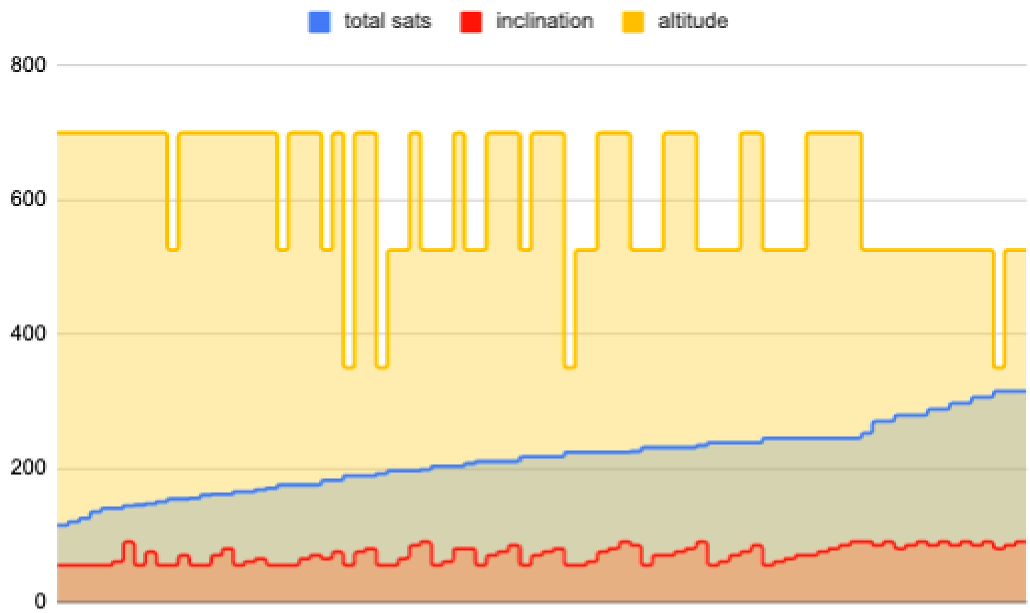

Figure 18 - Graph of the global coverage configurations

The results presented in Table 3 were plotted on a graph in Figure 18. As in the previous results graph, the configurations are ordered in increasing number of total satellites from left to right. It is clear from the graph that in case of scaling a configuration already in orbit at an altitude of 700 km by adding new satellites, a lower altitude of 500 km would become rapidly available.

The two configurations that will be taken into further consideration will be the polar orbits with 144 satellites and the one with 115 satellites in 23 orbital planes at 55 degrees of inclination.

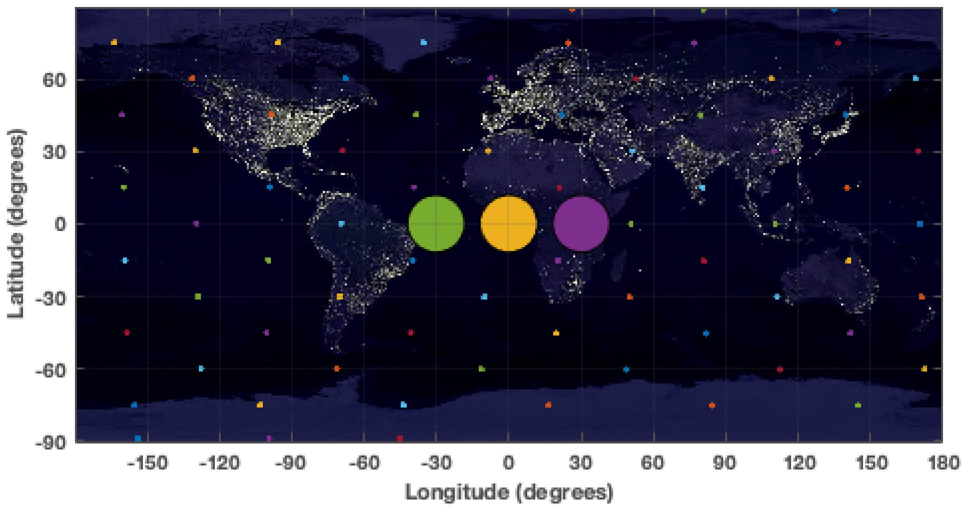

Figure 19 - Selected polar orbits on Mercator projection

The first of the two, already satisfies both the requirements of two satellites being visible at any time and global coverage, while the second one does not automatically cover the polar areas. The second configuration would therefore require one or more polar orbits to cover the space from 55 degrees to 90 degrees. The high number of orbital planes makes the "V"-shaped gaps due to the intersecting coverage circles near the polar areas between two satellites in adjacent orbits acceptably small.

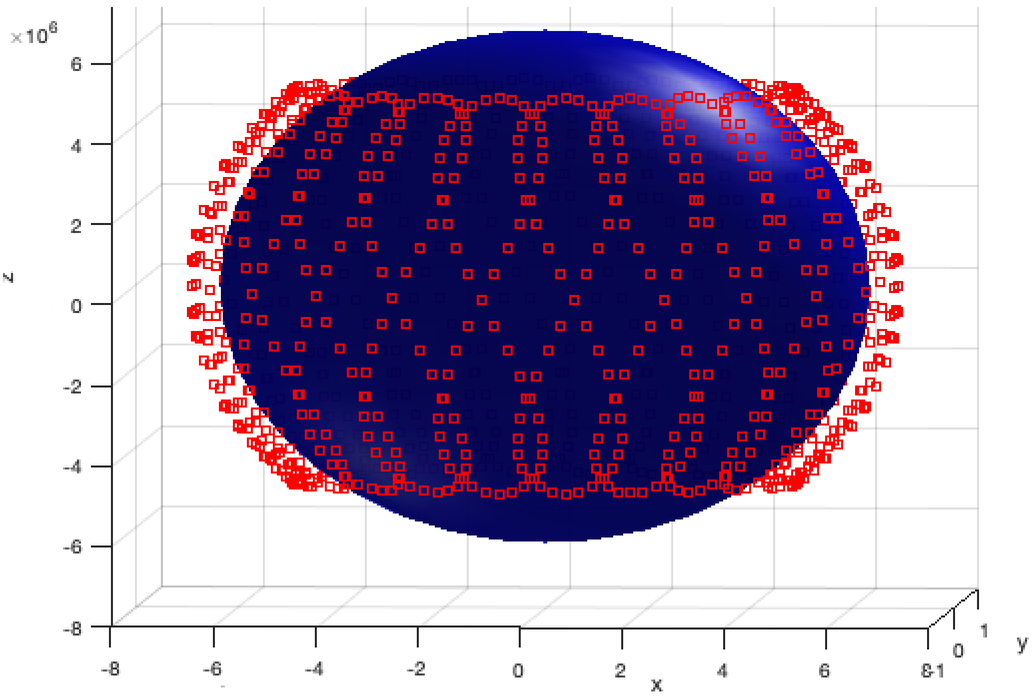

Figure 20 - Chosen orbits at 55 degrees of inclination

From Figure 19, it is clear that a large, principally uninhabited area is not covered by this solution. To cover it, a solution with two polar orbits, with 90 degrees of inclination, with planes perpendicular to each other, were hypothesized. It has to be noted, that throughout this paper, orbits defined as having an inclination of 90 degrees are approximately polar orbits as the true inclination would be closer to 87 degrees.

From a section-view, as portrayed in Figure 21, the shade area is the field of view from the northern-most point to the area already covered by the lower latitude constellation. This section, of approximately 35 degrees in latitude, needs to be covered constantly by at least two satellites. To achieve this through a polar orbit, at the minimum altitude to keep the latency at its lowest, a number of satellites of at least 16 has to be used to cover the 70 degree span with at least two satellites always in the field of vision. The minimum altitude is given by the intersection of lines CD and AF in

Figure 21, at a radius of 9344 km, yielding a resulting altitude of 2973 km. For better redundancy, two polar orbits with 8 satellites each can be used instead of a single plane with 16 satellites.

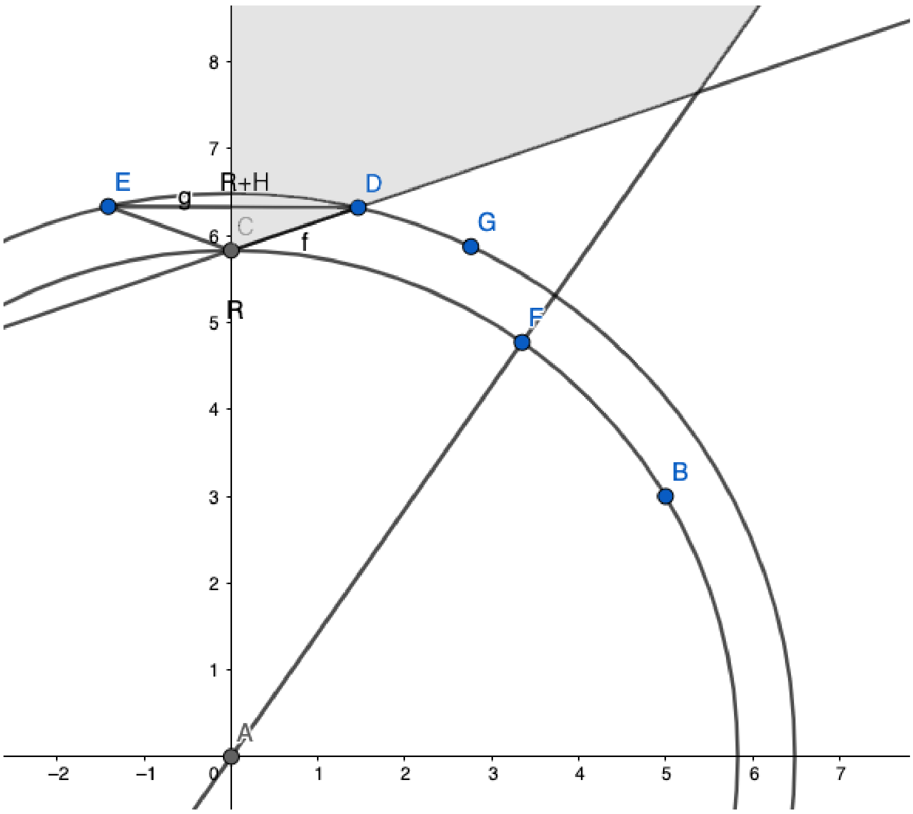

Figure 21 - Polar shade area of the chosen inclined orbits

Covering the polar areas brings the inclined configuration to 23 orbital planes at 55 degrees of inclination plus two higher-altitude polar planes with a total of 131 satellites.

The candidates for the economic study are summarized as follows.

Orbit Types	Orbital Planes	Total satellites	Altitude [km]
Polar & Inclined	2+23	131	2973, 700
Only Polar	16	144	700

Table 4 - Candidate configurations summary table

6. Latency in real-life configurations

The preliminary calculations using parts of arches to calculate the distance traveled by the signals in the inter-satellite links can be now refined to take into account the odd shapes that the paths would have to follow in the different configurations.

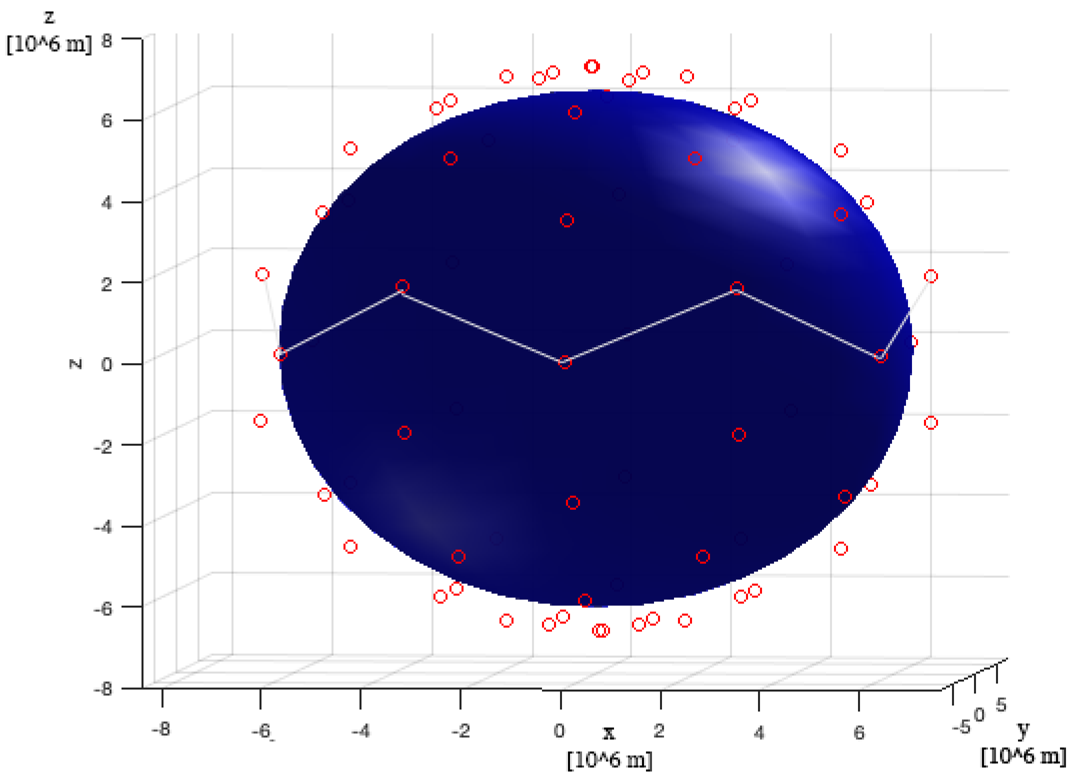

Figure 22 - Signal path in inter-satellite links

As displayed in Figure 22, the shortest East-to-West and West-to-East paths, as well as any diagonal paths for the signals will be longer than an arch of a circle described by the points at the orbital altitude.

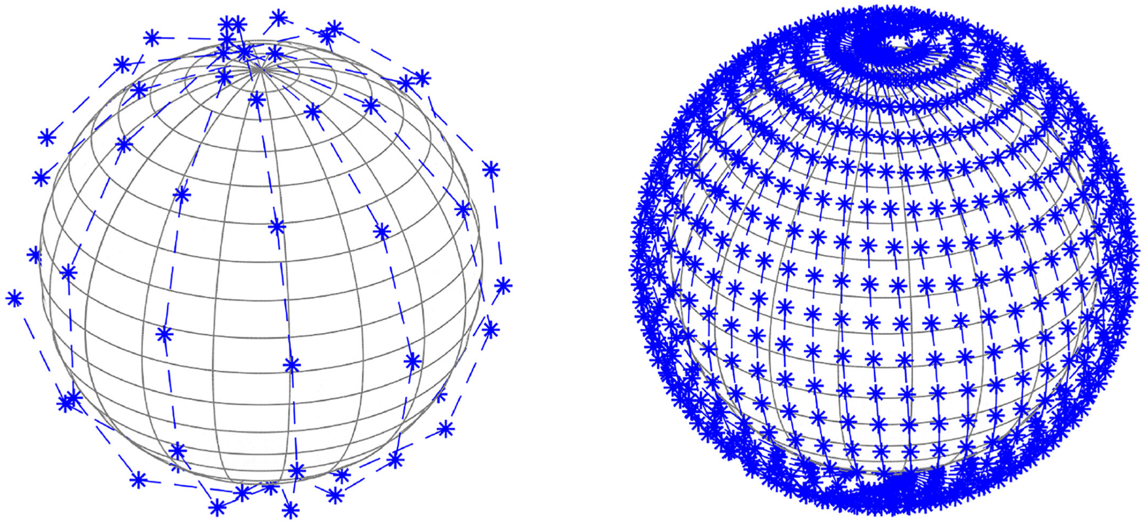

Figure 23 - Routing in Iridium vs Mega-constellations [26]

Extensive studies [26] on the best routing paths for varying numbers of satellites have been performed. Using the worst-case scenario model for the polar orbits and transmission from two points near the equator on opposites sides of the Earth, as indicated in Figure 22, the total path from East-to-West following the satellite points instead of the initial rough estimation of a portion of arch, increases the total distance of the path at the orbital altitude by 6.4%. This is an acceptable value in terms of latency, adding 4 milliseconds of latency to a one-way 180 degree path around the equator.

7. Economic and competitor analysis for each configuration

The economic estimate will be based on the launch cost for each of the configurations. Although the launch is only one phase of the lifecycle of a satellite, it is the single largest upfront cost, making up for about 21% of the total cost to consider when creating a LEO constellation. Other factors include the cost of the satellites themselves and the engineering work in the development phases. Below is a breakdown of the costs associated with the Life Cycle Cost (LCC) of a typical NASA small satellite mission from an MIT study.

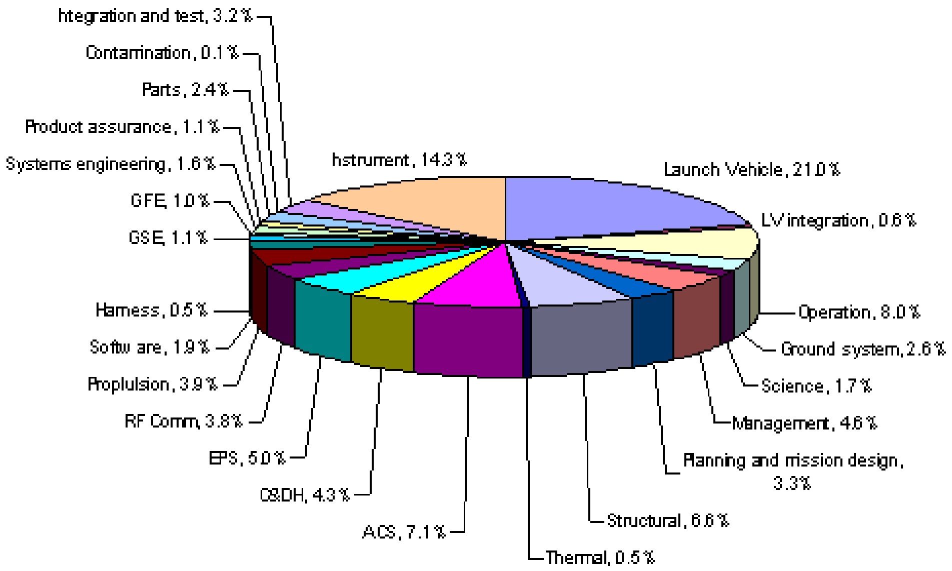

Average NASA small satellite mission	
Total mission cost	$145M
Mission development time	3 years
Dry mass	407 kg.
Spacecraft bus cost	41% TMC

Figure 24 - Average NASA small satellite mission cost breakdown

The launch cost is heavily dependent on the orbit and weight of the satellites. The most popular launch vehicle for medium payloads, considered between 2,000 kg and 20,000 kg, is the Russian-built Soyuz, with an average cost to LEO of $17,900/kg.

SpaceX has reported significantly lower costs, as their launch volume is increasing. The only relevant data from the first batches of Starlink launches, is that up to 60 were launched in each batch. The latest Soyuz launchers have fairing dimensions slightly smaller than SpaceX's Falcon 9, with a diameter of 4,1 m compared to 4,6 m, and an equal height. The cargo capacity in kilograms though, is about half that of the Falcon 9, meaning that the physical volume of the spacecraft can easily handle the payload, which will be limited by the weight capacity of 4,850 kg [24].

The weight of each satellite, considering that the most similar mesh network in production is Starlink, can be estimated to be similar at 250 kg/unit. The total cost of a launch, for a single batch of 20 satellites would therefore be of approximately 86 million USD. The unit cost of each satellite varies between 1 million USD and different sources report that SpaceX has been able to lower this cost down to 250,000 to 500,000 USD with volume.

Orbital change maneuvers would be needed in the case of launching multiple satellites in different orbital planes within the same launch, which is a technique which is still being perfected by the international space agencies. An alternative is to have one launch per orbital plane, but it would be cost-effective only if launching a higher number of satellites in each plane, or to initially scale the constellation to higher numbers of satellites to allow for higher bandwidth to the end users.

At 1 million USD, without considering the maneuvers to deliver the satellites in two different orbital planes within the same launch and the altitude difference for the polar orbits, the cost of the launch and the satellites can be summarized in Table 5. For the polar orbits in the second configuration, a single launch per orbit was considered, as the orbits would be significantly spaced apart.

Orbit Types	**Orbital Planes and Total Satellites**	**Launch types and numbers of launches**	**Cost [million USD]**
Polar & Inclined	2+23, 131	2 for polar orbits and 12 (at 2 planes/launch) for the inclined orbits	
		Launch cost	1204
		Satellite Unit Cost	131
		Total Cost	**1335**
Only Polar	16, 144	16 polar orbits	
		Launch cost	1376
		Satellite Unit Cost	144
		Total Cost	**1520**

Table 5 - Launch cost summary table

Considering this estimate takes into account less than half of the total costs displayed in Figure 22, a total estimate is significantly similar to the funds raised by companies such as OneWeb, totaling 3.4 billion USD, to deliver a working constellation.

The total estimates grant an advantage to the mixed setup of polar and inclined orbits, which provide a denser coverage at latitudes under 60 degrees. This is similar to the solution adopted by SpaceX for the Starlink constellation with 24 orbital planes at an inclination of 53 degrees with the recent addition of one polar orbital plane.

8. Conclusion

This thesis studied the creation of a LEO constellation and the optimization of the parameters to lower the number of necessary satellites to achieve global coverage and low latency internet connections.

A mathematical model was implemented in Matlab to analyze different configurations to achieve global coverage, while retaining limits in parameters such as the altitude of the constellation to respect latency requirements.

Optimal configurations were found for different altitudes, orbital planes, and numbers of satellites. Ranked by the number of satellites, two configurations were selected.

An economic estimate of the cost of production and launch of each configuration was performed, yielding results similar to the current industry standards and projects in development. In particular, the analysis finds that the cost estimates for polar orbits are higher than the corresponding costs for the mixed setup of inclined and polar orbits.

Regardless of orbit type, the cost of launch ranges between 1.2 and 1.4 billion USD, while the inclusion of satellites approximately adds an extra 150 million USD to each cost estimate. The cost estimates imply that the mixed setup of inclined and polar orbits are more advantageous, while also providing a denser coverage at latitudes under 60 degrees.

The study proved the feasibility of networks similar to the ones of OneWeb and SpaceX, in their different configurations. It concluded that inclined orbits, assisted by fewer polar orbits offer a denser coverage of the more populated regions while retaining global coverage with a saving in cost of 10 percent compared to using only polar orbits.

Further studies could be performed on the optimization of the routing between satellites to find optimal balances between guiding traffic through inter-satellite links or ground stations.

References

[1] Sakpod Tongleamnak and Masahiko Nagai - Simulation of GNSS Availability in Urban Environments Using a Panoramic Image Dataset - Volume 2017, Article ID 8047158, 26 Jan 2017

[2] Federal Communications Commission DA 21-34 January 8, 2021

[3] NASA Content - Article “July 12, 1962: The Day Information Went Global” published Jul 10, 2012 - accessed via NASA.gov website on 12 Feb 2021

[4] Patrick Daum, A Satellite Footprint Visualization Tool, Lancaster University, September 2005

[5] Shkelzen Cakaj, Bexhet Kamo, Algenti Lala, Alban Rakipi - The Coverage Analysis for Low Earth Orbiting Satellites at Low Elevation, (IJACSA) International Journal of Advanced Computer Science and Applications,Vol. 5, No. 6, 2014

[6] Marco Nugnes and Camilla Colombo, Coverage Area Determination for Conical Fields of View Considering an Oblate Earth, Politecnico di Milano, 2015

[7] Stephen R. Pratt and Others, An Operational And Performance Overview Of The Iridium Low Earth Orbit Satellite System. Communications Surveys, Second Quarter 1999.

[8] Analytical Graphics Inc., Systems Tool Kit (STK), Example of IGSO Orbit, www.agi.com

[9] The Iridium Constellation, Panorama de l'actualite de l'ASECNA - N Special - Oct. 2019

[10] Xin Yang, Low Earth Orbit Mega Constellations –Satellite and Terrestrial Integrated
Communication Networks, University of Surrey, November 2018

[11] Acceptable Latency in Video-Conference Applications, Twilio, Retrieved on 17 Feb 2021 from https://www.twilio.com/docs/glossary/what-is-latency

[12] Inigo del Portillo, A Technical Comparison of Three Low Earth Orbit Satellite Constellation Systems to Provide Global Broadband, MIT, October 1st 2018. 69th International Astronautical Congress 2018 Bremen, Germany

[13] Matt Williams, About 3% of Starlink satellites have failed so far, Universe Today, retrieved on 17 Feb 2021 from https://phys.org/news/2020-10-starlink-satellites.html

[14] NASA.gov, Definition of Two-line Element Set Coordinate System , Retrieved from https://spaceflight.nasa.gov/realdata/sightings/SSapplications/Post/JavaSSOP/SSOP_Help/tle_def.html on 17 Feb 2021

[15] Nicholas Zwiep Miura, Comparison And Design Of Simplified General Perturbation Models (Sgp4) and Code For NASA Johnson Space Center, Orbital Debris Program Office, California Polytechnic State University, May 2009

[16] Meg Noah, Satellite Orbit Coordinate Transformations, MATLAB Central File Exchange. Retrieved from https://www.mathworks.com/ on 17 Feb 2021.

[17] FAO Population Density Estimate world map 2005 - 2015

[18] Map of Sweden, Retrieved from mapsoftheworld.com in Feb 2021, Created in 2006

[19] Federal Communications Commission FCC 20-102, SAT-LOA-20190704-00057, Released July 30th, 2020

[20] Xin Yang, Low Earth Orbit (LEO) Mega Constellations – Satellite and Terrestrial Integrated Communication Networks, University of Surrey

[21] Delay is Not an Option: Low Latency Routing in Space, Mark Handley, University College London - retrieved on Nov 9, 2020

[22] Jonathan C. McDowell1, The Low Earth Orbit Satellite Population and Impacts of the SpaceX Starlink Constellation, Center For Astrophysics - Cambridge, MA, USA. Published April, 1 2020

[23] Olivier de Weck and Others, Communications Satellite Constellations MIT Industry Systems Study, Engineering Systems Learning Center (ESLC) Massachusetts Institute of Technology, Version 1.1, October 14, 2003

[24] Marta Massimiani, IOD/IOV LAUNCHERS ANALYSIS, IODISPLay, Ref. Ares (2015) 3712118 - 09/09/2015

[25] Caleb Henry, OneWeb raises $1.25 billion from returning investors, SpaceNews, March 18, 2019

[26] Manuel Roth, Hartmut Brandt,Hermann Bischl, Implementation of a geographical routing scheme for low Earth orbiting satellite constellations using intersatellite links, 26 June 2020 https://doi.org/10.1002/sat.1361

Glossary

ECI	Earth Centered Inertial, Reference Frame
FOV	Field of View
GEO	Geostationary Orbit
GLONASS	Global Navigation Satellite System
GPS	Global Positioning System
IGSO	Inclined Geosynchronous Orbit
LEO	Low Earth Orbit
LLC	Life Cycle Cost
MEO	Medium Earth Orbit
MIT	Massachusetts Institute of Technology
NASA	National Aeronautics and Space Administration
NORAD	North American Aerospace Defense Command
RGT	Repeat Ground Track
SGP4	Simplified General Perturbations 4
TEME	True Equator Mean Equinox
TLE	Two Line Element
USD	United States Dollar