Landing Trajectory Design Using Invariant Manifolds of Quasi Satellite Orbits

A Phobos Case Study

Miguel Avillez

Delft University of Technology

Landing Trajectory Design Using Invariant Manifolds of Quasi Satellite Orbits

A Phobos Case Study

by

Miguel Avillez

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday November 18, 2022 at 14:00.

Student number:	5386012	
Thesis committee:	Ir. R. Noomen,	Supervisor
	Dr. ir. E. Mooij	Committee chair
	Dr. A. Cervone	External examiner

This thesis is confidential and cannot be made public until November 18, 2024.

Cover image: Phobos over Mars, as captured by Mars Express' HRSC camera, on March 26, 2010
(original colors). Credits: ESA's Planetary Science Archive.

Preface

There is a slightly long but, hopefully, interesting and insightful report ahead, so I am going to keep this short. I want to thank Ron for the guidance, encouragement, and for always reminding me that progress was actually happening. Thank you to Dominic, for all the support when it came to using and implementing code in Tudat. I also wish to acknowledge Dr. Xuanyu Hu, for providing the gravity model used in the verification, and Maarten, for some very useful discussion regarding the uncertainty analysis. Finally, I would like to thank my family, none of this would be possible without their unconditional support.

Miguel Avillez
Delft, November 2022

Abstract

Quasi-satellite orbits (QSOs) are relative trajectories in three-body problems, in which a spacecraft appears to orbit the smaller body. QSOs constitute a major element of the proposals for orbiting several moons, e.g. Europa and Phobos; in particular, the upcoming MMX mission (developed by JAXA, to be launched in 2024) will use a series of two- and three-dimensional QSOs to orbit Phobos. In this work, the feasibility of using unstable invariant manifolds of QSOs for the design of landing trajectories is studied, using the Mars-Phobos system as a case study. Invariant manifolds are considered potentially beneficial due to the negligible ΔV required to insert the spacecraft into one. The analysis focuses on the invariant manifolds of three-dimensional QSOs as these are more favorable in terms of their scientific return (e.g. due to the better surface coverage).

A simplified model of the Mars-Phobos system was used: the circular restricted three-body problem formulation with the secondary's gravity described by a polyhedron. The usage of an autonomous model means that the QSOs are organized in monoparametric families, which were computed via pseudo-arclength continuation with a multiple-shooting differential corrector. These were first computed using an ellipsoidal model of Phobos and later continued to the irregular polyhedron model via homotopy. The computed three-dimensional families of QSOs, which bifurcate from the two-dimensional family, can be divided into two groups: families with approximately cylindrical orbits which can be continued up to high altitudes, and families with irregular orbits which start intersecting Phobos' surface at low altitudes.

The unstable invariant manifolds of the orbits in these families were propagated. The families of cylindrical orbits generally proved unsuitable for the generation of impact manifolds, since, due to Phobos' weak gravitational influence, most manifolds were observed to escape the system. Some exceptions were noted for the families located closer to Phobos. The families of irregular orbits, which pass significantly closer to the surface, were observed to be much more suitable for generating impact manifolds: these allow a large number of landing opportunities per revolution and make it possible to reach the entire surface of Phobos (except the Stickney crater). Minimum impact velocities of $13\ \mathrm{m\,s^{-1}}$ were found, which would not allow an unpowered descent for landing requirements similar to MMX's (around $1\ \mathrm{m\,s^{-1}}$). However, a completely unpowered landing might be possible for spacecraft with higher allowed impact velocities.

The trajectories were considered to require an initial maneuver for the manifold insertion and a final braking maneuver, to control the impact velocity and impact angle. The maneuvers were assumed to be fully impulsive. The first maneuver, which requires a negligible ΔV, was designed via multiple shooting. The second maneuver was designed using semi-analytical techniques and via optimization. The latter allowed meeting the strict impact-velocity constraints (maximum vertical and horizontal velocities of, respectively, $1\ \mathrm{m\,s^{-1}}$ and $0.15\ \mathrm{m\,s^{-1}}$), with a minimum ΔV in the order of $12.5\ \mathrm{m\,s^{-1}}$.

Finally, the sensitivity of the designed trajectories to navigation, maneuver, and gravity uncertainties was studied. An analysis based on a linearized stability index showed that the stability of the manifolds is mainly dependent on their time of flight and the instability of the orbit from which they emanate, with the increase of these two factors increasing the manifolds' instability. A Monte Carlo analysis showed that the stability index can be used to effectively predict the real stability of a manifold, and that there are robust manifolds that always guarantee impacts with the surface, without the need for corrective maneuvers, even in the presence of uncertainties. Applying a similar Monte Carlo analysis to the final braking maneuver showed that it would not allow robustly meeting the selected landing requirements, indicating that the design of this maneuver would require further study if it were to be used.

The invariant manifolds of three-dimensional QSOs allow the generation of robust and propellant-efficient landing trajectories. These require similar ΔVs to other methods in the case of low-velocity landings, but negligible ΔV in the cases where higher impact velocities are allowed, thus being particularly promising in the second case. However, it was found that, in the Mars-Phobos system, the application of the manifolds is limited to the families of cylindrical orbits very close to the surface and to the families of irregular orbits. Unfortunately, neither of those include the three-dimensional QSO that will be used by MMX. Nevertheless, the Mars-Phobos system is hypothesized to be particularly unfavorable for the application of landings via manifolds due to Phobos' weak gravity. It is possible that larger numbers of impact manifolds and families generating them would be observed for systems where the secondary has a stronger gravitational influence.

Nomenclature

List of Abbreviations

ABM	Adams-Bashfort-Moulton integrator
COM	Center of mass
CR3BP	Circular restricted three-body problem
DE	Differential Evolution algorithm
DOPRI	Runge-Kutta-Dormand-Prince integrator
ESA	European Space Agency
FTLE	Finite-time Lyapunov exponent
H3BP	Hill three-body problem
JAXA	Japan Aerospace Exploration Agency
L_i	Lagrange point i ($i = 1, \ldots, 5$)
MMX	Martian Moons eXploration
P_i	Primaries in the CR3BP ($i = 1, 2$)
QSO	Quasi-satellite orbit
RK	Runge-Kutta integrator
RKF	Runge-Kutta-Fehlberg integrator
SADE	Self-Adaptive Differential Evolution algorithm
STM	State transition matrix

Latin Symbols

a	Ellipsoid's largest semi-major axis	[m]
A_i	Amplitude of a periodic orbit in the coordinate i	[m]
b	Ellipsoid's intermediate semi-major axis	[m]
c	Ellipsoid's smallest semi-major axis	[m]
C	Jacobi integral	[−]
$C^{F/I}$	Rotation matrix between frame F and frame I	[−]
$C_{n,m}$	Unnormalized spherical harmonics cosine coefficient of degree n and order m	[−]
$D_{\boldsymbol{F}}$	Jacobian matrix of $\boldsymbol{F}$	Various
E	Polyhedron edge dyad	[−]
f	Generic function expressing the state derivative	Various
f	Deceleration factor (in continuation step-size control)	[−]
F	Polyhedron facet dyad	[−]
$\boldsymbol{F}$	Continuation constraints vector	Various
G	Universal gravitational constant	$[\mathrm{m^3\,kg^{-1}\,s^{-2}}]$
k	Stability index	[−]
L	Polyhedron per-edge factor	[−]
m	Mass	[kg]
M	Monodromy matrix	Various
$\hat{\boldsymbol{n}}$	Unit vector normal to a polyhedron feature	[−]
P	Poincaré map	Various
P_n	Legendre polynomial of degree n	[−]
$P_{n,m}$	Associated Legendre polynomial of degree n and order m	[−]
$\boldsymbol{r}$	Position vector with respect to the synodic frame	[m]
R	Spherical harmonics reference radius	[m]
$\boldsymbol{R}$	Position vector with respect to the inertial frame	[m]

$S_{n,m}$	Unnormalized spherical harmonics sine coefficient of degree n and order m	[−]
t	Time	[s]
t_1	Initial time	[s]
T	Orbit period	[s]
U	Gravitational potential	[$\mathrm{m^2\,s^{-2}}$]
V	Volume	[$\mathrm{m^3}$]
$\boldsymbol{V}$	Continuation free variables vector	Various
$\boldsymbol{V}$	Velocity vector with respect to the synodic frame	[$\mathrm{m\,s^{-1}}$]
x	First position coordinate with respect to the synodic frame	[m]
$\boldsymbol{x}$	Cartesian state vector with respect to the synodic frame	[m, $\mathrm{m\,s^{-1}}$]
y	Second position coordinate with respect to the synodic frame	[m]
z	Third position coordinate with respect to the synodic frame	[m]

Greek Symbols

ΔC	Change in Jacobi integral	[−]
Δs	Pseudo-arclength continuation step-size	[−]
Δt	Time interval	[s]
ΔV	Change in velocity	[$\mathrm{m\,s^{-1}}$]
ε	Multiple-shooting modified error	Various
ε	Norm of the manifold's initial state perturbation	[m, $\mathrm{m\,s^{-1}}$]
$\tilde{\varepsilon}$	Norm of the manifold's initial position perturbation	[m]
ζ	Third position coordinate with respect to the inertial frame	[m]
η	Second position coordinate with respect to the inertial frame	[m]
θ	Angular position of the secondary with respect to the inertial frame	[rad]
θ_{imp}	Impact angle	[rad]
λ	Longitude	[rad]
λ	Eigenvalue	[−]
λ	Continuation parameter	Various
μ	Gravitational parameter	[$\mathrm{m^3\,s^{-2}}$]
μ	Mass parameter of the CR3BP	[−]
ν	Stability index	[−]
ξ	First position coordinate with respect to the inertial frame	[m]
$\boldsymbol{\xi}$	Eigenvector	[−]
ρ	Density	[$\mathrm{kg\,m^{-3}}$]
τ	Bifurcations test function	[−]
ϕ	Latitude	[rad]
Φ	State transition matrix	Various
ψ	Continuation error-model function	Various
ω	Polyhedron per-facet factor	[−]
$\boldsymbol{\omega}$	Angular velocity vector	[$\mathrm{rad\,s^{-1}}$]

Mathematical Notation

$\boldsymbol{v}$	Vector (bold symbol)
v	Norm of vector $\boldsymbol{v}$
$\hat{\boldsymbol{v}}$	Unit vector with the direction of $\boldsymbol{v}$
d/dt	Time derivative with respect to the inertial frame
$\delta/\delta t$	Time derivative with respect to the synodic frame
$\dot{\square}$	Time derivative with respect to the synodic frame
$\square'$	Derivative with respect to the continuation arclength

Contents

1

Introduction

Phobos, a small and irregular moon of Mars, has for several decades been a topic of great interest in the scientific community. This interest stems from the lack of understanding about the origin and structure of this moon, with the answer to these unknowns being expected to give further insight into the processes involved in the formation of the Solar System and in the evolution of small moons and asteroids (Murchie et al., 2015). Additionally, Phobos has been proposed as an intermediate stop for refueling when transferring between the Earth and Mars, allowing a decrease of the launched propellant mass (Muscatello et al., 2012).

Several missions to explore Phobos have been designed. These include the failed Phobos-1 and Phobos-2, launched in 1988 (Sagdeev and Zakharov, 1989), and the also-failed Phobos-Grunt, launched in 2011 (Marov et al., 2004), all developed by Roscosmos. Furthermore, numerous other missions have been proposed by various space agencies, for example, the Phobos Sample Return mission, analyzed by ESA (Pickering, 2014). Finally, the Martian Moons eXploration (MMX) mission is currently being developed by JAXA, with its launch predicted for 2024. MMX, which is a sample return mission, will arrive in the Martian system in 2025. It will orbit Phobos using a series of increasingly lower two-dimensional quasi-satellite orbits (QSOs), as well as a mid-altitude three-dimensional QSO. To prepare for the landing operations, it will descend to an altitude of 40 m, deploy a small rover, and ascend back to orbit. Next, MMX will land on two different locations on Phobos, one on the near-Mars and one on the anti-Mars edge, and collect samples. To land, MMX will start from a two-dimensional low-altitude QSO and execute a ballistic descent to an altitude of 2 km; from there, it will land via a low-velocity vertical descent, with frequent maneuvers to counter the accelerations acting on the spacecraft (this last part is similar to what was executed by Hayabusa and Hayabusa2). After all operations are completed, the spacecraft will depart the Martian system in 2028, returning to Earth in 2029 (Kawakatsu et al., 2020; Kuramoto et al., 2022; Nakamura et al., 2021).

The dynamics in the proximity of Phobos correspond to those of the traditional three-body problem, with Mars and Phobos as main bodies, though highly perturbed by Phobos' non-spherical gravity field. As in the three-body problem, three main types of periodic orbits are possible: Lagrange point orbits, distant prograde orbits, and QSOs (Hénon, 1969). In the Phobos-Mars system, all distant prograde orbits intersect the surface, not allowing orbits of the former; instead, they can be used for transferring between different points on the surface (Liang et al., 2022). Meanwhile, Lagrange point orbits, though highly unstable, do allow orbiting Phobos, and have been studied with different dynamics models by Scheeres et al. (2019) and Zamaro (2015). QSOs, which constitute the core of most proposals to orbit Phobos, have been studied in much more detail; in fact, a very large part of the research on QSOs has been executed with respect to the Mars-Phobos system, likely motivated by the numerous mission proposals. Recent developments include the works by Chen et al. (2020) and Oshima and Yanao (2019), who studied families of two- and three-dimensional QSOs using continuation, by Baresi et al. (2021) and Scheeres et al. (2019), who analyzed two-dimensional QSOs through their computation as quasi-periodic tori, and by Canalias et al. (2017), who developed a semi-analytical method for generating quasi-periodic three-dimensional QSOs. Extensive analyses of QSOs in other systems have also been executed; in particular, in the Jupiter-Europa system, three-dimensional QSOs have been studied by Lara et al. (2007), using continuation-based techniques, and by Russell (2006), using a gridded search. Finally, it is worth mentioning Robin and Markellos (1980, 1983), as some of the first works to investigate three-dimensional QSOs.

The design of landing trajectories using invariant manifolds is a promising topic, as, starting from some operational orbit, they allow entering a landing trajectory with negligible propellant expenditure. These have been studied for various planet-moon and binary-asteroid systems using manifolds departing from Lagrange points or their orbits, often assuming small landers or science packages, as those allow more flexible landing requirements (Bury and McMahon, 2020; Ferrari and Lavagna, 2018; Herrera-Sucarrat et al., 2014; Zamaro, 2015). However, Lagrange point orbits might not be the most favorable science orbits, for example, due to their limited surface coverage or, in the case of Phobos, due to their high instability (Zamaro, 2015). Hence, the design of landing trajectories using invariant manifolds of more attractive orbits, namely QSOs, is more appealing.

While the usage of invariant manifolds of Lagrange point orbits for trajectory design has been often studied, this is not the case when it comes to QSOs. Using fast Lyapunov indicator maps, Scott and Spencer (2010) and Villac (2008) considered the design of transfers using invariant manifolds of unstable two-dimensional QSOs, respectively in the Sun-Earth and Jupiter-Europa systems. More recently, in a particularly relevant work, Oshima and Yanao (2019) focused on the usage of invariant manifolds of three-dimensional unstable QSOs for different applications, executing preliminary analyses on their usage for landing on Phobos, transferring from the Earth to a Moon-centered QSO, and for the computation of long-term stable QSOs. Finally, Pushparaj et al. (2021) used manifolds of QSOs to design transfers between two- and three-dimensional QSOs, in the Mars-Phobos system.

Mainly inspired by Oshima and Yanao (2019), this work focuses on the design of landing trajectories using the invariant manifolds of three-dimensional unstable QSOs. The focus is placed on three-dimensional QSOs (instead of two-dimensional ones) as these are preferred from a scientific point of view: they are better for mapping the surface, since they allow observing the poles, and are more suited for measuring the physical parameters of the moon, e.g. libration amplitude and high-order gravity coefficients (Chen et al., 2020; Nakamura et al., 2021). Compared to Oshima and Yanao, a significantly deeper analysis is executed here, placing particular emphasis on the relationship between the invariant manifolds and the families of three-dimensional QSOs from which they emanate. Additionally, the families of orbits considered here are different from the ones analyzed by Oshima and Yanao. As a case study, this analysis is executed with respect to the Mars-Phobos system.

Besides the analysis of the landing trajectories themselves, some more new elements are worth highlighting. These include an improved algorithm for the computation of the altitude with respect to polyhedra, a new algorithm for executing the shape homotopy between ellipsoids and polyhedra, the first analysis of the stability of invariant manifolds of QSOs, and the first analysis of the stability of invariant manifolds (of any type of orbit) in relation to the orbit from which they emanate.

1.1. Research Questions

As discussed, the main goal of this work is to analyze the usage of invariant manifolds emanating from three-dimensional QSOs as a tool for the design of landing trajectories. Therefore, the main research question this work aims to answer is formulated as:

- Can invariant manifolds of three-dimensional QSOs be used to design landing trajectories?

To structure the research, this main question has been divided into a series of sub-questions:

- How does the behavior of the invariant manifolds relate to the families of QSOs that create them?
- In the Mars-Phobos system, is it possible to land safely (i.e. complying with the landing requirements) using unpowered invariant manifolds?
- In case the answer to the previous question is negative: what is the minimum ΔV required to meet the landing requirements?
- Do invariant manifolds offer landing trajectories robust to uncertainties?

1.2. Landing Requirements

Throughout this report, the landing requirements specified in ESA's Phobos Sample Return mission proposal are considered. This was also proposed to be a sample return mission, with a similar mission profile to MMX. The following requirements are considered (Pickering, 2014):

- The vertical impact velocity shall be lower than $1.0\,\mathrm{m\,s^{-1}}$.
- The horizontal impact velocity shall be lower than $0.15\,\mathrm{m\,s^{-1}}$.
- The landing shall end with a free-fall phase starting at an altitude of at least 20 m.

The velocity requirements are also similar to the ones for MMX's rover, for which the impact velocity is required to be under 1.0 $\mathrm{m\,s^{-1}}$ (Nakamura et al., 2021).

1.3. Report Outline

This report is structured as follows. The used dynamics models are presented in Chapter 2. The applied numerical techniques are discussed in Chapter 3. In Chapter 4, the unpowered orbital motion in the vicinity of Phobos is studied, including the analysis of the three-dimensional QSOs and their invariant manifolds; focus is given to manifolds that allow reaching the surface of Phobos. In Chapter 5, the design of a braking maneuver is executed, to allow reaching the surface with low-enough velocities. An uncertainty analysis follows, in Chapter 6, considering both the cases of powered and unpowered landings. The verification of the used and developed software is presented in Chapter 7. The selection and tuning of several model properties, such as gravity model resolution, integrator, equations of motion, and optimizer, are discussed in Chapter 8. Conclusions and recommendations for future research are presented in Chapter 9. Finally, the equations used to execute the conversion between different frames are listed in Appendix A.

2

Dynamics Models

Phobos is the largest moon of Mars, orbiting the latter in a low-eccentricity (0.01511; Jacobson and Lainey, 2014), low-inclination (1.076°; Jacobson and Lainey, 2014) orbit. Phobos is tidally locked with Mars and is characterized by a heavily cratered surface, with an irregular shape similar to an ellipsoid. The dynamical environment in the proximity of Phobos (the region relevant for the analysis of landing trajectories) is dominated by the gravity induced by Phobos and Mars. The major perturbations are Phobos' irregular gravity field and orbital eccentricity, with the remaining perturbations having effects that are several orders of magnitude lower (Zamaro, 2015). For short propagations (7 days, which is longer than most propagations executed here), a model considering these perturbations is observed to very accurately reproduce a high-accuracy ephemerides model (Chen et al., 2020). Phobos' eccentricity induces a periodic perturbation, simply causing the different orbits to oscillate around the solutions obtained under the assumption of zero eccentricity, without changing its stability properties (Scheeres et al., 2019; Zamaro, 2015). As such, and considering that taking into account the eccentricity would make the analysis significantly more difficult, Phobos is here assumed to orbit Mars in a circular orbit. Thus, the only perturbation considered is Phobos' irregular gravity field. This model setup is common when analyzing the Mars-Phobos system, e.g. Chen et al. (2020), Pushparaj et al. (2021), and Scheeres et al. (2019).

In this section, a detailed description of the used dynamics models is presented. Phobos is considered to have an irregular gravity field, and to orbit Mars in a circular orbit, with Mars being modeled by a point mass. Phobos' gravity field is modeled using both an ellipsoid and a polyhedron; the importance of using two different models for Phobos is analyzed in Chapter 3. The used model parameters are summarized in Tab. 2.1.

Table 2.1: Physical parameters of the Mars-Phobos system. Data from Scheeres et al. (2019)[a], Andert et al. (2010)[b], Willner et al. (2014)[c], and Jacobson and Lainey (2014)[d].

Body	Quantity	Value
Mars	Gravitational parameter [$\text{km}^3\,\text{s}^{-2}$]	[a]4.2828372854×10^4
Phobos	Gravitational parameter [$\text{km}^3\,\text{s}^{-2}$]	[b]7.127×10^{-4}
	Largest semi-major axis of mean ellipsoid [km]	[c]13.03
	Intermediate semi-major axis of mean ellipsoid [km]	[c]11.40
	Smallest semi-major axis of mean ellipsoid [km]	[c]9.14
	Orbital semi-major axis [km]	[d]9375.0

2.1. Circular Restricted Three-Body Problem

The information presented in this section is based on Wakker (2015).

The motion of a small body in the vicinity of two large bodies (known as primaries) whose gravity has an important effect on the former's dynamics is often modeled using the circular restricted three-body problem (CR3BP). In the case of the Mars-Phobos system, the primaries are Mars and Phobos, the third body is the spacecraft. The more and less massive primaries (respectively, Mars and Phobos) are denoted, respectively, by primary and secondary.

The CR3BP corresponds to the three-body problem subject to three simplifying assumptions. Firstly, the mass of the primaries is assumed to be much larger than the mass of the third body; thus, the third body moves under the effect of the gravity of the primaries, but does not exert an acceleration on them. Consequently, the two primaries move in a two-body configuration, the two sharing the same orbital plane. Secondly, the two primaries move in circular orbits with respect to the barycenter of the system (also moving in circular orbits with respect to each other). Thirdly, the two primaries have point-mass gravity fields.

When analyzing the CR3BP, two reference frames are usually considered (Fig. 2.1). Firstly, a barycentric pseudo-inertial frame (ξ,η,ζ): the ζ-axis is parallel to the angular momentum of the primaries and perpendicular to their orbital plane; ξ and η are contained in the primaries' orbital plane. The second frame (x,y,z), known as the synodic frame, is also centered on the barycenter. The z-axis coincides with the ζ-axis, the x-axis goes from the primary P_1 to the secondary P_2, and the y-axis is defined to form a right-handed frame. The synodic frame rotates with constant angular velocity $\boldsymbol{\omega} = \dot{\theta}\hat{\boldsymbol{z}} = \dot{\theta}\hat{\boldsymbol{\zeta}}$ (the hats indicate unit vectors, the dot a time derivative), with (Wakker, 2015)

$$\dot{\theta} = \sqrt{\frac{G(m_1+m_2)}{\overline{P_1P_2}^3}} \tag{2.1}$$

where m_1 represents the mass of P_1, m_2 the mass of P_2 and $\overline{P_1P_2}$ the distance between P_1 and P_2. Notice that in the synodic frame P_1 and P_2 are fixed.

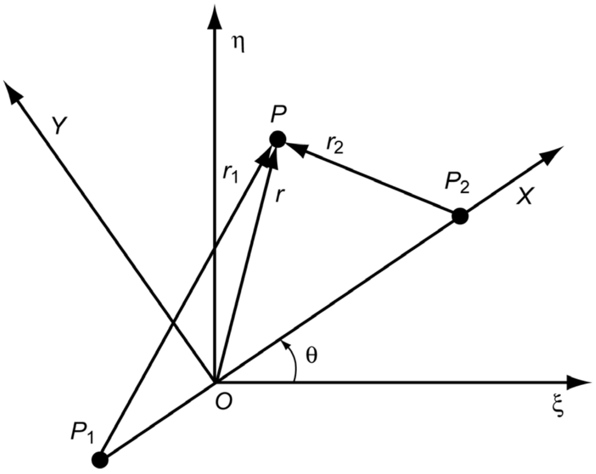

Figure 2.1: Pseudo-inertial (ξ,η,ζ) and synodic (x,y,z) frames used in the CR3BP. The ζ- and z-axes are not represented, being perpendicular to the figure. The primaries (P_1 and P_2) and the spacecraft (P) are represented (Wakker, 2015).

Consider the following notation: $\boldsymbol{r}$ represents the position of the spacecraft with respect to the barycenter O, $\boldsymbol{r_1}$ the position with respect to P_1 and $\boldsymbol{r_2}$ with respect to P_2. Additionally, assume d/dt and $\delta/\delta t$ to represent, respectively, the time derivative with respect to the inertial and synodic frame. The acceleration of the spacecraft with respect to the inertial frame is given by (Wakker, 2015)

$$\frac{d^2\boldsymbol{r}}{dt^2} = -\frac{Gm_1}{r_1^3}\boldsymbol{r_1} - \frac{Gm_2}{r_2^3}\boldsymbol{r_2} \tag{2.2}$$

and with respect to the synodic frame by (Wakker, 2015)

$$\frac{\delta^2\boldsymbol{r}}{\delta t^2} = -\frac{Gm_1}{r_1^3}\boldsymbol{r_1} - \frac{Gm_2}{r_2^3}\boldsymbol{r_2} - 2\boldsymbol{\omega}\times\frac{\delta\boldsymbol{r}}{\delta t} - \boldsymbol{\omega}\times(\boldsymbol{\omega}\times\boldsymbol{r}) \tag{2.3}$$

In the inertial frame, the acceleration depends on the position of the spacecraft and on the time, the latter because the position of the primaries (from which $\boldsymbol{r_1}$ and $\boldsymbol{r_2}$ start) depends explicitly on the time. Meanwhile, in the synodic frame, the acceleration depends solely on the state of the spacecraft, since $\boldsymbol{\omega}$ is constant and the position vectors $\boldsymbol{r}$, $\boldsymbol{r_1}$ and $\boldsymbol{r_2}$ all depart from fixed points (the primaries and barycenter are fixed in the synodic frame). Therefore, when expressed in a synodic frame, the system's dynamics are time-invariant (i.e. a spacecraft with a given state will always suffer the same acceleration, independently of the time). This property is essential for the applied techniques; as such, throughout this report, unless otherwise stated, all quantities are defined with respect to the synodic frame.

The equations of motion with respect to the synodic frame are usually simplified by expressing them in dimensionless coordinates, using the following units

$$\begin{aligned}&\text{Unit of mass:} \quad m_1 + m_2\\&\text{Unit of distance:} \quad \overline{P_1 P_2}\\&\text{Unit of time:} \quad 1/\dot{\theta}\end{aligned}$$

With this normalization, the equations depend on a single parameter known as the mass parameter μ, which is defined as

$$\mu = \frac{m_2}{m_1 + m_2} \tag{2.4}$$

All trajectories are here propagated using dimensionless coordinates (further details about the propagation set-up are presented later).

A final remark: the Hill three-body problem (H3BP), although not used, is sometimes referred to throughout this report. This model corresponds to a linearization of the CR3BP around the secondary.

2.1.1. Equilibrium Points

In the CR3BP, five equilibrium points exist, known as Lagrange points. By definition, the acceleration (with respect to the synodic frame) at these points is zero. All five Lagrange points are located in the xy-plane (Fig. 2.2). Three equilibrium points (L_1, L_2, and L_3) are located on the x-axis, thus being known as collinear Lagrange points. The other two Lagrange points (L_4 and L_5) are located in symmetric positions with respect to the x-axis, each of them forming an equilateral triangle with the primaries.

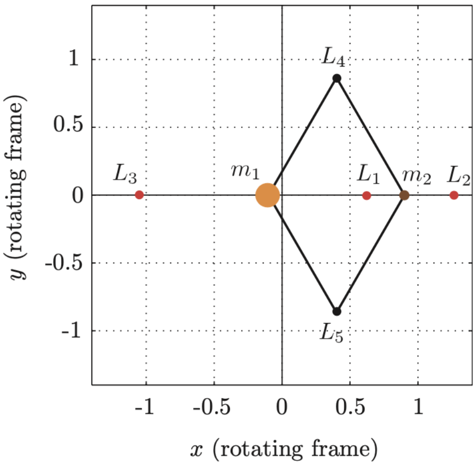

Figure 2.2: Lagrange points in the CR3BP, represented in the barycentric synodic frame (Koon et al., 2011).

The computation of the Lagrange points is executed by solving $\ddot{\boldsymbol{r}}(\boldsymbol{r}) = 0$. Here, this is done using Broyden's method (Press et al., 2007), which corresponds to a generalization of the secant root-finding method for higher dimensions. Broyden's method was selected because, similarly to the secant method, it does not require evaluating the Jacobian matrix of $\ddot{\boldsymbol{r}}(\boldsymbol{r})$. Contrary to what is often done, here no a priori knowledge about the Lagrange points' position (e.g. the fact that L_1 and L_2 are located on the x-axis) is used to simplify their computation. The initial guess for the position of the Lagrange points is generated according to Wakker (2015).

2.1.2. Integral of Motion

Since the CR3BP is time-invariant (with respect to the synodic frame), it admits one integral of motion, known as Jacobi's integral, given by (Wakker, 2015)

$$C = (x^2 + y^2) + 2\left(\frac{1-\mu}{r_1} + \frac{\mu}{r_2}\right) - (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) \tag{2.5}$$

The value of Jacobi's integral is determined by the initial state of the third body and represents the conservation of its energy. The first term of Jacobi's integral is associated with the centrifugal potential, the second with the gravitational potential, and the third with the kinetic energy.

2.1.3. Arbitrary Gravity Field

When modeling systems consisting of irregular bodies (e.g. Mars-Phobos, binary asteroids), it is common to use the CR3BP with primaries modeled by irregular gravity fields (as opposed to point-mass gravity fields). In this case, the two primaries are still considered to move in circular orbits with angular velocity given by Eq. 2.1; thus, the attraction between the primaries still corresponds to that of point masses. Only the motion of the third body is affected by the irregularity of the primaries' gravity.

Assuming that the irregular primaries have an axis of rotation parallel to their angular momentum and that they rotate synchronously with their orbital motion, the CR3BP with irregular primaries is also time-invariant (this is the case for Phobos, which as already mentioned, is tidally locked with Mars). Thus, the CR3BP with irregular primaries also has five equilibrium points (although slightly displaced from the Lagrange points in the CR3BP) and admits an integral of motion. The Jacobi's integral is given by (Scheeres, 2012)

$$C = (x^2 + y^2) + 2\left(U_{P_1} + U_{P_2}\right) - (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) \tag{2.6}$$

where U_{P_1} and U_{P_2} represent, respectively, the gravitational potential of the primary and of the secondary. If U_{P_1} and U_{P_2} are taken to be the gravitational potentials of point masses, the previous equation simplifies to the traditional Jacobi's integral (Eq. 2.5).

2.1.4. Symmetries

The CR3BP is characterized by a series of symmetries that will become relevant later on; their description is based on Russell (2006). Firstly, the equations of motion are invariant under the transformation $\{z \to -z\}$, which corresponds to the change in state

$$\boldsymbol{x}_1 = [x, y, z, \dot{x}, \dot{y}, \dot{z}]^T \to \boldsymbol{x}_2 = [x, y, -z, \dot{x}, \dot{y}, -\dot{z}]^T \tag{2.7}$$

Thus, the trajectories obtained by forward integrating $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are symmetric with respect to the xy-plane.

Secondly, the equations of motion are invariant under the transformation $\{t \to -t, y \to -y\}$, which corresponds to the change in state

$$\boldsymbol{x}_1 = [x, y, z, \dot{x}, \dot{y}, \dot{z}]^T \to \boldsymbol{x}_2 = [x, -y, z, -\dot{x}, \dot{y}, -\dot{z}]^T \tag{2.8}$$

As a result, the trajectory obtained by forward integrating $\boldsymbol{x}_1$ is symmetric, with respect to the xz-plane, to the trajectory obtained by backward integrating $\boldsymbol{x}_2$.

Finally, the equations of motion are invariant under the transformation $\{t \to -t, y \to -y, z \to -z\}$, which corresponds to the change in state

$$\boldsymbol{x}_1 = [x, y, z, \dot{x}, \dot{y}, \dot{z}]^T \to \boldsymbol{x}_2 = [x, -y, -z, -\dot{x}, \dot{y}, \dot{z}]^T \tag{2.9}$$

Therefore, the trajectory obtained by forward integrating $\boldsymbol{x}_1$ is symmetric, with respect to the x-axis (i.e. rotation of 180° around the x-axis), to the trajectory obtained by backward integrating $\boldsymbol{x}_2$.

During this report, these symmetries are used to characterize the different computed orbits. The symmetries of an orbit can be easily identified by projecting it into different planes. An xz-symmetric orbit is symmetric with respect to $y = 0$ in the xy-plane, and $y = 0$ in the yz-plane. An x-symmetric orbit is symmetric with respect to $y = 0$ in the xy-plane, $y = -z$ in the yz-plane, and $z = 0$ in the xz-plane. Do note that an orbit can display both types of symmetries.

2.2. Gravity Model

The models used to represent the gravity field of Phobos (ellipsoid and polyhedron) are described here. The gravity field of a given body is usually expressed by its potential U, defined in a body-fixed frame. To compute the gravitational acceleration, one evaluates the gradient of the potential, usually with respect to the body-fixed frame; if necessary, this acceleration is then converted to some other frame of interest (e.g. an inertial frame). In this case, since Phobos is considered to rotate synchronously with Mars, Phobos' body-fixed frame is always aligned with the synodic frame.

2.2.1. Spherical Harmonics Model

The most common way to represent the gravity field of an irregular body is by expanding the potential in a spherical-harmonics series, according to (Montenbruck and Gill, 2001)

$$U(r,\phi,\lambda) = \frac{\mu}{r}\sum_{n=0}^{\infty}\sum_{m=0}^{n}\left(\frac{R}{r}\right)^{n} P_{n,m}(\sin\phi)\left(C_{n,m}\cos(m\lambda) + S_{n,m}\sin(m\lambda)\right) \tag{2.10}$$

$$P_n(x) = \frac{1}{2^n n!}\frac{d^n}{dx^n}(x^2-1)^n \tag{2.11}$$

$$P_{n,m}(x) = (1-x^2)^{m/2}\frac{d^m P_n(x)}{dx^m} \tag{2.12}$$

This potential is expressed in a body-fixed frame using spherical coordinates, with radius r, latitude ϕ, and longitude λ. The function $P_n(x)$ is the Legendre polynomial of degree n, the function $P_{n,m}(x)$ is the associated Legendre polynomial of degree n and order m, $C_{n,m}$ and $S_{n,m}$ are the spherical harmonics coefficients. The Legendre functions and spherical harmonics coefficients are here presented in their unnormalized form. When calculating the gravity potential, the series is approximated by truncating it at a maximum degree $n_{\max}$ and maximum order $m_{\max}$.

Finally, R is the radius of the reference sphere. This sphere is known as the Brillouin sphere, which is usually defined to coincide with the sphere that circumscribes the body of interest (encompassing all of the mass of the body). The spherical-harmonics series is only guaranteed to converge outside the Brillouin sphere. Consequently, using spherical harmonics for $r \leq R$, in general, leads to (large) errors in the calculated acceleration. This can be problematic for bodies with non-spherical shapes (e.g. Phobos), as in that case there might be large regions between the surface of the body and the circumscribing sphere where the spherical-harmonics series does not converge.

2.2.2. Constant-Density Ellipsoid Model

A first gravity model that can be valid down to the surface (depending on the implementation) is a constant-density ellipsoid. Additionally, this model is symmetric, which is essential for some of the used techniques; considering that Phobos' shape is approximately ellipsoidal, this is the most-accurate symmetric gravity model available. Phobos' ellipsoid is characterized by the semi-major axes listed in Tab. 2.1. The largest semi-major axis is aligned with the Mars-Phobos direction, the intermediate one with Phobos' orbital velocity, and the minor one with Phobos' angular momentum (Pushparaj et al., 2021).

Consider the ellipsoid's semi-major axes to be represented by a, b, c, such that $a \geq b \geq c$. Assume a body-fixed frame centered in the center of mass of the ellipsoid, with the x-axis aligned with a, the y-axis aligned with b, and the z-axis aligned with c; therefore the shape of the ellipsoid is described by $(x/a)^2 + (y/b)^2 + (z/c)^2 = 1$. The mass of the ellipsoid is given by (Scheeres, 2012)

$$m = \frac{4\pi}{3}\rho abc \tag{2.13}$$

with ρ representing the density. Thus, its gravitational parameter is $\mu = Gm$.

The potential of such an ellipsoid can be described in closed form by (Scheeres, 2012)

$$U(\boldsymbol{r}) = \frac{-3\mu}{4}\int_{\lambda(\boldsymbol{r})}^{\infty}\frac{\phi(\boldsymbol{r},u)}{\Delta(u)}du \tag{2.14}$$

$$\text{with} \qquad \phi(\boldsymbol{r},u) = \frac{x^2}{a^2+u} + \frac{y^2}{b^2+u} + \frac{z^2}{c^2+u} - 1 \tag{2.15}$$

$$\Delta(u) = \sqrt{(a^2+u)(b^2+u)(c^2+u)} \tag{2.16}$$

with $\boldsymbol{r}$ expressed in the body frame. The function $\lambda(\boldsymbol{r})$ is found by solving $\phi(\boldsymbol{r},\lambda) = 0$, where λ is defined as the maximum real root of the equation. The gradient of the potential can be calculated through (Scheeres, 2012)

$$\frac{dU}{dx} = \frac{-3\mu x}{2}\int_{\lambda(\boldsymbol{r})}^{\infty}\frac{1}{(a^2+u)\Delta(u)}du \tag{2.17}$$

$$\frac{dU}{dy} = \frac{-3\mu y}{2}\int_{\lambda(\boldsymbol{r})}^{\infty}\frac{1}{(b^2+u)\Delta(u)}du \tag{2.18}$$

$$\frac{dU}{dz} = \frac{-3\mu z}{2}\int_{\lambda(\boldsymbol{r})}^{\infty}\frac{1}{(c^2+u)\Delta(u)}du \tag{2.19}$$

The potential and its gradient are well defined at any point exterior to the ellipsoid. However, computing the gradient requires solving a series of integrals, which are written in the form of Carlson's elliptic integrals; to efficiently and robustly solve these, one needs to implement adapted integration techniques (Press et al., 2007).

Alternatively, the potential of a constant density ellipsoid might be described by a spherical-harmonics expansion (Eq. 2.10) with coefficients given by (Boyce, 1997)

$$S_{n,m} = 0 \tag{2.20}$$

$$C_{n,m} = 0 \qquad \text{if } n \text{ odd or } m \text{ odd} \tag{2.21}$$

$$\begin{aligned} C_{2n,2m} = & \frac{3}{R^{2n}} \frac{n!(2n-2m)!}{2^{2m}(2n+3)(2n+1)!}(2-\delta_{0,m}) \times \\ & \sum_{i=0}^{\text{floor}(\frac{n-m}{2})} \frac{(a^2-b^2)^{m+2i}\left[c^2-1/2\left(a^2+b^2\right)\right]^{n-m-2i}}{16^i(n-m-2i)!(m+i)!i!} \qquad \text{for } n,m=0,1,2,\ldots \end{aligned} \tag{2.22}$$

where $\delta_{i,j}$ represents the Kronecker delta, which is defined as

$$\delta_{i,j} = \begin{cases} 0, & \text{if } i \neq j \\ 1, & \text{if } i = j \end{cases} \tag{2.23}$$

The reference radius is calculated through (Balmino, 1994)

$$R = \left(\frac{3}{1/a^2+1/b^2+1/c^2}\right)^{1/2} \tag{2.24}$$

Being defined through a spherical-harmonics expansion, the previous representation of the gravity field of a constant-density ellipsoid is only valid outside the Brillouin sphere. Nevertheless, it was opted to use it, as its implementation was readily available, but applying it solely for the computation of the periodic orbits (located outside the Brillouin sphere) from which the invariant manifolds are generated. An ellipsoidal spherical-harmonics model is used instead of a general spherical-harmonics model because, as mentioned, some of the used techniques require a symmetric gravity model.

The implementation of the CR3BP with ellipsoidal secondary is verified in Section 7.1.

2.2.3. Polyhedron Model

The propagation of landing trajectories requires a gravity model which is well-defined down to the surface of the central body, including inside the Brillouin sphere. Thus, alternatives to the spherical-harmonics formulation have to be considered.

The most commonly used way to represent the gravity of small bodies is through constant-density polyhedra, with the computation of the gravitational acceleration being based on the shape of the body. The body is assumed to have a constant density (often there is no information regarding its mass distribution) and its surface is approximated by a collection of polygons, usually triangles (Fig. 2.3). Since this gravity model is based on the shape of the body, it inherently allows determining whether a spacecraft is inside or outside the body (Werner and Scheeres, 1997), which is useful for detecting collisions. When compared with spherical harmonics, the polyhedron model has the disadvantage of being very computationally intensive (Tardivel, 2016). The computational demand is closely related to the resolution of the model used, a topic which is further analyzed in a following section.

A second alternative is the usage of mass concentrations (mascons), in which the mass distribution of the body is approximated by a collection of point masses. When modeling high-altitude orbits, a mascon model is less computationally demanding than a polyhedron model for similar accuracy (Tardivel, 2016). However, describing the gravity via discrete masses produces errors that become important at low altitudes; in order to achieve high accuracy near the surface, mascon models require a very large number of point masses, leading the computational advantage with respect to the polyhedron model to disappear (Tardivel, 2016). Additionally, a mascon model gives no information regarding whether the spacecraft is inside or outside the body (Werner and Scheeres, 1997).

A third alternative would be the usage of internal spherical harmonics (Takahashi et al., 2013). In this case, the potential is expanded in a spherical-harmonics series with respect to a point outside the body; the expansion is valid within a sphere centered on that point, extending down to the surface. This approach is very accurate and computationally efficient. However, it is not well suited for a global representation of the gravity

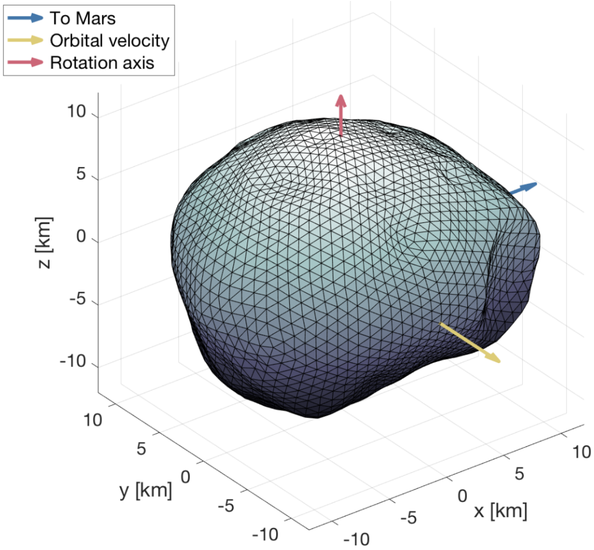

Figure 2.3: Polyhedron model of Phobos (Scheeres et al., 2019).

field, as that would require determining the internal spherical harmonics coefficients with respect to a series of different points distributed around the central body.

Based on the disadvantages associated with mascons and internal spherical harmonics, the usage of a polyhedron model is preferred here.

Mathematical Formulation

In a polyhedron model, the surface of the body is approximated by a series of triangular facets; each facet is bound by three edges and three vertices. The total number of facets (n_f), edges (n_e) and vertices (n_v) can be related through (Scheeres, 2012)

$$n_f = 2(n_v - 2) \tag{2.25}$$

$$n_e = 3(n_v - 2) \tag{2.26}$$

These relations are used for internal verification.

The gravitational potential, its derivative (i.e. the gravitational acceleration), and its Hessian (used for propagating the variational equations, see Subsection 3.1.2) are given, respectively, by (Werner and Scheeres, 1997)

$$U(\boldsymbol{r}) = \frac{1}{2} G\rho \left[\sum_{e \in \text{edges}} (\boldsymbol{r}_e^T E_e \boldsymbol{r}_e L_e) - \sum_{f \in \text{facets}} (\boldsymbol{r}_f^T F_f \boldsymbol{r}_f \omega_f) \right] \tag{2.27}$$

$$\nabla U(\boldsymbol{r}) = \frac{\partial U(\boldsymbol{r})}{\partial \boldsymbol{r}} = -G\rho \left[\sum_{e \in \text{edges}} (E_e \boldsymbol{r}_e L_e) - \sum_{f \in \text{facets}} (F_f \boldsymbol{r}_f \omega_f) \right] \tag{2.28}$$

$$\nabla\nabla U(\boldsymbol{r}) = \frac{\partial^2 U(\boldsymbol{r})}{\partial \boldsymbol{r}^2} = G\rho \left[\sum_{e \in \text{edges}} (E_e L_e) - \sum_{f \in \text{facets}} (F_f \omega_f) \right] \tag{2.29}$$

where G represents the universal gravitational constant and ρ the (constant) density of the polyhedron. The potential, acceleration, and Hessian are computed by summing the contributions produced by each individual edge and facet. The remaining terms in these equations are discussed in the following paragraphs.

The Laplacian of the potential can be calculated through (Werner and Scheeres, 1997)

$$\nabla^2 U(\boldsymbol{r}) = -G\rho \sum_{f \in \text{facets}} \omega_f \tag{2.30}$$

The term $-\nabla^2 U(\boldsymbol{r})/(G\rho)$ takes the value 0 for points outside the body, 4π for points inside it, 2π for points on a facet, and a solid angle in $]0, 4\pi[$ for points on edges or vertices (Park et al., 2010). Thus, the value of the

Laplacian can be used to efficiently detect whether a point is inside or outside the polyhedron and, as a result, to detect collisions with the body. When propagating a trajectory, the impact point is detected by applying a bisection method to detect the point with $-\nabla^2 U(\boldsymbol{r})/(G\rho) = 2\pi$, which would correspond to the exact impact point in case it happened at a facet (and would not exist in case it happened at an edge or vertex). In practice, however, the bisection method will never converge to the exact impact point, it will simply get gradually closer to the transition between 0 and 4π, eventually converging (within the specified tolerance) to a point at a very small distance from the surface, either inside or outside the body; this is valid for both impacts at facets, and at edges or vertices.

Continuing to the definition of the terms used in the computation of the potential: first, the terms associated with the potential of the facets are presented. Assume each facet to be described by three vertices i, j, and k with body-fixed positions $\boldsymbol{r}_i^o$, $\boldsymbol{r}_j^o$ and $\boldsymbol{r}_k^o$. Additionally, consider $\hat{\boldsymbol{n}}_f$ to represent the outward-pointing unit vector normal to the facet. The three vertices are assumed to be ordered according to the right-hand rule with respect to $\hat{\boldsymbol{n}}_f$, that is, ordered counterclockwise when seen from the outside of the polyhedron (Fig. 2.4). Given the positions of the vertices, the facet normal $\hat{\boldsymbol{n}}_f$ is calculated through

$$\hat{\boldsymbol{n}}_f = \frac{(\boldsymbol{r}_j^o - \boldsymbol{r}_i^o) \times (\boldsymbol{r}_k^o - \boldsymbol{r}_j^o)}{\left\| (\boldsymbol{r}_j^o - \boldsymbol{r}_i^o) \times (\boldsymbol{r}_k^o - \boldsymbol{r}_j^o) \right\|} \tag{2.31}$$

The facet dyad F_f is defined as

$$F_f = \hat{\boldsymbol{n}}_f \left(\hat{\boldsymbol{n}}_f\right)^T \tag{2.32}$$

and the per-facet factor ω_f as[1]

$$\omega_f = 2 \arctan_2 \left(\frac{\boldsymbol{r}_i \cdot (\boldsymbol{r}_j \times \boldsymbol{r}_k)}{r_i r_j r_k + r_i(\boldsymbol{r}_j \cdot \boldsymbol{r}_k) + r_j(\boldsymbol{r}_k \cdot \boldsymbol{r}_i) + r_k(\boldsymbol{r}_i \cdot \boldsymbol{r}_j)} \right) \tag{2.33}$$

where $\boldsymbol{r}_i$, $\boldsymbol{r}_j$ and $\boldsymbol{r}_k$ are the position vectors of the three vertices with respect to the field point (in this case, the spacecraft), which can be calculated through

$$\begin{aligned} \boldsymbol{r}_i &= \boldsymbol{r}_i^o - \boldsymbol{r} \\ \boldsymbol{r}_j &= \boldsymbol{r}_j^o - \boldsymbol{r} \\ \boldsymbol{r}_k &= \boldsymbol{r}_k^o - \boldsymbol{r} \end{aligned} \tag{2.34}$$

where $\boldsymbol{r}$ represents the position of the field point. Finally, $\boldsymbol{r}_f$ is a vector between the field point and any point on the facet f; in this case, it was selected $\boldsymbol{r}_f = \boldsymbol{r}_i$.

Note that, for a field point located on a facet, ω_f becomes singular, as both the numerator and denominator of Eq. 2.33 take the value zero. Nevertheless, it can be demonstrated that the limit of ω_f for a field point approaching the facet is zero (Werner, 2017). Thus, for a field point located on facet f, the corresponding per-facet factor is selected to be $\omega_f = 0$.

Now continuing to the definition of the terms associated with the potential of an edge: consider an edge to be defined by two vertices, with positions relative to the field point represented by $\boldsymbol{r}_i$ and $\boldsymbol{r}_j$. The variable $\boldsymbol{r}_e$ is a vector between the field point and any point on the edge e; in this case, it was selected as $\boldsymbol{r}_e = \boldsymbol{r}_i$. The edge vector is defined as $\boldsymbol{e}_{ij} = \boldsymbol{r}_i - \boldsymbol{r}_j$, and the per-edge factor as

$$L_e = \ln \left(\frac{r_i + r_j + e_{ij}}{r_i + r_j - e_{ij}} \right) \tag{2.35}$$

Observe that $(r_i + r_j - e_{ij})$ takes the value zero when the field point is located on an edge, meaning that L_e tends to infinity. Meanwhile, the term $E_e \boldsymbol{r}_e$ of Eq. 2.27 tends to zero, thus the potential has a singularity. For a field point approaching an edge, it can be shown that the limit of $E_e \boldsymbol{r}_e L_e$ is zero (Werner, 2017). Therefore, for a field point located on edge e, the corresponding per-edge factor is taken to be $L_e = 0$ (this approach is valid when computing the potential and its derivative, but not for the Hessian matrix).

Assume edge e to be shared by facets A and B, respectively, with facet normals $\hat{\boldsymbol{n}}_A$ and $\hat{\boldsymbol{n}}_B$. Furthermore, consider that for the vertices defining the edge (i.e. $\boldsymbol{r}_i$ and $\boldsymbol{r}_j$), the sequence $\boldsymbol{r}_i \rightarrow \boldsymbol{r}_j$ is followed counterclockwise with respect to facet A and clockwise with respect to facet B. The edge dyad is defined as

$$E_e = \hat{\boldsymbol{n}}_A \left(\hat{\boldsymbol{n}}_{ij}^A\right)^T + \hat{\boldsymbol{n}}_B \left(\hat{\boldsymbol{n}}_{ji}^B\right)^T \tag{2.36}$$

---

[1] The per-facet factor is computed using the arctangent function defined in $[-\pi, \pi]$.

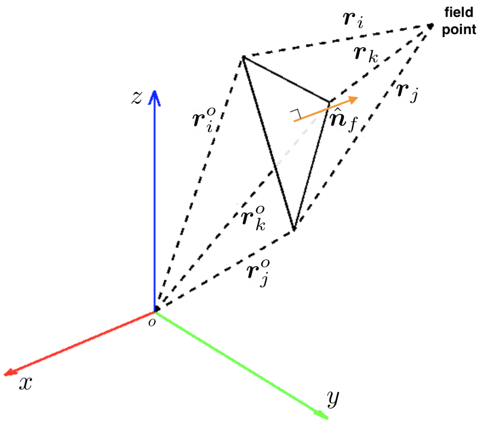

Figure 2.4: Representation of a facet, its outward-pointing normal, and the various position vectors used. Adapted from Hazra (2019).

where $\hat{\boldsymbol{n}}_{ij}^A$ and $\hat{\boldsymbol{n}}_{ji}^B$ represent the edge normals. The edge normal $\hat{\boldsymbol{n}}_e^f$ is defined as the unit vector normal to edge e and to $\hat{\boldsymbol{n}}_f$, and pointing away from the center of facet f (Fig. 2.5). These are computed according to

$$\hat{\boldsymbol{n}}_{ij}^A = \frac{(\boldsymbol{r}_j - \boldsymbol{r}_i) \times \hat{\boldsymbol{n}}_A}{\left\|(\boldsymbol{r}_j - \boldsymbol{r}_i) \times \hat{\boldsymbol{n}}_A\right\|} \tag{2.37}$$

$$\hat{\boldsymbol{n}}_{ji}^B = \frac{(\boldsymbol{r}_i - \boldsymbol{r}_j) \times \hat{\boldsymbol{n}}_B}{\left\|(\boldsymbol{r}_i - \boldsymbol{r}_j) \times \hat{\boldsymbol{n}}_B\right\|} \tag{2.38}$$

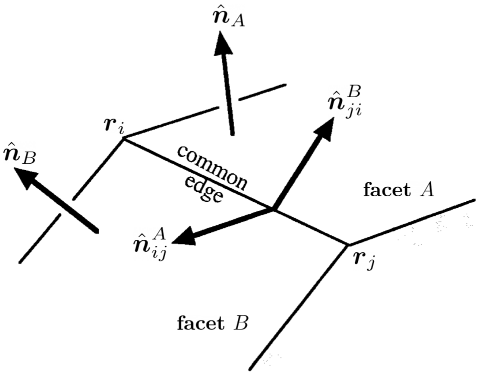

Figure 2.5: Representation of edge and facet normals. Adapted from Werner and Scheeres (1997).

Note that some of the terms used in the computation of the potential and acceleration depend solely on the geometry of the polyhedron, and not on the position of the field point. Thus, these terms (edge and facet dyads) are computed only once, before the start of the orbit propagation.

While the previous equations were presented in their original formulation, with the $G\rho$ term, here the gravitational parameter and volume are used instead, according to

$$G\rho = \frac{\mu}{V} \tag{2.39}$$

The gravitational parameter is a known quantity and the volume can be easily computed from the polyhedron. Once again, consider each facet to be described by three vertices with body-fixed positions $\boldsymbol{r}_i^o$, $\boldsymbol{r}_j^o$ and $\boldsymbol{r}_k^o$. As can be noticed (Fig. 2.4), each facet defines a tetrahedron, with the fourth vertex corresponding to the origin of the body-fixed reference frame. The volume of the tetrahedron associated with each facet can be calculated through (Dobrovolskis, 1996)

$$V_f = \frac{\boldsymbol{r}_i^o \cdot (\boldsymbol{r}_j^o \times \boldsymbol{r}_k^o)}{6} \tag{2.40}$$

Thus, the total volume of the polyhedron is given by (Dobrovolskis, 1996)

$$V = \sum_{f \,\in\, \text{facets}} V_f \tag{2.41}$$

The implementation of the CR3BP with polyhedral secondary is verified in Section 7.2, and that of the impact detection using the Laplacian of the polyhedron in Section 7.6.

Phobos Model, Achievable Accuracy and Resolution

Here, the shape model of Phobos developed by Willner et al. (2010) is used; the model is represented in Fig. 2.6, left. Two shape models of Phobos developed by the same authors are available: Willner et al. (2010) and Willner et al. (2014); however, only the former is directly available as a polyhedron model, through the 3D Asteroid Catalogue website (Frieger, n.d.). Considering that the two models are very similar and that the resolution of the polyhedron model will be reduced in the present work (discussed next), it was opted to use the older model.

Since, to date, only Phobos' $C_{2,0}$ and $C_{2,2}$ spherical harmonics coefficients have been measured (Yang et al., 2019), the polyhedron model effectively corresponds to the most accurate gravity model available. Nevertheless, the accuracy of this model is still limited by the assumption of constant density. Analyzing the asteroid 433 Eros, which has approximately homogeneous density and similar size to Phobos (433 Eros has a mean diameter of 16.84 km, Phobos of 22.16 km; NASA, n.d.), Miller et al. (2002) show that the assumption of constant density leads to errors in the computation of the surface acceleration in the order of 1%. There is no clear consensus regarding the internal structure of Phobos; Le Maistre et al. (2019) point to a homogeneous density distribution in the majority of the body, but increased density under the Stickney crater, while Guo et al. (2021) indicate an inhomogeneous density distribution, with the density decreasing with depth. Given this lack of consensus, one can hypothesize that the assumption of constant density would lead to errors in Phobos' surface acceleration of at least 1% (this error level would be associated with the case of approximately homogeneous density, similar to 433 Eros).

When using high-resolution polyhedron models, the gravity computation tends to be very computationally intensive. To tackle that issue, one can opt to reduce the model resolution (i.e. reduce the number of facets and vertices), according to the required model accuracy. Based on the 1% surface-acceleration accuracy threshold, a model with 1000 facets (Fig. 2.6, right) is selected (details regarding the selection of the number of facets are given in Section 8.1).

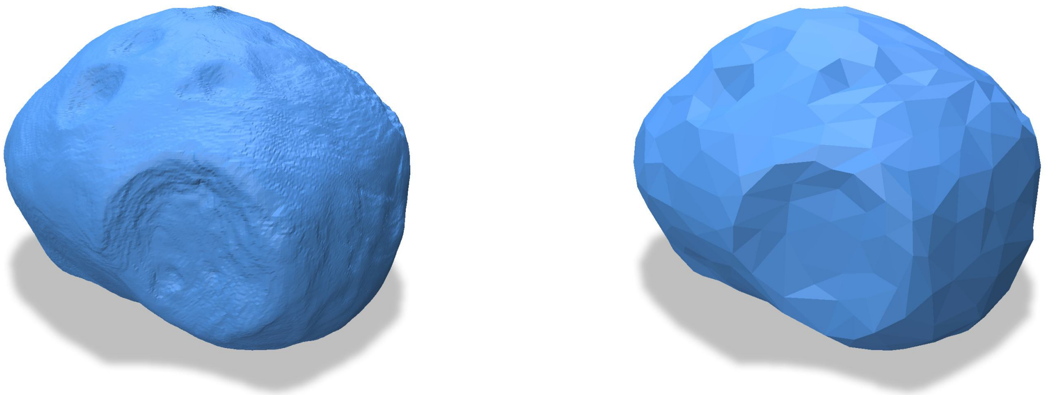

Figure 2.6: High-resolution polyhedron model of Phobos (developed by Willner et al., 2010) with 256000 facets (left), and reduced-resolution model with 1000 facets (right).

The reduction of the model's resolution is done using the Python library PyMeshLab (Muntoni and Cignoni, n.d.). As recommended by Van wal (2018), the model is simplified using a quadric-error based edge simplification method (Garland and Heckbert, 1997), which simplifies the model while trying to preserve the boundary of the body and the orientation of the surface normals, and allowing the position of the used vertices to be modified.

Finally, a note regarding the position of the polyhedron's center of mass (COM). The model developed by Willner et al. (2010) is provided with respect to a reference frame which is not centered on the COM (the difference between the two is in the order of hundreds of meters). Thus, when applying the model, its position is corrected such that the COM coincides with the origin of the body-fixed frame. This adjustment is also applied to correct the shift in the COM's position produced by the change in the resolution (although that shift is very minor, in the order of decimeters). Again, consider each facet to be described by the three vertices $\boldsymbol{r}_i^o$, $\boldsymbol{r}_j^o$, and $\boldsymbol{r}_k^o$, with each facet defining a tetrahedron. To compute the position of the COM, one first computes the COM of each tetrahedron, via (Dobrovolskis, 1996)

$$\boldsymbol{r}_f = \frac{\boldsymbol{r}_i^o + \boldsymbol{r}_j^o + \boldsymbol{r}_k^o}{4} \tag{2.42}$$

The position of the COM of the whole polyhedron is given by (Dobrovolskis, 1996)

$$\boldsymbol{r}_{COM} = \sum_{f \in \text{facets}} \frac{\Delta V_f \boldsymbol{r}_f}{V} \tag{2.43}$$

2.3. Shape Model

As mentioned, collisions with the surface are detected based on the value of the Laplacian of the polyhedron's gravity field. However, the Laplacian only gives information about whether a point is inside, outside, or on the surface of the polyhedron, not about the altitude with respect to the surface. To determine the altitude, the surface is taken to be the one defined by the polyhedron. The computation of the altitude is executed by determining the minimum distance to each of the elements that constitute the polyhedron (vertices, facets, and edges), following Van wal (2014).

Distance to a Vertex

Starting with the simplest case, the distance between a field point $\boldsymbol{r}$ and a vertex $\boldsymbol{r}_i^o$ is given by

$$d = \left\| \boldsymbol{r} - \boldsymbol{r}_i^o \right\| \tag{2.44}$$

Distance to a Facet

Consider a field point $\boldsymbol{r}$ and a facet with vertices $\boldsymbol{r}_i^o$, $\boldsymbol{r}_j^o$ and $\boldsymbol{r}_k^o$ and outward-pointing normal $\hat{\boldsymbol{n}}_f$. The perpendicular distance between the field point and the plane defined by the facet is

$$d = \left| (\boldsymbol{r} - \boldsymbol{r}_i^o) \cdot \hat{\boldsymbol{n}}_f \right| \tag{2.45}$$

This distance might be small (or zero) even if the field point is far from the facet. Thus, the computed distance is only considered valid if the field point is above/below the facet, meaning that the projection of the field point onto the plane defined by the facet should be contained by the facet. This projection, represented by $\boldsymbol{h}$, is computed via

$$\boldsymbol{h} = \boldsymbol{r} - d\hat{\boldsymbol{n}}_f \tag{2.46}$$

Now consider the vectors going from each vertex to the projected point, given by

$$\begin{aligned} \boldsymbol{u}_i &= \boldsymbol{h} - \boldsymbol{r}_i^o \\ \boldsymbol{u}_j &= \boldsymbol{h} - \boldsymbol{r}_j^o \\ \boldsymbol{u}_k &= \boldsymbol{h} - \boldsymbol{r}_k^o \end{aligned} \tag{2.47}$$

and the vectors defined using the cross product of these vectors

$$\begin{aligned} \boldsymbol{v}_{ki} &= \boldsymbol{u}_k \times \boldsymbol{u}_i \\ \boldsymbol{v}_{ij} &= \boldsymbol{u}_i \times \boldsymbol{u}_j \\ \boldsymbol{v}_{jk} &= \boldsymbol{u}_j \times \boldsymbol{u}_k \end{aligned} \tag{2.48}$$

This geometry is represented in Fig. 2.7. If the projected point $\boldsymbol{h}$ is located inside the facet, then each of the previous cross products is associated with a counterclockwise rotation, i.e. all $\boldsymbol{v}$ vectors point in the same direction; if $\boldsymbol{h}$ is located outside the facet, one of the $\boldsymbol{v}$ vectors will point in the opposite direction. Thus, $\boldsymbol{h}$ is located inside a facet if and only if $\boldsymbol{v}_{ki} \cdot \boldsymbol{v}_{ij} > 0$ and $\boldsymbol{v}_{ij} \cdot \boldsymbol{v}_{jk} > 0$. As mentioned, the distance to the facet d is only considered valid if $\boldsymbol{h}$ is inside the facet. Therefore, the minimum distance to the facets is determined by selecting, from the facets with valid d, the one that minimizes d.

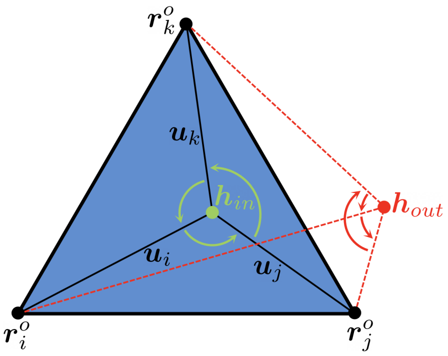

Figure 2.7: Representation of a single facet (blue triangle) and the projection of two field points onto the plane defined by it: one projection is contained by the facet ($\boldsymbol{h}_{in}$) and one is not ($\boldsymbol{h}_{out}$). Adapted from Van wal (2014).

Distance to an Edge

Finally, the distance to an edge is computed similarly to the distance to a facet, with the distance only being considered valid if the projection of the field point onto the line defined by the edge is contained by the edge (i.e. the projection is located between the two vertices defining the edge). Again, consider a field point $\boldsymbol{r}$ and an edge with vertices $\boldsymbol{r}_i^o$ and $\boldsymbol{r}_j^o$. The projection $\boldsymbol{h}$ of the field point onto the line defined by the edge is computed according to

$$\begin{aligned} \boldsymbol{h} &= \boldsymbol{r}_i^o + t(\boldsymbol{r}_j^o - \boldsymbol{r}_i^o) \\ t &= \frac{(\boldsymbol{r} - \boldsymbol{r}_i^o) \cdot (\boldsymbol{r}_j^o - \boldsymbol{r}_i^o)}{\left\| \boldsymbol{r}_j^o - \boldsymbol{r}_i^o \right\|^2} \end{aligned} \tag{2.49}$$

The projection is contained by the edge if $0 < t < 1$; the case $t = 0$ corresponds to $\boldsymbol{r} = \boldsymbol{r}_i^o$ and $t = 1$ to $\boldsymbol{r} = \boldsymbol{r}_j^o$. For the cases in which the projection is contained by the edge, the distance is computed according to

$$d = \|\boldsymbol{r} - \boldsymbol{h}\| \tag{2.50}$$

If the projection is not contained by the edge, then the distance is not considered valid.

Resulting Altitude and Simplifications

Finally, the altitude is determined by selecting the minimum value from the distances to the vertices and the valid distances to the edges and facets. Do note that the altitude corresponds to the unsigned distance to the surface; the information detailing whether a field point is inside or outside the polyhedron is obtained from the Laplacian.

Since the evaluation of the altitude requires assessing the distance to very large numbers (thousands in this case) of polyhedron elements, its application during the orbit propagation leads to a substantial increase in the computational time. Van wal (2014) deals with this issue by limiting the region with respect to which the altitude is computed using the latitude and longitude of the field point. However, here different simplifications are introduced. Firstly, when at high altitudes, the altitude is taken to be the minimum distance to the vertices of a simpler polyhedron (the distance to facets and edges is not computed). This method was applied for altitudes above 10.0 km using a polyhedron with 250 facets; maximum altitude errors in the order of 600.0 m were observed, which are acceptable considering the specified switchover altitude. For altitudes below 10.0 km the full polyhedron model (with 1000 facets) is used. The computational time is reduced by restricting the number of edges and facets with respect to which the distance is computed. Firstly, the distance to all vertices is computed, allowing the identification of the vertex closest to the field point. Next, the vertices connected to the closest vertex by some edge are selected. Finally, the distance is computed with respect to the edges and facets that contain either the closest vertex or one of the vertices connected to it (Fig. 2.8). This means that the altitude is computed with respect to the facets and edges connected to the closest vertex, but also with respect to a safety region around those. A safety region is used because computing the altitude just with respect to the facets and edges connected to the closest vertex does not ensure zero error in the computed altitude. In fact, even the usage of a safety region, no matter how many facets and edges it contains, does not guarantee zero error; this could occur, for example, if several facets with very acute angle(s) (i.e. very thin triangles) are located together. Fortunately, this type of extreme cases is not expected to be common in polyhedron models: the most-usual method for generating polyhedra (Delaunay triangulation) avoids these acute triangles, and so does the method used to reduce the model resolution. As such, the considered safety region should be sufficient in the majority of situations; for the tested cases, the described algorithm produced no error in the computation of the altitude.

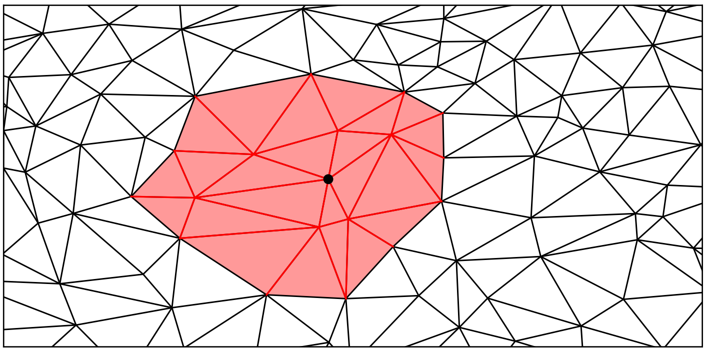

Figure 2.8: Representation of the geometry used when computing the altitude with respect to a polyhedron. The vertex closest to the field point is represented by a black dot; the facets and edges with respect to which the altitude is computed are colored in red.

Impact Angle

The same functions used to compute the altitude also allow determining the facet at which a given impact occurred, which in turn is used to compute the impact angle. Given a final velocity $\dot{\boldsymbol{r}}_f$ and a facet normal $\hat{\boldsymbol{n}}_f$, the impact angle θ_{imp} is computed according to

$$\cos\left(\frac{\pi}{2} - \theta_{imp}\right) = \frac{-\dot{\boldsymbol{r}}_f \cdot \hat{\boldsymbol{n}}_f}{\|\dot{\boldsymbol{r}}_f\| \, \|\hat{\boldsymbol{n}}_f\|} \tag{2.51}$$

The impact angle is measured with respect to the surface plane, taking the value 0° for impacts tangent to the surface and 90° for impacts normal to the surface (Fig. 2.9).

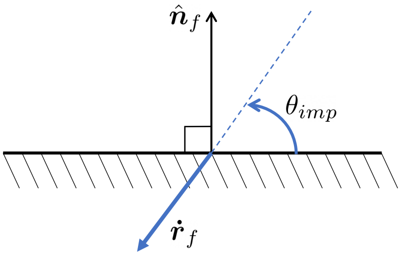

Figure 2.9: Representation of the impact angle. The horizontal black line corresponds to the surface.

2.4. Propagation Set-Up

Although the previous results were presented with respect to different frames, in practice two frames were used: a Mars-centered inertial frame and a Phobos-centered synodic frame.

The various orbits were propagated using Tudat, an open-source C++ astrodynamics toolbox developed at Delft University of Technology, with a Python interface named TudatPy (Dirkx et al., 2022; Tudat, n.d.). Since Tudat only supports propagations with respect to frames with inertial orientation, and considering that the dominant acceleration in the proximity of Phobos is produced by Mars, orbits were propagated with respect to an inertially oriented Mars-centered frame. All trajectories were propagated in dimensionless coordinates, using a variable-step Runge-Kutta-Dormand-Prince integrator, with eight-order integration steps and seventh-order steps for error control. Absolute and relative tolerances of 10^{-10} were used. The equations of motion were integrated using the Encke formulation. The selection of all these settings is presented in Section 8.2.

After integrating the trajectories, their analysis is executed with respect to a Phobos-centered synodic frame (as orbits in the proximity of Phobos are being analyzed). This frame can be obtained from the barycentric synodic frame through a translation along the synodic x-axis of $1 - \mu$, with μ corresponding to the mass parameter of the CR3BP (Wakker, 2015). Note that the Phobos-centered synodic frame is also time-invariant, thus all the results presented for the barycentric synodic frame are still valid (after applying the appropriate shift in the x-coordinate). The equations used for transforming between the Mars-centered inertial frame and the Phobos-centered synodic frame are presented in Appendix A.

Finally, it is worth mentioning that the implemented polyhedron shape and gravity models are already available for public use through Tudat.

3
Numerical Methods

Firstly, assuming a known periodic orbit, this chapter introduces the analysis of an orbit's stability and the computation of its invariant manifolds. Next, the methodology for computing families of periodic orbits is presented.

3.1. Isolated Solutions: Periodic Orbits

3.1.1. Periodic Orbits

The following discussion is based on Parker and Chua (1989). A system with state $\boldsymbol{x} \in \mathbb{R}^n$ described by

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}) \tag{3.1}$$

does not depend explicitly on the time t, being called an autonomous system. The CR3BP with synchronously rotating ellipsoidal or polyhedral secondaries (the systems considered here) are examples of autonomous systems when viewed from the synodic frame. A periodic orbit is a solution $\boldsymbol{x}(t)$ of the previous equation such that $\boldsymbol{x}(t) = \boldsymbol{x}(t+T)$ for all t. The minimum value $T > 0$ for which the previous relationship is valid is the period of the orbit. For autonomous systems, the period T can take any value.

A system

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}, t) \tag{3.2}$$

which depends explicitly on the time $t \in \mathbb{R}$ is non-autonomous. If there is a time period $T_f > 0$ such that $f(\boldsymbol{x}, t) = f(\boldsymbol{x}, t + T_f)$ for all t and $\boldsymbol{x}$, then the system is time periodic with period T_f. The CR3BP with ellipsoidal or polyhedral secondaries is an example of a non-autonomous system when viewed from the inertial frame (T_f corresponds to the primaries' period). In these systems, periodic orbits are only possible if $T = kT_f$, with k an integer. Thus, in general, the orbits analyzed here, which are periodic with respect to the synodic frame, are non-periodic when viewed from the inertial frame.

3.1.2. State Transition Matrix and Stability

The computation of periodic orbits and analysis of their stability is usually done using the state transition matrix (STM), which describes the motion in the proximity of a reference periodic orbit. Consider that the system $\dot{\boldsymbol{x}} = f(\boldsymbol{x})$ has a reference solution $\boldsymbol{x}_R(t)$ (not necessarily periodic), to which a perturbation $\delta\boldsymbol{x}(t_1)$ is applied at the initial time t_1. The evolution of the perturbed trajectory $\boldsymbol{x}(t)$ is given by

$$\boldsymbol{x}(t) = \boldsymbol{x}_R(t) + \delta\boldsymbol{x}(t) \tag{3.3}$$

which, considering the linearized equations of motion, can be related to the initial perturbation via

$$\delta\boldsymbol{x}(t) = \Phi(t, t_1)\delta\boldsymbol{x}(t_1) \tag{3.4}$$

where $\Phi(t, t_1)$ represents the STM. The STM is defined as (Montenbruck and Gill, 2001)

$$\Phi(t, t_1) = \frac{\partial\boldsymbol{x}(t)}{\partial\boldsymbol{x}(t_1)} \tag{3.5}$$

and is calculated through the first-order variational equations (Montenbruck and Gill, 2001)

$$\frac{d\Phi(t,t_1)}{dt} = \left.\frac{\partial f}{\partial \boldsymbol{x}}\right|_{\boldsymbol{x}_R(t)} \Phi(t,t_1) \tag{3.6}$$

$$\Phi(t_1,t_1) = I_{6\times 6} \tag{3.7}$$

where $\partial f / \partial \boldsymbol{x}$ represents the Jacobian of the dynamics. The variational equations are numerically integrated together with the state vector $\boldsymbol{x}_R$ (leading to a total integrated state with dimension 42).

There are some properties of the STM which are important to note. Firstly, the STM of a trajectory divided into arcs can be computed by multiplying the STM of each individual arc, according to (Bosanac, 2016)

$$\Phi(t_f,t_1) = \Phi(t_f,t_{f-1})\,\Phi(t_{f-1},t_{f-2})\,...\,\Phi(t_3,t_2)\,\Phi(t_2,t_1) \tag{3.8}$$

This property is used when applying the multiple-shooting differential corrector (Subsection 3.2.2). Secondly, the STM is symplectic, meaning that it has determinant det $\Phi(t_f,t_1) = 1$ and that its eigenvalues appear in reciprocal pairs (λ, $1/\lambda$).

The stability of periodic orbits is usually analyzed under the definition of Lyapunov stability, according to which an orbit is stable if a nearby solution remains close to the periodic orbit over time, and unstable if it diverges from the periodic solution (Bosanac, 2016). Thus, stability can either be associated with perturbed solutions that converge to the periodic orbit or that oscillate around it (bounded motion).

The stability of a periodic orbit can be assessed through the eigenvalues of its monodromy matrix M, which is defined as the STM evaluated over one orbital period, $M = \Phi(t_1+T,t_1)$. The monodromy matrix has two trivial eigenvalues $\lambda = 1$, associated with the periodicity of the orbit (Koon et al., 2011); the other four eigenvalues are associated with the linear dynamics in the vicinity of the orbit (Bosanac, 2016). Considering λ to represent an eigenvalue of the monodromy matrix and $\boldsymbol{\xi}$ the associated eigenvector, by definition, the following holds

$$\Phi(t_1+T,t_1)\boldsymbol{\xi} = \lambda\boldsymbol{\xi} \tag{3.9}$$

Comparing the previous equation with Eq. 3.4, it is possible to note that a perturbation in the direction of an eigenvector with eigenvalue $\|\lambda\| > 1$ will grow over time, corresponding to an unstable mode. A perturbation in the direction of an eigenvector with eigenvalue $\|\lambda\| \leq 1$ will remain bounded (either decreasing or oscillating), corresponding to a stable mode. A periodic orbit is considered stable if all its eigenvalues have $\|\lambda\| \leq 1$. Since the eigenvalues of the monodromy matrix appear in reciprocal pairs, an eigenvalue $\|\lambda\| < 1$ is always associated with an eigenvalue $\|\lambda\| > 1$; thus, stability is only possible if all eigenvalues have norm $\|\lambda\| = 1$ (i.e. eigenvalues located on the unit circle). The magnitude of the eigenvalues gives a measure of the rate at which a perturbation converges to or diverges from the periodic orbit, thus giving a measure of the latter's stability.

Since the eigenvalues of periodic orbits can take complex values, a straightforward assessment of the stability is often done using the stability indices, which, for periodic orbits, always take real values (complex values are associated with quasi-periodic motion; Seydel, 2010). Each stability index is defined as the sum of a pair of reciprocal eigenvalues (Broucke, 1969)

$$k = \lambda_1 + \lambda_2 \tag{3.10}$$

A pair of eigenvalues is associated with stable periodic motion if $\|k\| \leq 2$. The case $k = 2$ corresponds to an eigenvalue $\lambda = 1$ with multiplicity two, $k = -2$ corresponds to $\lambda = -1$ with multiplicity two, and $k \in]-2,2[$ corresponds to a pair of complex conjugate eigenvalues on the unit circle ($\lambda \neq -1, 1$). Unstable orbits are associated with $\|k\| > 2$, where the magnitude of $\|k\|$ gives a measure of the orbit's instability (similarly to the magnitude of unstable eigenvalues).

The stability index k (here used for the determination of bifurcations, as will be seen in the following sections) allows an easy assessment of the stability and evaluation of the structure of the eigenvalues. However, it does not allow an efficient comparison of the stability of orbits with different numbers of revolutions. Thus, when comparing the stability of different orbits, the normalized stability index introduced by Chen et al. (2020) is used. This stability index, defined as

$$\nu = \sqrt[n]{\max\|\lambda\|} \tag{3.11}$$

with n the number of revolutions of the periodic orbit, corresponds approximately to the stability of the orbit evaluated over a single revolution. An orbit is stable if $\nu = 1$ and unstable if $\nu > 1$.

As will be discussed, periodic orbits are organized in families, each family having different numbers of revolutions. The normalization of the stability index allows a fair comparison of the stability between orbits within the same family (with the same number of revolutions), but it is debatable whether the same applies to comparisons between orbits of different families, as the normalization likely has different effects depending on the propagation time. Additionally, since the orbits are propagated for different numbers of revolutions, the validity of assessing the stability based on the linearized dynamics (i.e. through the monodromy matrix) also varies between families. Nevertheless, the normalization does allow the different families to have similar stability indices, allowing the relative evolution of the stability along each family to be analyzed.

3.1.3. Invariant Manifolds

Invariant manifolds are structures that depart/approach unstable equilibrium points, periodic orbits, and quasi-periodic orbits, thus being useful to design transfers into/away from them. Since only manifolds of periodic orbits are used here, only these are presented below (although the principle governing the other two cases is analogous).

Invariant manifolds are regions of the state space such that a trajectory starting on one of them will remain on the manifold as time goes to infinity. Trajectories on a stable manifold approach the object that creates the manifold (e.g. periodic orbit) as time $t \to \infty$, while trajectories on an unstable manifold approach the object as time $t \to -\infty$. Thus, manifolds have an asymptotic nature, with the motion becoming increasingly slow as one approaches the creating object (Seydel, 2010).

Consider a system $\dot{\boldsymbol{x}} = f(\boldsymbol{x})$, for which an unstable periodic orbit is known. Its manifolds are usually obtained using the linearized equations of motion, and considering the orbit to be discretized into a series of nodes. Consider one of those nodes to be $\boldsymbol{x}^*$, and assume that node to have a monodromy matrix (i.e. the monodromy matrix of the periodic orbit with initial state $\boldsymbol{x}^*$) with eigenvalues λ_i and eigenvectors $\boldsymbol{\xi}_i$. From these eigenvectors, assume $\boldsymbol{\xi}_i,\ i = 1, ..., n_s$ to be stable (associated with $\|\lambda_i\| < 1$) and $\boldsymbol{\xi}_i,\ i = n_s + 1, ..., n_s + n_u$ to be unstable (associated with $\|\lambda_i\| > 1$). The set of stable and unstable eigenvectors spans, respectively, the following subspaces:

$$E_s = span\{\boldsymbol{\xi}_1, ..., \boldsymbol{\xi}_{n_s}\} \tag{3.12}$$

$$E_u = span\{\boldsymbol{\xi}_{n_s+1}, ..., \boldsymbol{\xi}_{n_s+n_u}\} \tag{3.13}$$

It is possible to show that at $\boldsymbol{x}^*$ the stable and unstable manifolds are tangent to, respectively, E_s and E_u (Seydel, 2010). Therefore, an initial state on the manifold can be obtained by perturbing $\boldsymbol{x}^*$ in the directions defined by these subspaces. Assuming that the monodromy matrix has a single pair of stable-unstable eigenvalues with eigenvectors $\boldsymbol{\xi}_s$ (stable) and $\boldsymbol{\xi}_u$ (unstable), the initial state on the two manifold is given by (Koon et al., 2011)

$$\boldsymbol{x}_s(t_1, \boldsymbol{x}^*) = \boldsymbol{x}^* \pm \varepsilon \boldsymbol{\xi}_s \tag{3.14}$$

$$\boldsymbol{x}_u(t_1, \boldsymbol{x}^*) = \boldsymbol{x}^* \pm \varepsilon \boldsymbol{\xi}_u \tag{3.15}$$

For each manifold, $\boldsymbol{x}^*$ can be perturbed in two directions (±), corresponding to the two directions each subspace takes at $\boldsymbol{x}^*$. Knowing the initial states on each manifold, the unstable and stable manifolds are then approximated by, respectively, forward and backward propagating the initial state (this only defines an approximation of the manifold as the initial state is obtained via linearization). The used perturbation ε should be selected to be small enough for the linear approximation to be valid, but not too small, otherwise, due to the asymptotic nature of manifolds, a very long propagation time would be needed to leave the neighborhood of $\boldsymbol{x}^*$, leading to the accumulation of numerical errors (Koon et al., 2011). When selecting the perturbation value, it is common to redefine it as a perturbation in position, such that (Vaquero, 2013)

$$\boldsymbol{x}(t_1, \boldsymbol{x}^*) = \boldsymbol{x}^* \pm \widetilde{\varepsilon} \frac{\boldsymbol{\xi}}{\sqrt{\boldsymbol{\xi}_x^2 + \boldsymbol{\xi}_y^2 + \boldsymbol{\xi}_z^2}} \tag{3.16}$$

The position perturbation $\widetilde{\varepsilon}$ corresponds to the distance between the point $\boldsymbol{x}^*$ and the initial state on the manifold. Koon et al. (2011) recommend a dimensionless position perturbation $\widetilde{\varepsilon} = 10^{-6}$.

The described procedure allows computing the manifold branch departing from a single node of a periodic orbit; the full manifold is computed by repeating the previous procedure for a series of $\boldsymbol{x}^*$ points distributed along the orbit. Since each initial state of the manifold is determined using the linearized equations of motion, it is not necessary to determine the monodromy matrix associated with each $\boldsymbol{x}^*$ to calculate its eigenvectors.

Instead, one can use the STM to propagate the eigenvectors associated with an initial state $\boldsymbol{x}(t_1)$ to the remaining states $\boldsymbol{x}(t)$ in the periodic orbit (Koon et al., 2011)

$$\boldsymbol{\xi}_s(\boldsymbol{x}(t)) = \Phi(t, t_1)\boldsymbol{\xi}_s(\boldsymbol{x}(t_1)) \tag{3.17}$$

$$\boldsymbol{\xi}_u(\boldsymbol{x}(t)) = \Phi(t, t_1)\boldsymbol{\xi}_u(\boldsymbol{x}(t_1)) \tag{3.18}$$

The multiplication by the STM modifies the norm of the eigenvectors, thus, they have to be re-normalized afterward.

So far it has been assumed that the system has a single pair of stable-unstable eigenvectors. If that were not the case, with $n_s > 1$ stable eigenvectors and $n_u > 1$ unstable eigenvectors, the initial guess for each manifold would be given by (Masdemont and Mondelo, 2004)

$$\boldsymbol{x}_s(\boldsymbol{x}^*) = \boldsymbol{x}^* \pm \sum_{i=1}^{n_s} \varepsilon_i \boldsymbol{\xi}_{s_i} \tag{3.19}$$

$$\boldsymbol{x}_u(\boldsymbol{x}^*) = \boldsymbol{x}^* \pm \sum_{i=1}^{n_u} \varepsilon_i \boldsymbol{\xi}_{u_i} \tag{3.20}$$

However, according to Masdemont and Mondelo (2004), for most values of ε_i, one converges to the trajectory that would be obtained by perturbing $\boldsymbol{x}^*$ just in the direction of the most unstable eigenvector (largest value of $\|\lambda_i\| > 1$) or the most stable eigenvector (smallest value of $\|\lambda_i\| < 1$); to obtain a more accurate estimate, one would need to use an expansion with higher-order terms (instead of just linear). Alternatively, one can use just the linear perturbation in the direction of the most stable and unstable eigenvectors, as although the approximation of the manifold is not as accurate, in terms of the dynamics it still represents the most dominant direction of departure from or arrival to the periodic orbit (Zamaro, 2015). This last approach is used here.

The implementation of the computation of the unstable invariant manifolds is verified in Section 7.5.

3.2. Families of Solutions and Their Continuation

In the CR3BP (possibly with an irregular secondary, as is the case in this work), periodic orbits are organized in families. Here the methodology for computing each of these families and switching between families is presented. Additionally, the transformation of orbits between a simplified dynamics model and a more complex one is discussed.

3.2.1. Implicitly Defined Curves

Consider a system of equations

$$\boldsymbol{F}(\boldsymbol{V}) = \boldsymbol{0}, \qquad \boldsymbol{F}: \mathbb{R}^{n+1} \to \mathbb{R}^n \tag{3.21}$$

This system can also be written as $\boldsymbol{F}(\boldsymbol{Y}, \lambda) = 0$, with $\boldsymbol{Y} \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$. The separation of $\boldsymbol{V}$ into $(\boldsymbol{Y}, \lambda)$ will only be represented when relevant to the discussion, otherwise, it will be omitted. Consider the Jacobian matrix of $\boldsymbol{F}$, which has dimensions $n \times (n+1)$, to be represented by

$$D_{\boldsymbol{F}}(\boldsymbol{V}) := \left.\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{V}}\right|_{\boldsymbol{V}} \tag{3.22}$$

Regarding the system $\boldsymbol{F}(\boldsymbol{V}) = \boldsymbol{0}$, assume that: (1) $\boldsymbol{F}$ has continuous derivatives, (2) there is a solution $\boldsymbol{V}^0$ such that $\boldsymbol{F}(\boldsymbol{V}^0) = \boldsymbol{0}$, and (3) the null space of the Jacobian $D_{\boldsymbol{F}}(\boldsymbol{V})$ (i.e. the space generated by the non-trivial linearly independent solutions $\boldsymbol{X}$ of $D_{\boldsymbol{F}}(\boldsymbol{V})\boldsymbol{X} = 0$) has dimension 1. Then, according to the implicit function theorem, the solutions of $\boldsymbol{F}(\boldsymbol{V}) = \boldsymbol{0}$ are organized in a one-dimensional curve of solutions (also known as branch or family of solutions[1]) that passes through $\boldsymbol{V}^0$ and that can be locally parameterized by λ, where λ can be chosen to be any of the elements of $\boldsymbol{V}$ (Allgower and Georg, 2003). Starting from a known solution of $\boldsymbol{F}(\boldsymbol{V}) = \boldsymbol{0}$, the process of obtaining the remaining points in the curve of solutions, which is known as continuation, requires two main tools. Firstly, one needs a predictor, which given a solution $\boldsymbol{V}^0$, generates the initial guess for another

---

[1] The terms curve and family of solutions are both used in this report. Curve of solutions is used when discussing the generic mathematical formulation of continuation, while family of solutions is used when applying it to the computation of orbits. This follows the generally used nomenclature.

solution; this initial guess, in general, will not be a solution to the equation. Secondly, one needs a corrector, which starting from the initial guess produces a solution.

When applied to the computation of periodic orbits, the system $\boldsymbol{F}(\boldsymbol{V}) = 0$ is constituted by a vector of constraints $\boldsymbol{F}(\boldsymbol{V})$, which have to be met in order produce a periodic orbit. The vector $\boldsymbol{V}$ is a vector of free variables which defines each periodic orbit through its period and its states at different times. To apply the implicit-function theorem to the CR3BP (with ellipsoidal or polyhedron secondary), some points are important to retain: firstly, $\boldsymbol{F}(\boldsymbol{V})$ should be formulated such that there is one more free variable than constraints; secondly, in general, the application of the implicit-function theorem is always valid, except at some singularities (known as bifurcations and discussed later).

3.2.2. Differential Correction

Newton's Method

The corrector step is discussed first, assuming an initial guess for a periodic orbit to be available. In general, the correction step is executed using Newton's method, presented following Bosanac (2016).

Consider $\boldsymbol{V}_d$ to represent a solution of $\boldsymbol{F}(\boldsymbol{V}) = 0$ (i.e. a vector of free variables defining a periodic orbit). Furthermore, consider $\boldsymbol{V}_0$ to represent the initial guess from which $\boldsymbol{V}_d$ is to be obtained. Expanding $\boldsymbol{F}(\boldsymbol{V})$ as a Taylor series around the initial guess $\boldsymbol{V}_0$ one obtains for the linear expression

$$\boldsymbol{F}(\boldsymbol{V}) \approx \boldsymbol{F}(\boldsymbol{V}_0) + \left.\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{V}}\right|_{\boldsymbol{V}_0} (\boldsymbol{V} - \boldsymbol{V}_0) = \boldsymbol{F}(\boldsymbol{V}_0) + D_{\boldsymbol{F}}(\boldsymbol{V}_0)(\boldsymbol{V} - \boldsymbol{V}_0) \tag{3.23}$$

Assuming the initial guess $\boldsymbol{V}_0$ to be close to $\boldsymbol{V}_d$, it is valid to write

$$\boldsymbol{F}(\boldsymbol{V}_d) = \boldsymbol{F}(\boldsymbol{V}_0) + D_{\boldsymbol{F}}(\boldsymbol{V}_0)(\boldsymbol{V}_d - \boldsymbol{V}_0) \Leftrightarrow \boldsymbol{F}(\boldsymbol{V}_0) + D_{\boldsymbol{F}}(\boldsymbol{V}_0)(\boldsymbol{V}_d - \boldsymbol{V}_0) = 0 \tag{3.24}$$

Since the previous equation is based on a linear approximation, solving it for $\boldsymbol{V}_d$ only allows a first-order correction. Therefore, to obtain convergence of $\boldsymbol{V}_0$ to $\boldsymbol{V}_d$, one needs to solve the equation iteratively, according to

$$\boldsymbol{F}(\boldsymbol{V}_i) + D_{\boldsymbol{F}}(\boldsymbol{V}_i)(\boldsymbol{V}_{i+1} - \boldsymbol{V}_i) = 0 \tag{3.25}$$

where i denotes the solution obtained during the i^{th} iteration and $i+1$ the new solution. If the Jacobian matrix $D_{\boldsymbol{F}}(\boldsymbol{V}_i)$ is square (i.e. the number of constraints and free variables is the same) and invertible, the system has a unique solution, which can be obtained via

$$\boldsymbol{V}_{i+1} = \boldsymbol{V}_i - (D_{\boldsymbol{F}}(\boldsymbol{V}_i))^{-1}\boldsymbol{F}(\boldsymbol{V}_i) \tag{3.26}$$

If the Jacobian $D_{\boldsymbol{F}}(\boldsymbol{V}_i)$ is square yet not invertible, or if it is not square, the Penrose-Moore pseudo-inverse is used instead of the inverse to solve the equation. In this work, as the number of constraints is smaller than the number of free variables, the right pseudo-inverse would be used

$$\boldsymbol{V}_{i+1} = \boldsymbol{V}_i - D_{\boldsymbol{F}}(\boldsymbol{V}_i)^T \left(D_{\boldsymbol{F}}(\boldsymbol{V}_i) D_{\boldsymbol{F}}(\boldsymbol{V}_i)^T\right)^{-1} \boldsymbol{F}(\boldsymbol{V}_i) \tag{3.27}$$

While analytically the previous equations are used to obtain $\boldsymbol{V}_{i+1}$, this is not what is done in practice. It is generally advised against inverting matrices numerically to solve systems of equations, as doing it is computationally less efficient and induces larger numerical errors (Cheney and Kincaid, 2008). The larger numerical errors become especially important when inverting ill-conditioned matrices, for which the inverse is numerically unstable. This is the case for the Jacobian $D_{\boldsymbol{F}}(\boldsymbol{V}_i)$ of periodic orbits near singularities of the solutions curve (known as bifurcations, which are discussed in Subsection 3.2.6), as in the singularity itself the Jacobian matrix in non-invertible (Allgower and Georg, 2003). Instead, the problem is solved by iteratively applying

$$\boldsymbol{V}_{i+1} = \boldsymbol{V}_i - (D_{\boldsymbol{F}}(\boldsymbol{V}_i))^{+}\boldsymbol{F}(\boldsymbol{V}_i) \tag{3.28}$$

The superscript $^+$ denotes the pseudo-inverse; it was determined using Python's NumPy library, which computes the pseudo-inverse of the matrix using its singular-value decomposition.

One of the advantages of using Newton's method is that it converges very fast to the solution; for a good enough initial guess, the method is known to converge quadratically (Cheney and Kincaid, 2008), thus

$$\frac{\|\boldsymbol{V}_d - \boldsymbol{V}_{i+1}\|}{\|\boldsymbol{V}_d - \boldsymbol{V}_i\|^2} \leq c \tag{3.29}$$

where c is a constant that depends on the distance between the initial guess $\boldsymbol{V}_0$ and the solution $\boldsymbol{V}_d$.

Single Shooting

The process of applying Newton's method to correct an initial guess for a periodic orbit is generally known as shooting. In the case of single shooting, the equations of motion and variational equations are integrated from an initial state $\boldsymbol{x}(t_1)$, over a time interval Δt (corresponding to the initial guess for the orbit's period). From the integration, one obtains a final state $\boldsymbol{x}(t_1+\Delta t) \neq \boldsymbol{x}(t_1)$. The single-shooting method, named as such because the orbit is propagated over a single arc, consists of iteratively applying Newton's method to correct the initial state and the integration time, eventually resulting in a periodic orbit (Fig. 3.1, left).

For single shooting, as free variables one uses the initial state and the integration time:

$$\boldsymbol{V} = [\boldsymbol{x}(t_1), \Delta t]^T \tag{3.30}$$

and defines $\boldsymbol{F}$ so as to ensure continuity between the initial and final state. When applying single shooting in the CR3BP, it is common to reduce the number of free variables by taking advantage of the symmetries inherent to the CR3BP (e.g. symmetry with respect to the xz-plane). However, those symmetries disappear when using an asymmetrical gravity model (e.g. polyhedron model).

Single shooting has the advantage of being very simple to implement, however, it is often sensitive to the quality of the initial guess, making its convergence problematic for unstable regions (Stoer and Bulirsch, 2002). This is particularly relevant here, as the computation of unstable invariant manifolds requires identifying unstable periodic orbits.

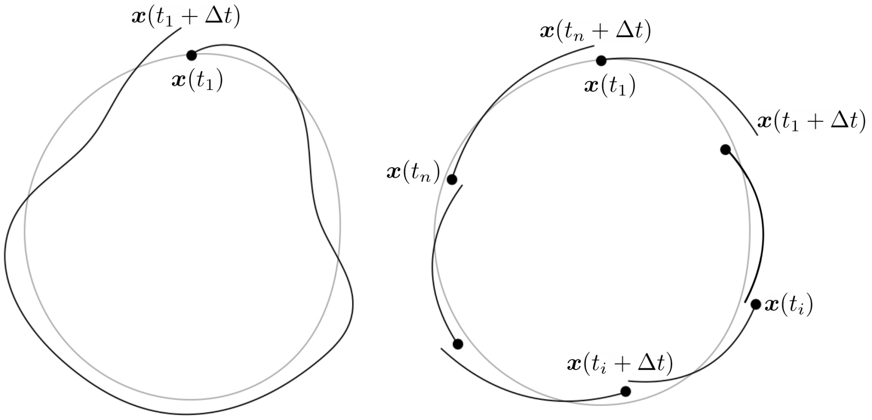

Figure 3.1: Representation of single shooting (left) and multiple shooting (right). The black curves correspond to the initial guess and the gray curves to the converged periodic orbit. Adapted from Baresi (2017).

Multiple Shooting

A second option for the application of the shooting method is through multiple shooting, where instead of using a single arc to describe the orbit multiple arcs are used (Fig. 3.1, right). Since it requires each arc to be propagated over a shorter amount of time, multiple shooting reduces the error between the initial-guess orbit and the periodic one; furthermore, it reduces the effect of the nonlinearities over each arc, meaning that the linearization made when deriving Newton's method is more accurate (Stoer and Bulirsch, 2002; Vaquero, 2013).

Since multiple arcs are used, the free variables vector consists of the initial state of each arc and integration time of the arcs:

$$\boldsymbol{V} = [\boldsymbol{x}(t_1), \boldsymbol{x}(t_2), ..., \boldsymbol{x}(t_n), \Delta t]^T \tag{3.31}$$

where $\boldsymbol{x}(t_i)$ represents the initial state of the i^{th} arc. All arcs are considered here to have the same integration time Δt, although one could also use different integration times for each arc.

Having defined $\boldsymbol{V}$, it is now necessary to construct $\boldsymbol{F}$. Firstly, one needs to use constraints to ensure continuity between the various arcs, which can be defined as

$$\begin{aligned} \boldsymbol{x}(t_i+\Delta t) &= \boldsymbol{x}(t_{i+1}), \quad i = 1, ..., n-1 \\ \boldsymbol{x}(t_n+\Delta t) &= \boldsymbol{x}(t_1) \end{aligned} \tag{3.32}$$

where $\boldsymbol{x}(t_i)$ and $\boldsymbol{x}(t_i+\Delta t)$ represent, respectively, the initial and final state of the i^{th} arc. Secondly, one needs a phase constraint. Since a periodic orbit can be described by infinite solutions $\boldsymbol{V}$ (with initial states in different positions of the orbit), the phase constraint serves to fix the initial state of the orbit, leading each periodic orbit to be described by a single solution $\boldsymbol{V}$. Phase constraints are scalar, having dimension 1.

Constructing $\boldsymbol{F}$ with the mentioned constraints and free variables, one would find $\boldsymbol{F}(\boldsymbol{V}) : \mathbb{R}^{6n+1} \rightarrow \mathbb{R}^{6n+1}$, with $6n+1$ free variables and $6n+1$ equations ($6n$ continuity constraints and 1 phase constraint). As can be observed, these dimensions are not compatible with the conditions necessary for the application of the implicit function theorem (Subsection 3.2.1), which requires the number of free variables to be equal to the number of equations plus 1. Two solutions for this problem are generally used.

The first option is to increase the number of free variables, by including an additional parameter λ, known as unfolding parameter. This requires reformulating the equations of motion to include λ, usually done via (Doedel et al., 2003)

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}) + \lambda \nabla C \tag{3.33}$$

where ∇C represents the gradient of the conserved quantity, in this case, the Jacobi integral. Upon convergence to a periodic orbit, the unfolding parameter would take the value $\lambda = 0$, and consequently, the equations of motion would correspond to their original formulation. The formulation with the unfolding parameter is not used here, as its inclusion in the equations of motion (propagated with Tudat) would not be straightforward.

A second option to solve the issue of the number of equations and free variables is to use the conserved quantity (Jacobi integral) to reduce the number of constraints, and to select a Poincaré surface of section to define the phase constraint (Doedel et al., 2003). In practice, this can be done by fixing one of the elements of the initial state of the first arc and by using continuity constraints solely between five elements of the final and initial arc; the conservation of the Jacobi integral throughout the continuous orbit should ensure continuity in the non-constrained element. Note, however, that the selection of which element to fix and which not to constrain is not innocuous. According to Eqs. 2.5 and 2.6, given a value of the Jacobi integral and of five state elements, there are multiple solutions for the sixth state element; therefore, continuity of the Jacobi integral and of five state elements effectively does not guarantee that the orbit is continuous. Thus, the continuity constraints and phase constraint should be chosen in parallel and taking into account the type of orbit being searched for, in order to decrease the likelihood of the differential corrector converging to a non-continuous solution. Since two- and three-dimensional QSOs always cross the xz-plane, the following, respectively, phase and final-arc continuity constraints are used

$$y(t_1) = 0 \tag{3.34}$$

$$\boldsymbol{q}(t_n + \Delta t) = \boldsymbol{q}(t_1) \qquad \text{with} \qquad \boldsymbol{q} = [x, y, z, \dot{x}, \dot{z}]^T \tag{3.35}$$

Although these conditions increase the likelihood of converging to a $\dot{y}(t_n + \Delta t)$ value such that the orbit is continuous, this is not guaranteed (e.g. for a poor initial guess). As such, even though $\dot{y}(t_n + \Delta t)$ is not constrained, after the convergence of Newton's method, it is checked whether the transition between $\dot{y}(t_n + \Delta t)$ and $\dot{y}(t_1)$ is indeed continuous.

The system of equations that is solved in order to determine periodic orbits can be summarized as

$$\boldsymbol{F}(\boldsymbol{V}) = \begin{bmatrix} \boldsymbol{x}(t_1 + \Delta t) - \boldsymbol{x}(t_2) \\ \boldsymbol{x}(t_2 + \Delta t) - \boldsymbol{x}(t_3) \\ \vdots \\ \boldsymbol{x}(t_{n-1} + \Delta t) - \boldsymbol{x}(t_n) \\ \boldsymbol{q}(t_n + \Delta t) - \boldsymbol{q}(t_1) \\ y(t_1) \end{bmatrix} = 0 \tag{3.36}$$

The Jacobian matrix of the system, with dimensions $6n \times (6n+1)$, can be easily calculated via

$$D_{\boldsymbol{F}}(\boldsymbol{V}) = \begin{bmatrix} \Phi(t_1 + \Delta t, t_1) & -I_{6\times6} & 0_{6\times6} & \cdots & 0_{6\times6} & \dot{\boldsymbol{x}}(t_1 + \Delta t) \\ 0_{6\times6} & \Phi(t_2 + \Delta t, t_2) & -I_{6\times6} & \cdots & 0_{6\times6} & \dot{\boldsymbol{x}}(t_2 + \Delta t) \\ 0_{6\times6} & 0_{6\times6} & \Phi(t_3 + \Delta t, t_3) & \cdots & 0_{6\times6} & \dot{\boldsymbol{x}}(t_3 + \Delta t) \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ -H_{5\times6} & 0_{5\times6} & 0_{5\times6} & \cdots & \dfrac{\partial \boldsymbol{q}(t_n + \Delta t, t_n)}{\partial \boldsymbol{x}(t_n)} & \dot{\boldsymbol{q}}(t_n + \Delta t) \\ [0,1,0,0,0,0] & 0_{1\times6} & 0_{1\times6} & \cdots & 0_{1\times6} & 0 \end{bmatrix} \tag{3.37}$$

where I corresponds to the identity matrix, 0 to the null matrix and $\Phi(t_j+\Delta t, t_j)$ to the STM between $\boldsymbol{x}(t_j)$ and $\boldsymbol{x}(t_j+\Delta t)$. The H matrix is given by

$$H_{5\times6} = \frac{\partial \boldsymbol{q}(t_n)}{\partial \boldsymbol{x}(t_n)} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{3.38}$$

3.2.3. Natural Parameter Continuation

The simplest way to compute the curve of solutions is via natural parameter continuation, which is based on the fact that the curve can be locally parameterized by λ (usually taken to be some kind of physical parameter, e.g. Jacobi integral or one of the state elements). Given a known solution $(\boldsymbol{Y}^j, \lambda^j)$, corresponding to the j^{th} solution along the solutions curve, one uses as initial guess for the $(j+1)^{th}$ solution $(\boldsymbol{Y}_0^{j+1}, \lambda_0^{j+1}) = (\boldsymbol{Y}^j, \lambda^j + \Delta\lambda)$, with $\Delta\lambda$ sufficiently small. Thus, the selected corrector is used to solve the extended system (Dichmann et al., 2003)

$$\boldsymbol{F}(\boldsymbol{Y}^{j+1}, \lambda^{j+1}) = \boldsymbol{0} \tag{3.39a}$$

$$\lambda^{j+1} = \lambda^j + \Delta\lambda \tag{3.39b}$$

This formulation of natural parameter continuation has the advantage of being very simple to implement; however, the generation of the initial guess does not take into account the shape of the solutions curve, which can make the convergence of the corrector more difficult, potentially leading it to diverge. A more robust type of natural parameter continuation generates the initial guess $(\boldsymbol{Y}_0^{j+1}, \lambda_0^{j+1})$ using the tangent to the curve of solutions at the known solution $(\boldsymbol{Y}^j, \lambda^j)$, according to (Dichmann et al., 2003; Seydel, 2010)

$$\begin{aligned} \lambda_0^{j+1} &= \lambda^j + \Delta\lambda \\ \boldsymbol{Y}_0^{j+1} &= \boldsymbol{Y}^j + \Delta\lambda \left.\frac{\partial \boldsymbol{Y}}{\partial \lambda}\right|_{\boldsymbol{Y}^j,\lambda^j} \qquad \text{with} \quad \frac{\partial \boldsymbol{Y}}{\partial \lambda} = -\left(\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{Y}}\right)^{-1} \frac{\partial \boldsymbol{F}}{\partial \boldsymbol{\lambda}} \end{aligned} \tag{3.40}$$

The equation for $\partial \boldsymbol{Y}/\partial\lambda$ is derived by taking the derivative of $\boldsymbol{F}(\boldsymbol{Y},\lambda) = \boldsymbol{0}$ with respect to λ. After generating the initial guess, the solution is once again found by solving Eq. 3.39. Although this method for generating the initial guess is more robust, it still fails in some situations, due to the singularity of $\partial \boldsymbol{Y}/\partial\lambda$.

3.2.4. Pseudo-Arclength Continuation

To overcome the limitations of natural parameter continuation, the more robust pseudo-arclength continuation is often used. This method considers the arclength s of the solutions curve (i.e. the length along the curve) to be the continuation parameter. Similarly to the second type of natural parameter continuation, the initial guess $\boldsymbol{V}_0^{j+1}$ is generated using the tangent to the solutions curve, according to (Doedel et al., 2003)

$$\boldsymbol{V}_0^{j+1} = \boldsymbol{V}^j + \Delta s \boldsymbol{V}^{j'} \tag{3.41}$$

where $\boldsymbol{V}^j$ represents a previously known solution and Δs the continuation step size (analogous to the $\Delta\lambda$ step used in natural parameter continuation). The superscript $'$ represents the derivative with respect to s, therefore $\boldsymbol{V}^{j'}$ corresponds to the vector tangent to the curve of solutions at $\boldsymbol{V}^j$. To find $\boldsymbol{V}^{j'}$, one takes the derivative of $\boldsymbol{F}(\boldsymbol{V}) = \boldsymbol{0}$ with respect to s, producing

$$D_{\boldsymbol{F}}(\boldsymbol{V}^j)\boldsymbol{V}^{j'} = \boldsymbol{0} \tag{3.42}$$

Thus, $\boldsymbol{V}^{j'}$ is defined in the null space of $D_{\boldsymbol{F}}(\boldsymbol{V}^j)$. As mentioned in Subsection 3.2.1, the null space of $D_{\boldsymbol{F}}(\boldsymbol{V})$ has dimension 1, meaning that there is a single (besides a scaling constant) vector $\boldsymbol{V}^{j'}$. To uniquely select the tangent $\boldsymbol{V}^{j'}$, it is assumed to be the vector that satisfies (Doedel et al., 1991)

$$\begin{aligned} &\left\|\boldsymbol{V}^{j'}\right\| = 1 \\ &\left(\boldsymbol{V}^{j-1'}\right)^T \boldsymbol{V}^{j'} > 0 \end{aligned} \tag{3.43}$$

This definition allows selecting each tangent vector $\boldsymbol{V}^{j'}$ such that the solutions curve is followed with a constant orientation (i.e. without ever going back). For the first element of each curve, the tangent vector was selected to satisfy

$$\det \begin{bmatrix} D_{\boldsymbol{F}}\left(\boldsymbol{V}^{0}\right) \\ \left(\boldsymbol{V}^{0'}\right)^{T} \end{bmatrix} > 0 \tag{3.44}$$

which effectively defines a “positive” direction of following the curve. The motivation for choosing the direction based on the determinant of the presented matrix should become clear in Subsection 3.2.6; nevertheless, the selection of which is the positive and which is the negative direction is completely arbitrary. Note that the curve is followed in the positive direction when using a continuation step $\Delta s > 0$, and in the negative direction when using $\Delta s < 0$.

Starting from the presented initial guess, one finds the next solution by solving the extended system (Doedel et al., 2003)

$$\boldsymbol{G}\left(\boldsymbol{V}^{j+1}\right) := \begin{bmatrix} \boldsymbol{F}\left(\boldsymbol{V}^{j+1}\right) \\ \left(\boldsymbol{V}^{j+1} - \boldsymbol{V}^{j}\right)^{T} \boldsymbol{V}^{j'} - \Delta s \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ 0 \end{bmatrix} \tag{3.45}$$

The second equation, known as pseudo-arclength constraint, ensures that the step along the arclength s has a size of Δs (analogous to Eq. 3.39b, for natural parameter continuation); in fact, the step size Δs only approximates the step along the arclength of the curve, thus the name *pseudo*-arclength (Dichmann et al., 2003). Geometrically, this constraint ensures that the corrector algorithm converges to a solution along a direction perpendicular to the initial guess (Fig. 3.2), improving the convergence properties of the algorithm. As previously discussed, the system $\boldsymbol{G}\left(\boldsymbol{V}^{j+1}\right) = \boldsymbol{0}$ is solved by iteratively applying Newton’s method, which requires using the Jacobian matrix of the system, given by (Doedel et al., 2003)

$$D_{\boldsymbol{G}}\left(\boldsymbol{V}^{j+1}\right) = \begin{bmatrix} D_{\boldsymbol{F}}\left(\boldsymbol{V}^{j+1}\right) \\ \left(\boldsymbol{V}^{j'}\right)^{T} \end{bmatrix} \tag{3.46}$$

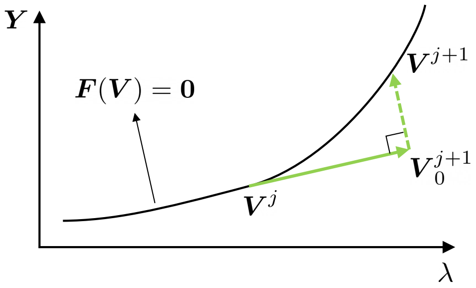

Figure 3.2: Representation of pseudo-arclength continuation. The black curve represents the curves of solutions, the solid arrow the initial guess, and the dashed arrow the direction of convergence.

3.2.5. Continuation Step-Size Control

When following a solutions curve via continuation, the convergence speed (i.e. the required number of iterations) and convergence success/failure of the differential corrector is heavily dependent on the quality of the initial guess, which depends on the continuation step size. The appropriate value of the step size depends on the local shape of the solutions curve. Indeed, a step size that performs well in a section of the curve might not in a different section; this was particularly evident when computing curves of three-dimensional orbits near the secondary of the CR3BP. To circumvent this issue, it is possible to adapt the continuation step size, based on the local performance of the differential corrector. Here, a step-size control mechanism based on an error model for the corrector, presented by Allgower and Georg (2003), is used. This method was selected based on its simple implementation and perceived ease of use. For examples of other step-size control algorithms, the reader is referred to Allgower and Georg (2003) and Seydel (2010).

The following description of the used step-size control method is based on Allgower and Georg (2003). This method assumes that the corrector converges in a direction orthogonal to the initial guess, which is true in the case of multiple shooting with pseudo-arclength constraint. The method is based on the existence of an error model function ψ, which describes, for the used corrector method, the evolution of the error in the free variables vector with the number of iterations. Knowing the number of corrector iterations that a continuation step took to converge, this error model is used to predict the step size which would be required for the corrector to converge in a user-specific number of iterations (which can either correspond to a step-size increase or decrease).

Considering $\boldsymbol{V}_i$ to represent the guess for the solution obtained at the i^{th} iteration and $\boldsymbol{V}_d$ the desired solution, the current modified error is represented by

$$\varepsilon_i(\Delta s) = \gamma \, \| \boldsymbol{V}_d(\Delta s) - \boldsymbol{V}_i(\Delta s) \| \tag{3.47}$$

where $\gamma > 0$ is a constant (ε_i is termed modified error because it is multiplied by the factor γ). The dependence of each variable on the used step size Δs is represented explicitly in the previous equation. The error model function ψ is a function that satisfies

$$\varepsilon_{i+1}(\Delta s) \leq \psi(\varepsilon_i(\Delta s)) \tag{3.48}$$

Since Newton's method converges quadratically, a possible error function is

$$\psi(\varepsilon) = \varepsilon^p, \qquad p = 2 \tag{3.49}$$

The previous equations use the modified error (which depends on the γ constant) instead of the error because the evolution of the error also depends on a constant; see Eq. 3.29 for the case of Newton's method.

Assume that for the step size Δs the corrector converged in m iterations. Furthermore, assume there is an unknown step size $\widetilde{\Delta s}$, which allows the corrector to converge in a chosen number of iterations $\widetilde{m}$. The deceleration factor f is defined as

$$f = \frac{\Delta s}{\widetilde{\Delta s}} \tag{3.50}$$

and can be approximated according to

$$f^2 \approx \frac{\varepsilon_0(\Delta s)}{\varepsilon_0(\widetilde{\Delta s})} \tag{3.51}$$

where $\varepsilon_0(\Delta s)$ and $\varepsilon_0(\widetilde{\Delta s})$ correspond, respectively, to the modified error of the initial guess when using steps Δs and $\widetilde{\Delta s}$. The approximated equation for f is derived using a Taylor series of the solutions curve with respect to a previously known solution. If estimates for $\varepsilon_0(\Delta s)$ and $\varepsilon_0(\widetilde{\Delta s})$ are determined, it is possible to use the previous equations to compute the step size $\widetilde{\Delta s}$ which allows the corrector to converge in the desired number of iterations.

Consider the quotient

$$w = \frac{\| \boldsymbol{V}_m(\Delta s) - \boldsymbol{V}_{m-1}(\Delta s) \|}{\| \boldsymbol{V}_m(\Delta s) - \boldsymbol{V}_0(\Delta s) \|} \approx \frac{\| \boldsymbol{V}_d - \boldsymbol{V}_{m-1}(\Delta s) \|}{\| \boldsymbol{V}_d - \boldsymbol{V}_0(\Delta s) \|} = \frac{\varepsilon_{m-1}(\Delta s)}{\varepsilon_0(\Delta s)} \tag{3.52}$$

The value of w is computed using $\boldsymbol{V}_m(\Delta s)$, $\boldsymbol{V}_{m-1}(\Delta s)$, and $\boldsymbol{V}_0(\Delta s)$, which correspond, respectively, to the free variables vectors obtained at the m^{th} (i.e. the final value), $(m-1)^{th}$, and 0^{th} (i.e. the initial guess) iterations of the multiple-shooting corrector with continuation step Δs. Using the approximate equation for w, the definition of error model function (Eq. 3.48), and the quadratic error function (Eq. 3.49), it is possible to show that

$$\varepsilon_0(\Delta s) \leq w^{1/(p^{m-1}-1)} \tag{3.53}$$

$$\varepsilon_0(\widetilde{\Delta s}) \leq \varepsilon_0(\Delta s)^{p^{m-\widetilde{m}}} \tag{3.54}$$

Hence, the two modified errors can be estimated according to

$$\varepsilon_0(\Delta s) = w^{1/(p^{m-1}-1)} \tag{3.55}$$

$$\varepsilon_0(\widetilde{\Delta s}) = \varepsilon_0(\Delta s)^{p^{m-\widetilde{m}}} \tag{3.56}$$

Note that the variables required to compute the two modified errors are either design parameters (the desired number of iterations $\widetilde{m}$), model parameters ($p = 2$), or can be determined as a byproduct of the application of the multiple-shooting corrector with step Δs (m and w). Having estimated $\varepsilon_0(\Delta s)$ and $\varepsilon_0(\widetilde{\Delta s})$, it is now possible to compute the value of f.

Having presented the computation of the f, it is finally possible to describe the application of the step-size control method. Starting from a known orbit, an initial guess for a new orbit is generated using step Δs and the multiple-shooting corrector is applied to correct that initial guess. If the corrector did not converge in a chosen maximum number of iterations, the step is updated according to

$$\widetilde{\Delta s} = \frac{\Delta s}{2} \tag{3.57}$$

and, starting from the previously known orbit, the updated step is used to generate a new initial guess for the orbit for which the corrector had failed to converge. If, on the other hand, the corrector did converge within the maximum number of iterations, the following procedure is followed: firstly, the deceleration factor f is computed; secondly, the modified deceleration factor f^* is computed via

$$f^* = \max\left\{\min\left\{f, 2\right\}, 0.5\right\} \tag{3.58}$$

Note that f^* is defined in $[0.5, 2]$. The modified deceleration factor f^* is used instead of f to prevent the step size from changing too quickly. Finally, the step is updated according to

$$\widetilde{\Delta s} = \frac{\Delta s}{f^*} \tag{3.59}$$

The updated step $\widetilde{\Delta s}$ is then used to predict the initial guess for the next orbit.

The algorithm was used with the desired number of iterations between 3 and 5; for values higher than that, the multiple-shooting corrector was observed to start struggling to converge often. Additionally, the corrector was applied with the maximum allowed number of iterations set to 10, as the cases that required larger numbers than that, in general, would not converge.

3.2.6. Bifurcations

Bifurcations are singularities of the continuation procedure along a curve of solutions, and generally represent a change in the qualitative behavior of the orbits in the solutions curve; this can correspond to a change of stability, to the intersection of the current curve of solutions with a new one, or to the end of the current curve (Seydel, 2010). Since, in this project, the goal of applying continuation is to find different families of periodic orbits (from which the invariant manifolds will depart), the current description of bifurcations is focused on bifurcations associated with the crossing of solution curves. The study of bifurcations is here divided into two: tangent bifurcations (in which two single-periodic families cross) and period-multiplying bifurcations (in which a higher-period family crosses a single-periodic family). In this section the different types of bifurcations are characterized; the exact computation of the bifurcation points and the switching between different curves of solutions are presented in subsequent sections.

Tangent Bifurcations

Once again consider the system $\boldsymbol{F}(\boldsymbol{V}) = 0$. Tangent bifurcations[2], which correspond to singularities of the continuation procedure, can be further subdivided into turning points and simple bifurcation points.

Turning points[3] occur when the solution curve has a local extremum with respect to the continuation parameter λ. This is the issue that motivated the introduction of the pseudo-arclength continuation; thus, turning points do not occur when applying this type of continuation (Doedel et al., 1991; Seydel, 2010).

The second type of bifurcation, which does occur when using pseudo-arclength continuation, is usually termed simple bifurcation point[4]. At one of these bifurcations, two different curves of solutions intersect each other (Seydel, 2010). It is common to further divide this type of bifurcation into transcritical and pitchfork bifurcations, depending on the way the two solution curves intersect; nevertheless, that distinction is not relevant in the context of the present work.

Before continuing, consider a generic matrix A with dimensions $m \times n$. According to the rank-nullity theorem

$$\mathrm{rank}(A) + \mathrm{nullity}(A) = n \tag{3.60}$$

where the rank of A corresponds to the number of linearly independent rows of the matrix, which coincides with the number of linearly independent columns; thus $\mathrm{rank}(A) \leq \min(m, n)$. The nullity of A is the dimension of its null space.

---

[2] Tangent bifurcations are also known as single-period bifurcations or stationary-point bifurcations.

[3] Turning points are also known as fold or flip bifurcations.

[4] Simple bifurcation points are also known as simple singular points.

Let the modified Jacobian H be defined as

$$H(\boldsymbol{V}) = \begin{bmatrix} D_{\boldsymbol{F}}(\boldsymbol{V}) \\ (\boldsymbol{V}')^T \end{bmatrix} \tag{3.61}$$

In a non-singular point (i.e. a point that is consistent with the assumptions of the implicit function theorem), $D_{\boldsymbol{F}}(\boldsymbol{V})$, which has dimensions $n \times (n+1)$, has a nullity of 1. Hence, $D_{\boldsymbol{F}}(\boldsymbol{V})$ has rank n, thus having n linearly independent rows, to which $\boldsymbol{V}'$ is orthogonal (since it is defined in the null space of $D_{\boldsymbol{F}}(\boldsymbol{V})$). Thus the modified Jacobian H, which has dimensions $(n+1)\times(n+1)$, has rank $n+1$. Therefore, for a curve without singular points, the modified Jacobian H is invertible, meaning that its determinant is never zero, always having the same sign (Allgower and Georg, 2003). Meanwhile, at a simple bifurcation point, two families intersect, meaning that $D_{\boldsymbol{F}}(\boldsymbol{V})$ has a null space with dimension 2 (i.e. two different tangent vectors), thus rank $n-1$. Consequently, H has rank n, meaning it is not invertible and has determinant zero (Allgower and Georg, 2003). Thus, the determinant of H is zero at the bifurcation, and non-zero in regions without bifurcations; it is possible to show that, when following the solutions curve, the determinant of H changes sign when passing over the simple bifurcation point (Doedel et al., 1991). As such, it is possible to detect a simple bifurcation point by tracking the sign of the test function

$$\tau(\boldsymbol{V}) = \det H(\boldsymbol{V}) \tag{3.62}$$

Evidently, this detection method only works if the continuation step size is small enough not to skip over two bifurcations at once. Note that the H matrix (Eq. 3.61) can be efficiently obtained as a byproduct of the continuation procedure, as the latter requires the computation of both $D_{\boldsymbol{F}}(\boldsymbol{V})$ and $\boldsymbol{V}'$.

Tangent bifurcations can also be analyzed from the point of view of maps (e.g. Poincaré map). The following discussion on bifurcations of maps tries (to some extent) to justify the presented results, but omits several points; for more formal and detailed derivations the reader is referred to Howard and MacKay (1987), Scheeres (2012), and Seydel (2010). A Poincaré map $P(\boldsymbol{x})$ is defined as the function that maps a state $\boldsymbol{x}$ on a given Poincaré section (e.g. the plane $y=0$), onto the state at the next crossing of the Poincaré section. Thus, the system of equations that guarantees a periodic orbit can be written as (Seydel, 2010)

$$\boldsymbol{F}(\boldsymbol{V}) := P(\boldsymbol{x}(t_1)) - \boldsymbol{x}(t_1) = 0 \qquad \text{with} \quad \boldsymbol{V} := \boldsymbol{x}(t_1) \tag{3.63}$$

In the case of a periodic orbit, and by definition of the Poincaré map, $P(\boldsymbol{x}(t_1)) = \boldsymbol{x}(t_1+T)$, with T representing the period of the orbit. As such, for periodic orbits (Seydel, 2010)

$$D_{\boldsymbol{F}}(\boldsymbol{V}) = \left.\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{V}}\right|_{\boldsymbol{V}} = \frac{\partial P(\boldsymbol{x}(t_1))}{\partial \boldsymbol{x}(t_1)} - \frac{\partial \boldsymbol{x}(t_1)}{\partial \boldsymbol{x}(t_1)} = M - I \tag{3.64}$$

where M and I represent, respectively, the monodromy and identity matrix. As indicated previously, at a tangent bifurcation, the Jacobian $D_{\boldsymbol{F}}(\boldsymbol{V})$ is singular, thus having zero as an eigenvalue; therefore, at the bifurcation, the monodromy matrix has one non-trivial pair of $\lambda = 1$ eigenvalues[5,6]. It is possible to show that at the tangent bifurcation that pair of non-trivial eigenvalues coalesces at $\lambda = 1$ (stability index $k=2$) and splits along the real axis (Fig. 3.3, left), or vice-versa (Howard and MacKay, 1987). Given that information, one can easily note that, when passing a tangent bifurcation, a family of orbits changes stability.

Period-Multiplying Bifurcations

In period-multiplying bifurcations, a family of T-periodic orbits intersects a new family of nT-periodic orbits, with n an integer. Since the family of T-periodic orbits can also be interpreted as a family of nT-periodic orbits (where each nT-periodic orbit corresponds to a T-periodic orbit repeated n times), the period-multiplying bifurcation of the original T-periodic family is, in fact, a tangent bifurcation of the equivalent nT-periodic family, where it intersects a new family with a similar period.

---

[5] The eigenvalues λ and eigenvectors $\boldsymbol{\xi}$ of a matrix A satisfy

$$A\boldsymbol{\xi} = \lambda\boldsymbol{\xi} \Rightarrow (A + cI)\boldsymbol{\xi} = (\lambda + c)\boldsymbol{\xi}$$

with c a real constant. Therefore a matrix $(A + cI)$ has eigenvalues $(\lambda + c)$.

[6] The reader may note that M always has a trivial pair of eigenvalues $\lambda = 1$, which would indicate that $D_{\boldsymbol{F}}(\boldsymbol{V})$ always has an eigenvalue $\lambda = 0$. One of the points that was omitted is that the Poincaré map should have dimension 4 and not 6, with this reduction resulting from the conservation of energy and from the selection of the Poincaré section. As such, the Poincaré map would only have the two non-trivial pairs of eigenvalues, having lost the trivial pair (Scheeres, 2012).

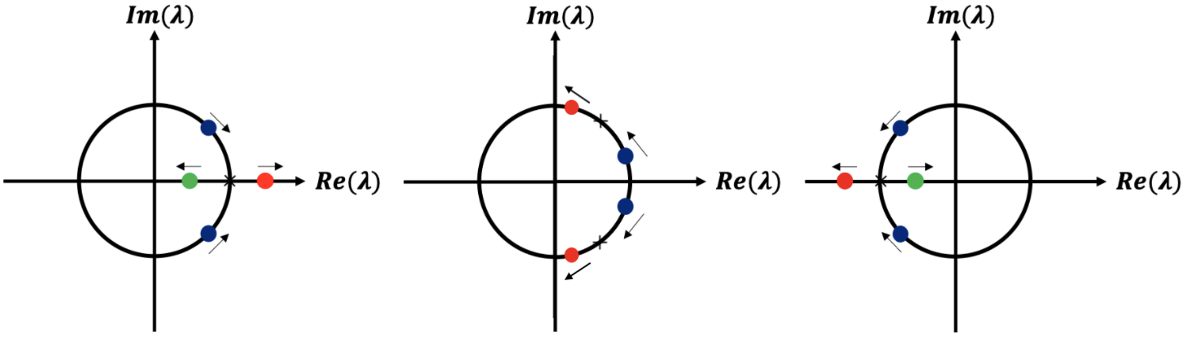

Figure 3.3: Representation of different types of bifurcations: tangent bifurcation (left), period-multiplying bifurcation (center) and period-doubling bifurcation (right). The black circle corresponds to the unit circle, the colored circles represent the eigenvalues of the monodromy matrix, and the crosses represent the bifurcation points. Adapted from Gupta (2020).

Since period-multiplying bifurcations are not singularities of the continuation procedure of the T-periodic family, they cannot be detected using the continuation Jacobian (as is done for tangent bifurcations). They can, however, be identified via a map-based analysis similar to the one presented before. Once again consider a Poincaré map $P(\boldsymbol{x})$; an nT-periodic orbit corresponding to a T-periodic orbit repeated n times verifies (Seydel, 2010)

$$\boldsymbol{F}(\boldsymbol{V}) := P^n(\boldsymbol{x}(t_1)) - \boldsymbol{x}(t_1) = 0 \qquad \text{with} \quad \boldsymbol{V} := \boldsymbol{x}(t_1) \tag{3.65}$$

where $P^n(\boldsymbol{x})$ represents the Poincaré map applied n times. Taking the derivative of $\boldsymbol{F}(\boldsymbol{V})$ and applying the chain rule, one finds (Seydel, 2010)

$$D_{\boldsymbol{F}}(\boldsymbol{V}) = \frac{\partial P^n(\boldsymbol{x}(t_1))}{\partial \boldsymbol{x}(t_1)} - \frac{\partial \boldsymbol{x}(t_1)}{\partial \boldsymbol{x}(t_1)} = \left(\frac{\partial P(\boldsymbol{x}(t_1))}{\partial \boldsymbol{x}(t_1)}\right)^n - \frac{\partial \boldsymbol{x}(t_1)}{\partial \boldsymbol{x}(t_1)} = M^n - I \tag{3.66}$$

where M represents the monodromy matrix of the T-periodic orbit, and M^n the monodromy matrix of the equivalent nT-periodic orbit. At a tangent bifurcation of the nT-periodic orbit, the Jacobian $D_{\boldsymbol{F}}(\boldsymbol{V})$ has zero as an eigenvalue. Therefore the monodromy matrix of the nT-periodic orbit (i.e. M^n) has one non-trivial pair of eigenvalues $\lambda = 1$; this means that the monodromy matrix of the T-periodic orbit (i.e. M), has a pair of eigenvalues $\lambda = \sqrt[n]{1} = \sqrt[n]{\exp(j2\pi d)}$, with j the imaginary unit and d an integer[7]. Consequently, a period-multiplying bifurcation occurs when the eigenvalues of the monodromy matrix of the T-periodic orbit cross the values (Lara et al., 2007)

$$\lambda, 1/\lambda = \cos\left(\frac{2\pi d}{n}\right) \pm j \sin\left(\frac{2\pi d}{n}\right) \Rightarrow \tag{3.67}$$

$$k = 2\cos\left(\frac{2\pi d}{n}\right), \qquad n \in \mathbb{N}^+,\ d = 1, 2, ..., \text{floor}\left(\frac{n}{2}\right) \tag{3.68}$$

Note that d and n are such that $d/n \leq 1/2$; since the eigenvalues appear in conjugate pairs, for larger values of the d/n ratio the critical eigenvalues would start repeating. This bifurcation can happen in any location of the unit circle besides $\lambda = 1$, with the stability index taking a value $k \in [-2, 2[$. The case $k = -2$ ($\lambda = -1$) corresponds to a special case: a period-doubling bifurcation; thus, a period-multiplying bifurcation is usually considered to correspond just to the cases $k \in]-2, 2[$.

In a period-multiplying bifurcation, a pair of eigenvalues of the original family crosses one of the critical points defined by Eq. 3.67, and remains on the unit circle (Fig. 3.3, center), thus not resulting in a change of stability of the original family (Bosanac, 2016). As such, the stability index is real before, in, and after the bifurcation, meaning the bifurcation can be trivially identified using a test function

$$\tau(k) = k - k^* \tag{3.69}$$

where k represents the current stability index and k^* the stability index of the period-multiplying bifurcation being searched. As already mentioned, at a period multiplying-bifurcation a new nT-periodic family is formed, usually named a $d : n$ resonant family. In each period, an orbit of this family executes n in-plane revolutions

---

[7] If a matrix A has eigenvalues λ, the matrix A^n has eigenvalues λ^n.

and $(n-d)$ out-of-plane revolutions (Lara et al., 2007), i.e. the orbit crosses the xz-plane in the $\dot{z} > 0$ direction n times, and crosses the xy-plane in the $\dot{y} > 0$ direction $(n-d)$ times.

Finally, in a period-doubling bifurcation, which is a special case of the period-multiplying bifurcation, the original T-periodic intersects a new $2T$-periodic family. In a period-doubling bifurcation, a non-trivial pair of eigenvalues coalesce at $\lambda = -1$ $(k = -2)$ and split along the real axis (Fig. 3.3, right), or vice-versa (Howard and MacKay, 1987). Consequently, and contrary to the period-multiplying bifurcation, the period-doubling bifurcation does produce a change in the stability of the original family. Similarly to period-multiplying bifurcations, period-doubling ones can be detected using a test function based on the value of the stability index (Eq. 3.69).

3.2.7. Computation of Bifurcations

Assume a test function $\tau(\boldsymbol{V})$ which has opposite signs for each side of a certain type of bifurcation. Computing the bifurcation point corresponds to determining the root of the test function. Therefore, after finding two consecutive solutions $\boldsymbol{V}$ which produce a change in the test function sign, the bifurcation point is determined by executing the pseudo-arclength continuation procedure with a step-size adaptation based on the secant root-finding method, according to (Allgower and Georg, 2003)

$$\Delta s_{j+1} = -\frac{\tau\left(\boldsymbol{V}^{j}\right)}{\tau\left(\boldsymbol{V}^{j}\right) - \tau\left(\boldsymbol{V}^{j-1}\right)} \Delta s_{j} \tag{3.70}$$

where $\boldsymbol{V}^{j}$ and $\boldsymbol{V}^{j-1}$ are two previously known solutions, and Δs_{j+1} is the step size to be used when determining the solution $\boldsymbol{V}^{j+1}$. Allgower and Georg (2003) also propose using Newton's method, which has faster convergence than the secant method, but the latter is preferred here as it does not require computing the derivative of the test function.

Simple bifurcation points are detected using as test function the determinant of the modified Jacobian (Eq. 3.62). Since these bifurcations are characterized by a stability index $k = 2$, one could also use a test function based on the value of the stability index. However, the determinant of the extended Jacobian has the advantage of only identifying simple bifurcation points, and not turning points (as is the case for the stability index), which is preferred here, as turning points do not indicate the crossing of solution curves. When computing period-doubling and period-multiplying bifurcations, the stability indices are used as test function (Eq. 3.69).

As previously indicated, and repeated here for clarity, at simple bifurcation points, the modified Jacobian

$$H(\boldsymbol{V}) = \begin{bmatrix} D_{\boldsymbol{F}}(\boldsymbol{V}) \\ \left(\boldsymbol{V}'\right)^{T} \end{bmatrix}$$

is singular. As such, the Jacobian of the continuation procedure (which was previously presented in Eq. 3.46)

$$D_{\boldsymbol{G}}\left(\boldsymbol{V}^{j+1}\right) = \begin{bmatrix} D_{\boldsymbol{F}}\left(\boldsymbol{V}^{j+1}\right) \\ \left(\boldsymbol{V}^{j'}\right)^{T} \end{bmatrix}$$

even though non-singular, becomes ill-conditioned near simple bifurcation points. Thus, when approaching the bifurcation, the secant method tends to become unstable, not allowing a very accurate determination of the bifurcation point (Allgower and Georg, 2003). Nevertheless, the achieved accuracy was observed to be sufficient for executing the branch-switching procedure (presented in the following subsection).

Since period-doubling and period-multiplying bifurcations do not constitute singularities of the continuation procedure (do recall that the continuation procedure is applied with respect to the single-periodic orbits), the secant method does not suffer from numerical instabilities, allowing the bifurcations to be determined very accurately.

3.2.8. Branch Switching

At simple bifurcation points two curves of solutions intersect, with the Jacobian $D_{\boldsymbol{F}}(\boldsymbol{V})$ having a null space with dimension 2. Consider the two base vectors of the null space to be represented by $\boldsymbol{V}'$ and $\widetilde{\boldsymbol{V}}'$, indicating, respectively, the vector tangent to the original branch and the vector tangent to the new branch. In principle, one could switch to the new branch by applying the traditional pseudo-arclength continuation procedure, with a perturbation in the $\widetilde{\boldsymbol{V}}'$ direction. However, the location of the bifurcation point cannot be determined exactly, only approximately; as such, at the approximate bifurcation point, the null space of the Jacobian

$D_{\boldsymbol{F}}(\boldsymbol{V})$ has dimension 1, with tangent vector $\boldsymbol{V}'$ (Allgower and Georg, 2003). Therefore, to switch branches it is first necessary to obtain an approximation of $\widetilde{\boldsymbol{V}}'$.

The tangent vector $\widetilde{\boldsymbol{V}}'$ is here approximated using the orthogonal-direction method, which selects $\widetilde{\boldsymbol{V}}'$ to be orthogonal to $\boldsymbol{V}'$ (Fig. 3.4). In principle, one could select $\widetilde{\boldsymbol{V}}'$ to have any direction within the null space of $D_{\boldsymbol{F}}(\boldsymbol{V})$; however, considering that no information about the ideal direction of $\widetilde{\boldsymbol{V}}'$ is available, selecting it to be orthogonal to $\boldsymbol{V}'$ is the logical option, as it ensures that the initial guess for the new branch is as distant as possible from the previous branch. Based on the orthogonal selection of $\widetilde{\boldsymbol{V}}'$, at the exact bifurcation point, one would have (Allgower and Georg, 2003)

$$\begin{bmatrix} D_{\boldsymbol{F}}(\boldsymbol{V}) \\ (\boldsymbol{V}')^T \end{bmatrix} \widetilde{\boldsymbol{V}}' = 0 \Leftrightarrow H\widetilde{\boldsymbol{V}}' = 0 \tag{3.71}$$

which specifies $\widetilde{\boldsymbol{V}}'$ to be contained in the null space of H (i.e. to be in the null space of $D_{\boldsymbol{F}}(\boldsymbol{V})$ and orthogonal to $\boldsymbol{V}'$). Since the computed location of the bifurcation point is an approximation, this equation has no solution. Nevertheless, an approximate solution is given by the unit eigenvector of $H^T H$ associated with the eigenvalue closest to zero (Allgower and Georg, 2003). This method for determining the solution derives from the singular-value decomposition, according to which the null space of a matrix A is spanned by the eigenvectors of $A^T A$ associated with zero eigenvalues; thus, the null space of a matrix A is approximately spanned by the eigenvectors of $A^T A$ associated with approximately zero eigenvalues (Press et al., 2007; Zimovan-Spreen, 2021).

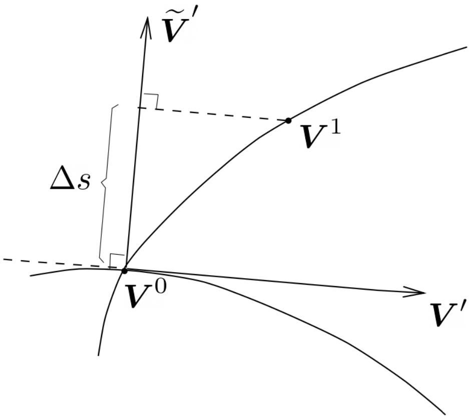

Figure 3.4: Representation of branch switching using the orthogonal-direction method, with convergence to the new branch using pseudo-arclength continuation. The solid lines represent the two intersecting branches of solutions and the solid vectors the (approximately) tangent direction to each branch. Adapted from Doedel (2011).

There are other more accurate methods for switching branches, namely by selecting $\widetilde{\boldsymbol{V}}'$ to be tangent to the new branch (involves obtaining an approximation for an equation known as the bifurcation equation, for details see Doedel et al., 1991). However, this would require evaluating the second derivative of $\boldsymbol{F}(\boldsymbol{V})$, which the simpler orthogonal-direction method does not; furthermore, the latter is expected to work well in most situations (Doedel et al., 1991).

To summarize, when encountering a simple bifurcation point, the vector tangent to the new branch is approximated using the orthogonal-direction method; that vector is then used to initiate the continuation procedure of the new branch. An analogous procedure is used to switch branches at period-multiplying bifurcations (and period-doubling ones). As described previously, a period-multiplying bifurcation of a simple-periodic family corresponds to a tangent bifurcation of an nT-periodic family. Therefore, to switch branches, the simple-periodic bifurcating orbit is propagated over n revolutions, and the branch-switching procedure for simple bifurcation points is applied to the nT-periodic bifurcating orbit (Zimovan-Spreen, 2021). When switching to the nT-periodic family, the number of multiple-shooting nodes used to describe each orbit is increased by a factor n (Wulff and Schebesch, 2006); similarly, the tolerance used to assess the convergence of the multiple-shooting corrector is also increased by a factor n.

3.2.9. (Shape) Homotopy

A final point to discuss is the technique known as homotopy, which is a form of natural parameter continuation (Subsection 3.2.3) used to transform solutions between different systems of equations. Consider the system of equations $\boldsymbol{F}(\boldsymbol{V}) = \boldsymbol{0}$, where a solution is difficult to find, and a simplified system $\overline{\boldsymbol{F}}(\boldsymbol{V}) = \boldsymbol{0}$, where it is easier to find a solution (for example, through analytical techniques). Now consider a homotopy function defined as (Seydel, 2010)

$$\boldsymbol{G}_{hom}(\boldsymbol{V},\alpha) = \alpha\boldsymbol{F}(\boldsymbol{V}) + (1-\alpha)\overline{\boldsymbol{F}}(\boldsymbol{V}), \qquad \text{with } \alpha \in [0,1] \tag{3.72}$$

Starting from a solution to the simpler system of equations ($\alpha = 0$), one applies continuation using the homotopy parameter α. This parameter is gradually increased, producing a series of systems of increasing complexity and their solutions; eventually, one is able to obtain a solution to the more complex system of equations ($\alpha = 1$). In astrodynamics, homotopy is often used to transform solutions between simpler and more complex dynamics models. For example, in the context of the Mars-Phobos system, Scheeres et al. (2019) use this technique to transform periodic orbits between the H3BP with ellipsoidal secondary and the more-complex formulation with polyhedral secondary; Zamaro (2015) uses it to transform periodic orbits from the CR3BP to the elliptic restricted three-body problem (similar model to the CR3BP, but with the primaries having elliptical orbits).

A different form of homotopy is used by Capannolo et al. (2019), who transforms orbits between the CR3BP and the CR3BP with polyhedral secondary by gradually changing from a sphere-shaped polyhedron to an asteroid-shaped polyhedron. An analogous technique is proposed here, with the difference that the shape homotopy is started from a polyhedron with the shape of an ellipsoid. The cartesian position of points on the surface of an ellipsoid, as a function of the latitude ϕ and longitude λ, is given by (Poelaert et al., 2004)

$$x(\phi,\lambda) = r\cos\phi\cos\lambda \tag{3.73}$$

$$y(\phi,\lambda) = r\cos\phi\sin\lambda \tag{3.74}$$

$$z(\phi,\lambda) = r\sin\phi \tag{3.75}$$

$$\text{with} \qquad r(\phi,\lambda) = \frac{abc}{\left(bc\cos\phi\cos\lambda\right)^2 + \left(ac\cos\phi\sin\lambda\right)^2 + \left(ab\sin\phi\right)^2} \tag{3.76}$$

where a, b and c correspond to the ellipsoid's semi-major axes. To transition between the ellipsoid and the Phobos-shaped polyhedron, the shape of the polyhedron is gradually changed according to

$$\boldsymbol{r}_{hom}(\phi,\lambda,\alpha) = \alpha\boldsymbol{r}_P(\phi,\lambda) + (1-\alpha)\boldsymbol{r}_E(\phi,\lambda) \tag{3.77}$$

where $\boldsymbol{r}_E$ and $\boldsymbol{r}_P$ denote, respectively, the position of each polyhedron vertex according to the ellipsoid and Phobos' shape. When determining the ellipsoidal polyhedron of Phobos, the mean semi-major axes listed in Tab. 2.1 were used. The polyhedron model of Phobos is represented in Fig. 3.5 for different values of the homotopy parameter.

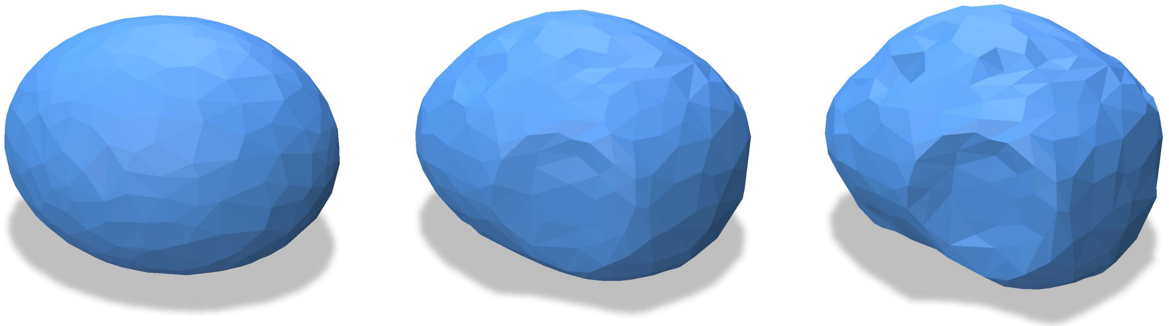

Figure 3.5: Shape homotopy of Phobos' polyhedron model: homotopy parameter $\alpha = 0$ (left), $\alpha = 0.5$ (center) and $\alpha = 1.0$ (right).

A few differences between the proposed shape homotopy (ellipsoidal polyhedron to small-body polyhedron) and the traditional model homotopy (namely, spherical harmonics of an ellipsoid to polyhedron) are worth highlighting. Firstly, the former corresponds to changing the shape of the body such that each intermediate model has constant density, but each two consecutive models (usually) have different densities; the latter corresponds to transitioning between the ellipsoid and polyhedron in a way that the areas that are included in both have the same density in all intermediate models, while the density of the areas included in only one of them gradually increases/decreases between intermediate models.

The second point worth mentioning is that the shape homotopy is effectively not as “smooth” as the model homotopy. With the model homotopy, $\alpha = 0$ corresponds exactly to the simplified model used. That is not the case for the shape homotopy, where $\alpha = 0$ corresponds to a slightly different model with respect to the simplified one, as an ellipsoidal polyhedron (constituted by flat facets) is not a perfect representation of a curved ellipsoid. This difference will tend to decrease as the number of polyhedron facets increases. For the models used, this initial discontinuity in the shape homotopy was not observed to be problematic.

When applying the shape homotopy, three intermediate models were used (α taking the values 0.25, 0.5, and 0.75). During the homotopy, the described multiple-shooting algorithm was applied, but without the pseudo-arclength constraint.

3.2.10. Final Points

Here a summary of the described continuation procedure is presented, and some additional points of attention are highlighted. In the following chapter, the continuation algorithm will be used to compute families of three-dimensional QSOs; starting from these orbits, unstable invariant manifolds will be propagated and used for designing landing trajectories. In the following chapter, for additional clarity, the relevant parts of the algorithm are mentioned again when analyzing the respective results. The implementation of the described methods is verified in Section 7.3 (multiple shooting and continuation) and Section 7.4 (computation of bifurcations).

The workflow required for the application of the continuation algorithm is represented in Fig. 3.6. The first step in the continuation procedure is the generation of an initial guess for an orbit belonging to a family of interest, which can be done using some analytical approximation. The initial guess is then corrected, to obtain a periodic orbit, using multiple shooting (without the pseudo-arclength constraint, as using this constraint requires knowing another orbit in the family). Next, all the orbits in that family are computed: using pseudo-arclength continuation, each new orbit is used to generate an initial guess for the next one, and that initial guess is corrected via multiple shooting (with the pseudo-arclength constraint). By tracking the evolution of the eigenvalues of the monodromy matrix along the computed family, one detects where a bifurcation was jumped over, thus identifying one orbit on each side of the bifurcation. The bifurcating orbit is then computed by applying the continuation algorithm with a step-size selection based on some root-finding algorithm (the

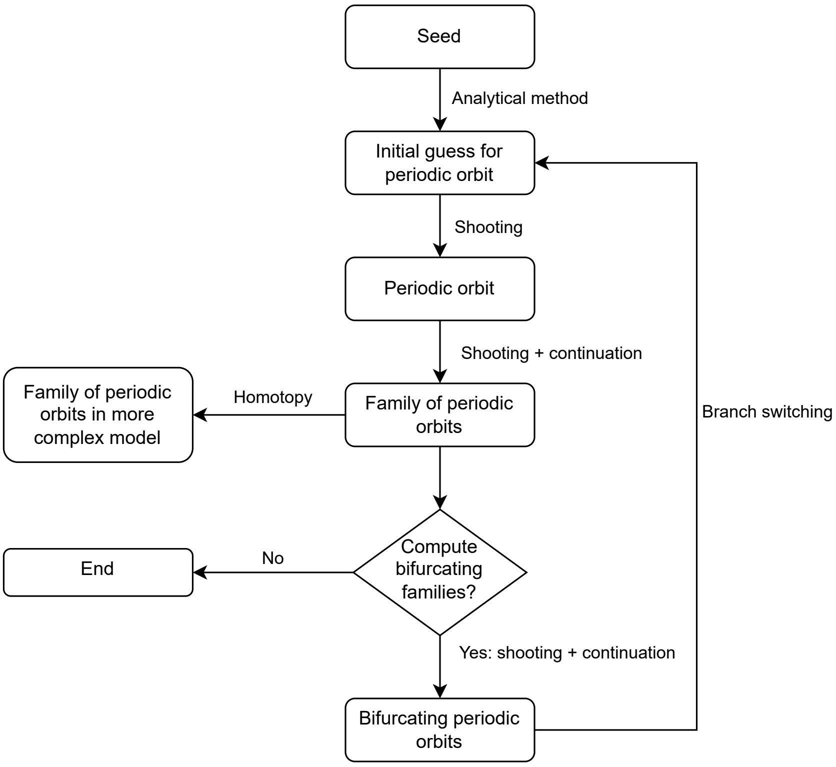

Figure 3.6: Flowchart summarizing the used continuation algorithm.

secant method in this case). At the bifurcation point, the orthogonal-direction branch-switching method is applied, allowing the continuation algorithm to jump to the newly found family. The continuation and bifurcation detection procedures are then repeated along this new family. Finally, if necessary, a homotopy procedure might be applied to convert the obtained solutions between different models.

When analyzing a symmetric system, for example, the CR3BP with ellipsoidal secondary, the different families of orbits intersect each other at the bifurcation points. However, when analyzing an asymmetric model, e.g. CR3BP with polyhedral secondary, the different families of orbits no longer intersect each other (Scheeres et al., 2019). Thus, if a symmetric model similar to the asymmetric one is available, one can first apply continuation with the symmetric model to obtain examples of periodic orbits in all families of interest, and then convert those orbits to the asymmetric model using homotopy. This is the reason why two gravity models for Phobos, i.e. ellipsoid and polyhedron, were introduced.

When analyzing two-dimensional orbits, there is one stability index associated with the in-plane motion and another associated with the out-of-plane motion (Lara, 2003). As such, if the initial family of periodic orbits is two-dimensional (which is the case in this work), and one is only interested in the two- or three-dimensional bifurcating families, then one would only execute the branch-switching procedure at the bifurcations associated with, respectively, the horizontal or vertical stability index. Here only three-dimensional bifurcating families are analyzed, thus only the bifurcations associated with the vertical stability index are taken into account.

4

Orbital Motion Around Phobos

In this chapter, the families of three-dimensional QSOs in the vicinity of Phobos are computed and analyzed. The invariant manifolds of the unstable members of this family are then computed and analyzed, paying attention to the relationship between the characteristics of the manifolds and the family from which they originate.

4.1. Quasi-Satellite Orbits

4.1.1. Two-Dimensional Quasi-Satellite Orbits with Ellipsoid

The two-dimensional QSOs presented in this section were computed using the CR3BP with ellipsoidal secondary. Only very general results are presented, which are valid for other models, e.g. CR3BP, and other planetary systems.

Depending on the location, Phobos' sphere of influence is either below or very close to its surface (Zamaro, 2015), thus not being possible to orbit it in the sense of the two-body problem (i.e. with the moon as the central body). Instead one can use QSOs, which are a type of relative orbits around the secondary that are present in three-body systems (e.g. CR3BP). QSOs are generally known for their stability. When viewed from an inertial frame, QSOs can be interpreted as low-eccentricity prograde orbits around the primary, though perturbed by the secondary. However, these trajectories remain close to the secondary throughout its orbit around the primary, thus appearing to be secondary-centered "elliptic orbits" when viewed from a synodic frame (Fig. 4.1). With respect to the synodic frame, QSOs are followed in a retrograde direction; for this reason, these orbits are also called distant retrograde orbits, though this term is more often applied to members of the family which are far away from the secondary, having a kidney-like shape. For high altitudes, i.e. $A_x > 100$ km for the Mars-Phobos system (where A_i denotes the amplitude of the QSO in the coordinate i), QSOs have a period similar to that of the secondary; for lower altitudes, the period decreases (Scheeres et al., 2019).

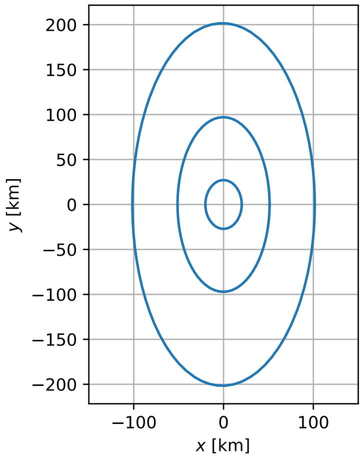

Figure 4.1: Examples of two-dimensional QSOs in the Mars-Phobos system. The reference frame is centered on Phobos.

To compute the family of two-dimensional QSOs via continuation, one can start with the analytical initial guess for a high-altitude QSO given by (Hénon, 1969)

$$\begin{aligned} &y(t_1) = z(t_1) = \dot{x}(t_1) = \dot{z}(t_1) = 0 \\ &\dot{y}(t_1) = -2x(t_1) \\ &T = 2\pi \end{aligned} \tag{4.1}$$

where $\boldsymbol{x}(t_1)$ and T correspond, respectively, to the dimensionless initial state and dimensionless orbit period. An initial x-coordinate of 100 km was used. Starting from this initial guess, the procedure outlined in the previous chapter was applied to continue the initial family of two-dimensional QSOs. The QSOs represented in Fig. 4.1 were selected from this family. In the case of the Mars-Phobos system, QSOs are always stable (Fig. 4.2), therefore not having unstable manifolds. Since the stability indices are defined in]-2,2[, the two-dimensional QSO family only has period-multiplying bifurcations.

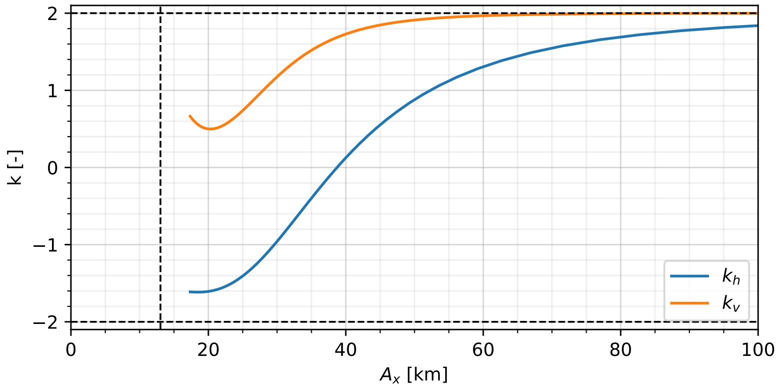

Figure 4.2: Vertical and horizontal stability indices (respectively, k_v and k_h) as a function of the x-amplitude, for the orbits constituting the two-dimensional QSO family. The vertical dashed line represents Phobos' surface (mean ellipsoid), while the horizontal lines demarcate the boundaries of the stability region.

4.1.2. Three-Dimensional Quasi-Satellite Orbits with Ellipsoid

In this section, the families of three-dimensional QSOs bifurcating from the two-dimensional QSO family are analyzed, using the CR3BP model with ellipsoidal secondary.

Since this work is focused on the usage of invariant manifolds of three-dimensional QSOs, only the bifurcations of the two-dimensional QSO family associated with the vertical stability index are considered. Families of three-dimensional QSOs are computed for vertical bifurcations up to $A_x = 50.0$ km, coinciding with the amplitude of the three-dimensional QSO that will be used by MMX (Nakamura et al., 2021); the considered bifurcations are represented in Fig. 4.3. Each family is continued until either A_z reaches 100.0 km or the minimum distance to the center of Phobos gets lower than 14.0 km (since usually the closest approach to the surface occurs near the largest semi-major axis of Phobos, this corresponds to an altitude of approximately 1.0 km; c.f. Tab. 2.1).

While a single new family is created at single-period and period-doubling bifurcations, families of three-dimensional periodic orbits bifurcating from two-dimensional ones at period-multiplying bifurcations always occur in pairs (Robin and Markellos, 1980, 1983). At a $d:n$ bifurcation with n odd a family of x-symmetric orbits and one of xz-symmetric orbits are formed; at a $d:n$ bifurcation with n even, two families of doubly-symmetric orbits (simultaneously x-symmetric and xz-symmetric) are formed.

As here only three-dimensional families appearing from period-multiplying bifurcations are analyzed, at each bifurcation, two full families are computed, with two branches per family (initial perturbation in the positive and negative direction). While continuing each family, the continuation procedure was observed to have some issues when encountering bifurcations: it was expected/desired that it would simply skip over the bifurcations, thus remaining on the same family, however, in some cases it was observed to switch to a different family at the bifurcation point. A database of families of three-dimensional QSOs was generated by selecting from the computed branches the ones that always appeared to remain in the same family. The bifurcations represented in Fig. 4.3 already include just the ones for which valid branches were found (in particular, the $1:10$ and $1:11$ families are not represented). Since orbits originating from the same bifurcation but with

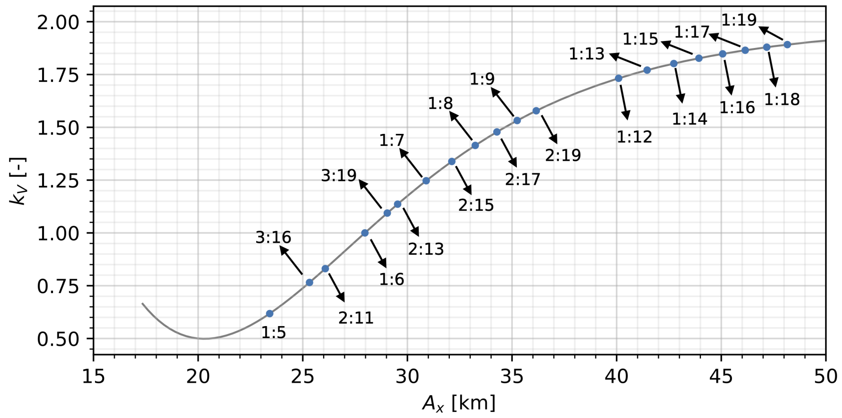

Figure 4.3: Vertical stability index as a function of the x-amplitude, for the orbits constituting the two-dimensional QSO family. The period-multiplying bifurcations for which three-dimensional families were computed are marked with circles.

different symmetries often have very similar energies and amplitudes, it is sometimes difficult to distinguish them based on the evolution of these characteristics throughout the family. Instead, the evolution of the difference between each orbit's maximum and minimum pseudo angular-momentum ("pseudo" because it is defined with respect to the synodic frame)

$$\max \|\boldsymbol{r} \times \dot{\boldsymbol{r}}\| - \min \|\boldsymbol{r} \times \dot{\boldsymbol{r}}\| \tag{4.2}$$

was observed to be particularly helpful in recognizing the different families.

Different factors might have contributed to this undesired behavior at bifurcations. First and foremost, the selected algorithm, a priori, offers no guarantee of remaining in the same family when skipping over a bifurcation; perhaps other variants of this method exist that allow dealing with this situation more robustly. Secondly, the fact that the multiple-shooting differential corrector was formulated without taking advantage of symmetries (so that it was possible to use it with the polyhedron model) makes the continuation more susceptible to switching to different families, as the differential corrector is able to converge to any type of orbit; meanwhile, a differential corrector formulated using symmetries would only be able to converge to orbits displaying the desired symmetry. The usage of a variable continuation step size might also have contributed to the observed problems: one can hypothesize that when approaching bifurcations, which are singularities of the Jacobian matrix, the differential corrector will have more difficulties converging, leading the step size to be reduced, and thus increasing the likelihood of the generated initial guess being very close to the bifurcation, and as a result very close to both families, making it easier for the differential corrector to converge to the newly encountered family. Do note that using a fixed step size would probably be even more problematic: firstly, it would require a manual selection of the step size for each family; secondly, to ensure convergence, it would likely require small step sizes to be selected, again meaning that the generated initial guesses would often be too close to the bifurcations, while unnecessarily increasing the computational time.

These families of three-dimensional QSOs in the Mars-Phobos system were previously computed by Chen et al. (2020) and Pushparaj et al. (2021), though using different variants of the pseudo-arclength algorithm. Chen et al., who focus on computing high-A_z QSOs, suffer the same bifurcation-skipping issues reported here. They overcome the issue by forcing the continuation procedure to use a larger step size when approaching a bifurcation; this is defined manually, based on the observed curvature of each family in an A_x vs. A_z plot. Pushparaj et al. do not mention any issues.

Analyzing the entire set of computed families, one observes that the orbits are stable for low A_z, but become increasingly unstable as A_z increases (Fig. 4.4, left). This behavior is particularly noticeable for the families with the bifurcating orbit below $A_x = 40.0$ km; for the other families ($A_x \geq 40.0$ km, corresponding to the $1:n$ families, $n \geq 11$), although the instability does increase with A_z, it does so very slowly. The variation in the normalized stability index appears to generally be related to the evolution of the minimum distance to Phobos throughout each family (Fig. 4.4, right). The $1:n$, $n \geq 11$ families tend to grow in A_z with small variations in A_x and A_y, each orbit forming approximately the surface of an elliptic cylinder; as such, there are only small variations in the minimum altitude throughout the family and small variations in the stability. Similarly, the families that bifurcate for $A_x < 40.0$ km become more unstable as the minimum altitude decreases. Two exceptions can be observed: the two families that bifurcate for $A_x < 40.0$ km (left side of the plots) and for which A_z grows above 60.0 km; these are members of the 1 : 5 and 1 : 6 families.

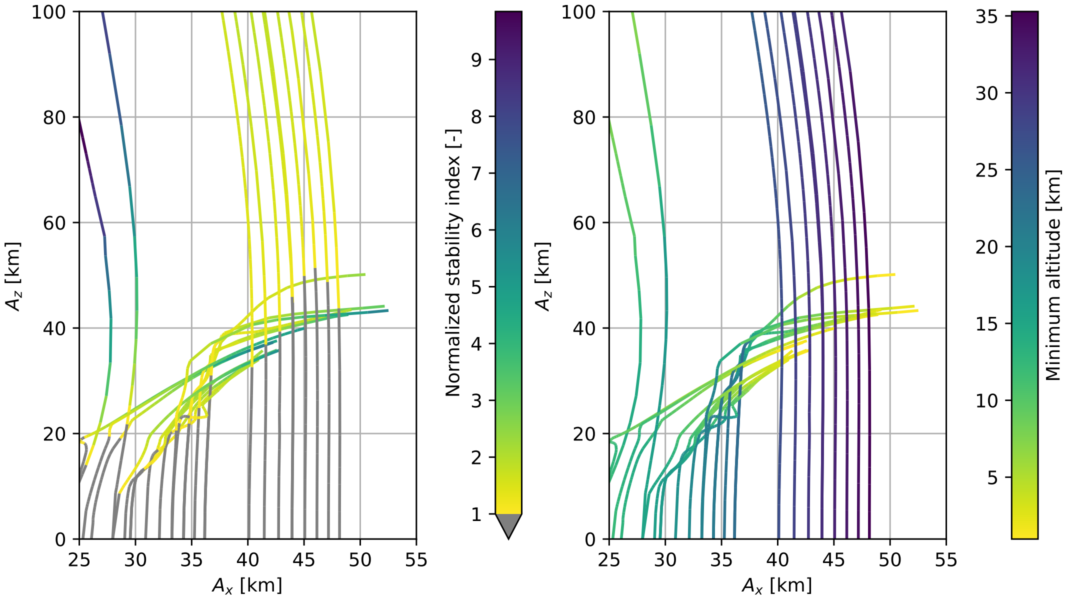

Figure 4.4: Normalized stability index (left) and minimum altitude with respect to the mean ellipsoid (right), as a function of the x-amplitude and z-amplitude of the originating orbit. Each line corresponds to a family of orbits. In the stability-index plot (left), gray lines represent stable orbits.

For all computed families the Jacobi integral decreases for increasing A_z (Fig. 4.5). Therefore, manifolds departing from orbits with higher A_z will have higher impact velocities on Phobos' surface than those departing from lower A_z.

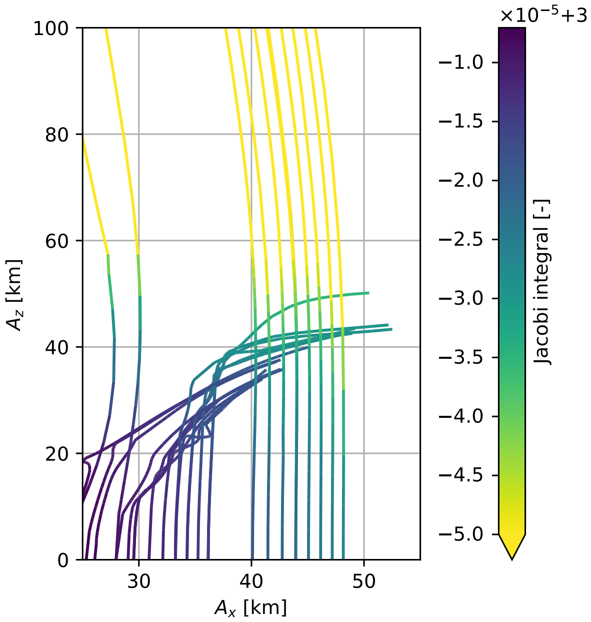

Figure 4.5: Jacobi integral as a function of the x-amplitude and z-amplitude of the originating orbit. Each line corresponds to a family of orbits.

The computed families of QSOs, in the A_x vs. A_z plot, appear very similar to the ones determined by Pushparaj et al. (2021). Exceptions are the high-altitude 1 : 5 and 1 : 6 families (left side of the previous plots). The first reason for that is that Pushparaj et al. only compute xz-symmetric families, while the high-altitude 1 : 5 family computed here is x-symmetric. Secondly, Pushparaj et al. compute a single family per bifurcation;

thus, it is possible they only computed the low-altitude $1:6$ family (also computed here). The occurrence of high-altitude $1:n$ families, $n \geq 11$ appears to be in line with Chen et al. (2020). However, Chen et al., who only compute xz-symmetric families, are able to compute the high-altitude $1:5$ family, which here appears to be x-symmetric. This might be related to differences in the dynamical structure induced by differences in the dynamics models: Chen et al. uses the CR3BP, while here the CR3BP with ellipsoidal secondary is used. Similar discrepancies in the dynamical structure were previously observed by Scheeres et al. (2019), who noted differences in the bifurcations of Lagrange point orbits, between the H3BP and the H3BP with an ellipsoidal secondary.

In the following pages, representative members of the computed families of three-dimensional QSOs are plotted. The represented families are named according to $d:n-symmetry[(code)]$, where $code$ is only used in the cases where there is more than one family with the same type of resonance and symmetry. The type of symmetry characterizing each orbit can be identified using the rules in Subsection 2.1.4.

The $1:5-xz$ and $1:5-x$ families are represented in Figs. 4.6 and 4.7, and the $1:6-xz/x(BE)$ and $1:6-xz/x(R)$ families in Figs. 4.8 and 4.9. These are representative of the $1:n$ families, $n < 11$, for n odd and n even. Looking at the three-dimensional plots, it is possible to identify the value n of a given family by, for example, counting the number of crossings with the xz-plane for $x < 0$. The $n-d$ value, which corresponds to the number of crossings of the xy-plane with $\dot{z} > 0$, is harder to identify visually. The $1:6-xz/x(BE)$ family is coded with BE because in the yz-plane the orbits appear to be a combination of a butterfly (B) and an ellipse (E); the $1:6-xz/x(R)$ family is coded with R because in the yz-plane the orbits have a more regular (R) shape. Note that despite the different symmetries, the $1:5-xz$ family is similar to the $1:6-xz/x(BE)$ family (both in shape and in maximum A_z, with $A_z < 60.0$ km), and the $1:5-x$ family to the $1:6-xz/x(R)$ family (both have similar shapes, and both continue up to $A_z \geq 60.0$ km). The shape of the $1:5-xz$ and $1:6-xz/x(BE)$ families becomes increasingly irregular as A_z increases (top to bottom, when looking at the trajectory plots), as a result of the previously discussed decrease in the minimum distance to Phobos.

The $1:12-xz/x$ and $1:13-xz$ are represented, respectively, in Figs. 4.10 and 4.11; these are illustrative of the $1:n$ families, $n \geq 11$. The members of these families grow in A_z while maintaining the rest of the shape approximately identical, with the projection onto the xy-plane always remaining very similar to the shape of the bifurcating two-dimensional QSO. Thus, the members of these families display a very regular shape, always appearing to be approximately contained on the surface of an elliptic cylinder. This is likely due to the relatively weak effect of Phobos' gravity (these families have high A_x, also having large minimum distances to Phobos' surface).

Finally, the $2:13-xz$ and $2:13-x$ families are represented, respectively in Figs. 4.12 and 4.13, illustrating the $d:n$ families, $d \geq 2$. These display very irregular shapes, with large variations with respect to the bifurcating two-dimensional QSO. Nevertheless the chaotic appearance, similarities can be identified with respect to the $d = 1$ families; in particular, the $2:13-xz$ family is clearly identifiable with the $1:5-xz$ family. None of the $d \geq 2$ families was observed to grow up to high A_z.

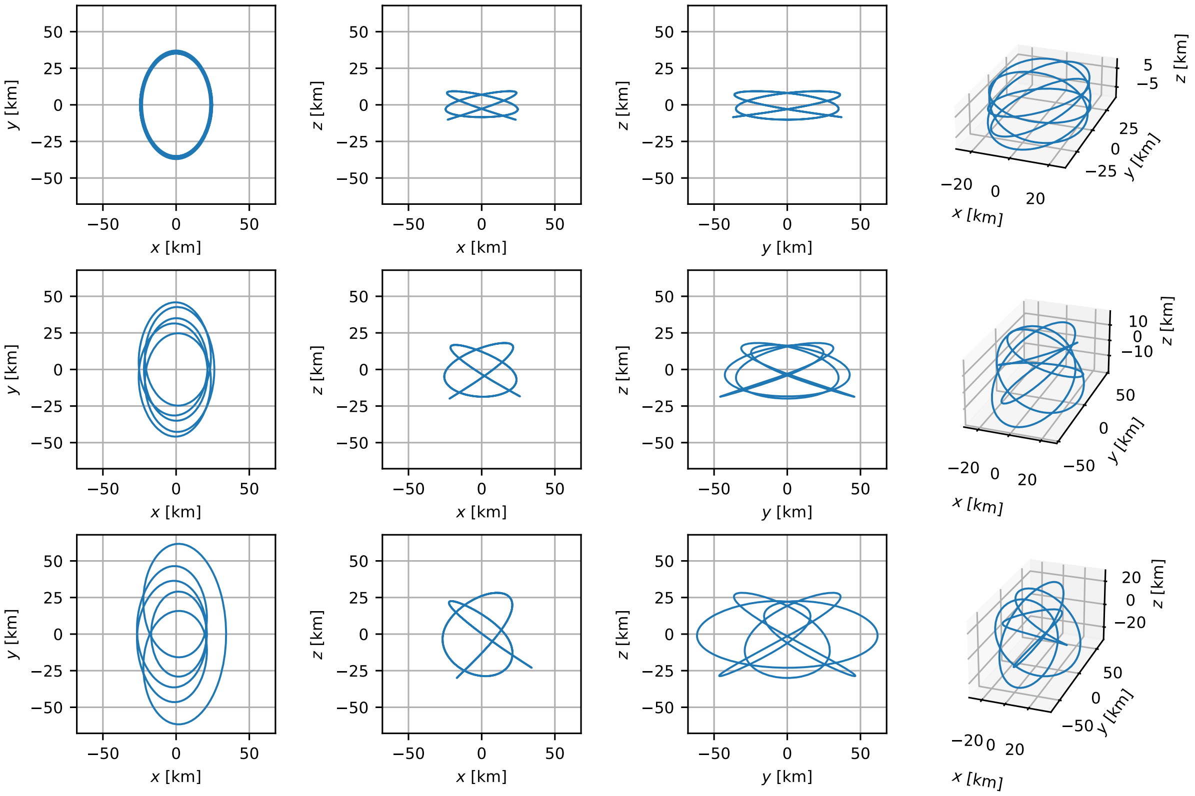

Figure 4.6: Trajectory (x, y and z coordinates) of representative orbits of the family $1:5-xz$. Each row corresponds to one orbit.

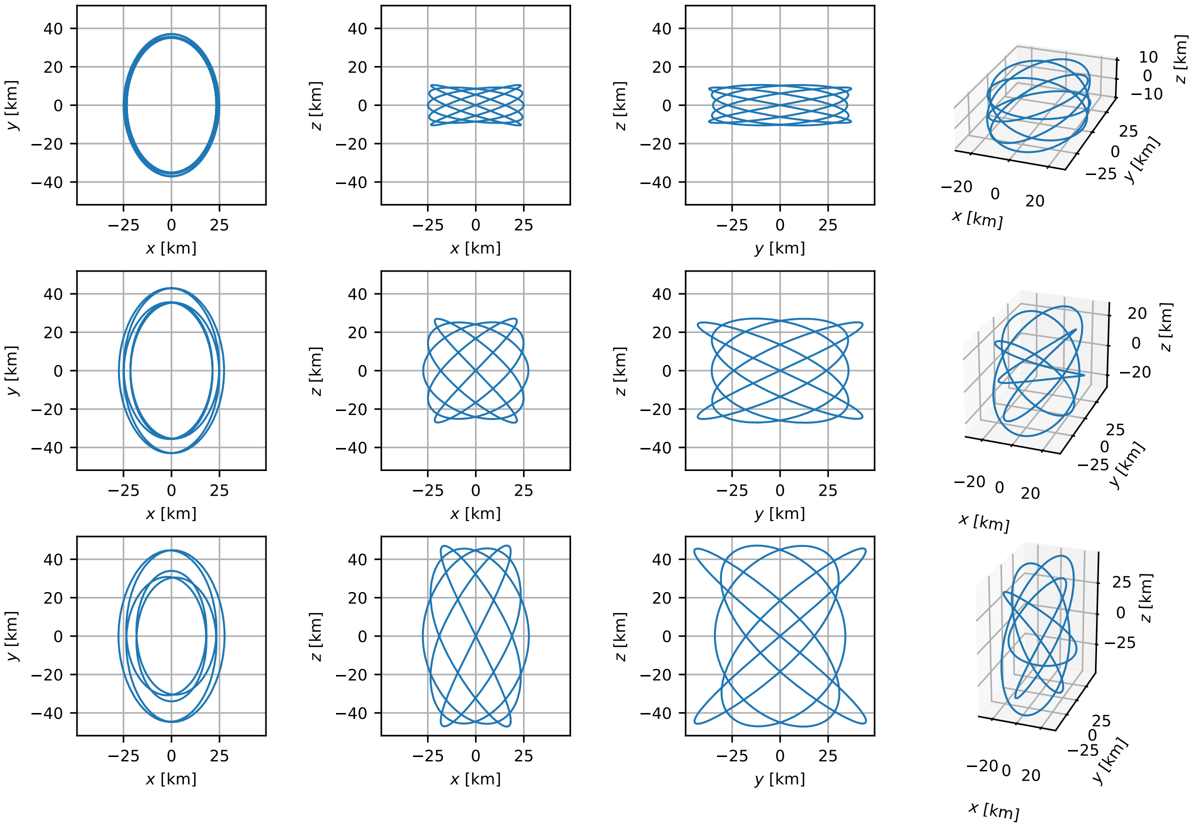

Figure 4.7: Trajectory (x, y and z coordinates) of representative orbits of the family $1:5-x$. Each row corresponds to one orbit.

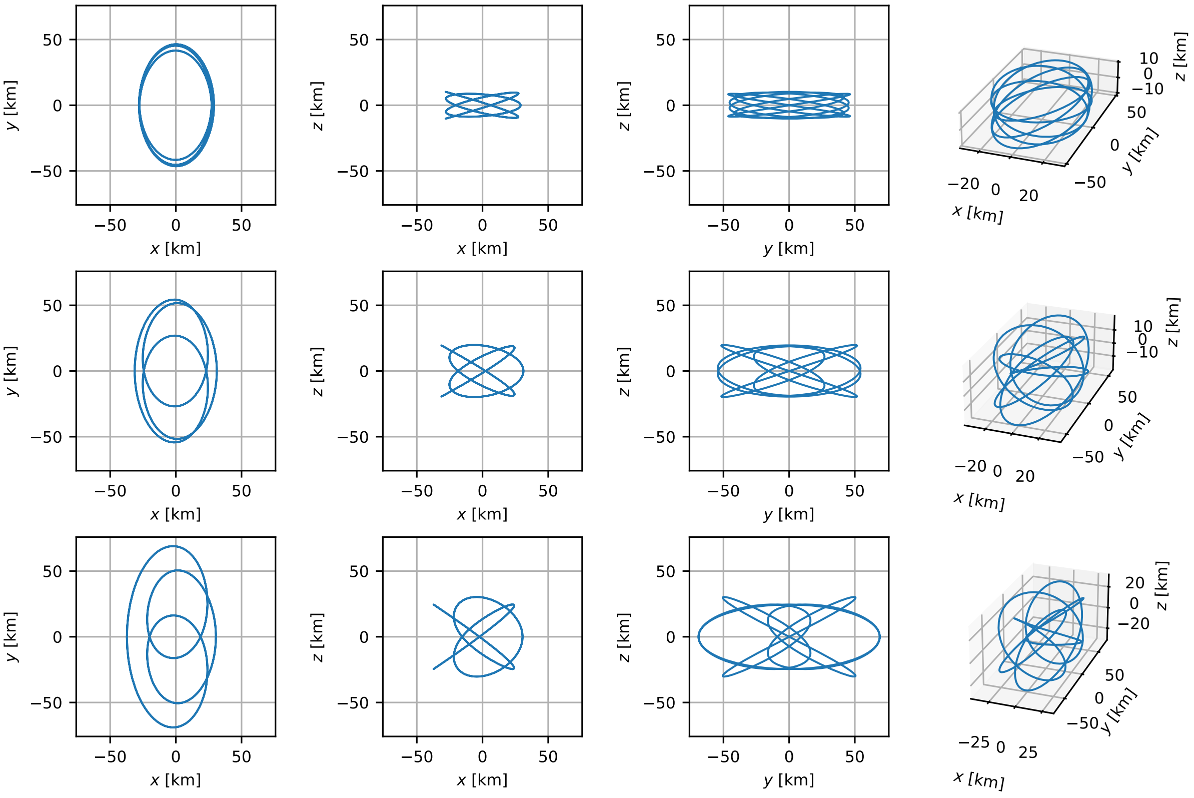

Figure 4.8: Trajectory (x, y and z coordinates) of representative orbits of the family $1:6-xz/x(BE)$. Each row corresponds to one orbit.

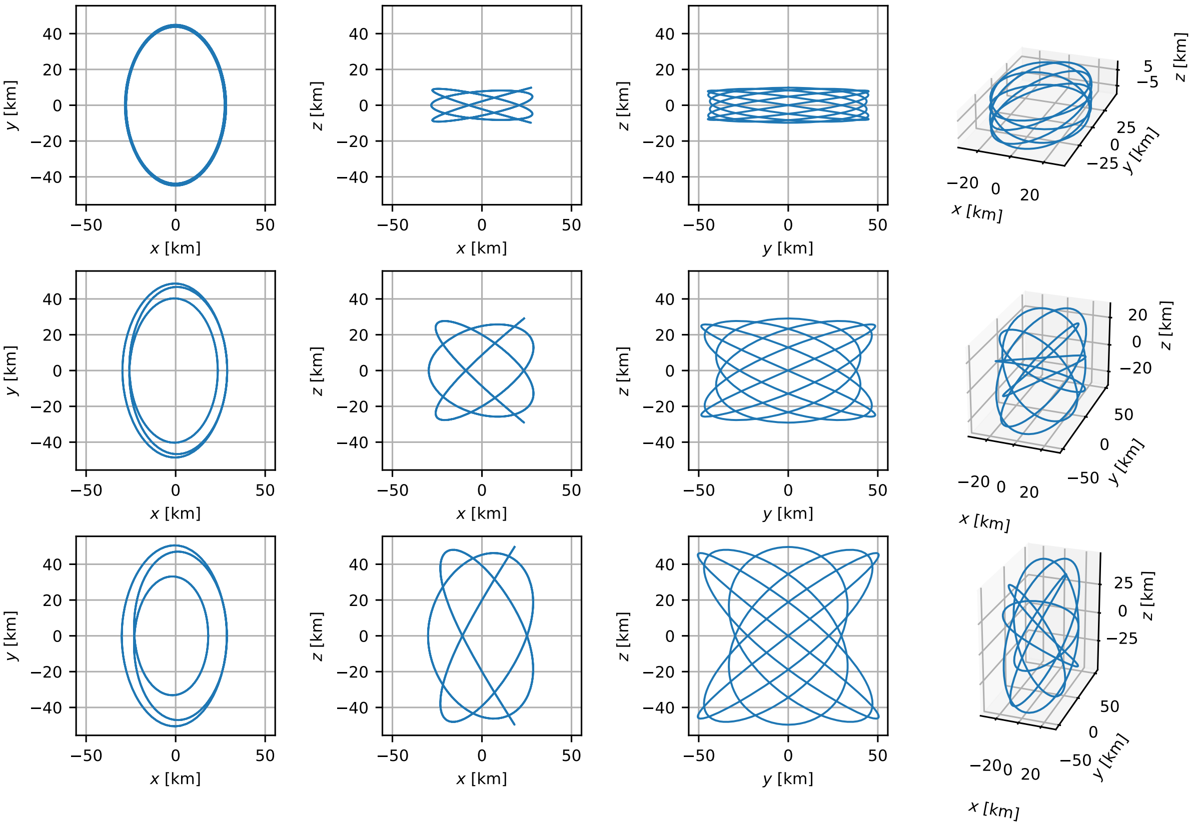

Figure 4.9: Trajectory (x, y and z coordinates) of representative orbits of the family $1:6-xz/x(R)$. Each row corresponds to one orbit.

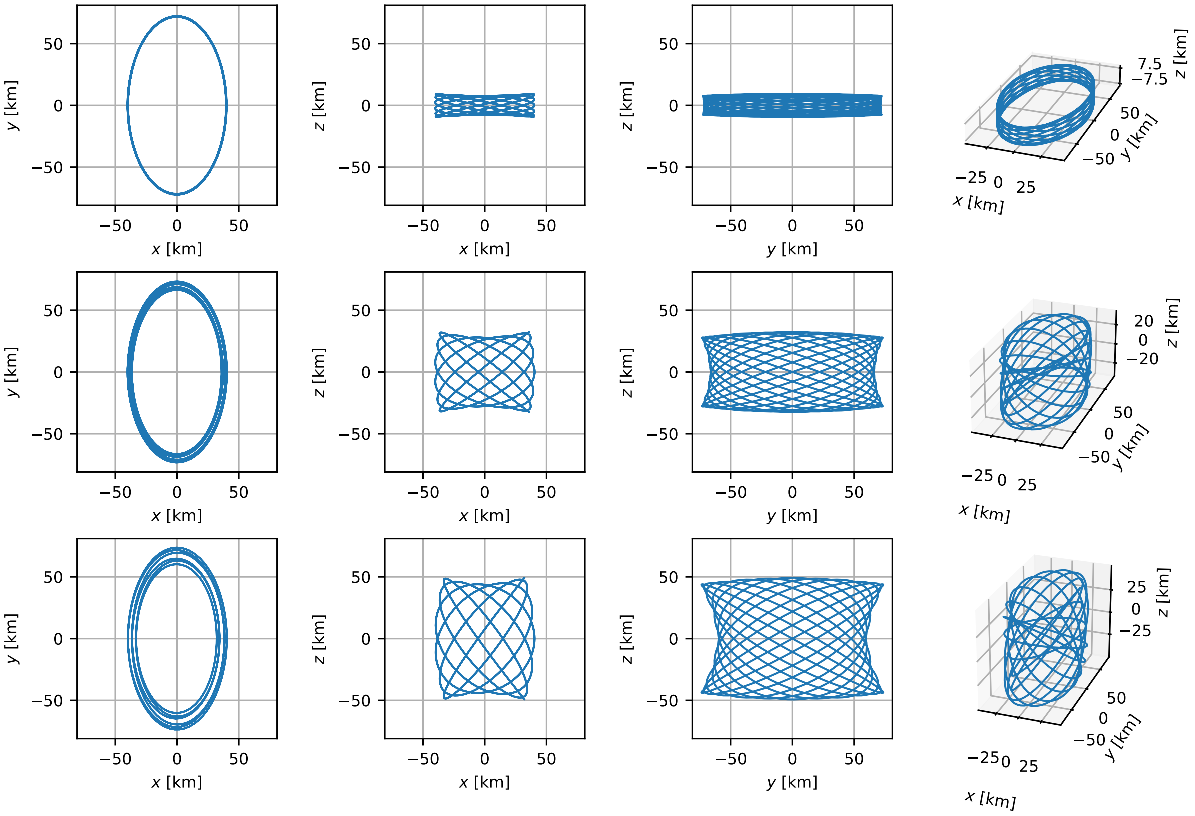

Figure 4.10: Trajectory (x, y and z coordinates) of representative orbits of the family $1:12 - xz/x$. Each row corresponds to one orbit.

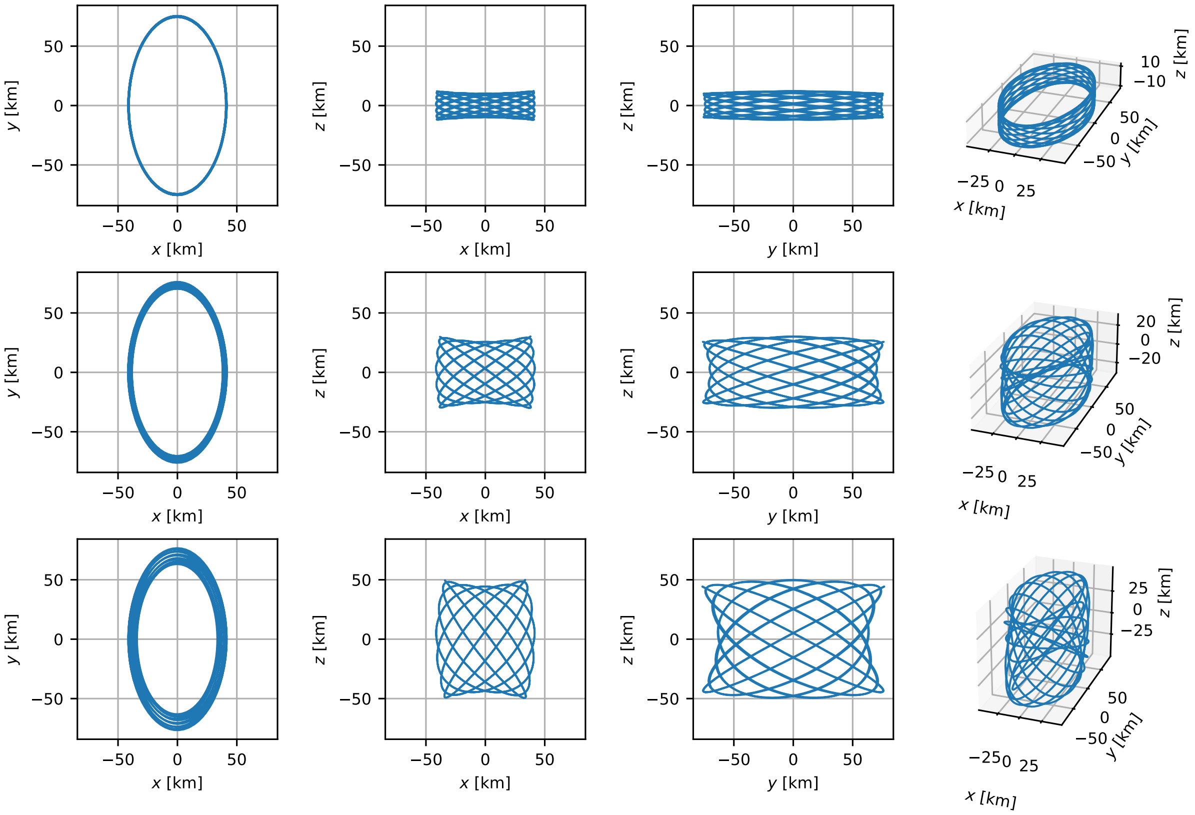

Figure 4.11: Trajectory (x, y and z coordinates) of representative orbits of the family $1:13 - xz$. Each row corresponds to one orbit.

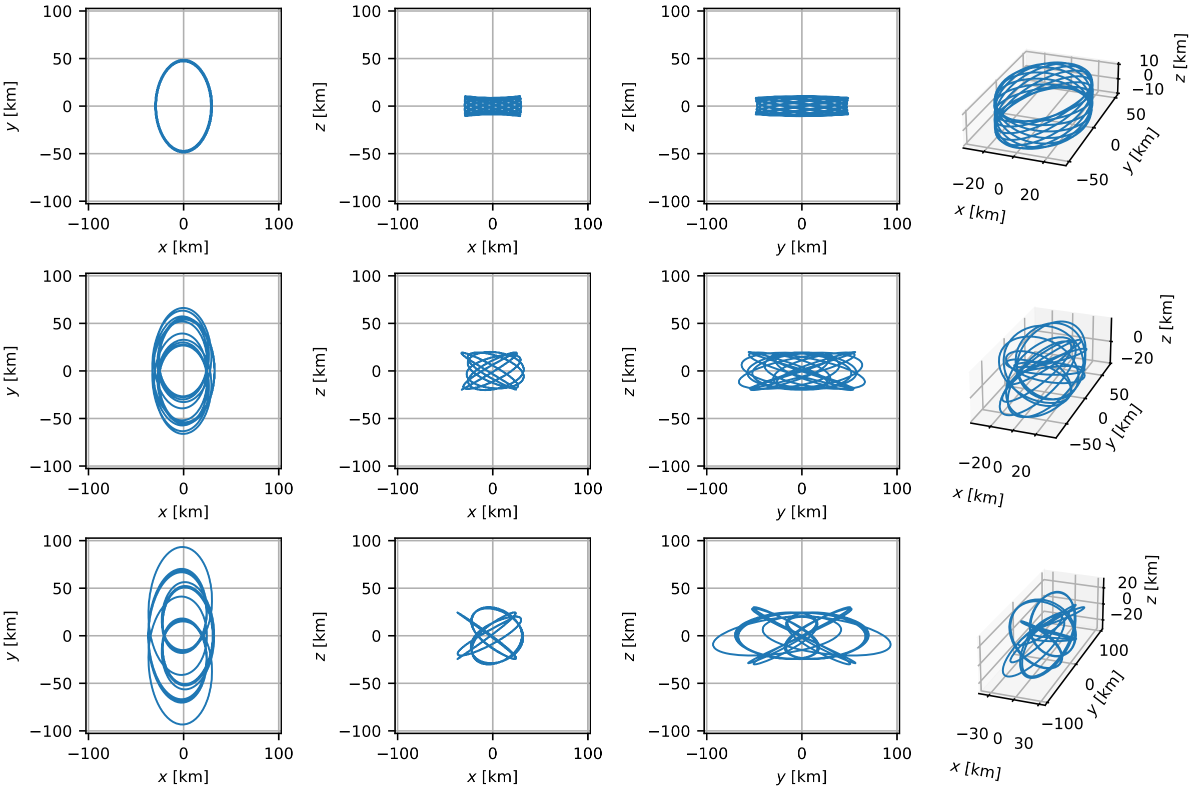

Figure 4.12: Trajectory (x, y and z coordinates) of representative orbits of the family $2:13-xz$. Each row corresponds to one orbit.

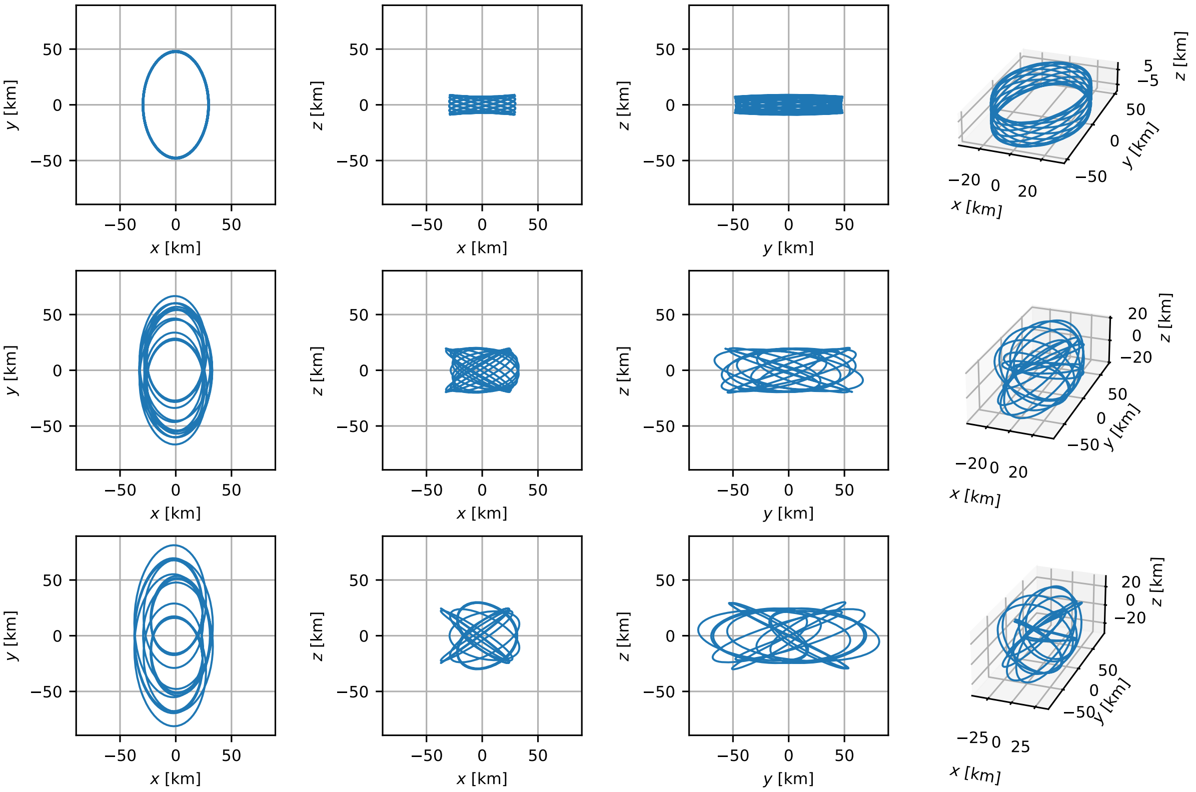

Figure 4.13: Trajectory (x, y and z coordinates) of representative orbits of the family $2:13-x$. Each row corresponds to one orbit.

4.1.3. Three-Dimensional Quasi-Satellite Orbits with Polyhedron

Having created a database of families using the CR3BP with ellipsoidal secondary, these were then continued to the CR3BP with polyhedral secondary by applying the previously described homotopy procedure to each individual three-dimensional orbit. It was opted to compute all orbits via homotopy instead of computing a single orbit per family via homotopy and then obtaining the rest via continuation (as is done by Scheeres et al., 2019) to reduce the computational time.

The families of three-dimensional QSOs computed with the polyhedron are very similar to the ones obtained with the ellipsoid, displaying the same qualitative characteristics, and only very small differences in terms of A_x, A_y, A_z, stability, and Jacobi integral. As an example, the evolution of the amplitude and minimum distance to Phobos' origin along the $1:5-xz$ family (one of the families closest to the surface, hence being more affected by changes in the gravity model) is represented in Fig. 4.14 for the two gravity models; it can be observed that the family behaves very similarly for the two models. Given the similarity of the three-dimensional QSOs between the ellipsoidal and polyhedral models, the results presented for the former are also valid for the latter. Thus, no further discussion about the different families is presented.

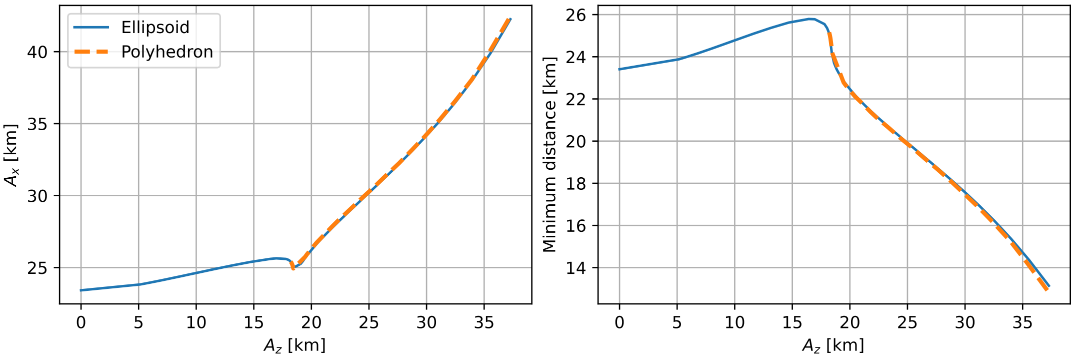

Figure 4.14: x-amplitude (left) and minimum distance to the origin of the reference frame (right) as a function of the z-amplitude, for the orbits of the $1:5-xz$ family computed using the ellipsoidal and polyhedral gravity models of Phobos. The full family is represented for the ellipsoid, while only the unstable part of the family is represented for the polyhedron.

4.2. Energy Analysis

Since energy is conserved in the CR3BP, one can in a straightforward manner obtain some conclusions about the performance of landings via manifolds, even before any propagation is executed.

The following plots, which represent the distribution of different quantities over Phobos' surface (e.g. Jacobi integral, in Fig. 4.15), are plotted using the Mollweide projection. This projection was used as it allows avoiding the strong distortion near the poles that occurs when creating Cartesian plots of the latitude versus longitude. However, it is worth noting that, contrary to the usual case of spherical bodies, here the Mollweide projection does not preserve the relative surface area, since Phobos has a shape similar to an ellipsoid. Each of these plots represents the whole surface of Phobos. The dashed lines represent latitudes of $-60°$, $-30°$, $0°$, $30°$, and $60°$ (horizontal lines) and longitudes of $-120°$, $-60°$, $0°$, $60°$, and $120°$ (curved lines). The vertical line corresponds to a longitude of $0°$ (as marked above it), which corresponds to looking at Phobos from the Mars perspective; in some instances of this report, the central longitude is instead taken to be $180°$, corresponding to looking at Phobos from the anti-Mars perspective. Even though the x-axis of the synodic frame points from Mars to Phobos, the longitude is presented using the usual convention for moons, i.e. assuming the x-axis to point from the moon to the planet and the z-axis to be aligned with the angular momentum.

The distribution of the Jacobi integral throughout Phobos' surface (i.e. the Jacobi integral at each surface point assuming the velocity to be zero) is represented in Fig. 4.15. A detailed analysis of the distribution of the Jacobi integral on Phobos has been executed by Scheeres et al. (2019), as such here only the features relevant to the design of landing manifolds are highlighted. In general, one can note that the distribution of the Jacobi integral is closely related to the surface topography (Fig. 4.16); for example, the Stickney crater, located around a longitude of $-50°$ and latitude of $0°$, can be easily identified from the Jacobi integral. The differences between the Jacobi integral and the surface topography are related to the contributions to the former besides Phobos' gravity. The most relevant point to retain is that the Jacobi integral takes larger values around the Mars and anti-Mars edges (respectively, around longitudes of $0°$ and $180°$), than around the leading and trailing edges (respectively, longitudes of $-120°$ to $-60°$, and $60°$ to $120°$). Therefore, the manifolds departing from a certain periodic orbit will have a higher impact velocity at the Mars and anti-Mars edges than at the leading and trailing edges (the manifolds departing from a given orbit all have approximately the same Jacobi integral; thus, the ones impacting points which have larger components of the Jacobi integral associated with the position will have higher impact velocities, according to Eq. 2.5).

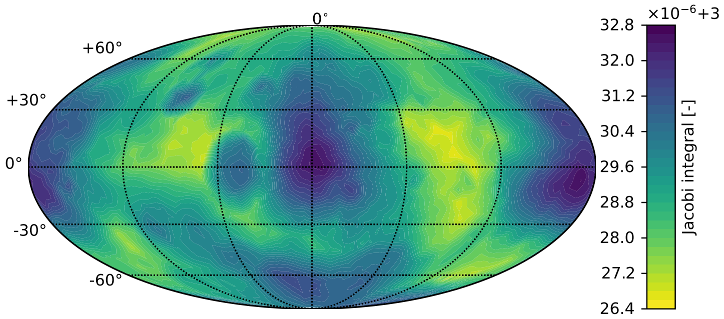

Figure 4.15: Jacobi integral at Phobos' surface, as a function of the latitude and longitude, seen from the Mars perspective.

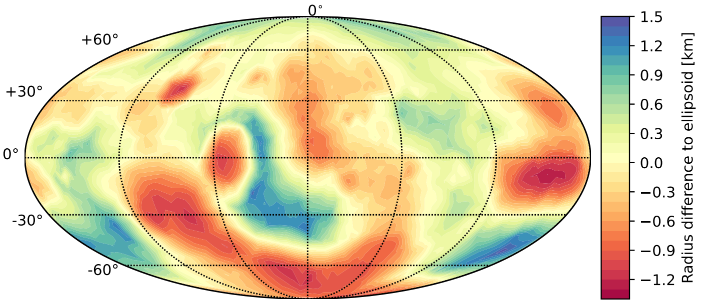

Figure 4.16: Radius difference between Phobos' surface and its mean ellipsoid, as a function of the latitude and longitude, seen from the Mars perspective.

Further conclusions can be drawn by analyzing the zero-velocity curves in the vicinity of Phobos (Fig. 4.17). Each of these curves (or surfaces, when analyzing the three-dimensional space) is associated with a specific Jacobi value and is determined by assuming the velocity at each point to be zero. Analyzing the zero-velocity curves, it is possible to observe that, for low-enough altitudes (below the location of the Lagrange points) the zero-velocity Jacobi integral increases with decreasing altitude. Thus, an impact manifold will generally increase its velocity as the altitude decreases (assuming the final portion of the manifold is approximately normal to the surface), reaching its maximum value at the impact point.

4.3. Invariant Manifolds

Having determined multiple three-dimensional QSO families in the CR3BP with a polyhedral secondary, their unstable invariant manifolds can now be initialized and propagated. For each orbit, the manifolds were propagated starting from a series of points equally spaced in time. For each revolution of each orbit, 35 starting points were used, corresponding to an angular separation between consecutive points of approximately 10°; thus, for each $d:n$ resonant orbit, a total of $35n$ manifolds were propagated. Each manifold was propagated until one of the following conditions was met: (1) impact with Phobos, (2) distance to the origin of Phobos larger than 200.0 km (corresponding to the maximum distance at which MMX will be from Phobos while in a QSO), or (3) propagation time larger than 30.0 periods of Phobos (corresponding to approximately 10 days). This propagation time was considered reasonable as, in most cases, it was large enough for the spacecraft to either impact or escape Phobos.

Since the distribution of orbits along each family is very unequal, with a larger density of orbits near bifurcations or in regions where the family curves, the manifolds were only computed for orbits for which the

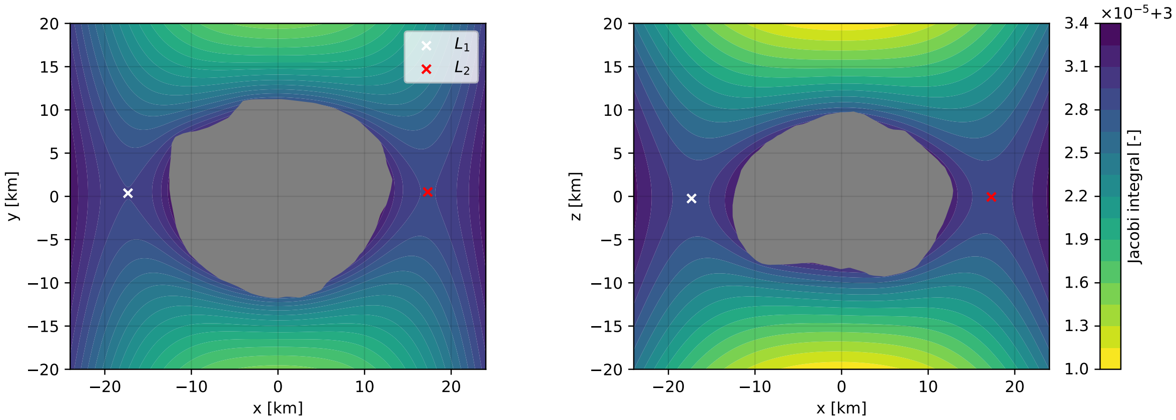

Figure 4.17: Zero-velocity curves in the vicinity of Phobos. The locations of the Lagrange points are marked.

root sum square of the difference in amplitudes with respect to the last computed orbit

$$\Delta A = \sqrt{(A_{x,i} - A_{x,i-1})^2 + (A_{y,i} - A_{y,i-1})^2 + (A_{z,i} - A_{z,i-1})^2} \tag{4.3}$$

was larger than 2.0 km. This value allows a very significant decrease in the number of tested orbits (and, consequently, in the computational time), while still allowing the reduced families to maintain the shape of the original ones.

It is important to note that while the invariant manifolds of Lagrange point orbits form well-defined tube-like structures (an example can be seen in Section 7.5), that is not the case for resonant orbits (Vaquero, 2013). As such, the invariant manifolds considered here often have a somewhat chaotic appearance, with the visualization of their trajectories, in general, not being particularly insightful. Some examples of invariant manifolds that impact Phobos' surface can be observed in Fig. 4.18, and an example of a representative altitude profile in Fig. 4.19. In general, the manifolds show an altitude profile similar to a three-dimensional QSO up to a certain point (approximately 5.5 revolutions in this case). Afterward, the minimum and maximum altitudes start to change, eventually leading to a flyby-like approach (here at approximately 6.5 revolutions), with the spacecraft reaching its maximum altitude approximately 0.75 revolutions later, and impacting the surface approximately 0.25 revolutions after that. Some irregularities can be observed in the evolution of the altitude near the regions with minimum altitude: these result from the irregularities in Phobos' surface, mainly associated with the topography of the Stickney crater.

The following sections analyze the properties of the invariant manifolds that do impact the surface, relating them to the characteristics of the families from which they originate.

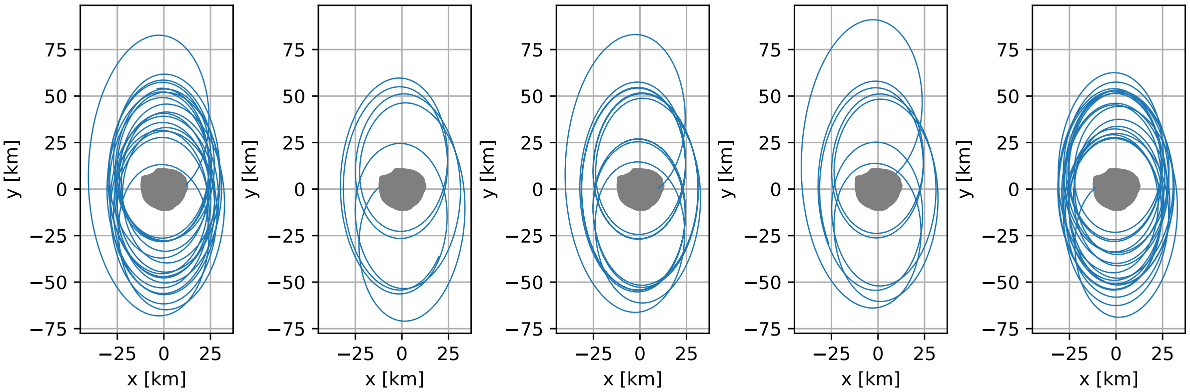

Figure 4.18: Projection onto the xy-plane of different impact manifolds originating from a single orbit of the $1:6 - xz/x(BE)$ family.

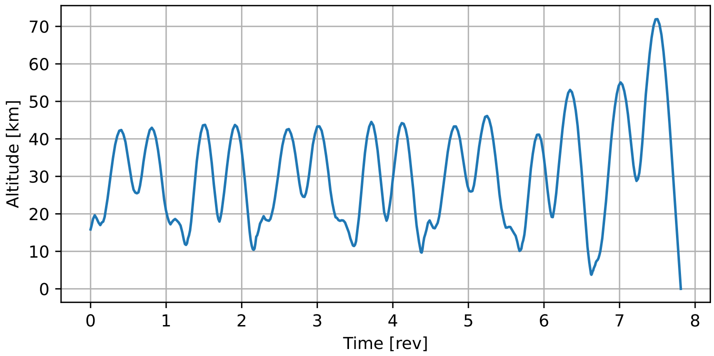

Figure 4.19: Altitude as a function of time for one of the manifolds originating from an orbit of the $1:6-xz/x(BE)$ family (the third manifold represented in Fig. 4.18).

4.3.1. Families With a Large Number of Impacts

Analysis of a Representative Family

To analyze the properties of the impact manifolds, the $1:6-xz/x(BE)$ family is first analyzed as a representative example, and (when possible) the conclusions are generalized to the remaining families.

As will be shown, impacts can occur in most surface locations, the distribution of impacts along the surface being a function of the shape of the originating orbit and of Phobos' surface geometry. The former has a particularly important influence, creating differences even for orbits in the same family. The impact patterns for the $1:6-xz/x(BE)$ family (Fig. 4.20), analyzed here, are representative of the patterns observed for most families.

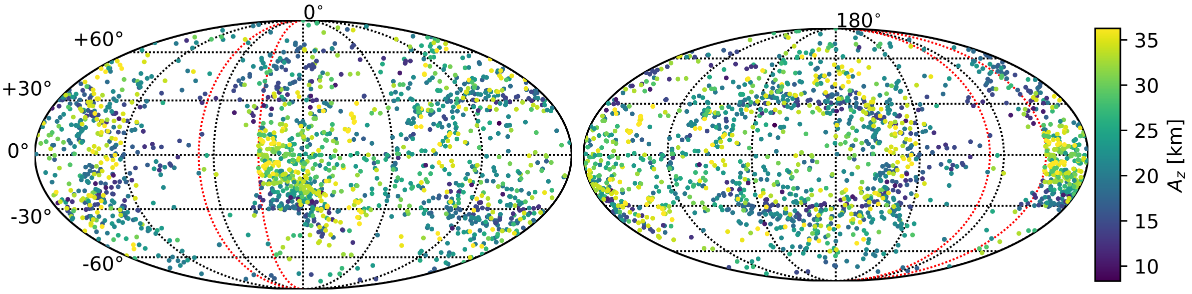

Figure 4.20: Location of impacts on Phobos' surface originating from the $1:6-xz/x(BE)$ family, with the z-amplitude of the originating orbit marked. The red lines indicate the approximate longitudes covered by the Stickney crater. The impacts are represented from the Mars (left) and anti-Mars (right) perspectives; both plots represent the full surface.

A first characteristic to note is the existence of a region for longitudes of −60.0° to −30.0° (longitudes marked by dashed red lines in the plot) and latitudes of −20.0° to 20.0° for which no impacts were registered. In fact, no impacts were observed in this region for any of the computed families of orbits. This region corresponds to the Stickney crater, having a lower altitude with respect to the surrounding area; additionally, the higher-latitude edge of the crater is continued by a region with higher altitude than the reference ellipsoid, corresponding to the crater rim (Fig. 4.16, blue region at a latitude of approximately −30.0°). As the orbits are retrograde, the mentioned features shield this region against impacts. Nevertheless, impacts do occur in the parts of the crater with the lowest longitude, i.e. between −70.0° and −60.0°.

Also noteworthy is the fact that there is a higher density of impacts in the trailing edge (longitude around 90°) than in the leading edge (longitude around −90°); this was observed for approximately half of the tested families, for the other impacts appeared equally distributed along the two edges. Again, this behavior is related to Phobos' surface geometry: the leading edge is, in most places, below the mean ellipsoid, while the trailing edge is usually above it (Fig. 4.16). Thus, trajectories passing near the trailing edge will be more likely to impact the surface than those passing near the leading edge. This type of pattern seems to be mostly associated with manifolds that approach the leading or trailing edge after having a point with maximum altitude in the field of view of the impact edge. Examples of such manifolds impacting and missing the leading edge are given in Fig. 4.21; observe that both manifolds have their maximum distance to Phobos at approximately $x = 0$ km and $y = 90$ km, and after that point approach the surface while always being in sight of the leading edge. For manifolds that approach the leading/trailing edge from another direction (e.g. from above or below Phobos), the asymmetry between the two edges is less noticeable or even absent.

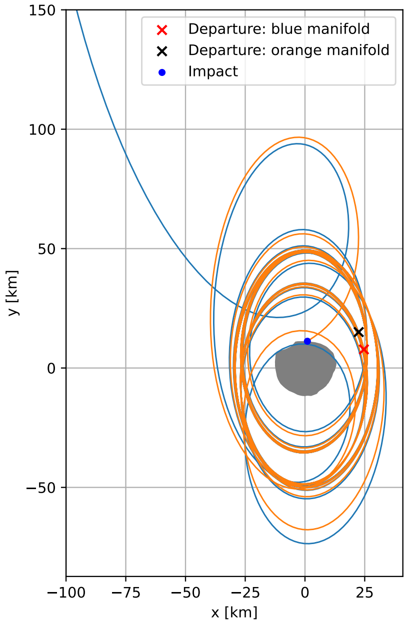

Figure 4.21: Projection of two manifolds with similar departure conditions onto the xy-plane. One manifold impacts the surface (orange line) and one does not (blue line). The departure locations of the two manifolds are marked with crosses.

A third feature to observe in Fig. 4.20 is the large concentration of impacts in the vicinity of the Mars (longitude of 0.0°) and anti-Mars (longitude of 180.0°) edges. The majority of the computed families display this pattern; the others either displayed a very uniform coverage of the surface, without any region with higher concentration, or displayed large concentrations of impacts at some other location (only observed for the 1:5-x family). The origin of the increased impacts in the Mars/anti-Mars edges is twofold: firstly, three-dimensional QSOs maintain to some extent the elliptical shape of two-dimensional QSOs, having the closest approach to the origin of the reference frame when crossing the xz-plane; secondly, Phobos has an ellipsoidal shape, with the major semi-major axis aligned with the x-axis. The combination of the two effects means that the orbits have their lowest altitudes at the Mars/anti-Mars edges. Since unstable manifolds tend to, at least initially, follow the shape of the originating orbit, their closest approach to Phobos and subsequent impact will also tend to occur at the Mars/anti-Mars edges.

Besides the higher concentration of impacts near the Mars and anti-Mars edges, one can also note some differences between these two edges: in the former, the impacts are concentrated on the edge itself, while in the latter the impacts are concentrated around the edge and extend to the trailing edge. The differences are produced by the shapes of the orbits. Near the Mars edge ($x < 0$), the periodic orbits cross the yz-plane at low altitudes ($z < 10$ km), leading the manifolds to impact the edge itself (Fig. 4.22, top, dashed lines). Meanwhile, near the anti-Mars edge ($x > 0$), the closest approach of the orbits to the surface occurs at the yz-plane crossing with high altitude ($z > 20$ km); as such, the manifolds will tend to impact further away from the edge, possibly in the trailing edge (Fig. 4.22, bottom, solid lines). The asymmetry of the impact pattern around the anti-Mars edge (between longitudes larger or smaller than 180°) is related to the part of the orbit that was being approximately followed when the impact occurred, depending on whether the impact occurs shortly after approaching the orbit's $x > 0$ yz-plane crossing (impacts at longitudes below 180°; represented in Fig. 4.22, bottom) or approximately one revolution after that (impacts at longitudes above 180°).

Finally, still analyzing Fig. 4.20, it is possible to observe that the z-amplitude of the originating orbit does not affect significantly the impact location. This is possibly related to the large range of times of flight associated with the manifolds of each orbit. For the family at hand, this range takes values between 5 and 20 revolutions, for minimum times of flight below 5 revolutions; this range allows the manifolds of a given orbit to have different behaviors depending on their times of flight, hence the z-amplitude does not produce any clear impact pattern.

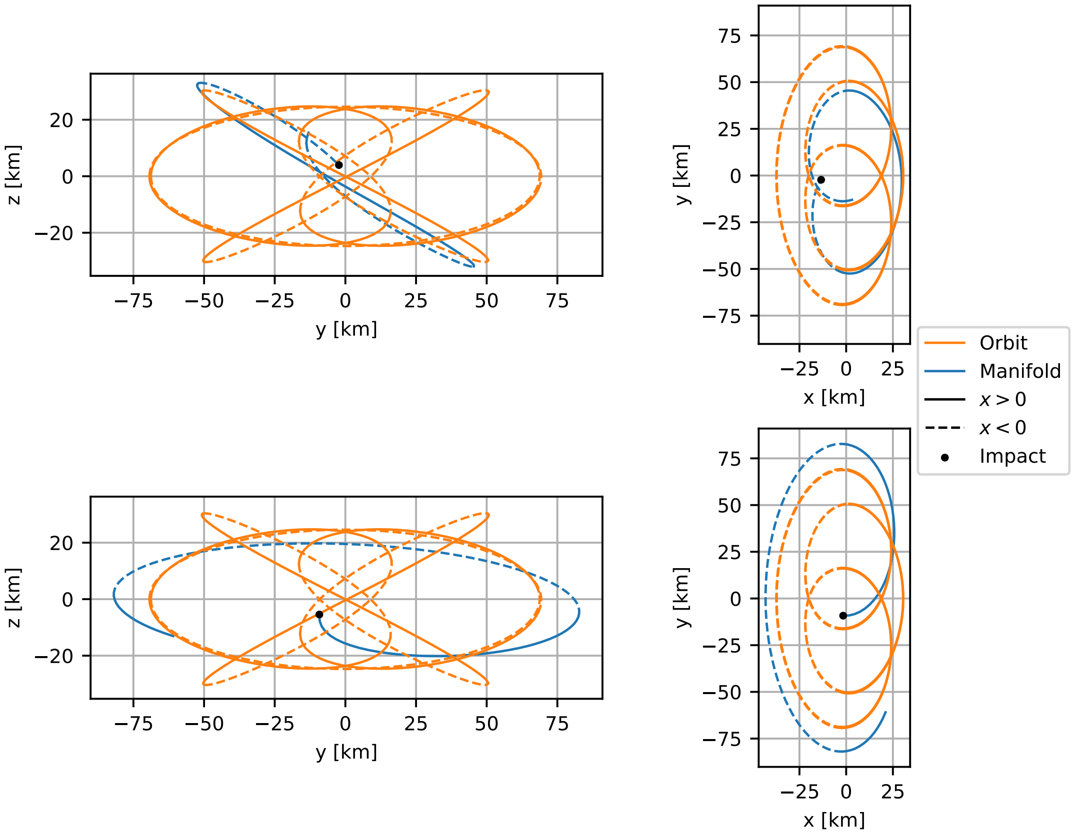

Figure 4.22: Example of two manifolds impacting the Mars edge (top) and trailing edge (bottom), along with the periodic orbit originating them. The trajectories are represented through their projections onto the yz-plane (left) and xy-plane (right). Only the final revolution of each manifold is represented.

Analyzing the impact angles (Fig. 4.23), it is possible to note that near the Mars and anti-Mars edges the impact angle generally increases for increasing longitude (recall that the impact angle is defined with respect to the surface plane, Fig. 2.9, meaning that an increase in the impact angle corresponds to a steeper impact). The increase in the impact angle results from the fact that QSOs are retrograde: when following the final portion of the manifold, impacts that occur before passing the edge (i.e. for higher longitude than the edge) will tend to curve downwards, producing larger impact angles, while impacts that occur after passing the edge (i.e. for lower longitude than that of the edge) will have to curve around the edge, producing shallower impacts. This variation of the impact angle resulting from Phobos' shape is noticeable for the majority of families with a sufficiently large number of impacts. Besides the variation with longitude, it is also possible to observe a variation with latitude: on the Mars edge, the impact angle decreases with decreasing latitude. This effect is likely also related to the local topography of Phobos, which near the Mars edge also displays an approximately diagonal variation (Fig. 4.16). One could argue that the diagonal variation might be associated with the shape of the manifold, with the impact angle varying in the direction traveled by the manifold; however, the manifolds of the considered family both cross the Mars edge for decreasing and increasing latitude (and decreasing longitude in both cases). Thus, a clear diagonal variation would not be expected, strengthening the hypothesis that the observed variations are induced by Phobos' topography. The diagonal variation is usually not as visible as the variation with longitude, and is only present for part of the tested families, indicating that the local shape of the orbit and the resulting manifolds often have a larger influence on the impact angle than the mentioned variations in Phobos' shape.

Analysis of All Families With a Large Number of Impacts

Analyzing the mean time to impact associated with the various families of orbits (Fig. 4.24, left), one observes that, in general, for increasing A_z there is a decrease in the time to impact. Since three-dimensional QSOs become more unstable for higher A_z (Fig. 4.4, left), unstable manifolds will depart faster from the originating orbits as A_z increases, thus impacting the surface or escaping the system faster. Additionally, for most families,

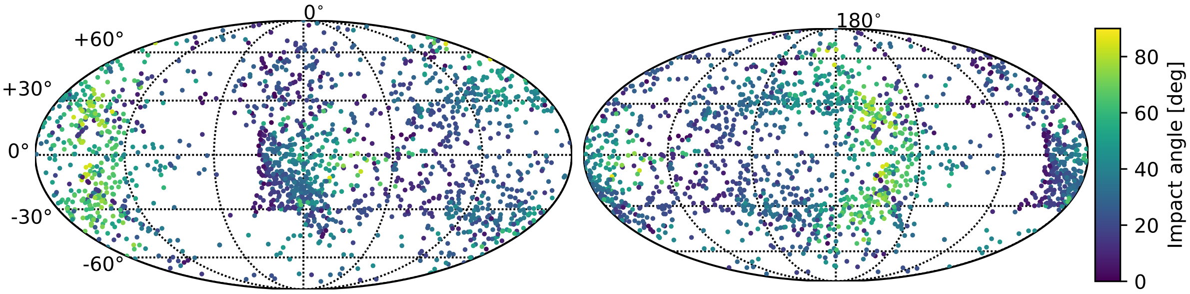

Figure 4.23: Impact angle as a function of the impact location (longitude and latitude) on Phobos' surface, for manifolds originating from the $1:6 - xz/x(BE)$-family. The impacts are represented from the Mars (left) and anti-Mars (right) perspectives; both plots represent the full surface.

the minimum distance to the surface decreases as A_z increases, also contributing to faster impacts. It is also possible to note that some families have the lowest times to impact for low A_z values; these correspond to orbits with very small numbers of impact manifolds, the associated propagation times not being representative.

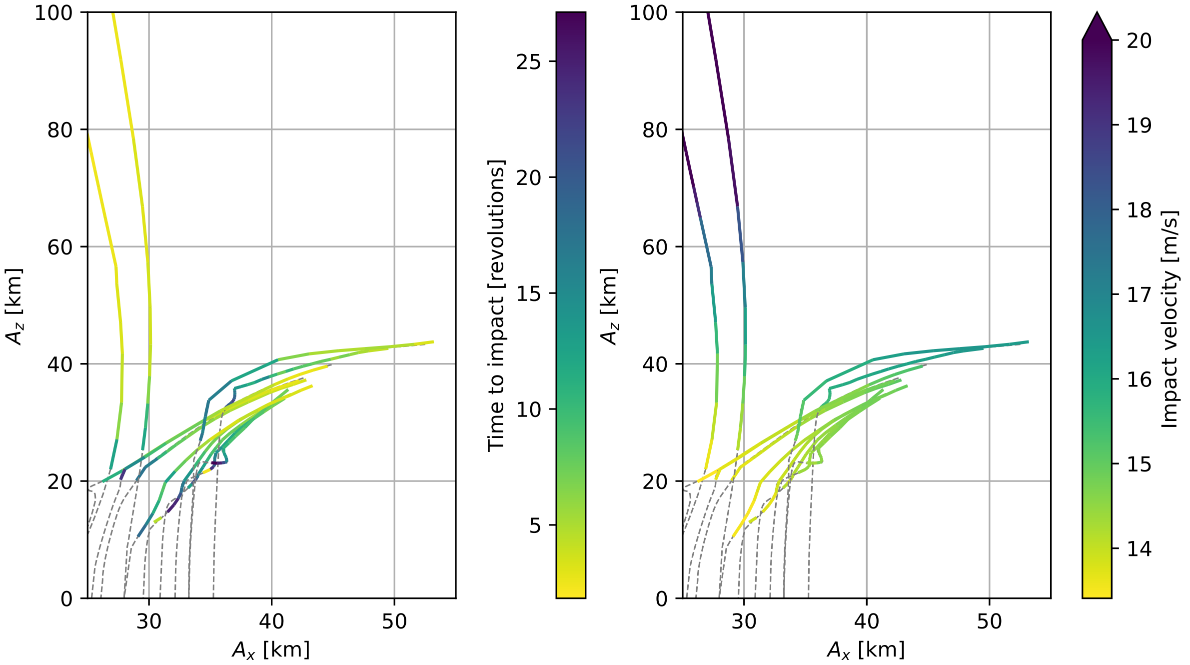

Figure 4.24: Mean propagation time (left) and mean final velocity (right), for the impact manifolds originating from families with large numbers of impact manifolds, as a function of the x-amplitude and z-amplitude of the originating orbit. The dashed lines correspond to the portions of the families that either are stable or did not produce impact manifolds.

It is possible to notice that the impact velocity increases as the z-amplitude of the initial periodic orbit increases (Fig. 4.24, right). This is simply a result of the conservation of energy. As previously discussed, as the z-amplitude increases along a family of orbits, the orbits' Jacobi integral decreases (Fig. 4.5), thus leading to higher impact velocities. For each orbit, the impact velocities also vary with the impact location: impacts at points with higher surface Jacobi (e.g. Mars and anti-Mars edges) have higher velocities than those at points with lower surface Jacobi (e.g. leading and trailing edges); c.f. Fig. 4.15. The variation of the impact location was observed to produce maximum variations in the impact velocity in the order of $1\,\mathrm{m\,s^{-1}}$. The minimum observed impact velocity is approximately $13\,\mathrm{m\,s^{-1}}$, significantly higher than the requirements for the horizontal and vertical impact velocities (respectively, $0.15\,\mathrm{m\,s^{-1}}$ and $1.0\,\mathrm{m\,s^{-1}}$).

Finally, it is worth analyzing the surface coverage of the different types of orbits, which could be relevant for missions that require some flexibility in the selection of the impact location. For example, ESA's Phobos Sample Return mission required access to 20% of the surface as a minimum, with a goal of 50% (Pickering, 2014). To compare the surface coverage properties across the different families, the percentage of the surface area covered by impacts and the number of manifolds leading to impacts per revolution of the originating orbit were computed. The percentage of the surface covered by impacts was estimated by dividing the surface into hexagons and verifying in which hexagons impacts had been detected. This whole process was executed using the location of the impacts with respect to the Mollweide projection. The size of the used hexagons was tuned so that visually it allowed an accurate representation of the regions with impacts; 1068 hexagons were used, each corresponding to, on average, approximately 0.1% of the surface area. As previously mentioned, the Mollweide projection does not preserve the relative area in the case of Phobos, due to its ellipsoidal shape (the Mollweide projection only preserves the area for spherical bodies). Therefore, the estimated surface area with impacts does not coincide with its true value. However, since Phobos' semi-major axes have relatively similar dimensions, it was considered that the estimated area should still be a reasonable approximation of the truth, thus allowing a valid comparison between different orbits. A more accurate estimate of the surface area covered by impacts could be obtained via an area-preserving projection of a triaxial ellipsoid, according to Nyrtsov et al. (2015). Do note that the process of comparing the covered surface area between families is biased in favor of families with more revolutions, since more revolutions means more manifolds, potentially increasing the covered surface area. Nevertheless, the results do not appear to be too heavily influenced by this factor, as no clear trend favoring orbits with a higher number of revolutions is observed.

The mean percentage of the surface area covered by impacts and the mean percentage of manifolds leading to impacts per revolution of the originating orbit are listed in Tab. 4.1, for each of the families being analyzed (the mean is taken across the orbits of each family). It is possible to note that, in general, the number of impacts per revolution tends to decrease for increasing A_x of the bifurcating orbit. A particularly noticeable exception to both trends is the $1:8-xz/x(BE)$ family, which appears to offer a poor geometry for producing impact manifolds: most impacts occur for a small number of the tested orbits, and are located in small regions of the surface. The d value of the $d:n$ families does not appear to significantly influence the number of impacts per revolution, but does influence the surface coverage, with the $d \geq 2$ families being able to reach larger parts of the surface. As previously discussed, these families produce more irregular orbits, covering a larger region of space; as a result, their manifolds also tend to cover a larger region of space, consequently being able to reach a larger portion of the surface.

Table 4.1: Mean percentage of Phobos' surface covered by impacts and mean percentage of impacts per revolution (from a total of 70 manifolds per revolution) for the families with a large number of surface impacts (mean taken over the values of each orbit); for each family, its name and the A_x value for the bifurcating orbit are listed. The table is divided into three blocks, corresponding to (from top to bottom): $1:n$ families that go up to high A_z, $1:n$ families that do not go up to high A_z, $2:n$ families, and $3:n$ family (neither go up to high A_z).

Family	Bifurcating A_x [km]	Surface area [%]	Impacts per revolution [%]
$1:5-x$	23.4	3.5	12.6
$1:6-xz/x(R)$	28.0	3.3	11.0
$1:5-xz$	23.4	5.7	23.0
$1:6-xz/x(BE)$	28.0	6.8	22.3
$1:7-xz$	30.9	5.1	13.8
$1:8-xz/x(BE)$	33.2	1.8	5.5
$1:9-xz$	35.3	3.6	7.0
$2:11-x$	26.1	11.5	21.0
$2:13-xz$	29.5	11.5	18.8
$2:15-xz$	32.1	10.4	13.6
$3:16-xz/x(BE)$	25.3	12.8	17.0

Comparing the $1:n$ families that grow up to high A_z ($1:5-x$ and $1:6-xz/x(R)$) and the ones that do not, it is clear that the latter offer a larger number of landing opportunities per revolution and better surface coverage. This is likely a result of the fact that the two families that do grow to high A_z always have a shape similar to an elliptic cylinder, not changing their minimum distance to Phobos significantly throughout the family. Meanwhile, the low-A_z families have a more irregular shape, allowing closer approaches to the surface, which leads the manifolds to more easily impact the surface. Still analyzing the $1:n$ families, the optimum in terms of surface coverage and landing opportunities appears to occur for the $1:6-xz/x(BE)$ family. Likely it offers the best balance in terms of distance to the surface and number of revolutions: a larger distance to the surface decreases the number of impacts per revolution, but a larger distance to the surface is associated with

$1:n$ families with larger n, which cover a larger region of space, thus in principle also being capable to reach a larger region of the surface.

The evolution of the percentage of covered surface area and the percentage of manifolds leading to impacts along some representative families are represented in Fig. 4.25. The covered surface area and the number of manifolds leading to impacts display a very similar behavior, as one would expect. For the represented families that do not grow to high A_z ($1:6-xz/x(BE)$, $2:13-xz$, $3:16-xz/x(BE)$), as A_z increases, the surface coverage and number of impacts initially increase, reach a maximum value, decrease, and finally display either a plateau or a small increase. Although the reasons explaining this behavior are not completely clear, some potential explanations can be suggested. The initial increase and then decrease are likely related to the evolution of the stability within each family. For weak instabilities (which occur at low A_z values), a large number of manifolds remains in the vicinity of Phobos; increasing the instability, those manifolds will tend to deviate faster from the originating orbit, either escaping or impacting the surface, thus increasing the number of impacts and the covered surface area. As previously discussed, manifolds tend to generally follow their originating orbit, with surface impacts being triggered by a flyby-like approach to the surface, which happens near the regions where the originating orbit has a low altitude. It is hypothesized that further increasing the instability, making the manifolds depart even faster from the originating orbit, will create more situations where the manifolds deviate from the orbit before the latter's close approach to the surface, therefore not triggering an impact. The final increase or plateau in the percentage of impacts is probably related to the orbits' decreasing distance to the surface. Since the orbits pass very close to the surface, even if the manifolds do depart faster from the orbit, in a large number of cases the manifolds will still pass close enough to the surface for an impact to be triggered. Note that the mentioned effects are all present over the entirety of each family, simply having different relative importance depending on the region of the family.

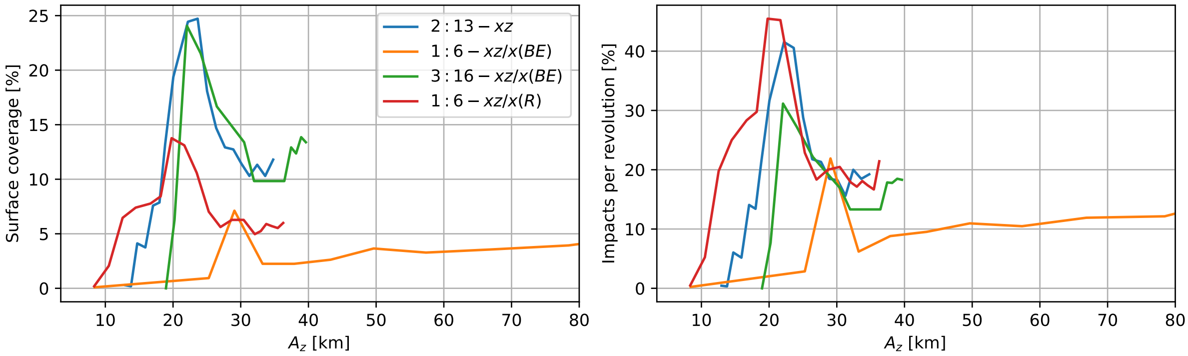

Figure 4.25: Percentage of Phobos' surface covered by impacts (left) and percentage of manifolds leading to impacts per revolution of the originating orbit (right) as a function of the z-amplitude. The results are presented for some representative families with large numbers of impacts; each curve corresponds to one family.

Looking at the family that reaches high-A_z values ($1:6-xz/x(R)$ family): the slightly different shape in the evolution of the surface coverage and number of impacts with respect to the low-A_z families should be associated with the also different shape of the orbits. However, the evolution is likely explained by the same factors. Firstly, there is a peak in the surface coverage for the "ideal" instability (high enough for the number of trajectories that do not escape or impact the surface to be low, but low enough such that the manifolds do not deviate from the originating orbit too fast); secondly, there is an increase in the surface coverage as a result of the decreasing distance of the orbit to the surface.

Analyzing the surface-coverage values, it is possible to note that the $2:13-xz$ and $3:16-xz/x(BE)$ families show larger values than the $1:6-xz/x(BE)$ one; as mentioned, this results from the different number of revolutions executed by the orbits of each family (the $2:13$ and $3:16$ orbits execute a significantly larger number of revolutions than the $1:6$ ones). It is also possible to observe that the $1:6-xz/x(BE)$ family would not be able to meet the 20% requirement specified for ESA's Phobos Sample Return mission (Pickering, 2014). However, it is worth mentioning that, firstly, the final impact location could be controlled via some low-altitude maneuver, which would allow increasing the accessible area. Secondly, Pickering (2014) does not analyze the surface accessibility from a trajectory-design point of view, meaning that it is not clear whether the Phobos Sample Return mission would be able to satisfy the 20% requirement.

4.3.2. Families With (Almost) No Impacts

As discussed, a large part of the families with low A_x produce a large number of manifolds that impact the surface. However, this is not the case when one increases the A_x of the bifurcating two-dimensional QSO. For the $1:n,\ n \geq 11$ families (which appear at bifurcations with $A_x \geq 40.0$ km) the number of impacts is negligible. For low-A_z members of these families, part of the manifolds escaped and part remained in the vicinity of Phobos (without impacting it); as A_z increases (and the instability increases; Fig. 4.4, left) the number of manifolds that remain close to Phobos quickly decreases, with almost all manifolds escaping (Fig. 4.26). For each orbit of these families, the number of manifolds that impact the surface is usually under 5 (recall that $35n$ manifolds were computed for each family). Additionally, the origins of these impact manifolds are usually spaced along the orbit, with no two consecutive manifolds impacting the surface (an example is presented in Fig. 4.27), indicating that those would not constitute a robust landing solution. The lack of impacts is likely related to the shape of the orbits in these families. As already mentioned, these are orbits that are always well contained on the surface of an elliptical cylinder, with the dimension of its base (i.e. A_x and A_y) remaining approximately constant throughout each family. Thus, the minimum distance to the surface also remains approximately constant, always having a relatively large value. As such, given the weakness of Phobos' gravity and its small size, the unstable manifolds originating from these orbits depart in chaotic directions which in general do not intersect the surface.

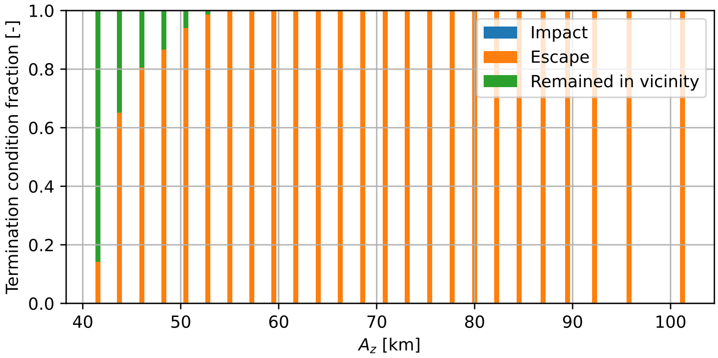

Figure 4.26: Fraction of unstable manifolds which impacted the surface, escaped or remained in the vicinity of Phobos at the final propagation time, as a function of the z-amplitude of the originating orbit, for the $1:14 - xz/x$ family. The fraction of impact trajectories (blue) is represented but is not clearly visible due to their very small number.

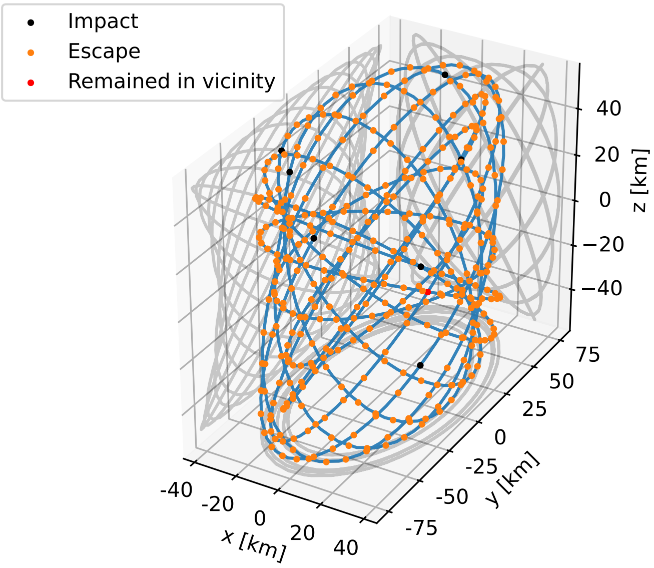

Figure 4.27: Trajectory (x, y, and z coordinates) of orbit of the $1:12 - xz/x$ family. The departure location of each invariant manifold is marked with a circle, colored according to the manifold's termination condition (impacting the surface, escaping, or remaining in the vicinity of Phobos).

4.3.3. Conclusions

Finally, some overall conclusions can be presented. Regarding the number of impacts and surface coverage, the main point is: the closer a family is to the surface, the larger the number of impact manifolds that its orbits will be able to generate. Firstly, this means that, for a given $d : n$ resonance, the low-A_z irregular orbits are more suitable to generate impacts than their high-A_z cylinder-like counterparts, as the former get closer to the surface. Secondly, for a given shape of orbits (irregular versus cylinder-like), the lower the A_x of the bifurcating orbit, the more impacts the family will be able to generate. Thirdly, and considering just the cylinder-like families: these are able to produce large numbers of impacts, but only for low-enough A_x; as A_x increases, the number of impacts quickly decreases, being negligible for $n \geq 11$ (and $d = 1$). The poor performance of the $n \geq 11$ families is likely related to the weakness of Phobos' gravity field; potentially, these offer larger numbers of impact manifolds in systems where the secondary has a stronger gravity. Fourthly, and considering just the irregular families: all of these were able to generate a large number of impacts; however, it is also worth noting that all of them have relatively low A_x values, thus, it is possible their performance would degrade if families with larger A_x were tested. Finally, all families display a small set of orbits with the best properties in terms of the number of impacts and surface coverage; the location of these orbits along each family is a function of the evolution of the stability and the minimum altitude along the family.

From the tested manifolds, minimum impact velocities in the order of 13 $\mathrm{m\,s^{-1}}$ were determined, which is significantly higher than the specified impact-velocity requirements. Therefore, it can be concluded that purely ballistic trajectories are unsuitable for the landing of a spacecraft, with the application of some kind of braking maneuver being essential. The design of such is studied in the following chapter. It is worth noting, however, that the observed impact velocities could be acceptable in some situations; for example, also studying landings on Phobos, Çelik et al. (2019) consider maximum impact velocities of 15 $\mathrm{m\,s^{-1}}$ for the deployment of small scientific payloads. In this case, it would be possible to execute a landing using the invariant manifolds without any kind of maneuver (except for the insertion into the manifold itself). However, for such high values of the impact velocity, other problems become important, namely the possibility of bouncing off the surface, which is out of the scope of this work.

5

Maneuver Design

The usage of unstable invariant manifolds allows reaching the surface of Phobos, but not with velocities satisfying the specified requirements (maximum vertical and horizontal velocities of, respectively, $1.0\ \mathrm{m\,s^{-1}}$ and $0.15\ \mathrm{m\,s^{-1}}$). Thus, the execution of a braking maneuver is required. The application of two impulsive ΔVs is assumed: one initial maneuver executed at the originating periodic orbit, to insert the spacecraft into the desired unstable manifold, and one braking maneuver, executed when approaching the surface.

The initial maneuver is determined by applying a multiple-shooting procedure (similar to Subsection 3.2.2) with the initial position constrained to be the position along the periodic orbit from which the initial state of the unstable manifold was created, and the final position constrained to be the final position of the unstable manifold; applying this procedure, one obtains a slight modification of the unstable manifold which departs from the periodic orbit (without any position discontinuity) and impacts the surface of Phobos. With this methodology, mean and maximum departure ΔVs in the order of $3 \times 10^{-2}\ \mathrm{m\,s^{-1}}$ and $2 \times 10^{-1}\ \mathrm{m\,s^{-1}}$ were found (negligible when compared with the ΔVs applied during the braking maneuver, as will be seen).

This chapter discusses the design of the braking maneuver. This is done in two phases: firstly, the maneuver is selected semi-analytically, based on energy considerations; secondly, it is optimized. A minimum maneuver altitude of 40 m is considered, corresponding to the minimum required free-fall altitude with an additional safety factor associated with the position uncertainty (Chapter 6).

All maneuvers are represented with respect to a TNW reference frame, the unit vectors being defined as

$$\hat{\boldsymbol{t}} = \frac{\dot{\boldsymbol{r}}}{\|\dot{\boldsymbol{r}}\|} \qquad \hat{\boldsymbol{n}} = \hat{\boldsymbol{w}} \times \hat{\boldsymbol{t}} \qquad \hat{\boldsymbol{w}} = \frac{\boldsymbol{r} \times \dot{\boldsymbol{r}}}{\|\boldsymbol{r} \times \dot{\boldsymbol{r}}\|} \tag{5.1}$$

This reference frame is spacecraft-centered, with $\hat{\boldsymbol{t}}$ parallel to the velocity vector, $\hat{\boldsymbol{w}}$ parallel to the pseudo angular-momentum vector, and $\hat{\boldsymbol{n}}$ forming a right-handed frame.

To decrease the computational time, a reduced set of invariant manifolds was used for the maneuver design. For each originating periodic orbit, only manifolds belonging to blocks of at least three consecutive impact manifolds were considered, as those were expected to produce more robust trajectories than isolated impact manifolds. In each block, only the manifold in the middle of the block was selected for further analysis. For example, the orbit presented in Fig. 4.27 does not display any consecutive impacts manifolds, therefore no manifold of this orbit was used for the maneuver design. The described procedure resulted in a total of approximately 1200 manifolds.

5.1. Semi-Analytical Maneuver

Since the energy is conserved in each unpowered part of the trajectory, it is possible to determine, in a straightforward manner, a quasi-optimal ΔV to apply in each manifold that allows reaching the surface with the desired velocity. This has the advantage of being much faster than the optimization, but only allows targeting the final velocity magnitude, not its components (horizontal and vertical). A target final velocity of $\sqrt{1.0^2 + 0.15^2} \approx 1.01\ \mathrm{m\,s^{-1}}$ was considered. The semi-analytical design of the maneuver also allows gaining insight into the behavior of the subsequent optimization.

5.1.1. Methodology

The goal of the maneuver is to reduce the impact velocity, which corresponds to increasing the Jacobi integral (i.e. decreasing the energy), as efficiently as possible. Consider the velocity before and after the application of the maneuver to be represented, respectively, by

$$\begin{aligned} \boldsymbol{V}_- &= [V_T, 0, 0]^T \\ \boldsymbol{V}_+ &= [V_T + \Delta V_T, \Delta V_N, \Delta V_W]^T \end{aligned} \tag{5.2}$$

Note that, by definition, the velocity before the maneuver only has a component in the T-direction. Using Eq. 2.6, the variation of the Jacobi integral associated with the maneuver is

$$\Delta C = C_+ - C_- = V_-^2 - V_+^2 = 2V_T\Delta V_T - \Delta V_T^2 - \Delta V_N^2 - \Delta V_W^2 \tag{5.3}$$

Since only the ΔV_T component allows increasing the Jacobi integral, a maximum increase of the Jacobi integral is obtained when applying the maneuver in the direction of the velocity, thus $\Delta V_N = \Delta V_W = 0$. Therefore, ΔC takes the value

$$\Delta C = 2V_T\Delta V_T - \Delta V_T^2 \tag{5.4}$$

This corresponds to a parabola with zeros at $\Delta V_T = 0$ and $\Delta V_T = -2V_T$ and a maximum of $\Delta C = V_T^2$ (Fig. 5.1, orange curve).

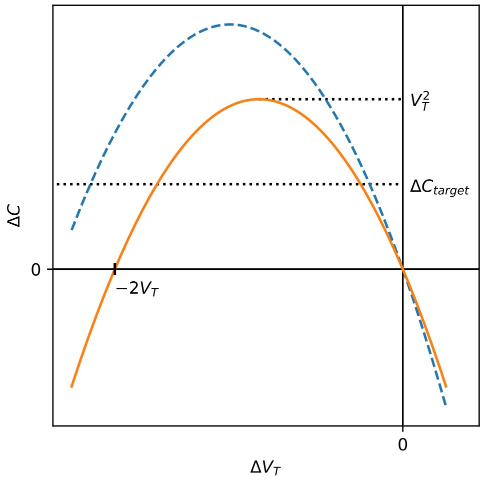

Figure 5.1: Change in Jacobi integral as a function of the applied ΔV_T, for two values of V_T. The blue curve corresponds to the application of ΔV_T at a higher V_T than for the orange curve.

Assuming the minimum change in the Jacobi integral that allows meeting the desired final velocity to be $\Delta C_{target} > 0$, then one needs to apply a maneuver

$$\Delta V_T = -V_T \pm \sqrt{V_T^2 - \Delta C_{target}} \tag{5.5}$$

Assuming the post-maneuver velocity to be in the same direction as the pre-maneuver velocity (so as to continue the motion toward the surface), the maneuver to apply is

$$\Delta V_T = -V_T + \sqrt{V_T^2 - \Delta C_{target}} \tag{5.6}$$

To reduce the required ΔV_T, the maneuver should be applied at the highest possible V_T, as can be noticed by comparing the orange and blue curves in Fig. 5.1 (the latter has a higher V_T, thus requiring a lower ΔV_T to meet an arbitrary $\Delta C_{target} > 0$). As previously discussed, the component of the Jacobi integral associated with the position increases with decreasing altitude, leading to an increase in the velocity (Section 4.2). Therefore, in general, for maximum efficiency, the maneuver should be applied at the minimum allowed altitude; nevertheless, in terms of implementation, it is always checked whether the point with maximum velocity indeed coincides with the lowest altitude.

The desired change in Jacobi integral ΔC_{target} corresponds to the difference between the desired final Jacobi integral and the Jacobi integral of the unpowered manifold. The desired final Jacobi integral is computed using the desired final velocity (i.e. 1.01 $\mathrm{m\,s^{-1}}$) and the impact position of the unpowered manifold (since the impact position of the powered trajectory is not known a priori). Propagating the trajectory with the computed ΔV_T leads to a slightly (assuming the ΔV_T is applied at low altitude) different impact position than the one predicted when using the unpowered manifold, thus the final velocity will also be slightly different than the desired one. Therefore, the ΔV_T to apply is recomputed using the new impact position, and, as a result, the new ΔC_{target}. This procedure is applied iteratively until ΔV_T converges, which usually occurs within 3 to 4 iterations for a tolerance of 10^{-4} $\mathrm{m\,s^{-1}}$.

While this method works well for the majority of the cases, it has one limitation: the computation of the ΔV_T is not defined if $\Delta C_{target} > V_T^2$ (Fig. 5.1), which occasionally occurs for shallow impacts. Assuming 1 to represent the maneuver point, 2 to represent the impact point, and C_r the component of the Jacobi integral associated with the position, then

$$C_1 = C_2 \Leftrightarrow C_{r,1} - V_1^2 = C_{r,2} - V_2^2 \Rightarrow V_2 = \sqrt{C_{r,2} - C_{r,1} + V_1^2} \tag{5.7}$$

The minimum value of V_2 (i.e. the impact velocity) is

$$\min V_2 = \begin{cases} \sqrt{C_{r,2} - C_{r,1}}, & \text{if } C_{r,2} - C_{r,1} \geq 0 \\ 0, & \text{otherwise} \end{cases} \tag{5.8}$$

Using these equations, and given the maneuver location, it is possible to compute the minimum impact velocity at each surface point. Consider an example of the distribution of minimum impact velocities along Phobos' surface for a given maneuver point (Fig. 5.2). Landings are feasible in the regions with minimum impact velocity below 1.01 $\mathrm{m\,s^{-1}}$, unfeasible otherwise. Since the maneuver is executed at very low altitudes, the impact point will tend to be directly below it (approximately), and located in a region of feasible landings. However, for very shallow impacts, the impact point may be located at larger distances to the maneuver point; in particular, it might be located in one of the regions with an impact velocity larger than 1.01 $\mathrm{m\,s^{-1}}$. In that case, the computation of ΔV_T is undefined.

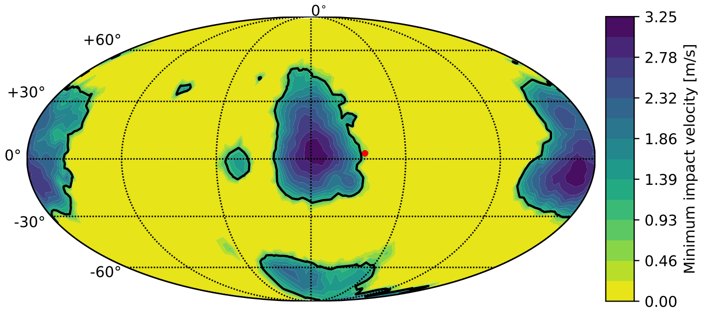

Figure 5.2: Minimum impact velocity as a function of the latitude and longitude, for a maneuver executed at an altitude of 40.0 m and latitude and longitude marked by the red dot. The black line is the level curve associated with 1.01 $\mathrm{m\,s^{-1}}$ impacts.

5.1.2. Results

As mentioned, the method fails when trying to compute the ΔV_T for impact positions that, for a given maneuver location, have minimum impact velocities above 1.01 $\mathrm{m\,s^{-1}}$, which was expected to occur for shallow impacts. Indeed, the majority of failed computations were observed to occur for manifolds with impact angles below 20.0° (Fig. 5.3). No clear pattern between the method failures and the associated impact location was observed. It is worth noting that the majority (81.5%) of the failures happened after the first iteration, meaning that the occurrence of an impact in an unfeasible region is often produced by the various ΔV_T guesses, and not by the unpowered manifolds themselves.

In most cases, the maneuver altitude is selected to have the minimum allowed value, i.e. 40.0 m, usually associated with the maximum velocity (Fig. 5.4, orange curve). Only a small number of cases (0.7% out of 1213 manifolds) was observed to have slightly higher maneuver altitudes, between 50.0 m and 110.0 m. These are cases where the final approach to the surface happens in a shallow way (impact angles under 20.0°), which in some instances allows the altitude of the spacecraft to decrease while it remains in regions with approximately constant position-components of the Jacobi integral (i.e. the trajectory is approximately parallel

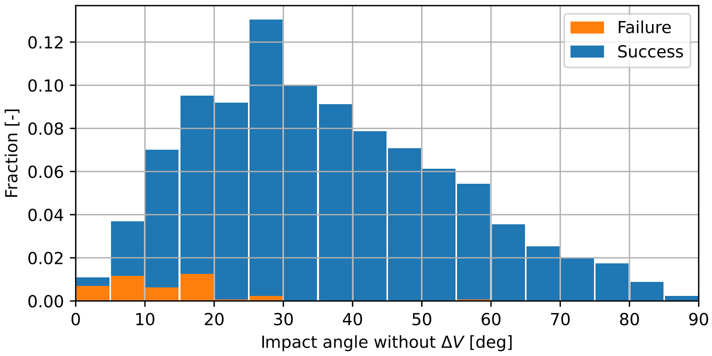

Figure 5.3: Distribution of successful and failed computations of the semi-analytical maneuver, as a function of the impact angle of the associated unpowered manifold (i.e. without application of any maneuver).

to the zero-velocity curves, c.f. Fig. 4.17), hence not showing a strong variation of the velocity with the altitude in the final portion of the trajectory. In these situations, the maximum velocity sometimes occurs above the minimum altitude (Fig. 5.4, blue curve). For the plotted low impact-angle manifold, the behavior of the velocity with the altitude changes at an altitude of approximately 1.5 km, as a result of a change in the relative direction between the trajectory and the zero-velocity curves: the trajectory is approximately parallel to the zero-velocity curves for the lower altitudes, thus displaying a slower evolution of the velocity with the altitude, and approximately perpendicular to them for the higher altitudes.

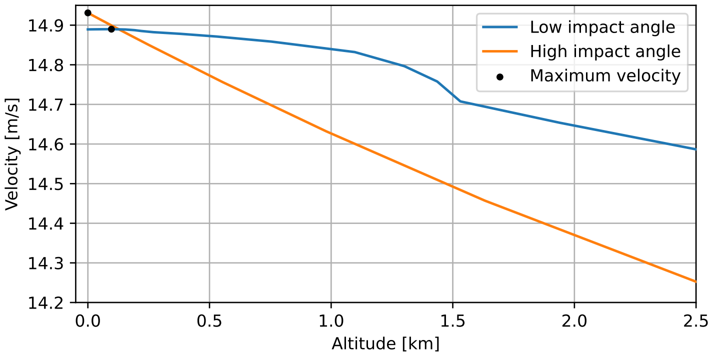

Figure 5.4: Velocity as a function of the altitude during the final part of the trajectory for two impact manifolds, one with high and one with low impact angle, departing from the same orbit of the $1:6-xz/x(R)$ family. The maximum-velocity point of each manifold is marked.

The application of the semi-analytical maneuver tends to increase the impact angle (Fig. 5.5, left), with this effect being stronger for manifolds with shallower unpowered impacts. This happens because, after the maneuver, the spacecraft is accelerated nearly vertically (since the maneuver is applied at very low altitudes, Phobos' gravity is the dominating acceleration). This vertical acceleration will have a larger influence for trajectories with low unpowered impact-angles, as in those cases the angle between the acceleration and the post-maneuver velocity is larger (recall that the pre- and post-maneuver velocities have the same direction). While the maneuver increases the observed impact angles, with the majority of cases having values above 40.0° (Fig. 5.5, right; compare with Fig. 5.3), these values are still generally lower than 81.5°, for vertical and horizontal impact velocities of, respectively, $1.0\ \mathrm{m\,s^{-1}}$ and $0.15\ \mathrm{m\,s^{-1}}$.

In general, the applied maneuvers were observed to have only small effects on the landing position with respect to the unpowered manifolds (Fig. 5.6); most maneuvers suffered impact-position changes under 300.0 m, with a mean position change of 37.07 m. The position change increases as the impact angle of the unpowered manifolds decreases: lower unpowered impact angles mean larger changes in the impact angle after applying the maneuver, thus leading to larger position changes. For the families of orbits with a high number of impacts computed with a maneuver, similar patterns in the impact locations to the ones described for the unpowered manifolds were observed. For the families with low numbers of impacts, no patterns were distinguishable; however, considering the relatively small (when compared to Phobos' size) changes in the impact position induced by the maneuvers, one can infer that the same patterns would appear if a larger

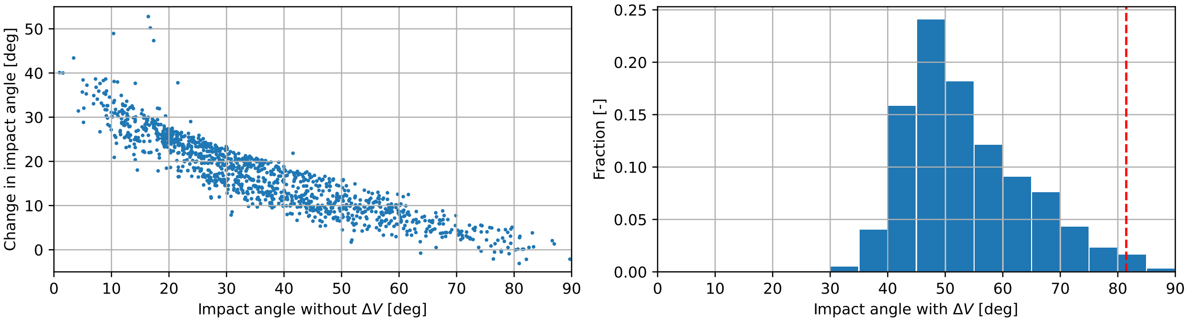

Figure 5.5: Change in the impact angle produced by the applied maneuver, as a function of the impact angle of the unpowered manifold (left). Distribution of the impact angle for the powered trajectories; the dashed red line represents the impact angle associated with the required impact velocities (right).

number of powered trajectories were computed. As in the case of the unpowered manifolds, depending on the used family/orbit, here it is also possible to land on the entire surface of Phobos (except the Stickney crater), though impacts still tend to be concentrated near the Mars and anti-Mars edges.

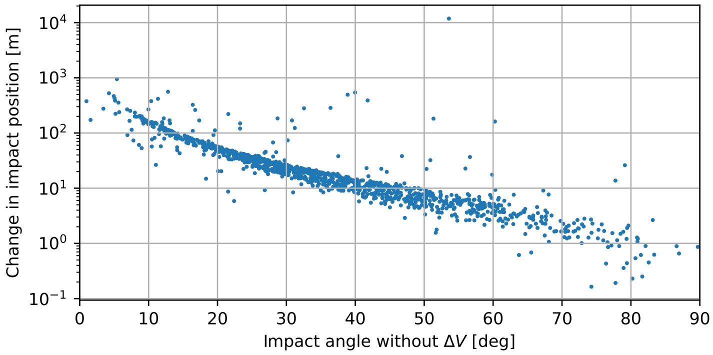

Figure 5.6: Impact-position change produced by the applied maneuver, as a function of the impact angle of the unpowered manifold.

The distribution of the minimum ΔV computed for each orbit is represented in Fig. 5.7, right. It is a similar distribution to the one presented for the impact velocity of the unpowered manifolds (Fig. 4.24, right), with the minimum ΔV increasing as the z-amplitude increases. Some irregularities are present in this distribution, which are related to differences in the impact locations between orbits: depending on the number of computed manifolds per orbit, some orbits have impacts in the regions associated with the lowest impact velocities, while others do not. For each individual orbit, a maximum difference between the minimum and maximum ΔV of $0.831\ \mathrm{m\,s^{-1}}$ was observed. Similarly to the unpowered impact velocities, the minimum values of the ΔV tend to be associated with impacts in the leading and trailing edges, while the maximum ΔV values are associated with the Mars and anti-Mars edges. The minimum time of flight obtained for each orbit is represented in Fig. 5.7, left. The time of flight tends to decrease as one follows each family of orbits, displaying the same trend discussed for the unpowered manifolds (Fig. 4.24, left). Observe how a large number of families allows landings with relatively short times of flight, many of them under five revolutions. Considering the small difference between the minimum and maximum ΔV per orbit, it is possible to execute a landing with a short time of flight while incurring only a small ΔV penalty with respect to the optimal value.

While the application of a semi-analytical maneuver allows meeting a constraint in the final-velocity norm, it does not allow selecting the final impact angle; indeed, most of the computed maneuvers violate the vertical and horizontal impact-velocity constraints. To deal with this situation, it is instead possible to optimize the braking maneuver; this topic is analyzed in the following section.

5.2. Optimized Maneuver

As discussed in the previous sections, the transfer from the initial orbit to the surface is conceptualized to have two maneuvers: an initial maneuver to insert the spacecraft into the appropriate invariant manifold, and a second maneuver to reduce the impact velocity. Determining the first maneuver via multiple shooting

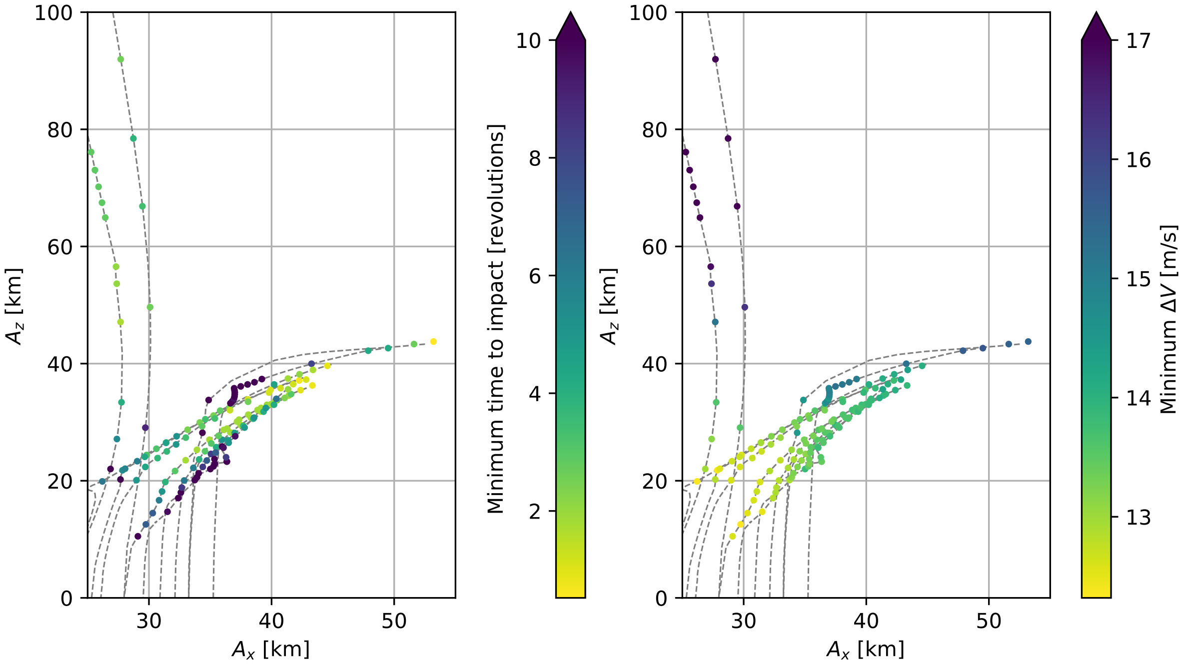

Figure 5.7: Minimum time to impact and minimum ΔV as a function of the x-amplitude and z-amplitude of the originating orbit (each circle corresponds to one orbit). The dashed lines correspond to full families, including stable orbits and unstable orbits for which no maneuvers were computed.

produced values in the order of 3×10^{-2} m s^{-1} (worst case), while the semi-analytical computation of the second maneuver produced values above 12 m s^{-1}. While optimizing the transfers, in order to meet the specified impact-angle constraint, it was opted to only optimize the second maneuver, maintaining the first fixed. Firstly, assuming a transfer that takes advantage of the invariant manifolds, one would always expect the ΔV associated with the first maneuver to be much smaller than the second one; thus, the optimization of the first maneuver would significantly increase the design space likely without allowing significant improvements to the total ΔV. Secondly, optimizing the first maneuver would not only increase the design space, but also significantly increase the propagation time: if the first maneuver were to be optimized, each trajectory would have to be propagated from the departure orbit; if only the second maneuver is optimized, each trajectory only needs to be propagated from the boundary of the possible maneuver locations. Thirdly, as the transfer from the initial orbit to the surface often takes a large number of revolutions, it is not inconceivable that optimizing the first maneuver would make the optimization unstable.

5.2.1. Grid Search

Considering the design parameters to consist of the altitude at which the second maneuver is executed and the applied ΔV, a grid search was executed to help in the selection of the parameter bounds. For a single impact manifold, the maneuver altitude was varied between 40.0 m (the same minimum altitude as considered previously) and 8.0 km, and each ΔV component between −20.0 m s^{-1} and 20.0 m s^{-1} (components specified with respect to the synodic frame). Analysis of the influence of the ΔV on the impact velocity (Fig. 5.8) confirmed that, during the optimization, the ΔV should indeed be specified with respect to a TNW frame, as in that case, the influence of each ΔV component on the impact velocity was clear. The optimal ΔV_N and ΔV_W values appeared to be located close to zero, thus, these components' bounds were selected as −5.0 and 5.0 m s^{-1}. For minimum impact velocity, the majority of the velocity change is executed via ΔV_T; based on these results and on the ΔV values computed semi-analytically, it was opted to allow ΔV_T to vary between −30.0 and 0.0 m s^{-1} (both conservative bounds). No impact velocities respecting the constraints were found, which should be expected considering that the grid search was executed with a relatively small number of grid points.

Analyzing the influence of the maneuver altitude on the minimum impact velocity per altitude (Fig. 5.9, left), one notes that below approximately 0.8 km there is a clear reduction in the minimum impact velocity for decreasing maneuver altitude; for higher altitudes, the minimum impact velocity does not appear to vary

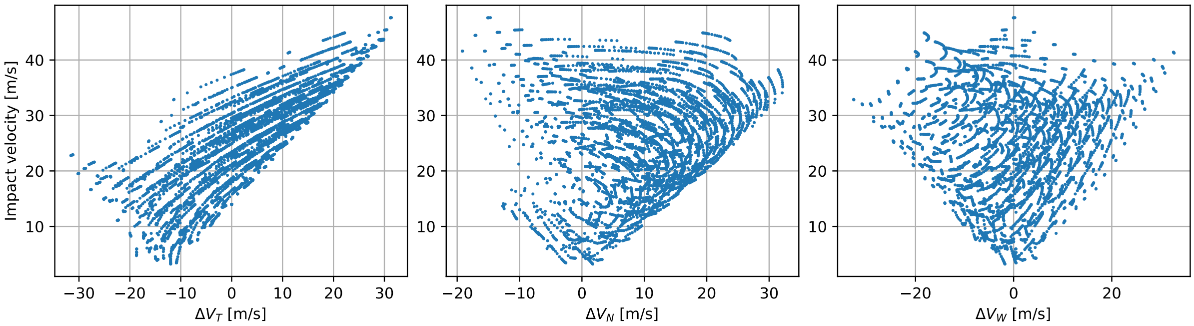

Figure 5.8: Impact velocity as a function of the ΔV components, with respect to a TNW frame. The points are organized in sets of small lines: each line corresponds to one maneuver value (specified with respect to the synodic frame) and the points in the line correspond to different maneuver altitudes.

strongly with the maneuver altitude. Firstly, this is a consequence of the maneuver velocity. For low altitudes (where Phobos' gravity dominates), in general, there is a very clear increase in the spacecraft's velocity, leading the maneuver to be more effective in reducing the energy, and consequently the impact velocity. For higher altitudes, the pattern of the variation of the velocity with the altitude is not as clear, and as a result, neither is the variation of the impact velocity. Thus, and in accordance with the executed energy-based analysis, the application of a ΔV is most effective at low altitudes.

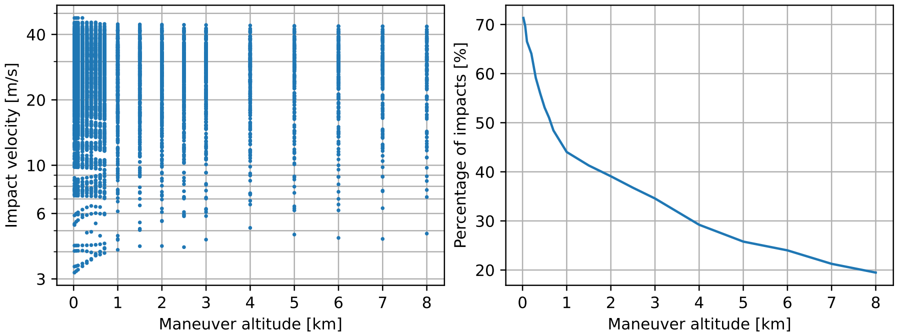

Figure 5.9: Impact velocity as a function of the maneuver altitude (left). Percentage of surface impacts as a function of the maneuver altitude (right).

It is possible to observe that the fraction of maneuvers that lead to impacts with the surface clearly decreases with the altitude (Fig. 5.9, right): higher-altitude maneuvers mean that the spacecraft has a larger amount of time to deviate from the reference manifold, in this case often leading it to escape the system, as the maneuvers were not optimized. However, even after optimization, lower-altitude maneuvers are expected to be more robust to ΔV uncertainties, as in this case the errors would have less time to affect the trajectory. Considering the higher effectiveness and (expected) higher robustness of executing the maneuvers at low altitudes, the maximum allowed value for the maneuver altitude is selected to be 500.0 m.

5.2.2. Optimization Setup

The optimization is subject to the following constraints: horizontal impact velocity $V_{f,h} \leq 0.15\,\mathrm{m\,s^{-1}}$, vertical impact velocity $V_{f,v} \leq 1.0\,\mathrm{m\,s^{-1}}$, and final state consisting of a surface impact. These were implemented as penalties in the cost function. The following cost function J (which the optimizer minimizes) is used, defined using a series of penalty functions P_i

$$J = \Delta V + P_1 + P_2 + P_3 \tag{5.9}$$

$$P_1 = \begin{cases} 0, & \text{if } V_{f,h} \leq 0.15\,\mathrm{m\,s^{-1}} \\ 100 V_{f,h}, & \text{otherwise} \end{cases} \tag{5.10}$$

$$P_2 = \begin{cases} 0, & \text{if } V_{f,v} \leq 1.0\,\mathrm{m\,s^{-1}} \\ 100 V_{f,v}, & \text{otherwise} \end{cases} \tag{5.11}$$

$$P_3 = \begin{cases} 0, & \text{if the spacecraft impacted the surface} \\ 10^5, & \text{otherwise} \end{cases} \tag{5.12}$$

The penalty functions were selected to ensure that a trajectory that meets the impact velocity constraints always has a lower cost than one that does not and that a trajectory that impacts the surface always has a lower cost than one that does not.

The optimization minimizes the cost function by selecting four variables: the altitude at which the maneuver is executed, and the applied ΔV (specified with respect to the TNW reference frame). The selected bounds for each parameter are repeated here:

$$\begin{aligned} \text{altitude} &\in [40, 500]\ \mathrm{m} \\ \Delta V_T &\in [-30, 0]\ \mathrm{m\,s^{-1}} \\ \Delta V_N &\in [-5, 5]\ \mathrm{m\,s^{-1}} \\ \Delta V_W &\in [-5, 5]\ \mathrm{m\,s^{-1}} \end{aligned} \tag{5.13}$$

Specifying the maneuver location using the time since departure instead of the altitude was initially considered. However, the former resulted in a very poor optimizer performance. This is likely related to the fact that while specifying the maneuver altitude ensures that the maneuver is executed (as long as the spacecraft gets close enough to the surface), that is not the case when specifying the maneuver time. As such, when specifying the maneuver time, the cost function minimum would likely be located between two regions with significantly larger costs (maneuver executed too early, resulting in too high impact velocity, and maneuver time after the impact time, resulting in even higher impact velocity), with the optimizer struggling to converge to the feasible region.

Each optimization run uses one of the manifolds for which the initial maneuver has been computed (via multiple shooting) as a reference. This manifold is propagated from the departure QSO to the maximum maneuver altitude (500 m), thus determining the initial state used during the optimization. Each time the cost function is evaluated, the state is propagated from the altitude of 500 m until either the spacecraft impacts the surface, the propagation time reaches 2 Phobos' periods or the distance to the center of Phobos grows above 200.0 km.

5.2.3. Optimizer

The optimization was executed using Pygmo, the Python interface of Pagmo, a C++ optimization library developed by ESA (Biscani and Izzo, 2020). The optimization was executed using a Differential Evolution (DE) algorithm, a form of global optimization that has been shown to be effective in dealing with high-thrust trajectories (Musegaas, 2012; Olds et al., 2007). The behavior of the DE algorithm depends on four parameters: the population size, the crossover probability (CR), the mutation scaling factor (F), and the mutation variant. The population used by the algorithm consists of a series of design-variables vectors $\boldsymbol{x}_{i,G}$ (the i^{th} vector of the G^{th} generation), each vector being constituted by a set of values assigned to the parameters being optimized. The initial population is obtained by assigning to the design-variables vectors random values obtained according to a uniform distribution between the lower and upper bound of each parameter. In each generation, to find a better solution, each $\boldsymbol{x}_{i,G}$ is modified via three operations: (1) mutation, in which a certain number of vectors $\boldsymbol{x}$ is selected and combined to form a mutated vector $\boldsymbol{v}_{i,G}$ (this operation depends on F); (2) crossover, in which $\boldsymbol{x}_{i,G}$ and $\boldsymbol{v}_{i,G}$ are combined to form a new vector $\boldsymbol{u}_{i,G}$ (this operation depends on CR); (3) selection, where the best vector between $\boldsymbol{x}_{i,G}$ and $\boldsymbol{u}_{i,G}$ is selected to be kept in the following generation (Olds et al., 2007).

Pagmo includes three versions of DE: the standard DE (described above), Self-Adaptive Differential Evolution (SADE), and DE1220. DE was rejected a priori, as according to Pagmo's documentation the other two versions tend to perform significantly better. After testing SADE (without tuning its parameters) and DE1220 (which does not depend on any parameters) with two different impact manifolds, the former was selected as it showed better performance, even without any parameter tuning.

SADE automatically adapts CR and F, thus to tune the optimizer one only needs to select the population, the method used to mutate the population, and the method used to adjust CR and F. After tuning (Section 8.3), the SADE-jDE version was selected, with population size 7, mutation variant *best/2/bin*, and a maximum of 1000 function evaluations per run (corresponding to 143 generations). The SADE-jDE version, introduced by Brest et al. (2006), adapts CR and F by extending each vector $\boldsymbol{x}_{i,G}$ with the CR and F values. The mutation variant *best/2/bin* means that (Storn and Price, 1997): firstly, the mutated vector $\boldsymbol{v}_{i,G}$ is generated by modifying the best vector $\boldsymbol{x}_{i,G}$ (*best*); secondly, $\boldsymbol{v}_{i,G}$ is generated by modifying the best $\boldsymbol{x}_{i,G}$ with the differences between two pairs of $\boldsymbol{x}_{i,G}$ vectors (2); thirdly, each value in $\boldsymbol{u}_{i,G}$ is obtained by executing independent binomial experiments with the values of $\boldsymbol{x}_{i,G}$ and $\boldsymbol{v}_{i,G}$ (*bin*).

5.2.4. Results

Tuning Data

While tuning the optimizer, two manifolds, one with high impact angle (62.995°) and one with low impact angle (8.011°), were optimized using a series of different optimizer settings, producing in the process data that allows gaining further insight into the problem. Only optimization runs that converged to a valid solution (i.e. respecting all the constraints) are considered here.

Analyzing the influence of the maneuver altitude on the required ΔV (Fig. 5.10), one notes that the minimum ΔV value is obtained for the minimum allowed maneuver altitude. However, it is possible to observe that the maneuver altitude has a much stronger effect on the required ΔV for the high impact-angle manifold than for the low impact-angle one (notice the red lines represented in the two plots and their slopes). This happens because, in the final phase of the descent, the altitude has a relatively weak effect on the velocity for the low impact-angle manifold, which is not the case for a high impact-angle. For example, between the altitudes of 100 and 40 m, the low impact-angle manifold increases the velocity by 0.02 $\mathrm{m\,s^{-1}}$, while the high-impact angle manifold by 0.3 $\mathrm{m\,s^{-1}}$. Due to the low angle of impact, the manifold travels a larger distance close to the surface, which allows it to remain in regions with approximately constant C_r levels (resulting in an almost constant velocity), while the altitude changes as a result of the changing local topography. In the case of the high impact-angle manifold, the local topography is approximately constant (since only a small distance is traveled), with the velocity changing as a result of a change in altitude which coincides with a changing C_r. As discussed in Section 5.1, the higher the velocity at which the maneuver is applied, the more efficient the reduction of the spacecraft's velocity (ΔV-wise). Thus, if the velocity has a weak variation with the altitude, so does the ΔV required to meet the target final velocity.

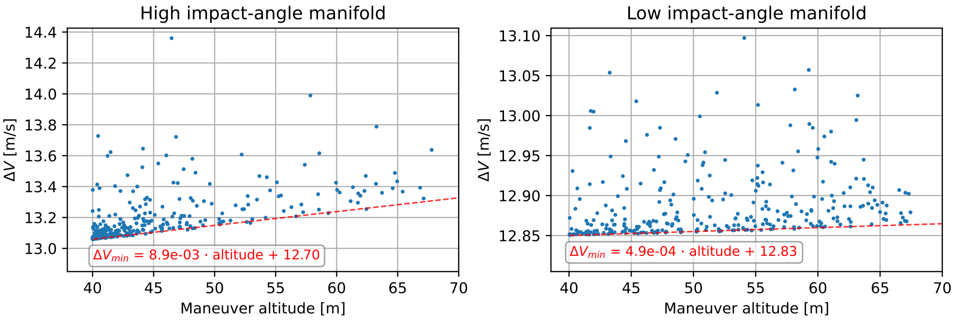

Figure 5.10: ΔV as a function of the maneuver altitude, for a series of converged optimization runs (executed for tuning the optimizer), for high (left) and low (right) impact-angle manifolds. The minimum ΔV as a function of the altitude is represented by a dashed red line, with its equation annotated next to it.

The weak dependence of the required ΔV on the maneuver altitude indicates that low impact-angle manifolds might be more robust to errors and uncertainties. Firstly, the weak dependence on the altitude immediately indicates that the trajectory is more robust to position navigation errors. Secondly, as a result of the robustness to changing altitude and the larger amount of time traveled at low altitudes (the low and high impact-angle manifolds take, respectively, 48 s and 9 s to travel between the altitudes of 100 and 40 m), a low impact angle would allow more time to apply the maneuver without noticeable effects on the ΔV. Thus, in the case of a maneuver failure, one would have a longer time to execute another maneuver, e.g. with backup thrusters.

Analyzing the effect of the components of the maneuver on its norm (Fig. 5.11), it is first possible to note that ΔV varies approximately linearly with ΔV_T. This happens simply because ΔV_T is the only component that contributes to reduce the spacecraft's energy, thus having a significantly higher value than the other two.

Analyzing ΔV_N, one sees that it is in general positive, and takes larger values for the low impact-angle manifold than for the high impact-angle one; since ΔV_N points roughly in the down direction, it is the component that plays the largest influence on the impact angle. As such, it takes a higher value for the low impact-angle manifold, where the final impact angle requires a larger change in order to meet the specified constraints. It is also possible to observe that for the high impact-angle manifold the ΔV_N values appear to have converged to a single optimal point, while for the low impact-angle manifold the ΔV_N values are much

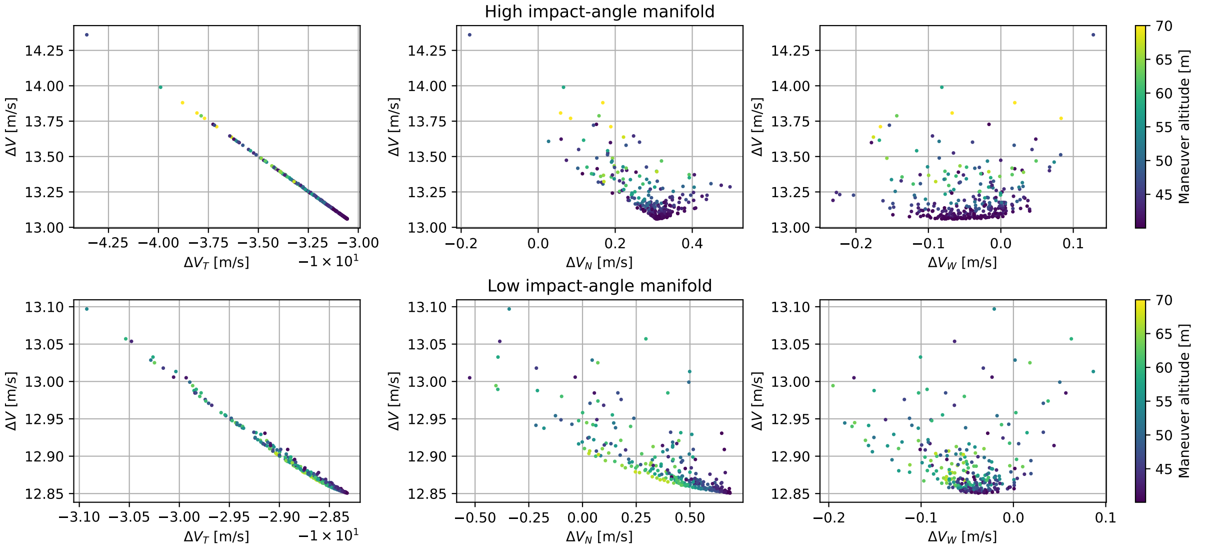

Figure 5.11: ΔV norm as a function of the altitude and the ΔV components (specified with respect to a TNW frame), for a series of converged optimization runs (executed for tuning the optimizer), for high (top) and low (bottom) impact-angle manifolds.

more spread. This is also a result of the weak influence of the altitude on the ΔV, with different “optimal” ΔV_N values for each altitude value. The necessary ΔV_N increases as the maneuver altitude decreases because the lower altitude results in a shorter time before impact, with Phobos’ acceleration (which is approximately normal to the surface in this final phase) having less time to naturally increase the impact angle.

Finally, ΔV_W has lower values than the other two components, is generally negative, and quite spread, for both manifolds. The ΔV_W component, which either makes the spacecraft go to the right or to the left, mainly influences the impact velocity by slightly modifying the impact location. The obtained ΔV_W value is evidently limited by the fact that while it can lead the spacecraft to regions with lower C_r (resulting in lower impact velocity), increasing ΔV_W will require an increase in ΔV_T (as the former always increases the energy). For the two tested manifolds, a negative ΔV_W tends to lead the spacecraft to regions with lower associated impact velocities.

Tuned-Optimizer Data

Using the tuned optimizer, the braking maneuver was optimized for the reduced set of approximately 1200 manifolds. The optimizer was observed to converge to a valid solution (i.e. a solution that respects all constraints) in 98.04% of the cases. The low percentage of failures is considered acceptable, as the present research only constitutes a feasibility study, not being particularly important to analyze all individual manifolds. Furthermore, by running the optimization with a different seed or a larger number of function evaluations, one would likely be able to obtain valid solutions for all manifolds. The convergence failures do not appear to be related to specific values of the unpowered impact angle, with the distribution of the impact angle for the failed cases following the distribution for the total number of cases. Finally, it is worth noting that the high percentage of converged optimizations indicates that generally it is possible to obtain a valid landing solution, even with strict impact-velocity constraints.

In most cases, the optimizer converged to maneuver altitudes close to the minimum, with 52.41% of the cases having maneuver altitudes between 40 and 45 m. For the remaining cases, maneuver altitudes up to 150 m were found. No clear pattern between the maneuver altitude and unpowered impact angle was observed; it was only noticeable that higher-altitude maneuvers do not occur for very-high unpowered impact angles (above 75°). Note that the convergence to maneuver altitudes well above 40 m might happen due to an imperfect optimizer convergence (especially in the cases where the altitude does not influence strongly the velocity), but also because that is indeed the true optimum.

Analyzing the optimized ΔV components (Fig. 5.12), it is first possible to note that, as expected, ΔV_T always has a much higher value than the other two components. Secondly, ΔV_W always takes significantly smaller values, distributed around 0, which would be expected based on the hypothesis that it moves a trajectory to the left/right in order to decrease the Jacobi integral of the impact point; thus, its sign depends on the specific location of the impact, with no global trend expected. Finally, ΔV_N also takes much smaller values than ΔV_T, generally decreasing as the unpowered impact angle increases (since the required impact angle correction would be lower) and decreasing as the maneuver altitude increases (since Phobos' gravity acts for a longer period time, thus exerting a larger influence in the increase of the impact angle). However, it is also possible to observe some points with negative ΔV_N, often with a significantly higher norm than the positive values; in particular, there is a series of points with $\Delta V_N < -0.5\ \mathrm{m\,s^{-1}}$ for unpowered impact angles below 20°. This is probably related to cases where increasing the impact range is particularly beneficial. Assuming a constant gravitational acceleration after the maneuver, using a negative ΔV_N will decrease the post-maneuver vertical velocity (either making it less negative or more positive), in turn leading to a longer time of flight and to a longer distance covered between the maneuver and the impact. This might allow an impact in a region with a lower surface Jacobi value, and consequently, a lower impact velocity, which would allow decreasing the applied ΔV_T. However, these adjustments to ΔV_T and ΔV_N will also influence the impact angle, which would require further adjustments to meet the impact angle constraint, making it difficult to predict the final effect on the total ΔV. Nevertheless, it can be hypothesized that for some situations with low unpowered impact angle, where decreasing ΔV_N would allow a more noticeable increase in the distance between the maneuver and impact, the obtained decrease in the impact velocity is sufficient to compensate for the ΔV changes to correct the impact angle.

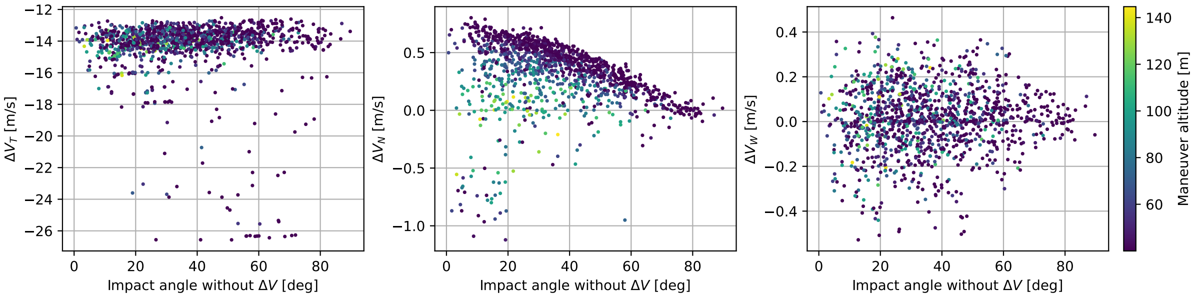

Figure 5.12: $\boldsymbol{\Delta V}$ components (specified with respect to a TNW frame) and maneuver altitude as a function of the unpowered impact-angle, for the full set of optimized manifolds.

The effect of the maneuvers on the landing position was observed to have a similar pattern to the one presented for the semi-analytical maneuver, with larger position changes for smaller unpowered impact angles. However, larger position changes were observed here, with a mean value of 170.25 m. Evidently, the optimized maneuver tends to produce larger changes in the impact angle (in order to meet the constraints), thus leading to larger position changes. A small number of cases with position changes above 1 km were observed. These either corresponded to maneuvers applied at relatively high altitudes, or maneuvers applied significantly before the impact of the unpowered manifold, for cases where the unpowered manifold got very close to the surface during the revolution preceding the impact.

Impacts were observed on the entire surface (except the Stickney crater), though with a higher concentration on the Mars and anti-Mars edges. Similarly to the results presented when analyzing the unpowered manifolds, $d:n$ families with $d \geq 2$ and the $1:6-xz/x(BE)$ family appeared particularly favorable in terms of surface coverage, with some individual orbits being able to produce manifolds that can reach the entire surface. Thus, when designing a mission for which flexibility in the landing location is an important factor, attention should be paid to the members of these families in particular.

Comparison with Semi-Analytical Maneuver

Although the maneuvers computed semi-analytically do not allow meeting the target impact angles, they do allow obtaining a reasonable initial guess for the ΔV required to do so (Fig. 5.13). In the vast majority of cases, the semi-analytical ΔV shows errors below $0.6\ \mathrm{m\,s^{-1}}$, significantly lower than the minimum observed optimal ΔV of $12.5\ \mathrm{m\,s^{-1}}$. No clear pattern was seen between the ΔV error and the norm of the optimal ΔV; however, the ΔV error increases for decreasing impact angle of the unpowered manifold. In general, a semi-analytical maneuver predicts a lower ΔV than an optimized one, as the former only takes into account the impact velocity, while the latter also determines the ΔV required to correct the impact angle. In a very small number of cases, the optimized ΔV is lower than the semi-analytical one. This happens because the optimization has some degree of control over the impact location, allowing the selection of regions with lower C_r, associated with lower impact velocities, and consequently lower ΔVs. These situations only occur for relatively high unpowered impact-angles, above 75°.

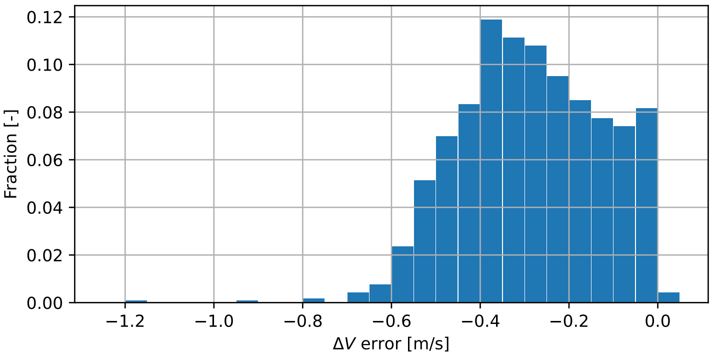

Figure 5.13: Distribution of the error of the semi-analytical ΔV with respect to the optimized ΔV.

Since the maneuver is applied at a very low altitude, the minimum obtained times of flight are also very similar to the ones predicted using the semi-analytical maneuver (the two types of trajectories only differ in the portion after the maneuver). Overall, the distributions of the minimum time of flight and ΔV along each family of optimized trajectories are nearly identical to the ones presented for the semi-analytical case (Fig. 5.7); hence, the plots are not repeated here.

6

Uncertainty Analysis

In the previous chapters, invariant manifolds departing from QSOs were computed, and the maneuvers necessary for the manifolds to impact the surface with low-enough velocities were designed. However, it is important to analyze whether the designed transfers would be robust to the various uncertainties present in the system. This chapter discusses that analysis.

Three types of uncertainties are considered: uncertainty in the application of the ΔVs, navigation uncertainties, and gravity uncertainties. Unless otherwise stated, the following uncertainties are all considered to be normally distributed and are specified using their 3σ values. The application of each ΔV is assumed to have uncertainties of 5.0 % in the norm and 0.1° in the direction (half-cone angle); the second angle required to define the uncertain ΔV is considered to have a uniform distribution in $[0, 2\pi[$ (i.e. the uncertain ΔVs are uniformly distributed around the cone with the specified half-cone angle). These values are used by Chen et al. (2020) for the analysis of three-dimensional QSOs around Phobos. Assuming autonomous navigation (as proposed by ESA's Phobos Sample Return mission proposal; Pickering, 2014), navigation uncertainties of 36.73 m in each position component and $0.693\ \mathrm{cm\,s^{-1}}$ in each velocity component are considered; these values are estimated by Chen et al. (2022) for a mission profile similar to MMX, using autonomous navigation. Uncertainties in Phobos' gravitational parameter of 0.024 % are considered, corresponding to the uncertainties that Chen et al. (2022) predict would be obtained after a geodesy campaign around Phobos. Just for reference, the uncertainty in Mars' gravitational parameter has a value of 1.13×10^{-7} % (Genova et al., 2016).

Each transfer trajectory is constituted by two phases: firstly, the application of the ΔV to enter the invariant manifold and the subsequent unpowered trajectory to the second maneuver point; secondly, the application of the second ΔV and unpowered trajectory to the surface. Since the error in the application of the two maneuvers is considered independent, and assuming that the second maneuver would be recomputed taking into account the errors produced by the first one, the uncertainty analysis is also divided into the two mentioned phases, though with one small difference: the first phase is considered to consist of the transfer from the orbit to the surface (instead of from the orbit to the maneuver location). A few reasons justify this: firstly, this allows a more general analysis, also applicable to unpowered transfers; secondly, since the maneuvers are applied at very low altitude, the difference between trajectories propagated until the impact or until the maneuver altitude is very small, thus an analysis of the former also allows insight into the latter; thirdly, propagating the trajectories until a surface impact instead of until a specified altitude is faster, as the former does not require evaluating the altitude throughout the trajectory.

Monte Carlo analyses are executed for the two cases, with different representative manifolds being tested. For each transfer phase and each manifold, Monte Carlo runs are executed considering all the uncertainties, just ΔV uncertainties, just navigation uncertainties, and just Phobos' gravity uncertainties. In each Monte Carlo analysis, 10^4 simulations were executed, with errors sampled from the described distributions. Since the propagations were executed using dimensionless units, in each one the initial state was renormalized according to the uncertain value of Phobos' gravitational parameter.

6.1. Orbit to Surface

6.1.1. Stability Index Definition

The initial analysis of the orbit-to-maneuver phase using Monte Carlo data proved difficult, as the results appeared to behave chaotically. Thus, the analysis of the manifolds using some stability index was considered appropriate, as it would allow a faster analysis of a larger number of manifolds, helping to make the patterns in the data clearer. The preliminary Monte-Carlo analysis showed that the navigation uncertainty dominates this phase of the transfer, with the ΔV and Phobos' gravity uncertainties having negligible effects (this behavior is further discussed in the following section). As such, the usage of a stability index based on the STM would appear particularly suitable, since it gives information about the effect of deviations in the initial state.

The usage of a stability index similar to Eq. 3.11, based on the eigenvalues of the monodromy matrix, is not possible: the assessment of the stability using the eigenvalues of the monodromy matrix is based on Floquet's theorem, which is only valid under the assumption of a periodic solution (Verhulst, 2000). As such, using the eigenvalues of the STM to assess the stability of the transfer trajectories, which are not periodic, would likely not give trustworthy results.

Instead, the singular-value decomposition of the STM can be used, which gives information about the direction and magnitude of the maximum expansion as a result of a perturbation. The singular-value decomposition of the STM is given by (Muralidharan, 2021)

$$\Phi(t, t_1) = U\Sigma V^T \Rightarrow \Phi(t, t_1)V = U\Sigma \tag{6.1}$$

where U and V are 6×6 orthogonal matrices (thus $V^{-1} = V^T$), and Σ is a 6×6 diagonal matrix, with the elements on the diagonal corresponding to the singular values of $\Phi(t, t_1)$. The singular values of the STM correspond to the square root of the eigenvalues of $\Phi(t, t_1)^T \Phi(t, t_1)$. Since $\Phi(t, t_1)^T \Phi(t, t_1)$ is a symmetric matrix, its eigenvalues, and consequently the singular values, are positive real numbers. Based on the previous equation, it is possible to observe that a unitary perturbation at time t_1 in the direction of one of the columns of V, results in a deviation at time t in the direction of the corresponding column of U with the magnitude of the associated singular value. The maximum singular value gives the magnitude of the expansion or contraction as a result of a perturbation in the direction that ensures the maximum expansion of the initial perturbation (Muralidharan, 2021). As such, the stability of the trajectory can be assessed by evaluating the maximum singular value.

Given the physical significance of the maximum singular value, the following stability index was selected

$$\nu = \ln \max \sqrt{\lambda_{\Phi^T \Phi}} \tag{6.2}$$

This index is similar to the one used by Ceriotti and Sanchez (2016), who analyze the stability of transfers designed using invariant manifolds of Lagrange point orbits. Ceriotti and Sanchez use as stability index the maximum singular value of the STM divided by the transfer time, to allow transfers with different times of flight to be compared. However, some numerical tests using this time-normalized stability index showed that it benefited shorter transfers, likely because the error growth should be approximately exponential for unstable trajectories (Alizadeh and Villac, 2010). A time normalization using the n^{th} root of the maximum singular value, with n the number of revolutions (not necessarily an integer), was also considered. This would be similar to the time normalization used by Chen et al. (2020), in the stability index of Eq. 3.11. However, this normalization was also considered inappropriate, as it was observed to benefit longer transfers. Overall, it was considered that normalizing the stability index using the transfer time would not allow transfers with different times of flight to be compared accurately. Hence, it was opted to use the stability index in two ways: firstly, it is compared at the end of the transfers, thus measuring the stability of the whole transfer, independently of its duration; secondly, it is compared at fixed epochs throughout the transfers, allowing the error evolution to be compared between trajectories with different times of flight.

It is worth noting that the introduced stability index is very similar to the finite-time Lyapunov exponent (FTLE), with the difference that the latter is divided by the propagation time. Lyapunov exponents give a measure of the average expansion/contraction rate of an initial perturbation as time goes to infinity (thus, each trajectory is associated with a specific value of the Lyapunov exponents), with the FTLE giving an approximation of that rate for long-enough transfer times (Alizadeh and Villac, 2010). Lyapunov exponents are also defined using the natural logarithm, to account for the exponential error growth; this is the same reason why the natural logarithm was used in Eq. 6.2, though in this case it simply leads the stability indices of the different trajectories to have similar magnitudes. A similar discussion to the one presented above can be used to explain why the division by the propagation time used in FTLE was not applied here: the FTLE only converges to the Lyapunov exponent for long-enough propagation times, however, since transfers with very different times of flight are analyzed here, using the exact definition of the FTLE would lead to the comparison of trajectories in different stages of their convergence to the Lyapunov exponent, thus not giving a trustworthy account of the difference in their stability.

Finally, note that the used stability index is limited by the linearization it is based on; nonetheless, it should still allow a valid comparison of the relative instabilities of the different invariant manifolds (Ceriotti and Sanchez, 2016).

6.1.2. Stability Index Analysis

The uncertainty analysis using the stability index was executed for the full set of impact manifolds determined in Section 4.3, instead of the reduced set used during the maneuver design.

The first important effect on the manifolds' stability is the propagation time. According to the principle of exponential divergence, chaotic (i.e. unstable) nearby trajectories diverge exponentially with time, while regular (i.e. stable) ones diverge linearly with time (Villac, 2008). Thus, the stability index should grow approximately linearly and logarithmically, respectively for unstable and stable manifolds. Examples of the evolution of the stability index as a function of time for different manifolds are represented in Fig. 6.1. Most manifolds display a similar evolution to the blue, orange, and purple curves: the stability index increases approximately linearly with time (thus corresponding to unstable trajectories), and with similar rates for the different manifolds originating from a given QSO; later, the error growth suffers a faster increase when executing the flyby-like approach that precedes the impact. The impact manifolds with a long time of flight generally behave similarly to the red curve (note that this behavior is less common than the previous one): initially (in the plotted example, below a time of 14 revolutions), the error displays a slow evolution with time, approximately logarithmic (corresponding to a stable trajectory); eventually, the stability index starts evolving linearly with time (corresponding to an unstable trajectory), leading to an impact with the surface. Although these are the generally-observed patterns, a large number of manifolds display a different behavior, for example, not showing any increase in the error growth rate after executing the pre-impact flyby, or even displaying an error decrease after executing the flyby. In all cases, right before impact, there is a very sharp increase in the stability index, likely related to the influence of Phobos' irregular gravity field (though the size of that increase also varies significantly between trajectories). To identify the effects governing some of the mentioned differences in the evolution of the stability index, different manifolds originating from the same orbit and with similar times of flight (within 0.2 revolutions) were compared. No clear pattern between the error growth and the departure location along the originating periodic orbit, flyby altitude, impact location, or impact angle was found.

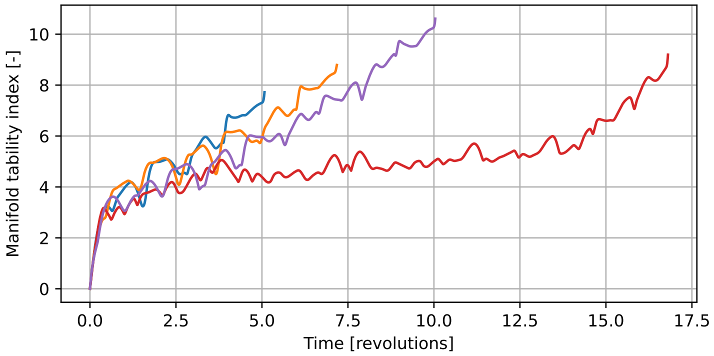

Figure 6.1: Manifold stability index as a function of time, for four different manifolds originating from the same orbit of the $1:6-xz/x(BE)$ family. Each curve corresponds to a single manifold. The point with the largest time for each curve corresponds to the impact.

To study the effect of the instability of the originating orbit, the stability index of the different manifolds of a single family was plotted as a function of the orbit's stability index; the stability indices of the manifolds were plotted at fixed epochs, i.e. before the manifolds impacted the surface (Fig. 6.2). This analysis was executed for all the computed families of orbits. Since different stability indices were used for the orbits and the manifolds, no direct comparison between the magnitudes of the two should be done. As can be noted, the instability of the manifolds tends to increase as the instability of the orbits increases. This likely happens because, at least initially, the manifolds tend to be similar to the orbit from which they depart; as such, the same dynamics that act on the originating orbit, making it more or less unstable, also act on the manifold, thus leading the stability of the two to generally evolve in the same way. Interestingly, for lower instabilities of the orbits, the minimum instability of the manifolds appears approximately constant; for larger instabilities of the originating orbits, this phenomenon disappears, with the minimum instability of the manifolds increasing with the instability of the orbits. The approximately constant minimum instability of the manifolds across different orbits appears to be related to the occurrence of initially stable manifolds (similar to the red curve in Fig. 6.1), for which the error displays a logarithmic evolution. For larger orbit's instabilities, all manifolds are unstable from their start, thus leading to the observed linear growth.

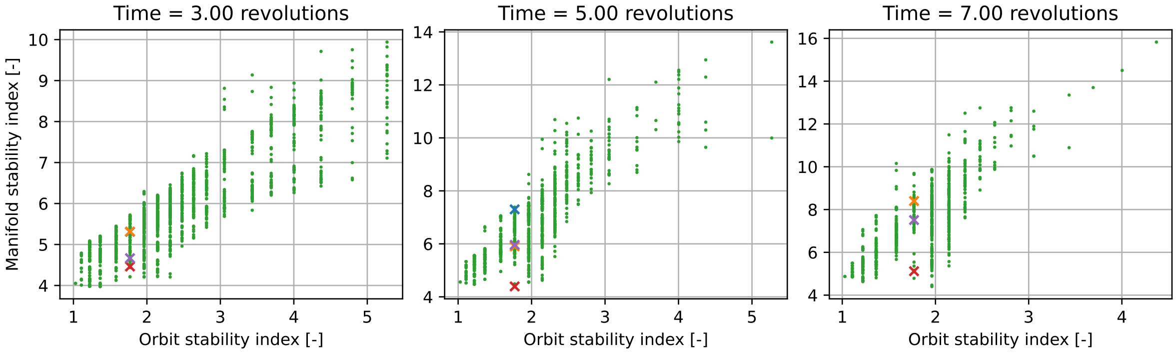

Figure 6.2: Manifold stability index as a function of the originating orbit's stability index, evaluated at different epochs, for the manifolds originating from the $1:6-xz/x(BE)$ family. Each point corresponds to a single manifold. The manifolds represented in Fig. 6.1 are highlighted with a cross, having the same color in both plots. In the left plot the blue and orange overlap, in the center one the purple and orange curves overlap, in the right one the blue cross is not represented (after its impact).

Comparing the effect of the orbits' instability and the propagation time on the manifolds' instability (Fig. 6.3), one observes that for short times the instability of the manifold is dominated by the time, with a more reduced effect from the orbit's instability; meanwhile, for longer propagation times the opposite is true, with the variation of the manifold's stability being influenced mainly by the orbit's stability. The weaker effect of the time for larger times should be related to the fact that at that point only stable manifolds with the slower logarithmic error-growth remain in the system, the unstable manifolds having either escaped or impacted the surface earlier. Considering that the times of flight decrease with increasing instability of the initial orbit (Section 4.3), and assuming shorter transfer times are preferred, using a more unstable initial orbit might be beneficial, as it would offer more opportunities for landings with short times of flight, without significantly influencing the stability of the used manifolds.

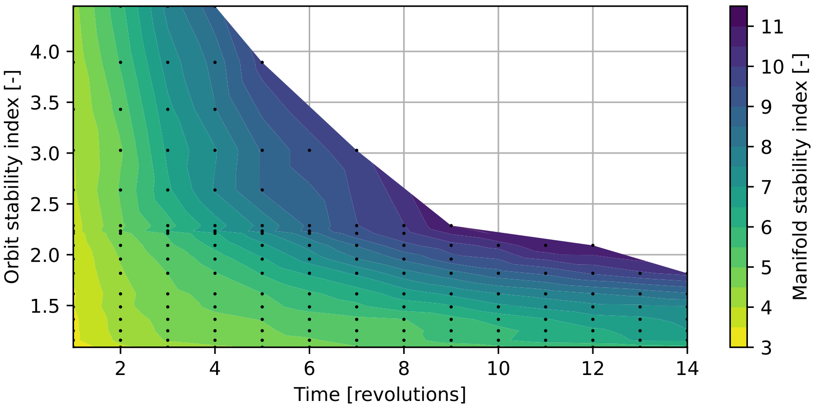

Figure 6.3: Mean manifold stability index as a function of the time and originating orbit's stability index, for the manifolds produced by the $2:13-xz$ family. The used data points are represented by black dots; each point corresponds to the mean stability index for all manifolds originating from a certain orbit evaluated at a certain time.

Finally, no clear pattern was observed when analyzing the relationship between the impact location (latitude and longitude) and the stability of the manifolds; this analysis was executed for manifolds with similar transfer times and originating from orbits with similar stability. This lack of influence of the impact location likely happens because the manifolds tend to be very similar independently of the impact location, only becoming different in the final portion of the trajectory (approximately during the last half revolution).

6.1.3. Monte Carlo Analysis

In an initial analysis, 10^4 Monte Carlo runs were executed for several different manifolds with different properties (e.g. time of flight, initial orbit, impact location) but, as previously mentioned, this did not allow much insight into the factors governing the stability of the different manifolds (which motivated the analysis with the stability index). However, the Monte Carlo analysis did show that the navigation uncertainties dominate the effect of the total uncertainties, with the maneuver and Phobos' gravitational uncertainties having a negligible contribution. This was observed for all tested manifolds, with an example of the impact-location dispersion being given Fig. 6.4. The same was observed when analyzing the percentage of impacts: for the represented manifold, when considering just the maneuver or gravity uncertainties, all trajectories impacted the surface, while when analyzing the full uncertainties and the navigation uncertainties, respectively, 62.19% and 62.23% did.

Analyzing the presented example for the impact-location dispersion with all uncertainties, it is possible to note that most impact locations are distributed according to four groups: a first group centered around a latitude of −30° and longitude of 0° (near the nominal impact location); a second one centered around a latitude of −60° and longitude of −120°; a third one distributed over the anti-Mars and trailing edges, with latitudes between −50° and 20°; a final group distributed along the trailing, Mars, and leading edges, with latitudes between 0° and 50°. Each of these groups is associated with trajectories with a relatively narrow range of times of flight. The first and second groups correspond to trajectories with times of flight similar to the nominal case: the first group is constituted by trajectories similar to the nominal one, impacting the surface near the Mars edge; the second group is constituted by similar trajectories, but which pass near the Mars edge without impacting it, having their impact slightly later along their trajectory. The third and fourth groups correspond, respectively, to trajectories that impact the surface approximately one and two revolutions after the nominal impact time. Other impact locations that appear randomly distributed are associated with larger times of flight.

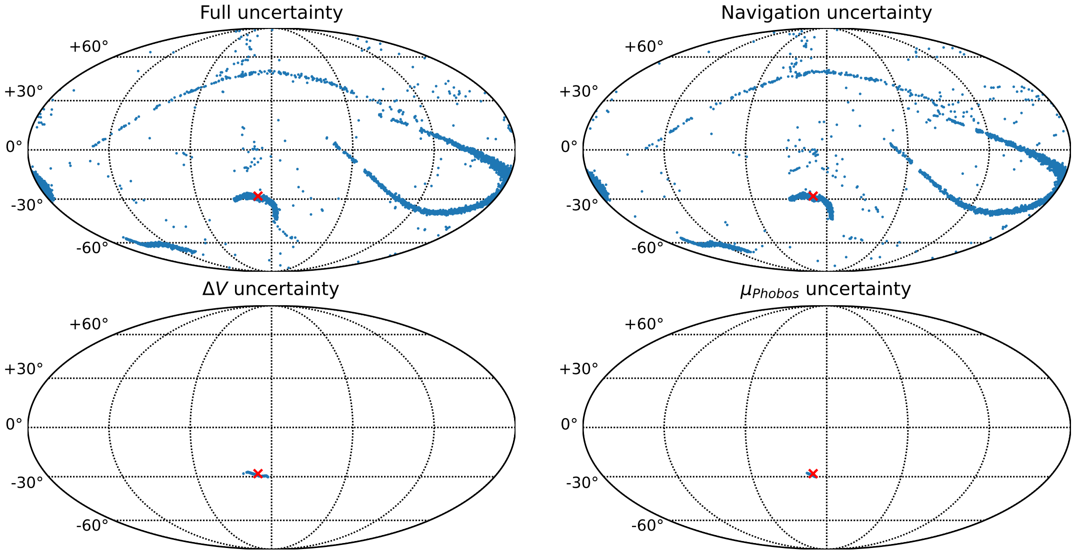

Figure 6.4: Distribution of the impact locations, as seen from the Mars perspective, for the uncertainty analysis of the orbit-to-surface phase. Each plot represents the full surface of Phobos. The data is represented for a single manifold originating from the $1:6-xz/x(BE)$ family, using different uncertainties: all uncertainties (top left), only navigation uncertainties (top right), only ΔV uncertainties (bottom left), and only Phobos' gravitational parameter uncertainties (bottom right). The nominal impact location is marked with a red cross.

The lack of influence of the maneuver uncertainties is likely related to the fact that they were specified as a fraction of the norm of the nominal ΔV required to enter the manifold, which has very low values (around $2\times10^{-1}\ \mathrm{m\,s^{-1}}$, in the worst cases). The weak effect of the uncertainty in Phobos' gravity results from the fact that this uncertainty has a small value (percentage-wise), and the fact that Phobos' gravity is much smaller than Mars', only becoming important when the trajectories are already close to the surface. Thus, the gravity uncertainties only become relevant close to the impact epoch, not having enough time to influence the trajectory significantly.

To assess the validity of the conclusions presented when studying the stability index, these were compared with the results of various Monte-Carlo runs. There are, however, some differences with respect to the previously described Monte-Carlo runs. Firstly, these were executed for all manifolds for which the maneuvers have been designed (approximately 1200 manifolds). Secondly, to obtain an acceptable computational time, only 10^3 runs per manifold were executed, a number which was considered to still provide representative results. Thirdly, these runs were executed considering just the navigation uncertainties, since these are dominant and also the uncertainties that can be directly correlated to the stability index (which depends on the STM).

The percentage of impacts as a function of the manifold's stability index is represented in Fig. 6.5. There is a clear correlation between the two values, with the percentage tending to decrease as the stability index increases. However, the relationship is not completely linear, especially for higher stability indices, likely because the stability index is based on linearized dynamics. The performance of the stability index is expected to degrade for larger times of flight and larger instabilities, which indeed appears to be the case here. Nevertheless, the stability index does approximate reasonably well the percentage of impacts, showing that the previously presented conclusions are valid. Throughout the following analysis, a percentage of impacts above 99.73% (corresponding to 3σ) is considered acceptable. As can be noticed, manifolds with at least 3σ impacts only occur for stability indices below 6.5; however, it is worth noting that there are also some manifolds with lower values than 6.5 which have less than 3σ impacts.

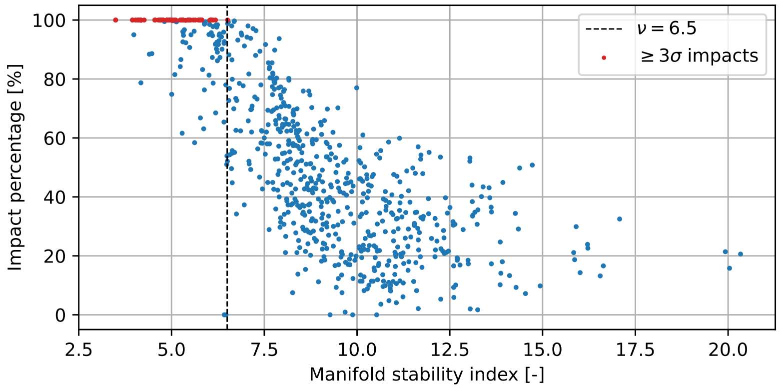

Figure 6.5: Percentage of impacts obtained from the Monte-Carlo analysis as a function of the stability index of the nominal manifold.

As would be expected, the manifolds with 3σ impacts tend to have short times of flight, most being lower than 2 revolutions. However, there are some cases with higher values, up to 10 revolutions. For the tested families, there always appears to be an optimal region for producing manifolds with 3σ impacts. These regions usually occur somewhere in the middle of each family, with manifolds with 3σ impacts being produced by orbits with z-amplitude between 22 and 37 km (these are the bounds for all tested families, within each family the values vary, though still being similar to these). This behavior corresponds to a balance between the instability associated with the propagation time and the instability associated with the dynamics in the proximity of the periodic orbit, in accordance with the results presented in Fig. 6.3. Low-altitude orbits are not very unstable, with their manifolds tending to spend a long time in the vicinity of the orbit, thus increasing the instability of the manifold; meanwhile, for high-altitude orbits, which are more unstable, the manifolds quickly get unstable due to the dynamics, even for relatively short propagation times.

Analyzing the nominal impact locations for the manifolds with more than 3σ impacts, one notes that these tend to be concentrated near the Mars edge, with some impacts in the proximity of the anti-Mars edge (Fig. 6.6). As discussed when analyzing the stability index, no strong effect of the impact location on the manifolds' stability was observed, so this distribution might simply be a result of the larger number of tested manifolds with impact locations on the Mars edge. Regarding the impact locations for the manifolds propagated with uncertainties (i.e. the results of the Monte-Carlo analysis for the cases with more than 3σ impacts): for approximately half the manifolds, all the impacts were observed to be located on the same edge as the nominal impact location; for the other half, most impacts would be located on the nominal edge and the others on the opposite edge.

For the manifolds with more than 3σ impacts, nominal impact velocities between 14 $\mathrm{m\,s^{-1}}$ and 16.5 $\mathrm{m\,s^{-1}}$

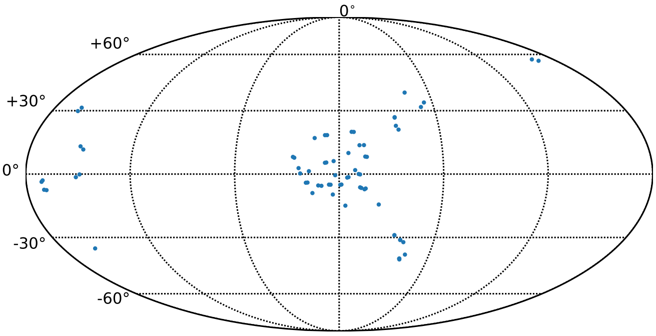

Figure 6.6: Nominal impact location for the manifolds with more than 99.73% of successful impacts.

were found. Thus, these manifolds would allow a very robust deployment of scientific payloads with higher allowed impact velocities than what was used as requirement here (e.g. 15 $\mathrm{m\,s^{-1}}$, as considered by Çelik et al., 2019), without the need for any kind of orbit control or braking maneuvers.

Overall, there are several invariant manifolds that allow a very robust transfer from the initial QSO to the surface, which could be used even without the execution of corrective maneuvers while on the manifold. There is also a very large number of manifolds associated with lower impact percentages; depending on the mission requirements, these could also be used, but they would likely require active control of the spacecraft's trajectory while on the manifold to ensure impacts with the surface. It is possible that the same techniques applied to systems where the secondary has a larger gravitational influence, with a sphere of influence above the surface, would result in a larger number of robust impact manifolds.

Finally, it was shown that the stability index can be used effectively as an auxiliary tool for mission design, as it indeed allows straightforwardly selecting robust manifolds (although a high stability index does not necessarily mean a robust manifold).

6.2. Braking-Maneuver to Surface

Four manifolds were tested, including high and low impact-angle unpowered manifolds (considered, respectively, to correspond to angles above 60° and below 30°), and impacts on the Mars and trailing edge. Impacts on the anti-Mars and leading edges were not tested, as those were expected to behave similarly to the other two edges. All manifolds were selected to have impacts at low latitudes (below 30°). Additionally, all manifolds were selected to depart from orbits of the same family ($1:6-xz/x(BE)$ family) and with similar z-amplitude, even though the departure orbit does not influence directly the final part of the trajectory. The properties of the tested manifolds are summarized in Tab. 6.1.

Table 6.1: Unpowered impact angle, powered impact latitude, powered impact longitude, maneuver altitude, and time between the maneuver and impact, for the manifolds used in the uncertainty analysis for the braking-maneuver to surface case.

Impact angle [deg]	Edge	Latitude [deg]	Longitude [deg]	Maneuver altitude [m]	Time [rev]
16.666	Mars	-28.08	-10.08	107.6	0.00648
68.094	Mars	15.11	18.75	40.9	0.00164
18.158	Trailing	5.16	96.82	82.5	0.00467
61.176	Trailing	0.26	73.09	40.5	0.00168

In all executed Monte-Carlo runs (10^4 per manifold), all trajectories impacted the surface, likely as a result of the low altitude at which the maneuvers are applied. Thus, the likelihood of impacting the surface is governed solely by the part of the trajectory between the initial periodic orbit and the maneuver location.

Impact Location

The effect of the uncertainties on the impact location is represented in Fig. 6.7, for the two manifolds impacting the Mars edge (Tab. 6.1). The effect of the gravitational-parameter uncertainty is not represented individually, as it was observed to have a negligible effect, with latitude and longitude changes in the order of 10^{-5} to 10^{-4} deg. This weak influence results from the short times of flight and the small magnitude of the uncertainty.

Regarding the low impact-angle manifold, both the state and ΔV uncertainties have a relevant contribution to the impact dispersion, with the total dispersion resulting from the combination of these two individual

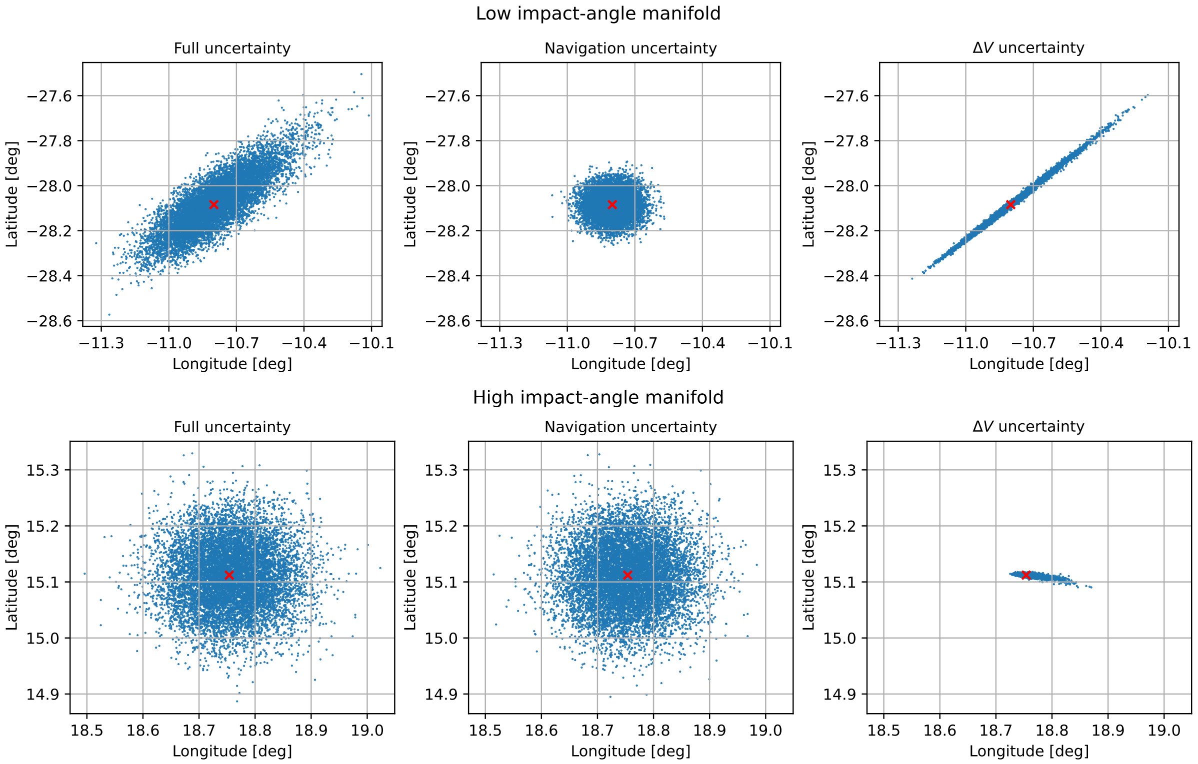

Figure 6.7: Distribution of the impact locations for the uncertainty analysis of the braking-maneuver to impact phase. Data is represented for a low (top) and a high (bottom) impact-angle manifold, with impacts near the Mars edge. Different uncertainties are represented: all uncertainties (left), only navigation uncertainties (center), and only ΔV uncertainties (right). The nominal impact location is marked with a red cross.

effects. The state uncertainty produces an approximately circular impact dispersion, as both the initial position and velocity are normally distributed around zero (the initial errors effectively forming a sphere). Meanwhile, the ΔV uncertainty produces impacts distributed along a thin ellipse, which approximately follows the ground track of the nominal manifold.

Regarding the high impact-angle manifold, the total impact dispersion is dominated by the state uncertainty (which has a circular distribution), with the ΔV error having a much smaller effect. The relatively small effect produced by the ΔV error is likely related to the lower maneuver altitude (allowing less time for the errors to affect the trajectory), and to the fact that the direction of the applied ΔV is closer to the direction of the pre-maneuver velocity than in the low impact-angle case, meaning that ΔV errors will have less effect in the direction of the post-maneuver velocity.

Interestingly, even though the low and high impact-angle manifolds have very different maneuver altitudes and times of flight (between the maneuver and the impact), the state uncertainty produced very similarly sized impact-location distributions for the two: the high impact-angle manifold has longitude and latitude 3σ values of, respectively, $0.173°$ and $0.182°$, and the low impact-angle one of, respectively, $0.157°$ and $0.189°$. This similarity is likely induced by the short propagation times, with the final position dispersion appearing to be a direct result of the initial position error (i.e. the propagation times are not long enough for the different initial states to diverge/converge from/to each other).

The impact patterns observed for the manifolds impacting the trailing edge are identical to the ones discussed for the Mars edge, with only small differences in the size of the impact position dispersion. This is likely related to the fact that the maneuver is applied at low altitudes, meaning that the dynamical environment should be similar in the two edges, being dominated by Phobos' gravity (which is approximately downward pointing in both cases).

Overall, the impact location appears to be less sensitive for high impact-angle manifolds, due to the lower influence of ΔV uncertainties. Thus, if the impact location is an important factor, it might be worth focusing on manifolds that naturally produce high impact angles.

Impact Velocity

The effect of the uncertainties on the horizontal and vertical impact velocities is represented in Fig. 6.8 for the two manifolds impacting the Mars edge. Since the horizontal impact velocity is defined as the component of the velocity parallel to the surface, it always takes values larger or equal to zero. The effect of the uncertainty in Phobos' gravitational parameter is not represented in the plots as, again, it was observed to have a negligible influence, in the order of 10^{-5} $\mathrm{m\,s^{-1}}$ to 10^{-4} $\mathrm{m\,s^{-1}}$.

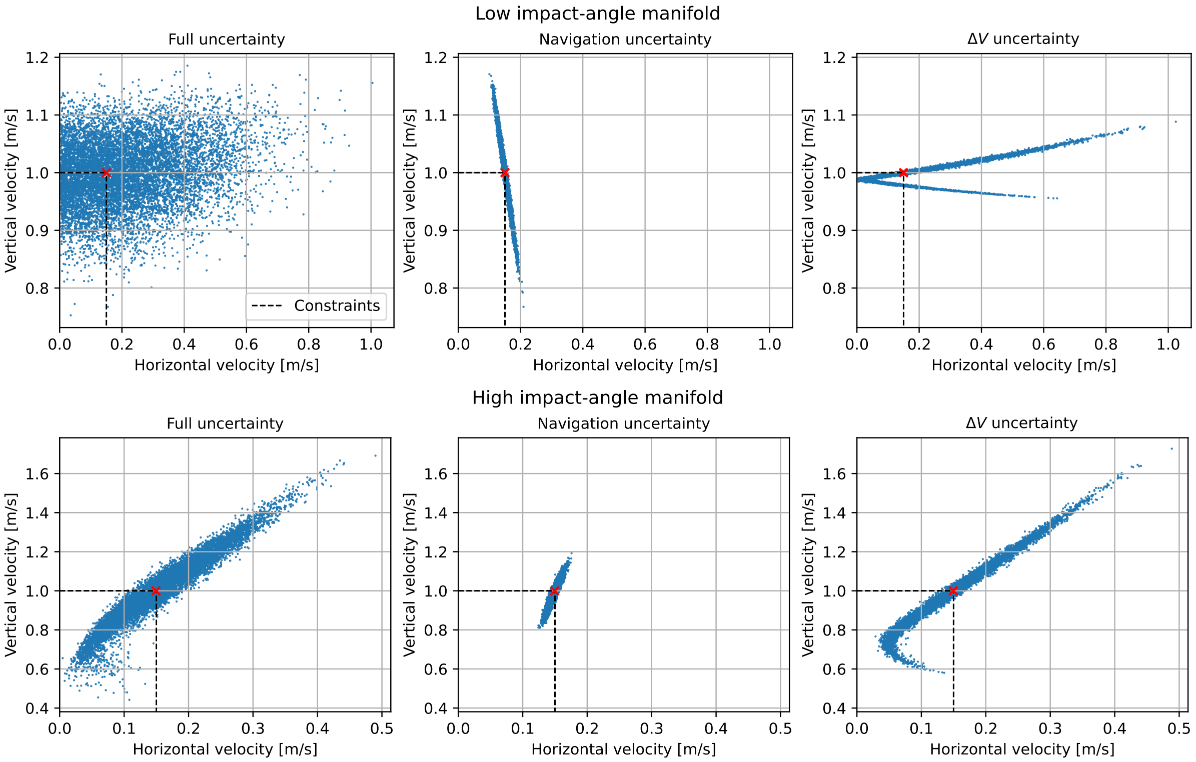

Figure 6.8: Distribution of the horizontal and vertical impact velocities for the uncertainty analysis of the braking-maneuver to impact phase. Data is represented for a low (top) and a high (bottom) impact-angle manifold, with impacts near the Mars edge. Different uncertainties are represented: all uncertainties (left), only navigation uncertainties (center), and only ΔV uncertainties (right). The impact-velocity constraints are represented by black dashed lines. The nominal impact velocities are marked with a red cross.

Analyzing the low impact-angle manifold with full uncertainties, one observes that the impact velocities are distributed over half of an ellipse, with a large percentage of constraint violations, of 78.4%. There are relatively small violations of the vertical-velocity constraint (below 0.2 $\mathrm{m\,s^{-1}}$), but much larger violations of the horizontal-velocity one (up to 0.85 $\mathrm{m\,s^{-1}}$). The high impact-angle manifold also displays a large (but lower) percentage of constraint violations, 54.5%. However, the impact velocities have a significantly different distribution, which is ellipse-shaped and has significantly larger vertical-velocity constraint violations (up to 0.7 $\mathrm{m\,s^{-1}}$), but lower violations on the horizontal velocity (below 0.5 $\mathrm{m\,s^{-1}}$).

In both cases, the dispersion in the impact velocities is mainly driven by the ΔV uncertainties, with the navigation uncertainties having a smaller effect. The navigation uncertainties produce similar constraint violations for both manifolds and mainly influence the vertical-impact velocity (violations of around 0.2 $\mathrm{m\,s^{-1}}$ for the vertical velocity, and 0.05 $\mathrm{m\,s^{-1}}$ for the horizontal). However, the distributions have different orientations for the two manifolds: for the high impact-angle one, the two velocity components increase in the same direction, while for the low impact-angle case the vertical velocity decreases as the horizontal one increases. In fact, this difference is not related to the impact angle of the manifold, but to the shape of the trajectory relative to the direction of the local acceleration field. Due to the proximity to Phobos, the vertical acceleration always points down, meaning that the vertical velocity increases with increasing time of flight (Fig. 6.9, right). In cases where the horizontal acceleration has a direction contrary to the horizontal velocity, the horizontal velocity decreases for increasing time of flight (Fig. 6.9, left; blue, green, and red curves), hence leading to the pattern observed here for the low impact-angle manifold. Meanwhile, in cases where the horizontal acceleration is aligned with the horizontal velocity, the latter increases over time (Fig. 6.9, left; orange curve), thus leading to the pattern observed for the high impact-angle manifold.

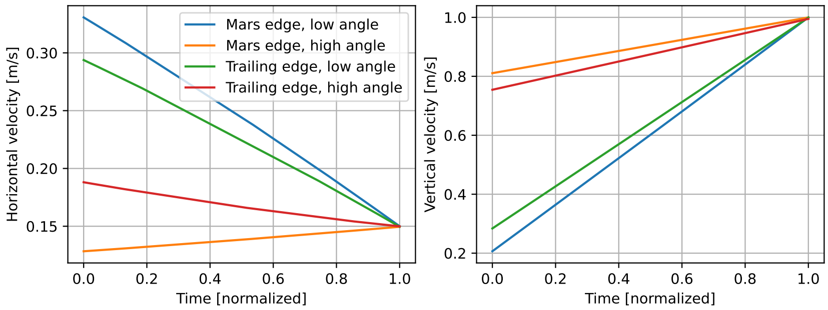

Figure 6.9: Horizontal and vertical velocity as a function of time, for each of the tested nominal manifolds for the uncertainty analysis of the braking-maneuver to impact phase. The time is normalized using each manifold's time of flight, therefore steeper lines do not necessarily mean stronger accelerations.

As mentioned, the ΔV uncertainties have a much stronger effect on the impact velocities than the navigation uncertainties. For the low impact-angle manifold, the ΔV uncertainties mainly influence the horizontal impact velocity (constraint violations in the vertical and horizontal velocities up to, respectively, $0.1\,\mathrm{m\,s^{-1}}$ and $0.85\,\mathrm{m\,s^{-1}}$), while for the high impact-angle manifold the uncertainties mainly influence the vertical impact velocity (constraint violations in the vertical and horizontal velocities up to, respectively, $0.6\,\mathrm{m\,s^{-1}}$ and $0.35\,\mathrm{m\,s^{-1}}$). Thus, the ΔV uncertainties mainly affect the velocity in the direction in which they were applied, with the ΔV having a larger horizontal component for the low impact-angle manifold, and a larger vertical component for the high-impact angle one.

It is also possible to note that the ΔV uncertainties form two branches in the distribution of the velocities, in one part with the vertical velocity increasing as the horizontal velocity increases, and in another with the contrary variation. This is also induced by the direction of the acceleration during the descent, with the observed distribution corresponding to the balance between the acting vertical and horizontal accelerations over the time of flight. In the case of the low impact-angle manifold, the two branches reach a horizontal velocity of zero (corresponding to an impact angle of 90°), which does not occur for the high impact-angle manifold. This happens because, for the low impact-angle manifold, the horizontal velocity decreases over time (Fig. 6.9, left), crossing zero for perturbed trajectories with long-enough time of flight; meanwhile, for the high impact-angle manifold the horizontal velocity increases with time, never reaching zero. The same type of behavior was observed for the two manifolds impacting the trailing edge.

Finally, some discussion regarding the distribution of the impact velocities with full uncertainties is required. As already pointed out, the majority of the tested cases violate the constraints, with a larger number of constraint violations for the low impact-angle manifolds (78.4% and 76.6% of constraint violations, respectively, for the Mars and trailing edge) than for the high impact-angle ones (54.5% and 62.2% of constraint violations, respectively, for the Mars and trailing edge). Part of this issue results from the fact that the maneuvers were designed without any safety margin, targeting the maximum values allowed by the constraints. Hence, taking into account safety margins in the selection of the target impact velocity would contribute to a lower number of constraint violations.

However, even with safety margins, it appears that in this case there would still be a large number of constraint violations. The ΔV uncertainty is observed to have a particularly large effect, and would thus be something to further analyze when designing a real mission. Additionally, the post-maneuver velocity is usually very low, below $1\,\mathrm{m\,s^{-1}}$, thus having the same order of magnitude as the norm of the usual ΔV errors, which results in large relative errors in the post-maneuver velocity. Therefore, it might be necessary to either use a more accurate thruster or to divide the application of the ΔV into multiple maneuvers, with the subsequent maneuvers correcting the errors of the previous ones.

It is also worth noting that, since the dispersion patterns significantly depend on the used manifold, a careful selection of the manifold might allow shaping the dispersion patterns to better fit with the used constraints. In particular, using high impact-angle manifolds would be more appropriate for cases with more strict horizontal-velocity constraints, while low impact-angle manifolds would be more appropriate for strict vertical-velocity constraints. Furthermore, the selection of the impact location, and consequently the acceleration field relative to the trajectory, allows influencing the impact-velocity pattern produced by the navigation uncertainties. Compare Fig. 6.8, bottom, with Fig. 6.10: the two have similar ΔV-induced impact distributions, and similar distributions produced by the navigation uncertainties, though with opposite orientations, leading to quite different impact patterns for the full uncertainties. For a design taking into

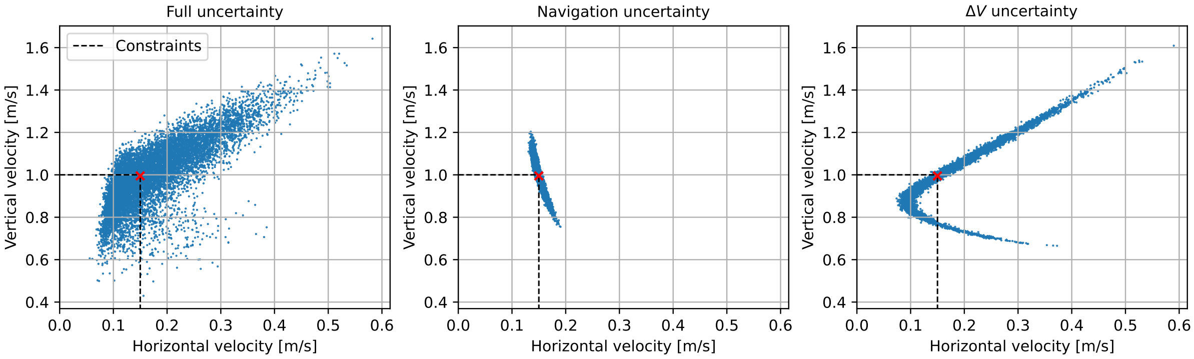

Figure 6.10: Distribution of the horizontal and vertical impact velocities for the uncertainty analysis of the braking-maneuver to impact phase. Data is represented for a high impact-angle manifold, with impacts near the trailing edge. Different uncertainties are represented: all uncertainties (left), only navigation uncertainties (center), and only ΔV uncertainties (right). The impact-velocity constraints are represented by black dashed lines. The nominal impact velocities are marked with a red cross.

account safety margins, the case of Fig. 6.10 would allow a lower number of constraint violations. The selection of the impact location so as to produce favorable navigation-uncertainty dispersion patterns could be especially relevant for cases where the navigation uncertainties are more important than here (for example, due to lower ΔV uncertainties).

No significant differences in the distribution of the impact velocities were observed between the Mars edge and trailing edge, indicating that both are equally robust to the uncertainties. As previously mentioned, this is likely related to the fact that the maneuver is applied at a low altitude and at a short distance from the impact location.

6.3. Conclusions

The effect of different uncertainties (navigation, ΔV, Phobos' gravity) on the designed transfers, which follow invariant manifolds and use two impulsive maneuvers, was tested. This analysis was divided into two phases, corresponding to the phases of the transfer initiated by each maneuver: transfer from the periodic orbit to the surface or maneuver location, and transfer from the maneuver location to the surface. In both phases, the uncertainty in Phobos' gravity proved to be negligible, while the other two uncertainty sources showed varying importance.

The transfer between the orbit and the maneuver location was analyzed using a stability index (based on the STM) and Monte-Carlo data. This phase of the transfer, which is mostly governed by the navigation error, is the phase that drives the impact likelihood. The stability of the manifolds was observed to be influenced by both their time of flight and the stability of the orbit from which they emanate. It was shown that there are invariant manifolds that allow robust transfers to the surface, without requiring any kind of orbit control; this is important for any spacecraft but could be especially relevant if deploying to the surface small scientific packages or CubeSats with limited to no capabilities of orbit control. These robust invariant manifolds were observed to allow impacts in the vicinity of the Mars and anti-Mars edges, which, as previously mentioned, are the regions where MMX is planned to land.

The second phase of the transfer, i.e. from the braking-maneuver location to the surface, was observed to not influence the likelihood of impacting the surface, only changing the impact position and velocity. The considered uncertainties would not allow meeting the specified constraints on the impact velocity, with the uncertainties in the ΔV being particularly impactful. The selection of the ΔV with safety margins and the selection of a nominal invariant manifold with favorable impact conditions (angle and location) would allow shaping the impact dispersion, resulting in a higher number of cases meeting the selected landing requirements. However, based on the presented results, this would likely still not be sufficient to satisfactorily comply with the landing requirements. Therefore, if further analyzing the possibility of using invariant manifolds to land, it would be important to give special attention to this final part of the trajectory, for example by using thrusters with lower uncertainties or by designing this final phase differently (e.g. application of the braking maneuver at higher altitudes, allowing smaller corrective maneuvers to be applied afterwards).

7

Verification

Here the steps followed in the verification of the code used for obtaining the presented results are described. The validation of these results is not possible, as Phobos has never been successfully orbited.

7.1. CR3BP with Ellipsoidal Secondary

The verification of the CR3BP with ellipsoidal secondary was executed in three steps. Firstly, the computation of the equilibrium points of a uniformly rotating ellipsoid executed by Scheeres (1994) was reproduced. Scheeres uses a dynamics model including solely the uniformly rotating ellipsoid, therefore the computation of the equilibrium points (which are defined with respect to a synodic frame) serves to verify the correctness of the ellipsoidal gravity model and rotation model, but not its integration into the CR3BP. The equilibria were calculated for the asteroid Vesta, using the following parameters: $G = 6.672 \times 10^{-11}\ \mathrm{m^3\,kg^{-1}\,s^{-2}}$, $a = 265$ km, $b = 250$ km, $c = 220$ km, rotational period of 5.3 h, and $\rho = 3.5\ \mathrm{g\,cm^{-3}}$. The obtained equilibria can be found in Tab. 7.1; as can be observed, the computed positions coincide with the theoretical ones up to the given number of digits.

Table 7.1: Position of the equilibrium points of the asteroid Vesta (modeled by a uniformly rotating ellipsoid), with respect to its body-fixed frame, according to Scheeres (1994) and computed.

Equilibrium point	Coordinate	Scheeres (1994)	Computed	Error
Saddle	x [-]	1.94097	1.940966962	-3.038×10^{-6}
	y [-]	0.0	$2.198955746 \times 10^{-17}$	$2.198955746 \times 10^{-17}$
	z [-]	0.0	0.0	0.0
Center	x [-]	0.0	0.0	0.0
	y [-]	1.92377	1.923771674	1.674×10^{-6}
	z [-]	0.0	0.0	0.0

In a second step, the computation of the Lagrange points L_1 and L_2 executed by Pushparaj et al. (2021) was reproduced. Pushparaj et al. use a dynamics model consisting of the H3BP with the primary (Mars) modeled by a point mass and the secondary (Phobos) by a uniformly rotating ellipsoid; thus, the only difference with respect to the model used here is the linearization of the CR3BP (i.e. the H3BP). The following parameters were used: $G = 6.674 \times 10^{-11}\ \mathrm{m^3\,kg^{-1}\,s^{-2}}$, $a = 13.03$ km, $b = 11.40$ km, $c = 9.14$ km, $1/\dot{\theta} = 4387.63$ s, semi-major axis of 9377 km, and density of Phobos $\rho = 1.860\ \mathrm{g\,cm^{-3}}$. The computed equilibria are very close to the ones predicted by Pushparaj et al. (Tab. 7.2), with differences in the order of meters; the discrepancy can be explained by the fact that, due to the relatively small distance between Phobos and Mars, the H3BP is not a very accurate approximation of the CR3BP (Scheeres et al., 2019). As such, at the initial simulation epoch (for which the Lagrange points were calculated), the model of the CR3BP with ellipsoidal secondary is considered to be set up correctly.

As a final verification step, a QSO generated using Eq. 4.1 for $x(t_1) = 80$ km, with Mars and Phobos as primaries, was propagated (as a single arc) over three periods of Phobos' orbit. As expected, the Jacobi integral remains approximately constant throughout the trajectory (Fig. 7.1), with only small oscillations resulting from

Table 7.2: Position of the L_1 and L_2 Lagrange points of the Mars-Phobos system (with Phobos modeled by a uniformly rotating ellipsoid), in a synodic Phobos-centered frame, according to Pushparaj et al. (2021) and computed.

Point	Coordinate	Pushparaj et al. (2021)	Computed	Error
L_1	x [m]	-17340.0	-17328.8184	11.1816
	y [m]	0.0	$-1.51973028 \times 10^{-7}$	$-1.51973028 \times 10^{-7}$
	z [m]	0.0	0.0	0.0
L_2	x [m]	17340.0	17348.1607	8.1607
	y [m]	0.0	$7.54019786 \times 10^{-19}$	$7.54019786 \times 10^{-19}$
	z [m]	0.0	0.0	0.0

numerical errors. The steps in the variation of the Jacobi integral are a result of the limited accuracy of the floating-point representation, which, for the orbit's Jacobi integral of $C = 2.99993$, has a value of 4.441×10^{-16} (value corresponding to Python's native 64-bit floating-point representation). The constant value of the Jacobi integral serves as a final indication of the correctness of the used model, showing that the system is indeed autonomous (i.e. the ellipsoid rotates uniformly); furthermore, this constant value inherently verifies the computation of the Jacobi integral itself (and by extension the computation of the gravitational potential).

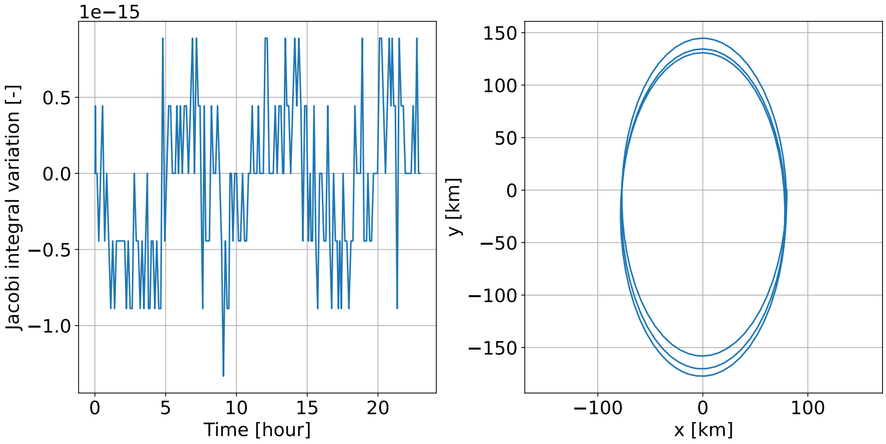

Figure 7.1: Jacobi integral variation with respect to its initial value as a function of time (left) and y-position as a function of x-position (right), for a Phobos-centered QSO (with Phobos modeled by a uniformly rotating ellipsoid).

Other elements of the developed code, for example, the transformation of the state between an inertial and a synodic frame (and vice-versa) and the transformation between dimensional and dimensionless coordinates (and vice-versa), are verified by the previous steps.

7.2. CR3BP with Polyhedral Secondary

As a first verification step, a polyhedron model of the cuboid represented in Fig. 7.2 (with dimensions 20 × 10 × 10 m) was used. The potential, its gradient, and its Hessian were computed at the points $P_a = (0,0,0)$ (point in a vertex), $P_b = (5,0,0)$ (point in an edge), and $P_c = (0,3,2)$ (point in a facet), and the results compared with D'Urso (2014). Additionally, the Laplacian of the potential was computed at these points and at the points $P_d = (-5,5,5)$ (point outside the body) and $P_e = (10,5,5)$ (point inside the body), and the results compared with their theoretical values (presented in Subsection 2.2.3). The obtained results, using $G = 6.67259 \times 10^{-11}\ \mathrm{m^3\,kg^{-1}\,s^{-2}}$ and $\rho = 2.670\ \mathrm{g\,cm^{-3}}$, are listed in Tab. 7.3. It is possible to observe that the values computed for the potential, its gradient and its Hessian are consistent with D'Urso (2014) up to the number of digits provided by the latter; the values of the Laplacian coincide with the theoretical values, up to the floating-point representation accuracy. Thus, the equations describing the polyhedron gravity can be concluded to have been implemented correctly.

Next, to verify the correctness of the polyhedron model of Phobos (with 1000 facets), its acceleration was compared with a spherical harmonics gravity model. The spherical harmonics coefficients up to degree and

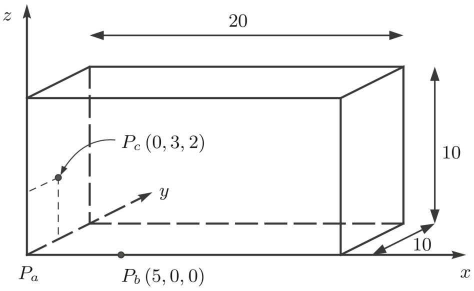

Figure 7.2: Representation of the cuboid used in the verification of the polyhedron model. Adapted from D'Urso (2014).

Table 7.3: Potential of a cuboid, as well as its gradient, Hessian and Laplacian, at different points, according to D'Urso (2014), theory, and computed.

Location	Quantity	D'Urso (2014)[a]/Theoretical[b]	Computed	Error
P_a	U [$\mathrm{m^2\,s^{-2}}$]	[a]$3.19403761604211 \times 10^{-5}$	$3.1940376160421150 \times 10^{-5}$	5.42×10^{-20}
	U_x [$\mathrm{m\,s^{-2}}$]	[a]$2.31329148957265 \times 10^{-6}$	$2.3132914895726513 \times 10^{-6}$	1.27×10^{-21}
	U_y [$\mathrm{m\,s^{-2}}$]	[a]$1.91973919943187 \times 10^{-6}$	$1.9197391994318700 \times 10^{-6}$	0.0
	U_z [$\mathrm{m\,s^{-2}}$]	[a]$1.91973919943187 \times 10^{-6}$	$1.9197391994318700 \times 10^{-6}$	0.0
	$-\nabla^2 U/(G\rho)$ [-]	[b]$\in [0, 4\pi]$	0.5000000000000001π	-
P_b	U [$\mathrm{m^2\,s^{-2}}$]	[a]$3.99993558939122 \times 10^{-5}$	$3.9999355893912276 \times 10^{-5}$	7.45×10^{-20}
	U_x [$\mathrm{m\,s^{-2}}$]	[a]$9.90115534890074 \times 10^{-7}$	$9.9011553489007030 \times 10^{-7}$	-3.60×10^{-21}
	U_y [$\mathrm{m\,s^{-2}}$]	[a]$3.24128042248715 \times 10^{-6}$	$3.2412804224871538 \times 10^{-6}$	3.81×10^{-21}
	U_z [$\mathrm{m\,s^{-2}}$]	[a]$3.24128042248715 \times 10^{-6}$	$3.2412804224871538 \times 10^{-6}$	3.81×10^{-21}
	$-\nabla^2 U/(G\rho)$ [-]	[b]$\in [0, 4\pi]$	1.0000000000000004π	-
P_c	U [$\mathrm{m^2\,s^{-2}}$]	[a]$4.03528375471853 \times 10^{-5}$	$4.0352837547185350 \times 10^{-5}$	4.74×10^{-20}
	U_x [$\mathrm{m\,s^{-2}}$]	[a]$4.73368592565013 \times 10^{-6}$	$4.7336859256501306 \times 10^{-6}$	8.47×10^{-22}
	U_y [$\mathrm{m\,s^{-2}}$]	[a]$9.68164362892554 \times 10^{-7}$	$9.681643628925530 \times 10^{-7}$	-1.06×10^{-21}
	U_z [$\mathrm{m\,s^{-2}}$]	[a]$1.59674500375495 \times 10^{-6}$	$1.5967450037549484 \times 10^{-6}$	-1.69×10^{-21}
	U_{xx} [$\mathrm{s^{-2}}$]	[a]$-4.02204713784183 \times 10^{-8}$	$-4.022047137841809 \times 10^{-8}$	2.12×10^{-22}
	U_{yy} [$\mathrm{s^{-2}}$]	[a]$-5.01781942367494 \times 10^{-7}$	$-5.0178194236749376 \times 10^{-7}$	2.12×10^{-22}
	U_{zz} [$\mathrm{s^{-2}}$]	[a]$-5.77398275537941 \times 10^{-7}$	$-5.7739827553794116 \times 10^{-7}$	-2.12×10^{-22}
	U_{xy}, U_{yx} [$\mathrm{s^{-2}}$]	[a]$1.87140408935899 \times 10^{-7}$	$1.8714040893589887 \times 10^{-7}$	-1.32×10^{-22}
	U_{xz}, U_{zx} [$\mathrm{s^{-2}}$]	[a]$3.51261972418670 \times 10^{-7}$	$3.5126197241867027 \times 10^{-7}$	2.65×10^{-22}
	U_{yz}, U_{zy} [$\mathrm{s^{-2}}$]	[a]$8.58712984897779 \times 10^{-8}$	$8.5871298489778694 \times 10^{-8}$	7.94×10^{-22}
	$-\nabla^2 U/(G\rho)$ [-]	[b]2π	2.000000000000001π	8.88×10^{-16}
P_d	$-\nabla^2 U/(G\rho)$ [-]	[b]0	$6.184411880623507 \times 10^{-17}\pi$	6.18×10^{-17}
P_e	$-\nabla^2 U/(G\rho)$ [-]	[b]4π	4.000000000000001π	8.88×10^{-16}

order 20, with reference radius $R = 14.0$ km, determined by Hu and Jekeli (2015) were used[1]. These coefficients were obtained based on the shape model developed by Willner et al. (2010), which is also the model that serves as the basis for the simplified polyhedron used here. To compare the polyhedron and spherical harmonics models, the position of the former in the body-fixed frame was corrected to ensure that both models have the same COM. The position of the COM according to the spherical harmonics model was computed through (Scheeres, 2012)

$$\begin{aligned} x_{COM} &= C_{1,1} R \\ y_{COM} &= S_{1,1} R \\ z_{COM} &= C_{1,0} R \end{aligned} \tag{7.1}$$

where, as previously, R represents the spherical harmonics reference radius, and $C_{n,m}$ and $S_{n,m}$ the unnor-

---

[1] The spherical harmonics coefficients were kindly provided by Dr. Xuanyu Hu via personal communication.

malized spherical harmonics coefficients. The relative acceleration error between the two gravity models (represented in Fig. 7.3) was computed through Eq. 8.1 (considering the polyhedron gravity to correspond to the "truth"), for points located in a sphere with 15.0 km radius, and using the gravitational parameter of Phobos listed in Tab. 2.1. Maximum and mean errors of, respectively, 2.0546% and 0.4481% were obtained; the maximum error is observed in the region with approximate longitude and latitude of, respectively, −35° and −5°, corresponding to the rim of the Stickney crater. The error between the polyhedron and spherical harmonics model is likely related to differences in the sampling of the initial shape model determined by Willner et al. (2010): Hu and Jekeli (2015) determine the spherical harmonics coefficients from a polyhedron model with a 3° resolution in latitude and longitude, giving a total of 7200 vertices, while the polyhedron model used here was obtained by simplifying a model with 128002 vertices (256000 facets). Based on these considerations, and on the fact that the observed errors are relatively low, the used polyhedron model is considered to be correct in its description of Phobos.

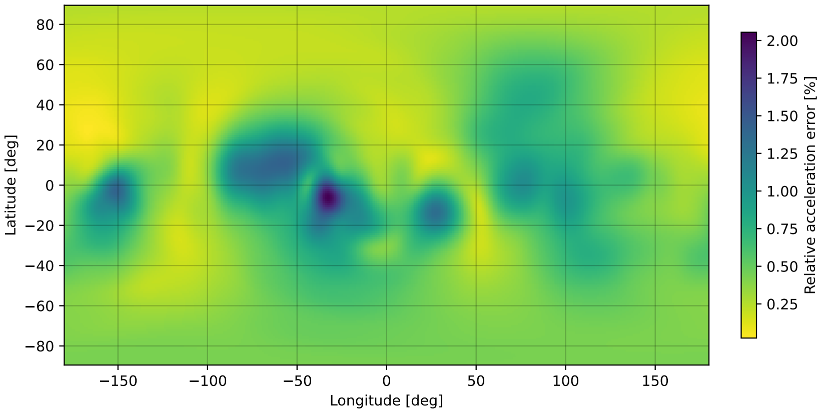

Figure 7.3: Relative acceleration error between the polyhedron model of Phobos with 1000 facets, and the spherical harmonics gravity model up to degree and order 20 determined by Hu and Jekeli (2015). The relative acceleration error is presented as a function of the latitude and longitude, for points in a sphere with 15.0 km radius.

To verify the correct integration of the polyhedron model into the CR3BP, the computation of the Lagrange points L_1 and L_2 executed by Çelik et al. (2019) was reproduced. Çelik et al. use a dynamics model consisting of the H3BP, with the primary (Mars) modeled by a point mass and the secondary (Phobos) by a uniformly rotating polyhedron. The differences with respect to the dynamics model used here are the linearization of the CR3BP (i.e. the H3BP), and the polyhedron model used: Çelik et al. used the Phobos model developed by Gaskell (2006), with approximately 50000 facets, while here a simplification with 1000 facets of the model developed by Willner et al. (2010) is used. When computing the Lagrange points, the following parameters were used: $G = 6.674 \times 10^{-11}\ \mathrm{m^3\,kg^{-1}\,s^{-2}}$, semi-major axis of 9377.2 km, mass of Mars of 6.42×10^{23} kg, mass of Phobos of 1.07×10^{16} kg and density of Phobos $\rho = 1.872\ \mathrm{g\,cm^{-3}}$. The computed equilibria are very close to the ones predicted by Çelik et al. (Tab. 7.4), with differences in the order of tens of meters. This discrepancy is likely related to the differences in the polyhedron models used (which were generated based on different data), and to the fact that, as already mentioned, the H3BP is not a very accurate approximation of the Mars-Phobos CR3BP.

Table 7.4: Position of the L_1 and L_2 Lagrange points of the Mars-Phobos system (with Phobos modeled by a uniformly rotating polyhedron), in a synodic Phobos-centered frame, according to Çelik et al. (2019) and computed.

Point	Coordinate	Çelik et al. (2019)	Computed	Error
L_1	x [m]	-17298.9	-17334.9811	36.0811
	y [m]	394.5	369.3403	25.1597
	z [m]	-258.1	-246.1494	-11.9506
L_2	x [m]	17295.5	17311.5963	-16.0963
	y [m]	452.6	503.1576	-50.5576
	z [m]	-32.5	-68.9829	36.4829

Finally, to verify that the rotation model of Phobos is set up correctly, a QSO generated with Eq. 4.1 for $x(t_1) = 80$ km, with Mars and Phobos as primaries, was propagated over three periods of Phobos' orbit (Fig. 7.4).

As can be observed, the Jacobi integral remains approximately constant throughout the orbit, with only small oscillations resulting from numerical errors; as before, the regular steps in the variation of the Jacobi integral are associated with the limits in the accuracy of the floating-point representation. The conservation of the Jacobi integral proves that the simulated system is indeed autonomous.

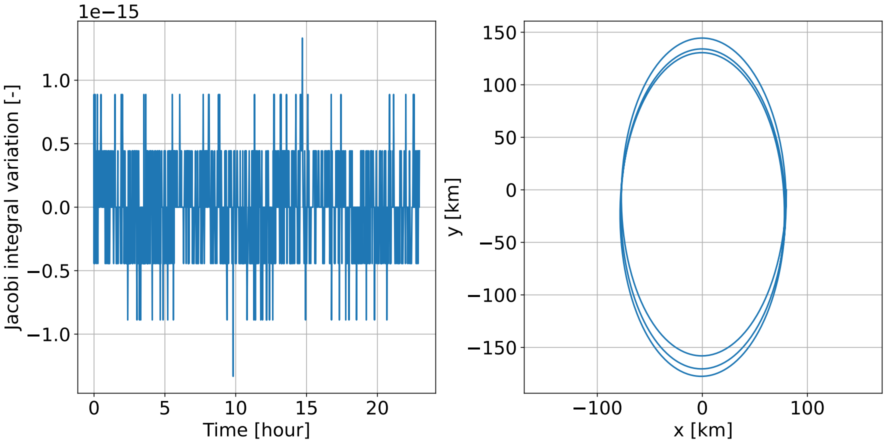

Figure 7.4: Jacobi integral variation with respect to its initial value as a function of time (left) and y-position as a function of x-position (right), for a Phobos-centered QSO (with Phobos modeled by a uniformly rotating polyhedron with 1000 facets).

7.3. Multiple Shooting and Continuation

In order to verify the implementation of the continuation procedure coupled with the multiple-shooting corrector, the family of two-dimensional QSOs of the Mars-Phobos system was computed in the CR3BP with ellipsoidal secondary and CR3BP with polyhedral secondary, using the parameters listed in Tab. 2.1. Both models were used because the verification steps presented below not only verify the continuation procedure but also the propagation of the variational equations.

To check whether the differential corrector produced continuous orbits, the root sum square of the Jacobi error between consecutive arcs (i.e. the difference between the Jacobi integral at the end of one arc and the Jacobi integral and the start of the following arc) was computed for each orbit in the family (Fig. 7.5). Throughout the family, and for both models (polyhedron and ellipsoid), the root sum square of the error shows values in the order of 10^{-15}, indicating that the orbits are indeed continuous.

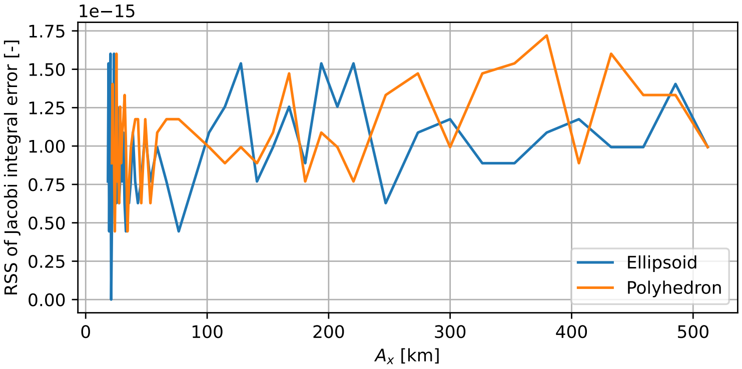

Figure 7.5: Root sum square of the Jacobi integral error between consecutive arcs of each orbit, as a function of the x-amplitude of each orbit, for the two-dimensional QSO family, computed using ellipsoidal and polyhedral models of Phobos.

Since the multiple-shooting corrector was implemented using Newton's method, which is known to converge quadratically (i.e. rate of convergence equal to or larger than 2), its performance can be assessed by estimating the convergence rate observed in practice. The convergence rate α can be estimated according to

(Senning, 2019)

$$\alpha \approx \frac{\log \dfrac{\|\boldsymbol{V}_i - \boldsymbol{V}_{i-1}\|}{\|\boldsymbol{V}_{i-1} - \boldsymbol{V}_{i-2}\|}}{\log \dfrac{\|\boldsymbol{V}_{i-1} - \boldsymbol{V}_{i-2}\|}{\|\boldsymbol{V}_{i-2} - \boldsymbol{V}_{i-3}\|}} \tag{7.2}$$

where, as previously, $\boldsymbol{V}_i$ represents the free variables vector obtained in the i^{th} multiple-shooting iteration. As can be observed (Fig. 7.6), for a large part of the computed orbits in the two-dimensional QSO family, the estimated convergence rate remains above $\alpha = 2$; there are, however, some regions (orbits with x-amplitudes between 30 km and 100 km), for which the convergence rate takes values between 1.5 and 2. Nevertheless, it is important to note that Newton's method is only guaranteed to have a converge rate larger than 2 for an initial guess that is close enough to the true solution (the values that would allow establishing whether an initial guess is good enough were not estimated here). The fact that the convergence rate tends to increase with the x-amplitude (up to approximately 300 km) should be a result of the decreasing influence of the perturbation produced by the secondary's gravity, leading to smoother dynamics and consequently a faster convergence as the altitude increases. Above 300 km, the convergence rate remains approximately constant; this is likely a region where the influence of the secondary's gravity is already negligible, and, as a result, the dynamics are approximately constant even with changing x-amplitude. Overall, the fact that the observed convergence rates are either higher than or close to 2 serves as an indication that Newton's method is performing as expected.

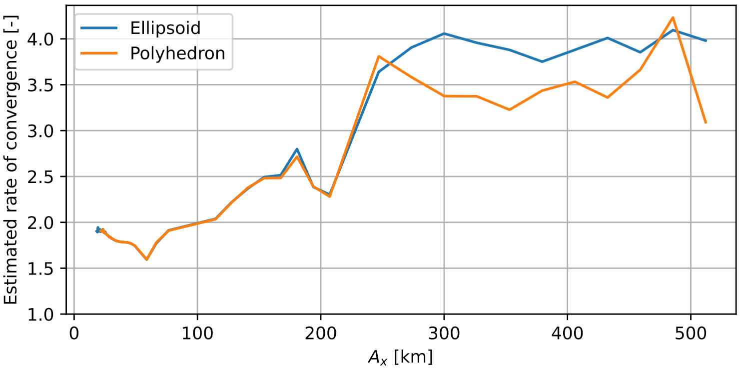

Figure 7.6: Estimated convergence rate of the multiple-shooting corrector as a function of the x-amplitude of each orbit, for the orbits in the two-dimensional QSO family, computed using an ellipsoidal and a polyhedral model of Phobos.

Finally, the eigenvalues of the monodromy matrix are verified. This serves as an additional verification of the multiple-shooting procedure and verifies the computation of the monodromy matrix itself (the propagation of each arc's STM was inherently verified by the previous steps). As stated in Subsection 3.1.2, the monodromy matrix of periodic orbits has a pair of unit eigenvalues (Fig. 7.7, left). The decrease in the error to unity as the x-amplitude increases is, as before, likely related to the decrease in the influence of the secondary's perturbation, making it easier for the multiple-shooting algorithm to achieve periodicity. Since the monodromy matrix is symplectic, its eigenvalues also appear in conjugate pairs (Fig. 7.7, right). The monodromy matrix is symplectic because it is determined from the multiplication of the STM (also symplectic) of each arc; therefore, the error between conjugate eigenvalues is simply a result of the integration accuracy, not being related to the periodicity (or even continuity) of the orbit. Thus, the approximately constant behavior of the error with changing x-amplitude can be explained by the usage of a variable step-size integrator applied to relatively similar orbits, resulting in a similar number of integration steps between orbits. The difference in magnitude between the error for the pairs $i = 2$ and $i = 3$ should be a result of the fact that the former is associated with planar motion (in the xy-plane) and the latter with vertical motion; since the family of computed orbits is either two-dimensional (with ellipsoid) or nearly two-dimensional (with polyhedron), the eigenvalues associated with three-dimensional motion have a lower error. Do note that although the eigenvalues of the STM are related to the response to perturbations, their computation depends on the Jacobian of the system's dynamics, which in turn depends on the state itself; hence the influence of the state on the accuracy of the eigenvalues. Overall, both types of errors in the monodromy matrices' eigenvalues show very low values, indicating that the latter's computation is being executed correctly.

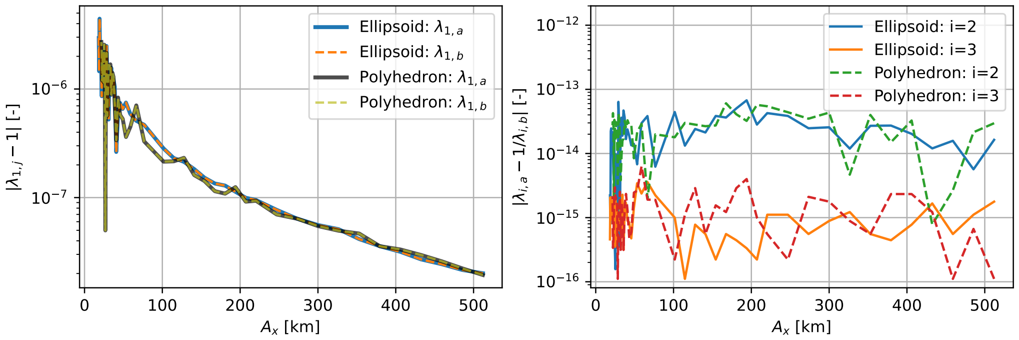

Figure 7.7: Error associated with the eigenvalues of the monodromy matrix as a function of the x-amplitude of each orbit: error to unity for the trivial eigenvalue pair (left) and error between conjugate eigenvalues for the non-trivial pairs (right). The errors refer to the two-dimensional QSO family, computed using ellipsoidal and polyhedral models of Phobos.

7.4. Bifurcations

To verify the detection of bifurcating orbits, the computation of (some of the) vertical period-multiplying bifurcations of the two-dimensional QSO family in the Jupiter-Europa system executed by Lara et al. (2007) was reproduced. The CR3BP with mass parameter $\mu = 2.528 \times 10^{-5}$ was used as dynamics model. The period of the bifurcating orbits and the x and $\dot{y}$ coordinates at the intersection with the $y = 0$ plane are listed in Tab. 7.5. The computed values agree very well with the ones presented by Lara et al.; the small discrepancies are likely related to differences in the tolerance of the multiple shooting procedure and the tolerance of the root-finding method used to compute the bifurcations. Hence, the computation of bifurcations is considered to be implemented correctly. Note that, since the detection of the bifurcations is based on the stability indices, this also further verifies the computation of the monodromy matrix and its eigenvalues.

Table 7.5: Dimensionless state and period of different vertical period-multiplying bifurcations of the two-dimensional QSO family in the Jupiter-Europa system, according to Lara et al. (2007) and computed.

Bifurcation	Quantity	Lara et al. (2007)	Computed	Error
2:11	x [-]	0.012584843180049 71	0.01258484317569530	-4.3544×10^{-12}
	$\dot{y}$ [-]	-0.05967592114004498	-0.0596759211418033	-1.7583×10^{-12}
	T [-]	1.465846672906467	1.465846672246029	-6.6044×10^{-10}
1:7	x [-]	0.009833151434219369	0.00983315148325535	4.9036×10^{-11}
	$\dot{y}$ [-]	-0.06188292069217829	-0.0618829206292504	6.2928×10^{-11}
	T [-]	1.058008159074479	1.058008166143251	7.0688×10^{-9}
1:22	x [-]	0.003951211985392241	0.00395121196917648	-1.6216×10^{-11}
	$\dot{y}$ [-]	-0.0841082661756476	-0.0841082663219096	-1.4626×10^{-10}
	T [-]	0.2967692219374928	0.296769220184971	-1.7525×10^{-9}

7.5. Invariant Manifolds (Unstable)

The propagation of invariant manifolds is based on the computation of the eigenvalues and eigenvectors of the monodromy matrix, and on the propagation of the unstable eigenvector over the different arcs that constitute the relevant orbit, using the STM associated with each of those arcs. The computations of each arc's STM, and of the monodromy matrix and its eigenvalues have already been verified. Since the eigenvalues and eigenvectors are computed using the same function from Python's NumPy library, the computation of the eigenvectors is also considered verified. Therefore, it is only left to prove that the computation of the perturbation applied at each node to obtain the initial state of the manifold is executed correctly. To do this, the propagation of the unstable manifolds of a planar-Lyapunov orbit near the L_1 point of the Earth-Moon system executed by Vaquero (2013) was reproduced. The CR3BP model was used, with mass parameter $\mu = 0.01215057$, distance between primaries 384.0×10^3 km and Earth's gravitational parameter 3.986×10^5 km^3 s^{-2}. The Lyapunov orbit was propagated using initial state

$$\boldsymbol{x}(t_1) = [3.230816 \times 10^5, 0, 0, 0, -0.036946, 0]^T \quad \text{km, km s}^{-1} \tag{7.3}$$

and period 11.69886 days (multiple shooting was applied to correct a small final-state discontinuity). The manifolds were propagated over 21.7 days, starting from 40 nodes equally spaced in time along the periodic orbit, and using a position perturbation of 30.0 km. It is possible to observe that the computed unstable manifolds are very similar to the ones presented by Vaquero (Fig. 7.8). This, associated with the fact that a large part of the necessary elements of the code was already verified, gives a good indication that the unstable manifolds are being propagated correctly.

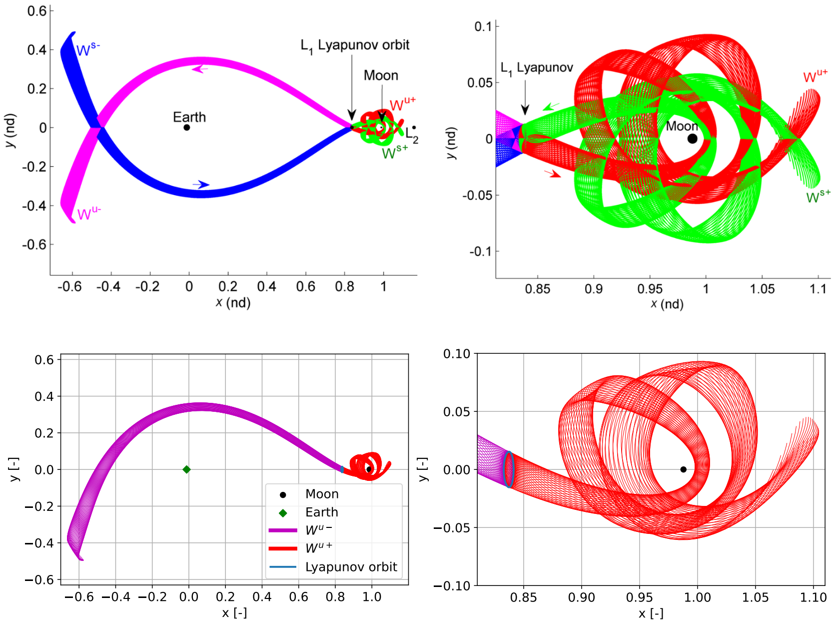

Figure 7.8: Stable (W^s) and unstable (W^u) manifolds associated with a planar Lyapunov orbit near the L_1 point of the Earth-Moon system: as computed by Vaquero (2013) (top), and computed here (bottom, only unstable manifolds).

7.6. Impact Detection

No values that could be used as a reference for verifying the detection of impacts with the polyhedron's surface were found in literature. Instead, since both the computation of the Laplacian of the polyhedron's potential and the model of the CR3BP with polyhedral secondary have been verified, the detection of impacts was verified by comparing the Laplacian of the potential between trajectories propagated using as termination condition the time (i.e. without checking for collisions) and the value of the Laplacian associated with impacts. This analysis was executed for the initial states

$$\begin{aligned} \text{Point 1: } \boldsymbol{x}(t_1) &= [-17.2, 0, 0, 0, 0, 0]^T \quad \text{km, km s}^{-1} \\ \text{Point 2: } \boldsymbol{x}(t_1) &= [17.2, 0, 0, 0, 0, 0]^T \quad \text{km, km s}^{-1} \\ \text{Point 3: } \boldsymbol{x}(t_1) &= [0, 15.0, 0, 0, 0, 0]^T \quad \text{km, km s}^{-1} \\ \text{Point 4: } \boldsymbol{x}(t_1) &= [0, 0, 16.0, 0, 0, 0]^T \quad \text{km, km s}^{-1} \end{aligned} \tag{7.4}$$

All points are located a few kilometers from Phobos' surface: points 1 and 2 are close to, respectively, L_1 and L_2, point 3 is close to Phobos' leading edge, and point 4 is located above Phobos. The Laplacian history is represented in Fig. 7.9 for each of these points, with and without detection of impacts; the trajectories were propagated using the parameters listed in Tab. 2.1, for a maximum of half of Phobos' period. For all points, impacts are detected at the first crossing of $-\nabla^2 U/(G\rho) = 2\pi$ (the value associated with the polyhedron facets), and all impact points have very low distances to the surface (in the order of 10^{-3} m or smaller). Thus, the detection of impacts with the polyhedron is considered verified. One final note: for the propagations without impact detection, it is sometimes possible to observe multiple changes of $-\nabla^2 U/(G\rho)$ between 0 and 4π, these correspond to trajectories that cross the surface multiple times.

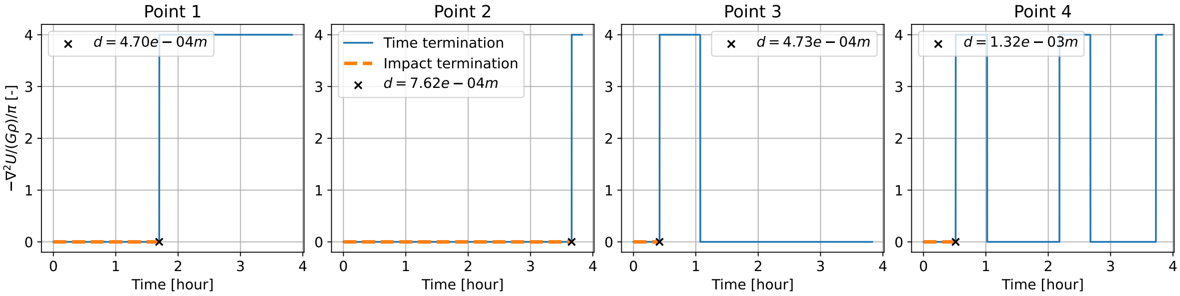

Figure 7.9: Phobos' term $-\nabla^2 U/(G\rho)$ as a function of time, for points propagated with (orange, dashed line) and without (blue, solid line) detection of impacts with Phobos' surface. The value of the distance d to the polyhedron at the detected impact point is marked with a cross.

7.7. General Results

As discussed in Section 4.1, though different dynamics models and continuation setups were used, the computed families of three-dimensional QSOs appear to be consistent with Chen et al. (2020) and Pushparaj et al. (2021).

The magnitude of the observed impact velocities appears to be similar to the ones determined by Oshima and Yanao (2019). However, a direct comparison is difficult, as Oshima and Yanao use the CR3BP, and compute manifolds originating from families of QSOs which were not analyzed here (families of three-dimensional QSOs bifurcating from families of two-dimensional multi-periodic QSOs).

8

Selection of Model Properties

The propagation of a spacecraft's trajectory requires the selection of a series of properties, such as the number of facets of the used polyhedron model, the selection of the numerical integrator, and the formulation of the equations of motion which are integrated. Similarly, the optimization of those trajectories also requires the selection of the optimizer and its parameters. The selection of all of these is here presented.

8.1. Number of Facets of the Polyhedron Model

As argued in Subsection 2.2.3, due to the assumption of constant density, the high-resolution polyhedron model of Phobos (determined by Willner et al. (2010), with 256000 facets) is assumed to produce surface-acceleration errors of at least 1%. Considering that the computation of a gravity field using a high-resolution polyhedron is very computationally intensive, here the resolution of the used model is tuned using as reference the 1% surface acceleration error threshold; the methodology to reduce the number of facets presented in Subsection 2.2.3 is used. The relative acceleration error is computed according to

$$\varepsilon = \frac{\|\ddot{\boldsymbol{r}} - \ddot{\boldsymbol{r}}_T\|}{\|\ddot{\boldsymbol{r}}_T\|} \tag{8.1}$$

where $\ddot{\boldsymbol{r}}$ represents the acceleration produced by the lower-resolution model and $\ddot{\boldsymbol{r}}_T$ the one produced by the high-resolution model (considered to be the true acceleration). The relative acceleration error is computed for a series of field points sampled from the surface of the high-resolution model.

The mean and maximum surface-acceleration errors are represented in Fig. 8.1 as a function of the number of facets. The maximum error is represented for the 99.73% field points with the lowest error, corresponding to a 3σ threshold; this was used instead of the maximum for all field points as the latter often presents an irregular evolution, with regions where it stagnates due to small numbers of outlying field points. It can be observed that both errors decrease for increasing number of facets.

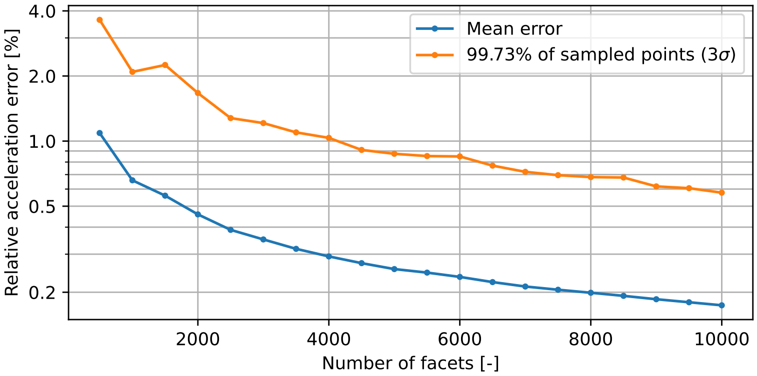

Figure 8.1: Mean and maximum (maximum for 99.73% of the sampled field points, corresponding to 3σ) relative surface-acceleration error between a high-resolution Phobos model with 256000 facets, and lower-resolution models with different numbers of facets.

A model with 1000 facets was selected, which allows a mean acceleration error below the 1% error threshold. One could argue that a maximum acceleration error below 1% would be more appropriate, corresponding to approximately 5000 facets, more in line with the number of facets used by Scheeres et al. (2019). However, 1000 facets was selected as it was considered to offer a better balance between accuracy and computational time. Firstly, when using polyhedron models, the computational time scales approximately linearly with the number of facets (Tardivel, 2016), meaning that a 1000 facets model offers a very significant computational-time improvement with respect to the 5000 facets one. Secondly, since Phobos' gravitational model is not very well known, using a very high accuracy model is not particularly important in this case, as most of the presented analyses would likely have to be rerun with a more accurate model, after a geodesy campaign around Phobos.

8.2. Integrator and Formulation of the Equations of Motion

The Phobos Sample Return mission proposal, developed by ESA (Pickering, 2014), considered a similar mission profile to the one which will be used by MMX, including a static landing on the surface of Phobos. This study required a minimum landing accuracy of 50 m (3σ). Based on this requirement, an orbit integration accuracy in the order of 0.1 m to 1 m is here considered reasonable and is thus used as reference for the selection of the integrator and formulation of the equations of motion.

Since different types of orbits are analyzed, the selection of the integrator and equations of motion is executed using six trajectories considered representative: a three-dimensional low-altitude QSO (belonging to the $1:5-x$ family, with approximate dimensions $28\times52\times24$ km) and one of its manifolds, a three-dimensional medium-altitude QSO (belonging to the $1:14-xz/x$ family, with approximate dimensions $42\times80\times43$ km), and a vertical Lyapunov orbit (around the L_2 point, z-amplitude of approximately 1.5 km) and one of its manifolds[1].

8.2.1. Benchmark

The selection of the integrator and formulation of equations of motion requires the definition of a high-accuracy reference solution (a benchmark). The benchmark is generated using a fixed step-size eight-order Runge-Kutta (RK) method and the Cowell formulation of the equations of motion (i.e. propagation of the Cartesian coordinates).

When integrating differential equations, two types of error are present: rounding errors, which are associated with the finite accuracy of the floating-point representation (and as a result have an erratic behavior), and truncation errors, which are a result of the approximation of the dynamics executed by the integrator (Press et al., 2007). When using a fixed step-size integrator, the truncation error dominates the total error for large step sizes. As the step size decreases, so does the truncation error (the two vary approximately linearly); for sufficiently small step sizes, the accumulation of the rounding error starts dominating, and the total error stops decreasing. Since the benchmark should correspond to the solution with the maximum accuracy, its step size is selected to correspond to the transition between the regions of truncation-dominated and rounding-dominated error.

The integration error associated with a given step size is estimated by comparing the obtained trajectory with the trajectory obtained using half of that step size. Although this estimate for the total error is only valid in the region where the truncation error dominates, it allows an easy identification of the transition between the two error regimes. The evolution of the error with the integration step size is represented in Fig. 8.2. As expected, the error evolves approximately linearly with the step size in the region dominated by the truncation error; nonetheless, in some cases, some oscillations can be observed in this region. In general, a smaller step size leads to a lower total error, however, it is possible that in some parts of the trajectory a slightly larger step actually produces a smaller accumulated error (e.g. due to differences in the sign of the local truncation error). As the propagation time grows, the importance of these effects decreases and the lower local truncation error associated with a lower step size would tend to prevail. However, here relatively short propagation times are used and similar step sizes are tested, thus leading to the observed oscillations. Depending on the trajectory, step sizes between 100.0 s and 400.0 s were selected for the benchmark.

8.2.2. Integrator

The various variable step-size integrators implemented in Tudat are here tested. These include the single-step methods Runge-Kutta-Fehlberg (RKF) and Runge-Kutta-Dormand-Prince (DOPRI) of different orders (RKF4(5),

---

[1] When the selection of the integrator and equations of motion was executed, the study of Lagrange points orbits was being considered, though it was later rejected.

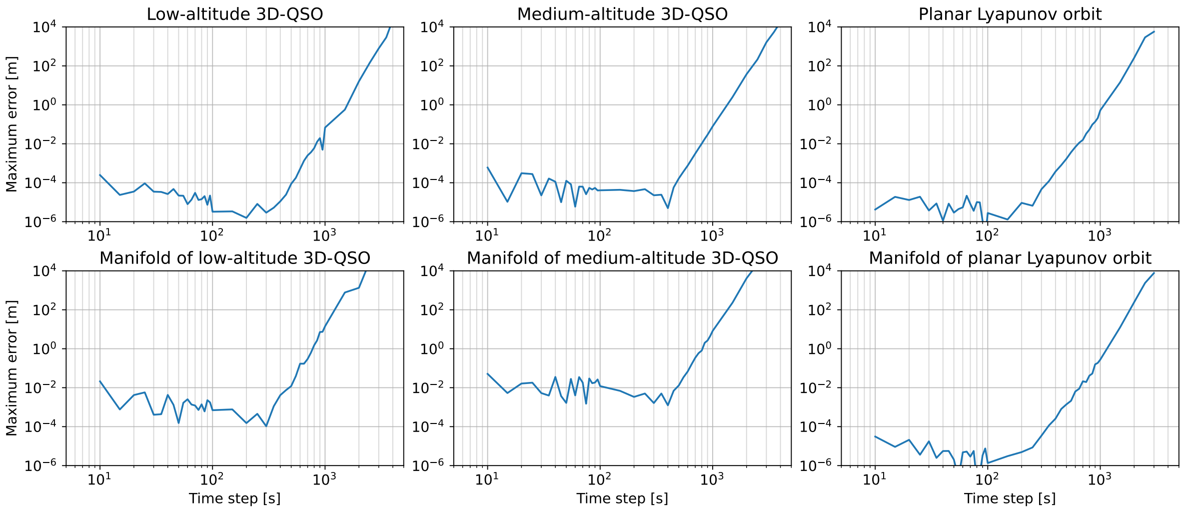

Figure 8.2: Maximum error throughout the trajectory (with respect to a more accurate orbit propagation) as a function of the integrator step size, for an RK8 integrator.

RKF5(6), RKF7(8) and DOPRI8(7)), and the variable-order multi-step method Adams-Bashfort-Moulton. The latter has been tested using different configurations: variable-step variable-order (ABM), variable-step fixed-order (ABMk, with k the order), and fixed-step variable-order (ABM-f).

Other methods implemented in Tudat, as Euler and Bulirsch-Stoer are rejected a priori. The former is based on a first-order Taylor series, thus being known for producing high integration errors; the latter, being an extrapolation method, is characterized by a very sparse output, making it inappropriate for the propagation and analysis of orbits over short periods of time (as is the case here). Fixed step methods (as Runge-Kutta) are also rejected a priori, as a large part of the propagated trajectories go through large changes in the dynamics, from being in-orbit (where the irregular gravity field of the secondary has a low impact) to being on the surface of the secondary. For more details on the mentioned methods, the reader is referred to Montenbruck and Gill (2001) and Press et al. (2007). The evolution of the error with the number of function evaluations is represented in Fig. 8.3, for the different orbits and integrators.

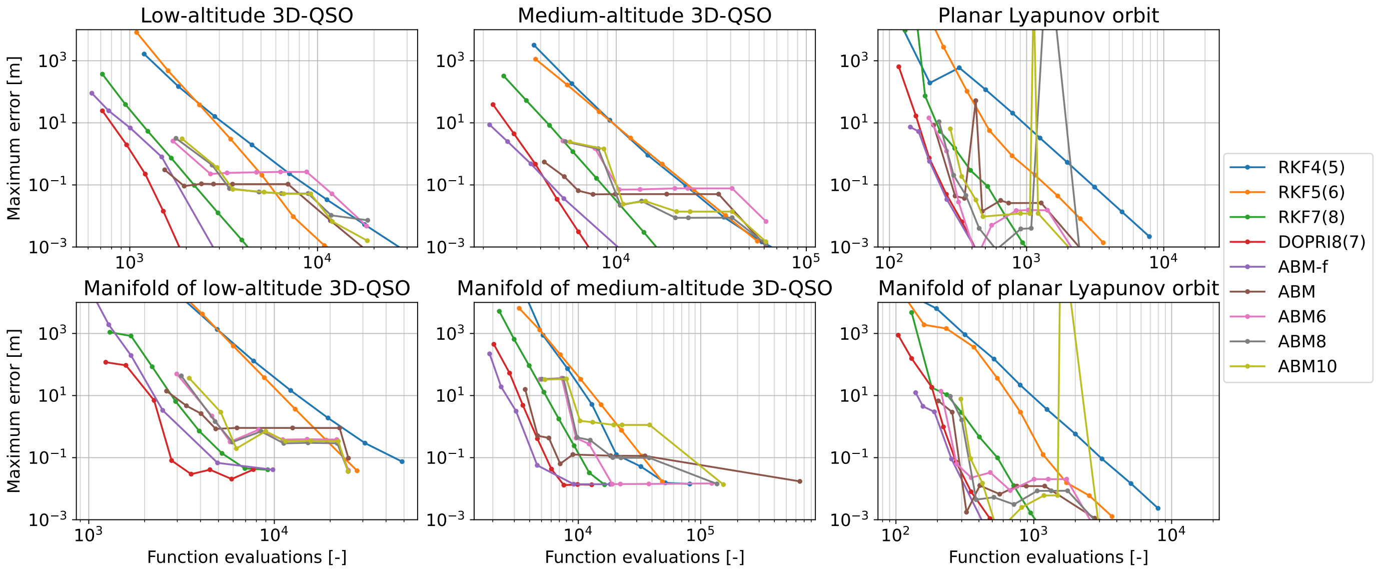

Figure 8.3: Maximum error throughout the trajectory (with respect to the benchmark) as a function of the number of function evaluations, for different integrators.

Variable-step ABM methods (ABMk and ABM) show a clear performance deficit with respect to the fixed-step ABM-f, with the variable-step ones showing regions where the maximum error remains constant with increasing number of function evaluations. This is a result of the poor step-size control algorithm currently implemented in Tudat for ABM methods, which only allows the step size to be halved or doubled in each integration step[2]. Since each ABM integration step uses the information from previous states, ABM methods have to be initialized with another method, DOPRI8(7) in the case of Tudat. In general, DOPRI8(7) uses much larger step sizes than ABM, thus a very sharp decrease in step size is required when transitioning between the two methods; however, since Tudat only allows halving the step size, the first ABM step uses a larger step size than would be required to meet the integration tolerance. Thus, the integration error with these methods is often dominated by the truncation error resulting from that first step. The areas with approximately constant error correspond to the integration tolerances for which the first ABM step uses the same step size.

It is also possible to observe that in the case of the planar Lyapunov orbit and its manifold, for some tolerances, the error with variable-step ABM methods has a very sharp peak. This is also related to the step-size control algorithm currently implemented in Tudat, which is documented to have some tendency to reduce the step size excessively (Tudat, n.d., version 0.3.1). To avoid that reduction, it is necessary to specify a realistic minimum step size (1 s was used); the peaks in the error correspond to situations that required a very large reduction in the step size, which was not allowed by the specified minimum step size, leading to a very large integration error in the relevant integration step.

Finally, in some cases, e.g. the manifold of the medium-altitude three-dimensional QSO, it is possible to observe that above a certain number of function evaluations, the integrator error stops improving. This corresponds to the integrators and step sizes for which the produced error is of the same order of magnitude or smaller than the benchmark error.

The fixed-step variable-order ABM-f method, not being dependent on the mentioned step-size control algorithm, has a much smoother behavior, similar to the variable-step RK methods. As would be expected, the performance of the RK methods improves with the order of the method, with the highest-order method, DOPRI8(7), offering the best performance. Overall, ABM-f and DOPRI8(7) are the best-performing methods; in some instances, the two perform very similarly, in others the latter offers slightly better performance. As such, and considering the superior flexibility of the step-size control of the DOPRI8(7) integrator with respect to the order control of the ABM-f method, DOPRI8(7) was selected.

8.2.3. Formulation of the Equations of Motion

Using the selected integrator, different formulations of the equations of motion can be used when propagating the spacecraft's trajectory. The formulations implemented in Tudat are Cowell, Encke, Gauss' planetary equations, modified equinoctial elements, and unified state model (with quaternions, modified Rodrigues parameters, and exponential mapping). Details about these formulations are given by Betts (1994), Vittaldev et al. (2012), and Wakker (2015). The usage of Gauss' planetary equations is rejected a priori as this formulation has several singularities relevant to the considered trajectories, namely zero eccentricity and zero inclination (with respect to the central body, i.e. Mars); the unified state model formulations are also rejected, as the propagation of the variational equations with these is not implemented in Tudat.

The evolution of the error with the number of function evaluations is represented in Fig. 8.4, for the different orbits and formulations of equations motion tested. In general, Encke shows the best performance: it tends to be either similar or slightly better (for shorter propagations) than modified equinoctial elements and is always better than Cowell. With Encke, the deviation with respect to a reference Kepler orbits is integrated numerically, thus Encke's good performance likely has two origins: firstly, for orbits far from the surface (i.e. medium- and low-altitude QSOs), the Phobos-induced perturbation is periodic and significantly smaller than Mars' gravity, leading the spacecraft to remain close to the reference orbit, oscillating around it; secondly, for orbits near the surface (i.e. impact manifolds), Phobos exerts a strong perturbation, but for a very short period of time, meaning that the spacecraft still remains relatively close to the reference orbit. Thus, the equations of motion are propagated using the Encke formulation. Relative and absolute integration tolerances of 10^{-10} are selected, allowing the specified accuracy requirements to be met for all tested trajectories.

8.3. Optimizer Tuning

As discussed in Subsection 5.2.3, the SADE optimizer was selected. The algorithm was tuned using a simplification of the approach employed by Musegaas (2012). The tuning of the algorithm requires the selection of the population size, the mutation variant, and the algorithm version. There are 18 available mutation variants, listed in Tab. 8.1; for a description of the meaning of each variant, the reader is referred to Storn and Price

---

[2] There are plans to address the issues with the ABM step-size algorithm implemented in Tudat in the near future.

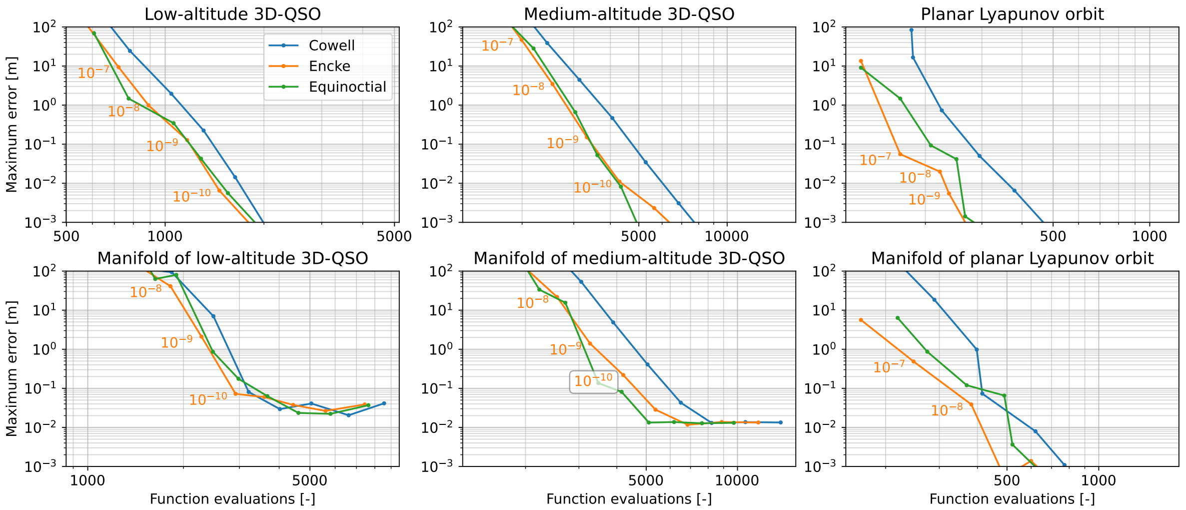

Figure 8.4: Maximum error throughout the trajectory (with respect to the benchmark) as a function of the number of function evaluations, for different formulations of the equations of motion. Some of the tolerances used with the Encke formulation are marked next to the corresponding point.

Table 8.1: Tested SADE mutation variants and corresponding identification numbers.

No.	Variant	No.	Variant	No.	Variant
1	*best/1/exp*	7	*rand/1/bin*	13	*best/3/exp*
2	*rand/1/exp*	8	*rand-to-best/1/bin*	14	*best/3/bin*
3	*rand-to-best/1/exp*	9	*best/2/bin*	15	*rand-to-current/2/exp*
4	*best/2/exp*	10	*rand/2/bin*	16	*rand-to-current/2/bin*
5	*rand/2/exp*	11	*rand/3/exp*	17	*rand-to-best-and-current/2/exp*
6	*best/1/bin*	12	*rand/3/bin*	18	*rand-to-best-and-current/2/bin*

(1997). There are two available SADE versions: iDE (introduced by Elsayed et al., 2011) and jDE (introduced by Brest et al., 2006).

As discussed by Musegaas (2012), the selection of the population size is strongly dependent on the selection of the remaining parameters. As such, the tuning process is here divided into two phases: in the first phase, the mutation variant and algorithm version are selected; in the second phase, for a fixed mutation variant and algorithm version, the population size is selected. During the first phase, optimizations with each possible pair of mutation variant and algorithm version (total of 36 pairs) are executed. To reduce the effect of the population size on the selection of the two, each pair of mutation variant and algorithm version is run for three population sizes (10, 20, and 40); for each pair, only the best result across the different population sizes is considered. Finally, since SADE is a stochastic optimizer, to reduce the effect of the used seed in the optimizer tuning, each optimization run is repeated with three different seeds. To obtain a robust tuning valid for most manifolds, the previous analysis is executed for two manifolds, one with high impact angle (63.00°) and one with low impact angle (8.01°). All optimizations were run for 1520 function evaluations, i.e. 1520 orbit propagations executed during the optimization.

To select the best optimizer settings, two performance measures are considered. Firstly, the number of function evaluations required for the optimizer to converge to the true optimum. The number of function evaluations is used instead of the computational time as the latter is dependent on the computer being used, and on the number of processes running simultaneously; thus, the number of function evaluations allows a more robust comparison, even if the computational time varies between function evaluations. Since the true optimum is not known a priori, it is considered to be the best value found from all optimization runs. A given run is considered to have converged to the true optimum when its fitness function is within a certain tolerance of the true optimum, here selected to be 0.1 $\mathrm{m\,s^{-1}}$. The best result obtained in each run is not taken into account when selecting the best algorithm settings because it is considered more important that a set of settings consistently converges to the vicinity of the global optimum, rather than whether it is able (possibly just for some seed values) to produce the very best result(s) found. The number of function evaluations to reach the optimum is taken to be the mean value across the different tested seeds. The second performance measure used is the percentage of runs in which the optimizer converged to the true optimum: for each combination of mutation variant, algorithm version, and population size, only the cases in which the algorithm converged to the true optimum with all three tested seeds are considered valid.

The results from this tuning process are represented in Fig. 8.5 (high impact-angle manifold) and Fig. 8.6 (low impact-angle manifold). Each circle represents the mean (across the different seeds) number of function evaluations to reach the optimum for a single combination of mutation variant, algorithm version, and population size. For each combination of mutation variant and algorithm version, there are three circles (some overlap), each representing one value of the population size. The circles are colored in green, orange, or red depending on the number of runs that converged to the true optimum; the circles are black if any of the optimization runs failed to meet the optimization constraints.

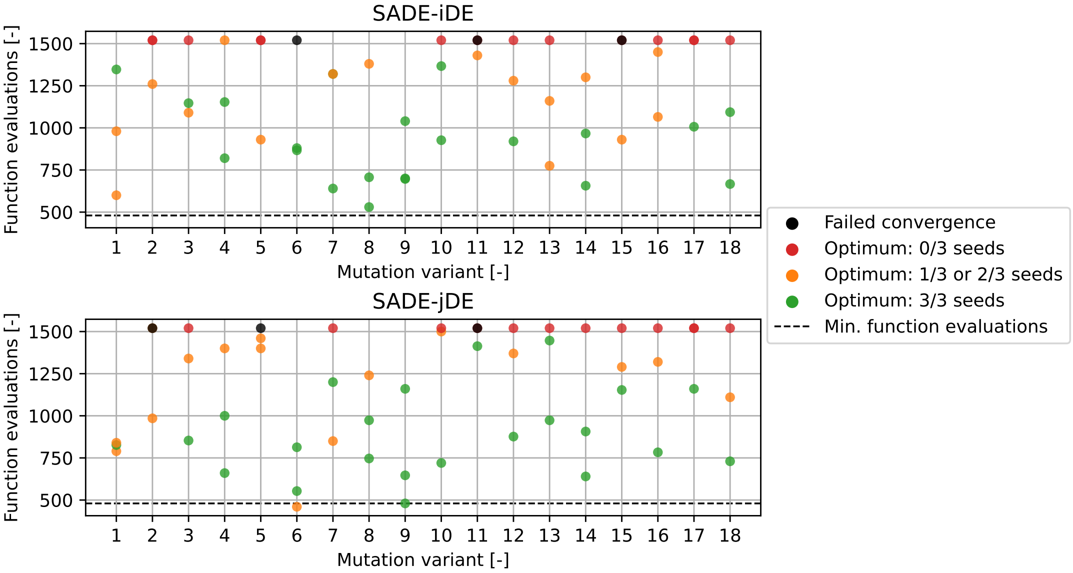

Figure 8.5: Mean number of function evaluations necessary for the optimizer to converge to the global optimum (within tolerance) when analyzing a manifold with high impact-angle (62.995°). Results for different SADE versions and mutation variants. Each circle corresponds to the mean number of function evaluations for a single population size, with varying seed. The minimum number of function evaluations is represented by a black dashed line.

The combination of mutation variant and algorithm version that offers the best general performance for the two tested manifolds is SADE-jDE with *best/2/bin* (number 9). In the case of the high impact-angle manifold, this is the combination that produced the lowest number of function evaluations. In the case of the low impact-angle manifold, it allowed obtaining a number of function evaluations very close to the minimum. Additionally, for both manifolds, SADE-jDE with *best/2/bin* was able to converge to the true optimum in all runs (even for different population sizes), indicating its robustness. It is worth mentioning that in the original presentation of the DE algorithm, Storn and Price (1997) highlight the *best/2/bin* mutation variant as tending to perform very well in a variety of cases; that indeed appears to be the case here. Thus, SADE-jDE with *best/2/bin* is the selected version of the algorithm.

In the second phase of the tuning process, and having selected SADE-jDE with *best/2/bin*, the population was tuned; additionally, the maximum number of generations was also selected. During this phase, each manifold was optimized using ten different population values: 7 (minimum value allowed in SADE), 10, 15, 20, 25, 30, 40, 50, 60, and 80. The optimization for each population size was repeated using five different seeds; the number of seeds was increased in this second phase to allow a slightly more robust selection of the maximum number of generations. The obtained results are represented in Fig. 8.7; again, the color of the circles represents the number of runs per seed that converged to the true optimum (within the specified tolerance). It is possible to note that the optimizer performance tends to improve with decreasing population size, with the best performance happening for a population of 7. This behavior is likely related to the small number of variables being optimized, since the optimal population size tends to increase as the number of decision variables increases (Musegaas, 2012).

For a population size of 7, the high and low impact-angle manifolds required on average, respectively,

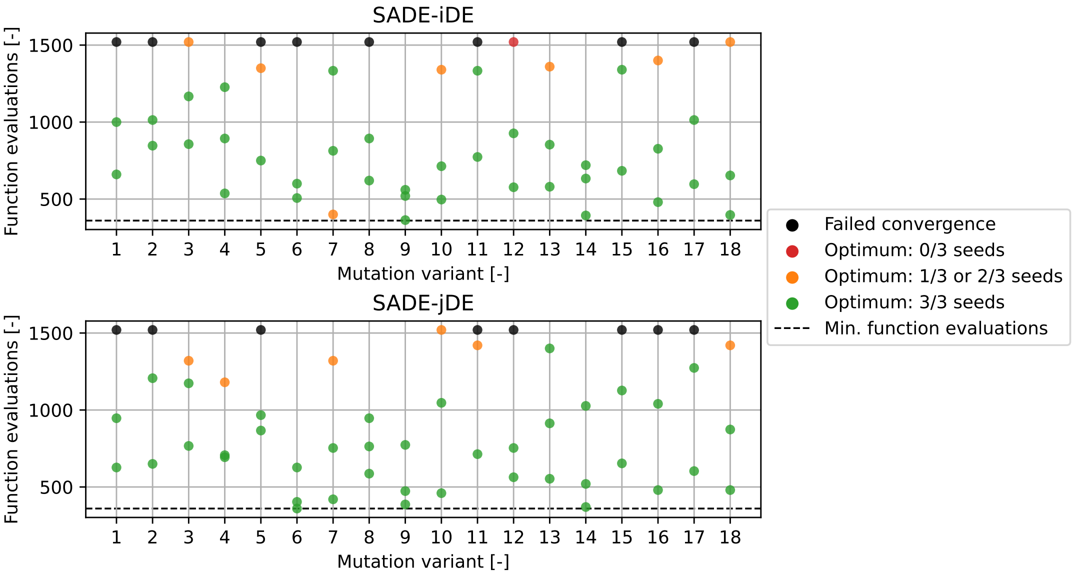

Figure 8.6: Mean number of function evaluations necessary for the optimizer to converge to the global optimum (within tolerance) when analyzing a manifold with low impact-angle (8.011°). Results for different SADE versions and mutation variants. Each circle corresponds to the mean number of function evaluations for a single population size, with varying seed. The minimum number of function evaluations is represented by a black dashed line.

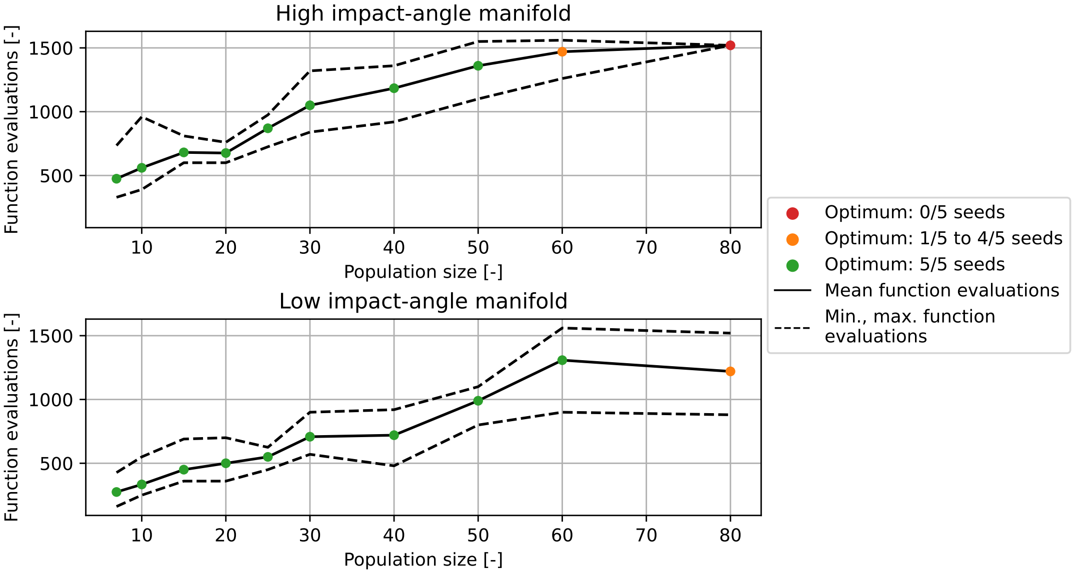

Figure 8.7: Mean, minimum and maximum number of function evaluations necessary for the optimizer to converge to the global optimum as a function of the population size, for a manifold with high impact-angle (top) and low impact-angle (bottom).

475 and 275 function evaluations for the optimizer to converge to the optimum. For this population size, the maximum observed number of function evaluations was approximately 700, for one of the high impact-angle runs. As such, and to be conservative, a maximum number of function evaluations of 1000 is selected.

9

Conclusions and Recommendations

In this chapter, the conclusions of this work are presented as answers to the research questions and recommendations for future research work are proposed.

9.1. Conclusions

To analyze the outcomes of this work, the research questions that initiated it are repeated and answered. The subquestions are answered first, and the main research question is analyzed at the end.

> **How does the behavior of the invariant manifolds relate to the families of QSOs that create them?**

Invariant manifolds were computed starting from orbits belonging to families of three-dimensional QSOs, with two main groups of families: families of approximately cylindrical orbits, going up to high A_z values, and families of irregular orbits, which start intersecting Phobos' surface at relatively low A_z values (under 50 km). From the cylindrical families, only the ones very close to the surface (1 : 5 and 1 : 6 resonances) produced large numbers of impacts; the ones further away from it ($1 : n$, with $n \geq 11$) proved unsuitable for the generation of impact manifolds, with most escaping the system, likely due to the weak gravitational influence of Phobos. The families with more irregular orbits, passing closer to the surface, produced much larger numbers of impact manifolds and better surface coverage. The number of impacts per revolution tends to increase as the distance to the surface decreases, while the surface coverage improves as the number of revolutions in the orbits increases (i.e. as the distance to the surface increases); the first of these two effects is dominant, with the 1 : 6, 2 : 11, and 2 : 13 families offering the best properties in terms of surface access.

The families with high numbers of impacts allow accessing the entire surface (except the Stickney crater), though with higher concentrations of impacts in the Mars and anti-Mars edge, the two regions where MMX will land. Thus, the invariant manifolds of QSOs would allow a more flexible selection of the landing location than other alternatives. Zamaro (2015), analyzing invariant manifolds of Lagrange point orbits, only finds impacts in the vicinity of the Mars and anti-Mars edges. Çelik et al. (2019), considering trajectories departing from two-dimensional QSOs via an impulsive maneuver, finds the impact locations to be concentrated in the leading and trailing edges, with very limited landing opportunities on the other two edges.

The distribution of impact angles and impact locations along the surface were observed to be strongly dependent on Phobos' surface geometry and on the shape of the invariant manifolds, which, in turn, generally follow the shape of the orbits from which they emanate. The distribution of the impact velocities also results from the influence of Phobos and the initial orbit (i.e. the value of the Jacobi integral on the surface and on the orbit).

> **In the Mars-Phobos system, is it possible to land safely (i.e. complying with the landing requirements) using unpowered invariant manifolds?**

The minimum computed impact velocity was approximately 13 $\mathrm{m\,s^{-1}}$, with the low-A_z families ($A_z \leq$ 50 km) having impact velocities up to 17 $\mathrm{m\,s^{-1}}$. Therefore, assuming maximum impact velocities in the order of 1 $\mathrm{m\,s^{-1}}$, similar to the requirements of the Phobos Sample Return mission and of MMX's rover, it would not be possible to land safely. However, if the goal was to land a spacecraft or payload with higher allowed impact velocities, for example, the 15 $\mathrm{m\,s^{-1}}$ considered by Çelik et al. (2019), then it would be possible to use the manifolds without the need for any braking maneuver.

What is the minimum ΔV required to meet the landing requirements?

To determine the ΔV necessary to meet the specified landing requirements, transfer trajectories were designed with two impulsive maneuvers: a maneuver to insert the spacecraft into the invariant manifold and a braking maneuver executed at low altitude. The ΔV budget is dominated by this second maneuver, the first one being much smaller. The second maneuver was designed both semi-analytically and via optimization, the two resulting in similar ΔV values, though only the latter allows meeting all the impact constraints. Depending on the selected departure orbit, minimum ΔV values between 12.5 $\mathrm{m\,s^{-1}}$ and 26.6 $\mathrm{m\,s^{-1}}$ were found, although all orbits of the low-A_z families have minimum ΔVs below 16.3 $\mathrm{m\,s^{-1}}$.

Starting from a two-dimensional QSO, the proposal for ESA's Phobos Sample Return mission considers a direct transfer to the surface, applying a single ΔV of $5-10$ $\mathrm{m\,s^{-1}}$ to insert the spacecraft into an impact trajectory and executing a 4 $\mathrm{m\,s^{-1}}$ corrective maneuver to increase the accuracy of the impact location (Pickering, 2014); however, the size of the initial QSO and the impact velocities are unclear. Starting from a two-dimensional QSO with $A_x = 22$ km, Çelik et al. (2019) finds a minimum impact velocity of 5.5 $\mathrm{m\,s^{-1}}$, after an initial maneuver of 7.2 $\mathrm{m\,s^{-1}}$; thus, a total ΔV of approximately 12 $\mathrm{m\,s^{-1}}$ would be expected for similar impact-velocity requirements to the ones considered here. Overall, the ΔV values available in the literature for landings departing from QSOs appear to be similar to the ones determined here, however, a direct comparison is difficult due to the differences in terms of the departure and impact conditions.

Analyzing landings via invariant manifolds of Lagrange point orbits, Zamaro (2015) finds impact velocities of $3-8$ $\mathrm{m\,s^{-1}}$, which would correspond to ΔVs smaller than the ones found here. However, to execute such a landing, one would need to first transfer to the Lagrange point orbit (likely from an initial QSO) and probably require additional propellant for station keeping before entering the manifold. As such, the overall propellant expenditure would likely be larger than when departing from a QSO.

Do invariant manifolds offer landing trajectories robust to uncertainties?

The analysis of the sensitivity of the invariant manifolds to uncertainties was executed in two phases: one corresponding to the transfer from the orbit to the surface (for unpowered landings) or to the maneuver point (for powered landings), and one corresponding to the transfer between the maneuver point and the surface (only for powered landings).

The first phase showed that the stability of a manifold depends on its time of flight and on the stability of the originating orbit. It was shown that indeed there are manifolds (usually with short times of flight) that allow very robust landings, with all families having a number of orbits with manifolds that almost always (i.e. in at least 99.73% of the cases) allow reaching the surface without the application of any corrective maneuver, even in the presence of uncertainties. These robust manifolds allow landing on the Mars and anti-Mars edges of Phobos, the two regions where MMX is currently planned to land. Due to the lack of literature on the topic, it is not clear how the robustness of the manifolds compares to other landing strategies.

The second phase of the uncertainty analysis showed that for the selected braking-maneuver set-up, in general, it would not be possible to meet the impact-velocity constraints. Nevertheless, this does not point to a lack of robustness of the manifolds, simply indicating that the braking maneuver itself is not robust and that the methodology used in this final part of the trajectory should be further studied.

Can invariant manifolds of three-dimensional QSOs be used to design landing trajectories?

This work has analyzed the usage of invariant manifolds of QSOs as landing trajectories, with a focus on three-dimensional QSOs due to their advantages as science orbits. The Mars-Phobos system was used as a case study. The usage of invariant manifolds proved to allow robust landing trajectories able to access the Mars and anti-Mars edges of Phobos (the landing locations of MMX). These either require similar ΔVs to other strategies, if only low impact velocities are allowed, or a negligible ΔV (for the manifold insertion), if the impact-velocity constraints are less strict (in which case no braking maneuver would be required). Thus, the usage of invariant manifolds is particularly promising in the second case, which could be relevant for the deployment of small scientific packages or CubeSats to the surface. An additional advantage of using invariant manifolds of three-dimensional QSOs relates to the fact that they would allow entering a landing trajectory directly from an orbit that is optimal from a scientific point of view. This is opposed to the strategy that will be employed by MMX, in which a transfer to a two-dimensional QSO will be executed before the landings, and to the strategy analyzed by Zamaro (2015), according to which the landing trajectories would depart from very unstable Lagrange point orbits.

In the Mars-Phobos system, landings were observed to only be possible when departing from orbits in low-A_z families or orbits in high-A_z families very close to the surface. In the case of MMX, which will use a three-dimensional QSO member of a high-A_z family with $A_x = 50$ km (Nakamura et al., 2021), landings via invariant manifolds would not be possible. Nevertheless, since MMX will use two-dimensional QSOs as low as $A_x = 20$ km, in case a three-dimensional QSO with similar dimensions were considered, the usage of invariant manifolds would probably be feasible (depending on the selected A_z value). In this case, using invariant manifolds to transfer the spacecraft from the QSO to an altitude of 2 km and executing a low-velocity vertical descent from there (i.e. replacing the final braking maneuver designed here with the final portion of the transfer currently being considered for MMX) could be a robust and fuel-efficient option.

Finally, it is worth highlighting that Phobos' gravity is very weak, likely making the application of the manifolds as landing trajectories particularly difficult. It is possible that systems where the secondary has a stronger gravitational influence would allow larger numbers of robust impact manifolds and larger numbers of families that are able to generate them (e.g. members of high-A_z families more distant from the surface).

9.2. Recommendations

Several recommendations are presented for future research work, either highlighting aspects of this work that deserve further attention or proposing possible research topics motivated by the applied techniques.

As mentioned, although the Mars-Phobos system was analyzed as a case study here, most results should be valid for other systems. Hence, a first recommendation would be the study of landing trajectories using invariant manifolds of QSOs in other three-body systems, for example, Jupiter-Europa or Saturn-Enceladus (both systems where the usage of QSOs has been previously studied).

Although the used dynamics models are expected to provide trustworthy results, these are nonetheless simplified. Therefore, the analysis could be improved by using higher-accuracy models. The first major perturbation to include would be Phobos' eccentricity, which could be modeled using the elliptic restricted three-body problem with an irregular secondary. In this case, due to the periodic perturbation, all orbits would become quasi-periodic; as such, an analysis based on the computation of quasi-periodic tori (e.g. Baresi et al., 2021) and their manifolds could be beneficial. Finally, other smaller perturbations could be included, namely Mars' non-spherical gravity, the solar and Martian radiation pressures, and fourth-body perturbations. It is not expected that these perturbations would strongly affect the obtained results, at least in qualitative terms. Nevertheless, if this concept were to be applied in a real mission, it would be important to analyze it considering the most-accurate dynamics model available.

In the same way that unstable invariant manifolds of unstable QSOs can be used to design landings, their stable manifolds can be used to design take-offs (with the difference that the manifolds would be computed by backward integrating the equations of motion). This could be particularly useful for MMX, assuming it goes back to a QSO between the execution of its two landings.

When designing the landings, here the possibility of bouncing on the surface was ignored. While this is a reasonable assumption considering the used impact-velocity constraints (maximum velocity norm of approximately 1.01 $\mathrm{m\,s^{-1}}$), it would be more questionable if higher impact velocities were allowed. For example, when studying the deployment of small scientific payloads on Phobos, Çelik et al. (2019) consider maximum impact velocities of 15 $\mathrm{m\,s^{-1}}$. In those situations, the design of landing trajectories coupled with the analysis of the bouncing motion would be essential.

Here, the landing trajectories were designed by assuming that the spacecraft would follow an invariant manifold until very close to the surface, and apply a single braking maneuver at the end. Evidently, this is not a very robust option, being sensible to thruster failures or errors. Thus, other possibilities for the design of the final portion of the trajectory should be considered, for example, applying the braking maneuver at higher altitudes and using multiple small maneuvers to correct the errors resulting from the initial maneuver, deal with discrepancies between the used dynamics model and reality, and further control the impact velocity. Additionally, for more accurate results, these maneuvers should be modeled as non-impulsive.

The design of landing trajectories using invariant manifolds is only possible for unstable orbits, as otherwise no manifolds exist, and requires sufficiently unstable orbits, so that the time required to depart the vicinity of the orbit is not too large. Even though the considered orbits are weakly unstable, in general, it would be preferable to use stable orbits or orbits with as low instabilities as possible (to reduce the requirements in terms of station-keeping), which would render the usage of manifolds either impossible or excessively slow. As an alternative, it would be possible to study the design of landing trajectories by applying a maneuver in the direction of largest expansion, as determined using the singular-value decomposition of the monodromy matrix (Subsection 6.1.1) of the initial periodic orbit. This type of approach has been considered by Muralidharan (2021) for designing station-keeping maneuvers and general transfer trajectories.

Studying single-periodic orbits, Doedel et al. (2007) mapped all existing families of Lagrange point orbits and their connections. Considering the very high interest in the application of multi-periodic orbits for mission design (e.g. this work, Chen et al. (2020), and Pushparaj et al. (2021), who used QSOs in the Mars-Phobos system, Zimovan-Spreen (2021), who studied near-rectilinear halo orbits and QSOs in the Earth-Moon system, and Broschart et al. (2009), who studied terminator orbits near asteroids), one assumes that a global mapping of the available multi-periodic families could be very useful as a design tool. Some first steps in that direction were executed by Bury et al. (2022), but such a process is made difficult by the sheer number of families, and by the fact that one does not know a priori how families intersect each other, where each family ends, or which families are actually connected; as such, a continuation algorithm might inadvertently compute the same family multiple times. All of these make the process of analyzing the data quite laborious. To automatically generate a global mapping of the available families, one could apply some graph-search algorithm (e.g. depth- or width-first search), with each bifurcation point corresponding to a node of the graph. This would require being able to uniquely identify each node of the graph, which might not be a completely trivial task. Having determined such a global database of families, one could then use the orbits computed in the process as interpolation nodes, allowing a mission designer to straightforwardly generate periodic orbits in a given region or family of interest.

A

Transformation Between Frames

The equations of motion are propagated, using Tudat, with respect to a primary-centered (i.e. Mars-centered) inertial frame. Nevertheless, all the analyses in this report are done with respect to a secondary-centered (i.e. Phobos-centered) synodic frame; additionally, a barycentric synodic frame is used for verification purposes. As such, it becomes essential to define the transformations between the different frames.

The generic transformation between an initial frame I and a final frame F is here presented; the transformations are then simplified according to the CR3BP assumptions, assuming I to correspond to the primary-centered inertial frame and F to the barycenter/secondary-centered synodic frame. Consider the following notation: an upper-case vector $\boldsymbol{V}$ represents a vector defined with respect to frame I, a lower-case vector $\boldsymbol{v}$ represents a vector defined with respect to frame F; a superscript ^{I} indicates a vector written using the coordinates of frame I, a superscript ^{F} indicates a vector written using the coordinates of frame F. For example, $\boldsymbol{R}^F$ represents a position vector defined with respect to frame I, using the coordinates of frame F.

The coordinates of frame F are obtained from the coordinates of frame I using the rotation matrix $C^{F/I}$, according to

$$\begin{bmatrix} \hat{\boldsymbol{x}}^F \\ \hat{\boldsymbol{y}}^F \\ \hat{\boldsymbol{z}}^F \end{bmatrix} = C^{F/I} \begin{bmatrix} \hat{\boldsymbol{x}}^I \\ \hat{\boldsymbol{y}}^I \\ \hat{\boldsymbol{z}}^I \end{bmatrix} \tag{A.1}$$

Thus, for any vector $\boldsymbol{v}^I$ or $\boldsymbol{V}^I$, the following applies: $\boldsymbol{v}^F = C^{F/I}\boldsymbol{v}^I$ and $\boldsymbol{V}^F = C^{F/I}\boldsymbol{V}^I$. The inverse transformation is obtained via

$$\begin{bmatrix} \hat{\boldsymbol{x}}^I \\ \hat{\boldsymbol{y}}^I \\ \hat{\boldsymbol{z}}^I \end{bmatrix} = \left(C^{F/I}\right)^{-1} \begin{bmatrix} \hat{\boldsymbol{x}}^F \\ \hat{\boldsymbol{y}}^F \\ \hat{\boldsymbol{z}}^F \end{bmatrix} \tag{A.2}$$

Since $C^{F/I}$ defines the rotation between two sets of orthogonal unit vectors, $C^{F/I}$ is an orthogonal matrix, therefore (Vallado, 2013)

$$\left(C^{F/I}\right)^{-1} = \left(C^{F/I}\right)^T \tag{A.3}$$

In the case of the CR3BP, the rotation matrix is given by

$$C^{F/I}(t) = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{A.4}$$

$$\theta(t) = \dot{\theta} \cdot (t - t_i) - \theta(t_i) \tag{A.5}$$

where t_i represents the initial time and t the time at which the rotation matrix is calculated. In the following equations, the dependence of the rotation matrix on the time is omitted.

A.1. Position and Velocity

A.1.1. Position

Considering $\boldsymbol{R}_0^I$ to represent the position of the origin of frame F with respect to frame I, written using the coordinates of frame I, one obtains the following equation (Vallado, 2013)

$$\boldsymbol{R}^I = \boldsymbol{R}_0^I + \boldsymbol{r}^I \Leftrightarrow \tag{A.6a}$$

$$\boldsymbol{r}^F = C^{F/I}\left(\boldsymbol{R}^I - \boldsymbol{R}_0^I\right) \tag{A.6b}$$

For the CR3BP, this corresponds to

$$\boldsymbol{r}^F = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \left(\boldsymbol{R}^I - \boldsymbol{R}_0^I\right) \tag{A.7}$$

A.1.2. Velocity

The time derivatives with respect to frame I and with respect to frame F are denoted, respectively, by

$$\frac{d}{dt} \quad \text{and} \quad \frac{\delta}{\delta t} \tag{A.8}$$

In the following equations, the time derivative is alternatively represented using the dot notation, which is considered to indicate the time derivative with respect to the reference frame in which the vector is defined; in the case of the position and velocity vectors:

$$\dot{\boldsymbol{R}}^I = \frac{d\boldsymbol{R}^I}{dt} \qquad \dot{\boldsymbol{R}}^F = \frac{d\boldsymbol{R}^F}{dt} \qquad \dot{\boldsymbol{r}}^I = \frac{\delta \boldsymbol{r}^I}{\delta t} \qquad \dot{\boldsymbol{r}}^F = \frac{\delta \boldsymbol{r}^F}{\delta t} \tag{A.9}$$

Considering an arbitrary vector $\boldsymbol{a}$, the derivatives with respect to the two frames can be related through (Wakker, 2015)

$$\frac{d\boldsymbol{a}}{dt} = \frac{\delta \boldsymbol{a}}{\delta t} + \boldsymbol{\omega}^I \times \boldsymbol{a} \tag{A.10}$$

where $\boldsymbol{\omega}$ represents the angular velocity of frame F with respect to frame I, written in the coordinates of frame I.

Taking the time derivative d/dt of Eq. A.6b, and applying Eq. A.10 with $\boldsymbol{a} = \boldsymbol{r}^I$, one obtains (Vallado, 2013)

$$\begin{gathered} \frac{d\boldsymbol{R}^I}{dt} = \frac{d\boldsymbol{R}_0^I}{dt} + \frac{\delta \boldsymbol{r}^I}{\delta t} + \boldsymbol{\omega}^I \times \boldsymbol{r}^I \Leftrightarrow \\ \dot{\boldsymbol{R}}^I = \dot{\boldsymbol{R}}_0^I + \dot{\boldsymbol{r}}^I + \boldsymbol{\omega}^I \times \boldsymbol{r}^I \end{gathered} \tag{A.11}$$

Finally, the velocity with respect to frame F, written with the coordinates of frame F, is obtained by using the coordinate rotation in Eq. A.1

$$\dot{\boldsymbol{r}}^F = C^{F/I}\left(\dot{\boldsymbol{R}}^I - \dot{\boldsymbol{R}}_0^I\right) - C^{F/I}\left[\boldsymbol{\omega}^I \times \left(\boldsymbol{R}^I - \boldsymbol{R}_0^I\right)\right] \tag{A.12}$$

In the case of the CR3BP, $\boldsymbol{\omega}^I = [0, 0, \dot{\theta}]^T$, therefore the previous equation simplifies to

$$\begin{aligned} \dot{\boldsymbol{r}}^F &= \dot{\theta} \begin{bmatrix} -\sin\theta & \cos\theta & 0 \\ -\cos\theta & -\sin\theta & 0 \\ 0 & 0 & 0 \end{bmatrix} \left(\boldsymbol{R}^I - \boldsymbol{R}_0^I\right) + \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \left(\dot{\boldsymbol{R}}^I - \dot{\boldsymbol{R}}_0^I\right) \\ &= \dot{C}^{F/I}\left(\boldsymbol{R}^I - \boldsymbol{R}_0^I\right) + C^{F/I}\left(\dot{\boldsymbol{R}}^I - \dot{\boldsymbol{R}}_0^I\right) \end{aligned} \tag{A.13}$$

A.1.3. Full State

According to Eq. A.7 and Eq. A.13, for the CR3BP, the transformation of the full state between frames I and F can be represented by

$$\begin{bmatrix} \boldsymbol{r}^F \\ \dot{\boldsymbol{r}}^F \end{bmatrix} = C_{\text{full}}^{F/I} \begin{bmatrix} \boldsymbol{R}^I - \boldsymbol{R}_0^I \\ \dot{\boldsymbol{R}}^I - \dot{\boldsymbol{R}}_0^I \end{bmatrix} \qquad \text{with} \qquad C_{\text{full}}^{F/I} = \begin{bmatrix} C^{F/I} & 0_{3\times3} \\ \dot{C}^{F/I} & C^{F/I} \end{bmatrix} \tag{A.14}$$

The inverse transformation is represented by

$$\begin{bmatrix} \boldsymbol{R}^I \\ \dot{\boldsymbol{R}}^I \end{bmatrix} = \left(C_{\text{full}}^{F/I}\right)^{-1} \begin{bmatrix} \boldsymbol{r}^F \\ \dot{\boldsymbol{r}}^F \end{bmatrix} + \begin{bmatrix} \boldsymbol{R}_0^I \\ \dot{\boldsymbol{R}}_0^I \end{bmatrix} \qquad \text{with} \qquad \left(C_{\text{full}}^{F/I}\right)^{-1} = \begin{bmatrix} \left(C^{F/I}\right)^T & 0_{3\times3} \\ \left(\dot{C}^{F/I}\right)^T & \left(C^{F/I}\right)^T \end{bmatrix} \tag{A.15}$$

A.2. Acceleration

Similarly to what was done for the velocity, one applies the time derivative d/dt to Eq. A.11, and uses Eq. A.10 with $\boldsymbol{a} = \delta \boldsymbol{r}^I/\delta t$ and $\boldsymbol{a} = \boldsymbol{r}^I$ to simplify the equation, resulting in (Vallado, 2013)

$$\begin{aligned} \frac{d^2\boldsymbol{R}^I}{dt^2} &= \frac{d^2\boldsymbol{R}_0^I}{dt^2} + \frac{\delta^2\boldsymbol{r}^I}{\delta t^2} + 2\boldsymbol{\omega}^I \times \frac{\delta\boldsymbol{r}^I}{\delta t} + \frac{d\boldsymbol{\omega}^I}{dt} \times \boldsymbol{r}^I + \boldsymbol{\omega}^I \times \left(\boldsymbol{\omega}^I \times \boldsymbol{r}^I\right) \Leftrightarrow \\ \ddot{\boldsymbol{R}}^I &= \ddot{\boldsymbol{R}}_0^I + \ddot{\boldsymbol{r}}^I + 2\boldsymbol{\omega}^I \times \dot{\boldsymbol{r}}^I + \dot{\boldsymbol{\omega}}^I \times \boldsymbol{r}^I + \boldsymbol{\omega}^I \times \left(\boldsymbol{\omega}^I \times \boldsymbol{r}^I\right) \end{aligned} \tag{A.16}$$

The acceleration with respect to frame F, written with coordinates of frame F, is obtained by using the coordinate rotation in Eq. A.1, thus

$$\ddot{\boldsymbol{r}}^F = C^{F/I}\left(\ddot{\boldsymbol{R}}^I - \ddot{\boldsymbol{R}}_0^I\right) - C^{F/I}\left(2\boldsymbol{\omega}^I \times \dot{\boldsymbol{r}}^I\right) - C^{F/I}\left(\dot{\boldsymbol{\omega}}^I \times \boldsymbol{r}^I\right) - C^{F/I}\left[\boldsymbol{\omega}^I \times \left(\boldsymbol{\omega}^I \times \boldsymbol{r}^I\right)\right] \tag{A.17}$$

The second term is known as Coriolis acceleration and the fourth as centripetal acceleration.

Under the CR3BP assumptions (including $\dot{\boldsymbol{\omega}}^I = \boldsymbol{0}$), the previous equation reduces to

$$\begin{aligned} \ddot{\boldsymbol{r}}^F = &\begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \left(\ddot{\boldsymbol{R}}^I - \ddot{\boldsymbol{R}}_0^I\right) + 2\dot{\theta} \begin{bmatrix} -\sin\theta & \cos\theta & 0 \\ -\cos\theta & -\sin\theta & 0 \\ 0 & 0 & 0 \end{bmatrix} \left(\dot{\boldsymbol{R}}^I - \dot{\boldsymbol{R}}_0^I\right) + \\ &\dot{\theta}^2 \begin{bmatrix} -\cos\theta & -\sin\theta & 0 \\ \sin\theta & -\cos\theta & 0 \\ 0 & 0 & 0 \end{bmatrix} \left(\boldsymbol{R}^I - \boldsymbol{R}_0^I\right) \end{aligned} \tag{A.18}$$

A.3. State Transition Matrix

Consider the STM with respect to an arbitrary frame A, written with the coordinates of frame A, to be represented by $\Phi^A(t, t_i)$. Using this notation, the STMs in the initial frame I and final frame F can be related through (Dang, 2017)

$$\Phi^F(t, t_i) = C_{\text{full}}^{F/I}(t)\ \Phi^I(t, t_i)\ \left[C_{\text{full}}^{F/I}(t_i)\right]^{-1} \tag{A.19}$$

Bibliography

Alizadeh, I. and Villac, B. F. (2010). “Targeting Requirements and Stability Characterization for a Class of Ballistic Transfers”. In: *Journal of Guidance, Control, and Dynamics*, 33 (5), pp. 1473–1489.

Allgower, E. L. and Georg, K. (2003). *Introduction to Numerical Continuation Methods*. 1st ed. Society for Industrial & Applied Mathematics.

Andert, T. P., Rosenblatt, P., Pätzold, M., Häusler, B., Dehant, V., Tyler, G. L., and Marty, J. C. (2010). “Precise Mass Determination and the Nature of Phobos”. In: *Geophysical Research Letters*, 37 (9).

Balmino, G. (1994). “Gravitational Potential Harmonics From the Shape of an Homogenous Body”. In: *Celestial Mechanics and Dynamical Astronomy*, 60, pp. 331–364.

Baresi, N. (2017). *Spacecraft Formation Flight on Quasi-Periodic Invariant Tori*. PhD thesis, University of Colorado Boulder, Boulder, Colorado.

Baresi, N., Dei Tos, D. A., Ikeda, H., and Kawakatsu, Y. (2021). “Trajectory Design and Maintenance of the Martian Moons eXploration Mission Around Phobos”. In: *Journal of Guidance, Control, and Dynamics*, 44 (5), pp. 996–1007.

Betts, J. T. (1994). “Optimal Interplanetary Orbit Transfers by Direct Transcription”. In: *The Journal of the Astronautical Sciences*, 42 (3), pp. 247–268.

Biscani, F. and Izzo, D. (2020). “A Parallel Global Multiobjective Framework for Optimization: Pagmo”. In: *Journal of Open Source Software*, 5 (53), p. 2338.

Bosanac, N. (2016). *Leveraging Natural Dynamical Structures to Explore Multi-Body Systems*. PhD thesis, Purdue University, West Lafayette, Indiana.

Boyce, W. (1997). “Comment on a Formula for the Gravitational Harmonic Coefficients of a Triaxial Ellipsoid”. In: *Celestial Mechanics and Dynamical Astronomy*, 67, pp. 107–110.

Brest, J., Greiner, S., Bošković, B., Mernik, M., and Zumer, V. (2006). “Self-Adapting Control Parameters in Differential Evolution: A Comparative Study on Numerical Benchmark Problems”. In: *IEEE Transactions on Evolutionary Computation*, 10 (6), pp. 646–657.

Broschart, S. B., Scheeres, D. J., and Villac, B. F. (2009). “New Families of Multi-Revolution Terminator Orbits Near Small Bodies”. In: *Astrodynamics Specialist Conference*. Pittsburgh, Pennsylvania.

Broucke, R. (1969). “Stability Of Periodic Orbits In The Elliptic, Restricted Three-Body Problem”. In: *AIAA Journal*, 7 (6), pp. 1003–1009.

Bury, L. and McMahon, J. (2020). “Landing Trajectories to Moons from the Unstable Invariant Manifolds of Periodic Libration Point Orbits”. In: *AIAA Scitech Forum*. Orlando, Florida.

Bury, L., McMahon, J., and Lo, M. (2022). “A Study of Periodic Orbits Near Europa”. In: *Celestial Mechanics and Dynamical Astronomy*, 134 (27).

Canalias, E., Lorda, L., Martin, T., Laurent-Varin, J., Charles, J., and Mimasu, Y. (2017). “Trajectory Analysis for the Phobos Proximity Phase of the MMX Mission”. In: *International Symposium on Spaceflight Dynamics*. Matsuyama, Japan.

Capannolo, A., Ferrari, F., and Lavagna, M. (2019). “Families of Bounded Orbits near Binary Asteroid 65803 Didymos”. In: *Journal of Guidance, Control, and Dynamics*, 42 (1), pp. 189–198.

Çelik, O., Baresi, N., Ballouz, R. L., Ogawa, K., Wada, K., and Kawakatsu, Y. (2019). “Ballistic Deployment From Quasi-Satellite Orbits Around Phobos Under Realistic Dynamical and Surface Environment Constraints”. In: *Planetary and Space Science*, 178.

Ceriotti, M. and Sanchez, J. P. (2016). “Control of Asteroid Retrieval Trajectories to Libration Point Orbits”. In: *Acta Astronautica*, 126, pp. 342–353.

Chen, H., Canalias, E., Hestroffer, D., and Hou, X. (2020). “Effective Stability of Quasi-Satellite Orbits in the Spatial Problem for Phobos Exploration”. In: *Journal of Guidance, Control, and Dynamics*, 43 (12), pp. 2309–2320.

Chen, H., Rambaux, N., Lainey, V., and Hestroffer, D. (2022). “Mothership-Cubesat Radioscience for Phobos Geodesy and Autonomous Navigation”. In: *Remote Sensing*, 14 (7).

Cheney, W. and Kincaid, D. (2008). *Numerical Mathematics and Computing*. 6th ed. Thomson Brooks/Cole.

D’Urso, M. G. (2014). “Analytical Computation of Gravity Effects for Polyhedral Bodies”. In: *Journal of Geodesy*, 88, pp. 13–29.

Dang, Z. (2017). “New State Transition Matrix for Relative Motion on an Arbitrary Keplerian Orbit”. In: *Journal of Guidance, Control, and Dynamics*, 40 (11), pp. 2917–2927.

Dichmann, D. J., Doedel, E. J., and Paffenroth, R. C. (2003). “The Computation of Periodic Solutions of the 3-Body Problem Using the Numerical Continuation Software AUTO”. In: *Libration Point Orbits and Applications*. Ed. by G. Gómez, M. W. Lo, and J. J. Masdemont. 1st ed. World Scientific, pp. 489–528.

Dirkx, D., Fayolle, M., Garrett, G., Avillez, M., Cowan, K., Cowan, S., Encarnacao, J., Lombrana, C. F., Gaffarel, J., Hener, J., Hu, X., Van Nistelrooij, M., Oggionni, F., and Plumaris, M. (2022). “The Open-Source Astrodynamics Tudatpy Software – Overview for Planetary Mission Design and Science Analysis”. In: *Europlanet Science Congress*. Granada, Spain.

Dobrovolskis, A. R. (1996). “Inertia of Any Polyhedron”. In: *Icarus*, 124 (243), pp. 698–704.

Doedel, E. J., Romanov, V. A., Paffenroth, R. C., Keller, H. B., Dichmann, D. J., Galán-Vioque, J., and Vanderbauwhede, A. (2007). “Elemental Periodic Orbits Associated With the Libration Points in the Circular Restricted 3-Body Problem”. In: *International Journal of Bifurcation and Chaos*, 17 (8), pp. 2625–2677.

Doedel, E. J. (2011). “Numerical Analysis of Nonlinear Equations”. In: *Summer School on Advanced Topics in Numerical and Computational Bifurcation Analysis*. Technical University of Denmark, Copenhagen, Denmark.

Doedel, E. J., Keller, H. B., and Kernevez, J. P. (1991). “Numerical Analysis and Control of Bifurcation Problems (I): Bifurcation in Finite Dimensions”. In: *International Journal of Bifurcation and Chaos*, 1 (3), pp. 493–520.

Doedel, E. J., Paffenroth, R. C., Keller, H. B., Dichmann, D. J., Galán-Vioque, J., and Vanderbauwhede, A. (2003). “Computation of Periodic Solutions of Conservative Systems With Application to the 3-Body Problem”. In: *International Journal of Bifurcation and Chaos*, 13 (6), pp. 1353–1381.

Elsayed, S. M., Sarker, R. A., and Essam, D. L. (2011). “Differential Evolution With Multiple Strategies for Solving CEC2011 Real-World Numerical Optimization Problems”. In: *2011 IEEE Congress of Evolutionary Computation*. New Orleans, Lousiana, pp. 1041–1048.

Ferrari, F. and Lavagna, M. (2018). “Ballistic Landing Design on Binary Asteroids: the AIM Case Study”. In: *Advances in Space Research*, 62 (8), pp. 2245–2260.

Frieger, G. (n.d.). *3D Asteroid Catalogue*. URL: `https://3d-asteroids.space` (visited on May 2, 2022).

Garland, M. and Heckbert, P. S. (1997). “Surface Simplification Using Quadric Error Metrics”. In: *Conference on Computer Graphics and Interactive Techniques*. New York, New York.

Gaskell, R. (2006). *Gaskell Phobos Shape Model V1.0. VO1-SA-VISA/VISB-5-PHOBOSSHAPE-V1.0*. NASA Planetary Data System.

Genova, A., Goossens, S., Lemoine, F. G., Mazarico, E., Neumann, G. A., Smith, D. E., and Zuber, M. T. (2016). “Seasonal and Static Gravity Field of Mars From MGS, Mars Odyssey and MRO Radio Science”. In: *Icarus*, 272, pp. 228–245.

Guo, X., Yan, J., Andert, T., Yang, X., Pätzold, M., Hahn, M., Ye, M., Liu, S., Li, F., and Barriot, J. P. (2021). “A Lighter Core for Phobos?” In: *Astronomy and Astrophysics*, 651 (A110).

Gupta, M. (2020). *Finding Order In Chaos: Resonant Orbits and Poincaré Sections*. MSc thesis, Purdue University, West Lafayette, Indiana.

Hazra, S. (2019). *Autonomous Guidance for Asteroid Descent Using Successive Convex Optimisation: Dual Quaternion Approach*. MSc thesis, Delft University of Technology, Delft, Netherlands.

Hénon, M. (1969). “Numerical Exploration of the Restricted Problem. V. Hill’s Case: Periodic Orbits and Their Stability”. In: *Astronomy and Astrophysics*, 1, pp. 223–238.

Herrera-Sucarrat, E., Palmer, P. L., and Roberts, R. M. (2014). “Asteroid Observation and Landing Trajectories Using Invariant Manifolds”. In: *Journal of Guidance, Control, and Dynamics*, 37 (3), pp. 907–920.

Howard, J. E. and MacKay, R. S. (1987). “Linear Stability of Symplectic Maps”. In: *Journal of Mathematical Physics*, 28 (5), pp. 1036–1051.

Hu, X. and Jekeli, C. (2015). “A Numerical Comparison of Spherical, Spheroidal and Ellipsoidal Harmonic Gravitational Field Models for Small Non-Spherical Bodies: Examples for the Martian Moons”. In: *Journal of Geodesy*, 89 (2), pp. 159–177.

Jacobson, R. A. and Lainey, V. (2014). “Martian Satellite Orbits and Ephemerides”. In: *Planetary and Space Science*, 102, pp. 35–44.

Kawakatsu, Y., Kuramoto, K., Usui, T., Ikeda, H., Yoshikawa, K., Sawada, H., Ozaki, N., Imada, T., Otake, H., Maki, K., Otsuki, M., Muller, R., Satoh, Y., Mary, S., Grebenstein, M., Tokaji, A., Yuying, L., Gonzalez Franquesa, F., Pushparaj, N., and Chikazawa, T. (2020). “System Definition of Martian Moons eXploration (MMX)”. In: *International Astronautical Congress*. Online.

Koon, W. S., Lo, M. W., Marsden, J. E., and Ross, S. D. (2011). *Dynamical Systems, the Three-Body Problem and Space Mission Design*. 1st ed. Marsden Books.

Kuramoto, K., Kawakatsu, Y., Fujimoto, M., Araya, A., Barucci, M. A., Genda, H., Hirata, N., Ikeda, H., Imamura, T., Helbert, J., Kameda, S., Kobayashi, M., Kusano, H., Lawrence, D. J., Matsumoto, K., Michel, P., Miyamoto, H., Morota, T., Nakagawa, H., Nakamura, T., et al. (2022). “Martian Moons Exploration MMX: Sample Return Mission to Phobos Elucidating Formation Processes of Habitable Planets”. In: *Earth, Planets and Space*, 74 (12).

Lara, M. (2003). “Repeat Ground Track Orbits of the Earth Tesseral Problem as Bifurcations of the Equatorial Family of Periodic Orbits”. In: *Celestial Mechanics and Dynamical Astronomy*, 86, pp. 143–162.

Lara, M., Russell, R., and Villac, B. (2007). “Classification of the Distant Stability Regions at Europa”. In: *Journal of Guidance, Control, and Dynamics*, 30 (2), pp. 409–418.

Le Maistre, S., Rivoldini, A., and Rosenblatt, P. (2019). “Signature of Phobos’ Interior Structure in Its Gravity Field and Libration”. In: *Icarus*, 321, pp. 272–290.

Liang, Y., Pushparaj, N., and Kawakatsu, Y. (2022). “Point-to-Point Jumping Transfer on Phobos by Prograde g’ Family”. In: *Journal of Guidance, Control, and Dynamics*, 45 (6), pp. 1172–1183.

Marov, M. Y., Avduevsky, V. S., Akim, E. L., Eneev, T. M., Kremnev, R. S., Kulikov, S. D., Pichkhadze, K. M., Popov, G. A., and Rogovsky, G. N. (2004). “Phobos-Grunt: Russian Sample Return Mission”. In: *Advances in Space Research*, 33 (12), pp. 2276–2280.

Masdemont, J. and Mondelo, J. (2004). “Notes for the Numerical and Analytical Techniques Lectures”. In: *Advanced Topics in Astrodynamics Summer Workshop*. Barcelona, Spain.

Miller, J. K., Konopliv, A. S., Antreasian, P. G., Bordi, J. J., Chesley, S., Helfrich, C. E., Owen, W. M., Wang, T. C., Williams, B. G., Yeomans, D. K., and Scheeres, D. J. (2002). “Determination of Shape, Gravity, and Rotational State of Asteroid 433 Eros”. In: *Icarus*, 155, pp. 3–17.

Montenbruck, O. and Gill, E. (2001). *Satellite Orbits: Models, Methods, Applications*. 2nd ed. Springer.

Muntoni, A. and Cignoni, P. (n.d.). *PyMeshLab*. URL: `https://pymeshlab.readthedocs.io/` (visited on Mar. 31, 2022).

Muralidharan, V. (2021). *Stretching Directions in Cislunar Space: Stationkeeping and an Application to Transfer Trajectory Design*. PhD thesis, Purdue University, West Lafayette, Indiana.

Murchie, S. L., Thomas, P. C., Rivkin, A. S., and Chabot, N. L. (2015). “Phobos and Deimos”. In: *Asteroids IV*. Ed. by P. Michel, F. E. Demeo, and W. F. Bottke. 1st ed. University of Arizona Press, pp. 451–467.

Muscatello, A. C., Mueller, R., Sanders, G. B., and Larson, W. E. (2012). “Phobos and Deimos Sample Collection and Prospecting Missions for Science and ISRU”. In: *Concepts and Approaches for Mars Exploration Meeting*. Houston, Texas.

Musegaas, P. (2012). *Optimization of Space Trajectories Including Multiple Gravity Assists and Deep Space Maneuvers*. MSc thesis, Delft University of Technology, Delft, Netherlands.

Nakamura, T., Ikeda, H., Kouyama, T., Nakagawa, H., Kusano, H., Senshu, H., Kameda, S., Matsumoto, K., Gonzalez-Franquesa, F., Ozaki, N., Takeo, Y., Baresi, N., Oki, Y., Lawrence, D. J., Chabot, N. L., Peplowski, P. N., Barucci, M. A., Sawyer, E., Yokota, S., Terada, N., et al. (2021). “Science Operation Plan of Phobos and Deimos From the MMX Spacecraft”. In: *Earth, Planets and Space*, 73 (227).

NASA (n.d.). *Solar System Dynamics Database*. URL: `https://ssd.jpl.nasa.gov` (visited on Mar. 31, 2022).

Nyrtsov, M. V., Fleis, M. E., Borisov, M. M., and Stooke, P. J. (2015). “Equal-Area Projections of the Triaxial Ellipsoid: First Time Derivation and Implementation of Cylindrical and Azimuthal Projections for Small Solar System Bodies”. In: *Cartographic Journal*, 52 (2), pp. 114–124.

Olds, A. D., Kluever, C. A., and Cupples, M. L. (2007). “Interplanetary Mission Design Using Differential Evolution”. In: *Journal of Spacecraft and Rockets*, 44 (5), pp. 1060–1070.

Oshima, K. and Yanao, T. (2019). “Spatial Unstable Periodic Quasi-Satellite Orbits and Their Applications to Spacecraft Trajectories”. In: *Celestial Mechanics and Dynamical Astronomy*, 131 (23).

Park, R. S., Werner, R. A., and Bhaskaran, S. (2010). “Estimating Small-Body Gravity Field From Shape Model and Navigation Data”. In: *Journal of Guidance, Control, and Dynamics*, 33 (1), pp. 212–221.

Parker, T. S. and Chua, L. O. (1989). *Practical Numerical Algorithms for Chaotic Systems*. 1st ed. Springer.

Pickering, A. (2014). *CDF Study Report: Phobos Sample Return*. European Space Agency.

Poelaert, D., Schniewind, J., and Janssens, F. (2004). “Surface Area and Curvature of the General Ellipsoid”. In: arXiv: `1104.5145`.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical Recipes: the Art of Scientific Computing*. 3rd ed. Cambridge University Press.

Pushparaj, N., Baresi, N., Ichinomiya, K., and Kawakatsu, Y. (2021). “Transfers Around Phobos via Bifurcated Retrograde Orbits: Applications to Martian Moons eXploration Mission”. In: *Acta Astronautica*, 181, pp. 70–80.

Robin, I. and Markellos, V. (1980). “Numerical Determination of Three-Dimensional Periodic Orbits Generated From Vertical Self-Resonant Satellite Orbits”. In: *Celestial Mechanics and Dynamical Astronomy*, 21, pp. 395–434.

Robin, I. and Markellos, V. (1983). “The Mechanism of Branching of Three-Dimensional Periodic Orbits from the Plane”. In: *Dynamical Trapping and Evolution in the Solar System*. Ed. by V. Markellos and Y. Kozai. 1st ed. Springer, pp. 213–224.

Russell, R. P. (2006). “Global Search for Planar and Three-Dimensional Periodic Orbits Near Europa”. In: *The Journal of the Astronautical Sciences*, 54 (2), pp. 199–226.

Sagdeev, R. Z. and Zakharov, A. V. (1989). “Brief History of the Phobos Mission”. In: *Nature*, 341, pp. 581–585.

Scheeres, D. J. (1994). “Dynamics About Uniformly Rotating Triaxial Ellipsoids: Applications to Asteroids”. In: *Icarus*, 110, pp. 225–238.

Scheeres, D. J. (2012). *Orbital Motion in Strongly Perturbed Environments: Applications to Asteroid, Comet and Planetary Satellite Orbiters*. 1st ed. Springer.

Scheeres, D. J., Van wal, S., Olikara, Z., and Baresi, N. (2019). “Dynamics in the Phobos Environment”. In: *Advances in Space Research*, 63, pp. 476–495.

Scott, C. J. and Spencer, D. B. (2010). “Calculating Transfer Families to Periodic Distant Retrograde Orbits Using Differential Correction”. In: *Journal of Guidance, Control, and Dynamics*, 33 (5), pp. 1592–1605.

Senning, J. R. (2019). *Computing and Estimating the Rate of Convergence*. Gordon College, Wenham, Massachusetts. URL: `http://www.math-cs.gordon.edu/courses/ma342/handouts/rate.pdf`.

Seydel, R. (2010). *Practical Bifurcation and Stability Analysis*. 3rd ed. Springer.

Stoer, J. and Bulirsch, R. (2002). *Introduction to Numerical Analysis*. 3rd ed. Springer.

Storn, R. and Price, K. (1997). “Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces”. In: *Journal of Global Optimization*, 38, pp. 341–359.

Takahashi, Y., Scheeres, D. J., and Werner, R. A. (2013). “Surface Gravity Fields for Asteroids and Comets”. In: *Journal of Guidance, Control, and Dynamics*, 36 (2), pp. 362–374.

Tardivel, S. (2016). “The Limits of the Mascons Approximation of the Homogeneous Polyhedron”. In: *AIAA/AAS Astrodynamics Specialist Conference*. Long Beach, California.

Tudat (n.d.). *Tudat Documentation*. URL: `https://docs.tudat.space/` (visited on June 1, 2022).

Vallado, D. A. (2013). *Fundamentals of Astrodynamics and Applications*. 1st ed. Microcosm Press.

Van wal, S. (2014). *The Ballistic Deployment of Asteroid Landers*. MSc thesis, Delft University of Technology, Delft, Netherlands.

Van wal, S. (2018). *High-Fidelity Simulation of Small-Body Lander/Rover Spacecraft*. PhD thesis, University of Colorado Boulder, Boulder, Colorado.

Vaquero, M. (2013). *Spacecraft Transfer Trajectory Design Exploiting Resonant Orbits in Multi-Body Environments*. PhD thesis, Purdue University, West Lafayette, Indiana.

Verhulst, F. (2000). *Nonlinear Differential Equations and Dynamical Systems*. 2nd ed. Springer.

Villac, B. F. (2008). “Using FLI Maps for Preliminary Spacecraft Trajectory Design in Multi-Body Environments”. In: *Celestial Mechanics and Dynamical Astronomy*, 102, pp. 29–48.

Vittaldev, V., Mooij, E., and Naeije, M. C. (2012). “Unified State Model Theory and Application in Astrodynamics”. In: *Celestial Mechanics and Dynamical Astronomy*, 112, pp. 253–282.

Wakker, K. F. (2015). *Fundamentals of Astrodynamics*. 1st ed. Delf University of Technology, Institutional Repository Library. URL: `http://resolver.tudelft.nl/uuid:3fc91471-8e47-4215-af43-718740e6694e`.

Werner, R. A. (2017). “The Solid Angle Hidden in Polyhedron Gravitation Formulations”. In: *Journal of Geodesy*, 91, pp. 307–328.

Werner, R. A. and Scheeres, D. J. (1997). “Exterior Gravitation of a Polyhedron Derived and Compared With Harmonic and Mascon Gravitation Representations of Asteroid 4769 Castalia”. In: *Celestial Mechanics and Dynamical Astronomy*, 65, pp. 313–344.

Willner, K., Oberst, J., Hussmann, H., Giese, B., Hoffmann, H., Matz, K. D., Roatsch, T., and Duxbury, T. (2010). “Phobos Control Point Network, Rotation, and Shape”. In: *Earth and Planetary Science Letters*, 294 (3-4), pp. 541–546.

Willner, K., Shi, X., and Oberst, J. (2014). “Phobos’ Shape and Topography Models”. In: *Planetary and Space Science*, 102, pp. 51–59.

Wulff, C. and Schebesch, A. (2006). “Numerical Continuation of Symmetric Periodic Orbits”. In: *SIAM Journal on Applied Dynamical Systems*, 5 (3), pp. 435–475.

Yang, X., Yan, J. G., Andert, T., Ye, M., Pätzold, M., Hahn, M., Jin, W. T., Li, F., and Barriot, J. P. (2019). “The Second-Degree Gravity Coefficients of Phobos From Two Mars Express Flybys”. In: *Monthly Notices of the Royal Astronomical Society*, 490 (2), pp. 2007–2012.

Zamaro, M. (2015). *Natural and Artificial Orbits Around the Martian Moon Phobos*. PhD thesis, University of Strathclyde, Glasgow, Scotland.

Zimovan-Spreen, E. (2021). *Trajectory Design and Targeting for Applications to the Exploration Program in Cislunar Space*. PhD thesis, Purdue University, West Lafayette, Indiana.