Ariane 6 Engine Bay Recovery

A Trajectory Optimisation for a Guided Semi-Ballistic Re-entry

AE5822: Master Thesis Aerospace Engineering

Lucas van der Poll

Delft University of Technology

Ariane 6 Engine Bay Recovery

A Trajectory Optimisation for a Guided Semi-Ballistic Re-entry

by

Lucas van der Poll

To obtain the degree of Master of Science for Aerospace Engineering at the Delft University of Technology, to be defended on Tuesday June 30th.

Student Number:	5009235
Instructor:	M. Naeije
Thesis Committee:	M. Naeije
	E. Mooij
	W. Simons
Project Duration:	September, 2025 - June, 2026
Faculty:	Faculty of Aerospace Engineering, Delft

Cover: A mock-up of the Ariane 6 guided ballistic engine bay during re-entry.

TU Delft

Preface

This thesis represents the conclusion of my Master's Degree in Aerospace Engineering in the Space Track, following the Spaceflight Dynamics Profile, at The Delft University of Technology. The work for this thesis was done from September 2025 to June 2026.

This research represents the interest I have developed over the years in spaceflight, and specifically in launch and re-entry vehicles, over the years I have studied at TU Delft. Starting with my BSc in Aerospace Engineering in 2019, my love for spaceflight has only grown, making my choice for the Space Track a natural one. I am now proud to present this work as the pinnacle of the knowledge and expertise I have developed over my years of studying here. The subject matter at hand specifically combines the two aspects of space flight that interest me most and aims to improve the sustainability of a vehicle, which in our current day is growing into an ever more pressing matter.

I would like to take this opportunity to thank those that have lent me their support throughout the past months. I specifically would like to thank my supervisor, Marc Naeije, who, despite personal setbacks, has provided me with excellent guidance throughout the project, closely monitoring my progress and offering advice whenever any issues arose. I would also like to thank Dominic Dirkx and the other people on the discussion forums of TudatPy, who have helped me find the issues in my code that would have otherwise caused me days of extensive searching to find. I would like to thank Erwin Mooij for his suggestion of a guidance algorithm, allowing me to move forward quickly with my research project. Furthermore, I would like to thank Henk Cruijssen and Niels van Mierlo for our discussions, which helped me figure out what the best way to proceed was when I hit a wall. Finally, I would like to thank my friends and family for standing with me throughout these last years, no matter how tough things became. I could not have done it without your support and love. I would like to thank my mother for making the 3D models of the vehicle for this research and for printing them to have a souvenir of all this. Finally, I'd like to thank Nikola for all the love you have given me over the years. I love you very much.

Lucas van der Poll
Delft, June 2026

Summary

To improve the cost and sustainability of the Ariane 6, the engine bay of the first stage could be recovered for refurbishment. In this MSc research, the feasibility and cost of this mission profile are assessed from a trajectory point of view, making use of the new reusable Prometheus engine under development. To do this, a 3-DOF TudatPy simulation has been set up, and a guidance algorithm based on Apollo re-entry guidance is implemented into it. Five landing locations were selected to investigate the possibility of crossing the Atlantic during the re-entry following a LEO launch, which was deemed possible if a correction to the initial state is performed after the engine bay is disconnected from the first stage. Following an MC analysis, an MDO was performed to find the set of best identified trajectories to minimise the propellant mass needed for this initial state correction. A trajectory arriving at Cabo Verde was selected as the best option for this mission profile. This reusable mission profile was estimated to be cheaper per launch and over its life cycle and could be more sustainable if the boosters were also made reusable. This is because additional boosters may often be necessary to deal with the significant payload penalty that is incurred, although there are possibilities to work around this penalty. An issue that certainly must be addressed is raising the launch cadence to at least 30 launches per year, as below this number reusability may not be a viable option. The economic viability depends on reaching a sufficient launch cadence as much as on the technical recovery architecture. Furthermore, as this research investigated the trajectory, it must still be seen if the required hardware fits within the engine bay. If these issues can be addressed, the reusable engine bay mission profile may prove worthy of more thorough investigation, to determine if it can be adopted for the Ariane 6 and future vehicles.

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
A62	Ariane 62
A64	Ariane 64
DRL	Downrange Landing
DOF	Degree of Freedom
ESA	European Space Agency
FOSTRAD	Free Open Source Tool for Re-entry of Asteroids and Debris
GTO	Geostationary Transfer Orbit
IAU	International Astronomical Union
IERS	International Earth Rotation Service
IHS	Improved Harmony Search
ITRS	International Terrestrial Reference System
L/D	Lift-to-Drag ratio
LEO	Low Earth Orbit
MACO	Multi-objective Ant Colony Optimisation
MAIT	Manufacture, Assembly, Integration and Test
MATLAB	Matrix Laboratory
MC	Monte Carlo
MDO	Multi-disciplinary Design Optimisation
MEE	Modified Equinoctial Elements
MOEAD	Multi-Objective Evolutionary Algorithm by Decomposition
MOEAD_GEN	Multi-Objective Evolutionary Algorithm by Decomposition Generational
MSc	Master of Science
NASA	National Aeronautics and Space Administration
NRLMSISE-00	Naval Research Laboratory Mass Spectrometry and Incoherent Scatter radar extended to Exosphere, year 00
NSGA2	Non-dominated Sorting Genetic Algorithm
NSPSO	Non-dominated Sorting Particle Swarm Optimisation
PyGMO	Python Parallel Global Multiobjective Optimiser
RCS	Reaction Control System
RK	Runge-Kutta
RKDP	Runge-Kutta-Dormand-Prince
RKF	Runge-Kutta-Fehlberg
RTLS	Return-To-Launch-Site
SPICE	Spacecraft Planet Instrument C-matrix Elements
SSO	Sun-Synchronous Orbit
STL	Stereolithography
TU	Technical University
Tudat	TU Delft Astrodynamics Toolbox
US76	United States 1976 standard atmosphere model
USM6	Unified State Model using Rodrigues parameters
USM7	Unified State Model using quaternions

Abbreviation	Definition
USMEM	Unified State Model using Exponential Map

Symbols

Symbol	Definition	Unit
b	Scale factor of exponential atmosphere	[m^{-1}]
C	Cost	[MY]
C_D	Drag coefficient	[-]
C_L	Lift coefficient	[-]
C_p	Radiation pressure coefficient	[-]
c_0, c_1	Guidance deadband coefficients	
D	Drag force	[N]
d	Distance to target location	[m]
F_1, F_2, F_3	Guidance proportional control gain values	
f	Cost correction factor	[-]
f_{m_p}	Propellant mass fitness	[-]
f_r	Number of bank angle reversals fitness	[-]
g	Gravitational acceleration	[m/s^2]
H	Number of work hours	[-]
h	Height	[m]
K	Guidance overcontrol gain value	[-]
k	Heat flux constant	[-]
L	Lift force	[N]
M	Mass	[kg]
m	Mass	[kg]
m_p	Propellant mass	[kg]
N	Number of refurbishments	[-]
n, x	Cost regression values	[-]
n_p	Production number	[-]
O/F	Oxidiser-to-fuel ratio	[-]
R	Radial position	[m]
R_e	Radius of Earth	[m]
R_N	Nose radius	[m]
r	Number of bank angle reversals	[-]
S	Drag area	[m^2]
s	Downrange distance	[m]
T_1	First unit cost	[€]
t	Time	[s]
V	Velocity	[m/s]
V_E	Entry velocity	[m/s]
v	Velocity	[m/s]
W	Weight	[N]
$\mathbf{X}$	2-DOF guidance state	
β	Ballistic coefficient	[kg/m^2]
γ	Flight-path angle	[rad]
γ_E	Entry flight-path angle	[rad]
μ	Standard gravitational parameter	[m^{3}s^{-2}]
ψ	Heading angle	[rad]
ψ_d	Heading deadband value	[rad]
ϕ	Bank angle	[rad]
ρ	Atmospheric density	[kg/m^3]

Symbol	Definition	Unit
ρ_0	Atmospheric density at sea level	[kg/m^3]

1

Introduction

In recent years, the space industry has experienced a variety of changes. Perhaps the most famous of these is the push towards reusability. SpaceX prides itself on the return and reuse of its Falcon 9 first stage, reducing its launch costs and improving its sustainability. [39] Blue Origin's New Shepard is fully reusable, both the launcher and the capsule, although this is only a suborbital launch vehicle. [4] Other companies and agencies are now following suit, seeking to develop their own reusable launch vehicles, such as the European Space Agency (ESA) with Themis [19] and Rocket Lab's Neutron. [35]

In 2014, ESA made the decision to start the development of Ariane 6, the next generation of their heavy launch vehicles, to compete with the evolving space market. [14] This launch vehicle entered service in 2024 but has received some criticism even before its first launch, specifically on its lack of reusability. This often makes the cost of launching on an Ariane 6 rocket higher than its competition and provides ESA with bad optics, as it is not as sustainable as other, similar vehicles. [3] Nonetheless, Ariane 6 is now in use and will remain so for a number of years, considering the investment made into it. Already there are plans in development to make improvements, focused on reducing launch costs and improving sustainability by introducing reusability into the design. [28]

One of the possible means to introduce reusability into the design of the Ariane 6 vehicle is by devising a method for the return of the engine bay of the first stage. The research project presented in this report seeks to investigate this option for feasibility and to compare its performance to that of the non-reusable Ariane 6 and other reusable mission profiles to determine if implementing such a change is worthwhile. First, in chapter 2, an overview of the literature research for this project is presented. After this, the research question and subquestions are presented in chapter 3. This is followed in chapter 4 by an overview of methods and tools that were expected to be used during the research. After this, in chapter 5, the actual work done during the research is laid out. Then, in chapter 6, the produced results are presented and discussed. Finally, in chapter 7, the conclusion and recommendations for future research can be found.

2

Literature Review

In this chapter, the relevant information from the literature review is summarised. First of all, section 2.1 contains an overview of the Ariane 6 launch vehicle and the developments ESA is making for the next generation of its fleet. This is followed by an overview of reusable launch vehicles in section 2.2. Then, in section 2.3, an overview is given of the guidance algorithm that would be used in the research. Finally, section 2.4 presents the model that would be used to estimate the costs of the various launch vehicle configurations.

2.1. Ariane 6

The Ariane 6 is the latest generation of heavy launch vehicle used by ESA. It consists of three stages: solid rocket boosters, ignited along with the main stage, and the upper stage, which takes the payload to its final orbit. Two variants are available, Ariane 62 (A62) and Ariane 64 (A64), which differ in the number of boosters used and in their payload capabilities to Geostationary Transfer Orbit (GTO) and Low Earth Orbit (LEO). How these variants compare to one another can be found in Table 2.1. An illustration of these two variants can also be found in Figure 2.1.

Table 2.1: A comparison of the A62 and A64 variants.[14][33][38]

Ariane Variant	A62	A64
Number of boosters	2	4
Payload mass capability to GTO (t)	4.5	11.5
Payload mass capability to LEO (t)	10.4	21.7
Estimated cost per launch (2026M€)	105.0	115.0

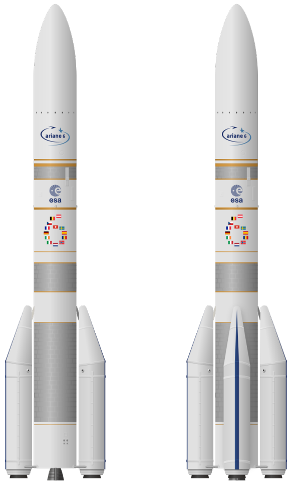

Figure 2.1: The Ariane 62 (left) and 64 (right).[36]

The design of the Ariane 6 was chosen to be as flexible as possible, with its variants offering options to deploy payloads of various masses to various orbits. This importantly would also allow ESA to launch medium-lift payloads without needing to rely on the Russian Soyuz rockets, which the agency had been using for this purpose. It was also seen as important to retain the capability to launch multiple satellites, an option that was available on the Ariane 5. [44] The launches are performed from ESA's launch complex in French Guiana.

The first launch of the Ariane 6 occurred on the 9th of July 2024. This launch was a partial failure, as the upper stage failed to re-ignite and deorbit in an effort to limit space debris. Even before its maiden flight, however, the programme suffered from criticism. This is mainly due to the high development costs the Ariane 6 programme has suffered, as well as the fact that the vehicle is entirely single-use. [3] Already, development has begun within ESA to develop the technologies needed to incorporate reusability in its launch vehicles, in the form of the Ariane NEXT project. [28] Since 2017, development has started on Prometheus, a low-cost reusable engine ESA hopes to use going forward. Hot-fire tests started in 2022, and several successful tests have been performed. [15] This engine is intended to be used on Themis, ESA's testbed for reusability technologies. [19] The technologies and expertise developed here will likely be used when the Ariane 6 is replaced by the next generation of launch vehicle, but this will not occur for many years. Significant gains could therefore already be made by introducing some of these developments on the Ariane 6 already. This is something that is indeed already under consideration for the new Prometheus engine.

The engine bay of the Ariane 6 is the most valuable part of the first stage, amounting to approximately $10-15\%$ of the total costs of the vehicle, according to the numbers provided to the project that preceded this research. [11] The engine frame is manufactured by Airbus in the Netherlands, [16] and then transported to a facility outside of Paris, where the Vulcain engine is incorporated into it. [18] From here it is then transported to the launch site in French Guiana.

The Prometheus engine, which is undergoing testing, is set to be the engine for future ESA launchers. The comparison between it and the Vulcain engine, which is currently in use, can be found in Table 2.2. It can be seen that the Vulcain engine does perform slightly better than the Prometheus engine, but the latter is about a tenth of the cost, is re-ignitable, and uses liquid methane, which is more dense and requires less cooling than liquid hydrogen. The Prometheus engine is also designed to be reusable, up to five times.

If any modifications are made to the vehicle, these may add mass, which in turn results in the payload capacity being reduced. The first stage has a payload penalty of 0.25 [-], meaning that for every [kg] added to the first stage, the payload capacity is reduced by 0.25 [kg]. For the upper stage, the payload penalty is 1.0 [-]. [11]

Table 2.2: A comparison of the Vulcain and Prometheus engines.[15][18][11][48][47]

Engine	Vulcain	Prometheus
Re-ignitable	No	Yes
Propellant	LOX/LH2	LOX/CH4
Cost (2026M€)	± 10.0	± 1.0
Vacuum Thrust (kN)	± 1325	± 1200
Vacuum Specific Impulse (s)	± 430	± 360

Figure 2.2: The Prometheus Engine. [15]

2.2. Reusable Launch Vehicles

A reusable launch vehicle can be defined as one that is used to perform multiple flights. First, in subsection 2.2.1, the advantages and disadvantages of a reusable launch vehicle are laid out to determine in what scenario it makes sense to use a reusable launch vehicle. Then, an overview is provided in subsection 2.2.2 of the considerations that must be made in the architecture and mission profile of the reusable vehicle.

2.2.1. Advantages and Disadvantages of Reusable Launch Vehicles

A reusable launch vehicle has certain advantages and disadvantages compared to a non-reusable one. [9][46] This must be considered before it is considered if using a reusable launch vehicle is actually worth it. The most important of these are laid out below:

The main advantages are the following:

- Reusable launch vehicles have a reduced manufacturing cost per launch compared to single-use vehicles. This is because not all the hardware must be manufactured once again. The cost to recover and refurbish a piece of hardware should be less than the cost to produce it anew, else it makes little sense to introduce reusability.
- Reusable launch vehicles reduce the manufacturing load. As not all components must be manufactured for every flight, this reduces the workload of the facilities manufacturing these components. This can be a significant advantage if the output of facilities is a limiting factor in the number of flights that can be taken. This can also improve turnaround times between launches.
- Reusable launch vehicles are more sustainable. This is especially important in our current day, as many industries are doing what they can to reduce their environmental impact. Being able to not produce a new component for every flight would reduce the environmental footprint of the launch vehicle.

The main disadvantages are the following:

- Reusable launch vehicles have an increased development cost. This is because the additional complexity and reliability implemented in these designs typically calls for specific expertise, more stringent requirements, and additional testing. All of these increase development costs.

- Reusable launch vehicles are generally less efficient. This is because they are designed for additional goals beyond the mission parameters they have in common with a non-reusable equivalent, namely that parts of them must return safely for recovery. The section to return will have to carry additional fuel, heat shields, aerodynamic control surfaces, etc. Those elements of the design must also be more reliable and be designed with recovery, refurbishment and service in mind. This can result in a payload mass penalty, that is, a reduction in the payload mass that can be delivered to the target orbit as a result of using the reusable vehicle.
- Reusable launch vehicles are more failure-prone. This is due to the increased complexity introduced by the need to recover parts of the vehicle. The additional operations and the design choices made to allow for these make the design more complex and therefore increase the chance for some failure during the mission.

Designing and using a reusable launch vehicle is mainly advantageous when a high frequency of launches is needed: the benefit of the reduced manufacturing costs is maximised, as is the reduction in environmental impact, and the reduced manufacturing load allows the vehicles to still be supplied with hardware without running into supply chain issues. This also spreads out the increased development costs. The reduced efficiency and increased failure risks are downsides that must simply be accepted, although increasing the expertise in reusable vehicles may also reduce these eventually.

2.2.2. Architecture and Mission Profile Considerations

There are several considerations to make for the architecture of a reusable launch vehicle. The first of these is to determine the degree of reusability. A vehicle can be fully reusable or partially reusable. [9] Most launch vehicles are partially reusable, with typically the first stages being returned. The main reason for this is that the greater the altitude and velocity of a rocket stage, the more effort and energy will be required to return it safely to Earth. [6] First stages, which rise to lower altitudes and travel at lower velocities, are less difficult and expensive to return for reuse, making them more economically viable. Which stages or components are to be reused is an important consideration to make before determining the architecture and mission profile of the launch vehicle and the elements that are to be recovered.

The first consideration is the type of trajectory the returning stage will follow. To improve clarity, the trajectory is divided into four stages:

- Orbital flight, which takes from the separation of the returning stage to when it hits the atmosphere.
- Re-entry, which takes from when the returning stage hits the atmosphere to when it no longer needs additional protection from atmospheric effects.
- Atmospheric flight, which takes from the end of the re-entry phase until the start of the landing and recovery phase.
- Landing and recovery, which takes from when the stage makes its final preparation for landing and/or recovery until the stage is recovered.

The considerations in each of these phases are described in more detail below. Generally, three types of trajectories have been identified, these are:

- Unpowered trajectory
- Powered return to launch site (RTLS) trajectory.
- Powered downrange landing (DRL) trajectory.

An unpowered trajectory is a trajectory that does not use the engine during the return. The RTLS trajectory is a trajectory where the stage performs retroburns with its engine to return to a location close to the launch site. The DRL trajectory is a trajectory where the stage performs retroburns during the return but does not return to the launch site, arriving instead at some location downrange of the launch location. [7] An example of each of these trajectories can be found in Figure 2.3, Figure 2.4 and Figure 2.5 respectively.

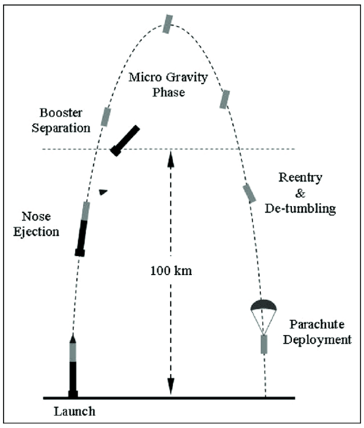

Figure 2.3: An example of an unpowered return trajectory.[29]

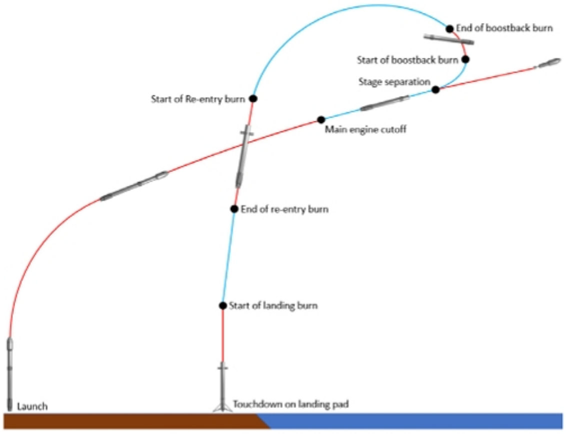

Figure 2.4: An example of an RTLS trajectory.[7]

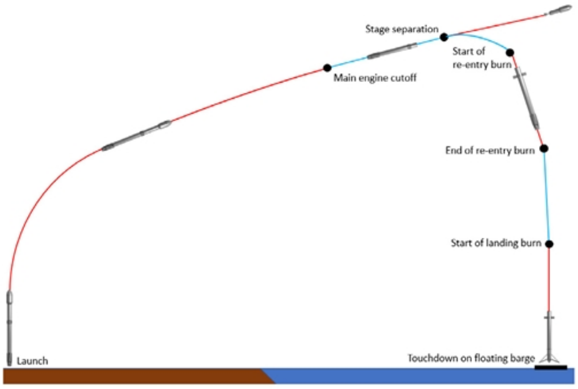

Figure 2.5: An example of a DRL trajectory. [7]

First of all, the returning stage disconnects from the rest of the launch vehicle and commences its orbital flight phase. During this time, the stage may perform a burn with its engine to adjust its trajectory. This is the case for the RTLS trajectory, which at this point performs a retroburn to counteract its horizontal velocity. Both the unpowered and DRL trajectories are purely ballistic in this stage, with no thrust applied. Regardless of the trajectory type, this phase typically includes an adjustment of the attitude of the returning stage to ensure the vehicle has the attitude required for the next phase. As this occurs at high altitudes, where atmospheric density is too low for aerodynamic surfaces to have an effect, a reaction control system (RCS) is needed on the returning stage for any needed manoeuvres. [7]

As the atmospheric density rises, the returning stage will begin to suffer the effects of re-entry. The stage will have to endure high heat loads during this phase and must decelerate from its very high orbital velocity. For the RTLS and DRL trajectories, this is done with a retroburn. This burn reduces the velocity of the stage, and the exhaust plumes shield the stage from the effects of re-entry. This reduces the need for thermal protection and, if the components are designed accordingly, can remove the need for components such as a heat shield. [7] The unpowered trajectory, of course, does not use retroburns during this phase and so must decelerate aerodynamically and must be protected from the harsh re-entry environment with more extensive thermal protection. [43] An option that is of special interest for reusable vehicles is deployable heat shields. [8] These are interesting because they are contained within the stage and deployed after separation, limiting the impact the thermal protection has on the shape of the launch vehicle before separation. An illustration of such a heat shield can be found in Figure 2.6.

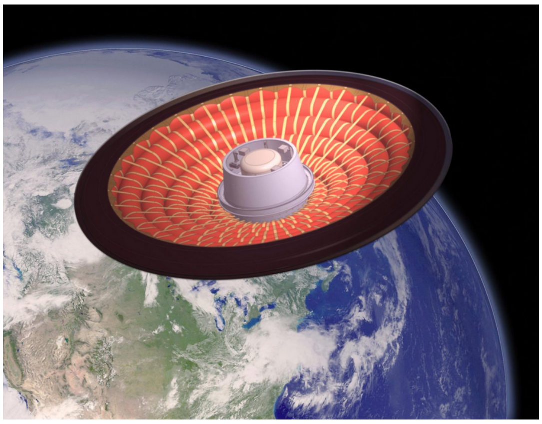

Figure 2.6: A deployable heat shield. [32]

Once the harshest conditions of the re-entry environment have passed, the returning stage is in its atmospheric flight stage. Regardless of the trajectory type, the situation here is the same: typically the returning stage will not be performing any burns, instead simply gliding to a lower altitude. To control the attitude and, with it, the gliding trajectory of the vehicle, aerodynamic surfaces can be used, of which there are generally three types in use: petals are convenient due to their lack of external components but require very high actuator forces to control and contribute a significant amount of mass. Grid fins offer an attractive option, as these are compact, effective for their mass over a wide range of flow conditions, and result in small hinge moments. If greater mobility is desired, planar wings can be implemented, at the cost of a greater volume when stowed and greater actuation forces. [26] Each of these options is illustrated in Figure 2.7. An alternative option that has a great history of use is parachutes or parafoils. These do, however, have the disadvantage of less control over the descent, as a longer time in the air will result in more significant deviations due to wind. [30]

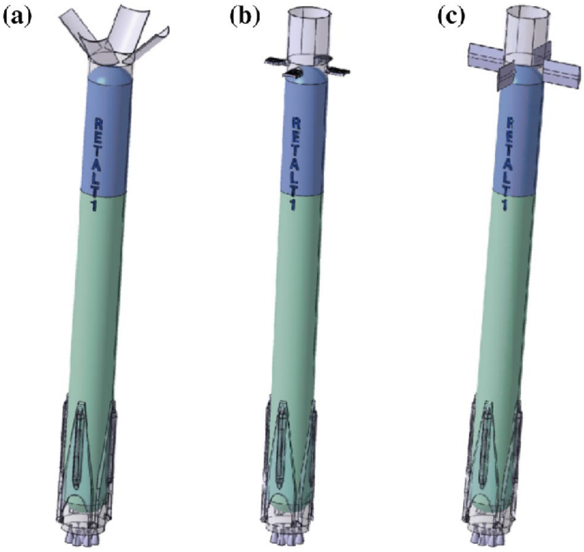

Figure 2.7: Examples of possible aerodynamic control surfaces on reusable rockets. These are a) petals, b) grid fins and c) planar fins. [26]

Finally, the returning stage approaches the surface of the Earth, and it enters its landing phase. In the case of the RTLS and DRL trajectories, a final retroburn is typically initiated here to decelerate the returning stage to a velocity the landing method can handle. The landing can be a splashdown, a mid-air capture, or a landing. Splashdowns, as water landings are called, require the stage to be buoyant by itself in water or require the stage to deploy flotation devices. Splashdowns do often occur at fairly high velocities, incurring a significant shock at the moment of impact. [31] Furthermore, seawater, where most water landings occur due to the generally wide, calm expanse of the sea, has very significant adverse effects on spacecraft components. Corrosion due to salt water begins to show visible signs within 30 minutes of splashdown, significantly increasing the costs of refurbishment. [5]

There are multiple options when it comes to captures of the stage in the air. One such approach is a catch by helicopter. This does require the stage to travel at a velocity the helicopter can match, which will generally mean it needs a parachute. The stage also cannot be too heavy, as the helicopter must be able to sustain the sudden shock of catching the stage. Another option for capture in the air is to capture by tower, which requires minimal adaptations to the returning stage, but this does require a very high degree of precision in the final approach.

Finally, a returning stage may perform a landing by touching down on a surface. This does require a high degree of precision in the landing, as typically a landing pad or floating landing platform is needed. The stage requires landing gear of some sort, which will typically be deployable, remaining folded until the last moment. The stage will have to reduce its velocity low enough so that the landing gear can sustain the landing loads and bring the stage to rest and ensure the final orientation is within acceptable limits. [42]

Although the three trajectory types have been treated as completely distinct, a mission may combine these by implementing the solutions of different trajectory types at various phases of the return. It is important to note, however, that if any retroburns are desired to be implemented in the return, a re-ignitable engine is strictly necessary, and the returning stage must carry sufficient propellant at the time of stage separation.

After the returning stage has been recovered, it must be refurbished before it can be used again. This requires the components to be inspected and tested, repairing or replacing components as necessary. This process can be sped up if the reusable stage is designed with an eye on easy access for maintenance. [9]

Previous research has been done on introducing reusability in the Ariane 6, during which a mission profile was selected. This project chose a partially reusable option, detaching the engine bay after the first stage separation. This was done because the engine bay at the aft end of the first stage contains several key components, amounting to the majority of the cost of the first stage. Discarding the rest of the first stage makes it easier to return these valuable components whilst not losing too much value in the components that are not reused. The engine bay would perform a passive atmospheric re-entry before further decelerating with a parachute. Once acceptable velocities were reached, a parafoil would be deployed, allowing for a controlled descent, after which the stage would be captured mid-air by a helicopter, which would land the engine bay on a recovery vessel. [10] An image of the vehicle designed during this research project can be found in Figure 2.8.

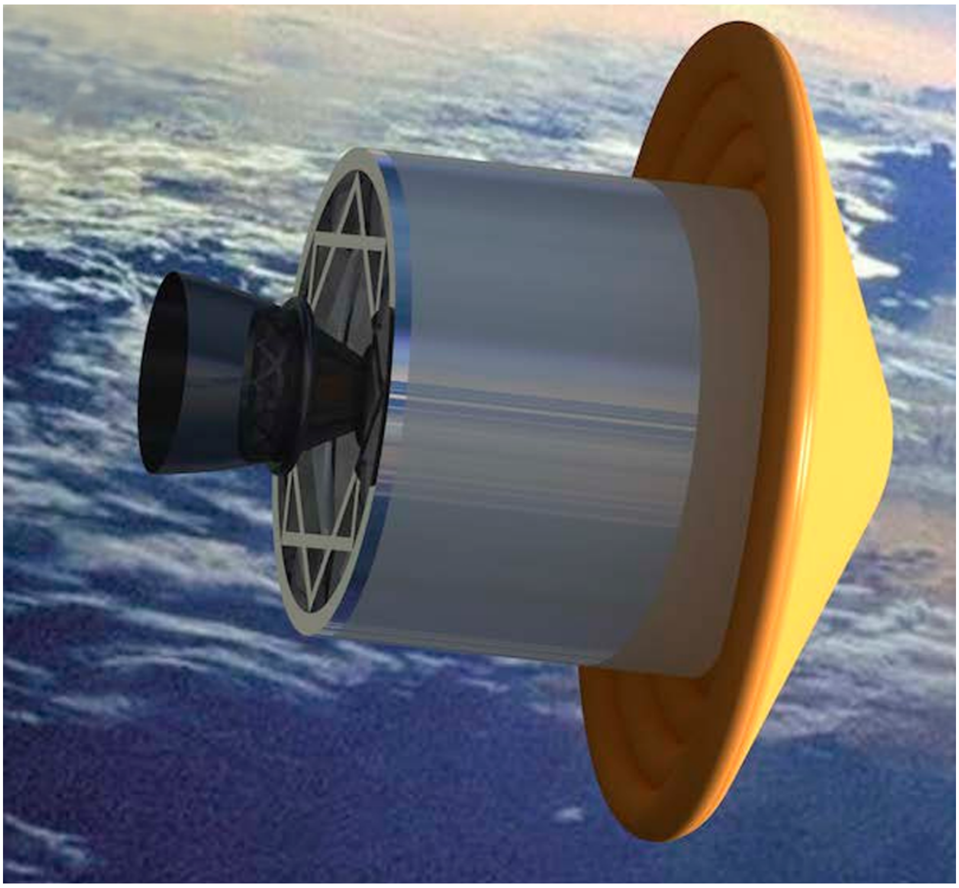

Figure 2.8: The vehicle designed during the previous research project on Ariane 6 engine bay recovery.

2.3. Re-entry Guidance

To have the vehicle travel to a specific desired location, guidance is needed. The vehicle has a lift vector for a skip entry to be feasible. This lift vector, which is perpendicular to the airflow, can be used to control the trajectory by rolling the vehicle about its velocity vector. The bank angle ϕ, which is the angle the lift vector makes with the vertical, can be commanded to the vehicle, which can perform a roll to match this commanded bank angle. In doing so, the lift vector is angled, reducing its vertical component and increasing its lateral component. The reduction in the vertical component results in the vehicle making a steeper descent into the atmosphere, increasing the loads but reducing the total downrange distance of the trajectory. Meanwhile, the lateral component of the lift vector causes the vehicle to veer away from its current heading. This can be a beneficial consequence if the vehicle is not going towards the desired location, but if it is already, the effect can be mitigated by performing bank reversals. This is a procedure where the bank angle is inverted once the vehicle leaves a certain predetermined envelope, resulting in the vehicle returning to the initial heading and then veering off course the other way. By repeatedly performing these bank reversals, the trajectory snakes its way downrange to the desired location. [27] As the bank angle reversals will be made using RCS thrusters with limited propellant, it is of interest to minimise the number of bank angle reversals.

To control the bank angle to manage the trajectory, a guidance algorithm must be used. There are multiple such algorithms available, but one is to be considered for this research: A version of the Apollo entry guidance algorithm that was used by NASA for the landing of the Curiosity rover on Mars. This algorithm was chosen for a number of reasons. The Apollo entry guidance algorithm is one with extensive heritage, having been successfully used in many missions. This particular implementation is a relatively recent one and so is quite up-to-date, while also having an implementation readily available. [1]. In addition, as the guidance algorithm is not extremely complex, it could be effectively verified, providing a high degree of confidence in its performance. The execution of the guidance algorithm is also not computationally intensive, which means it can be used by simpler on-board computers or free up space on more complex systems for other computations.

The guidance algorithm works by creating a reference trajectory using the expected entry interface state of the vehicle. The reference trajectory is propagated using an exponential atmosphere and a user-defined reference bank-angle profile. This reference trajectory is created before the re-entry is performed and provided to the vehicle, which will use it as its reference to command its bank angle. The reference trajectory and the guidance algorithm use a 2-DOF representation of the vehicle, neglecting lateral motion. This allows the spacecraft's position and velocity to be fully represented using four parameters: The altitude h, The downrange distance s, the velocity magnitude v, and the flight-path angle γ. This leads to the vehicle state vector $\mathbf{x}$ found in Equation 2.1. To find the equations of motion, the derivative of this state $\dot{\mathbf{x}}$ must be defined as well, which can be found in Equation 2.2.

$$\mathbf{x} = \begin{bmatrix} h \\ s \\ v \\ \gamma \end{bmatrix} \tag{2.1}$$

$$\dot{\mathbf{x}} = \begin{bmatrix} \frac{dh}{dt} \\ \frac{ds}{dt} \\ \frac{dv}{dt} \\ \frac{d\gamma}{dt} \end{bmatrix} = \begin{bmatrix} v\sin(\gamma) \\ v\cos(\gamma) \\ -\frac{D}{m} - g\sin(\gamma) \\ \frac{1}{v}\left(\frac{v^2\cos(\gamma)}{R_e+h} + \frac{L}{m}\cos(\phi) - g\cos(\gamma)\right) \end{bmatrix} \tag{2.2}$$

These are then used to propagate a reference trajectory and then calculate factors F_1, F_2 and F_3, which act as gain values in a proportional controller, as seen in Equation 2.3.

$$\phi = \phi^*(v) + K\frac{-(s - s^*(v)) - F_2(v)(\dot{h} - \dot{h}^*(v)) - F_1(v)(\frac{D}{m} - \frac{D}{m}^*(v))}{F_3(v)} \tag{2.3}$$

Where K is the overcontrol gain value specified by the user, and values denoted with a $*$ represent those quantities in the reference trajectories. Various quantities in Equation 2.3 are indicated as being a function of v. This is because during the re-entry, the vehicle will compare its current velocity to the reference trajectory and select the values for s^*, $\dot{h}^*$, $\frac{D}{m}^*$, F_1, F_2 and F_3. In this way the vehicle will compare its state to the state the vehicle has at the same velocity in the reference trajectory and then attempt to adjust its trajectory to match. [1]

This is of course only a longitudinal guidance algorithm, attempting to place the vehicle at the correct downrange distance by varying the bank angle. If no lateral guidance is implemented, this would result in the vehicle veering off-course. The lateral guidance algorithm that will be used is one making use of heading angle deadbands. Every guidance cycle, the current heading angle of the vehicle is determined and compared to the desired heading angle to fly directly at the target location. If this heading error is greater than the current deadband value, the sign of the bank angle commands is reversed. The heading angle deadband value is commonly defined as in Equation 2.4, and this expression will be used for this research as well. [37]

$$\psi_d = c_0 + c_1 V^2 \tag{2.4}$$

Where ψ_d is the deadband value, and c_0 and c_1 are the deadband coefficients, which are selected by the user.

This implementation of the guidance algorithm offers four decisions to be made by the user:

- The reference bank angle used in the generation of the reference trajectory.
- The overcontrol gain value K.
- The deadband coefficient c_0.
- The deadband coefficient c_1.

2.4. Launch Vehicle Cost Estimation

One of the goals in introducing reusability is to reduce the costs of the launch vehicle. To be able to do this, a method is needed to estimate the costs of a theoretical launch vehicle. As there is no available data on the specific components a theoretical launcher uses, a bottom-up approach is not possible, and so a top-down method is needed. One such option is the Transcost model. This divided the costs of the vehicle into three parts: development cost, vehicle cost and flight operations cost. Development costs include the development of each of the vehicle elements, the project system engineering, integration and testing, and testing. Vehicle cost is simply the cost of producing the vehicle. Operation costs comprise direct operation costs, indirect costs related to the launch, and refurbishment costs.

The method relies on parametric cost estimations in the form of Equation 2.5.

$$C = n \cdot M^x \cdot f_i \tag{2.5}$$

Where C is the cost, n and x are regression values specific to a spacecraft part based on statistics, and M is the mass of the spacecraft part. Additional correction factors f_i are often multiplied in to account for various effects, such as the development status of a part, its technical quality, or the learning factor of the production team. [25]

The production costs will be estimated using a method based on this cost estimation relationship, using the ratio of the masses of the reusable vehicle to the regular one. This can be found in Equation 2.6. The learning factor f_4, accounting for the learning factor of the production team is calculated as in Equation 2.7. [11]

$$C_{pr} = 2\left(\frac{M_{reusable}}{M_{expendable}}\right) f_4 C_{stage} \tag{2.6}$$

$$f_4 = 0.85^{\frac{log(n_p)}{log(2)}} \tag{2.7}$$

The refurbishment, transport and propellant costs are determined, as in a recent study into reusable options for Ariane 6 [34], is based on a method for the estimation of launch vehicle costs known as $OLSTICE. [13] The refurbishment costs are calculated as in Equation 2.8. Note this includes a factor representing the increase in the cost of refurbishment for each reuse, as the vehicle endures more wear. The transportation costs are calculated as in Equation 2.9, where t_{travel} is in hours, H is the number of work hours in a work year, l is the number of launches per year, and C_{MY} is the cost of a man-year. The propellant costs are calculated as in Equation 2.10, where O/F is the oxidiser-fuel ratio and C_f and C_{ox} are the costs of the fuel and propellant per kg.

$$C_{ref} = C_{pr} \cdot 0.25\left(N^{\frac{\ln(1.15)}{\ln(2)}}\right) \tag{2.8}$$

$$C_{trans} = \frac{1.2 \cdot t_{travel} \cdot 20}{H} + \frac{760000}{l \cdot C_{MY}} \tag{2.9}$$

$$C_{prop} = \frac{M_p}{O/F+1} C_f + \left(M_p - \frac{M_p}{O/F+1}\right) C_{ox} \tag{2.10}$$

The development costs are defined as in Equation 2.11, where C_{eng} and C_{MAIT} are the engineering and MAIT costs, respectively. These are each calculated as in Equation 2.12 and Equation 2.13, where T_1 is the cost of the first unit, created by filling in Equation 2.6 with $n_p = 1$. [13]

$$C_{dev} = C_{eng} + (C_{eng} + C_{MAIT}) \cdot 0.1 + C_{MAIT} \tag{2.11}$$

$$C_{eng} = 3.0 \cdot T_1 \cdot 0.9 \tag{2.12}$$

$$C_{MAIT} = T_1 \cdot 0.9 \cdot 3.1 \tag{2.13}$$

Using these, the cost per launch can be calculated. The life cycle cost can be calculated by adding the costs over the entire lifetime of the mission, as in Equation 2.14.

$$C_{LCC} = C_{dev} + C_{ops} \cdot \#launches + C_{prod,tot} \tag{2.14}$$

Here the operation costs also include the recovery and refurbishment of the vehicle. The cost per launch for the vehicle is then simply obtained by dividing C_{LCC} by the number of launches, resulting in the expression in Equation 2.15.

$$C_{launch} = \frac{C_{dev} + C_{prod,tot}}{\#launches} + C_{ops} \tag{2.15}$$

This is a useful metric that can be used to determine where the break-even point may lie between a reusable and non-reusable vehicle. The final metric by which the costs can be compared is the cost per payload mass kilogram, obtained by simply dividing the cost per launch by the maximum payload capacity of the vehicle.

3

Research Objective

In this chapter, the research objective is presented. In section 3.1, four possible mission profiles for a reusable Ariane 6 are presented, of which the most promising is selected to be the main focus of the research. This is then followed by section 3.2, which lays out the research question of the thesis. Finally, in section 3.3, the subquestions are presented that will be used to answer the main research question.

3.1. Mission Profiles

As the question of introducing reusability in the design of the Ariane 6 is a rather wide one, this section will narrow down the mission profiles a reusable Ariane 6 could follow.

To begin with, two decisions are made to constrain the problem somewhat:

- Only the reusability of the first stage is considered. The first stage constitutes a significant part of the manufacturing cost and is therefore likely worth reusing. Unlike the upper stage, its altitude and velocity at separation are still within a range where return of the stage would likely not be prohibitively expensive. The boosters may also be worth reusing, especially as these are jettisoned at much lower altitude and velocity, but this is considered to be outside the scope of this research.
- The mission profile of the Ariane 6 is not adjusted. The existing mission profile was selected to meet various requirements for its capabilities. The objective is not to redesign the vehicle, simply to introduce some level of reusability. Any changes made will therefore leave the Ariane 6 mission profile unchanged.

To narrow down the possible mission profiles, a design option tree is used. This can be found in Appendix A. To introduce reusability into the Ariane 6 first stage, the first choice is the degree of reusability. Namely, if the full first stage, or part of it, is reused. Both of these options may be feasible and are therefore kept to be investigated. For the fully reusable cases, a passive re-entry is removed as an option, as with the slender first stage it will be difficult to maintain stability during re-entry. An active re-entry should be a possibility here, though, as the vehicle should have ample space to store the required propellant. Both a down-range landing and a return-to-launch-site mission are feasible, and so these are both kept. Both the downrange landing and the return to launch site options require the vehicle to be guided during its re-entry, as they both need a significant degree of control over the vehicle.

It must then be decided what to do with the vehicle after it has survived re-entry. As using a parachute or parafoil is much more prone to being blown away from the landing site, this option is eliminated for the active re-entries, meaning that the vehicle must glide to the landing site before decelerating with its engine for a powered landing. This does require a small amount of fuel to be carried up to this point. For the landing itself, splashdowns are eliminated due to the physical shock incurred on the vehicle, and the adverse effects salt water has on spacecraft components. Mid-air captures are eliminated as well, as capture by a tower requires incredible precision whilst offering only a modest saving in mass. Capture by helicopter is also discarded, as in the previous research, considering only the engine bay, the loads were at the limits of what the helicopter could sustain, and so the entire core stages would be too heavy. This leaves landings by touchdown, be it on ground or on a floating platform.

For the partially reusable mission, the engine bay will be the section that is returned, as it contains the most valuable components of the first stage. For its re-entry, an active re-entry is seen as infeasible, as the detached engine bay will probably not be able to carry enough fuel to decelerate the stage enough and shield it from the atmosphere. A passive re-entry is possible here. Within passive re-entries, gliding re-entries are eliminated as an option. These require the vehicle to be developed specifically for it, which would interfere with the design of the Ariane 6 too much. Skipping re-entries is eliminated as well, as these require the vehicle to generate more lift than it likely is able to. Ballistic entries, or lifting ballistic entries, are kept as the feasible option. The ballistic entries can be either guided or unguided, and so these are both kept as options.

For these mission profiles, it must also be decided what to do with the vehicle after it has survived re-entry. The unguided mission is likely to incur significant uncertainty in the vehicle's final position, which causes a controlled glide descent to a landing pad to be infeasible. Splashdowns are eliminated here as well, once again due to the physical shock incurred by the vehicle and the adverse effects saltwater has on spacecraft components. This leaves the use of a parachute or parafoil. For these, once again, precise landings on a specific location are not feasible, as once again the vehicle is likely to be blown off-course during its descent. This leaves mid-air captures, where, again, capture by a tower requires incredible precision and so is rejected. This means the only feasible option is capture by helicopter, which is the mission profile presented in previous research. It is worth noting that there are concerns about safety during the capture as well as during flight due to instabilities incurred by carrying the stage and that the loads endured during capture were at the edge of what the helicopter could sustain.

This leaves four possible mission architectures and profiles that are considered feasible ways to introduce reusability in the Ariane 6:

1. **First Stage Return to Launch Site:** Following the first stage separation, the stage will flip and perform a burn to counteract its velocity, returning its trajectory to a landing site near the launch site. The stage then flips again, and once the effects of re-entry begin to present themselves, it performs another burn, slowing the vehicle down and shielding it from the effects of re-entry. The stage then glides through the atmosphere to the landing site before performing a retroburn for a final deceleration, landing on a landing pad in the launch complex in French Guiana, where it will be recovered by a ground crew.

   This mission will require the first stage to be outfitted with RCS thrusters to adjust the attitude before re-entry. For control during the glide through the atmosphere, grid fins can be used, which ESA is already developing with the Themis project. Landing gear is also under development in the Themis project, making the surface landing possible. [19] Finally, this mission will be outfitted with ESA's new Prometheus engine, which was designed for reusability. It crucially has the ability to re-ignite, allowing any of the necessary retroburns to occur.
2. **First Stage Downrange Landing:** Following the first stage separation, the stage will flip and wait until it hits the atmosphere. Once the effects of re-entry begin to present themselves, it performs a burn, slowing the vehicle down and shielding it from the effects of re-entry. The stage then glides through the atmosphere to the landing site before performing a retroburn for a final deceleration, landing on a floating platform in the Atlantic Ocean, from where it will be recovered by a recovery vessel.
   This mission requires the same hardware modifications present in mission profile 1.
3. **Unguided Ballistic Engine Bay:** Following first-stage separation, the engine bay detaches from the first stage. It will deploy a heat shield, which it will use to perform a ballistic re-entry. Afterwards, the heat shield is discarded, and the vehicle will deploy its parachute to slow its descent, followed by a parafoil, allowing it to be captured by a helicopter and delivered to a recovery ves-

sel.
This mission will require the vehicle to be outfitted with RCS thrusters to adjust the attitude before re-entry. A deployable heat shield should be used to not affect the design of the Ariane 6 too much. Conveniently, ESA is already investigating options for this technology. [17] Finally, the vehicle must bring a parachute and parafoil.

4. **Guided Ballistic Engine Bay:** Following the first-stage separation, the engine bay detaches from the first stage. It will deploy a heat shield, which it will use to perform a guided ballistic re-entry. Afterwards, the heat shield is discarded, and the engine bay will glide through the atmosphere to the landing site before performing a retroburn to decelerate. In this way it may be allowed to cross the Atlantic Ocean. If this is indeed possible, the engine bay could land on a landing site in Western Africa or Europe, where it will be recovered by a ground crew.
   This mission will require the engine bay to be outfitted with RCS thrusters to adjust the attitude before re-entry, as well as to modulate its bank angle. As with mission 3, the vehicle must be equipped with a deployable heat shield. For control during the glide through the atmosphere, planar fins will have to be used due to the reduced moment arm of the shorter engine bay. As with missions 1 and 2, landing gear must be included to allow it to safely touch down on the landing pad. Finally, a small fuel tank must be included, as the vehicle does not have access to the entire fuel tank of the first stage for its burns.

Especially option 4 is very promising, as it recovers the most valuable part of the first stage without the issues the unguided option presents, and the vehicle can be recovered closer to where it will be refurbished. Investigating whether this mission profile is feasible, and how worthwhile it is, will be the main focus of the research to be performed. The research will be limited to the re-entry phase of the vehicle, leaving the final approach and landing for potential future research. The vehicle's shape will be the same as that designed during the previous research.

3.2. Research Question

The main research question for the thesis project will be as follows:

Is a transatlantic re-entry a feasible and cost-effective method to recover the Ariane 6 Engine Bay?

Feasible and cost-effective here will be investigated considering the trajectory, guidance, mission architecture and a top-level estimation of the costs. The final segment after the re-entry is beyond the scope of the research, as are considerations on the basis of hardware, structure, thermal protection, aerodynamic stability and controllability, and operations.

3.3. Subquestions

To answer that research question, the work is divided into a number of subquestions. These are:

- *Is it possible to have the engine bay cross the Atlantic during its re-entry?*
- *What useful locations could be reached in such a flight?*
- *How can such a flight be guided?*
- *What is the recommended re-entry trajectory for the engine bay?*
- *How does the performance of the guided engine bay mission profile compare to alternative mission profiles and the non-reusable variant?*

How each of these subquestions will be researched is discussed in chapter 4.

4

Methods & Tools

In this chapter, the methods and tools that were used in the research are discussed. In section 4.1, the software used to determine the lift and drag coefficients of the vehicle is presented. This is followed by section 4.2, where it is discussed how the main simulation for the research project will be set up. Finally, section 4.3 lays out how the simulation will be used to find the optimal solutions for this case. In section 4.4, the link to the GitHub repository with the code and plots produced during the research can be found.

4.1. Lift & Drag Coefficients

To simulate the vehicle, it is necessary to know the lift and drag coefficients throughout the flight. Although the shape of the vehicle is known, namely that from the previous research done into the recovery of the engine bay, the lift and drag coefficients are not, and so these must be determined first. For this, the Free Open Source Tool for Re-entry of Asteroids and Debris (FOSTRAD) will be used. This is a tool that can be run on MATLAB to simulate objects undergoing re-entry. For this research, especially relevant is that the software uses an STL model of the vehicle as input and can then simulate its lift and drag coefficients at various altitudes and Mach numbers. It also has the ability to determine that parts of the vehicle are shadowed from the flow, which is the case for the vehicle being considered. [20]

4.2. TudatPy Simulation

To investigate the research questions, a simulation will be set up. This will be done in the TU Delft Astrodynamics Toolbox (Tudat). This is an open-source toolbox designed by TU Delft specifically for astrodynamic research, implemented in C++, but it can be interfaced with through Python, which is how the simulator will be constructed and run for the research. [41]

The simulation will run from the moment of first stage separation, to the end of the re-entry phase of the trajectory. The simulation will be a 3-DOF simulation, using three spatial degrees of freedom, and parametrising anything that is a function of any rotations, such as the lift vector. The objective is to investigate the feasibility of the trajectory, not an in-depth analysis of the vehicle, and so this 3-DOF model is deemed to be sufficient. The simulation is based on code provided by D. Dirkx and S. Gehly in the course Propagation & Optimisation at TU Delft. This code also dealt with a re-entry trajectory, and so could be easily modified to suit the purposes of this research. In the setup of the simulation, a number of choices have to be made with regards to the environment. For each of these choices, Tudat offers several options, which must be compared to investigate their effects. An overview of these can be found in Table 4.1. [40]

Table 4.1: Choices to be made in the simulation setup, and the options available in Tudat.

Simulation setup choice	Available options in Tudat
Degree and order of central body gravitational model	Point mass, spherical harmonic
Included third bodies, and the degree and order of their gravitational model	various solar system bodies (Sun, Moon, etc.), point mass and spherical harmonic depending on the body
Radiation Pressure	Cannonball, panelled radiation target, do not include
Atmosphere model	exponential atmosphere, US76 standard atmosphere, NRLMSISE-00 atmosphere model
Central body shape model	Spherical, Oblate sphere
Central body rotation model	Default SPICE, simple from SPICE, IERS Earth rotation model
Ephemerides source	Keplerian from SPICE, full ephemerides from SPICE
Propagator	Cowell, Encke, Kepler Elements, Modified Equinoctial Elements, Unified State Model
Integrator	various
integrator tolerance/step size	various

Once the simulation has been set up, it should be verified and validated. The model will be verified by performing unit tests on modules of the simulation. Furthermore, extreme value tests can be performed to determine if the model responds as expected. Finally, the produced outputs can be assessed to see if the simulation provides results that seem realistic. Validation of the model is challenging, as there is no independent test data for the mission to compare the model output to. Instead an analogous mission for which test data is available should be used by entering the conditions of this flight into the model and seeing if the outputs correspond to the data.

The simulation can then be used to answer the first subquestions. First, it will be run in a simple setup to see if it is feasible for the engine bay to cross the Atlantic. This will be followed by an iteration on various inputs to the simulation to see what range can be achieved during the re-entry. This will allow a destination to be selected where it is desired to deliver the engine bay. These initial inputs can then be used to perform an optimisation.

4.3. Optimisation

Selecting an optimum is not simple, as there are multiple objectives, inputs and constraints the design must take into account. This means the problem becomes a Multi-disciplinary Design Optimisation (MDO). A number of variables will be chosen as inputs to the problem, which the optimisation scheme will vary to produce outputs. A number of objectives are then selected, for which the optimisation scheme will attempt to maximise performance. The problem will be constrained by several factors that must not be exceeded, for re-entry trajectories these are commonly the maximum allowable g-load, heat load and heat flux. A penalty function will be created for these, penalising the performance of an input set if it exceeds the constraint levels, without discarding the information that set produces.

Each of the inputs is iterated across a range one at a time to investigate the envelope. Then, the MDO can take place, optimising initial designs generated with a Monte-Carlo analysis. The PyGMO library will be used to perform the optimisation, which contains a number of optimisers to choose from. [24] In the optimisation, a variable will be set up called fitness, which will exist as an array containing a scaled factor of its performance on each objective. The optimiser will attempt to reduce the fitness value in each of these array entries, resulting in a set of Pareto-optimal solutions. These are solutions where their performance on one objective cannot be improved without sacrificing their performance on another objective. The Pareto-optimal solutions together form the Pareto-front, the front on which the Pareto-optimal solutions can be found. The code used for this is based on code developed in collaboration with J. Dijkstra and J. Pedra during the course Propagation & Optimisation at TU Delft. [12]

Once the front of Pareto-optima has been found, a sensitivity analysis should be performed, to check the robustness of the solution. The relevant conditions of the vehicle will be varied in a range around the conditions used in the optimisation, and then the propagation is performed. This way it can be seen how much the results differ based on changes in the initial condition, which is an important step in determining how reliable the produced results are.

The set of optimal solutions will then be compared to the other reusable cases and the non-reusable baseline case. The various options will be compared in a trade-off, which will then be used to guide the conclusions on whether the proposed mission, to use a guided transatlantic re-entry to recover the Ariane 6 engine bay, is worthwhile.

4.4. GitHub Link

Throughout this research project, many pieces of code and plots were produced. These can be found in the following GitHub repository:

MSc Thesis GitHub Repository

5

Work Done

This chapter discusses the technical work done to produce the results in chapter 6. First, section 5.1 states some specifics on the mission vehicle, setting its initial state and estimating its mass. Then, section 5.2 shows how the C_L and C_D values of the vehicle were determined. Next, section 5.3 lays out how the settings used for the simulation were selected. In section 5.4, it is explained how the guidance algorithm was integrated into the simulation. After this, section 5.5 states how the verification and validation were performed. With this, the simulation is ready to use. Then, section 5.6 describes the initial investigations that were done with this completed simulation. The possible target locations for the vehicle are described in section 5.7. With these determined, it was time to perform the optimisation. In section 5.8 the inputs, objectives and constraints of the optimisation problem are laid out. Afterwards, section 5.9 shows the Monte Carlo analysis that was performed to gain insight into the problem. Finally, section 5.10 describes the optimisation itself.

5.1. Mission Vehicle

This section presents specifics of the mission vehicle. First, subsection 5.1.1 presents the initial state of the vehicle and its load limits. Then, subsection 5.1.2 shows how the vehicle mass was estimated.

5.1.1. Initial State and Load Limits

Before any other work can be done, some specifics of the mission under consideration should be established. The guided ballistic engine bay has been selected as the main mission profile to be researched. As there is a wide space into which the vehicle can deliver its payload, the decision was made to constrain this. The mission under consideration will be a LEO mission, with the launch vehicle departing eastward from French Guiana. Shortly to the east of the launch site, the first stage will burn out and separate, from where the upper stage will bring the payload to its desired orbit. The first stage will then disconnect the engine bay, which will return as described in section 3.1. The initial state of the engine bay results from an estimate of the trajectory the launch vehicle travelled on and can be found in Table 5.1.

Table 5.1: The estimated initial state of the engine bay after it has disconnected from the first stage.

Parameter	Value	unit
Altitude	157.7	km
Latitude	5.3	degrees north
Longitude	-50	degrees west
Velocity	6.93	km/s
Heading angle	95.25	deg
Flight path angle	-0.8	deg

The three loads to watch during re-entry are the g-load, heat flux and heat load. As the vehicle will not carry people, the g-load is set to 10 g's. For the heat flux constraint, a value of 1 [MW/m^2] is used, as this is a value that can be reasonably expected to be reached during re-entry by a vehicle re-entering at 7.5 [km/s], which is faster than the vehicle under consideration is likely to travel, making this a conservative estimate. [23] Similarly, a value of 200 [MJ/m^2] is taken as the total heat load constraint, which is rather conservative for the re-entry. [22]

5.1.2. Mass Estimation

As stated in chapter 3, the final approach is not being considered for this research. The cut-off point is set at 30 [km], after which the vehicle will flip engine-first, jettison its heat shield, and guide its descent to the landing zone using the control fins. The vehicle will, however, have to carry with it the necessary elements to perform this phase of the mission, and this mass contribution needs to be taken into account. In the previous research project on this subject, a vehicle was developed for the unguided ballistic entry mission, which therefore already included mass estimates for the separation mechanism, heat shield and RCS thrusters, which are relevant here, and for a parachute and parafoil, which are not. Including the useful parts and leaving out those that are not necessary results in a vehicle with a mass of 8667 [kg]. [11] This is not the complete picture, however, as the vehicle will have a few additional mass contributions:

- propellant for deceleration burn
- tanks for propellant
- control fins
- manoeuvring thruster propellant
- manoeuvring thruster propellant tanks
- landing gear

Without specific data on the required control authority of the post-re-entry phase, state of the vehicle after disconnecting, guidance bank angle profile and landing loads, it is challenging to determine accurate masses for the control fins, RCS propellant and landing gear.

As is described in section 3.1, planar fins are used to control the vehicle after re-entry. It was chosen to use four of these fins, each set at a 90 [deg] offset around the vehicle. As it was not entirely certain how much control authority the vehicle would need, a very rough estimate of the control fin mass was made. To be able to stow the fins flush against the body of the vehicle during re-entry, it was estimated that the fins could be approximately 0.8 [m] across and 1.4 [m] long and likely would be approximately 5 [cm] thick. For a fin made of aluminium, with a density of 2700 [kg/m^3], this results in a mass of roughly 150 [kg] per fin. Including an attachment and actuation mechanism of 50 [kg], this leads to an estimate of 200 [kg] per fin, for a total of 600 [kg].

The vehicle's final dry mass at landing would be roughly 8000 [kg]. The Falcon 9 first stage, which has been demonstrated to be able to land many times at this point, weighs 25600 [kg] dry. Assuming the landing procedure for these two vehicles is similar, that they have the same number of legs, and that landing loads scale linearly with the mass of the vehicle at landing, which will be close to the dry mass, the landing legs for the engine bay should be 3.2 times lighter than those of the Falcon 9. The legs of the Falcon 9 weigh 600 [kg] each, with four legs, which means the legs of this vehicle should weigh in at 187.5 [kg] each. To account for the uncertainty of this estimate, the mass of the legs is set to 200 [kg] each, totalling 800 [kg].

The vehicle's bank angle control will follow an algorithm based on the Apollo guidance algorithm. The RCS thruster propellant mass will therefore be estimated based on the Apollo command module re-entry. The command module weighs 5560 [kg] and carries 120 [kg] of propellant for its RCS system, which accounts for 2.2% of its mass. [1] As the engine bay will have a similar RCS propellant requirement for the guidance bank angle commands, this figure will be used. With a dry mass of approximately

8000 [kg], the required propellant mass would be 172.6 [kg]. It is further estimated that the tanks for propellant mass would weigh about 10% of the propellant mass, coming to 189.9 [kg]. This figure is rounded up to 200 [kg], as this is once again a rather rough estimate.

These are very rough estimates and could do with reassessment with better data in the future. Nonetheless, these are the values that will be assumed going forward. This brings the mass of the vehicle, excluding propellant mass, to 10267 [kg]. Now, only the propellant mass must be estimated. To do this, using the initial outcomes of the TudatPy simulation, a figure for the velocity at 30 [km] altitude was reached. It was then simulated with a simple forward Euler method simulation, assuming an exponential atmosphere and what the velocity would be at 2 [km] altitude. Here it was assumed that the vehicle behaves as a blunt cylinder, with a drag coefficient of 1.15. As in reality the vehicle will not be falling straight down and is using control fins to guide its descent, it will experience more drag, and so the result here will overestimate the final velocity. The rocket equation, seen in Equation 5.1, is then used to calculate the required propellant mass for the final deceleration to standstill. This propellant mass, with a margin of 10% to account for tankage, is then fed into the calculation for the velocity at 2 [km] altitude and used to determine a new propellant mass value. This is iterated until the wet mass of the vehicle differs by less than 1%. This results in a required propellant mass of 317 [kg], making the total wet mass of the vehicle at the start of the re-entry 10648.25 [kg].

$$m_p = m_{dry}\left(e^{\frac{\Delta V}{I_{sp} g_0}} - 1\right) \tag{5.1}$$

5.2. FOSTRAD

To be able to simulate the vehicle, first FOSTRAD was used to calculate the lift and drag coefficients of the vehicle. To do this, an STL model of the vehicle is created using the aeroshell dimensions developed during the previous research, which can be found in Figure 5.1. The remainder of the vehicle is modelled as a cylinder, 5.4 [m] in diameter and 5.134 [m] long, touching the back end of the heat shield. The control fins, landing legs, RCS thrusters and engine are not included for the model used in FOSTRAD, as these are located in areas where they will be shadowed from the flow by the aeroshell, as can be seen in Figure 5.2, and so were expected to have minimal effect.

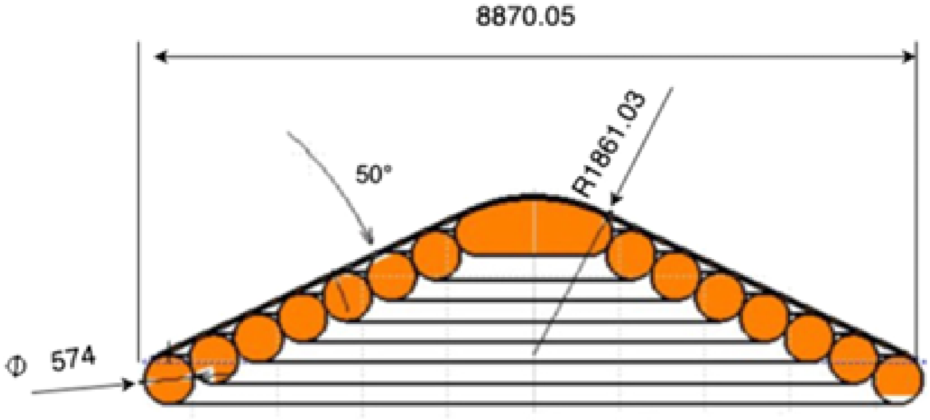

Figure 5.1: The aeroshell shape developed during the previous research, which is used for this research as well.

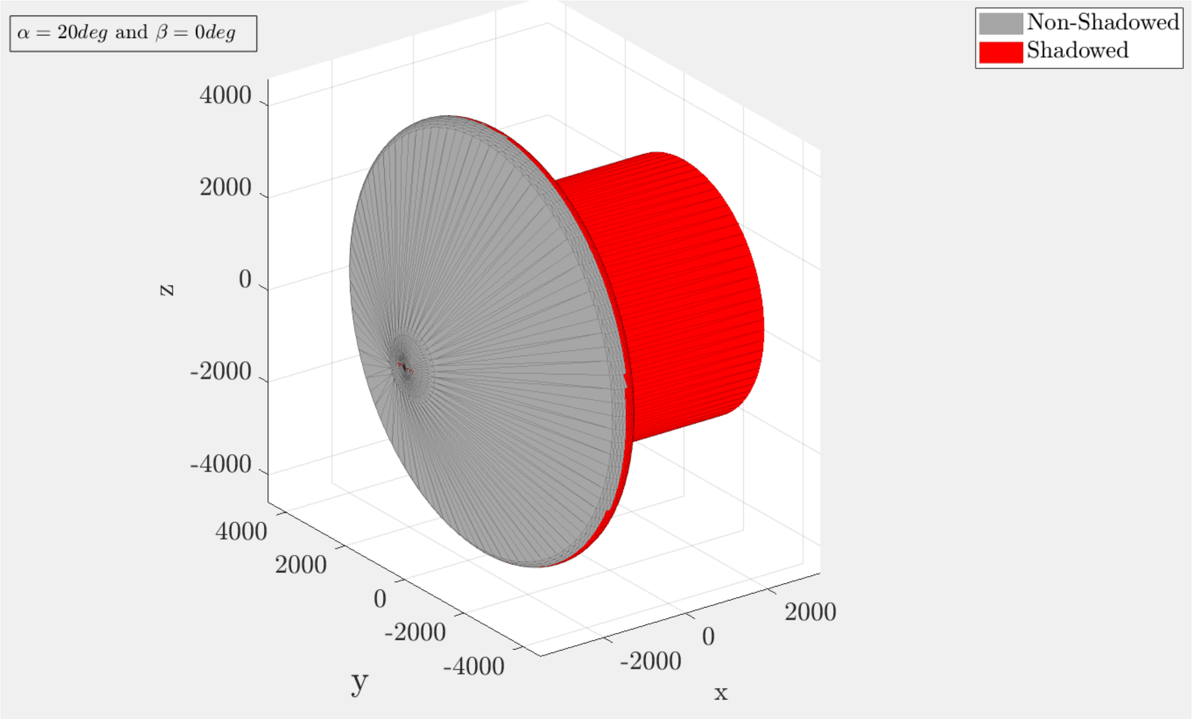

Figure 5.2: A window with one of the initial checks performed by FOSTRAD, where it checks the areas of the vehicle shadowed from the flow.

As the lift and drag coefficients of the vehicle can change throughout the re-entry, depending on the flight conditions, the calculation with FOSTRAD should be run for various parameters that might affect these values. The vehicle designed in the previous research was made to re-enter at a constant angle of attack of 20 [deg], and so the relevant parameters to vary were determined as the Mach number and the altitude. Lift and drag coefficients tend to change quickly with rising Mach numbers, but do not change dramatically at higher numbers. For this reason, the coefficients were evaluated at Mach numbers of 1, 2, 3, 4, 5, 7, 9, 12, 15, 20 and 30. The altitude was varied at increments of 5 [km] from 30 [km] to 200 [km]. The results can be found in Figure 5.3, Figure 5.4, Figure 5.5, Figure 5.6, Figure 5.7 and Figure 5.8.

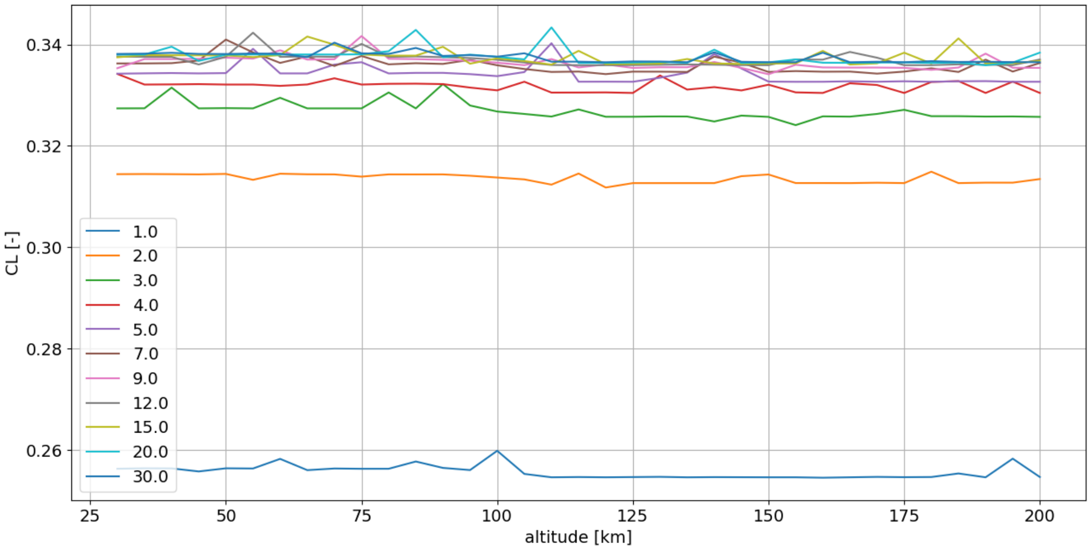

Figure 5.3: The lift coefficient as a function of altitude, for various Mach numbers, as produced by FOSTRAD.

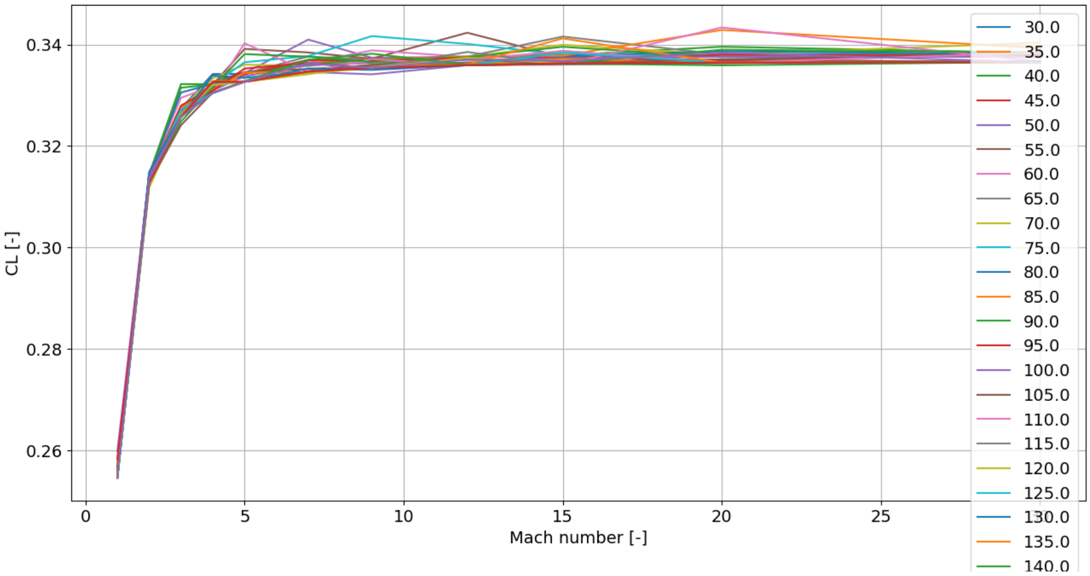

Figure 5.4: The lift coefficient as a function of Mach number, for various altitudes, as produced by FOSTRAD.

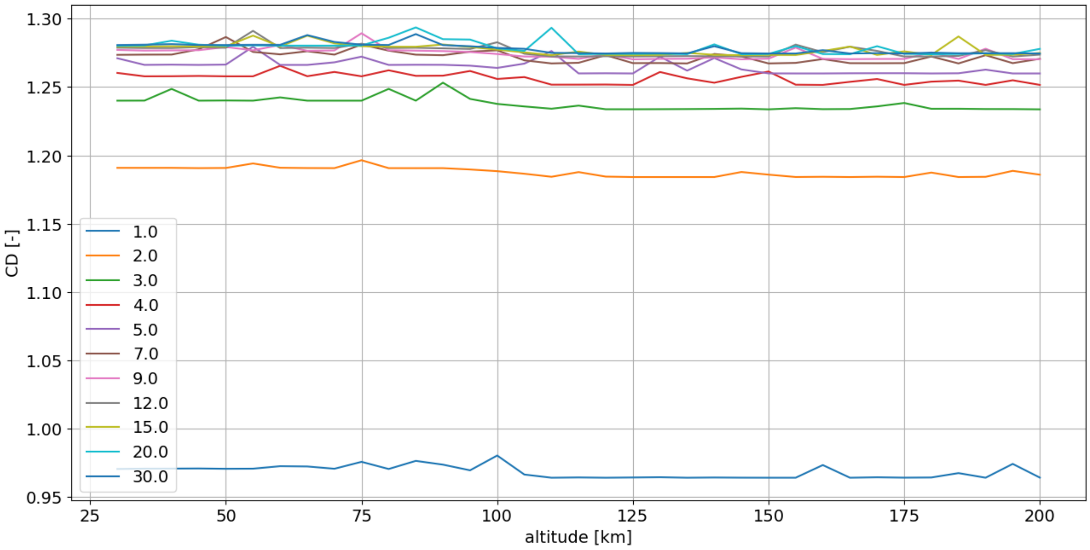

Figure 5.5: The drag coefficient as a function of altitude, for various Mach numbers, as produced by FOSTRAD.

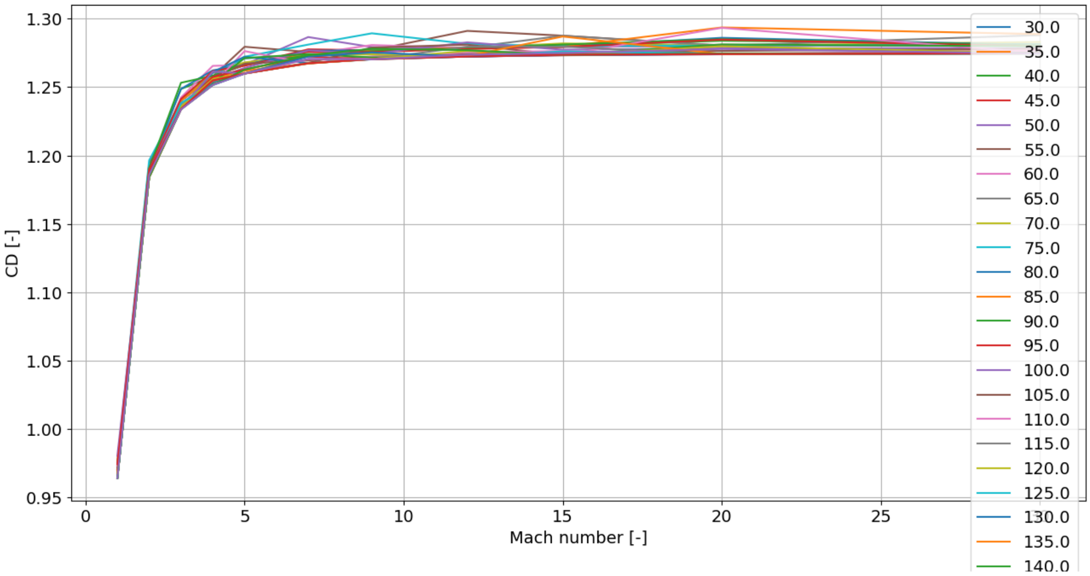

Figure 5.6: The drag coefficient as a function of Mach number, for various altitudes, as produced by FOSTRAD.

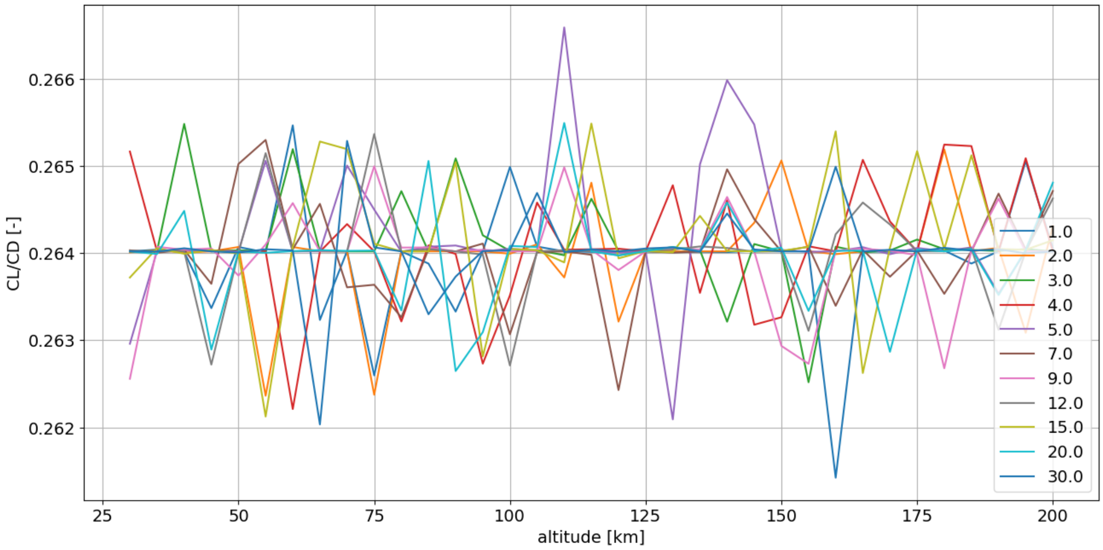

Figure 5.7: The lift-to-drag ratio as a function of altitude, for various Mach numbers, as produced by FOSTRAD.

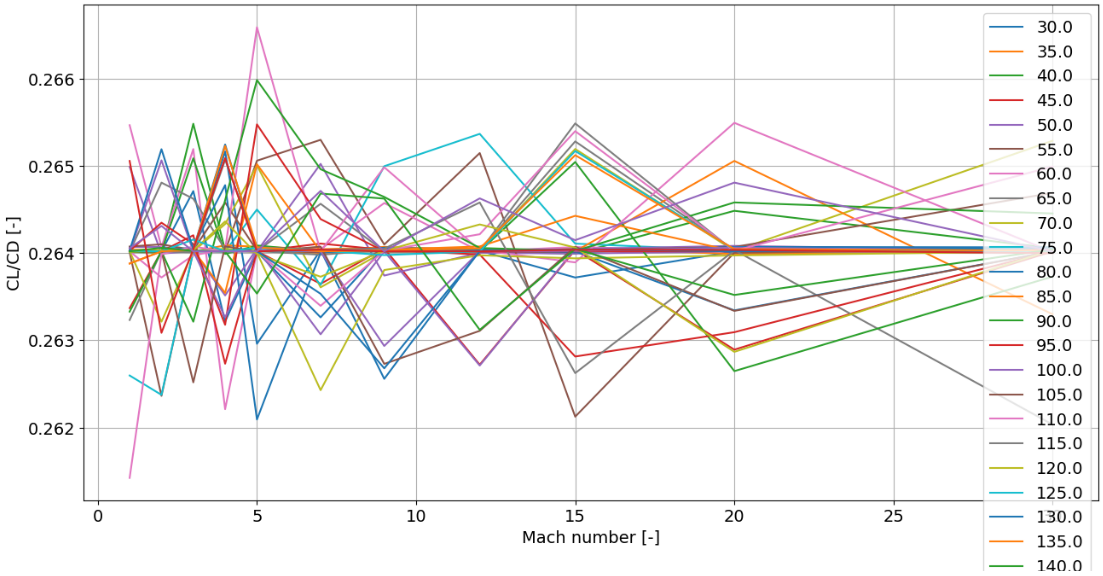

Figure 5.8: The lift-to-drag ratio as a function of Mach number, for various altitudes, as produced by FOSTRAD.

As can be seen in Figure 5.3 and Figure 5.5, the aerodynamic coefficients do not differ significantly depending on the altitude, as although they have some uncertainty, they remain close to the same value throughout. Due to these uncertainties, however, the altitude is kept as a dependent variable based on which to select a lift and drag coefficient. In Figure 5.4 and Figure 5.6, the expected behaviour of the coefficients with respect to the Mach number can be found: As the Mach number rises, they change quickly, after which they remain close to the same value at higher Mach numbers. The lift-to-drag ratio is plotted in Figure 5.7 and Figure 5.8, where it can be seen that this ratio does not change greatly throughout both ranges of assessed values, differing at most 0.0026 from the average value of 0.264. The values produced by this analysis are used in the TudatPy simulation, having the vehicle interpolate a lift and drag coefficient value based on its current altitude and Mach number.

5.3. Simulation Setup

With values for the mass and lift and drag coefficients of the vehicle determined, the TudatPy simulation can be set up. To do this, decisions must be made on the simulation and environment settings found in Table 4.1. To do this, a desired simulation accuracy must be defined. At the end of the re-entry, the vehicle will flip and control its return to a landing pad. It has been assumed that the size of this landing pad is 10 x 10 [m], and so to make sure the vehicle will accurately travel to the landing pad, the physical model accuracy should be better than 10 [m]. To ensure uncertainties are predominantly due to the selected physical model and not numerical in nature, the numerical model accuracy is set to one order of magnitude less than the physical model requirement and so is set to 1 [m]. In subsection 5.3.1 the process used to select the numerical model choices of the simulation is described. Then, subsection 5.3.2 shows how the physical model choices of the simulation were selected.

5.3.1. Numerical Model Selection

The first selection to be made is the integrator, propagator and timestep, which are the source of the numerical uncertainties and so will use the lower 1 [m] accuracy value. To be able to compare these, data of the mission trajectory would be desired, but naturally this is not available. Instead, a baseline simulation is generated, using an RKF5(6) integrator, a Cowell propagator and various time steps.

Numerical errors consist of two types of errors: truncation errors and rounding errors. The former are due to the discretisation of a non-discrete value, such as using discrete time steps. This error will be reduced with smaller time steps. Rounding errors are due to the representation of floating-point values in a computer. At a certain time step size, the truncation error will be reduced sufficiently that the rounding errors take over, but until that time, the truncation errors are dominant. At each time step, the error between the simulation using that time step and one with double that time step is used to determine the degree of accuracy of the higher time step. This can be done because truncation errors, which vary on higher orders of the time step, are assumed to be dominant. The baseline simulation should then be selected as the simulation with the time step with the smallest errors, where truncation errors are still dominant. The plot produced can be found in Figure 5.9. Plotted on a log-log scale, the regions where each of the error types is dominant can be identified: Truncation errors are dominant in the region where the maximum error decreases linearly in the plot with a decreasing time step value, rounding errors will likely somewhat level off or even increase with a further decrease in time step value. In this figure, it can be seen that truncation errors remain dominant for time steps of 0.1 [s] and greater and that below this value rounding errors become dominant. As this is a good measure of the accuracy of a simulation using twice that time step value, the baseline simulation will be using a value of 0.2 [s] as its time step.

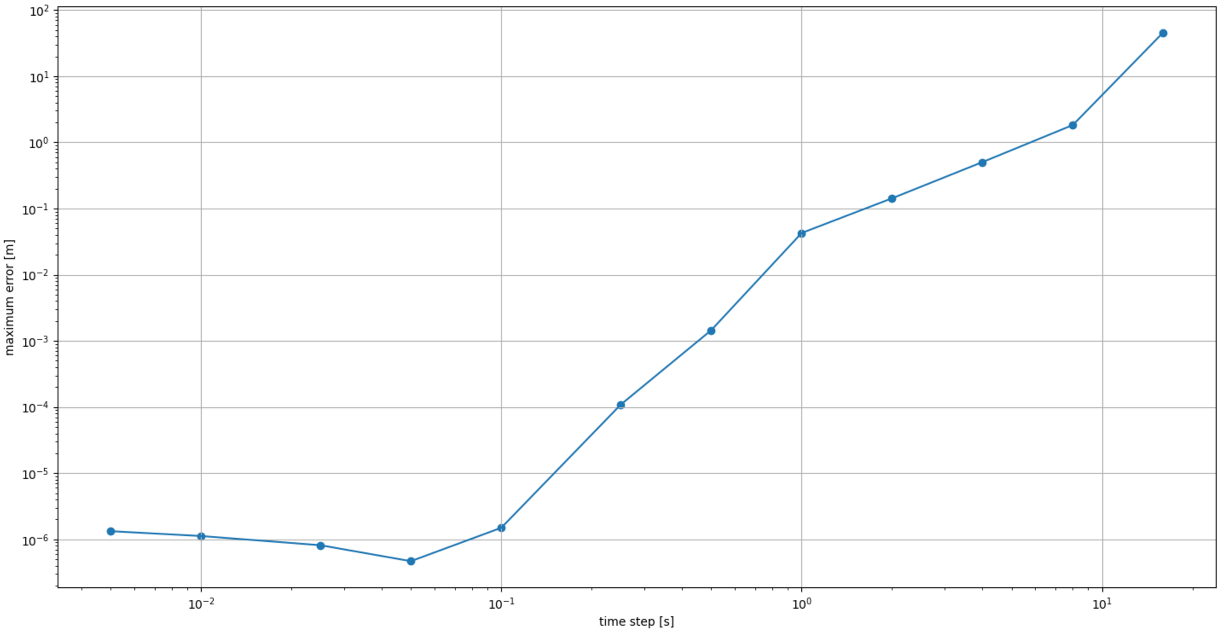

Figure 5.9: The maximum error between a simulation and a simulation using double its time step for various time step values.

With the baseline simulation generated, which is seen as the closest one can possibly get to the true value at this point, various integrators and propagators can be assessed. Seven propagators were selected for this:

- Cowell
- Encke
- Kepler
- MEE
- USM6
- USM7
- USMEM

Each of these propagators can be combined with an integrator, of which four are selected for this analysis:

- RK4(5)
- RK5(6)
- RKDP7(8)
- RKF12(10)

Finally, various time step values can be assessed for each combination. The integrators can either be of a fixed time-step size, in which case a time step value is used, or can be variable-time integrators, in which case a tolerance value is set which the integrator uses to determine the step sizes. For the fixed step versions, the time steps assessed are 0.5, 1.0, 2.0, 4.0, 8.0 and 16.0 [s], and for the variable time step versions, the tolerance values assessed are: 10^{-6}, 10^{-7}, 10^{-8}, 10^{-9}, 10^{-10}, 10^{-11} and 10^{-12}. For each of these combinations of propagator, integrator and step size or tolerance, the maximum error with respect to the baseline simulation is determined, along with the number of function evaluations. The goal here is to find the combination with the fewest function evaluations that still achieves errors less than the 1 [m] numerical simulation requirement. Figures showing the performance of each of the combinations can be found in Figure 5.10, Figure 5.11, Figure 5.12, Figure 5.13, Figure 5.14, Figure 5.15 and Figure 5.16.

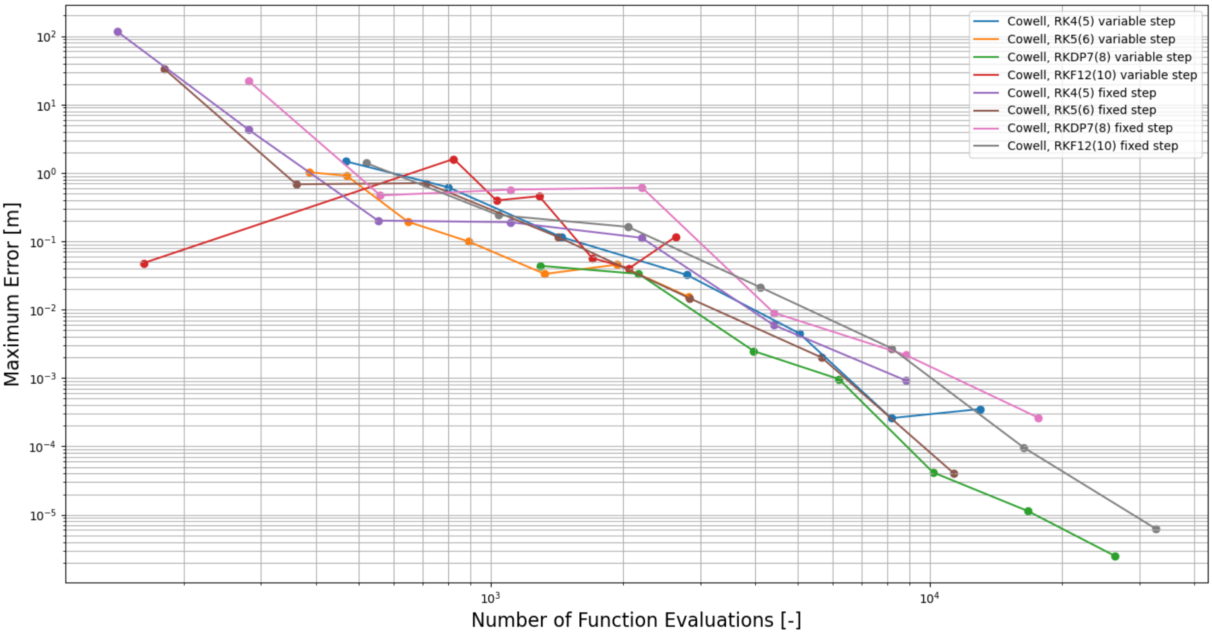

Figure 5.10: The maximum error and number of function evaluations for the Cowell propagator for the selected integrators and time steps or tolerances.

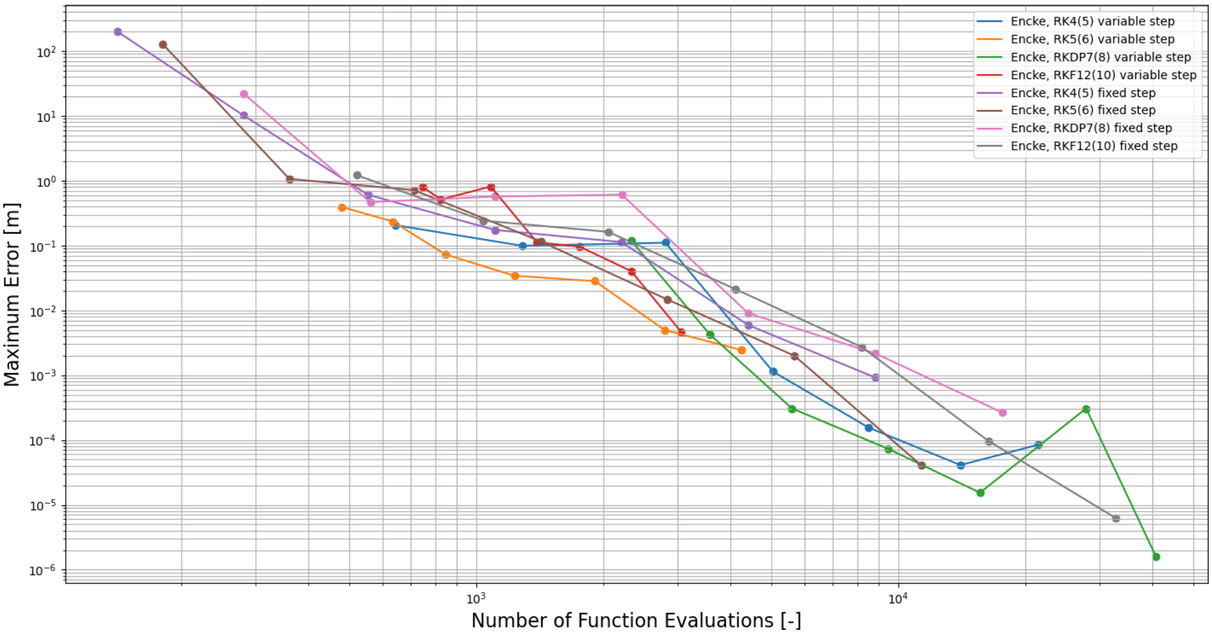

Figure 5.11: The maximum error and number of function evaluations for the Encke propagator for the selected integrators and time steps or tolerances.

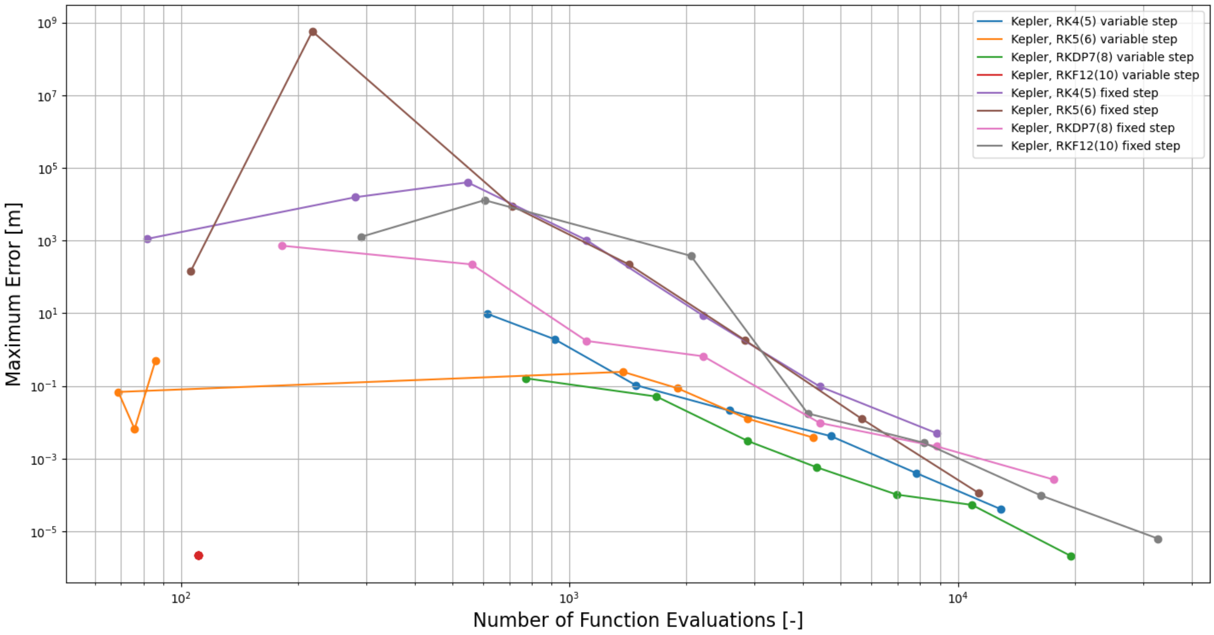

Figure 5.12: The maximum error and number of function evaluations for the Kepler propagator for the selected integrators and time steps or tolerances.

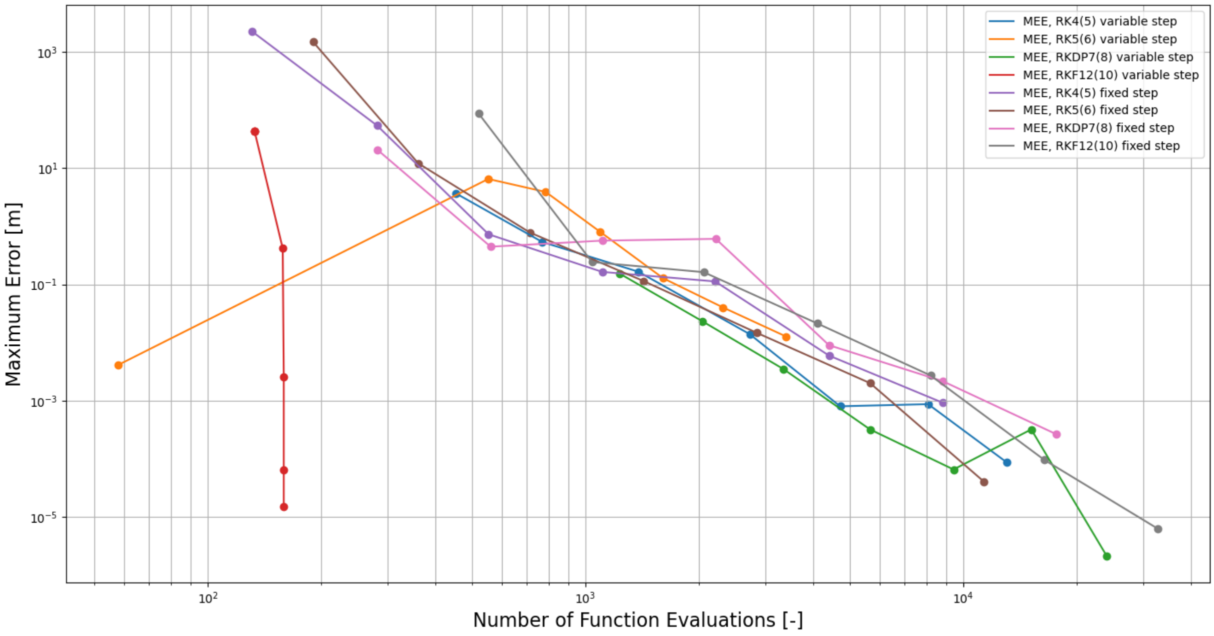

Figure 5.13: The maximum error and number of function evaluations for the MEE propagator for the selected integrators and time steps or tolerances.

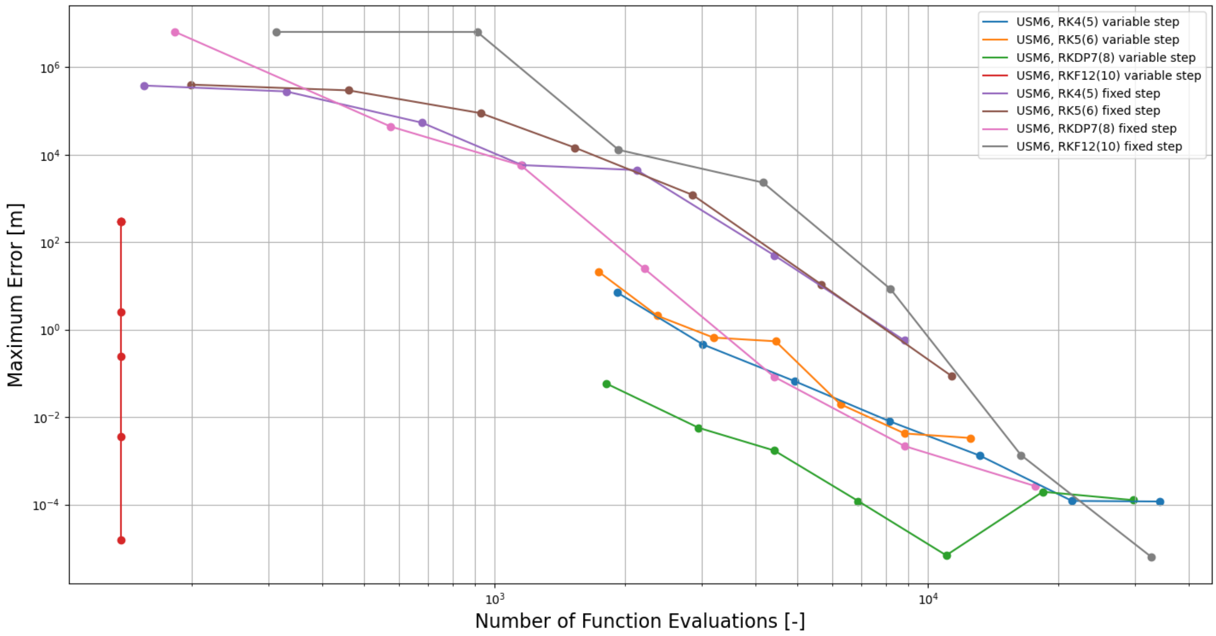

Figure 5.14: The maximum error and number of function evaluations for the USM6 propagator for the selected integrators and time steps or tolerances.

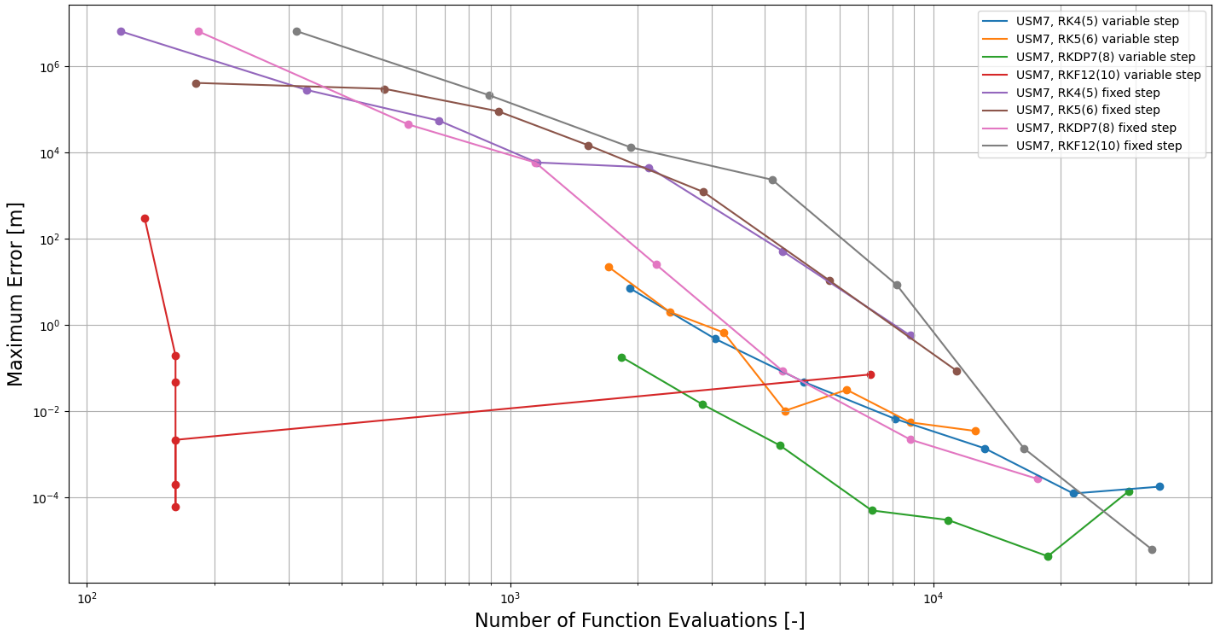

Figure 5.15: The maximum error and number of function evaluations for the USM7 propagator for the selected integrators and time steps or tolerances.

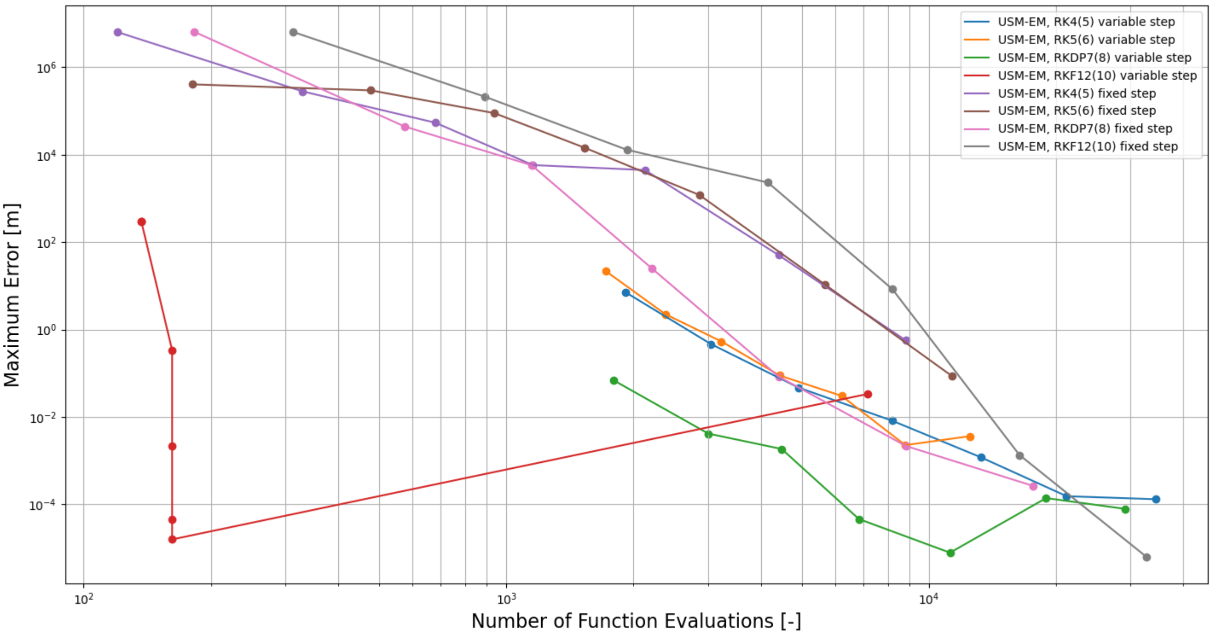

Figure 5.16: The maximum error and number of function evaluations for the USMEM propagator for the selected integrators and time steps or tolerances.

From these images, a combination can be selected. Immediately, a number of options are notable, such as the RKF12(10) variable step options for the MEE, USM6, USM7 and USMEM propagators, or the MEE RK5(6) variable step with a tolerance of $10-6$. These options show very good accuracy for very few function evaluations. Upon investigating these options, however, they behaved like the simulation in Figure 5.17. The number of function evaluations is short because the simulation did not successfully complete, and the maximum error is low because the error was never given the opportunity to diverge greatly. Once the unsuccessful simulations were removed, the combination that provided the least number of function evaluations while remaining within the 1 [m] accuracy limit is the Cowell propagator using an RK5(6) fixed-step integrator with a step size of 16 [s]. Upon investigating this simulation, it seemed like it behaved as expected, and so this option was used moving forward.

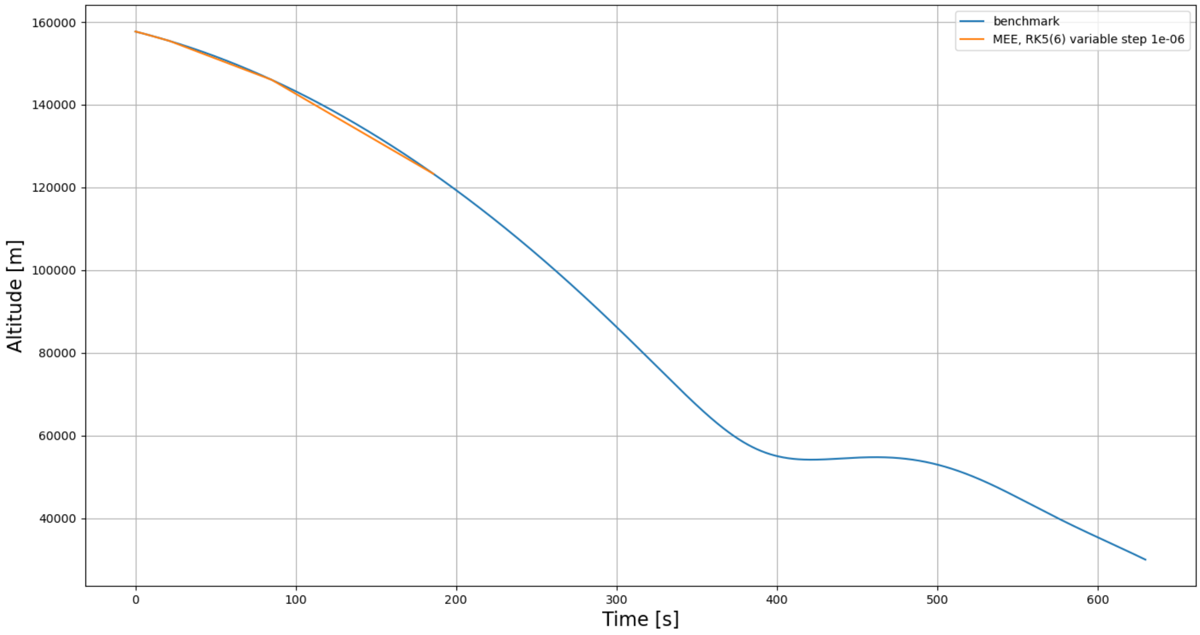

Figure 5.17: The simulation made with the MEE propagator, an RK5(6) integrator with variable step size with tolerance 10^{-6}, compared to the baseline benchmark simulation.

5.3.2. Physical Model Selection

With the settings determining the numerical error determined, it is time to look at the settings for the physical environment. Here the goal is to achieve a position accuracy of 10 [m]. As it is challenging to consider how each setting might interact with others, each individual model choice will seek to achieve an accuracy of 1 [m] so that all together they may achieve the model accuracy requirement. It is also relevant to quantify how much time a more complex option adds to the simulation, as eventually a great number of simulations must be run, meaning even small increases in run time can cause significant delays. First, the third bodies are considered. Five bodies were expected to have a potential impact: The Moon, Mars and Venus due to their relative proximity, The Sun and Jupiter due to their mass. In each case, a simulation including the body is compared to a simulation that includes no third bodies to get a measure of what the effect is of including the body. Here it was found that Jupiter, Mars and Venus did not have a significant effect and so would be safe to not include them. The Moon and the Sun have a more significant effect, and so the impact of including both was assessed as well. Here it appeared that together they caused a maximum position error of around 1 [m], which was considered significant enough to include in the simulation. As the impact was not very great, it was assumed a point mass representation for both of these would be sufficient. Including both these bodies adds 0.0051 [s] to the simulation time.

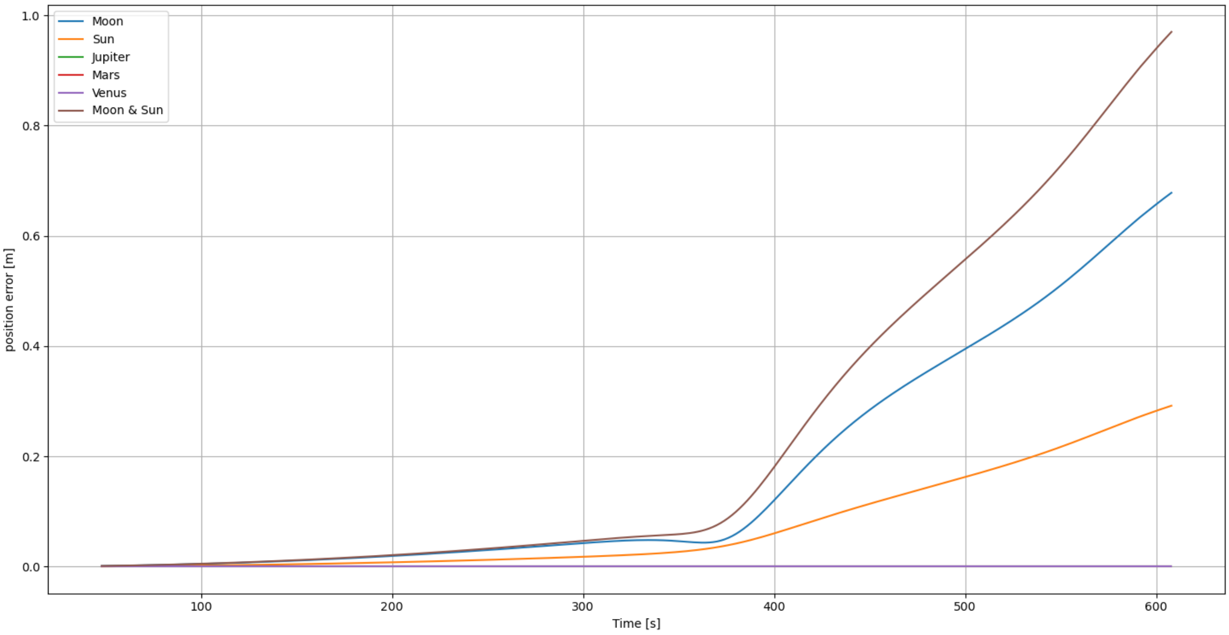

Figure 5.18: The position error over time of simulations including different third bodies with respect to a simulation without third bodies.

One setting of great relevance to the simulation is the atmospheric model. The default option is the US76 atmosphere, but a more complex option exists in the NRLMSISE-00 atmosphere model, and a simpler option is a predefined exponential atmosphere. These are compared in Figure 5.19. As anticipated, for this trajectory the atmospheric model has a very significant effect, with several kilometres in position error between each option. This means the NRLMSISE-00 model, the closest to reality of these models, is a necessity to achieve an accurate simulation. Using this option adds 0.0035 [s] to the simulation time.

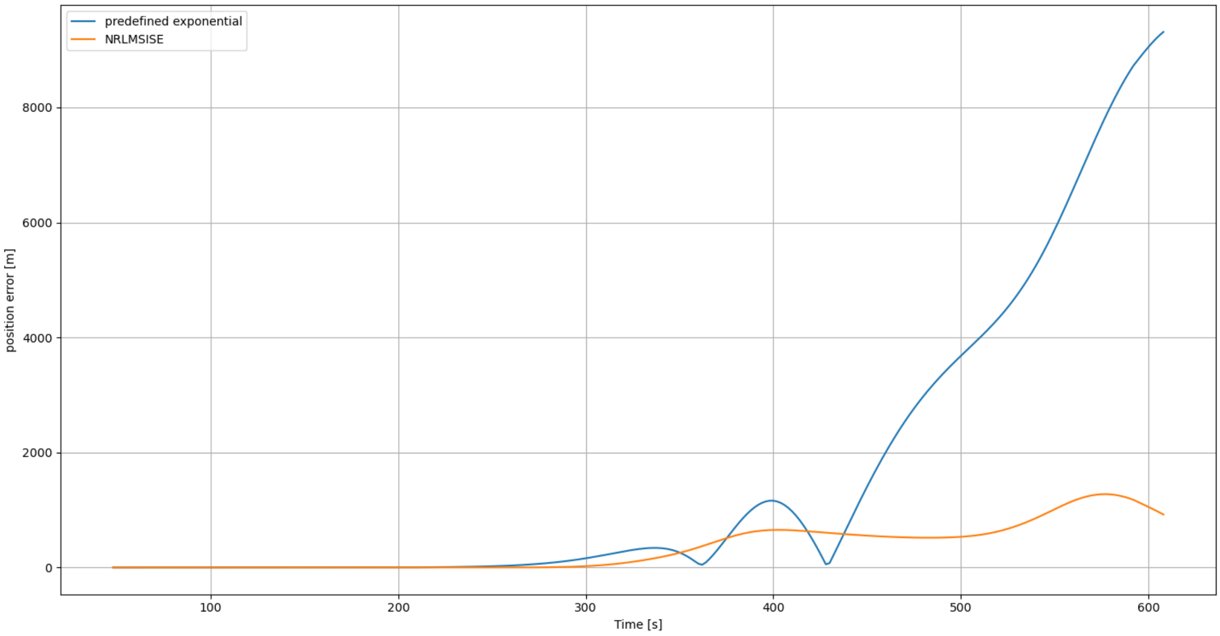

Figure 5.19: The position error over time of simulations including different atmospheric models.

Next up, the shape of the central body is considered. By default, this is modelled as a sphere, which is expected to cause quite a difference with respect to modelling the Earth as an oblate sphere. This is confirmed in Figure 5.20, where it can be seen there is a massive difference between the two. An oblate sphere is chosen going forward, which does not add any significant time to the simulation.

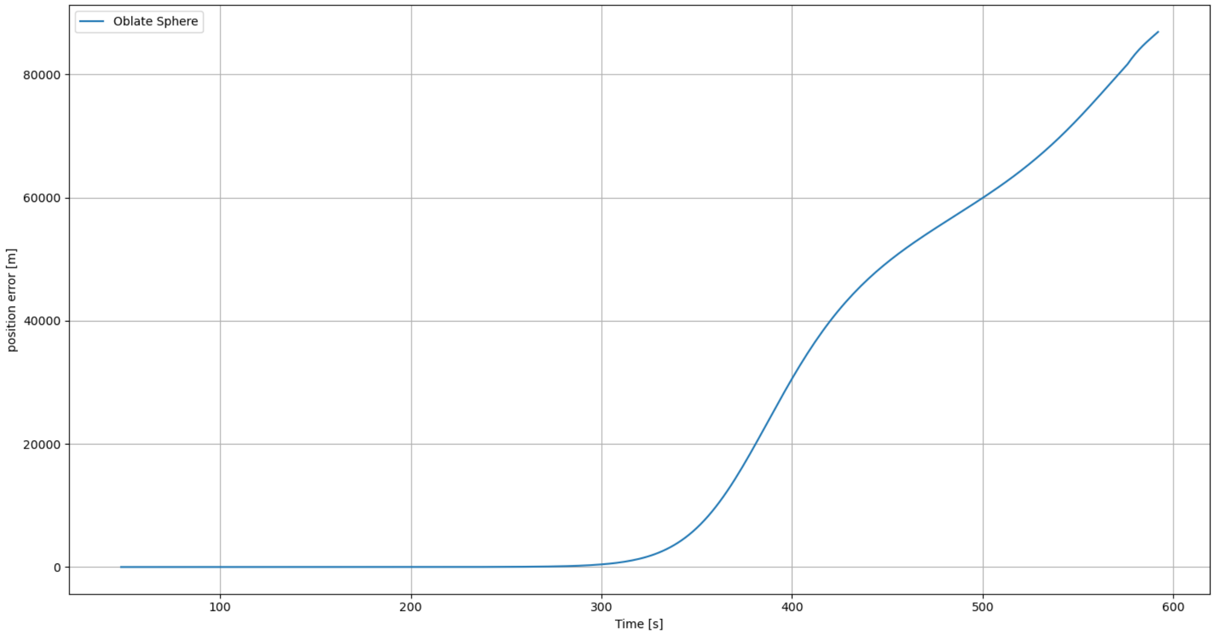

Figure 5.20: The position error over time of the simulation with Earth as an oblate sphere compared to a sphere.

The next option to consider was the spherical harmonic degree and order of the gravity field of Earth. Here an issue was encountered: upon assessing the position error of using a certain degree and order always encountered errors at least in the order of metres, no matter how many orders and degrees were used. Furthermore, using greater orders and degrees did not always improve the level of accuracy. This is believed to be because the differing spherical harmonics cause a slight deviation during the pre-entry phase of the trajectory, causing the vehicle to hit the upper atmosphere at differing velocities, positions and times. These errors then propagate out throughout the re-entry, causing differences between each of these, no matter the accuracy used. Instead of the normal method of analysis, it was considered that beyond the first few orders and degrees, the physical nature of the re-entry does not differ significantly, as the greater the degree and order, the less effect they have. A simulation was run with a degree and order of 15, which was deemed to be suitably close to the real value, and then simulations were run without the atmospheric forces acting on the vehicle, in an attempt to isolate the effect of the spherical harmonics themselves. The results can be found in Figure 5.21, from where it seems that without the atmosphere, the errors are far smaller, and behaved more as would be expected. Using a degree and order of 6 causes a maximum error of roughly 2.5 [m], which is significant but acceptable and does not add a significant amount of time to the simulation.

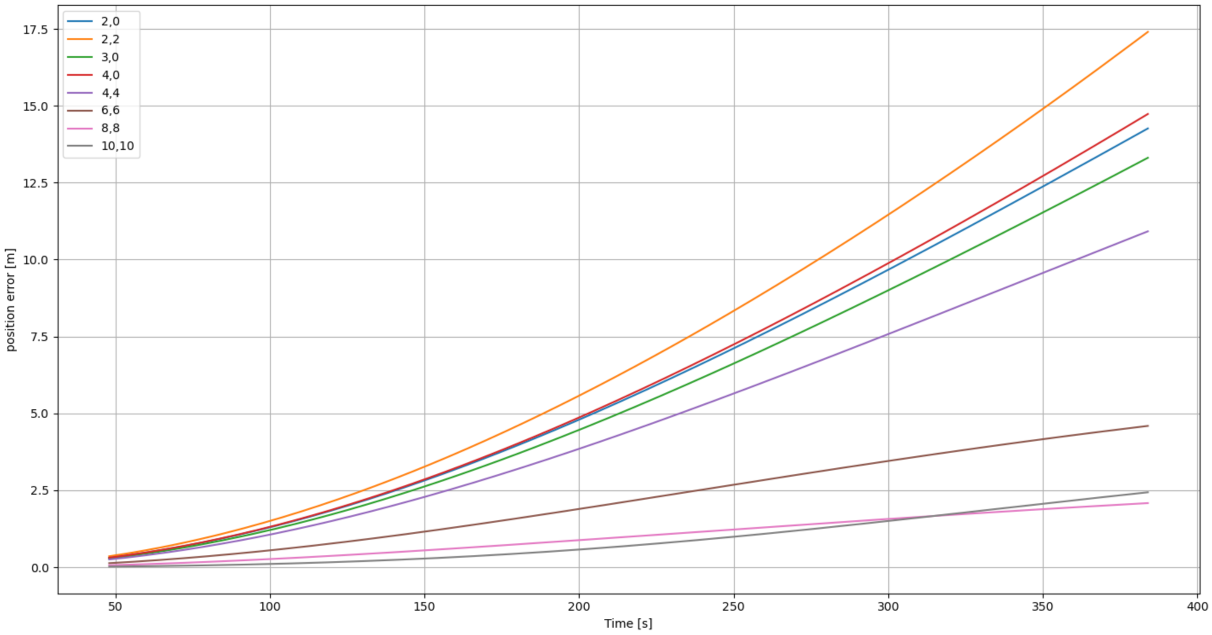

Figure 5.21: The position error over time of simulations with several degrees and orders of gravity field spherical harmonics, with respect to those of a simulation using degree and order 15. The atmospheric forces were disabled for this analysis in an attempt to isolate the effect of this setting.

The next setting to analyse is the solar radiation pressure. Here cannonball radiation pressure is considered, using an estimated radiation pressure coefficient of 1.3. Furthermore, to assess what the maximal effect could be, a simulation is run with a radiation pressure coefficient of 2.0, which would be a body that perfectly reflects all incoming radiation and so experiences the maximum possible solar radiation pressure. As can be found in Figure 5.22, the effect of radiation pressure is very limited, even for the worst-case scenario of a C_p of 2.0, and so it can be safely neglected.

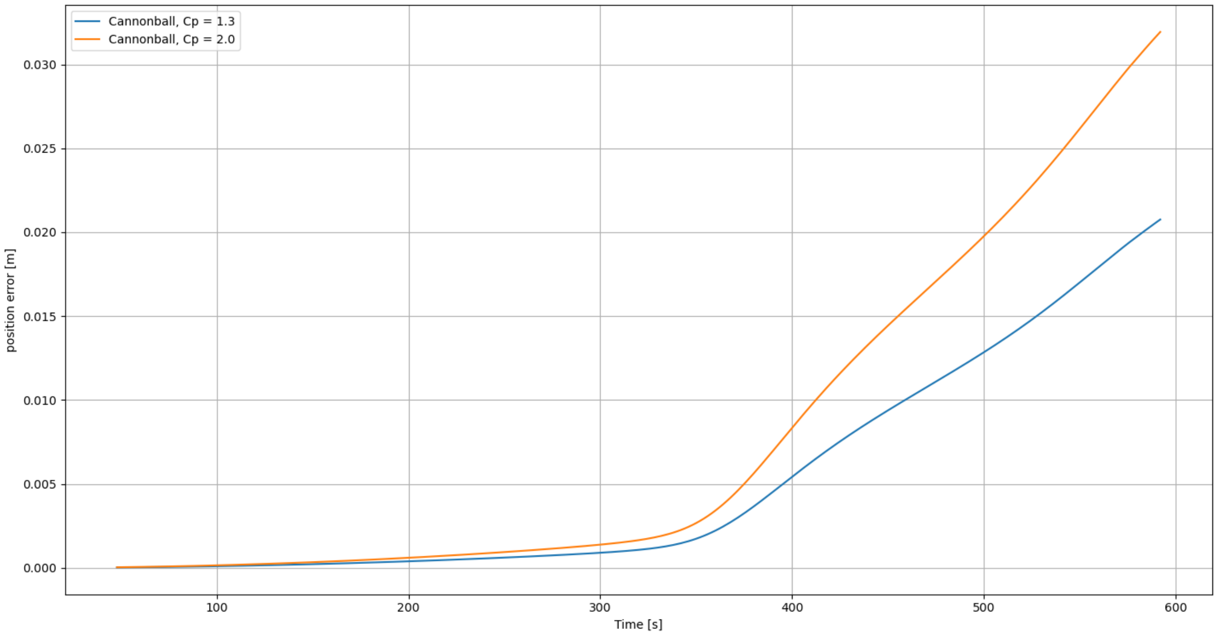

Figure 5.22: The position error over time of simulations with cannonball radiation pressure, using a C_p of 1.3 and 2.0, compared to a simulation not including it.

A model choice that could potentially save some time is the ephemerides source. Using Keplerian orbits to determine the position of bodies over time, rather than being extracted from SPICE. In the analysis in Figure 5.23, it can be seen that using these less accurate ephemerides has very little effect on the maximum error, even when sourcing all ephemerides this way. Implementing this reduces the simulation runtime by 0.006 [s].

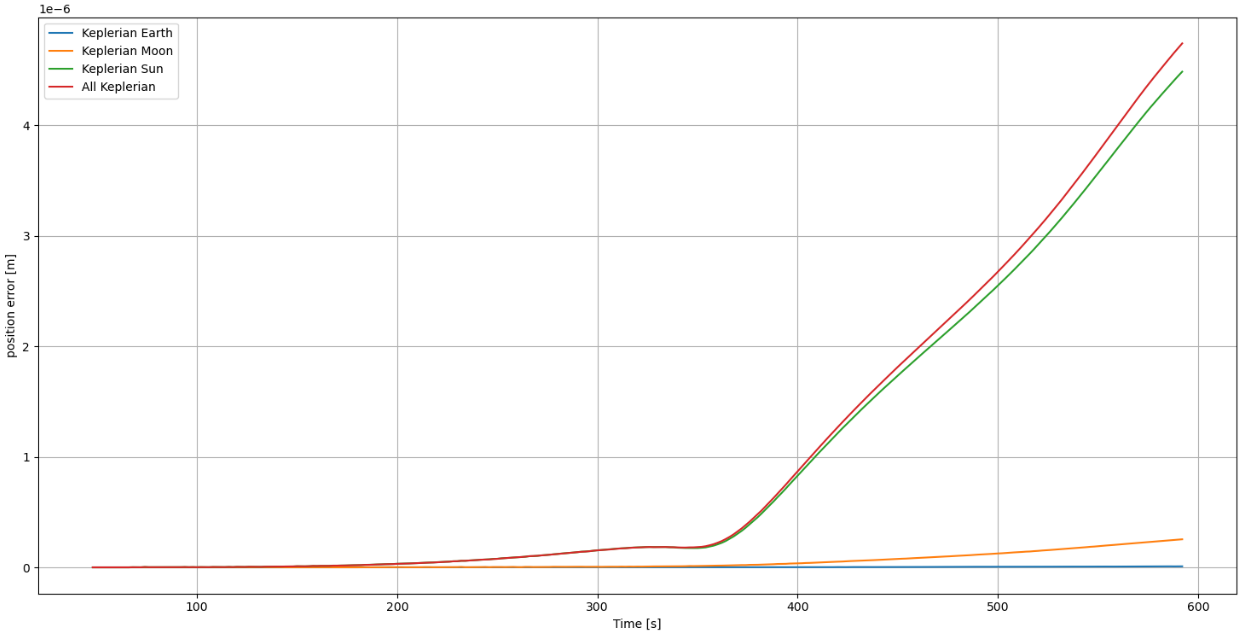

Figure 5.23: The position error over time of the simulation using Keplerian ephemerides for the various bodies, compared to one using the default ephemerides from SPICE.

Finally, different models can be used for the rotation of the Earth. In addition to the default model, a simple rotation model from the SPICE initial state can be used, as well as the more accurate IERS Earth rotation model. After investigating these options in Figure 5.24, it can be seen that the simple from SPICE option does not differ much at all, but that the IERS rotation model, indicated as ITRS, improves the accuracy by approximately 1 [m], so this option will be used. This results in the simulation being approximately 0.041 [s] longer.

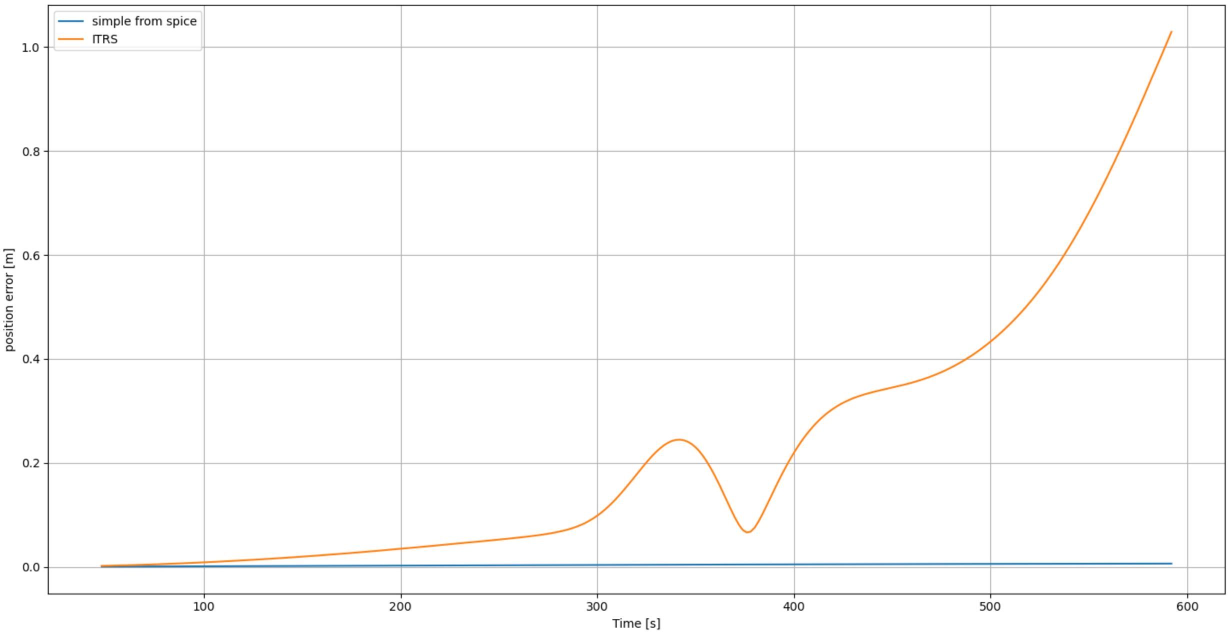

Figure 5.24: The position error over time of the simulation using various Earth rotation models, compared to the default model.

With all models assessed, the simulation settings have been selected. An overview of these can be found in Table 5.2. Using these options, it can confidently be said the model accuracy falls within the desired 10 [m] limit. The total increase to simulation time from the use of more accurate and computationally intensive models amounts to 0.0436 [s], most of which is attributed to the Earth rotation model.

Table 5.2: The choices made for the setup of the TudatPy simulation.

Simulation setup choice	**Simulation setup selection made**	Approximate runtime increase
Degree and order of central body gravitational model	6,6	± 0 [s]
Included third bodies, and the degree and order of their gravitational model	Sun & Moon, both point mass	0.0051 [s]
Radiation Pressure	not included	0 [s]
Atmosphere model	NRLMSISE-00 atmosphere model	0.0035 [s]
Central body shape model	Oblate sphere	± 0 [s]
Central body rotation model	IERS Earth rotation model	0.041 [s]
Ephemerides source	Keplerian from SPICE for all bodies	-0.006 [s]
Propagator	Cowell	-
Integrator	RK5(6) fixed step	-
integrator tolerance/step size	16 [s]	-

5.4. Guidance Algorithm

As stated in section 2.3, a guidance algorithm based on the Apollo entry guidance is used for longitudinal guidance, combined with heading-angle deadbands for the lateral guidance. This left four decisions to be determined:

- The reference bank angle used in the generation of the reference trajectory.

- The overcontrol gain value K.
- The deadband coefficient c_0.
- The deadband coefficient c_1.

Of these, the overcontrol gain K and deadband coefficients c_0 and c_1 are left as inputs for the optimisation, which leaves the reference bank angle profile used by the reference trajectory.

The main effect of the reference trajectory bank angle is on the range of the trajectory and the loads that are experienced. Increasing the bank angle will result in a steeper entry, shortening the range and increasing the loads. The decision was made to select a reference bank angle trajectory based on the expected downrange distance the reference trajectory would reach. The trajectory was split into three segments, based on the velocity, while keeping in mind the experienced loads. A typical plot of the loads simulated by the reference trajectory can be found in Figure 5.25 and Figure 5.26. The first segment is from the moment of entry till the velocity reaches 70% of the entry velocity, which should cover the peak loads experienced. The second segment is from 70% to 50% of the entry velocity, when the vehicle is likely to still have high control authority but does not yet experience the second peak that typically occurs in the g-load. The final segment covers the rest of the trajectory, from 50% entry velocity until the vehicle reaches 30 [km], which is the part of the trajectory with the least control authority, and which incorporates the second peak that typically occurs in the g-load.

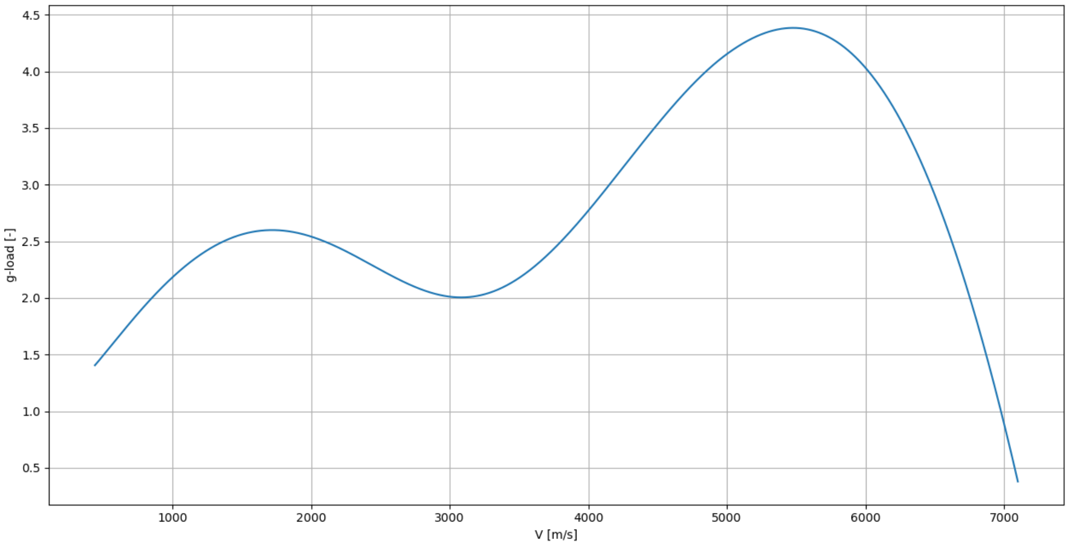

Figure 5.25: A typical plot of the g-load over velocity from the reference trajectory.

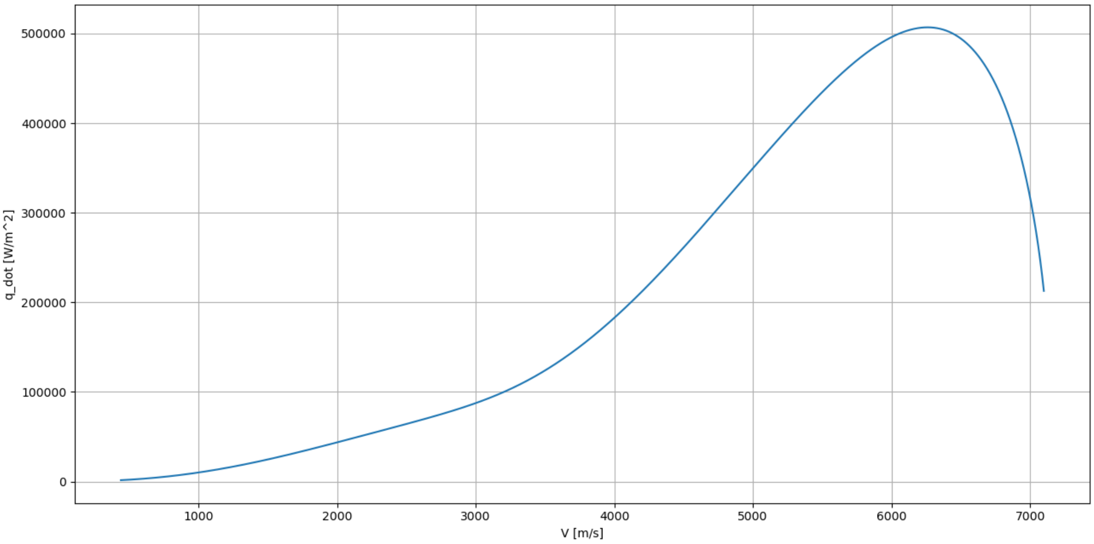

Figure 5.26: A typical plot of the heat flux over velocity from the reference trajectory.

To determine whether the reference trajectory achieves a sufficient downrange distance, a value must be set for the desired range to the target location. For this, a value is chosen of 5 [km], meaning that in the 30 [km] of the final descent, the vehicle must travel approximately 167 [m] laterally each kilometre of the descent. This corresponds to a glide slope of roughly 80 [deg], which is considered to be an achievable figure.

To generate the reference trajectory, the simulation is run once to generate the entry interface values, defined here as the values of the trajectory at which the drag force acceleration reaches 0.2 g's or the altitude reaches 60 [km], as these are the two conditions that can start the guidance operation. With the entry interface values determined, an initial guess is used to generate a reference trajectory, with an associated downrange distance. If this distance is not within 5 [km] of the range to the target location, the bank angle profile must be varied to achieve the desired range.

The selection of the reference trajectory starts with the first segment, between the entry and 70% of the entry velocity. If the trajectory undershot the target, the bank angle in the segment is decreased. If it overshot, the bank angle is increased, so long as the g-load and heat flux in this segment have not yet reached 80% of the limits set. These increases or decreases are done by increments of 1.0 [deg]. If the bank angle for this segment cannot be changed, due to reaching the minimum or maximum value of 5.0 and 185.0 [deg] or due to the loads exceeding 80% of the limit with an overshoot, the current bank angle is kept for this segment, and the code moves to the next segment. This also occurs if the trajectory was overshot or undershot by changes made in the current segment. Each segment uses the same logic as for the first segment. If the third segment still does not reach the desired downrange distance, another loop through the segments is performed, this time allowing loads up to 100% of the load limits. If this fails, or the number of iterations reaches 300, the loop is ended and the final, best guess bank angle profile is used. Once a bank angle profile is selected, this is used to generate the reference trajectory to be compared to by the guidance during the actual simulation. This is performed for every separate simulation, as the entry interface, and therefore the reference trajectory, may differ between runs.

The recommendation from the guidance algorithm was to run a guidance loop at least once per second. In accordance with this, the simulation time step is changed to 1 [s]. As it had been used so far and as it had produced no issues, the selected RK5(6) is kept as the integrator. 5 [deg/s] is seen as a feasible rotation rate, and so the limit to the bank angle change per time step is set to this value.

The implementation of the guidance algorithm assumes the vehicle has sufficient controllability to fly at a constant angle of attack, to have a roll rate of 5 [deg/s], and to cross 5 [km] of distance after re-entry. The verification of the aerodynamic control authority, and the sizing of the actuators and control surfaces is beyond the scope of the research.

5.5. Verification & Validation

With the simulation set up, it must be verified and validated. As TudatPy is frequently used for scientific research and continuously maintained by its developers, its functions are considered to be verified. The remaining code is unit tested to verify it is working as desired, seeing if inputs result in expected outcomes and using extreme value tests.

Validation would ideally require flight data of a mission done by the vehicle, but as no such mission has been flown, this is not an option. Instead, a similar mission is chosen, specifically the re-entry of the Apollo 10 command module. Although this vehicle is entering at greater velocity, they are both lifting ballistic entries performed by vehicles controlling their descent with their bank or roll angle. The vehicle's bank angle profile is recreated, as seen in Figure 5.27, and then the simulation is run with the Apollo vehicle's initial state, resulting in the figures in Figure 5.28, Figure 5.29 and Figure 5.30.

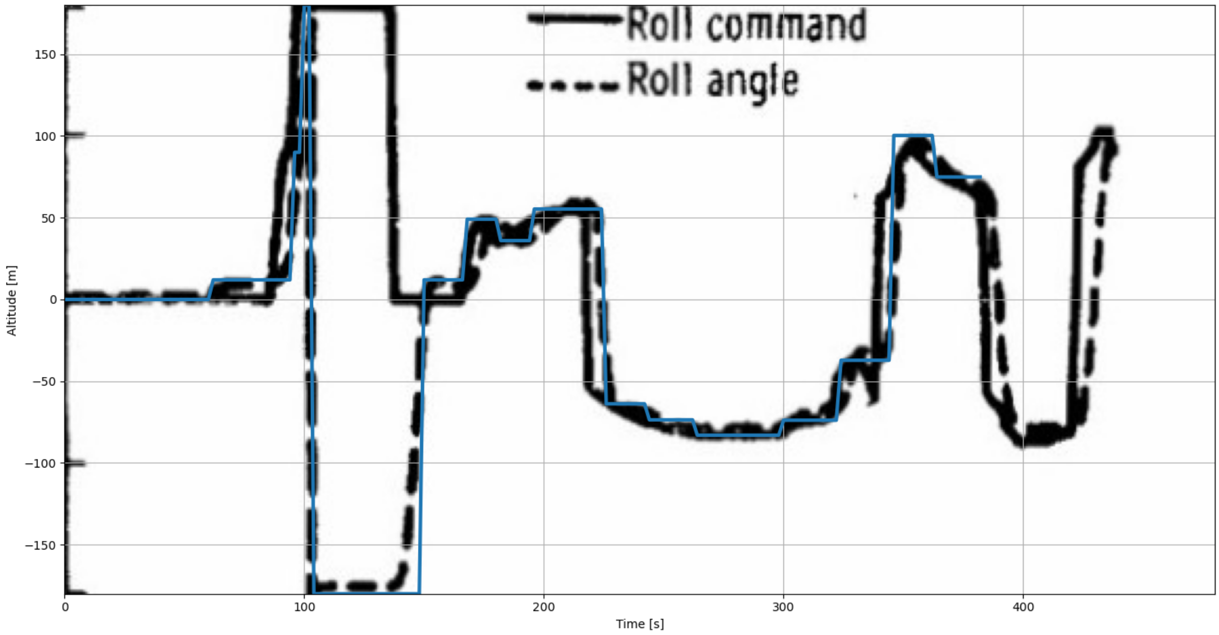

Figure 5.27: The roll angle profile of the Apollo 10 re-entry and the bank angle profile used by the validation trajectory. [21]

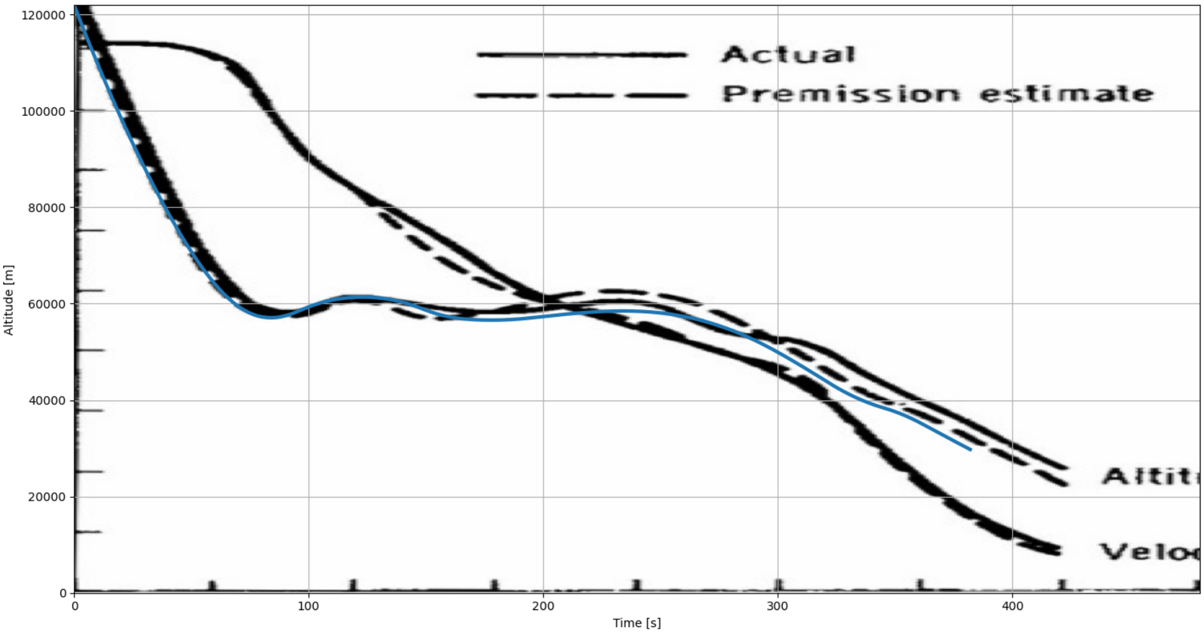

Figure 5.28: The altitude over time of the Apollo 10 re-entry and the validation trajectory.[21]

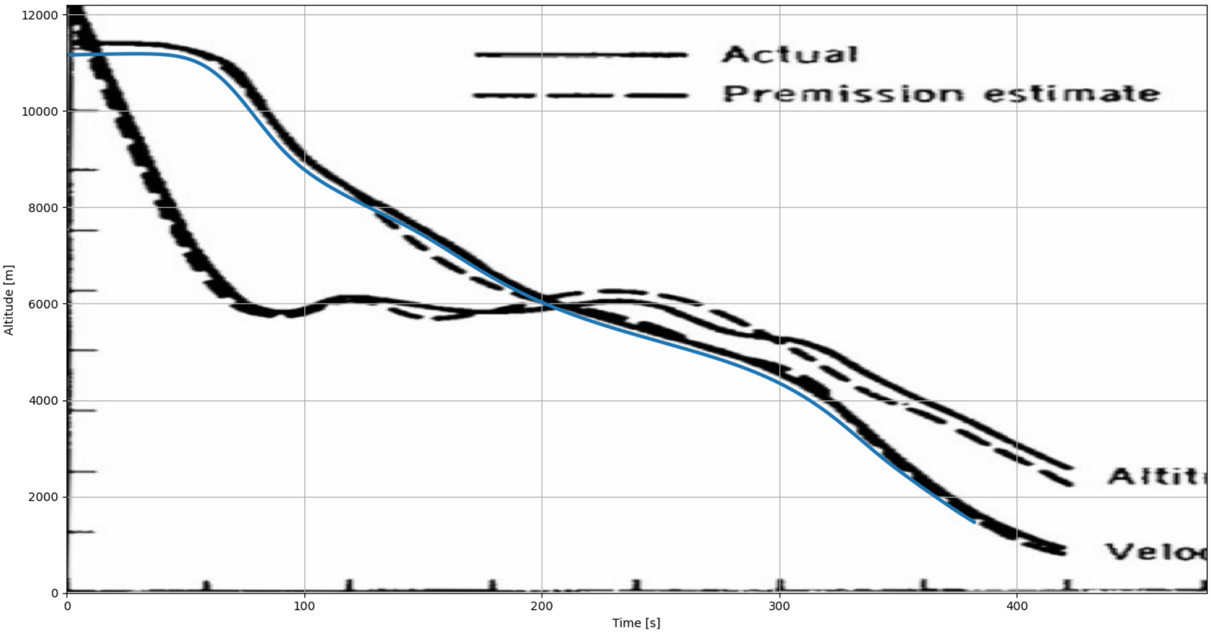

Figure 5.29: The velocity over time of the Apollo 10 re-entry and the validation trajectory. [21]

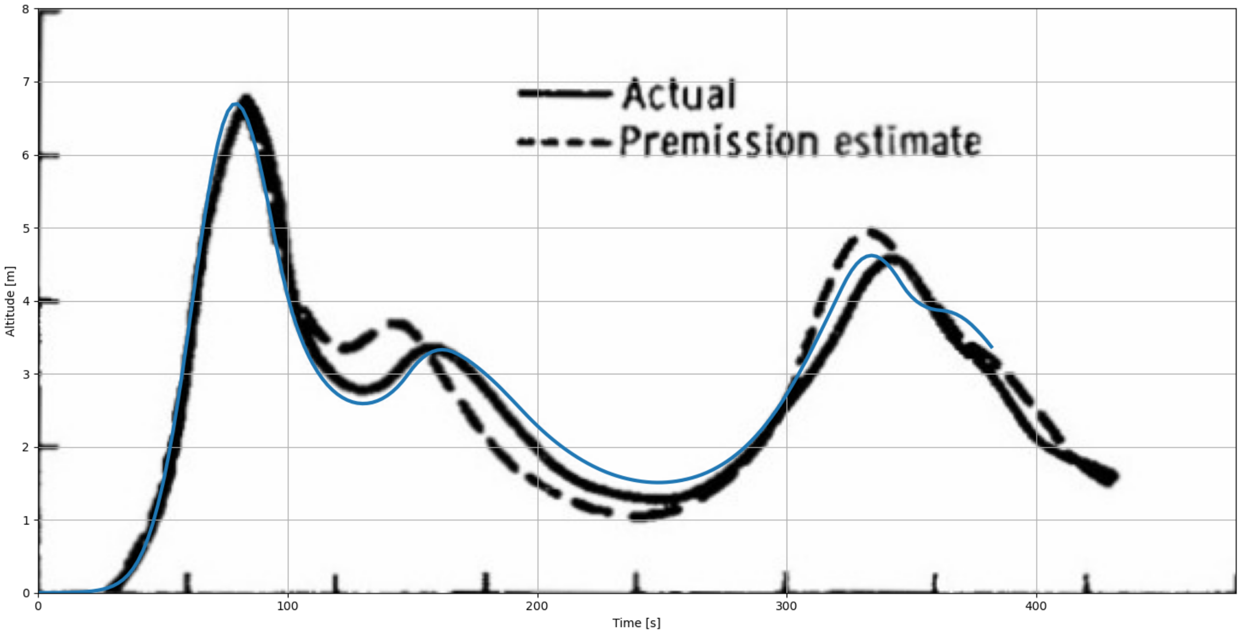

Figure 5.30: The g-load over time of the Apollo 10 re-entry and the validation trajectory. [21]

From these images, it can be seen that there is a reasonably good fit between the validation data and the produced results. There are some discrepancies, but these are attributed to inaccuracies that arose due to the discretisation of the continuous roll angle profile from the Apollo mission into numbers that could be used for the validation, as there was no data, and so this had to be done manually. Nonetheless, the behaviour of the simulation is quite similar to the mission data, and so there is reasonable confidence that the simulation sufficiently accurately represents the physical environment.

5.6. Initial Investigations

As is described in chapter 4, the completed simulation is used to perform some initial investigations into the capabilities of the vehicle. Here it appeared that indeed the vehicle was not capable of skipping re-entry. Using the initial conditions stated in Table 5.1, the trajectory was propagated without guidance, commanding full lift up, resulting in the ground track visible in Figure 5.31. The range achieved during this flight is approximately 3200 [km]. Furthermore, the limited lift and significant mass and drag of the vehicle limit somewhat the control authority of the vehicle, which causes it to not be able to reach any useful target locations under its own control, instead ending up in the ocean.

Figure 5.31: The ground track of the simulation using the vehicle's initial state and no guidance.

It is therefore necessary to perform an adjustment to the initial state after the engine bay has disconnected from the stage. This initial state correction will require some amount of propellant, which will in turn increase the cost of the vehicle and incur a payload mass penalty. It is therefore especially important to limit the amount of propellant needed for this initial correction.

5.7. Landing Sites

With a manoeuvre being needed after the disconnection, many options are opened for where the vehicle could travel. From these, a number of options must be picked to investigate. A number of factors must be considered when selecting the target locations to be considered:

- **Proximity to the manufacturing site:** It is assumed refurbishment of the engine and the engine bay will occur at the manufacturing site at Les Mureaux, near Paris. A landing site closer to Paris reduces the travel time and costs.
- **ESA member states:** To limit the influence other countries may have on Europe's independent access to space, it is useful to select a landing site that is on the soil of an ESA member state. A less attractive option exists in the form of countries outside of ESA that are on good terms with Europe, as these are likely willing to cooperate.
- **Heading:** The vehicle will be heading eastwards, under the initial conditions, at a heading of roughly 95 degrees. Target locations that are close to this flight path will require less deviation.
- **Distance:** The vehicle will be re-entering at a certain flight path angle and velocity, which will result in a certain downrange distance. In the initial investigation, this distance was 3200 km. Target locations that differ more from this will require greater deviations.
- **Safety:** As with any mission, there is a chance of failure. Furthermore, the vehicle will jettison its heat shield near the end of its journey. To ensure safe operations, the vehicle should travel over seas and sparsely populated areas if possible.

Using these factors, five target locations were selected. These can be found in Table 5.3, as well as on the map in Figure 5.32.

Table 5.3: The selected target locations and how they compare on the selection criteria.

Target Location	Proximity to Manufacturing Site	ESA Affiliation	Heading	Distance	Safety
Natal	On the north coast of Brazil, Natal is the furthest of all considered landing sites.	Although not an ESA member state, Brazil is a stable country on good terms with Europe.	128 [deg], 33 [deg] south of initial heading.	2287 [km], 913 [km] short of expected range.	Flight path mostly over the ocean.
Cabo Verde	Off the west coast of Africa, Cabo Verde is about one third the distance from launch to landing site.	Although not an ESA member state, Cabo Verde is a stable country on good terms with Europe.	67 [deg], 28 [deg] north of initial heading.	3364 [km], 154 [km] over expected range.	Flight path mostly over the ocean.
The Canarias	Off the west coast of Africa, The Canarias are about half the distance from launch to landing site.	Territory of Spain, an ESA member state.	58 [deg], 37 [deg] north of initial heading.	4674 [km], 1474 [km] over expected range.	Flight path mostly over the ocean.
The Azores	Out in the Atlantic, The Azores are about half the distance from launch to landing site.	Territory of Portugal, an ESA member state.	33 [deg], 62 [deg] north of initial heading.	4536 [km], 1336 [km] over expected range.	Flight path mostly over the ocean.
Les Mureaux	Just outside of Paris, Les Mureaux is where the engine bay is assumed to be refurbished.	Territory of France, an ESA member state.	37 [deg], 58 [deg] north of initial heading.	7056 [km], 3856 [km] over expected range.	The final part of the flight path is over a densely populated part of France.

Figure 5.32: The Ariane 6 launch site (red) and possible engine bay landing sites (green). Map via Google Maps.

Out of these, it seems like Cabo Verde may make for the most attractive option. It is close to the previously found range and not far off in terms of its heading. Although it is not part of an ESA member state, Cabo Verde has close ties with Europe. It is not especially close to the refurbishment site, which will bring in some further transport costs. How exactly this will cause this option to compare with the others will require the results of the optimisation.

5.8. Optimisation Inputs, Objectives and Constraints

To find the optimum trajectories to each of the target locations and to be able to compare their performance, an MDO is set up. To do this, choices have to be made on inputs, objectives and constraints for this problem. This is done in subsection 5.8.1, subsection 5.8.2 and subsection 5.8.3 respectively.

5.8.1. Inputs

First and foremost, there are 5 different target locations, but as the only possibility here is to select between them without ending anywhere in between, these are not used as inputs to the optimisation. Instead, five optimisations will be run separately, one for each target location, and compared afterwards. This leaves as inputs to the optimisation the specifics of the initial trajectory adjustment discussed in section 5.6, and the parameters of the guidance algorithm, as stated in section 2.3. This produces a set of six inputs:

- Magnitude of the velocity vector V.
- Heading angle of the vehicle ψ.
- Flight path angle of the vehicle γ.
- Guidance overcontrol value K.
- Guidance deadband value c_0.
- Guidance deadband value c_1.

5.8.2. Objectives

Three objectives have been identified, the most important of which is the propellant mass used for the initial correction. As the propellant mass needed may form the bulk of the payload mass penalty, it is important to minimise this value. This is used by calculating the change in the velocity vector before and after the correction and using that with Equation 5.1 to calculate the propellant needed for this adjustment. The second most important objective is the range to the target location. It was debated whether this would be more appropriate as an objective or as a constraint, as it is not certain if there is any added benefit to approaching the target more closely. For now, this is kept under the objectives, but this would be reconsidered after the Monte Carlo analysis. This distance is defined as the great-circle distance using the Earth's average radius between the vectors from the centre of the Earth to the target location and the final location of the vehicle. The final objective is the number of bank angle reversals. Minimising this would lead to less propellant mass needed for large reversals of the bank angle and is therefore a relevant but less important objective. The number of bank angle reversals is kept track of during the re-entry in the guidance algorithm. The list of objectives, therefore, is:

- The propellant mass used during the initial trajectory correction m_p.
- The final distance to the target location d.
- The number of bank-angle reversals r.

5.8.3. Constraints

Three constraints are easily selected: the peak g-load, peak heatflux, and total heatflux values are values that must be considered. The g-load is an output that is directly extractable from the simulation. Heat flux must be calculated, which is done using Equation 5.2.

$$\dot{q} = k\sqrt{\frac{\rho}{R_N}}V^3 \tag{5.2}$$

Where k is a constant with a value of $1.74153 \cdot 10^{-4}$, and R_N is the nose radius of the vehicle, which has a value of 1.861 [m] for this vehicle. The total heat load is then calculated by integrating the stagnation point heat flux. The final constraint is set to the velocity at the end of the trajectory, so the vehicle can survive the loads for the final descent after 30 [km]. This makes the list of constraints as follows:

- Peak g-load must be less than 10.
- Peak heat flux must be less than 1.0 [MW/m^2].
- Total heat load must be less than 200 [MJ/m^2].
- Final velocity must be less than 600 [m/s].

With the inputs, objectives and constraints determined, the process of optimisation can be started.

5.9. Monte Carlo Analysis

Before the full optimisation is run, a Monte Carlo (MC) analysis is performed to gain insight into what effect the inputs have and what ranges should be considered for them during the optimisation. For this, an initial guess is made by manually varying the parameters to achieve a reasonable solution. These will then be varied randomly within a range. The initial guesses, and the variation range of the parameters, can be found in Table 5.4. First, subsection 5.9.1 presents the one-at-a-time MC analysis, then subsection 5.9.2 presents the all-at-once MC analysis.

Table 5.4: The initial guess and variation range used for the MC analysis.

Input Parameter	Initial Guess for Natal	Initial Guess for Cabo Verde	Initial Guess for The Canarias	Initial Guess for The Azores	Initial Guess for Les Mureaux	Variation Range
V [m/s]	6450	6970	7275	7375	7505	± 50
ψ [deg]	126.0	68.5	50.5	31.0	35	± 1.0
γ [deg]	-0.8	-0.8	-0.8	-0.8	-0.8	± 0.2
K [-]	1.0	1.0	1.0	1.0	1.0	$-1.0, +9.0$
c_0 [deg]	2.0	2.0	2.0	2.0	2.0	$-1.5, +3.0$
c_1 [deg/(m/s^2)]	$1.633 \cdot 10^{-7}$	$1.633 \cdot 10^{-7}$	$1.633 \cdot 10^{-7}$	$1.633 \cdot 10^{-7}$	$1.633 \cdot 10^{-7}$	$\pm 1.633 \cdot 10^{-7}$

5.9.1. One-at-a-Time Monte Carlo Analysis

First, a one-at-a-time MC analysis is performed. One input is varied using a uniform distribution, while the remaining inputs are kept at the value of the initial guess. These inputs are then used to run the simulation, and the outputs and constraint values are plotted. The results for Cabo Verde, which was expected to be the most promising location, can be found in Figure 5.33, Figure 5.34, Figure 5.35, Figure 5.36, Figure 5.37, Figure 5.38, Figure 5.39, Figure 5.40, Figure 5.41, Figure 5.42, Figure 5.43, and Figure 5.44. The figures for the remaining target locations can be found in Appendix B. This entire process was performed for two seeds, namely 42 and 22. The plots shown are for seed 42, the plots for seed 22 follow the same sort of patterns visible in these plots.

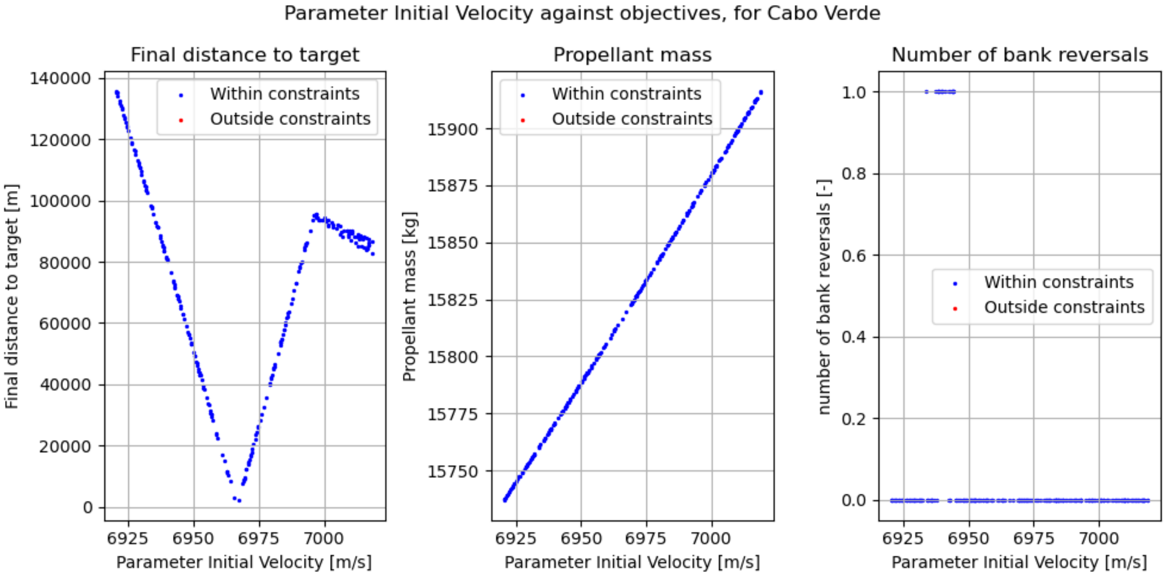

Figure 5.33: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial velocity for Cabo Verde.

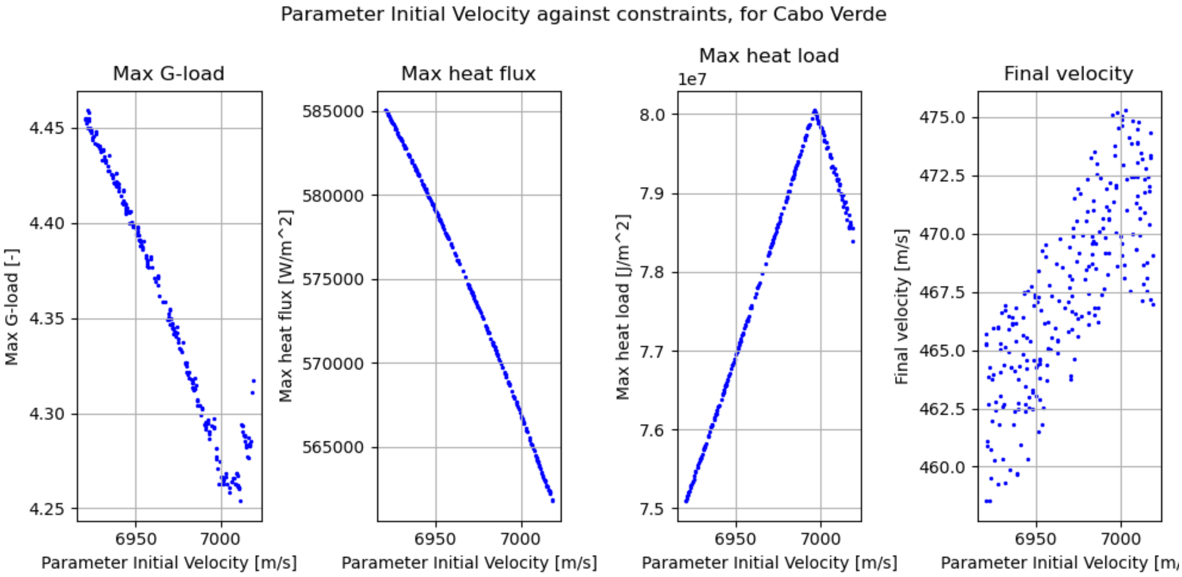

Figure 5.34: The maximum g-load, maximum heat flux and total heat load produced by simulations in the single MC variation of the initial velocity for Cabo Verde.

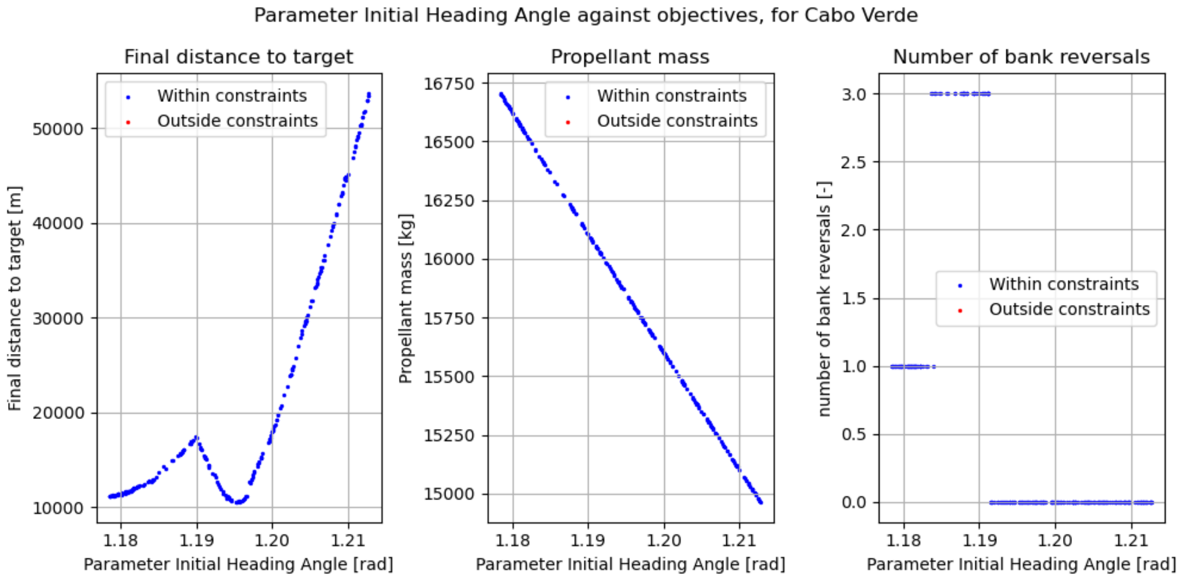

Figure 5.35: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial heading angle for Cabo Verde.

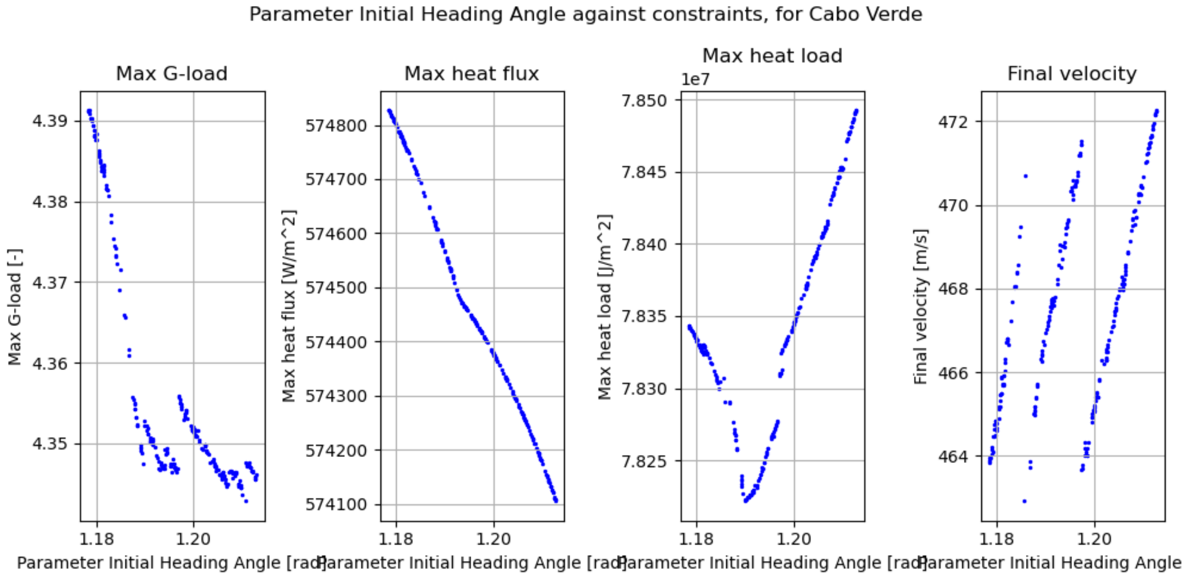

Figure 5.36: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial heading angle for Cabo Verde.

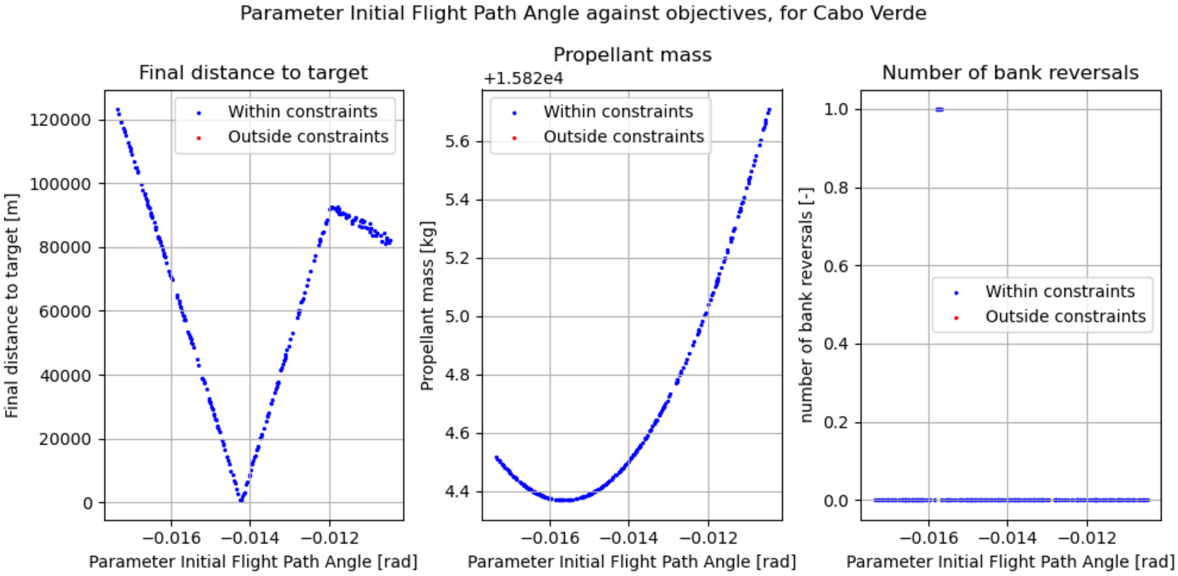

Figure 5.37: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial flight path angle for Cabo Verde.

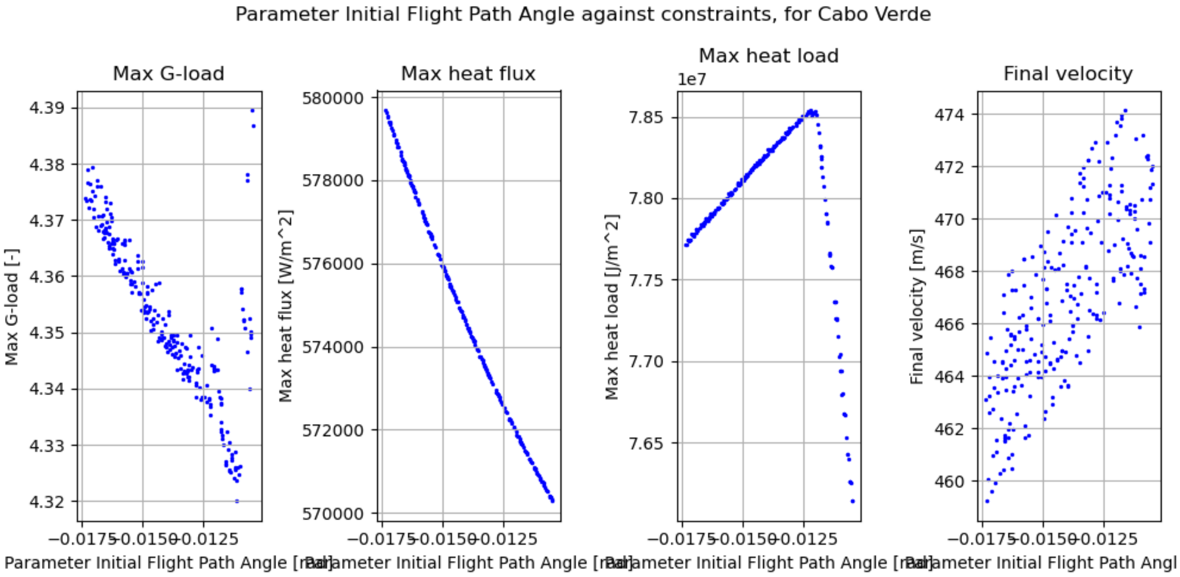

Figure 5.38: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial flight path angle for Cabo Verde.

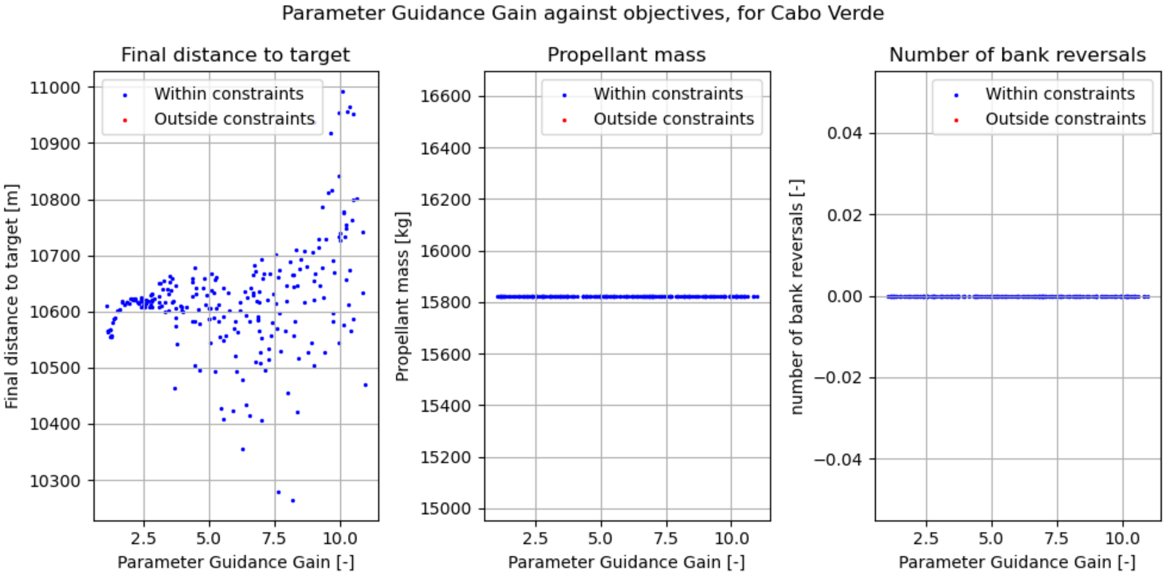

Figure 5.39: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the overcontrol guidance gain for Cabo Verde.

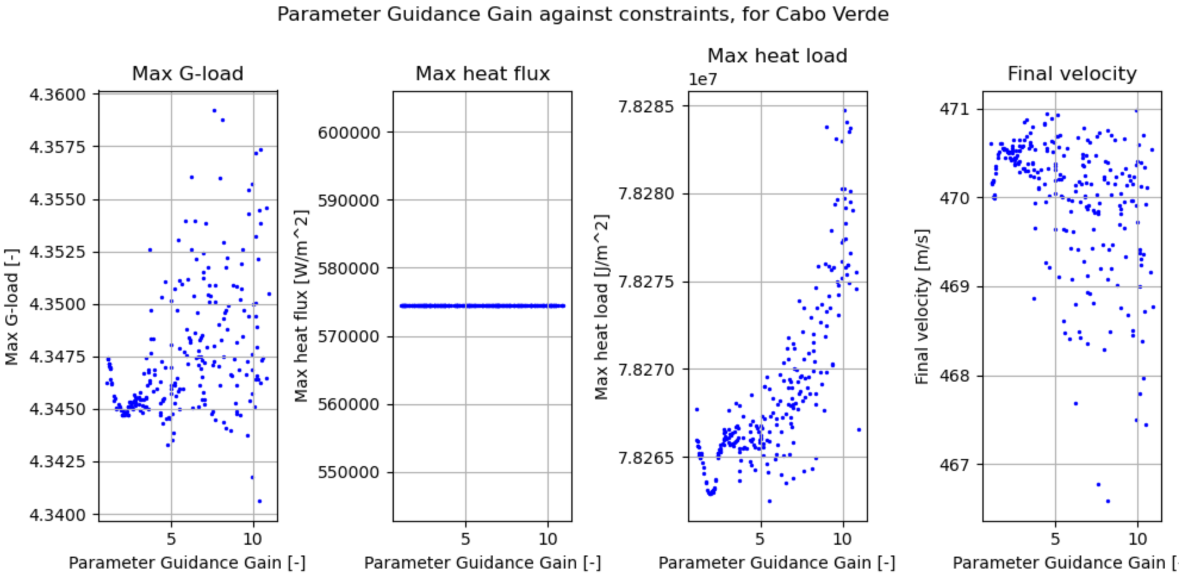

Figure 5.40: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the overcontrol guidance gain for Cabo Verde.

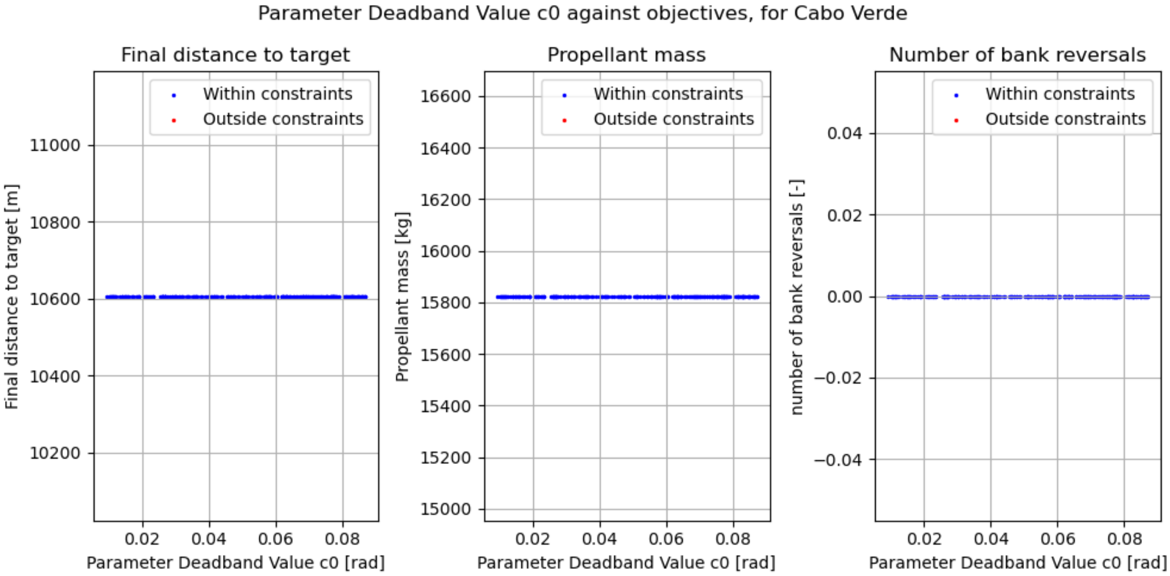

Figure 5.41: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c0$ for Cabo Verde.

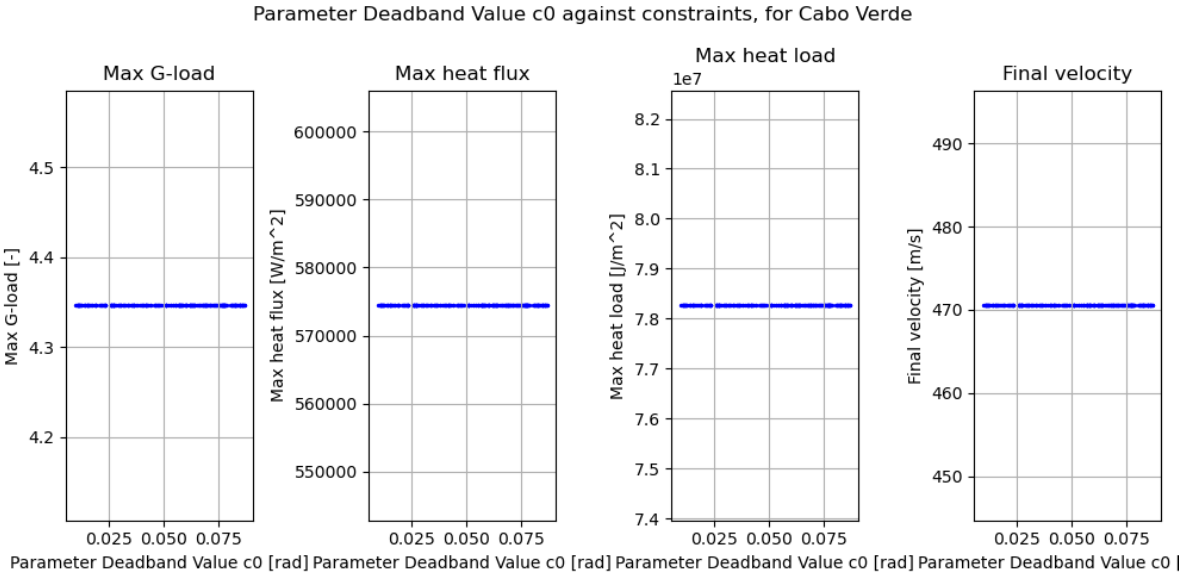

Figure 5.42: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c0$ for Cabo Verde.

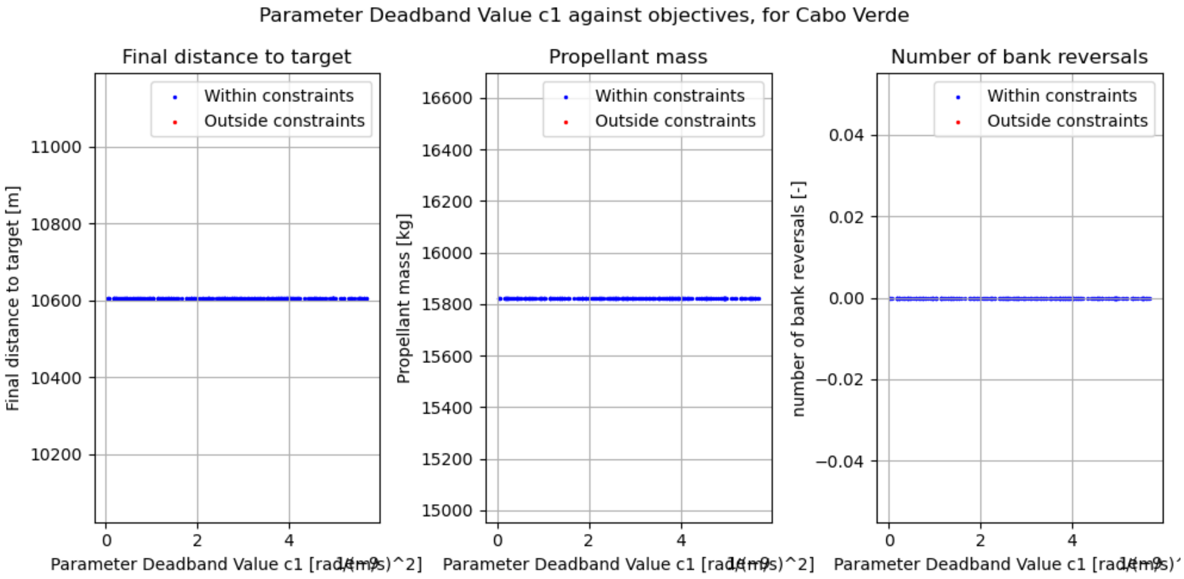

Figure 5.43: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c1$ for Cabo Verde.

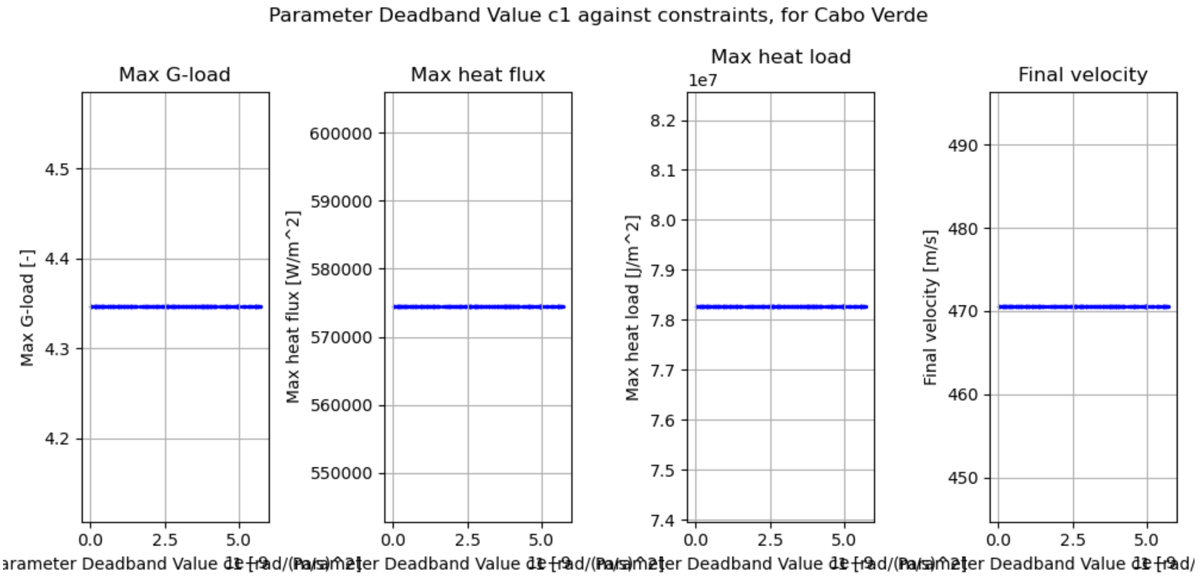

Figure 5.44: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c1$ for Cabo Verde.

From these plots a number of conclusions can be drawn for each of the input parameters:

- As initial velocity increases, it reaches a point where the final distance to the target is minimised before it increases again, as the vehicle is now more likely to overshoot. At higher velocities, the distance starts to drop again, as the vehicle uses the higher control authority to dive much more steeply into the atmosphere. Increasing the velocity naturally increases the amount of propellant mass needed. The number of bank angle reversals was mostly unaffected, all remaining within acceptable levels. Surprisingly, increasing initial velocity actually reduces the peak g-load and heat flux, which, upon investigation, seemed to be because due to the increase in lift the vehicle drops to lower altitudes less quickly, and so its peak deceleration and heat flux are slightly reduced. This corresponds to the heat load increasing with initial velocity, as the vehicle spends more time in the re-entry and so accumulates more heat, up to the point where it can dive deeper into the atmosphere with higher velocity, which drops the heat load once again. Finally, the final velocity is generally increased by increasing the initial velocity.

- Similar to the initial velocity, there is an optimal heading angle for the distance to target, beyond which the vehicle misses the target by more. Diverting more from the initial heading before the trajectory correction naturally increases the required propellant mass, for Cabo Verde this means using a higher heading angle requires less propellant. The number of bank angle reversals does seem affected by the heading angle, with a more precisely directed trajectory actually requiring more reversals, while one that is coming at the target from one side requires zero or one reversal. Reducing the heading angle interestingly enough raises the g-load and heat flux, but not by very much. The heat load seems to be lowest for a trajectory that is aimed more precisely towards the target. The final velocity seems to respond somewhat erratically to the initial heading angle.
- The initial flight path angle seems to have a very similar effect compared to the initial velocity. There is an optimum where the distance to the target location is minimised. The number of bank angle reversals is not greatly affected by it. The effect on the constraints is once again similar to that of the initial velocity. The main difference is that the propellant mass is not significantly affected by the flight path angle. This makes it potentially an interesting option to compensate for a reduction in initial velocity if the effects can cancel each other out.
- The effect of the guidance overcontrol gain value is difficult to quantify. It has no effect on the propellant mass, as expected, but only for some target location does it affect the number of bank angle reversals. Depending on the target location, the other constraints and objectives are affected, but not in a consistent manner. Generally, the spread of the performance on any of these increases as the gain value increases.
- The deadband values had no effect on the constraints and objectives for Cabo Verde, but they did affect some of the other target locations. Decreasing the values seems to somewhat reduce the distance to the target location, as the lateral corridor becomes tighter, leading to more bank angle reversals and higher loads as the vehicle uses less of its lift to stay in the upper atmosphere and decelerate more gradually. This is very general, however, as not all of these were affected.

It is also notable that effectively none of the single MC managed to come within the desired 5 [km]. This was not yet seen as a major concern.

5.9.2. All-at-once Monte Carlo Analysis

Although the one-at-a-time MC is good for isolating the effect of one input parameter on the objectives and constraints, it does not capture the interactions between input parameters. To fully explore the design space, an all-at-once MC analysis is performed. Similarly to the one-at-a-time MC, the inputs are varied using a uniform distribution across the range specified in Table 5.4, but now all inputs are varied at the same time. This was again done for two seeds, 42 and 22. The results for Cabo Verde can be found in Figure 5.45, Figure 5.46, Figure 5.47, Figure 5.48, Figure 5.49, Figure 5.50, Figure 5.51, Figure 5.52, Figure 5.53, Figure 5.54, Figure 5.55, and Figure 5.56. The figures for the remaining target locations can be found in Appendix B. The plots shown are for seed 42, the plots for seed 22 follow the same sort of patterns visible in these plots.

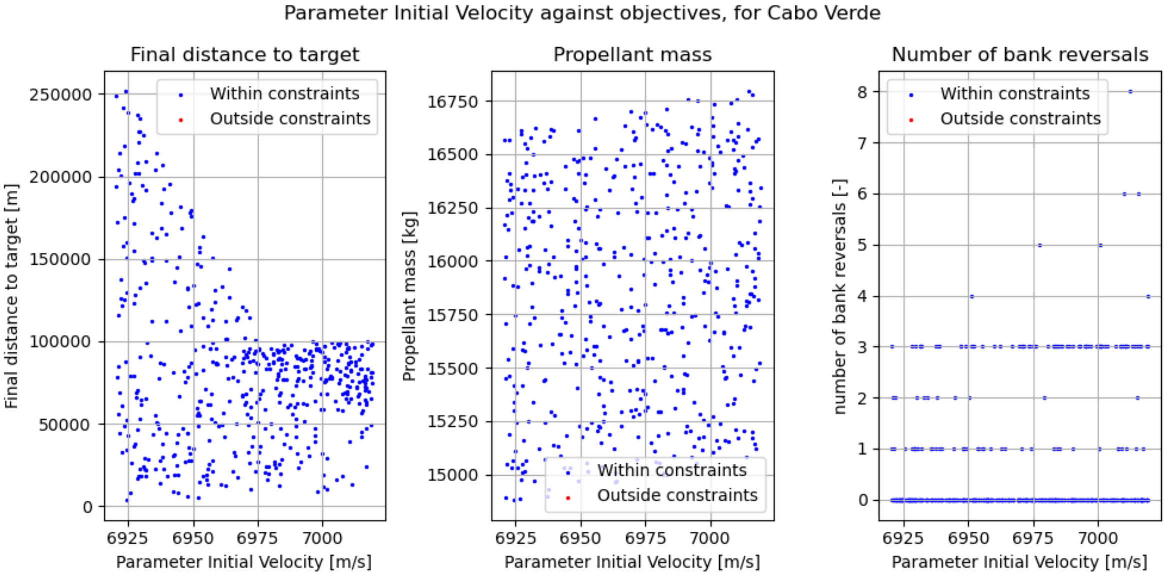

Figure 5.45: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial velocity produced by simulations in the all-at-once MC variation for Cabo Verde.

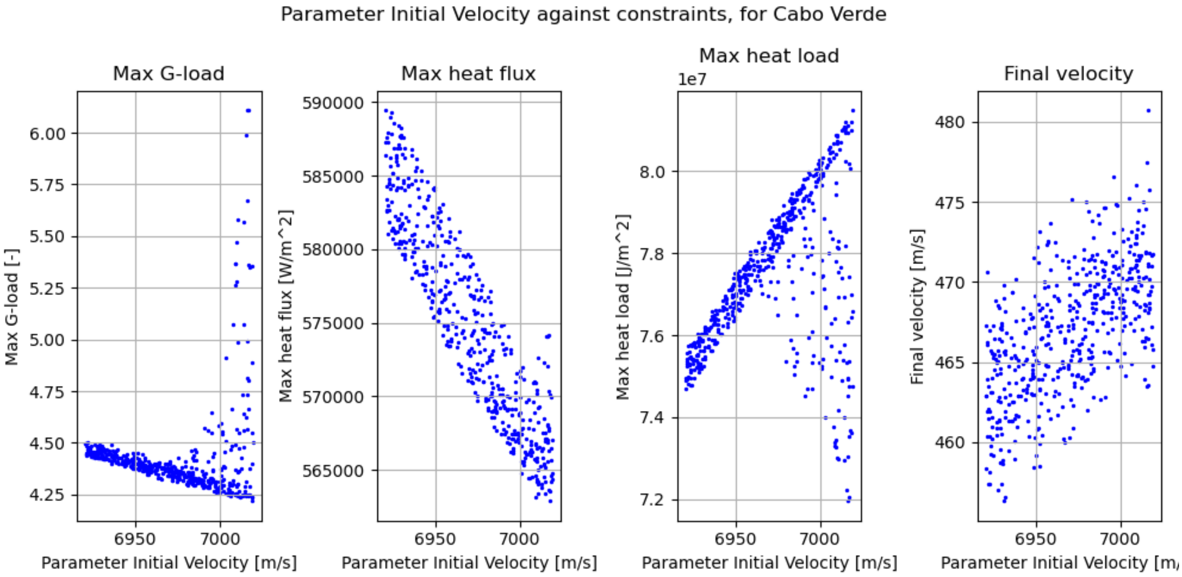

Figure 5.46: The maximum g-load, maximum heat flux and total heat load over the value of the initial velocity produced by simulations in the all-at-once MC variation for Cabo Verde.

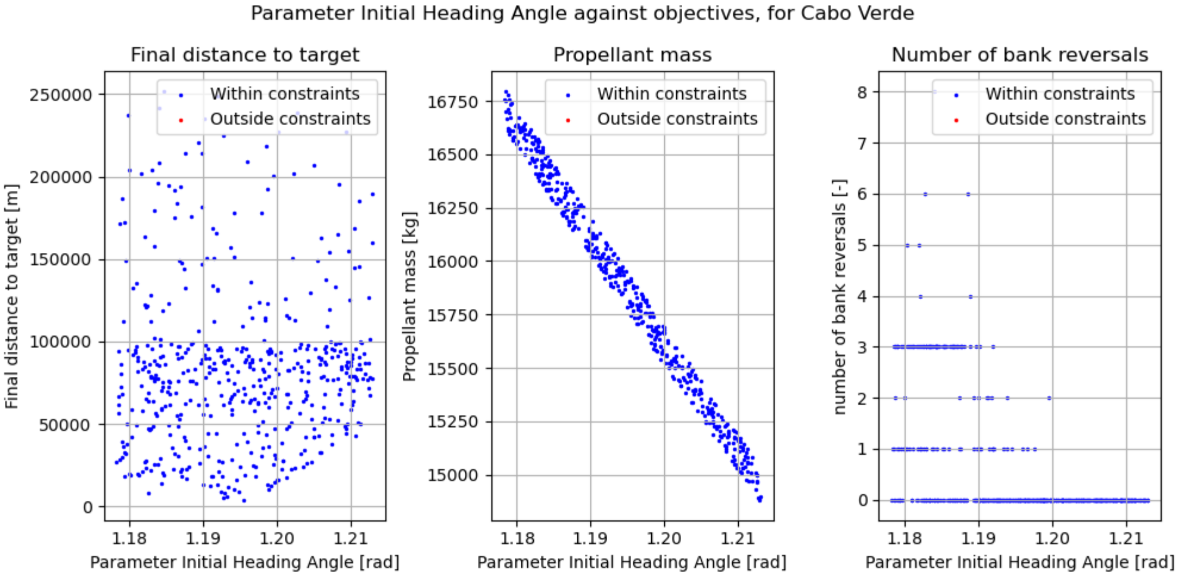

Figure 5.47: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial heading angle produced by simulations in the all-at-once MC variation for Cabo Verde.

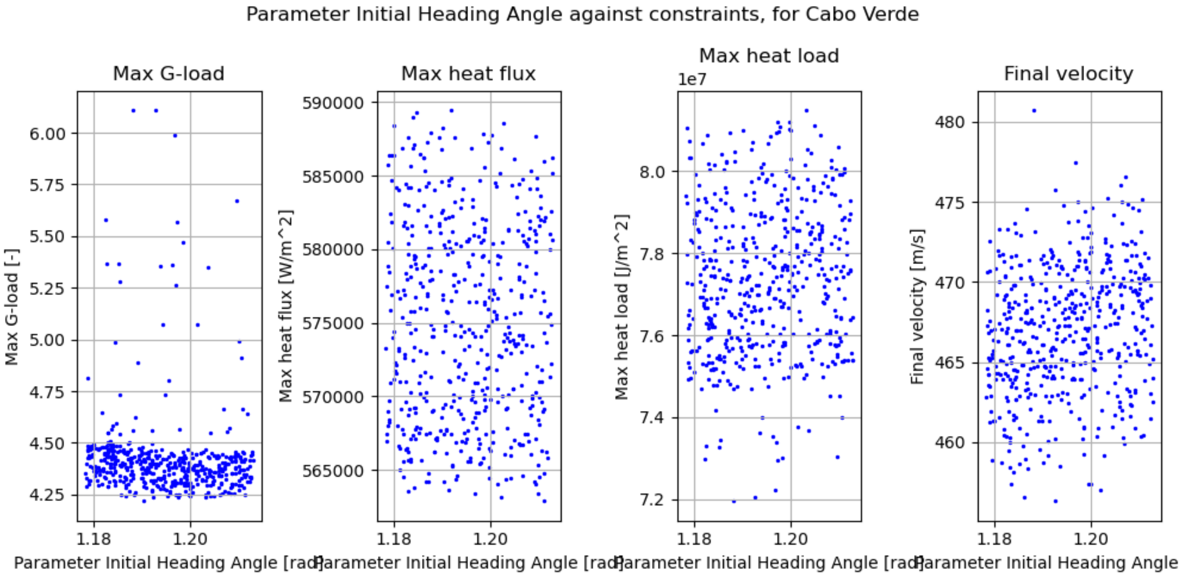

Figure 5.48: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial heading angle produced by simulations in the all-at-once MC variation for Cabo Verde.

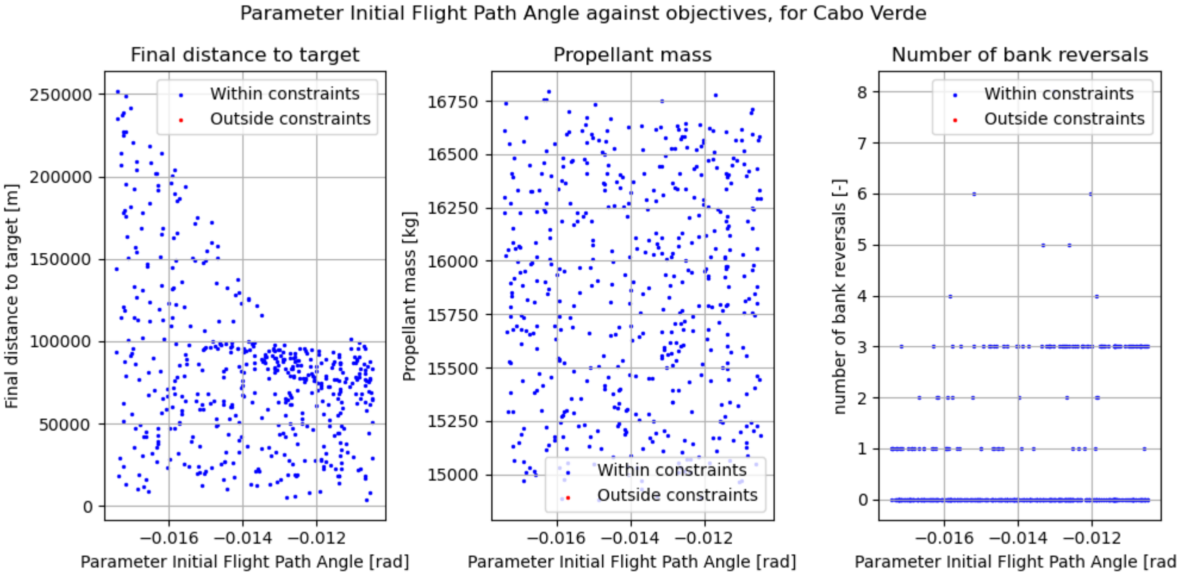

Figure 5.49: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for Cabo Verde.

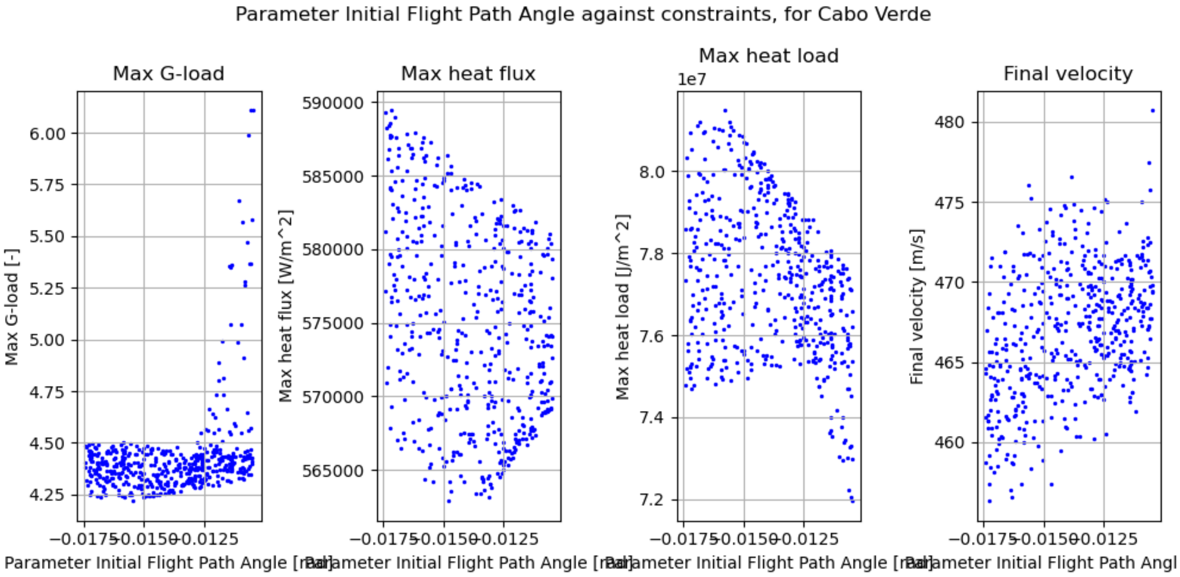

Figure 5.50: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for Cabo Verde.

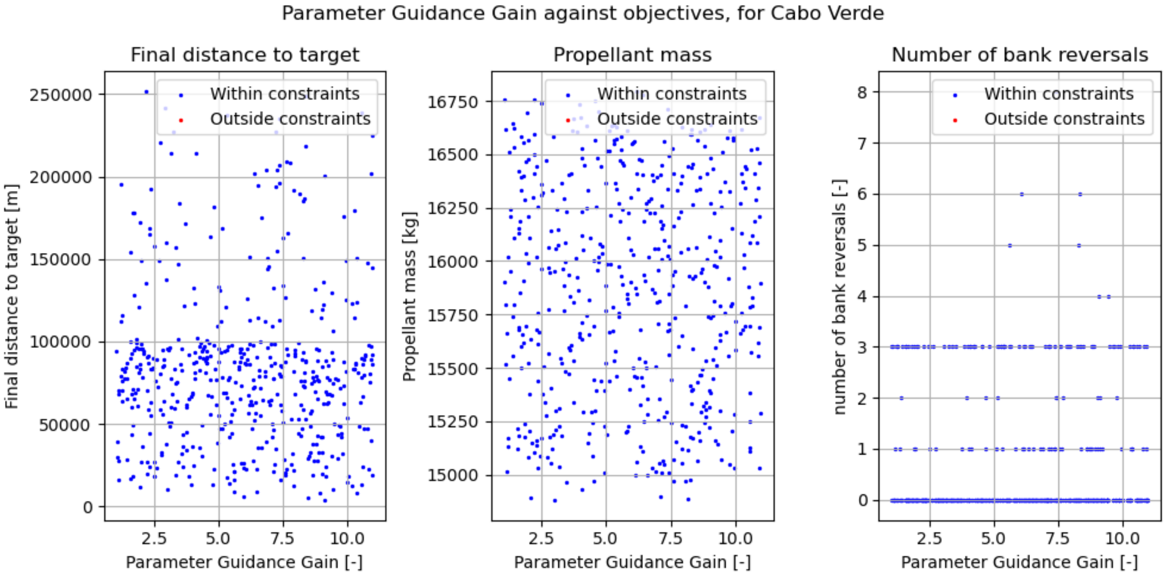

Figure 5.51: The final distance to target, propellant mass and number of bank angle reversals over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for Cabo Verde.

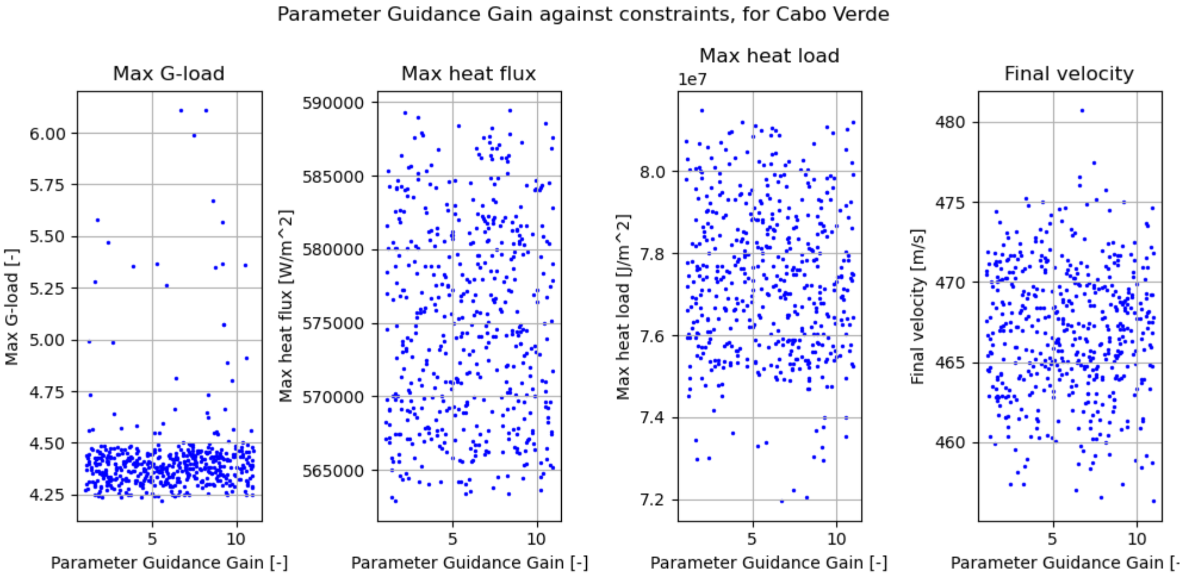

Figure 5.52: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for Cabo Verde.

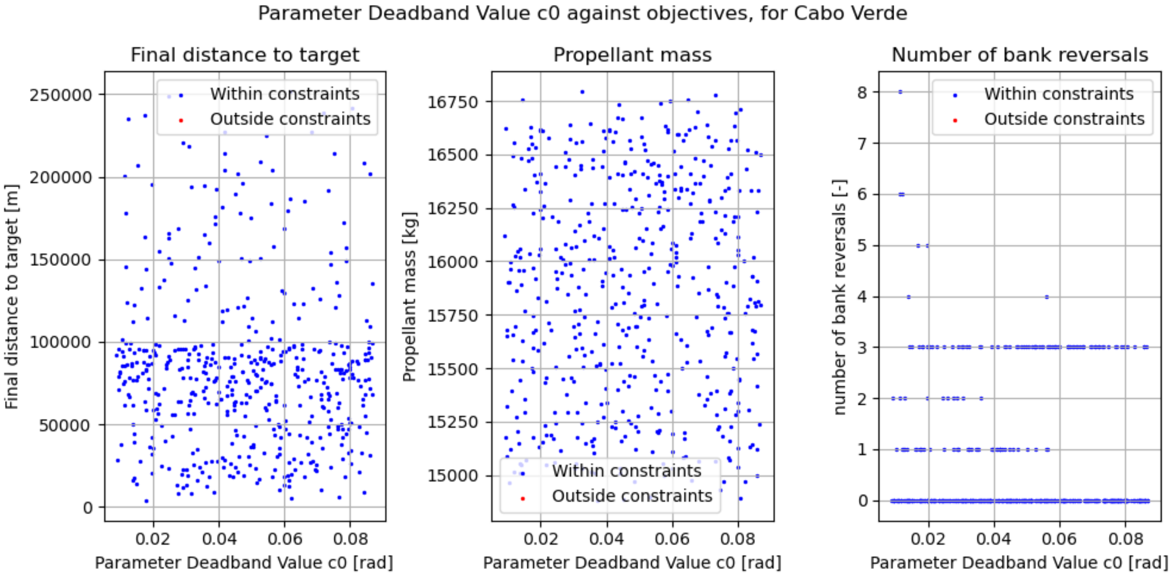

Figure 5.53: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for Cabo Verde.

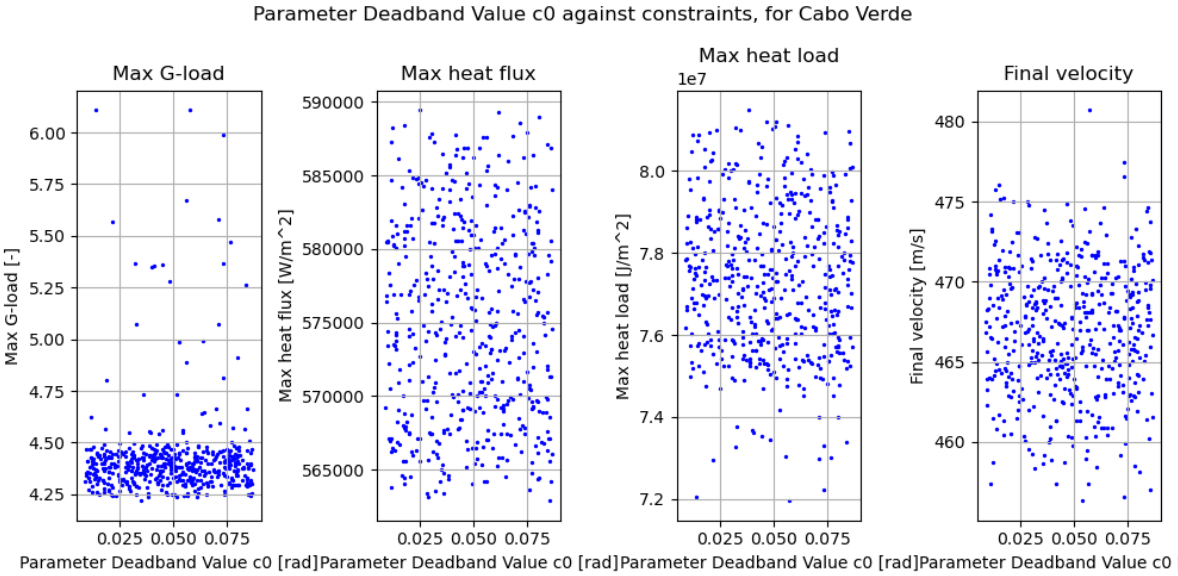

Figure 5.54: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for Cabo Verde.

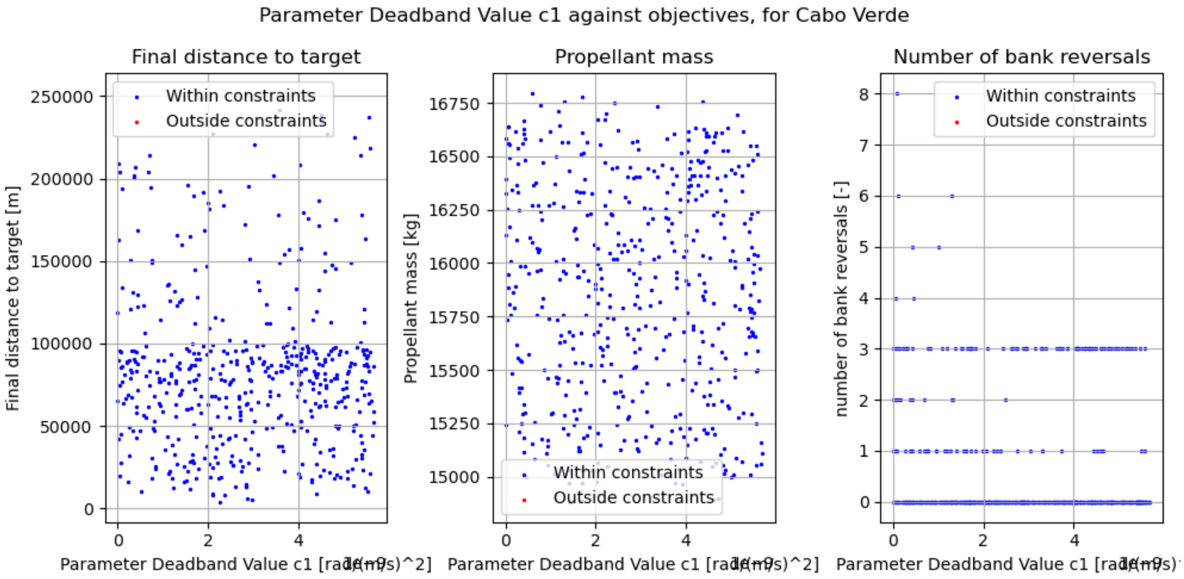

Figure 5.55: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for Cabo Verde.

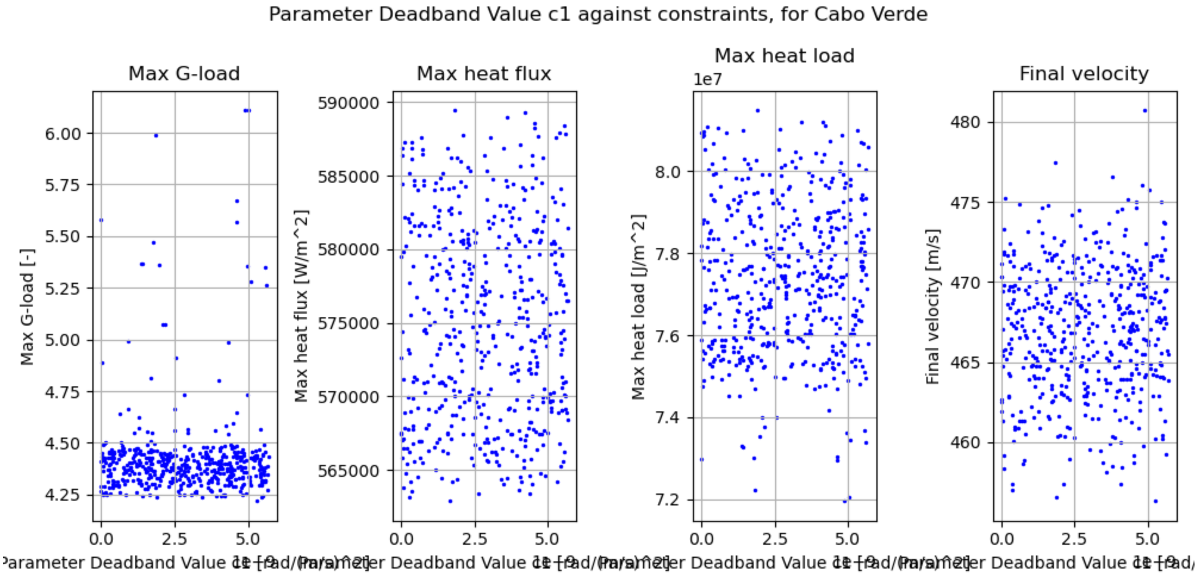

Figure 5.56: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for Cabo Verde.

Using these plots the conclusions that were drawn from the one-at-a-time MC analysis can be revisited for each of the input parameters:

- Increases to the initial velocity reduce the spread of distances to the target locations, which is believed to be because at higher velocities the vehicle has more control to guide itself to the target. The propellant mass generally increases with increasing velocity, but this is not the greatest effect, as the spread is quite loose. The effect on the number of bank angles is difficult to quantify. Increasing initial velocity initially drops the peak g-load and heat flux, but the former rises quickly at higher velocities due to the steeper dives the guidance performs. Heat load again has the opposite behaviour, increasing with increasing velocity before occasionally decreasing quickly at high velocities. Once again the final velocity increases with increasing initial velocity.

- The effect of the heading angle on the final distance is much less pronounced in this MC analysis, there still seems to be a minimum value, but there is a significant spread even there. The effect on the propellant mass is very clear, however, heading angle appears to be the dominant factor in the required propellant mass for the initial correction. It is expected the optimisation will attempt to minimise diversions from the initial heading angle. The number of bank angle reversals appears to increase with lower heading angles, with higher values tending towards zero. The effect on g-load, heat flux, heat load and final velocity does not appear to show much of a trend in this analysis.
- The initial flight path angle once again seems to have a very similar effect compared to the initial velocity, especially on the distance to target. An effect on the propellant mass or number of bank angle reversals cannot be distinguished. The max g-load increases in spread significantly at smaller values of the flight path angle. Conversely, the spreads of the heat flux and heat load seem to drop at smaller values, with the heat load increasing in range sharply towards the edge. The final velocity does seem to rise somewhat with smaller values of the flight path angle.
- The effect of the guidance overcontrol gain value is once again difficult to quantify. It has no distinguishable effect on any of the objectives or constraints in this analysis, which makes it difficult to anticipate what the optimisation may seek to do with it.
- The deadband values had no effect on most of the constraints and objectives, but now the expected behaviour can be seen for the number of bank angle reversals: reducing the value of the deadband values increases the number of bank angle reversals, as the vehicle bounces between the edges of the narrower lateral corridor.

During this MC analysis, it was found that some of the random combinations of inputs did achieve a range within 5 [km], indicating that it is possible for the simulation to achieve such accuracies. Furthermore, it was found that for Les Mureaux, the majority of input combinations resulted in trajectories that exceeded a constraint value. It is therefore doubtful if this location will result in a suitable outcome.

To summarise, it seems that the initial velocity, initial heading angle and initial flight path angle have clear effects on the final distance to the target location and the required propellant mass. Each of these have a certain value that achieves the best distance performance, where deviating in either direction from this value increases the range to the target. Each of these will likely be shifted away from these values to reduce the propellant mass required, while still remaining within the 5 [km] range. Of these three parameters, the heading angle has the most significant effect on the propellant masses required, with angles closer to the heading angle of the vehicle before the correction having dramatically lower propellant masses, in the order of 1000 [kg] throughout the investigated range. It is expected that the optimisation will attempt to use such heading angles, and compensate for it using the other parameters. This is followed by the initial velocity, which has less of an effect on the propellant mass compared to the heading angle, but still a significant one, in the order of 100 [kg] throughout the investigated range. Where possible, the optimisation should seek to decrease the velocity, so long as this does not require worse heading angles. Finally, the flight path angle has an insignificant impact on the propellant mass, in the order of only 1 [kg] throughout the investigated range. It does however significantly impact the constraints and the range to the target location, and so this parameter likely can be used to compensate for other parameters taking a value for better propellant mass. The mission will therefore likely have a more shallow entry, lower velocity, and a heading angle closer to the initial state before the correction. Finally, the effect of the guidance inputs is difficult to quantify. It appears these typically, but not always, strongly impact the number of bank angle reversals, with especially the deadband values causing more of these for lower values. These parameters likely have a best value for each specific case, which could cause errors if the vehicle's initial state is subjected to uncertainties.

5.10. Optimisation

With the MC analysis completed, it is time to move on to the optimisation. To begin with, the range to the target location is changed into a constraint at this point, as this felt more appropriate for the mission: no substantial gain is likely to be made by approaching the target more closely, so long as the vehicle ends within the range, it can reach the target during its final descent. This produces a multi-objective optimisation problem, with six inputs, five constraints, and two objectives. In subsection 5.10.1 the used input bounds and fitness and penalty functions are described. Finally, subsection 5.10.2 shows how the optimisation algorithm was selected.

5.10.1. Input bounds and Fitness and Penalty Functions

For the optimisers available in PyGMO to perform the optimisation, they must be provided with input bounds and a fitness function. The algorithm will then produce a set of inputs within the bounds, assess their fitness using the fitness function, and based on the relative fitnesses of various inputs, produce new inputs. Each of these cycles is known as a generation, and each input set producing an output is known as a population, or pop. Using the MC analysis performed in section 5.9, input bounds are determined for each target location such that they capture a reasonably wide decision space while eliminating what does not seem to produce feasible outcomes. The input bounds can be found in Table 5.5.

Table 5.5: The input bounds used in the optimisation.

Input Parameter	Input Bound for Natal	Input Bound for Cabo Verde	Input Bound for The Canarias	Input Bound for The Azores	Input Bound for Les Mureaux
V [m/s]	$[6400, 6500]$	$[6925, 7000]$	$[7225, 7325]$	$[7325, 7425]$	$[7475, 7540]$
ψ [deg]	$[124.9, 127.2]$	$[67.61, 69.33]$	$[49.27, 51.57]$	$[30.37, 32.09]$	$[33.80, 36.10]$
γ [deg]	$[-0.974, -0.573]$	$[-0.974, -0.573]$	$[-0.974, -0.573]$	$[-0.974, -0.573]$	$[-0.974, -0.573]$
K [-]	$[1.0, 11.0]$	$[1.0, 11.0]$	$[1.0, 11.0]$	$[1.0, 11.0]$	$[1.0, 11.0]$
c_0 [deg]	$[0.286, 5.157]$	$[0.286, 5.157]$	$[0.286, 5.157]$	$[0.286, 5.157]$	$[0.286, 5.157]$
c_1 [deg/(m/s^2)]	$[2.864 \cdot 10^{-8}, 3.438 \cdot 10^{-7}]$	$[2.864 \cdot 10^{-8}, 3.438 \cdot 10^{-7}]$	$[2.864 \cdot 10^{-8}, 3.438 \cdot 10^{-7}]$	$[2.864 \cdot 10^{-8}, 3.438 \cdot 10^{-7}]$	$[2.864 \cdot 10^{-8}, 3.438 \cdot 10^{-7}]$

In addition to the input bounds, the fitness function must be specified. This function takes the inputs and performs a simulation of the vehicle trajectory, then returns a fitness value for each objective. The algorithms attempt to minimise the value of fitness, and so a better performance should correspond to a lower fitness value. The definition of the fitness values for the propellant mass and number of bank angle reversals can be found in Equation 5.3 and Equation 5.4 respectively. These definitions should keep the fitness values away from zero, which was a concern for the number of bank angle reversals, avoiding possible errors due to the penalty functions having no effect. Furthermore, these should keep the fitness values roughly within the same order of magnitude, which is convenient.

$$f_{m_p} = \frac{m_p}{10000} \tag{5.3}$$

$$f_r = r + 1 \tag{5.4}$$

Finally, the constraint values should be incorporated. This is done by adding a penalty value to the fitness if the vehicle violates a constraint. Constraint penalty values can be absolute, where exceeding the constraint at all causes a complete invalidation of the solution, modelled by adding a very large value to the fitness function. This does mean that any information on that solution is lost entirely and not used by the optimisation. This is only applied in one case for this optimisation and not actually for one of the set constraints: if the simulation did not complete, the solution is invalidated in this way.

Potentially more useful are proportional constraint penalties, where violating one of the constraints imposes a proportional penalty on the solution. In this way, solutions that exist just over the limit of the constraints are still considered in the optimisation. The downside here is that it is possible for solutions to be produced which actually exceed a constraint. This method of applying penalties is used for each constraint. For the peak g-load, peak heat flux, total heat load and final velocity, a fairly strict penalty is applied, as exceeding these could be fatal for the mission. For each of these, if the constraint is violated, the fitness is multiplied by the ratio of the value attained to the value of the constraint, multiplied by 100. The distance to the target is seen as a less critical constraint, as exceeding it does not necessarily mean the mission will be lost, and so this penalty is computed simply by the ratio of the final range to the target to the range constraint. This means the range constraint can be more easily violated by the optimisation.

5.10.2. Optimiser selection

The final step to take before the optimisation can be performed is to select the optimisation scheme to be used. Six optimisation algorithms were considered for this optimisation problem:

- Improved Harmony Search (IHS).
- Non-dominated Sorting Genetic Algorithm (NSGA2).
- Multi-Objective Evolutionary Algorithm by Decomposition (MOEAD).
- Multi-Objective Evolutionary Algorithm by Decomposition Generational (MOEAD_GEN).
- Multi-objective Ant Colony Optimisation (MACO).
- Non-dominated Sorting Particle Swarm Optimisation (NSPSO).

To choose an optimisation algorithm to use, a smaller optimisation batch is run, using 40 pops for 15 generations, for two different seeds, namely 42 and 22. The average fitness at each generation is then plotted, and the results are compared to one another. The desired outcome here is to have an average fitness that is reduced smoothly and quickly without rising again or displaying any strange behaviour, reaching a low final fitness. A good agreement between seeds is also desirable, as this would indicate the outcome is not dependent on the seed arbitrarily selected at the start. The optimisation selection plots can be found in Figure 5.57, Figure 5.58, Figure 5.59, Figure 5.60 and Figure 5.61.

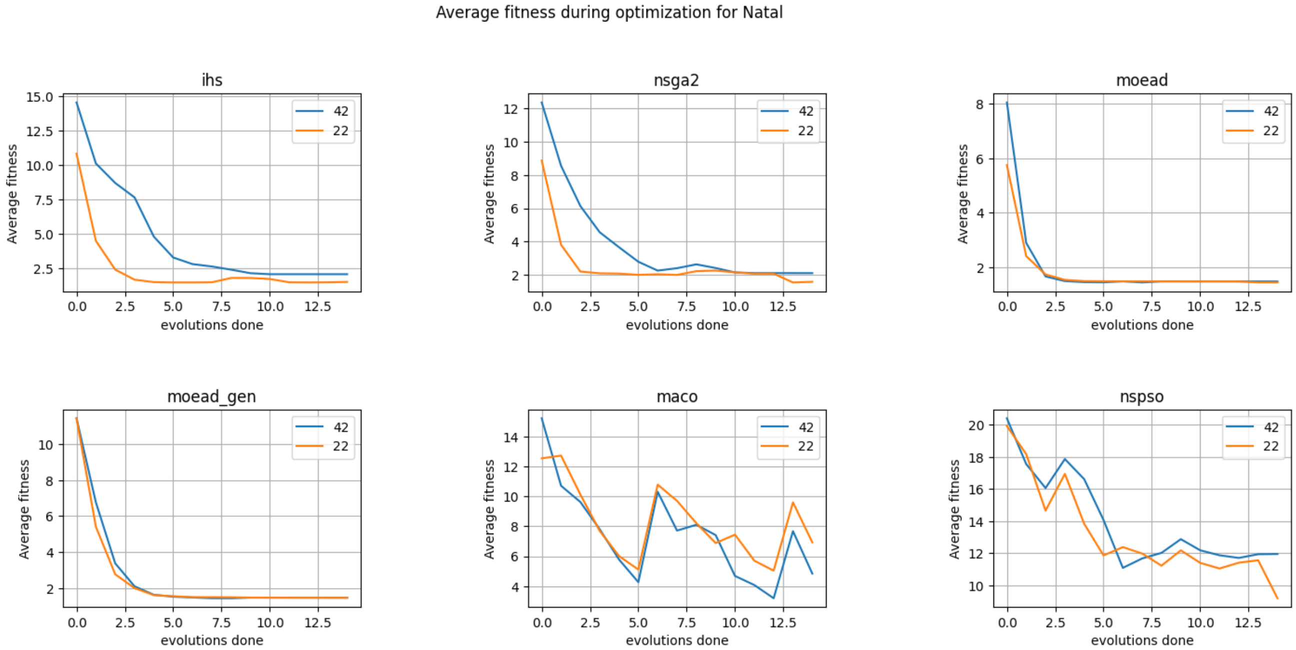

Figure 5.57: The average fitness over each generation, or evolution, for each of the six considered optimisation algorithms for Natal.

As can be seen from Figure 5.57, the MOEAD and MOEAD_GEN algorithms seem to perform well for Natal, smoothly and quickly settling to low values. The IHS and NSGA2 algorithms also have reasonable performance but show some rises and falls and notably seem to also be seed-dependent. MACO and NSPSO show erratic behaviour and do not reach fitness values as low as the others and so seem unsuitable for this location.

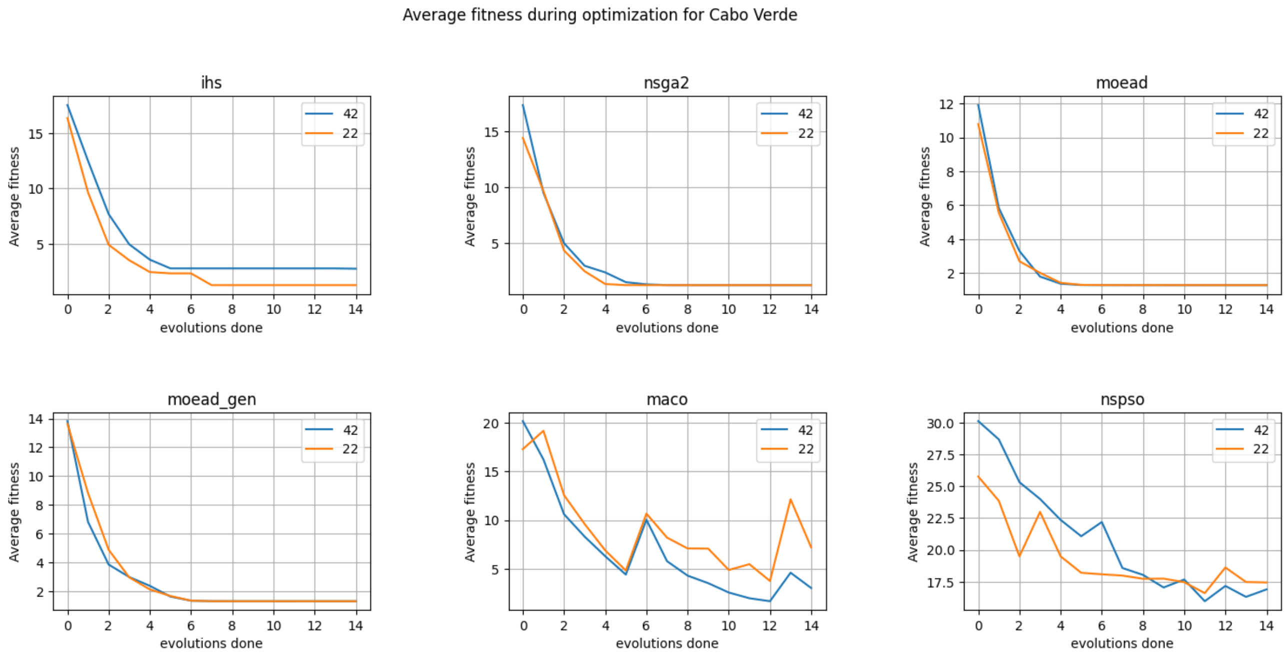

Figure 5.58: The average fitness over each generation, or evolution, for each of the six considered optimisation algorithms for Cabo Verde.

Figure 5.58 shows that once again the MOEAD and MOEAD_GEN algorithms perform well, and for this location, NSGA2 behaves better than for Natal. IHS once again does not have good agreement between seeds but does not rise in later evolutions for this location. Once again MACO and NSPSO seem unsuitable.

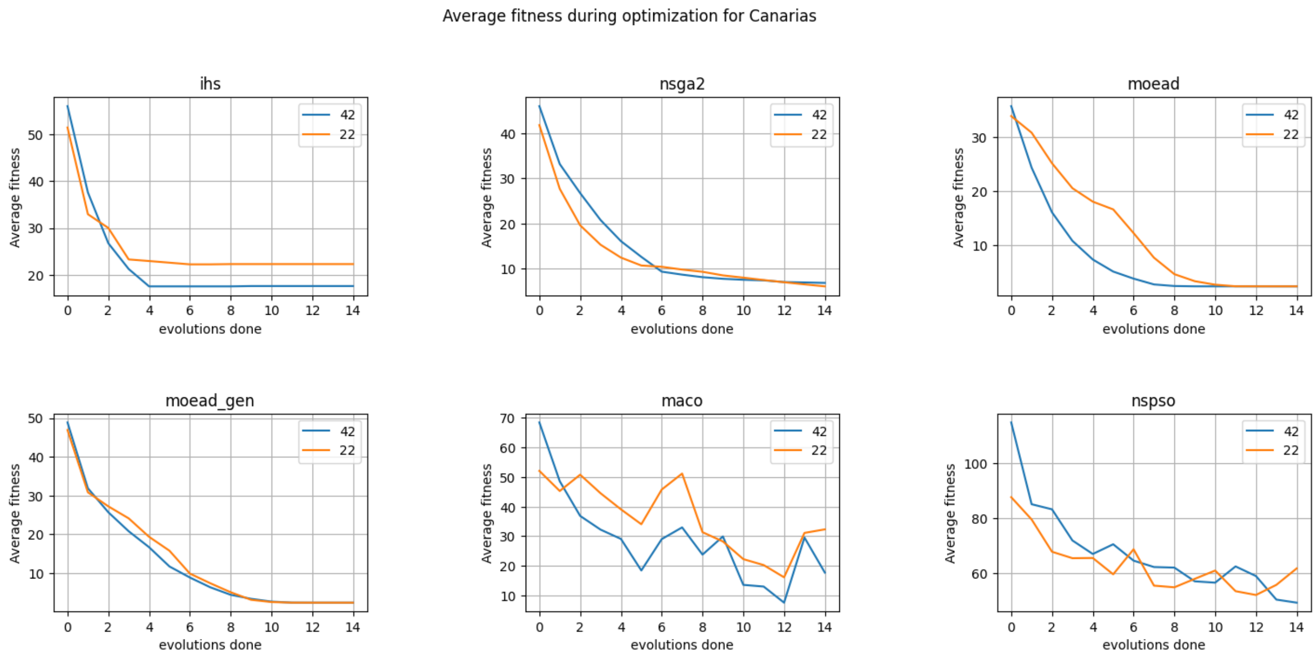

Figure 5.59: The average fitness over each generation, or evolution, for each of the six considered optimisation algorithms for The Canarias.

In Figure 5.59, it can be seen that MOEAD_GEN and MOEAD perform well again, but this time MOEAD has a significant split between its seeds at first, only later settling to the same value. NSGA2 does reasonably as well here, but shows some differences between seeds and is notably still reducing somewhat at the end. IHS shows good performance for only one seed. MACO and NSPSO once again show erratic behaviour and bad fitness scores, which is not desirable.

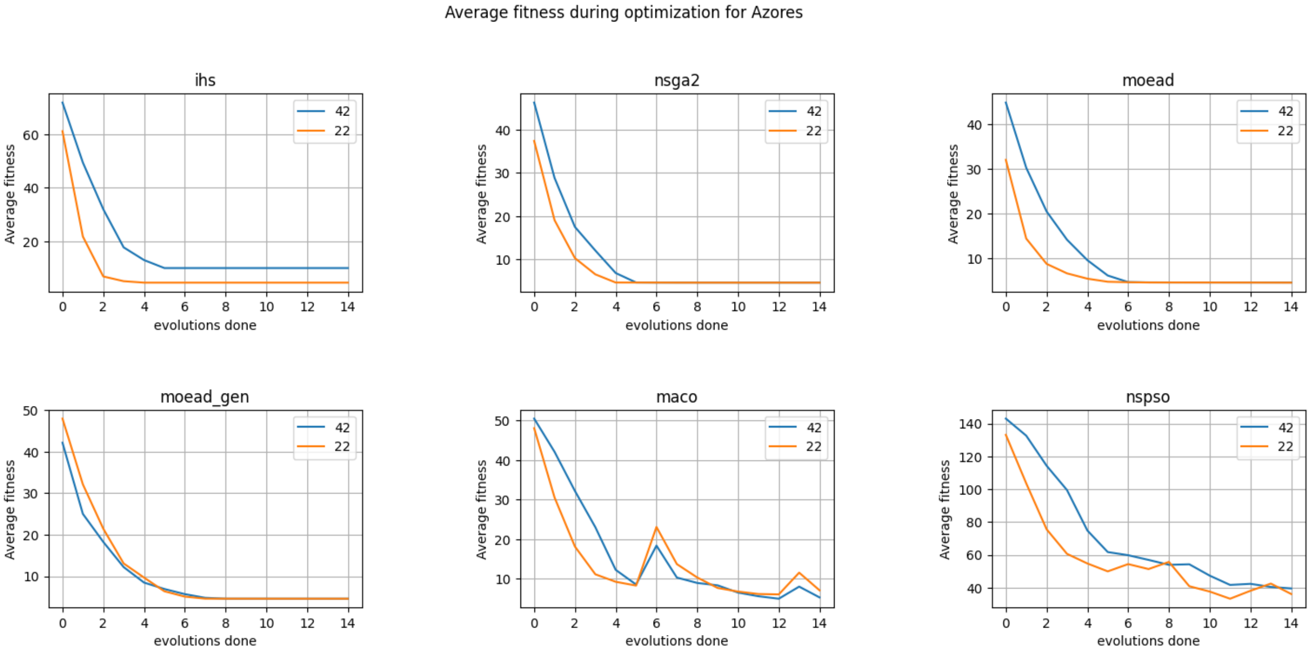

Figure 5.60: The average fitness over each generation, or evolution, for each of the six considered optimisation algorithms for The Azores.

For The Azores, MOEAD, MOEAD_GEN and NSGA2 show good performance, as can be seen in Figure 5.60. IHS once again has good performance depending on the seed used. MACO and NSPSO are less erratic compared to other locations, but still not as good as the other algorithms.

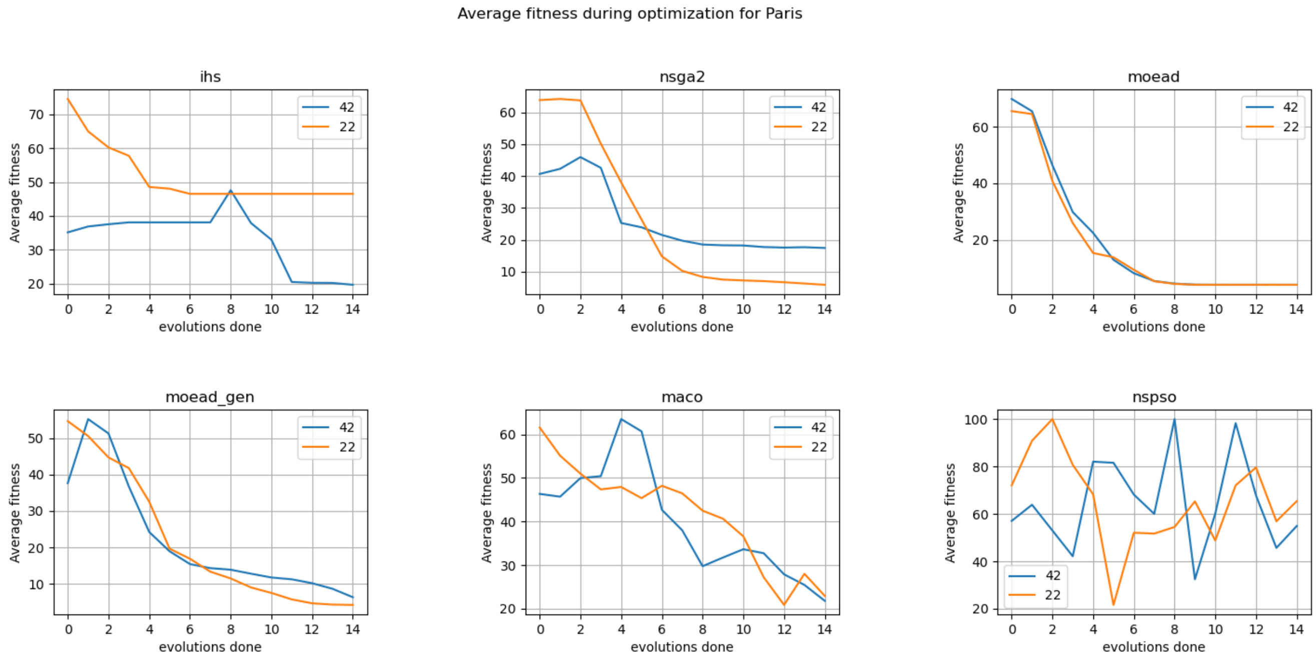

Figure 5.61: The average fitness over each generation, or evolution, for each of the six considered optimisation algorithms, for Les Mureaux.

Finally, as can be found in Figure 5.61, the algorithms generally show the worst behaviour for Les Mureaux. Really only MOEAD behaves as desired. MOEAD_GEN shows somewhat different performance based on seeds but does eventually reach similar low fitness values. NSGA2 shows very different behaviour depending on the seed. IHS, MACO and NSPSO do not work well at all for this location.

Upon considering the behaviour of the algorithms, the decision was made to use the MOEAD_GEN algorithm. It showed the best behaviour for each of the target locations except for Les Mureaux, generally achieving slightly better average fitness values. As for Les Mureaux, the MC analysis showed that the vast majority of simulations violated some constraint, doubts existed whether this location would yield a feasible option, and so the poor behaviour for this location is deemed to be acceptable.

5.10.3. Optimisation

With all the preparation done, the optimisation can be performed. For the full optimisation, 400 pops are evolved over 20 generations for five different seeds: 42, 22, 96, 35 and 11. The resulting fitnesses from the final generation, the average fitness over time and the best propellant mass fitness at each generation can be found in Figure 5.62, Figure 5.63, Figure 5.64, Figure 5.65, Figure 5.66, Figure 5.67, Figure 5.68, Figure 5.69, Figure 5.70 and Figure 5.71.

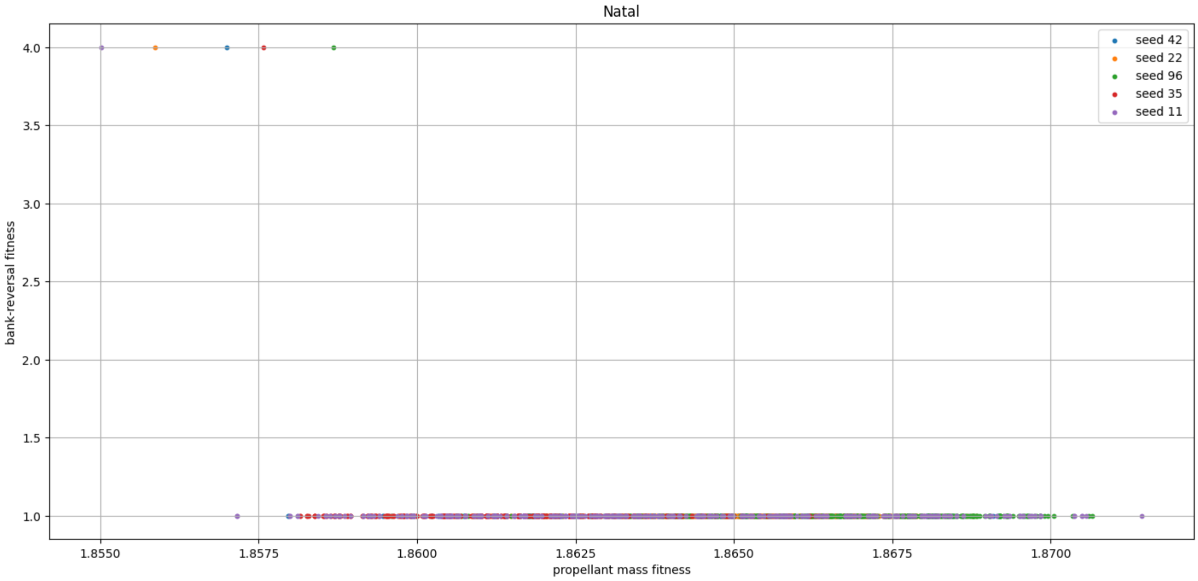

Figure 5.62: The propellant mass and number of bank angle reversals fitness of the final generation for Natal.

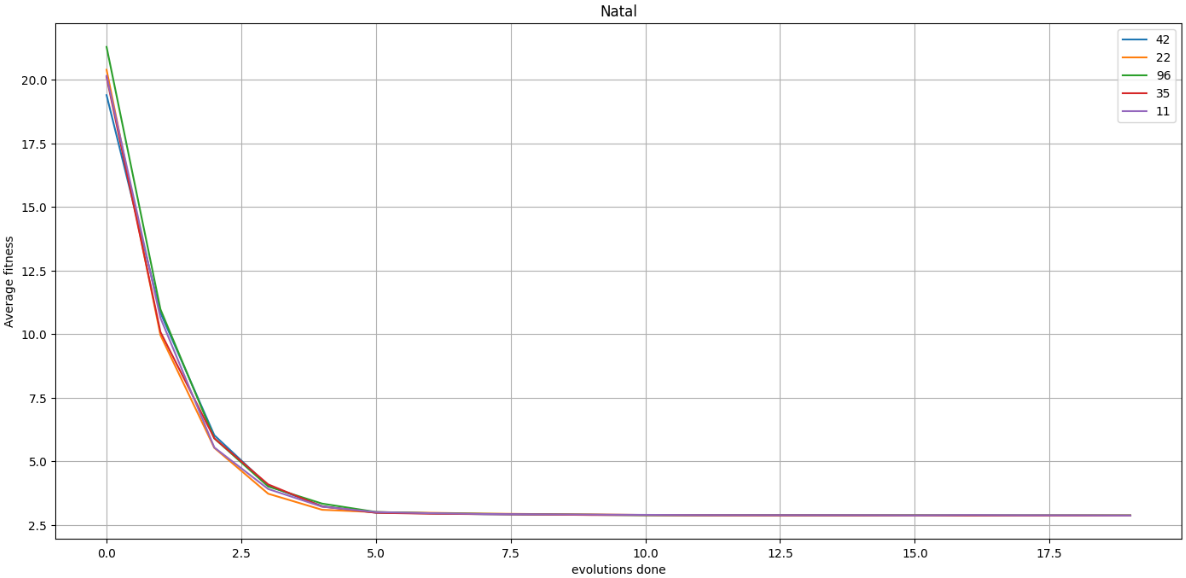

Figure 5.63: The average fitness over generations for Natal.

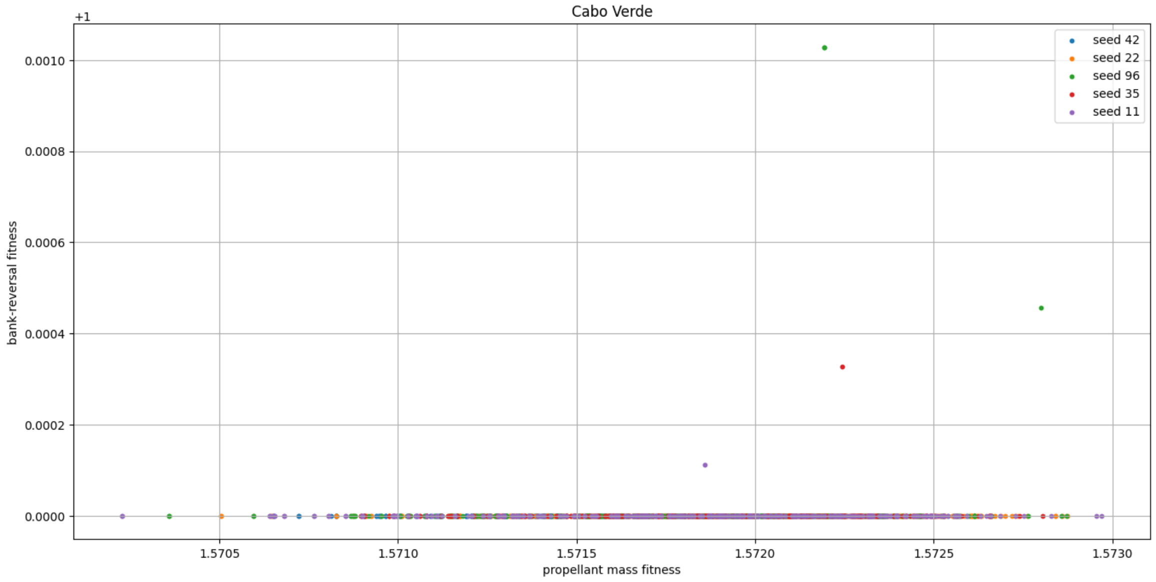

Figure 5.64: The propellant mass and number of bank angle reversals fitness of the final generation for Cabo Verde.

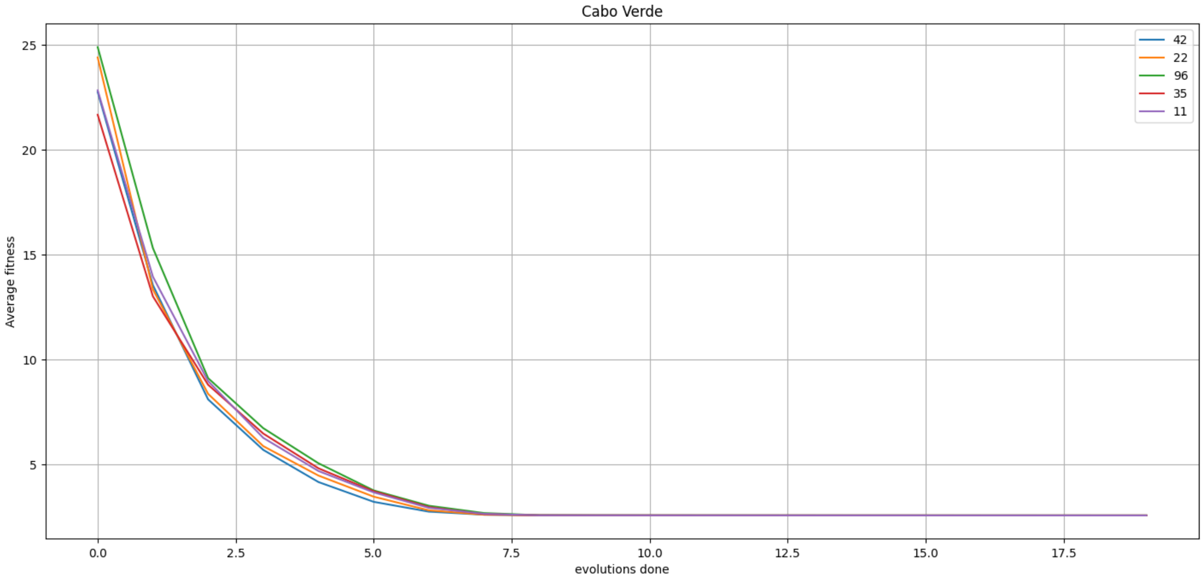

Figure 5.65: The average fitness over generations for Cabo Verde.

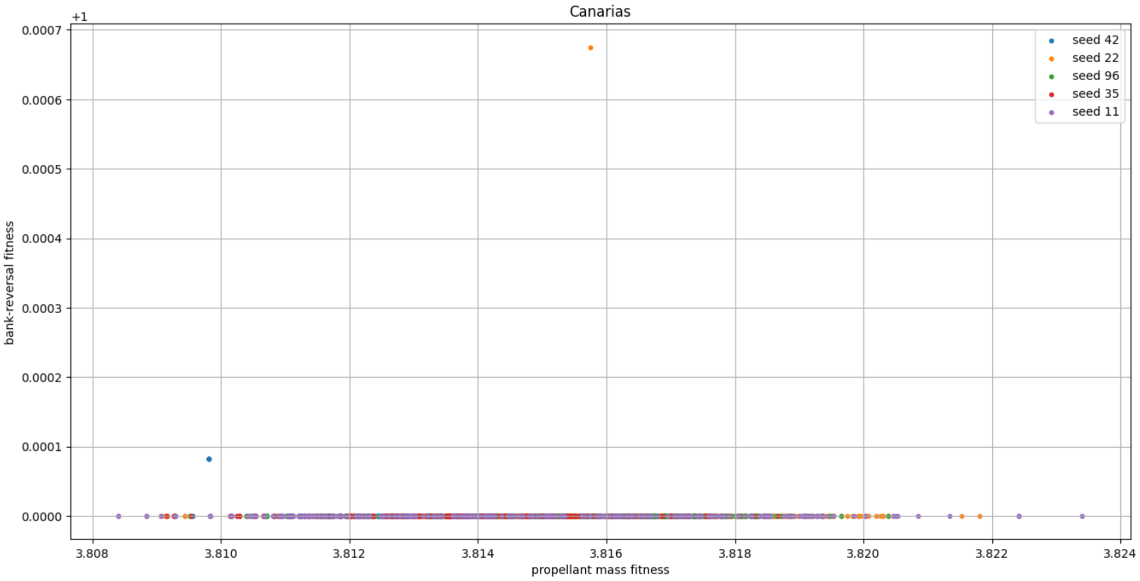

Figure 5.66: The propellant mass and number of bank angle reversals fitness of the final generation for The Canarias.

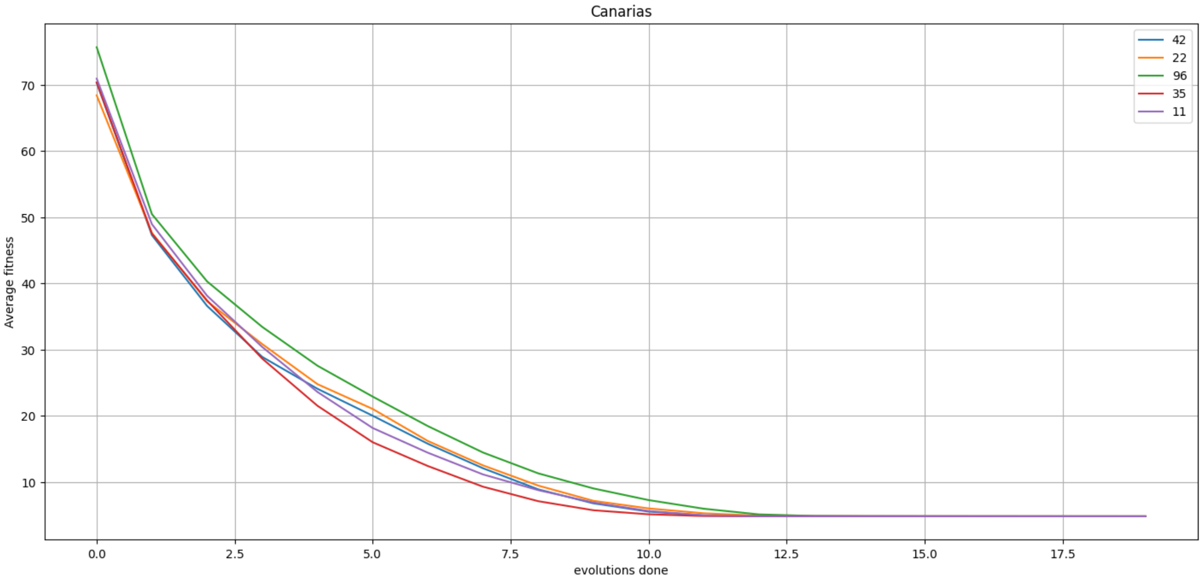

Figure 5.67: The average fitness over generations for The Canarias.

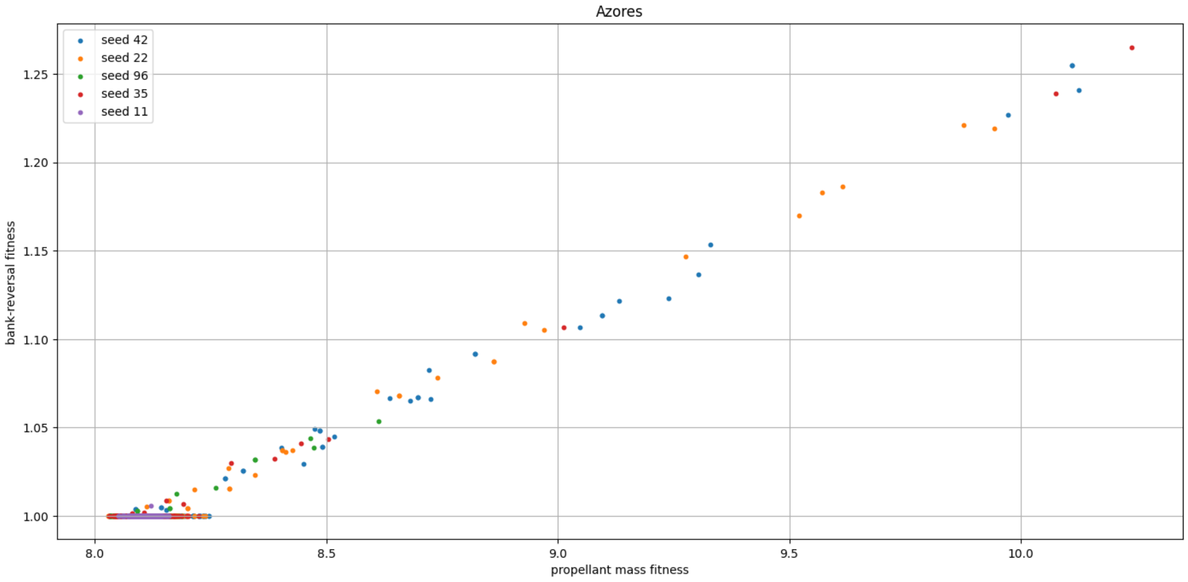

Figure 5.68: The propellant mass and number of bank angle reversals fitness of the final generation for The Azores.

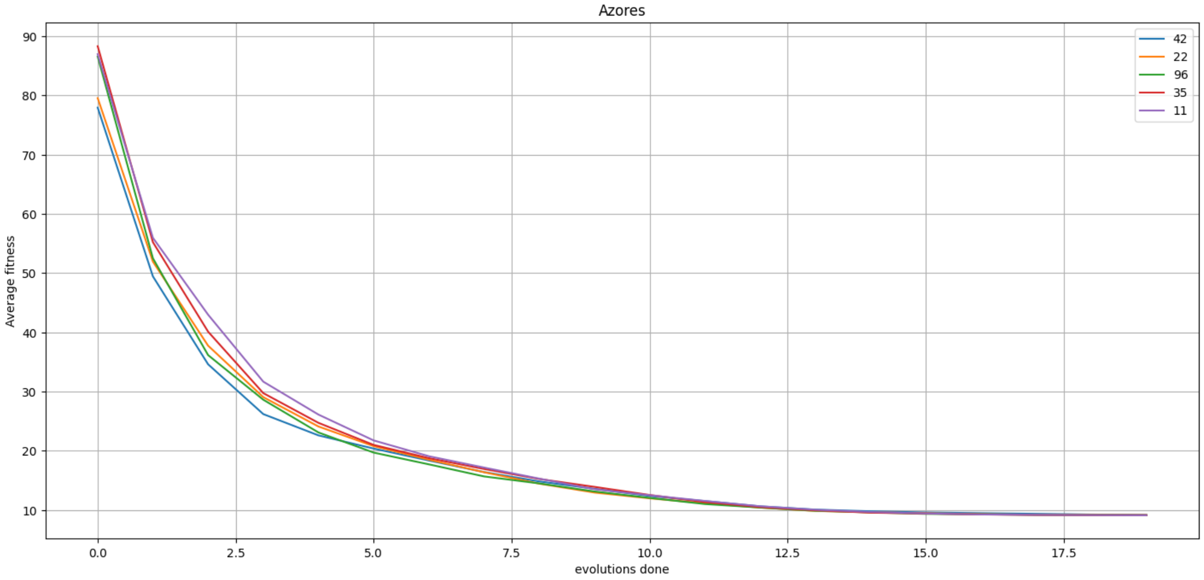

Figure 5.69: The average fitness over generations for The Azores.

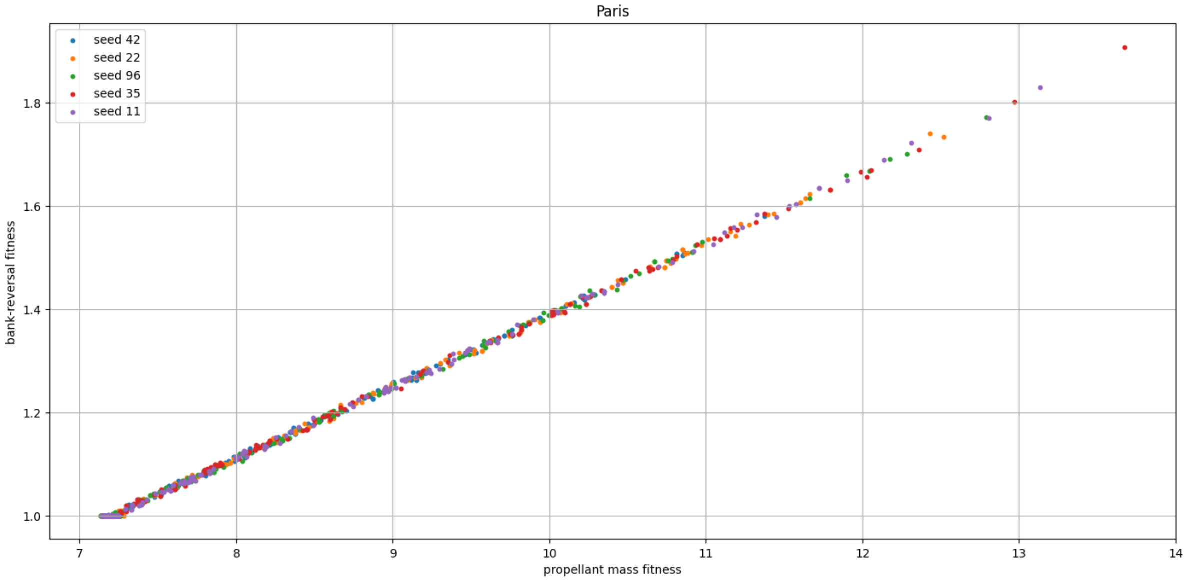

Figure 5.70: The propellant mass and number of bank angle reversals fitness of the final generation for Les Mureaux.

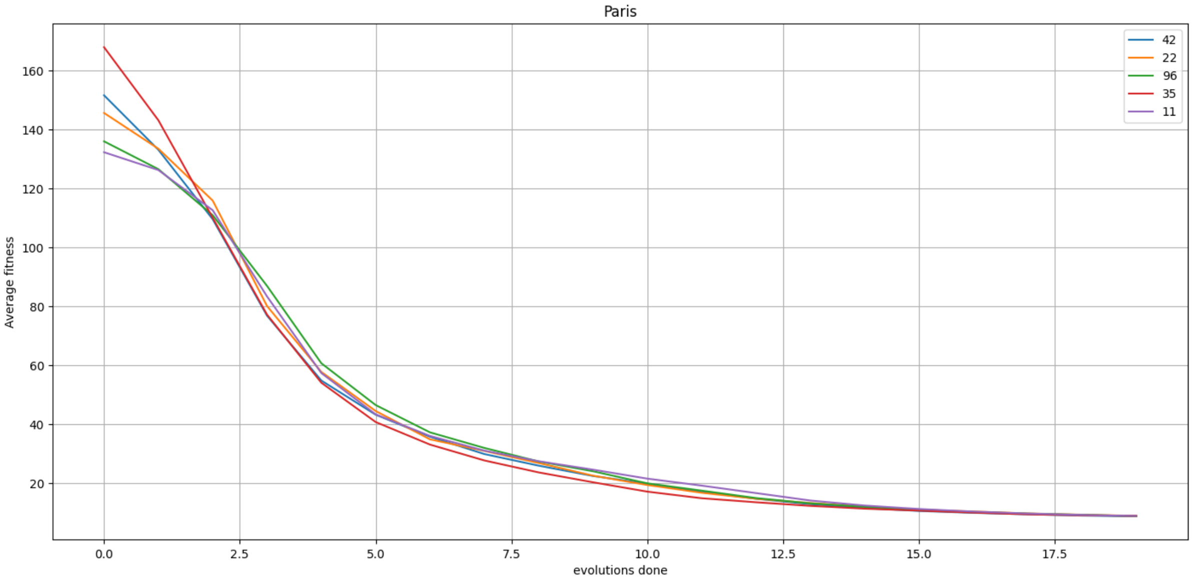

Figure 5.71: The average fitness over generations for Les Mureaux.

From the fitness plots Figure 5.62, Figure 5.64, Figure 5.66, Figure 5.68 and Figure 5.70, it can be seen that the pops do not form a Pareto front as one would typically expect. It would appear that the objectives do not in fact compete with one another very much, often allowing the best propellant mass fitness to achieve the best possible bank angle reversal fitness value. Even for Natal, where the best propellant mass option has a bank angle fitness of 4, it appears they do not compete with one another. In Figure 5.72, it can be seen that the actual reversal of the bank angle is very minor compared to the bank angle adjustments made by the guidance algorithm. It seems that the vehicle does not actually perform major bank angle reversals throughout its trajectory. With the benefit of hindsight, the number of bank angle reversals may have been better suited as a constraint, making the process a single-

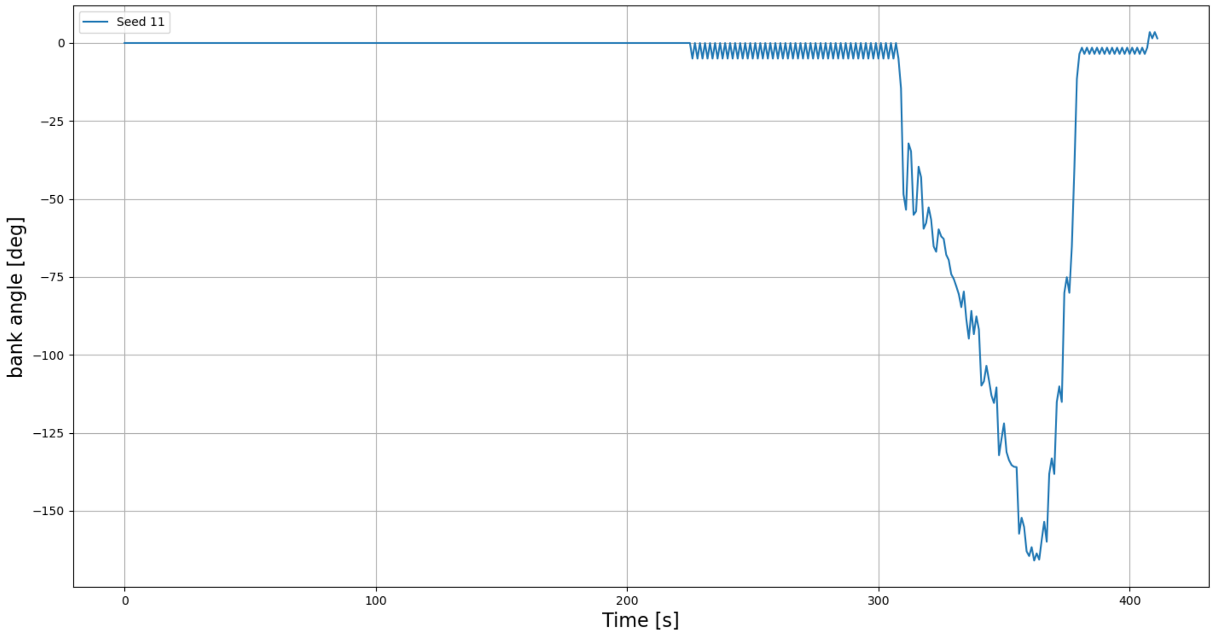

Figure 5.72: An example of a trajectory for Natal with higher bank angle reversal fitness.

objective optimisation.

Regardless, the outcomes produced still seem like good solutions for the propellant mass fitness, which is the objective that truly matters. This can be seen from the plots of the average fitness in Figure 5.63, Figure 5.65, Figure 5.67, Figure 5.69 and Figure 5.71. For each seed, the average value quickly drops to lower values where they settle, with each of the seeds settling to roughly the same value. Upon investigating this more closely, it appears the average value is still decreasing slightly. To investigate whether this is an improvement in the best value, the best propellant mass fitness at each generation is plotted in Figure 5.73, Figure 5.74, Figure 5.75, Figure 5.76 and Figure 5.77.

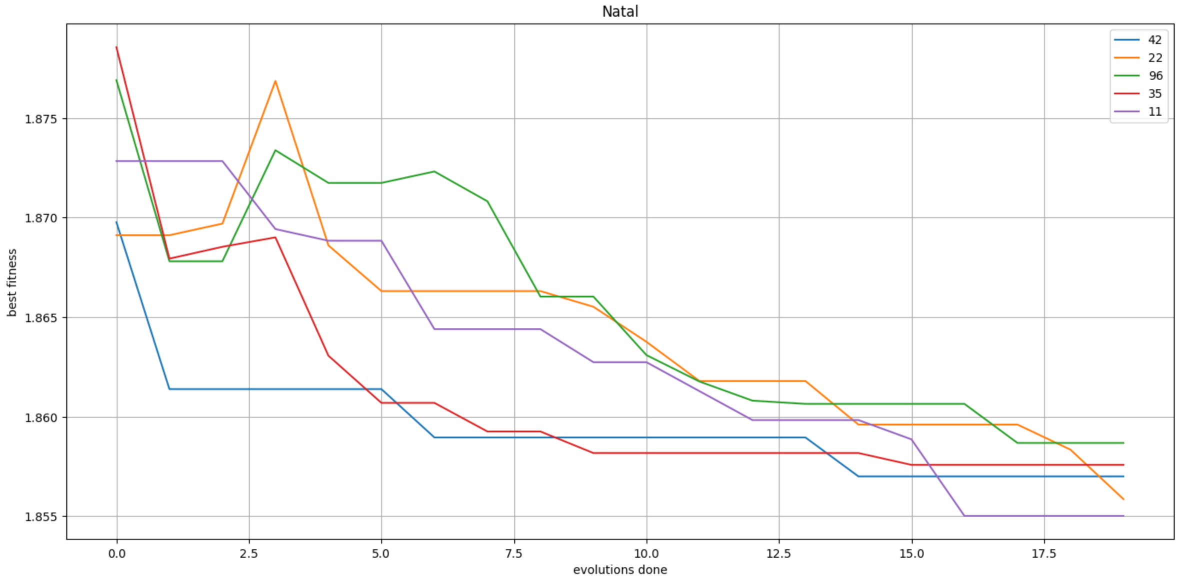

Figure 5.73: The best propellant mass fitness over generations for Natal.

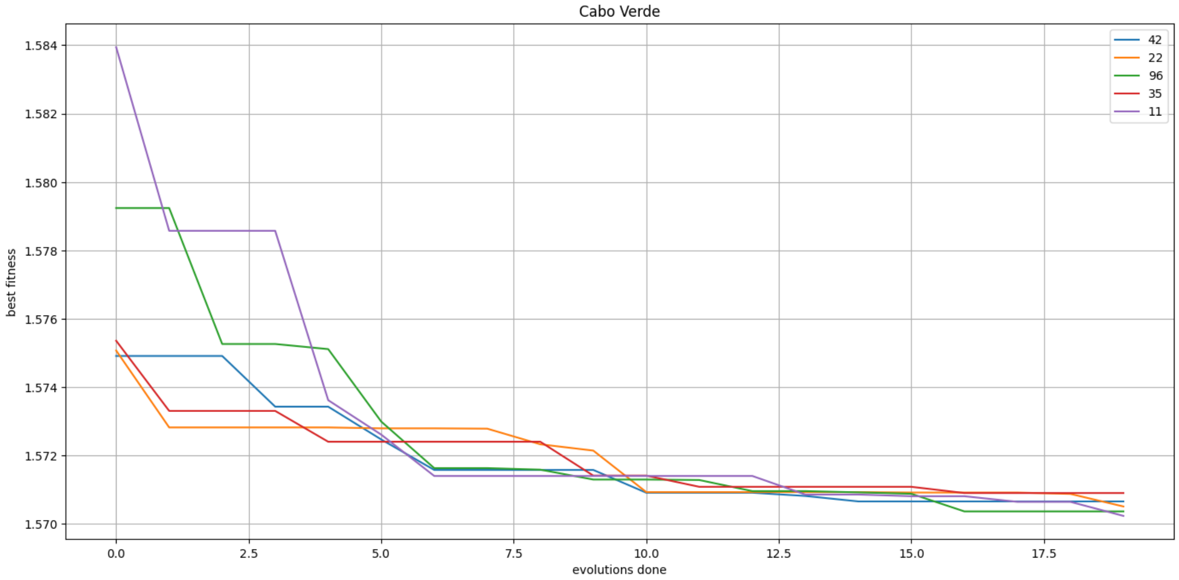

Figure 5.74: The best propellant mass fitness over generations for Cabo Verde.

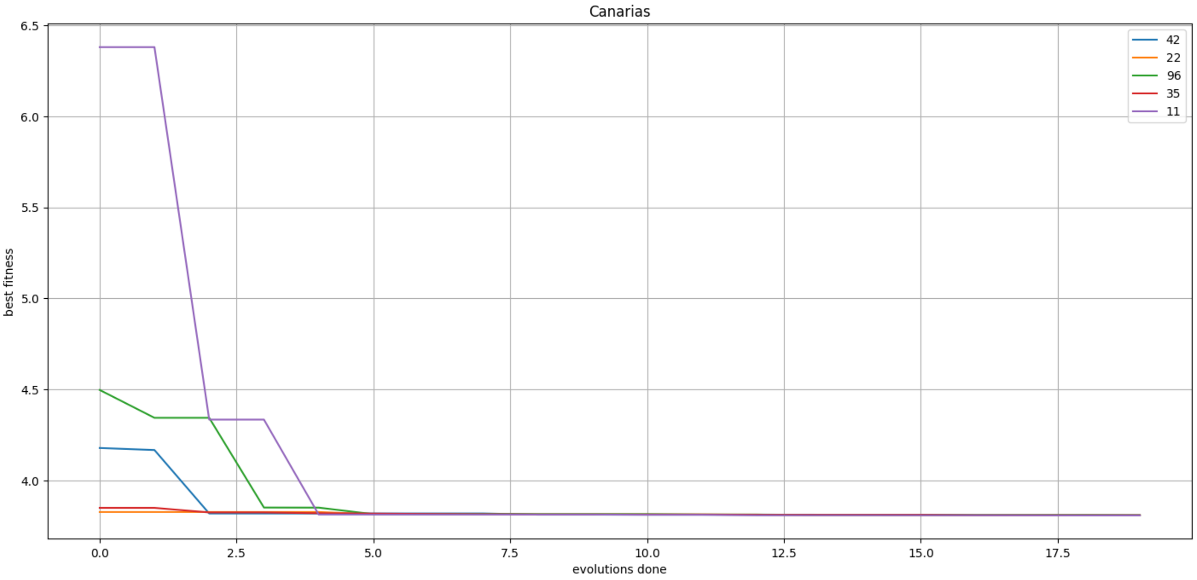

Figure 5.75: The best propellant mass fitness over generations for The Canarias.

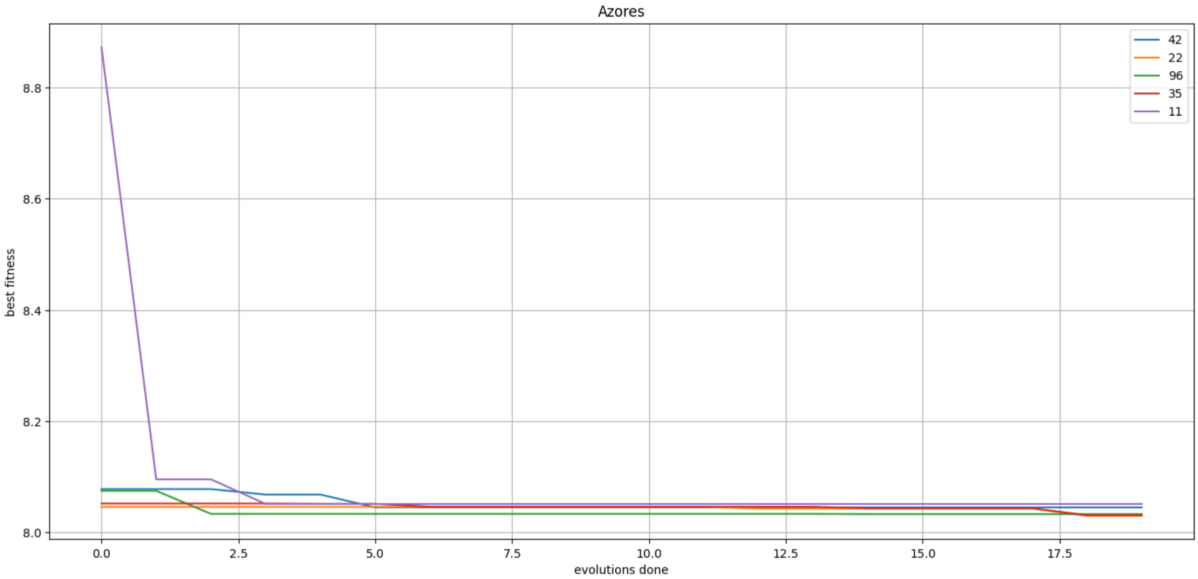

Figure 5.76: The best propellant mass fitness over generations for The Azores.

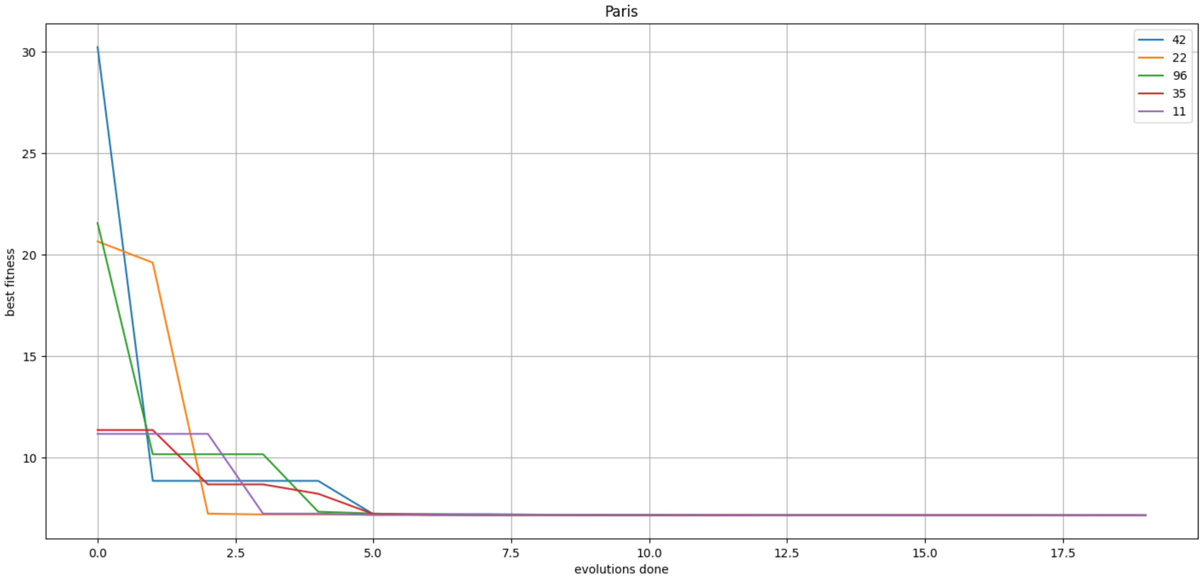

Figure 5.77: The best propellant mass fitness over generations for Les Mureaux.

It seemed like the best value for the fitness drops rapidly as well to a certain value, but does show minor improvements as the optimisation continues. To investigate whether it is worth it to add more generations to the optimisation, the optimisation is performed again for only one seed, 42, for 40 generations. The results for each target location were that the improvement in fitness was less than 0.1% of the value at 20 evolutions, with the most significant change being 72 [kg] for Azores. This is significant but not a major issue because, as can be seen in section 6.1, The Azores will be rejected as a target location even if this 72 [kg] improvement would be taken into account. The next most significant improvement was 11.2 [kg] for Natal, which is likely negligible compared to the uncertainty that exists around the mass estimates in subsection 5.1.2. The outcomes for 20 generations will therefore be used.

6

Results & Discussion

With the results generated by the optimisation, these can now be analysed. In section 6.1, the specifics of the outcomes of the optimisation are laid out. Then, in section 6.2, the sensitivity of the outcomes to uncertainties in the initial conditions is investigated. Following this, the costs of the missions from the optimisation will be estimated. Afterwards, in section 6.4, the outcomes from the optimisation are compared to alternative mission profiles presented in chapter 3, to determine if implementing the guided ballistic engine bay is worth considering. Finally, section 6.5 discusses the limitations of the research performed, which should be kept in mind with the results.

6.1. Optimal Trajectories

With the optimisation completed, a set of the five best identified solutions has been generated for each target location. These trajectories are compared to one another and to the initial guess for that target location in subsection 6.1.1, subsection 6.1.2, subsection 6.1.3, subsection 6.1.4 and subsection 6.1.5. At the end, subsection 6.1.6 compares the results from each target location and explains which will be used carrying forward.

6.1.1. Natal

Plots of various quantities of the trajectories for Natal can be found in Figure 6.1, Figure 6.2, Figure 6.3, Figure 6.4, Figure 6.5. An overview of how these trajectories compare can be found in Table 6.1.

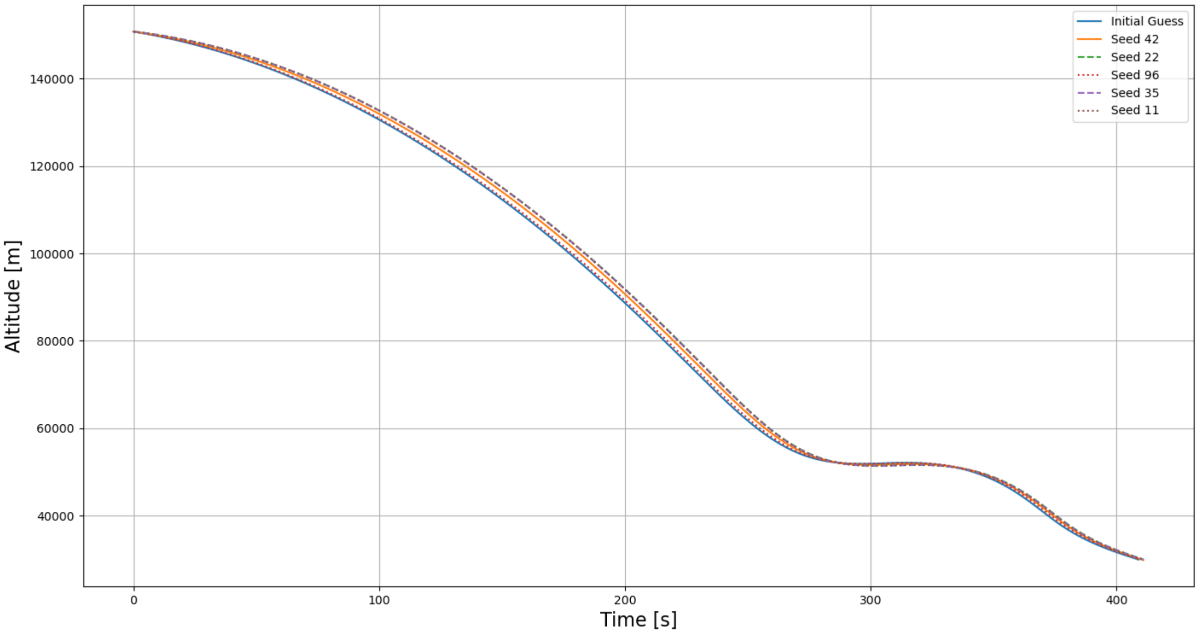

Figure 6.1: The altitudes over time of the best identified trajectories for each seed and the initial guess for Natal.

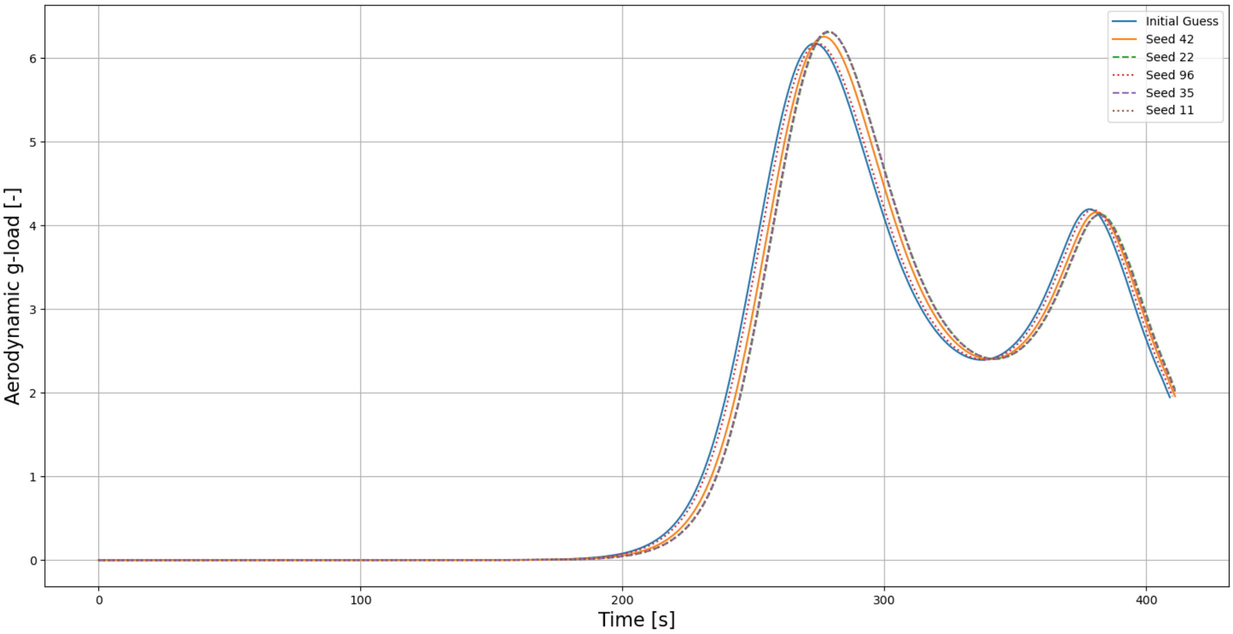

Figure 6.2: The g-load over time of the best identified trajectories for each seed and the initial guess for Natal.

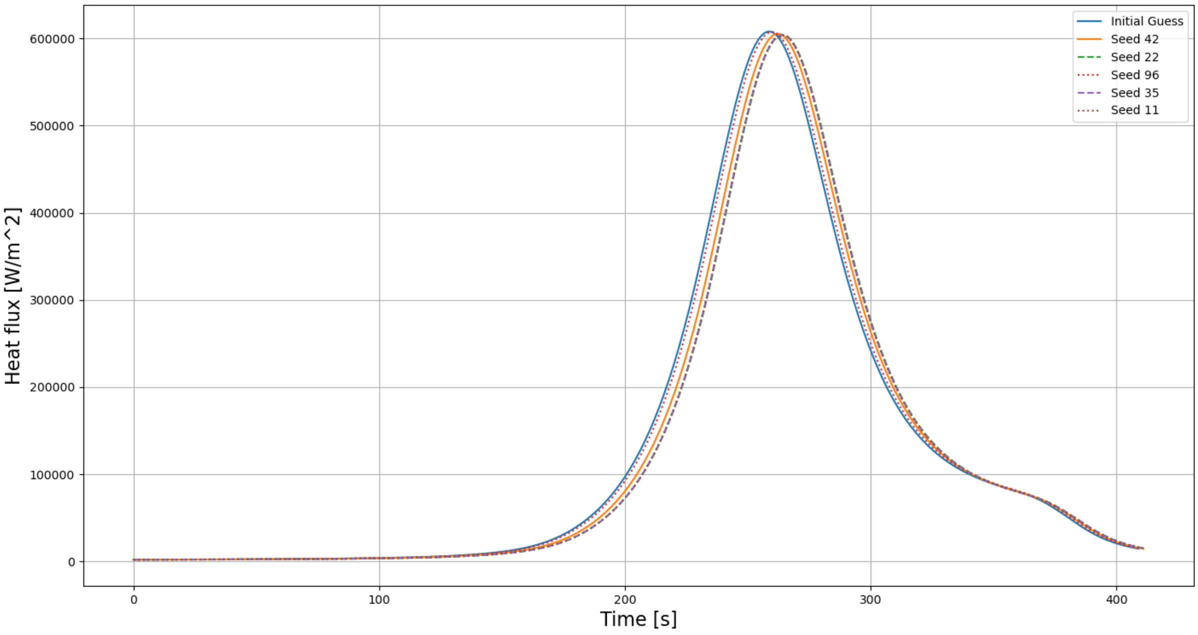

Figure 6.3: The heat flux over time of the best identified trajectories for each seed and the initial guess for Natal.

Figure 6.4: The latitude-longitude map of the best identified trajectories for each seed and the initial guess for Natal.

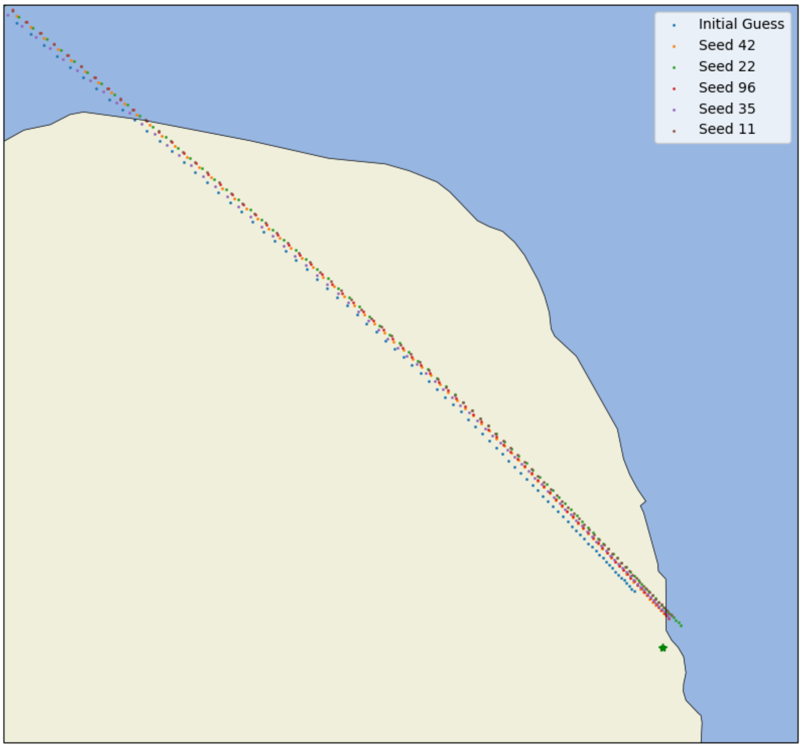

Figure 6.5: A close-up of the latitude-longitude map of the best identified trajectories for each seed and the initial guess for Natal.

Table 6.1: A comparison of values associated with the best identified trajectories for each seed and the initial guess for Natal. For the inputs, the value with respect to the initial guess is given.

Seed	Initial Guess	42	22	96	35	11
V [m/s]	6450	-29.59	-47.70	-1.394	-49.70	-49.63
ψ [deg]	126.0	-0.0870	-0.0795	-0.102	-0.059	-0.087
γ	-0.8	$+0.142$	$+0.227$	$+0.029$	$+0.226$	$+0.226$
K [-]	1.0	$+1.888$	$+3.988$	$+4.952$	$+7.758$	$+4.208$
c_0 [deg]	2.0	-1.527	-0.655	-1.094	-1.686	-1.713
c_1 [deg/(m/s^2)]	$1.633 \cdot 10^{-7}$	$-1.272 \cdot 10^{-7}$	$-1.341 \cdot 10^{-7}$	$-1.286 \cdot 10^{-7}$	$-1.171 \cdot 10^{-7}$	$-1.344 \cdot 10^{-7}$
peak g-load [-]	6.173	6.254	6.307	6.170	6.314	6.313
peak heat flux [MW/m^2]	0.6079	0.6054	0.6041	0.6074	0.6040	0.6040
total heat load [MJ/m^2]	51.07	50.20	49.67	51.01	49.62	49.62
final velocity [m/s]	562.6	559.75	574.9	565.2	570.2	570.5
Final range to target [km]	9.65	4.564	4.401	4.590	4.797	5.088
Propellant mass [kg]	18682.2	18570.0	18558.6	18586.8	18575.8	18550.2

As expected, compared to the initial guess, the best identified trajectories use a smaller initial velocity, a smaller diversion from the post-disconnection flight path (which is a negative heading angle difference for Natal) and a more shallow entry angle. The guidance overcontrol value of each solution is increased, although there is some variance in the amount by which that is done. Interestingly, the deadband values are both increased and decreased with respect to the initial guess. The peak g-load and heat flux are higher for each solution, whereas the total heat load and velocity are lower. Each of the load constraints stays well within the limit values. Each seed except for 11 resulted in a solution that arrived within 5 [km] of the target. Seed 22 resulted in the solution among these with the lowest propellant mass and so was selected as the final best identified trajectory for this location.

6.1.2. Cabo Verde

Plots of various quantities of the trajectories for Cabo Verde can be found in Figure 6.6, Figure 6.7, Figure 6.8, Figure 6.9, Figure 6.10. An overview of how these trajectories compare can be found in Table 6.2.

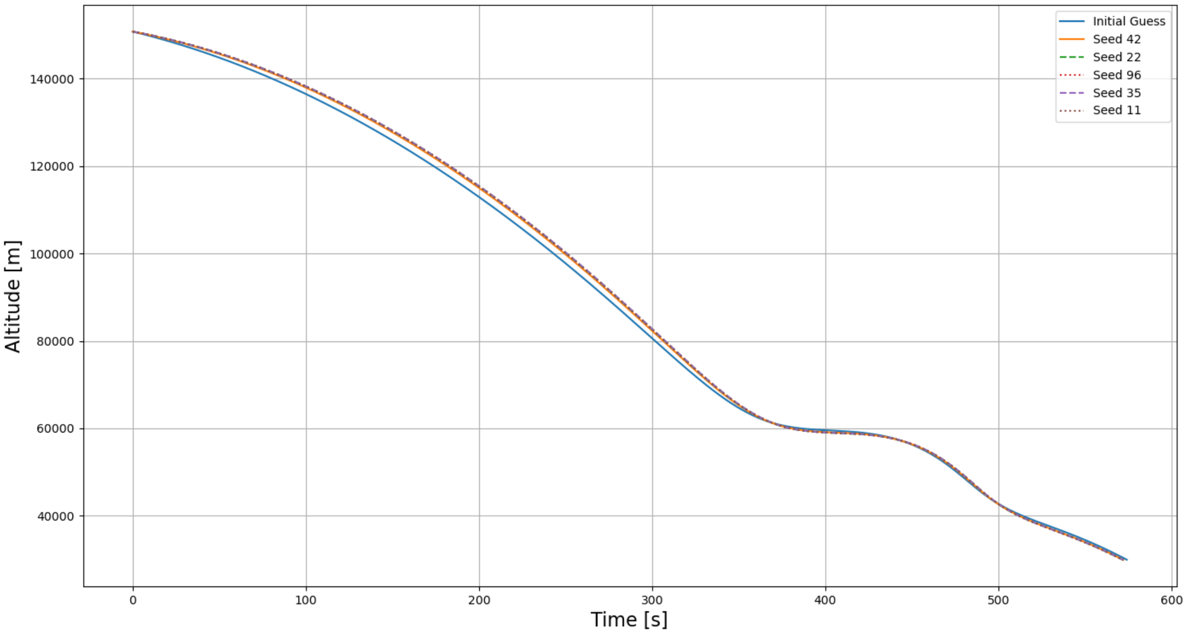

Figure 6.6: The altitudes over time of the best identified trajectories for each seed and the initial guess for Cabo Verde.

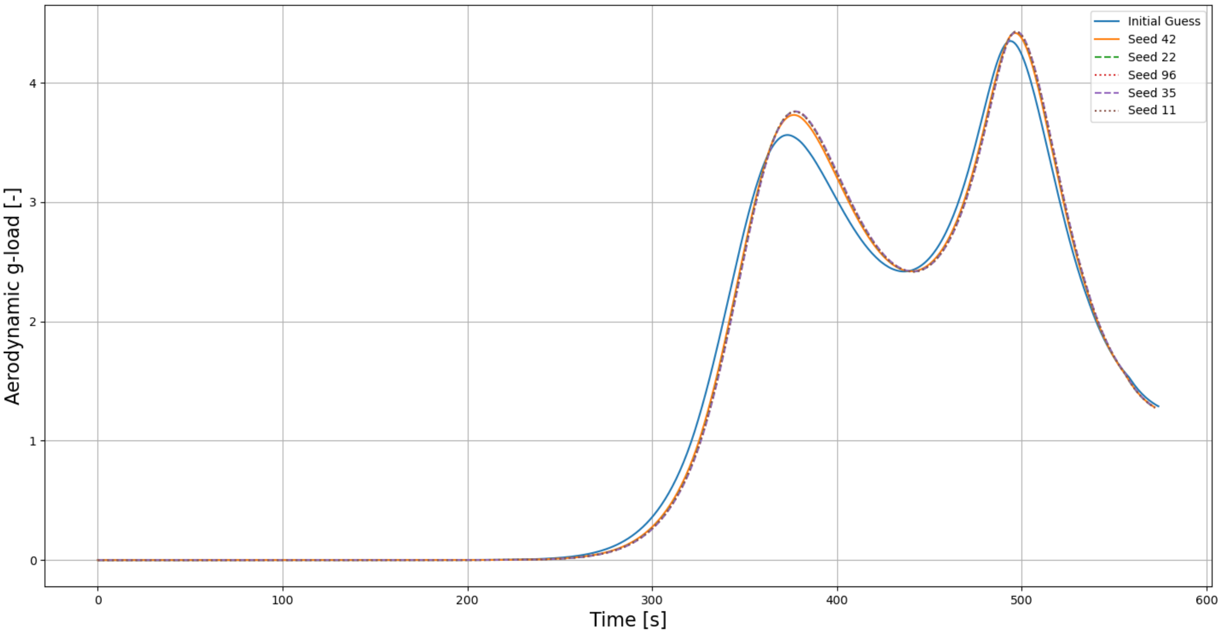

Figure 6.7: The g-load over time of the best identified trajectories for each seed and the initial guess for Cabo Verde.

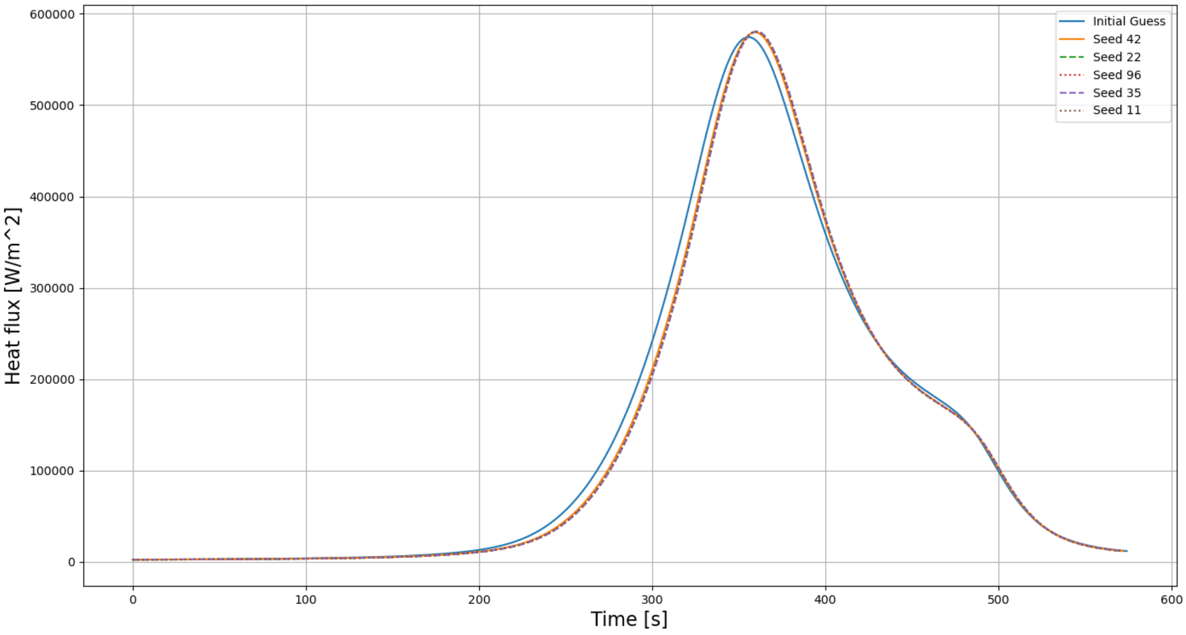

Figure 6.8: The heat flux over time of the best identified trajectories for each seed and the initial guess for Cabo Verde.

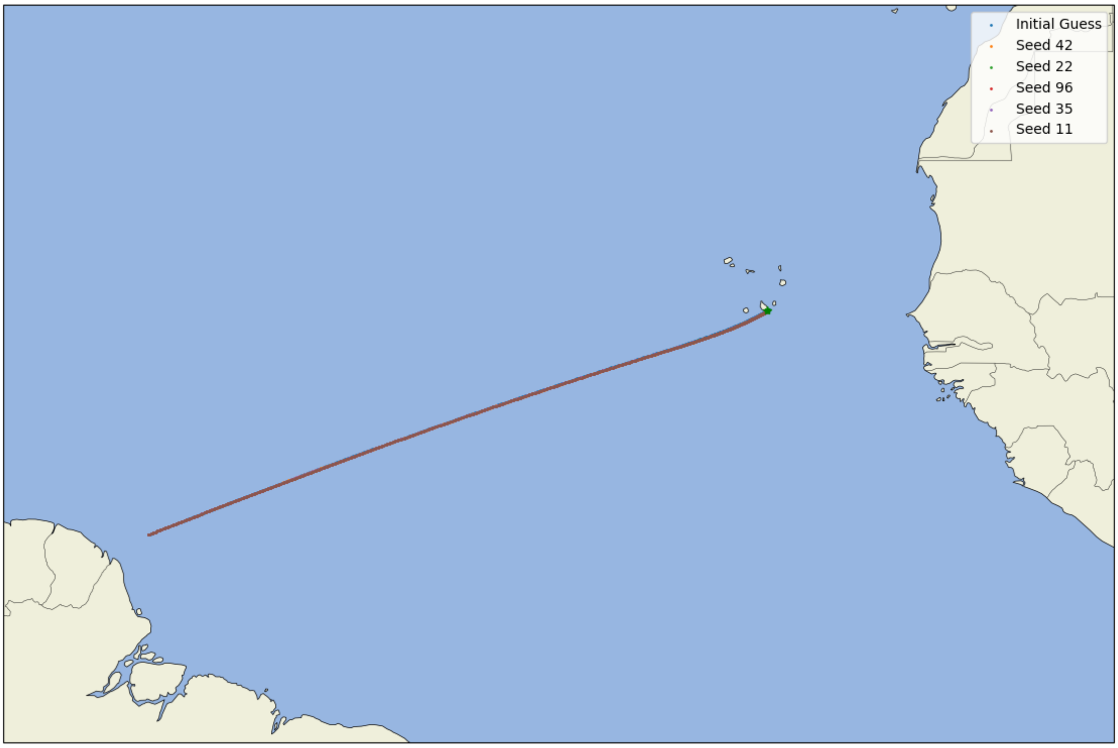

Figure 6.9: The latitude-longitude map of the best identified trajectories for each seed and the initial guess for Cabo Verde.

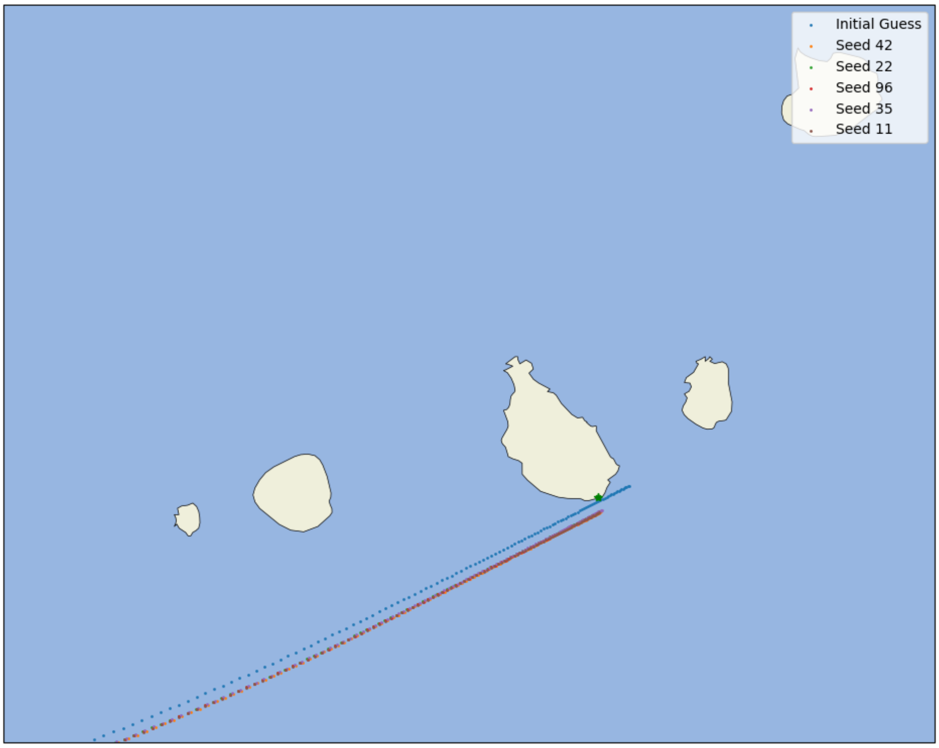

Figure 6.10: A close-up of the latitude-longitude map of the best identified trajectories for each seed and the initial guess for Cabo Verde.

Table 6.2: A comparison of values associated with the best identified trajectories for each seed and the initial guess for Cabo Verde. For the inputs, the value with respect to the initial guess is given.

Seed	Initial Guess	42	22	96	35	11
V [m/s]	6970	-37.44	-43.70	-43.63	-44.02	-43.41
ψ [deg]	68.5	$+0.0593$	$+0.0490$	$+0.0508$	$+0.0439$	$+0.0527$
γ	-0.8	$+0.151$	$+0.178$	$+0.179$	$+0.182$	$+0.177$
K [-]	1.0	$+0.3436$	$+2.057$	$+1.108$	$+1.300$	$+1.431$
c_0 [deg]	2.0	-1.575	-0.5938	$+2.840$	$+2.841$	-1.513
c_1 [deg/(m/s^2)]	$1.633 \cdot 10^{-7}$	$-8.128 \cdot 10^{-8}$	$+2.603 \cdot 10^{-8}$	$-1.269 \cdot 10^{-8}$	$+1.284 \cdot 10^{-8}$	$-1.338 \cdot 10^{-7}$
peak g-load [-]	4.350	4.417	4.427	4.427	4.428	4.429
peak heat flux [MW/m^2]	0.5745	0.5797	0.5806	0.5806	0.5805	0.5806
total heat load [MJ/m^2]	78.26	76.18	75.84	75.84	75.82	75.86
final velocity [m/s]	469.9	466.80	465.4	465.5	466.5	465.0
Final range to target [km]	10.54	4.983	4.877	4.917	4.477	5.097
Propellant mass [kg]	15824.5	15706.5	15705.0	15703.6	15709.0	15702.3

For Cabo Verde the velocity, heading angle and flight path angle show the same expected behaviour as for Natal, with, notably, the heading angle change being positive here. The guidance overcontrol value of each solution is increased, though for the solution with seed 42 this is by only a very small amount. For this location, the deadband values are all decreased with respect to the initial guess, narrowing the lateral corridor. The peak g-load and heat flux are generally higher for each solution, but not for all.

The total heat load is lower for each of these, while the velocity is higher for most, but not all solutions. Once again, each of the load constraints is not exceeded by a sizeable margin. Each seed except for 11 resulted in a solution that arrived within 5 [km] of the target. Seed 96 resulted in the solution among these with the lowest propellant mass and so is selected as the final best identified trajectory for this location.

6.1.3. The Canarias

Plots of various quantities of the trajectories for The Canarias can be found in Figure 6.11, Figure 6.12, Figure 6.13, Figure 6.14, Figure 6.15. An overview of how these trajectories compare can be found in Table 6.3.

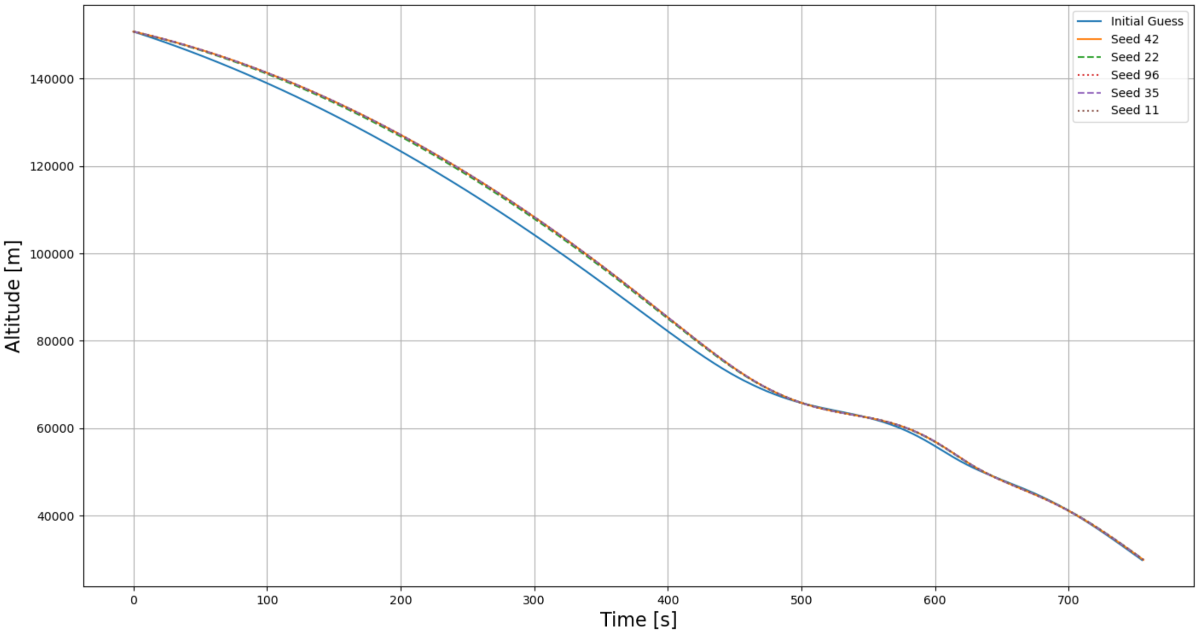

Figure 6.11: The altitudes over time of the best identified trajectories for each seed and the initial guess for The Canarias.

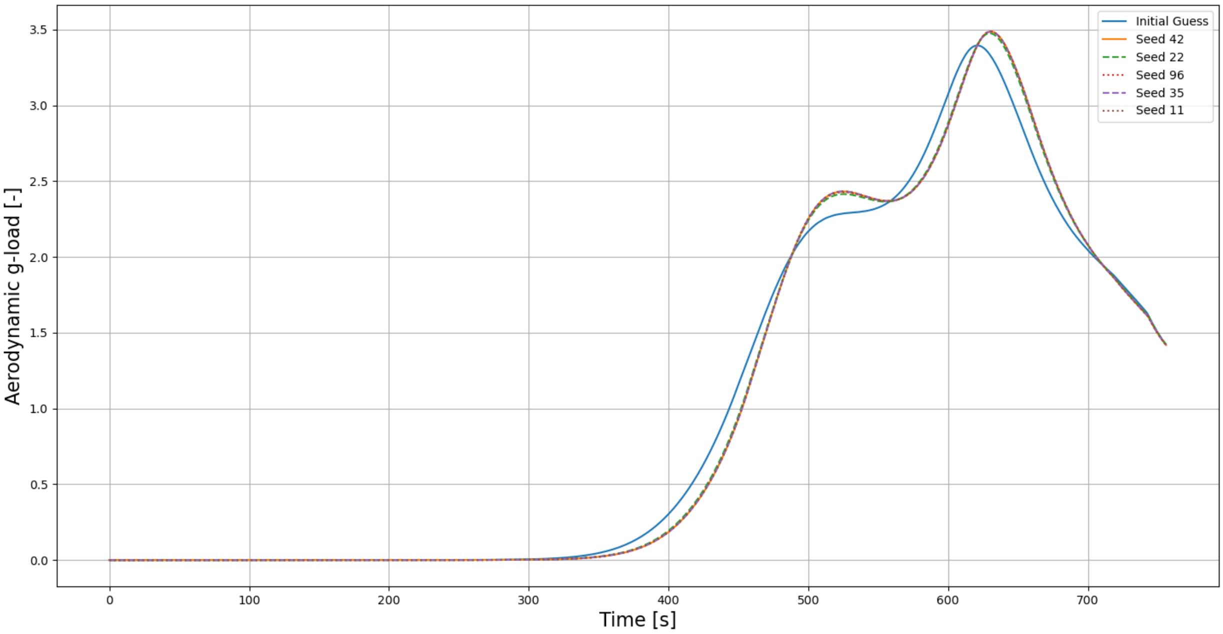

Figure 6.12: The g-load over time of the best identified trajectories for each seed and the initial guess for The Canarias.

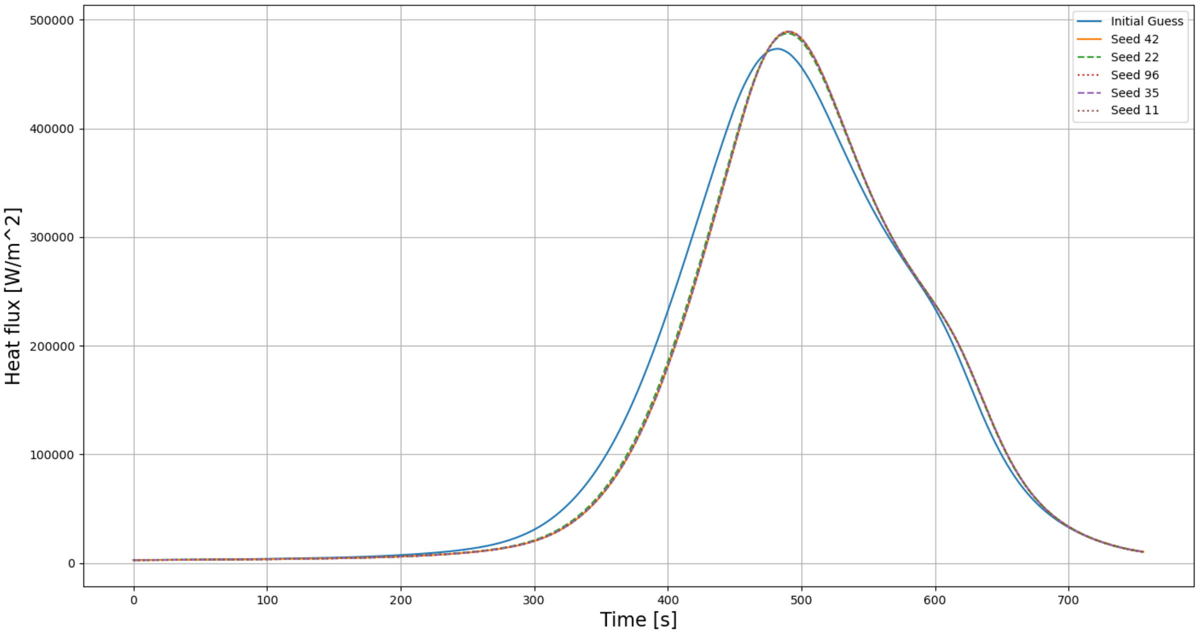

Figure 6.13: The heat flux over time of the best identified trajectories for each seed and the initial guess for The Canarias.

Figure 6.14: The latitude-longitude map of the best identified trajectories for each seed and the initial guess for The Canarias.

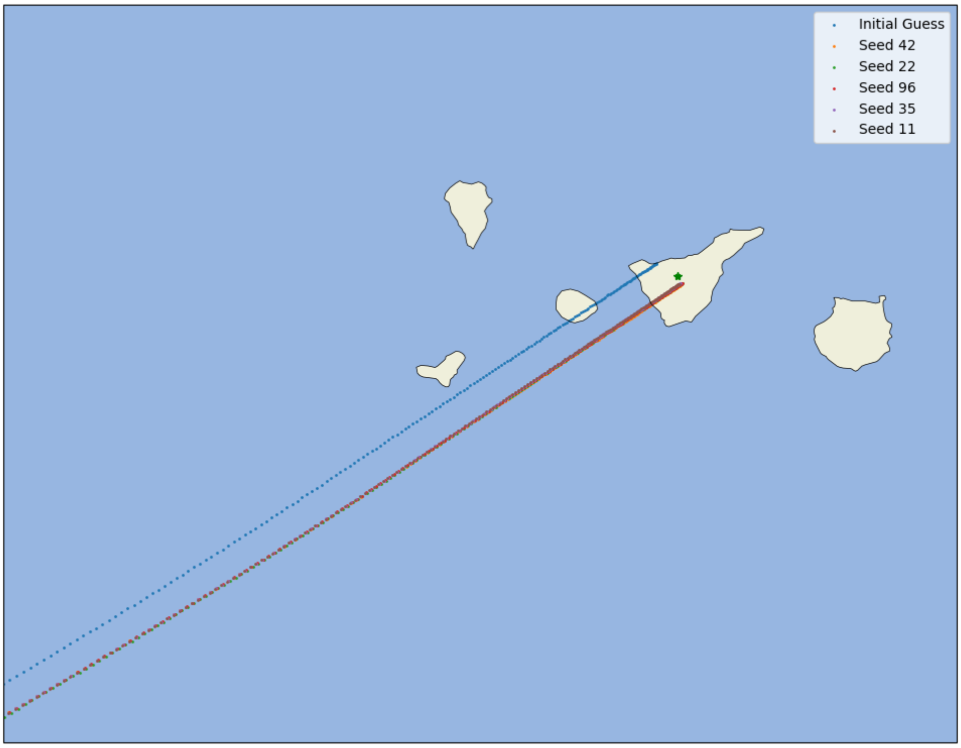

Figure 6.15: A close-up of the latitude-longitude map of the best identified trajectories for each seed and the initial guess for The Canarias.

Table 6.3: A comparison of values associated with the best identified trajectories for each seed and the initial guess for The Canarias. For the inputs, the value with respect to the initial guess is given.

Seed	Initial Guess	42	22	96	35	11
V [m/s]	7275	-42.79	-38.44	-41.75	-42.92	-42.88
ψ [deg]	50.5	$+0.2403$	$+0.2591$	$+0.2444$	$+0.2424$	$+0.2473$
γ	-0.8	$+0.222$	$+0.200$	$+0.217$	$+0.229$	$+0.222$
K [-]	1.0	$+8.312$	$+2.260$	$+5.152$	$+4.851$	$+3.655$
c_0 [deg]	2.0	$+0.451$	-1.073	-0.672	-0.329	-1.126
c_1 [deg/(m/s^2)]	$1.633 \cdot 10^{-7}$	$-6.911 \cdot 10^{-9}$	$+1.639 \cdot 10^{-7}$	$+5.556 \cdot 10^{-8}$	$+1.693 \cdot 10^{-7}$	$-7.071 \cdot 10^{-8}$
peak g-load [-]	3.396	3.487	3.476	3.487	3.487	3.487
peak heat flux [MW/m^2]	0.4732	0.4990	0.4872	0.4885	0.4890	0.4890
total heat load [MJ/m^2]	98.47	94.71	95.11	94.81	94.70	94.70
final velocity [m/s]	484.9	485.6	487.5	486.9	485.7	484.5
Final range to target [km]	14.82	5.694	5.115	5.106	4.592	4.718
Propellant mass [kg]	38730.4	38095.6	38091.3	38095.1	38091.6	38084.1

For this location, similar performance can again be seen for the velocity, heading angle, flight path angle and guidance overcontrol value, although this last one shows significantly larger values than before. The deadband values show inconsistent behaviour, being either increased or decreased somewhat. Once again, the solutions showed an increase in peak g-load and heat flux and a decrease in the total heat load, with all of them remaining solidly within the margins. The final velocity shows both increases and decreases for these solutions. Unfortunately only two of the solutions actually achieve the desired range. Of these, the one for seed 11 requires the least propellant mass, and so it is selected for this target location.

6.1.4. The Azores

Plots of various quantities of the trajectories for The Azores can be found in Figure 6.16, Figure 6.17, Figure 6.18, Figure 6.19, Figure 6.20. An overview of how these trajectories compare can be found in Table 6.4.

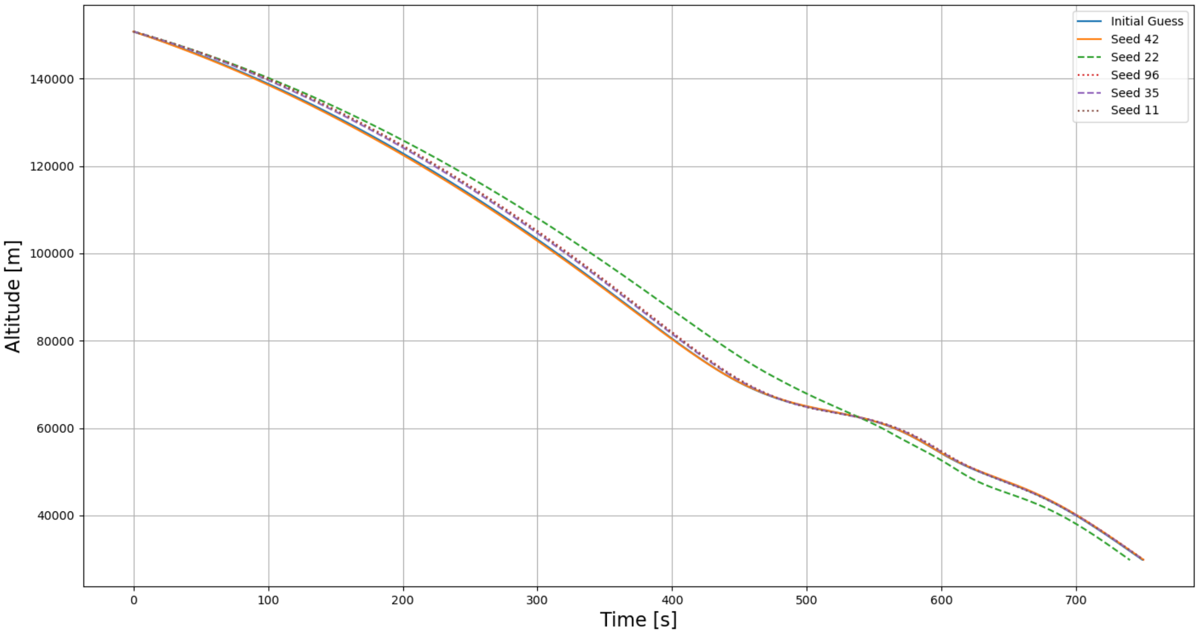

Figure 6.16: The altitudes over time of the best identified trajectories for each seed and the initial guess for The Azores.

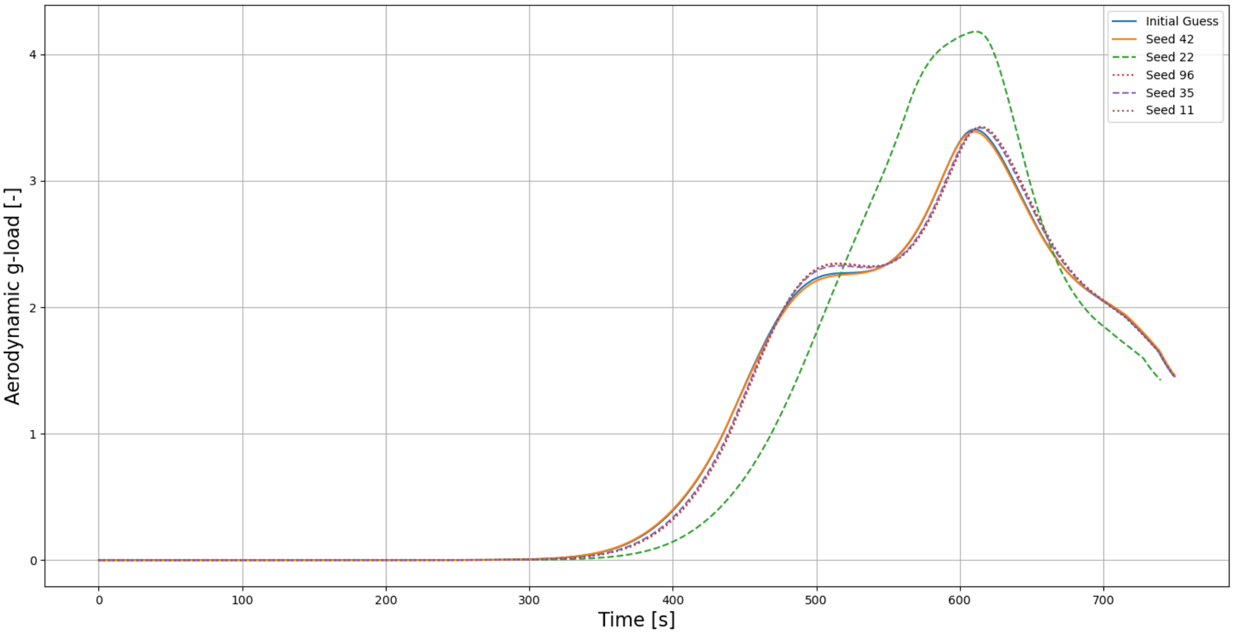

Figure 6.17: The g-load over time of the best identified trajectories for each seed and the initial guess for The Azores.

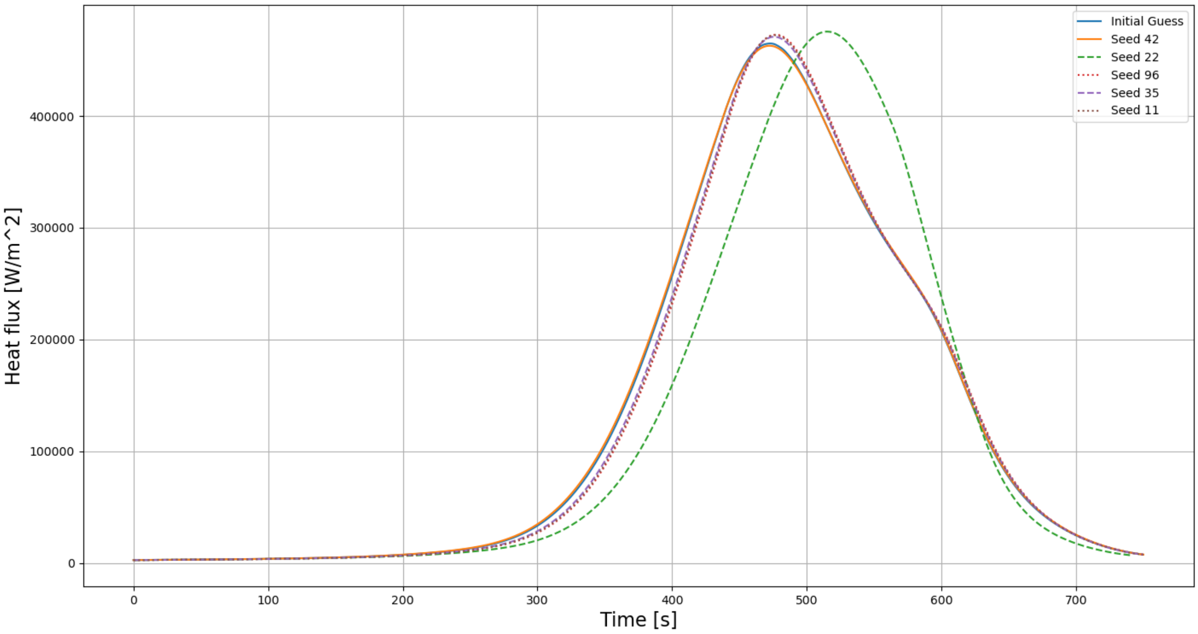

Figure 6.18: The heat flux over time of the best identified trajectories for each seed and the initial guess for The Azores.

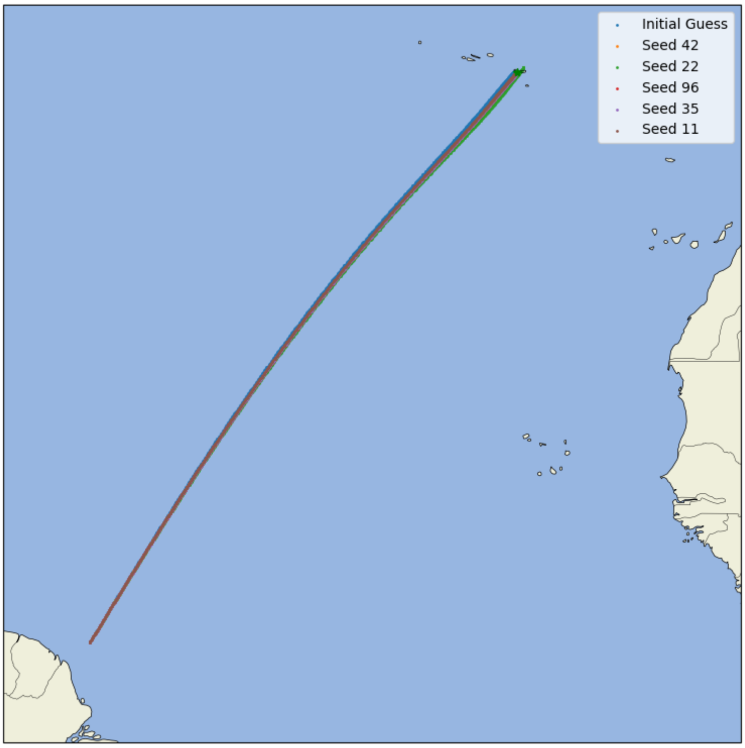

Figure 6.19: The latitude-longitude map of the best identified trajectories for each seed and the initial guess for The Azores.

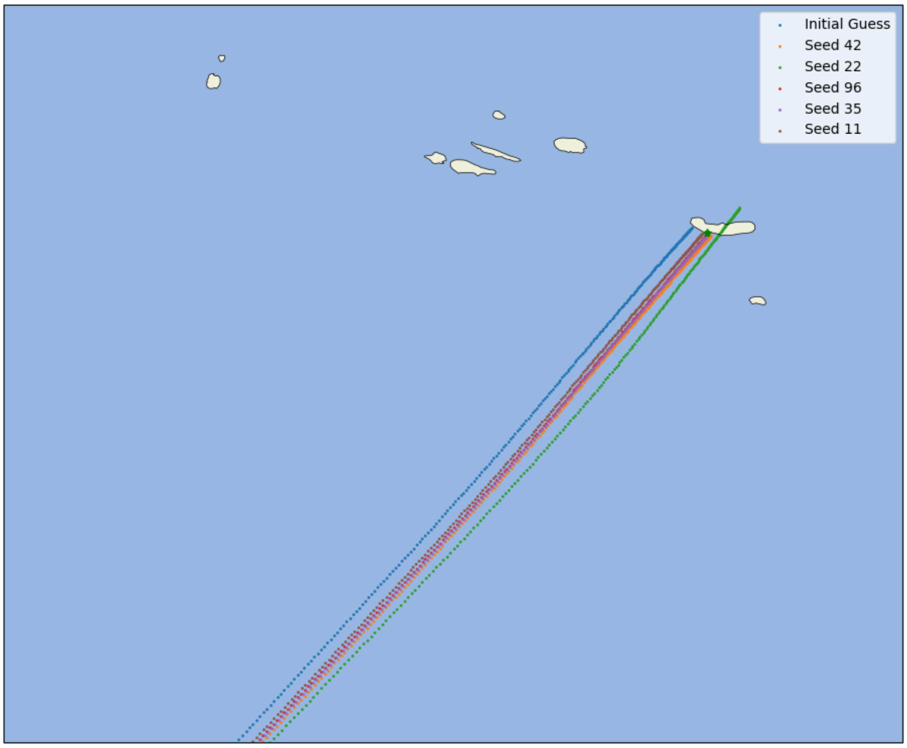

Figure 6.20: A close-up of the latitude-longitude map of the best identified trajectories for each seed and the initial guess for The Azores.

Table 6.4: A comparison of values associated with the best identified trajectories for each seed and the initial guess for The Azores. For the inputs, the value with respect to the initial guess is given.

Seed	Initial Guess	42	22	96	35	11
V [m/s]	7375	$+3.522$	$+9.178$	-21.25	-16.65	-19.89
ψ [deg]	31.0	$+0.3700$	$+0.4576$	$+0.2815$	$+0.3134$	$+0.2220$
γ	-0.8	-0.025	-0.095	$+0.103$	$+0.076$	$+0.095$
K [-]	1.0	$+4.315$	$+9.052$	$+8.453$	$+4.826$	$+8.289$
c_0 [deg]	2.0	$+1.444$	$+0.954$	$+1.025$	$+0.416$	$+2.194$
c_1 [deg/(m/s^2)]	$1.633 \cdot 10^{-7}$	$+1.436 \cdot 10^{-7}$	$-1.311 \cdot 10^{-8}$	$+1.746 \cdot 10^{-7}$	$-5.291 \cdot 10^{-8}$	$+8.907 \cdot 10^{-8}$
peak g-load [-]	3.406	3.388	4.180	3.424	3.419	3.426
peak heat flux [MW/m^2]	0.4649	0.4628	0.4755	0.4727	0.4708	0.4723
total heat load [MJ/m^2]	95.46	95.98	90.52	93.71	94.14	93.79
final velocity [m/s]	491.5	491.6	485.3	492.1	488.2	487.2
Final range to target [km]	16.17	5.151	44.90	1.620	4.184	4.577
Propellant mass [kg]	81412.9	80451.0	80298.4	80330.9	80311.0	80512.9

The solutions for The Azores show some unexpected behaviour, most notably the solution for seed 22. This one shows a very different behaviour and ends up far from the target but does so with a much lower propellant mass, compensating for the fitness penalty. As the remaining options were close or within the range, this is not seen as a significant problem, this option is just discarded. One of the solutions, seed 42, did not succeed in reaching within 5 [km] and also showed somewhat different behaviour, but the three that succeeded showed the expected behaviour for the velocity, heading angle and flight path angle. The overcontrol value tends to assume considerable values for this location, while the deadbands were increased or decreased slightly, generally widening the lateral corridor. Each of the successful simulations showed the same behaviour as with the other location with regards to the g-load, heat flux and heat load. Of all the optimisations that succeeded, the solution for seed 35 has the least propellant mass and so is selected for this location. It should also be noted that the required propellant mass for this location is very high, likely making it unsuitable.

6.1.5. Les Mureaux

Plots of various quantities of the trajectories for Les Mureaux can be found in Figure 6.21, Figure 6.22, Figure 6.23, Figure 6.24, Figure 6.25. An overview of how these trajectories compare can be found in Table 6.5.

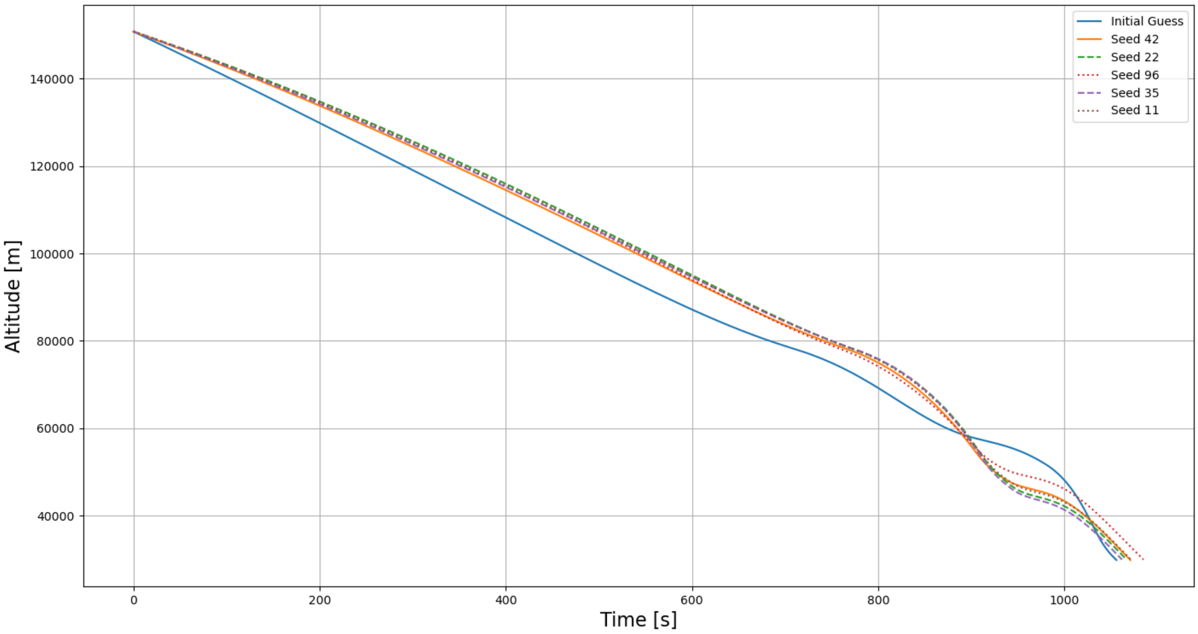

Figure 6.21: The altitudes over time of the best identified trajectories for each seed and the initial guess for Les Mureaux.

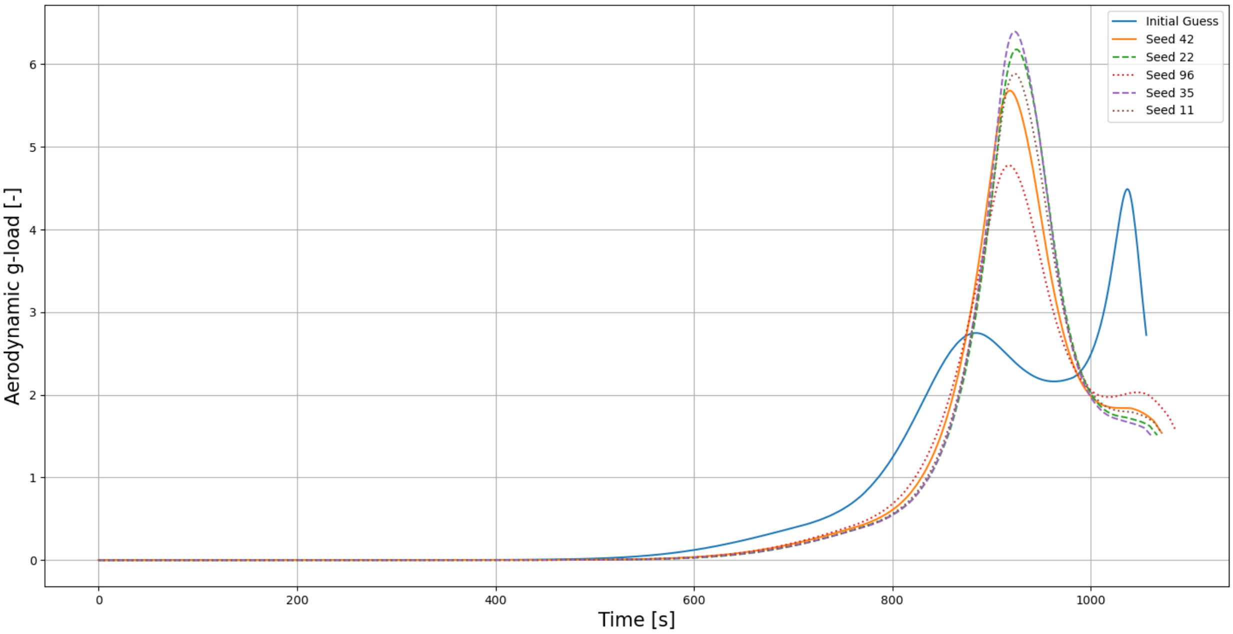

Figure 6.22: The g-load over time of the best identified trajectories for each seed and the initial guess for Les Mureaux.

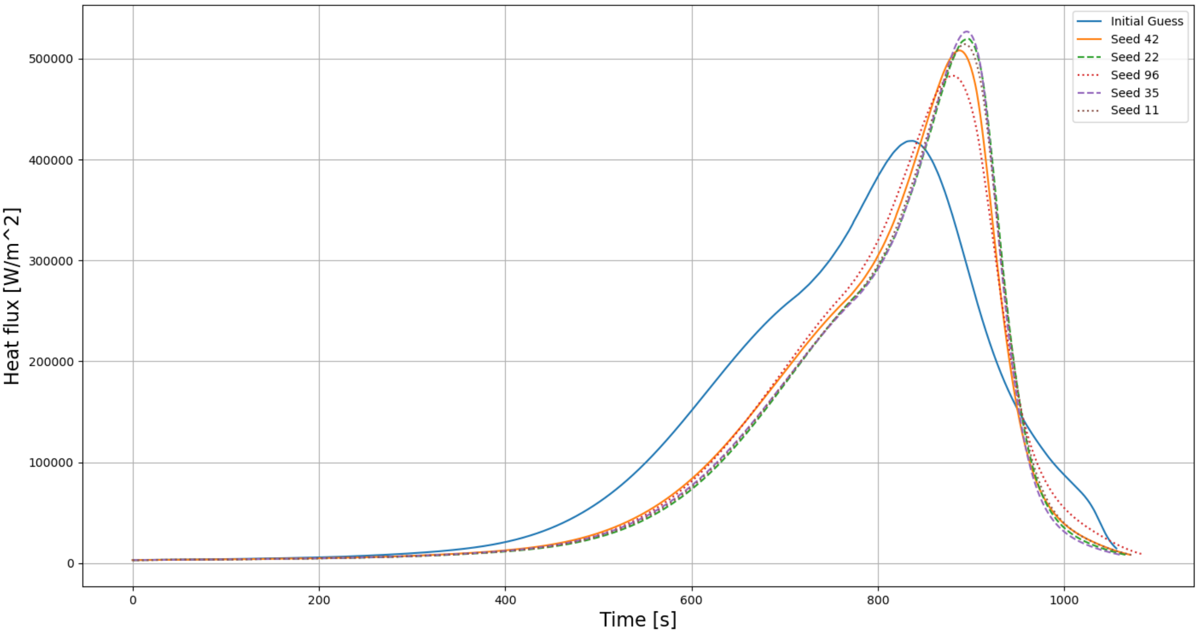

Figure 6.23: The heat flux over time of the best identified trajectories for each seed and the initial guess for Les Mureaux.

Figure 6.24: The latitude-longitude map of the best identified trajectories for each seed and the initial guess for Les Mureaux.

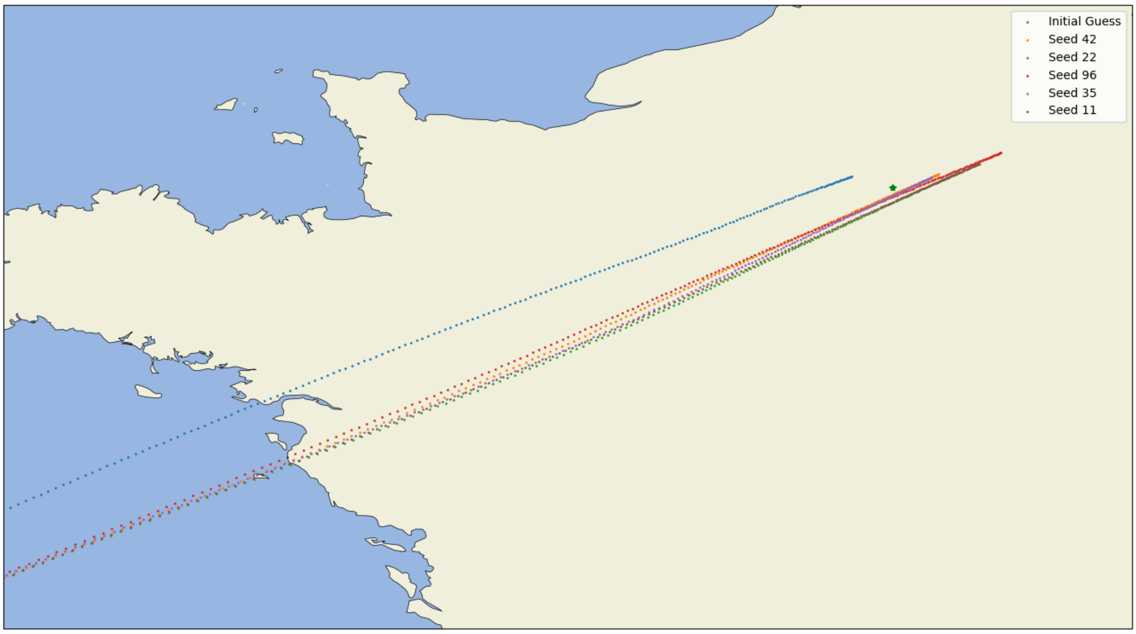

Figure 6.25: A close-up of the latitude-longitude map of the best identified trajectories for each seed and the initial guess for Les Mureaux.

Table 6.5: A comparison of values associated with the best identified trajectories for each seed and the initial guess for Les Mureaux. For the inputs, the value with respect to the initial guess is given.

Seed	Initial Guess	42	22	96	35	11
V [m/s]	7505	-16.83	-21.74	-22.86	-17.60	-19.76
ψ [deg]	35.0	$+0.3933$	$+0.3852$	$+0.3822$	$+0.3671$	$+0.3898$
γ	-0.8	$+0.176$	$+0.223$	$+0.215$	$+0.194$	$+0.208$
K [-]	1.0	$+5.656$	$+0.129$	$+9.160$	$+4.801$	$+8.255$
c_0 [deg]	2.0	$+0.399$	$+0.439$	-0.380	$+0.855$	-1.406
c_1 [deg/(m/s^2)]	$1.633 \cdot 10^{-7}$	$+1.613 \cdot 10^{-8}$	$+1.579 \cdot 10^{-8}$	$-1.653 \cdot 10^{-8}$	$+5.897 \cdot 10^{-9}$	$+1.777 \cdot 10^{-7}$
peak g-load [-]	4.490	5.68	6.180	4.777	6.395	5.885
peak heat flux [MW/m^2]	0.4185	0.5082	0.5198	0.4831	0.5269	0.5144
total heat load [MJ/m^2]	125.7	111.0	108.8	112.6	108.9	109.7
final velocity [m/s]	667.0	519.1	516.9	529.8	514.9	517.8
Final range to target [km]	23.88	27.17	38.60	65.49	22.62	51.00
Propellant mass [kg]	72619.6	71404.0	71362.2	71355.1	71460.5	71375.7

For this location, it is immediately evident that none of the best identified solutions met the 5 [km] range constraint. Based on this fact, combined with the substantial propellant mass requirement for the solutions that were achieved, Les Mureaux is rejected as a target location and will not be considered further.

6.1.6. Optimisation Outcomes

Of the five target locations, only Les Mureaux failed to produce a viable solution, and so it is rejected. Furthermore, considering the payload mass penalty of 0.25, the best solution for The Azores still incurs a 20077.75 [kg] penalty on the payload from the propellant mass alone. Considering that for a LEO launch the Ariane 64 configuration has a payload capacity of 21650 [kg], this would leave only 1522.25 [kg], which is too little to be useful. For this reason, The Azores is eliminated as well. This leaves the three remaining options of Natal, Cabo Verde and The Canarias.

For each of these, the best solution used a lower velocity and a more shallow entry, as well as a heading angle that diverges as little as possible from the vehicle state after it disconnects from the rest of the stage. There appears to be no consistency of changes to the guidance overcontrol and deadband values among the best solutions. Of the remaining locations, Cabo Verde shows the lowest propellant mass requirement, which was in line with the expectations. How this result impacts the costs will be seen in section 6.3.

6.2. Sensitivity Analysis

With the outcomes of the optimisation produced, it should be assessed how sensitive the results are to uncertainties. As the environment was looked at extensively in section 5.3, the main concern here is the uncertainties in position and velocity after the vehicle has disconnected from the core stage. To quantify these, a sensitivity analysis is performed. The initial states are perturbed using a normal distribution. As unfortunately there is no data on what the uncertainty of the position and velocity of the vehicle is after the disconnection, the standard deviation of this distribution is set to various uncertainty values ranging from 10 to 0.1. The position is perturbed using this standard deviation, the velocity is perturbed using a distribution using $1/10th$ of this value. After the perturbation, the vehicle performs the re-entry as normal but using the reference trajectory file generated for the unperturbed trajectory. This is done 500 times for five different seeds, the same five as before. The question of interest here is whether the best identified trajectories still reach the 5 [km] range requirement and, if they do not, by how much they miss it. These results can be found in Figure 6.26, Figure 6.27, Figure 6.28, Figure 6.29, Figure 6.30, Figure 6.31, Figure 6.32, Figure 6.33, Figure 6.34, Figure 6.35, Figure 6.36 and Figure 6.37.

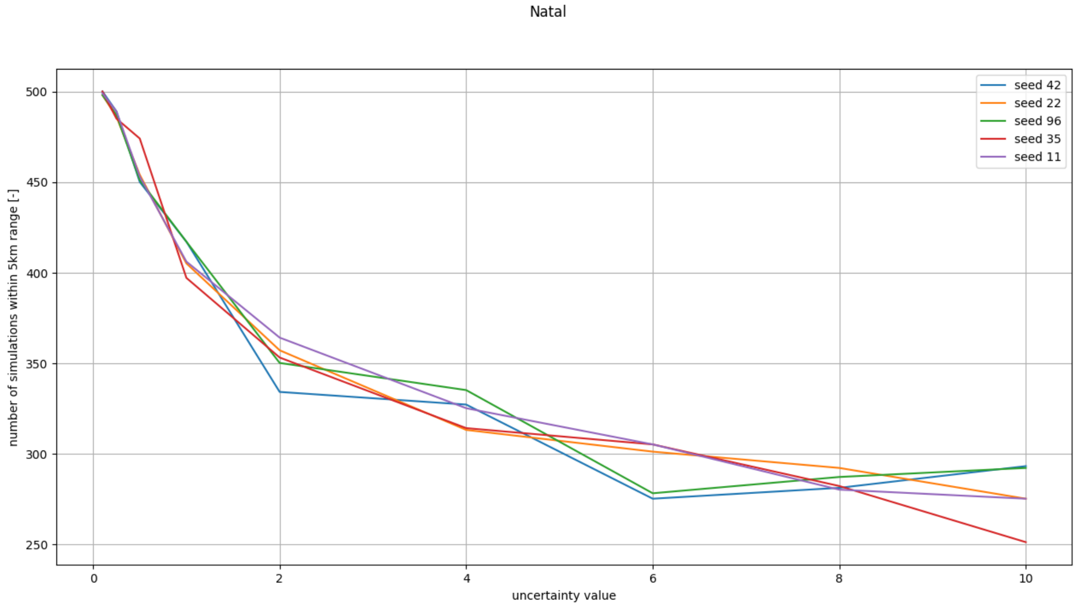

Figure 6.26: The number of simulations that ended within the 5 [km] range when perturbed using various uncertainties for each seed for Natal.

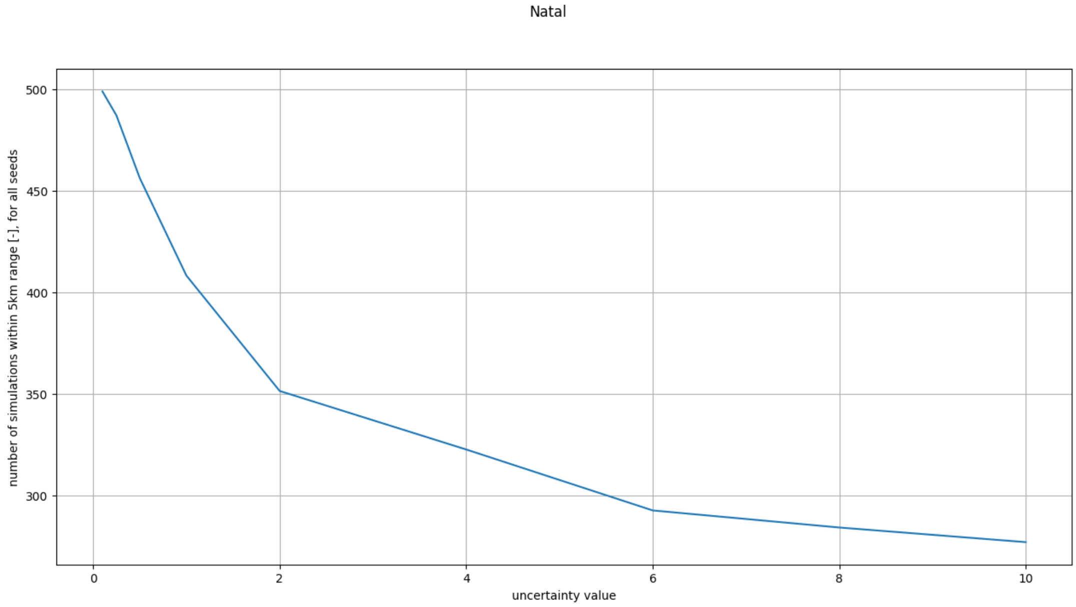

Figure 6.27: The average number of simulations that ended within the 5 [km] range when perturbed using various uncertainties for all seeds for Natal.

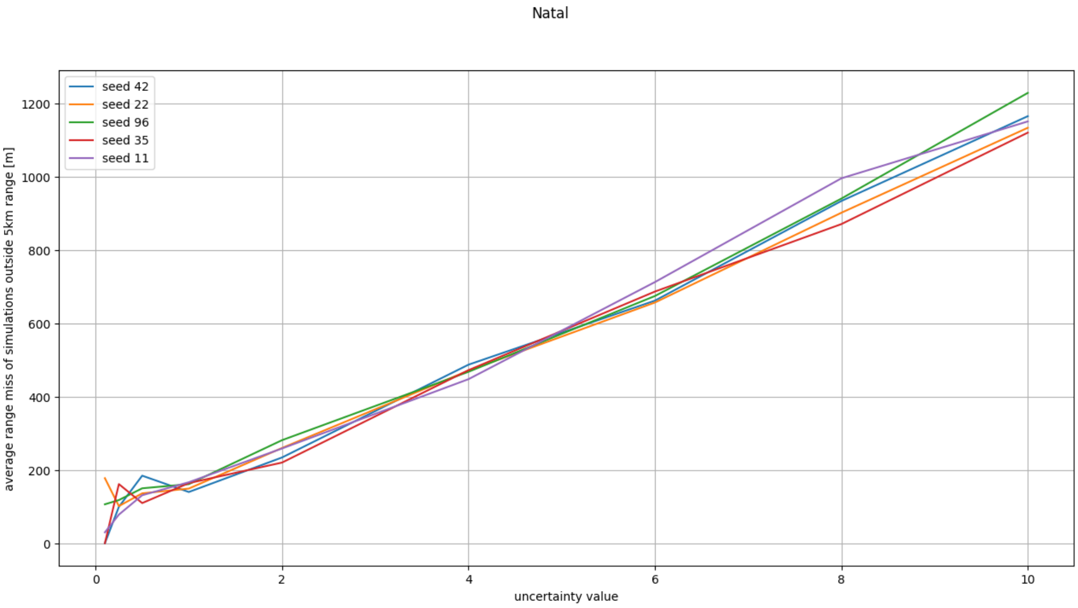

Figure 6.28: The average distance by which the simulations that end outside the range constraint miss that range constraint for each seed for Natal.

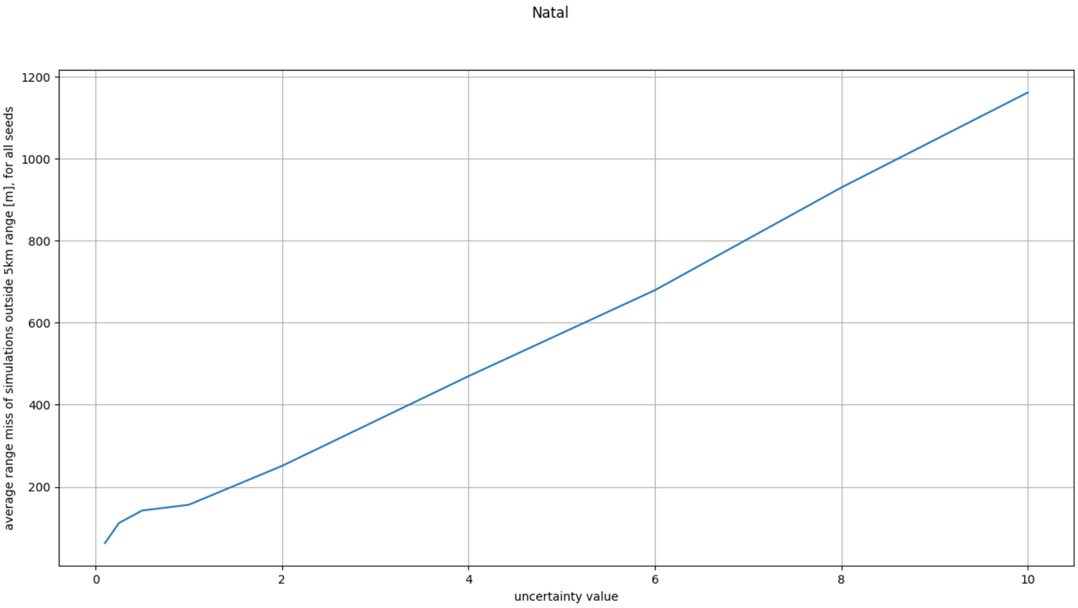

Figure 6.29: The average distance by which the simulations that end outside the range constraint miss that range constraint for all seeds for Natal.

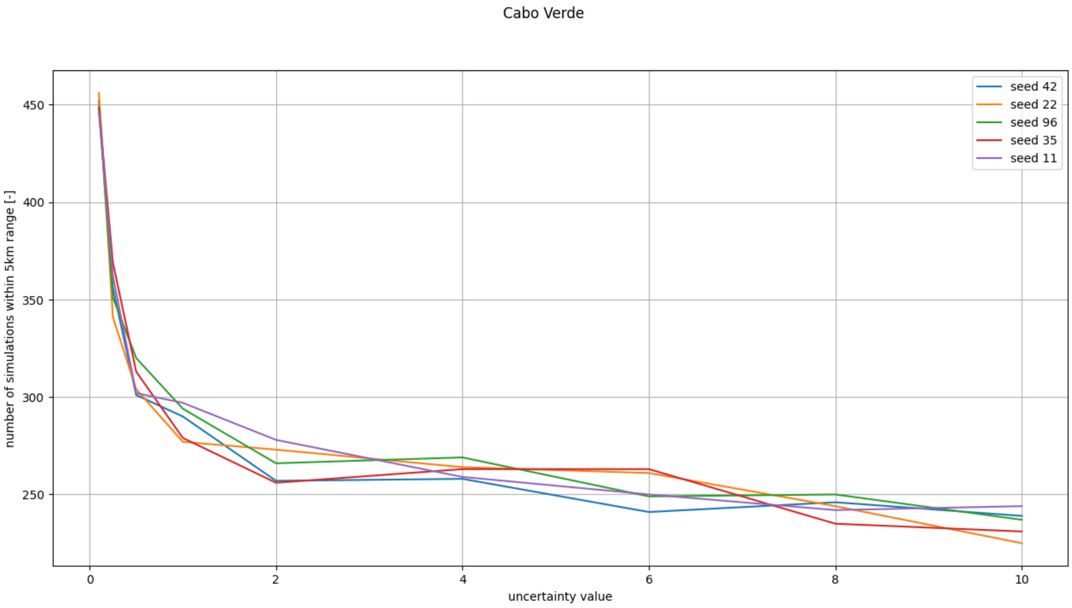

Figure 6.30: The number of simulations that ended within the 5 [km] range when perturbed using various uncertainties for each seed for Cabo Verde.

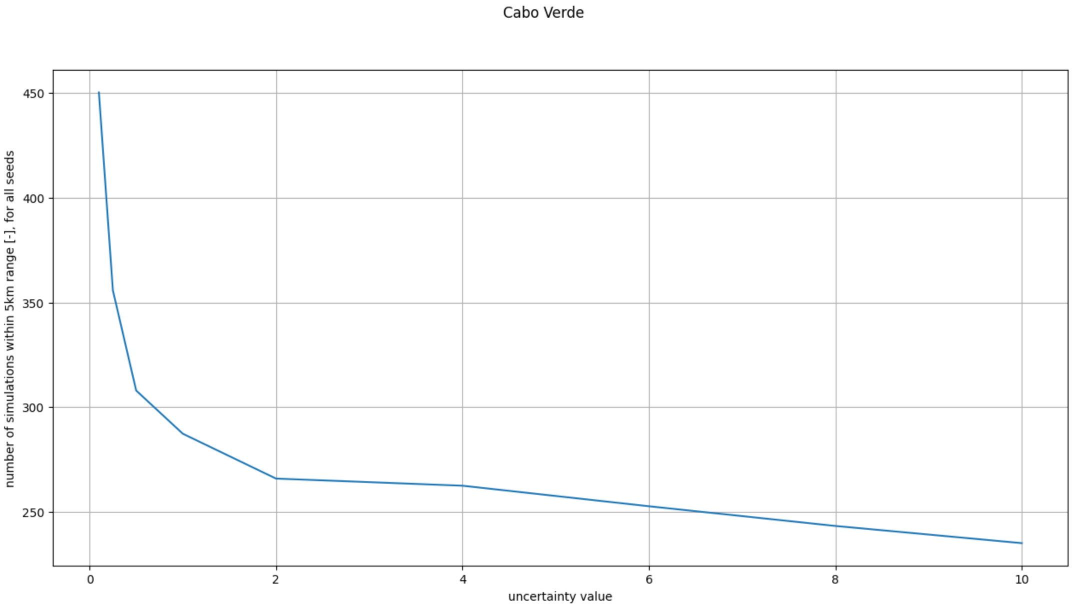

Figure 6.31: The average number of simulations that ended within the 5 [km] range when perturbed using various uncertainties for all seeds for Cabo Verde.

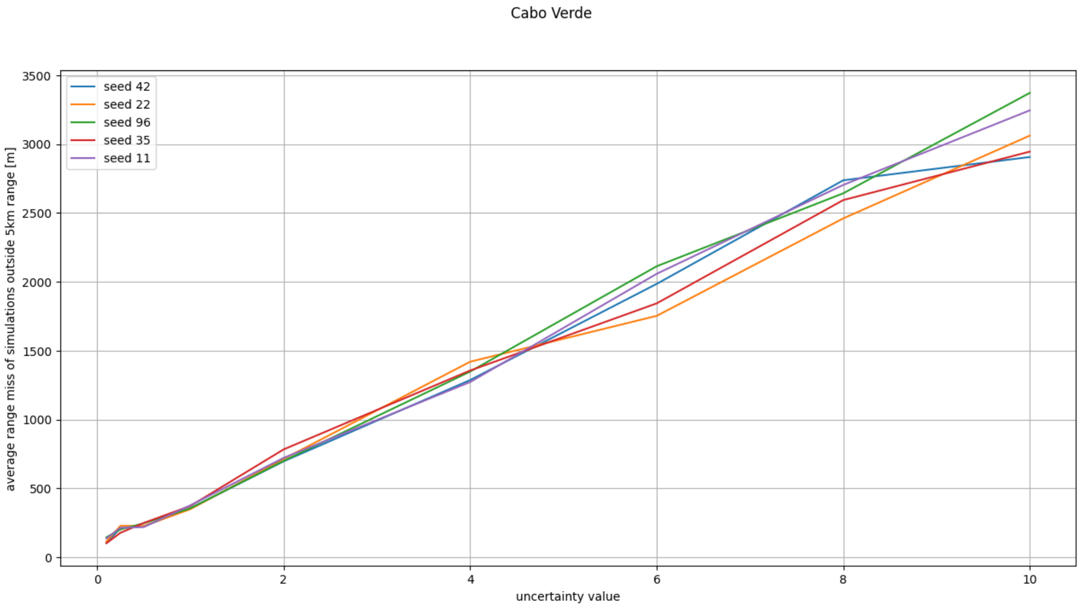

Figure 6.32: The average distance by which the simulations that end outside the range constraint miss that range constraint for each seed for Cabo Verde.

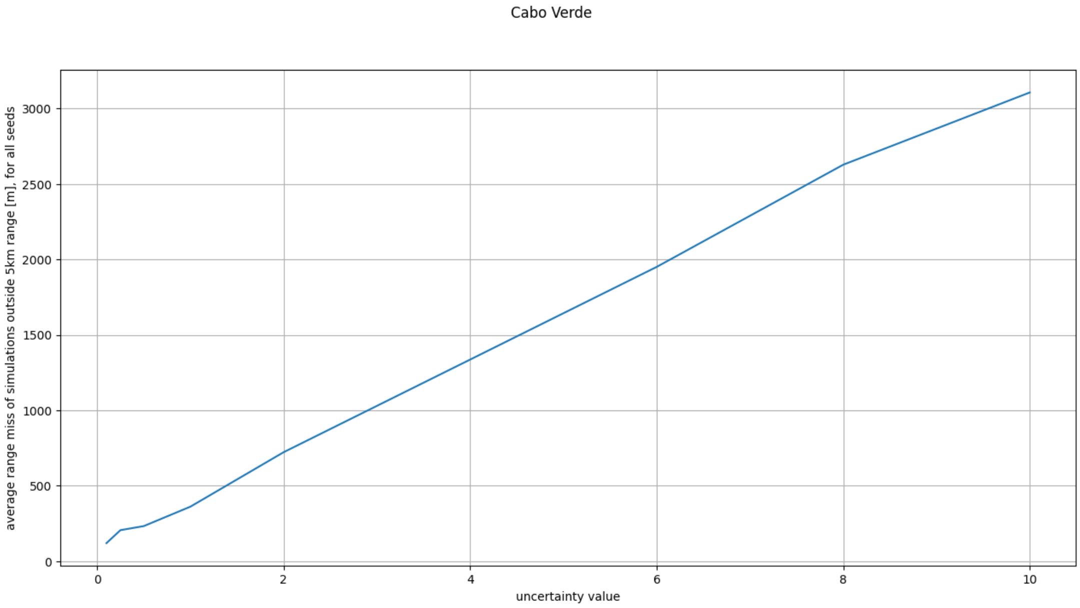

Figure 6.33: The average distance by which the simulations that end outside the range constraint miss that range constraint for all seeds for Cabo Verde.

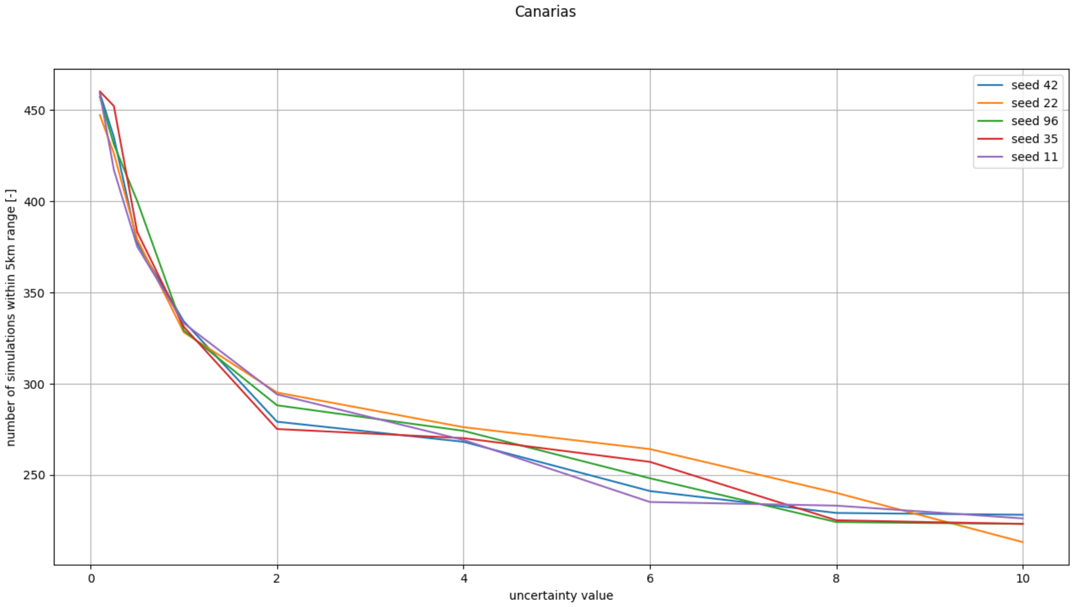

Figure 6.34: The number of simulations that ended within the 5 [km] range when perturbed using various uncertainties for each seed for The Canarias.

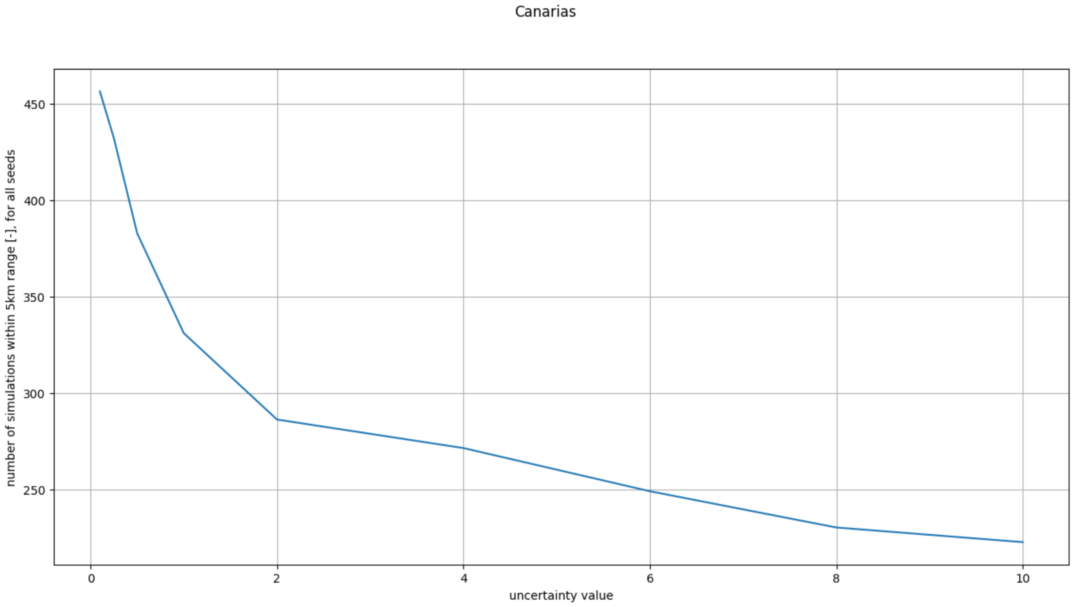

Figure 6.35: The average number of simulations that ended within the 5 [km] range when perturbed using various uncertainties for all seeds for The Canarias.

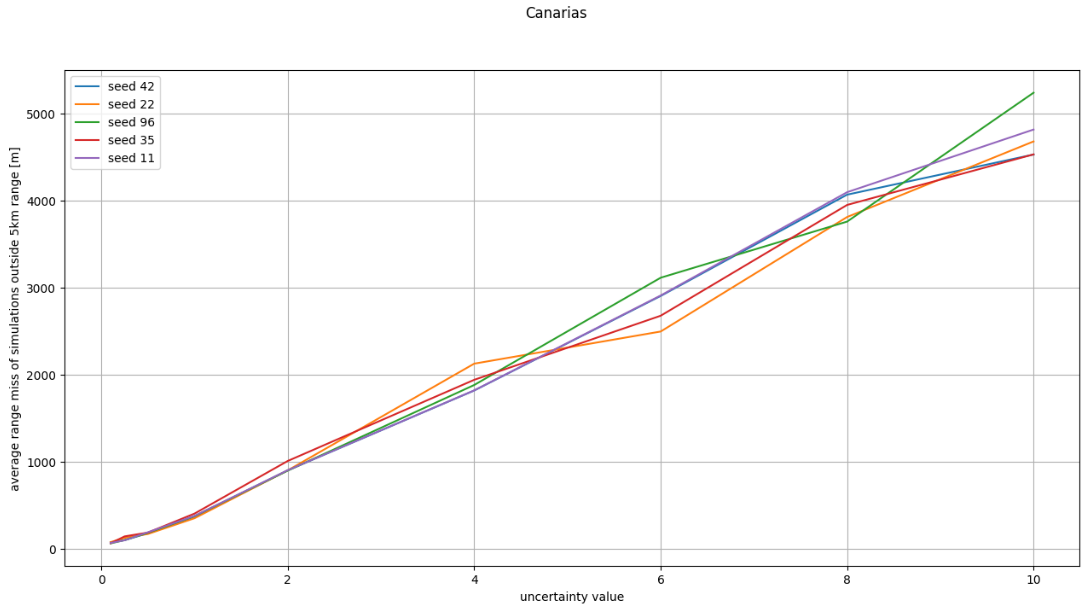

Figure 6.36: The average distance by which the simulations that end outside the range constraint miss that range constraint for each seed for The Canarias.

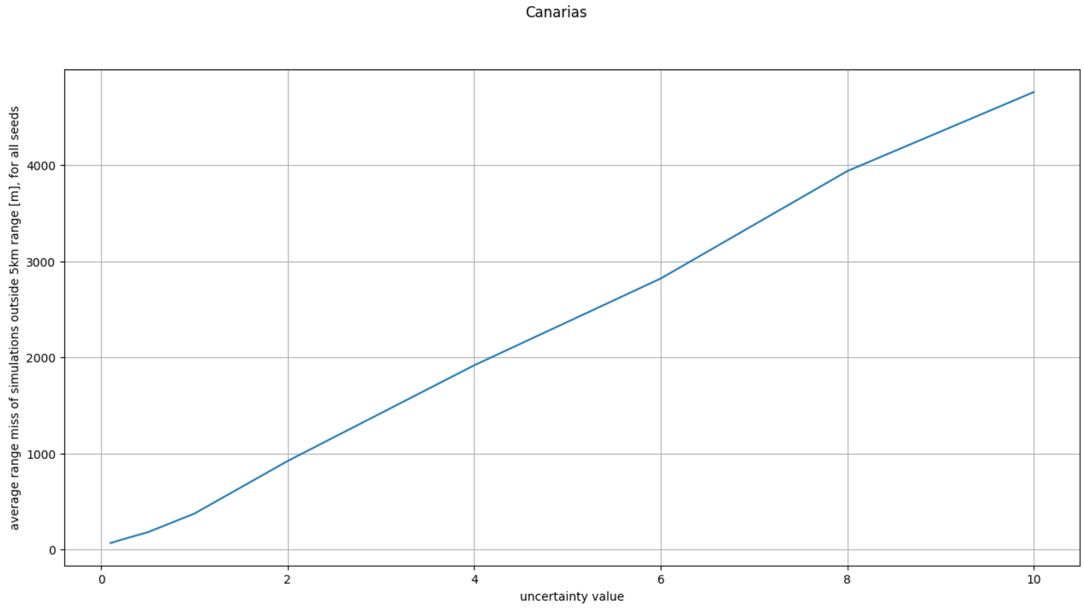

Figure 6.37: The average distance by which the simulations that end outside the range constraint miss that range constraint for all seeds for The Canarias.

For each of the considered locations, all five seeds have quite similar behaviour to one another, lending more credibility to these results. The location with the least number of simulations that miss the range is Natal, then Cabo Verde, with The Canarias missing the most, likely because this is also the order from shortest to longest trajectories. As the trajectory lasts longer, it seems the errors can propagate larger and larger. For all three locations, the number of simulations that end within range drops quickly at first, then tapers out to a slower descent. Even at the highest uncertainty value of 10, which corresponds to a standard deviation of 10 [m] and 1 [m/s], Natal has more than half of the simulations ending within range, while Cabo Verde and The Canarias have slightly less than half.

For the simulations that miss the target, once again the average miss is lowest for Natal, then Cabo Verde, then The Canarias. For most of the uncertainties, the average range miss seems to increase linearly with increasing uncertainty values, except for at the lower values, where they tend to be less consistent. The average range miss is within a few kilometres of the target location, which may still be feasible for the vehicle to overcome on its final approach, as the 5 [km] range was an assumption. Although it seems that the results are quite sensitive to perturbations in the initial state, the resulting differences in the outcomes may well be manageable.

6.3. Cost Estimation

With the final solutions for each target location determined, the costs of the reusable launch vehicle can be estimated. The reusable launch vehicle costs are comprised of a number of factors, which can be found in section 2.4. As the Prometheus engine can be used up to five times, this will be used as the number of flights before the vehicle must be produced again. This means that for every five launches, the vehicle is produced once and recovered, refurbished and filled with propellant four times. In the previous research, the approach was taken to pay off the development costs fully during the first five years of launches, and this will be done for this vehicle as well. [11] It should be noted that this model assumes the remainder of the Ariane 6, and therefore also the expendable case, have reached their optimal cost, and therefore will not be improved over time. It also neglects the chance of failure occurring, as there is no estimate for the odds at this point.

The equations in Equation 2.6, Equation 2.8, Equation 2.9, Equation 2.10 and Equation 2.11 take a number of values as inputs, which are then used to calculate the cost per launch and life cycle costs as per Equation 2.15 and Equation 2.14. An overview of these inputs can be found in Table 6.6. The travel times were determined by estimating the distance to Le Havre, France, from where it is assumed the vehicle would be transported the short distance to Les Mureaux upriver by barge or rail, and then dividing this by the expected travel speed of 13.5 [kn]. [3] The figure of 30 launches per year is the number given by the chief of ArianeGroup as the number of launches per year from which point reusability could be considered for the Ariane 6, as below that launch cadence the engineering crews would have a great deal of downtime. [2] The propellant mass figures include a 10% margin, as well as the propellant mass needed for the landing burn. An inflation calculator was used to convert all costs into costs for 2026.

Table 6.6: An overview of the inputs used to calculate the costs.

Input Parameter	Value
$M_{reusable}$ [kg]	10330.45
$M_{expendable}$ [kg]	6875
C_{stage} [M€, 2017]	3.7
$t_{travel,Natal}$ [hr]	574.1
$t_{travel,CaboVerde}$ [hr]	370.36
$t_{travel,Canarias}$ [hr]	232.8
H [hr]	2080
l [-]	30
C_{MY} [k€, 2025]	286.425
$M_{p,Natal}$ [kg]	20763
$M_{p,CaboVerde}$ [kg]	17622
$M_{p,Canarias}$ [kg]	42241
O/F [-]	6
C_f [€/kg, 2015]	1.35
C_{ox} [€/kg, 2015]	0.14
T_1 [M€, 2026]	14.48

Using these values, the cost per launch is calculated. This can be found in Figure 6.38, Figure 6.39 and Figure 6.40, where it is compared to the cost of the expendable Ariane 6 launchers but with the cost reduction of using the Prometheus engine instead of Vulcain.

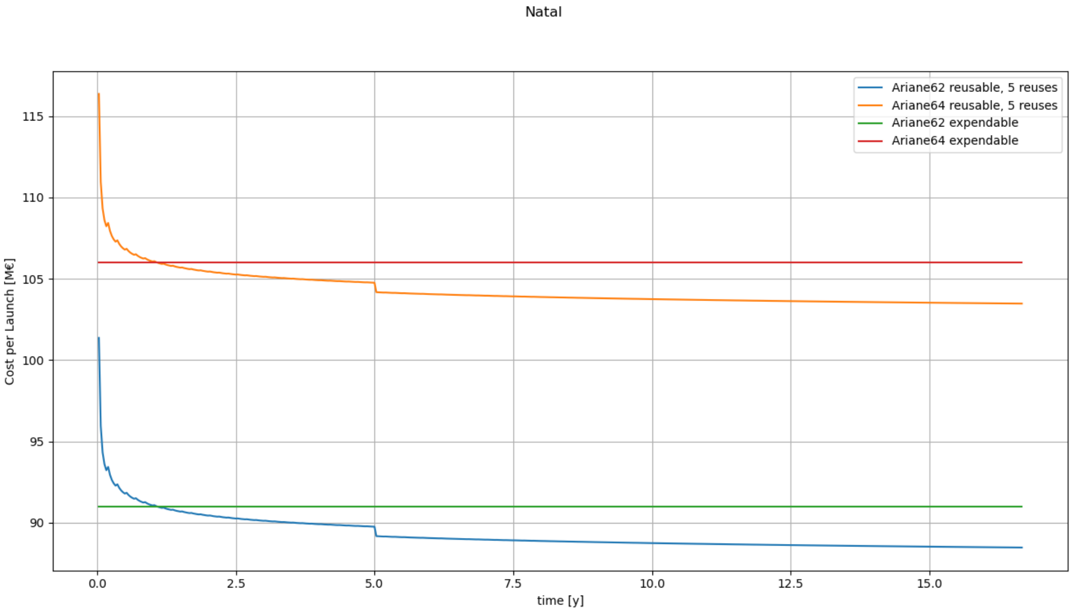

Figure 6.38: The cost per launch over the years for the reusable and expendable Ariane 6 variants for Natal.

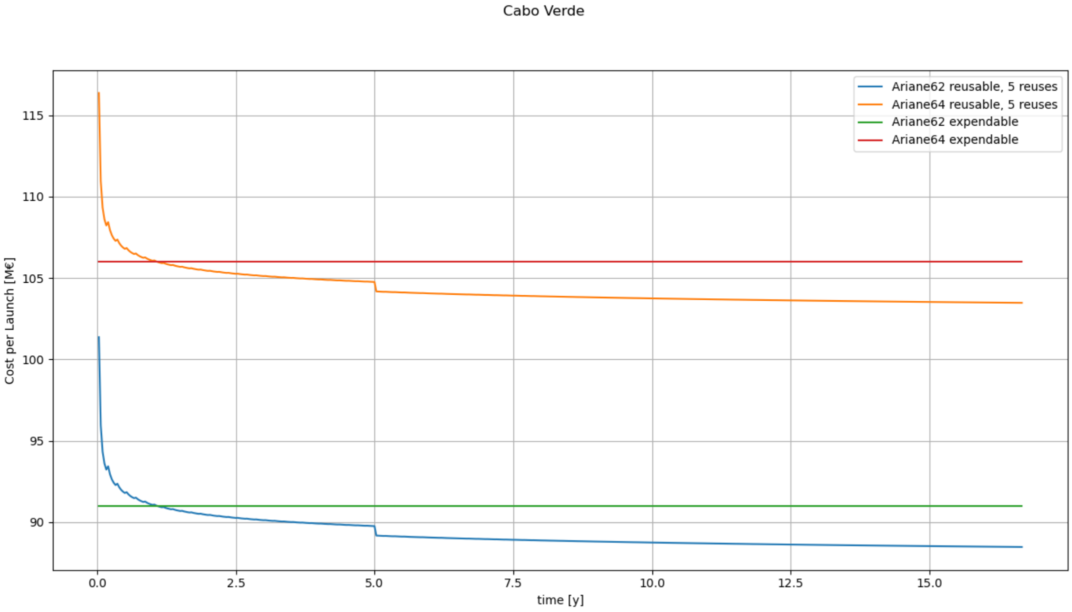

Figure 6.39: The cost per launch over the years for the reusable and expendable Ariane 6 variants for Cabo Verde.

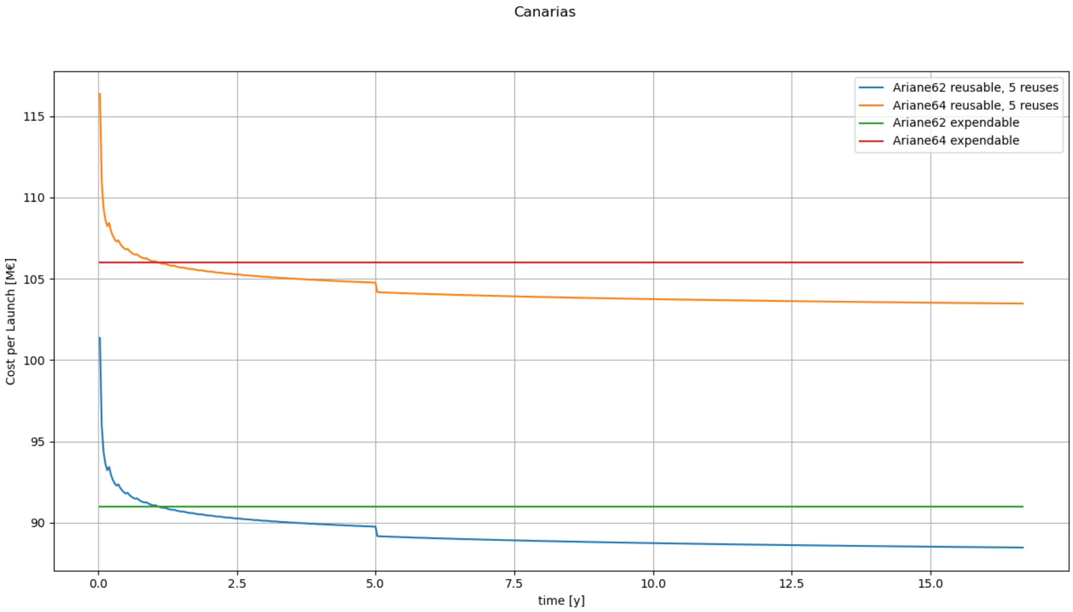

Figure 6.40: The cost per launch over the years for the reusable and expendable Ariane 6 variants for The Canarias.

In these figures, it can immediately be seen that each of the target locations behaves very similarly, so much so that they cannot be easily distinguished from one another. This is because at this point, the only differences between these are the propellant and transportation costs, which account for only a small part of the launches, which are dominated by production and refurbishment costs. As the costs are calculated by adding the reusable engine bay to the Ariane vehicle without the engine bay, the graphs for the A62 and A64 variants are simply shifted up or down copies of one another. In each of these, it can be seen that the costs per launch start significantly higher but drop quickly and cross below the cost per launch of the expendable vehicle. For all three locations, this happens at launch 33 for both A62 and A64, occurring early within the second year for a launch cadence of 30 launches per year. Each of these missions reaches a cost per launch of approximately 88.47 [M€] for A62 and 103.47 [M€] for A64. The life-cycle costs, or cumulative costs, can be found in Figure 6.41. As each of the graphs is once again very similar to one another, only the figure for Cabo Verde is displayed.

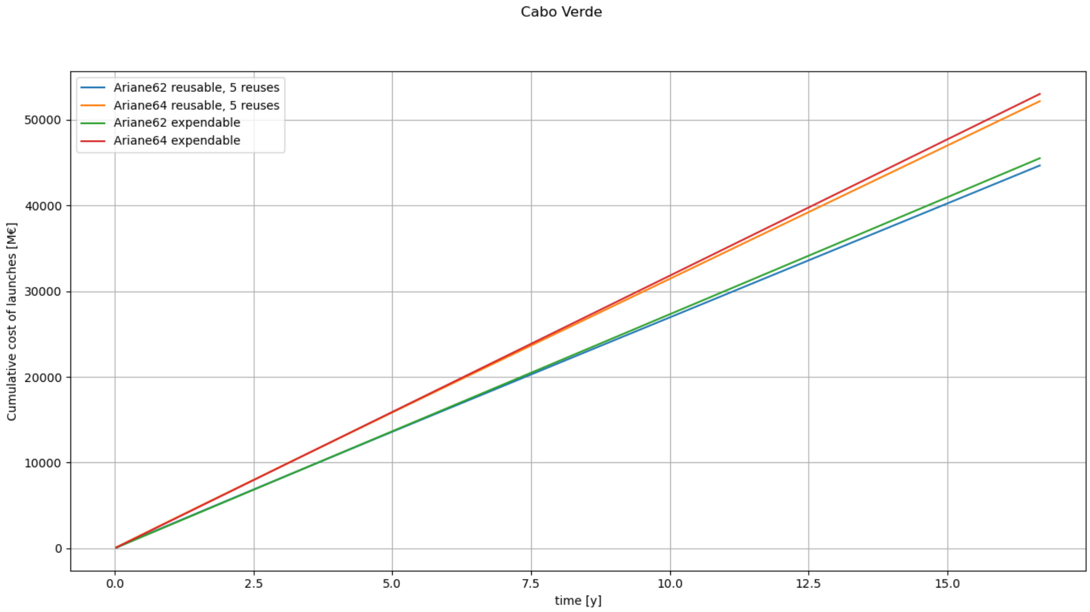

Figure 6.41: The life cycle cost over the years for the reusable and expendable Ariane 6 variants for Cabo Verde.

In this figure, it can be seen that there is a significant difference between the expendable and reusable missions, amounting to approximately 845 [M€] over the course of the 15 years modelled. For Natal and Cabo Verde, launch number 103 is the moment when the life cycle cost of the reusable variant becomes less than the life cycle cost of the expendable version for both the A62 and A64 versions. For The Canarias, this happens one launch later, at number 104. This moment occurs in the first half of the third year of operations using this launch cadence. Where the differences between the target locations truly become apparent is in the cost per payload mass figures, obtained by dividing the cost per launch from Figure 6.38, Figure 6.39 and Figure 6.40 by the maximum payload capacity available for these missions. The expendable vehicles have a capacity of 10350 [kg] and 21650 [kg] to LEO for A62 and A64, respectively. The payload mass penalty for the reusable system for each target location, along with the resulting maximum payload masses, can be found in Table 6.7. The cost per payload mass can be found in Figure 6.42, Figure 6.43 and Figure 6.44.

Table 6.7: The payload mass penalty of the reusable variant, as well as the payload capability to LEO, for each target location.

Location	Payload mass penalty [kg]	A62 payload capability [kg]	A64 payload capability [kg]
Natal	6510.9	3839.1	15139.1
Cabo Verde	5654.3	4695.6	15995.6
The Canarias	12368.5	not possible	9281.5

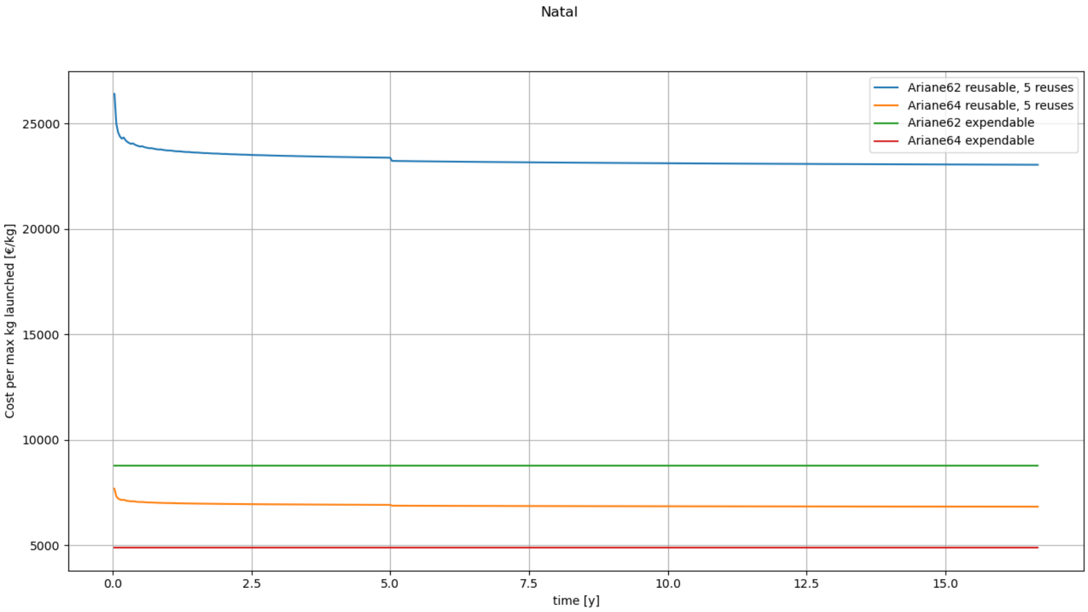

Figure 6.42: The cost per kg of payload over the years for the reusable and expendable Ariane 6 variants for Natal.

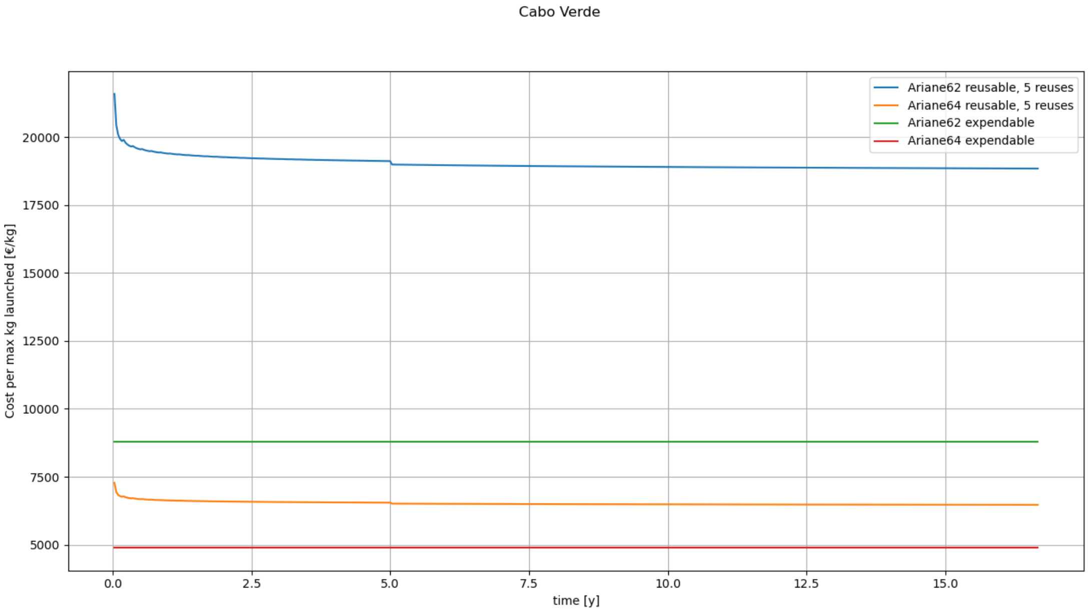

Figure 6.43: The cost per kg of payload over the years for the reusable and expendable Ariane 6 variants for Cabo Verde.

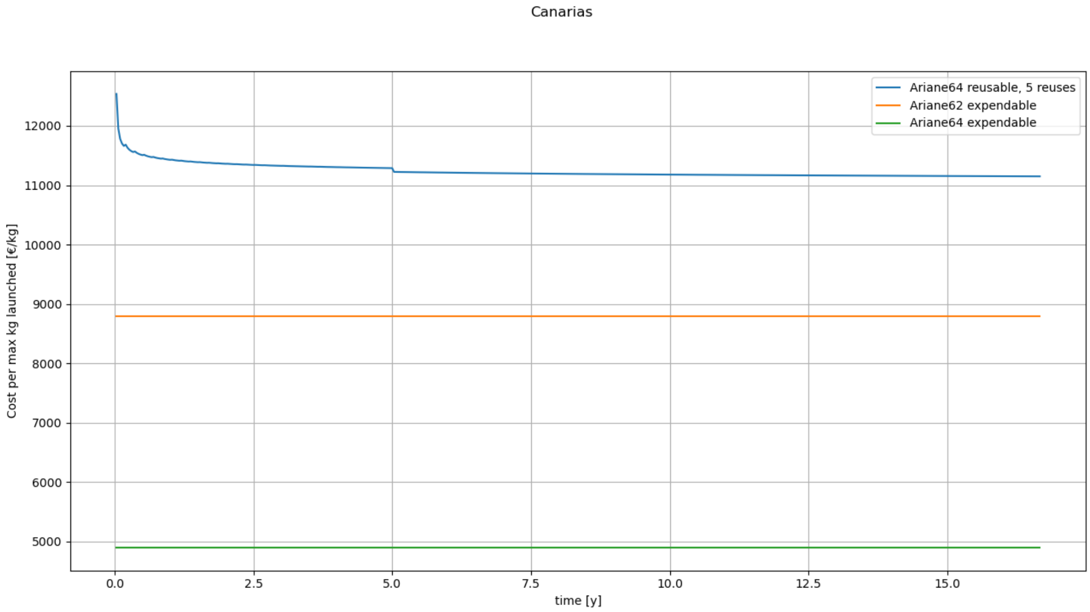

Figure 6.44: The cost per kg of payload over the years for the reusable and expendable Ariane 6 variants for The Canarias.

In these figures, it can be seen that per kg, the A62 reusable case is significantly more expensive, but that the A64 vehicle is less expensive than the expendable A62 launch for Cabo Verde and Natal. It can also be seen that for The Canarias, the A62 reusable case is not feasible, as the payload mass penalty incurred is greater than the payload capability, and that the A64 launch is still more expensive than either of the expendable launches per kg. It does seem however, that launching a reusable A64 mission landing in Cabo Verde or Natal may be a more attractive choice than using an expendable A62 vehicle. Between these two, Cabo Verde is indeed the better choice, costing less per payload kg launched. It should also be noted that over the launch history of the Ariane 6, most launches did not use the full payload mass capacity, and so depending on the missions that would wish to be launched, the payload mass penalty is not necessarily counterproductive.

The final consideration when it comes to costs is that the vehicle is reused every five launches, meaning that the fifth launch could go without the needed propellant or additional hardware. Even just leaving out the propellant mass raises the payload capacity to 9101.2 [kg] for a reusable A62 and 20401.3 [kg] for a reusable A64, landing in Cabo Verde. This notably also raises the payload capacity to a sufficient value to be able to launch the heaviest payload the Ariane 64 has launched to date. [45] The cost per payload mass, accounting for this, for Cabo Verde, can be found in Figure 6.45, where it can be seen that the cost per kg of every fifth launch gets close to the expendable case, especially for the A64 variant. With 30 launches per year, the fifth launch occurs roughly every second month, so if heavier launches can be scheduled smartly, the payload mass penalty issue can be alleviated further.

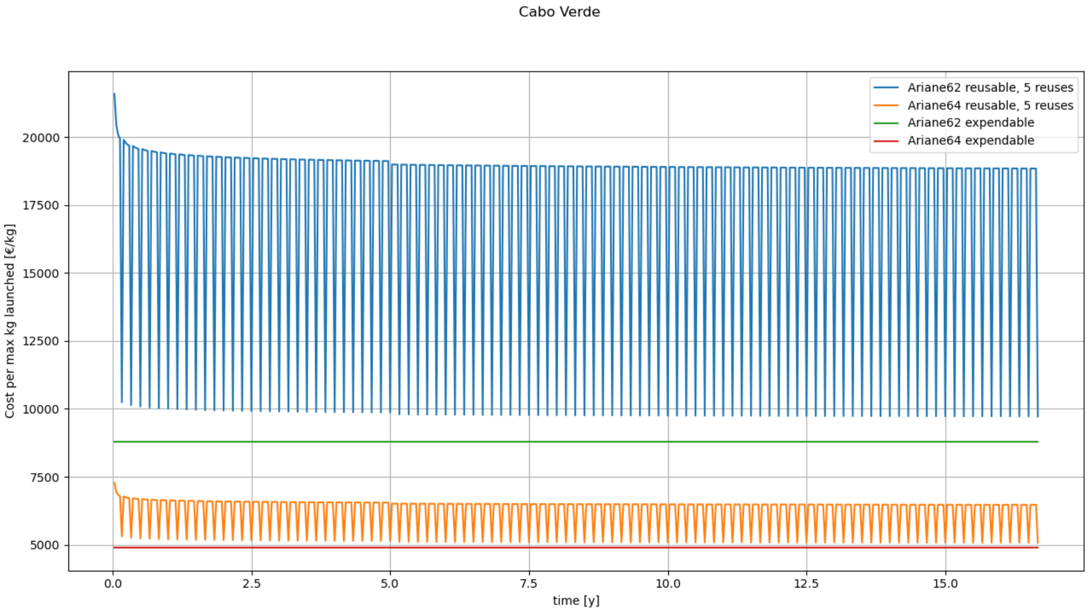

Figure 6.45: The cost per kg of payload over the years for the reusable and expendable Ariane 6 variants, accounting for the raised payload capacity of every fifth flight for Cabo Verde.

To see what its influence is on the costs, the number of refurbishments of the vehicle has been varied. The results can be found in Figure 6.46. Here it can be found that using this model, five uses is actually the best number for this launch vehicle, with more or fewer reuses raising the cost per launch. This is a combination of various factors: Every time the vehicle is refurbished, the refurbishment is more expensive, every time the vehicle is produced, the production costs decrease due to the learning factor, and more reuses mean the more expensive production has to occur less often. This comes together such that for five reuses, the more expensive refurbishments are prevented, and the production costs of the stage are reduced from the learning curve at a relatively good pace, but not too many refurbishments are performed. In this way, five reuses seems to be the best-case scenario for the vehicle.

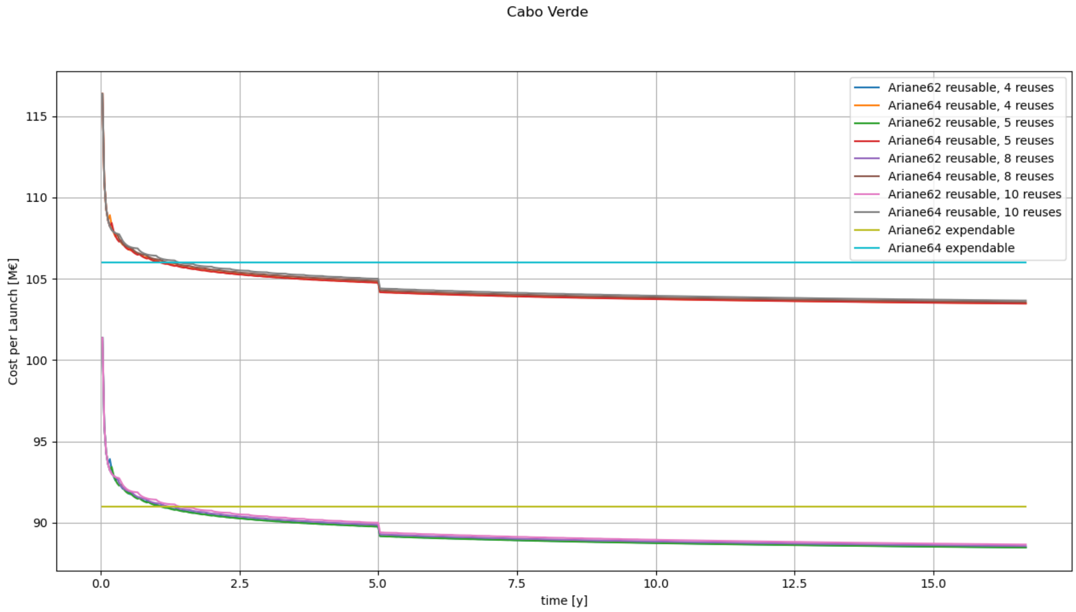

Figure 6.46: The cost per launch over the years for the reusable and expendable Ariane 6 variants for several numbers of reuses for Cabo Verde.

6.4. Comparison

With the analysis of the reusable engine bay mission completed, its performance should be compared to the alternative mission profiles to determine whether or not it is worthwhile to implement. There are four alternative mission profiles that it can be compared with:

- The standard expendable vehicle
- An unguided ballistic engine bay, entering using a deployable heat shield, then decelerating further by parachute, after which it glides on a parafoil to be caught by a helicopter before hitting the ocean. This is the mission profile selected during previous research. [10]
- Return-to-launch-site, flipping after the upper stage is disconnected, burning to cancel its velocity and return its trajectory to the launch site, burning to slow down before hitting the atmosphere, and then burning to slow to a landing.
- Downrange landing, performing a burn to slow down before hitting the atmosphere, and then burning to slow to a landing.

In subsection 6.4.1, a number of these mission profiles are rejected due to them being infeasible. Then, in subsection 6.4.2, the mission profile researched in this project is compared to the remaining mission profile.

6.4.1. Infeasible mission profiles

Upon investigating these, a number of mission profiles were deemed to be infeasible. Starting with the ballistic engine bay, there are strong concerns with the final segment of the mission, the catch by helicopter. This is a highly complex and high-risk manoeuvre, with loads at the limit of what the helicopter can sustain. For this reason, this option is rejected as infeasible.

For the return-to-launch-site, the vehicle must perform three burns. The first of these is cancelling the horizontal velocity, allowing the vehicle to drop back to the launch site. With a velocity of 6.93 [km/s] and a flight path angle of -0.8 [deg], the horizontal velocity becomes 6.86 [km/s]. Even if this were the only burn that would have to be performed, meaning the mass after the burn could go down to the first-stage dry mass of 23000 [kg], using Equation 5.1 it was calculated that this burn would require 98.3% of the propellant a fully fuelled first stage. This manoeuvre requires far too much propellant and so is infeasible.

Finally, the downrange landing is analysed using a simple forward Euler simulation. From the initial condition, the vehicle is propagated forward, with a burn starting at various moments and at various thrust angles. Even with the vehicle carrying enough propellant to fully take up the payload capacity with propellant, the vehicle cannot decelerate sufficiently before it hits the upper atmosphere, at which point the heat flux rapidly increases to unacceptable levels. The vehicle is travelling at too great a velocity at its initial conditions to use a downrange landing, and so this option is infeasible.

As it was assumed there would be no change to the Ariane 6 mission profile, a Space-X style recovery, with a lower altitude separation of the first stage for a downrange landing was not under consideration. Nonetheless, the engine bay recovery offers some advantages over such a mission. To begin with, it does not require an expensive redesign of the vehicle, being compatible with the current mission architecture. This also has a higher technology readiness, as lifting ballistic re-entries have a long history, whereas the components of a fully powered re-entry are still under development for ESA. Finally, this concept has a lower mass to be recovered.

6.4.2. Comparison

The rejection of the three infeasible missions simply leaves the expendable mission and the optimised guided engine bay mission analysed in this research. There are a number of aspects on which these will be compared to one another:

- **Cost:** The cost is one of the most important factors under consideration here, as driving down the costs of launching with the Ariane 6 is one of the main purposes of starting this research project. This will be done based on cost per launch and life cycle costs, as payload is considered separately.
- **Sustainability:** The other main reason for this research project to be started was because the Ariane 6 has received some criticism for being fully expendable, making it less sustainable in a time when this is something that requires more and more consideration.
- **Payload Capability:** Another important aspect is the payload capability. Lower capability means that certain missions will not be able to be launched with a vehicle or that some missions, which could have been launched together, must now be split up.
- **Reliability:** As each launch, and its payload, is very expensive, it is desirable that each element of the mission, or the whole mission, if possible, be flight-proven and reliable.
- **Complexity:** A more complex mission has more points of failure, raising the chance of an expensive mishap. Less complex missions are more desirable.

The two mission profiles will now be compared on each of these points. This comparison can be found in Table 6.8. From this comparison, it is determined that implementing the guided ballistic engine bay is worthwhile if a few caveats can be managed.

First of all, the reusable case has a reduced payload capacity, which is especially impactful for the A62 case, which is left with less than half its initial payload capability and therefore makes launches very expensive per kg launched. This may be compensated for by launching in the A64 configuration more often. This does, however, raise the costs again somewhat and will counteract any sustainability gains made. The way to deal with this is to make the boosters reusable as well, reducing their costs and environmental impact, making using more boosters much less impactful.

Secondly, using the reusable case impacts the payload capacity quite significantly, causing a reduction in capability to 45.37% and 73.88% of the expendable cases for A62 and A64, respectively. This can be somewhat compensated for by launching the A64 more often, especially if the boosters are made reusable. Furthermore, the payload capacity is not always used fully, although the missions where significantly less than the payload mass is launched have until now generally been for SSO or MEO missions. In addition, as every fifth launch is not recovered and therefore does not need to bring on board the propellant needed for the correction of the initial state, the payload capability of every fifth mission is raised to 87.93% and 94.23% of the expendable case for A62 and A64, respectively, assuming the propellant is the only saving that can be made. If other pieces of hardware can be removed for this final mission, the situation may be even better for this fifth launch. Notably, this fifth launch would have sufficient payload capability to launch the heaviest payload an Ariane 6 has ever launched, which is approximately 20000 [kg]. As the fifth launch would occur every other month, the payload capability issue can also be alleviated with clever scheduling of payloads.

Third of all, this is only worth it if a sufficient launch cadence can be reached. 30 launches per year was stated to be the number from which reusability would be achievable concerning the downtime of the manufacturing teams. In addition, having a lower launch cadence delays the pay-off of the development costs and the break-even points for the cost per launch and life cycle costs. Until now, Ariane 6 vehicles have made seven launches in approximately two years, for roughly 3.5 launches per year. If the current schedule is kept, 2026 will see 6 total launches, with 9 launches already scheduled for 2027. Although the number is rising, this is not yet anywhere near the required number of launches per year. On the other hand, the reduction in launch costs and the improved sustainability may make Ariane 6 a more attractive launcher option, drawing more customers and filling some of the demand that is needed to operate the reusable vehicle.

Furthermore, this research represents an investigation into the feasibility and cost of this trajectory but has not taken a close look at the vehicle itself. Performing a full 6-DOF simulation of the vehicle should provide further insight into matters like the controllability of the bank angle using the RCS thrusters and the stability of the vehicle during re-entry. The vehicle has been assumed to have sufficient aerodynamic authority and stability, but this has not been demonstrated. This is a matter that requires a good deal of attention, as the vehicle flying at a constant angle of attack to produce lift, and the guidance algorithm to rotate that lift vector are critical to guiding the vehicle to the target location. It should also be investigated what the uncertainties of the state of the vehicle may be following the disconnection from the core stage, and whether the resulting error in the final distance to the target can be overcome during the final descent.

In addition, the hardware required for this mission is not yet available. Several elements are under development, but these must be finished and tested before this mission profile can be considered. This is namely regarding the new Prometheus engine, which makes it possible to even consider this mission profile, but this also regards control fins, landing legs and more.

Finally, this research project has worked to demonstrate the feasibility of achieving a desired trajectory but has not looked into hardware. The propellant tanks, tanks for RCS propellant and deployable heat shield will all need to be stored within the vehicle. The control fins, landing legs and RCS thrusters will need to have a place on the outside of the vehicle, where they do not interfere with the operation of the launch vehicle or the attachment of the boosters.

Table 6.8: The comparison between the expendable Ariane 6 and the guided ballistic engine bay variant.

Mission Profile	Expendable Ariane 6	Guided Ballistic Engine Bay
Cost	100 [M€] for A62, 115 [M€] for A64.	Significantly more expensive at first, for a landing at Cabo Verde with 30 launches per year, this case becomes cheaper per launch after just over a year. With these same parameters, the life cycle costs become cheaper in 3.5 years. After 15 years, the cost reaches 88.47 [M€] for A62 and 103.47 [M€] for A64. The life cycle costs come out to approximately 845 [M€] less for each version compared to the expendable case.
Sustainability	A fully expendable mission, which is a point of criticism levelled at this design. In today's age, sustainability is very important, also in spaceflight, and so a fully expendable mission is possibly adding more pollution and using more materials that could otherwise be reused.	Debatably more sustainable. The engine bay is recovered for refurbishment, which is certainly a point for this profile, but this does require a somewhat intensive manoeuvre in space. Furthermore, it may require launching in an A64 configuration when otherwise an A62 launch would suffice, meaning that recovering the engine bay may need two additional boosters to be added to the vehicle. If these are also kept expendable, this is likely detrimental to sustainability as a whole. If these could also be recovered, this would not be as impactful and likely would make the mission more sustainable compared to the fully expendable case.
Payload Capability	10350 [kg] for A62, 21650 [kg] for A64	For the best-case scenario, a landing in Cabo Verde, a total payload mass penalty of 5654.3 [kg] is incurred, bringing the maximum A62 payload to 4695.6 [kg] and the maximum A64 payload to 15995.6 [kg], 45.37% and 73.88% of the expendable cases, respectively. It should be noted that not every launch uses up the whole payload capability, and so this may not be as impactful as would initially seem. Furthermore, every fifth launch is expendable and therefore does not require the propellant for the correction of the initial state, raising the payload capability significantly, high enough that the heaviest payload launched on the Ariane 6 could still be launched.
Reliability	Has been demonstrated to work several times now, flying seven missions, with six successes and one, the test launch, partial failure.	Has not been flight proven as a whole. ESA is currently developing or testing individual components that would be needed for the mission: the Prometheus engine, landing legs and deployable heat shields. Control fins are not yet under development. The guidance algorithm, however, has extensive heritage.
Complexity	No added complexity, the current mission forms the baseline.	The disconnection, re-entry, flip, final approach and landing all add complexity, and therefore risk of failure.

6.5. Limitations

There are several assumptions that were made in the process of setting up the simulation, which impose some limitations on the analysis performed in this research.

To begin with, a 3-DOF model was used for the simulation, not a 6-DOF model. This therefore does not capture any dynamics related to the three neglected degrees of freedom. As this research investigated the feasibility of the trajectory, this does not pose significant issues, but for future research a full 6-DOF should be used to also investigate the rotational dynamics and the stability of the vehicle.

Fairly general estimations were made on the mass of the vehicle, the range to the target location the vehicle can cross during the final descent and the propellant needed for the final deceleration. The aerodynamic analysis in FOSTRAD was done using a simplified model of the vehicle which excludes external components shadowed from the flow. This causes some uncertainty on the results, which should be addressed if more accurate values can be obtained.

During the setup of the simulation, it was assumed the vehicle could stably fly at a constant angle of attack throughout the flight. It was also assumed the vehicle could achieve a roll rate of 5 [deg/s]. Furthermore, the used atmospheric model was the best available, but certainly any inaccuracies there affect the controllability of the vehicle. The optimisation assumes the vehicle has sufficient aerodynamic authority. Verification of this, and the sizing of the control fins and actuators fall outside the scope of this research.

In terms of the estimation of the costs, these are naturally reliant on the accuracy of the used models. These are based on statistics, but the actual relations may be somewhat different. Furthermore, the input values used, such as the learning factors, expendable vehicle cost and man year cost, all greatly impact the estimated costs. In addition, the simplification is made that the expendable Ariane 6 costs no longer improve over time.

Finally, as this research only investigated the feasibility of the trajectory, no attention has been placed on the structure or interior of the vehicle. Once it is clear how the introduction of the new Prometheus engine has affected the layout of the vehicle, it should be seen if sufficient volume is available to accommodate the propellant tanks, control fin actuators, heat shield etc.

7

Conclusion & Recommendations

7.1. Conclusion

To act on criticism that has been levelled against the Ariane 6 in the past, this research project aimed to investigate whether it is feasible and worthwhile to recover its engine bay for refurbishment. Following the establishment of several possible mission profiles, the option to use a guided ballistic re-entry for the engine bay is selected. This led to the main research question to be answered in this research:

- ***Is a transatlantic re-entry a feasible and cost-effective method to recover the Ariane 6 engine bay?***

To answer this question, a 3-DOF TudatPy simulation is set up to achieve a physical model accuracy of 10 [m]. A guidance algorithm based on the Apollo entry guidance is implemented to guide the vehicle to a target location during re-entry. After investigating the behaviour of the vehicle, it was determined that a burn would be necessary after disconnecting from the first stage to correct the initial state in order to send the vehicle to a useful target location. Five such potential locations were selected as targets for the trajectory. After this, a multidisciplinary design optimisation is set up, with inputs, objectives and constraints selected for it. A Monte Carlo analysis is performed, both one-at-a-time and all-at-once, to gain insight into the effect of the inputs on the outputs, after which the full optimisation is performed. The results were assessed, eliminating some of the target locations and determining which performed best. A sensitivity analysis was performed, finding that the results are quite sensitive to uncertainties in the initial state but that these may still be overcome. The results of the optimisation were used to estimate the costs of the mission, and this was taken into the comparison of the selected mission profile to the alternatives laid out earlier. During the research, each of the subquestions laid out in section 3.3 was answered:

- ***Is it possible to have the engine bay cross the Atlantic during its re-entry?***
  It appears to be possible from a trajectory point of view, to cross the Atlantic during the re-entry, although to reach a useful location, a burn is needed after the engine bay has disconnected from the first stage, as can be seen in section 5.6.
- ***What useful locations could be reached in such a flight?***
  Five useful target locations were selected in section 5.7: Natal, Cabo Verde, The Canarias, The Azores and Les Mureaux. Of these, Les Mureaux and The Azores were rejected after the optimisation in section 6.1, as the missions to these locations did not result in a feasible option. All three remaining options are feasible, although The Canarias has a very significant payload penalty associated with it, making the other two more desirable. Cabo Verde was identified as the best target location in section 6.3.
- ***How can such a flight be guided?***
  A guidance algorithm based on the Apollo re-entry guidance algorithm is selected to control the bank angle for downrange guidance, with heading angle deadbands for lateral guidance, as shown in section 2.3. A system was set up to generate the reference trajectory this guidance

algorithm uses in section 5.4. The best values of the guidance overcontrol and deadband values, K, c_0 and c_1 vary strongly based on the target location and the initial conditions, as seen in section 6.1.

- ***What is the recommended re-entry trajectory for the engine bay?***
  The best identified trajectory is one with Cabo Verde as its target. From the initial state, a correction is performed on the vehicle to set the initial velocity $V = 6926.37$ [m/s], the initial heading angle 68.551 [deg] and the initial entry angle -0.621 [deg]. The guidance algorithm then uses guidance overcontrol value $K = 2.108$ [-] and deadband values $c_0 = 4.84$ [deg] and $c_1 = 1.506 \cdot 10^{-7}$ [deg/(m/s^2)]. This trajectory, selected in section 6.1 as the best option, minimises the propellant mass needed for the initial trajectory while staying within the load constraints, the constraint on the final velocity, and the constraint on the final distance to the target location.
- ***How does the performance of the guided engine bay mission profile compare to alternative mission profiles and the non-reusable variant?***
  In section 6.4, it was found that the return to launch site and downrange landing mission profiles appear to be infeasible for the Ariane 6 first stage. The unguided ballistic engine bay is also rejected due to the significant concerns with the catch by a helicopter segment of the mission. This leaves the comparison between the case investigated in this research and the current fully expendable case. The reusable case seems more expensive initially but becomes cheaper per launch after just over a year and cheaper over its life before 3.5 years, with a launch cadence of 30 launches per year. It is arguably more sustainable, as the engine bay is recovered and refurbished, but this does require a burn to correct the initial state and for landing and may need additional boosters when this may not have been necessary for the expendable case. There is a significant payload penalty, but this can potentially be managed by adding extra boosters more often and scheduling heavier payloads for every fifth launch, which is not recovered and so has more payload capacity. As it is only conceptual currently, it cannot be given the same reliability as the expendable case, although elements of the hardware are undergoing development and testing currently. The reusable mission is more complex, potentially adding more risk of failure.

With each of the subquestions answered, a return can be made to the main research question this project set out to answer:

- ***Is a transatlantic re-entry a feasible and cost-effective method to recover the Ariane 6 engine bay?***

The conclusion drawn from the results is that on a trajectory-level analysis, this indeed appears to be a feasible and cost-effective method to recover the Ariane 6 engine bay. This option may therefore be worth adopting, if a number of concerns can be dealt with:

- The hardware to achieve the mission is not all mission ready. The landing legs, control fins, but especially the Prometheus engine must be available and tested to achieve this mission profile.
- The vehicle must have sufficient aerodynamic control authority and stability, both during re-entry and during the final descent. This should be demonstrated, so that it can be evaluated what is actually possible with values that were not estimated. Specifically the final descent is important considering what uncertainties may exist in the initial state of the vehicle, as these may cause the final distance to the target location to increase.
- The boosters are not reusable, raising the cost and lowering the sustainability, as the reusable mission may use more of them more often. Making them reusable improves this issue significantly.
- The payload penalty is substantial, and so it must be considered if this is acceptable. Not every launch uses up the entire payload capacity, launching in the four-booster configuration more often can add more payload capacity, which is not a significant issue if they are reusable as well, and scheduling heavier payloads for the final launch, when the reusable engine bay is expended, can all help deal with this issue.
- A sufficient launch cadence is reached, at least 30 launches per year. If this is not done, the development costs, spread over the first five years of operations, will drive up the costs, significantly delaying the return on investment. This will also cause significant downtime for the manufacturing

teams, which is not desirable. The economic viability of this mission profile depends highly on the ability to reach a sufficient launch cadence.

- The engine bay must be able to fit the additional hardware inside and outside of itself. In this research project, it was only investigated if it would be feasible and worthwhile to use this mission profile, but for it to work, it must also be investigated if there is sufficient space within the vehicle and if the outside hardware does not interfere with the operations of the launch vehicle. If this is not the case, this may require a redesign of the engine bay, at which point this approach may no longer be worth the costs invested in it.

Should these concerns be resolved, the mission profile investigated in this research offers an attractive option for introduction of reusability. In this case it may be worth analysing in more depth, to consider adopting it for the Ariane 6 and future vehicles.

7.2. Recommendations for future research

This final section contains recommendations for future research using the ballistic guided engine bay mission.

First and foremost, to determine if further research into this topic is warranted, the available volume and layout of the engine bay using the new Prometheus engine should be determined, and with that, it can be seen if sufficient space is available to fit the hardware needed for this mission. If this is not the case, a redesign of the engine bay likely adds more development costs than the gains could effectively pay off, making this not a feasible option to consider for the Ariane 6.

This research focused on the re-entry segment of the mission and so did not look at the final approach after hitting 30 [km] of altitude. The controllability of the vehicle and the guidance used to do this should be investigated, specifically to determine if the 5 [km] range requirement used during the optimisation is achievable or if even a wider range can be taken. It should also be looked at what precision can be achieved and to what degree the deviations caused by uncertainties in the initial state can be compensated for. If these provide sufficient accuracy, it may be considered to dispense with guiding the vehicle to a landing site on land and instead having it attempt a landing on a barge at sea.

The boosters being expendable currently significantly counteracts gains made on costs and sustainability by using the reusable engine bay. As the boosters reach a lower altitude and velocity, these are less challenging to recover in terms of their trajectory, but because they are solid rocket boosters, they do not have the capability to re-ignite. This problem is worthwhile to investigate, as it could further reduce costs and improve sustainability.

The mass estimates made in this project were necessarily quite rough, especially when it came to the control fins, landing legs and RCS propellant. Revisiting the analysis performed here with better mass estimates will shed more light on what exactly the costs and payload penalty are for these missions.

The analysis in this research was performed for a LEO mission, which accounts for three of the seven Ariane 6 missions launched so far. The feasibility of this mission profile should also be investigated for other mission profiles. Specifically, Sun-Synchronous Orbits (SSO's) and Geostationary Transfer Orbits (GTO's) are common among past and planned mission payloads, and these will have different initial conditions and feasible target locations. If these also prove viable, this would lend itself to this mission profile being more worthwhile, whereas if this does not work for these missions, it likely is not worth making this change for LEO missions alone.

Finally, with the development of the next generation of Ariane vehicles in mind, this is an attractive option if the mission profile of the launch vehicle is kept similar to the Ariane 6. The first stage reaches quite a significant altitude and velocity, making other options for recovery infeasible. If it is decided that the first stage should once again reach such a state, this is an effective way to recover the engine bay. In addition, even if this mission profile would not be worth it for the Ariane 6 due to requiring a significant redesign, this would not be a problem for the next generation of vehicles, as these could be designed with this mission segment in mind. If, on the other hand, the mission profile is altered significantly, with the first stage reaching lower altitudes and velocities at the time of separation, a different mission profile may be preferable, such as a propulsive downrange landing.

References

[1] T. Anthony. *An Apollo entry guidance implementation in Python*. 2021. URL: `https://github.com/thomasantony/msl-apollo-entry-guidance` (visited on 04/19/2026).

[2] E. Berger. *Ariane chief seems frustrated with SpaceX for driving down launch costs*. 2018. URL: `https://arstechnica.com/science/2018/05/ariane-chief-seems-frustrated-with-spacex-for-driving-down-launch-costs/` (visited on 05/30/2026).

[3] E. Berger. *Europe's Ariane 6 rocket is turning into a space policy disaster*. 2023. URL: `https://arstechnica.com/science/2023/04/europes-ariane-6-rocket-is-turning-into-a-space-policy-disaster/` (visited on 06/30/2025).

[4] Blue Origin. *New Shepard*. 2025. URL: `www.blueorigin.com/new-shepard` (visited on 06/30/2025).

[5] M.J. Bozack. "SRB seawater corrosion project". In: *Contractor Report* (1991).

[6] T. Bruno. *The Secrets of Rocket Design Revealed*. 2023. URL: `https://medium.com/%40ToryBrunoULA/the-secrets-of-rocket-design-revealed-e2c7fc89694c` (visited on 07/11/2025).

[7] T. Bykerk et al. "Retro-propulsion in rocket systems: Recent advancements and challenges for the prediction of aerodynamic characteristics and thermal loads". In: *Progress in Aerospace Sciences* (2024).

[8] V. Carandente, G. Zuppardi, and R. Savino. "Aerothermodynamic and stability analyses of a deployable re-entry capsule". In: *Acta Astronautica* (2013).

[9] R. Childress-Thompson, P. Farrington, and D. Thomas. "A Framework for Assessing the Reusability of Hardware (Reusable Rocket Engines)". In: *JANNAF Programmatic and Industrial Base (PIB) Meeting* (2016).

[10] C. Dek et al. "A recovery system for the key components of the first stage of a heavy launch vehicle". In: *Aerospace Science and Technology* (2020).

[11] C. Dek et al. "VuAB Recovery". In: *TU Delft DSE report* (2018).

[12] J. Dijkstra, J. Pedra, and L. van der Poll. *PropagationandOptimization*. 2024. URL: `https://github.com/jeroen1066/PropagationandOptimization` (visited on 05/17/2026).

[13] N. T. Drenthe. "$OLSTICE". MA thesis. TU Delft, 2016.

[14] ESA. *Ariane 6*. 2025. URL: `https://www.esa.int/Enabling_Support/Space_Transportation/Launch_vehicles/Ariane_6_overview` (visited on 06/30/2025).

[15] ESA. *Ariane 6*. 2025. URL: `https://www.esa.int/Enabling_Support/Space_Transportation/Future_space_transportation/Prometheus` (visited on 07/11/2025).

[16] ESA. *Ariane 6 – made in the Netherlands*. 2025. URL: `https://www.esa.int/Enabling_Support/Space_Transportation/Ariane/Ariane_6_made_in_the_Netherlands` (visited on 04/18/2026).

[17] ESA. *Development of a Flexible Thermal Protection System for Deployable Inflatable Heatshields and Hypersonic Decelerators*. 2023. URL: `https://nebula.esa.int/content/development-flexible-thermal-protection-system-deployable-inflatable-heatshields-and` (visited on 08/04/2025).

[18] ESA. *The engines of Ariane 6*. 2024. URL: `https://www.esa.int/Enabling_Support/Space_Transportation/Ariane/The_engines_of_Ariane_6` (visited on 04/18/2026).

[19] ESA. *Themis*. 2025. URL: `https://www.esa.int/Enabling_Support/Space_Transportation/Themis` (visited on 06/30/2025).

[20] Alessandro Falchi et al. "FOSTRAD: An advanced open source tool for re-entry analysis". In: *15th Reinventing Space Conference*. 2017.

[21] C. A. Graves and J. C. Harpold. *Apollo Experience Report Mission Planning For Apollo Entry*. Tech. rep. NASA, 1972.

[22] G. A. Harvey. “SIMULATED REENTRY HEATING BY TORCHING”. In: *NASA Langley Research Center* (2008).

[23] D. Isakeit et al. *The Atmospheric Reentry Demonstrator*. Tech. rep. ESA, 1998.

[24] D. Izzo and F. Biscani. *Welcome to PyGMO*. 2015. URL: `https://esa.github.io/pygmo/` (visited on 08/04/2025).

[25] D. E. Koelle. “THE TRANSCOST-MODEL FOR LAUNCH ESTIMATION AND ITS APPLICATION SYSTEMS ANALYSIS”. In: *Acta Astronautica* (1984).

[26] A. Krammer, L. Blecha, and M. Lichtenberger. “Fin actuation, thrust vector control and landing leg mechanisms design for the RETALT VTVL launcher”. In: *CEAS Space Journal* (2021).

[27] M. A. Miller. “Managing Energy and Mode Transitions in Skip Entry Guidance for Lunar Return Trajectories”. In: *S.B., Aeronautical and Astronautical Engineering Massachusetts Institute of Technology* (2007).

[28] A. Patureau de Mirand, J. Bahu, and E. Louaas. “Ariane Next, a vision for a reusable cost efficient European rocket”. In: *8TH EUROPEAN CONFERENCE FOR AERONAUTICS AND SPACE SCIENCES (EUCASS)*. 2019.

[29] A. Mohammadi and M. Tayefi. “Moving mass control system in conjunction with brain emotional learning-based intelligent control for rate regulation of suborbital reentry payloads”. In: *Proceedings of the Institution of Mechanical Engineers Part I Journal of Systems and Control Engineering* (2012).

[30] E. Mooij. *AE4-870B_10 Planetary Entry and Descent*. `https://brightspace.tudelft.nl/d2l/le/content/498858/viewContent/3063297/View`. 2023. (Visited on 07/18/2025).

[31] E. Nakano et al. “Water Landing Impact of Recovery Space Capsule: A Research Overview”. In: *The Journal of Space Technology and Science* (2013).

[32] NASA. *Illustration of Low-Earth Orbit Flight Test of an Inflatable Decelerator (LOFTID)*. 2022. URL: `https://scitechdaily.com/nasas-successful-launch-deployment-and-retrieval-of-loftid-an-innovative-inflatable-heat-shield/` (visited on 07/19/2025).

[33] A. Parsonson. *CNES Boss Blames Contractors for Ariane 6 Being Too Expensive*. 2024. URL: `https://europeanspaceflight.com/cnes-boss-blames-contractors-for-ariane-6-being-too-expensive/` (visited on 04/18/2026).

[34] M. A. Rafiq et al. “ArianeNEXT Preliminary Design Report”. In: *TU Delft Collaborative Space System Design Project* (2025).

[35] Rocket Lab. *Neutron*. 2025. URL: `https://rocketlabcorp.com/launch/neutron/` (visited on 06/30/2025).

[36] SkywalkerPL. *Ariane 62 (left) and Ariane 64 (right), final design*. 2017. URL: `https://en.wikipedia.org/wiki/Ariane_6#/media/File:Ariane_62_and_64.svg` (visited on 07/14/2025).

[37] K. Smith. *PREDICTIVE LATERAL LOGIC FOR NUMERICAL ENTRY GUIDANCE ALGORITHMS*. 2016. URL: `https://www.google.com/url?sa=t&source=web&rct=j&opi=89978449&url=https://ntrs.nasa.gov/api/citations/20160001182/downloads/20160001182.pdf&ved=2ahUKEwjtlaPDsPqTAxWehv0HHRE3LlgQFnoECAwQAQ&usg=AOvVaw2Fi-JywxHiNYjQlsZmyLgP` (visited on 04/19/2026).

[38] R. Smith. *Europe Complains: SpaceX Rocket Prices Are Too Cheap to Beat*. 2018. URL: `https://www.fool.com/investing/2018/06/02/europe-complains-spacex-rocket-prices-are-too-chea.aspx` (visited on 04/18/2026).

[39] SpaceX. *Falcon 9*. 2025. URL: `https://www.spacex.com/vehicles/falcon-9/` (visited on 06/30/2025).

[40] Tudat Team. *TudatPy API Reference*. 2025. URL: `https://py.api.tudat.space/en/latest/index.html#` (visited on 08/04/2025).

[41] Tudat Team. *Welcome to Tudat Space!* 2025. URL: `https://docs.tudat.space/en/latest/` (visited on 08/04/2025).

[42] C. Thies. "Investigation of the landing dynamics of a reusable launch vehicle and derivation of dimension loading for the landing leg". In: *CEAS Space Journal* (2022).

[43] O. Uyanna and H. Najarfi. "Thermal protection systems for space vehicles: A review on technology development, current challenges and future prospects". In: *Acta Astronautica* (2020).

[44] Cyrille Vanlerberghe. *Ariane 6 : la version de la dernière chance*. 2014. URL: `https://www.lefigaro.fr/sciences/2014/09/05/01008-20140905ARTFIG00351-ariane-6-la-version-de-la-derniere-chance.php` (visited on 07/11/2025).

[45] G. Waschinski and C. Sohier. *Arianespace successfully launches 32 Amazon Leo satellites with the first Ariane 64*. 2026. URL: `https://press.ariane.group/arianespace-successfully-launches-32-amazon-leo-satellites-with-the-first-ariane-64/?lang=eng` (visited on 05/30/2026).

[46] J. R. Wertz. "ECONOMIC MODEL OF REUSABLE VS. EXPENDABLE LAUNCH VEHICLES". In: *IAF Congress, Rio de Janeiro, Brazil Oct. 2–6, 2000*. 2000.

[47] Wikipedia. *Prometheus (rocket engine)*. 2026. URL: `https://en.wikipedia.org/wiki/Prometheus_(rocket_engine)` (visited on 04/18/2026).

[48] Wikipedia. *Vulcain (rocket engine)*. 2026. URL: `https://en.wikipedia.org/wiki/Vulcain_(rocket_engine)` (visited on 04/18/2026).

A

Design Option Tree

To narrow down the many options for reusability in section 3.1, a design option tree is used. In this tree, all options that are thought of are noted, before pruning the non-feasible options. This way the chance of overlooking a design option is minimised. The design option tree used can be found in Figure A.1.

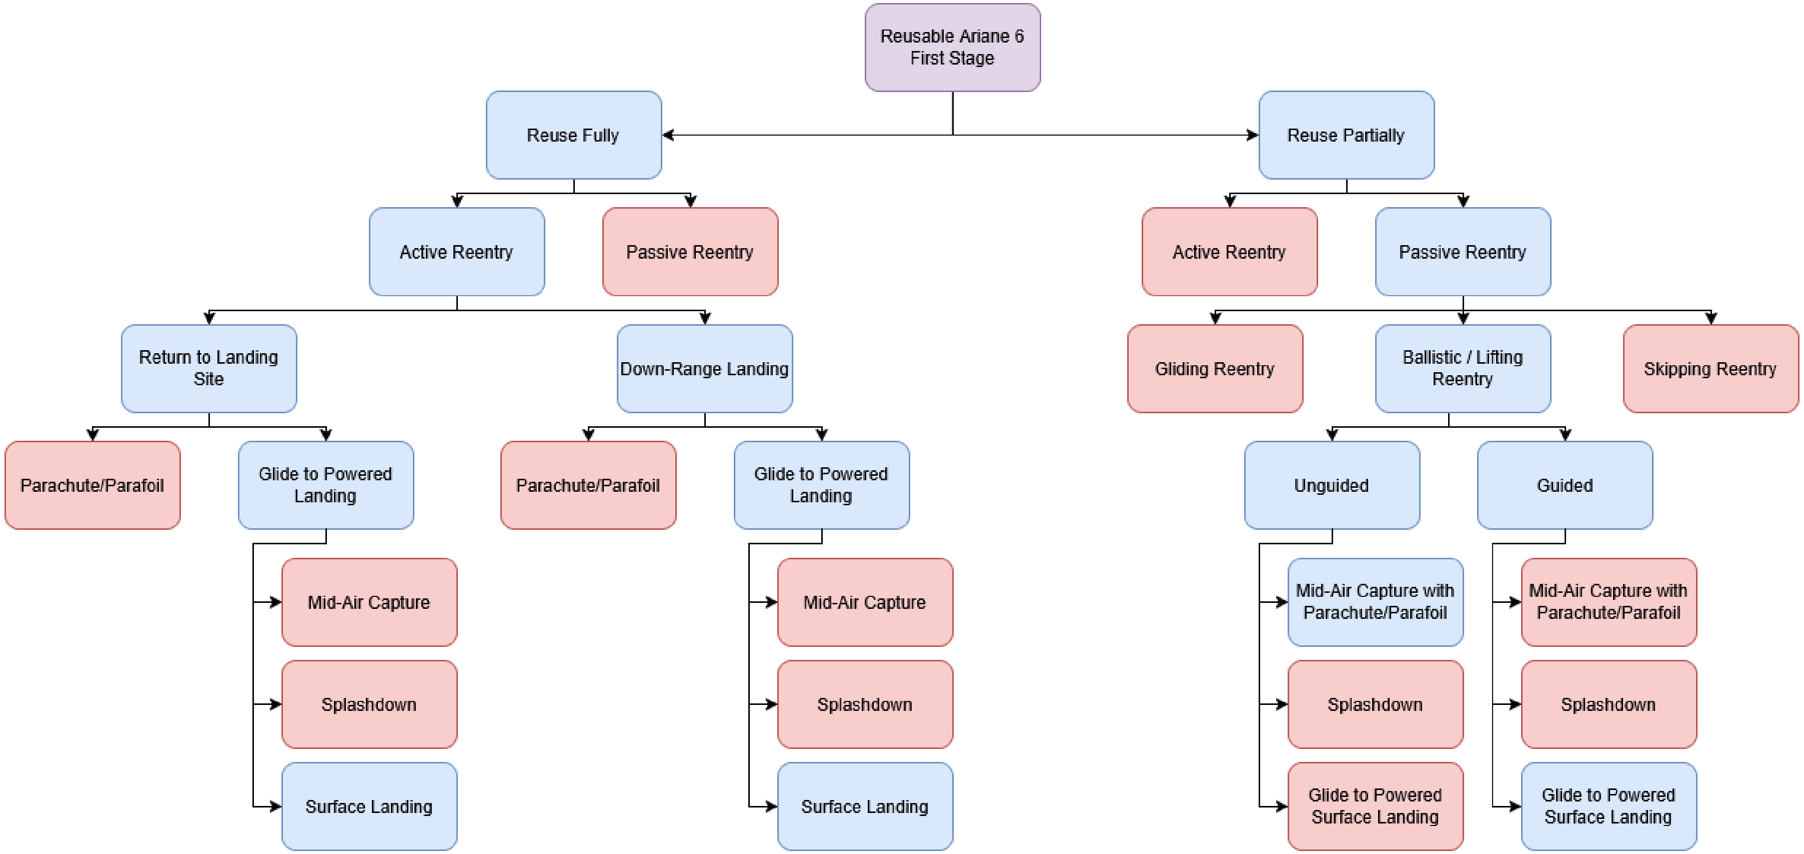

Figure A.1: The Design Option Tree used to select mission profiles under consideration for the research project. Purple represents the objective, red the pruned, non-feasible branches, and blue the feasible options.

B

Monte Carlo Plots

To gain insight into what effect the input parameters have on the outputs and constraints, an MC analysis is performed, as described in section 5.9. In that section, the outcomes for Cabo Verde are displayed, in this appendix, the results for the remaining target locations can be found. In section B.1, the plots for the one-at-a-time MC analysis can be found, while in section B.2, the plots for the all-at-once MC analysis can be found.

B.1. One-at-a-Time Monte Carlo

This section contains the plots produced during the one-at-a-time MC using seed 42 for all target locations except for Cabo Verde, as these can already be found in subsection 5.9.1.

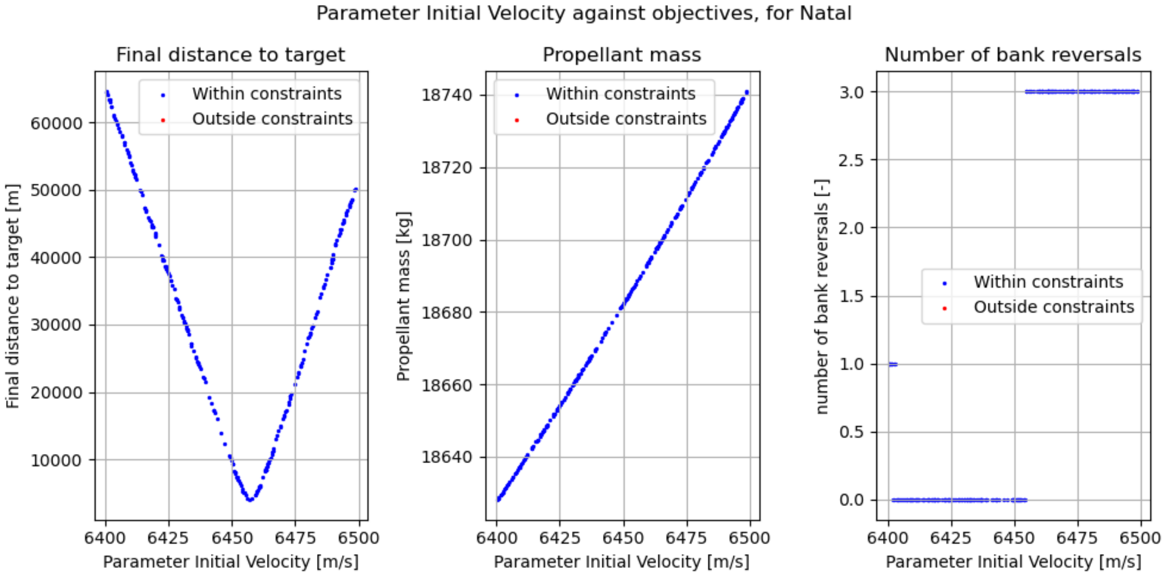

Figure B.1: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial velocity for Natal.

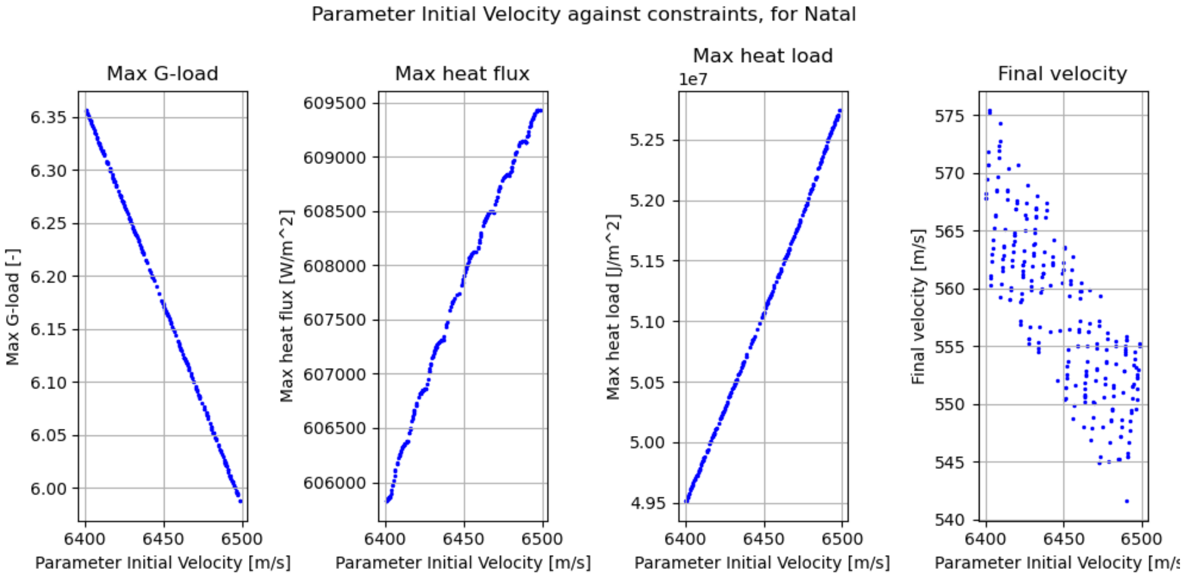

Figure B.2: The maximum g-load, maximum heat flux and total heat load produced by simulations in the single MC variation of the initial velocity for Natal.

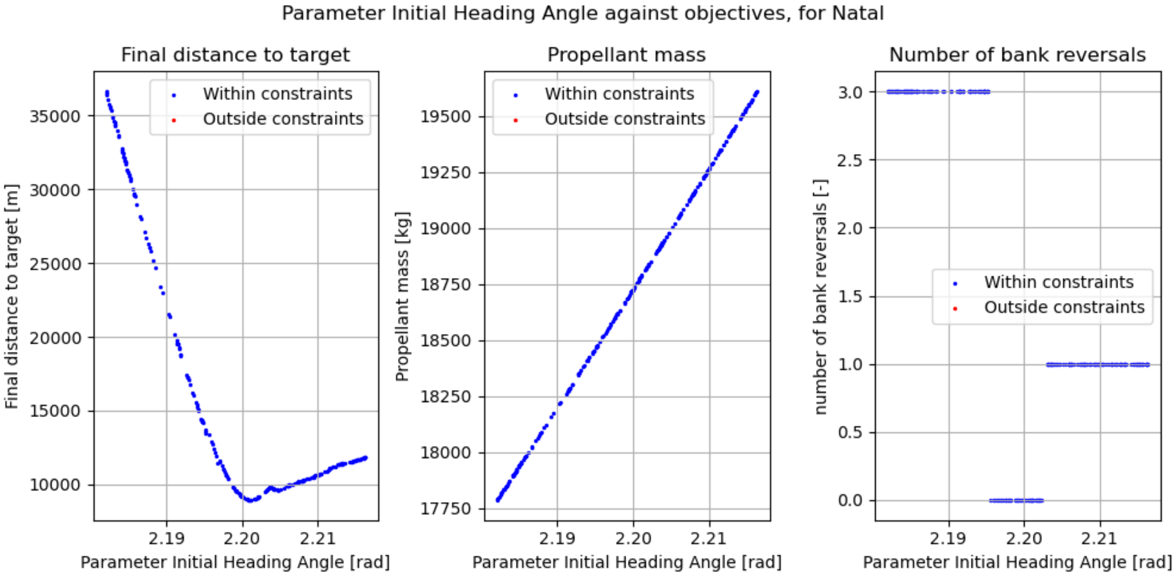

Figure B.3: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial heading angle for Natal.

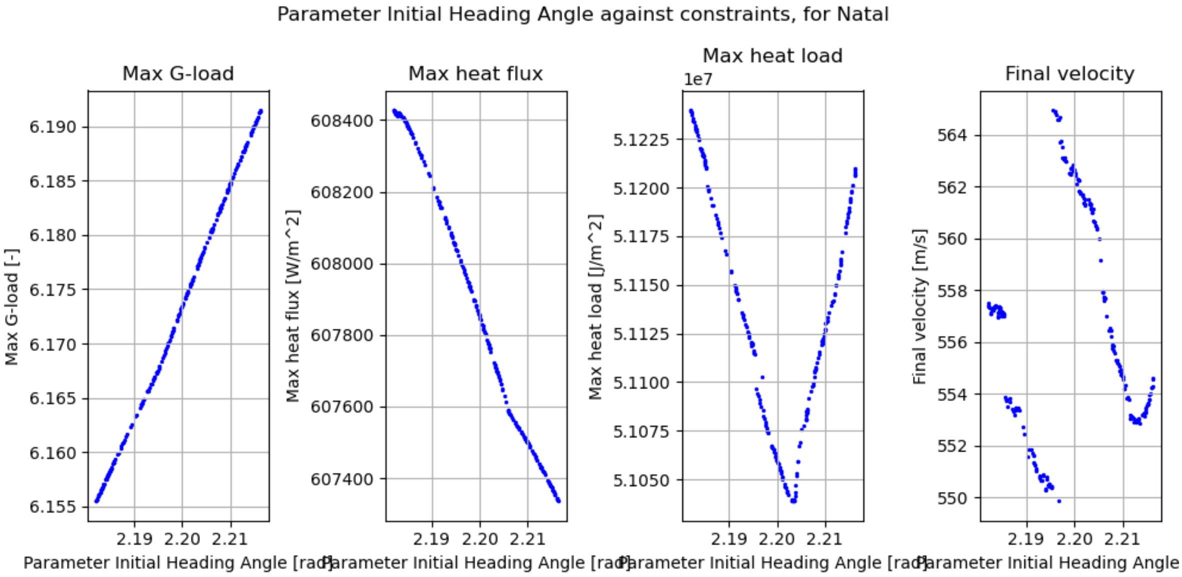

Figure B.4: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial heading angle for Natal.

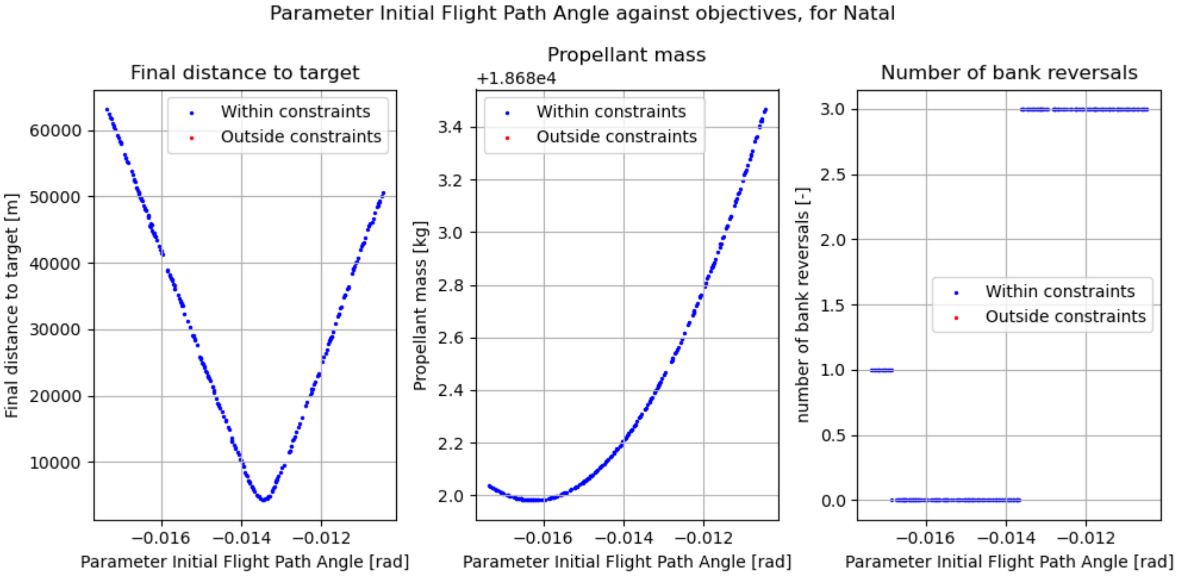

Figure B.5: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial flight path angle for Natal.

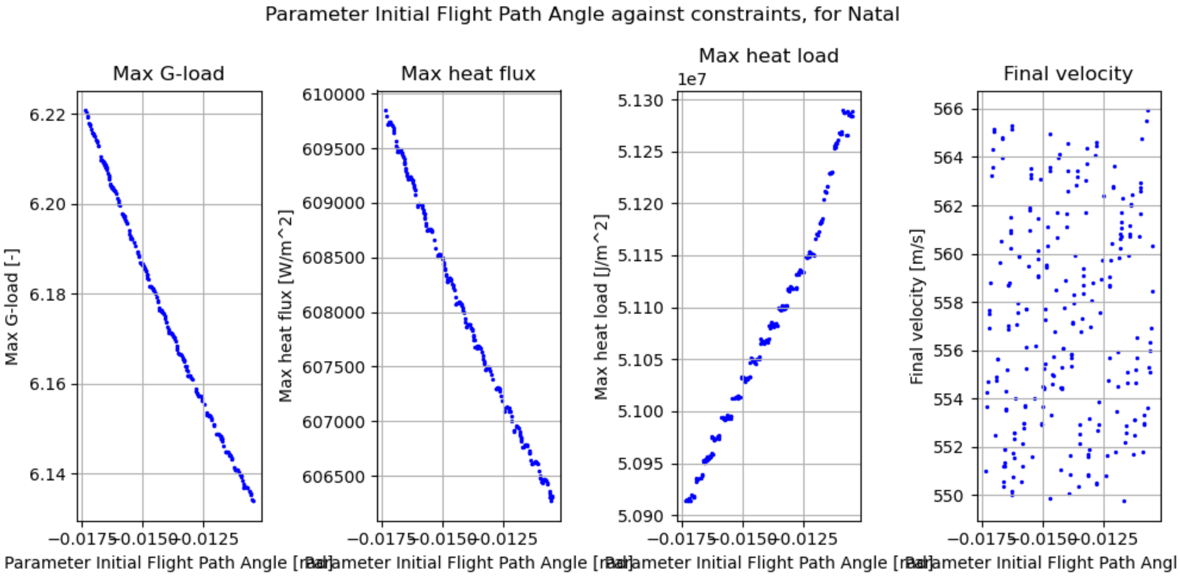

Figure B.6: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial flight path angle for Natal.

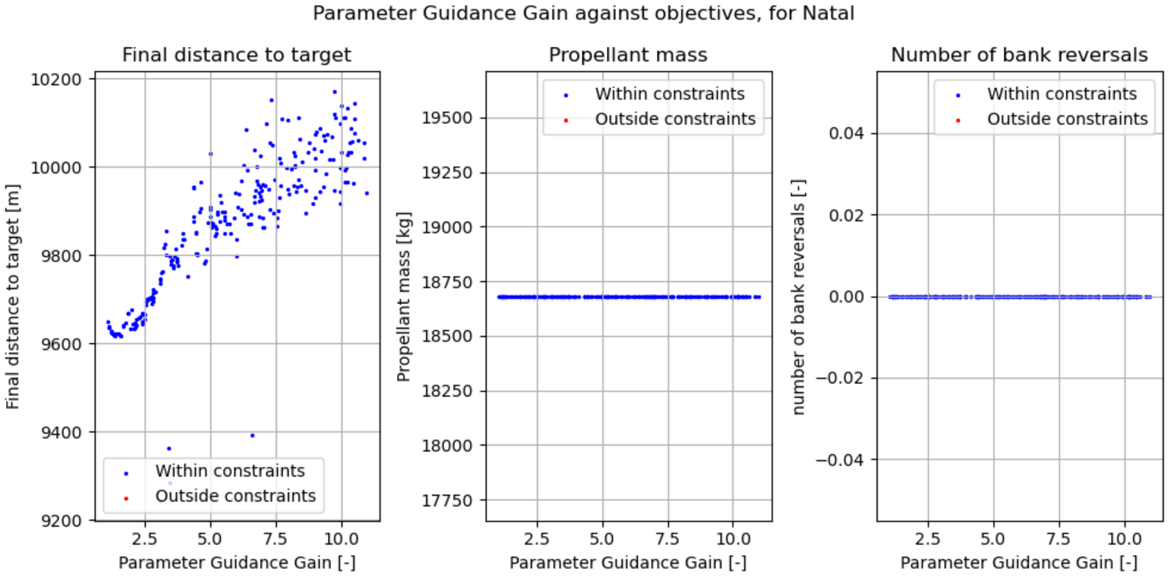

Figure B.7: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the overcontrol guidance gain for Natal.

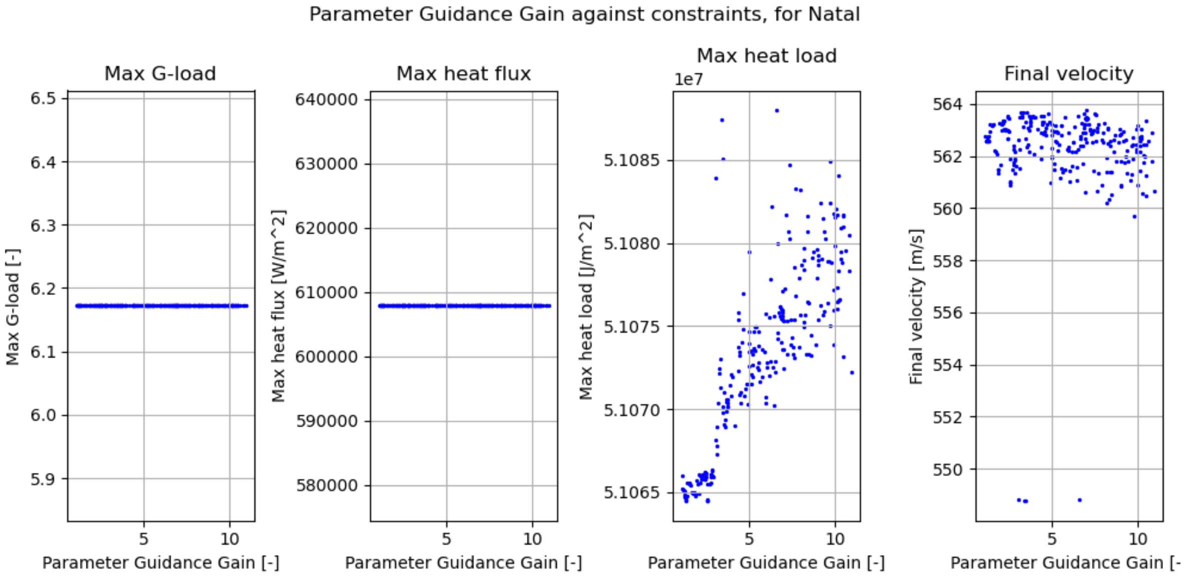

Figure B.8: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the overcontrol guidance gain for Natal.

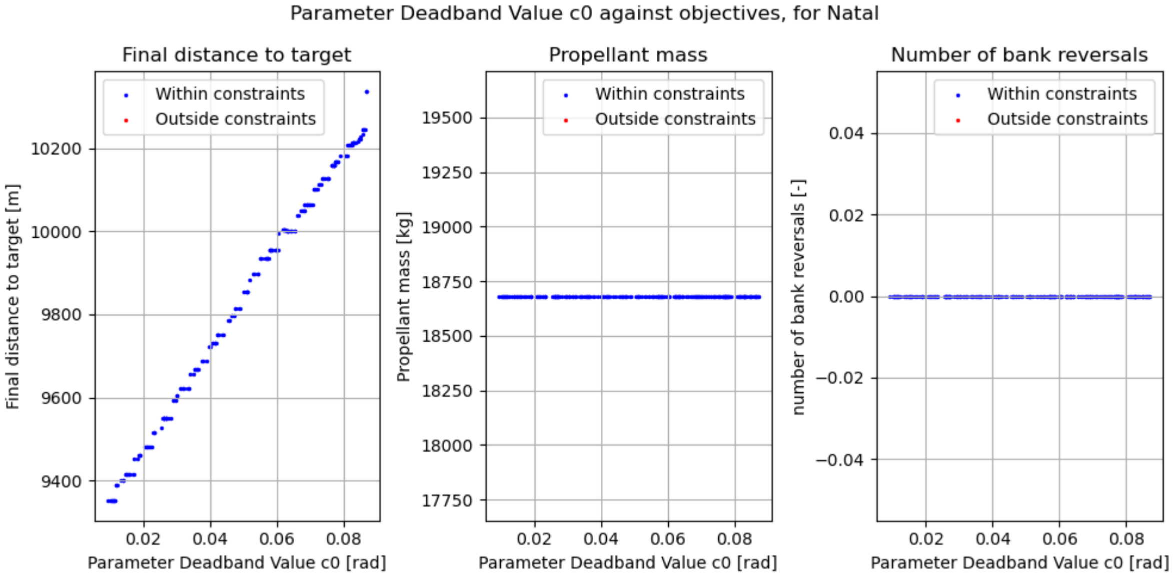

Figure B.9: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c0$ for Natal.

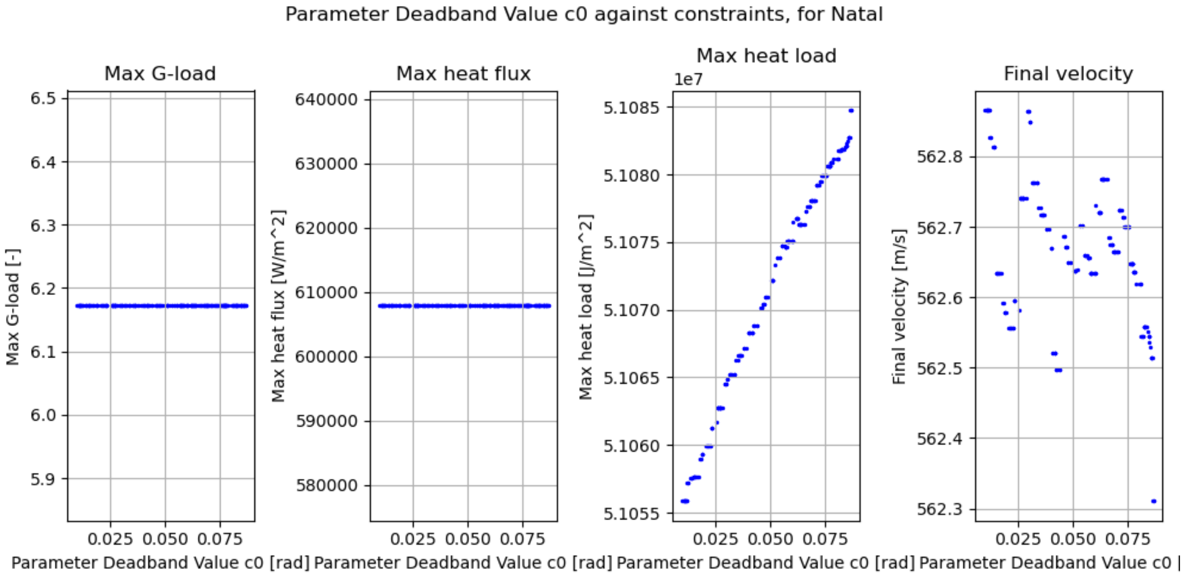

Figure B.10: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c0$ for Natal.

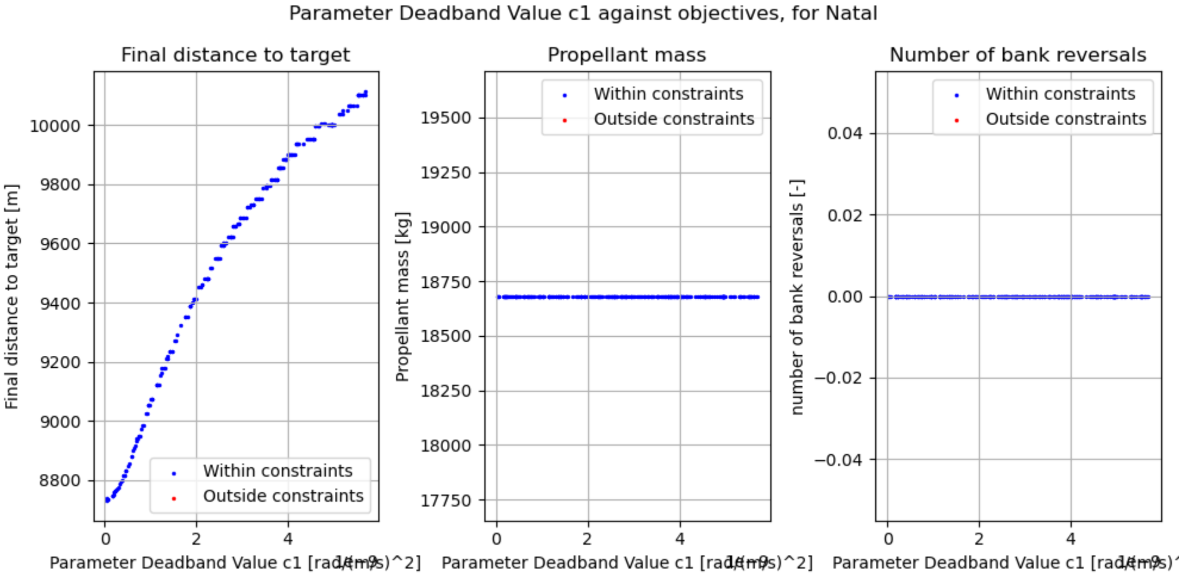

Figure B.11: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c1$ for Natal.

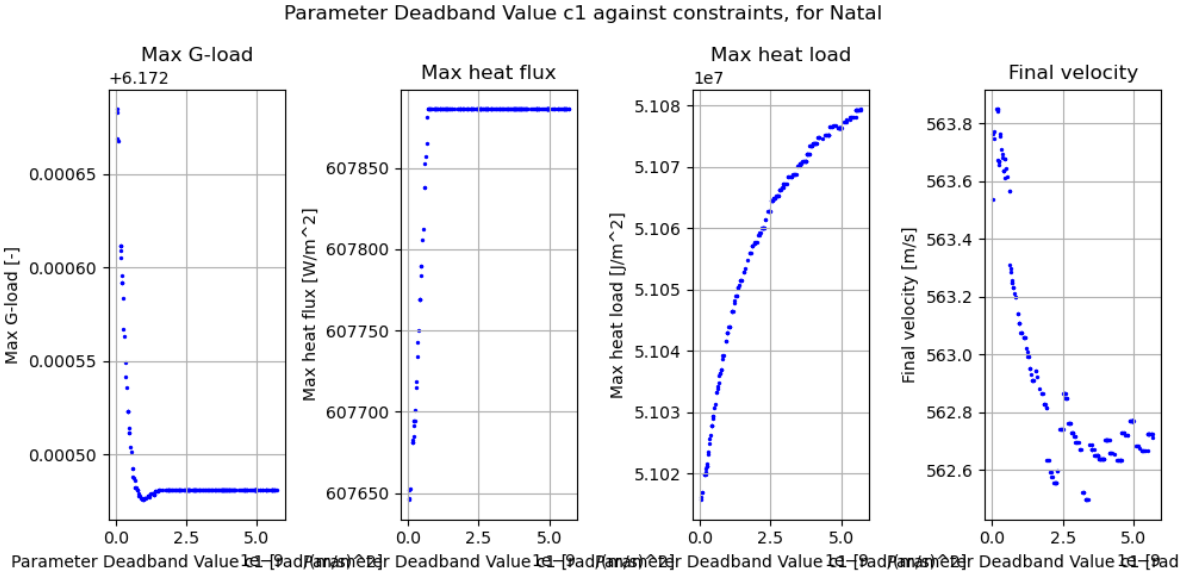

Figure B.12: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c1$ for Natal.

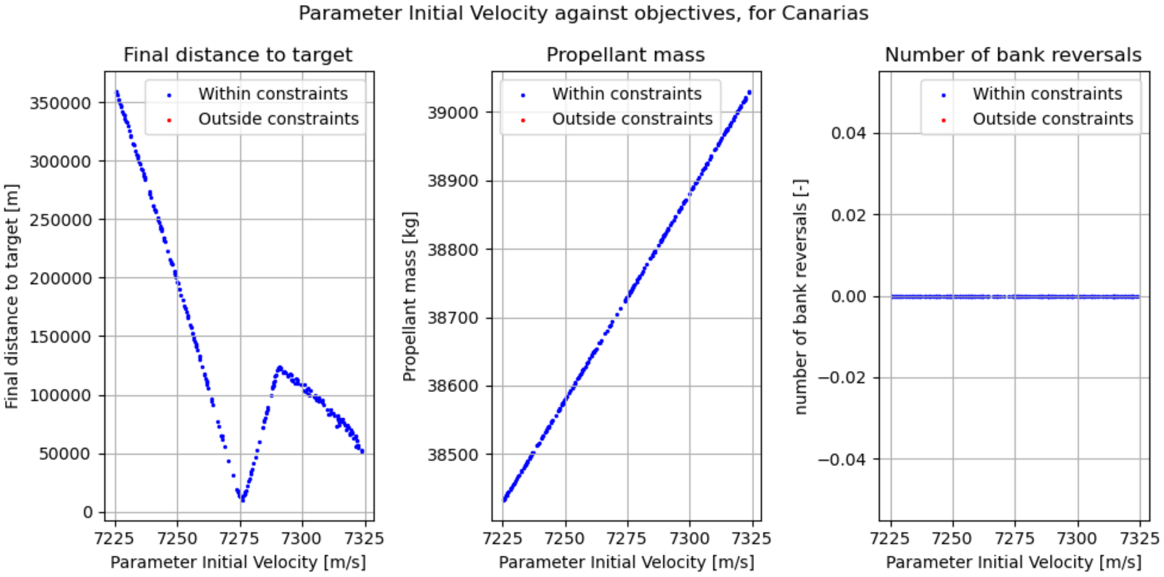

Figure B.13: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial velocity for The Canarias.

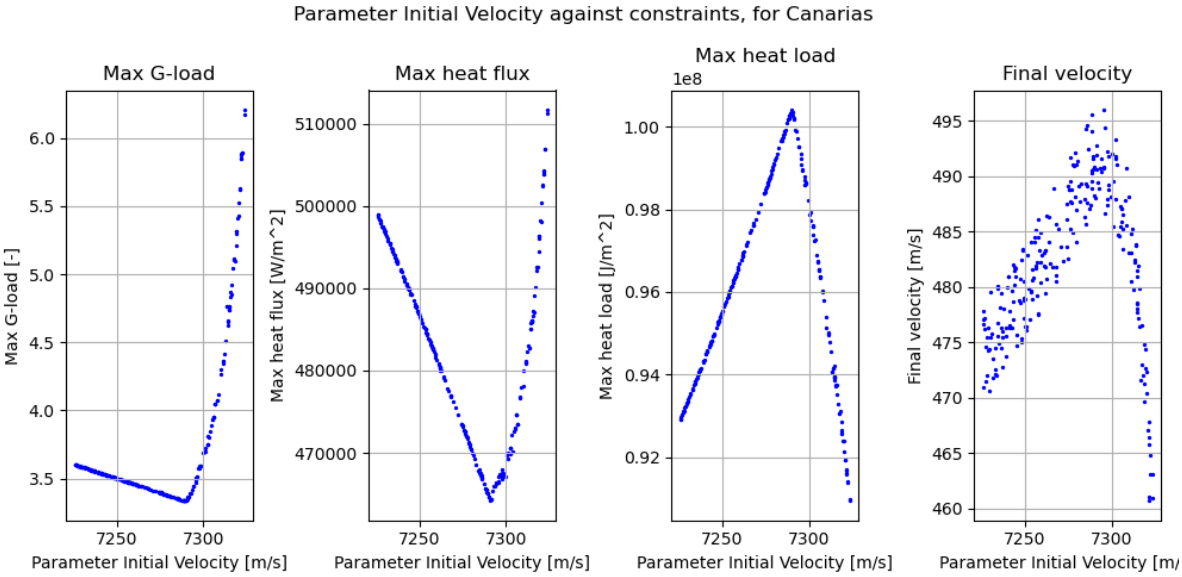

Figure B.14: The maximum g-load, maximum heat flux and total heat load produced by simulations in the single MC variation of the initial velocity for The Canarias.

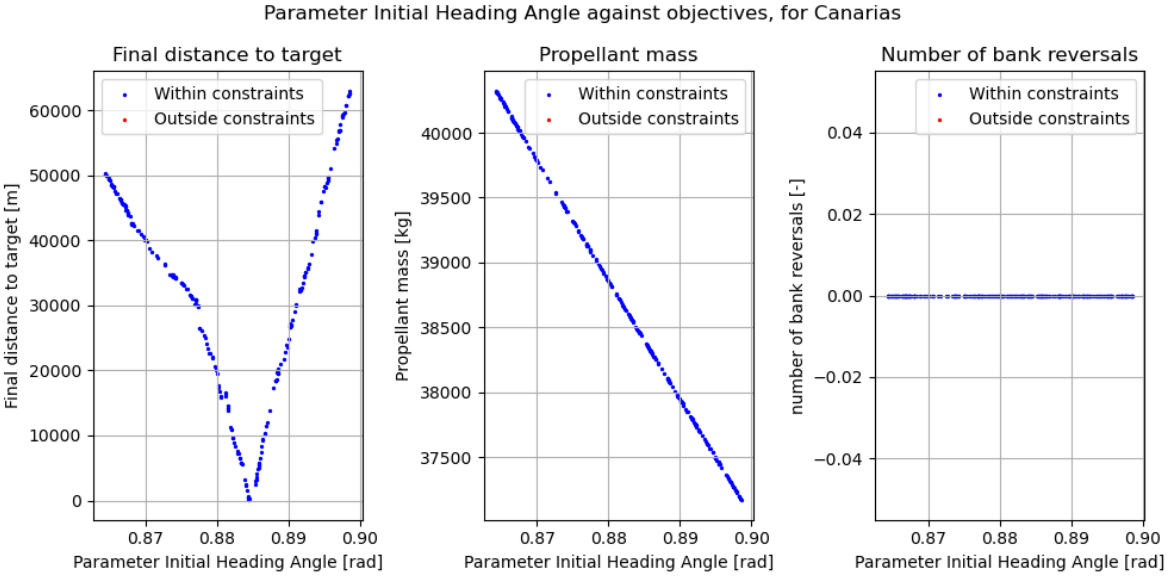

Figure B.15: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial heading angle for The Canarias.

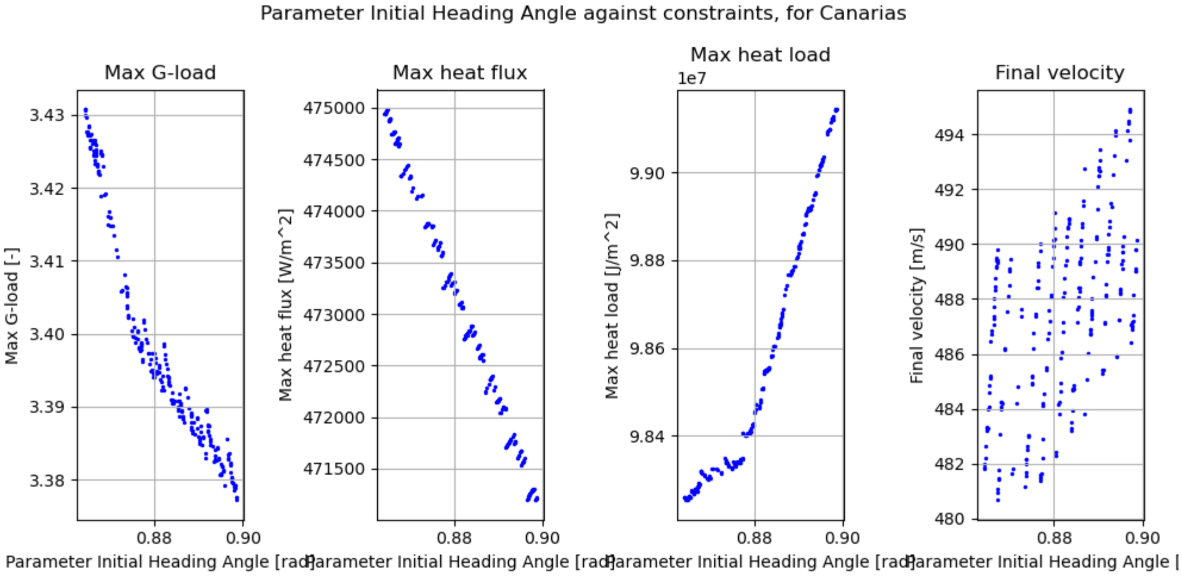

Figure B.16: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial heading angle for The Canarias.

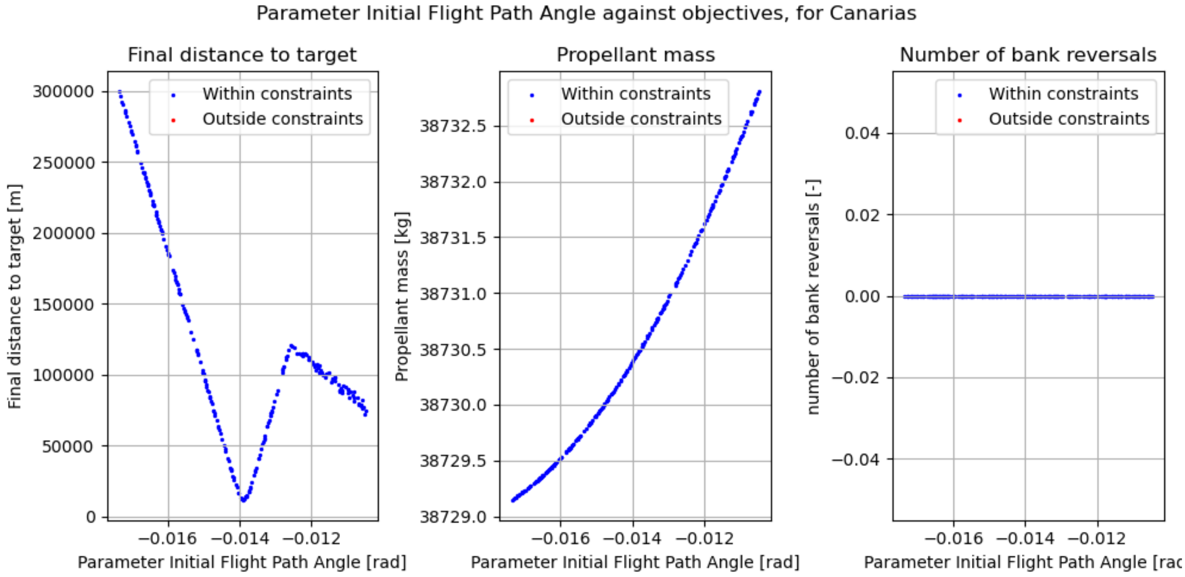

Figure B.17: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial flight path angle for The Canarias.

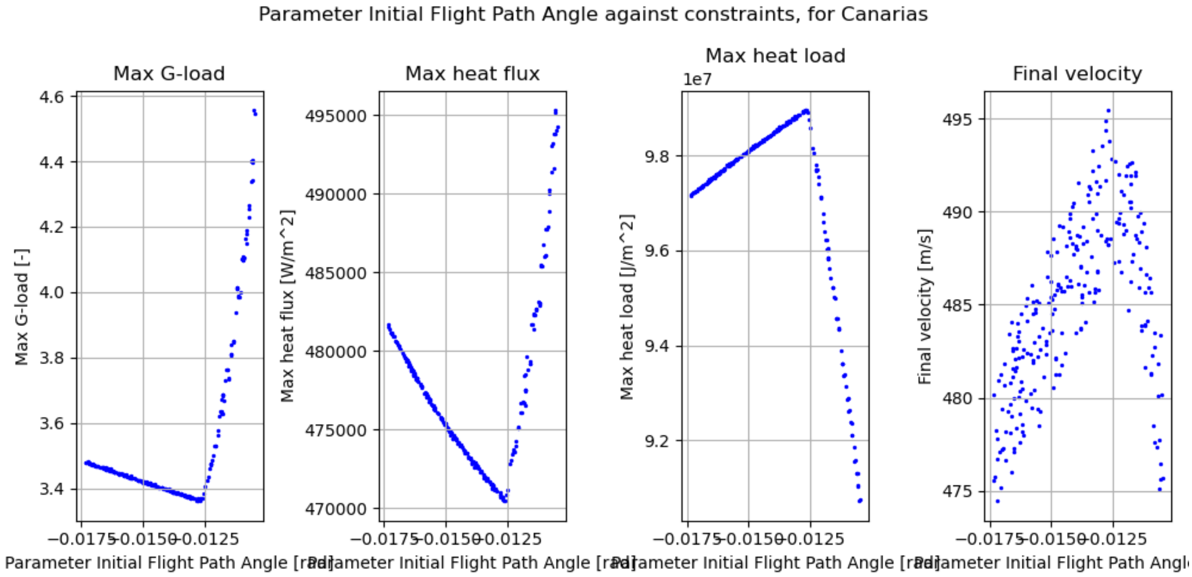

Figure B.18: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial flight path angle for The Canarias.

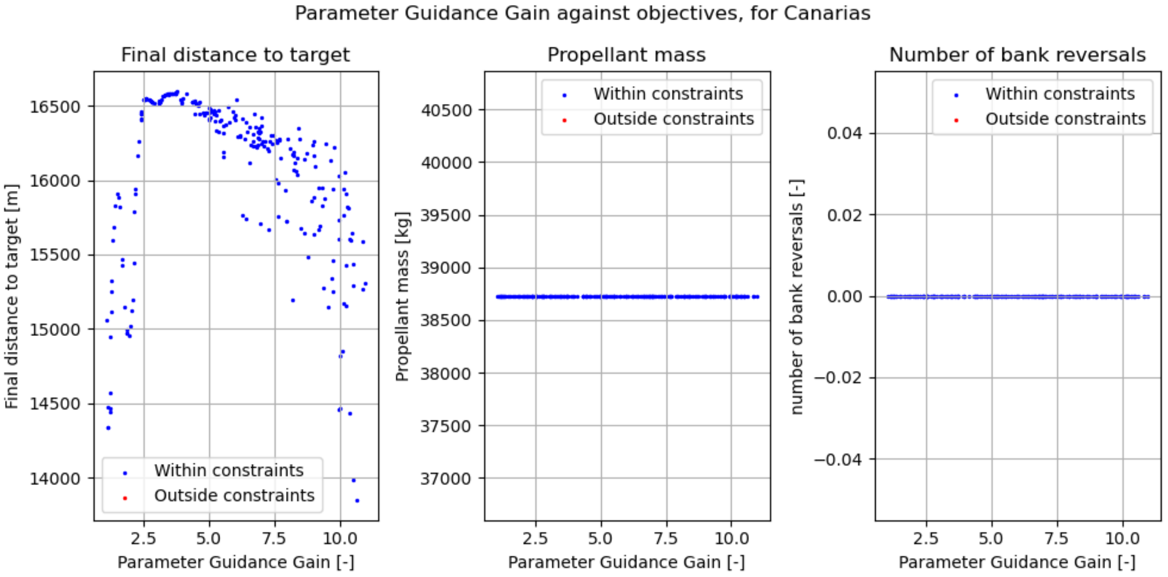

Figure B.19: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the overcontrol guidance gain for The Canarias.

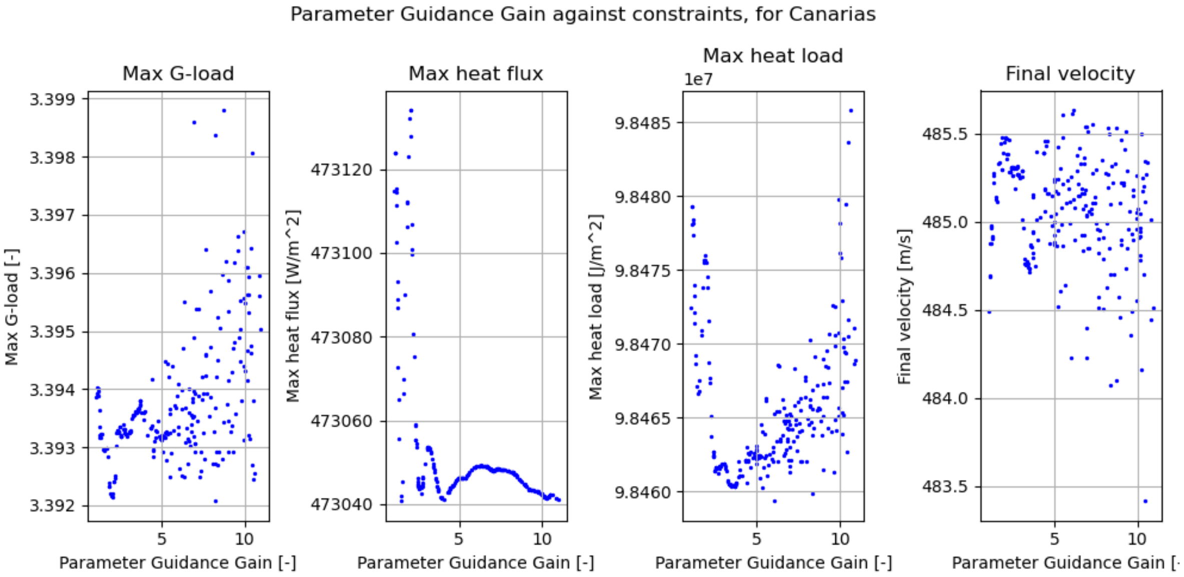

Figure B.20: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the overcontrol guidance gain for The Canarias.

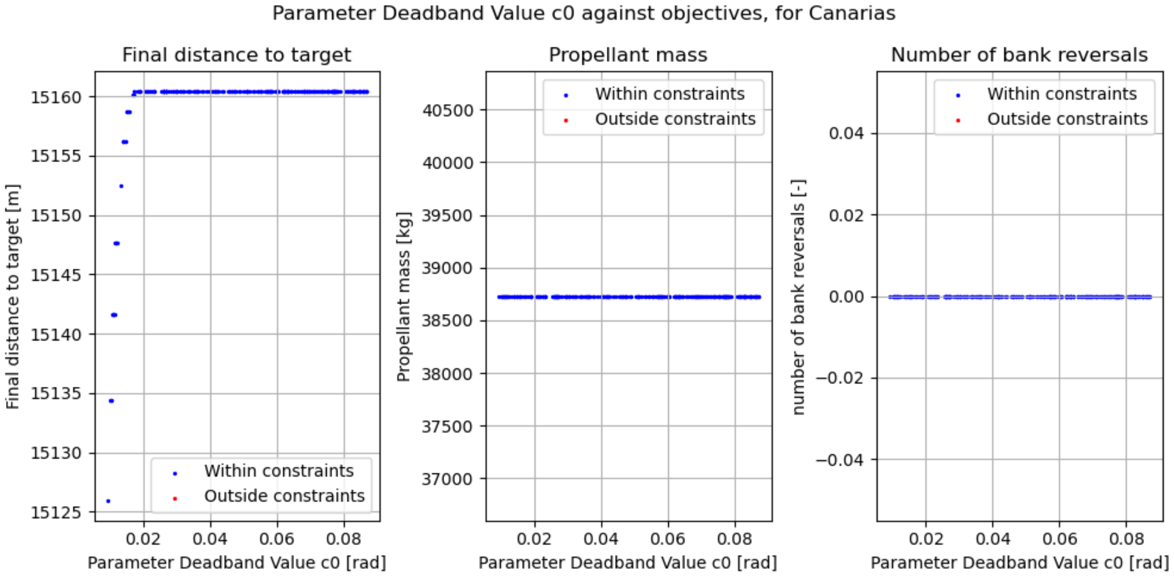

Figure B.21: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c0$ for The Canarias.

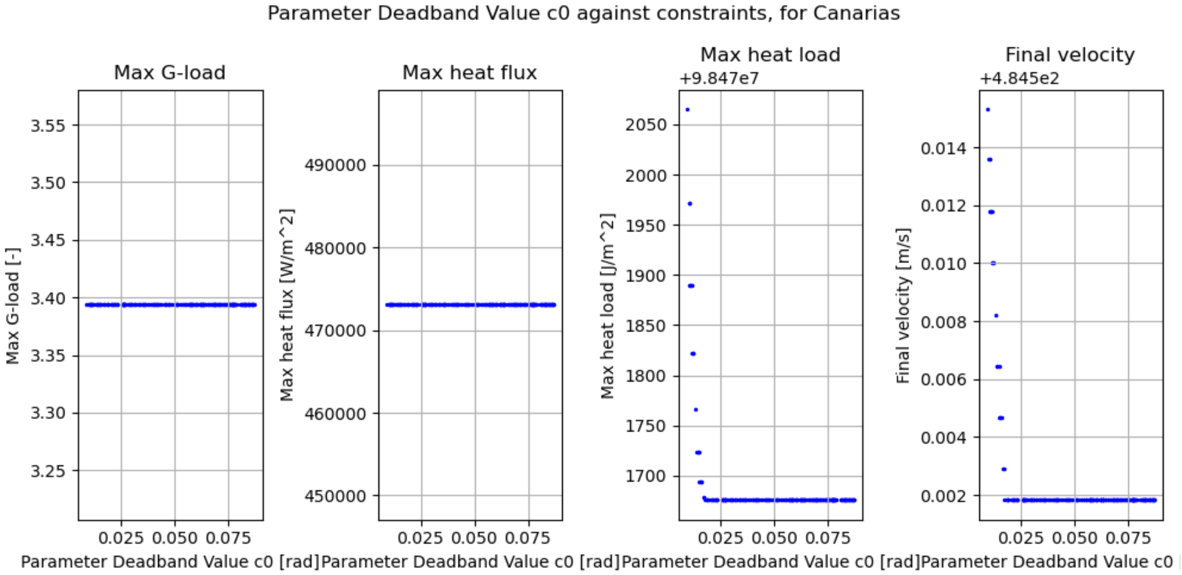

Figure B.22: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c0$ for The Canarias.

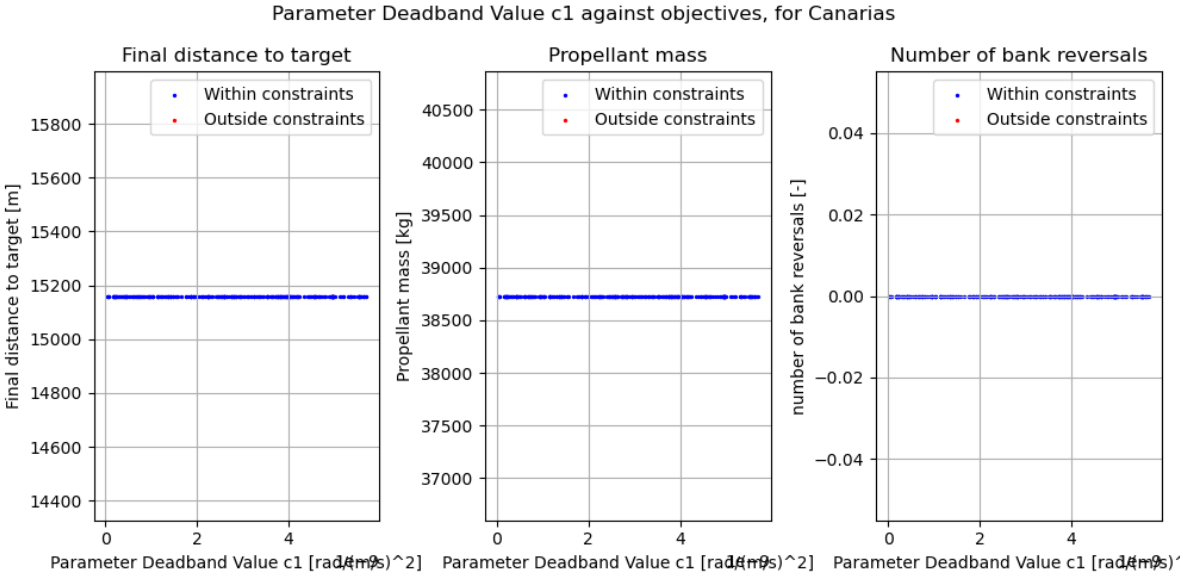

Figure B.23: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c1$ for The Canarias.

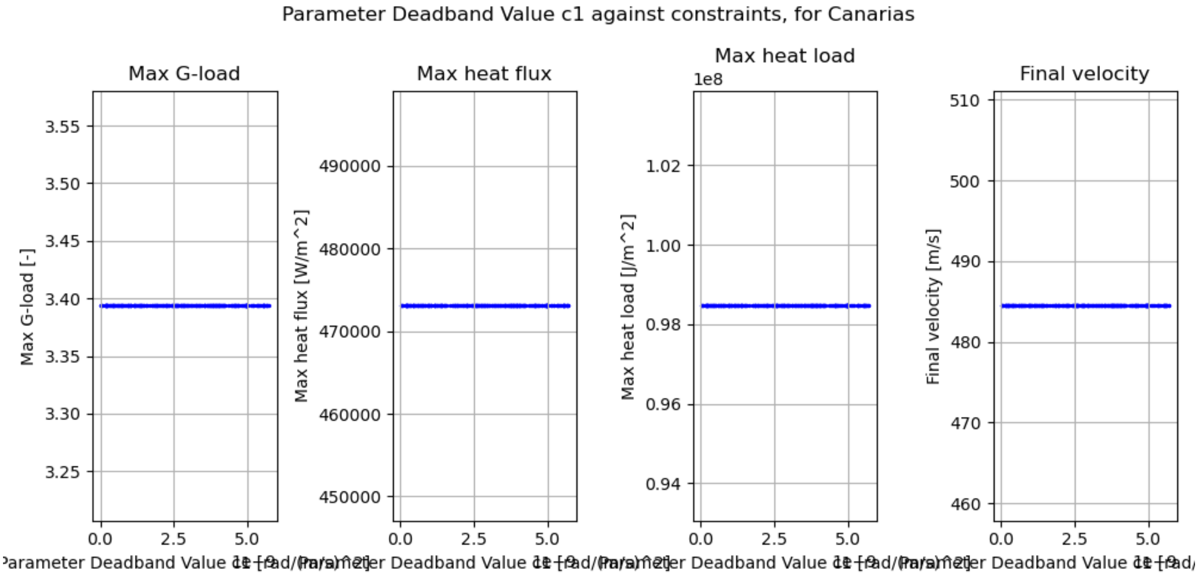

Figure B.24: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c1$ for The Canarias.

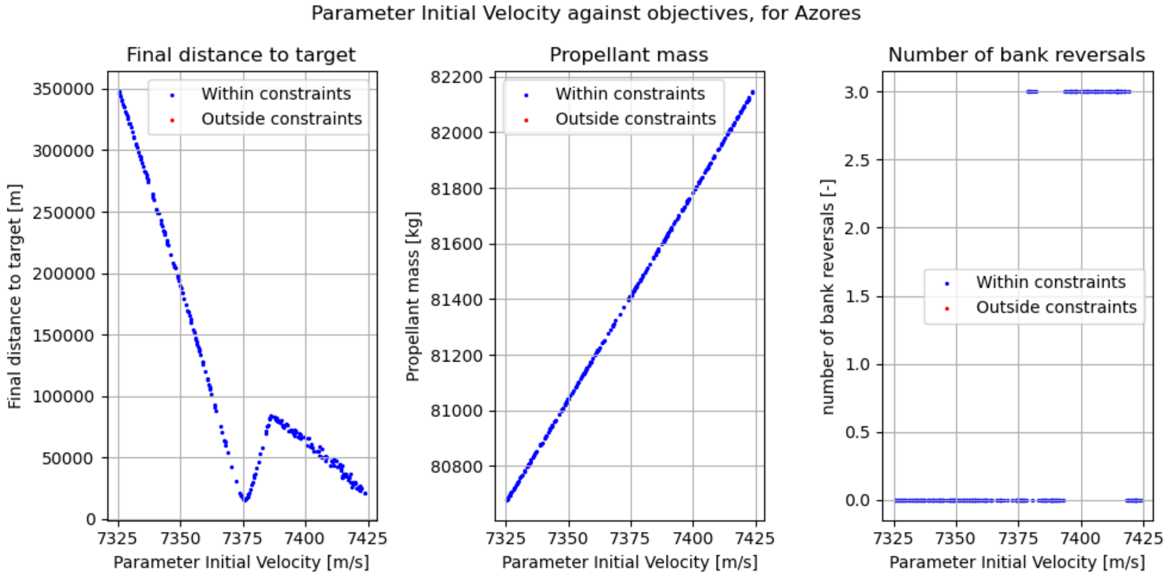

Figure B.25: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial velocity for The Azores.

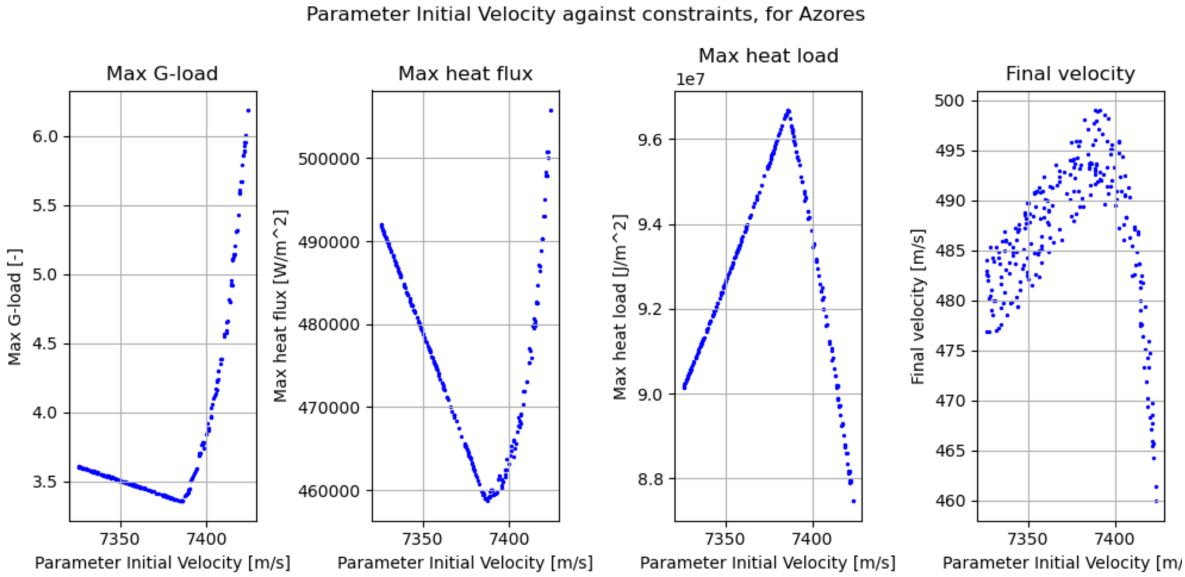

Figure B.26: The maximum g-load, maximum heat flux and total heat load produced by simulations in the single MC variation of the initial velocity for The Azores.

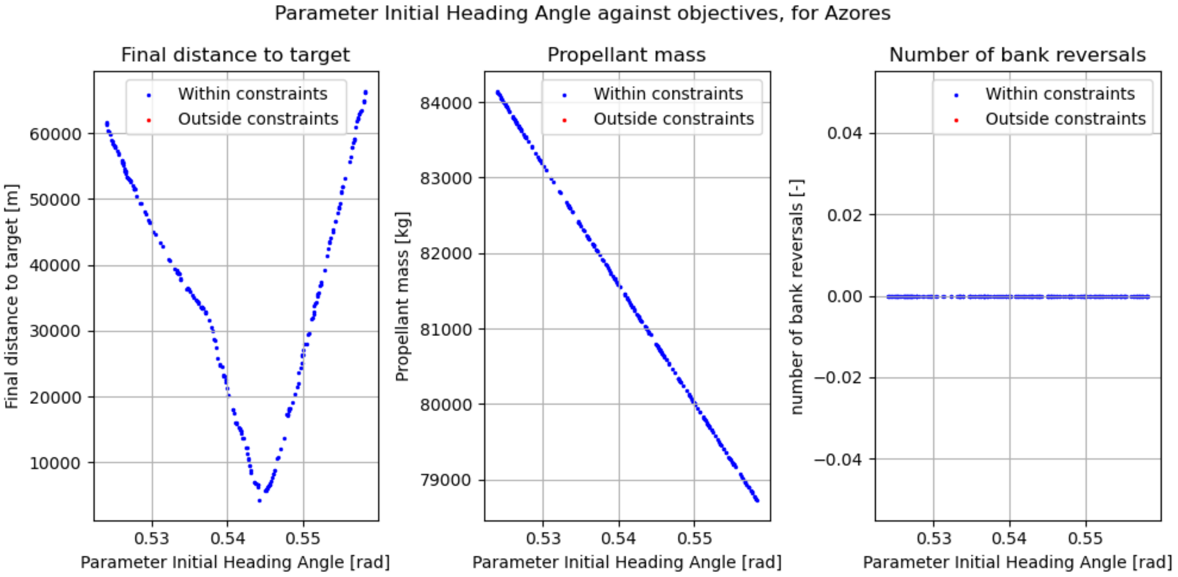

Figure B.27: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial heading angle for The Azores.

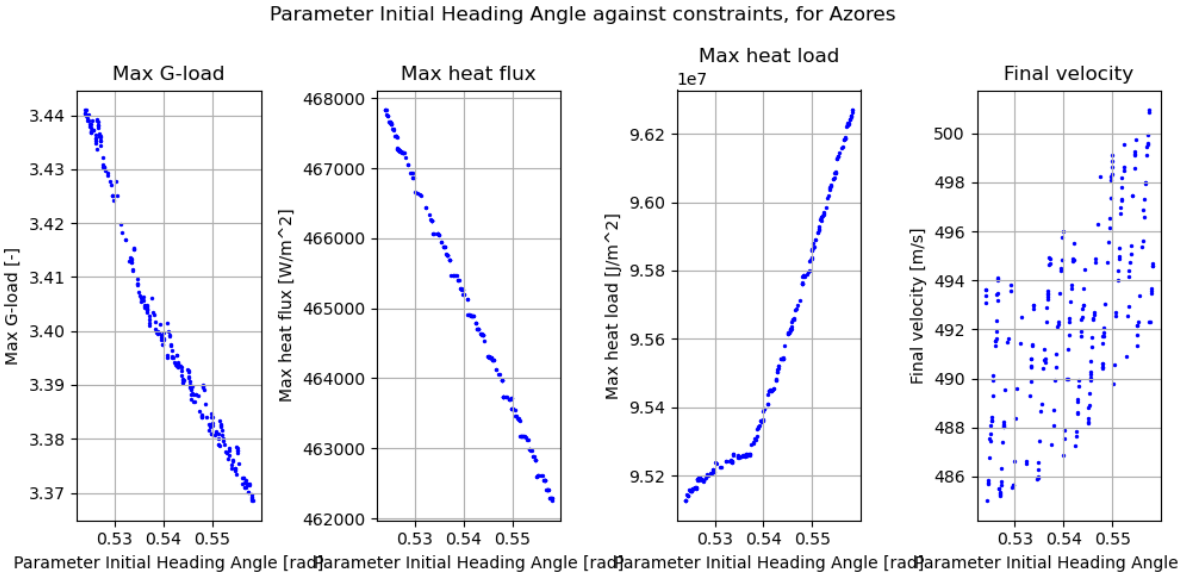

Figure B.28: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial heading angle for The Azores.

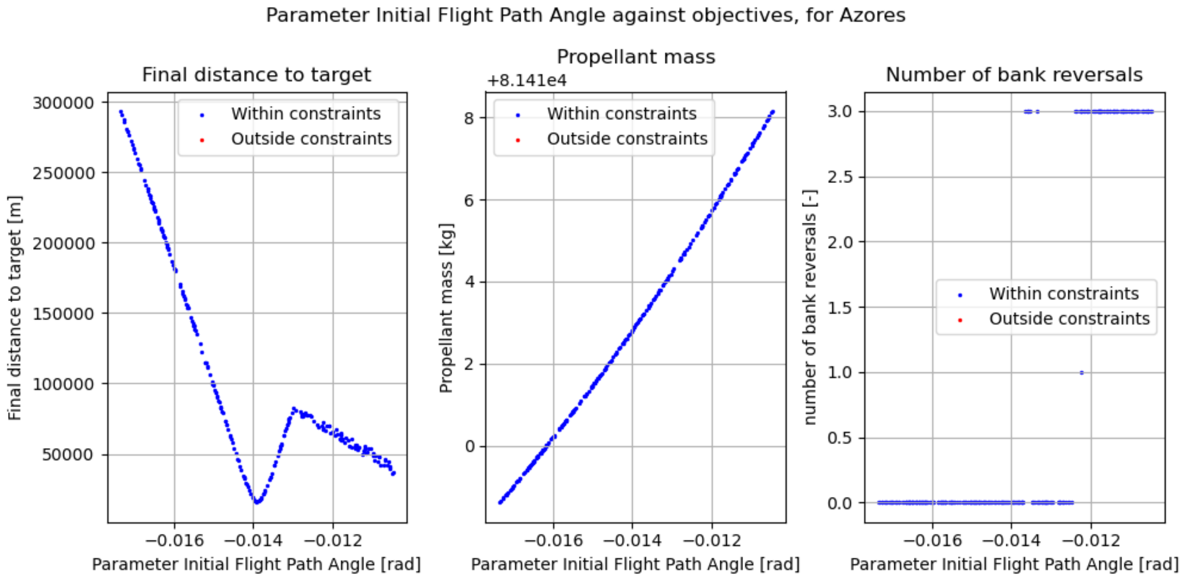

Figure B.29: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial flight path angle for The Azores.

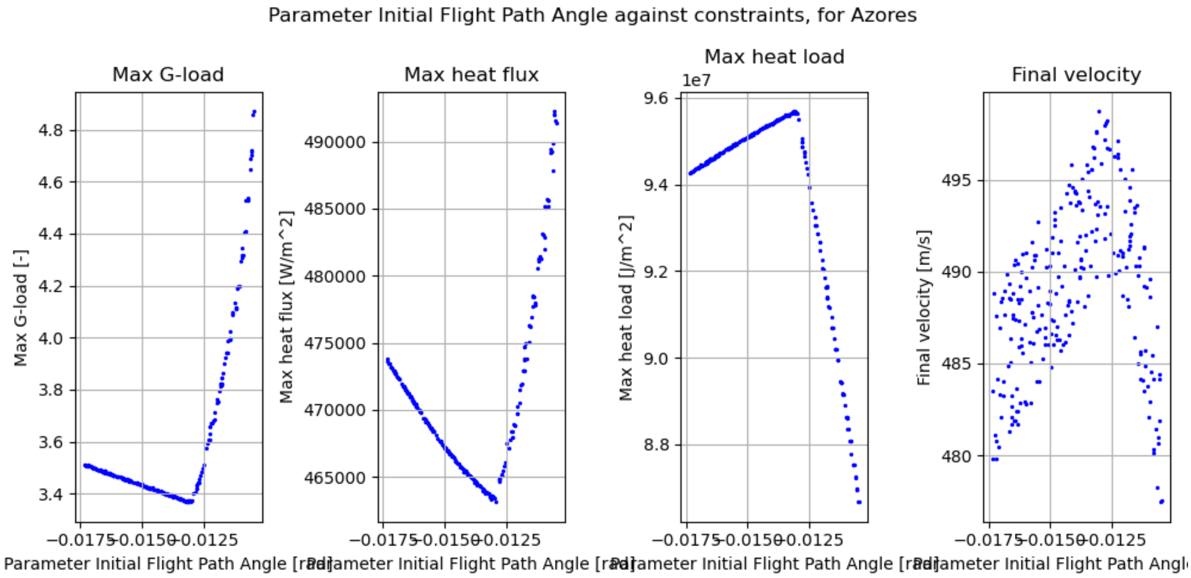

Figure B.30: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial flight path angle for The Azores.

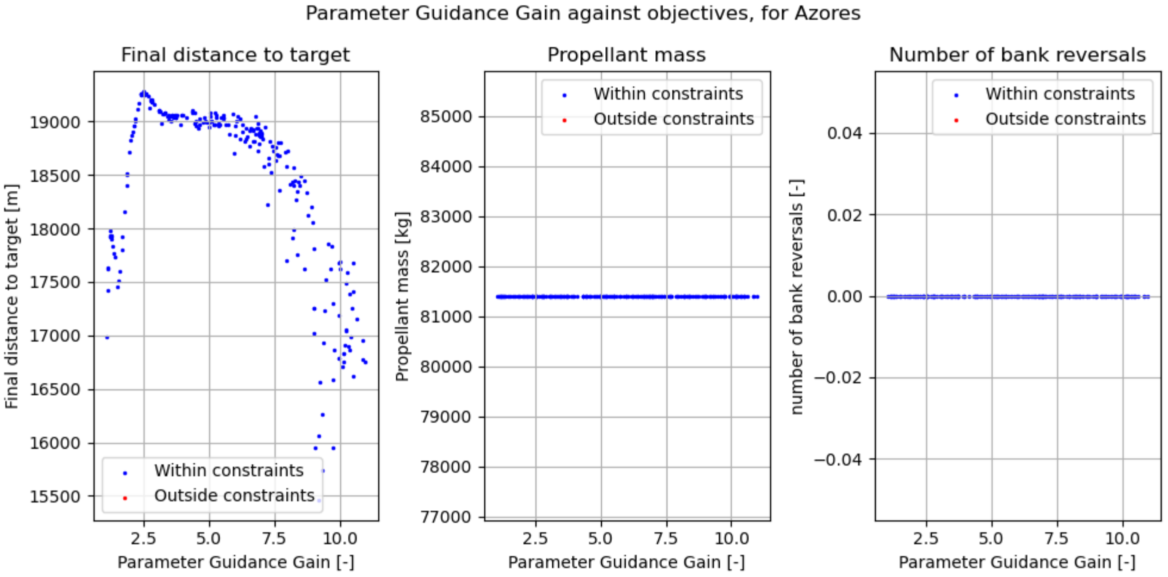

Figure B.31: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the overcontrol guidance gain for The Azores.

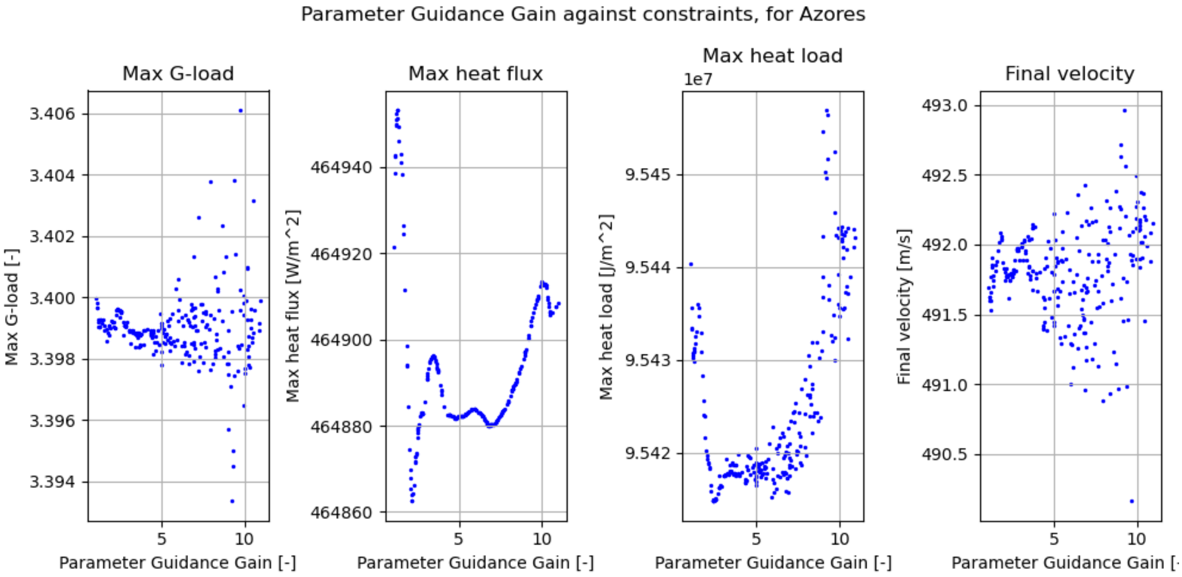

Figure B.32: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the overcontrol guidance gain for The Azores.

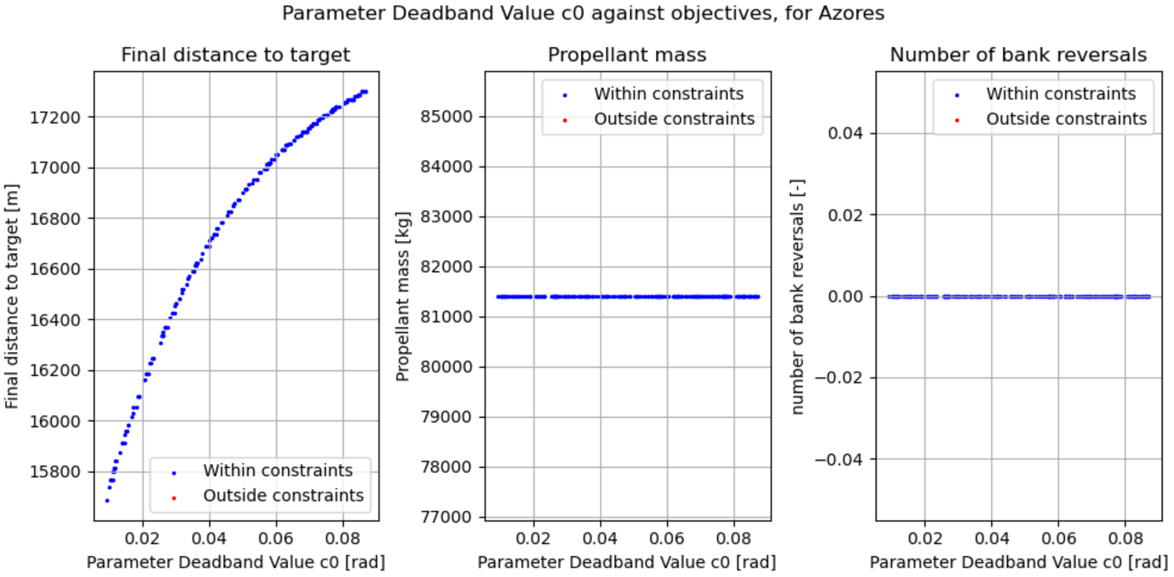

Figure B.33: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c0$ for The Azores.

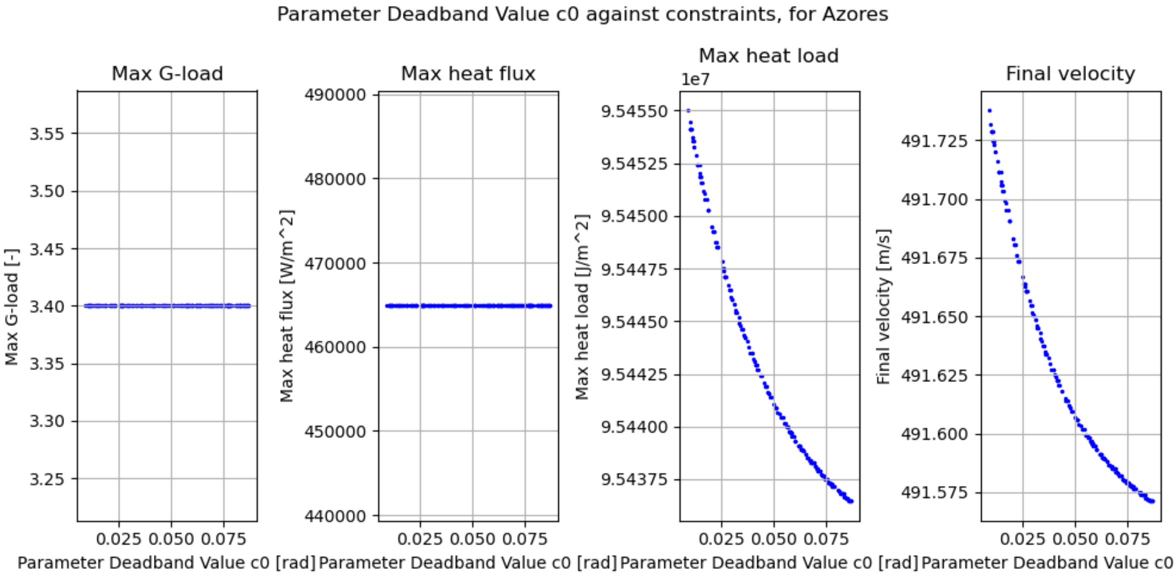

Figure B.34: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c0$ for The Azores.

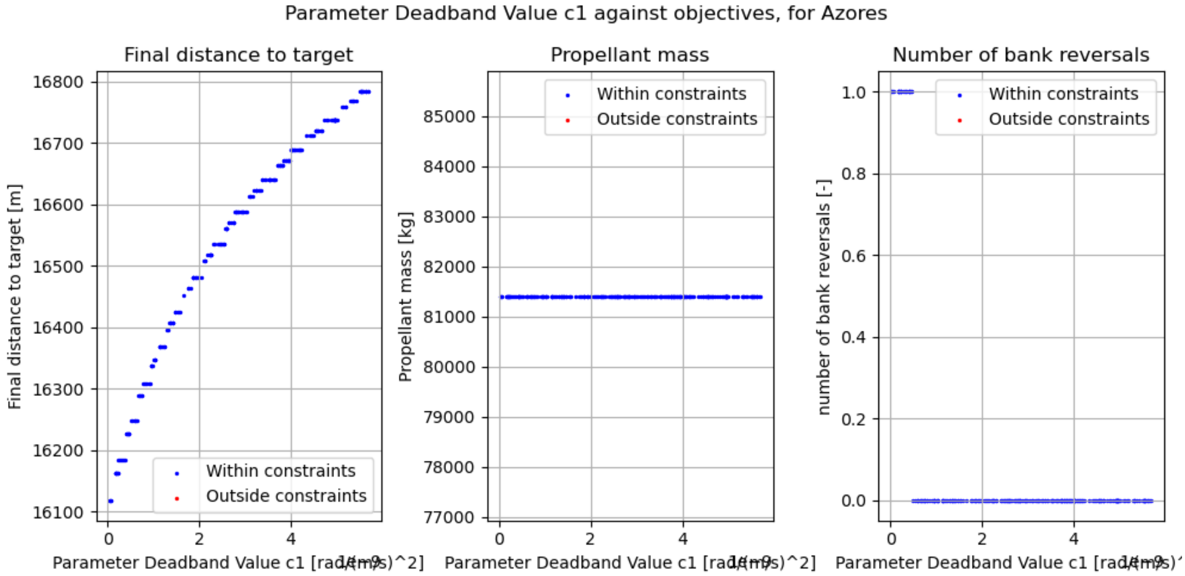

Figure B.35: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c1$ for The Azores.

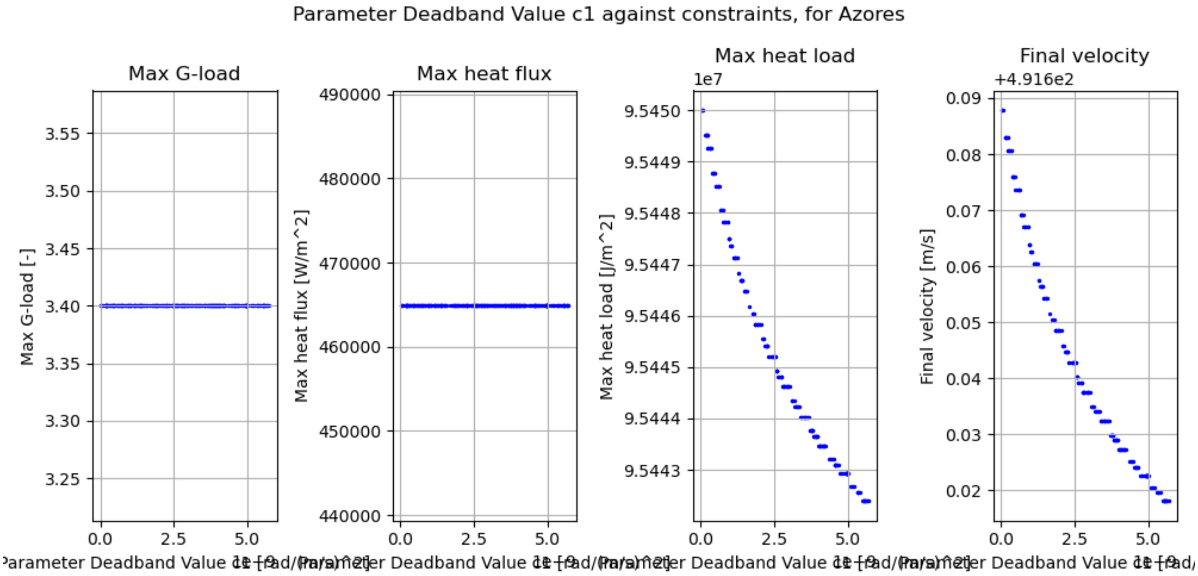

Figure B.36: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c1$ for The Azores.

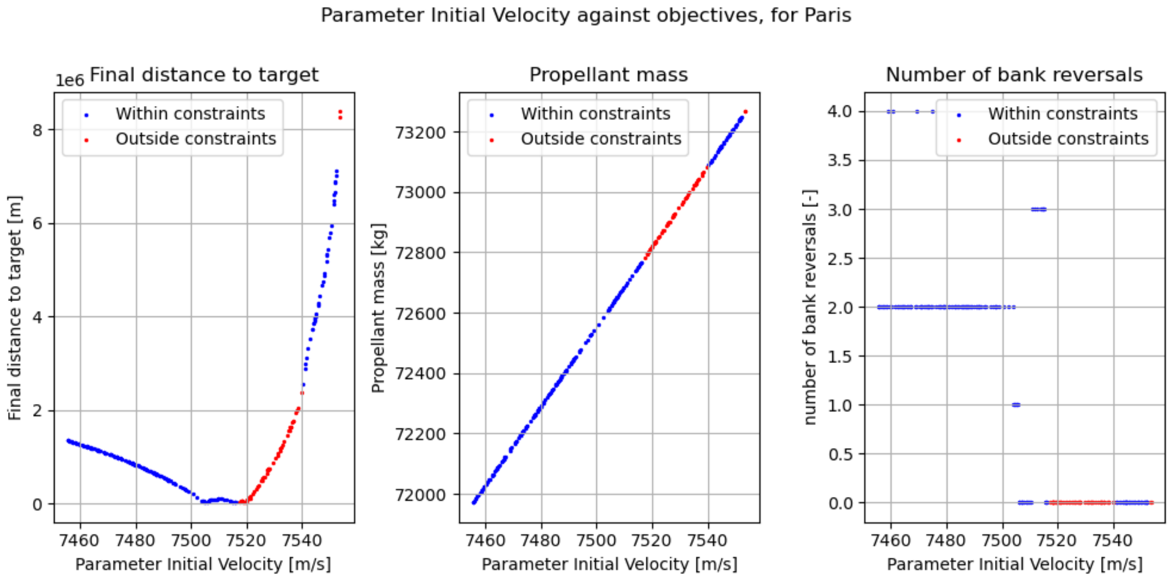

Figure B.37: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial velocity for Les Mureaux.

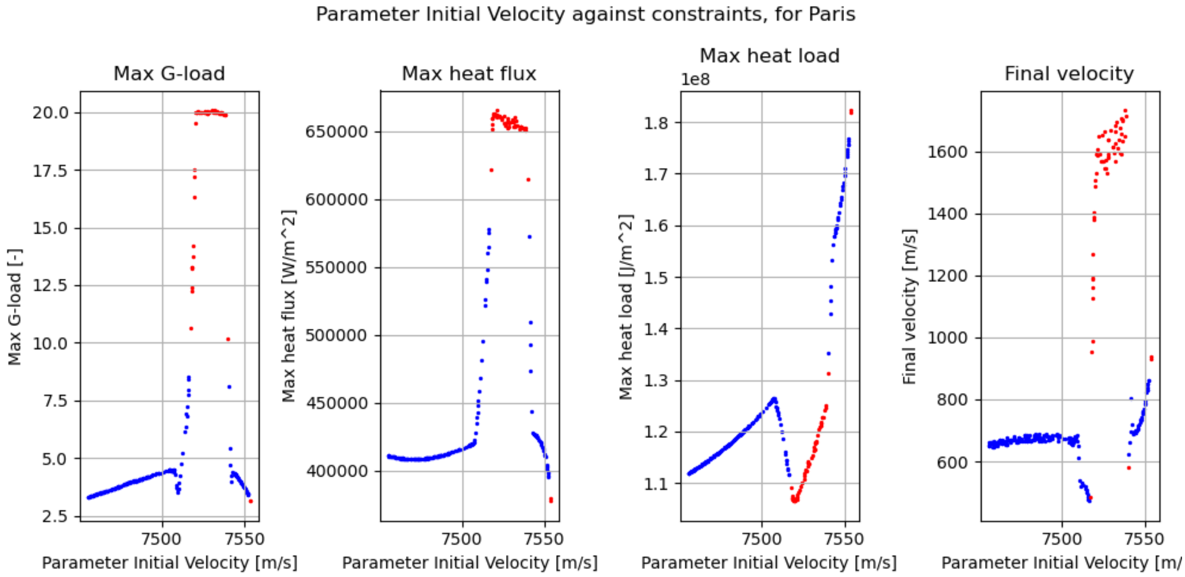

Figure B.38: The maximum g-load, maximum heat flux and total heat load produced by simulations in the single MC variation of the initial velocity for Les Mureaux.

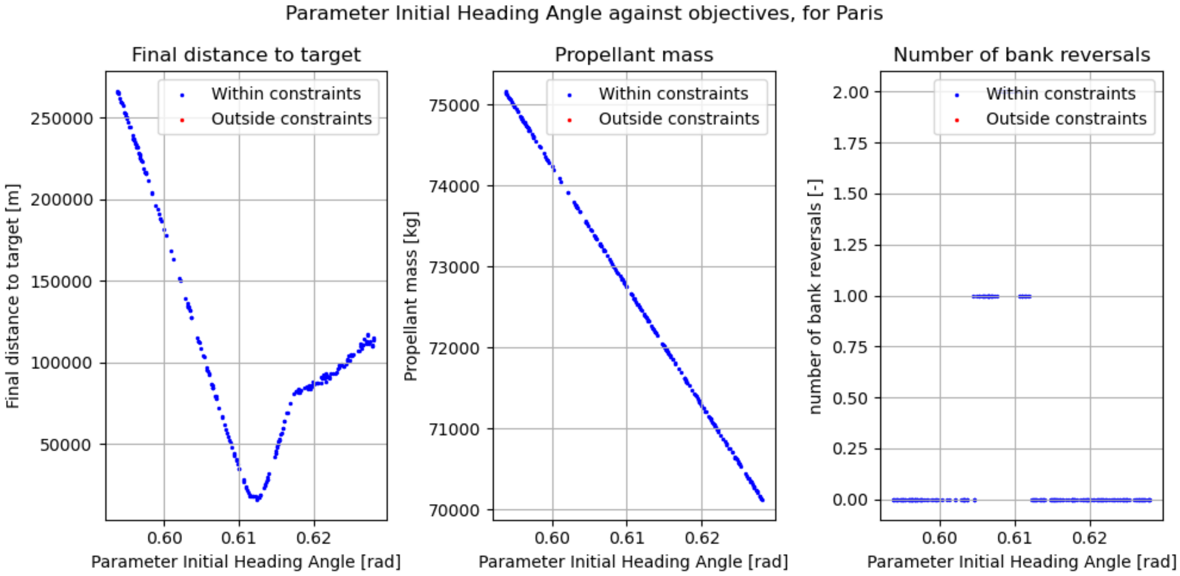

Figure B.39: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial heading angle for Les Mureaux.

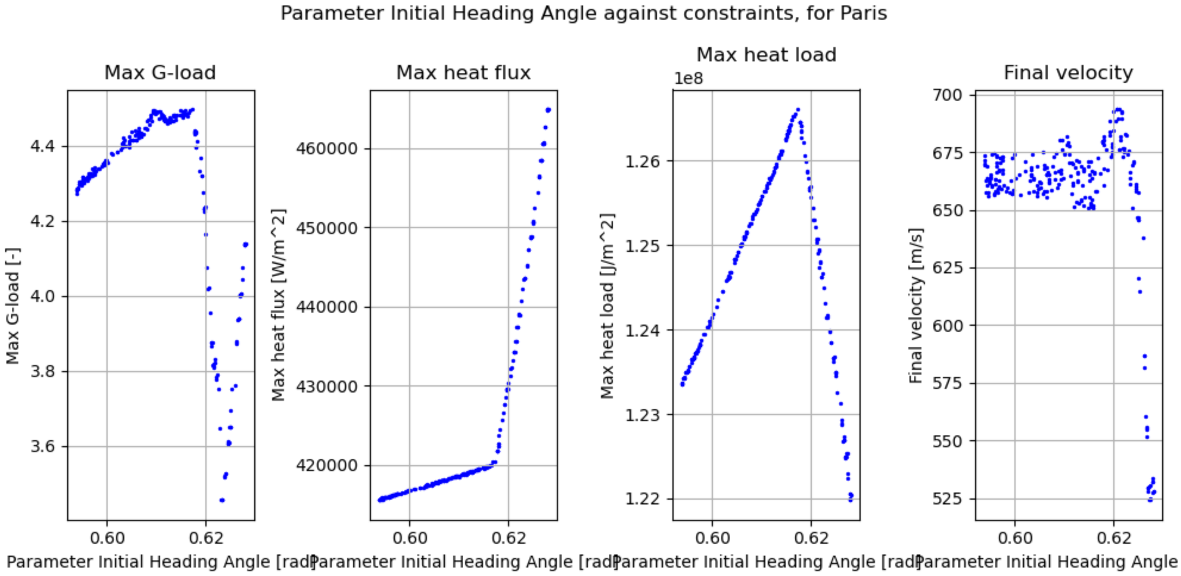

Figure B.40: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial heading angle for Les Mureaux.

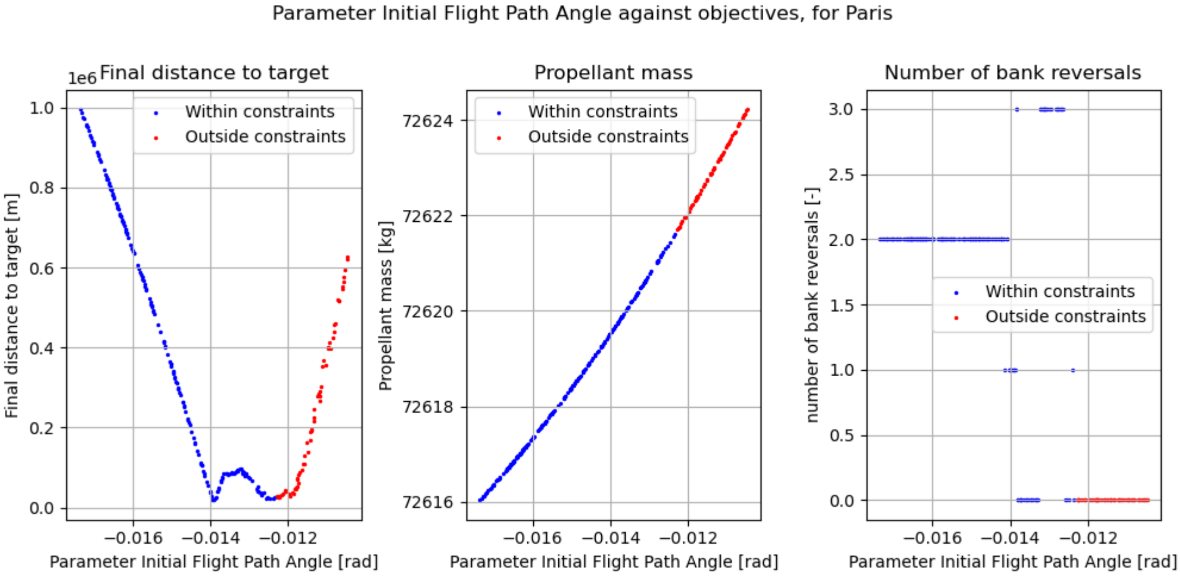

Figure B.41: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the initial flight path angle for Les Mureaux.

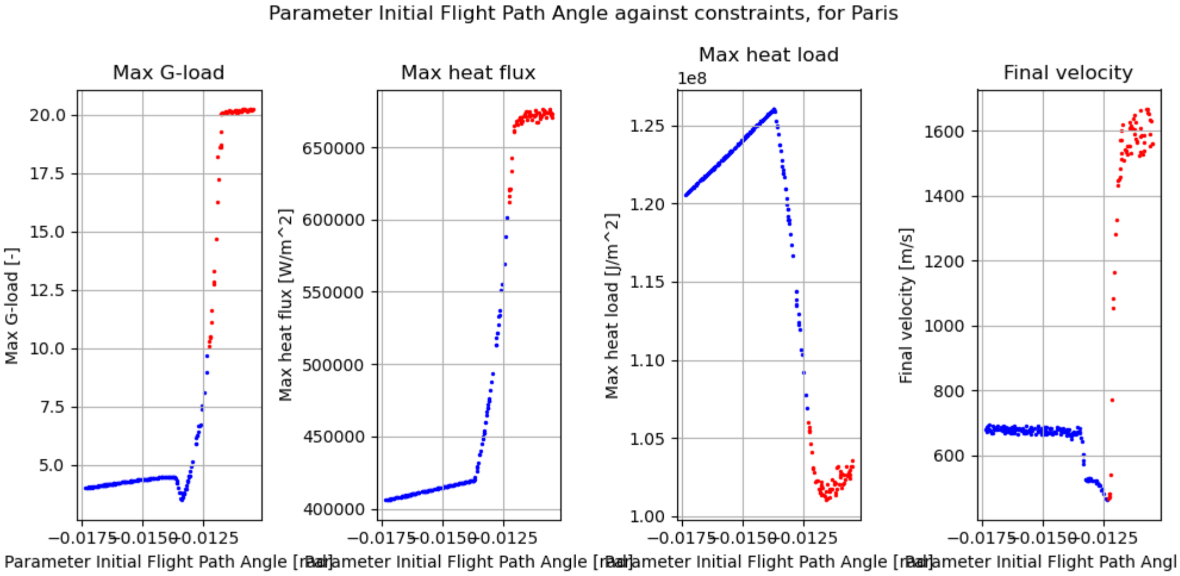

Figure B.42: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the initial flight path angle for Les Mureaux.

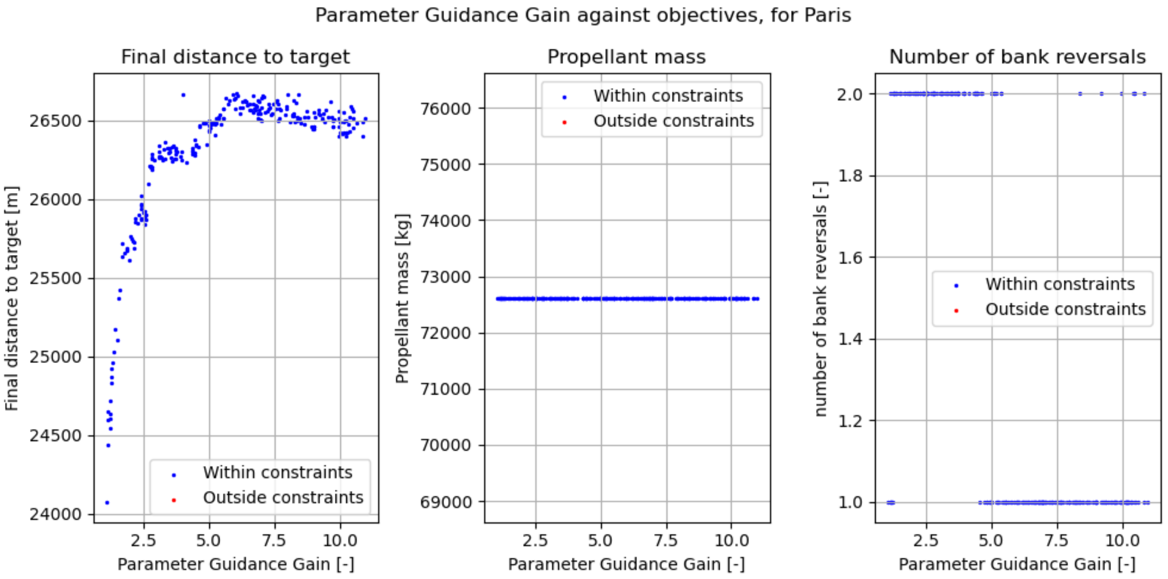

Figure B.43: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the overcontrol guidance gain for Les Mureaux.

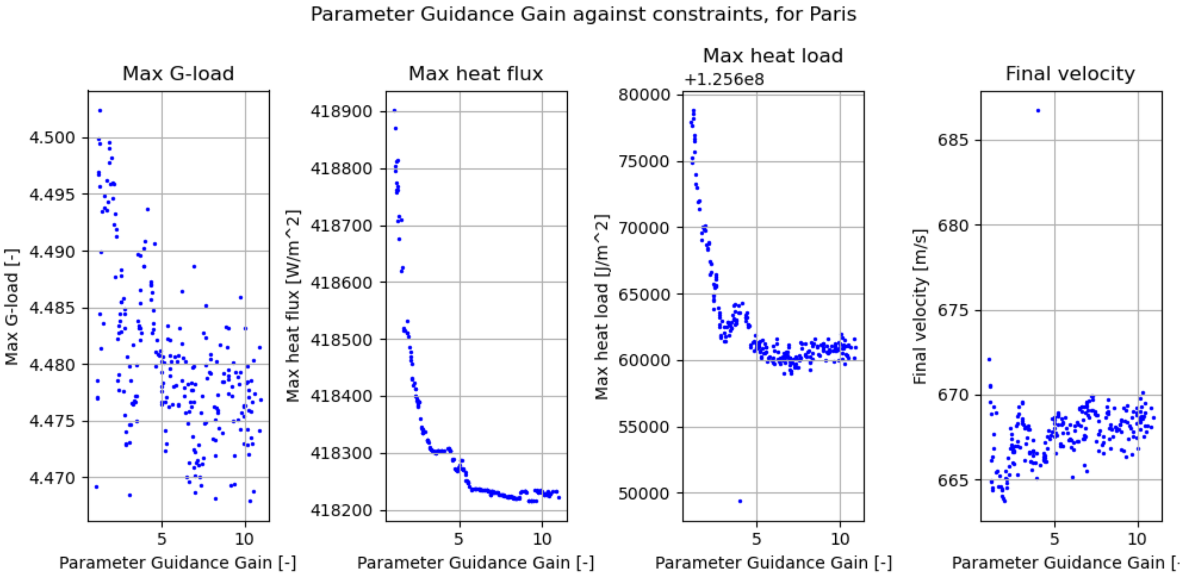

Figure B.44: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the overcontrol guidance gain for Les Mureaux.

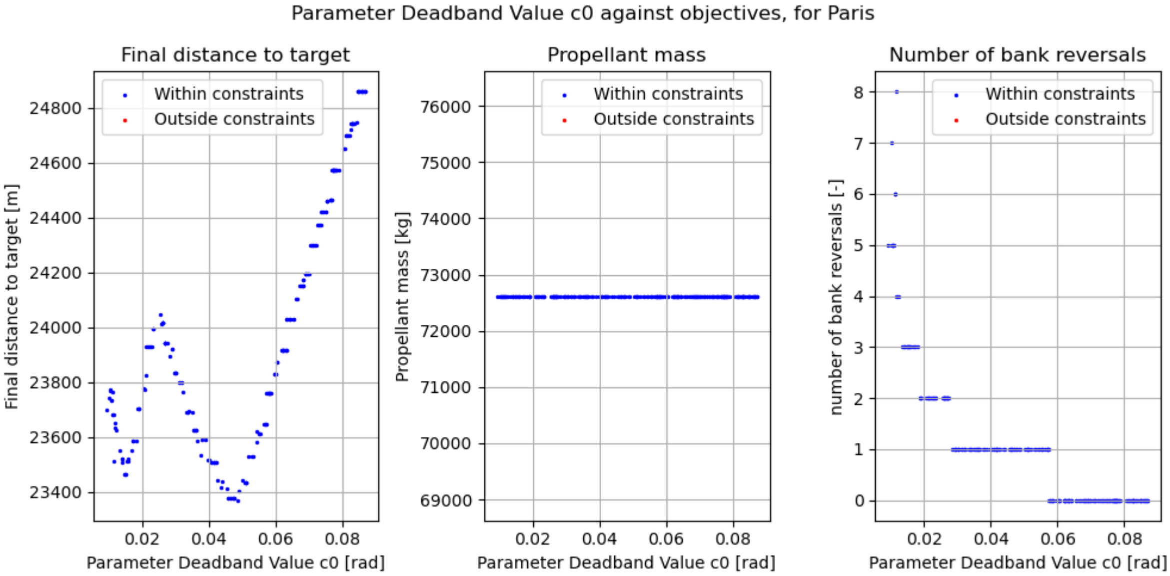

Figure B.45: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c0$ for Les Mureaux.

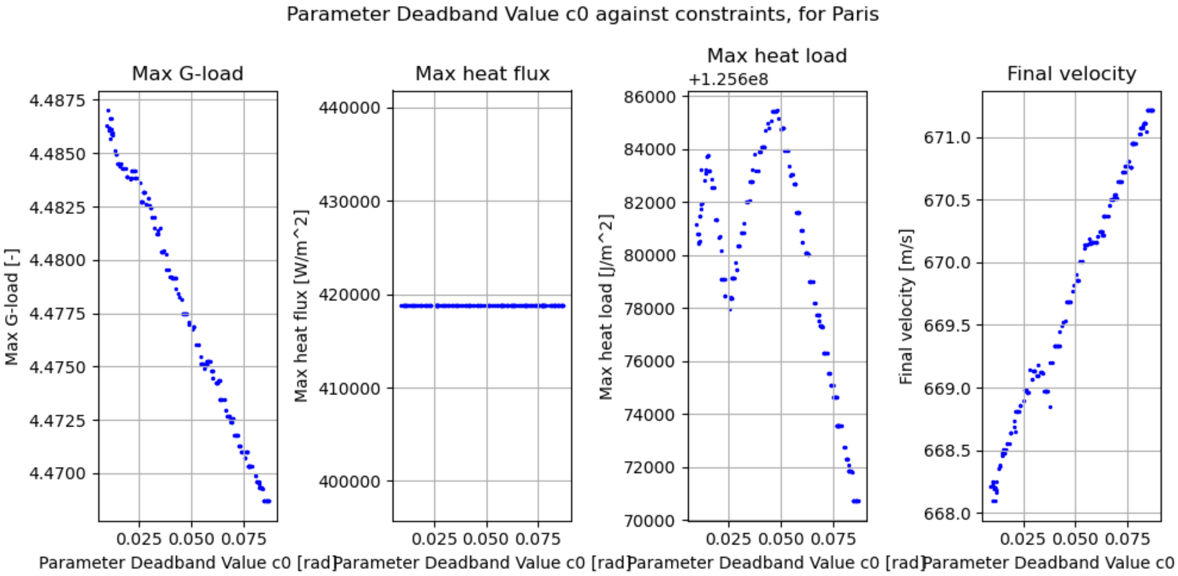

Figure B.46: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c0$ for Les Mureaux.

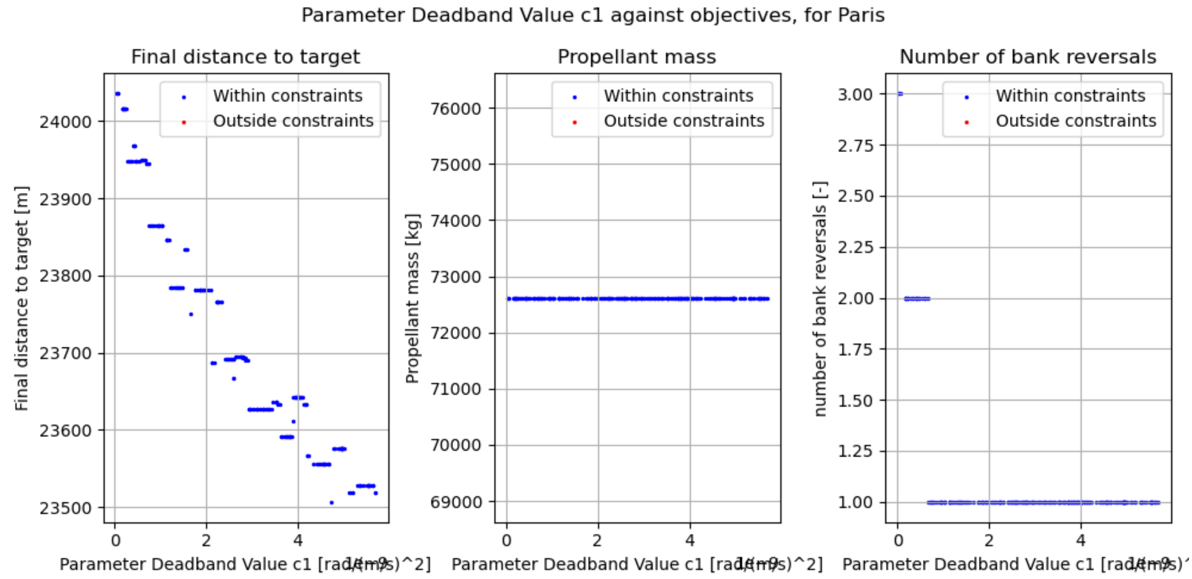

Figure B.47: The final distance to target, propellant mass and number of bank angle reversals produced by simulations in the single MC variation of the deadband value $c1$ for Les Mureaux.

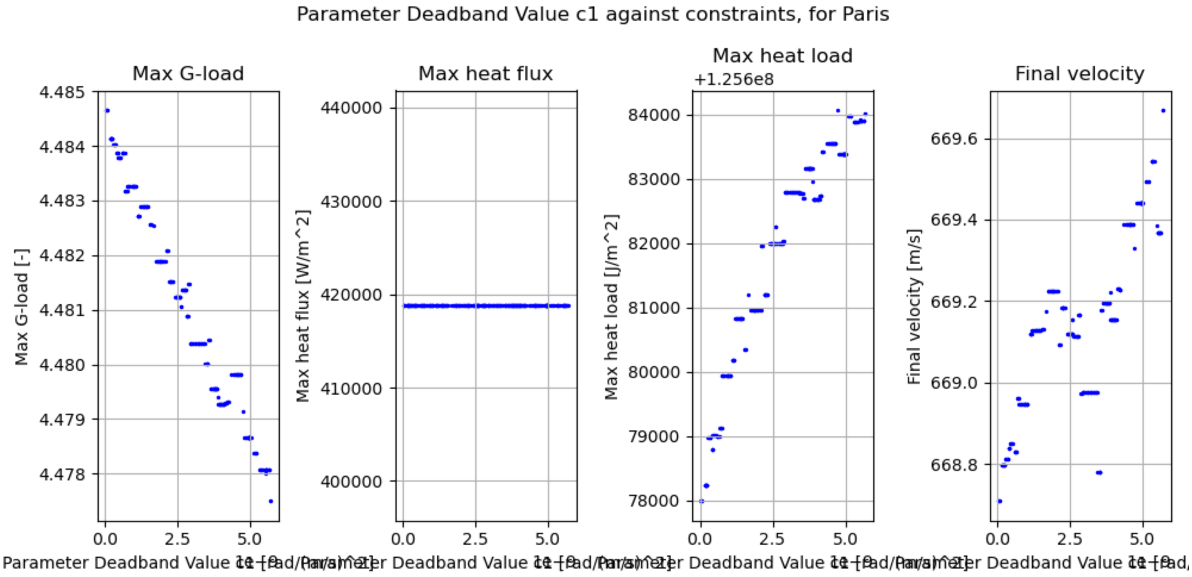

Figure B.48: The maximum g-load, maximum heat flux, total heat load and final velocity produced by simulations in the single MC variation of the deadband value $c1$ for Les Mureaux.

B.2. All-at-once Monte Carlo

This section contains the plots produced during the all-at-once MC using seed 42 for all target locations except for Cabo Verde, as these can already be found in subsection 5.9.2.

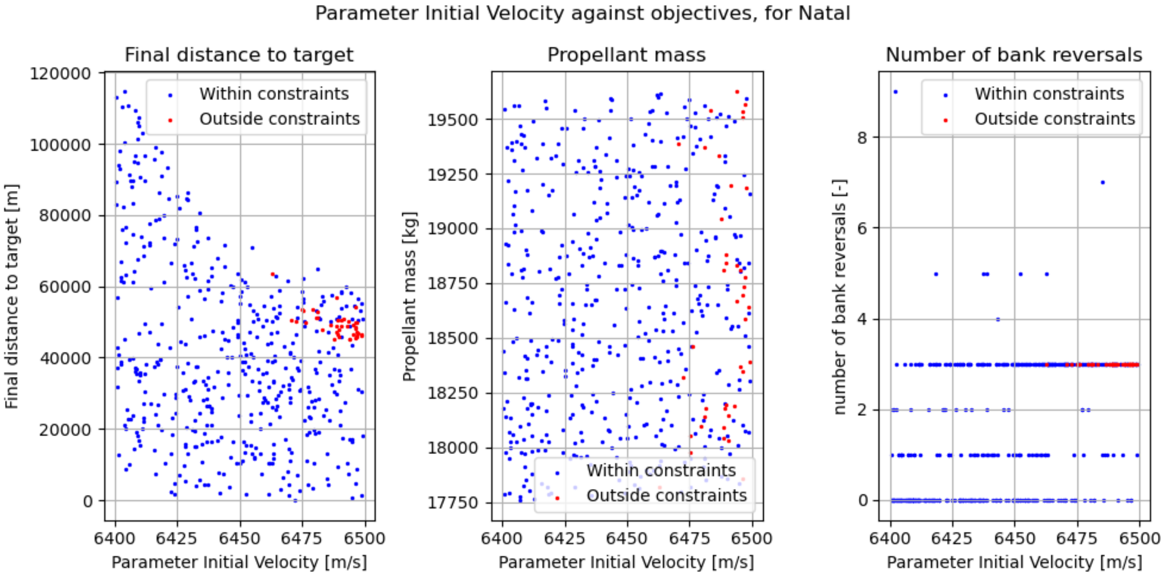

Figure B.49: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial velocity produced by simulations in the all-at-once MC variation for Natal.

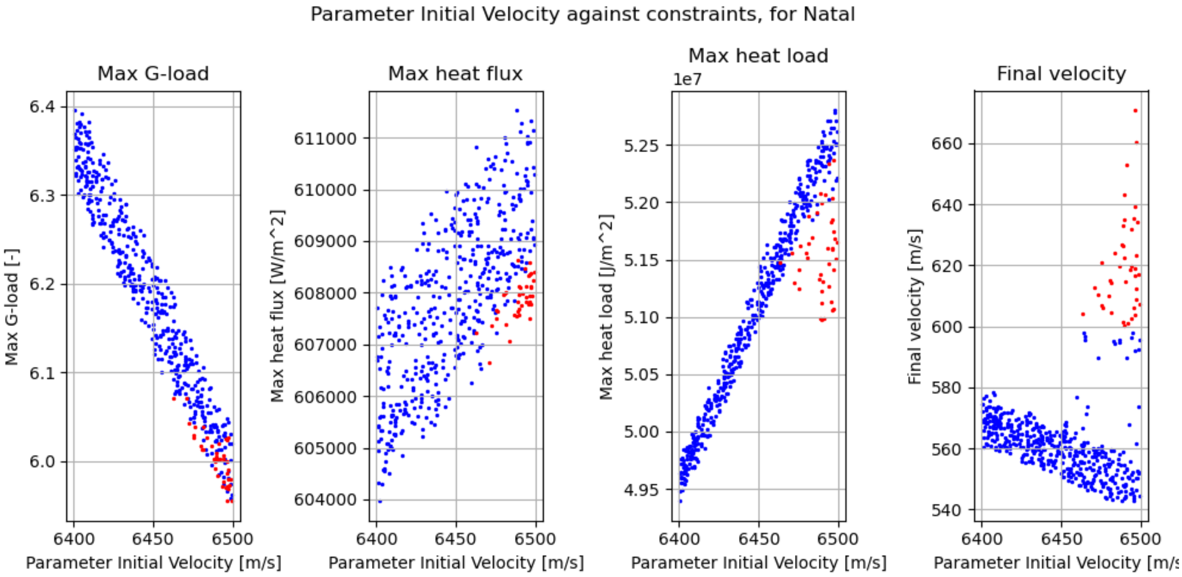

Figure B.50: The maximum g-load, maximum heat flux and total heat load over the value of the initial velocity produced by simulations in the all-at-once MC variation for Natal.

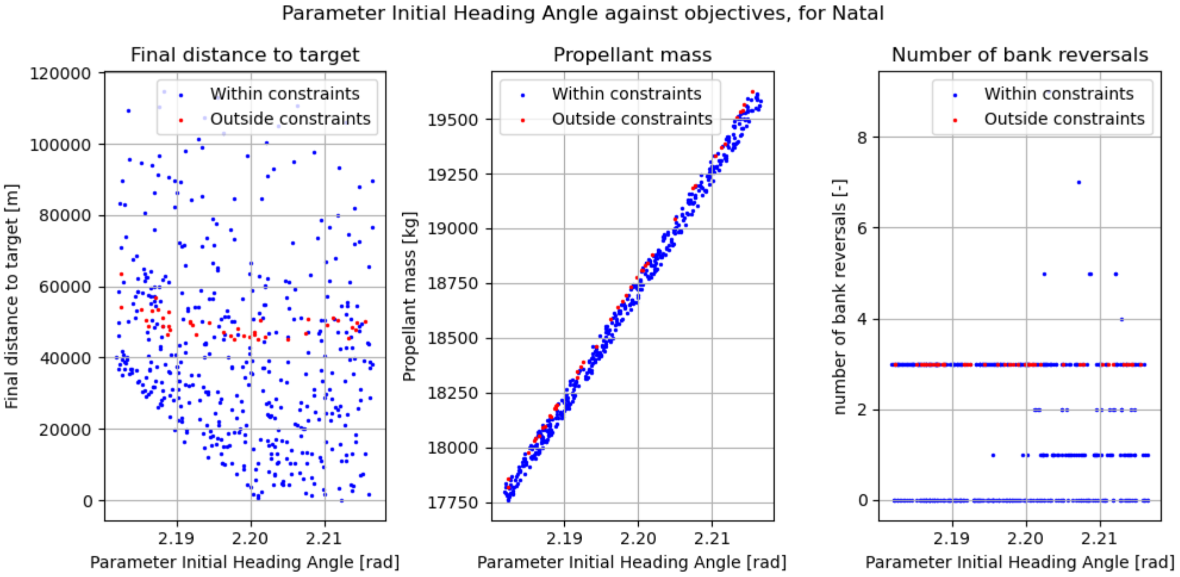

Figure B.51: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial heading angle produced by simulations in the all-at-once MC variation for Natal.

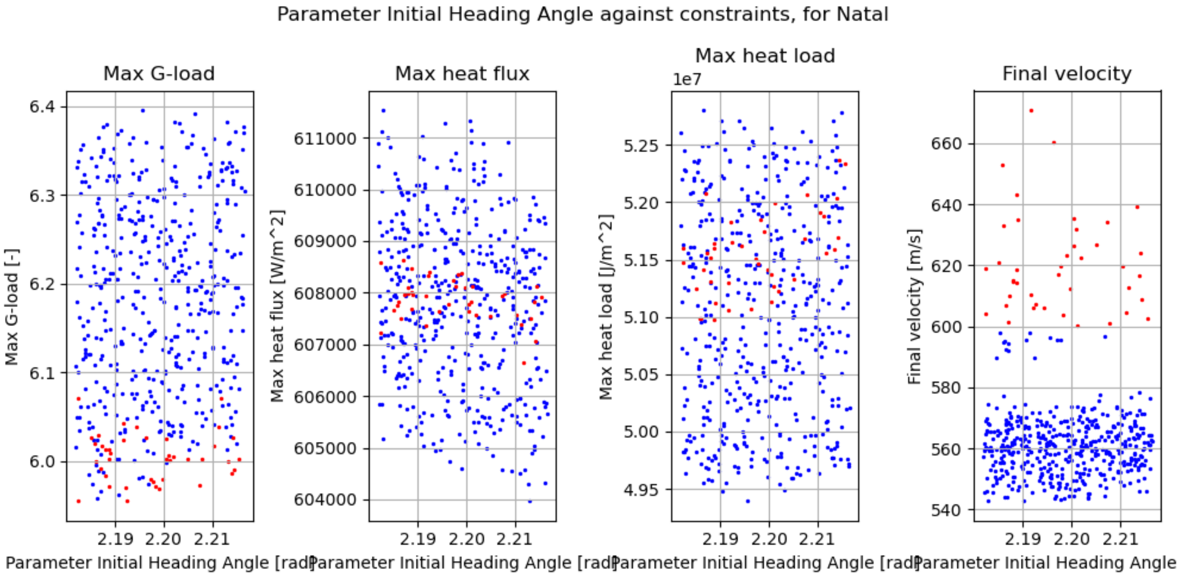

Figure B.52: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial heading angle produced by simulations in the all-at-once MC variation for Natal.

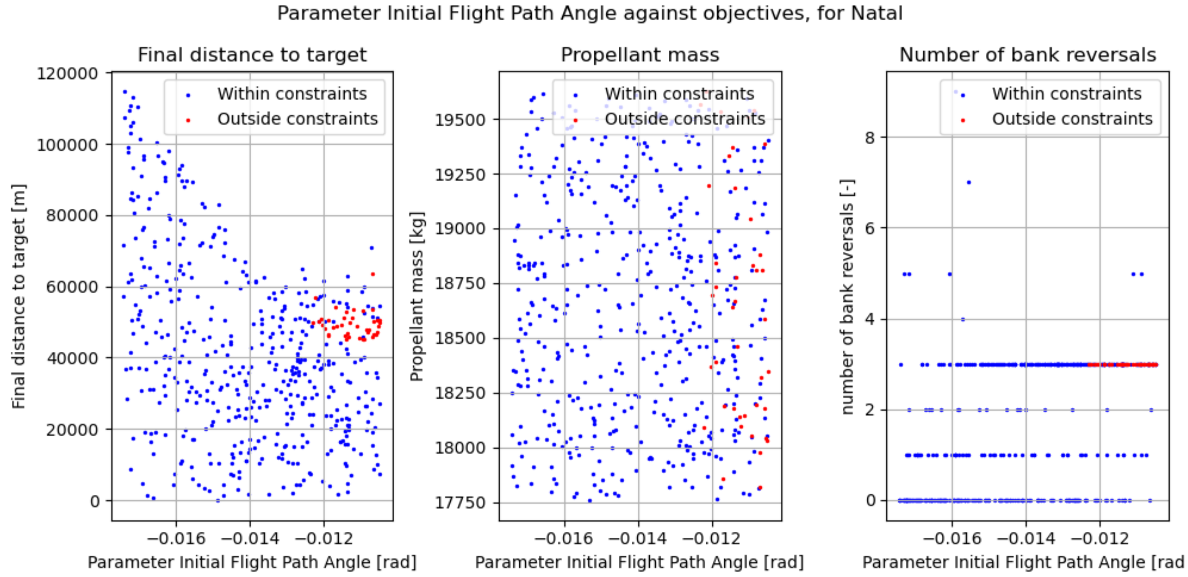

Figure B.53: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for Natal.

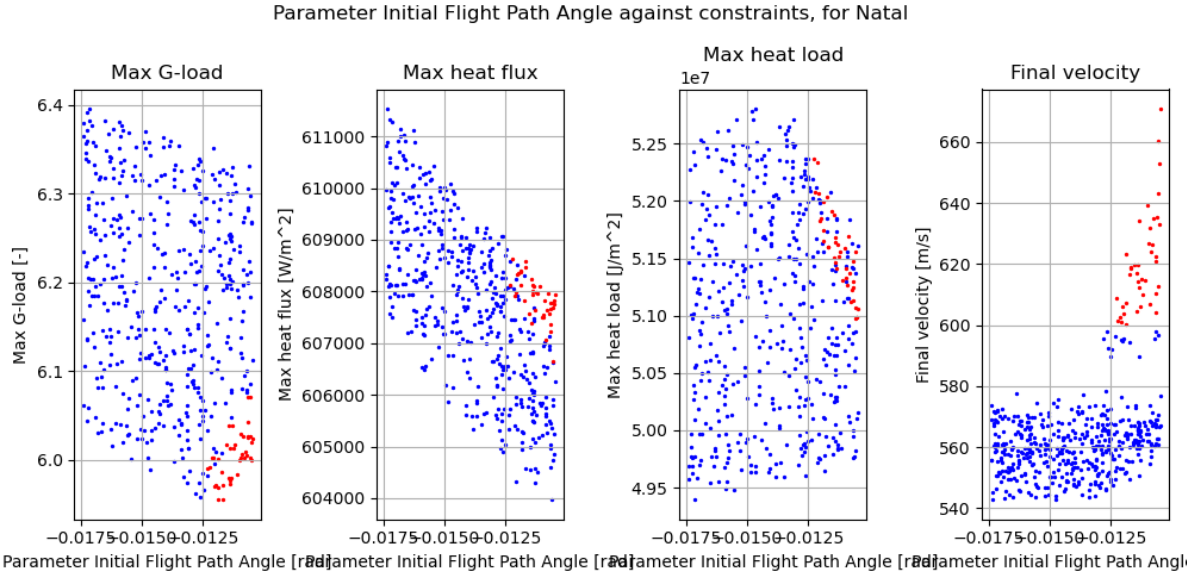

Figure B.54: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for Natal.

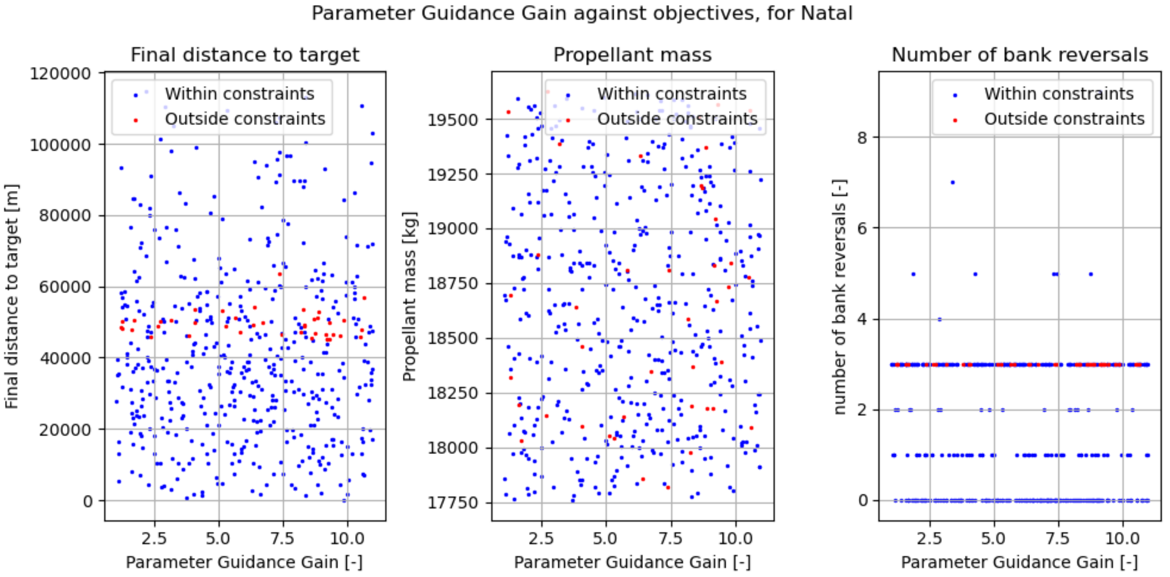

Figure B.55: The final distance to target, propellant mass and number of bank angle reversals over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for Natal.

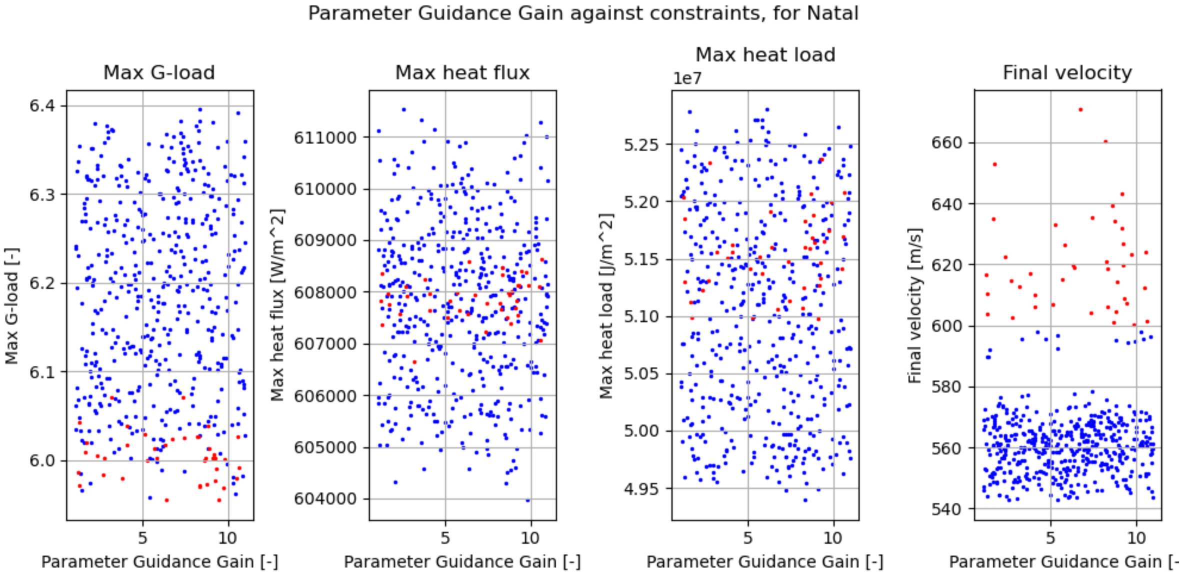

Figure B.56: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for Natal.

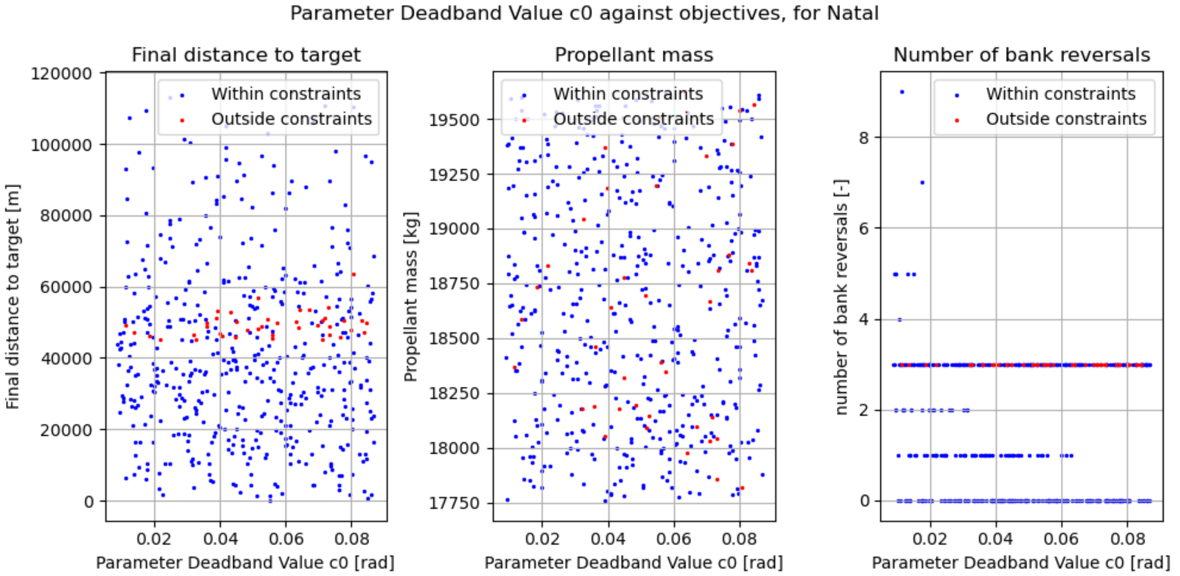

Figure B.57: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for Natal.

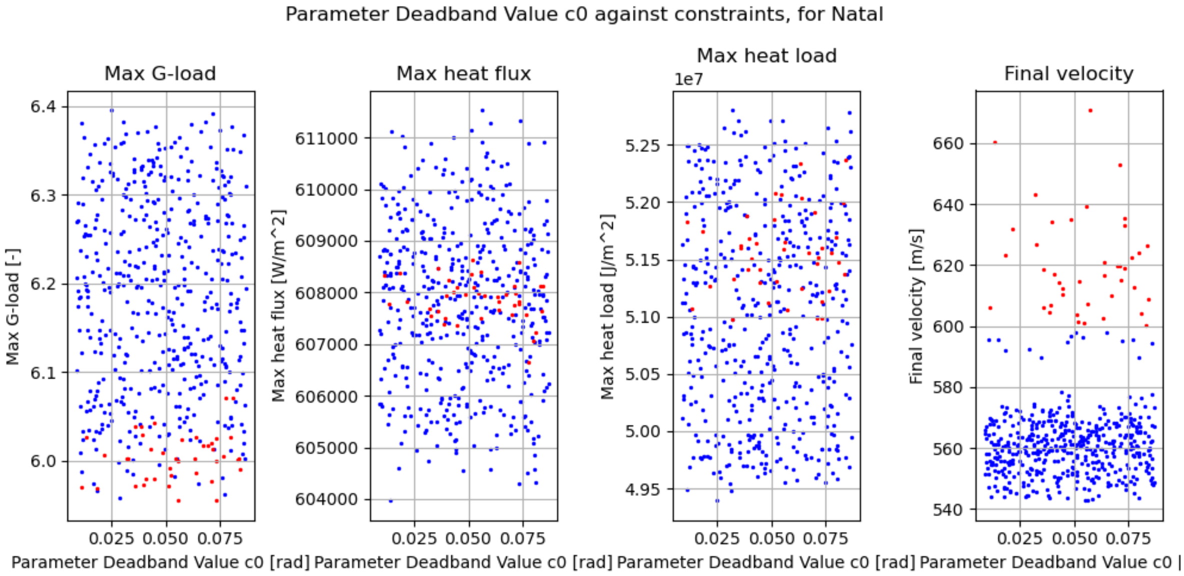

Figure B.58: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for Natal.

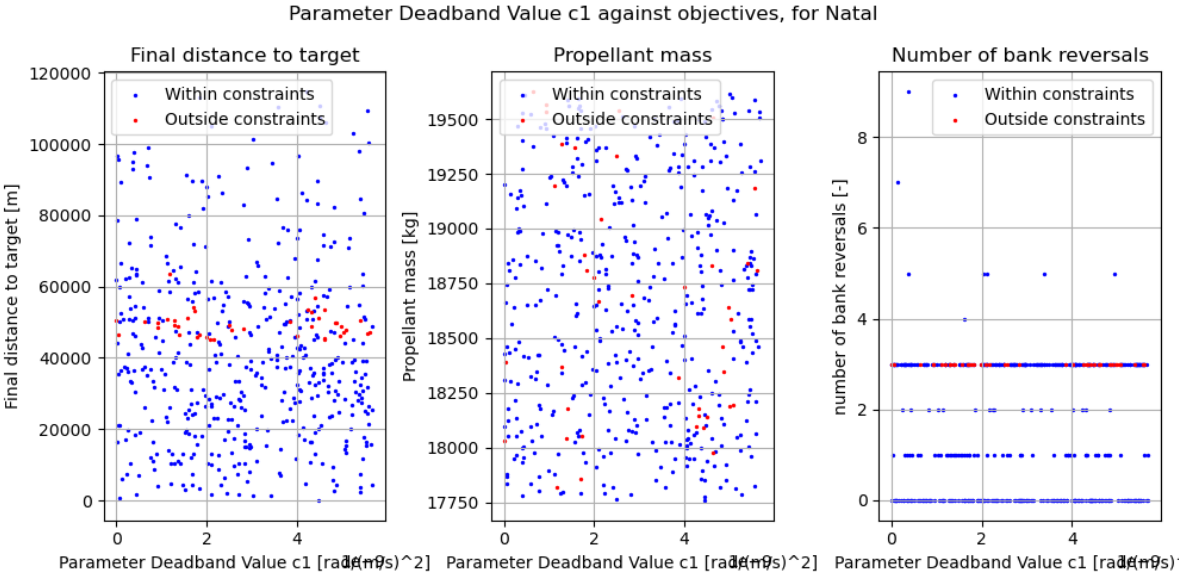

Figure B.59: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for Natal.

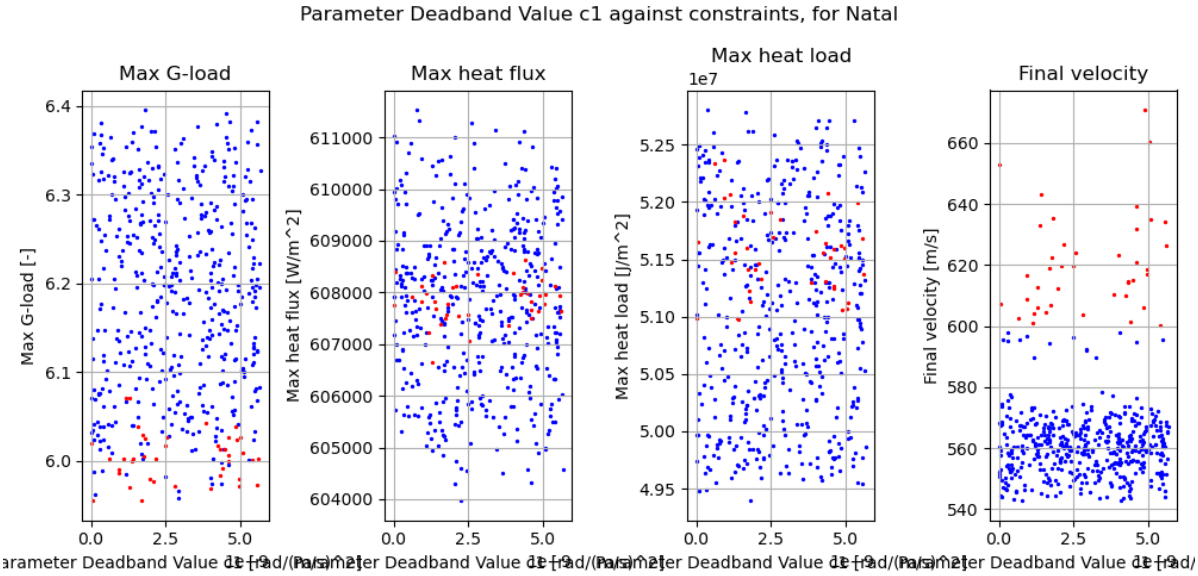

Figure B.60: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for Natal.

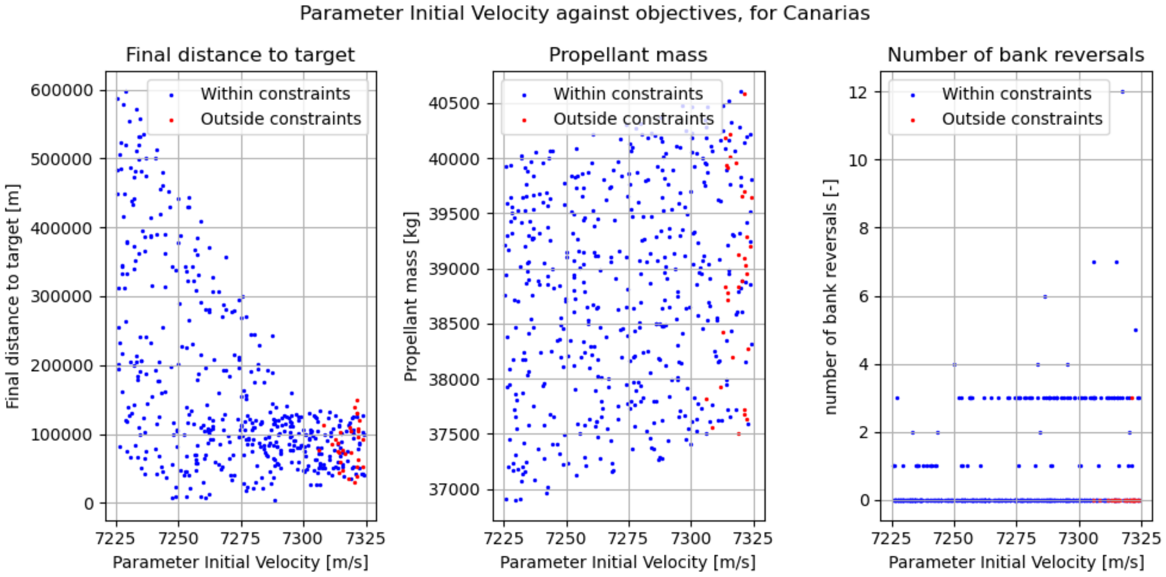

Figure B.61: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial velocity produced by simulations in the all-at-once MC variation for The Canarias.

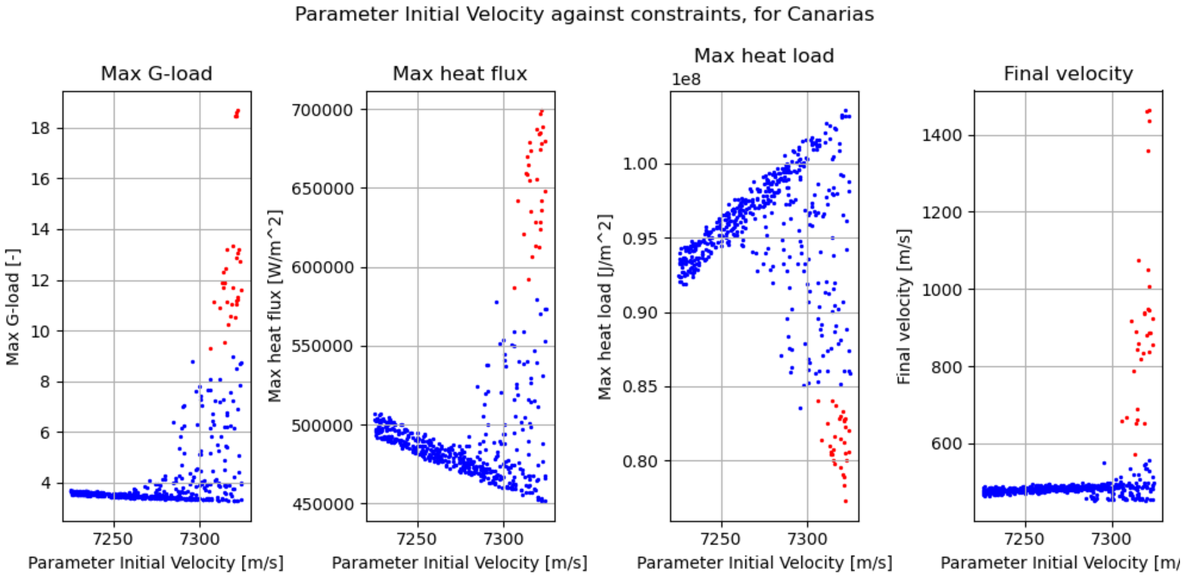

Figure B.62: The maximum g-load, maximum heat flux and total heat load over the value of the initial velocity produced by simulations in the all-at-once MC variation for The Canarias.

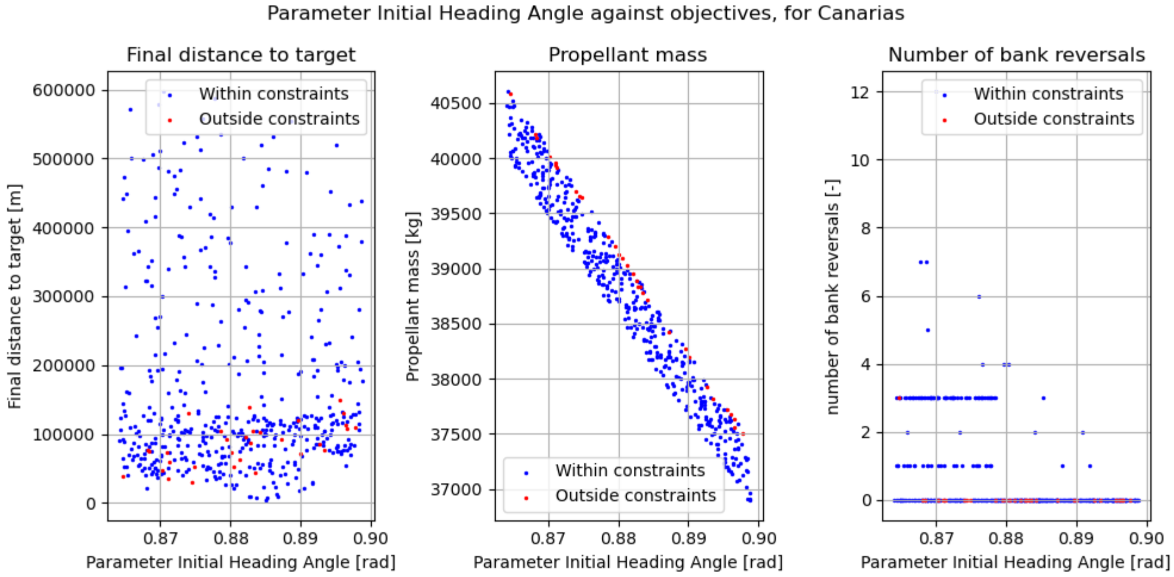

Figure B.63: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial heading angle produced by simulations in the all-at-once MC variation for The Canarias.

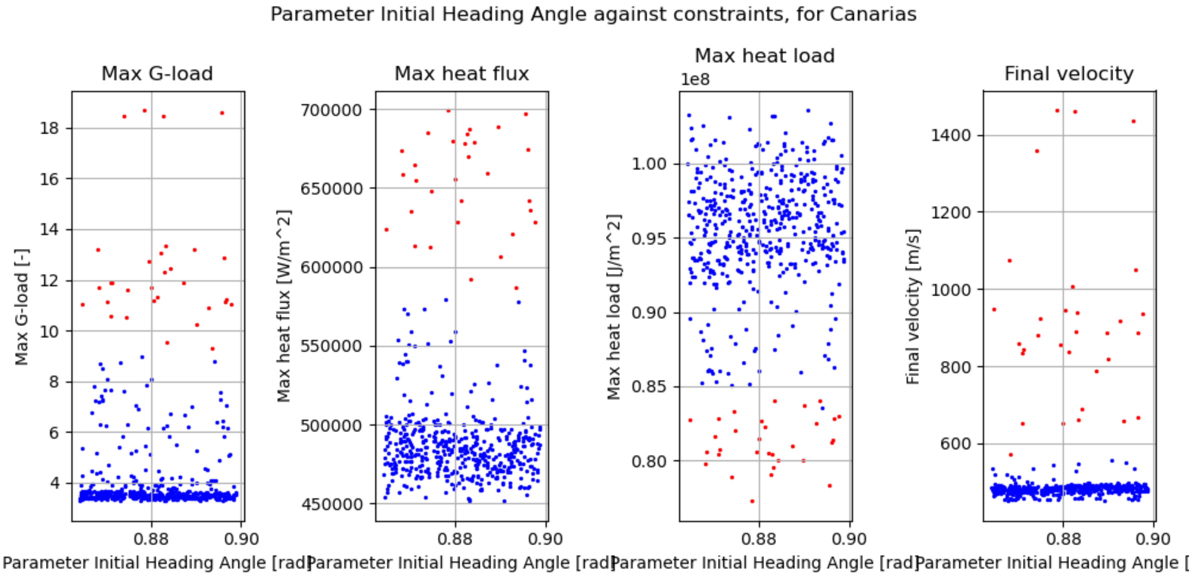

Figure B.64: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial heading angle produced by simulations in the all-at-once MC variation for The Canarias.

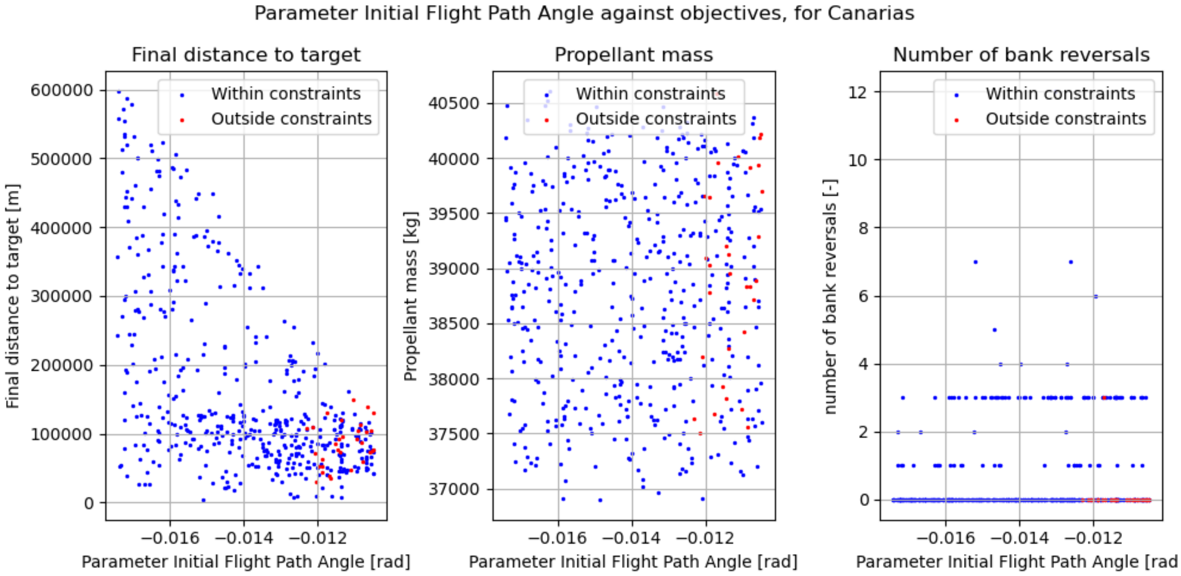

Figure B.65: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for The Canarias.

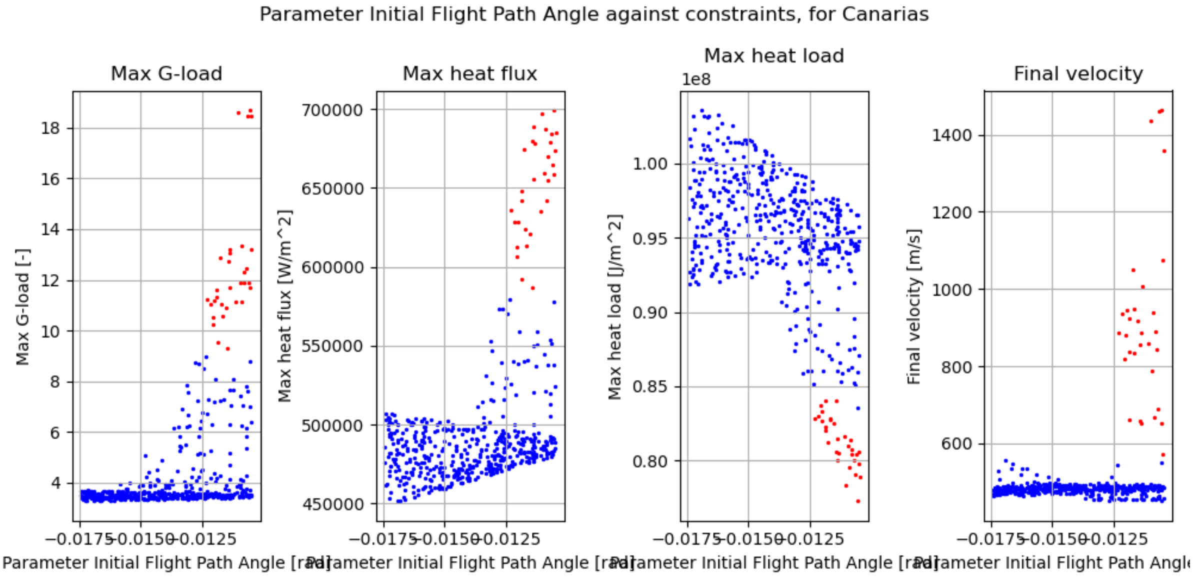

Figure B.66: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for The Canarias.

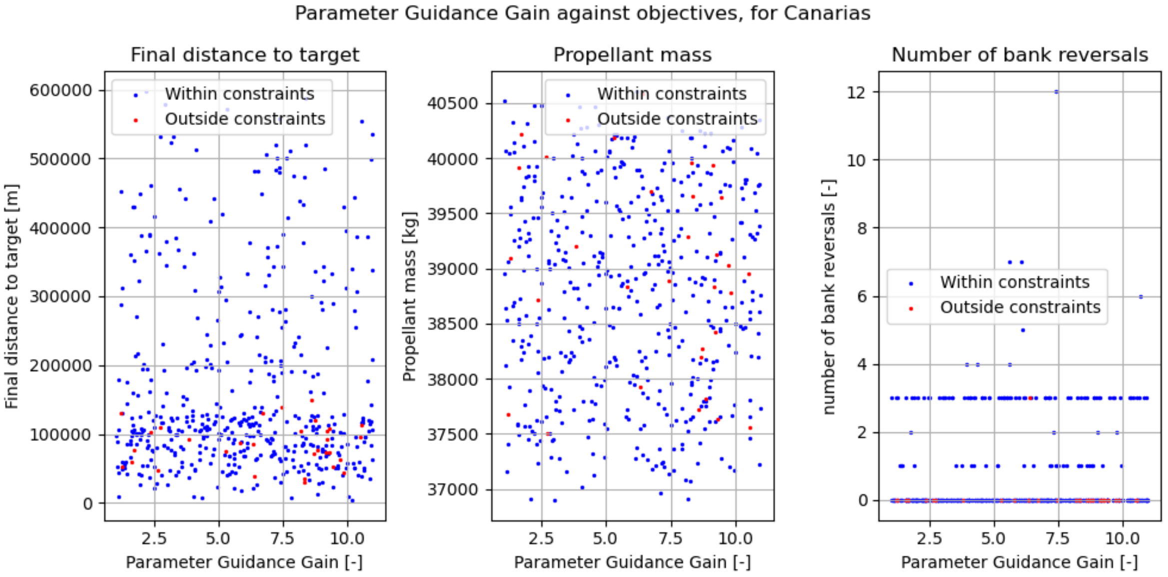

Figure B.67: The final distance to target, propellant mass and number of bank angle reversals over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for The Canarias.

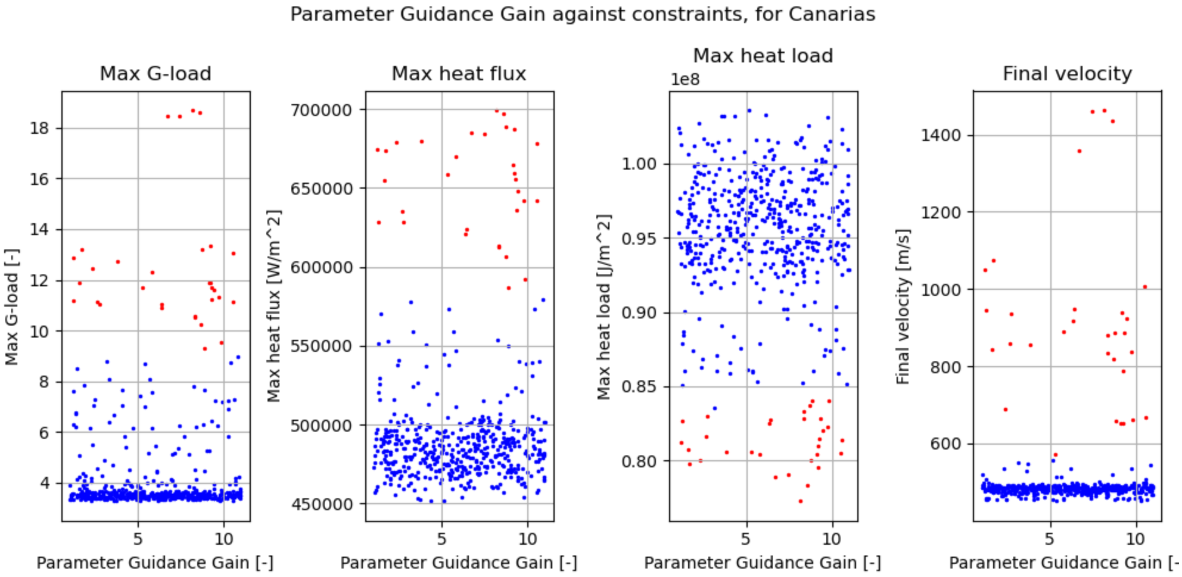

Figure B.68: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for The Canarias.

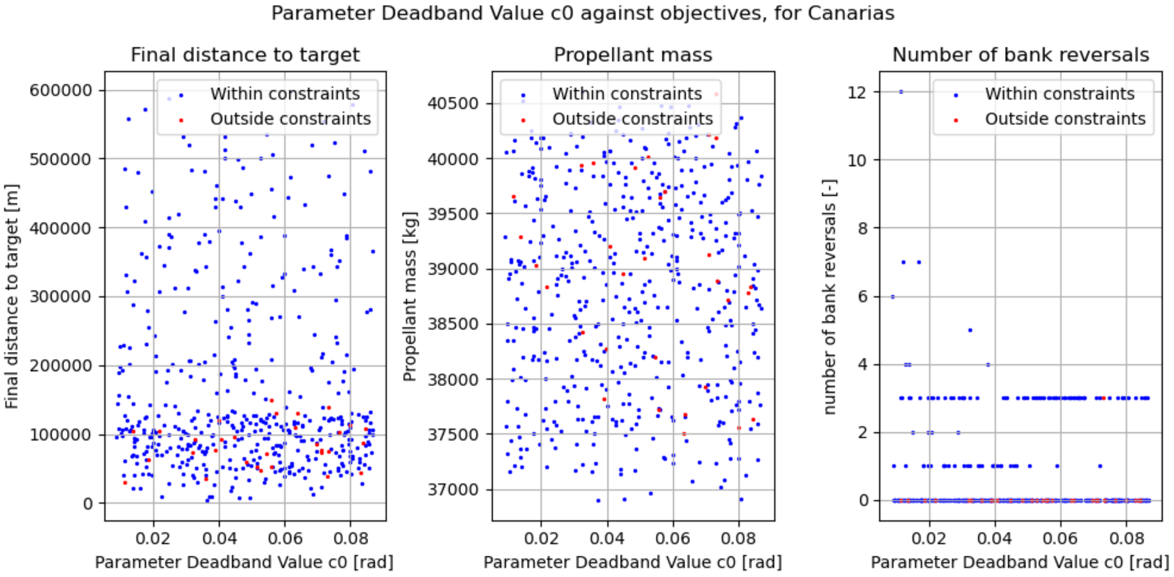

Figure B.69: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for The Canarias.

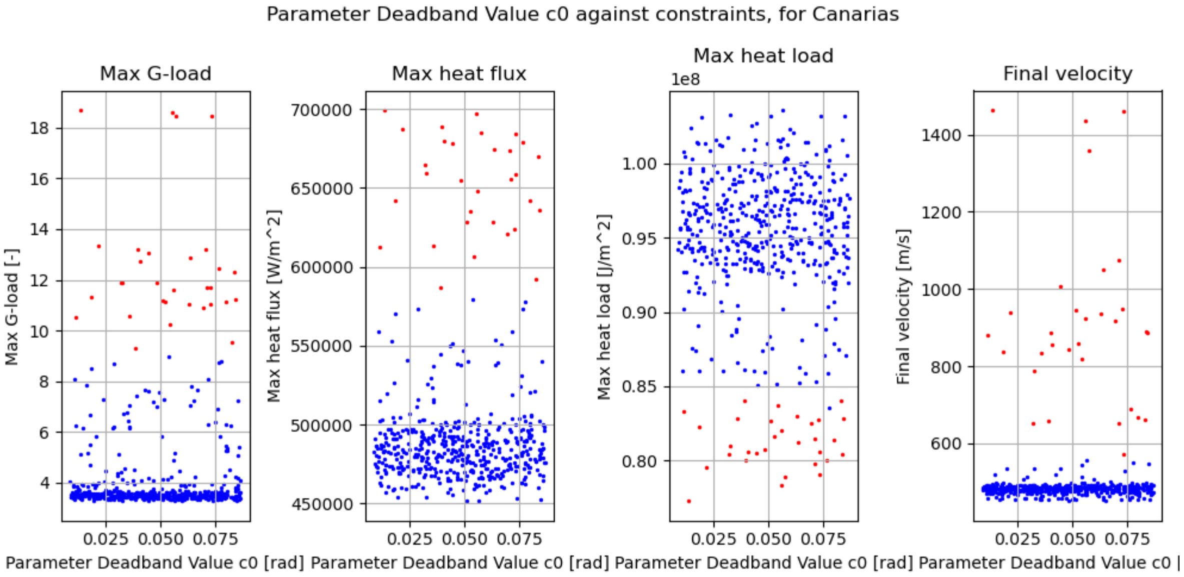

Figure B.70: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for The Canarias.

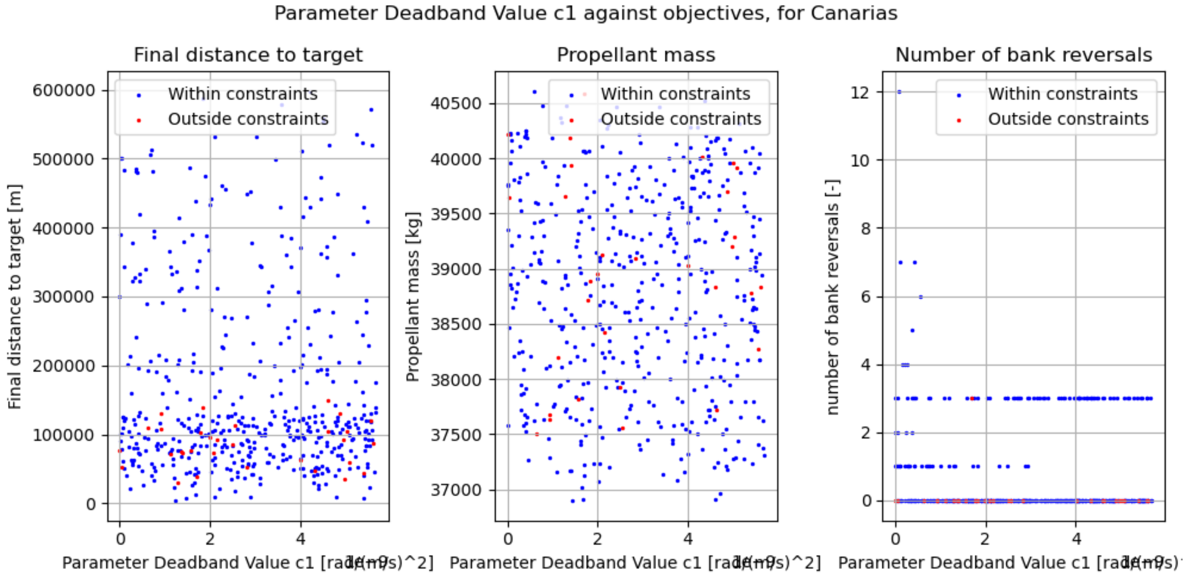

Figure B.71: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for The Canarias.

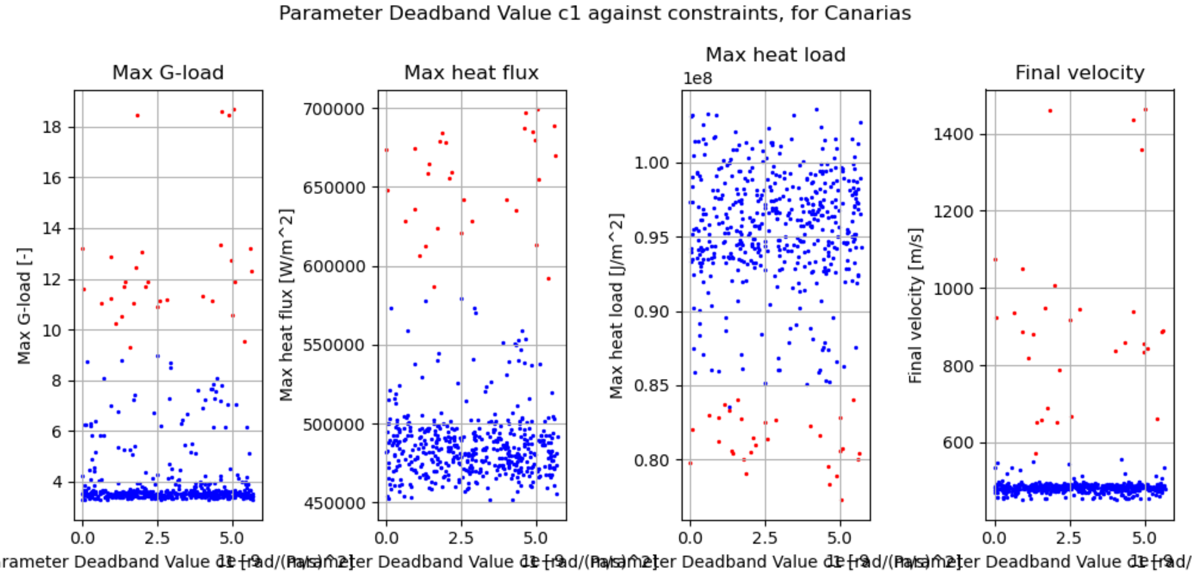

Figure B.72: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for The Canarias.

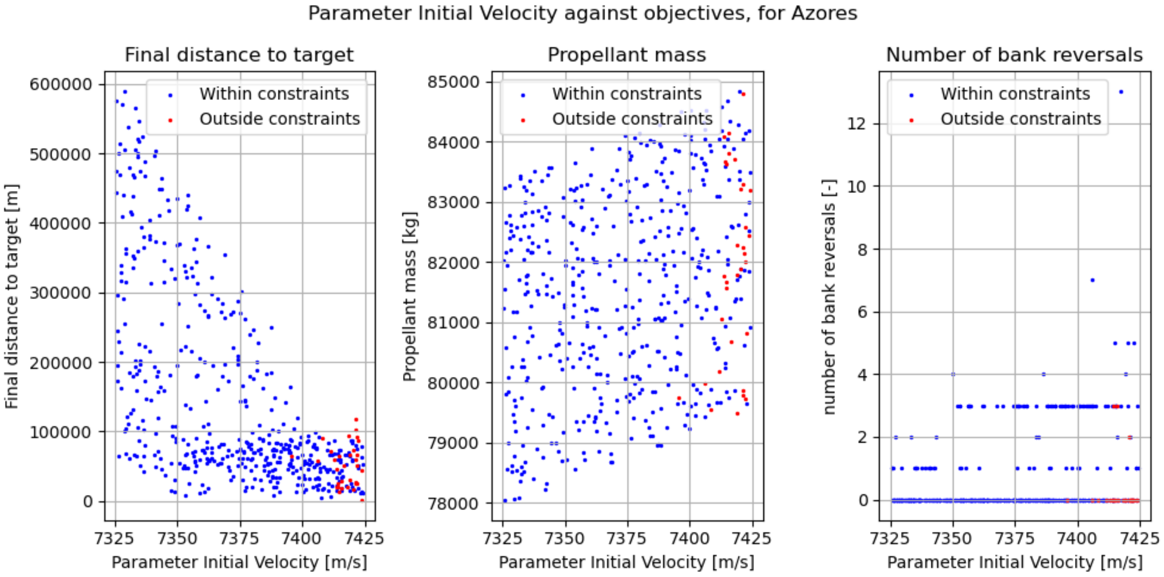

Figure B.73: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial velocity produced by simulations in the all-at-once MC variation for The Azores.

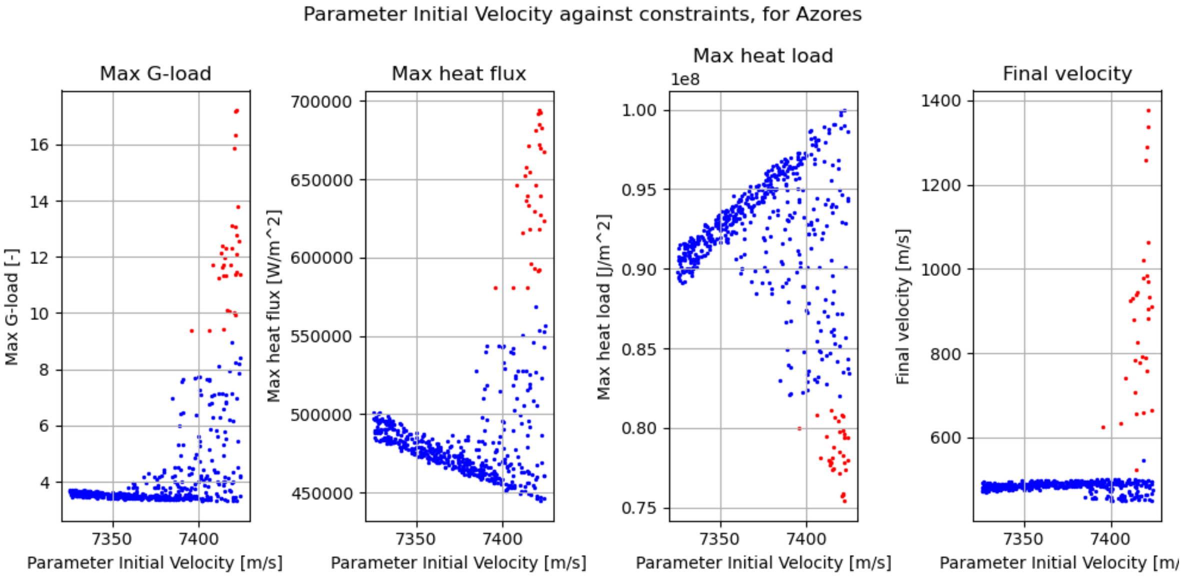

Figure B.74: The maximum g-load, maximum heat flux and total heat load over the value of the initial velocity produced by simulations in the all-at-once MC variation for The Azores.

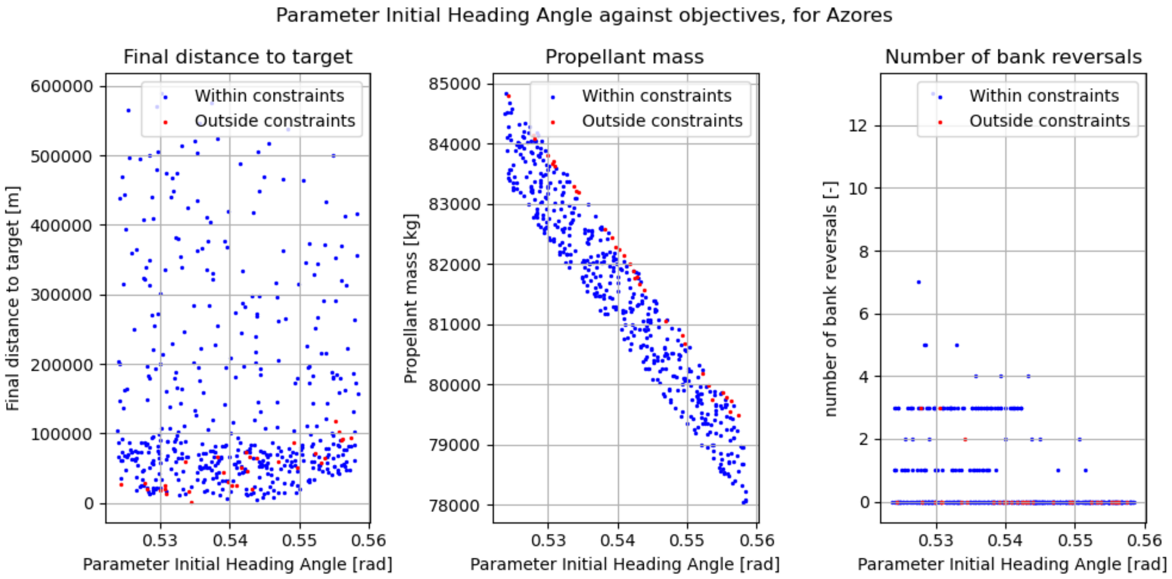

Figure B.75: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial heading angle produced by simulations in the all-at-once MC variation for The Azores.

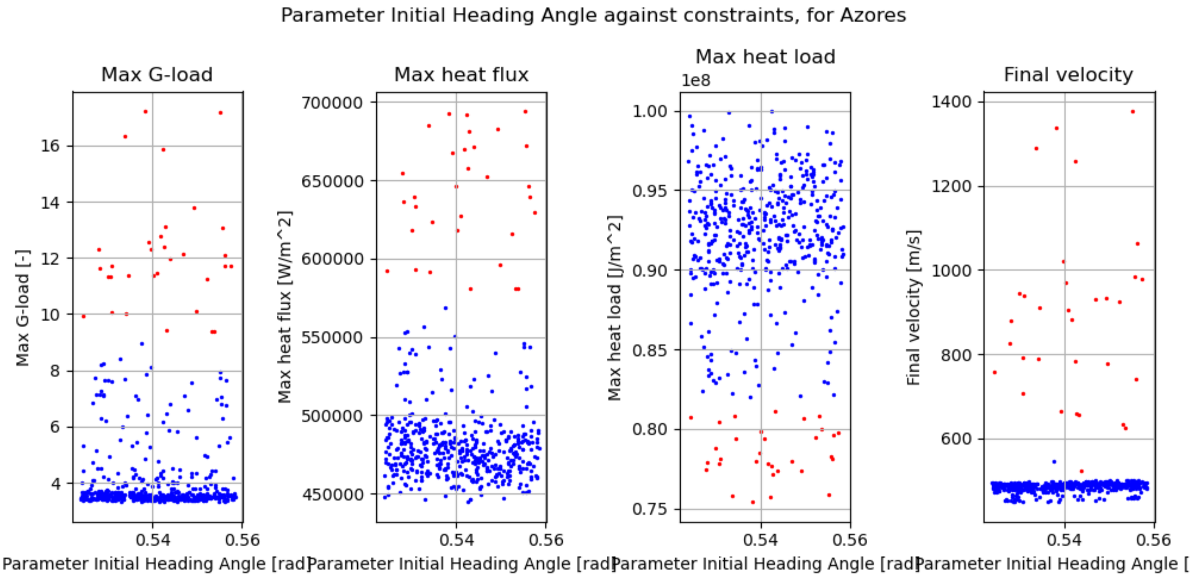

Figure B.76: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial heading angle produced by simulations in the all-at-once MC variation for The Azores.

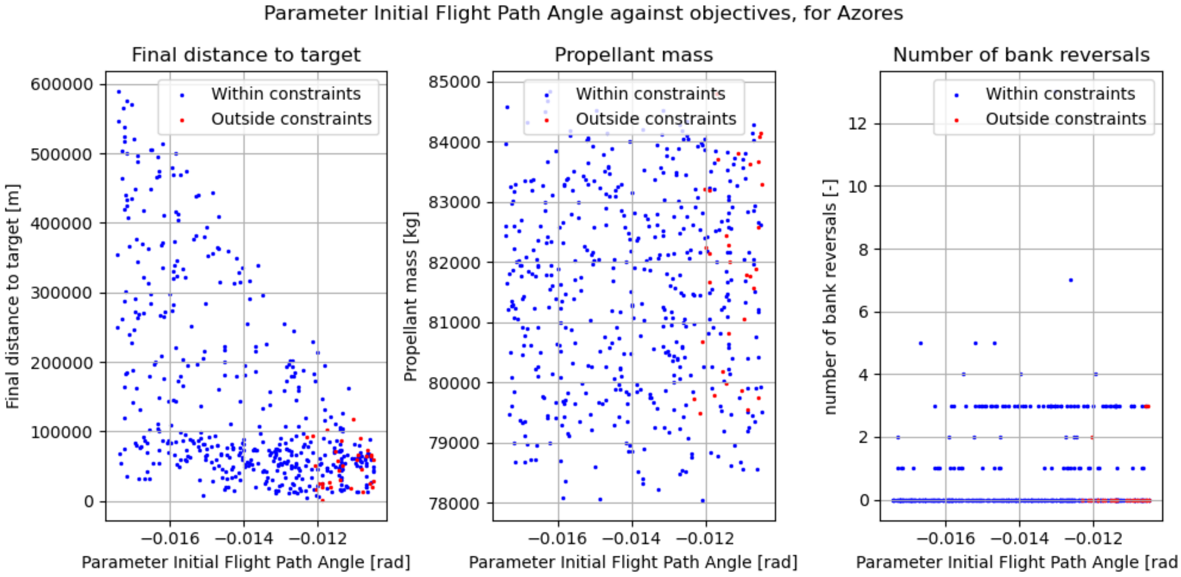

Figure B.77: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for The Azores.

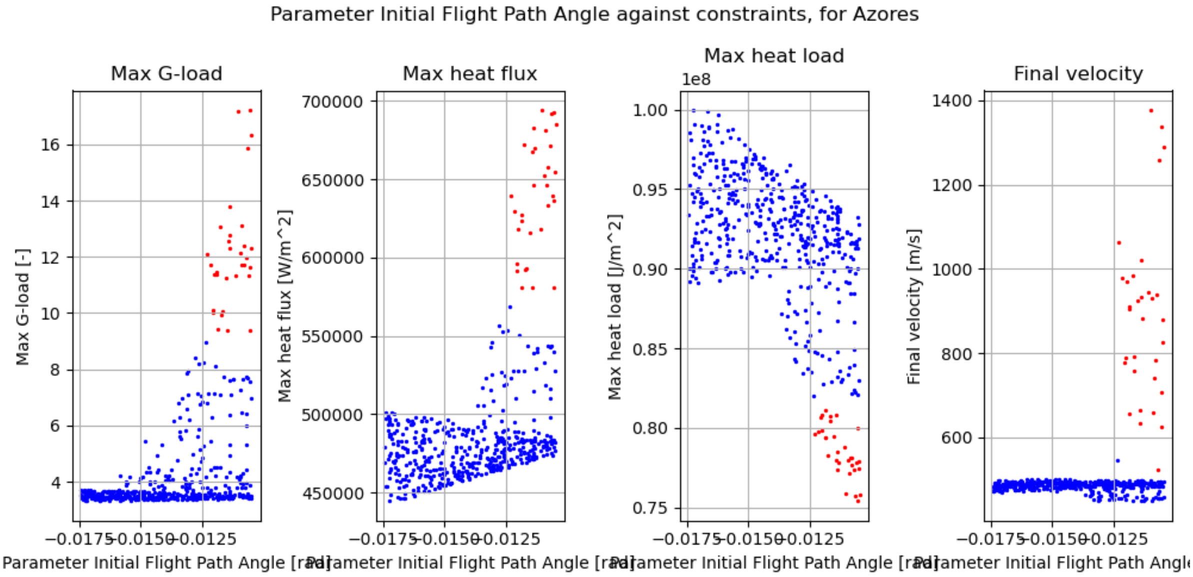

Figure B.78: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for The Azores.

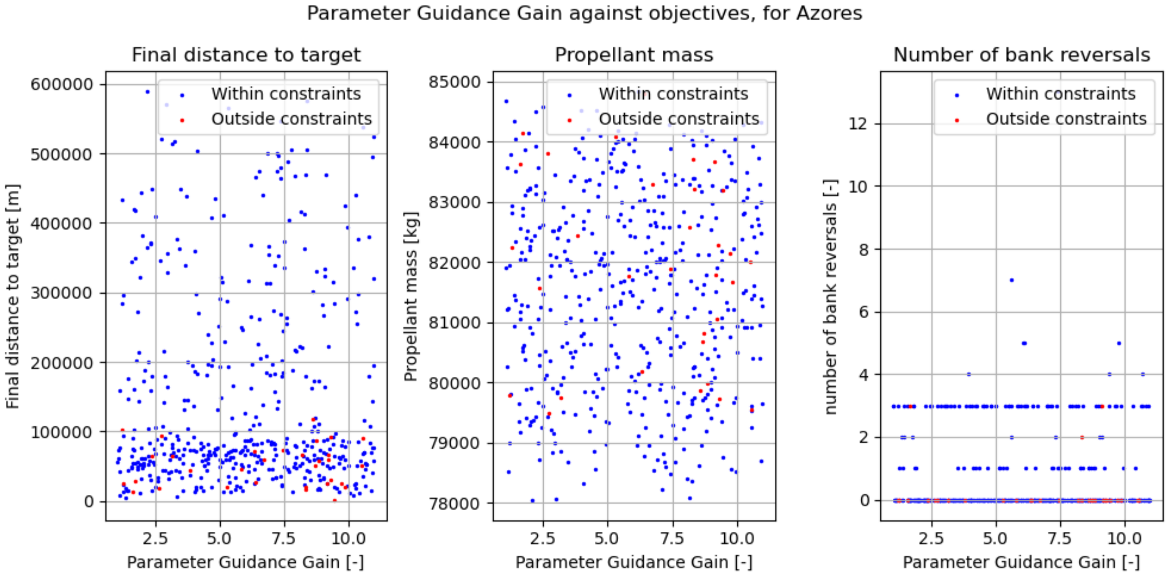

Figure B.79: The final distance to target, propellant mass and number of bank angle reversals over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for The Azores.

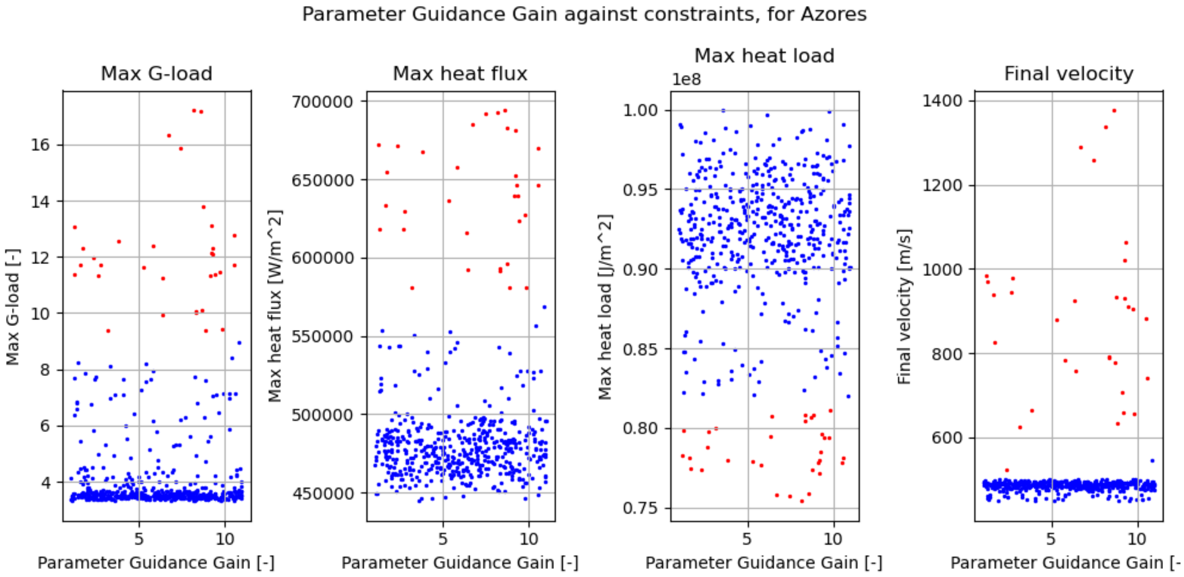

Figure B.80: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for The Azores.

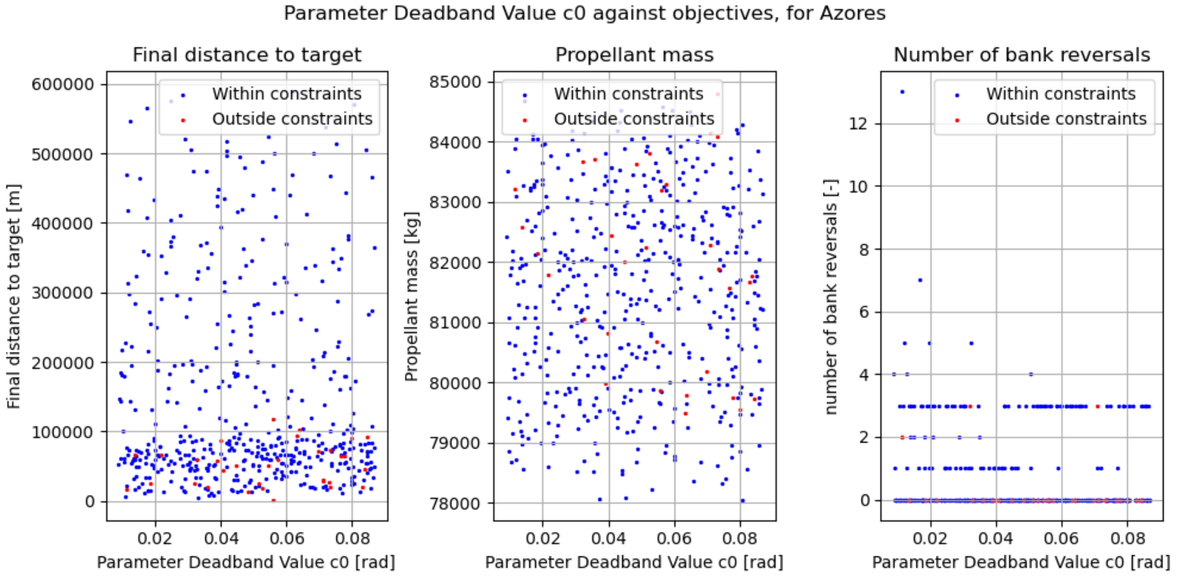

Figure B.81: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for The Azores.

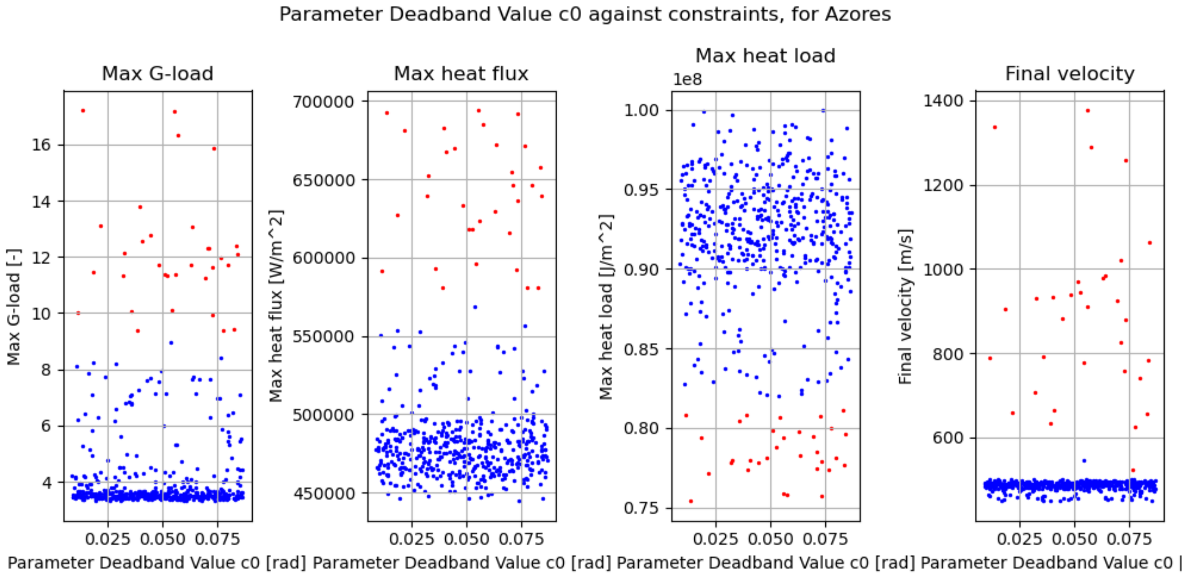

Figure B.82: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for The Azores.

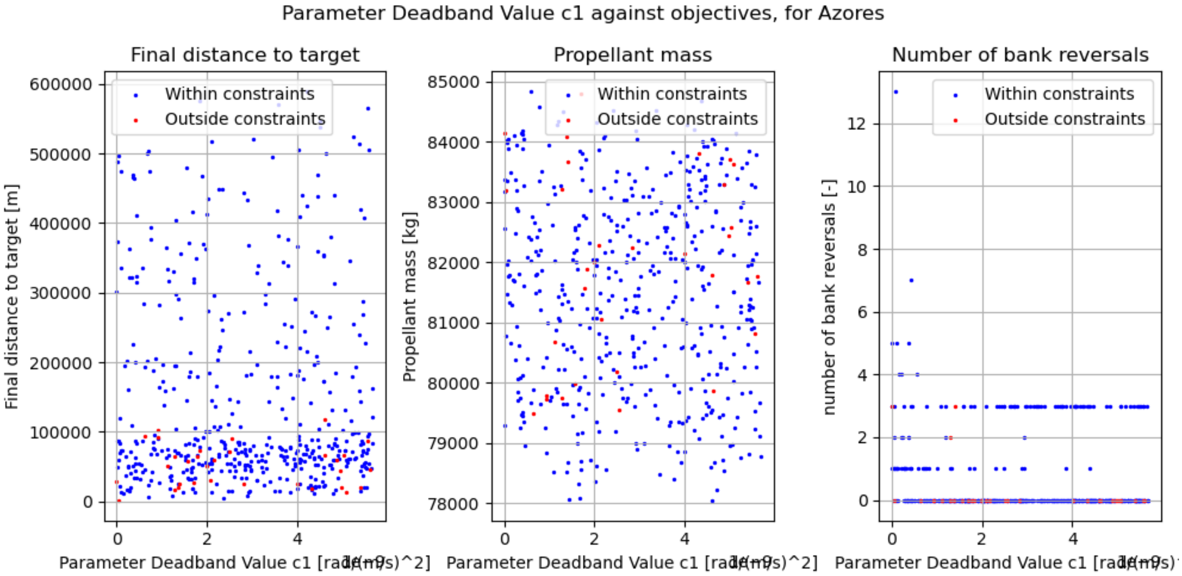

Figure B.83: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for The Azores.

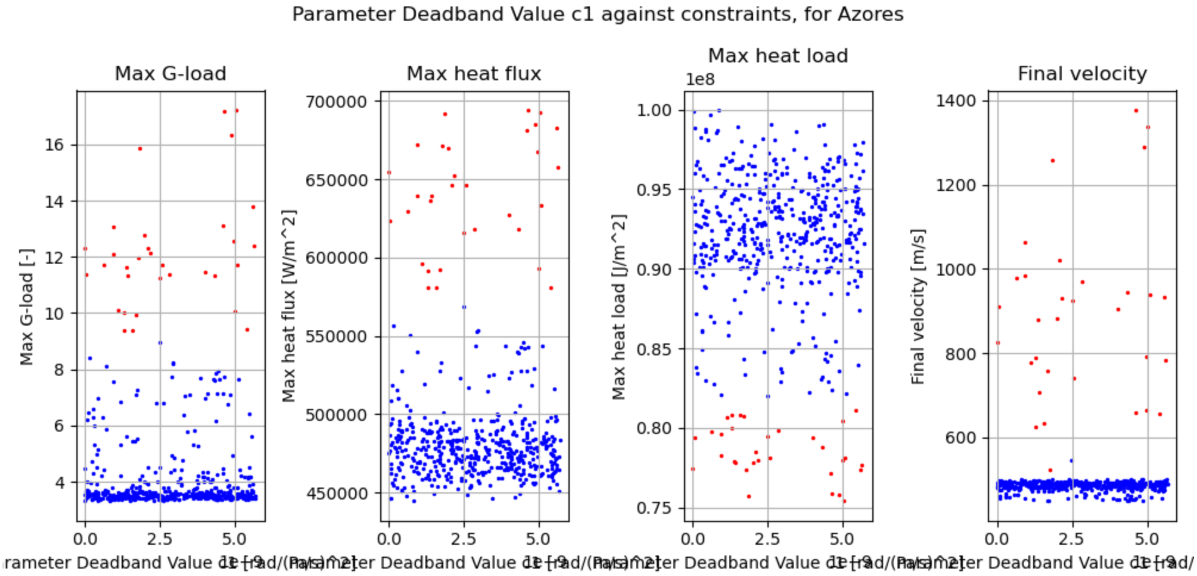

Figure B.84: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for The Azores.

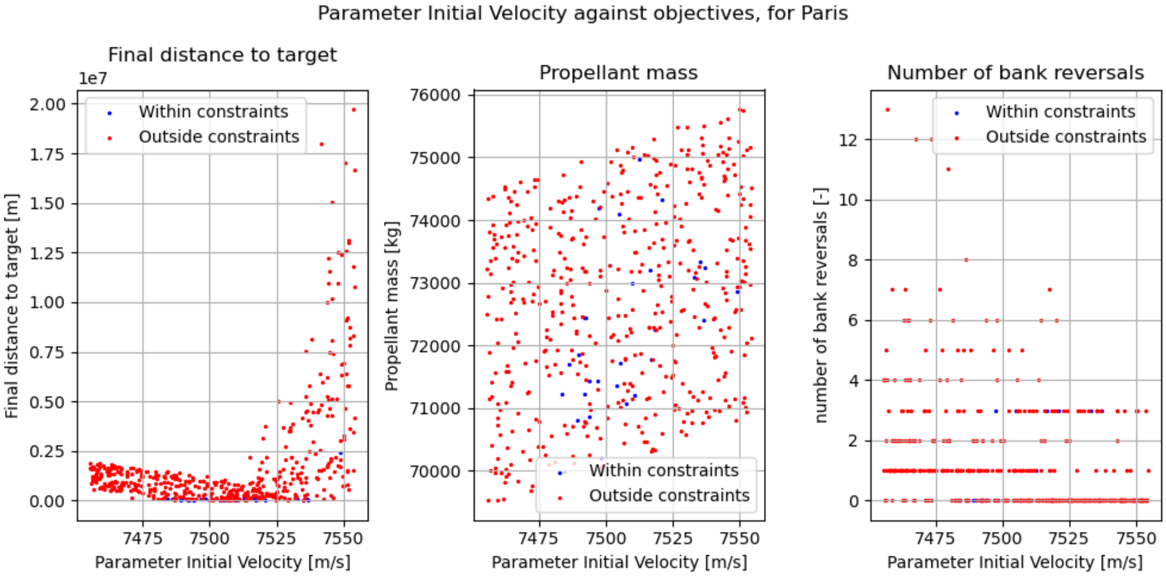

Figure B.85: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial velocity produced by simulations in the all-at-once MC variation for Les Mureaux.

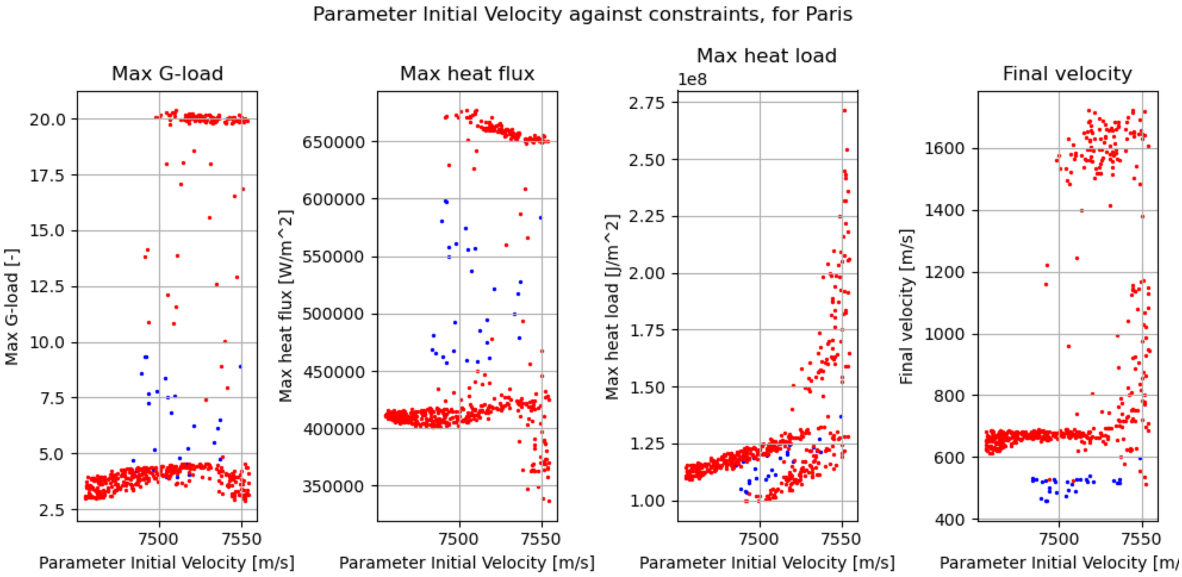

Figure B.86: The maximum g-load, maximum heat flux and total heat load over the value of the initial velocity produced by simulations in the all-at-once MC variation for Les Mureaux.

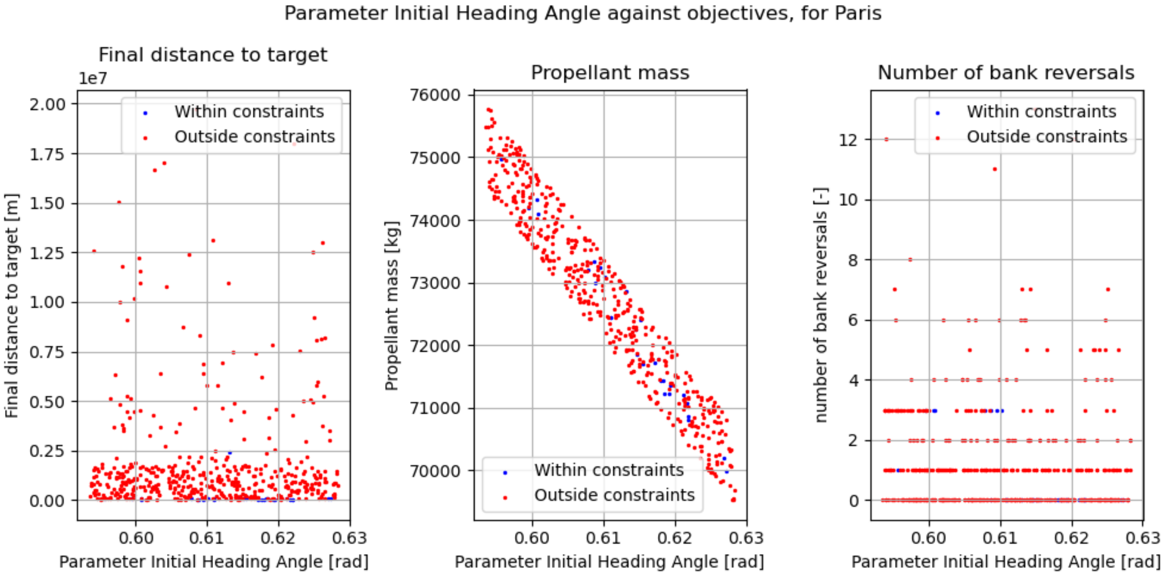

Figure B.87: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial heading angle produced by simulations in the all-at-once MC variation for Les Mureaux.

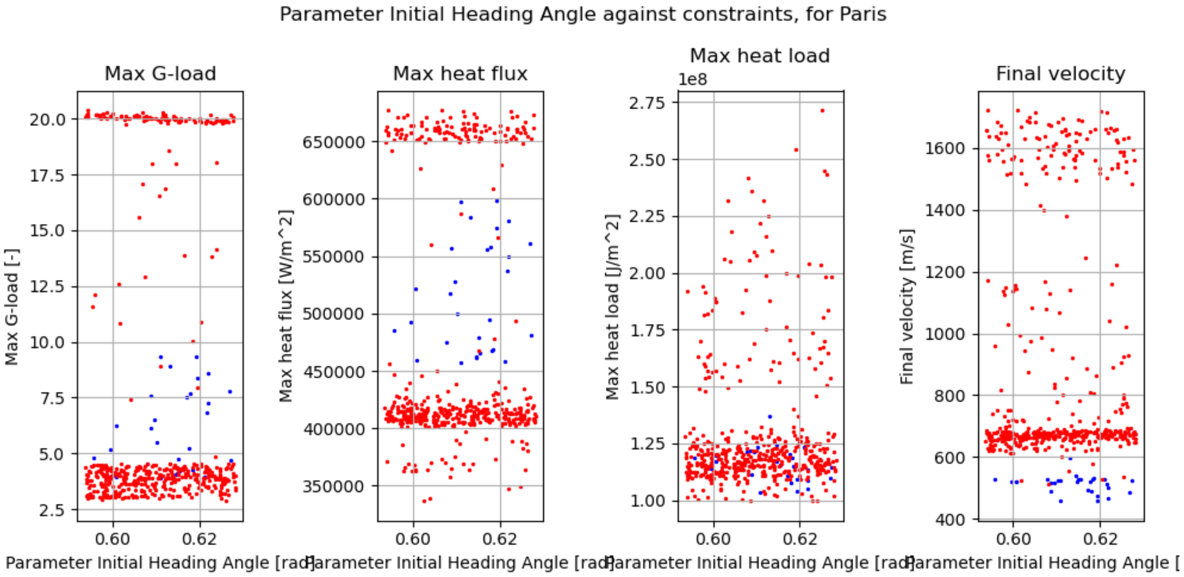

Figure B.88: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial heading angle produced by simulations in the all-at-once MC variation for Les Mureaux.

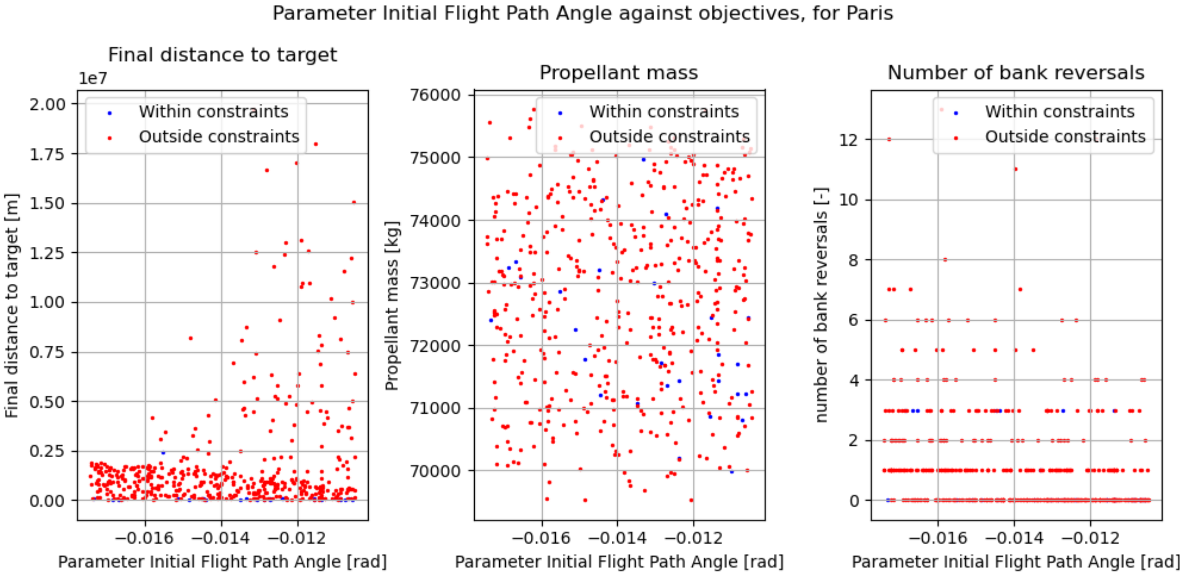

Figure B.89: The final distance to target, propellant mass and number of bank angle reversals over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for Les Mureaux.

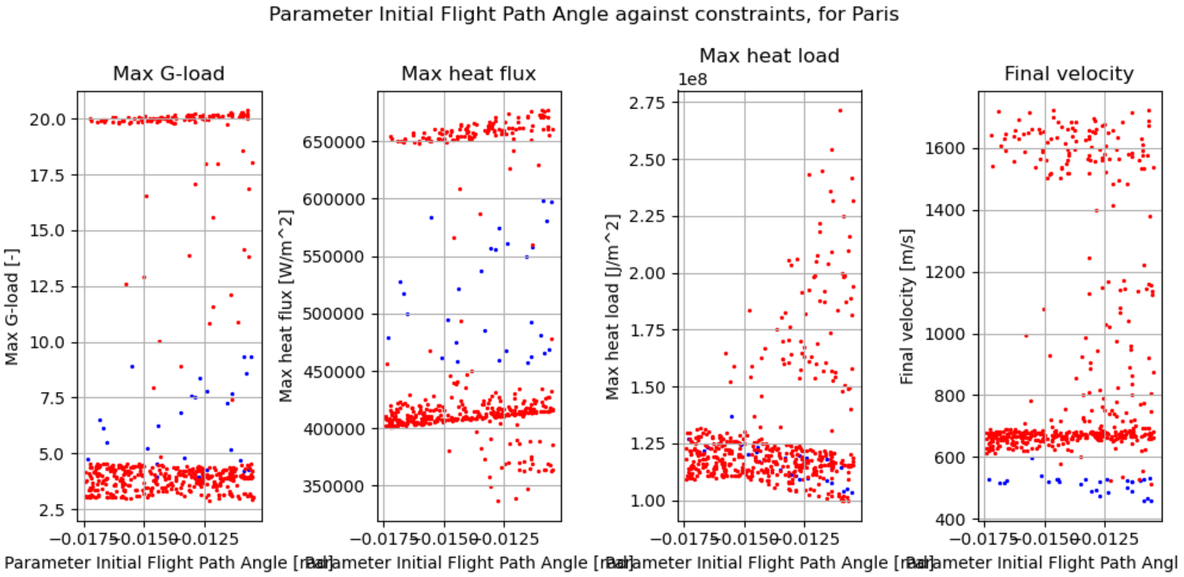

Figure B.90: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the initial flight path angle produced by simulations in the all-at-once MC variation for Les Mureaux.

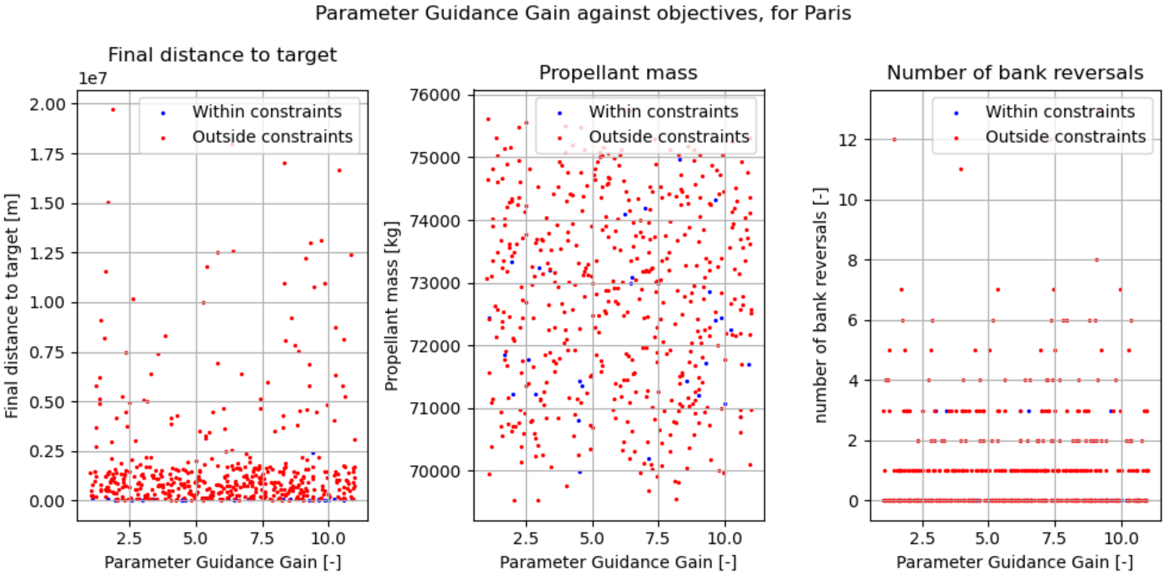

Figure B.91: The final distance to target, propellant mass and number of bank angle reversals over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for Les Mureaux.

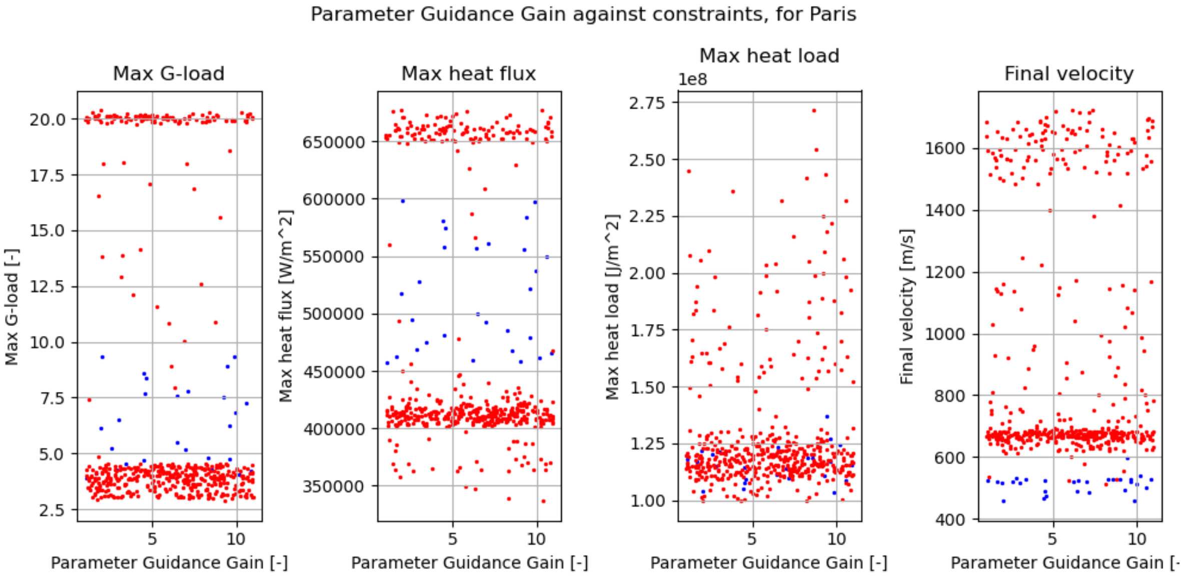

Figure B.92: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the guidance overcontrol gain produced by simulations in the all-at-once MC variation for Les Mureaux.

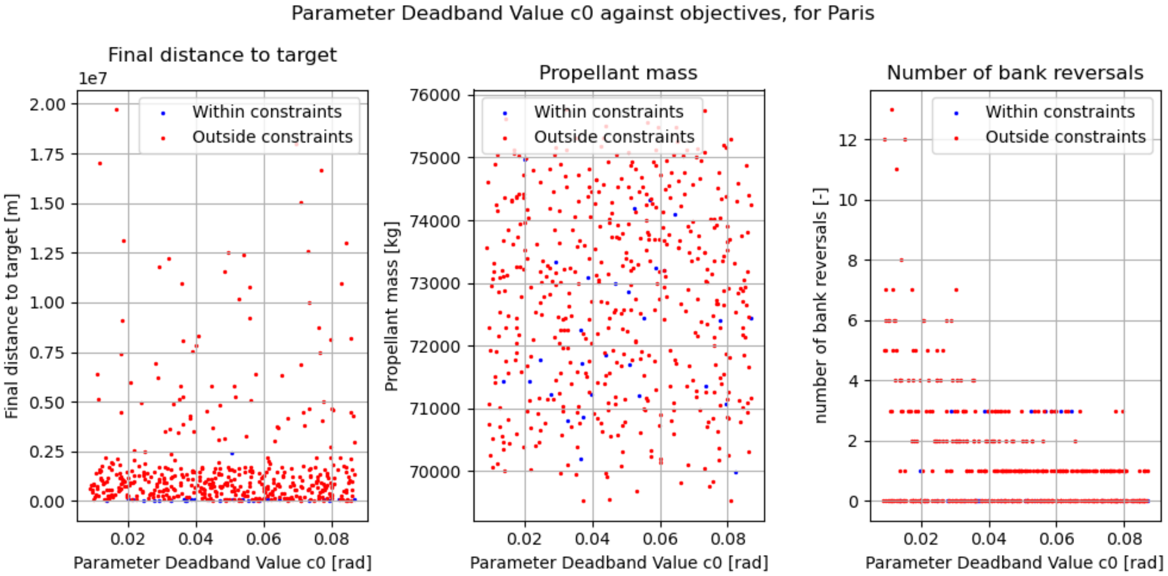

Figure B.93: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for Les Mureaux.

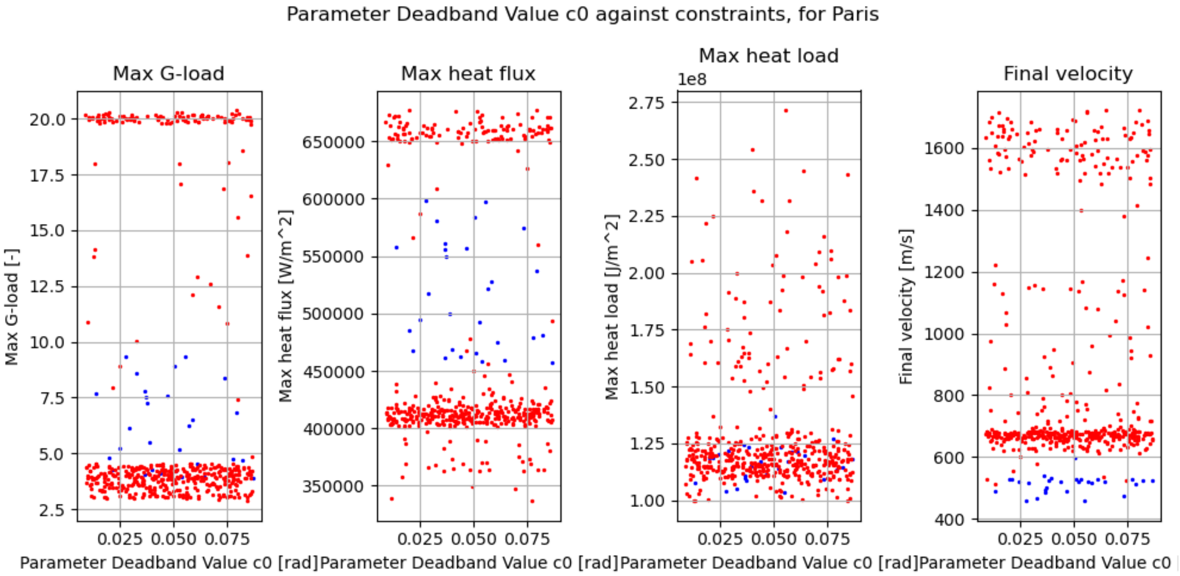

Figure B.94: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c0$ produced by simulations in the all-at-once MC variation for Les Mureaux.

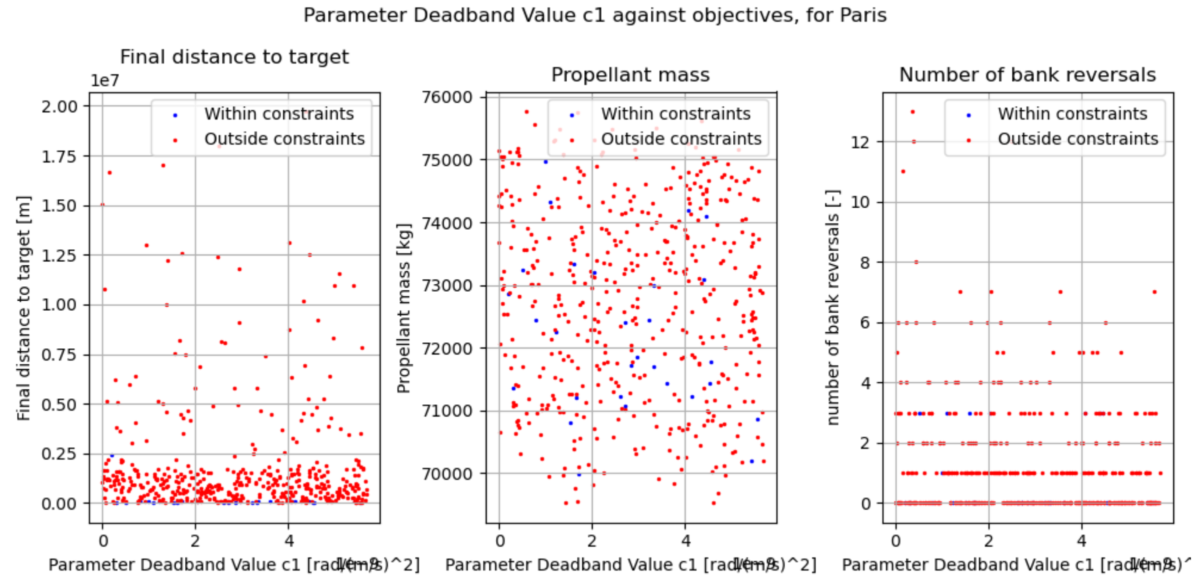

Figure B.95: The final distance to target, propellant mass and number of bank angle reversals over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for Les Mureaux.

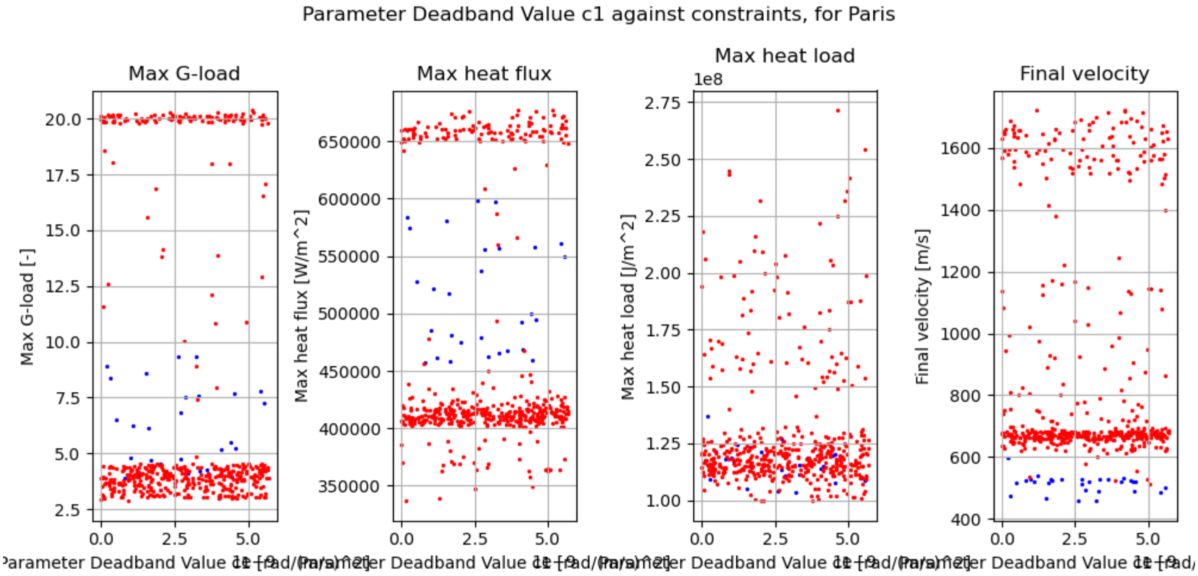

Figure B.96: The maximum g-load, maximum heat flux, total heat load and final velocity over the value of the deadband $c1$ produced by simulations in the all-at-once MC variation for Les Mureaux.