Comet Interceptor

Robust navigation under high uncertainties and corrupted measurements

MSc Thesis

Tiago Cunha de Araújo

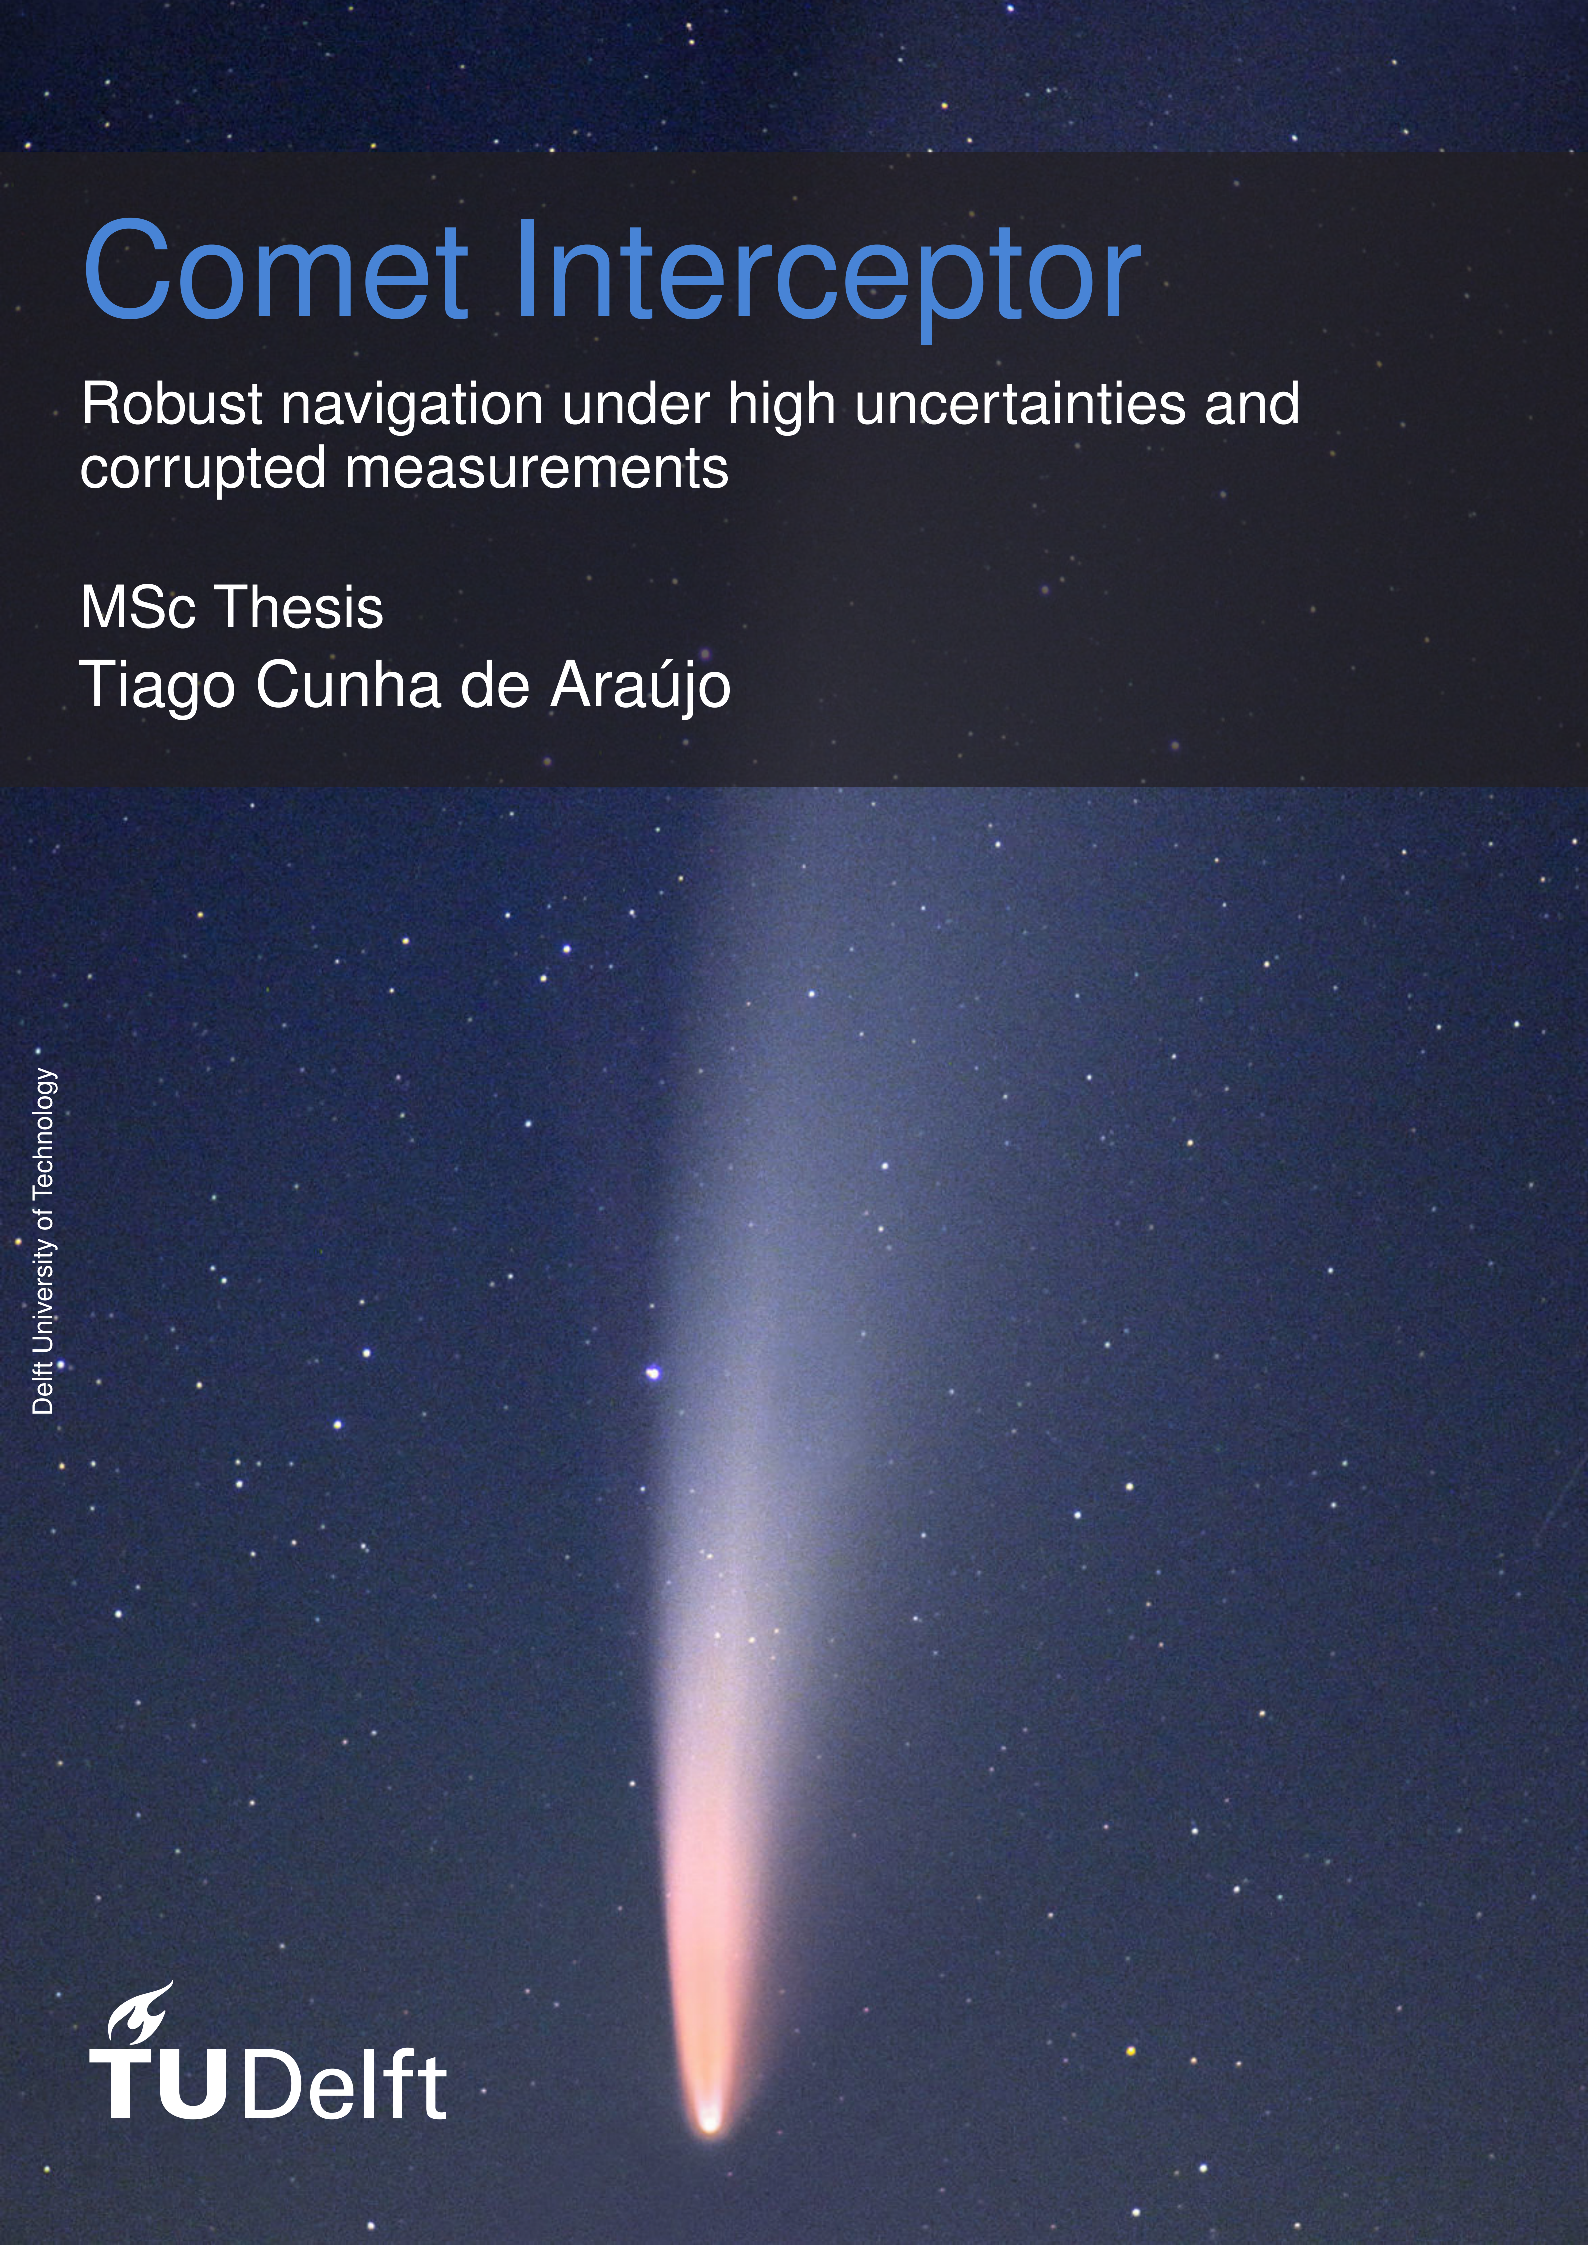

Delft University of Technology

TU Delft

DELFT UNIVERSITY OF TECHNOLOGY

MSC THESIS

Comet Interceptor: robust navigation under high uncertainties and corrupted measurements

---

Author:
Tiago Cunha de Araújo (5124824)

Supervisors:
Erwin Mooij – TU Delft
Francisco Cabral – GMV

Wednesday 6th July, 2022

Cover Image: Comet Neowise, captured at 3:36 a.m. on July 11, 2020.
Courtesy of Goldendale Observatory State Park.

Acknowledgements

This thesis project would not have been possible without GMV Portugal, and especially Francisco Cabral, who proposed this theme. He saw in me a student who was interested in pursuing a career in the field of space mission analysis and issued me this challenge, which gave me an opportunity to gain valuable further experience in the area, while also pushing me to deepen my knowledge in certain topics which were not the main focus of my studies, especially navigation, but that are so essential for this type of work.

Another central piece in the elaboration of this thesis was the consistent and good-humoured supervision by professor Erwin Mooij (who I am sure hates that I added "professor" before his name). Throughout the year, our meetings were a systematic source of good energy, excitement for the results of the work being performed, as well as much-needed feedback and motivation to keep going. Not only that, but during some tougher and more isolated times resulting from the pandemic, which took place during the elaboration of much of this thesis, our talks were a counterbalance for the solitude that everyone working at home was feeling, myself included.

Keeping me company were also all my friends and colleagues back at GMV, with whom I contacted almost daily. Not only were they always helpful when my work needed some of their aid or expertise, but their high spirits helped me keep on pursuing my work even when facing troubles. Pedro Lima and José Carvalho, especially, were my "partners in crime", and their great senses of humour helped me immensely during this project. Additionally, Pedro Cachim was a huge help during some critical parts of this thesis, so I cannot thank him enough for his contribution.

I could not have made it this far without my parents. They have throughout my life strived to make sure that I was happy and had everything I needed to pursue my goals, and during my MSc at TU Delft this was no exception. For the last year, they have always tried to stay up to date with how the project was going, and help me keep a positive and healthy mindset about my work, especially when I was feeling down.

Finally, this thesis is dedicated to all my family and close friends who represent a major block of who I am as a person and who I am constantly striving to be. Each in their own way is a part of me, and I am immensely grateful to every single one of them.

Summary

This thesis was elaborated at TU Delft in collaboration with GMV. The work was developed in the framework of the Comet Interceptor mission, an ESA endeavour aimed at sending a spacecraft to fly-by a comet. To achieve this, an autonomous optical navigation solution was proposed for the encounter, where the on-board camera system snaps pictures of the comet as it approaches, which are then processed and fed to a filtering system. The main goal of this study is to design the filtering system by studying different possibilities, especially in terms of robustness. This is because the mission is expected to be characterised by high uncertainties, related not only with environmental parameters, but also connected to the spacecraft's characteristics and its measurement system. Additionally, the possibility of unpredictable events such as outbursts happening and severely altering the comet's centre of brightness as seen in the captured images must be studied, and the navigation made robust against such phenomena.

With such a background in mind, this thesis first gives some more detail regarding the mission's description and its objectives, and then analyses the proposed research question. After that, various chapters are introduced to give insight into the many aspects of the mission which need to be studied in preparation for the performed analysis in the final sections of the document, attempting to answer the posed research objectives.

Firstly, the mission heritage is looked into, and some past endeavours to comets employing similar strategies are mentioned. This then leads into a discussion about the already established requirements for the mission, either originating from GMV, or added to those to aid the thesis work in particular, keeping in mind that every single one should make sense given the physical description of the mission.

Secondly, the cometary environment is studied, which helps the description of the analysis that follows, both in terms of the forces to be considered, as well as regarding possible disruptive events that might affect measurements, and consequentially the navigation.

After that, a discussion about the dynamics takes place, which mainly decides which forces are relevant and thus should be considered when simulating the fly-by scenarios, and those which do not have a big enough impact to be added to that list. Small sections regarding the reference frames and the propagation scheme are also included in this chapter.

Then, the most important topic of the study is considered, the navigation system. Describing the different filtering possibilities, their pros and cons, as well as considering some techniques which will could be added to aid in recognising measurements corrupted by the aforementioned disruptive events are the focus of this chapter.

Sections on numerical methods and on the software environment where the work will take place are then included. In the former, various numerical processes, which will be crucial for the navigation system's implementation, are discussed and their maths laid out, such as integrators. Then, the latter regards the architectural design of the software, where diagrams of the algorithms being used are presented to introduce the reader into the software framework. Following that, verification and validation of the developed code is discussed, and a small section dedicated to the optimisation of the code.

With all the previous topics out of the way, the characteristics of the scenario are laid down, by defining its duration, time step, nominal orbit geometry, and the expected values and respective uncertainties of all the parameters influencing either the system's dynamics or camera measurements. The goal is to have all the scenario-defining information and values condensed in one place to facilitate reproduction of the performed work.

Finally, two analyses are carried out. The first aims at understanding the scenario better, and figuring out how each variable affects the results of the simulations. Right after, the main study of this thesis is performed, where the scenario is run multiple times and the performance of the various proposed navigation system solutions evaluated.

The first conclusion drawn from this study is that, in this mission, by far the most demanding period happens right before the spacecraft flies by the target, as the relative speed is so high and the vehicle so close to it that keeping the payload correctly pointed is challenging. Despite this, most of the filters analysed are able to always handle the navigation task without failing numerically, even in the presence of disruptive outbursts corrupting camera measurements, with the notable exception of the unscented ones. However, no solution was found to be capable of permanently estimating the target's position with such accuracy that the pointing error would allow for the payload to always be correctly pointed towards the comet's nucleus. In general, all simulations end up with an average of 10-20 seconds of payload downtime due to this. To mitigate this issue, alterations are proposed, such as adapting the measurement frequency over time to match how demanding the navigation is,

or using other cameras or sensors to acquire better information about the comet's position from further away. The ultimate goal is to decrease this downtime to as close to zero as possible, such that the navigation can reliably keep the payload pointing at all times during the mission, regardless of the conditions.

Contents

Abbreviations

Abbreviation		Definition
a-*b* outburst	–	outburst starting *a* minutes and ending *b* minutes prior to the closest approach time
67P	–	67P/Churyumov–Gerasimenko
ADCS	–	Attitude determination and control system
C/A	–	Closest approach
Cam	–	Camera
CCD	–	Charge-coupled device
CDF	–	Concurrent design facility
CMOS	–	Complementary metal-oxide-semiconductor
CoCa	–	Comet Camera
CometEclipticJ2K	–	Comet Ecliptic J2000 reference frame
DNC	–	Dynamically new comet
DS1	–	Deep Space 1
EKF	–	Extended Kalman filter
ENV	–	Environmental requirement
ESA	–	European Space Agency
FCS	–	Flight Control System
FDS	–	Flight Dynamics Systems
FoV	–	Field of view
GNC	–	Guidance, Navigation and Control
GNSS	–	Global Navigation Satellite System
HRI	–	High Resolution Imager
iFoV	–	Instantaneous field of view
IP	–	Image processing
IPF	–	Image processing with unresolved target
IMU	–	Inertial measurement unit
JAXA	–	Japan Aerospace Exploration Agency
KF	–	Kalman filter
Lidar	–	Light Detection And Ranging
LoS	–	Line of sight
MIS	–	Mission requirement
MRI	–	Medium Resolution Imager
NAC	–	Narrow-angle Camera
NASA	–	National Aeronautics and Space Administration
NExT	–	New Exploration of Tempel 1
Osiris-REx	–	Origins, Spectral Interpretation Resource Identification, Security-Regolith Explorer
NMA	–	Non-modelled acceleration
RBSA	–	Regression-based sensitivity analysis
RK4	–	Runge-Kutta 4th order
rng()	–	Random-number generating seed
S/C	–	Spacecraft
SEL2	–	Sun-Earth Lagrangian point L2

Abbreviation		Definition
SO	–	Study objective
SRP	–	Solar radiation pressure
TCM	–	Trajectory correction manoeuvre
TRL	–	Technology readiness level
UKF	–	Unscented Kalman filter
V&V	–	Verification and validation
WAC	–	Wide-angle Camera

List of symbols

Symbol		Definition	Units
$\boldsymbol{a}$	–	acceleration	m/s^2
$\hat{\boldsymbol{a}}$	–	vector of Chebyshev coefficients	–
$\bar{\boldsymbol{a}}$	–	auxiliary vector used in the rank-one update	–
A_{dust}	–	effective area of the spacecraft on the side of the incoming dust	m^2
$\boldsymbol{A}_i$	–	outgassing force direction components	N
A_{Sun}	–	effective area of the spacecraft on the sunlit side	m^2
$\boldsymbol{A}$	–	Jacobian of the system's dynamics function with respect to the estimated states	–
$\boldsymbol{B}$	–	Jacobian of the system's dynamics function with respect to the considered states	–
c	–	speed of light	m/s
C_r	–	radiation pressure coefficient	–
C/A	–	closest-approach distance	m
d	–	distance	–
$\boldsymbol{D}$	–	diagonal matrix resulting from the UDU factorisation	–
D	–	number of dimensions of the state vector	–
dm_d	–	dust mass infinitesimal	kg
$\bar{\boldsymbol{f}}$	–	auxiliary vector used in the modified rank-one update	–
$\boldsymbol{F}$	–	force vector	N
$\boldsymbol{F}(x)$	–	system's dynamics function	–
FoV	–	camera's field of view	rad
$\boldsymbol{G}(x)$	–	system's measurement function	–
$g(r)$	–	outgassing radial magnitude	–
$\tilde{\boldsymbol{H}}$	–	Jacobian of the system's measurement function with respect to the system's states	–
$\boldsymbol{H}$	–	product of $\tilde{\boldsymbol{H}}$ and $\boldsymbol{\phi}$ (time updated Jacobian)	–
$\boldsymbol{H^F}$	–	regression-based sensitivity analysis' hat matrix	–
h	–	integrator time step	s
$\boldsymbol{I}$	–	identity matrix	–
$\boldsymbol{I_M}$	–	moment of inertia matrix	kg·m^2
j	–	number of observations in the processed batch	–
$\boldsymbol{J}$	–	square matrix of ones	–
$k_{1\ldots4}$	–	RK4 integrator coefficients	–
$\boldsymbol{K}$	–	Kalman gain matrix	–
k_g	–	outgassing radial magnitude exponent constant	–
L	–	regression-based sensitivity analysis least squares' fit number of fitting terms	–
$\boldsymbol{l_n}$	–	column vectors of the Cholesky decomposition of the initial covariance matrix	–
m	–	spacecraft mass	kg
m'	–	order of magnitude of dust particle's mass	–
m_g	–	outgassing radial magnitude exponent constant	–
MD	–	Mahalanobis distance	–

Symbol		Definition	Units
N_d	–	cometary dust spatial density	kg/m^3
n_g	–	outgassing radial magnitude exponent constant	–
nP	–	number of parameter states considered by the filters	–
nS	–	number of estimated states considered by the filters	–
$\boldsymbol{N}$	–	normal vector	–
N_{mea}	–	number of independent measurements	–
N_{states}	–	number of states	–
p, p_{total}	–	number of pixels and pixel total	–
$\boldsymbol{P}$	–	filter's state covariance matrix	–
$\boldsymbol{P}^-$, $\boldsymbol{P}^+$	–	predicted and estimated covariances of the system's state	–
P_{Sun}	–	total power emitted by the Sun	W
$\boldsymbol{P_S}$	–	estimated states' covariance matrix block	–
$\boldsymbol{P_p}$	–	considered states' covariance matrix block	–
$\boldsymbol{P_{Sp}}$	–	estimated and considered states' cross covariance blocks	–
q_{SC}	–	spacecraft's orientation in space	–
$\boldsymbol{Q}$	–	unscented filters' Kalman matrix equivalent	–
Q_d	–	dust production rate, or cometary activity	kg/s
$\boldsymbol{r}$,$\dot{\boldsymbol{r}}$,$\ddot{\boldsymbol{r}}$	–	position, velocity, and acceleration vectors	m,m/s,m/s^2
$\boldsymbol{ref_{x,y,z}}$	–	first, second and third unit vectors of the Comet Ecliptic J2000 reference frame triad	–
r_0	–	outgassing radial magnitude constant	au
$\boldsymbol{R_{X2Y}}$	–	rotation matrix from the reference frame X to Y	–
$\boldsymbol{R_{SD2CAM}}$	–	rotation matrix between the sunward/Sun perpendicular frame and the camera's FoV frame	–
R_{comet}	–	target comet's radius	m
$\boldsymbol{R}$	–	measurement covariance matrix	–
R^2	–	regression-based sensitivity analysis least squares fit coefficient of determination	–
R^2_{adj}	–	regression-based sensitivity analysis least squares fit adjusted coefficient of determination	–
$\boldsymbol{R'}$	–	Cholesky decomposition of the measurement covariance matrix	–
$\bar{R}$	–	independent measurements' covariance matrix	–
$\hat{\boldsymbol{R}}$, $\hat{\boldsymbol{S}}$, $\hat{\boldsymbol{T}}$	–	first, second and third unit vectors of the B-plane reference frame triad	–
$\boldsymbol{s}$, $\boldsymbol{s}^-$	–	sigma points and their propagations, respectively	–
$\bar{\boldsymbol{S}}$	–	auxiliary matrix used in the rank-one update	–
$\boldsymbol{S_{cons}}$	–	consider covariance analysis' sensitivity matrix	–
t	–	time	s
T	–	measurement rejection threshold	–
$\boldsymbol{T}$	–	auxiliary vector in Chebyshev interpolation	–
t_f	–	total encounter time	s
u_d	–	terminal speed of the ejected dust flow	m/s
$\boldsymbol{U}$	–	upper triangular matrix with diagonal of ones resulting from the UDU factorisation	–
$\boldsymbol{v}$	–	velocity	m/s
$\bar{\boldsymbol{v}}$	–	auxiliary vector used in the modified rank-one update	–
$\boldsymbol{v}_0$	–	relative velocity of the spacecraft relative to the incoming dust	m/s

Symbol		Definition	Units
$\boldsymbol{V}$	–	cross covariance matrix	–
$\boldsymbol{W}$	–	sigma-point weights vector	–
$\boldsymbol{W}$	–	covariance of the measurement residual / innovation covariance	–
$\boldsymbol{X}$, $\boldsymbol{x}$	–	true filter state, and difference between the predicted and reference filter states	–
$\boldsymbol{X^*}$	–	filter's reference state	–
$\boldsymbol{X^-}$, $\boldsymbol{X^+}$	–	filter's *a priori* and *a posteriori* state estimates	–
$\boldsymbol{X^F}$	–	regression-based sensitivity analysis least squares' input matrix	–
$\hat{\boldsymbol{X}}$, $\hat{\boldsymbol{Y}}$, $\hat{\boldsymbol{Z}}$	–	first, second and third unit vectors of the vehicle's reference frame triad	–
$\boldsymbol{y^F}$	–	regression-based sensitivity analysis least squares' real output vector	–
$\hat{\boldsymbol{y}}^{\boldsymbol{F}}$	–	regression-based sensitivity analysis least squares' predicted output vector	–
$\boldsymbol{Z}$, $\boldsymbol{z}$	–	true measurement, and difference between the true and predicted measurements	–
$\bar{\boldsymbol{Z}}$	–	sigma-point observation	–
$\hat{\boldsymbol{Z}}$	–	weighted average of sigma-point observations	–
α	–	angular half-width of a dust jet	rad
α_g	–	outgassing radial magnitude constant	–
$\alpha_{R_{comet}}$	–	target comet's apparent radius	rad
$\bar{\boldsymbol{\alpha}}$	–	auxiliary vector used in the modified rank-one update	–
β	–	angle between the relative velocity vector and the LoS direction	rad
$\hat{\gamma}$	–	regression-based sensitivity analysis least square's matrix of coefficients	–
$\boldsymbol{\delta}$	–	measurement residuals	–
Δ	–	difference between two scalars or vectors	–
ϵ	–	momentum enhancement factor	–
$\boldsymbol{\epsilon_{IP}}$	–	image processing error	–
ϵ_{SD}, ϵ_{SP}	–	sunward and Sun perpendicular components of the image processing error, respectively	–
$\boldsymbol{\epsilon}^F$	–	regression-bases sensitivity analysis least squares' residuals vector	–
η	–	scanning mirror angle	rad
$\boldsymbol{\theta}$	–	consider states' transition matrix	–
$\boldsymbol{\Lambda}$	–	information matrix	–
$\bar{\boldsymbol{\lambda}}$	–	auxiliary vector used in the modified rank-one update	–
μ	–	gravitational parameter of the considered body	m^3/s^2
$\boldsymbol{\mu}$	–	average state for sigma point computation	–
σ	–	standard deviation	–
$\boldsymbol{\phi}$	–	state transition matrix	–
$\boldsymbol{\Psi}$	–	full transition matrix	–
$\boldsymbol{\omega}$	–	angular velocity	rad/s

Chapter 1

Introduction

The first chapter is dedicated to analysing the mission that this study will contribute to. Firstly, a thorough description of the mission phase by phase is put forward, followed by the objectives that it aims to achieve. Secondly, a research question is established, which serves a motivation for the thesis work moving forward. Finally, study objectives are set, which are simpler questions stated as the motivation to develop the work of each of the chapters necessary to analyse the problem and gather the data to answer the main research question. Thus, the planning for the coming work is established here.

1.1 Mission description

As remnants of the formation of the Solar System, or other stellar systems, comets have for millennia astonished mankind, being seemingly inexplicable phenomena. Depending on their characteristics, such astronomical bodies can be seen even during the day, spanning more than 45 degrees on the sky with their tails. These awe-inducing spectacles unsurprisingly led populations all around the world to come up with explanations for their existence, usually supernatural or religious in nature. In fact, a few centuries after it was noted that comet Halley had a 75-76 year long period, which happened near the turn of the 18th century, records have been found for 29 out of the 30 apparitions it has had from 240 BCE up until 1986.

Eventually, however, using Tycho Brahe's observations and Newton's recently developed gravitational theory, Edmond Halley was able to identify that the comet, which ended up being named after him, performed periodic visits of the inner Solar System, as the orbital parameters of each passage were roughly the same each time, just like the period separating the perihelia. With this, it became clear that comets were astronomical bodies with long – when compared to the Earth's, certainly – orbital periods, which became visible once their orbits brought them close to the Sun, forming big tails in the sky discernible with the naked eye.

As their creation is now understood, it is known that they hold some of the very last pieces of raw material remnant from the formation of the Solar System, or even other similar star systems. Thus, comets are remarkably interesting from a science perspective, as astronomical fossils of mankind's home in the universe.

The Comet Interceptor mission has been announced by the European Space Agency in June 2019 as a fast-class mission aimed at being Europe's next effort for studying comets, following the successes of Giotto and Rosetta. Composed by three spacecraft (S/C), a main one and two smaller probes, these will travel to the orbit of a pristine, yet to discover comet and perform a fly-by from three different points of view, allowing for a 3D characterisation of this object[1]. Because of this division, do note that, during this introduction, when "spacecraft" is used in the singular form, it generally refers to the conjoined vehicles operating as one, whereas the plural form refers to each operating on its own after separation.

Being a fast-track development mission, or F-class mission, Comet Interceptor takes approximately eight years to launch. With a mass smaller than one metric ton, the three spacecraft will be launched in a standard dual launch configuration configuration, typically known as "piggybacking", taking advantage of another ESA mission's, ARIEL, which is planned to launch in late 2028 in an Ariane 62 from Kourou, French Guiana. Not only will this launch the mission into space, but also boost it to the Sun-Earth Lagrangian point L2 (SEL2), on the opposite side of the Earth as seen from the Sun, where the first major phase of the mission begins.

The target body, a dynamically new comet (DNC), differs from the ones studied in the aforementioned missions for not having ever before visited the inner Solar System, which is why it is also referred to as pristine. Whereas these former short-period comets perform multiple fly-bys of our planetary system, the Comet Interceptor spacecraft will encounter a body with material that has undergone little transformation since its formation and the formation of the Solar System itself, allowing for the study of the evolution of these bodies as they migrate inwards. Another possible target is an interstellar body originating from a distinct star system, which, although much rarer and thus less likely to appear in the first place, would allow the exploration of the formation and evolution of these other systems.

---

[1] European Space Agency. ESA's new mission to intercept a comet. Retrieved March 16, 2021 from `http://www.esa.int/Science_Exploration/Space_Science/ESA_s_new_mission_to_intercept_a_comet`.

In either case, these types of bodies are typically only found a few months to years before their perihelion, which does not leave enough time for a whole mission to be planned, constructed and launched to meet the target body. This is why the Comet Interceptor's first phase is located at L2. Here, the spacecraft will wait for one of these suitable targets to be detected, evaluated and selected. Then, using its own propulsion, the spacecraft will rendezvous with the comet, with each separate component detaching from one another 24-48h prior to the closest approach, and performing the fly-by on their own [GMV, 2020a].

During this encounter phase, each spacecraft, equipped with complementary science payloads, will perform measurements of the comet's nucleus, its gas, dust and plasma environment from the different angles, allowing for a 3D description of the interaction of the comet with the incoming solar wind. After this, the smaller probes will relay the acquired data to the main spacecraft, which will in term downlink all the measurements back to Earth. From launch to encounter, the mission is expected to be completed within 5 years.

The mission is divided into phases as follows:

- Launch & Early Operation Phase – takes 2 days and ends upon the first trajectory correction manoeuvre being performed.
- Transfer to SEL2 – cruise until reaching the pseudo-periodic parking orbit around SEL2, expected to take 90-120 days.
- Waiting at SEL2 – station keeping around SEL2, may take up to 3 years.
- Transfer to encounter – from SEL2 to the fly-by region, the spacecraft uses its on-board propulsion system, taking between 1/2 and 4 years, for most considered transfers. During the transfer, it is possible to fly-by an asteroid so as to test its systems and the operational timeline of the actual comet. Science may also be acquired but its not the focus of such an endeavour.
- Approach and Encounter – in the last 1-2 months before the closest approach, the spacecraft navigates relying on optical measurements, and engages targeting manoeuvres to release the small probes to the desired trajectory before diverting the main spacecraft to its own fly-by trajectory.
- Post-encounter – following the fly-by, the mission focuses on the downlink of the acquired data, which is expected to take up to 6 months.
- Decommissioning – operations to decommission the spacecraft in heliocentric orbit, expected to take approximately 2 weeks.

There are a few key aspects, arising from the basic design of the mission, which pose some of the greatest challenges to its success [GMV, 2020a]. Firstly, since the spacecraft is launched previously to the detection of the target comet, the whole mission must be designed with this lack of information in mind. Without knowing the target *a priori*, the comet's physical characteristics, as well as the fly-by conditions are unknown, and a range of possibilities must always be accounted for in the performed analyses.

Secondly, depending on the chosen target, the fly-by speed of the spacecraft relative to the comet can vary between 10 and 70 km/s, which means that the encounter phase is too short for any failure to be recoverable. Additionally, the actuators aboard the spacecraft must be designed to account for this short fly-by, meaning that the tracking speeds should be high enough to keep the instruments pointed in the right direction. This requirement is further aggravated by the fact that the scientific measurements are performed by an instrument with a small field of view (FoV), which demands for correspondingly small pointing errors.

Finally, the fly-by will delve into the coma and gas regions of the comet, which are irregular and unpredictable, adding perturbations to the dynamics of the spacecraft that are amplified by the high relative fly-by speed. Due to this as well as all the aforementioned reasons, robustness of all the spacecraft's systems is crucial, even more so than usual for space missions.

One important aspect to note is that the mission is being planned with backup targets in mind. In the eventuality of a DNC not being found within the first three years of waiting at SEL2, the objective is switched to one of two pre-selected comets – 73P/Schwassmann-Wachmann or 26P/Grigg-Skjellerup. These are, in the earliest stages of mission development, where the project currently stands, the chosen backups, and the transfers from SEL2 to the respective encounter zones have been determined. However, should further analyses determine that there are reasons to believe that the targets are sub-optimal in terms of mission restrictions or objectives, then there is an additional list of backups being studied.

1.2 Mission objectives

The three spacecraft, A (the main spacecraft) and B1 and B2 (the two probes), all of which initially cruise as one, eventually separate prior to the closest approach with the selected target comet, ending up with three

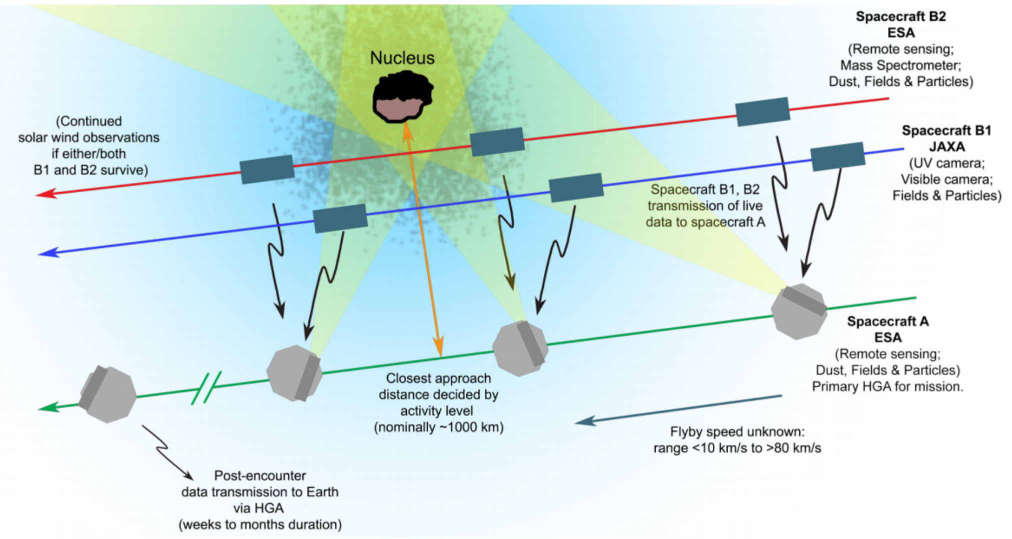

Figure 1.1: Schematic representing the fly-by trajectories and the instruments aboard the different spacecraft. Retrieved from [GMV, 2021a].

separate and distinct fly-by orbits of the body. Using the instruments on-board of each spacecraft – which will not be specified here as the focus of this study is the navigation, but are indicated in Figure 1.1 – the mission aims at studying the following aspects [GMV, 2021a]:

1. The bulk shape of the nucleus and its surface's morphology.
2. The bulk composition of the nucleus' surface.
3. The comet's activity in terms of material ejection.
4. The molecular composition across multiple locations of the coma.
5. The characteristics of the dust in the coma.
6. The structure of the plasma environment around the comet, as well as its interaction with the dust.
7. The transfers of energy, mass and momentum throughout the coma.

1.3 Research question

As the title of the document suggests, the focus of the study is not the science aspect of the mission, but the navigation during the closest-approach phase. Since the encounter with the comet will take place at an extremely high relative speed, even in astrodynamical terms, the most crucial part of the fly-by happens too fast for the navigation solution to be able to rely on communication with ground. Because of this, one of the mission requirements is that spacecraft A must be able to perform optical autonomous navigation in the vicinity of the comet.

Given the above, the research question proposed can be stated as

> **RSQ – 1.** How do outbursts present in a cometary environment during the encounter phase affect the optical translation navigation solution?

The most significant challenges standing in the way of the success of this endeavour are the high uncertainties arising from the poor *a priori* knowledge of the comet and its respective environmental characteristics, as well as, just as importantly, the possibility of disruptive events taking place, which can corrupt the navigation sensor's measurements, such as material outbursts from the comet.

With the goal of answering this question in mind, the next section performs an overview of the topics which need to be investigated, analysed or discussed to establish the necessary bases to give an answer, and to delve into the specifics of how such can be achieved. By devising equally simple sub-questions for each topic, an objective is thus established for each of the following chapters.

1.4 Study objectives

To tackle the research question at hand, the first step is to analyse previous missions, ESA's or otherwise, that have explored comets in the past. This has the objective of both gaining knowledge of how these past missions have achieved their goals, in terms of the pre-flight mission analysis performed or the actual systems employed; as well as shedding some light on the discoveries that they have made about the comets themselves and their environment, as a bonus. This can be summarised by the question

SO – 1. Have previous cometary exploration missions employed similar navigation solutions?

In comparison to these past endeavours, the Comet Interceptor will be discussed, in terms of the systems that have been chosen in the planning, as well as regarding all the requirements which have already been established and must be obliged. These will significantly narrow down the analysis up ahead, and should thus be as meticulously specified as possible. This results in

SO – 2. What are the navigation-related requirements for the Comet Interceptor mission?

Having studied previous missions and the gained knowledge on cometary environments, the focus turns to consolidating this information, and other gained through different indirect means, so as to describe, with as much detail a possible, the environment around a comet, which will affect the spacecraft as they fly-by. Not only is the objective knowing what sort of environment surrounds the body, mainly in terms of particles, their sizes, velocities and mass, as a function of the position, but how this environment evolves over time as the comet approaches the Sun and interacts with the incoming solar wind. This allows for a prediction of the perturbing forces caused by particle impacts against the spacecraft.

Additionally, a description of possible disruptive events happening in the vicinity of the comet is a necessary step, as these, however unlikely, can have a major impact in the measurements and make the navigation diverge, causing the mission to fail. As such, they should be accounted for and properly described, so as to develop tests and assess whether any designed navigation solution can handle them. As such, the goal will be answering the following

SO – 3. How can the cometary environment affect the navigation?

After that, the conditions are met to study the flight dynamics. Having discussed the reference frames of interest where the relevant variables are defined, the focus turns to describing all the forces that act – significantly – on the spacecraft. The main forces are, as in most astrodynamics problems, gravitational, and an analysis of which bodies to consider for this, and which model to use for each, needs to be performed. Other forces to consider are the cometary dust impacts and the solar radiation pressure, as the main perturbations. With the set of relevant forces described, the best propagation method to apply the dynamics is discussed. The question posed thus becomes

SO – 4. What forces are required to describe the spacecraft's dynamics adequately?

With that, the description of the problem at hand is concluded, and the focus can be turned towards the navigation systems employed to answer the research question. Firstly, an assortment of different filtering algorithms are proposed, and a preliminary trade-off between them is performed, with performance, robustness and computational effort being the primary criteria. Finally, additional techniques, which might help handling the high uncertainties and corrupt sensor data, are taken into consideration.

SO – 5. Which are the best *a priori* filtering candidates for solving the task at hand?

Having described the primary candidates to tackle the problem at hand, a description of the methods used to simulate the mission's approach and test the designed solutions must be performed. Starting with the numerical methods, the integration method, as well as any other required computations, are chosen in accordance with the required accuracy of the results or the filters that will be implemented, and described. Then, an analysis of the software used is performed, starting with its architecture. This is to make sure that the developed code is consistent with the framework in which it is included. This exercise will also crucially help catalogue the already available functionalities included in the software, as these will be tools which do not need to be developed or verified and validated (V&V), saving precious time and effort. And, finally, the verification and validation of all the developed scripts and associated models must be performed, since without this process the results cannot be trusted. In summary

SO – 6. What numerical algorithms will be needed to properly simulate the mission?

and

SO – 7. How can the developed code be tested to perform V&V within the software framework?

After the developed functionalities have been tested, the focus shifts towards simulating the encounter and answering the primary research question **RSQ – 1**. For that, the scenario needs to be set first. Whereas before the cometary environment and the dynamics have been analysed, now the respective parameters that define all the factors impacting the mission must be established, resulting in

SO – 8. How is the Comet Interceptor's scenario defined?

Arising from this definition is the question regarding the effect of each of the parameters considered. Because the mission is defined by multiple dynamics and measurement-related biases and noises, knowing which one or ones are the driving factors can help better understand the results gathered in the analysis performed afterwards. Thus

SO – 9. How much does each scenario uncertainty affect the simulation results?

Finally, the main question is answered by analysing the multiple proposed navigation solutions in the defined scenario. Firstly, the goal is knowing how these handle the baseline, unaltered scenario, before introducing the measurement disruption events and comparing the results.

SO – 10. How do the proposed navigation solutions handle the baseline and disrupted scenarios?

With these questions being pursued, the mission should be understood well enough to answer the main research question. Not only that, but the knowledge gained should help give advice for future work, as well as determine whether the proposed requirements are adequate and the mission objectives feasible with the established restraints. That not being the case, suggestions can then be put forward to either relax some of those restrictions, or change them entirely in an attempt to ensure that the objectives can be met.

Chapter 2

Mission heritage

In this chapter, the first goal is to analyse some past missions to comets and asteroids in search of similar optical navigation efforts, especially those who rely on autonomy from ground. These will establish a base to then, looking at the sensors used by these missions, discuss those used by Comet Interceptor in an attempt to draw parallels between the mission, and understand why such strategies have been chosen.

Having established the components to be used, the mission and system requirements are set as a way of narrowing down on the specific scenario conditions that have to be analysed. The mission requirements are more general and establish the characteristics of the target, trajectory, operations, and objectives. Then, firstly, the system requirements lay down what the GNC block must be able to do, especially in terms of navigation, and, secondly, the cameras used in this mission are characterised.

2.1 Previous missions

The heritage is a crucial part of most space missions, when such information is available. By studying previous missions to similar targets, in this case, comets, one can not only find valuable scientific knowledge which these have acquired, and might be helpful for the development of new endeavours, but also details about how said missions have been executed, and as such draw conclusions from their successes and failures. Following this methodology ensures that the mission being developed uses all the past information, giving it the best chance of succeeding, by being designed with robustness to all possible disruptive elements in mind. Since space missions are such huge efforts, both in monetary and total work terms, and are by nature hard to recover in case of a failure, this is highly desirable.

As such, the first part of this chapter is dedicated to analysing previous fly-by related space missions, with the objective of discovering how their proximity navigation solutions have been implemented. In particular, cases of autonomous optical navigation are the real focus, and assessing how these have been implemented in the past can give some clues as to the best strategy for the Comet Interceptor mission.

It is worth noting two aspects at this point. Firstly, the missions being analysed will mostly be related to cometary exploration. Even though it is true that this sort of navigation can be employed during rapid fly-bys of similar bodies, such as asteroids, comets have certain particularities which pose a challenge to the data acquisition and image processing systems which need to be accounted for. This refers to the aforementioned gas outbursts and other such phenomena, for example. As such, these cases will likely provide the most useful information for the case at hand for including these particularities.

Secondly, the focus of the heritage review will be on the navigation systems. Even though previous missions have acquired invaluable information on the cometary environment, which will be crucial to properly design the mission for its task, the environment used to develop Comet Interceptor will be described in detail its respective chapter. Nevertheless, understanding how the volume around a comet is composed and interacts with the spacecraft is crucial, as is knowing how to model any disruptive events, so any information regarding this is certainly welcome. Nevertheless, the navigation systems will be the major focus of the following analysis.

With this in mind, looking at NASA's website[1], a list of past and present missions to comets can be found, and from these, the ones following NASA's Deep Space 1 (DS1) will be those of interest. This is because DS1 was the first effort to employ a strategy of autonomous optical navigation in a deep-space mission [Riedel et al., 2000], named AutoNav. This system, or some similar variation, would later see usage in following missions, namely NASA's Stardust [Bhaskaran et al., 2004] and Deep Impact [Mastrodemos et al., 2005].

On Europe's side of space exploration, Giotto, which performed a fly-by of comet Halley in 1986, was actually planned to beat the aforementioned American efforts by 13 years and use a camera to autonomously track the comet's nucleus. Sadly, though, the on-board camera responsible for this tracking malfunctioned right before the closest approach, likely due to cometary dust impacts [Bhaskaran et al., 2004]. Nevertheless, with the recent Rosetta mission, ESA returns to cometary exploration, and even though the main part of the mission's profile consists of a prolonged study of the body as opposed to a swift fly-by, the spacecraft still used optical methods

---

[1] National Aeronautics and Space Administration. Missions to Comets: Current and Past Missions. Retrieved March 23, 2021 from `https://nssdc.gsfc.nasa.gov/planetary/planets/cometpage.html`.

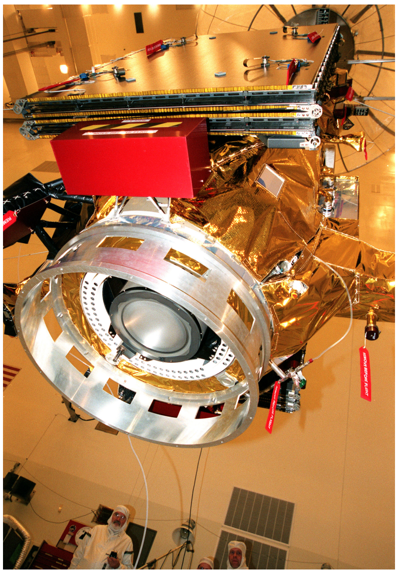

Figure 2.1: View of the Deep Space 1 spacecraft from the bottom, exposing the ion propulsion system, one of the main technology tests for this mission. Credit to: NASA/JPL.

to obtain information and navigate in this phase. Thus, even though no autonomous system is required, due to the prolonged nature of the visit to the comet, a lot can still be learned about the equipment. Adding to this, before entering this main section of the mission, Rosetta flew by asteroid 2867 Šteins, where indeed such a system was put to use.

These are the primary sources of knowledge, since they present the closest mission scenarios relative to the Comet Interceptor. However, other endeavours have used optical navigation solutions to visit similar bodies, such as asteroids, and other will in the future, meaning that this technology is crucial for missions to smaller bodies, whose ephemerides are poorly known [Leonard et al., 2017].

2.1.1 Mission description

Going chronologically, the first mission of interest is DS1. Launched in 1998, this endeavour had a primary objective of testing new technologies in deep-space missions, including the autonomous navigation of interest, as well as others that are nowadays taken as granted, such as ion propulsion. In doing so, the spacecraft would fly-by the asteroid 9969 Braille, at an altitude of just 27 km. After having succeeded, the mission was extended to include an additional 2700 km fly-by of comet 19P/Borrelly. Curiously, following the first fly-by, the star-tracker aboard the vehicle was not functional, and, to extend this mission, NASA had to re-purpose the on-board camera spectrometer to substitute the previous instrument, despite the FoV of the latter being 100 times smaller, and the process taking 80 times the original time to estimate the orientation[2].

Then, Stardust launched in 1999, with the main goal of gathering dust and volatiles from the coma of the P/Wild 2 comet, returning them to Earth afterwards. In doing so, the spacecraft first flew by the 5535 Annefrank asteroid, on its way to the main target. After achieving the mission's goal, Stardust saw a mission extension named New Exploration of Tempel 1 (NExT), and the spacecraft was redirected perform a fly-by with the goal of observing crater alterations following the impact of Deep Impact's impactor on comet 9P/Tempel 1[3].

This extension would not have happened without the next mission, Deep Impact, which aimed at performing a rendezvous with the already mentioned comet 9P/Tempel 1, and launching a projectile directly at the nucleus. With this, pristine material from beneath the body's surface is exposed and studied by the main spacecraft, shedding light on the properties of the surface and interior of this type of celestial bodies. Following the trend of the previous missions, Deep Impact also saw an extension of its objective. Known as EPOXI, this phase was composed of two parts. Firstly, the spacecraft performed a fly-by of 103P/Hartley 2, another comet, returning images and data, and then it used its imaging system to observe stars with known exoplanets, characterising

---

[2] ESA - Earth Online. Deep Space 1. Retrieved March 25, 2021 from `https://earth.esa.int/web/eoportal/satellite-missions/d/deep-space-1#foot12%29`.

[3] NASA Space Science Data Coordinated Archive. Stardust/NExT. Retrieved March 25, 2021 from `https://nssdc.gsfc.nasa.gov/nmc/spacecraft/display.action?id=1999-003A`.

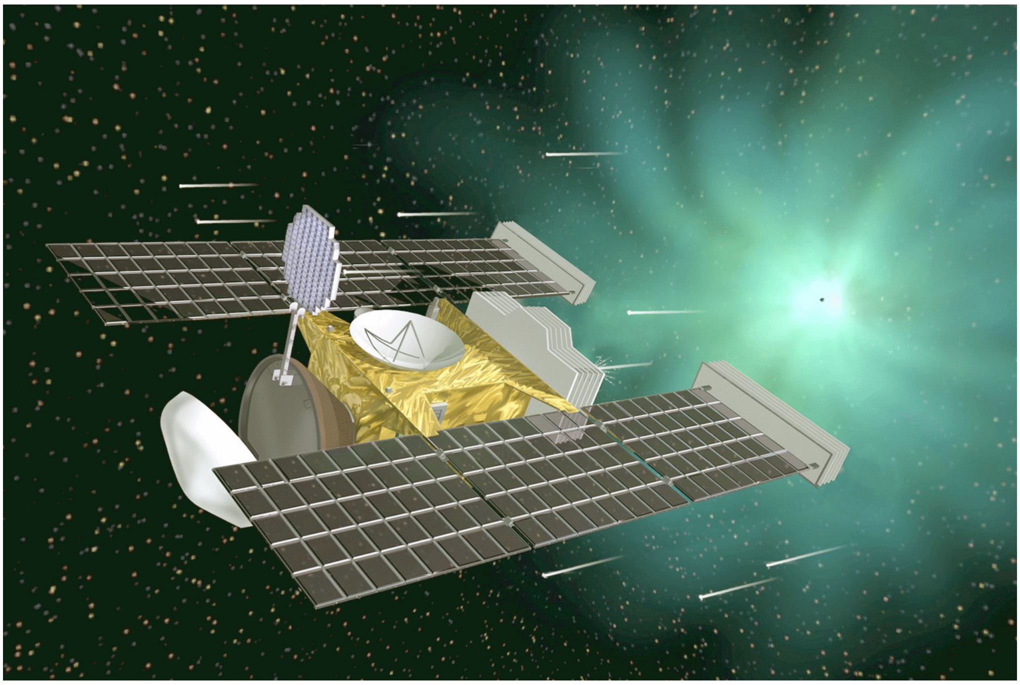

Figure 2.2: Artistic depiction of the Stardust spacecraft on a cometary environment. The protective shielding, on the right of both the main part of the spacecraft and the solar panels, is clearly recognisable. Credit to NASA/JPL.

said planets and discovering new ones[4].

Leaving the domain of cometary missions, others are worth mentioning that have employed similar strategies, all dedicated to studying asteroids. NASA's Origins, Spectral Interpretation Resource Identification, Security-Regolith Explorer – Osiris-REx – mission launched in 2016 with the goal of approaching asteroid Bennu, navigate to its surface, gather a sample and then return it to Earth [Lorenz et al., 2017].

From JAXA come two equally important missions, Hayabusa and its successor Hayabusa 2. The first endeavour approached and landed on asteroid Itokawa in late 2005, with the goal of gathering and returning asteroid material back to Earth, similarly to Osiris-REx [Hashimoto et al., 2010]. Hayabusa 2 was quite similar, approaching asteroid Ryugu nearly 13 years after its predecessor to recover and return material samples from beneath the surface. This was done by creating a crater with an impactor and gathering material from the resulting impact site. Other remarks include landing and deploying two rovers to explore an asteroid's surface for the first time[5].

Not only is using optical navigation a common solution for missions of the past, but it is also planned to be used in the near future, with examples such as HERA. Dedicated to study the binary asteroid system, Didymos, this ESA venture will operate in close proximity to the targets, meaning that high autonomy of the on-board GNC is a major mission goal [Pellacani et al., 2019]. Currently being developed, this mission is schedule to launch starting late 2024.

2.1.2 Used systems

In all three of NASA's missions described above, the spacecraft used some version of AutoNav, the system used for autonomous optical navigation. To make this work, each would need two central pieces: a camera, with which to obtain said optical information, that is, capture light from the body of interest and other nearby sources; as well as star trackers, or some similar device, capable of observing the stars surrounding the spacecraft and matching the obtained images with databases of stars, providing the spatial orientation of the instruments, and thus of the spacecraft.

In the case of DS1 and Stardust, this was precisely the case. The former had an on-board camera composed of a 1024×1024 pixel array charge-coupled device (CCD), with a 0.77 deg FoV. Each pixel was represented by a 12 bit number, meaning that brightness was discretised as a value between 0 (black) and $2^{12} = 4096$ (white).

---

[4] NASA Space Science Data Coordinated Archive. Deep Impact/EPOXI. Retrieved March 25, 2021 from `https://nssdc.gsfc.nasa.gov/nmc/spacecraft/display.action?id=2005-001A`.

[5] NASA's Science Mission Directorate. Hayabusa 2. Retrieved March 30, 2021 from `https://solarsystem.nasa.gov/missions/hayabusa-2/in-depth/`.

Figure 2.3: Image captured by Deep Impact's main spacecraft just 67 s after it collided with the mission's projectile vehicle. Brightness saturation in the image is caused by light scattering off the ejected material. Credit to NASA/JPL.

Additionally, the camera was fixed to the spacecraft's body, meaning that to re-orient its lens, the whole S/C had to be adjusted [Bhaskaran et al., 2004].

Stardust employed a similar strategy, this time with an FoV of nearly 3.5 deg, and the data number of each pixel compressed to take only 8 bits. On top of that, despite also being fixed to the body of the spacecraft, the light path included an adjustable mirror, spanning an angle of 200 deg, effectively turning the FoV of the camera to 200×3.5 deg [Bhaskaran et al., 2004].

In both cases, star trackers were present aboard the respective spacecraft, yet for distinct reasons they were not used during fly-by. In the case of DS1, as mentioned in the previous section, this instrument ended up becoming damaged and not being functional, forcing the usage of another sensor as a substitute. Stardust, on the other hand, purposefully disabled its trackers in the closest stages of approach, as it was feared that the coma's opacity could obscure the stars, and other particles in the vicinity causing signals that could be mistaken by stars. As such, these were replaced during this phase by the gyros of the inertial measurement unit.

Deep Impact followed a similar trend to its predecessors, using two CCD camera systems named Medium Resolution Imager (MRI) and High Resolution Imager (HRI). They both employed 1024×1024 pixel sensors, with 14 bit digitisation. Then, whereas MRI had an FoV of approximately 0.6 deg, the HRI system had an FoV of nearly 0.1 deg. These two imagers were aboard the main spacecraft, which performed a fly-by of the asteroid. However, the impacter spacecraft, which was detached from it and targets the body itself, had a separate camera system, which was in all metrics equal to MRI [Mastrodemos et al., 2005]. Attached to a common platform, together with the cameras, were two star trackers and an inertial reference unit, for attitude measurements.

Similar to this double-camera setup was Rosetta's own OSIRIS system, which employed a Narrow-angle Camera (NAC) and a Wide-angle Camera (WAC). Whereas the NAC had a higher resolution, allowing images to be obtained from a higher distance and detailed surface studying from its vicinity, the WAC allowed for a wider coverage, which made it possible for it to monitor the entire nucleus from close distances. Both were composed by 2048×2048 pixel CCD sensors, with NAC having an FoV of 2.20×2.22 deg, and WAC having an FoV of 11.35×12.11 deg. Not surprisingly, Rosetta's attitude determination also employed star trackers, just as the previous examples have.

Focusing now on Osiris-REx, this mission used two WAC composed by 2592×1994 pixel CMOS sensors, with a field of view of 44×32 deg. Additionally, two CCD-based with 1024×1024 pixel-sized sensors, one with 0.8 deg and another with 4 deg as their fields of view, were taken to serve as NACs [Golish et al., 2020]. From afar, during the approach phase, the narrow angle imagers were used to perform star-based optical navigation, where an image of both the asteroid and the surrounding star would used to determine the camera's orientation and with it the target's location. Then, once the target was close enough to be resolved, the WACs were used to include both the target and the stars in the measured frames. Finally, when the spacecraft was able to approach Bennu, the navigation turns from star-based to landmark tracking, where a map of the surface was created and constantly updated, and the features of the target used to guide the vehicle [Williams et al., 2018].

JAXA's Hayabusa, on the other hand, took advantage of a star tracker camera, as well as one narrow 5.7 deg FoV and two wider 60 deg FOV cameras, each with 1000×1024 sensors. In a similar fashion to the previous described mission, during approach Hayabusa used the optical and radiometric – from contact with ground – information to navigate. Then, when near the target, lidar was used for better radial information, and the optical tracking uses target markers, released by the spacecraft and shinning light from the surface, as landmarks to guide the vehicle towards the surface [Hashimoto et al., 2010].

Analysing this set of missions, a pattern becomes clear that fly-by missions have, in the past, used a combination of at least 1 or 2 CCD sensor cameras to capture images of the target body, in combination with star trackers, as well as inertial measurement units, in certain cases, for the attitude aspect of GNC. It also shows that using a combination of NACs and WACs is beneficial, as the former is mostly useful when approaching the target, whereas the latter becomes essential in its close proximity.

Additionally, a difference between fly-by and slow-approach missions must be established. Whereas for fly-bys the speed of the spacecraft usually means that the encounter takes little time, requiring on-board autonomous GNC, for slow approaches this is not required, as the operational timelines of the mission allow for the mission's ground segment to process the measurements and command GNC, except for the most critical of manoeuvres, such as landings. This last type of missions also require a broader sensor suite, usually employing lidar or other means of measuring the distance to the target radially, as well as requiring more complex GNC solutions, usually landmark tracking.

2.2 Mission's expected systems

As the Comet Interceptor mission currently stands at the earliest stages of development, a definitive design of the systems aboard the spacecraft is not yet established. Instead, as the analysis evolves and multiple iterations are performed, the performance resulting from the made choices is evaluated and such choices revised. This is precisely the case with the GNC system.

In terms of optical equipment, the Comet Interceptor is equipped with a Comet Camera (CoCa), a camera system whose primary objective is making scientific measurements, but might end up playing a relevant role in GNC. Additionally, the spacecraft possesses a NAVCAM, which is, as the name suggests, responsible for capturing images to be used by the GNC system. The inclusion of distinct systems for science and navigation, which is actually a mission requirement, allows for more versatility, as these distinct objectives require equally distinct camera characteristics, orientations, or even operational timelines. On top of that, since one of the requirements is that the spacecraft must be able to obtain line of sight (LOS) measurements simultaneously with reference star measurements, to determine attitude, star trackers are, unsurprisingly, included [GMV, 2020a].

There is, however, an open trade-off regarding the mounting strategies of the NAVCAM and CoCa cameras. Borrowing the strategy used by mission Stardust, it is possible to mount either camera with the addition of a scanning mirror, allowing for a 180 deg coverage around one axis without exposing the payload to mission compromising dust impacts. However, there is the drawback that, because two separate systems would require two distinct scanning mirrors, with each bearing the considerable mass of 15 kg and equally considerable size, adding these devices would increase mission mass and complexity. On the other hand, not installing such devices means that the cameras are fixed, making it necessary for the whole spacecraft to be oriented to image different regions of space. Doing so during the encounter with the target comet, and within its dusty surroundings, would require additional shielding to protect the vehicle and its payload from incoming impacts, adding mass and cost to the mission.

Faced with this trade-off, three strategies were proposed:

1. Fixed NAVCAM and CoCa – with both systems fixed, the GNC becomes simpler, but at the cost of requiring heavy shielding and payload protection, as well as fast and precise spacecraft attitude adjustments.
2. Fixed NAVCAM and CoCa with scanning mirror – this arrangement simplifies the attitude control requirements, while balancing performance and cost, because it does not require additional shielding. As a result, however, the NAVCAM needs protection, and it eventually loses the target body from its FoV, without spacecraft attitude adjustments.
3. Scanning mirrors for NAVCAM and CoCa – this option also simplifies the attitude control requirements, but at the cost of having to carry two cumbersome mechanisms.

The first and third options have been preliminarily discarded. The former is a result of the increased mass and cost arising from the additional shielding required to adjust the vehicles attitude to point the cameras while keeping it and its payload safe from dust impacts; whereas the latter suffers the same issue due to the scanning mirrors. As such, even though all three options might require testing, the second one will be the main focus of the mission analysis. Nevertheless, it must be kept in mind that the possibility of it being proven unfeasible is a reality and the preliminary choice altered to one of the other two, or even a new option needing to be considered.

After choosing the details of preliminary approach, the actual composition of the NAVCAM system must be sorted to comply with the imposed navigational requirements while respecting the restriction of having its cameras fixed. Since, for the initial portion of the approach, the GNC is ground-based and the main focus of the camera is resolution, thus allowing for an earlier detection of the target comet, a NAC is considered the best choice for the job. On the other hand, during the spacecraft-based GNC encounter phase, FoV becomes more important than resolution, to keep the target in sight for as long as possible, and as such a WAC is chosen.

These are the sensors to be used for navigation purposes during the Comet Interceptor's fly-by. The choice is adequate, since this is an F-class mission, which requires that the used components be cheap and with a high TRL. As became clear from the analysis of the previous section, using a combination of cameras and star trackers on space missions to small astronomical bodies is the most common approach, which means that these components are not only backed by the mission's heritage, but also implies that they are readily available without any need for technological development.

As a result, this analysis will use these as the tools to back the measurements required for the navigation estimation. It is, of course, possible that this is not the best option in terms of performance, but in terms of cost and TRL it seems to be the right direction. As such, other sensor suites could have been chosen, including, for example, some sort of altimeter measurement giving information in the radial direction relative to the target comet, since, as it stands, the measurements provide bearings-only information. However, the study will stick to the baseline chosen for the mission at this point of its development. The discussion of inclusion of other sensors will be left for the thesis work, in case the results indicate that the current setup is unable to withstand reasonable disturbances and uncertainties.

2.3 Mission requirements

Before moving forward, establishing all already existing mission requirements is crucial, as is proposing additional conditions to add to this list. In doing so, it becomes clear what the boundaries are when a solution to the research question is pursued, which facilitates this effort. However, even more important than compiling a list of requirements is understanding the reasoning behind these, and where they stem from given the physical environment in which the mission takes place.

> **MIS – 1.1.** The Comet Interceptor mission shall be designed for a fly-by with a dynamically new comet.

The first mission requirement derives directly from the mission objectives, which aims at studying these types of comets because of their pristine state. For the purposes of this analysis, this requirement implies that the cometary environment will be slightly different from that of a periodic comet. The differences will be analysed in Chapter 3.

> **MIS – 1.2.** Spacecraft A shall be compatible with a nominal closest approach (C/A) of 1000 km.

Note that Spacecraft A refers to the mission's main vessel, as opposed to the two smaller probes released from it. This will be the case for the whole analysis, as spacecraft A is indeed the one of interest for this study.

This requirement establishes the minimum distance from the comet that the S/C achieves during its fly-by nominally, which means that the mission is designed according to this value, but it is expected that a small enough margin of error exists for nothing crucial to change relative to the nominal scenario. To establish such a value, a few factors come into play, the most important of which are likely safety and science. Safety refers to the dangers of passing too close to the comet. For 1000 km, it is not expected that a meaningful risk of collision with the comet itself exists, but as the S/C approaches the nucleus, the density of dust is expected to increase – this will be discussed in the following chapter – and thus not decreasing this distance ensures that the S/C can survive the dust impacts with the desired amount of shielding. The science aspect is related to the amount and quality of the data obtained. Since the mission has explicit science requirements, this distance must be established such that the on-board instruments are able to play their role and fulfil the requirements. As such, making the mission fly closer to the comet is likely to yield better scientific results, at the cost of requiring more shielding to withstand the harsher environment.

For the case of this study, this requirement affects the geometry of the fly-by trajectory, and thus constrains the navigation and observation performances.

> **MIS – 1.3.** Spacecraft A shall be compatible with a B-plane impact vector (see Section 4.1) aligned with the projection of the target-Sun direction onto the plane, with the vehicle passing between the Sun and the target.

With this, the mission must be aimed at flying-by the target while crossing the Sun-target orthogonal projection. This is such that the input power from the vehicle's solar panels can be maximised, as these will have a fixed orientation during the whole encounter. Because one of the sides of the spacecraft must be turned towards the target, such that the sensor suite can take measurements, then the opposite side of the spacecraft should contain the power system with the solar panels pointing towards the Sun, and forcing the spacecraft to cross this line ensures maximum solar exposure given the remaining fly-by restrictions.

> **MIS – 1.4.** Spacecraft A shall be compatible with a B-plane delivery error (see Section 4.1) as high as 150 km (1 sigma).

MIS – 1.4. establishes the margin of error discussed after the previous requirement. This can be useful to comprehend the degree of variation that the nominal trajectory can have, with respect to its passage at the B-plane, and simulate it, to ensure that the developed navigation solution is able to cope with the different scenarios. The reasoning behind this requirement likely stems from the precision of ground-based GNC performed before the encounter starts. This determines the deviation from the nominal during this phase, as well as, again, the same factors as before, science and safety, since allowing this error to be too high can put the S/C or the science objectives at risk.

> **MIS – 1.5.** Spacecraft A shall be capable of performing a fly-by with heliocentric distance between 0.9 and 1.25 au. For the sake of this analysis, the lower bound will be used.

The distance to the Sun set in this requirement is likely due to activity, power or fuel reasons. Activity refers to the material released from the comet's nucleus as it approaches the Sun, and setting an interval for this distance means that the phase of the comet's fly-by of the Sun that the mission aims to study happens at this heliocentric distance. In terms of power, since the S/C will be solar-powered, establishing these bounds makes it so that the mission lies within a spherical section with adequate energy production relying on the on-board solar cells. Finally, the bounds are also likely related to the amount of fuel that the spacecraft can carry and use to make its way from SEL2 to the fly-by location. Allowing this distance to deviate too much from the thresholds, and thus from 1 au, could require more fuel than the mission has available.

For the purposes of the study, this distance to the Sun mostly influences its gravitational influence and the magnitude of the solar radiation perturbation upon Comet Interceptor. Since these effects are more intense near the Sun, as well as other phenomena such as outbursts, which are described in Chapter 3, the lower bound – 0.9 au – is used for the nominal trajectory, to take these into account where their impacts are the strongest.

> **MIS – 1.6.** Spacecraft A shall be capable of performing a fly-by with relative speed to the comet between 10 and 70 km/s.

Restricting the fly-by speed relative to the target body also restricts the catalogue of bodies at the mission's disposal to be chosen as targets. As the mission will launch, and then the target body will be discovered and chosen, the data available on this is in the form of the typical characteristics of the DNC population. Thus, restricting the fly-by will also exclude a group of bodies that do not abide by the restriction, thus lowering the probability that a suitable body is found within the expected time of the mission. On the other hand, allowing even higher fly-by speeds will obviously affect the mission's design, with the S/C needing significantly more shielding from dust impacts. A balance between these two factors results in this requirement.

Relative to the study of encounter navigation, the fly-by speed varying this much will mean that the developed solution must be more robust to account for the different scenarios.

> **MIS – 1.7.** Spacecraft A shall be capable of performing a fly-by with a solar aspect angle between 45 and 135 deg.

As for the previous requirement, **MIS – 1.7.** restricts the range of comets available to be chosen at targets for the Comet Interceptor mission. However, this is essential to guarantee that the spacecraft has appropriate visibility conditions to perform GNC and science operations. For the purposes of navigation, the fly-by solar aspect angle will determine the position of the spacecraft relative to the illuminated side of the target comet, which obviously influences the measurements and the way that the captured information is used to compute the centre of the body with the image processing algorithm. This angle is defined as the one between the fly-by relative velocity vector and the target-Sun vector at the time of encounter.

> **MIS – 1.8.** The encounter phase of spacecraft A shall start immediately after the last trajectory correction manoeuvre (TCM) is performed.

This requirement establishes when the encounter phase actually begins, which determines when the phase of interest for this study begins. As it currently stands, GMV proposes that the last TCM be performed approximately 20 hours before closest approach, so this is effectively when the developed autonomous system on-board the S/C assumes control of navigation.

> **MIS – 1.9.** Spacecraft A shall be capable of autonomously controlling its attitude, basing this on on-board processing of images obtained from the NAVCAM system, ensuring that the instruments keep track of the target during the encounter phase.

This requirement is central for the case of this study. By requiring the S/C to autonomously control its attitude with respect to the comet, an equally autonomous navigation system must be developed to determine the relative position with respect to the target and adjust the attitude accordingly. The reason to rely on autonomous on-board GNC stems from the fact that the planned fly-bys will have very high relative velocities. This implies that the GNC operations do not have time to rely on the ground segment to be performed, as the S/C can require angular velocities of up to a few degrees per second to keep the target in frame of the CoCa system near the C/A [GMV, 2020a].

All this leads to the next requirement:

> **MIS – 1.10.** The spacecraft A shall allow the CoCa instrument to continuously observe the target before, during, and after the closest approach.

Such a requirement is obvious, as the CoCa instrument is the one responsible for achieving S/C A's science objectives, namely by taking high resolution images of the comet's nucleus. For that to happen, the navigation system must be accurate enough to ensure that the attitude of the S/C or/and of the scanning mirror – in case this ends up being used, as has been preliminarily determined – keeps the target within the instrument's FoV, which is 0.69×0.92 deg, or 180×0.92 deg, with the addition of the scanning mirror.

2.4 System requirements

Having established the relevant mission requirements, as a representation of the operational restrictions that must be followed given the very nature of the mission itself and its objectives, the focus is now turned toward the system requirements. These constrict how the systems are designed such that the mission's objectives are met.

2.4.1 GNC requirements

> **GNC – 1.1.** The navigation subsystem shall be capable of acquiring images of the target with the dedicated navigation camera system.
>
> **GNC – 1.2.** The navigation subsystem shall be able to estimate the impact vector in the B-plane by processing the obtained images from the navigation camera system.

These two requirements targeting the navigation subsystem of Comet Interceptor are necessary to ensure that **MIS – 1.9.** can be fulfilled. To autonomously navigate relying on the processing of images taken from the on-board cameras, the S/C must have such an imaging system and must be able to perform these computations on the acquired data.

> **GNC – 1.3.** The navigation subsystem shall be able to provide updated estimations of the fly-by conditions – impact vector and time of closest approach – to the remaining payload during the encounter phase.

The above requirement forces the navigation subsystem, in particular its navigation estimation part, which is the concern of this study, to work continuously during the encounter phase to provide estimations of the fly-by characteristics with ever increasing precision. As such, the way that the filtering is implemented will be restricted by this, since acquiring a big batch of data over the encounter to then use all the information in one go is out of question. The strategy to implement must then be somewhat sequential, using the information captured by NAVCAM and subsequently processed to update the estimate of the impact vector and time of C/A as the S/C closes its distance to target and these estimates become clearer. An alternative is using a system of smaller batches, which is updated sequentially as time passes. These details will be further analysed in Section 5.2.

This is crucial for the mission to succeed, because, as the vehicle approaches, the estimate of these parameters becomes increasingly more accurate, whereas, from afar, the comet will barely be visible, thus making the measurements more prone to uncertainties. The system should then take advantage of this and use the data as it is acquired to always give the payload the most up to date estimates possible, thus ensuring that they are kept pointing in the right direction.

GNC – 1.4. The GNC autonomous navigation function shall be compatible with a comet nucleus size of 10 km

GNC – 1.5. The GNC autonomous navigation function shall be compatible with a comet nucleus density of 500 kg/m^3.

Establishing the comet's nucleus size, which will be the element that must be imaged to determine the spacecraft's location along its fly-by path, is obviously important to perform an analysis of the navigation system's performance, as the comet's diameter determines the size of the nucleus as captured by the camera. The value chosen is a reference for comet sizes, as is discussed further ahead in Section 3.1, given that nucleus typically vary from 1 to 50 km in diameter.

The comet's density allows for, in combination with the nucleus' size, an estimation of its mass to be made. This will be useful when computing the comet's gravity influence during the fly-by. The value chosen stems from the work of [Weissman and Lowry, 2008].

2.4.2 NAVCAM requirements

This section is dedicated to establishing the requirements for the cameras used in the optical navigation performed by the Comet Interceptor mission. These requirements must balance, on one hand, the goals of the mission, which will require certain characteristics from the cameras to be satisfied, and the cost of the cameras and their physical limitations, on the other hand, as these cameras not only need to have a high enough TRL to be used in this fast-track mission, but also need to take into account its budget. As such, the main requirements which define the cameras are the size of the sensors, their field of view, as well as the range of possible exposure times available. With this information, a model of the cameras can be established to simulate their real behaviour, and thus their presence integrated into the simulation to be performed during the final chapters of this document.

However, due to confidentiality reasons, these characterising values for the navigation cameras have been removed from this public version.

Chapter 3

Cometary environment

As Comet Interceptor aims at flying-by a comet, understanding the environment surrounding the mission's target is crucial. Knowing this will help list the voids in cometary knowledge, which the mission wishes to fill, stated as the science goals, but it will also give the mission analysis process the best available information related to what the spacecraft must be designed to endure, as interplanetary space can be a harsh environment.

In the specific case of this study, to develop an autonomous navigation system, the way that the environment affects the spacecraft's dynamics and the measurements taken of the comet in the form of photographs must be understood to simulate the scenarios and assess whether the navigation performance is up to par with the mission's requirements.

Because of this necessity, the next sections contain an overview of the general structure around a comet, followed by an analysis of the relevant zones as seen from the mission analysis point of view. Finally, a description of the possible phenomena surrounding the comet's nucleus that can affect the image capture and correct processing of the centre of the body is made, with the goal of adding these to the list of mishaps relative to which the navigation system must be robust. In that regard, outbursts, cosmic rays, and outgassing are the discussed phenomena.

3.1 General model

A comet's core is its nucleus, a chunk of material typically ranging from 1 to 50 km in diameter – despite not being necessarily spherical, as can be seen in Figure 3.1 [Lissauer and de Pater, 2013]. Around it is the coma, a cloud made up of gas and dust, spanning up to 10^4 – 10^5 km in diameter, which usually obscures the nucleus itself. Even further around the coma is the invisible hydrogen coma, which as the name implies is mostly made up of hydrogen, and can have a size up to 10^6 – 10^7 km in diameter. From the centre, and in the direction opposite to that of the Sun, two tails can form, a dust tail and an ion tail, the former typically curved and yellowish, and the latter straight and blue. This can be seen in the schematic in Figure 3.2.

The coma and the tails are not a constant presence around the comets. Instead, they are formed once their orbits bring them closer to the Sun, where its radiation sublimates gases from the surface of the nucleus. As these gases evolve off from the surface, dust from the surface follows them, thus forming these surrounding structures. In fact, since the nuclei themselves are so small, comets can only be seen with the naked eye or small optical aid when these form, and they are indeed what is visible in the first place. The majority of the gases ejected from the nucleus are exposed to the sunlight, suffering photodissociation and photoionisation over time, forming ever simpler radicals and ions, with the hydrogen eventually ending up forming the aforementioned hydrogen coma.

The dust that is ejected with the gas varies in size, from a few centimetres down to the sub-micrometre scale. The smallest particles are capable of reaching near-escape velocities of about 1 km/s, whereas the larger, more massive, grains barely gain any speed at all from this ejection, in the order of magnitude of 1 m/s. The pressure from the sublimating gases keep these bigger particles from falling back to the surface, but as the comet recedes from the Sun, they collapse and form a dust crust. Upon returning to the inner Solar System, the Sun's influence restarts the sublimation process, which is eventually able to build enough pressure to blow off a part of this crust, exposing the surface once again in a spike of cometary activity typical of this type of old, periodic comet. Over multiple passages, these comets can have their crusts become so thick that eventually no coma is able to form, even when near their perihelia, becoming extinct [Lissauer and de Pater, 2013].

In the case of dynamically new comets, however, their surfaces have never been exposed to a star as they enter the Solar System for the first time. As such, their surfaces are rich in highly volatile species which have never been sublimated, as well as other unstable species formed in the period since the formation of the Solar System by irradiation by galactic cosmic rays. These two factors leads to these types of comets being comparatively much more active at higher distances from the Sun [Lissauer and de Pater, 2013].

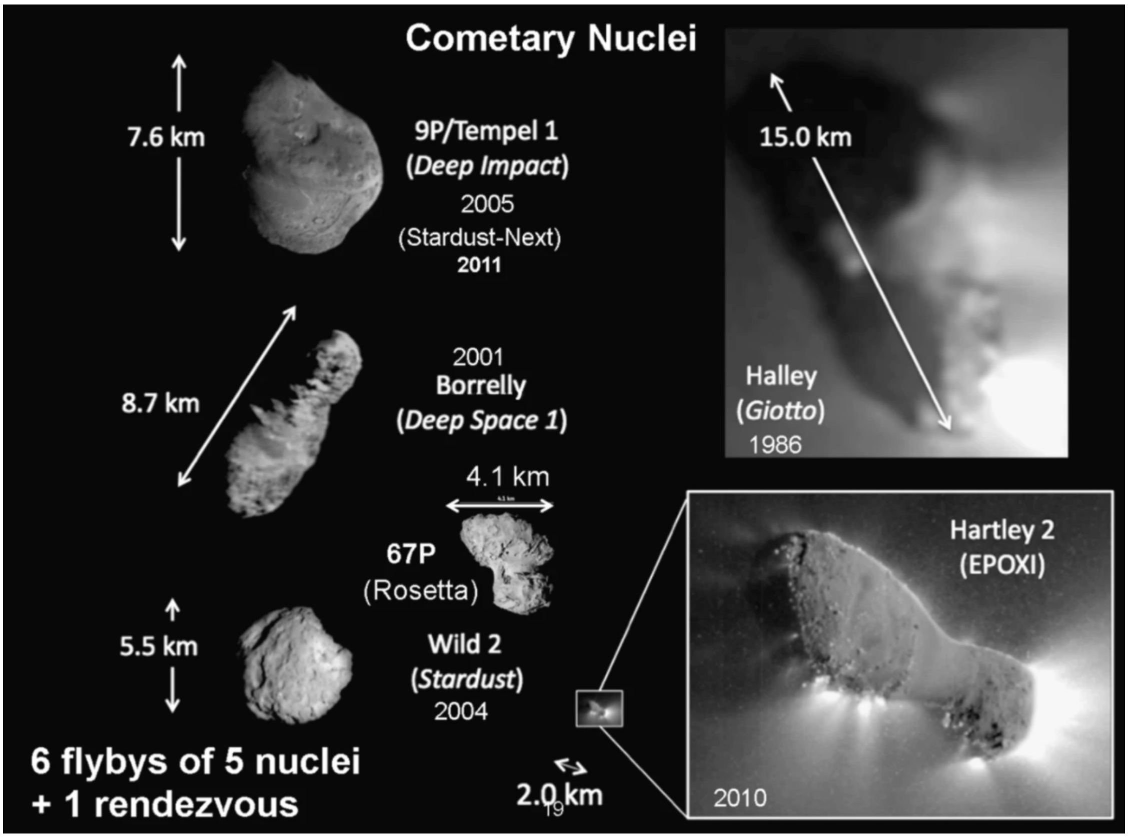

Figure 3.1: Cometary nuclei imaged by a multitude of previous space missions. Retrieved from [Keller and Kührt, 2020].

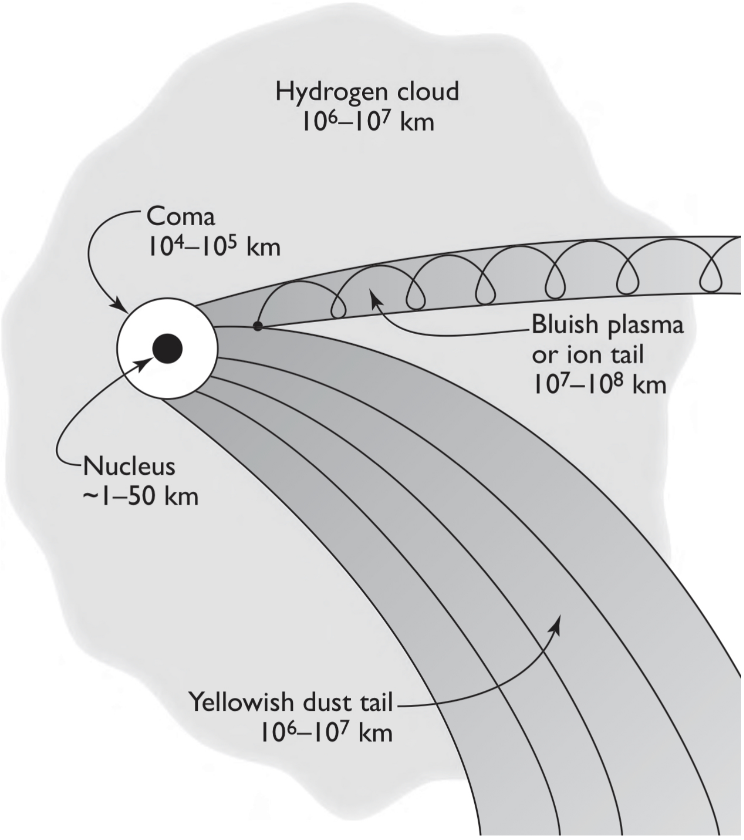

Figure 3.2: Schematic representing the various parts of a comet's surroundings. Retrieved from [Lissauer and de Pater, 2013].

3.2 Regions of interest

From the above description, the cometary environment can now be understood with all its distinct parts and their characteristics. However, for the purposes of the mission, not all regions are equally important, and, in fact, the main focus will fall on the coma and the nucleus. As the spacecraft flies-by, it is obviously desirable that the trajectory falls on the sunward side of the comet, such that the imaging instruments can perform their tasks. And by keeping the mission on this side of the target, the dust and plasma tails will not be relevant factors.

Thus, from a more operational point of view, the comet's structure can be divided into:

- Coma – defined as the region that extends up to the reference threshold of 10^5 km from the nucleus, the coma is, from an operational stand-point, the region of the fly-by where all in-situ experiments can operate continuously.
- Dust cloud – located inside the coma itself, in this zone no unprotected equipment should be exposed to the dust flow. To be unprotected means that the payload is outside the footprint of the dust shield, and thus the spacecraft's attitude should remain adjusted such that the shield can protect all instruments from the incoming dust flow. Its size is dependant on the comet's activity, the fly-by speed, and the vehicle's shielding. It is defined by the flux of dust particles (#/m^2) of a certain energy – which depends on particle mass and fly-by speed – surpassing a certain threshold, defined to reflect a danger to the payload.
- Central fly-by region – this is the section of the trajectory where the angle β between the relative velocity vector and the LoS direction is in the range 5 deg < β < 175 deg. This marks the transition between a nearly stationary LoS direction to a rapidly changing one.
- Detailed imaging region – defined as the zone where the imaging instruments are able to resolve the nucleus of the target comet, that is, take images where it occupies more than 2 pixels on the sensor. Since the fly-by C/A distance is defined, this depends mostly on the characteristics of the chosen imaging system, but on the size of comet as well.

The autonomous navigation system will be designed to function during the encounter phase, spanning from 20 hours before the closest approach up until 20 hours after that. In a real mission scenario, this time period would be determined by the best available spacecraft's state estimation, considering its distance to the comet and the speed relative to it. In the case of this study, however, 20 hours prior to the C/A is determined with respect to the nominal trajectory, since it is with respect to this objective that the real trajectory is then generated with some degree of random error. Considering, again as an approximation, a nearly straight trajectory with constant speed, due to the short period in question and the typically low mass of comets, the length of the encounter can be computed to be

$$d = v \cdot t \iff d = 14.4 \cdot 10^5 \text{ km (if v = 10 km/s)} \vee d = 100.8 \cdot 10^5 \text{ km (if v = 70 km/s)} \tag{3.1}$$

From the rough calculation above, it is plainly obvious that for all the possible fly-by velocities, the encounter phase spans all the above regions, starting and ending well outside the coma. Nevertheless, during the whole span of this phase, the spacecraft flies with its shielded face in the direction of incoming dust, with the only attitude adjustments allowed being around the direction of the velocity, so as not to compromise the shield's protection of the payload. Additionally, the central fly-by region, where, due to the increasingly fast rotation speed required to keep track of the comet, the autonomous navigation's performance will be stressed the most, only takes

$$t = d/v = 2 \cdot (d_{C/A}/\tan\beta)/v \iff t \approx 38.1 \text{ min (if v = 10 km/s)} \vee t \approx 5.4 \text{ min (if v = 70 km/s)} \tag{3.2}$$

The figures computed above demonstrate how the problem at hand is really concentrated in the minutes preceding the closest approach. Even more so considering that the information obtained from afar is not as precise in determining the time of C/A and the B-plane vector as that captured just before the C/A, which further emphasises the absolute necessity of designing a robust navigation solution.

3.3 Possible disruptive events

3.3.1 Outbursts

With the central focus of this study being the robustness of the developed navigation system, as mentioned previously, one of the main points of interest regarding the cometary environment must be the disruptive events

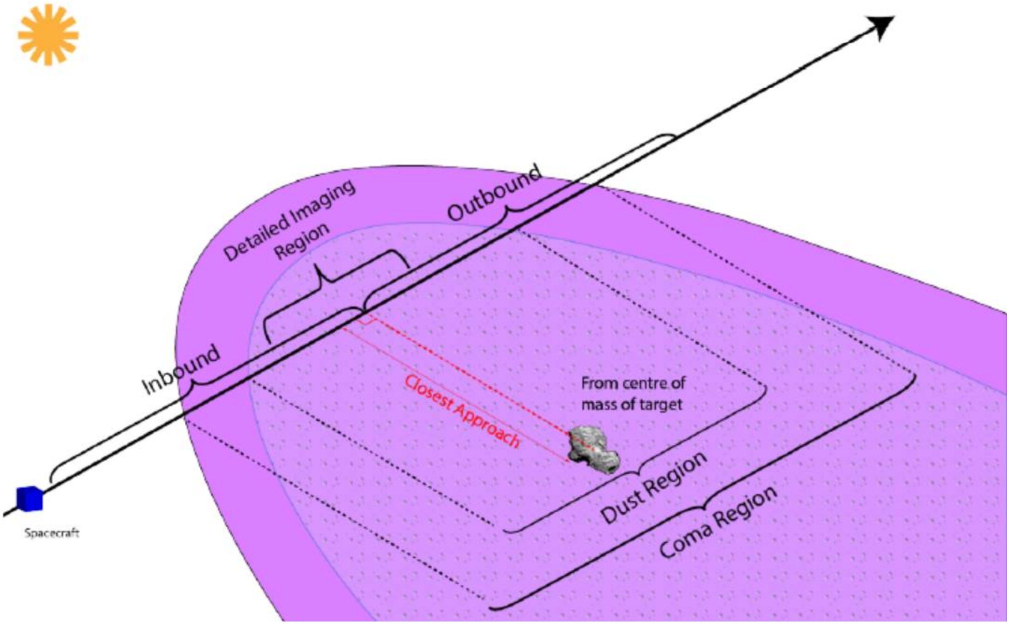

Figure 3.3: Near-comet regions of interest for the fly-by. Retrieved from [GMV, 2020a].

that take place. This refers to any temporary and out of the ordinary episode of cometary activity, which might, nonetheless, be expected to happen during a fly-by. Thus, their behaviour needs to be modelled to assess the respective impact on the acquired optical data.

To take such phenomena into account, available data from previous missions is extremely valuable. Using images acquired from these past endeavours, it is possible to understand the nature, frequency and impact of these events, always considering, of course, that the mission conditions are not the same, and thus some adaptation might be required to establish an adequate comparison.

Firstly, it is necessary to understand the main processes that cause ejections of mass from cometary nuclei. Either gas gently sublimates from the surface of the nucleus dragging the dust along or it is outgassed from cavities present under the surface [Wesołowski et al., 2020]. There are other proposed mechanisms for this, but they are either marginal in terms of mass ejected, or not yet very well established, so these remain as the most consensual means.

The Rosetta mission was crucial to gain information on such phenomena as, unlike most other missions to comets, Rosetta actually had the objective of remaining around the nucleus for a long time, exploring the diverse environment of the target and taking pictures of it for a span of months. Regardless of the process, the observed results lie in two categories. These can manifest as long-lasting collimated jets, lasting for up to several hours, but being relatively faint, presenting approximately 10% of the nucleus' brightness. These are also repetitive, being dependent on the energy emitted by the Sun, which means that they evolve over time in a predictable manner and show the same features daily [Vincent et al., 2016]. However, due to the small optical impact, these can hardly be considered as disruptive.

Outbursts, on the other hand, represent short and relatively unpredictable bursts of gas and dust. Lasting up to a few minutes – this is hard to pin down further due to the low frequency of image captures for Rosetta – these phenomena can achieve a brightness comparable to that of the nucleus itself. As a result, these should definitely be taken into account when simulating the comet's environment, as the sudden bursts can change the centre of brightness of the image quite suddenly and significantly, and thus disrupt the navigation.

There are a few important remarks to make about the observations presented by [Vincent et al., 2016]. Firstly, the duration of the events can only be said to be of a couple of minutes due to the capture of multiple images of one particular outburst, which happened to last this long. However, in this specific case the spacecraft snapped pictures with a frequency much higher than normal – 5 to 30 minutes per image – which means that this duration can vary between events. Despite this, the authors mention that these are similar in time terms to outbursts observed in other comets. Thus, this will be the order of magnitude considered for the duration when simulating these phenomena.

Secondly, this data was obtained from comet 67P/Churyumov–Gerasimenko, explored by the Rosetta mission. This comet is not a DNC, which means that it frequently visits the inner regions of the Solar System

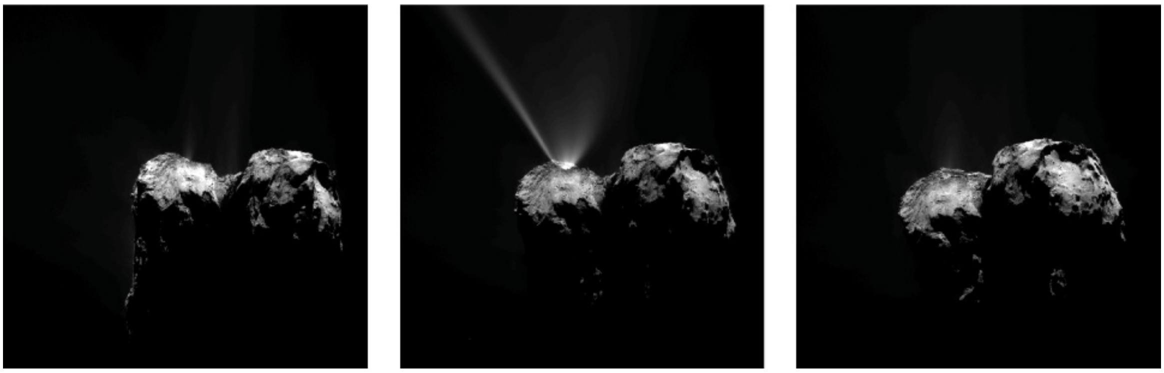

Figure 3.4: Images captured by Rosetta showing a outburst event taking place on comet 67P, during perihelion. These are separated by 1/2 h each and were taken from 332 km from the nucleus. Retrieved from [Vincent et al., 2016].

and never, in fact, gets too far from the Sun, with its aphelion lying just outside the orbit of Jupiter[1]. Given that the target of Comet Interceptor will be – ideally – a DNC, characterised by never having flown-by the Solar System, then it is fully possible that the differences between the two can result in the outbursts being manifested in significantly different ways, be it in temporal or optical terms. Such differences could arise from DNCs being covered by highly volatile species, as mentioned in a previous section. This results in an increase in activity at higher distances from the Sun, and likely means that the target of Comet Interceptor will show brighter and more frequent activity in terms of outbursts, given that 67P's perihelion, which coincided with the study of Rosetta, is located at about 1.2 au from the Sun, lying just near the upper bound for the established interval for heliocentric distances during fly-by. As a result of all this, running simulations to ensure the robustness of the navigation system should take these possible discrepancies in mind, and some scenarios with highly exacerbated activity should be tested to determine how the performance scales with bigger, more frequent, disruptions, which might be possible due to the difference in the target's nature.

Having discussed all the limitations of the information from Rosetta, its characteristics should now be discussed. Regarding the morphology of the emitted outbursts, these can be narrow jets, broad jets, or a combination of the two. Figure 3.4 shows the third case, where on the left a part of the outburst is very thin and nearly parallel, with slowly decaying brightness, whereas the rest presents a wider angle, with the centre of brightness much closer to the origin of the material. The observations showed that no particular option is more probable than the others.

In terms of frequency, the outbursts appear to happen about once every 2.4 nucleus rotations, which is in accordance with data from the Deep-Impact mission, reporting such events approximately every 2 rotations. Finally, the time of day when these events occurred was not random, as, again, they are powered by the Sun's incoming light. The data indicates that the majority of recorded outbursts lied either at Sunrise or around noon. The former is likely a result of the thermal stresses caused by the sudden absorption of solar energy as the Sun rises and rapid temperature changes, whereas the latter may be related to the heat received from the Sun reaching volatiles buried deeper into the nucleus.

Unfortunately, no complete models were found for simulating the physical aspects of the outbursts. Most of the available data was obtained by the Rosetta mission, and only a few tens of photographs of such phenomena are available. From that, solely a few equations describing some characteristics of the outbursts exist, but not any well established model. Given this, it seems like it would be a worth-while endeavour to do so, and use what available information from cometary observations to create a full model. From a scientific standpoint there is no doubt of the value of such a job. Taking the current study and its goals in mind, however, this is not the most practical nor useful approach. This is because of the image processing algorithm, which will be discussed in deeper detail in Section 8.5.3. Its work is done between an image of the nucleus being taken and using this information to know where it is with respect to the spacecraft. Its job is to autonomously analyse the image taken by the navigation camera and determine where in the light information captured the nucleus' light is. Only then, with this information and knowledge of the spacecraft's orientation in space provided by other instruments, can information about the direction to the comet be assessed and used by the state estimators.

Because of this disconnect between raw image data and its consequences on the simulation, a realistic physical model of the outbursts would be of little use in simulating their impact on the Comet Interceptor's mission, since it still would not explain the exact impact on the measurements processed by the algorithm. For that, either a simplified model is created for how these two entities interact, or outburst images are generated

---

[1] European Space Agency. Rosetta's Target: Comet 67P/Churyumov. Retrieved April 12, 2021 from `https://sci.esa.int/web/rosetta/-/14615-comet-67p`.

and fed to the algorithm as it is to analyse the results. In the current case, where the latter option is not available, the former will have to be chosen. The suggested approach is then to assume that an outburst causes a bias in the processed nucleus position with respect to the unperturbed value. This will be the framework used to introduce these phenomena into the simulations performed in this thesis, and the mathematical details of how this is done will be discussed in Chapter 8.

3.3.2 Cosmic rays

Cosmic rays are subatomic particles – mainly protons, but also electrons, neutrons, as well as others – travelling through space at relativistic speeds[2]. It is a well known phenomenon that these interact with the Earth's atmosphere creating a cascade of particles that can sometimes cause artefacts in camera sensors. However, this is also a very real possibility in space applications, as demonstrated by the data gathered by Hubble's Wide Field and Planetary Camera 2 [Space Telescope Science Institute, 2004]. This instrument, using a CCD sensor just like the NAVCAM system will, detects multiple occurrences of cosmic rays, and an attempt is made in the cited handbook to characterise the frequency and intensity of these events.

This study determines that the majority of cosmic rays create an image of 4 pixels in the sensor, with the vast majority of the cases not going above 9 pixels. This means that the effect of cosmic rays will most likely be felt before the spacecraft gets too close, when the size of the comet is thus similar to one of these events. However, the data obtained regarding brightness and frequency of cosmic rays affecting the measurements is relative to the characteristics of Hubble's sensor, making it hard to draw comparisons with the Comet Interceptor mission.

Apart from this, there was not much information readily available regarding this subject, and especially based on previous space missions. However, there are some reasons to consider the impact of these phenomena to be secondary to other more permanent ones such as outbursts. First of all, their very sporadic nature makes them less impactful, as the mission's phase is comprised of 40 hours, during which more than 1000 images of the comet's nucleus are taken. Thus, cosmic rays momentarily affecting some of these pictures should not see a relevant change in results, unless a significant number of images became affected, which, despite the lack of data regarding the frequency of such events, seems to be unreasonable.

Secondly, as will be discussed further ahead in this study, the only relevant images taken from the comet are the ones corresponding to the phase of the encounter when the spacecraft is near the closest approach, in the previous 1 hour period, approximately. Given this relatively small window, any cosmic rays affecting measurements taken before just do not matter as the images at those times do not in the first place. This does mean that the previous argument regarding the number of images being large must be taken with a grain of salt, as the number of relevant measurements is actually smaller than the whole, and their importance is increasingly higher as the spacecraft approaches the nucleus. Nevertheless, in this crucial period of the mission, the frequency of measurements is also increased, which should counteract this effect.

Additionally, all the above reasons did not take into account a crucial step of the data acquisition, which is the image processing algorithm, which was discussed above. This means that analysing the response of the mission's navigation to the presence of cosmic rays is a much more complex task than anticipated without a model of how these two agents interact. It might be the case that some rays cause a visible impact on the acquired images of the nucleus, but the image processing algorithm is able to discern the true position of the nucleus regardless, meaning that it ends up having no impact whatsoever. On the other side of the spectrum, a worst case would be a ray causing an artefact in the image which mimics the comet's signature, thus tricking the algorithm into picking an absolutely wrong direction to the nucleus. The bottom line, however, is that without further modelling of these phenomena and, more importantly, of the image processing algorithm itself, which is wildly out of the scope of this thesis, the study of their interaction cannot be done.

Finally, despite all of the arguments and possibilities put forward so far, there is a last line of defence of the navigation against these events. Also discussed ahead, in Section 5.4, are the outlier rejection techniques, which allow the navigation to put aside images that stand out from the expected given the best state estimation at the time. Then it stands to reason that would the cosmic rays be negatively impacting the measurements taken by the NAVCAM, these algorithms would also be able to detect this impact and reject the image if it is deemed too large.

As a result of all the state reasons, the study of the impact of these phenomena will not be performed in this thesis, but a more careful analysis, especially regarding their interaction with the image processing, is suggested for future mission analyses.

---

[2] NASA - Goddard Space Flight Center. Cosmic rays. Retrieved May 10, 2021 from `https://imagine.gsfc.nasa.gov/science/toolbox/cosmic_rays1.html`.

3.3.3 Outgassing

Not so much an event as a constant perturbation to keep in mind while studying comets, outgassing, that is, the sublimation of volatile substances from a comet's nucleus, can be a source of acceleration capable of slowly changing the target's trajectory. If this happens uniformly around the whole body, then the net force is zero. However, an uneven sublimation can cause the comet to drift over time. With a non-uniform surface and only one side being illuminated by the Sun, it is expectable that this phenomenon does, in fact, happen differently at different parts of the comet. Additionally, this outgassing is likely highly influenced by the characteristics of the target's surface. Given that a cometary surface can deviate from a perfectly symmetrical spherical shape quite significantly, this is another factor to take into account.

Whereas the previous two perturbations affected the mission more directly, by possibly corrupting the navigation images captures by the NAVCAM, the presence of outgassing is a more tenuous effect, changing the comet's expected state over time, influencing the dynamics of the system instead. As such, the description of this will be left for the next chapter, where the dynamics are analysed, and there the force's model will be discussed, as well as how this contribution will be considered.

3.4 Environment model requirements

As a result of the discussion in this chapter, environment model requirements are proposed as a way of summarising the main characteristics of the cometary environment. Regarding the nucleus' size and density, **GNC – 1.4** and **GNC – 1.5** already establish these parameters, so that the target's mass is defined. In the case of the comet's gravity, the dust impact forces, and the outgassing, all their impacts will be addressed in the following chapter, along with the remaining relevant forces. And since the cosmic rays will not be included, that leaves the outbursts to be characterised. These will, in fact, be the main disruptive events considered in the analysis performed later in the thesis. As such, these events must obey

ENV – 1.1. The Comet Interceptor mission shall be designed considering the presence of outbursts.

ENV – 1.2. The outbursts shall be defined by a given probability of happening, time of origin, duration, orientation in space, and intensity.

By establishing the first clause, the study ahead must analyse how these outbursts affect the performance. The way that they are defined was already discussed in Section 3.3.1, and how they are modelled in the performed simulations will be the topic of discussion further ahead during Section 10.3. Despite this, the parameters used to define these events are already enumerated in **ENV – 1.2**, meaning that all that is left is to do is choose adequate values for each parameter.

Chapter 4

Flight Dynamics

The flight dynamics chapter is dedicated to analysing the forces that are exerted on the spacecraft and thus determine its trajectory as it flies-by the target comet. For that, firstly a discussion on the relevant reference frames is required. Then, all the forces are analysed and filtered to determine whether they are required to simulate the scenario or can be discarded. Regarding this, the forces of gravity of the comet, Sun, and the various planets in the Solar System are considered. Afterwards, other perturbing forces are also discussed, including the solar radiation pressure, cometary dust impacts, and accelerations resulting from outgassing events. Following this, a brief discussion regarding the necessity of simulating the attitude of the vehicle is had, before the choice of propagator to be used is analysed and made.

4.1 Reference frames

The reference frames are the sets of origin points together with 3 oriented axes relative to which variables such as the state or the measurements are expressed. Defining these early on will make it easier to describe how such quantities are defined later, and will even help to characterise the central body.

The main reference frame will be Ecliptic J2000 centred on the comet, abbreviated as CometEclipticJ2K, meaning that the centre will always be the comet and its axes will be oriented in the same directions as Ecliptic J2000. Thus, $\boldsymbol{ref_x}$ points towards the mean vernal equinox at the J2000 epoch, $\boldsymbol{ref_z}$ is perpendicular to the Earth's ecliptic plane, and $\boldsymbol{ref_y} = \boldsymbol{ref_z} \times \boldsymbol{ref_x}$. This choice is due to the fact that most ephemerides of Solar System bodies are expressed in J2000 reference frames, thus facilitating the usage of these values for the scenario at hand. Centring the frame at the comet implies that the central body will be the comet, as, relative to it, the spacecraft will suffer the largest changes in state in absolute terms, helping reduce numerical errors during the computations.

The B-plane and the associated reference frame are useful in evaluating the fly-by's targeting accuracy. The plane is defined by passing though the target while being perpendicular to the asymptote defined by the initial relative fly-by velocity. In the case of Comet Interceptor, since the spacecraft performs a fly-by of the body under relatively low central gravity, as will be discussed in the next sections, the trajectory is well approximated by a straight line,, making the B-plane approximately contain the C/A point.

The B-plane reference frame is formed by the RST triad, where $\hat{\boldsymbol{S}}$ is perpendicular to the plane in the direction of the fly-by relative velocity, and $\hat{\boldsymbol{T}}$ is contained in the plane and points in the opposite direction of the orthogonal projection of the Sun onto the plane. Finally, $\hat{\boldsymbol{R}} = \hat{\boldsymbol{S}} \times \hat{\boldsymbol{T}}$ completes the right-handed triad. This can be seen in Figure 4.1.

Also on the figure, the intersection of the vehicle's trajectory with the B-plane is marked in red. The vector from the target to this point is called the impact vector of the trajectory, and its norm the impact parameter. Using these tools, requirements **MIS – 1.2.** and **MIS – 1.3.** can be characterised. Because of the aforementioned similarity of the encounter trajectory with a straight line, the C/A point should approximately coincide with the

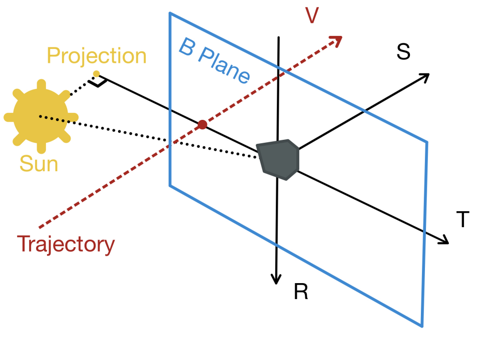

Figure 4.1: Representation of the B-plane, and the RST triad defining the respective reference frame.

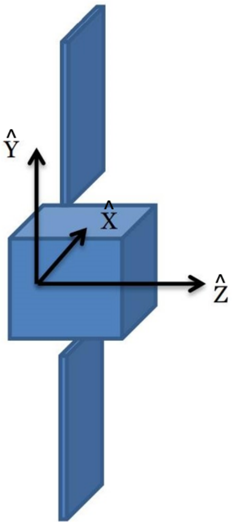

Figure 4.2: Reference frame of the vehicle. Retrieved from [ESA - Concurrent Design Facility, 2019].

vehicle's crossing of the B-plane, thus being equal to the tip of the impact vector. Additionally, the B-plane delivery error refers to the difference from the nominal impact parameter. Thus, the objective will be to aim the spacecraft such that the closest approach lies between 850 km and 1150 km from the target.

The previous two reference frames are trajectory-oriented tools for the design of the Comet Interceptor scenario. The next two will be more useful in describing the measurements. The first is the spacecraft's body frame, which is oriented with its axes perpendicular to the panels, as shown in Figure 4.2. The $\hat{\boldsymbol{Y}}$ axis is aligned with the solar panels, $\hat{\boldsymbol{Z}}$ is aligned with the fly-by's relative velocity vector, and $\hat{\boldsymbol{X}}$ is such that $\hat{\boldsymbol{X}} \times \hat{\boldsymbol{Y}} = \hat{\boldsymbol{Z}}$ [ESA - Concurrent Design Facility, 2019].

Since the trajectory is nearly straight, as will be discussed in Section 4.2, then the rotation between the B-plane and body reference frames is simple, and given by

$$\boldsymbol{R_{Body2BPlane}} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{4.1}$$

From this point, all that is left is to define the reference frame of the sensors, starting with the WAC, which will take the navigation measurements. This is not yet set in permanently defined, but the current preliminary choice is to have the WAC oriented slightly towards the side of the spacecraft facing the comet, to increase its coverage near the C/A zone, instead of pointing in the direction perpendicular to the ram face. Thus, it can be assumed that, using an FoV of 50 deg, according to **CAM – 1.4.**, this camera should then be rotated by 50/2 - 0.5 degrees, to allow the camera to still capture the direction of velocity with a margin of 0.5 deg [GMV, 2020b]. Given this description, the rotation matrix between the body and the WAC frames is the following

$$\boldsymbol{R_{Body2WAC}} = \begin{bmatrix} \cos\left(FoV/2 - 0.5\frac{\pi}{180}\right) & 0 & -\sin\left(FoV/2 - 0.5\frac{\pi}{180}\right) \\ 0 & 1 & 0 \\ \sin\left(FoV/2 - 0.5\frac{\pi}{180}\right) & 0 & \cos\left(FoV/2 - 0.5\frac{\pi}{180}\right) \end{bmatrix} \tag{4.2}$$

Another sensor to consider is the CoCa, which is located in the facet of the spacecraft perpendicular to $\hat{\boldsymbol{Z}}$. Given that it is mounted with a scanning mirror, which can rotate along the $\hat{\boldsymbol{Y}}$ axis with a angle η going from -90 deg to 90 deg. Thus, the rotation matrix for the reference frame of the CoCa is given by

$$\boldsymbol{R_{Body2CoCa}} = \begin{bmatrix} \cos\left(\frac{\pi}{2} + \eta\right) & 0 & -\sin\left(\frac{\pi}{2} + \eta\right) \\ 0 & 1 & 0 \\ \sin\left(\frac{\pi}{2} + \eta\right) & 0 & \cos\left(\frac{\pi}{2} + \eta\right) \end{bmatrix} \tag{4.3}$$

The reference frames of the star trackers are missing from this analysis, since they will not be necessary. As discussed further ahead in Chapter 7, these sensors are not simulated exactly as they function. Instead, since the order of magnitude of their accuracy is known, their error contribution is simply added to the true system attitude to obtain the measured attitude, meaning that their orientation is not required for the purposes of this study.

4.2 Dynamics

To properly simulate the environment of the Comet Interceptor's encounter phase, a discussion must be had regarding the relevant forces acting on the spacecraft. This means that, for all the external factors acting on the spacecraft, these must be analysed in order to determine whether it is worthwhile to consider their impact on the simulations. Despite this, the approach while simulating this mission will always be to include the forces and contributions related to the comet and its environment, regardless of their impact. As the model of the environment is iterated along the mission's development, this allows for the contributions' parameters or models to be continuously updated over time and their effects evaluated. While it is not the most efficient tactic computationally, this helps the mission design process.

Unlike regular missions, however, the Comet Interceptor does not have a clearly defined target whose characteristics can be studied before launch and used to simulate the actual mission. Not only that, but this lack of knowledge also means that the mission cannot know where and when the fly-by will take place, and thus it is hard to study the influence of third-body perturbations, mostly for the case of the planets in the Solar System. These can be in zones of their heliocentric orbits which are closer to the fly-by area, perhaps influencing the spacecraft significantly, or precisely on the opposite side of the Solar System from the Sun, meaning that their impact can certainly be disregarded.

To account for this, the mission analysis can be performed by having these variations into account, studying the various possible scenarios and how the mission would behave should they come to fruition. Another approach that can be pursued is establishing a nominal scenario: using characteristics expectable of a typical comet, and varying values for parameters such as fly-by speed or solar aspect angle – always keeping in mind the mission and system requirements – a nominal study case can be created, which represents what the mission's environment is supposed to be like when a suiting target is chosen.

Doing this implies that the whole mission must be developed with these possible variations in mind, thus striving to keep the systems as adaptable as possible to the different situations, which can be encountered. But it should also mean the opposite, that the target selection is performed keeping in mind the conditions which the mission is prepared to withstand. A combination of these two will be considered throughout the remaining analysis.

4.2.1 Gravity

As any space mission, the forces of gravity from the many intervening bodies are the main drivers of spacecraft dynamics. In the case of Comet Interceptor, this is no exception. As the main reference frame chosen for the mission is centred on the target, then the comet contributes with the central body gravity. As such, all the remaining neighbouring bodies, the Sun and planets, are seen as third-body perturbations.

$$\ddot{\boldsymbol{r}}_{\boldsymbol{central}} = -\frac{\mu_{central}}{r^2}\hat{\boldsymbol{r}} \qquad \ddot{\boldsymbol{r}}_{\boldsymbol{3rd}} = -\frac{\mu_{3rd}}{r^2}\hat{\boldsymbol{r}} + \frac{\mu_{3rd}}{r^2_{central}}\hat{\boldsymbol{r}}_{\boldsymbol{central}} \tag{4.4}$$

In the equation above, $\mu_{central}$ and μ_{3rd} represent the gravitational parameters of the central and third bodies being considered, respectively, $\boldsymbol{r}$ the vector pointing from the body exerting the gravitational influence to the spacecraft and $\boldsymbol{r_{central}}$ the vector from this very same body to the central body, the target comet in this case.

To assess the relevant sources of this force for the scenario, the Sun as well as all planets in the Solar System were considered. To do so, each case was taken separately, and the "best case" was analysed, where the comet is located as close to the body being studied as possible considering their orbits and the mission requirements. Integrating the force over time to get a deviation in position shows that the only body whose impact is big enough to be considered is the Sun, as all other planets even in this ideal case barely reach 1 km in difference over the unperturbed case. This implies that, in general, the gravity of these planets do not need to be considered for this mission. The same applies for the comet nucleus' own gravity, which, even considering the largest expectable diameter for a comet, is inconsequential. These results are summarised in Figure 4.3.

For the case of the Earth the test is slightly more complex, as the fly-by of the comet can lie arbitrarily close to the planet. However, the conclusion is similar: unless the Earth is 0.1 au or closer, its force can be disregarded. Should this not be the case, a whole layer of complexity would be added to this analysis, which will not be discussed.

Regarding the model being used, harmonic effects could be considered in the case of the comet's gravity. However, due to the paradigm of the mission, which establishes that the target is only chosen once the spacecraft is launched, it is impossible at the moment to pin down the geometry of the comet being studied. It could be the case that these terms do have a significant effect, despite the body's small mass, but for that a study regarding

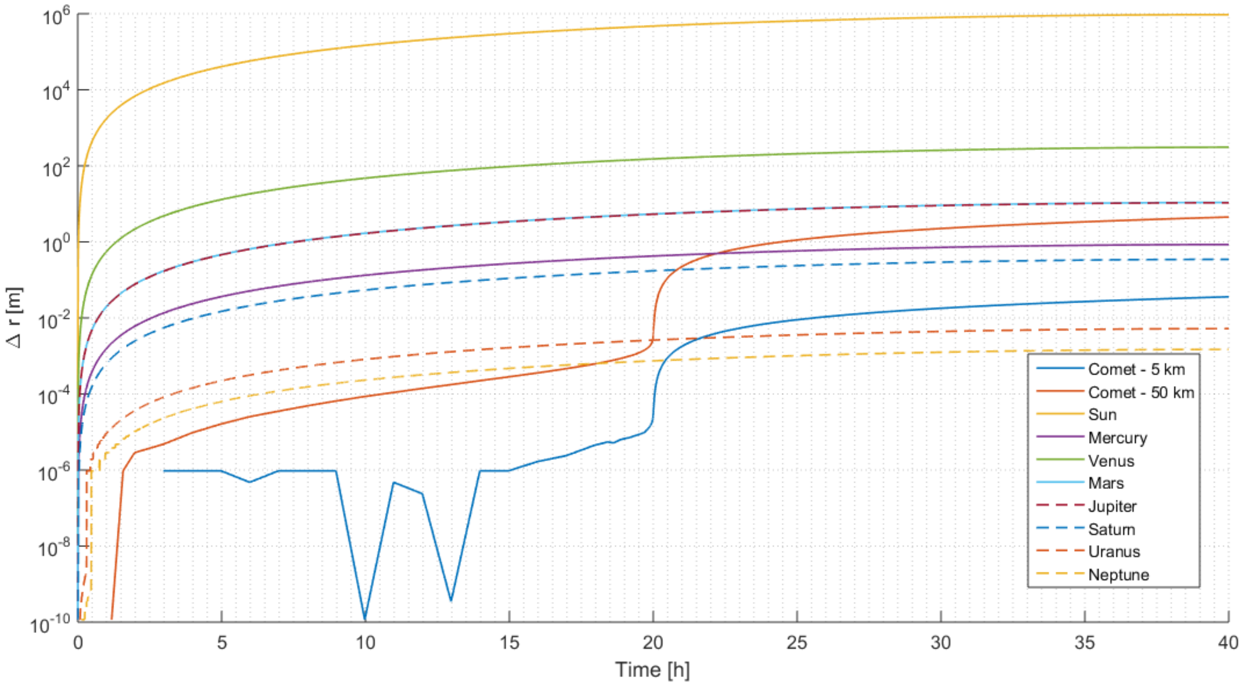

Figure 4.3: Difference between perturbed and unperturbed trajectories for each body being analysed considering a 70 km/s relative fly-by velocity. Different comet sizes are also considered.

the possible shapes of the comet and their impacts on the simulation should be performed. For the purposes of this thesis, however, these will not be considered.

4.2.2 Solar radiation pressure

The solar radiation pressure is a force that, except for cases where a vehicle is constantly in some body's shadow, is always present in Solar System exploration. Despite that, it is usually a residual contribution to the mission's dynamics, which can, nonetheless, change the expected trajectory significantly over time. This is the main reason why, despite taking it into account for the vehicle's dynamics, it is not required to go into much detail for the model of the solar radiation pressure, as the encounter phase spans only 40 hours.

Using the cannon-ball radiation pressure model, the force exerted by the incoming photons can be described by (4.5). In it, C_r is the radiation pressure coefficient, P_{Sun} the total power emitted by the Sun in the form of radiation, c the speed of light, A_{Sun} the effective area of the spacecraft from the incoming light's perspective, and $\boldsymbol{r_{Sun}}$ the vector from the Sun to the spacecraft.

$$\boldsymbol{F_{rad}} = C_r \cdot \frac{P_{Sun}}{c} \frac{A_{Sun}}{4\pi r_{Sun}^2} \cdot \boldsymbol{\hat{r}_{Sun}} \tag{4.5}$$

In the above equation, most terms are either constants or parameters referring to the vehicle's characteristics. The C_r coefficient is left undefined, however. This scalar is related to how much light is absorbed or reflected. Setting it to 1 represents a full absorption of the incoming light, whereas 2 is the number for a full (specular) reflection. It can be deduced from conservation of linear momentum that in case of a reflection, a photon exerts twice as much force as compared to just absorption, which explains the aforementioned values. In reality, this approximation is imperfect, as most surfaces reflect diffusely as well, which is less "effective" at pushing the surface than specular reflection. Thus, adjusting this parameter represents the contribution of each of the three interactions, absorption, and specular and diffuse reflection.

To make computations, the values of the parameters – and some others – are needed. The solar luminosity, P, is 3.839·10^{26} W, the speed of light, c, is 299792458 m/s, and one astronomical unit is equal to 149.597870700·10^{9} m [Wakker, 2015]. As for the vehicle's characteristics, an effective cross-section of 5 m^2, a C_r of 1.5, and a vehicle mass of 650 kg are used [GMV, 2020a].

Using these values and considering the closest distance to the Sun as permitted by the requirements of 0.9 au, a rough Δv and Δr computation can be made considering this approximately constant and radially outward force. Doing so yields a difference of approximately 700 metres, which is certainly noticeable but far smaller than the impact of the main driving force, the Sun's gravity. As such, this force shall be considered, and since its impact is not major, there is no need to explore further more precise models for its computation.

4.2.3 Comet dust impacts

The particle collisions with the spacecraft generate measurable and significant forces that must be taken into account. As analysed in [Edenhofer et al., 1987], for a duration of 100 seconds surrounding the point of C/A, the Giotto mission suffered a change in speed of approximately 23 cm/s due to dust impacts. Given that the Halley fly-by took place at 0.89 au from the Sun, with a C/A of about 600 km and a relative speed of approximately 68 km/s, it can be expected that the changes in speed suffered by Comet Interceptor during its fly-by be similar to these, given that the fly-by characteristics are somewhat similar. A basic model describing the forces exerted by the impact of particles can be developed from the analysis of Giotto's fly-by data done in the aforementioned article. In it, the spatial density of dust emitted by a jet can be computed as

$$N_d = \frac{Q_d}{\alpha^2 \pi r^2 u_d} \tag{4.6}$$

where Q_d is the dust production rate, α the jet's angular half-width, r is the radial distance from the comet, and u_d is the terminal speed of the dust flow. However, if α is set to 2, then the above equation describes an omnidirectional jet of dust, which can be considered as an approximation for the average dust distribution around the comet. The value of u_d = 400 m/s is henceforth used as mentioned in [Edenhofer et al., 1987].

To deduce the force acting on the spacecraft, the mass impacting it in a certain time interval is considered

$$\begin{aligned} dm_d &= N_d A_{dust} v_0 \cdot dt \\ F_d &= \frac{dp_d}{dt} = N_d A_{dust} v_0^2 (1+\epsilon) \\ \boldsymbol{F_d} &= N_d A_{dust} v_0 \boldsymbol{v_0} \end{aligned} \tag{4.7}$$

The above derivation assumes that the moment of the impacting dust is completely absorbed by the vehicle upon impact, which means that the impact is considered (nearly perfectly) inelastic. That is to say that the momentum enhancement factor, ϵ, is considered 0. A_{dust} represents the area of impact being considered, which will also be 5 m^2 as a first guess [GMV, 2020b] and $\boldsymbol{v}_0$ the impact velocity, which, given that the fly-by velocity is several orders of magnitude greater than the dust's velocity relative to the target, can be considered approximately equal to the fly-by velocity.

To use the derived expressions, only the dust production rate is missing. The analysis using the data from Giotto done by [ESA - Concurrent Design Facility, 2019] concluded that, using this model, it is impossible to come up to a single value for this term, as trying to adjust it to match the flux of particles measured by the mission's instruments results in Q_d = 33000 kg/s. On the other hand, the value required to explain the measured mass of particles would be 10 times less, Q_d = 3300 kg/s.

Thus, this is the interval of values that will be used for the simulation. Despite spanning one order of magnitude, a random value will be generated at the beginning of each different simulation, so that its impact can be studied. This discussion then continues in Chapter 8. However, before that, just to get an idea of the order of magnitude of the force being handled, a simple Δr computation is performed, by integrating this force over time. Considering for now the maximum value of 33000 kg/s for the dust production rate, the result is approximately 800 m, meaning that this force is on a par with the solar radiation pressure in terms of this metric, even though they point in different directions.

4.2.4 Outgassing and non-modelled accelerations

As mentioned in the last chapter, outgassing is a phenomenon describing the emission of volatiles from the nucleus of a comet, which, should it happen anisotropically, will result in a net acceleration. Described in [Królikowska, 2004], this force has three components, equal to

$$\boldsymbol{F}_{outgas} = \boldsymbol{A}_i \cdot g(r) \tag{4.8}$$

$$g(r) = \alpha_g \left(\frac{r}{r_0}\right)^{-m_g} \left[1 + \left(\frac{r}{r_0}\right)^{n_g}\right]^{-k_g} \tag{4.9}$$

where r is the distance to the Sun, in astronomical units, and α_g = 0.1113, r_0 = 2.808 au, m_g = 2.15, n_g = 5.093, and k_g = 4.6142. Despite the complicated expression, the graphical representation of $g(r)$ resembles an inverse proportionality curve, jumping to infinity around 0^+, being decreasing for all positive r values, and equalling 1 at 1 au. This makes physical sense, as with an increasing distance from the Sun, it is to expect that the decreasing energy intake from the comet results in lower sublimation rates, thus reducing the impact of outgassing.

After describing the force analytically, [Królikowska, 2004] then spends quite sometime discussing the values of $\boldsymbol{A}_i$, which thus determine the acceleration's strength, by analysing data from the observation of different comets. However, for the sake of this study, the value used for these three components will, unlike for the previous forces, be generated randomly at its beginning, since the direction of the outgassing is not known beforehand. For that, a normal distribution with mean 0 and a standard deviation of $5 \cdot 10^{-9}$ m/s^2 is used in accordance with [GMV, 2021b], as the value that currently, in the planning of the Comet Interceptor mission, better describes what the expected strength of this force will be. This magnitude is in accordance with that given by [Królikowska, 2004] for a comet such as 26P/Grigg-Skjellerup. This, as stated back in Section 1.1, is one of the two backup comets that will be chosen in case after three years no better target is found.

Because of its relatively small impact, however, this force is included under the banner of non-modelled accelerations, where, as aforementioned, it is treated as a dynamic bias generated randomly in each simulation. This is an easier way of including such small contributions into the simulation, especially in cases like this, when their description is unpredictable in terms of strength and direction. Additionally, this is an easy way of simulating any dynamical uncertainties in future steps of mission planning, as the value of the non-modelled acceleration's standard deviation can be easily changed to reflect the predicted order of magnitude of the perturbations.

Do note that it is being assumed that this contribution is constant over time. The reality is that the fly-by happens quite quickly, especially the encounter phase, lasting less than two full days. In that period, it is possible that the outgassing profile of the comet remains unchanging or varies significantly. Since this is unpredictable, it is considered that its impact is constant. This makes its effect larger than a force varying randomly over time, while also being much simpler to model, avoiding the introduction of process noises into the simulation.

4.3 Attitude

So far, the dynamics discussion has been restricted to the translational aspect of the vehicle, because it will be the most important to simulate during encounter. As per what was established by the mission requirements, this very phase begins as soon as the last trajectory correction manoeuvre is performed, which means that the spacecraft's fly-by from 20 h before the C/A until 20 h after is performed without any manoeuvres, and the path followed is only affected by external forces, being somewhat close to a straight line.

During this period, the major task that the attitude determination and control system (ADCS) has is ensuring that the shielded face of the spacecraft is pointed in the direction of its relative velocity to block dust impacts. And the only external force capable of influencing this mechanism is caused by these very same impacts, which, depending on the place where they hit the vehicle, their mass and velocity, can cause a varying effect on the attitude. However, due to the expected cumulative distribution of particles according to their mass, the majority of particles have such low mass that their effect barely affects the pointing of the payload. This is what is shown in Figure 4.4, as the expected amount of particles encountered by a fly-by mission decreases rapidly as particles of higher masses are considered. On the other hand, for these particles with higher masses, from 1 mg upward, the expected frequency of such impacts is significantly smaller throughout the mission, and for big enough masses, the effect on attitude eventually ceases to be the issue, as the structural integrity of the spacecraft is threatened.

Thus, the attitude dynamics will not considered as part of the simulation, since the ADCS error is, in general, small, and modelled as a white noise, as described in [GMV, 2020b]. This is done to represent the behaviour of the attitude determining star trackers, as discussed later during Section 8.5.2. Because of this, as mentioned in [ESA - Concurrent Design Facility, 2019], the absolute pointing error (APE) is driven by the accuracy of the relative position of the vehicle with respect to the comet. This is another reason for needing autonomous tracking, instead of a pre-conceived operational timeline, since the knowledge of the comet's ephemerides beforehand is not precise enough to allow for such a planning.

4.4 Propagation methods

The propagation methodology refers to the way that the state of the spacecraft is represented. Depending on this choice, the dynamics' differential equation obtained and its integration results can change. This can make the whole process more easy to understand, more efficient while remaining equally accurate or vice-versa,

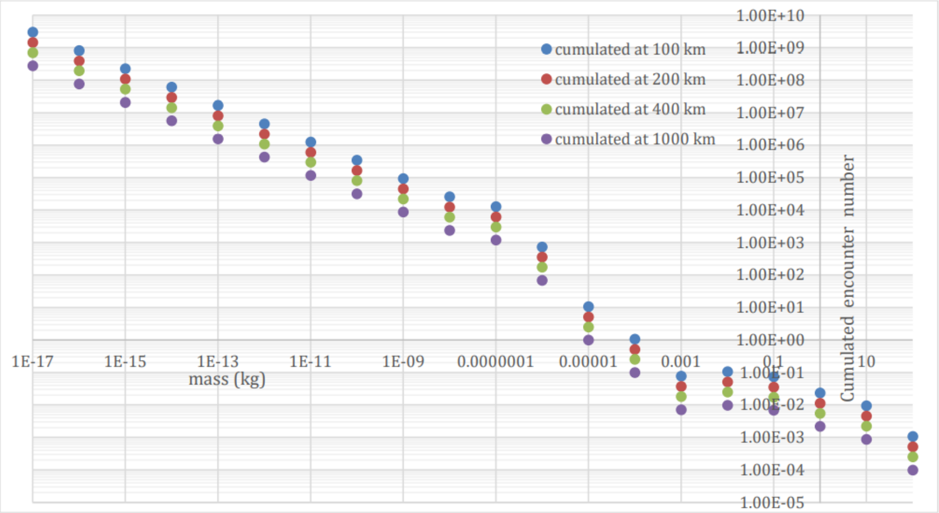

Figure 4.4: Relation between the cumulative frequency of particles and their masses. The data assumes a cometary activity of 33000 kg/s, a vehicle with a 1 m radius frontal face, and varying C/A distance. Retrieved from [GMV, 2021a].

or even more robust. Thus, a good deal of thought should go into choosing the appropriate propagation method given the problem being analysed and its characteristics.

In the case of Comet Interceptor, the spacecraft is travelling in a low gravity environment, with the comet being a tiny influence due to not being massive enough, and the Sun being the main third body to consider. Other nearby planets can, in very specific scenarios, come into the equation ever so slightly, but in the overwhelming majority of cases this should not represent a contribution big enough to dominate the Sun's role gravitational-wise, or even be big enough to consider. As perturbing forces, the solar radiation pressure and cometary dust impacts are the only ones considered, and as discussed in the section above, these represent small effects given the trajectory's relative fly-by speed order of magnitude and the Sun's gravitational pull.

As a result of this not only low-gravity, but overall low-force environment, and of the fact that the encounter phase itself is not long, in astrodynamical terms, at just 40 hours, the acting forces do not alter the orbit very far from its inertial straight-line path. Thus, propagation techniques which take advantage from using Keplerian elements have little in their favour, as this problem does not have a strong central body defining a Keplerian orbit, which is then perturbed by smaller contributions. This would be the case if a Sun-centred reference frame were to be considered, since the Sun is the dominating force, but this would represent a big disadvantage in terms of numerical accuracy.

Then, the Cowell propagation scheme should, in principle, be an adequate choice for the problem at hand. Since the state derivatives are low, then this method should not lead to any problematic numerical errors, while being the simplest approach. This is defined as

$$\boldsymbol{x} = \begin{pmatrix} \boldsymbol{r} \\ \dot{\boldsymbol{r}} \end{pmatrix}, \quad \dot{\boldsymbol{x}} = \begin{pmatrix} \dot{\boldsymbol{r}} \\ \ddot{\boldsymbol{r}} \end{pmatrix}, \quad \ddot{\boldsymbol{r}} = \frac{\sum \boldsymbol{F}}{m} - \boldsymbol{a_{comet}} \tag{4.10}$$

These are the equations which need to be solved to simulate the system's evolution, given a certain starting condition defining the state 20 hours prior to the closest approach.

Do note that the last term in (4.10) was added as a correction, since the Cowell formulation is based on Newton's 2$^{\text{nd}}$ law of motion. Since it only applies to inertial frames, then the acceleration of the frame with respect to an inertial observer must be included if it exists. In the case of Comet Interceptor, the main driver of the comet's dynamics is the Sun's gravity, and that is the acceleration considered in this term, and this can be used as the inertial centre relative to which the comet is accelerating. In exact terms, the real inertial reference should be centred on the centre of mass of the Solar System, and other planets' gravitation influences, as well as other minor perturbations, will also contribute to the comet's dynamics. These effects are, however, discarded given the little impact they have.

Chapter 5

Navigation system

This chapter is dedicated to describing the way that the navigation is performed aboard the spacecraft. Firstly, an analysis of the system's state is done, where the different variables that can be included in the estimation process are stated. Secondly, a brief discussion is had regarding the way that the uncertainty parameters are handled by these estimators. Then, the chosen options to use for the filtering process are laid out and their algorithms carefully described. Three formulations will be considered, and, for each, a sequential and batch form is used, for a total of six filters.

With the contenders established, a trade-off between the pros and cons of each formulation are discussed, in order to make an *a priori* prediction of what the best options might end up being. Finally, some strategies are introduced as possibilities to be used in aiding the robustness of these estimators against corrupted measurements. These include the alteration of navigation parameters, such as the size of the batches of measurements being processed, the usage of other sensors' information to discard corrupted measurements, as well as other algorithms capable of discarding outlier images.

5.1 System states and navigation description

The central part of this study is undoubtedly the navigation system, since it is the portion of the spacecraft that needs to be designed and tested to answer the posed research question (see Section 1.3). Thus, from the true state of the spacecraft to the system's best estimation, all elements of the chain must be discussed and their impact on the simulation of the scenario properly described.

Using the analysis of the mission's dynamics during encounter performed in the previous chapter, the true state of the system can be propagated forwards to obtain the vehicle's real position and velocity at all times. However, from the get-go, the on-board systems do not have access to this information. Instead, the best information available is the state estimation resulting from the last contact with ground 20 hours before the closest approach. This final communication step also provides the spacecraft with the state's covariance, that is, the degree of certainty of each of the state's components and how they are related with one another.

It is, therefore, important to define the state of the spacecraft before continuing. Although trivial sounding, this has a little more to it than it seems, as the state, for the purposes of navigation, is not only defined by the position and velocity, but this vector must be augmented with all the system variables or environment parameters which are not perfectly known – i.e, nearly all of them – and that are somehow related to the measurements or the dynamics. The state then is composed by:

Basic state entries:
1-3. $\boldsymbol{r}$ – vehicle's position vector with respect to the target comet.
4-6. $\dot{\boldsymbol{r}}$ – vehicle's velocity vector with respect to the target comet.
Augmented state entries:
7\. μ_{comet} – the comet's gravitational parameter.
8\. Q_d – the cometary activity value.
9-11. $\boldsymbol{r_{Sun}}$ – the Sun's position relative to the comet.
12\. SRP – solar radiation pressure, related to the uncertainty in the effective area and radiation pressure coefficient.
13\. R_{comet} – the comet's radius.
14-16. **NMA** – non-modelled acceleration caused by external sources not considered – planets or moons – or by errors in the used acceleration models.
17-18. $\mathbf{Cam}_{\text{bias}}$ – bias related to the camera's orientation misalignment.
19-20. $\mathbf{IP}_{\text{bias}}$ – bias contribution from the image processing.

Noises:

1-2. $\mathbf{IP}_{\text{noise}}$ – white noise contribution from the image processing.
3-5. $\boldsymbol{q}_{SC}$ – spacecraft's orientation in space.

The augmented state entries are what is called consider states. These consider states are the solution to the problem of not taking into account the uncertainty of the parameters affecting the problem's dynamics or measurement functions, and thus ending up with a filter with an overly optimistic covariance estimate [Carpenter et al., 2018]. To mitigate this issue, a distinction is made between estimated, or solve states, which correspond to the basic states above, and consider states, whose best estimate and uncertainty are considered, but never updated. This helps making the filter's estimates more realistic by taking these uncertainties into account, all while recognising that they are not crucial states to update unlike the vehicle's position and velocity. Additionally, for a majority of these parameters, it is also expectable that the estimation effort could do little to improve their values. Do note that the considered states' values will not be their nominal ones, but instead represent their differences with respect to said nominal. One aspect that also might stand out is the presence of the comet's gravitational parameter, despite the comet's gravity not being dynamically considered. This is a software-based choice, as FCS-ATOMIC, which is analysed further in Chapter 7, is formulated to always work with a central body gravity force, and thus the uncertainty in its gravitational parameter was included as a consider state. Despite this, it is not at all relevant, and thus does not play a role in the analysis to be performed as a part of the following thesis work.

This state definition is the one used for the previous mission analysis simulations done by GMV, and it does the job of encapsulating most of the known contributions of uncertainty in the system's state, its dynamics and the measurements it performs. From entries 1. through 6. the covariance matrix expresses the accuracy of the system's state. Entries 7. through 12. and 14. through 16. are related to the factors that can induce error in the computation of the forces, the Sun's gravity, its radiation pressure, and the cometary dust impacts, with an added contribution of non-modelled accelerations for all other dynamical effects not being considered, such as the outgassing. Finally, 13. and 17. through 20. are states related to the uncertainty in the measurement and image processing phases – which is analysed further ahead, along with all the rationale between these parameters, during Chapter 8.

Do note that a distinction is made between the biases included in the state entry, and the noises that also influence it but, in their condition of noises, are not included in the state vector, and are instead considered in the measurement covariance matrix, which the filter receives. Together, however, they make up all the factors affecting the scenario. And, all in all, their nominal values and standard deviations are established and discussed later in Chapter 8. Additionally, it is important to mention that the filters will have perfect information regarding the covariance of these parameters. As such, they will not know what the exact values of the biases or noises are, but they will know how these uncertainties are generated. Thus, the covariances of these parameters and the measurement noises will not be tuned, as the GNC system already knows all the uncertainties which are being considered by the simulation, as well as their characteristics.

With the initial values of all these state entries being provided at the beginning of the encounter phase, it is, as aforementioned, possible to propagate the basic state to obtain a prediction of the trajectory that the spacecraft undertakes during the closest phase of the fly-by. Yet, this prediction is never exact, as the initial state estimate is precisely that, an estimation, and the initial error is propagated over time. Additionally, there are all the consider states' uncertainties described above, which will slightly, but certainly, ensure that the true physical trajectory described by the vehicle differs from the *a priori* prediction.

This is where the GNC system enters the picture. According to the mission requirements, this needs to act on its own to ensure that the spacecraft uses the available information and instruments to determine its position and act accordingly to adjust the scanning mirror, striving to keep the comet inside the CoCa's FoV. As such, the system must be designed to automatically take measurements using the NAVCAM to capture the comet, and use the pictures to gain information about its location.

With this goal in mind, a schedule must be accorded beforehand so that the WAC can snap the pictures at the required times, or according to a preset frequency. This must take into account certain factors, these mainly being the exposure period required to capture the target, which will naturally become much smaller the closer to the target the vehicle is, and the internal specifications of the electronics on-board, which will determine the maximum amount of information that can be stored or even moved between components for usage, and the rate at which the images are able to be processed. These instrument-related details are, however, not discussed in this analysis, and a frequency within the interval established by **CAM – 1.6.** is assumed. This choice is discussed further ahead in Chapter 8.

This then leads to the second step, which is processing the taken picture. This procedure will not be studied in detail in the context of this study, yet understanding how the image processing technique for the Comet Interceptor mission works is required to model its functioning, as well as its shortcomings, such that the actual scenario can be simulated with utmost realism. This will be responsible for taking the snapped picture from the spacecraft's cameras, a series of brightness values for each pixel, and transforming these into information about the target comet's whereabouts in the picture. Knowing the estimated location of the comet relative to the camera's pointing axis, the camera's location and orientation relative to the frame of the vehicle, and the orientation of the spacecraft in space, it is possible to assess the location of the comet from the point of view of the spacecraft – or vice-versa, in this case – and use this information to aid the state-estimation effort.

It is this estimation effort that is the focus of the navigation system. It is performed by a filter, that is, an algorithm that, given the best estimate of the state system some time in the past, uses its dynamical propagation laws to make a prediction of the state in the present and confronts the result with the obtained measurements of the true state. Then, using the available information about the prediction's covariance and the measurements' uncertainties, it produces a new estimation of the current state of the spacecraft, and the cycle is repeated.

The main portion of this chapter will be dedicated to the description and trade-off between various algorithms available for this job, as there are a myriad of factors to take into consideration when attempting to make the best choice given the scenario being studied. Obviously, the robustness of the methods will be taken into consideration the most, and attention will also be put towards filters which have been used in the past, either for space missions or other aerospace applications, since it is a field that is usually very demanding, again, in terms of robustness. Regardless, since the goal of the thesis is to study different options, all the considered algorithms will end up being used.

5.2 Filtering algorithm

As previously mentioned, the filtering algorithm will be the brain of the navigation operation. Receiving information about the system's current state estimate and its underlying uncertainty, as well as external measurements regarding this very same state, the job of the filter is to combine these two pieces of information, which will inevitably disagree with one another, and find out what the new best estimate for the state's system is.

To achieve this, there are many possible algorithms to consider, which will be traded-off, such that a prediction of the best candidates for the task at hand can be made. However, there are more details than the method itself. Different formulations using the same method can also be discussed, where the underlying logic is the same but the equations are manipulated algebraically, which can yield some advantages in terms of computational efficiency or accuracy of the obtained results, due to being more resistant to numerical errors, for example.

Additionally, a distinction can be made as to how the observations are used between sequential and batch algorithms. On the one hand, the former is able to give state estimates right after every observation. The latter, on the other hand, takes advantage of a whole set of data, called a batch, to give an estimate at a determined epoch. This last technique might be advantageous for the case of the Comet Interceptor's encounter, as analysing a batch of data as one whole makes it easier to distinguish outliers in the measurements, corrupted by some external event, which can then be treated accordingly. Additionally, doing so might also help reduce the influence of the noises affecting the measurement data, as they should have the natural tendency of cancelling each other.

In the end, the six algorithms presented below were chosen for the advantages they show, and taking into consideration the heritage that the usage of Kalman filters has in the aerospace industry. Additionally, the number of options was not higher due to time and computational effort restrictions, as the number of variables being studied in this mission will quickly escalate over this thesis, and will be a cause of constant concern and analysis. As such, the choice of filters tried to balance including a myriad of different formulations such that these could be studied and discussed – one of the more informal goals of this thesis – while also making sure that these were relevant choices and would not overwhelm the number of tests needed down the road during the main analysis.

One aspect to consider is the available knowledge regarding the initial state of the system. Since all the proposed methods in the following sections take advantage of the previous state estimations, this information should be provided. In the case of the encounter of the Comet Interceptor mission, such information stems, as aforementioned, from prior estimations done during the previous phase, the approach. These also use the on-board cameras' information for navigation, yet instead of processing the measurements autonomously, these are sent back to Earth and processed by the mission's ground segment, and relayed back. Because the approach ends with the last TCM, according to **MIS – 1.8.**, then the effect of the manoeuvre must be added to the last available state estimation, and its dispersion added to the state's respective covariance.

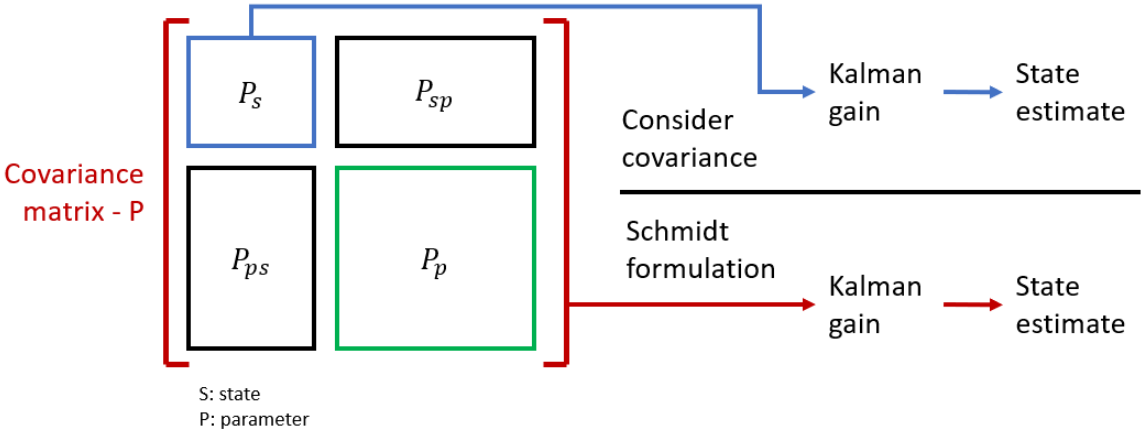

Figure 5.1: Schematic representing the different usage of the covariance matrix in the computation of the Kalman gain by the two described methods.

5.2.1 Consider formulation

Before the discussion regarding which filters were chosen for this job, their algorithms, and results, is had, an equally important and underlying formulation-related topic must be addressed, apart from the ones already mentioned above. The issue is related to how the consider states are taken into account when performing the filter estimates. As previously mentioned, these are present to make the filter's estimate more realistic, by considering all the multiple sources of error and uncertainty present in the knowledge it has of its systems and of the surrounding environment, all the while not having their values estimated, just the respective covariance values taken into account.

The initial plan was to follow the consider covariance analysis described by [Tapley et al., 2004], which is based on the sensitivity matrix, S_{cons}. This represents the derivative of each of the basic states with respect to each of the considered parameters. This works such that the regular filters without consider states are kept unchanged, and steps are added to update this matrix as time passes and new information is acquired, as well as to use it to correct the initial estimations.

The great advantage of this method is the presence of the sensitivity matrix, which gives the user a great deal of information regarding how much each of the parameters being considered is affecting the system's state over time. By comparing the orders of magnitude of each parameter combined with their respective derivatives against one another, the biggest contributors to the change in state estimation and its respective covariance become clear.

Nevertheless, as the name suggests, this method only directly affects the covariance of the estimated states, and not the states themselves. This is because it only uses the state covariance block when computing the Kalman gain, instead of the whole covariance that takes into account all the parameters being considered, as is shown in the schematic of Figure 5.1. One other formulation which does not see this problem is the Schmidt formulation. Simply put, this approach suggests treating all states equally, performing a regular filter estimation, before undoing all the changes made to the consider parameter entries in the state correction and covariance, since their goal is just to make the estimation result more realistic, and should not be estimated. This method is described by [Zanetti and D'Souza, 2013].

Using the initially suggested approach showed, however, some problems. Firstly, since it is fundamentally different from the Schmidt alternative, that is the one upon which the sequential unscented Kalman filter – that was already developed by GMV – and its batch counterpart are based, this would mean that the results would then compare filters with not only a different estimation method, but also a distinct approach to the considered parameters. This specific formulation was chosen for the unscented filters, as it is significantly simpler to implement, while still allowing for these uncertainties to be considered, albeit at the cost of a heavier computational load. Doing this introduces an extra variable in the problem, which makes drawing conclusions as to which filter is the best from the results less straight-forward.

Secondly, as was mentioned just above, its results are not as accurate as those provided by the Schmidt alternative, as the full covariance is not used in the computation of the Kalman gain, which determines the correction done after each measurement. And third, and most critically, the results of this implementation always showed signs of being incorrect, either due to an undiscovered error in the implementation itself, or because it was not suited to solve the problem at hand. This would manifest itself in the navigation results, with the covariance of the position variables decreasing a few hours before the closest approach, rapidly increasing again, before once again going down at the closest approach. This evolution made no physical sense, as the information of the images regarding the comet's position should always increase as the vehicles closes in on the target, and thus the confidence in the state predictions progressively increase as well, resulting in lower covariance values. All this lead to the abandoning of the consider covariance method in favour of the Schmidt formulation, since it is mathematically simpler, more accurate, and similar to the pre-existing filters.

5.2.2 Extended Kalman filter

No discussion of filtering algorithms would be complete without mentioning the linear Kalman filter. It relies on the linearisation of the dynamical and observation equations, which govern the movement of the spacecraft and the results of taken observations. This linearisation is performed about a reasonable reference trajectory, such that the problem becomes studying not the state of the system itself, but its deviation over time with respect to the reference, which, for small enough deviation values, should ensure that the approximation is feasible.

However, the algorithm of interest for this section is the extended version of the Kalman filter. The main difference between the two is that this version updates the reference trajectory after every state estimation such that it always reflects the best knowledge of the true trajectory, in an attempt to make the linearisation errors less significant. This change also allows for the state propagation to be done by numerical integration, instead of relying on the linearisation procedure. The derivation of this algorithm, is beyond the scope of this study, but is elegantly summarised by [Zanetti and D'Souza, 2013], especially regarding the changes made by introducing the Schmidt formulation. It can be implemented with following steps.

(0) Given initial values:

$$\boldsymbol{X}^{+}_{\boldsymbol{k-1}}, \boldsymbol{P}^{+}_{\boldsymbol{k-1}}, \boldsymbol{Z}_{\boldsymbol{k}}, \boldsymbol{R}_{\boldsymbol{k}}$$

(1) Integration from t_{k-1} to t_k:

$$\dot{\boldsymbol{X}}^* = \boldsymbol{F}(\boldsymbol{X}^*, t), \qquad \text{with } \boldsymbol{X}^*(t_{k-1}) = \boldsymbol{X}^{+}_{\boldsymbol{k-1}} \tag{5.1}$$

$$\dot{\boldsymbol{\phi}}(t, t_{k-1}) = \boldsymbol{A}(t)\boldsymbol{\phi}(t, t_{k-1}), \qquad \text{with } \boldsymbol{\phi}(t_{k-1}, t_{k-1}) = \boldsymbol{I} \quad \text{and } \boldsymbol{A}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X}_{\boldsymbol{s}}(t)} \tag{5.2}$$

$$\dot{\boldsymbol{\theta}}(t, t_{k-1}) = \boldsymbol{A}(t)\boldsymbol{\theta}(t, t_{k-1}) + \boldsymbol{B}(t), \qquad \text{with } \boldsymbol{\theta}(t_{k-1}, t_{k-1}) = \boldsymbol{0} \quad \text{and } \boldsymbol{B}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X}_{\boldsymbol{p}}(t)} \tag{5.3}$$

(2) Prediction (*a priori* update):

$$\boldsymbol{P}^{-}_{\boldsymbol{k}} = \boldsymbol{\Psi}(t_k, t_{k-1})\boldsymbol{P}^{+}_{\boldsymbol{k-1}}\boldsymbol{\Psi}^T(t_k, t_{k-1}), \qquad \text{with } \boldsymbol{\Psi} = \begin{bmatrix} \boldsymbol{\phi} & \boldsymbol{\theta} \\ 0 & \boldsymbol{I} \end{bmatrix} \tag{5.4}$$

(3) Correction (*a posteriori* update):

$$\boldsymbol{z}_k = \boldsymbol{Z}_k - \boldsymbol{G}(\boldsymbol{X}^*, t_k) \tag{5.5}$$

$$\tilde{\boldsymbol{H}}_{\boldsymbol{k}} = \frac{\partial \boldsymbol{G}(\boldsymbol{X}^*, t_k)}{\partial \boldsymbol{X}} \tag{5.6}$$

$$\boldsymbol{W}_{\boldsymbol{k}} = \tilde{\boldsymbol{H}}_{\boldsymbol{k}}\boldsymbol{P}^{-}_{\boldsymbol{k}}\tilde{\boldsymbol{H}}^{\boldsymbol{T}}_{\boldsymbol{k}} + \boldsymbol{R}_{\boldsymbol{k}} \tag{5.7}$$

$$\boldsymbol{K}_{\boldsymbol{k}} = \boldsymbol{P}^{-}_{\boldsymbol{k}}\tilde{\boldsymbol{H}}^{T}_{\boldsymbol{k}}\boldsymbol{W}^{-1}_{\boldsymbol{k}} \tag{5.8}$$

$$\boldsymbol{X}^{+}_{k} = {\boldsymbol{X}^{-}}^{*}_{\boldsymbol{k}} + \boldsymbol{K}_{\boldsymbol{k}}\boldsymbol{z}_{\boldsymbol{k}} \tag{5.9}$$

$$\boldsymbol{P}^{+}_{\boldsymbol{k}} = \boldsymbol{P}^{-}_{\boldsymbol{k}} - \boldsymbol{K}_{\boldsymbol{k}}\tilde{\boldsymbol{H}}_{\boldsymbol{k}}\boldsymbol{P}^{-}_{\boldsymbol{k}} - \boldsymbol{P}^{-}_{\boldsymbol{k}}\tilde{\boldsymbol{H}}^{\boldsymbol{T}}_{\boldsymbol{k}}\boldsymbol{K}^{\boldsymbol{T}}_{\boldsymbol{k}} + \boldsymbol{K}_{\boldsymbol{k}}\boldsymbol{W}_{\boldsymbol{k}}\boldsymbol{K}^{\boldsymbol{T}}_{\boldsymbol{k}} \tag{5.10}$$

(4) Undo corrections made to consider states (Schmidt formulation):

$$[\boldsymbol{X}^{+}_{\boldsymbol{k}}]_p = [\boldsymbol{X}^{+}_{\boldsymbol{k-1}}]_p \tag{5.11}$$

$$[\boldsymbol{P}^{+}_{\boldsymbol{k}}]_p = [\boldsymbol{P}^{+}_{\boldsymbol{k-1}}]_p \tag{5.12}$$

(5.13)

(5) Replace k with $k+1$ and return to step (1).

There is quite a lot of notation to unpack in the equations above. First and foremost, $\boldsymbol{X}$ and $\boldsymbol{Z}$ are the true state of the system and the performed measurement. $\boldsymbol{X}^*$ is the reference state of the system, the state according to the reference trajectory chosen beforehand, for the first iteration, which is then taken to be the latest state estimation by the filter. Then, the lower case z represents the difference between the measured values and the simulated measurements obtained considering the reference trajectory. $\boldsymbol{F}$ is the function that determines the system's dynamical evolution, and $\boldsymbol{G}$ outputs the measurement given a certain time and state. Additionally, the $^-$ (minus) and $^+$ (plus) symbols represent the predicted – or *a priori* – value and the estimated – or *a posteriori* value – respectively. The difference between the two is that the former is obtained by simply getting the initial values and propagating the system in time, whereas the latter is the correction of the prediction when the measurements are taken into account. Thus, for example, $\boldsymbol{P_k^-}$ is the prediction of the state's covariance at time k, that is, a measure of its uncertainty, whereas $\boldsymbol{P_k^+}$ is the corrected covariance at time k. $\boldsymbol{\phi}$ is called the state transition matrix, since, due to the linearisation performed, it transforms state deviations from one point in time to another, as seen in the prediction step, together with $\boldsymbol{\theta}$, the consider transition matrix, which is the analogue for the considered parameters, giving the change in state at a different time as the result of a change in these parameters. Then, $\boldsymbol{K}$ is the Kalman gain matrix, $\boldsymbol{R}$ the observation noise covariance matrix; and, finally, the subscripts p and s refer to the considered parameter states and to the estimated states, respectively, as per the notation utilised in Figure 5.1.

To solve this problem, the only values required are the initial state and its respective covariance, as well as the next measurement in line and its respective covariance, as well. Do note that, for the first iteration of this cycle there is a possibility that the system does not have access to an initial estimate of its state. In that case, this iteration can be skipped, and the initial guess of the next iteration taken as the one corresponding to the first measurement and its propagation in time as the reference trajectory, for lack of better information. And, of course, information about the system's dynamics and observation process are required, in the form of the function $\boldsymbol{F}$, $\boldsymbol{G}$, and their respective derivatives. Another aspect that must be kept in mind when using the EKF is that the reference trajectory should not be updated right after the first iteration of the algorithm, especially if the observations are noisy [Tapley et al., 2004]. Allowing a few observations to take place without updating the reference trajectory, thus defaulting to the regular Kalman filter, should be done to allow for the estimates to converge, correcting any possible issues with noise.

The EKF algorithm presents the advantage of not being as prone to linearisation errors as the regular Kalman filter, thus resulting in a quicker convergence to the best estimates [Tapley et al., 2004]. Additionally, being one of the most recognisable methods in filtering, it is well understood, and many formulations based on it have been developed to improve its performance or accuracy. On the negative side, it is important to take into consideration that this method is still reliant on linearisation, and thus is prone to failure especially in situations of rapid variation of the system's state or measurements relative to the expected reference. This aspect should definitely be taken into consideration when evaluating its robustness, as prolonged corrupted measurements might be able to make the filter diverge.

5.2.3 UDU extended Kalman filter

The UDU formulation, in turn, has seen widespread use in the aerospace industry for a long time. The idea behind this version is to decompose the covariance matrix into two factors, one being an upper triangular matrix with a diagonal of ones, $\boldsymbol{U}$, and the other a diagonal matrix, $\boldsymbol{D}$, such that $\boldsymbol{P} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^T$. The usage of this method intends to improve the filter's performance in terms of computational stability [Zanetti and D'Souza, 2013]. One of the results of its usage is that the symmetry of the covariance is ensured by default, and checking whether its semi-positive definite is as easy as observing if the elements of $\boldsymbol{D}$ are non-negative. Additionally, since $\boldsymbol{D}$ is diagonal and $\boldsymbol{U}$ has ones in its main diagonal, they can be stored together as one matrix, with $\boldsymbol{D}$ as the main diagonal do $\boldsymbol{U}$. Indeed, this is how this matrix is kept and used in real spacecraft's internal computers.

The following equations describe the algorithm as per [Zanetti and D'Souza, 2013], where the derivation for the majority of the expressions can also be found. Some other aspects will have their own separate reference.

(0) Initialisation:

Given $\boldsymbol{X_{k-1}^{+}}$, $\boldsymbol{P_{k-1}^{+}}$, $\boldsymbol{Z_k}$ and $\boldsymbol{R_k}$

$$\boldsymbol{U_{k-1}^{+}}\boldsymbol{D_{k-1}^{+}}\boldsymbol{U_{k-1}^{+T}} = \boldsymbol{P_{k-1}^{+}} \tag{5.14}$$

(1) Integration from t_{k-1} to t_k:

$$\dot{\boldsymbol{X}}^* = \boldsymbol{F}(\boldsymbol{X}^*, t), \qquad \text{with } \boldsymbol{X}^*(t_{k-1}) = \boldsymbol{X_{k-1}^{+}} \tag{5.15}$$

$$\dot{\boldsymbol{\phi}}(t, t_{k-1}) = \boldsymbol{A}(t)\boldsymbol{\phi}(t, t_{k-1}), \qquad \text{with } \boldsymbol{\phi}(t_{k-1}, t_{k-1}) = \boldsymbol{I} \ \text{ and } \boldsymbol{A}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X_s}(t)} \tag{5.16}$$

$$\dot{\boldsymbol{\theta}}(t, t_{k-1}) = \boldsymbol{A}(t)\boldsymbol{\theta}(t, t_{k-1}) + \boldsymbol{B}(t), \qquad \text{with } \boldsymbol{\theta}(t_{k-1}, t_{k-1}) = \boldsymbol{0} \ \text{ and } \boldsymbol{B}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X_p}(t)} \tag{5.17}$$

(2) Prediction (*a priori* update) – UDU time update algorithm [Thorton and Bierman, 1980]:

$$\bar{\boldsymbol{\Psi}}_{\boldsymbol{k}} = \boldsymbol{\Psi}(t_k, t_{k-1})\boldsymbol{U_{k-1}^{+}}, \qquad \text{with } \boldsymbol{\Psi} = \begin{bmatrix} \boldsymbol{\phi} & \boldsymbol{\theta} \\ 0 & \boldsymbol{I} \end{bmatrix} \tag{5.18}$$

```
for $i = n \rightarrow 1$, where $n$ is the number of states
  $\boldsymbol{D_k^{-}}(i,i) = \bar{\boldsymbol{\Psi}}_{\boldsymbol{k}}(i,:)\boldsymbol{D_{k-1}^{+}}\bar{\boldsymbol{\Psi}}_{\boldsymbol{k}}(i,:)^T$    (5.19)
  if $i > 1$    (5.20)
    for $j = 1 \rightarrow i-1$
      $\boldsymbol{U_k^{-}}(j,i) = \frac{\bar{\boldsymbol{\Psi}}_{\boldsymbol{k}}(j,:)\boldsymbol{D_{k-1}^{+}}\bar{\boldsymbol{\Psi}}_{\boldsymbol{k}}(i,:)^T}{\boldsymbol{D_k^{-}}(i,i)}$    (5.21)
      $\bar{\boldsymbol{\Psi}}_{\boldsymbol{k}}(j,:) = \bar{\boldsymbol{\Psi}}_{\boldsymbol{k}}(j,:) - \boldsymbol{U_k^{-}}(j,i)\bar{\boldsymbol{\Psi}}_{\boldsymbol{k}}(i,:)$    (5.22)
    end for
end for
```

(3) Obtain independent measurements [Bierman, 1977]:

$$\tilde{\boldsymbol{H}}_{\boldsymbol{k}} = \frac{\partial \boldsymbol{G}(\boldsymbol{X}^*, t_k)}{\partial \boldsymbol{X}_k} \tag{5.23}$$

$$\boldsymbol{z_k} = \boldsymbol{Z_k} - \boldsymbol{G}(\boldsymbol{X_k^*}, t_k) \tag{5.24}$$

$$\boldsymbol{R_k} \xrightarrow{Cholesky} \boldsymbol{R_k'}^T \boldsymbol{R_k'} \tag{5.25}$$

$$\bar{\boldsymbol{R}}_{\boldsymbol{k}} = \boldsymbol{R_k'}^{-1}\boldsymbol{R_k}\boldsymbol{R_k'}^{-T} = \boldsymbol{I} \tag{5.26}$$

$$\bar{\boldsymbol{H}}_{\boldsymbol{k}} = \boldsymbol{R_k'}^{-1}\tilde{\boldsymbol{H}}_{\boldsymbol{k}} \tag{5.27}$$

$$\bar{\boldsymbol{W}}_{\boldsymbol{k}} = \bar{\boldsymbol{H}}_k\boldsymbol{U_k^{-}}\boldsymbol{D_k^{-}}\boldsymbol{U_k^{-}}^T\bar{\boldsymbol{H}}_{\boldsymbol{k}}^T + \bar{\boldsymbol{R}}_{\boldsymbol{k}} \tag{5.28}$$

$$\bar{\boldsymbol{z}}_{\boldsymbol{k}} = \boldsymbol{R_k'}^{-1}\boldsymbol{z_k} \tag{5.29}$$

(4) Modified rank-one update:

```
for $j = 1 \rightarrow N_{mea}$, where $N_{mea}$ is the number of independent measurements
  $\bar{\boldsymbol{f}} = \boldsymbol{U}_{\boldsymbol{k}}^{-T} \bar{\boldsymbol{H}}_{\boldsymbol{k}}^{T}(j,:)$    (5.30)
  $\bar{\boldsymbol{v}} = \boldsymbol{D}_{\boldsymbol{k}}^{-} \bar{\boldsymbol{f}}$    (5.31)
  $\bar{\boldsymbol{K}}(1,1) = \begin{bmatrix} \bar{\boldsymbol{v}}(1) & 0 & \cdots & 0 \end{bmatrix}^{T}$    (5.32)
  $\bar{\boldsymbol{\alpha}}(1) = \bar{\boldsymbol{R}}_{\boldsymbol{k}}(j,j) + \bar{\boldsymbol{v}}(1)\bar{\boldsymbol{f}}(1)$    (5.33)
  $\bar{\boldsymbol{D}}(1,1) = \left[\bar{\boldsymbol{R}}_{\boldsymbol{k}}(1,1)/\bar{\boldsymbol{\alpha}}(1)\right] \boldsymbol{D}_{\boldsymbol{k}}^{-}(1,1)$    (5.34)
  for $i = 2 \rightarrow N_{states}$
    $\bar{\boldsymbol{\alpha}}(i) = \bar{\boldsymbol{\alpha}}(i-1) + \bar{\boldsymbol{v}}(i)\bar{\boldsymbol{f}}(i)$    (5.35)
    $\bar{\boldsymbol{\lambda}}(i) = -\left[\bar{\boldsymbol{f}}(i)/\bar{\boldsymbol{\alpha}}(i-1)\right]$    (5.36)
    $\bar{\boldsymbol{D}}(i,i) = \left[\bar{\boldsymbol{\alpha}}(i-1)/\bar{\boldsymbol{\alpha}}(i)\right] \bar{\boldsymbol{H}}_{\boldsymbol{k}}$    (5.37)
    $\bar{\boldsymbol{U}}(:,i) = \boldsymbol{U}_{\boldsymbol{k}}^{-}(:,i) + \bar{\boldsymbol{\lambda}}(i)\bar{\boldsymbol{K}}(:,i-1)$    (5.38)
    $\bar{\boldsymbol{K}}(:,i) = \bar{\boldsymbol{K}}(:,i-1) + \bar{\boldsymbol{v}}(i)\boldsymbol{U}_{\boldsymbol{k}}^{-}(:,i)$    (5.39)
  end for
  $\boldsymbol{K}_{\boldsymbol{k}}(:,j) = \bar{\boldsymbol{K}}(:,n)/\bar{\boldsymbol{\alpha}}(n)$    (5.40)
end for
```

(5) Rank-one update:

```
for $k = 1 \rightarrow N_{mea}$
  $\bar{\boldsymbol{S}} = \begin{bmatrix} 0(nS,nS), 0(nS,nP) \\ 0(nP,nS), I(nP,nP) \end{bmatrix}$    (5.41)
  $\bar{\boldsymbol{a}} = \bar{\boldsymbol{S}}\boldsymbol{K}_{\boldsymbol{k}}(:,k)$    (5.42)
  for $i = N_{states} \rightarrow 2$
    $\boldsymbol{D}_{\boldsymbol{k}}^{+} = \bar{\boldsymbol{D}}(i,i) + \sum_{b=i+1}^{n} \left[\bar{\boldsymbol{U}}^{2}(i,b)\bar{\boldsymbol{D}}(b,b)\right] - \sum_{b=i+1}^{n} \left[(\boldsymbol{U}_{\boldsymbol{k}}^{+})^{2}(i,b)\boldsymbol{D}_{\boldsymbol{k}}^{+}(b,b)\right] + \bar{\boldsymbol{W}}_{\boldsymbol{k}}(k,k)\bar{\boldsymbol{a}}^{2}(i)$    (5.43)
    for $j = i-1 \rightarrow 1$
      $\boldsymbol{U}_{\boldsymbol{k}}^{+}(j,i) = \left[\sum_{b=i}^{n} \left[\bar{\boldsymbol{D}}(b,b)\bar{\boldsymbol{U}}(j,b)\bar{\boldsymbol{U}}(i,b)\right] - \sum_{b=i+1}^{n} \left[\boldsymbol{D}_{\boldsymbol{k}}^{+}(b,b)\boldsymbol{U}_{\boldsymbol{k}}^{+}(j,b)\boldsymbol{U}_{\boldsymbol{k}}^{+}(i,b)\right] + \bar{\boldsymbol{W}}_{\boldsymbol{k}}(k,k)\bar{\boldsymbol{a}}(i)\bar{\boldsymbol{a}}(j)\right]/\boldsymbol{D}_{\boldsymbol{k}}^{+}(i,i)$    (5.44)
    end for
  end for
end for
```

(6) Correction (*a posteriori* update):

$$\boldsymbol{K}_{\boldsymbol{k}} = \begin{bmatrix} \boldsymbol{K}_{\boldsymbol{states}} \\ 0 \end{bmatrix} \tag{5.45}$$

$$\boldsymbol{X}_{k}^{+} = \boldsymbol{X}_{\boldsymbol{k}}^{-*} + \boldsymbol{K}_{\boldsymbol{k}}\boldsymbol{z}_{\boldsymbol{k}} \tag{5.46}$$

(7) Replace k with $k+1$ and return to step (1).

Although it might seem initially that this method differs greatly from the EKF analysed before, this is completely not the case. Indeed, the inputs of the function are precisely similar, only differing in the factorisation presented by the state's covariance matrix. Then, both steps (1) and (6), integration and state *a posteriori* correction, respectively, are similar to those of the EKF. All the alterations made to the filter are, then, used to handle the factorisation of the covariance matrix.

Whereas before the matrix was treated as one, and changes made to it used regular matricial operations, now the UDU factorisation must be kept, meaning that the transformations suffered by the matrix must be changed to reflect that. In (2), the time update of the matrix is absolutely equivalent to the prediction step (2) of the EKF, only adjusted numerically to keep the desired factorisation. The same can be said, then, for steps (4) and (5), where the covariance matrix is first updated given the new measurement information, similarly to what is performed in step (3) of the EKF. Then, in (5), the Schmidt part of the filter is introduced, and the changes made to the consider states are undone. These are the functions performed by the modified rank-one update and by the rank-one update, respectively. Additionally, the former computes the Kalman gain matrix while doing so, again in a similar fashion to before.

One important change made to this filter is that, since the updates to the covariance matrix are done in loops instead of using matricial operations, the measurement must be processed one component at a time. Since the camera measurements are composed of two entries, one for each axis, then the measurements must first be transformed to ensure that their components are linearly independent from one another, which is done by considering a diagonal measurement covariance matrix. Thus, the transformed measurement and all update-related variables which must then be transformed as well are represented with the $^-$ (bar) symbol. For this transformation to be done, the measurement covariance matrix initially undergoes a Cholesky decomposition, resulting in the $\boldsymbol{R_k'}$ matrix.

Nomenclature-wise, nS and nP represent the number of estimated states and the number of considered parameters, respectively, and all the remaining variables, such as $\bar{\boldsymbol{\Psi}}$, $\bar{\boldsymbol{f}}$, $\bar{\boldsymbol{v}}$, $\bar{\boldsymbol{\alpha}}$, $\bar{\boldsymbol{\lambda}}$, $\bar{\boldsymbol{S}}$, and $\bar{\boldsymbol{a}}$ are auxiliary, serving only the purpose of aiding the computations performed in the covariance updating loops contained in (2), (4) and (5). The relevant inputs for these loops are factorised covariance matrices, $\boldsymbol{U}$ and $\boldsymbol{D}$, the dynamics and measurement Jacobian matrices, $\boldsymbol{\Psi}$ and $\boldsymbol{H}$, as well as the $\bar{\boldsymbol{S}}$ matrix and the Kalman gain matrix $\boldsymbol{K}$, specifically for the case of the Rank-one update in (5). And, naturally, the important outputs are the transformed factorised covariance matrices. In the end, nonetheless, the algorithm outputs a state estimation and its respective covariance, in a similar fashion to all the remaining algorithms.

5.2.4 Batch extended Kalman filter

As previously mentioned, another way to filter the incoming measurements is to process them in groups, or batches, where more than one is considered at a time. Thus, instead of getting an estimate of the state at the time of a measurement, the batch algorithms instead use a batch of measurements to estimate the state at any point in time. As such, it will undoubtedly be necessary to make multiple transformations between distinct instants, using for that the linearisation process discussed above and the resulting state transition matrix, $\boldsymbol{\phi}$.

One interesting point is that the underlying maths is the same for the extended Kalman filter and for its batch counterpart, but the Kalman name usually only refers to the sequential formulation. This does not mean, however, that the results of the two are similar, since the reference trajectory is updated to the newest estimate after each step of the sequential filter, meaning that they slowly drift apart.

The algorithm described below is almost identical to the one derived in [Tapley et al., 2004], and uses the information matrix instead of its covariance counterpart. The only change made to reflect the Schmidt formulation is the correction of the changes made to the parameters' estimation and their covariance, returning their values back to the original ones.

(0) Initialisation:

$$\begin{aligned}
&\text{Given values for } \boldsymbol{X_0^*},\ \boldsymbol{P_0},\ t_0\\
&\Lambda = \boldsymbol{P_0}^{-1}\\
&\boldsymbol{N} = 0\\
&\boldsymbol{\phi}(t_0, t_0) = \boldsymbol{I}\\
&t_{k-1} = t_0
\end{aligned}$$

(1) Next measurement:

Read $t_k,\ \boldsymbol{Z_k},\ \boldsymbol{R_k}$

(2) Integration from t_{k-1} to t_k:

$$\dot{\boldsymbol{X}}^* = \boldsymbol{F}(\boldsymbol{X}^*, t), \qquad \text{with } \boldsymbol{X^*_{k-1}} \tag{5.47}$$

$$\dot{\boldsymbol{\phi}}(t, t_0) = \boldsymbol{A}(t)\boldsymbol{\phi}(t, t_0), \qquad \text{with } \boldsymbol{\phi}(t_{k-1}, t_0) \quad \text{and } \boldsymbol{A}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X_s}(t)} \tag{5.48}$$

$$\dot{\boldsymbol{\theta}}(t, t_{k-1}) = \boldsymbol{A}(t)\boldsymbol{\theta}(t, t_{k-1}) + \boldsymbol{B}(t), \qquad \text{with } \boldsymbol{\theta}(t_{k-1}, t_{k-1}) = \boldsymbol{0} \quad \text{and } \boldsymbol{B}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X_p}(t)} \tag{5.49}$$

(3) Accumulation:

$$\tilde{\boldsymbol{H}}_{\boldsymbol{k}} = \frac{\partial \boldsymbol{G}(\boldsymbol{X}^*, t_k)}{\partial \boldsymbol{X}} \tag{5.50}$$

$$\boldsymbol{z_k} = \boldsymbol{Z_k} - \boldsymbol{G}(\boldsymbol{X^*_k}, t_k) \tag{5.51}$$

$$\boldsymbol{H_k} = \tilde{\boldsymbol{H}}_{\boldsymbol{k}} \boldsymbol{\Psi}(t_k, t_0), \qquad \text{with } \boldsymbol{\Psi} = \begin{bmatrix} \boldsymbol{\phi} & \boldsymbol{\theta} \\ 0 & \boldsymbol{I} \end{bmatrix} \tag{5.52}$$

$$\Lambda = \Lambda + \boldsymbol{H_k^T} \boldsymbol{R_k^{-1}} \boldsymbol{H_k} \tag{5.53}$$

$$\boldsymbol{N} = \boldsymbol{N} + \boldsymbol{H_k^T} \boldsymbol{R_k^{-1}} \boldsymbol{z_k} \tag{5.54}$$

$$\tag{5.55}$$

(4) Fork 1:

k is not the last observation:
$k \rightarrow k-1$
Return to step (1)
k is the last observation:
Go to (5)

(5) Normal equation:

$$\boldsymbol{x_0^+} = \Lambda^{-1} \boldsymbol{N} \tag{5.56}$$

(6) Cleanup:

$$\begin{aligned} &\boldsymbol{X_0^+} \rightarrow \boldsymbol{X_0^*} + \boldsymbol{x_0^+} \\ &\boldsymbol{P_0^+} = \Lambda^{-1} \\ &[\boldsymbol{X_k^+}]_p = [\boldsymbol{X_{k-1}^+}]_p \\ &[\boldsymbol{P_k^+}]_p = [\boldsymbol{P_{k-1}^+}]_p \end{aligned} \tag{5.57}$$

There are some aspects to note about the above description. It assumes that no measurement is taken at the time that the estimate is considered, t_0. Should that, however, be the case, then $\Lambda = \boldsymbol{P}_0^{-1} + \boldsymbol{H}_0^T \boldsymbol{R}_0^{-1} \boldsymbol{H}_0$ and $\boldsymbol{N} = \boldsymbol{H}_0 \boldsymbol{R}_0^{-1} \boldsymbol{z}_0$ are the equations to be used in step (0). For this case, the state transition matrix is obviously equal to the identity matrix. It is also worth pointing out that, unlike what happened for the EKF, in this filter the state estimation is done for the initial time of each batch, and not for the time of the final measurement. Because of this, after the estimation is performed, this state and its covariance needs to be propagated forward to the beginning of the next batch.

Using a batch filter approach can have the advantage of grouping up a set of measurements to use at once, instead of using them as they are taken. Because each set is filtered, computations can be performed with the obtained data to, for example, identify outliers, that is, measurements disrupted by cometary events, such as the outbursts described in Section 3.3, or more general phenomena such as cosmic rays. By varying the frequency of taken measurements and the size of batches, it might be possible to ensure that these events are not captured by too many photographs and that there are enough of them to exclude the outliers. For example, picking a high frequency of measurements, with outbursts that last about a few minutes, can result in a batch having more than 10 images disrupted by the event. Likewise, if the batches are small, then these 10 pictures become a more significant portion of the batch, which can make it difficult to identify which are the outliers. Even though it would be desirable to snap as many images as possible during the encounter, this fact needs to be taken into account when defining the measurement rate. Additionally, a higher frequency would require more processing power for the IP process and for the filter itself, and this power is obviously limited and valuable in a space mission. On top of this is the fact that, as aforementioned, when processing a batch of data at once, the impacts of the noises in the measurements can cancel each other out, leading to better results.

Nevertheless, this method has its disadvantages. It does not incorporate some of the formulation changes as the UDU information filter, so its performance in terms of numerical precision is not expected to be as good, similarly to that of the EKF. Additionally, for the majority of the encounter, the state of the comet as seen from the spacecraft changes punily with time, whereas near the C/A point the situation is exactly the opposite. Thus, for most of the trajectory the filter can accumulate measurements over long periods of time before processing them to update the arc's estimate. But, as it approaches the comet, the need for more frequent estimations arises, and the fact that the batch formulations wait for a full batch of measurements to be collected before processing them all at once to update the state estimation means that the GNC system is left for longer without updated state estimations, only being allowed to propagate the previous available information. Because of this, between processing batches, the navigation performance should be worse than the sequential filters.

5.2.5 Batch UDU extended Kalman filter

In the spirit of presenting all the filter options both in a sequential and batch formulation, here the UDU factorisation changes are introduced to the batch EKF, with the intention of combining the advantages of the batch formulations with the computation stability of performing the changes to the state's covariance matrix using the UDU cycle operations seen in the previous UDU filter.

(0) Initialisation:

$$\begin{aligned}
&\text{Given values for } \boldsymbol{X_0^*},\ \boldsymbol{P_0},\ t_0\\
&\boldsymbol{U_0}\boldsymbol{D_0}\boldsymbol{U_0}^T = \boldsymbol{P_0}\\
&\boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^T = \boldsymbol{U_0}^{-T}\boldsymbol{D_0}^{-1}\boldsymbol{U_0}^{-1} = \boldsymbol{P_0}^{-1}\\
&\boldsymbol{N} = 0\\
&\boldsymbol{\phi}(t_0, t_0) = \boldsymbol{I}\\
&t_{k-1} = t_0
\end{aligned}$$

(1) Next measurement:

$$\text{Read } t_k,\ \boldsymbol{Z_k},\ \boldsymbol{R_k}$$

(2) Integration from t_{k-1} to t_k:

$$\dot{\boldsymbol{X}}^* = \boldsymbol{F}(\boldsymbol{X}^*, t), \qquad \text{with initial condition } \boldsymbol{X_{k-1}^*} \tag{5.58}$$

$$\dot{\boldsymbol{\phi}}(t, t_0) = \boldsymbol{A}(t)\boldsymbol{\phi}(t, t_0), \qquad \text{with initial condition } \boldsymbol{\phi}(t_{k-1}, t_0) \quad \text{and } \boldsymbol{A}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X_s}(t)} \tag{5.59}$$

$$\dot{\boldsymbol{\theta}}(t, t_{k-1}) = \boldsymbol{A}(t)\boldsymbol{\theta}(t, t_{k-1}) + \boldsymbol{B}(t), \qquad \text{with } \boldsymbol{\theta}(t_{k-1}, t_{k-1}) = \boldsymbol{0} \quad \text{and } \boldsymbol{B}(t) = \frac{\partial \boldsymbol{F}(\boldsymbol{X}^*, t)}{\partial \boldsymbol{X_p}(t)} \tag{5.60}$$

(3) Obtain independent measurements [Bierman, 1977]:

$$\tilde{\boldsymbol{H}}_{\boldsymbol{k}} = \frac{\partial \boldsymbol{G}(\boldsymbol{X}^*, t_k)}{\partial \boldsymbol{X}} \tag{5.61}$$

$$\boldsymbol{H}_{\boldsymbol{k}} = \tilde{\boldsymbol{H}}_{\boldsymbol{k}} \boldsymbol{\Psi}(t_k, t_0), \qquad \text{with } \boldsymbol{\Psi} = \begin{bmatrix} \boldsymbol{\phi} & \boldsymbol{\theta} \\ 0 & \boldsymbol{I} \end{bmatrix} \tag{5.62}$$

$$\boldsymbol{z}_{\boldsymbol{k}} = \boldsymbol{Z}_{\boldsymbol{k}} - \boldsymbol{G}(\boldsymbol{X}^*_{\boldsymbol{k}}, t_k) \tag{5.63}$$

$$\boldsymbol{R}_{\boldsymbol{k}} \xrightarrow{Cholesky} \boldsymbol{R}'^{\,T}_{\boldsymbol{k}} \boldsymbol{R}'_{\boldsymbol{k}} \tag{5.64}$$

$$\bar{\boldsymbol{R}}_{\boldsymbol{k}} = \boldsymbol{R}'^{\,-1}_{\boldsymbol{k}} \boldsymbol{R}_{\boldsymbol{k}} \boldsymbol{R}'^{\,-T}_{\boldsymbol{k}} = \boldsymbol{I} \tag{5.65}$$

$$\bar{\boldsymbol{H}}_{\boldsymbol{k}} = \boldsymbol{R}'^{\,-1}_{\boldsymbol{k}} \tilde{\boldsymbol{H}}_{\boldsymbol{k}} \tag{5.66}$$

$$\bar{\boldsymbol{z}}_{\boldsymbol{k}} = \boldsymbol{R}'^{\,-1}_{\boldsymbol{k}} \boldsymbol{z}_{\boldsymbol{k}} \tag{5.67}$$

$$\tag{5.68}$$

(4) Accumulate:

$$\boldsymbol{N} = \boldsymbol{N} + \bar{\boldsymbol{H}}^T_{\boldsymbol{k}} \bar{\boldsymbol{R}}^{-1}_{\boldsymbol{k}} \bar{\boldsymbol{z}}_{\boldsymbol{k}} \tag{5.69}$$

(5) Rank-one update:

```
for $l = 1 \to N_{mea}$
  $\bar{\boldsymbol{a}} = \bar{\boldsymbol{H}}^{\boldsymbol{T}}_{\boldsymbol{k}}(l, :)$    (5.70)
  $\bar{c} = \bar{\boldsymbol{R}}_{\boldsymbol{k}}(l, l)$    (5.71)
  for $i = N_{states} \to 1$
    $\hat{\boldsymbol{D}} = \boldsymbol{D}(i,i) + \sum_{b=i+1}^{n} \left[\boldsymbol{U}^2(i,b)\boldsymbol{D}(b,b)\right] - \sum_{b=i+1}^{n} \left[\hat{\boldsymbol{U}}^2(i,b)\hat{\boldsymbol{D}}(b,b)\right] + \bar{c} \cdot \bar{\boldsymbol{a}}(i)^2$    (5.72)
    for $j = i - 1 \to 1$
      $\hat{\boldsymbol{U}}(j,i) = \left[\sum_{b=i}^{n} [\boldsymbol{D}(b,b)\boldsymbol{U}(j,b)\boldsymbol{U}(i,b)] - \sum_{b=i+1}^{n} \left[\hat{\boldsymbol{D}}(b,b)\hat{\boldsymbol{U}}(j,b)\hat{\boldsymbol{U}}(i,b)\right] + \bar{c} \cdot \bar{\boldsymbol{a}}(i)\bar{\boldsymbol{a}}(j)\right] / \hat{\boldsymbol{D}}(i,i)$    (5.73)
    end for
  end for
  $\boldsymbol{U} = \hat{\boldsymbol{U}}$    (5.74)
  $\boldsymbol{D} = \hat{\boldsymbol{D}}$    (5.75)
end for
                                                                  (5.76)
```

(4) Fork 1:

k is not the last observation:
  $k \to k - 1$
  Return to step (1)
k is the last observation:
  Go to (5)

(5) Normal equation:

$$\boldsymbol{x}_{\boldsymbol{0}}^{+} = \boldsymbol{U}^{-T}\boldsymbol{D}^{-1}\boldsymbol{U}^{-1}\boldsymbol{N} \tag{5.77}$$

(6) Cleanup:

$$\begin{aligned}
&\boldsymbol{X}_{\boldsymbol{0}}^{+} \to \boldsymbol{X}_{\boldsymbol{0}}^{*} + \boldsymbol{x}_{\boldsymbol{0}}^{+}\\
&\boldsymbol{P}_{\boldsymbol{0}}^{+} = \boldsymbol{U}^{-T}\boldsymbol{D}^{-1}\boldsymbol{U}^{-1}\\
&[\boldsymbol{X}_{\boldsymbol{k}}^{+}]_p = [\boldsymbol{X}_{\boldsymbol{k-1}}^{+}]_p\\
&[\boldsymbol{P}_{\boldsymbol{k}}^{+}]_p = [\boldsymbol{P}_{\boldsymbol{k-1}}^{+}]_p
\end{aligned} \tag{5.78}$$

This filter is in most aspects equal to the batch EKF that was described in the previous section. The only change made is to the way that the covariance matrix is handled, which, in this case, is its inverse, or the information matrix, represented using the UDU factorisation. Because of that, the respective accumulation performed after each measurement is read is done according to the rank-one update algorithm, in an attempt to capture the computational stability of the previous filter but now associated with a batch rationale for processing NAVCAM information.

Whereas in the previous filter two loop operations are performed to update the covariance matrix, here only one is done, corresponding to the first one from the previous filter, where the new information is used to update its entries. The second one, which undoes the changes made to the consider parameters block of the matrix, is instead replaced by changing it to its initial state at the beginning of the iteration, since, as this filters works with the information matrix, the same operation cannot be performed.

5.2.6 Unscented Kalman filter

Up until this point, all the filtering algorithms analysed are based on the same principle: using the best known estimate for a certain time, propagate that until a new time, use the propagated state and the knowledge of the system to predict what a measurement taken at that moment in time would be, and compare the obtained result with actual measurements taken from the real world. In this process, the covariance of the system's state is also propagated based on a linearisation of the system's dynamics function, using the state transition matrix. The difference between the predicted and true measurements and the propagated covariance, together with the measurement's own covariance, are then used to determine the estimate that minimises the resulting covariance.

This is, however, not the only way of propagating the uncertainty of the state of the system from one time to another, and performing an estimation. A different class of filters, named particle filters, approach this problem differently. Instead of relying on linearisation, they make a selection of points surrounding the initial estimation, as a representation of the covariance of all the system's state variables. Then, by propagating these points along with the estimate, a new set of points is obtained, which can be used to draw information about how the uncertainties have evolved towards this new instant. To achieve this, the particle filters make a random selection of points around the initial estimate.

There are more deterministic and efficient methods than this, however. Described by [Julier and Uhlmann, 2004], the unscented Kalman filter is based on the idea that approximating the probability distribution of the state of a system is easier than approximating the function describing its dynamics, which can be arbitrarily complex and nonlinear. Thus, instead of creating an approximation for the dynamics of the system and propagating its state and covariance over time, a set of points, called sigma points, and a set of respective weights, is purposefully chosen to contain information about the mean and covariance at the initial state – possibly along with more parameters – and from their transformation, an estimate of the next state and its respective covariance can be drawn. The algorithm described in the equations below is based on the work of [Julier and Uhlmann, 2004], where the expressions are also derived.

(1) Sigma-point selection algorithm:

Use $\boldsymbol{X}_{\boldsymbol{k-1}}^{+}$ and $\boldsymbol{P}_{\boldsymbol{k-1}}^{+}$ to determine the set of sigma points $\boldsymbol{s}_{\boldsymbol{k}}$ and their respective weights $\boldsymbol{W}_{\boldsymbol{k}}(i)$.

(2) Propagation from t_{k-1} to t_k:

$$\dot{\boldsymbol{s}}_{\boldsymbol{k}}(i) = \boldsymbol{F}(\boldsymbol{s}_{\boldsymbol{k}}(i), t) \tag{5.79}$$

(3) Prediction (*a priori* update):

$$\boldsymbol{X}_{\boldsymbol{k}}^{-} = \sum_{i=0}^{p} \boldsymbol{W}_{\boldsymbol{k}}(i)\boldsymbol{s}_{\boldsymbol{k}}^{-}(i) \tag{5.80}$$

$$\boldsymbol{P}_{\boldsymbol{k}}^{-} = \sum_{i=0}^{p} \boldsymbol{W}_{\boldsymbol{k}}(i)[\boldsymbol{s}_{\boldsymbol{k}}^{-}(i) - \boldsymbol{X}_{\boldsymbol{k}}^{-}][\boldsymbol{s}_{\boldsymbol{k}}^{-}(i) - \boldsymbol{X}_{\boldsymbol{k}}^{-}]^T \tag{5.81}$$

$$\bar{\boldsymbol{Z}}_{\boldsymbol{k}}(i) = \boldsymbol{G}(\boldsymbol{s}_{\boldsymbol{k}}^{-}(i), t) \tag{5.82}$$

$$\hat{\boldsymbol{Z}}_{\boldsymbol{k}} = \sum_{i=0}^{p} \boldsymbol{W}_{\boldsymbol{k}}(i)\bar{\boldsymbol{Z}}_{\boldsymbol{k}}(i) \tag{5.83}$$

$$\boldsymbol{S}_{k} = \sum_{i=0}^{p} \boldsymbol{W}_{\boldsymbol{k}}(i)[\bar{\boldsymbol{Z}}_{\boldsymbol{k}}(i) - \hat{\boldsymbol{Z}}_{\boldsymbol{k}}][\bar{\boldsymbol{Z}}_{\boldsymbol{k}}(i) - \hat{\boldsymbol{Z}}_{\boldsymbol{k}}]^T \tag{5.84}$$

$$\boldsymbol{V}_{\boldsymbol{k}} = \sum_{i=0}^{p} \boldsymbol{W}_{\boldsymbol{k}}(i)[s_{\boldsymbol{k}}^{-}(i) - \boldsymbol{X}_{\boldsymbol{k}}^{-}][\bar{\boldsymbol{Z}}_{\boldsymbol{k}}(i) - \hat{\boldsymbol{Z}}_{\boldsymbol{k}}]^T \tag{5.85}$$

(4) Correction (*a posteriori* update):

$$\boldsymbol{Q}_{\boldsymbol{k}} = \boldsymbol{V}_{\boldsymbol{k}}\boldsymbol{S}_{\boldsymbol{k}}^{-1} \tag{5.86}$$

$$\boldsymbol{P}_{\boldsymbol{k}}^{+} = \boldsymbol{P}_{\boldsymbol{k}}^{-} - \boldsymbol{Q}_{\boldsymbol{k}}\boldsymbol{S}_{\boldsymbol{k}}\boldsymbol{Q}_{\boldsymbol{k}}^{T} \tag{5.87}$$

$$\boldsymbol{X}_{\boldsymbol{k}}^{+} = \boldsymbol{X}_{\boldsymbol{k}}^{-} + \boldsymbol{Q}_{\boldsymbol{k}}[\boldsymbol{Z}_{\boldsymbol{k}} - \hat{\boldsymbol{Z}}_{\boldsymbol{k}}] \tag{5.88}$$

Due to the inherent differences between this method and the ones discussed in the previous sections, the notation has changed slightly and new elements have been introduced. Here, $s(i)$ represents each of the sigma points, and $\boldsymbol{s}^{-}(i)$ the results of their dynamical propagation in time. Then, $\boldsymbol{X}^{-}$ and $\boldsymbol{X}^{+}$ represent the state *a priori* prediction and *a posteriori* correction, respectively. Additionally, $\bar{\boldsymbol{Z}}_k(i)$ are the observations expected to happen at each of the sigma point's propagated states, whereas $\hat{\boldsymbol{Z}}$ is the predicted observation resulting from their weighted average. Finally, $\boldsymbol{S}$ and $\boldsymbol{V}$ are the innovation covariance and cross covariance matrices, respectively, used in the correction step, and $\boldsymbol{Q}$ is a matrix with a similar role to that of the Kalman matrix in a linear Kalman filter. Regardless, it is plain that the overall steps taken to make an estimation are the same and follow a similar order, meaning that despite having a different approach to tackle the filtering problem, the underlying idea of this algorithm is the same as the ones that precede it, in the sense that an initial estimate is propagated to the time of the measurement, and then used to simulate a measurement, which is combined with the the real one to produce an estimation correction.

Despite the similarities, the unscented transformation does not suffer from some of the issues that limit the regular Kalman transformation. It does not rely on linearisation, which is an assumption whose validity is hard to verify. Additionally, employing the latter method requires the evolution function to be differentiable, which means that discontinuities pose an obstacle to its usage. And finally, the Kalman transformation, again, relies on differentiation, which is a complex process to deduce mathematically and then code, making it prone to errors [Julier and Uhlmann, 2004].

Instead of this, the unscented transformation presents itself as a more robust method, not needing to rely on the Jacobian of the dynamics functions, and thus working around discontinuities, whilst keeping a similar order of magnitude in terms of computational efficiency. This is highly dependent, of course, on the number of states present, and consequently on the number of sigma points requiring propagation and whose measurement then needs to be simulated. For the case of the Comet Interceptor, there are 20 states that need to be dealt with, meaning that this filter, along with its batch counterpart, should be significantly heavier. Nevertheless, since the sigma-point set used for the transformation can be chosen, it is possible to do so including any additional information known about the estimate or its error distribution, such as higher order statistical information than the mean and covariance [Julier and Uhlmann, 2004].

This, however, exposes one of the less advantageous features of the UKF. The choice of sigma-point set allows for extra information to be added to the transformation, as aforementioned, but this choice means that the algorithm is less general. Since the chosen set can affect the obtained results significantly [Julier and Uhlmann, 2004], a degree of tunning and experimenting can be necessary to ensure that the best set is picked for each particular scenario being studied, which adds to the complexity of the analysis.

Regarding the choice of sigma-point set, [Bitzer, 2016] performs a great analysis of the different options available and their interpretations, which will be summarised. As for the algorithms themselves, the derivations are not presented here. The first, simplest set is called the *min* set, and represents the smallest collection of points required to convey information about mean and covariance, which, in a state with D dimensions, is 1 and D, respectively, adding up to $D+1$ total points. However, much better and still relatively simple sets can be used, such as the base set. This is useful for symmetric distributions, and takes advantage of this symmetry by picking a set of symmetric points relative to the mean value. Thus, it ends up with $N=2D$ points, but also conveys information about all odd central moments, which are 0.

$$\boldsymbol{s_n} = \boldsymbol{\mu} + \sqrt{D}\boldsymbol{l_n}, \quad n = 1, ..., D \tag{5.89}$$

$$\boldsymbol{s_{n+D}} = \boldsymbol{\mu} - \sqrt{D}\boldsymbol{l_n}, \quad n = 1, ..., D \tag{5.90}$$

$$w_j = \frac{1}{2D} = \frac{1}{N}, \quad j = 1, ..., 2D \tag{5.91}$$

In the above equation, $\boldsymbol{l_n}$ represents each of the D columns of the Cholesky decomposition of the initial covariance matrix. This decomposition algorithm will be described in Chapter 6.

Then, more complex sets of points are introduced, like the *mean* set, which manipulates the weight given to each sigma point to capture higher order moments.

$$w_o \in [0, 1[$$

$$\boldsymbol{s_0} = \boldsymbol{\mu} \tag{5.92}$$

$$\boldsymbol{s_n} = \boldsymbol{\mu} + \sqrt{\frac{D}{1-w_0}}\boldsymbol{l_n}, \quad n = 1, ..., D \tag{5.93}$$

$$\boldsymbol{s_{n+D}} = \boldsymbol{\mu} - \sqrt{\frac{D}{1-w_0}}\boldsymbol{l_n}, \quad n = 1, ..., D \tag{5.94}$$

$$w_j = \frac{1-w_0}{2D} = \frac{1}{N}, \quad j = 1, ..., 2D \tag{5.95}$$

Finally, the scaled set is analysed, which introduces further parameters that aim at scaling down the influence of unknown higher-order moments. This adds even more complexity to the problem and would require a great level of tunning, so this option is discarded. And, as such, the mean set is taken as the baseline option as the UKF's sigma-point set generation algorithm, as it encapsulates the most information without the addition of too much complexity or tuneable parameters.

5.2.7 Batch unscented Kalman filter

Similarly to what was proposed in the case of the EKF and UDU formulation EKF, denoted henceforth as UDUEKF, the batch version of the UKF will also be studied. Due to the way that the unscented transformation works, the adaptations required to turn the UKF from its sequential form into a batch one are much more reduced than in the case of a regular Kalman filter. To do so, one needs only to consider multiple measurements at once and treat them equally, despite corresponding to different times. As such, the derivation and the equations of this filter will not be presented here. Instead, a schematic by [Park et al., 2010] exemplifies the process that this filter undergoes, which should be easy to understand by anyone who has read and understood the equations of the last section on the sequential UKF.

It can be seen in Figure 5.2 that the general idea of this algorithm is similar to that of its sequential counterpart. Apart from handling the propagation of sigma-points to multiple different times and simulating multiple measurements, the equations used by both are similar, which is why the adaptation from one form into the other is as easy as aforementioned.

Despite being similar, here a slightly different notation is used. Instead of $^-$ and $^+$ to denote the prediction, or *a priori* updates, and the correction, or *a posteriori* updates, respectively, the symbols $^-$ and $\hat{}$ are utilised. The sigma points are denoted by $\boldsymbol{\chi}$ instead of $\boldsymbol{s}$, and their respective observation simulations by $\boldsymbol{\gamma}$ instead of $\boldsymbol{\bar{Z}}$.

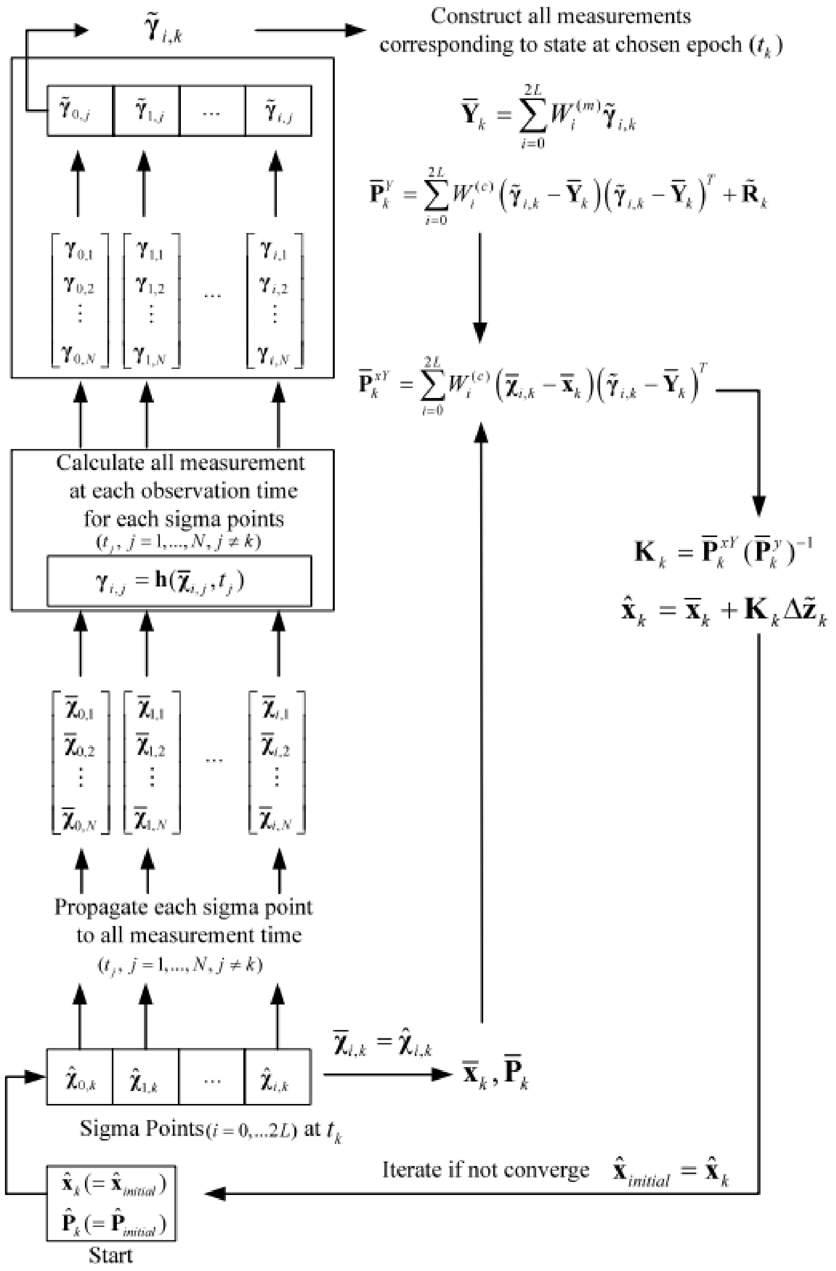

Figure 5.2: Schematic representing the algorithm undergone by the batch UKF. Retrieved from [Park et al., 2010].

$\hat{\boldsymbol{Z}}$, $\boldsymbol{S}$, $\boldsymbol{V}$ and $\boldsymbol{Q}$ are represented by $\bar{\boldsymbol{Y}}$, $\bar{\boldsymbol{P}}^{Y}$, $\bar{\boldsymbol{P}}^{xY}$ and $\boldsymbol{K}$, respectively. In spite of this difference in notation, the overall method of estimation is very similar to that of the sequential UKF.

Do note, additionally, that, in both the sequential and batch formulations of the UKF, no mention whatsoever of consider states is made. This is because these filters apply the Schmidt formulation not by changing the Kalman gain to leave the states and covariances of the consider parameters unchanged, but instead treats them as regular states, before undoing all changes made to them at the end of each step, as was discussed in Section 5.2.1. According to [Zanetti and D'Souza, 2013], these two methods are equivalent.

5.3 Trade-off

Even though the intent of this thesis will be to test all the six presented alternatives for filtering algorithms, a trade-off between them is always positive. This way, a prediction is made as to the expected best and worst methods, which can then be compared with the obtained results from the testing performed further ahead in Chapter 10.

The filter pool consists of four variations of the Kalman filter, all extended, and two of them using the UDU factorisation of the covariance matrix, whereas the other two rely on the unscented transform. From the former four filters, the UDU ones are supposed to have a higher computational stability, without anything else being different from the regular EKF and batch EKF, thus it is expected that their results are better or at least similar to their regular counterparts performance-wise. Regarding the computational load, however, the opposite should be true.

In terms of the sequential versus the batch formulations, both cases have pros and cons. The sequential will produce a new estimation at each time step with the most recently acquired image of the nucleus, meaning that regarding the available information at any instant, the sequential filters should have the upper hand. This is especially important during periods when the state changes very rapidly, as is the case right before the closest approach. Then, having frequent updates to the state estimation will be crucial, because the goal is to have the payload constantly pointing at the nucleus.

Despite this, the batch filters do have the advantage of processing a bunch of measurements all at once. What this can provide is a dilution of the effect of the noises affecting the measurements. Since these are random contributions to the measured values, the combination of a few of them can cancel their overall effect, making the final result less affected by these uncertainties. As such, in spite of having less frequent updates to their state estimations than the sequential filters, the batch filters should produce more reliable results.

With all this in mind, and considering that the most crucial and navigation-wise demanding section of the encounter phase is the period right before the closest approach, the sequential filters should have an edge over the batch formulations.

Finally, predicting whether the unscented filters will outperform the regular Kalman filters is no easy task. Of course that the former come with an increase in computation cost, but as aforementioned, these can convey even more information about the propagated set of points than just the mean and covariance thanks to the choice of sigma points. As such, the UKFs might be the favourites.

5.4 Additional techniques

As mentioned previously, there are strategies that can be employed to help the filters provide better state estimates, and be more robust against events such as outbursts. In this section, the three approaches suggested in the thesis proposal are discussed.

5.4.1 Altering navigation parameters

Here, the navigation parameters refer to the frequency of measurements, as well as the size of batches used, in case of the batch algorithms. The major factors to be considered are the minimum time to perform a measurement, which can be restricted by the hardware itself or by the exposure time required for each picture, the time to perform the image processing algorithm, and the duration of the filtering algorithm as well.

For the case of the batch algorithms, the batch size clearly impacts the results. Since a batch is a group of measurements collected to be processed all at once to produce one estimation update to the system's state, choosing a batch size too big will result in too much time passing between estimations. For a demanding period of the mission in terms of navigation, for example, right when the spacecraft is passing by the comet, such a choice can end up being pernicious for the overall performance, especially for the payload pointing capability.

Nevertheless, picking a batch size too small would just defeat the purpose of utilising batch filters in the first place, as processing a bunch of measurements simultaneously can even out the effects of noises, for example.

Despite this, with three batch filters plus all the variables which will be analysed further ahead in this thesis, introducing another degree of freedom in the form of varying batch sizes would increase the already high number of variables needing consideration. As such, in the analyses performed in this document, the batch size will always be equal to 5, meaning that these filters will perform corrections to their estimates once every 5 time steps, and process 5 measurement each time. This is seen to be a middle ground between a batch size too big and too small.

In terms of adjusting the measurement frequency, this strategy is indeed implemented. As discussed down the road during Chapter 8, when the time of the simulation reaches 20 hours, which is nominally when the closest approach to the comet takes place, the frequency of measurements is increased in the period right before the spacecraft flies by the comet's nucleus and is closer to it, meaning that more and better information can be extracted from the captured measurements, with a higher number of images being snapped. The goal is to use this interval of better images of the comet to the maximum, striving to achieve the best possible navigation estimation before the NAVCAM looses sight of the target once it leaves its FoV.

Regardless of the filtering method used, it would be relevant to study how changing these parameters alters the results of the simulations, and assess whether the predictions discussed here hold up when facing these scenarios.

5.4.2 Usage of other sensors to validate measurement

Whereas the last technique focused more on the interaction between the measurements and the filtering, relying on other sensors to validate the obtained data is strictly focused on the measurement part of navigation. The idea here is to use other instruments, which are not affected by the same disturbances as the NAVCAM, to confirm whether a significant change between images is caused by an actual change in the state of the system or by an unexpected perturbation phenomenon, for example.

The sensors which should be more relevant to assess this are the star trackers and the IMU. Since both of these can detect changes in the attitude of the spacecraft, their data can be crossed with the data from the obtained images to determine whether a sudden change in the position of the comet as given by the NAVCAM/IP route is also picked up by either of these. Should that not be confirmed, then a good case can be made for the existence of a disruptive phenomenon. To take this into account, the covariance associated with the measurement can be increased to reflect the corruption, or the image even removed completely, either leaving the batch with the remaining measurements, or, in the case of a sequential filter, skipping the correction step of that iteration.

Despite the discussion above, this method will not be used in this analysis, as it is similar to the one which is discussed right in the next section in terms of the results they might provide, while being significantly more complex. For this to work, other sensors would have to be modelled and implemented, and their information compared. This validation, however, as will be obvious in the next section, can be done in a much simpler way. Nevertheless, in future stages of the mission's development, this can certainly be a method to consider.

5.4.3 Filter measurement rejection technique

Capable of discerning whether a given measurement is adequate to be used by the filter or corrupted by an unpredictable event, a rejection technique can be critical to improve the robustness of the navigation. One possible method for this is known as innovation filtering, and is used in space applications such as in GNSS. It is simply based on the idea of comparing the residual of each measurement, that is, the difference between the predicted measurement and the actual value arriving from the sensors, with its own covariance, given by the innovation matrix [Groves, 2013]. The equations used to evaluate this are as follows.

$$\boldsymbol{\delta} = \boldsymbol{z} - \boldsymbol{G}(\bar{\boldsymbol{x}}, t) \tag{5.96}$$

$$\boldsymbol{W} = \boldsymbol{H}\bar{\boldsymbol{P}}\boldsymbol{H}^T + \boldsymbol{R} \tag{5.97}$$

$$\boldsymbol{z_i} = \frac{\boldsymbol{\delta_i}}{\sqrt{\boldsymbol{W}_{ii}}} < T \tag{5.98}$$

$$\tag{5.99}$$

Using the residual computed in the first equation and the innovation covariance resulting from the second, the inequality defined in the third equation is then used to determine whether the measurement fits within the chosen criteria. The left hand side of the expression corresponds to each normalised measurement innovation, which encapsulates the size of the residuals when compared to their respective covariances given by the innovation matrix. Thus, by adjusting the threshold T, it can be set to determine how big the residual needs to be for its associated measurement to be considered corrupted. Another alternative for this computation is to use the Mahalanobis distance, a multi-dimensional generalisation of the idea of measuring how far away a random variable is from its mean. In this case, it represents how much the residual vector is distanced from its mean – which is zero – again normalised by the respective innovation matrix, and is given by the equation below.

$$MD = \sqrt{\boldsymbol{\delta}^T \boldsymbol{W}^{-1} \boldsymbol{\delta}} < T \tag{5.100}$$

The Mahalanobis distance is powerful tool, typically used when studying rejection of outliers in a data set [Fung, 1999]. It thus makes sense that the two computations are similar, coinciding (that is, $\sum \boldsymbol{z_i} = MD$) for a diagonal innovation matrix. Even so, whereas the former analyses the measurement component by component, the Mahalanobis distance handles the whole vector as a whole, meaning that they can disagree in certain cases.

The tunning of this parameter T will be of the utmost importance. [Groves, 2013] mentions 3 as an example of a possible value to use, but in practice this value will have to be adjusted empirically. For this, the implemented values will have to be tested with the proposed navigation estimators and their performance evaluated for various thresholds, to determine whether the inclusion of this technique is beneficial, and if so which value improves the results the most.

Implementing this method requires just a small addition to the measurement correction phase of each of the filters, some of which already compute the residual and the associated innovation matrix, which can be used to make the process more efficient. In case of a positive result, the measurement should be disregarded, which means that the whole correction step is skipped entirely and the predicted state and covariance are used for the next iteration of the filter. However, in case too many images in a row end up being unusable, the system state prediction could become too inaccurate because of an absence of data. In the event of a persistent disruptive event, such as an outburst, right around the time of the closest approach, this could happen.

This will be the main procedure that will be implemented and tested to try and improve the navigation's performance by rejecting the worst measurements obtained by the NAVCAM.

Chapter 6

Numerical methods

When working on solving an engineering problem, it is always necessary to consider the numerical tools available to solve complicated pieces of maths. Whether it is solving a differential equation, thus requiring an integration method, or simply inverting a large matrix as efficiently as possible, the right choice of tools will help the problem at hand to be solved either more quickly or accurately, depending on what is most important. As such, this chapter is dedicated to laying out some of the most important numerical algorithms used in this study, especially the numerical integrator, the UDU decomposition algorithm, and the interpolator used to handle the ephemerides of the bodies in the scenario.

6.1 Integration

Almost no astrodynamics problem would be solvable without the aid of an integration method. Since the dynamics of these scenarios are defined by ordinary differential equations, some integration tool is required to solve it and determine the future state of the system given a certain initial condition. To assess what the best method is for each situation, a trade-off between speed and precision is always the central reason for the choice. With different orders, fixed or variable time step, there is a lot to choose from.

Nevertheless, as was discussed previously in Section 4.4, the Comet Interceptor mission presents quite an unusual set of acting forces, in the sense that most of them are small in magnitude, as well as being nearly constant for the duration of the trajectory. Compared to the distance between the spacecraft and the comet, the distance between the comet and the Sun is so large that for the whole encounter its gravitational pull changes very little in terms of magnitude and direction. The perturbation of the solar radiation pressure, similarly, only changes with the distance and direction of the Sun, both of which, again, do not change much during this phase of the fly-by. Finally, the cometary dust impacts are the origin of the force which varies the most over the whole of the trajectory, with the force at the beginning of the encounter being approximately 10^7 times smaller than at the C/A, thus the same cannot be said about this contribution. Yet, this should only be problematic around the closest approach, since the force varies with $1/r^2$.

Because of such simple dynamics, coupled with a small propagation period of just 40 hours – at most, which will be reduced, as discussed in the Chapter 8 – the highest order, variable time step methods should not be required for the simulation of this phase of the mission. The Runge-Kutta 4 (RK4) method is proposed as the baseline to be used to integrate the vehicle's dynamics, since these change very little over the integration period – with the exception of the cometary dust impacts, the impact of which is relatively small – thus not requiring a variable time step, as well as due to the fact that a fixed time step allows the integration to be performed between intervals that line up with the frequency of the measurements and the navigation system. To ascertain which time step should be used, a simple test can be performed, taking a variety of decreasing values and comparing the obtained trajectories with a baseline trajectory obtained with a tiny time step.

To clarify, in the context of the Comet Interceptor mission, the position error should be in the order of magnitude of a few metres, guaranteeing a margin between this value and the hundreds of kilometres considered in **MIS – 1.4.** for the B-plane delivery error. Additionally, it should be considerably smaller than the contributions of the forces considered in the problem. Since the solar radiation pressure was computed to perturb the spacecraft by approximately 600 m, then this requirement is verified.

The RK4 method is defined as follows [Süli and Mayers, 2003].

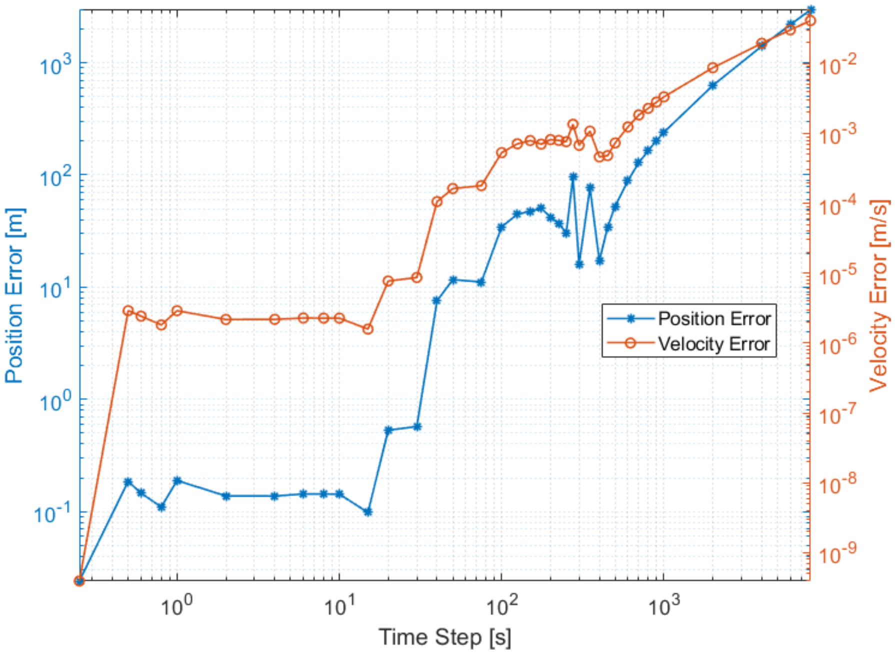

Figure 6.1: Maximum position and velocity error of RK4 integrations using multiple time steps, when compared to a benchmark time step of 0.1 seconds.

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{F}(t, \boldsymbol{x}), \quad \boldsymbol{x}(t_0) = \boldsymbol{x_0} \tag{6.1}$$

$$t_{n+1} = t_n + h \tag{6.2}$$

$$\boldsymbol{x_{n+1}} = \boldsymbol{x_n} + \frac{1}{6}h(k_1 + 2k_2 + 2k_3 + k_4) \tag{6.3}$$

$$k_1 = \boldsymbol{F}\left(t_n, \boldsymbol{x_n}\right) \tag{6.4}$$

$$k_2 = \boldsymbol{F}\left(t_n + \frac{h}{2}, \boldsymbol{x_n} + h\frac{k_1}{2}\right) \tag{6.5}$$

$$k_3 = \boldsymbol{F}\left(t_n + \frac{h}{2}, \boldsymbol{x_n} + h\frac{k_2}{2}\right) \tag{6.6}$$

$$k_4 = \boldsymbol{F}\left(t_n + h, \boldsymbol{x_n} + hk_3\right) \tag{6.7}$$

With the initial condition $\boldsymbol{x_0}$ and the defined time step h, for each step the coefficients k_1 through k_4 can be computed to determine the next state of the system $\boldsymbol{x_{n+1}}$, thus solving the problem posed by the first line of the equation.

Finally, do note that the choice of RK4 for this task was not only motivated by the previous analysis, because this is a method that is already implemented in FCS-ATOMIC, the software framework to be used for this thesis, which will be discussed in further detail in Chapter 7. As such, significant effort is saved in creating, verifying and validating the code.

Analysing now the time step used for the integration, Figure 6.1 shows that steps smaller than approximately 20 seconds show little improvements in maximum error. On the other end of the spectrum, the time step can be vastly increased until 10 minutes while still keeping the maximum error under about 100 metres. This is not a clear line of threshold for the position error, but since this work will deal mostly with distances in the order of kilometres or more, setting the maximum bar below that represents a cautious approach. Thus, in this large interval, many time step possibilities are available, depending on whether simulation time or accuracy are the focus. That discussion, however, will be left for later in Section 8.2

6.2 UDU decomposition

The UDU formulations of the EKF require the decomposition of the covariance matrix in the form $\boldsymbol{P} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^T$, where all matrices have dimensions $n \times n$. $\boldsymbol{U}$ is an upper triangular matrix with a diagonal of ones, whereas $\boldsymbol{D}$

is a purely diagonal matrix. The algorithm works as follows [Tapley et al., 2004].

```
for $j = n, ..., 1$ :
    $\boldsymbol{D}_j = \boldsymbol{P}_{jj} - \sum_{k=j+1}^{n} \boldsymbol{P}_{jk} \boldsymbol{U}_{jk}$    (6.8)
    for $i = 1, ..., j-1$ :
        $\boldsymbol{U}_{ij} = \frac{\boldsymbol{P}_{ij} - \sum_{k=j+1}^{n} \boldsymbol{P}_{jk} \boldsymbol{U}_{ik}}{\boldsymbol{D}_j}$    (6.9)
    end for
end for
                                                                    (6.10)
```

6.3 Interpolation

As done by JPL[1], the ephemerides of the various bodies being simulated can be represented by the Chebyshev coefficients that describe the polynomials used to fit their trajectories. In this way, the whole state of the body over time does not need to be saved. Instead, the trajectory is divided into arcs, a polynomial is fit to the data, and the resulting coefficients describing each arc contain the body's state over time. This also serves as an interpolation procedure, as the time at which the state is being accessed might have not been in the original information. But using the polynomial fit, an approximate state at each time of the interval can be known.

With information regarding the coefficients, the state can be retrieved as follows.

```
$\boldsymbol{T}(1) = 1$
$\boldsymbol{T}(2) = t$
for $i = 3, ..., n+1$ :
    $\boldsymbol{T}(i) = 2t \cdot \boldsymbol{T}(i-1) - \boldsymbol{T}(i-2)$
end for                                                             (6.11)
```

$$\boldsymbol{x} = \begin{pmatrix} \hat{a}(1) & \hat{a}(2) & \dots & \hat{a}(n+1) \end{pmatrix} \cdot \begin{pmatrix} \boldsymbol{T}(1) \\ \boldsymbol{T}(2) \\ \vdots \\ \boldsymbol{T}(n+1) \end{pmatrix}$$

In the equation above, t is the time corresponding to the desired state, scaled to the range [-1,1] with respect to the limits of the arc being considered; and $\hat{\boldsymbol{a}}$ is the vector of Chebyshev coefficients for the corresponding variable being interpolated. $\boldsymbol{T}$, in turn, is simply an auxiliary vector to aid the computation.

Do note that, similarly to the integrator method, RK4, the interpolation algorithm will not be discussed further throughout this thesis, nor will it be tested or go through V&V in the next chapter. These two tools are, then, included here as they play a major role during the propagation and integration of the vehicle's state over the mission's duration, without which the navigation performance could not be assessed. These were not developed or altered, and were already present in the software framework, but their understanding is, nevertheless, critical to establish the foundation upon which the future analysis is laid upon.

---

[1] Jet Propulsion Laboratory. JPL Planetary and Lunar Ephemerides. Retrieved June 4, 2021 from `https://ssd.jpl.nasa.gov/?planet_eph_export`.

Chapter 7

Simulation software FCS-ATOMIC

Developed by GMV under a contract with ESA, FCS-ATOMIC is a multi-functional modular framework, with the goal of allowing for feasibility studies to be performed along the early mission requirement design phase, thus saving cost. To do so, the toolbox is equipped with a suite of pre-developed functions to allow for a simple and quick analysis of the mission's performance in terms of spacecraft navigation, along with systems and operations constraints and cost.

The framework takes a combination of the Flight Dynamics Systems (FDS) team from ground and the spacecraft's own on-board autonomous GNC solution according to the authority hierarchy established by the user and the mission's operational constraints. Thus, analysing the impact on performance and cost resulting from different systems, operations and authority design choices also becomes possible. Additionally, FCS-ATOMIC allows for later mission simulations to be performed, with more numerous and detailed systems to be included, each capable of being turned on or off according to the user's required level of accuracy. Finally, Monte Carlo simulations are also included, to validate performances and assess the robustness of the developed Flight Control System (FCS) strategy.

Since the Comet Interceptor mission is being developed in this framework, the analysis performed to answer the research question that concerns this study will also be embedded in this environment. As such, this chapter is dedicated to, firstly, discussing the architectural design of the framework, where the major blocks involved in the navigation being studied are described schematically, to indicate what composes them and how the information flows within them. Secondly, the functionalities that had to be developed specifically for this thesis are enumerated.

Following that, a verification and validation effort is made to test whether these created blocks have been implemented correctly. This part mainly focus on the filters themselves as well as on the functions that compute the analytical derivative of the Jacobians of the dynamics' accelerations and camera measurements. Finally, a brief discussion is had regarding the optimisation of the simulation time with the usage of some of the developed tools.

7.1 Architectural design

In this section and its corresponding subsections, the primary goal is to present schematics that represent how the relevant parts of the code of FCS-ATOMIC for this specific thesis study function. This will be done in figures with the elements of the most important parts of the code represented, be they functionalities or

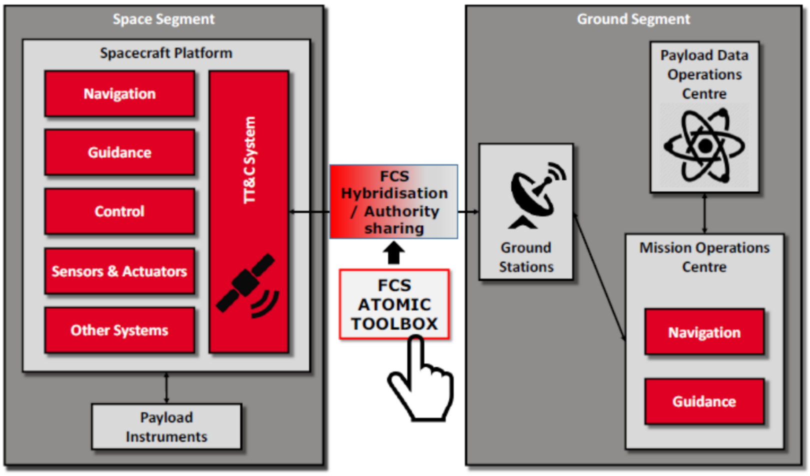

Figure 7.1: Representation of the hybridisation of the GNC and FDS components of the mission performed by FCS-ATOMIC. Retrieved from [GMV, 2020a].

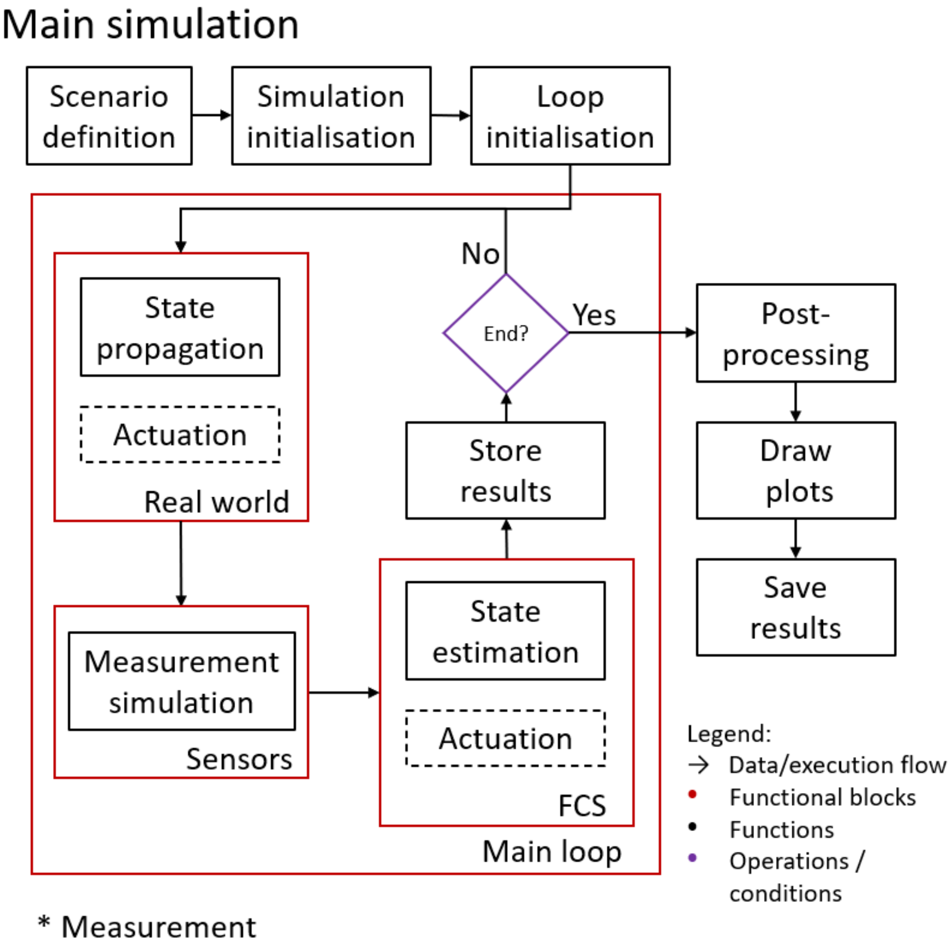

Figure 7.2: Schematic of the main loop of FCS-ATOMIC.

variables, and the flow of the algorithms represented by arrows connecting said elements. This will be done for the main loop of the code, which goes over the general functions executed along each time step of the simulation, as well as for the filter, dynamics, and measurement simulation functions. Since these are the main focus of the study done in this thesis, their inputs and outputs, as well as the flow of their execution will be represented to better explain how they were implemented.

7.1.1 Main loop

Before entering the specifics of FCS-ATOMIC, it is worth taking a step back and analysing the overall picture of its code as represented by very broad functionalities. In Figure 7.2 this is shown in the form of a schematic. After the scenario is defined and the simulation initialised, which is done by loading all the relevant parameters that characterise it from data files and initiating all the relevant variables that will be used during the simulation, the code's execution enters the main loop. Here, all major aspects of the simulation of a space mission can be modelled and studied, as is the modular nature of FCS-ATOMIC. However, for this study, there are only three relevant parts to this loop. Firstly, the real world is simulated, where the initial state of the spacecraft is propagated in time using the environment's characteristics and the spacecraft's parameters to simulate the forces acting upon the vehicle. In here, the actuation also happens, in case there is any. For the case of the scenario being studied, nonetheless, this is not what happens, which is why this block is surrounded by a dashed line.

After this, the real state of the spacecraft after a time period equal to the time step's duration is obtained, and this is used in the sensors block, where this state and the sensor's and environment's parameters are used to simulate what a true measurement would output if it were done in that environment and in that state. For the case of the Comet Interceptor, the NAVCAM's images are modelled, and the camera's characteristics, the spacecraft's position and orientation, and the various noises and biases affecting this measurement are put together to obtain the information that the vehicle itself then has access to, with the goal of determining its state and choose a course of action.

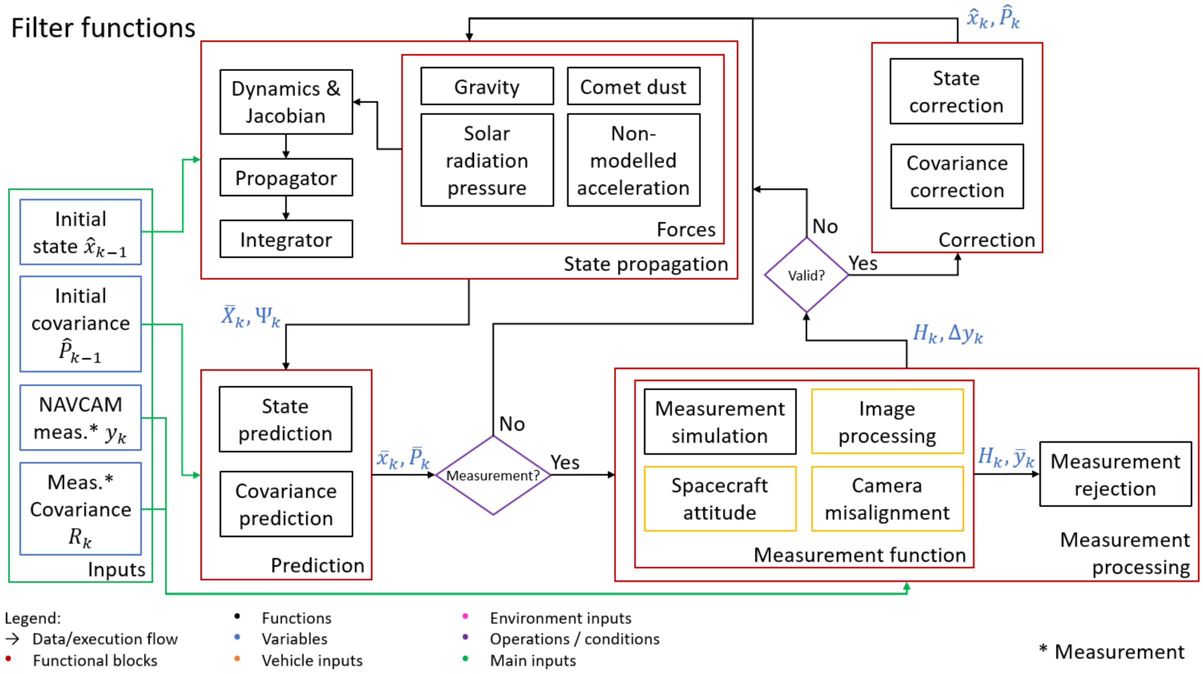

Figure 7.3: Schematic of the navigation filters used.

This is done by the FCS – flight control system – which has a last estimate of the vehicle's state and this new piece of information in the form of a camera measurement to work with and obtain a reliable estimate of where the spacecraft currently is, given not the complete set of parameters describing the vehicle and its surrounding environment, but instead the best available knowledge regarding these at the time of the estimation. With it, the navigation is performed, and obviously guidance and control would also be performed in this step to steer the vehicle into the desired trajectory according to the performed estimation, but since the Comet Interceptor does not actively change its course during the encounter, these blocks are not discussed.

Finally, after all this is done, the results of the relevant computations are stored, and the loop checks for its end condition, which is reaching the total stipulated simulation time, in the case of the Comet Interceptor. If it has not yet reached its end, the loop restarts at the real world block. If the final time has been reached, then a post-processing phase is entered to treat the simulated data according to the commands of the user, the relevant results are plotted, and all the information regarding the simulation's inputs, numerical results and corresponding plots is saved in data files.

7.1.2 Filter functions

Inside the FCS block, the star of the show are the filters used to perform the navigation estimates. Despite differing amongst each other on the details of how this estimation is performed – which can be studied in detail in Chapter 5 – all filters have a general outline that is common regarding the order of the operations which are undertaken. This is what is represented in Figure 7.3.

The main inputs of these functions are the previous state estimation and its respective covariance, together with any new measurement information and its covariance. With the previous state estimate, the filter uses the information about the vehicle and the environment to propagate its position and velocity to the time of the measurements, as well as the covariance of the state. This is the prediction step, or the *a priori* prediction. If there is no measurement to be added to this algorithm, then it stops here, returns the propagated information, and only returns during the next step.

In case there are measurements to be used, on the other hand, the propagated state, along again with the vehicle and environment information that the filter possesses is used to simulate what said measurement should look like. Before using it to update the estimation, however, the measurement rejection is performed to assess whether the simulated and real measurement disagree significantly more with one another considering the expected standard deviation of this residual.

Finally, if this check is passed, the second main phase of the filter is entered, where this measurement information is used to update the previous state and covariance estimations. This is called correction, or *a posteriori* prediction. After it is done, the filter returns the estimated values and the cycle restarts at the next

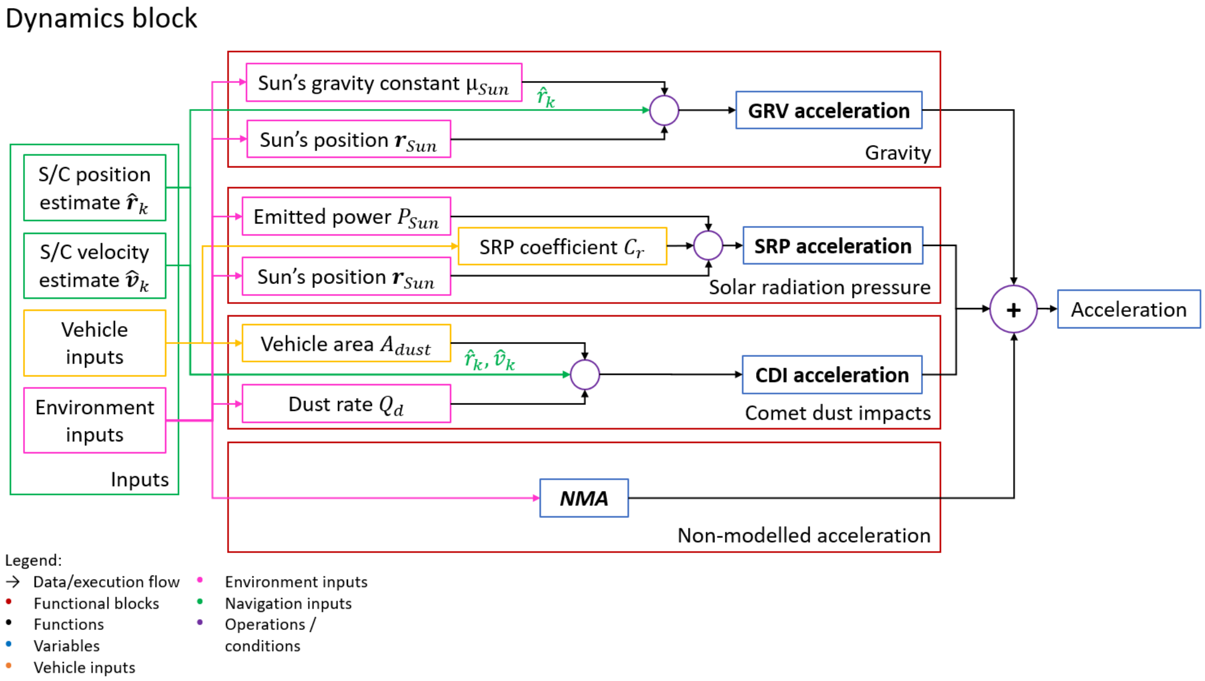

Figure 7.4: Schematic of the dynamics function used to compute the accelerations acting on the spacecraft.

time step.

7.1.3 Dynamics block

Within either the real world block of the main loop of the function, or inside the filters when the previous state estimate is being propagated in time, the dynamics block is crucial to determine the forces acting on the spacecraft, thus defining its future trajectory. In it, both the state of the spacecraft and multiple environmental and vehicle parameters are used to compute these forces, according to what is represented in Figure 7.4.

In this simpler schematic, the major variables involved in the computation of each of the forces acting on the vehicle are represented. The details of how these forces are modelled and what each variable represents are more carefully laid out in Section 4.2. In here are represented the different major inputs required to compute the value of each force that the function uses. Of course that, when running the real world block, for example, the dynamics function will have access to the real values of all these parameters, whereas when the same algorithm is run inside the FCS block, the vehicle has imperfect information about the majority of these parameters, meaning that the inputs vary according to the context where the function is called. This lack of perfect knowledge is, as previously discussed, why the Schmidt formulation of the filters was used.

7.1.4 Measurement block

The other crucial function to run the Comet Interceptor's navigation loop is the measurement simulation. This is where the NAVCAM's images are modelled, and from them the position of the target with respect to the spacecraft is then extracted. For this, not only are the vehicle's position and orientation needed, which along with the camera's characteristics can be used to derive the perfect measurement containing the true position of the nucleus, but also many other terms regarding the various noises and biases affecting the measurement are included. What terms are needed to obtain these, how they are computed, and then used altogether to obtain the final simulated measurement is the focus of Figure 7.5.

In this image, the misalignment bias of the camera, the bias resulting from the existence of outbursts, as well as the errors arising from the image processing algorithm are all represented, and the various inputs and parameters required to simulate them laid out. With them, the respective biases and noises are computed and then added to the perfect measurement that is the baseline.

Similarly to the dynamics block, the measurement block is invoked in different stages of each iteration of FCS-ATOMIC's main loop. After the real state is propagated forward by one time step, the real measurement is simulated, using all the accurate parameters regarding the vehicle and environmental variables necessary to describe the camera's behaviour. Inside the FCS block, on the other hand, where the measurement is simulated

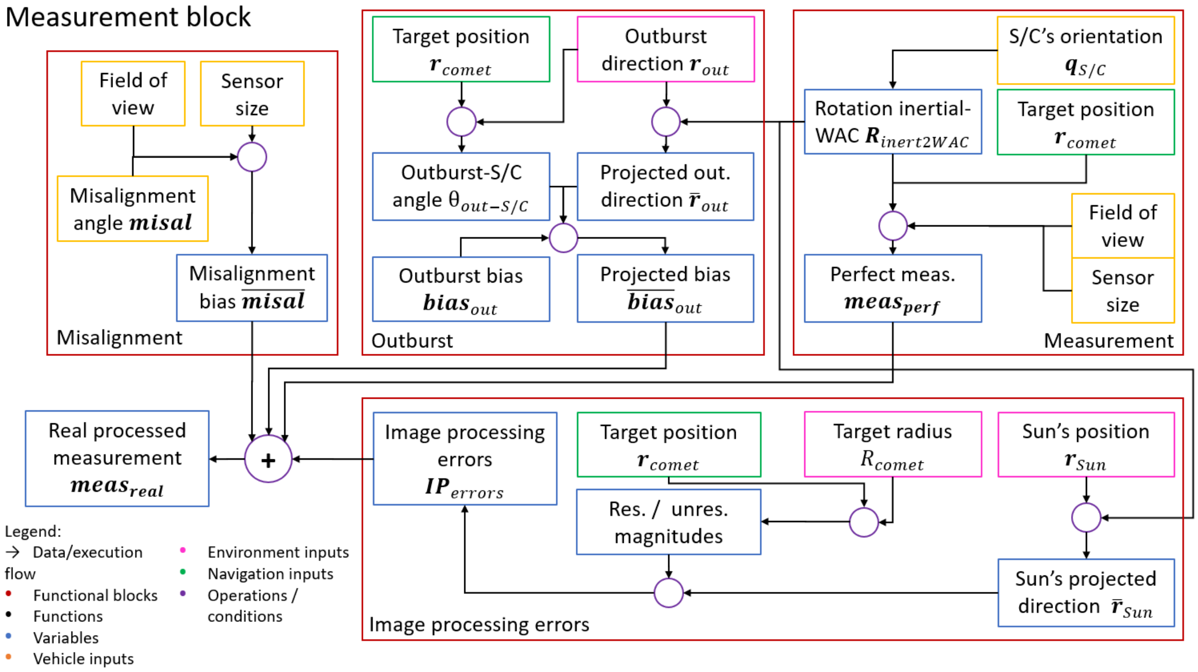

Figure 7.5: Schematic of the measurement function used to simulate the information captured by the NAVCAM.

to compare the result with the real input, the information available to this system is, as aforementioned, imperfect, meaning that the system either has an incorrect knowledge of the biases and noises, or has an idea of how large they can be in the form of a standard deviation and mean, which is the case of the Comet Interceptor, or even has no information about their existence at all, which is what happens in the case of the outbursts when measurements are being simulated within the estimators. In both cases, the function's logic is the same, and only the input databases change. Whereas the real world function has access to all the true values, the FCS only knows those values given to the spacecraft beforehand.

7.1.5 Developed functionalities

From the architecture described above, it is important to distinguish from the FCS-ATOMIC framework those blocks of coed which were specifically developed for the sake of this thesis. These blocks can be divided into three parts.

First, and most important, are the filters themselves. Apart from the already available UKF, all the remaining filters had to be developed from scratch. These fit inside the FCS block represented in Figure 7.2. They are tasked with receiving standardised inputs, transforming them into usable information about the current state estimation and new measurement information, using it to update the estimation, before then shaping the results back to fit the mold of the block's outputs.

Inside the filter block, there are two sub-blocks where the Jacobians of the dynamics and measurement functions are computed. These are not needed in the unscented filters, but for all other Kalman filters they are a key part of the algorithm, and thus had to be developed. As previously described, these are tasked with computing the derivative matrices of the acceleration and measurement vectors with respect to each estimated and considered state.

Finally, the last block to be developed is related to the outbursts and their simulation. Inside Figure 7.5, there is a sub-block responsible for simulating the impact of these phenomena. According to its direction, size, and the position off the spacecraft with respect to the target, a bias is computed and then added to the processed measurement.

For all these blocks of code, the algorithms need to be developed and integrated into FCS-ATOMIC, since all of them are sub-blocks of other pre-existing functions. Then, finally, these have to be properly tested.

7.2 Verification and validation

One of the most crucial efforts in developing the tools needed for this navigation study is performing their verification and validation. The aim of this is to assess whether the algorithms have been correctly implemented

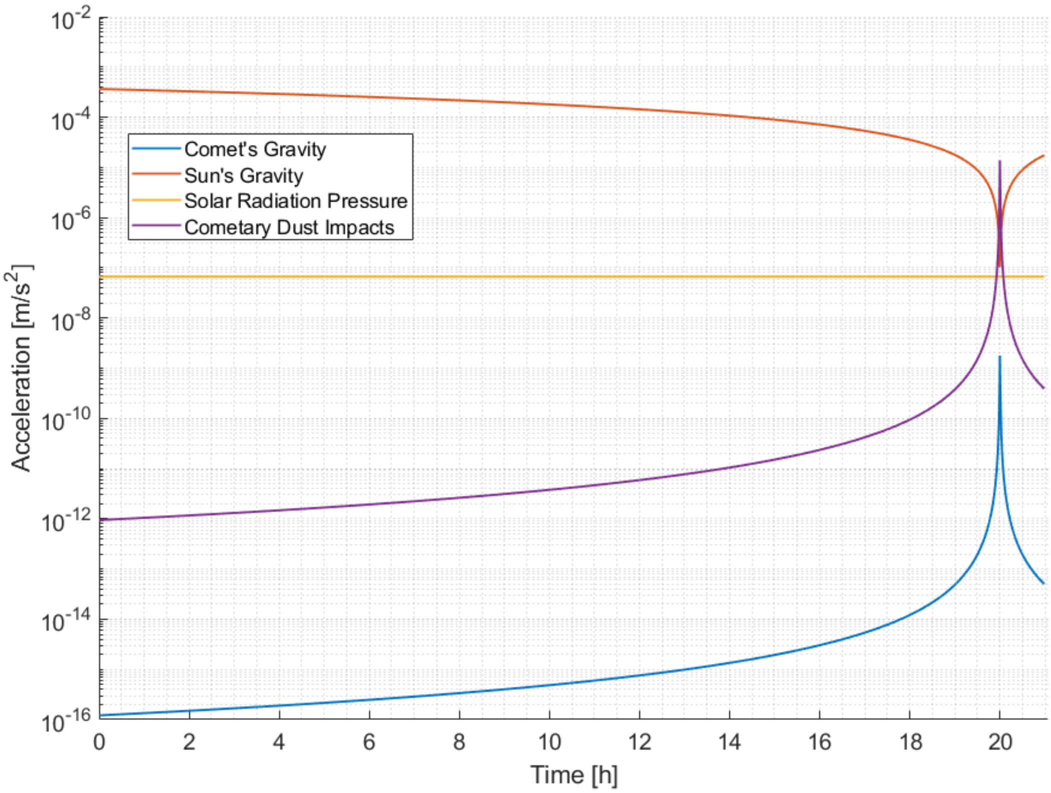

Figure 7.6: Magnitude of the acceleration caused by each of the four simulated forces over time. Only the first 21 hours are present, since the forces should be symmetric in the outgoing leg of the fly-by.

and work as intended. There are two main categories defining the major tools that were implemented. The first were the functions that compute the Jacobians of the dynamics, $\boldsymbol{A}$, and of the sensor observations, $\boldsymbol{H}$. These give the sensitivity of each component of the acceleration and pixel location of the comet, respectively, to changes in the system's states and parameters, and are required to implement the EKF and UDUEKF. The other functions are the filters themselves, their input and output management, as well as all the mathematical expressions within them, as described in Chapter 5. Unfortunately, for the majority of these, there is no dataset to be used to compare with the results of the implemented functions. As a result, these efforts of V&V will have to rely on other strategies to confirm the correct functioning of the created code.

7.2.1 Dynamics acceptance test

Despite having been developed previously to this study by GMV, the dynamics are a central part of the simulations. Because of this, they are here the target of some acceptance tests to ensure that the computed values are sane. To do so, the accelerations produced by each of the four forces described in Figure 7.4 is graphed here, as well as the real trajectory taken by the spacecraft. Do note that all results gathered in this section are produced from a simulation with seed 0.

Analysing Figure 7.6, the orders of magnitude of the forces and the shapes of the graph are in accordance to the established in Chapter 4. Firstly, both the cometary dust impacts and the comet's gravitational attraction vary with $1/r^2$, and the results reflect that, with both forces greatly increasing around the 20 hour mark, when the spacecraft flies by the comet. However, as expected, the magnitude of the comet's gravity is by far the smallest force of the bunch. Secondly, in terms of the solar radiation pressure, it can be seen that its effect on the vehicle is nearly constant, as the distance from the Sun varies very little in relative terms during the mission. In terms of impact, this force is the second biggest of the four considered. Finally, the Sun's gravity is the star of the bunch, being by far the most impactful. Its magnitude profile is almost the opposite of those of the comet's gravity and its dust impacts, because, as the spacecraft approaches the comet, the Sun's gravity is nearly the same on the spacecraft as it is on the comet, the central body being considered, meaning that the two forces cancel out partially and the resulting intensity lowers. This is precisely what is seen in the respective curve in the graph.

With these four forces acting on the vehicle over the fly-by duration, the real trajectory can be obtained and plotted.

In Figure 7.7, the real trajectory is represented as a blue line. In a dashed orange line is the trajectory tangent

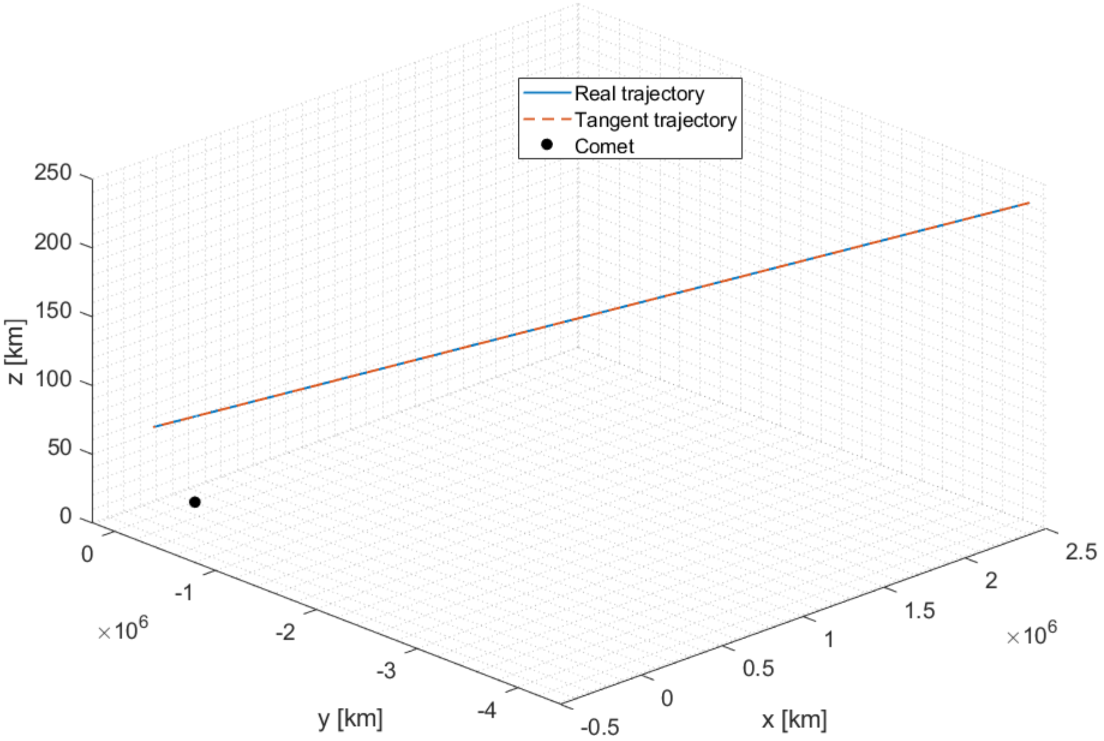

Figure 7.7: Real and tangent trajectories in 3D space. The comet is represented as a black dot at (0,0,0).

to the initial state of the spacecraft. This is to say that the orange line is generated by taking the initial position and propagating it in the direction of the initial velocity. This image, then, puts a perfect straight line trajectory against the real trajectory performed by the spacecraft. From studying this graph, some conclusions can be drawn. Firstly, if no distinction between the lines were to be made visually, it would be absolutely impossible to distinguish one from the other, as they are so close to being coincident. Secondly, even if the two were visually identified, without the legend it would also be hard to know which one is a perfect straight line and which one is not. What these two facts show is that, indeed as mentioned many times throughout the first chapters of this document, the real trajectory of the spacecraft can be well approximated by a straight line.

Nevertheless, the distance between the two curves, despite not being visible, is real, and can be computed. At the closest approach, the two trajectories differ by 622 km. Given that the real trajectory passes by the comet at 1314 km from it, this difference represents almost a 50% relative difference. This goes to show that, despite being very close to a straight line, the real trajectory must take the various forces acting on the vehicle into account, because, as it flies by the target, the difference made by the inclusion of these forces is relatively large.

7.2.2 Dynamics & measurement Jacobians

As already stated, these matrices refer to the derivative of each component of acceleration or of the measurements taken by the navigation camera with respect to all the states and parameters being considered. Whereas the accelerations have 3 distinct components, each camera measurement is comprised by the nucleus' position on the sensor plane, thus having 2 components. On the other hand, the system has 6 state entries being estimated, 3 for position and 3 for velocity, plus all the 14 dynamical and measurement-affecting parameters. Obviously, each set of parameters will only influence the corresponding matrix, meaning that the dynamical parameters will have no direct impact on the measurements, and thus their components in the Jacobian $\boldsymbol{H}$ are 0. With 3 and 2 outputs, and a total of 20 inputs, the Jacobians $\boldsymbol{A}$ and $\boldsymbol{H}$ are 3×20 and 2×20 in size, respectively.

The computation of these two matrices was implemented together with the computations of the accelerations and measurements they relate to, and to confirm that the created code runs as supposed, the simplest method is to compare the results against a numerical computation of the derivative. This can be done by using the finite differences method. This allows to compute an approximation of the true analytical derivative as follows

$$\frac{\partial f}{\partial x} = \lim_{\Delta \boldsymbol{x} \to 0} \frac{f(\boldsymbol{x} + \Delta \boldsymbol{x}) - f(\boldsymbol{x})}{\Delta x} \approx \frac{f(\boldsymbol{x} + \Delta \boldsymbol{x}) - f(\boldsymbol{x})}{\Delta x} \tag{7.1}$$

Starting with the definition of derivative of a function, a reasonable approximation of its analytical value can be obtained numerically. This will, of course, depend vastly on the specifics of the function itself and the term with respect to which the derivative is computed. Ideally, a small value of Δx should be used, to make the numerical value as close to the limit as numerically possible. However, depending on the variable being analysed, it must also be big enough to change the function's output, and not only that, but also make that change big enough to minimise the truncation error from dealing with numerical computations. This is a very particular concern for the case of the Comet Interceptor's analysis, since, as has been mentioned many times throughout this analysis, some of the parameters considered have a small impact on the results, and their derivatives will naturally be harder to approximate.

As a result, the implemented method was to try various values of Δx, spanning multiple orders of magnitude around the nominal value of each of the parameters, and comparing the resulting numerical approximation with the analytical derivatives. It could be argued that this approach is rather heuristic, but it would be extremely unlikely that a computed numerical value would show close proximity to a wrong analytical derivative value. That being said, attention should be paid to this happening with specific values, such as 0, as numerically it is easy to obtain a small enough Δx to cause no difference to be observed, which happens to coincide with a hypothetical 0 on the analytical derivative which is wrong. Luckily, these cases are easy to spot, as determining which variables alter the results and which do not is a simple task when the system is well known.

It is also important to point out that this experiment was performed at different points of the trajectory, as over time the spacecraft's state with respect to the nucleus changes quite significantly, thus altering the derivatives. To account for this, two situations were considered, one at the beginning of the simulation, that represents the majority of the simulation up until the closest approach – and after closest approach as well, due to symmetry – and another when the vehicle is passing by the comet.

Following this procedure yielded positive results, as the relative differences between the two methods of computation have small errors, up to only a few percent. This is seen in the aforementioned case of variables with reduced impact in the accelerations or the measurements, making the numerical errors more prevalent. In the case of other variables, such as the position of the spacecraft, which has a bigger impact on both the accelerations and measurements, its impact is several orders of magnitude smaller than 1%; or the camera misalignment biases, which alter the camera's information linearly, the obtained results perfectly coincide with the numerical approximation, since for a linear derivative, the approximation corresponds with the analytical results.

The analytical values also agree with the expectation that dynamical parameters would not affect measurement parameters and vice-versa. The only parameters that transverse both cases are the spacecraft's position components, which clearly influence both the accelerations felt by the vehicle and the measurements taken by the on-board camera. Even the velocity terms only affect the dynamics, being related to the strength of the cometary dust impact force.

To complement this discussion, an example of the Jacobian values is considered, as well as the numerical approximations' errors. Table 7.1 these two matrices are shown as computed at 19 hours and 55 minutes of simulation time. The seed used for the simulation was 0. Do note that the parameters are divided into three categories, visible by their colour. Green states only impact the dynamics, whereas blue states only impact the measurements. The red parameters, on the other hand, impact both, and thus are put in a different group.

Analysing this tables, firstly, as expected, the time derivative of the state position, $\boldsymbol{r}$, is equal to $\dot{\boldsymbol{r}}$, without any other components affecting it. Secondly, the groups of each parameters correspond perfectly with the results, with the green states showing no derivative values in $\boldsymbol{H}$, and the blue states showing no impact on $\boldsymbol{A}$. Additionally, the non-modelled accelerations have a dynamics derivative equal to the identity matrix, which makes sense, given that these are three terms that are directly added to the acceleration values. Finally, the camera misalignment derivatives are also diagonal, as expected, since it would not make sense for a misalignment in one axis to affect the other.

As for the other terms, however, there is very little analysis to be made, except to try and match the results against the respective numerical derivatives. Doing so yields the results in Table 7.2, which presents the relative differences between the analytical and numerical terms in percentage.

The first thing that stands out from the results is that some terms have zero relative error. While that is normal for terms which do not affect either the dynamics or the measurements, it also happens for some terms which do affect the Jacobian in question. For example, the NMA terms have zero error in their dynamics Jacobian components. This is because their influence is mathematically so simple, as they are just added to the total acceleration, that their derivatives are given exactly by the numerical approximation used here, making the final numerical result coincide with the analytical one. This, in fact, happens for the comet's gravitational parameter, μ_{comet}, for the cometary activity, Q_d, as well as for the camera misalignment and image processing biases, $\textbf{Cam}_{\text{bias}}$ and $\textbf{IP}_{\text{bias}}$, respectively.

For the case of the terms which end up non-zero errors, these vary significantly in terms of order of

Table 7.1: Dynamics and measurement Jacobians computed for seed 0, and 5 minutes before the 20 hour mark. Green parameters only affect the dynamics, whereas blue parameters only affect the measurements. The red entries impact both matrices.

		$\boldsymbol{A}$						$\boldsymbol{H}$		
Parameters		x	y	z	v_x	v_y	v_z		**mea**$_\text{x}$	**mea**$_\text{y}$
$\boldsymbol{r}$	1	0	0	0	1.1e-13	1.0e-15	5.2e-20	1	4.6e-05	7.7e-08
	2	0	0	0	1.2e-15	-5.7e-14	4.0e-20	2	2.2e-05	-1.2e-07
	3	0	0	0	-2.9e-18	5.5e-18	-5.4e-14	3	1.1e-08	-5.1e-05
$\dot{\boldsymbol{r}}$	4	1	0	0	-6.3e-13	2.1e-13	-4.6e-18	4	0	0
	5	0	1	0	2.1e-13	-9.1e-13	8.4e-18	5	0	0
	6	0	0	1	-4.6e-18	8.4e-18	-5.1e-13	6	0	0
μ_{comet}	7	0	0	0	-6.5e-16	1.2e-15	-3.4e-18	7	0	0
Q_d	8	0	0	0	1.7e-12	-3.1e-12	6.7e-17	8	0	0
$\boldsymbol{r_{Sun}}$	9	0	0	0	-2.8e-17	-2.8e-17	7.1e-20	9	-2.0e-23	-1.6e-22
	10	0	0	0	-2.6e-17	1.3e-17	-5.7e-18	10	-8.4e-13	-6.7e-12
	11	0	0	0	-1.1e-17	5.7e-18	1.3e-17	11	1.9e-12	1.6e-11
SRP	12	0	0	0	1.6e-02	-1.5e-06	7.0e-09	12	0	0
R_{comet}	13	0	0	0	0	0	0	13	-7.6e-05	5.0e-06
NMA	14	0	0	0	1	0	0	14	0	0
	15	0	0	0	0	1	0	15	0	0
	16	0	0	0	0	0	1	16	0	0
Cam$_\text{bias}$	17	0	0	0	0	0	0	17	1.1e+03	0
	18	0	0	0	0	0	0	18	0	1.1e+03
IP$_\text{bias}$	19	0	0	0	0	0	0	19	-1.0e+00	-2.2e-04
	20	0	0	0	0	0	0	20	-2.2e-04	1.0e+00

magnitude, with the majority of them being around the 10^{-5} mark or smaller. Some cases are bigger than that, with the maximum reaching 10^{-1}, but since these values are in percentage, all results end up with errors below 1%. Thus, these results strengthen the confidence in the implementation of the analytical version of these Jacobians.

7.2.3 Navigation filters

The most vital part of the navigation are the algorithms used to combine propagation of already existing information with newer data from camera measurements. These will be the determining factor in terms of navigation error and capability of pointing the payload at the target nucleus. The two most important parts of their implementation are the transcription of the equations that define these filters, as laid out in Section 5.2, as well as the integration of these computations into the FCS-ATOMIC framework, guaranteeing a consistency in the inputs and outputs.

7.2.3.1 UKF

Although not having been developed for the purposes of this thesis, the UKF is a central piece of the filters being studied. It represents a major alternative to the extended Kalman filters, along with its batch counterpart, which might be able to present better navigation results. Because of its importance, its results will be briefly tested and analysed. Because this filter has previously been implemented by GMV, and, in fact, used in the analysis of multiple space endeavours, it will, unlike the other filters, not be submitted to verification and validation.

To perform a simple acceptance test, this filter is run with a baseline scenario (meaning that no outbursts are present) generated with the seed 0. From the obtained results, the navigation error is shown, that is, the error between the estimated and real position and velocity of the spacecraft.

The graphs in Figure 7.8 show these errors. The vectors are represented in the along-track, cross-track, and radial components, respectively. The radial component points towards the target, the along-track component is perpendicular to it while being on the plane of the target's orbit around the Sun. Finally, the across-track component is perpendicular to these two.

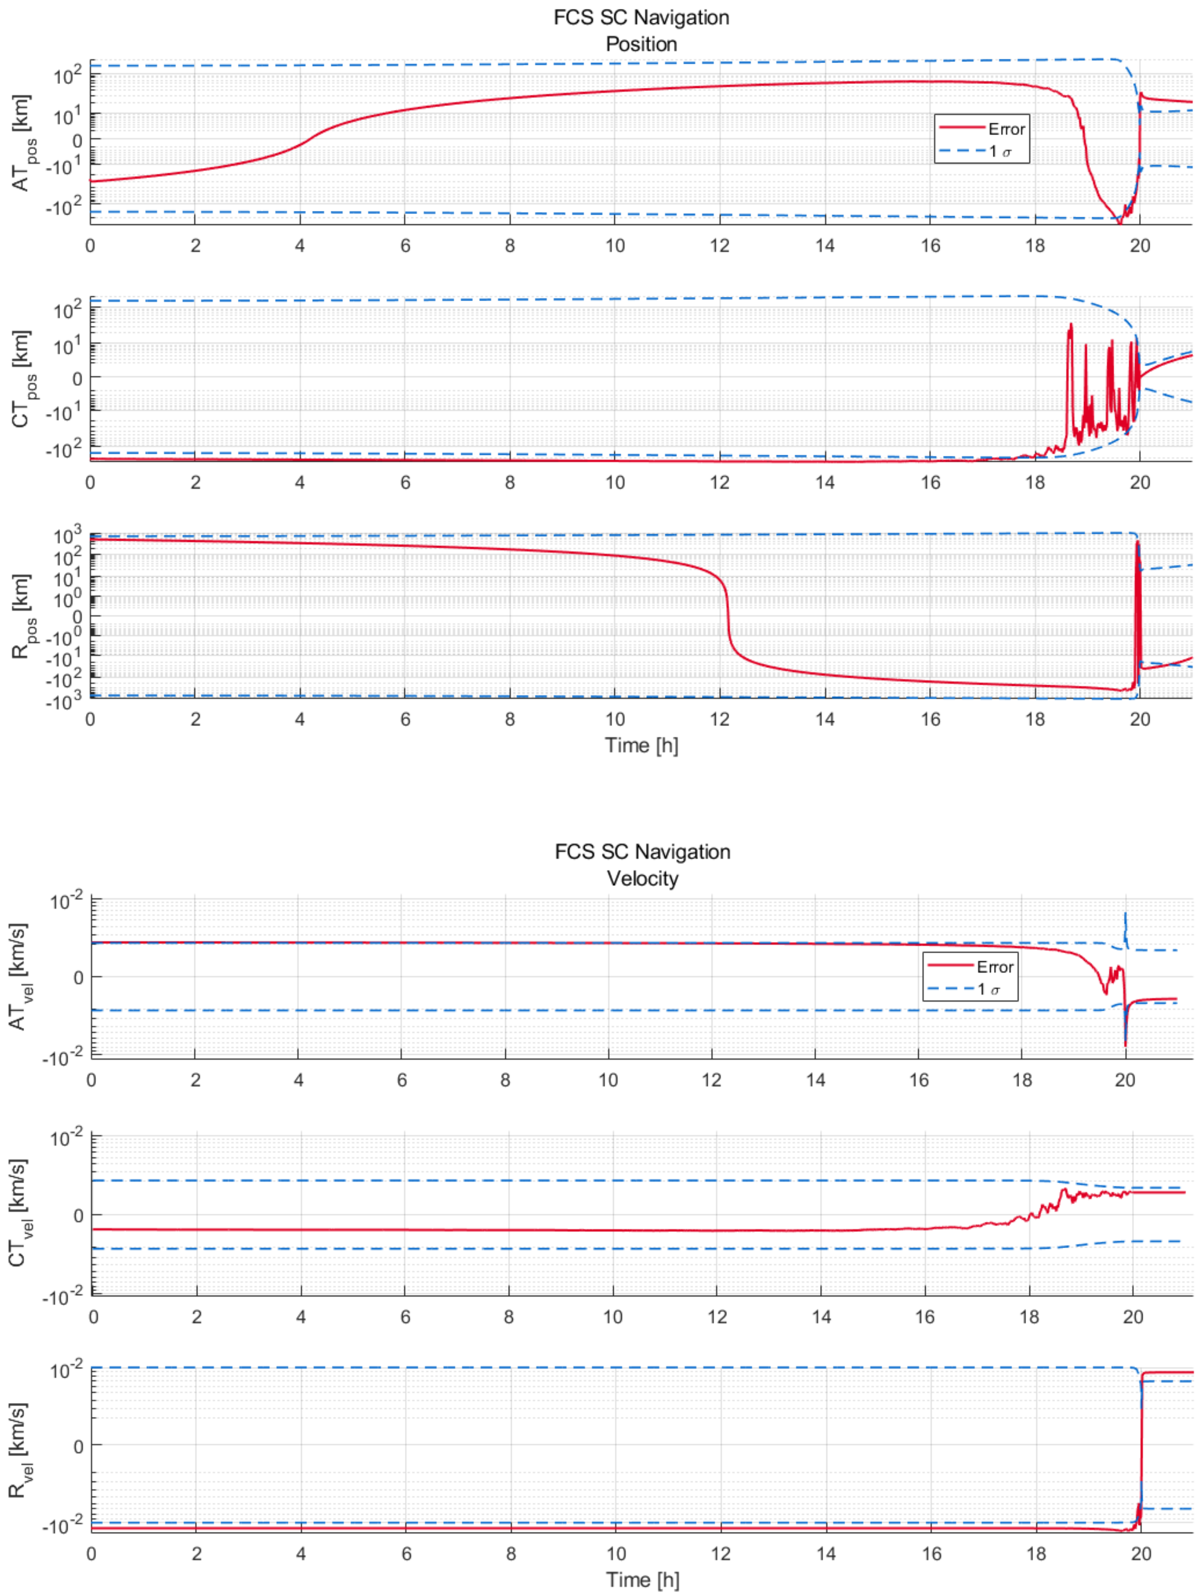

Figure 7.8: Navigation estimation error for the seed 0 scenario with the UKF. At the top is the position error, and the velocity error is at the bottom.

Table 7.2: Relative differences, in percentage, between the analytical dynamics and measurement Jacobians, and the numerical derivative approximation values.

		$\boldsymbol{A}$							$\boldsymbol{H}$	
Parameters		x	y	z	v_x	v_y	v_z		**mea**$_x$	**mea**$_y$
$\boldsymbol{r}$	1	0	0	0	4.9e-07	1.9e-07	7.7e-06	1	2.7e-10	1.7e-08
	2	0	0	0	9.7e-05	3.4e-07	1.5e-05	2	5.0e-08	1.0e-09
	3	0	0	0	3.3e-01	2.4e-05	4.6e-12	3	3.7e-07	4.1e-11
$\dot{\boldsymbol{r}}$	4	0	0	0	7.2e-08	1.4e-06	1.8e-05	4	0	0
	5	0	0	0	6.3e-08	2.4e-06	1.2e-05	5	0	0
	6	0	0	0	2.6e-03	6.9e-02	4.2e-10	6	0	0
μ_{comet}	7	0	0	0	0	0	0	7	0	0
Q_d	8	0	0	0	0	0	0	8	0	0
$\boldsymbol{r_{Sun}}$	9	0	0	0	1.3e-05	7.0e-08	9.0e-07	9	2.7e-02	1.8e-02
	10	0	0	0	7.9e-05	5.6e-08	3.2e-09	10	3.2e-08	4.3e-08
	11	0	0	0	2.7e-04	3.9e-07	3.4e-09	11	1.4e-08	1.0e-08
SRP	12	0	0	0	2.3e-08	3.0e-06	5.5e-04	12	0	0
R_{comet}	13	0	0	0	0	0	0	13	1.2e-10	1.9e-10
NMA	14	0	0	0	0	0	0	14	0	0
	15	0	0	0	0	0	0	15	0	0
	16	0	0	0	0	0	0	16	0	0
Cam$_{bias}$	17	0	0	0	0	0	0	17	0	0
	18	0	0	0	0	0	0	18	0	0
IP$_{bias}$	19	0	0	0	0	0	0	19	0	0
	20	0	0	0	0	0	0	20	0	0

The first aspect to note is the lack of evolution over the initial 18 hours or so. The position error does change signs for the along-track and radial components, but only because the position error and the velocity error had opposite signs. This means that either the FCS thought that the spacecraft was ahead of the real trajectory but with a slower velocity, meaning that eventually the real trajectory catches up to the estimation and the position sign is reversed; or that the spacecraft was behind the real trajectory but with a faster velocity, in which case the opposite phenomenon happens, but the sign inversion is still present.

Apart from these, very little happens over the first hours, with the covariance slowly rising and the error staying pretty much unchanged, in terms of magnitude, at least. This is because the filter knows that the spacecraft is still very far away from the target, and thus the NAVCAM's images are very sensitive to noises and biases. As a result, the UKF – as well as all other filters, for that matter – barely relies on the images until the vehicles approaches the comet enough for its information to be more reliable than the previous estimations. Then, around the 18 hour mark, these errors start evolving because of this.

Between then and the 20 hour mark, around which time the spacecraft reaches its closest approach to the target, the state estimation slowly improves, and the 1σ values for each component become smaller. After that, no more measurements are available, and the navigation becomes an open loop.

This represents the general evolution of the navigation error over time not only for this specific filter and seed, but for all the cases. The initial error determines the evolution over the majority of the first 20 hours, and then the measurement updates kick-in, reducing the magnitude of the error to a few kilometres. And it is in this period that, in fact, the different filter choices make a visible difference in the results. Whereas in the initial period all the filters correctly assess that the measurements are too unreliable to use, and thus all algorithms output results similar to an open-loop navigation case, for the final hours the different filters use the available camera information differently, and thus their performances vary.

7.2.3.2 Batch UKF

Having a pre-developed sequential UKF filter, the development of a batch version becomes quite simpler, as these share most of their formulation. The only difference between the two is the multiple measurements that need to be processed at once, meaning that many instances of the same computation are repeated. Thus, re-utilising the code from the already available filter leaves much less room for error. However, to test the result,

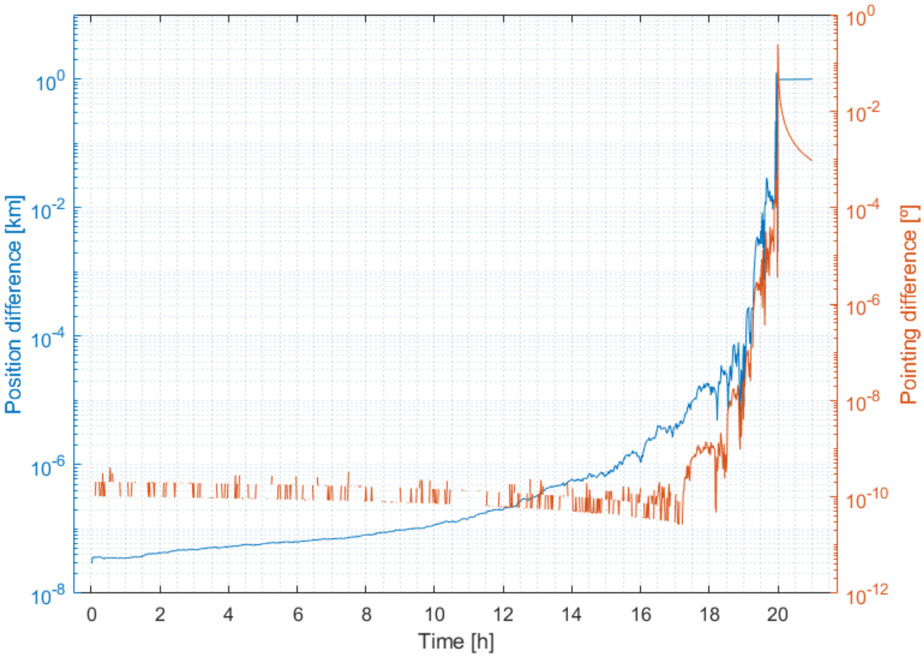

Figure 7.9: Position and pointing difference between the sequential UKF and batch UKF for seed 14.

the two filters can be compared by setting a batch size of 1, meaning that both filters should present the same results. By using the same random generating number seed it can be guaranteed that all aspects of the two scenarios are exactly the same apart from the filters themselves.

Despite these predictions, Figure 7.9 shows that the two filters have some different results in terms of both position and pointing accuracy. Despite being small in both metrics, reaching about 1 km and 0.2 degrees at the peak, these values are still small enough in the context of the mission to conclude that the observed disparity is too significant. In fact, it arises due to a difference between the way that the two filters are implemented. Whereas in the sequential filter, for simplicity, some measurement noises are taken to be a part of the state of the system, despite not being biases, in the batch formulation this is not done. As a result, in the former the biases and noises are mixed together when computing the measurement covariance, whereas in the latter only the biases are used, since otherwise the noises at the different times of the measurements would appear as though they were related, when in fact they cannot be as they are noises. Because of this, the noises' impact on covariance is computed just as it is done for the non-unscented filters, by using their respective covariance and the measurement's Jacobian matrix. This difference in the way these noises are treated is the source of the small difference observed between the two filters. Indeed, if this disparity is removed and both filters are made to handle these noises similarly, the results of this experiment are precisely equal between the two, ignoring some residual numerical error.

Re-establishing the batch size to 5, the results for navigation error are shown in Figure 7.10. Comparing these with those from Figure 7.8, the similarities are clear, as the filters show results close to one another. However, the batch filter does show more discrete results, as the estimations are only updated once every five iterations of the filter, resulting in graphs with "steps".

7.2.3.3 EKF and batch EKF

To develop the EKF filters, the spring-mass problem example from chapter 4.8.2 of [Tapley et al., 2004] was used. Even though these algorithms are different, if the "extended" portion is removed, meaning that the nominal state is used throughout the simulation instead of being updated at each step, they prove to be mathematically equivalent. Thus, doing so, both a batch and sequential linear Kalman filters were developed and tested against the example's data. With both results in agreement with the provided values, the filters were adapted to the FCS-ATOMIC framework, and the nominal state updates added, together the Schmidt portion of the filter, where the changes made to the considered parameters are undone, thus transforming them into their extended form.

Unfortunately, after making these changes and preparing the filters to operate with the framework, there is no data to assess if these were implemented correctly. To make up for that, as done before, the batch EKF can be run with a batch size of 1 and compared with its sequential counterpart. The results of that experiment are shown in Figure 7.11, where it can be seen that both filters almost coincide perfectly apart from some numeric

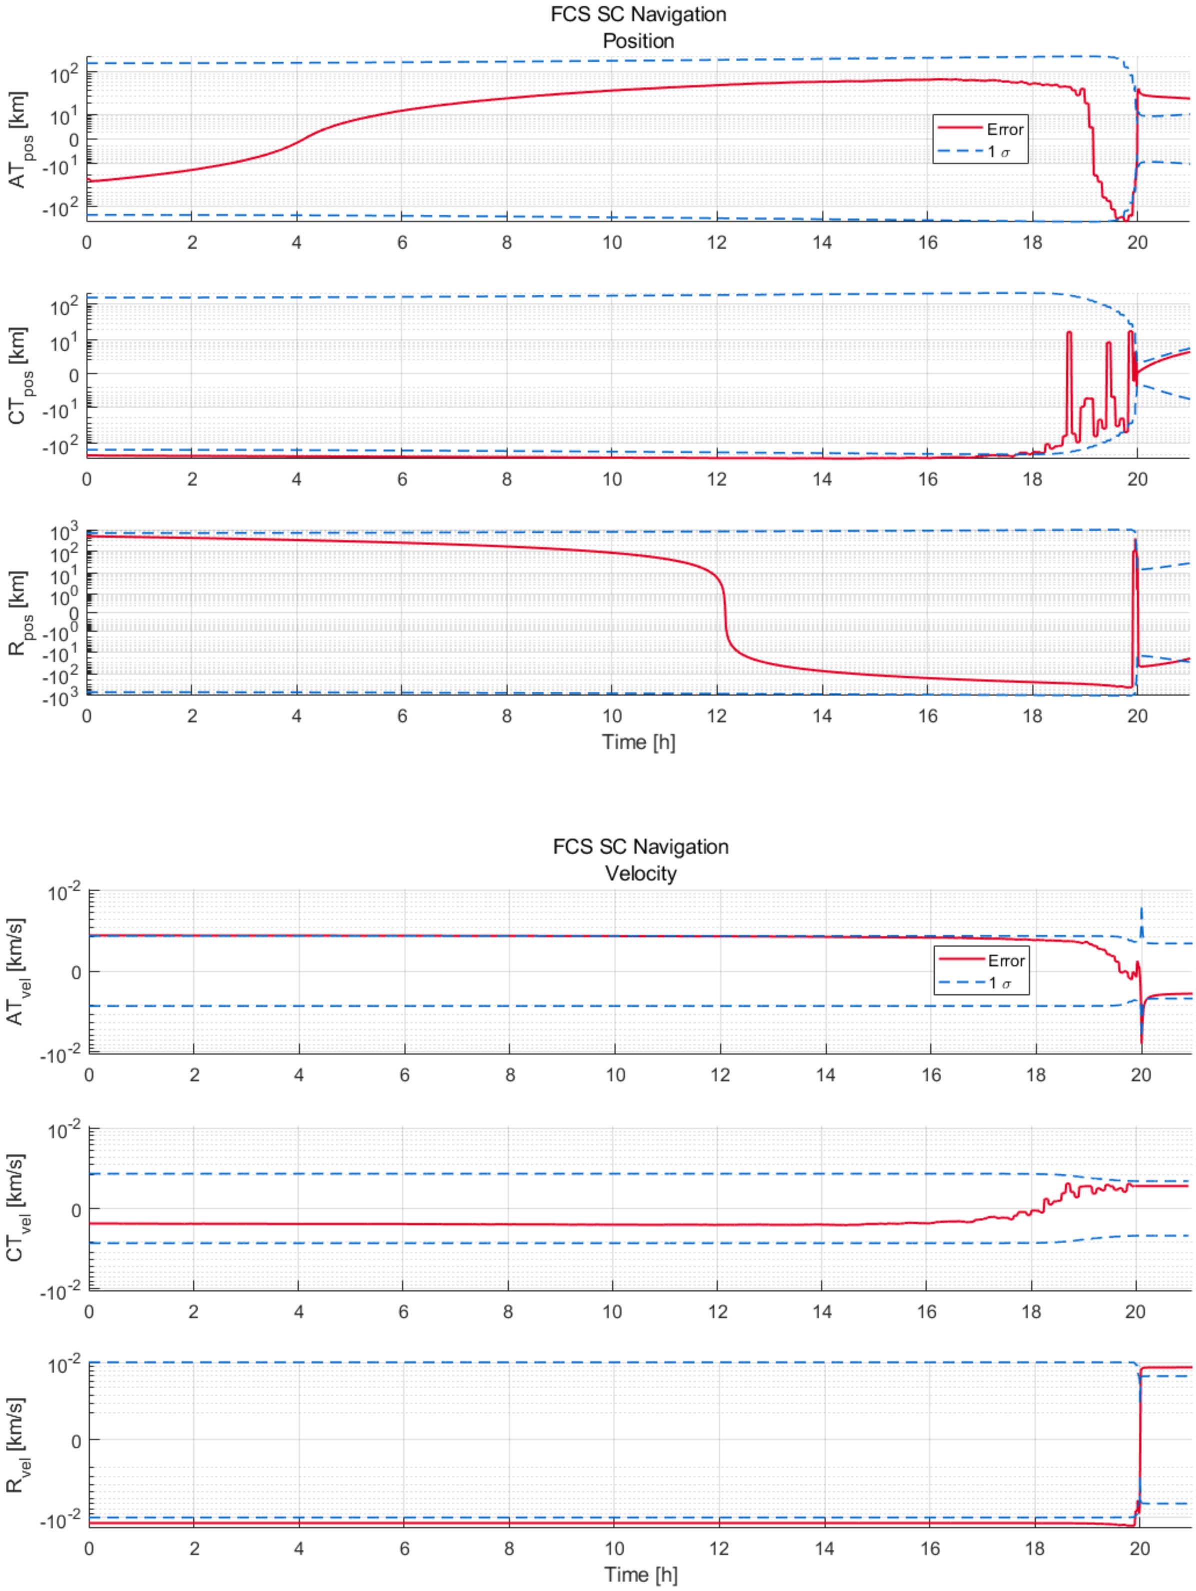

Figure 7.10: Navigation estimation error for the seed 0 scenario with the batch UKF. At the top is the position error, and the velocity error is at the bottom.

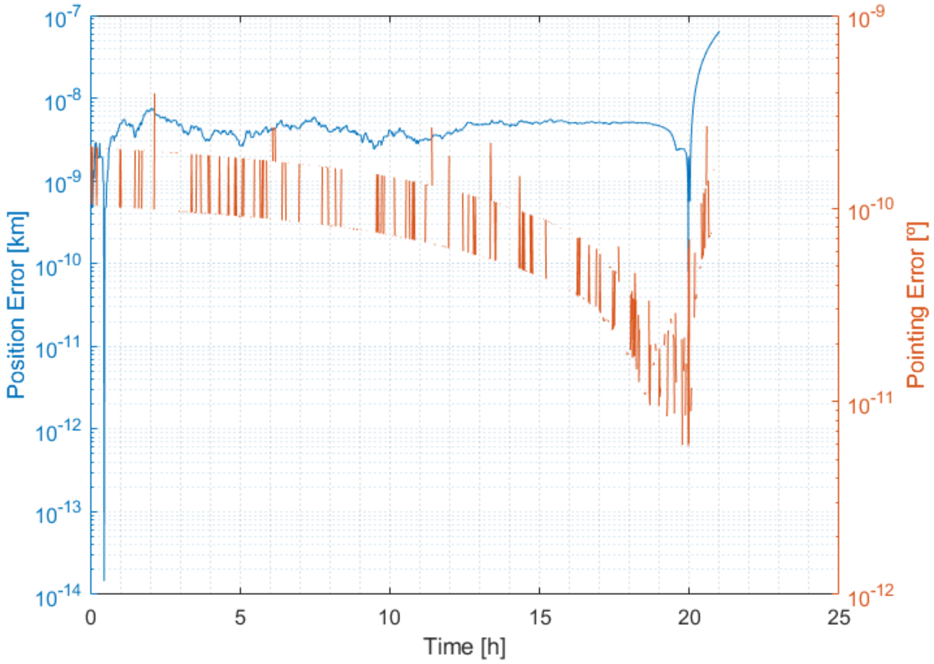

Figure 7.11: Position and pointing difference between the sequential EKF and batch EKF for seed 14.

error, as supposed. However, the batch and sequential formulations are significantly distinct, as the batch formulation relies on a different set of equations using the information matrix, instead of its inverse which is used by the sequential form. This reassures that the final transformations made to the filter were done so correctly, since these two formulations are in accordance with one another despite the major differences between the two algorithms.

Unlike what has been done for the previous filters, the navigation errors of the sequential EKF and batch EKF are not depicted here, as these graphs are very similar to those of the sequential UKF and batch UKF, respectively.

7.2.3.4 UDUEKF and batch UDUEKF

The fact that the UDU filters are built based on their EKF counterparts by adding these update cycles makes it easy to compare them. Since each operation replaces a matricial operation which is performed normally in the EKF filters, by replacing the cycles by the respective operations, the impact of each change made to these filters can be assessed. In other words, by only adding steps (2), (4) and (5) individually, while performing the other operations normally according to the EKF, the impact on the results by each of these steps can be quantified. The results of this experiment are shown in Figure 7.12.

In this graph, five lines show different comparisons of the sequential UDUEKF with the sequential EKF. The first is the base comparison, where the UDUEKF does not include steps (2), (4) nor (5). This makes sure that the two filters are indeed equal without these changes. For the majority of the duration, there is no blue line, meaning that the pointing difference is actually zero, thus is not visible in the logarithmic scale. The next two lines, orange and yellow, show the changes seen when steps (2) and (5) are introduced, and similarly to the previous example, these two steps do not add any meaningful changes to the results. Finally, the last two lines, purple and green, are almost perfectly coincident, which is why the former is not visible. This former shows the pointing impact of introducing step (4), whereas the latter emphasises the difference between the full sequential UDU filter and its EKF counterpart. Because these two coincide, the conclusion to draw from this analysis is that step (4) is the sole responsible by the difference between these two filters. While the other two steps do not introduce any meaningful difference in the results, the modified rank-one update makes the two filters differ by nearly a 0.1 deg pointing-wise. This value is quite high, considering that Figure 7.13 shows the average pointing difference between these two filters for seeds 0 through 5 to have a maximum of almost 10^{-2} deg. As such, these two filters are, in fact, different from one another in terms of the shown results, but on average their outputs should not be drastically different, as evidenced by these results.

Furthermore, whereas for the EKF and batch EKF the two filters perfectly converge in their estimates when the batch formulation is set to run with a batch size of 1, the same cannot be said for the UDUEKF and its

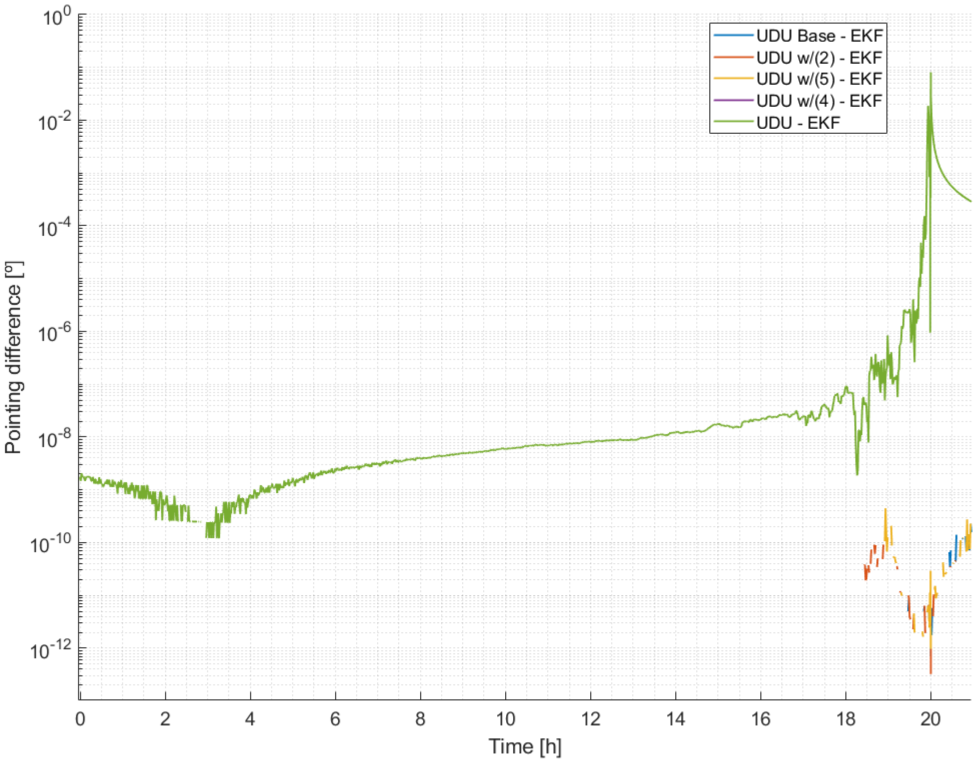

Figure 7.12: Pointing impact of the changes introduced by the sequential UDU filter with respect to its EKF counterpart.

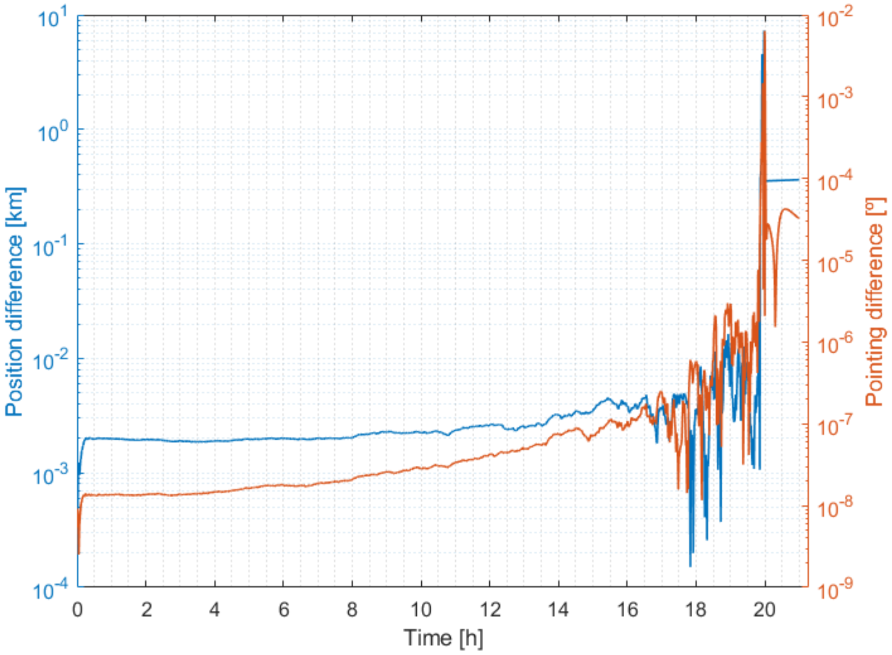

Figure 7.13: Average position and pointing difference between EKF and UDUEKF as simulated with seeds 0 through 5.

counterpart. For where the UDUEKF performs a regular *a posteriori* update on the covariance before then altering the changes made to the consider states with another "counter-update" according to the Schmidt formulation, the batch version, based on the batch EKF uses the information matrix instead, the inverse of the covariance. Thus, the same correction to the entries of the consider parameters would not be possible. Instead, the accumulation operation using the information matrix is done with a similar rank-one update method, the estimation performed, and the changes made to the obtained covariance matrix, as required by the Schmidt formulation, are made at the end, instead. As a result of this structural difference, the two filters do not share the same results, as these operations with the factored covariance are the source of difference within these formulations. As before, nonetheless, such changes are only observed right before the closest approach to the target.

Despite the lack of an ironclad methodology to assess the correct implementation of these filters, since they are based on the EKF and batch EKF, the only sections of the code that need to be verified are the aforementioned updates that replace matricial operations. And since they replace these well established computations, it becomes simple to compare the methods with the operations they substitute. Doing so shows that these different computations create very close results to one another, as expectable from navigation results so close for the majority of the simulation, with the differences between the two appearing more clearly in the smaller terms in absolute value, where the numerical error distinction between the methods matters the most in relative terms. It is due to these that the filters behave differently, and without such disparities, their results would be indistinguishable.

7.3 Simulation acceleration

Somewhat arbitrarily and certainly outside the scope of the work proposed for this thesis, the development and V&V of the dynamics' and measurement's Jacobian's analytical functions ended up proving a powerful tool to speed up the simulation performed in FCS-ATOMIC. This is because the computations of these derivatives do not happen exclusively during the execution of the navigation filters. In fact, after the real state of the system is propagated, the dispersion covariance with respect to the nominal trajectory is also propagated, and for that the derivative of the dynamics is required. At a later stage in the main cycle, after the measurements are simulated, their covariances are also computed, and for that, in turn, the Jacobian of these measurements with respect to the states and parameters is utilised.

In both these cases, the computations were done numerically, which makes the development of a general scenario easier, but at the cost of a substantial increase in processing power, since each derivative must be computed individually over a cycle. However, with the usage of these functions, these cycles could be removed. As such, a clause was introduced in the framework to skip these heavy computations in the declared presence of Jacobian-computing functions.

In doing so, all the numerical computations associated with these procedures are removed and replaced by much lighter ones, resulting in a simulation time cut by approximately 100 seconds, for the scenario being studied. In the context of the UKF and batch UKF, the usage of which make the simulations take nearly 8 minutes to run, this reduction accounts for 20% of the total time being reduced. However, for the remaining filters, whose times are much closer to the 5 minute mark, this is nearly a third of the run time removed.

Regardless of the filter being used, however, the truth is that 100 seconds saved during each simulation, in a study that, as will be clear in the chapters ahead, requires hundreds of runs to gather statistically relevant results from the data, results in a much shorter data acquisition process, or for more simulations to be performed during the same period. Take, for example, that, in total, the scenario being studied is run 1000 times. With 100 seconds saved for each single run, the total time exceeds 1 day. As such, these changes were crucial, not only to the process of analysing the Comet Interceptor scenario as it is, but also for GMV and future work using the FCS-ATOMIC framework.

Chapter 8

The scenario

This part of the study is dedicated to the definition of the scenario used to test the hypotheses put forward in the previous chapters. As such, all parameters that are needed to define the spacecraft's trajectory over time and its interaction with the environment are to be specified. In case of known or fixed parameters, these are given as a single value and corresponding units. On the other hand, for parameters which cannot be known ahead of time, the values used to define the probability density function that characterises their variation must be provided, so that with each new random number-generating seed the value changes accordingly. In general, though, normal distributions are used for most of the parameters. On top of the definition of these parameters, a brief discussion must be had regarding the justification for the choices of magnitude and variability. Do note that, in this discussion, similarly to what had been done previously during Chapter 4, the rationale will be to include all the contributions regarding to the comet and its environment, as it facilitates the mission analysis going forward by making the scenario more versatile to changes in future iterations.

The chapter will thus be divided into five parts. Firstly, considerations regarding simulation duration and the division of this period into adequate-sized steps is had. Then, the most important geometrical and orbit-defining aspects will be defined, namely the initial state of the spacecraft and the solar aspect angle of the orbit. Then, the dynamical parameters are set, which will influence the forces acting on the vehicle. After that, the variables regarding the observations taken by the camera system are put forward, which limit the accuracy with which the position of the comet's nucleus can be known from a distance. Only with these sets of values being determined can the simulations then be run and the results analysed in the framework of the defined restrictions. Finally, a discussion is had about the figures of merit used to characterise the scenarios being run, as a means to then identify which navigation solution is the best fit for the mission.

8.1 Simulation duration

One crucial parameter to define the simulation is the run duration. Assuming that it always starts at the beginning of the phase being considered – the encounter – the scenario should nominally take 40 hours to conclude. However, this duration can be significantly decreased, by simply realising that beyond the point of closest approach there is no longer much interest in the scenario results. After the last image is taken right before this point, the navigation becomes open-loop, as no more information is available to guide the state estimation of the spacecraft, and the state is merely propagated. Additionally, the focus of the mission is indeed around the closest-approach point, since this is when the vehicle is closest to the nucleus and able to obtain the majority of the scientific data, meaning that from there on the interest in the navigation is reduced. Finally, even if the last two arguments were not true, it still would be the case that the point that stresses the navigation the most is precisely around the point of closest approach. Thus, if the goal is to study the performance of the navigation effort, then this phase of the mission is the one that should be simulated, with less focus on what comes after.

Because of all the described reasons, reducing the simulation time from the nominal 40 hours to approximately 21 hours will benefit this study massively, as nearly half of the run time is saved and no significant information is lost as a result. Do note that the additional hour after C/A is merely added as a safeguard to guarantee that nothing of interest happens after this point, as hypothesised.

8.2 Time step size

The duration of the scenario is only half of the temporal definition required, since the simulation is discretised into time steps. Given the nature of the Comet Interceptor's approach distance and speed, which will be discussed in much greater detail in the section that follows, it turns out that for a good part of the first half of the encounter phase, the measurements taken by the navigation camera do very little to aid the state-estimation effort, and the results are nearly identical to an open-loop situation with no measurements. As such, during this period, the measurement frequency can be as high a possible without affecting the results. However, the frequency of the integration of the spacecraft's state must also be taken into consideration, since taking a higher period can mean that the results are not accurate enough for the analysis being performed.

It is, nevertheless, possible to make these two frequencies differ. That is to say that the state could be integrated twice as often as measurements were taken, for example, meaning that only one in every two steps would have a correction, or *a posteriori* phase with a measurement. On the other hand, the opposite could also happen, with twice as many measurements as integration steps, which would actually mean that only batch formulations would be possible for the filters. Despite all this, the approach taken in these simulations will be to match the two frequencies, since it vastly simplifies the setup required within FCS-ATOMIC.

Thus, the goal when setting the step duration is to make it as high as possible, since that reduces the total simulation duration quite substantially, while, at the same time, regarding the accuracy of integrating the dynamics and obtaining as many measurements as possible, especially when nearest to the nucleus, the opposite would be ideal. This has to be done, of course, within the interval stipulated in Chapter 6, which specified 10 minutes as the maximum time step duration.

The problem is that the needs of the simulation change quite drastically over its duration. For most of the first 20 hours not many measurements are needed, or useful anyway, as the position of the comet with respect to the camera sensor barely changes at all and the results are quite noisy, so ideally the time step would be as big as possible to reduce simulation time. But as the spacecraft approaches the comet, these two factors are no longer true, and thus frequency should be higher. To attend to these two competing requirements, a variable time step was implemented in FCS-ATOMIC. This allows for both to be satisfied, thus reducing the time required to simulate the scenario, while ensuring that enough information is gathered in the vicinity of the comet's nucleus.

As a result, four distinct phases are created. Starting at the beginning of the simulation and up until 3 minutes prior to the nominal closest-approach time, the frequency set for integration and measurement acquisition is 1/60 Hz, meaning that it is done once a minute. From 3 to 1 minutes prior to the 20 hour mark, this is increased to 1/5 Hz, or once every 5 seconds, increased again to 1 Hz after that, until closest approach. After that point, the frequency is then reduced back to the initial value of 1/60 Hz until the end of the simulation.

This choice of phases might seem to be somewhat arbitrary, but the increase in frequency around the 3 minutes before closest-approach mark is chosen to somewhat align with the comet becoming resolved in the camera measurements, at which point the image processing becomes more accurate at determining the nucleus' position in the photographs. As such, it is at this moment that an increase in measurement information is worth considering. Additionally, whereas for the majority of the encounter the comet's nucleus sits almost unmovable in the spacecraft's field of view, during these crucial minutes it zooms by the spacecraft rapidly, thus making this period the most demanding in pointing terms, both because of the proximity and the speed at which the fly-by takes place. Because of that, the additional measurements help the navigation capture the faster movement of the target during this short period, aiding the navigation to perform better.

8.3 Orbit geometry

As previously discussed in Section 4.2, due to the comet's very low mass, in astrodynamical terms at least, the trajectory of the spacecraft as it flies by the nucleus is mainly driven by solar gravitational attraction. In the vicinity of the nucleus, however, since both bodies feel this force, its effect is less significant than it otherwise would be, since the problem is being analysed from the frame of reference of the comet. As a result, the trajectory that the spacecraft describes could be approximately described by a straight line. Do note that the term "vicinity" should be taken with a grain of salt, given the relative speed of 10-70 km/s with respect to the nucleus and the 20 h time-frame considered in the first half of encounter scenario. Indeed, during this time the spacecraft, as well as after the closest approach, the vehicle travels an astronomical distance, but the conclusion is nonetheless unaltered, as this distance is but a small fraction of that of both bodies to the Sun.

Because of these facts, the trajectory of Comet Interceptor will mainly be defined by its initial position and velocity. These must be chosen with the requirements posed in Section 2.3 in mind, and such that the scenario represents an interesting study case to answer the posed hypotheses. For this, requirements **MIS – 1.2.** through **MIS – 1.8.** give all the restrictions that must be followed. Because of them, the initial state must allow for a nominal closest approach of 1000 km with respect to the target's nucleus, with a standard deviation of 150 km. Additionally, this closest approach should happen in the target-Sun plane, meaning that the speed in the perpendicular direction should nominally be 0, and the associated uncertainty low. Additionally, the relative fly-by speed is, as aforementioned, in the 10-70 km/s range, which couples with the phase's half duration of 20 h determines the distance from the comet's nucleus to the vehicle at the beginning of the scenario. Finally, the solar aspect angle lying in the 45-135 degree range determines the angle of the approximately straight trajectory of the vehicle with respect to the target-Sun line. All of these variations are summarised in Figure 8.1.

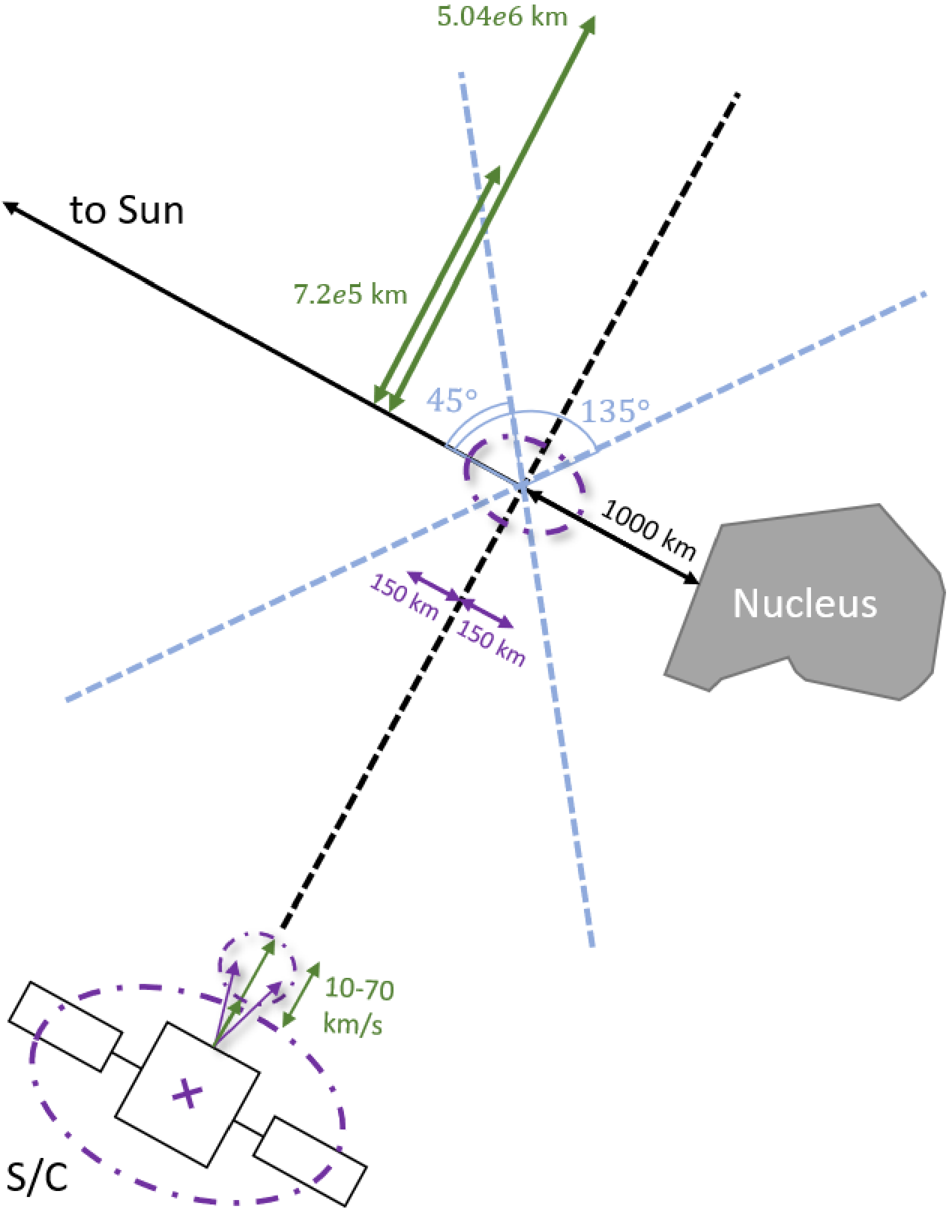

Figure 8.1: Geometry variations given the posed mission requirements.

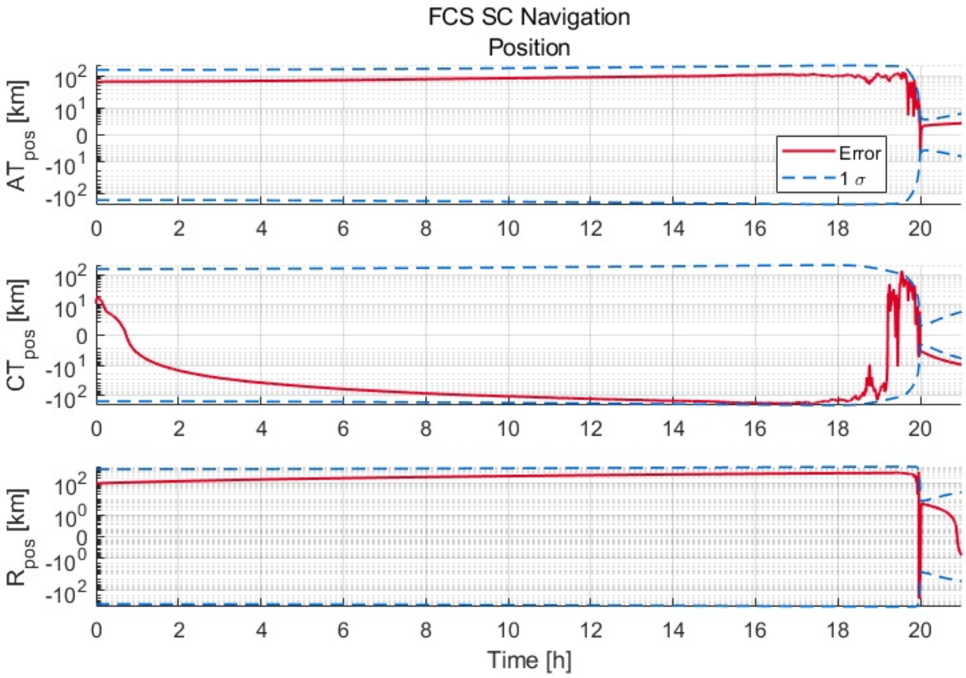

Figure 8.2: Position estimation results during an example baseline scenario.

In this schematic, the spacecraft, the nucleus, its nominal trajectory, and the sunward direction are indicated in black. Then, each variation with respect to this norm is indicated by a different colour. In light blue, the difference in solar aspect angle is represented, from 45 to 135 degrees. Do note that, in these cases, the closest-approach point is obviously different from the one indicated in the image. In green, the range of speed values is shown, from 10 to 70 km/s, which then results in smaller or bigger trajectory arcs, with length shown by the green arrows and respective values. Finally, in purple is the dispersion in initial position and velocity, which is such that there is a window of error at closest approach, which must not exceed 150 km. This window is shown as a circle, as the closest-approach point can also vary in the direction perpendicular to the image, not in the direction of the trajectory.

Analysing these degrees of freedom, the speed of the spacecraft in the 10-70 km/s range is most certainly the one with the greatest impact on the navigation performance. This is because in the case with the highest speed, for a given measurement frequency, the spacecraft is able to acquire seven times more images than for its lowest counterpart. Since this phase of the mission starts from far away, a significant portion of the trajectory is too far for accurate, position-improving measurements to be taken, and it is only in the last hours that the images start to portray the location of the comet with enough precision for the navigation estimate to start converging to its true value. This can be seen in Figure 8.2, which represents the position navigation error in the along-track, cross-track and radial directions, respectively, in red, as well as the respective standard deviations, in blue. As for the vast majority of the plot, the covariance increases slowly as a result of the lack of new information being acquired via the navigation camera, it can be seen that the early information is too noisy to be of any use for the filters. Not only that, but the majority of the estimation effort is concentrated in the last 15-30 minutes before the closest approach, which is when the covariance and the navigation error really start to tend towards 0.

With the importance of measurements being taken as close to the comet as possible to achieve the best navigation results being established, the number of such measurements clearly is a determining factor in the outcome of the experiments to come. In that sense, the scenario being analysed will reside near the highest speed threshold of 70 km/s. This is chosen to study the hardest possible scenario that the spacecraft has to face within the given requirement boundaries. Thus, the results obtained should represent a worst-case scenario in terms of performance for a given set of conditions, that should not worsen if the speed is decreased. Do note that the presence of outbursts might overrule the previous prediction, due to an increase in biased measurements.

One parameter whose importance should not be as large as the speed is the solar aspect angle. Its greatest impact on the system is by changing the measurements, in the sense that from different angles, the nucleus will have different illumination conditions. If the shape of the nucleus is assumed to be somewhat spherical, then only the half turned towards the Sun will be visible. With this in mind, and knowing that the spacecraft will approach the nucleus on this sunward side, then it can be concluded that the bigger the solar aspect angle, the better the images that the cameras will take, since more of the illuminated side will be visible.

However, there are two aspects which make this effect not a significant one. From the very beginning it should be known what the solar aspect angle for the trajectory is, which means that the image processing can be adjusted to take this into account. For example, in the case with an angle of 0, the image processing should know that the camera is turned towards a body, and that, of the half that could ever be visible, only half of it is, in turn, illuminated. This means that a bias must be considered of the order of magnitude of about half the nucleus' radius, helping to correct this issue. Additionally, either from very far away, at the beginning of the phase, or even near closest approach, this effect only corresponds to a bias of the order of magnitude of the nucleus' size, which is much smaller than its distance to the spacecraft. Thus, the error it might induce in the navigation should absolutely not be relevant enough to be a cause of concern. Even so, the value used will be 62.5 degrees. This is a result of the studies done for the approach phase that precedes the encounter in [GMV, 2020b], which determine that this is the lowest value possible given the considered conditions for that study. Additionally, the lowest value possible was chosen with the same rationale in mind as before, aiming to study the worst case scenario.

Finally, the initial position and velocity dispersion with respect to the nominal will also play an important role in the outcome of the spacecraft's autonomous navigation. Given the chosen initial speed of around 70 km/s and the mission phase's half duration of 20 h, it is possible to estimate the distance between the comet's nucleus and the vehicle at the start of the encounter. Then, around this initial state, there is some amount of uncertainty that is added to simulate the difference between the desired, nominal orbit, and the real orbit that the spacecraft ends up actually describing. This discrepancy arrives from a combination of the uncertainty in spacecraft state before this phase, as well as from the errors in the actuation that happens to steer the vehicle into the nominal orbit. That is to say that, because of the lack of perfect knowledge of the spacecraft's position and velocity with respect to the comet's nucleus, as well as imperfect actuation, the orbit that the mission is left with should differ from the one planned beforehand.

These differences between the nominal and actual trajectories will be generated randomly, according to a normal distribution, and will each have an average value of 0. However, given a small deviation 20 h prior to the closest approach, the final B-plane delivery error can be significant, with the threshold of 150 km in a nominal distance of 1000 km to the comet. This can have an impact on the measurements taken right before the C/A, as this difference in proximity will mean that the smallest distance will see a comet that is approximately 35% bigger than the biggest one, leading to better navigation results. At the same time, however, when the spacecraft is closer to the comet, a smaller position estimation error does lead to a higher pointing error, meaning that the navigation is also under more demanding conditions.

The dispersion errors are driven by the accuracy of the GNC performed in contact with ground before the data cut-off at 20 hours before the closest approach. In this regard, there is an asymmetry between the different directions of error. In the B-plane reference frame, it is easier to estimate the spacecraft's position in the two axes contained in the plane itself than it is in the direction perpendicular to the plane. According to [GMV, 2020b], for the two former terms, the estimated dispersion error is taken to have a standard deviation of 300 km, whereas for the latter, this error is given by the time-to-go error metric. This expresses how well the time of closest approach can be expressed, and given the mission's conditions, it is estimated that the standard deviation for time-to-go error is 10 seconds. The dispersion error's standard deviation for this radial component is then taken to be twice the distance the spacecraft travels in that time, so $20 \cdot 70$ km.

These are not the only parameters guided by this rationale, as the initial navigation error follows the same logic. Since the best state estimate at the beginning of the encounter phase is determined by the precision of the navigation up until that point, it is also the GNC accuracy that drives this error. As such, [GMV, 2020b] takes the initial navigation error standard deviations to be half of those of the dispersion error. This makes sense, as the dispersion error is the difference between the actual trajectory of the spacecraft and the nominal one established beforehand. As such, during the approach to the comet, the GNC's efforts are focused on converging the vehicle towards this desired trajectory, and any errors in the navigation, guidance or control will contribute to this dispersion. Thus, it is logical that the dispersion error has a higher standard deviation than the navigation error.

A similar rationale is used to obtain the velocity navigation error and dispersion error. On the one hand, for the former the in-plane components have a 2.1 m/s standard deviation, whereas the out-of-plane component's is 10 m/s. For the latter, on the other hand, the values are 2.8 m/s and 20 m/s, respectively.

8.4 Dynamical errors

With the initial conditions and dispersion that describe the rough geometry of the orbit being defined, the errors that affect the dynamics of the system must also be added. These reflect the uncertainties in the parameters that determine the forces that act on the spacecraft, and thus steer its movement.

Do note that, from this analysis, the comet's mass, and the associated gravitational parameter, will be left out, as their impact on the dynamics could easily be disregarded in the first place.

8.4.1 Comet's ephemerides

The most determining force acting on the spacecraft as it flies by the target comet in this mission is the Sun's gravity. Since this is a third-body force, its effect on the spacecraft must be computed while taking into account the effect on the comet's nucleus, which is the central body being considered. As such, imprecisions in the knowledge of the nucleus' position will lead to an error when computing this force.

This parameter will be determined by the precision of the orbit determining techniques used while observing the target comet from Earth. For the purposes of this study, the value used for this will be based on that given by [GMV, 2021a], which is 100 km, distributed spherically. As an approximation, the 100 km value for position will be used but without the spherical clause, for simplification, which will make it so that the results are statistically more pessimistic but not significantly different, especially since that this is rough estimate for this value in the first place.

As such, each scenario will generate a 3×1 error for the nucleus' position, one component for each direction, distributed normally centred in 0 and with a standard deviation of 100 km. This value of standard deviation is also accounted for in the covariance of the corresponding considered states.

8.4.2 Solar radiation pressure

One of the considered forces in this problem is the pressure exerted on the spacecraft by the incoming solar radiation. Analysing the formula that describes this force, as given in Section 4.2, this uncertainty can be related to an imperfect knowledge of the vehicle's effective area A_{Sun}, of the total power emitted as radiation by the Sun P_{Sun}, or even of the radiation pressure coefficient, C_r. Given that the vehicle can be thoroughly studied before launching in terms of its area and radiation pressure coefficient, the factor that should influence this variation the most should be the change in the radiation emitted by the Sun.

All in all, according to [GMV, 2020b], the standard deviation of this uncertainty is given to be 5% of the nominal value, as a typical guess considered for this force.

8.4.3 Comet activity

The impact of dust particles from the comet's environment surrounding the nucleus is also a source of force being considered. As discussed in Section 4.2, the way to describe this force mathematically and the value of the constants used to do so is still an open topic in the development of the Comet Interceptor mission. In this subsection, only the latter aspect will be addressed, as the former can be included in the non-modelled forces which will be discussed in the following subsection.

The comet activity describes the amount of mass ejected from the nucleus per second, which can be intersected by the spacecraft's trajectory. Thus, this parameter determines how much force on average the vehicle will feel as it encounters this material on its fly-by orbit. As aforementioned, according to [GMV, 2021a], the value used for this should be the very conservative 33000 kg/s, as it represents the highest value in the presented envelope. This value is, perhaps, more important in system designing than in mission planning, as the higher value represents the level of protection from particle impacts that the spacecraft must be ready to handle.

Because of this, it can be argued that a more realistic approach would be to consider an average value in the middle of the interval of possibilities and work with an unknown, randomly generated bias for each new simulation which represents the variation of about one order of magnitude that this value can have – 3,300 to 33,000 kg/s. With this range, the logarithmic midpoint is at about 10,000 kg/s, which is the nominal value being considered. Then, to consider the whole range of possible values, random biases are generated according to a normal distribution with 10,000 kg/s standard deviation, making sure that the upper bound of the interval is located approximately at the 3σ mark. However, this design choice makes it so that the generated values can end up below 0, so these cases must be rejected and computed again. Making this is more complex than perhaps could be achieved by simply considering a nominal value of 15,000 kg/s and a standard deviation of

5,000 kg/s, but ensures that the lower cometary activities are more common than the more extreme, unlikely values located towards the 33,000 kg/s mark.

8.4.4 Non-modelled acceleration

Finally, after touching all the parameters that can influence the forces being considered in the dynamical model of the system, it is a good practice to include an error term regarding possible imperfections in the models themselves. To take this into consideration, three bias terms will be considered along the previously mentioned consider states, one for acceleration in each direction in space. This is assuming that the forces being computed are not exactly the same as the true ones acting on the spacecraft, and that the corresponding total error is a constant over the fly-by, which might not be true for the latter statement.

Regardless, this approximation will introduce an uncertainty in the navigation that will change randomly from scenario to scenario. For a similar order of magnitude of error, its effect is greater than a randomly changing noise-like acceleration around 0. Thus, it is both more easy to implement than a process noise, and results in a more pessimistic result. Because of this, the approximation will be the one considered.

The major effect that make this acceleration a necessity is the outgassing from the comet, mentioned in Chapter 3 and described in Chapter 4. There, the acceleration is described to be equal to a constant A_i times a parameter that varies with the comet's heliocentric distance from approximately 1.25 to 0.6 with distances of 0.9 au to 1.25 au, respectively. The constant A_i will, on the other hand, be considered to vary randomly, thus generating three values of acceleration, one for each direction, centred at 0 and with a standard deviation of $5 \cdot 10^{-9}$ m/s^2, according to [GMV, 2021b].

8.5 Measurement errors

The last set of uncertainties related to the scenario definition are the ones affecting the measurements used to correct the navigation estimates. These are mostly related to attitude errors of the spacecraft and to the camera itself, as well as to the image processing algorithm that stands between the raw camera measurement and the information regarding the location of the comet's nucleus on the image.

8.5.1 Camera misalignment

One of the most simple problems associated with camera data acquisition in space missions arises from imperfect alignment of the camera with the desired orientation. This happens when the camera is pointing in a slightly different direction from the nominal, meaning that when images are obtained for navigation and the GNC system assumes a certain orientation to compare those images with a predicted result, this orientation does not correspond with reality. Such an effect can lead to a larger difference between these two metrics than what is really happening, introducing a bias into the system.

Despite being simple to explain, it is not simple to resolve, meaning that the bias is present and must be considered in this analysis. For its sake, two biases are considered, one in the "right-left" direction, and one in the "up-down" direction. Do note that a third component can also exist, as a boresight bias, that is, a rotation misalignment around the camera's pointing direction, but will not be considered in this study.

In terms of the two components considered, they are randomly generated at the beginning of each simulation and stay constant for its duration, working as biases. The generated values are done so according to a normal distribution with a mean of 0, and a standard deviation of 20 milliradians, according to [GMV, 2020b]. This number should be related to both the construction of the spacecraft and the rigidity of the materials that hold the camera in place, as well as to the precision of the calibration procedures done once the mission is in space, comparing data from different sensors to try and estimate possible errors of this sort. Additionally, the square of this value, 0.4 milliradians^2, is the variance of this bias that the system will use in the covariance matrix.

8.5.2 Spacecraft attitude

Just like camera misalignment, an error in vehicle attitude affecting measurements is an easy to understand phenomenon. As before, where a different orientation of the camera in the spacecraft's body leads to incorrect assumptions done during navigation calculations, this time its the whole orientation of the spacecraft being estimated incorrectly that leads to a similar effect. Instead of having two degrees of freedom, the attitude of the vehicle can be wrong in three different dimensions. And whereas the misalignment is a constant effect, thus being a bias, this error in attitude is a noise, as the mission uses its attitude sensors, mainly the star trackers, to track its orientation in space as it flies by the comet nucleus, being exposed to possible torques that change it over time, as a result of dust impacts, for example. Because the attitude readings coming from these sensors are not perfect, this noise must be considered.

As such, according to the values provided by [GMV, 2020b], each component of this attitude error has a standard deviation of 10 mdeg, which correspond to the accuracy of said sensors when determining the vehicle's orientation.

8.5.3 Image processing error

Other aspects that affect camera measurements but are not as easy to address as the previous two are all the errors associated with the image processing algorithm. This represents a crucial step between the images being captured by the NAVCAM and the usage of the captured information by the navigation. Despite this, its comprehension and thorough description is well beyond the scope of this thesis, and could certainly sprout studies of its own. To describe this algorithm, it is necessary to introduce the concept of the target being resolved in the taken pictures. This refers to the size of the comet on the camera sensor. For the purposes of Comet Interceptor, the target is considered resolved as soon as is surpasses 2 pixels in size [GMV, 2020a].

This distinction is relevant, as different image processing methods are utilised before and after resolution of the body. Before being resolved, which actually accounts for a majority of the encounter phase as can be seen by the approximation made in (8.3) for the WAC, a simple algorithm of analytical function fitting is used. This uses a point spread function that simulates how the light emitted by a point source is blurred in the camera's lens, which is fit to the data of the obtained picture to estimate the sub-pixel location of the comet.

This is the best available method from afar, as the spacecraft does not have access to any further information about the comet other than the faint light it is able to detect. Thus, computing and tracking the centre of brightness of the pixels considered to be the target is not only the best, but also the only available strategy.

$$r = \frac{\text{diameter}}{\frac{p}{p_{total}} FoV} = \frac{10}{\frac{2}{1024} 50 \frac{\pi}{180}} \approx 5867 \text{ km} \tag{8.1}$$

$$d_{to\ C/A} = \sqrt{r^2 - C/A^2} \approx 5781 \text{ km} \tag{8.2}$$

$$t_{before\ C/A} = \frac{d_{to\ C/A}}{v_{fly-by}} \approx 1\text{min } 23\text{s}, \text{ for 70 km/s fly-by speed} \tag{8.3}$$

On the other hand, as the spacecraft approaches the target comet, its size rapidly increases, and as a result more information becomes available. While this happens, using this method becomes increasingly less precise, because the comet is not evenly illuminated nor is its shape symmetric. Thus, the centre of brightness will have a considerable bias towards the illuminated side, which might be increased by its irregular shape.

Because of this, as the comet is resolved, the method used should be changed to better use the available information, and to take this factor into account. However, as it stands, the choice for this phase is still being traded-off between different options. Some of the approaches being considered are continuing on computing the centre of brightness but relying on the filter to estimate the underlying bias associated with this, or considering the different regions of the comet – nucleus, jets and coma, mainly – to improve accuracy.

As the analysis of the mission stands, the model of the IP error used is as described in [GMV, 2020b].

$$\boldsymbol{\epsilon_{IP}} = \boldsymbol{R_{SD2CAM}} \begin{bmatrix} \epsilon_{SD} + R_{comet} \\ \epsilon_{SP} \end{bmatrix} \tag{8.4}$$

$$\alpha_{R_{comet}} = \arctan\left(\frac{R_{comet}}{r}\right) \tag{8.5}$$

In (8.4), $\boldsymbol{\epsilon_{IP}}$ is the image processing error, and its computed using $\boldsymbol{R_{SD2CAM}}$, which is the 2-dimensional rotation matrix between the sunward/Sun perpendicular reference frame and the camera's FoV frame. It also requires ϵ_{SD} and ϵ_{SP}, the image processing errors in the sunward and Sun perpendicular directions, respectively. Each of these is assumed to be composed by a bias plus a white noise, varying with the resolved or unresolved status of the target. Finally, R_{comet} is the comet's radius.

Regarding the sunward and Sun perpendicular components, their noises and biases are constant, if the target is unresolved, but vary with the target's apparent radius, in case it is resolved. That is precisely what $\alpha_{R_{comet}}$, computed using the distance from the S/C to the comet r, is. With the apparent radius, the biases and noises are described by the values in Table 8.1.

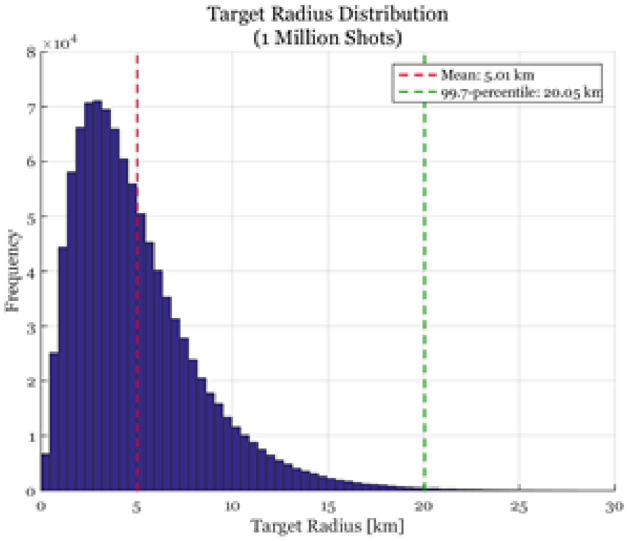

Figure 8.3: Distribution of comet radii in a sample of 1 million comets. Retrieved from [GMV, 2021b].

Table 8.1: Image processing errors for the Comet Interceptor mission. Retrieved from [GMV, 2020b].

Parameter			Nominal value	1-σ uncertainty
IP error	Resolved	sunward	0	0.1$\alpha_{R_{comet}}$ white noise, 0.5$\alpha_{R_{comet}}$ bias
		Sun perpendicular	0	0.1$\alpha_{R_{comet}}$ white noise, 0.1$\alpha_{R_{comet}}$ bias
	Unresolved	sunward	0	1 pixel white noise, 2 pixel bias
		Sun perpendicular	0	1 pixel white noise, 0.25 pixel bias

It is, nevertheless, important to keep in mind that, as the mission analysis of the Comet Interceptor is still in the earliest stages, there will certainly exist iterations changing these methods and the respective errors being considered.

In Table 8.1, values for the different stages of image processing and its different components can be found, in terms of the nominal value and 1σ uncertainties. For the biases, these values are generated randomly according to these parameters at the beginning of each scenario, whereas the white noises are updated at each step of the simulation.

8.5.4 Nucleus' radius

The final aspect interfering with the measurements is the size of the nucleus being photographed. Since once the nucleus is resolved the accuracy of the image processing improves, as the white noise and biases have significantly reduced standard deviations, the nucleus size can be quite preponderant in determining how many resolved images are taken of the body before it leaves the FoV of the cameras.

Using the data from Figure 8.3, the average comet radius is almost exactly 5 km. Whereas the distribution of the graph more closely resembles the form of a chi-squared curve, the generated values for the comet radius will be done according to a normal distribution around the average of 5 km, as a simplification. Doing so still encompasses lower frequencies for smaller comets as well as for larger comets, similarly to how the distribution in the figure does. The only downfall is that comets with higher radius appear more often than they should, whereas those with a smaller one are underrepresented. These should be seen the most often, since the mode of the distribution is around the 2-3 km range. Nevertheless, conclusions can still be drawn as to what happens when comet nuclei of different sizes are used, which is the important goal here.

As such, the radii will be generated according to a normal distribution centred in 5 km, with a standard deviation of 1 km. This will ensure that values below 0 are extremely unlikely, while studying the 2-8 km range, which is where the majority of the comets lie. Additionally, the mean value is in accordance with **GNC – 1.4.**.

8.6 Summary of dispersions and errors

In this section, all the dispersions and errors mentioned in this chapter are summarised in the form of Table 8.2.

Table 8.2: Summary of parameters, dispersion and error values in the considered scenarios.

	Name	Central value	Standard deviation
Geometry	Fly-by speed	70 km/s	–
	Solar aspect angle	62.5 deg	–
	Position navigation error	[0,0,0] km	[150,10·70,150] km
	Velocity navigation error	[0,0,0] km	[2.1,10,2.1] m/s
	Position dispersion error	[0,0,0] km	[300,20·70,300] km
	Velocity dispersion error	[0,0,0] km	[2.8,20,2.8] m/s
Dynamics	Comet's position	[0,0,0] km	[100,100,100] km
	Comet radius	5 km	1 km
	Solar radiation pressure	$4.54 \cdot 10^{-6}$ Pa	5%
	Cometary activity	10,000 kg/s	10,000 kg/s
	Non-modelled acceleration	[0,0,0] m/s^2	$[5 \cdot 10^{-9}, 5 \cdot 10^{-9}, 5 \cdot 10^{-9}]$ m/s^2
Measurements	Misalignment bias	[0,0] rad	[20,20] mrad
	Spacecraft attitude	[0,0,0] mdeg	[10,10,10] mdeg
	IP unresolved white noise	[0,0] pixel	[1,1] pixel
	IP resolved white noise	[0,0] pixel	[0.1,0.1]$\alpha_{R_{comet}}$ pixel
	IP unresolved bias	[0,0] pixel	[2,0.25] pixel
	IP resolved bias	[0,0] pixel	[0.5,0.1]$\alpha_{R_{comet}}$ pixel

One important aspect to note is the distinction between the real values of all these parameters and those known by the spacecraft's FCS. At the beginning of each simulation, the biases and noises are randomly generated, as aforementioned, but the vehicle's internal systems do not have information about the strength of these parameters. What they know are the statistical descriptions of each parameters described above. As such, for all the discussed terms, it is always assumed that the spacecraft knows exactly how they are generated, that is, it knows the central values of each, as well as their respective standard deviations, so at least the order of magnitude of these uncertainties is always known.

The filters' knowledge of the parameter's true definition again reinstates the argument made in Chapter 5 that these algorithms will not need any tunning. Since they already know the average and standard deviation of all the biases and noises affecting the dynamics and measurements, it stands to reason that the real values of these parameters be used, and represent the ideal strategy for the navigation effort.

8.7 Figures of merit

As important as carefully designing the scenarios being simulated to get relevant data given the mission requirements, is defining figures of merit that quantify the result of the simulation. In the case of this study for the Comet Interceptor, the goal is to evaluate the navigation performed by the spacecraft during fly-by, and determine whether it is able to provide good enough attitude and translation estimates to keep the CoCa pointing at the nucleus, to achieve one of the mission's scientific goals. Thus, as has been established in Section 4.3, since the absolute pointing error is driven mainly by the navigation error, this metric is important when evaluating the capability of the payload to be correctly pointed at the target.

However, since the navigation error in itself proves to be of little use, because from far away a huge navigation error is possible without much APE, whereas when crossing by the comet's nucleus the threshold for such error is much lower. It is still useful to compare the navigation error with the respective error covariance, as a value consistently over the covariance can indicate that either the navigation is being performed incorrectly or the covariance is wrongly estimated. Nevertheless, it would be handy to have a figure of merit that can, in a more straight-forward way, indicate whether the analysed navigation results do or do not allow the payload to keep the target within its FoV. And for that, the variable to use is the pointing error itself. This figure indicates the total misalignment between the position that the spacecraft believes the comet hold in its horizon and where it is actually located.

To then characterise whether the CoCa can be correctly pointed or not, the proposed approach is to establish a pointing error threshold, above which this instrument is considered to be unable to capture the comet's nucleus.

Establishing this threshold is difficult, as the information resulting from the simulations only provides the absolute pointing error and not the corresponding direction. With a FoV of 0.69 × 0.92 deg, in the worst case, an error of about just 0.35 deg will ruin the measurements, whereas with some luck, it could theoretically handle until about 0.58 deg. As an easy to spot threshold, then, the value of 0.50 deg, which lies between both cases, will be considered.

With this limit established, the concepts of uptime and downtime can be created. Uptime is considered to be a period during which the absolute pointing error is below this threshold, meaning that the CoCa is correctly pointed. Downtime is the opposite, which happens when the pointing error exceeds the 0.5 degree mark, meaning that the instrument looses the target from its field of view. These concepts will be major in the next chapters to characterise the performance of the navigation in the simulated scenarios.

Chapter 9

Parameter impact analysis

Before beginning the discussion about which navigation solutions work best with the scenario at hand, an additional piece of investigation is suggested. As the number of different parameters affecting the system were introduced in Chapter 8 is quite high, it would be desirable to know, before starting testing the different filters and drawing conclusions from the respective results, what are the impacts of each of the parameters. More specifically, knowing which errors represent the biggest effect in the simulation's performance, and which ones can be disregarded, will help unwrap the results of the navigation analysis performed later. To discuss this matter, this chapter is arranged in three sections. Firstly, the method of testing is described. Secondly, a preliminary analysis is carried out in order to cast aside the variables which are known to have little impact in the results before-hand. Then, with the remaining parameters, the procedure is carried out, the results shown, and conclusions drawn from them.

9.1 Analysis method

As a general idea, to understand how the multiple parameters affect the system, some way of changing the inputs must be put in place and the respective results then analysed in view of the values used. Many methods exist to achieve this, varying in difficulty and required run time. One such possibility would be using design of experiments and then performing analysis of variance – ANOVA – with the results. This could be a very effective way of studying this problem, but its fundamental functioning is not quite simple to implement in the framework of FCS-ATOMIC. This is because in this software, the scenario is defined in multiple files, which would need to be opened separately and changed between each simulation. This is quite complex given the number of parameters to change and how they are dispersed through different files.

Instead, a more adequate approach given the circumstances is to simply perform a Monte Carlo campaign. Since the framework is implemented with this in mind, with a simple number in one of the scenario files defining the seed to be used in generating all random parameters throughout the simulation, this method is certainly more advantageous in terms of set-up complexity. Choosing this method, it is important to once again mention that all the numbers are being generated pseudo-randomly using the default Mersenne Twister algorithm implemented in *MATLAB*. This is not ideal, as there are algorithms that serve a similar purpose while ensuring a more uniform distribution of the generated samples in the design space, such as the Sobol sequence, as can be seen in Figure 9.1. This hypothesis was considered, but unfortunately requires *MATLAB*'s Statistics and Machine Learning Toolbox to use this function, or must otherwise be implemented from scratch, which is undesirable. As a result, the default will be used, despite its disadvantages, and to understand how that can affect the analysis, the dispersion of the results within the design space will be briefly discussed.

Finally, with the scenario and the dispersion of the parameters being set, the navigation must be set up. For this, the extended Kalman filter is used, because it is the simplest, and thus fastest algorithm of the bunch, allowing the simulation to be run as many times as possible within a certain period of time to obtain the maximum amount of data. This is, however, disregarding the possibility that the different filters might show different results as to the impact of each parameter. While this is a possibility, it is unlikely, given that the rational expectation going into this analysis is that, while different parameters certainly have distinct impacts on the simulation performance, all the filters should handle the parameters similarly, meaning that a complicated set of values should results in overall worse performances and vice-versa. Nevertheless, the filter's algorithms have quite complex mathematical bases, so this possibility cannot be discarded completely. As such, a brief analysis of some results obtained with the other options for navigation will also be analysed, but not with the same amount of data as for the EKF, due to computation time constraints.

9.2 Variables of interest

Although in Chapter 8 a myriad of uncertainties related to the system were introduced into the scenario, as a way of simulating as many error factors affecting the result as possible, not all parameters have the same

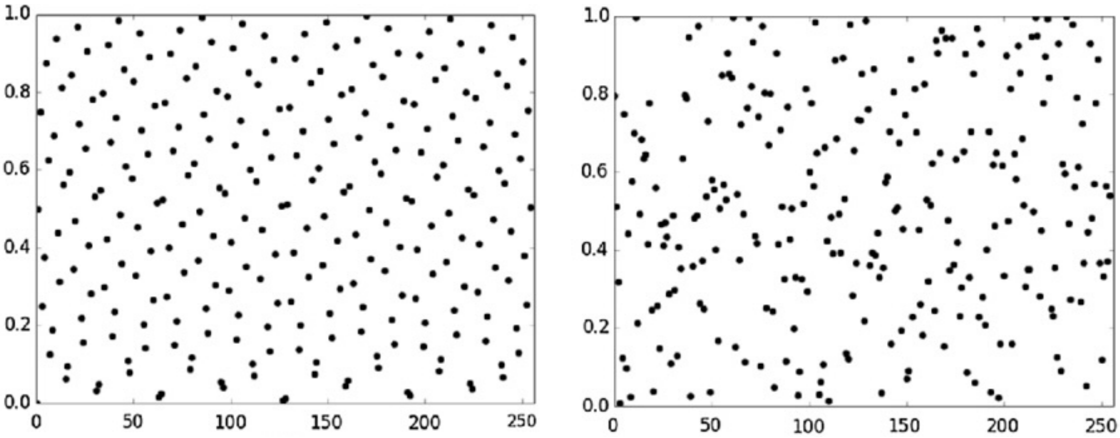

Figure 9.1: Comparison of 256 pseudo-randomly generated points according to Sobol sequence, on the left, and Mersenne Twister, on the right. Retrieved from [Di Gangi et al., 2019].

importance. While determining this importance is indeed the goal of this chapter's analysis, doing so assuming that all introduced uncertainties are relevant would be cumbersome, as to perform a Monte Carlo campaign, a number of runs should be performed per variable in the system. As such, and not knowing exactly how many runs are enough beforehand, the number of parameters being used should be kept as low as possible to minimise the number of times the scenarios must be simulated. Because of this, the goal of the present section is to determine which variables should be kept and which can be disregarded *a priori*.

The most obvious cases of discussion are the ones related to the dynamics of the mission. As mentioned in Chapter 4, the Comet Interceptor is peculiar in the sense that not many forces drive its dynamics. The most significant force establishing its trajectory is the Sun's gravitational pull, but given that the analysis is being made from the reference frame of the comet's nucleus, its impact is reduced as it also impacts the comet itself, dragging both bodies along simultaneously. The remaining interactions arising from the incoming Sun's radiation, the impact of dust on the spacecraft, the nucleus' own gravity, as well as the non-modelled accelerations being considered, are relatively unimportant. Thus, follows an analysis of each force and respective uncertainties.

Firstly, an uncertainty regarding the solar radiation pressure of around 5-10% – 1-2σ – would affect the nominal value of position displacement of around 650 m by about 100 m, meaning that not only is the force in question not largely consequential, but also its uncertainty is not big enough to change that. And all this is considering the whole 40 h of the encounter, whereas only the first half of this duration is relevant for navigation purposes.

In terms of the cometary dust impacts, the analysis can be simplified even further. Since the value computed in Chapter 4 already assumed a pessimistic value of cometary activity of 33000 kg/s, and the uncertainty generated will be done so around a smaller, more reasonable number of 10000 kg/s, then the computed sub-kilometre impact will – in the overwhelming majority of cases – only become smaller with the generated random values. As such, a similar conclusion can be drawn regarding this force and the associated uncertainty as was done for the previous case.

Considering the non-modelled accelerations, a_{NMA}, each axis has $5 \cdot 10^{-9}$ m/s^2 of standard deviation. These three contributions add up to $5\sqrt{3} \cdot 10^{-9}$, and the contribution in position variation due to these forces can be approximated as

$$\Delta r = \int \int_0^{t_f} a_{NMA} dt \tag{9.1}$$

$$\Delta r = \int \int_0^{t_f} \sqrt{3} \cdot 5 \cdot 10^{-9} dt \tag{9.2}$$

$$\Delta r = \sqrt{3} \cdot 5 \cdot 10^{-9} \cdot \frac{(40 \cdot 60 \cdot 60)^2}{2} \approx 90 \text{ m} \tag{9.3}$$

As for the previous cases, an effect of this dimension in a mission with a nominal closest approach of 1000 km can perfectly be ignored when performing the analysis of this chapter.

Finally, the error of the comet's ephemerides must be analysed. Changing the overall position of the system with respect to the Sun changes not only the gravity it exerts in the spacecraft and the nucleus' itself, but also increases the impact of the radiation pressure. For both these cases, the forces in question are affected by a dispersion with 100 km of standard deviation in position around the nominal, and a velocity dispersion of 1 m/s of standard deviation, which adds up to about another 100 km up until the closest approach. Do note that this deviation is not of the spacecraft with respect to the comet in any way, but of the whole spacecraft-nucleus system with respect to the Sun. Since the two mentioned forces scale with the inverse of the distance to the Sun squared, which is 0.9 au. With such a small perturbation in the scale of such a monumental distance, the expectations can only be that this impact will follow the other discussed contributions in terms of relevance.

A simple way to test these results is run a number of simulations where one group has all the parameters included in the navigation, and the other without the ones discussed in this section. Doing so yields consistent results showing a change in the navigation estimate of the position in the order of magnitude of tens of metres, and in the tenths of millimetres per second in velocity. This confirms the hypothesis that the theoretical uncertainties connected to the dynamics would not be crucial, similarly to the forces themselves in the first place. This conclusion greatly reduces the number of parameters to consider on this analysis of sensitivity, meaning that the much less resources are required to obtain results.

9.3 Results and analysis

With the previous analysis finished, the number of variables of interest is now limited to the initial state of the system and all the parameters regarding the observation used in the navigation. These are the two biases in image processing with an unresolved target, another two for when it is resolved, two additional contributions due to the misalignment of the camera itself, plus the radius of the comet's nucleus, which determines how soon the nucleus can be resolved. Added together, these amount to 7. Only left are the initial position and velocity of the spacecraft, as these determine how close the spacecraft gets on its fly-by, making the grand total of variables to analyse equal to 13.

One difficult task then is to decide how many samples should be used to draw conclusions, that is, how many differently generated simulations should be run. This can be a delicate question, because picking points at random and simulating is not an efficient way of studying the whole variable. This is why the choice of pseudo-random number generator was mentioned above. Adding to this is the fact that in a system like the one being studied in which little is known about the response surface, it is difficult to assess how sensitive it is to the variation of the inputs, meaning that picking to few points can lead to general conclusions which might not apply locally. With this in mind, a possible approach is to simply pick a number, study the results and draw conclusions, and then try again and check if the same results and conclusions hold.

For now, let the sample size be of 550 points – ranging from seeds 0 through 549 – corresponding to about 42 samples per variable. These were generated with *FCS-ATOMIC* in *Matlab*, by setting the *rng()* parameter to the values in the interval [0,549]. These in turn are used by the default Mersenne-Twister generator, as aforementioned, to create all the random values that define each specific simulation of the scenario. To characterise the results, the figure of merit used will be the maximum pointing error achieved during the simulation. This will help determine if the navigation was able to keep the payload pointed or not during the whole scenario. Doing as such yields the graphs in Figure 9.2, Figure 9.3, and Figure 9.4, where this pointing error is plotted against each variable individually.

The first shows the the error corresponding to each initial value of state dispersion with respect to the nominal in each of the 3 axes, both in position and velocity. Then. the camera misalignment and image processing errors are plotted, respectively, where IP represents the error when the comet is resolved and IPF when it is not – "F" stands for far. Finally, the same is done for all the values of comet radius considered.

The first thing that is apparent when looking at all these graphs is that the general shape of the distribution of the samples is very similar. They appear to be symmetric with respect to the line x = 0, with the majority of the points lying at the bottom, corresponding to low pointing error values, and their frequency becoming more and more scarce as this error increases. The same happens for higher and lower values of dispersion of each variable, which is expectable, since these dispersions are normally distributed around 0.

But most importantly, these graphs reveal that there is no strong correlation between the maximum pointing error and any of the isolated variables that strike the eye. This does not mean that there is no correlation between the variables and the figure of merit to be found, just that no individual variable is crucial on its own. As such, an analysis technique must be used to determine whether there might be correlations involving multiple variables.

To do so, one approach is to use the method of least-squares to fit the data to a response surface, and then use regression-based sensitivity analysis – RBSA – to determine the contribution of each of the regression terms to the obtained variability of the regression model. The analysis done in this study will be based on [Ridolfi and Mooij, 2016], where this process is elegantly explained and exemplified.

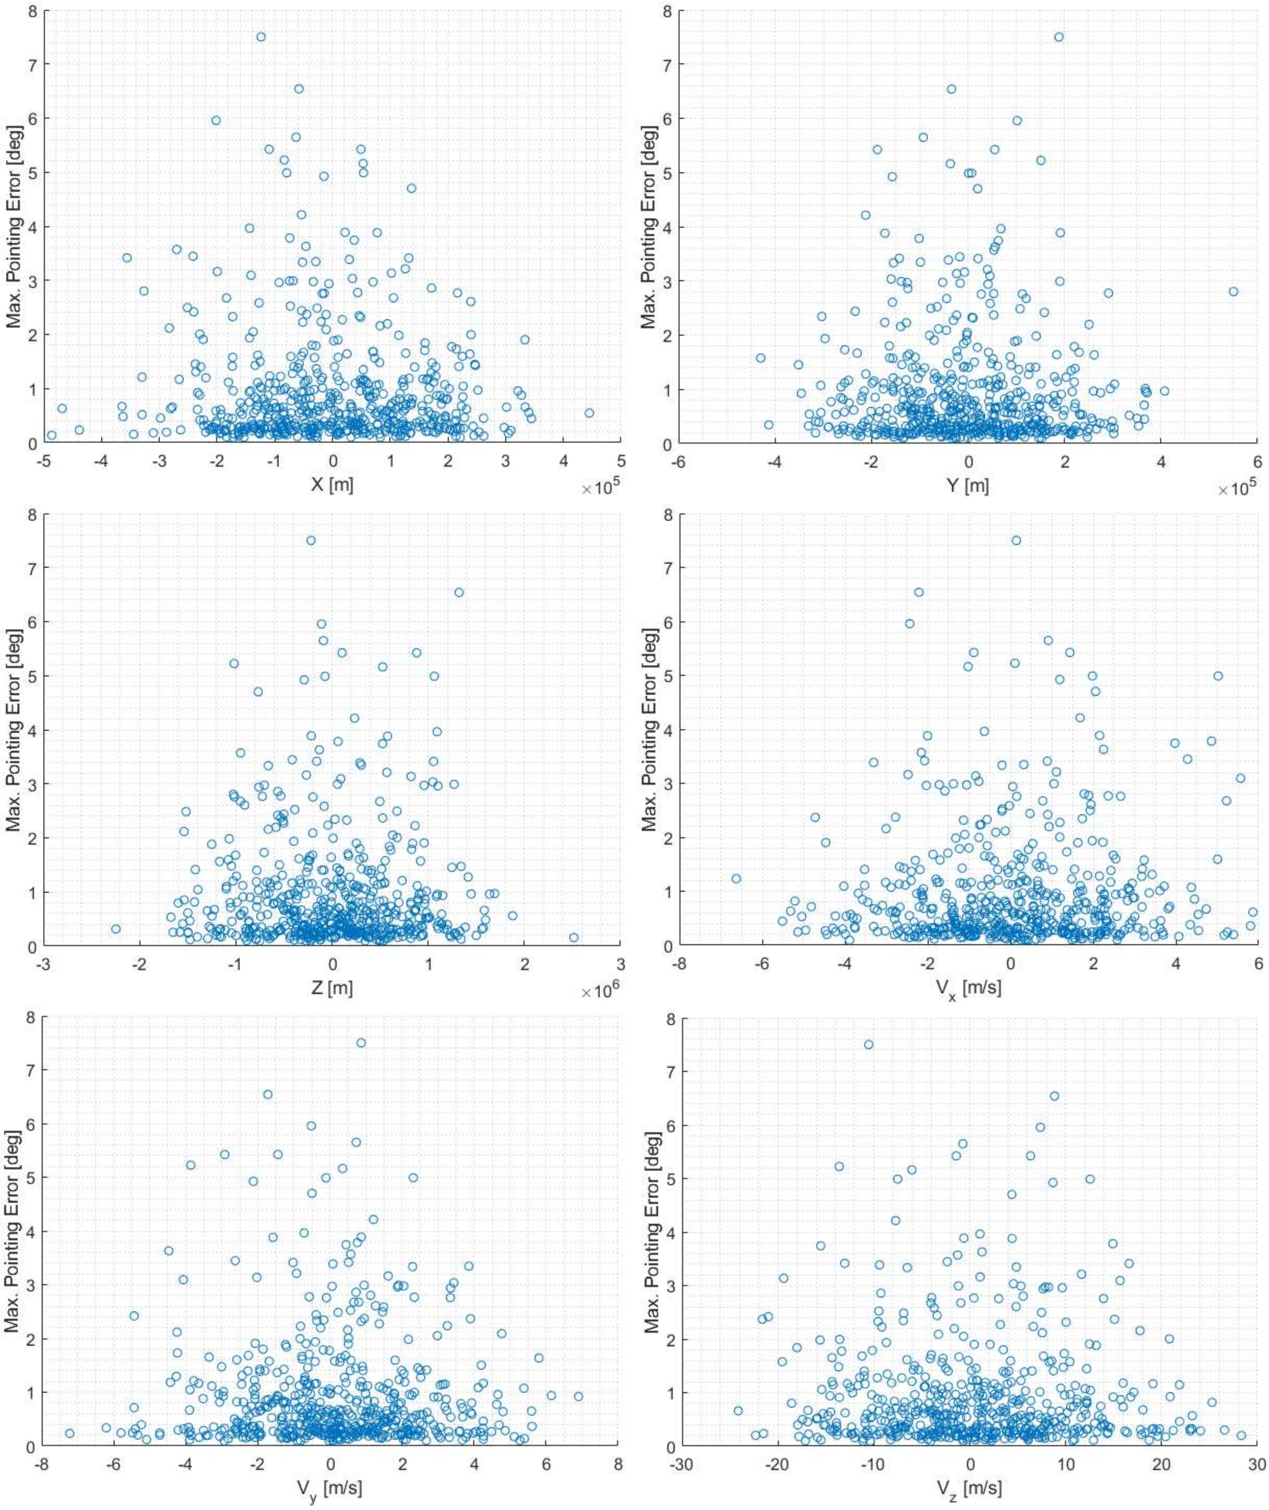

Figure 9.2: Maximum pointing error as a function of the dispersion of each position and velocity contribution in the initial state with respect to the nominal.

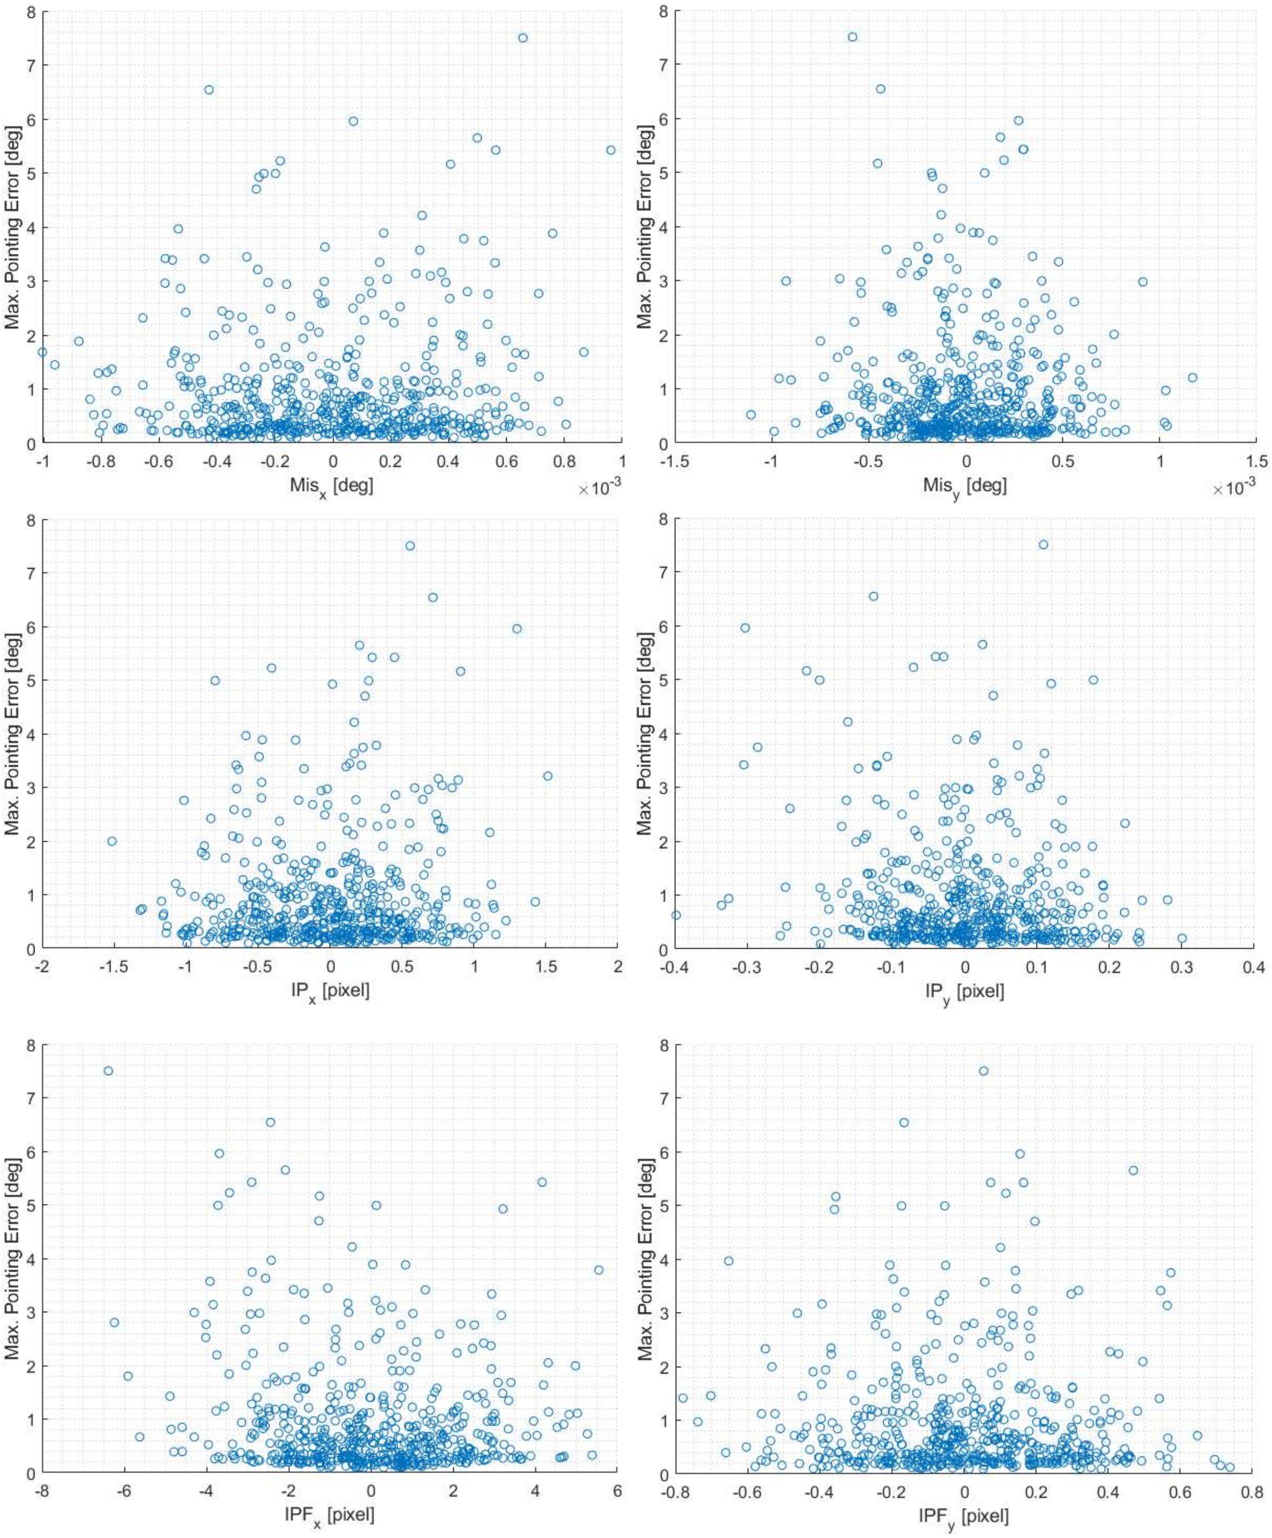

Figure 9.3: Maximum pointing error as a function of the measurement-related parameters. x and y refer to the two axes in the camera's frame.

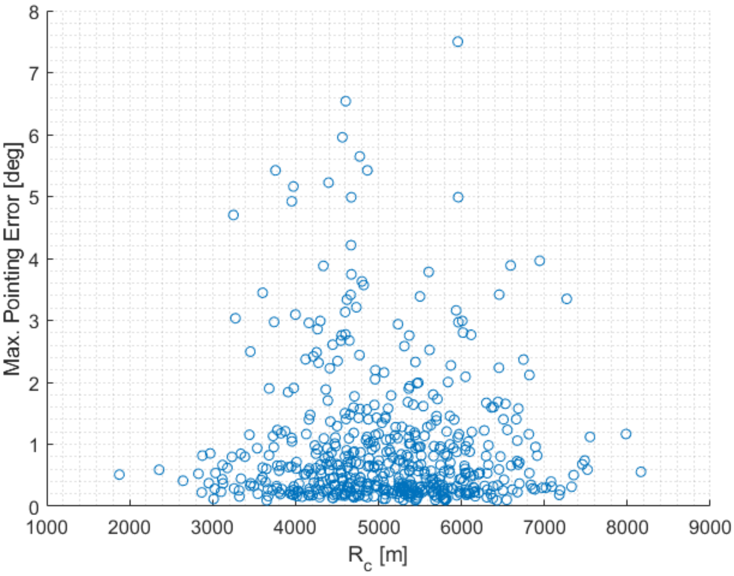

Figure 9.4: Maximum pointing error as a function of the comet's radius.

Before the sensitivity analysis, however, a fit is required. For that, an analytical expression needs to be put forward as an attempt to describe the available samples. In a case where a specific analytical model is being tested, this choice is more straight-forward, since the expression is given before-hand. In the case of this study, on the other hand, there is no preconception about what might be the best expression to fit available data. Because of that, the proposed method is to put forward a simple, conservative guess and then try to add terms and evaluate the evolution of the figures of merit that describe the quality of the fit. This discussion will be finished with the introduction of these concepts a bit further ahead in this section.

Consider now an initial guess to use to fit the data

$$y^F = \gamma_0 + \Sigma_i \gamma_{1i} x_i^F + \Sigma_i \Sigma_{j>i} \gamma_{1ij} x_i^F x_j^F + \Sigma_i \gamma_{2i} (x_i^F)^2 + \epsilon^F \tag{9.4}$$

where a constant term is added, linear and quadratic terms for each variable, as well as a terms for each combination of two variables, to account for the aforementioned correlations. For each of these terms there are coefficients, represented by γ, which need to be determined in accordance to the given dataset. This dataset is given by a series of inputs on the variables x^F and the respective outputs in y^F. However, unless the fit perfectly adjusts to this data, there will be a difference between the output values predicted by the fit, $\hat{y}^F$ and those obtained for the same inputs in reality, y^F. This difference is called the residual and is unique to each datum, since (9.4) applies to each sample's respective y^F and x^F values. As such, the goal of the least-squares method is to find the group of coefficients γ that fit the model to the data in the way that minimises the sum of the squares of the residuals ϵ^F. This can be done by considering the matrix form of (9.4), where all the samples are considered.

$$\boldsymbol{y^F} = \boldsymbol{X^F}\gamma + \boldsymbol{\epsilon^F} \tag{9.5}$$

$$\iff \begin{pmatrix} y_1^F \\ y_2^F \\ \vdots \\ y_N^F \end{pmatrix} = \begin{bmatrix} 1 & x_{11}^F & \dots & x_{1L}^F & x_{11}^F x_{12}^F & \dots & x_{1L-1}^F x_{1L}^F & (x_{11}^F)^2 & \dots & (x_{1L}^F)^2 \\ 1 & x_{21}^F & \dots & x_{2L}^F & x_{21}^F x_{22}^F & \dots & x_{2L-1}^F x_{2L}^F & (x_{21}^F)^2 & \dots & (x_{2L}^F)^2 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{N1}^F & \dots & x_{NL}^F & x_{N1}^F x_{N2}^F & \dots & x_{NL-1}^F x_{NL}^F & (x_{N1}^F)^2 & \dots & (x_{NL}^F)^2 \end{bmatrix}.$$

$$\cdot \begin{pmatrix} \gamma_0 & \gamma_{11} & \dots & \gamma_{1L} & \gamma_{112} & \dots & \gamma_{1L-1L} & \gamma_{21} & \dots & \gamma_{2L} \end{pmatrix}^T + \begin{pmatrix} \epsilon_1^F \\ \epsilon_2^F \\ \vdots \\ \epsilon_N^F \end{pmatrix} \tag{9.6}$$

Knowing this, the least-squares estimate can be computed as follows.

$$\hat{\gamma} = ((\boldsymbol{X^F})^T \boldsymbol{X^F})^{-1} (\boldsymbol{X^F})^T \boldsymbol{y^F} \tag{9.7}$$

Table 9.1: Coefficient of determination values for various cumulative modifications made to the fit model.

Terms	R^2_{adj}	R^2	Considered?
Baseline	0.15	0.31	Yes
$+x_i^F x_j^F x_k^F$	0.20	0.77	Yes
$+(x^F)^3$	0.22	0.79	Yes
$+(x^F)^4$	0.28	0.83	Yes
$+(x^F)^5$	0.25	0.84	No
$+(x^F)^6$	0.28	0.84	Yes
$+(x^F)^7$	0.29	0.86	Yes
$+(x^F)^8$	0.42	0.90	Yes
$+(x^F)^9$	0.43	0.92	Yes
$+(x^F)^{10}$	0.38	0.92	No
$+(x^F)^{11}$	0.39	0.93	No
$+\sqrt{\lvert x^F \rvert}$	0.45	0.93	Yes
$+\sqrt[3]{\lvert x^F \rvert}$	0.48	0.95	Yes
$+\log_{10} x^F$	0.55	0.97	Yes
$+\sin x^F$	0.52	0.98	No
$+\cos x^F$	0.43	0.97	No
$+\tan x^F$	0.49	0.97	No
$+1/x^F$	0.45	0.97	No
$+1/(x^F)^2$	0.47	0.97	No

With this list of coefficients, the predicted values given by the fit for each sample can be computed as

$$\boldsymbol{\hat{y}^F} = \boldsymbol{X^F}\hat{\gamma} \tag{9.8}$$

With these equations, the least-squares fit is performed and the coefficients are computed. However, there is no indication that the chosen model can be trusted to adequately describe the samples used. For that there is the coefficient of determination, R^2, which determines the fraction of the total variability that is due to the regression model. This can be computed using the following equation

$$R^2 = \frac{\sum_{i=1}^{N}(\hat{y}_i^F - E(\boldsymbol{y^F}))^2}{\sum_{i=1}^{N}(y_i^F - E(\boldsymbol{y^F}))^2} \tag{9.9}$$

$$= \frac{(\boldsymbol{y^F})^T \left[\boldsymbol{H^F} - \frac{1}{N}\boldsymbol{J}\right] \boldsymbol{y^F}}{(\boldsymbol{y^F})^T \left[\boldsymbol{I} - \frac{1}{N}\boldsymbol{J}\right] \boldsymbol{y^F}} \tag{9.10}$$

where $\boldsymbol{I}$ is the identity matrix, $\boldsymbol{J}$ a matrix of ones, and $\boldsymbol{H^F}$ is the hat matrix, which transforms the observed values to the fitted values.

$$\boldsymbol{\hat{y}^F} = \boldsymbol{H^F}\boldsymbol{y^F} = \boldsymbol{X^F}((\boldsymbol{X^F})^T\boldsymbol{X^F})^{-1}(\boldsymbol{X^F})^T\boldsymbol{y^F} \tag{9.11}$$

The value of R^2 varies from 0 to 1, and is a good place to start while studying the possible lack-of-fit of the used analytical model. However, there is a negative side in using this as the sole metric, as adding more terms to the model can make the coefficient of determination rise just due to the reduction of degrees of freedom, and not because the new term actually describes the distribution of the data. For example, 5 samples distributed along an exponential curve can be perfectly fitted by a polynomial of degree 4. This lack-of-fit will become obvious if more samples are added, as the polynomial will now struggle to fit all given values, and R^2 will naturally decrease. This will happen if more and more terms are added to the analytical model, meaning that there is no information to be found in this figure of merit as to how adequate the fitting model is for the problem.

A solution to this is to use R^2_{adj}, the adjusted coefficient of determination, which is given by

$$R^2_{adj} = 1 - \left(\frac{N-1}{N-(L+1)}\right)(1 - R^2) \tag{9.12}$$

This coefficient, unlike R^2, penalises the usage of more terms in the fitting model, since its expression has L, the number of fitting terms, in the denominator. Thus, this is preferable to determine whether the fit being used is adequate or not. Finally, regarding the threshold for distinguishing a good fit from a bad fit, the value put forward by [Ridolfi and Mooij, 2016] is of 0.9 for both R^2 and R^2_{adj}, meaning that above this value the vast majority of the variability is explained by the fit, which is desirable.

With this explanation, the ground is now set to analyse the given 550 samples gathered with the Comet Interceptor encounter scenario and try to find a suitable fit. Starting with the initially proposed fit – henceforth referred to as "baseline" – the values of these coefficients will be evaluated, and more terms will be considered. If these increase the value of R^2_{adj}, then they are added to the model, and if not they are discarded.

Do note that all attempted strategies and their respective results are not explicitly laid out here as is done in Table 9.1. Not only were countless other functions, like the exponential, for example, tested with the others

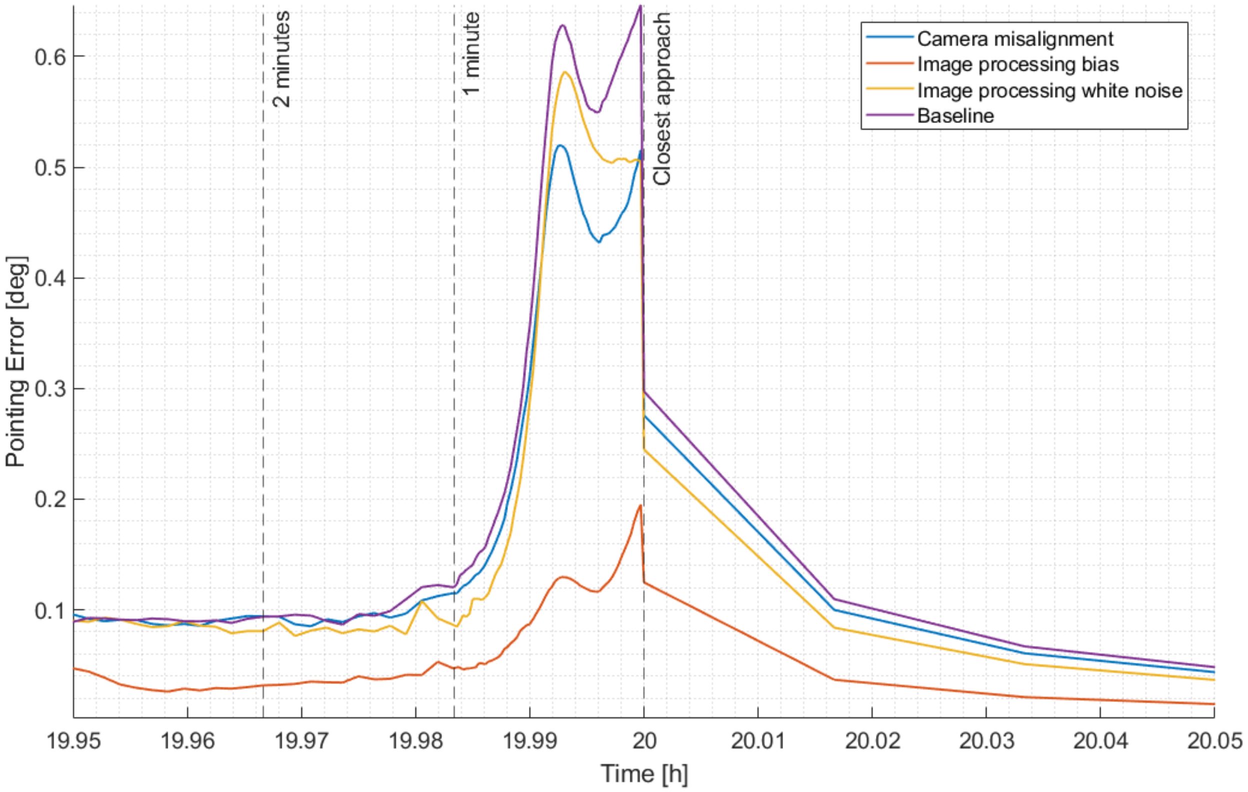

Figure 9.5: Average pointing error for each of the batches of EKF simulations run with seeds 0 through 19. In each case, the mentioned variable is removed from the simulation.

considered, but the data used for the fit model was also normalised. However, just as all the strategies described so far, the results did not improve significantly. There are infinitely more terms that could be added to the model, and at some point it could even be possible that the adjusted coefficient of determination reached the posed 0.9 threshold. However, with the 550 samples used, little more terms could be added, as the number of samples always needs to be higher than the number of coefficients in the fit model, otherwise the system of equations becomes underdetermined. However, there is a simpler reason as to why this fit did not achieve a good enough R^2_{adj} and it is hidden in the terms that are not considered here. In fact, two white noise terms corresponding to attitude determination error and image processing were mentioned in the last chapter, but their impacts have not been considered until now. And it is likely their presence that is "muffling" possible correlations between these biases and the pointing error. Imagine a case where one variable does have a meaningful impact on this error. This would mean that a large absolute value of the bias should increase the observer error in some way. Yet, if there is a significant white noise error contribution, its impact can have the effect of lowering some of the points where the pointing error should be higher and vice-versa, thus hiding the relation.

By doing a simple test can give more insight into what is happening. Twenty batches of simulations are run for the EKF multiple times. Once, nothing is changed, and acts as a baseline. In the other batches, each measurement affecting variable being studied is dialled down to 0, to assess the average results of these runs without it as compared with each other and with the baseline. The results of this experiment showed in Figure 9.5 show the average pointing error around the closest-approach time for each of the tested cases. It is clear that the removal of each variable decreases the pointing error and thus improves performance by a certain amount. However, by far the greatest improvement is seen when the image processing bias is excluded, thus indicating that this variable meaningfully affects the navigation performance, leaving the question of why that effect cannot be seen.

This impact of the measurement noises in the results can be, nonetheless, an indication that perhaps the variable being used thus far is not the most adequate to study the navigation's performance. As will be frequently discussed throughout the next chapters, it is always possible to utilise either the maximum pointing error, or another measure that better describes its evolution over time, instead of just capturing its maximum. This might be advantageous, as the maximum can certainly be affected by a random peak in noise intensity, thus making it a less reliable evaluation variable. To assess this possibility, the proposed approach is to use the concept of downtime, as described at the end of Chapter 8.

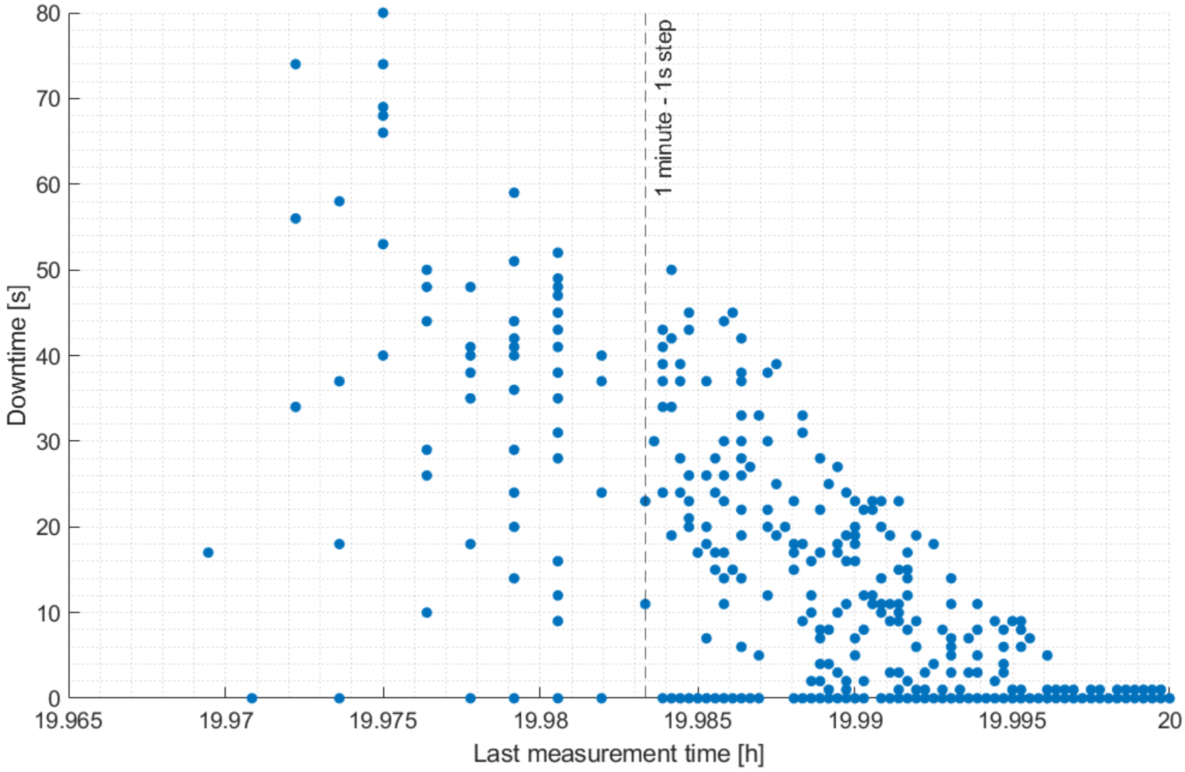

Figure 9.6: Representation of the total downtime for each tested cases as a function of the time of the last viable measurement.

Repeating the process laid out in the previous paragraphs, but now with this new figure of merit yields very similar results, but one certain variable that has not yet been represented does stand out. Figure 9.6 shows the total downtime for each of the tested seeds – again 0 through 549 – as a function of the time of the last valid measurement of the comet's nucleus. Also represented is a vertical line at 1 minute to the nominal 20 hour closest approach, which is relevant as it represents the point at which the simulation time step decreases from 5 seconds to 1 second, thus increasing the frequency of incoming measurement information.

In this graph, it can clearly be seen that there is a trend indicating that cases with measurement information ending sooner have a higher tendency of showing longer downtimes. Whereas the majority of simulations stop capturing valid images of the target close to the 20 hour mark, there is still a relevant number of cases with measurement cut-off times significantly before that, extending to almost 2 minutes prior, but with an increasingly lower number of cases as this difference becomes higher. And although these cases are not guaranteed to show a worse performance, as for the whole space of cases being studied there are plenty of examples with 0 seconds of downtime throughout, a clear trend can be seen where the higher downtime values are only seen for simulations where the measurement cut-off happens sooner. So, whereas the minimum of 0 seconds can be found in almost all of the variable space being studied, the maximum increases in a seemingly linear fashion as the cut-off moves back in time. This behaviour is likely due to the impact that the other uncertainties in the simulation have, which in some cases might allow for good results to happen in spite of sooner cut-off times.

This trend is mainly a consequence of the fact that the time step for the simulation is set with respect to the nominal trajectory. Nominally, the spacecraft's trajectory is such that it passes by the comet 20 hours after the encounter begins, marking the closest approach. Because the NAVCAM's orientation is fixed in the vehicle's body and its field of view is only 50 deg, it looses sight of the target before that happens. However, since the real trajectory in each distinct simulation presents a different dispersion of the initial position and velocity. These two effects make it so that the spacecraft in reality passes by the comet slightly sooner or later than it nominally would. And this is what makes the time of the last measurement change in the various cases, with some happening sooner and other later than the nominal depending on whether the vehicle is ahead or behind of the pre-planned trajectory.

However, since the time step of the simulation is adjusted with respect to the nominal, and some cases deviate significantly with respect to it, the result is that some instances get less measurements in the 1 second time step zone, which is where these are captured with the most frequency and thus help the navigation more accurately estimate the vehicle's position with respect to the target. Some cases even see their influx of measurement being cut prior to this change 1 minute before the 20 hour mark, meaning that the last images received by the filters have a 5 second interval between one another, further exacerbating the issue.

Indeed, plotting the same metric against the value of the initial radial distance and velocity dispersion – corresponding to the $\hat{S}$ direction in the B-plane reference frame – does confirm this hypothesis, as both graphs in Figure 9.7 show a similar trend. In either case, a positive value of the dispersion, corresponding to the vehicle being ahead of the nominal trajectory, shows an increase in the number of cases with higher and higher downtimes, in a very similar fashion to that of the already discussed graph in Figure 9.6.

What this analysis points to is towards the idea that the performance of the navigation will overall depend more on the geometry of the trajectory taken by the spacecraft than by the many uncertainties surrounding the dynamics and the measurement acquisition used by the filters. Of course that these parameters are still important, as all the analysed graphs only showed trends towards an increase in downtime in certain conditions but not perfect relations, in the sense that not only points followed said trends. For all zones of the variable space, with higher or lower dispersions, with sooner or later data cut-off times, there are cases with near zero downtime, meaning that, despite these variables being useful in understanding if there is a propensity for the simulation to have less than perfect results, they cannot predict exactly what the outcome will be. And, again, this likely happens due to the combined effect of all variables.

This prevalence of the geometric variables is further solidified by the data in Figure 9.8, where the downtime as function of the closest-approach distance of the vehicle to the nucleus is represented. This, once again, shows a trend that cases with higher downtime are almost exclusively characterised by a closest-approach distance below 1500 km, whereas above that value nearly all simulations show near flawless downtime results. This makes sense from a physical standpoint, as the nearer the spacecraft gets to the comet's nucleus, the harder it is to keep the target in the FoV of the cameras, since smaller errors in position estimation result in higher pointing inaccuracies. And it is again the initial dispersions which determine the variation of this closest-approach distance, since nominally the spacecraft would fly at about 1000 km from the target. The change in initial position and velocity with respect to the nominal are the reason the closest-approach distance deviates so much from the nominal as is seen in Figure 9.8. In fact, this variation is significantly above the threshold of 150 km 1σ established in **MIS – 1.4.**, with some values reaching well into the 2000 km. This will be discussed further in Section 11.2.

Thus, it stands to reason that the conclusion to this parameter impact analysis is that the initial parameters considered and studied in this chapter, despite certainly affecting the outcome of the simulation, are harder to understand, perhaps due to the presence of the image processing noise, which masks possible relations, or because their impact is smaller. On the other hand, the geometric parameters defining the orbit of the spacecraft end up being those which give the most reliable indication as to whether the navigation will perform well or can present some downtime along the fly-by. In conclusion, a sooner data cut-off of the incoming measurements means that there is a higher chance that the results show higher and higher downtimes, rather than none. This happens more noticeably with higher radial components of the position and velocity dispersions. Simultaneously, the minimum distance to the comet during the spacecraft's trajectory also has a similar effect, with smaller distances meaning that the vehicle gets closer to the target and thus more easily looses sight of it.

Do note that despite these correlations, these metrics do not fully predict whether a certain simulation's performance will be good or bad, as some cases with very early measurement cut-off times or low C/A distances to the target show absolutely no downtime. This is expected, as here only two variables are being analysed. Despite no quantitative relationship between the variables analysed earlier and the downtime results being found via the RBSA method employed, Figure 9.5 shows that the measurement-impacting variables related to the image processing and to the camera misalignment have a significant effect on the pointing error. And understanding the interaction between these variables might explain the high degree of variability in the results. For example, a case of a scenario with an early cut-off time can still see a good performance if the last measurements of the comet it has access to are relatively unperturbed by the remaining biases and noises in the system. However, as the cut-off happens later and later, it becomes increasingly harder to see all measurements negatively affected by these uncertainties, meaning that it is easier for the scenarios to present a better performance, and lower downtime, as can be seen in Figure 9.6. An equivalent argument can be made for the case of the closest-approach distance.

In terms of how these issues can be mitigated, the first one, as was already discussed, likely stems from using the nominal time frame to establish the intervals when the time step varies. This could be mitigated by trying to find a strategy to autonomously change the frequency of measurements and navigation estimations solely relying on the spacecraft's information of the system. It could use its state estimations to determine its distance to the closest-approach point and base the change in time step upon such predictions instead, or even do the same by studying the brightness of the comet's nucleus as an indicator of when this frequency should increase. The more simple approach, however, would be to establish bigger intervals when the simulation reduces the time step to first 5 seconds and then 1 second. This is done in the last two minutes before the

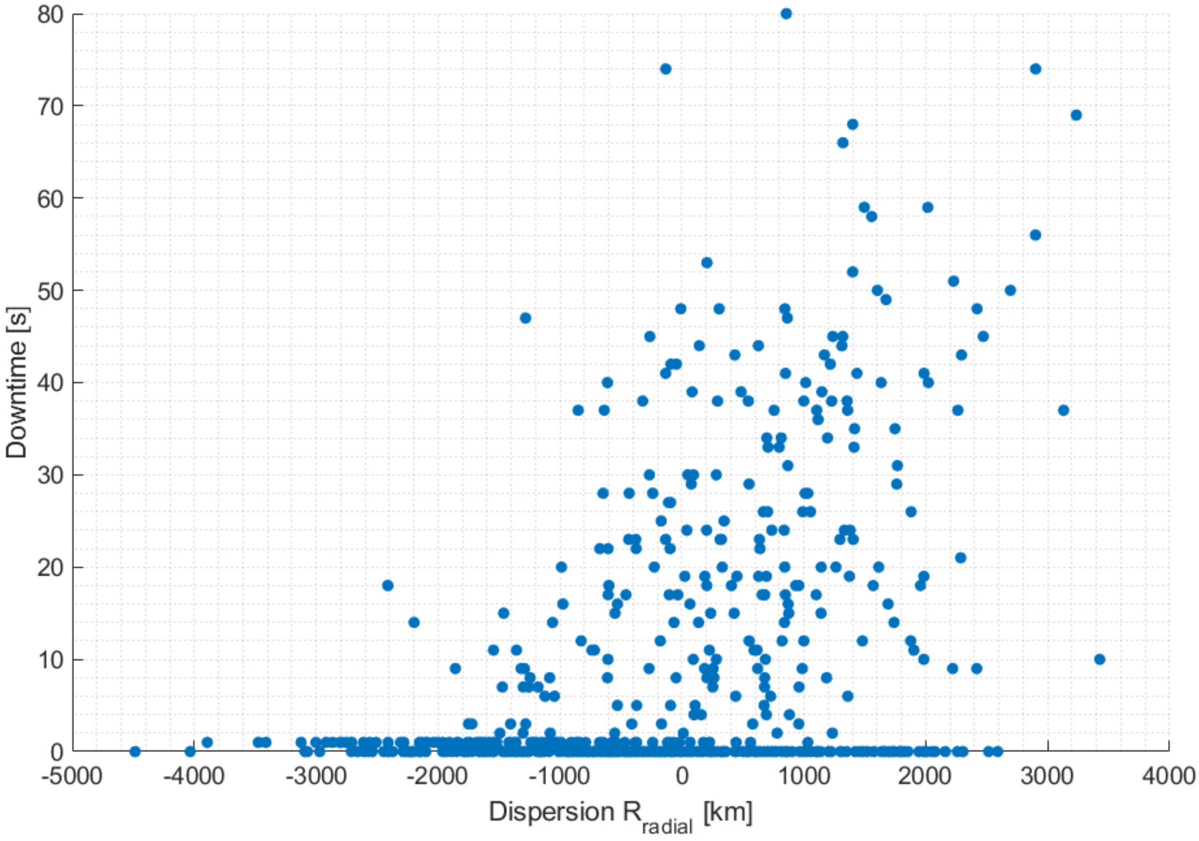

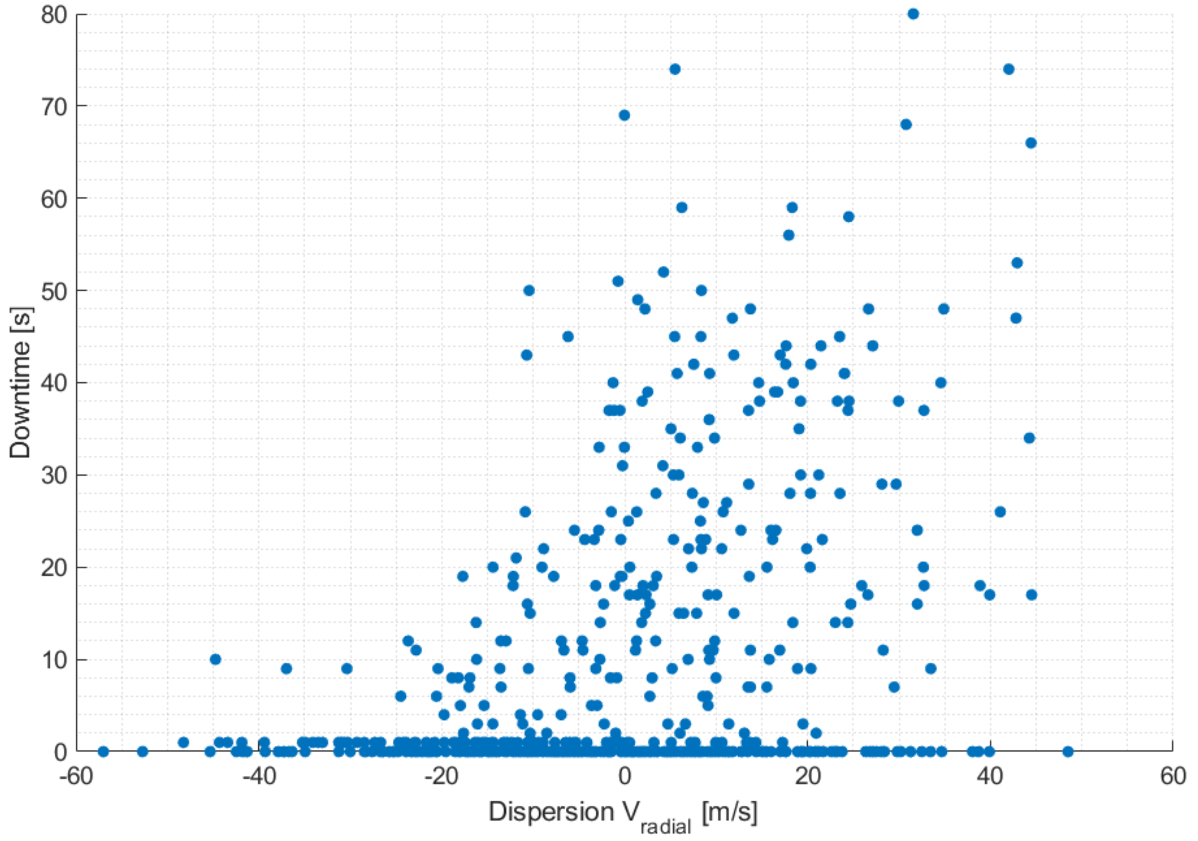

Figure 9.7: Values of initial dispersion of the position – at the top – and velocity – on the bottom – of the real trajectory of each simulation, in the radial direction.

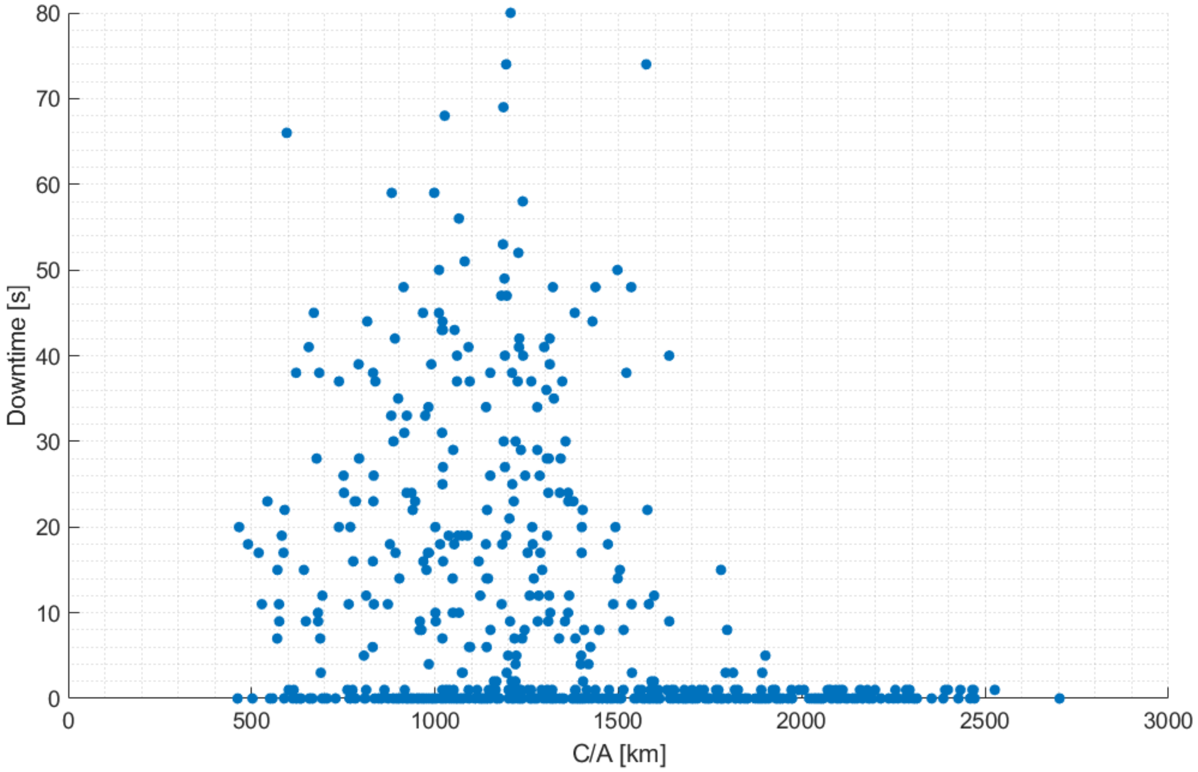

Figure 9.8: Downtime for each simulation being considered as a function of the closest-approach distance to the target comet's nucleus.

closest approach, but could perhaps be extended for the last 5 to avoid these issues, at a cost of a significantly increased simulation computational load. Regardless, either of these approaches would be simple enough to implement and help mitigate this issue by focusing the increase in measurement frequency around the time of the closest approach and not after.

The harder problem to solve is the closest-approach distance variation. Since during the encounter of the Comet Interceptor no control of the trajectory is used, the trajectory is defined *a priori* after the last manoeuvre during the preceding approach phase. As such, there is little to do after that to change the outcome in terms of the distance to the nucleus. And even more so because, in terms of the navigation performance, it would be ideal to increase this value to above 1500 km, where according to Figure 9.8 there is little chance for any downtime at all, but there is also the desired proximity to the target to take into account. Whereas from further away better navigation performances are achieved, as the distance to the comet increases, less information can be captured of the nucleus, which is tied to the mission's objectives. Thus, in this sense, it is not only much more difficult to "solve" this problem than the previous one, as it also should not be done without much deliberation, because it deeply affects other aspects of the mission which must be considered as well.

Chapter 10

Navigation performance

This section of the study will be dedicated to gathering the data to answer the core questions posed in Section 1.3. For that, firstly the tests being carried will be described. The analysis is divided into a baseline section, with no disruptive phenomena, and an outburst section, where the impact of these events is studied. Then, within each section, multiple analyses are made, discussing possible methods to address the results and drawing conclusions from them. This is particularly the case in the outburst analysis, where the multiple variables that define these events are studied to assess how much they impact the navigation.

Adding to this is a section where the failure cases of the unscented filters, which is an issue that appears and becomes quite relevant throughout the analysis, are discussed. There, possible methods to reduce the failure rate are also discussed.

10.1 Analysis method

Before drawing any conclusions, the methods used to test the various navigation solutions will be established. In total, there are six filters being tested, three different approaches with a sequential and a batch formulation. Adding to this is the innovation filtering, which will also be tested with each filter. These represent all the cases for the baseline scenario. They aim at determining which filters work better, in terms of performance, but also processing load, without any unexpected event-perturbing measurements. Not only this, but the innovation filter will also be analysed, as its functioning depends on a threshold value that can be adjusted to best fit the scenario. However, in the vicinity of comets, disruptive phenomena are to be expected, and as such, the basics introduced in Chapter 3 will serve to implement cometary outbursts into the scenario. Then with the impact of the outbursts being considered, the aim is to see whether the conclusions of the baseline scenario hold up in this case, and test if the innovation filter is able to mitigate the impact of these events.

For the former analysis, the method is as straightforward as running multiple cases with each filter, with repeating $rng()$ values, so that the results can be compared between one another and conclusions be drawn regarding the effectiveness of the estimators. To test the innovation filter, the only difference is this analysis must be repeated with it turned on and adjusted to various different values, to determine which range yields the best results.

As for the outburst analysis, there are more variables added to the scenario, mainly their intensity, orientation, duration and starting time. This adds to the complexity of the problem, as not only are the parameters of the scenario being considered, as well as the different filters and the innovation filtering threshold, but also various possible outbursts to be tested. All these variables make the process of obtaining data and drawing conclusions much harder, naturally, so some generalisation will have to happen in this part, especially regarding the variation of scenario-defining parameters. The method, then, will be to vary all variables and test them with a limited set of seeds, in an attempt to discard right away those which obviously cause little impact on the results, before pursuing a deeper analysis with as few variables as possible. This thematic will be developed further in the respective outburst impact analysis section.

The characteristics of both the baseline and outbursts analyses is summed up in Table 10.1. For the former, the filters, seeds, and innovation filtering threshold used are laid out. For the latter, the outburst intensities and directions are also added to this list.

10.2 Baseline analysis

10.2.1 Maximum pointing error analysis

In this section, as aforementioned, the goal is to establish which filter works best under the regular scenario without perturbations to the measurements. For that, all available filters were put against the same scenarios and the same pseudo-random-generating seeds. In this cases, the seeds are those from 0 through 50. To evaluate the performance of the navigation solutions, the pointing error will be the go-to variable. More specifically, the

Table 10.1: Characteristics of the baseline and outbursts analyses.

Baseline

Filters	EKF, Batch EKF, UDUEKF, Batch UDUEKF, UKF, Batch UKF
Seeds	[0,50]
Thresholds	$\infty\sigma$, 3σ, 4σ, 5σ

Outbursts

Outburst	Intensities $[R_{comet}]$	Seeds	Filters	Directions	Thresholds
5-2	4,6,10	[0,5]	EKF, Batch EKF, UDUEKF, Batch UDUEKF, UKF, Batch UKF	N, NE, E, SE, S, SW, W, NW	$\infty\sigma$, 3σ, 4σ
2-1	4	[0,5]			
	6	[0,20]			
	10	[0,5]			
1-0	4	[0,20]			

pointing error considered will initially be the maximum value over the simulation, which is what is represented in Figure 10.1.

This graph shows the results of the proposed baseline tested with each of the filters being considered for random-generating seeds considered. The boxes represent the 25th and 75th percentiles, as well as the minimum, maximum and median values for the gathered data of each estimator. Looking at these results, some aspects are clearly visible. Firstly, the minima of all filters are all close to one another, in the [0.1,0.2] degree interval, as well as the 25th percentiles, which lie around the [0.3,0.4] degree range. At the other end of the data, the 75th percentiles show considerably more variation, with the batch UDUEKF and UKF solutions almost coinciding at 1.6 deg, followed by the batch EKF and sequential UDUEKF at 1.5 deg and 1.45 deg, respectively. In this metric, it is visible that the EKF and batch UKF have the best results.

Where the filters behave the most differently, however, is beyond the 75th percentile, with the maximum results ranging from nearly 3 deg for the UKF to over twice that much on the sequential forms of the EKF and UDUEKF formulations. This shows that, robustness-wise, the batch formulations are more resilient than their sequential counterparts. Finally, in terms of median, all filters show results between 0.6 and 0.8 degrees, with the sequential forms here showing equal or slightly worse results than the batch equivalents.

Overall, analysing all the discussed details above, the filter that shows the most promising performance is the batch UKF, both in terms of general results and robustness. Nevertheless, this conclusion must be taken with a grain of salt, as both UKF filters present limitations in these testing conditions. The batch UKF and the UKF were not able to conclude their estimation efforts during all the simulations, as occasionally the covariance matrices would cease being positive definite, meaning that the generation of sigma points using the Cholesky factorisation would no longer be possible, stopping the navigation. This happened in 3/50 cases for the batch formulation, and 4/50 for the sequential, whereas no other filter faced such problems in the tested samples, including those to come in the next sections. This means that the UKF filters are, at least in these test conditions, not contenders for the category of most robust filters, as the others have a 100% simulation completion performance. This issue with the unscented filters is tested and discussed in detail later in Section 10.4.

10.2.2 Run-time analysis

These results only take into account the worst value of performance of the filters, but not the associated computing power. This can be done by graphing the average time taken by each run of the filters during each step of the simulation. In other words, being analysed is the average time that each filter function, and its respective children functions, takes to run during each individual step.

Figure 10.2 emphasises the different computational loads required to run each filter per time step, from beginning to end of the simulation. The graphs show the constancy of the sequential filters versus the peaks and troughs of their batch counterparts. This is because the former always perform the same operation regardless of iteration, whereas the latter alternate from a simple propagation, when no new data is available, to processing a whole batch of data at once, which naturally translates into a heavier computational load than doing a it for a single measurement. As a result, however, it is hard to assess from the graphs whether the batch alternative is worth it from a temporal perspective. For that, the average total run time is the correct metric. This is the sum of all the time spend running each of the filter functions, and again all their respective children functions, averaged

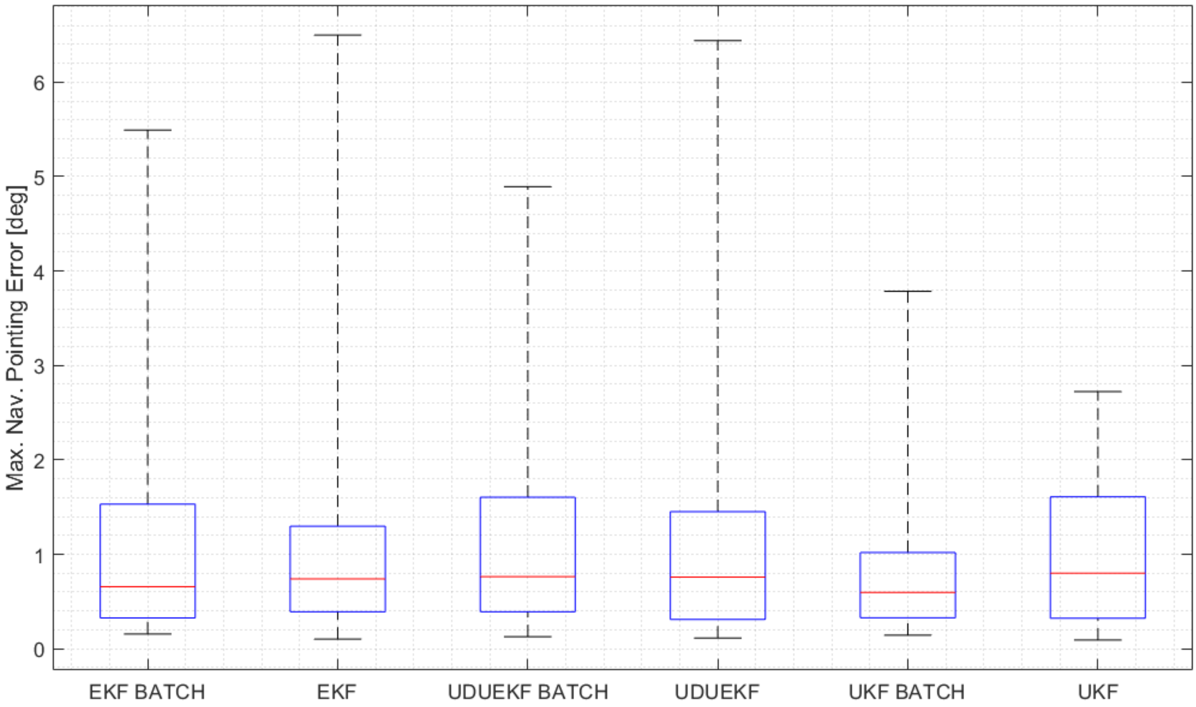

Figure 10.1: Box-plot of the maximum pointing error results of 51 simulations carried out for each of the filters being studied.

Table 10.2: Average total run time for each filter and nested functions being analysed. Data obtained from seeds 0 through 50, and only considering the initial 100 steps.

Filter	Average total duration [s]	Standard deviation [s^2]
Batch UKF	122.6	36.3
UKF	123.5	9.3
Batch EKF	21.7	3.8
EKF	11.6	1.1
Batch UDUEKF	27.0	8.8
UDUEKF	15.3	3.6

out over the available data set.

The data in Table 10.2 shows that, despite being faster in the majority of the steps, the batch EKF and UDUEKF filters are significantly slower than the sequential formulations. Also, in general the UDU formulations take longer than their corresponding EKF versions, as the UDU update algorithm requires updating the covariance matrix via a loop operation, which naturally takes longer than a simpler matricial operation. Surprisingly, the trend of batch versions being more cumbersome than sequential filters does not apply to the UKF, where the sequential form takes nearly the same time as the batch. Regardless of this, the main difference is how much longer these UKF filters take with respect to the EKF and UDUEKF ones, which is approximately 5 to 10 times more, given the big amount of sigma points that have to be propagated, in the former, and their respective measurement simulated at each time step, whereas for the latter only the current state estimate is propagated and its respective measurement considered. Since the vast majority of time spent performing estimations in these filters is allocated towards doing these exact two tasks, it is no wonder the UKF filters are so much heavier time-wise.

As to why some batch formulations are faster than the corresponding sequential ones, the answer again lies with the number of propagations being performed. In the batch EKF and UDUEKF, while waiting for a batch of measurements to be completed and processed, the filter merely propagates the best available knowledge forward in time. However, at the time step when the batch is processed, instead of keeping this knowledge, the filter instead starts from the beginning of the arc and propagates again, starting after the last measurement batch was processed. By re-doing these propagations, the filters naturally take longer, which is why the difference is seen, but it is not significant enough in the case of these two filters to require these intermediate states to be saved over time to save computational power, since only the previous estimated state is propagated. The same

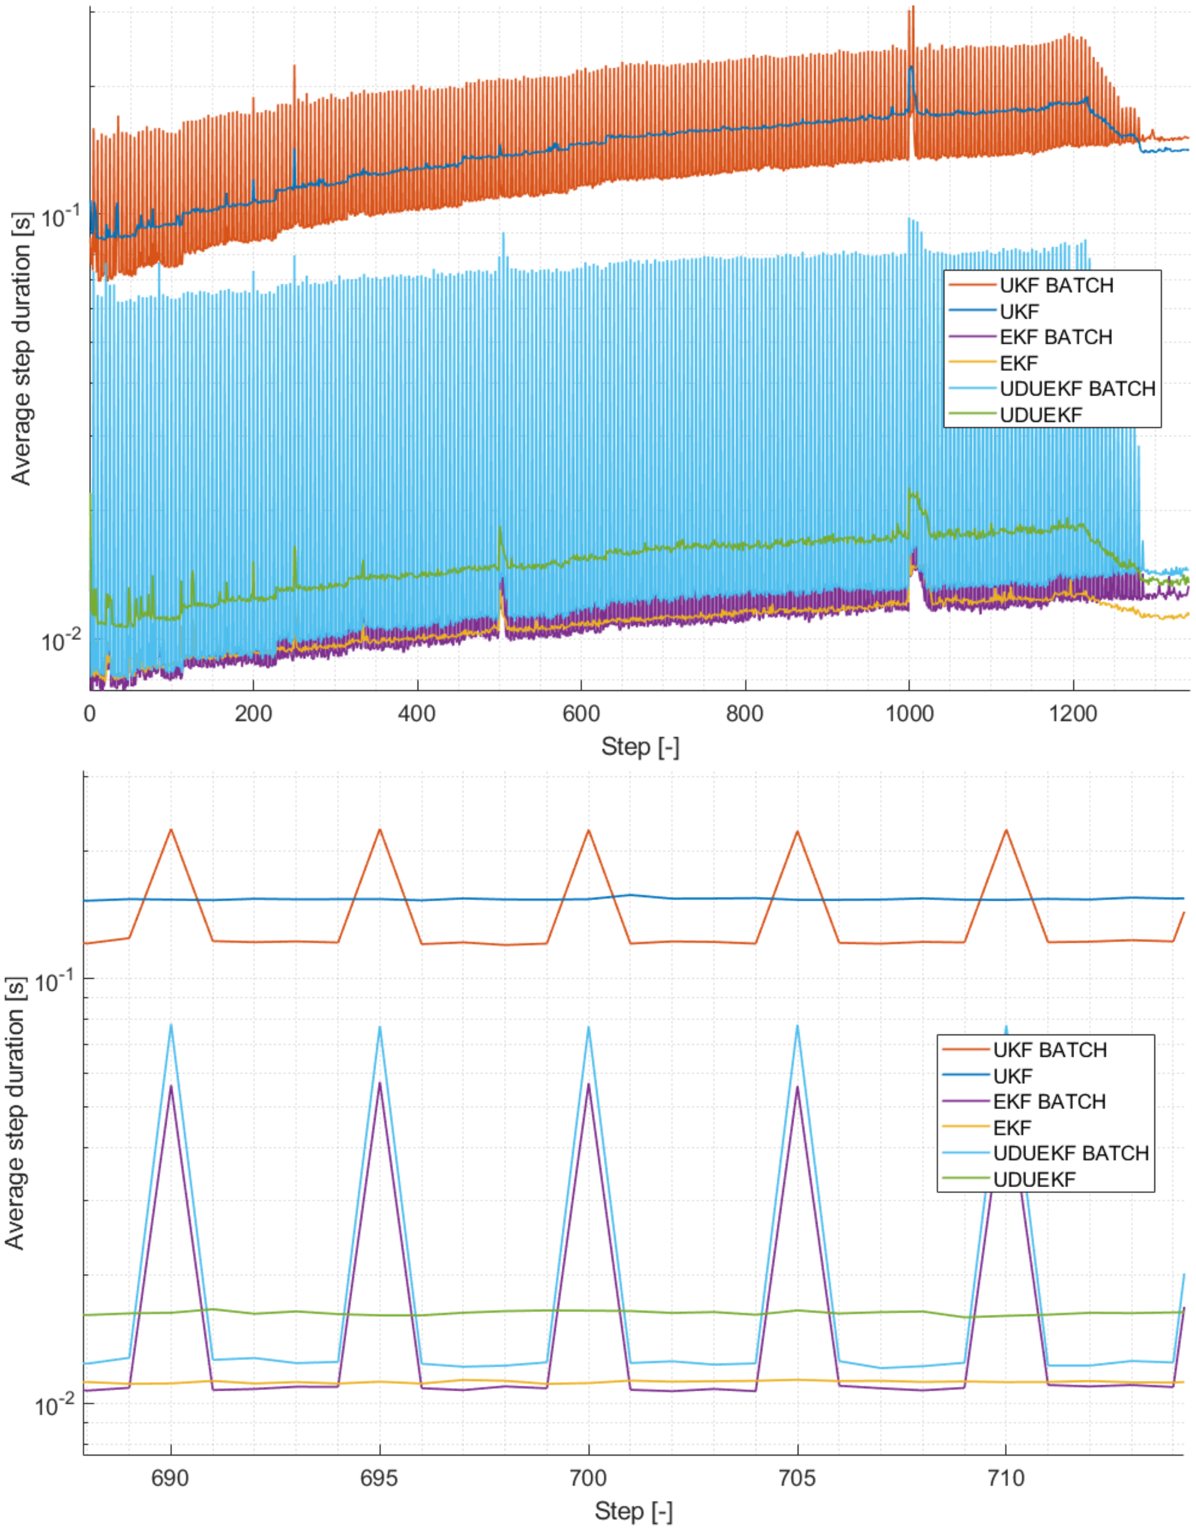

Figure 10.2: Average step duration of each filter along the baseline scenario.

cannot be said, on the other hand, for the case of the batch UKF, where, since there is such a large number of sigma points, their repeated propagation would most likely represent a steep rise in run time. Also, because the intermediate propagations are saved and reused, it makes sense that the average run times for the UKF and batch UKF are so close to one another. In fact, it is expectable that, if a similar rationale were to be applied to the batch EKF and batch UDUEKF by saving the propagations from the steps without batches being processed, the times of these filters and their sequential counterparts would be much closer, such as what can be seen in the case of the UKF. Thus, the main conclusion from this analysis is that the major difference in computational load is observed between the extended filters and the unscented filters, with the latter being by far the most cumbersome.

Another relevant aspect to point out is the average nearly linear increase in step duration as the simulation unfolds for every single filter. This is likely the manifestation of the increase in data size with each step, and the lack of optimisation to avoid the usage of big data structures within the code. In general, within FCS ATOMIC, the initial data and simulation parameters as well as the results thereof are saved within data structures, which naturally grow in size as more results are produced. Thus, carrying these variables into and out of the filter functions – as well as those functions nested within them – at their beginning and end, respectively, results in an increasingly larger computational burden which is then seen in Figure 10.2. This is a poor coding practice, that could with some work be avoided in this study, but that was not done as always having the majority of the scenario's variables and results within reach when developing the code, debugging it and analysing the results proved to be a huge advantage. To mitigate this effect on the step and total duration analysis of the filters, however, the results shown in Table 10.2 comprise just the duration measurements for the initial 100 steps, a period during which this problem has a minimal effect. Then, the duration of these 100 steps is multiplied by inverse of the fraction of the total time that they represent, so that a more accurate estimate of the comparative computation times taken by each filter in an optimised environment can be obtained.

Also in Figure 10.2, the step duration curves drop significantly somewhere between steps 1200 and 1300. This is a consequence of the absence of new measurement information after the spacecraft gets too close to the comet's nucleus for the NAVCAM to be able to capture the target. As such, from this moment on, which depending on the randomness of the scenarios can happen sooner or later, but always around the 20 hour mark, the filters stop receiving new information to perform the *a posteriori* updates on the estimates, thus reducing the computational load of the whole algorithm.

Regardless of the duration graph's details, it is clear that the sequential EKF and UDUEKF are the fastest to run, closely followed by their batch forms, because of the discussed reasons, whereas the UKF filters, in both formulations, take considerably longer to run than the rest. As a result of this and of the lack of robustness, the UKF filters cannot be said to present the best performance to computational load ratio of the bunch. Additionally, in terms of the standard deviation of the analysed times, the EKF and UDUEKF present the lowest obtained values, with a trend that bigger average total durations also have higher standard deviations. Notably, the batch UKF does present a value significantly above all the others. However, relative to the respective average total duration, it is about equal to that of the batch UDUEKF, for example, at nearly 1/3 of the mean.

Because of that, the EKF is likely the best choice for this baseline case, with very good percentile and median values, a solid computational efficiency, only lacking in the maximum leg of the box-plot. However, since for the purposes of pointing the payload, 2 degrees of error is just as bad as 6, this maximum value does not really change much, since the threshold for data acquisition is so much lower than it.

10.2.3 Pointing error threshold analysis

The previous analysis raises an important question regarding the temporal relevance of this data. Looking at the maximum pointing error may not show the full picture. For example, it could be case that one filter has higher maxima and thus shows worse results according to this analysis, but has a better overall performance, with lower pointing error for a greater portion of the mission. To assess this possibility, ranking the results according to the maximum pointing error is futile. Instead, the notion of pointing error threshold is again reintroduced to evaluate the results of the navigation over time.

Analysing Figure 10.3, the conclusions differ quite significantly from those gathered from the box-plots of the previous analysis. Whereas in the previous graphical representation it is impossible to show the failed cases of the sequential and batch UKF, in this successful scenarios percentage graph this can be shown, with the batch and sequential UKF lines starting below 100% before the closest approach, thus considering these cases as unsuccessful. As such, it is possible to see not only how each filter performs on average over time, but also how the failed cases for the unscented filters affect their comparative performance. Whereas the sequential UKF appears to be one of the worst performers for the majority of the analysed period, the batch UKF holds quite well when compared to the other filters, even being the best performer during certain intervals.

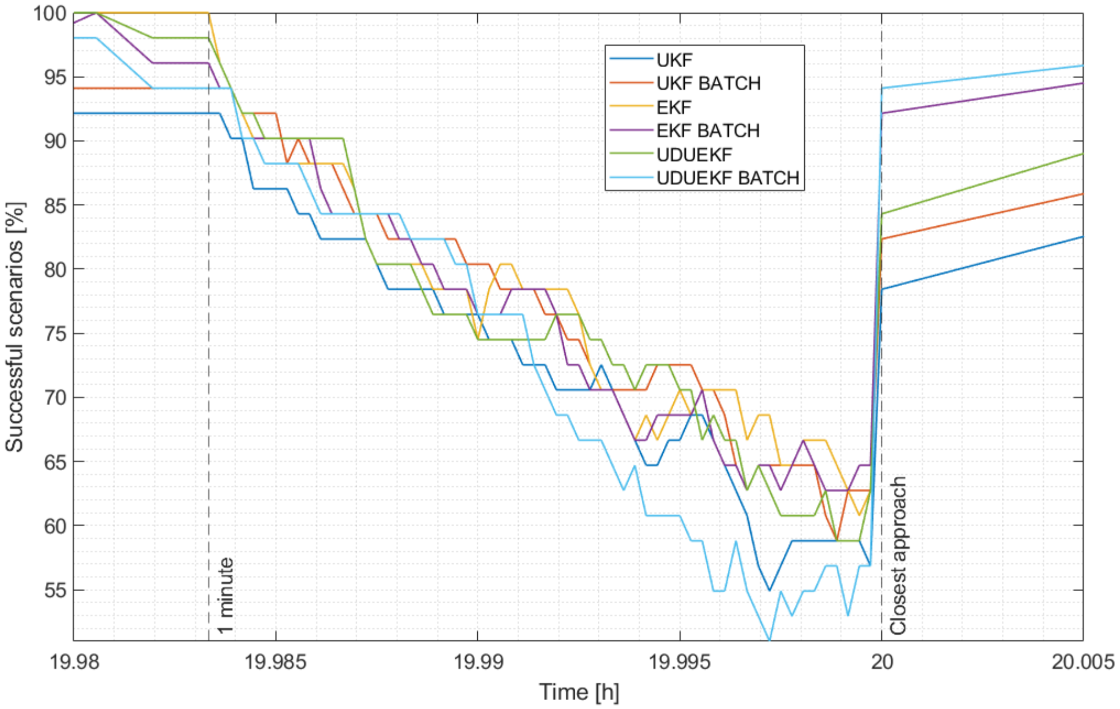

Figure 10.3: Percentage of scenarios with a pointing error below the threshold of 0.5 deg. Emphasised here is the section of the mission right before the C/A, since this is when loss of sight realistically has the potential to happen.

Nevertheless, what this graph more plainly shows is that, taking into account the UKF failed cases, the conclusions about the best filter for the navigation of the Comet Interceptor become much harder to draw than previously. However, analysing the graph closely, it can be concluded that, overall, the filters that represent the best result most consistently are the EKF, batch EKF, UDUEKF and batch UKF. Whereas the batch UDUEKF shows promise in the early stages of the closest approach only to then be the worst as the nucleus approaches, and the UKF is almost never as good as the former filters or as bad as the latter, these four compete most of the time for the highest rate of successful scenarios. Thus, from this graph it is easier to delineate losers rather than winners.

There is, nonetheless, another way of inspecting the data. Whereas in the previous paragraph the best performing filter at all times was the metric being considered, it can instead be argued that a good filter should have the highest possible successful scenario rate minimum. From this point of view, the UKF and batch UDUEKF again have no saving grace, as their results are significantly below the other contenders, with minima at approximately 55% and 50%, respectively. However, for the other four filters, the batch EKF now behaves the best, followed by the batch UKF and then the EKF and UDUEKF, with minima at about 62%, 60% and 58%, respectively. All these filters show very close results, and these complement the previous analysis in concluding that, despite not existing a clear best choice, there are two worst choices in the batch UDUEKF and UKF filters.

Finally, the question that still needs to be answered regarding this baseline scenario is whether these results are changed by introducing the innovation filtering, which will discard measurements it deems too bad by comparing their residuals with their respective covariance. For that, the analysis will be shifted, looking now at the downtime metric, which represents for how long a certain scenario presents a pointing error above the 0.5 degree established threshold.

In Figure 10.4 the average downtime for all the scenarios tested, again with seeds 0 through 50, for all six filters being analysed, and three different outlier rejection thresholds, is represented. Do note that $\infty\sigma$ means that the threshold is at ∞, which results in no outlier rejection, corresponding to the case being studied up until this point.

The shown results are mixed. For some filters, like the EKF and batch UDUEKF, the variation of the threshold parameter does very little to the outcome. For other, such as the UKF, the story is quite the opposite, with a 3σ threshold reducing downtime by 2-3 seconds with respect to the other options. Thus, in terms of the effect of introducing this parameter into the scenario, it can be said that it has a very positive impact on the performance of the UKF, whereas for the other filters, in general, the impact is either pretty insignificant, or detrimental.

Regardless of the impact of this technique, Figure 10.4 clearly shows that the EKF, for whatever value is chosen as the threshold, is the best alternative by having the lowest average downtime of the bunch by almost 1

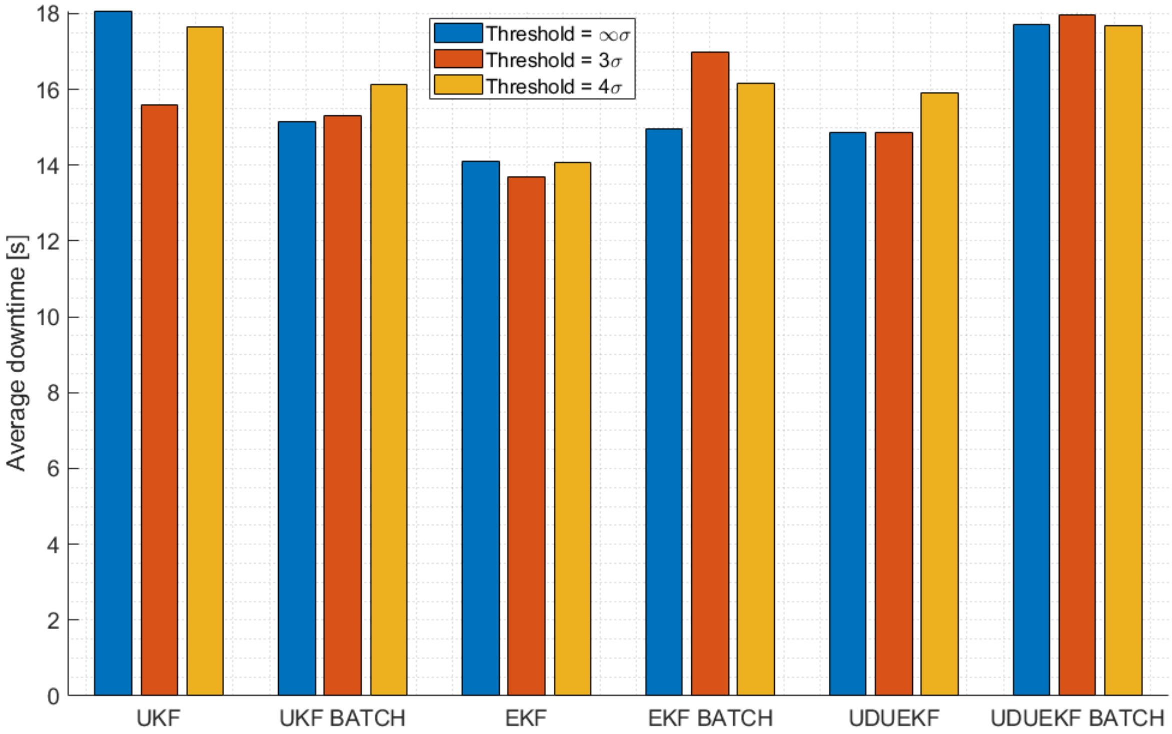

Figure 10.4: Average downtime for all the tested filters, with varying measurement filtering thresholds.

seconds with respect to the other closest filters. It also varies the least – similarly to the batch UDUEKF – to this parameter, meaning that the EKF can be chosen as the best performer for the baseline scenario without much concern for the outlier rejection threshold.

10.3 Outburst analysis

10.3.1 Outburst definition

Now that the baseline scenario has been studied, the impact on the navigation caused by outbursts is analysed here. As aforementioned at the beginning of the chapter, the goal is to check whether the same conclusions hold up for this case as for the regular undisturbed baseline scenario. Here, the filters will be exposed to a situation in which the measurements are perturbed for a certain time period, and unlike all the other parameters regarding the dynamics or the acquisition of measurements, the spacecraft has absolutely no knowledge of these events. As such, the filters – as well as the innovation filtering process – will be tested in their capability to recognise that measurements deviate significantly more than expected from the predicted, and thus adjust their estimation accordingly, and how this affects the final results.

Before doing so, however, one discussion is required. This is because of the four degrees of freedom mentioned in Section 10.1 to specify an outburst. These were duration, intensity, orientation and starting time. For the first three parameters, there are some guidelines to follow from studies made about previous missions based on the images these took. However, as for the time at which the outbursts happen, there is really no good way to predict it. The study space can, nonetheless, be significantly shortened. This is because the size of an outburst will be somewhere around the order of magnitude of the nucleus' size itself, as can, for example, be seen in the observations catalogued in [Vincent et al., 2016].

Assuming this, for an outburst that normally lasts a few minutes and has an intensity as described, the effects of its presence will not be felt for the vast majority of the mission. First of all, for the second half of the scenario, starting after the closest approach, the nucleus is no longer in sight, so there is no use in even considering it. In turn, during the first half the spacecraft approaches the nucleus for 20 h at 70 km/s, meaning that up until the last two hours before the closest approach, approximately, the spacecraft will be too far away to gain any useful information from the measurements, and the same thus applies to any outbursts. Again, this can be seen in Figure 8.2.

As a result, it is during this period of the simulation that the outburst can have the most impact, and it is the only one that needs to be studied. Regarding the rest of the 20 hours, running some simple tests does verify

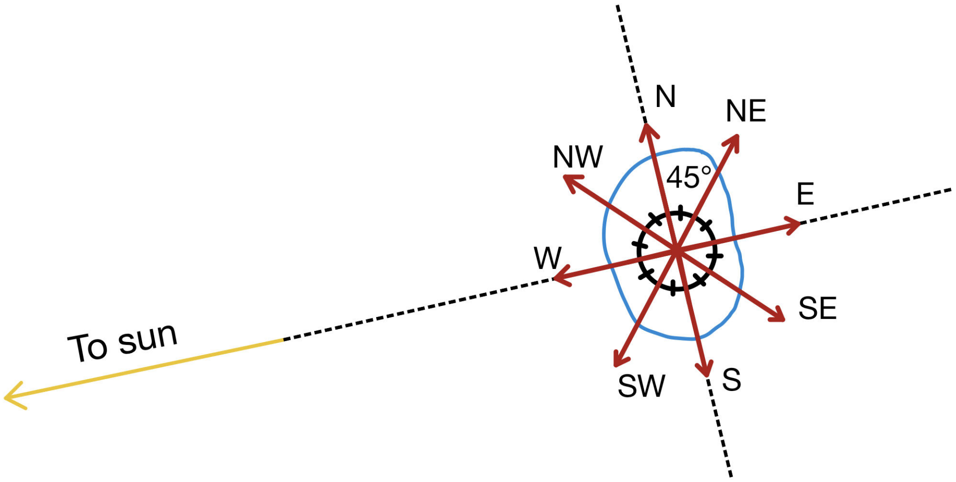

Figure 10.5: Representation of the various considered directions for the studied outbursts, labelled similarly to the directions on a compass.

that, indeed, the impact of any outburst of the considered order of magnitude of the nucleus before this period is imperceptible by looking at the results. Thus, to evaluate the consequences of outbursts, only those happening right before closest approach are considered relevant, and will be the focus of this section. In this period of time, there is still quite some margin as to when these events can occur. However, following a similar train of thought as in other parts of this thesis, the considered cases will be pessimistic, and happen right as the spacecraft is closest to the target. Initially, only two cases are considered, the first starting 5 minutes before the closest approach and vanishing 2 prior to it, and the second starting when the first ends and ending at 1 minute to the 20 hour mark. These shall, for simplicity's sake, occasionally be named 5-2 and 2-1 outbursts, respectively. After this initial analysis, a better idea of the impact of these times will allow for more examples to be chosen and studied. Importantly, these two cases decay in intensity over time linearly, from 100% strength when they appear to 0% at the instant that their duration is over.

With the starting time and duration defined, the outbursts only need their intensity and orientation to be fully defined. The intensity can have a high degree of variation, but for the purpose of this study, it is not useful to analyse the cases where the bias is too small to cause a meaningful impact, or too big in a way that becomes unrealistic, or goes beyond what the navigation can reasonably handle. Within this frame, the values chosen were {4,6,10} times the size of the nucleus' radius. These were taken as multiples of the radius as this makes the magnitude of the bias more comprehensible.

In terms of orientation, these phenomena are assumed to be radially outward-pointing from the target's centre, as a simplification, despite the nucleus not necessarily being spherical. Thus, this becomes a question of choosing where the phenomena originate from on the surface of a sphere. Clearly, if the origin is behind the comet or facing the spacecraft, then their effect will be mitigated by the geometry of the problem. When facing the camera, their impact will only magnify the nucleus' brightness, whereas when facing the opposite way, the nucleus itself will block part of the outbursts. As such, these perturbations are chosen to be oriented outwards from the horizon of the nucleus as seen from the spacecraft at the moment they begin. This choice ensures that only the most affecting outbursts are chosen, since they do not suffer from either of these problems. In doing so, the problem of choosing an outburst direction is also transformed from a 3D space to a 2D one, massively reducing the variable space. Then, eight different orientations are chosen, starting with the sunward direction and going around the terminator at 45 degree intervals, as depicted in Figure 10.5. For simplification, these are named according to the directions in a compass.

Considering the topics discussed above, the outbursts starting at the established time, with the defined intensities and directions can be added to the baseline scenario and affect the measurements. Then, the different variations of outbursts can be tested with the different filters being used, varying values of measurement rejection threshold, and changing random generating seeds. Unfortunately, due to the number of variability in this whole experiment – 3 intensities times 8 directions times 6 filters – means that per seed and threshold value combination defined 144 simulations must be run. Thus, the goal will be to run a small initial set of data, analyse the results, and try to see whether the analysis can be pursued in a smaller variable space.

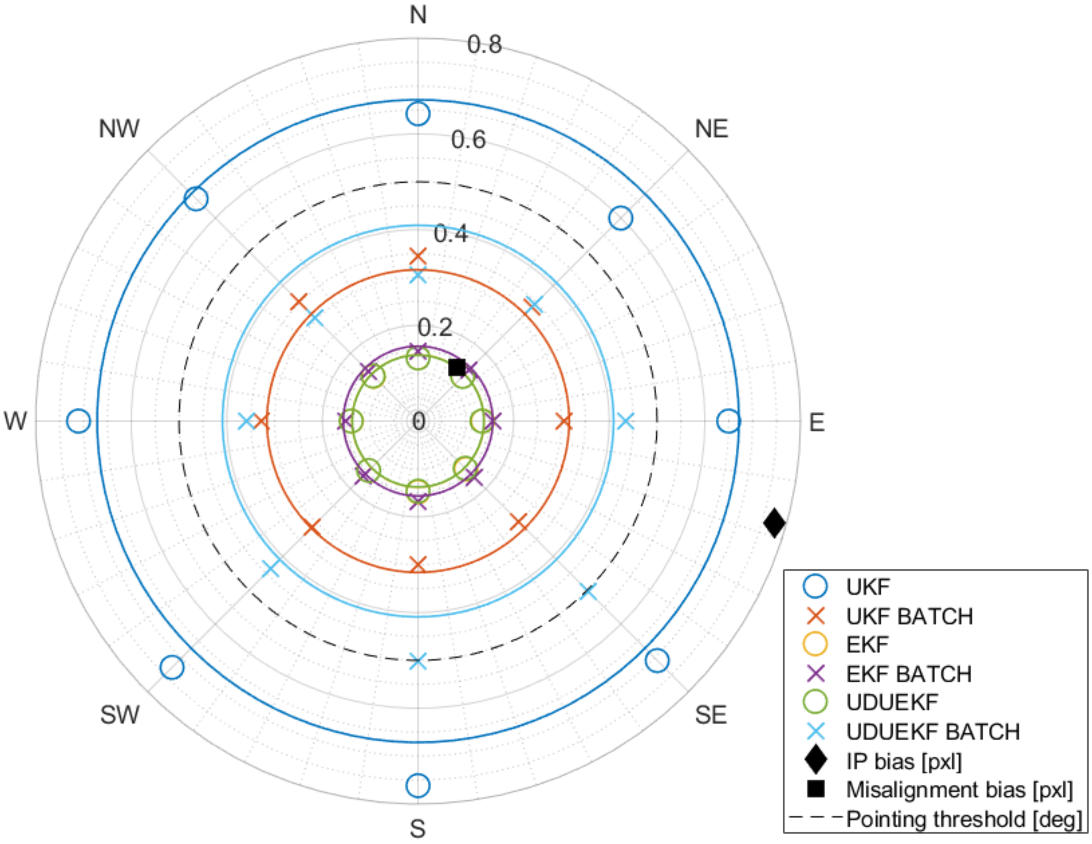

Figure 10.6: Polar plot of the maximum pointing error, in degrees, for the simulations with a 5-2 outburst, an intensity of 10 radii, and various directions. No measurement rejection is considered.

10.3.2 Maximum pointing error analysis

The initial dataset is gathered using seeds 0 through 4 and measurement rejection thresholds of $\infty\sigma$ and 3σ. The 4σ and 5σ values were also tested, but it quickly became clear that little to no change was seen with respect to using no rejection at all, in similarity with what was concluded from Figure 10.4. These first results are represented in the polar plots in Figure 10.6 through Figure 10.9.

These polar plots represent how each of the filters perform under various outburst conditions. In all these images, the same random number generating seed is used, and no measurement rejection is considered, for simplification. What does change between the four is the starting and ending time of the outbursts being considered, as well as their intensity. From top to bottom, firstly 5-2 outbursts are considered, with 10 radii of intensity, followed by 4, 6, and 10 radii of intensity but for 2-1 outbursts. Then, for each of the cases, the eight directions of outbursts are studied, named accordingly with the directions on a compass. For each direction, a symbol indicates the maximum pointing error of the simulation performed under the respective conditions and with navigation done by each filter. Additionally, for comparison, with the same colour used to represent each filter is a circle representing how that same filter performs for the identical conditions but without any outbursts. In other words, the coloured circles represent the baseline performance of the respective filters. Finally, a black dashed circle is present in all graphs at 0.5, representing the established threshold of 0.5 degrees above which the CoCa looses sight of the comet's nucleus. Do note that in all four graphs, the baseline for the EKF cannot be seen as its results are very close to those of the UDUEKF, so the lines coincide. The same also happens in Figure 10.6 but for the maximum pointing error marks.

What these graphs show is then the height of the pointing error peak reached during various outburst and navigation conditions, to study the effect of these events on the overall mission. The first thing that stands out is that, separating the figures into two sets, one with the 5-2 minute outbursts and the other with 2-1 minute outbursts, is that the former shows small variation with respect to the baseline case without any outburst represented by the coloured circles. Regardless of the considered directions or intensities – of which only the highest one is represented here due to how similar the graphs end up – the results change little with respect to the unperturbed, meaning that this time period before 2 minutes prior to the C/A is not significantly sensitive to outbursts, at least of the considered intensities.

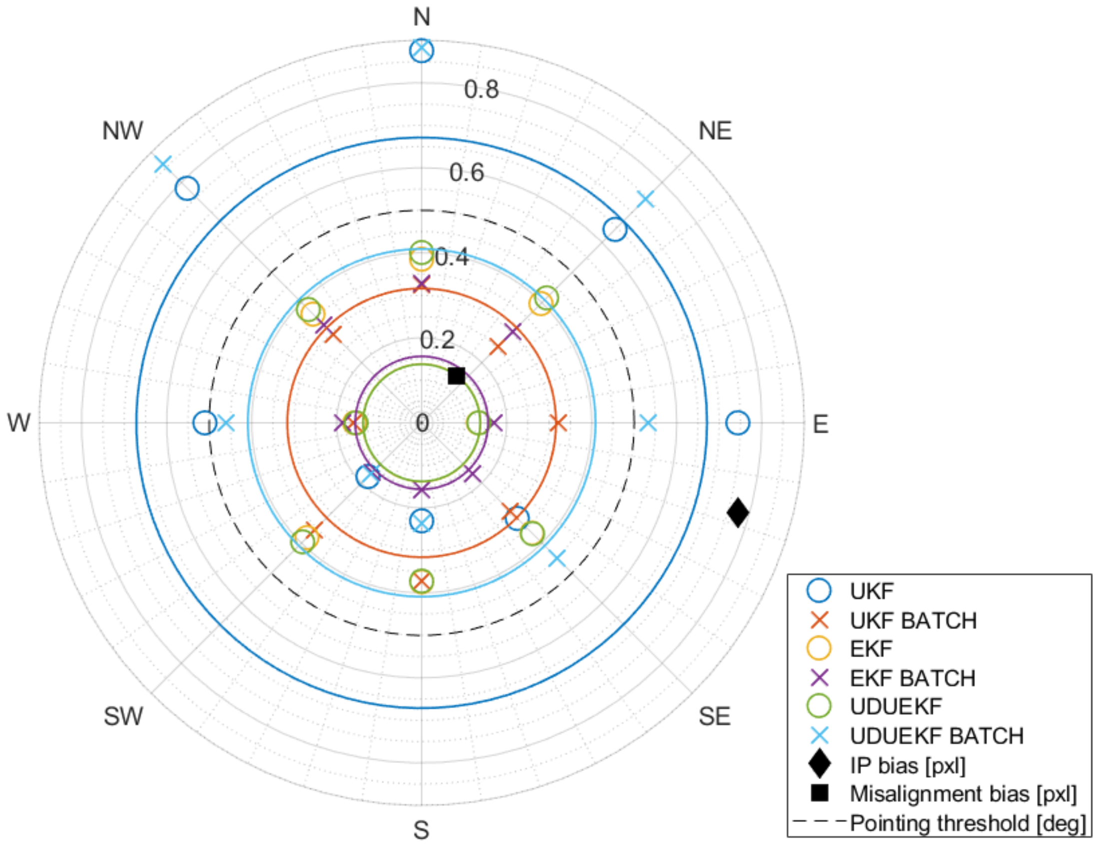

Figure 10.7: Polar plot of the maximum pointing error, in degrees, for the simulations with a 2-1 outburst, an intensity of 4 radii, and various directions. No measurement rejection is considered.

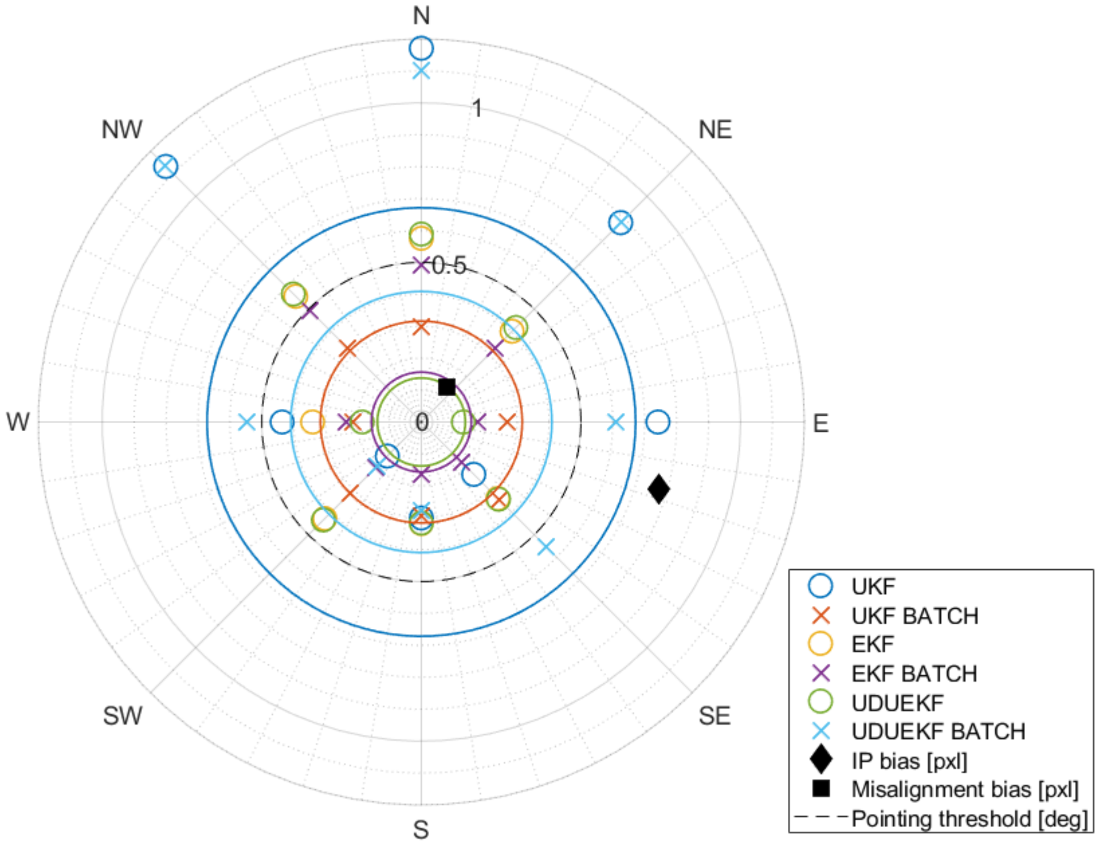

Figure 10.8: Polar plot of the maximum pointing error, in degrees, for the simulations with a 2-1 outburst, an intensity of 6 radii, and various directions. No measurement rejection is considered.

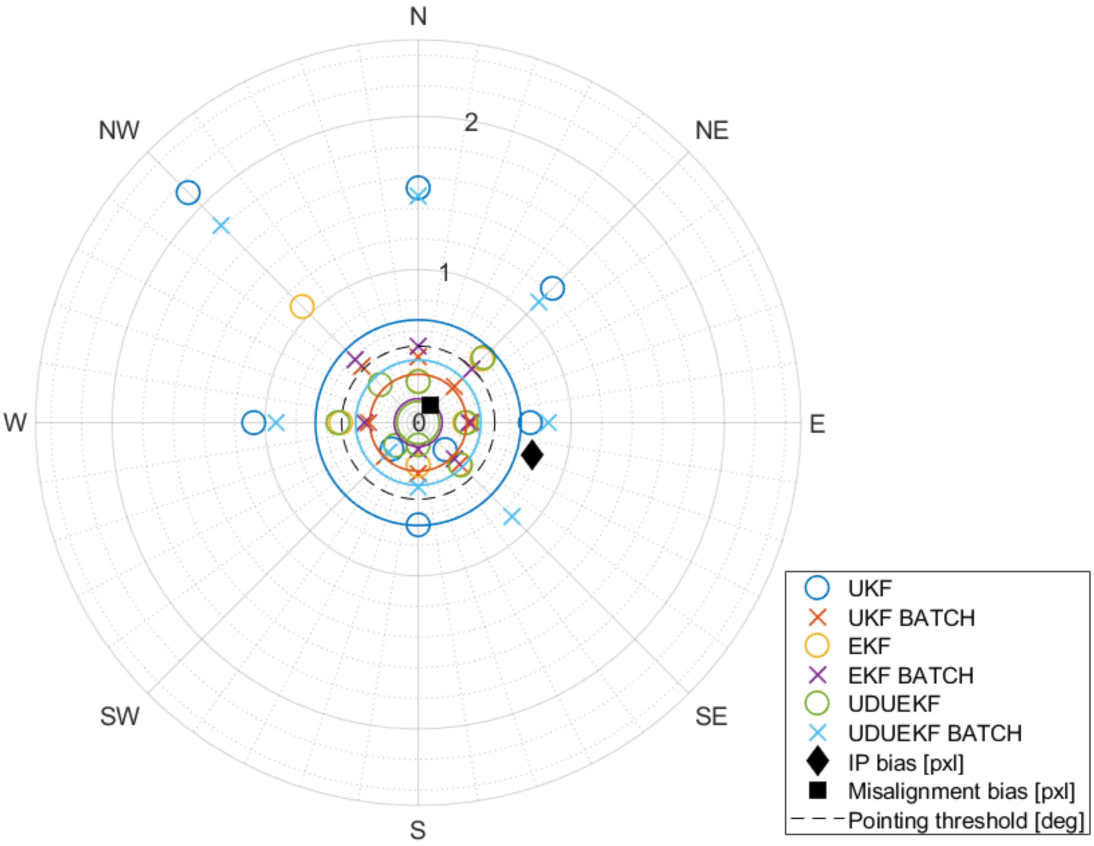

Figure 10.9: Polar plot of the maximum pointing error, in degrees, for the simulations with a 2-1 outburst, an intensity of 10 radii, and various directions. No measurement rejection is considered.

In the case of the latter, however, the conclusions are quite different. Immediately at 4 radii of intensity, the lowest value, it can be seen that, whereas certain directions of outbursts lead to results somewhat in accordance with those of the baseline case, for the rest these maxima behave quite unpredictably. Whereas in the previous type of outbursts the direction and intensity appeared to have little to no change on the results, now different directions certainly see the effects of these perturbations differently, and analysing the three graphs referring to the 2-1 outbursts, it becomes apparent that a larger intensity does seem to emphasise these disparities.

Of the many results obtained from this analysis similar to those shown, these – obtained with seed 4 – were chosen as they depict the trends shown by the majority of the cases. These trends were:

- The 5-2 outbursts do not meaningfully impact the results – not even with 10 radii of intensity, as can be seen in Figure 10.6. This confirms the hypothesis that the measurements closest to the point of C/A are the most crucial, meaning that outbursts happening right before the spacecraft swings by the nucleus cause the most impact.
- For the outbursts nearer to the C/A, the intensity of the event has the tendency to exacerbate the disparities between the baseline and observed results. For some cases, this means increasing the observed maximum pointing error, but the opposite is sometimes also observed. Again, this is aligned with the expectation that bigger biases caused by these events would provoke bigger changes in the results.
- Not shown above are the results using 3σ as measurement rejection threshold, because these show the tendency of vastly worsening the results in certain cases, instead of the opposite. What this suggests is that, despite being corrupted by these phenomena, the measurements taken in the few minutes or seconds before the 20 hour mark are so important to navigation that the intensities used for these perturbations were not large enough to justify discarding the images. This is exemplified in Figure 10.10, where the right graph shows much higher maximum error values than the left for certain filters and outburst directions, due to the introduction of measurement rejection.

The analysis done up until this point, although limited in terms of the small number of random scenarios considered, is already enough to gather some important conclusions. Most crucially, the effect of intensity on the results is fairly clear, meaning that, going forward this variable will be removed from the problem, and only

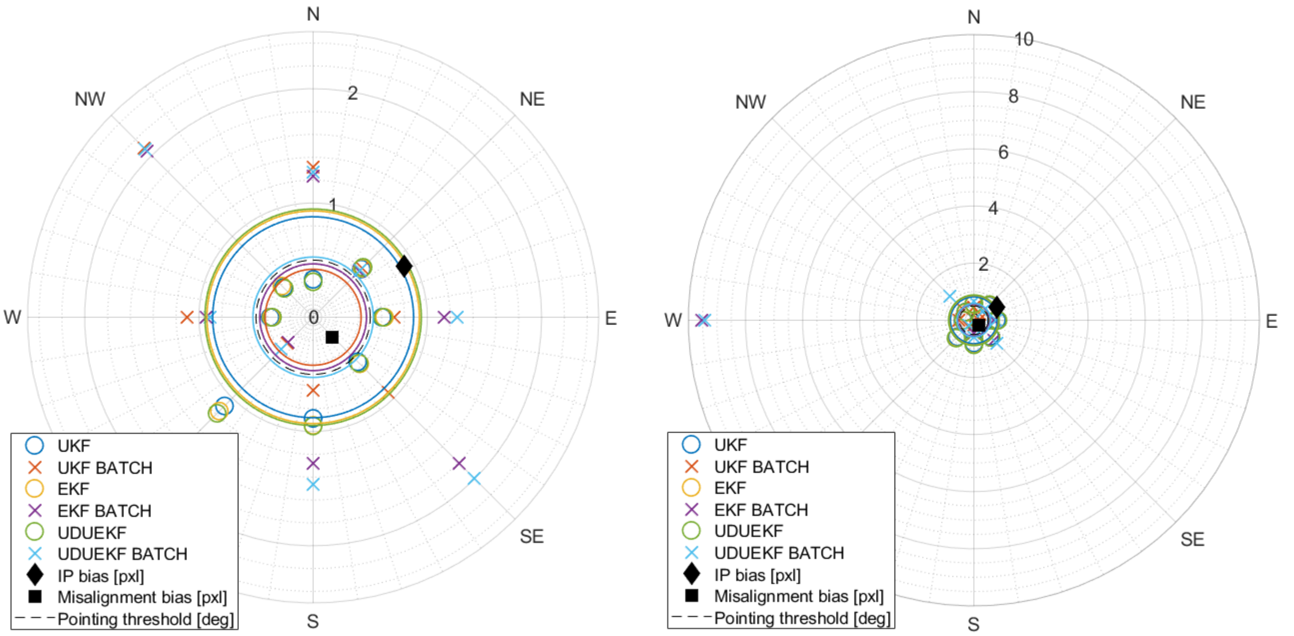

Figure 10.10: Maximum pointing error for the same scenario seed and outburst intensity. On the left, no measurement rejection is used, and on the right, a threshold of 3σ is used for that.

one value of intensity will be considered per each outburst being studied. Additionally, since the 5-2 outbursts had little impact, whereas the same cannot be said about the 2-1 ones, the zone of interest to be studied can be somewhat limited, with outburst starting 2-3 minutes or more before the closest approach not affecting the results in any relevant way.

This does lead to the question of what happens with an outburst in the 1-0 period. For this, the vary same simulations are run again, considering now this outburst, just considering an intensity of 4 radii of the nucleus. The lowest intensity is chosen because, for this outburst, the spacecraft is half as close to the nucleus as for the 2-1 case, thus a lower intensity should lead to similar results when compared to a more intense outburst starting sooner.

Doing so does, in fact, re-confirm the hypothesis that the later the outburst takes place, the worse its impact on the navigation. Surprisingly, the maximum error levels with outbursts of 4 radii of intensity were comparable or greater than those of 10 radii but happening just a minute earlier for the 2-1 outbursts. Nevertheless, another less positive conclusion can also be drawn from this comparison. Whereas the general order of magnitude of the maximum pointing error of the analysed outbursts was comparable, as just stated, the performances of the individual filters did not follow any clear trend. That is to say that, despite showing similar or worse results, the filters did not behave similarly in the 1-0 outburst with respect to the 2-1 case, with certain filters showing better results in one case than in the other. This means that there might not exist a relation between the period during which the outburst happens and the performance of each of the filters. Instead, these results hint at the possibility of not existing a clear favourite when it comes to choosing which filter is most suited for the mission at hand, instead indicating that the results can be unpredictable.

10.3.3 Analysis extension

To assess this possibility, this analysis must be extended. Now, only the 2-1 and 1-0 outbursts will be considered, with intensities 6 and 4 radii, respectively, as these cover the period of interest before the closest approach. Additionally, the intensity of the latter is chosen to be smaller to compensate for the temporal disparity, as explained above. Then, to compensate for the previous lack of random generating seeds considered, the number of scenarios run is also increased, with seeds spanning from 0 to 20. All these scenarios are then run with all six filters and the two thresholds for measurement rejection being considered, $\infty\sigma$ and 3σ.

However, this is not the single concern raised by the previous analysis. As done for the previous baseline study, here too it is important to not only draw conclusions from the maximum pointing error metric, but also to include a temporal representation of how the navigation performs over time. Thus, the method of using polar plots representing the maximum pointing error for each filter and outburst considered, as briefly used above, will now be replaced by a more fitting description, similar to the one performed previously, where the percentage of

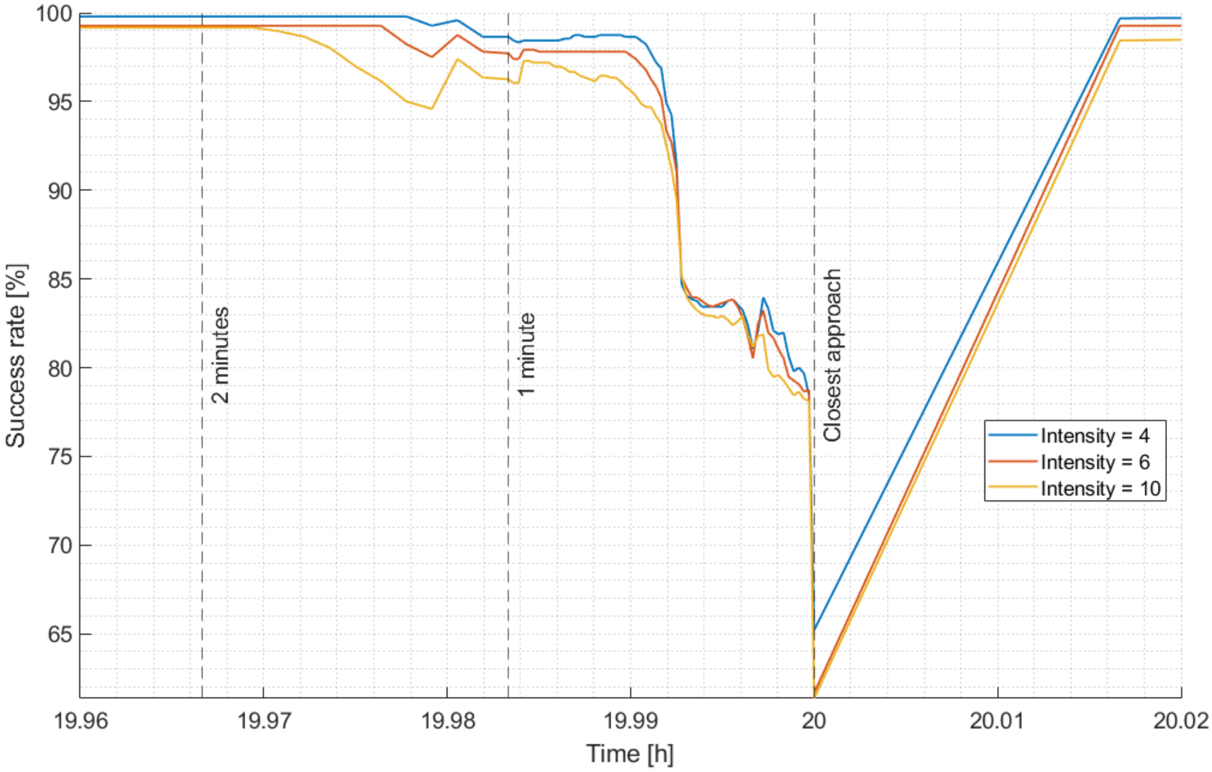

Figure 10.11: Percentage of cases for outbursts 5-2 and 2-1 and seeds 0 through 5 with pointing error below the 0.5 degree threshold.

scenarios with a pointing error below the established 0.5 degree threshold – or success rate – is graphed over time and as a function of the different variables of interest for this problem.

10.3.4 Intensity analysis

Before performing this analysis on the new dataset, however, let the previous data, with outbursts 5-2 and 2-1 with seeds 0 through 4 be analysed through the scope of the success rate, as is done in Figure 10.11. In the graph, the value for this metric corresponding to each intensity of the outbursts being tested is plotted over the simulation's duration. Despite the existence of other variables, such as the filters or measurement rejection thresholds being used, here all these variations are ignored. Thus, each line represents all the aggregated success rates for outbursts of intensity 4, 6 and 10, regardless of any other characteristics.

From the results in Figure 10.11, the conclusion established before can again be confirmed, with the higher intensities of outbursts lowering the success rate not only further but also starting sooner. One surprising result from this graph is that the outbursts with 6 and 10 radii of size show very little distinction at the closest approach. This highlights another shortcoming of the previous maximum pointing error analysis, which emphasises more the size of these maxima than the fact that they are above or below the threshold, that is what indeed determines if the CoCa is able to capture the comet's nucleus or not. Do note that, in this graph, as per the ones that follow in the next analysis, some of the curves do not reach the time zone being studied here with a 100% success rate. This is because of the aforementioned divergences in the UKF and batch UKF filters, which are counted as failures for the whole duration of the scenario, meaning that, when they exist, the success rates never reach 100% throughout the whole simulation period.

10.3.5 Outburst period analysis

With intensity out of the way, now remains the analysis of all the other variables. For that, firstly the impact of each of the outbursts is analysed. Here, all the data concerning the 5-2, 2-1 and 1-0 outbursts is considered simultaneously, despite these sets containing different numbers of seeds. The runs with different filters, measurement rejection thresholds, outburst intensities and orientations are considered all simultaneously here as well, just like previously, meaning that the only variable is not just the period of the outburst. The hope is that doing so will expose the trends in the success rate of scenarios dictated by each of the outbursts periods being studied.

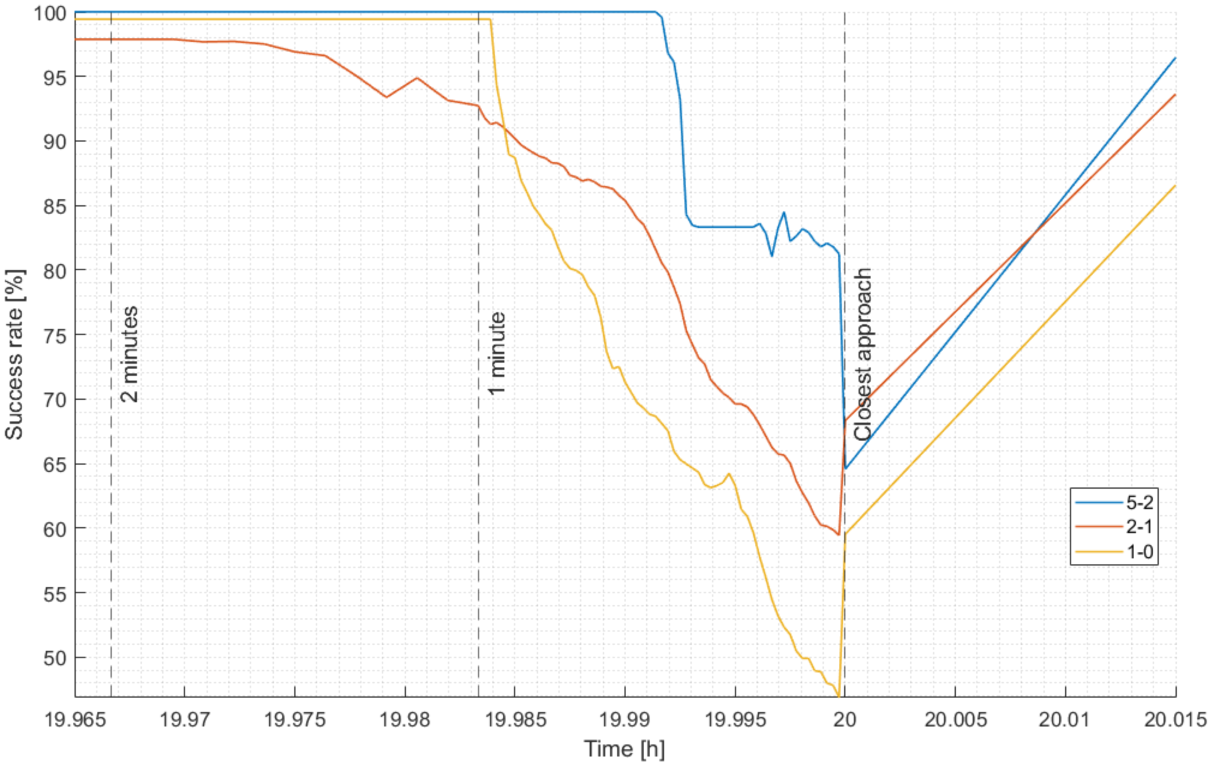

Figure 10.12: Percentage of cases for outbursts 5-2, 2-1, and 1-0 with pointing error below the 0.5 degree threshold, with all filters, measurement rejection thresholds, outburst intensities and directions combined.

In Figure 10.12 this information is conveyed. Despite of all the caveats mentioned in the paragraphs above, this graph clearly shows the effect that each outburst has on the respective success rate. For the 5-2 outbursts, despite the time zone where it happens not being shown in the graph, its null effect can be seen on the blue line, as the percentage of cases that are above the threshold at the beginning of this time period is nearly 100%. With the outburst being finished, this trend continues up until approximately 30 seconds to the C/A, when the success rate indeed rapidly drops, reaching a lowest value of about 65%. For the 2-1 outbursts, the impact is felt much sooner and during the 2-1 period, as expected, where the success rate more gradually drops and continues to do so the outburst ends, reaching a minimum of nearly 60%. Finally, the 1-0 cases show a similar behaviour, where up until the outburst begins, the success rate remains untouched, but right after the 1 minute mark, the curve sharply drops, becoming, by a significant margin, the worst of the three cases, reaching below 50% success at the closest-approach point. This is, as such, the expected behaviour, in the sense that if an outburst does indeed affect the success of the scenarios, its effect is seen as the outburst starts rather than later around the 30 seconds mark, and results in significantly lower performances in terms of the fraction of scenarios which are successful throughout the whole duration.

10.3.6 Outlier rejection threshold analysis

The next variable worth analysing now is the measurement rejection threshold. In this analysis, only $\infty\sigma$ and 3σ are considered, where the former represents simulations where no rejection is employed. Following a similar tactic as before, the success rates can be plotted for each threshold value. Doing so yields quite unremarkable results, with the two options showing similar curves, and thus this result is not shown. What is represented here is, on the other hand, the impact of this variable but with separate curves for the 2-1 outburst, for the 1-0 outburst, and for the baseline scenario with no outburst.

In Figure 10.13 the differences between measurement rejection thresholds between different outburst periods becomes clear. First of all, the two sets of curves for the cases with outbursts follow the expected success rate behaviour given the respective outburst period, as discussed above. Secondly, the two outbursts show almost opposite behaviours. On the one hand, the 2-1 case, not using any measurement rejection proves to have a higher success rate for the majority of the time, and, most importantly, a higher minimum at the closest approach. On the other hand, the 1-0 outbursts show that rejecting measurements with a 3σ threshold yields significantly better results for most of the closest-approach duration. And, critically, these two curves hit quite different

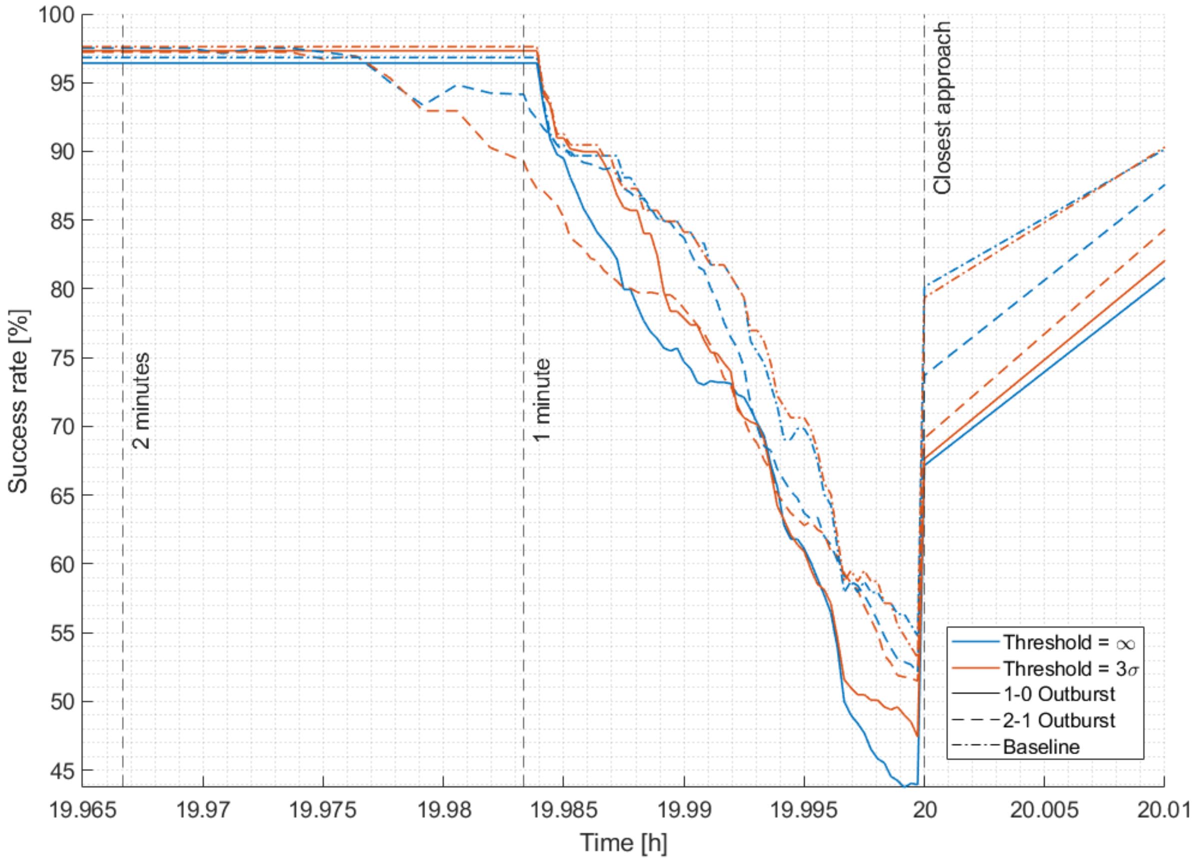

Figure 10.13: Percentage of cases, for outbursts 2-1, 1-0, and for the baseline scenario, with pointing error below the 0.5 degree threshold, as a function of the measurement rejection threshold utilised.

minima, with a nearly 5% improvement when the rejection is employed. As such, it can be concluded that for 2-1 outbursts the usage of this rejection method seems to be unnecessary, and perhaps slightly detrimental, whereas for the 1-0 ones it is advisable.

In addition to the two previous sets of curves, the baseline results are also included. Similarly to the behaviour under a 1-0 outburst, the success rate only starts dropping after the 1 minute mark is passed, and for this case, choosing either $\infty\sigma$ or 3σ shows little difference in the results, in alignment with what had been concluded by the baseline analysis. What is surprising, nonetheless, is that the baseline results actually do not show a much higher success rate than the set of results including 2-1 outbursts. Throughout most of the time period being analysed, the two baseline curves stay ahead of the 2-1 curves in terms of success rate, but right at the end, around the 20 hour mark, the two results do not differ much, further emphasising that the later outbursts are the ones with the greatest impact.

10.3.7 Filter analysis

An additional analysis that complements the previous one is that of the filters. As before, the temporal success rate of the scenarios is analysed, with the two distinct measurement threshold values and two different outbursts being plotted separately, to assess the impact of these variables on each of the filter's performances.

The two graphs in Figure 10.14 show the success rate of each of the filters tested across the scenarios generated with seeds 0 through 20. However, they separate additional variables. Whereas for the top graph the outbursts considered are the 2-1 kind, for the bottom graph they are 1-0. And each graph has coloured full lines, representing how each filter behaves without the measurement rejection, as well as coloured dashed lines, with the sames colours corresponding the same filter but with a 3σ measurement rejection threshold. This way, the hope is to evaluate not only which filters work the best, but how these two added variables influence that performance. Also on each graph are small boxes which zoom in on the zone of the curves in the seconds right before the closest-approach point, where most filters reach their minima, so that their performances in that area can more easily be compared against one another.

Starting with the 2-1 outbursts on the top graph, its results seem reminiscent of those of the previous section where the same analysis was made regarding the choice of the best filter for the baseline scenario, where multiple filters are, at different times, the ones with the best success rate. In this case, it appears that the majority of the time is spent with the EKF, UDUEKF and the batch UKF filters ahead of the others. There are, nonetheless, two more aspects to this comparison. Firstly, the minimum of each filter is perhaps the most important information conveyed by each plotted line, as it represents an expectation of the percentage of cases with the given conditions that would always stay below the pointing error threshold. Thus, a filter with a higher value in its minimum might be better than another which has the best success rate most of the time, but then drops right at the closest approach. In this regard, the aforementioned EKF, UDUEKF and batch UKF are the favourites by a landslide, with a nearly 3-4Δ% gap to the next best.

Another aspect to emphasise is the measurement rejection threshold effect. By comparing the full and dashed lines of each colour, the effect of the introduction of this mechanism can be evaluated. For the majority of filters being analysed here, the change from $\infty\sigma$ to 3σ causes the results to become worse, as the respective dashed lines are most of the time below their full line counterparts. It is, thus, safe to say that for these filters and under the 2-1 outburst it is detrimental to reject measurements at 3σ. However, whereas the others see negative or little change to the results, the batch UKF stands out as the rejection of measurements majorly improves its performance. At some points in time, the difference between the average success rates reaches 15Δ%, which is a massive improvement. Most crucially, right at the point of the closest approach where most filters are near 50-52% success rate, the batch UKF with measurement rejection sees 56% of its cases stay below the pointing threshold, showing similar performances to that of the EKF and UDUEKF. As such, for the 2-1 outburst data, the filters that show the best performance are either the EKF or UDUEKF, without any outlier rejection, or the batch UKF, with a 3σ threshold for this technique.

Shifting the focus now to the bottom graph of Figure 10.14, the overall picture changes quite significantly. Here, again, the behaviour of the curves is expectable given the nature of the outburst starting at 1 minute to C/A, which is clearly where the performance curves start dropping. Additionally, the lowest values observed for this case are below those of the 2-1 case. This again emphasises that later outbursts, even ones with lower intensity, do cause greater damage to the navigation performance, with the minima now ranging from 40% to 53%, a nearly 10Δ% decrease with respect to the top graph. Comparing now the filter results against each other, in this case there is a noticeable gap between the batch filters and the sequential filters for a majority of the analysis duration, with the batch formulations significantly out-performing their batch counterparts during the majority of the period being discussed. This is only true, however, for the case of the full lines with no measurement rejection. In the case of the dashed lines, on the other hand, no such clear trend can be seen.

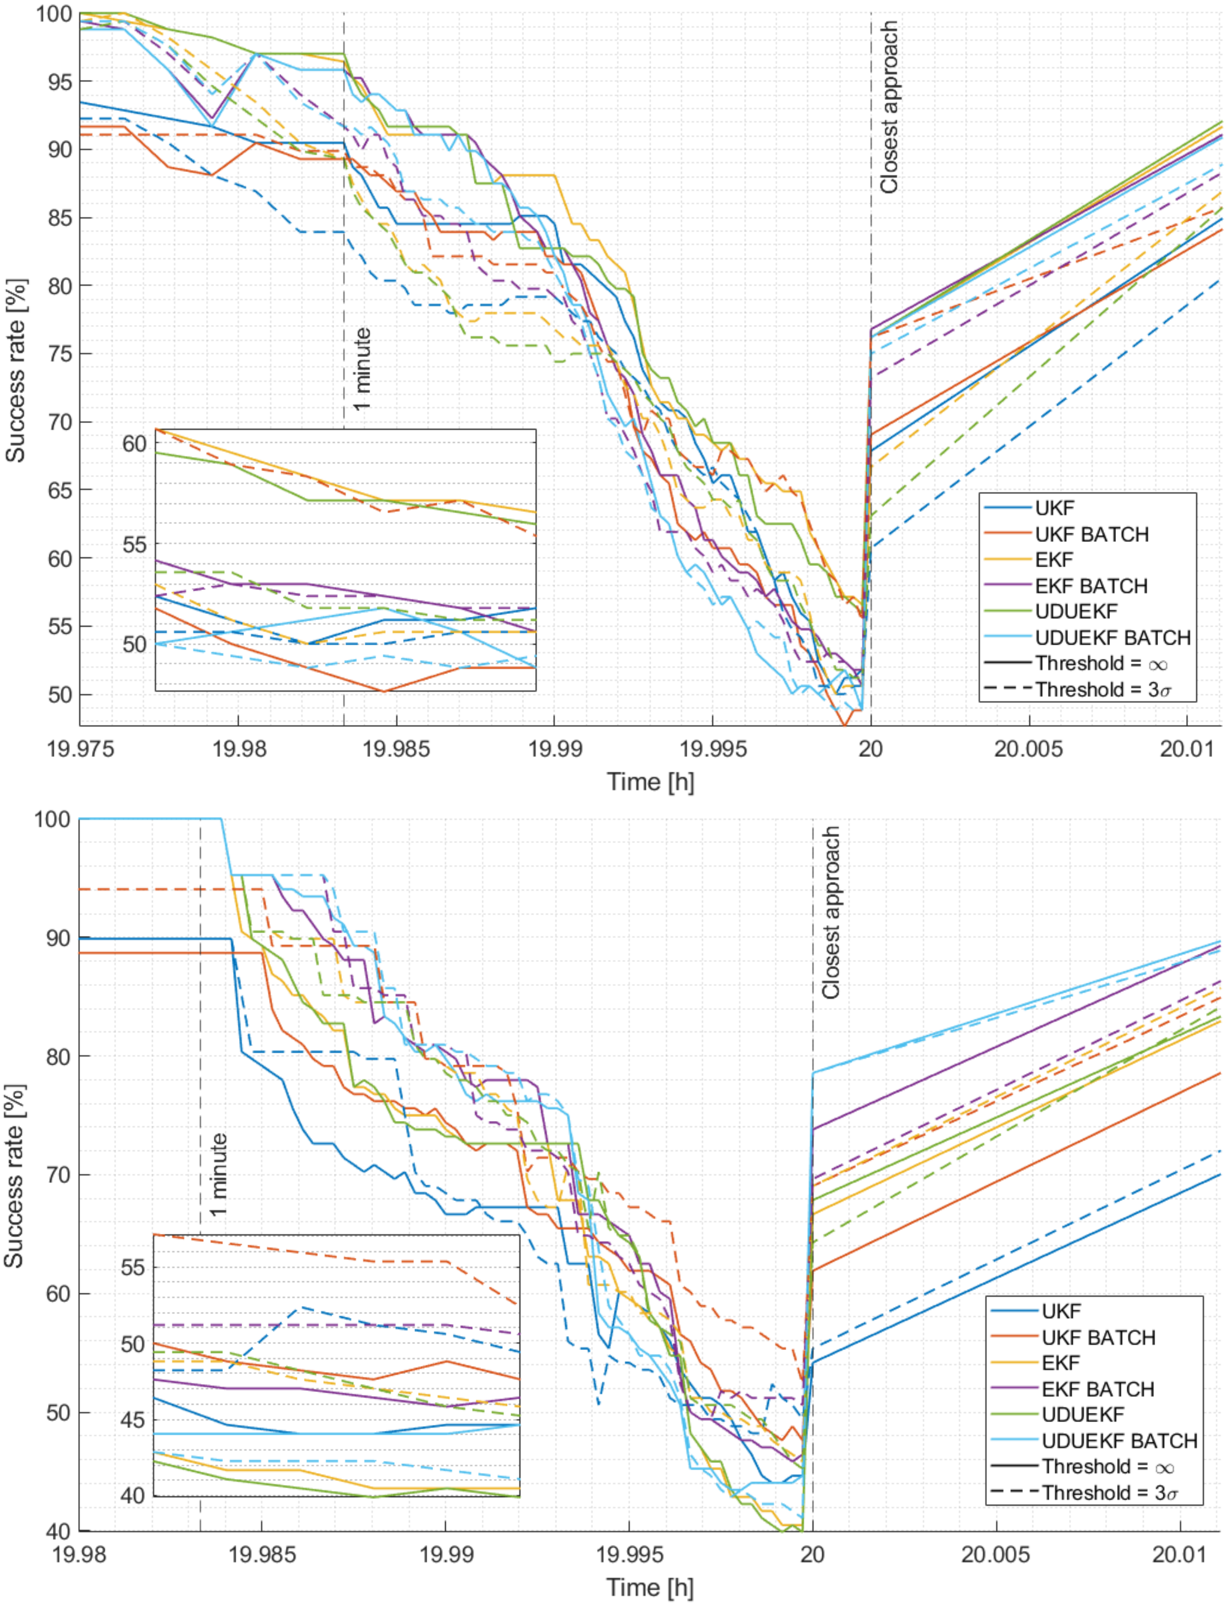

Figure 10.14: Percentage of cases, for outbursts 2-1 (at the top) and 1-0 (on the bottom) and varying measurement rejection threshold, with pointing error below the 0.5 degree threshold as a function of the filters used.

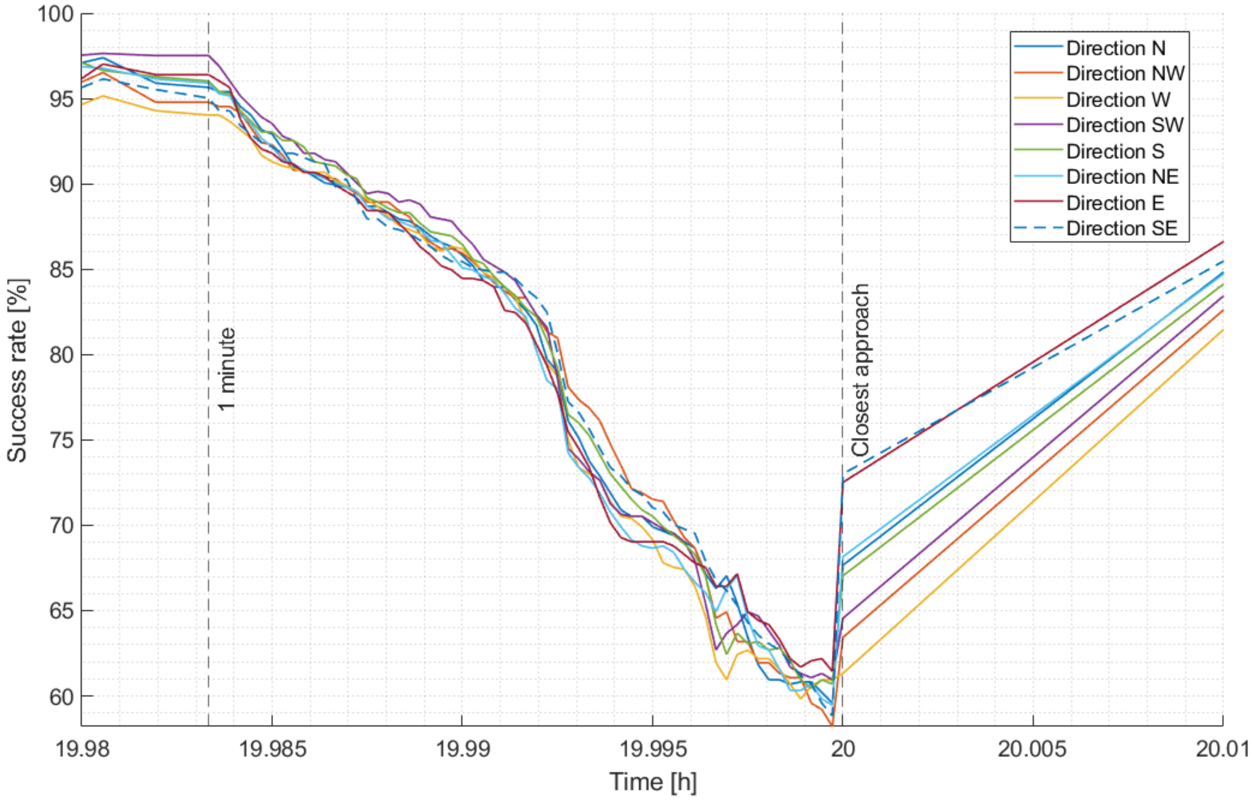

Figure 10.15: Percentage of cases for each outburst direction with pointing error below the 0.5 degree threshold over time.

Comparing now the case with $\infty\sigma$ with the 3σ, the results are, for most of the analysed period, mixed. The UKF curves wildly oscillate, with one being vastly better than the other at certain periods, only to then quickly reverse that trend. The batch EKF and batch UDUEKF curves are always close to one another, again without any case clearly outperforming the other. For the EKF, UDUEKF and batch UKF, however, the dashed curves are noticeably and almost constantly above their full curves, showing that these benefit the most from the 3σ outlier rejection.

Yet, there is a period when the results are not mixed, which coincidentally turns out to be the most crucial – the closest approach. Right at the 20 hour mark, where most of the curves hit their minima, the graph shows that for all filters, except the batch UDUEKF, the introduction of measurement rejection improves the results. Not only that, but this improvement is nothing but significant, with increases in success rate of about 5Δ% per each filter. This remarkable result leads to the conclusion that, for outbursts of the 1-0 kind, this outlier rejection technique can make quite a difference. Finally, perhaps the most notable result to be drawn from this graph is that, again, the batch UKF is among the best performing filters for these scenarios, and again with the measurement rejection turned on. Close second and third are now, however, the batch EKF and UKF, respectively, and both also with 3σ outlier rejection.

10.3.8 Direction analysis

There is, nonetheless, a variable which has yet to be scrutinised. The direction in which the outbursts are oriented is the last possible factor impacting the navigation results. To analyse its impact, the same approach is pursued here as has been done throughout the second part of this section. With all the gathered data put together, each direction's performance is plotted against the time for the last moments before C/A, resulting in Figure 10.15.

Contrary to the case of the other variables, there seems to be little conclusions to draw from this graph. For the majority of the time period being studied, most directions stay within a few percent points of one another, differing by around 3Δ% at their minima nearest to the closest approach. There are a few factors that could theoretically change these results. Firstly, the E-W directions are aligned with the Sun, which is the direction along which the measurement biases and white noises are greatest, statistically. For that reason, one could expect these zones to show some of the worst results, as outbursts aligning in these directions could add to these uncertainties. However, on the other side of the coin, it is also possible that these two phenomena happen

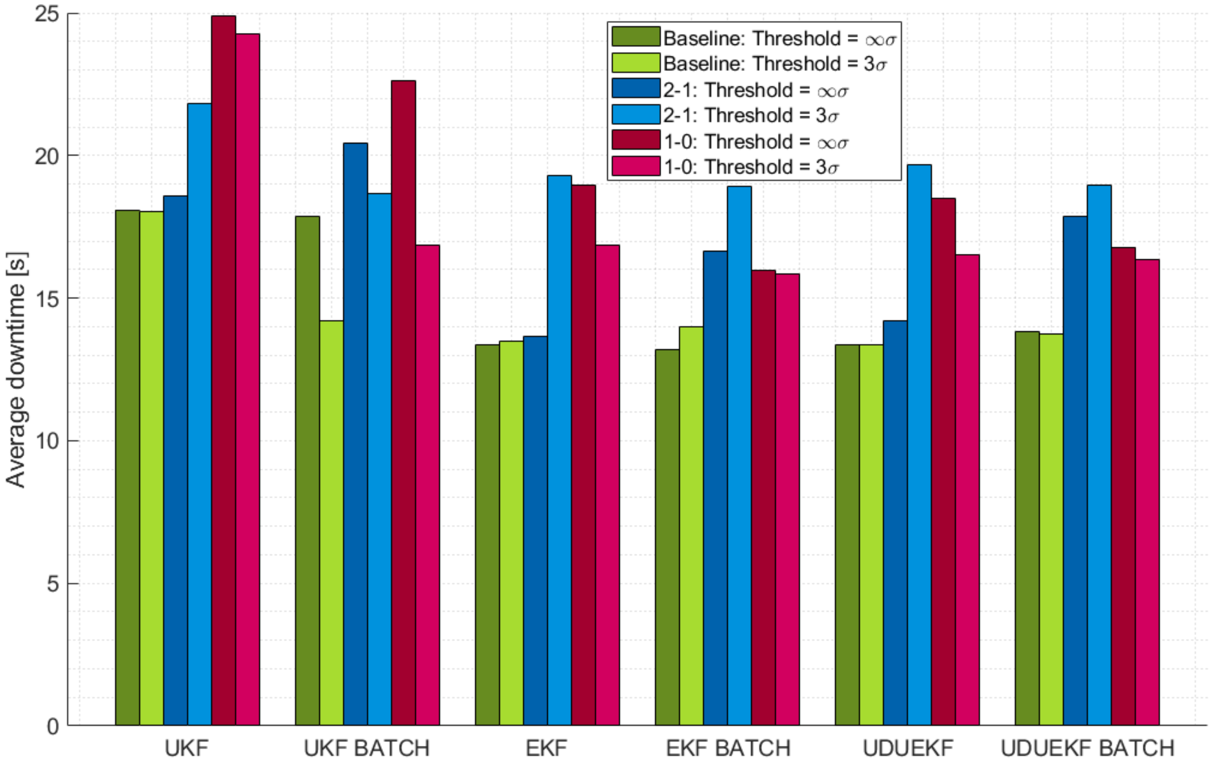

Figure 10.16: Average downtime during the last minute and a half before closest approach, considering a pointing error threshold of 0.5 degrees for each filter, with varying outburst period and outlier rejection threshold.

to cancel out, resulting in the average not seeing significant differences, which is perhaps why these directions seem unremarkable in the graph.

Additionally, there is the fact that, as the outbursts happen, the spacecraft is quickly flying by the comet's nucleus, meaning that their orientation can change quite rapidly. As the spacecraft flies in the plane of the equator, the polar outbursts should not experience this phenomenon, meaning that their impact is constant in intensity and orientation. Those on the equator plane, and the remaining 4, which are not one or the other alone, are seen by the spacecraft as changing orientation over time. Those on the sunlit side progressively rotate and orient themselves facing towards the vehicle, whereas those on the opposing side orient facing away from the vehicle and are hidden by the nucleus itself. This would seem to suggest that those oriented in the N-S directions and thus not suffering from this shift, would have the most systematic, greatest impact, especially in the 1-0 outbursts, which include the moment when the spacecraft has the target right by its side. Despite this rationale, even separating the data into 2-1 and 1-0 plots does not emphasise this theoretical effect on the N-S direction causing the respective results to present a lower success rate.

It seems, as a result, that there is no connection to be derived between the outburst directions and the navigation performances, at least from the limited set of data used for this analysis.

10.3.9 Downtime analysis

So far, the focus has been put on the success rate of the different scenarios and navigation solutions with time. Instead of that, let the downtime now be the figure of merit. This corresponds to the time that each simulated case spends with a pointing error above the established 0.5 degree threshold

The bar graph in Figure 10.16 shows the average downtime during the last 1.5 minutes before the 20 hour mark for each of the performed simulations. The results for each filter are separated, and each filter has baseline results, as well as results for the 2-1 and 1-0 outbursts. Additionally, each of these three categories is divided into $\infty\sigma$ and 3σ in terms of the outlier rejection threshold used. And it is indeed a great way of summarising many of the results obtained in the past few graphs in a simpler and visually more appealing way. This comes, however, at the cost of some loss of information due to the usage of the downtime average.

In this graph, it is plainly visible, for example, the impact that the failed cases for the UKF and batch UKF have on this metric, as the downtimes for these two filters are clearly higher than the rest in most categories. Additionally, the effect of using $\infty\sigma$ or 3σ as the outlier rejection threshold can be seen, changing little for the baseline case, except for the batch UKF, being mostly detrimental with the 2-1 outbursts, again except for the batch UKF, and improving all the cases when 1-0 outbursts are visible.

Nevertheless, it can also convey some misleading information, especially regarding the batch UKF filter. As was concluded from Figure 10.14, the batch UKF filter, with a 3σ threshold, is, for both outbursts being analysed, one of if not *the* best option available, at least in terms of the minimum success rate observed. Despite this, the average downtime metric over-punishes this filter, as well as its sequential counterpart, because the cases which diverge add a full 1.5 minutes of downtime to the total sum. As a result, the bars of the batch UKF do not really show that this filter is among the best, but rather implies that it is among the worst.

10.4 The unscented failures

Throughout the analysis performed in this chapter, the failure cases of the unscented filters, both batch and sequential, appear as quite remarkable. For certain instances, these filters reach a point, right before the closest approach, when their covariance matrices, as aforementioned, cease being positive semi-definite, meaning that the generation of sigma points in the next iteration fails, as the Cholesky cannot be performed. These failure cases, although relatively rare, do affect the performance and robustness of the respective unscented filters. What is even more remarkable is that, despite this negative effect on the results, the batch UKF is still one of the best performing filters among the tested, especially when rejecting measurements with a threshold of 3σ.

However, these failures do leave the question of what makes the filters behave this way. This is an important issue, because, were it to be understood, it could become possible to correct it, thus making the batch UKF the obvious choice, at least in terms of downtime and success rate performance, but still not so much in computational time terms. In an attempt to answer that, this section is dedicated to analysing the characteristics of these few scenarios that fail.

Starting first with the baseline analysis, 51 total seeds are run each with 3 different outlier rejection thresholds. For the case with no rejection, the UKF fails 4/51 cases, whereas the batch UKF does so for only 3/51. Interestingly, across all the simulations performed, seeds 15 and 16 seem to be the most problematic for these two filters, as both fail in this case, whereas for the other cases, the failures are more scattered, with the UKF failing in two instances that the batch UKF handles well, and the batch UKF failing in yet another instance that does not make the UKF fail.

In terms of what happens when 3σ or 4σ thresholds are used to reject measurements, the results do improve, with 4σ removing one of the failed instances of the UKF, and 3σ being by far the most advantageous for these filters, reducing the number of failed cases of each to 2. This is, nonetheless, what has slowly been arising from the observed data, be it for the baseline or for the outburst-perturbed case, where the inclusion of 3σ outlier rejection significantly improves the performance of these two filters.

Analysing these failure cases in terms of the most impactful parameters in the simulation's performance, the measurement cut-off time and the closest-approach distance, can perhaps tell more about what is causing these instances to diverge. In terms of the former variable the results are pretty mixed, with two cases experiencing loss of sight of the target before the 1 minute to 20 hours mark, one of them at about 50 seconds, one at 30 seconds, and finally one just 5 seconds before. As such, it cannot be said that all the cases suffer from a very early loss of target from the NAVCAM's FoV, but most of these cases are in the zone where some downtime is observed, so, except for that last case, this can be an influencing factor. Regarding now the closest-approach distance, the results range from approximately 670 km to 1200 km, which when compared to the array of values represented in Figure 9.8, lies on the left side of the graph, meaning that all the cases are located in the zone that experiences possible downtime values above 0, and not nearer to 2000 km, where the opposite happens.

Another piece of information, that comes from the outburst analysis, is the fact that when outbursts are present, the number of failing cases for these filters varies quite a bit, with some seeds that previously did not fail now failing for certain outbursts, and vice-versa, meaning that trying to predict when the filters cannot handle the scenario with these perturbations present becomes even harder. On the other hand, however, this again leads to the conclusion that the measurements play an important role in determining the failure or success of the unscented filters.

Regarding the failure rates of these filters, in the baseline analysis each is run with 51 seeds and 3 different outlier rejection thresholds, for a total of 153 simulations. Out of these, the batch UKF fails to conclude 8 and the UKF 9, resulting in rates of 5.2% and 5.9% respectively. During the scenario perturbed by outbursts, the number of simulations run is quite higher. Given all variables, each has 1072 simulations to its name, and each fails to conclude 63 of them, for a failure rate of 5.9%. From this, it can be concluded that, under the conditions used in this study, the presence or absence of outbursts does not meaningfully affect the success rate of the unscented filters, and that about 5-6% of the cases end up failing regardless.

It is known that these filters fail right around the time of the closest approach, which makes sense given that it is during this time that the state of the system varies the fastest, and thus the navigation is the most stressed. Additionally, the failing cases lie on the tougher side of scenarios, with earlier measurement cut-offs and smaller closest-approach distances. Finally, by rejecting some of the measurements with the highest residuals, the results seem to improve significantly by removing some of the failures.

The logic question is, then, since changing the closest-approach distance is out of the control of this thesis, as will be discussed ahead in Section 11.3, if there is a better outlier rejection threshold value to be used such that the number of cases not reaching the end of the simulation can be decreased further; and whether adjusting the measurement cut-off time can also improve the robustness of these filters. To answer the first question, the 4 simulations that fail with the UKF and batch UKF even when using a 3σ threshold were rerun for different thresholds, varying from 2σ to 4σ in 0.1σ increments. Out of the four cases, only one – seed 15 of the UKF – reached the end of the simulation, with thresholds 2σ through 2.4σ. Whereas this experiment indicates that tightening the threshold for outlier rejection can improve these filters' robustness, it is not enough to ensure that they always work as intended.

To test the second hypothesis, a change to the scenario's time step is made. Whereas, up until this point, at 19 hours and 57 minutes the time step would lower from 60 seconds to just 5 seconds, now the change is more drastic and goes right to 1 second at this mark. This means that, for the cases which suffer due to the navigation being impacted by very early measurement cut-off times, the frequency of measurements being taken is increased to 1 each second 2 minutes earlier, which should help mitigate the impact this factor has on the results. As before, the seven sequential and batch UKF cases that failed in the baseline analysis were simulated again using this new time-step definition. Surprisingly, all seven were able to finish the simulation, meaning that adjusting this parameter is able to greatly solve the unscented filters' robustness problem, at least in the baseline scenario.

However, these filters also show difficulties in finishing the simulation when outbursts are present. And repeating the same test for the cases with outbursts that show failures does not yield such positive results. In fact, none of the tests were able to finish. This is an indication that, whereas for the baseline scenario the change in time-step rationale was able to mitigate, if not solve, the unscented filters' robustness problems, the same cannot be said for the perturbed outburst-including scenarios. In these cases, this change might help, but it is not enough to correct the failed cases. This, then, means that the outbursts add another layer of perturbation to the navigation, affecting the robustness of these filters even further.

These results indicate that, indeed, some measures can be considered to try to mitigate the issue of unscented filters failing to finish the simulations. They perform better when outlier rejection is employed, especially around the 3σ threshold. Additionally, by increasing the measurement frequency before the measurement cut-off time, all the observed baseline cases were solved, despite not making a difference for the outburst scenarios. Even so, this leads to the conclusion that increasing the state estimation and measurement frequency even further is the most promising strategy to try and solve the robustness issue of these filters.

10.5 Analysis summary

In this chapter, many complex analyses were developed, and from the discussion of the results, there are multiple conclusions that were drawn. As such, this section serves as a short summary of the main lessons learned.

Baseline:

- The UKF and batch UKF fail to finish their simulations in certain cases, making them the least robust of the bunch.
- The EKF is the filter that is computationally the lightest.
- Using 3σ innovation filtering in the baseline scenario only improves the UKF's performance significantly, not altering or worsening the results of the other filters. The same can be said for 4σ, but in this case the improvement to the UKF's results are negligible.
- In general, the EKF is the best performing filter of the six, with lower average downtimes, according to Figure 10.4, and the fastest run-times. Nevertheless, the performance difference to the other filters is not significant, meaning that any of the extended Kalman filters could be used with little downside.

Outbursts:

- The 5-2 outbursts show very little change in results with respect to the baseline case, regardless of intensity or direction. This indicates that outbursts before the 3 or 2 minute mark do not meaningfully impact the outcome of the mission.
- The closer to the comet the spacecraft is when an outburst happens, the higher its impact on the navigation.
- The intensity of the outbursts is also reflected on the performance of the filters, with higher intensities increasing the observed downtimes, but to a lesser extent than the outburst time.
- The usage of outlier rejection with a threshold of 3σ shows positive results when applied to the 1-0 outbursts, but detrimental ones in the 2-1 case.
- Despite the robustness shortcomings, the batch UKF is the best performing filter for both the 2-1 and 1-0 outbursts, specifically when a 3σ outlier rejection is put into place, as it presents the highest success rate minima.
- The outburst directions seem to not affect the navigation results significantly.
- The presence of outbursts makes the unscented filters' robustness issues harder to solve than for the baseline case.

Chapter 11

Conclusions and recommendations

In this final chapter, the goal is to outline the conclusions drawn from the analyses performed in the previous chapters. Then, changes to be made to the mission's requirements in order to either allow for different results, or to relax the established limitations, are suggested. In addition to this, a section will be dedicated to enumerating the aspects related to the mission's description and to the developed analysis that should be considered in future works relating to the Comet Interceptor. Finally, the relevance of this work in the context of this mission's development, as well as that of other missions, is discussed.

11.1 Research questions and study objectives

Posed in Section 1.3 was the research question giving motivation to the whole analysis performed during this thesis project. Along with it, as well, were various study objectives also stated as questions, each corresponding to a chapter or section of the document. As such, it stands to reason that, as a means of concluding the performed work, these questions be answered objectively with the gathered results. Following the order chosen for this study, firstly the study objectives will be addressed one by one, ending with the research question, which should be addressed with all the information gathered along this analysis in mind.

SO – 1. Have previous cometary exploration missions employed similar navigation solutions?

Starting with **SO - 1**, after the discussion presented in Chapter 2, it became clear that the sort of very fast fly-by strategy employed by the Comet Interceptor has been used in multiple missions to smaller astronomical bodies in the past, including comets. And, among these, some have indeed relied on the acquisition of photographs taken from on-board cameras as a way of navigating in the vicinity of the comet. This is indicative that utilising autonomous navigation for this type of missions is an adequate strategy, and also one that has been analysed and implemented before.

SO – 2. What are the navigation-related requirements for the Comet Interceptor mission?

Then, in a later section of Chapter 2, the requirements available in the current stage of the mission's development were laid out. These were picked from multiple documents, from GMV, establishing the strategy for the Comet Interceptor case being studied and its respective limitations, especially in terms of the main mission's geometry and the navigation constraints, being thus divided into mission and system requirements.

These requirements are enumerated and their rationale briefly discussed in Section 2.3 and Section 2.4. In general, however, the navigation requirements for this mission establish that a camera be used to snap images of the target, using that information to update the vehicle's state estimation over time, so that the payload can be correctly pointed at the comet's nucleus. Additionally, all this process is to be done on-board of the spacecraft and be completely autonomous from communication and commands from the mission's ground segment.

SO – 3. How can the cometary environment affect the navigation?

By studying the comet's environment in Chapter 3, knowledge as to how the elements of the target's surrounding affect the mission was gained. In terms of the dynamics, the presence of dust, especially in the closest vicinity of the comet, has proved before to be a relevant force acting on the spacecraft. Luckily, the lessons learned from ESA's Giotto have allowed the Comet Interceptor to have a well-established model for the distribution of dust particles around the nucleus, which can then be used to derive a force model. Additionally, by analysing the typical sizes of comets and using known density values for their nuclei, it was concluded that, in general, their forces of gravity can be ignored as their impact is quite minuscule, especially for comet sizes around 10 km in diameter, which is the size being considered. Despite this, both these forces' descriptions and analyses are performed further ahead in Chapter 4.

Measurement-wise, the comet's coma can be a source of background light capable of affecting the images captured by the NAVCAM. These errors are considered to be dealt with by the image processing algorithm, nonetheless, so they are not studied here. The major impacts caused by the comet's environment are the outbursts originating from the target's nucleus. Described in some detail in this chapter, these light-reflecting and scattering jets of material are modelled as simple biases with a certain orientation, strength and duration, which thus affect the perceived position of the nucleus.

SO – 4. What forces are required to describe the spacecraft's dynamics adequately?

Chapter 4 discussed the methods used to analyse each force acting on the spacecraft during the Comet Interceptor's encounter to determine if these were worth considering in the simulations to come. The forces of gravity of the Sun, the target itself, as well as of all the planets in the Solar System were studied, and the conclusion was that the Sun's is the only major force among them capable of changing the trajectory meaningfully, especially during such a short 40 hour period – later shortened even further to 21 hours. The Earth can be an exception to this conclusion, in case the comet's fly-by happens in its close vicinity – due to **MIS – 1.5** – but that case would have to be analysed separately, or otherwise a target comet with a different trajectory must be chosen.

The forces caused by solar radiation pressure and the impacts of cometary dust particles are also discussed in this chapter, concluding that, despite being much less impactful than the Sun's attraction, they should still be considered relevant perturbations. As such, these are the three main forces determining the spacecraft's trajectory.

SO – 5. Which are the best *a priori* filtering candidates for solving the task at hand?

Because there are so many different estimation algorithms, and each algorithm can have multiple formulations, be they sequential or batch, for example, this question was one of the most difficult to address. Additionally, the chosen filters must take into account the characteristics of the mission being analysed, and should ideally be well-known algorithms with some heritage in the space industry.

The goal in choosing the six filters for this study was to compare different formulations while simultaneously assessing whether they work better in sequential of batch form. Attempting to answer this chapter's study objective, it was hard to predict whether the unscented or extended Kalman filters would perform better. However, the batch filters should have the advantage of being less affected by noises in the measurements, as their impacts are naturally mitigated by processing a batch simultaneously. On the other hand, the sequential filters produce an updated state estimation each time step, which, given the very demanding navigation effort that the spacecraft undertakes right as it passes by the comet, should be a big advantage over the batch formulations. These wait a few time steps before using the measurements, thus leaving the FCS for longer without newer information, which is why the sequential filters were predicted to have the upper hand.

SO – 6. What numerical algorithms will be needed to properly simulate the mission?

After defining the filters to be used, it was then time to analyse the software framework to determine what tools were required to perform the simulations, and which were already available in FCS-ATOMIC. As such, the most critical numerical methods needed to support the simulations are described in Chapter 6, and FCS-ATOMIC, GMV's software framework that was used for this thesis, was briefly analysed in Chapter 7. The main focus there was put on the filters themselves, as well as on the algorithms computing the Jacobians of the dynamics and measurement simulation functions, since these were the major tools that needed to be developed for the study.

SO – 7. How can the developed code be tested to perform V&V within the software framework?

For the case of the blocks of code discussed in the paragraph above, there were no available sets of validation data to be used to compare with the results of the created functions. And even if this were the case, since these functions had to be integrated into the FCS-ATOMIC framework, ideally the data set would have be include the integration portions of the code, pretty much meaning that the data would have to have originated from previous simulations within the framework.

For the case of the Jacobian computing functions, then, the strategy used was to compare the determined derivatives with approximated values obtained via numeric methods. As for the filter functions, a combination of data sets for similar filters, comparison between the sequential form and their batch counterparts working with a set batch size of 1, and analysis of the altered blocks of code distinguishing them from other filters were some of the methods used to assess whether these work as intended. The final confirmation of the implementation of these functions was that, when they are all put together and integrated into FCS-ATOMIC, the results are realistic and comparable to those of the already implemented UKF. Additionally, throughout the performed analyses, each filter was, for certain scenarios or at certain points of the simulation, the best performer of the bunch, indicating that there is no obvious implementation flaw with any of them.

SO – 8. How is the Comet Interceptor's scenario defined?

Chapter 8 is dedicated to exhaustively discussing all the parameters required to describe the scenario. It is mainly divided into three parts, with the first being focused on the main aspects of the simulation, such as the total time, time step size, and the geometric definition of the orbit. Then, two sections were dedicated to the parameters that influence the dynamics and the measurements, respectively, as well as to the observed variability of these parameters. With all these values summarised in Table 8.2, the initial state of the spacecraft was defined, as well as all the uncertainties in the parameters that characterise the environment, the vehicle, and the NAVCAM system obtaining the measurements.

SO – 9. How much does each scenario uncertainty affect the simulation results?

Similarly with the organisation of the scenario defining parameters laid out in the paragraph above, so here can the factors be divided into orbital geometry, dynamical and measurement-related. With a simple analysis performed at the beginning of Chapter 9, the conclusion that the dynamical uncertainties of the system changed the results very little was reached, thus being decided that these were to be left out of the following analysis.

By analysing the remaining variables, the result is that, despite measurably affecting the simulation results, which can be seen by removing their influence and rerunning the same cases, the measurement-related biases, especially the image processing bias, are not good predictors of the simulation's outcome in terms of performance. Instead, looking at the measurement cut-off time and at the closest-approach distance will give an idea of how bad the pointing error results can be for a certain scenario. And both of these are majorly determined by the initial state dispersion of the spacecraft with respect to the nominal trajectory, as, nominally, the spacecraft has a 1000 km closest approach to the comet, and the measurement cut-off happens a few seconds before the 20 hour mark. Results loosing sight of the target sooner and/or straying too close to it see significantly worse results.

SO – 10. How do the proposed navigation solutions handle the baseline and disrupted scenarios?

In Chapter 10, the established Comet Interceptor scenario was analysed with the six proposed filters to determine whether the navigation solution they provide is adequate for the mission. The chapter is divided into two parts. Firstly, the baseline scenario, fully described in Chapter 8, was studied, and, after that, a modified, or perturbed scenario was introduced with outbursts, and their effects discussed.

Starting with the baseline analysis, all the filters were analysed in terms of their maximum pointing error, computational load, success rate over time, and average downtime, for multiple seeds and with varying outlier rejection thresholds. The main conclusion from all these was that the EKF is, in general, the best suited filter both in terms of performance and computational efficiency. All the other extended Kalman filters were either operationally heavier and/or showed slightly worse results, although it can be said that any of them could be picked without a significant downside.

The same cannot be said for the unscented filters, as these showed a tendency to fail in certain situations, with the simulations not finishing due to numerical errors. This hurt both their performance and robustness when compared to the other four filters. Additionally, the usage of 3σ or 4σ thresholds did slightly improve their results, which is not observed for the extended filters.

Including, then, the outbursts in the perturbed scenario scenario, a similar analysis was performed again. This time, however, there were more variables in play, with the addition of the outburst period, intensity, and direction. Considering these three variables first, the conclusions were that the later an outburst happens and the more intense it is, the larger the impact it has on the navigation for all filters, worsening the results. The direction, on the other hand, seemed to have little to no effect on the navigation results.

As to which filter is best for handling this perturbed navigation, Figure 10.14 sums the results the best, and shows that the batch UKF with an outlier rejection threshold of 3σ has the overall lowest downtimes. These results even take into consideration the cases where either the UKF or batch UKF diverge, meaning that, when the latter is able to finish the simulation without failure, its results are significantly better than the competition.

Answering, then, the question posed in **SO – 10**, all the proposed filters can handle the baseline and disrupted scenarios fairly well, with a notable period in time right before the closest approach during which all candidates perform their worst, which is expected, since it is the most demanding region of the encounter in terms of pointing performance. Despite this, most filters show downtimes in the order of magnitude of around 10-20 seconds, with that metric worsening in the presence of outbursts. There are the exceptions of the UKF and batch UKF, which occasionally diverge during the closest approach, but even so their performances are reasonable. However, because of that, it would be deeply inadvisable to ever use just one of these filters to perform navigation in this mission. If these are to be included, a backup filter should be considered, or the strategies to improve their robustness that were discussed in Section 10.4 implemented.

> **RSQ – 1.** How do outbursts present in a cometary environment during the encounter phase affect the optical translation navigation solution?

To conclude this study, it can be said that the presence of outbursts on top of the baseline Comet Interceptor scenario can have an unpredictable effect on the results. If they happen when the spacecraft is closer to the target and have a high intensity, the observed behaviour is that their effects make the performance deviate further from the baseline, be it in a positive or negative way. Luckily, for the non-unscented filters, their presence did not cause the navigation to diverge in any simulated circumstance, meaning that their robustness was not affected. Even though this is also the case for the unscented filters, since the failure rate stays about the same, their robustness was already much worse to begin with. Additionally, the failure-mitigating strategies proposed do not seem to work as well under the influence of these events.

Despite this, the batch UKF seems to be the filter that can handle their presence in the measurements the best, especially when innovation filtering is used with a 3σ threshold. Because, however, it shows a tendency to fail under certain situations, it should always have a backup filter – an EKF, for its solid baseline performance and lowest computational power – to pick up the navigation effort in case the of a divergence. Nevertheless, since these events are so rare, the likelihood of one happening right as the spacecraft is closest to the comet with an intensity large enough to have a meaningful impact on the navigation is quite minimal.

11.2 Requirements update and recommendations

With the analyses performed during the previous chapters, the gathered data now provides a picture of how the Comet Interceptor system performs under the simulated mission scenario. With that, it is possible to discuss whether the posed requirements in Chapter 2 are viable or not, and if what they impose on the mission is realistic given all the other constraints being considered. To address this issue, firstly the requirements which are not met are addressed, which are

> **MIS – 1.4.** Spacecraft A shall be compatible with a B-plane delivery error (see Section 4.1) as high as 150 km (1 sigma).

and

> **MIS – 1.10.** The spacecraft A shall allow the CoCa instrument to continuously observe the target before, during, and after the closest approach.

Regarding **MIS – 1.4**, Figure 9.8 shows that there are scenarios with closest-approach distances way above 2000 km, which exceeds many times the standard deviation set for the B-plane delivery error of 150 km. As such, this requirement is clearly not met. The problem is adjusting the scenario or the mission to solve this issue. Since the closest-approach distance depends on the initial state of the spacecraft and on the scenario's dynamics, which are both aspects that cannot be changed, at least within the scope of this study, then there is very little to do change these results.

The dynamics are modelled fairly closely to what the respective forces are in reality, and even if something about them changed, their effects on the trajectory for all forces except the Sun's gravity are already minor. The other option would be to attempt to decrease the initial dispersion of the true trajectory with respect to the nominal trajectory. Nevertheless, this dispersion is dictated by the navigation and actuation accuracies of the previous phase of the mission, the approach, during which the GNC is performed by the ground segment. These parameters have not been studied in this thesis, and thus commenting on them or suggesting any alterations regarding them, in an attempt to align the results with **MIS – 1.4**, would be ludicrous. Thus, if nothing else changes in this regard, this requirement will have to be relaxed to the observed distribution of closest-approach distances, or discarded completely.

With respect to **MIS – 1.10**, significantly more can be discussed in terms of how it is broken and this can be solved. Regarding the observation of the target before and after the closest approach, there seems to be no problem with the current scenario, as the pointing issues only arise right in the few minutes/seconds before closest approach and during it. Thus, it is in this phase that the majority of the analysis will focus on.

As could be seen throughout Chapter 10, in general all simulations of the considered scenario showed average downtime roughly between 10-20 seconds. Despite these values, the average does not contain all the information regarding the distribution of downtime. In the graphs similar to Figure 10.3, it can be seen that, for the baseline case, without any outbursts – which represents the most likely scenario to happen, given the low chance of these events happening during the closest-approach period – only about 40%-50% of cases ever surpass the 0.5 degree pointing error threshold. As such, despite presenting a 10-20 second average of downtime, more than half of the cases are expected to never actually break this requirement.

Nevertheless, to completely fulfil **MIS – 1.10**, the pointing error would have to only rarely surpass the established threshold, and a high level of confidence would have to exist in this being a very unlikely event. Unfortunately, as it stands, this is not the case, thus the navigation's effort must be facilitated. From the conclusions drawn in Section 10.2, the maximum expected downtime observed in a certain scenario can be predicted best by analysing the measurement cut-off time and the closest-approach distance. In terms of the former, it was already discussed in that section that this issue could be mitigated by employing a smarter navigation step adjustment algorithm. The latter, on the other hand, curiously links back to **MIS – 1.4**, and thus the conclusion is that not much can be done to change this.

There might be other aspects which could also help this problematic. If a different image processing algorithm is developed with smaller biases or noises than the one being modelled, that would likely help improve the overall performances of the simulations. This is, however, dependent on a mission planning development which might not happen. One possibility which can be assessed is whether using the NAC – which was mentioned in the requirements in Chapter 2 – together with the already used WAC to obtain better navigation information from afar would improve the state-estimation accuracy enough to solve this problem. The same idea could be applied to any other sensor capable of obtaining navigation information in the first hours of the encounter.

Another option that might be worth considering is changing the status of the image processing error bias from a considered state to an estimated state. As shown in Figure 9.5, this variable has a significant impact on the magnitude of the pointing error, thus trying to estimate it can be a strategy capable of improving the navigation results. Since it affects the measurement results directly, its effect on the results should have good observability. Similarly, one other option that is easy to implement is increasing the measurement and estimation frequency. Since the fly-by happens so fast, increasing the amount of information the estimators have access to in the comet's vicinity should improve their accuracies significantly. Of course, this implementation should always take **CAM – 1.6** into account.

Finally, one method that can soften the burden on the navigation is to reduce the maximum relative fly-by speed of 70 km/s. As discussed multiple times, its high value makes it so that the spacecraft spends a very reduced period of time near the comet, where good images of the target can be captured and the navigation estimation thus corrected. By lessening this effect with a reduction in maximum allowed fly-by speed, it is expectable that the navigation results would improve significantly, which can be used to reduce the downtime experienced by the simulation results in addition to the other strategies suggested here. Since only this speed was studied in this thesis, however, other speed values would have to be analysed to determine the necessary reduction to this 10-70 km/s interval. This is, in fact, one of the topics of discussion in the following section.

From all these options, the ones that seem the most fitting for potentially solving this issue are the implementation of NAC measurements for the encounter phase in addition to the already present WAC ones, since it represents a sensor that is already included in the navigation system, and thus would require no changes to the spacecraft; as well as the estimation of the image processing biases and measurement frequency, as these are changes that can easily be made to the current simulations. Then, should these not suffice, the other options, which would require making bigger changes to the mission can be explored to try and eliminate all the downtime of the simulated scenarios. Not considered here is changing this requirement to not include the closest-approach phase, since it would violate one the mission's science goals of observing the comet's nucleus from a close distance. Thus, it should only be taken into account in case everything else fails.

11.3 Advice for future work

In this section, the areas of work related to the simulation of the Comet Interceptor mission that were glanced over or simplified are enumerated, particularly those which might have a meaningful impact on the results. As such, these are left as suggestions for future work regarding the analysis and development of this endeavour.

11.3.1 Instruments and modelling

As the goal of the thesis was to study the navigation's capability of keeping the payload pointing at the target throughout the mission, a great deal of focus was put into analysing the pointing error originating from this effort. However, this was only done considering the CoCa's FoV. One additional aspect to consider would be the scanning mirror assembly associated with the CoCa. During the majority of the encounter this is not problematic as the spacecraft is so far away that the nucleus is, for all practical purposes, immobile as seen by the vehicle. Nevertheless, in the few minutes preceding the closest approach, the situation is quite the opposite, and thus considering the actuation process of the scanning mirror adjusting the CoCa's orientation might be another relevant source of pointing inaccuracy.

11.3.2 Cometary environment

Moving now from the instruments to the comet's environment, there are also some aspects to mention which could deepen future studies. Starting with the star of the show, the outbursts as represented in this thesis are quite simplified. Instead of being considered by their physical nature, they are represented solely by the impact they have on the measurements. However, as the study of the mission progresses, eventually this method will have to be replaced by its counterpart, as simulating images of the nucleus and then putting them through the image processing algorithm itself takes over the current simulation methods. This shift in simulation practice seems to be one of the next steps in this mission's development, as it replaces the currently used image processing model with just biases, white noises, and whether the nucleus is resolved in the camera images or not. Because of this shift, the outburst models must be extended to ensure that they are detailed enough to properly simulate these phenomena physically when generating images.

There are some other aspects related to the cometary environment which were not developed. Firstly, the mentioned cosmic rays were briefly discussed in Chapter 3, but since little information was found about observations of these events made by space missions, pursuing this issue quickly proved to be difficult. Additionally, even if a reputable model was found describing how often these phenomena happen in space within the Solar System, and most importantly, how meaningful of an impact they presented to the camera measurements, there still would be a lack of information as to how they would influence the processed measurements. Being that these algorithms acted as an unknown layer for this thesis, apart from the simplified model used to simulate measurements, without any further information this issue could not be pursued further. However, with the simulations swapping this model with the method of simulating images and running them through the image processing, the impact of these stray particles can in the future be adequately assessed.

Additionally, there is one problematic regarding the nucleus' surroundings that has not yet been mentioned. Since the coma is a zone that is littered with gas and dust, these materials cause a meaningful impact on the mission. On top of the considered force of the dust striking the spacecraft, the light that is scattered or reflected by these will certainly affect the measurements taken by the optical instruments. In fact, [ESA - Concurrent Design Facility, 2019] states as an attitude and orbit control system requirement that the mission should consider the possibility of a failure of the star trackers for up to one hour due to this light. During an outage of this duration, there are either other instruments capable of estimating the attitude of the spacecraft, or the IMU will have to be used to integrate the changes in attitude from the last valid measurement onward. As a result, the absolute pointing error might become not only driven by the navigation error, but also by the integration error over the outage period, worsening the results of this study.

Finally, one aspect of the simulation which was simplified was the momentum enhancement factor ϵ, used in the computation of the cometary dust impact force. During this study, it was assumed to be equal to 0, which models perfectly plastic collisions, that is, the material is fused onto the vehicle's body. However, depending on particle mass, this value can go beyond 10, according to [Edenhofer et al., 1987], due to phenomena such as the ionisation of the vehicle's exposed faces due to the impacts. Because of this, it is possible that this force should be modelled as being several times larger than what is considered in this thesis, which can alter the trajectory and the navigation quite significantly, and thus should be the subject of further analysis.

11.3.3 Navigation

In terms of navigation being studied, due to the high number of variables present, some shortcuts had to be considered to make the amount of results being generated and the time to do so manageable with the available computational resources. Because of this, only six filters were considered, and for the half that is formulated in batches, only a batch size of 5 was taken. A higher value was not analysed, nor would it be advisable, since the fly-by happens so fast that, despite the time step being reduced to 1 second, waiting for more measurements to be taken before performing a state estimation would leave the filter more and more time without new corrections, leading to a pointing error curve with spikes between batches of measurements. Realistically, this could mean that the instruments, when nearest to the comet's nucleus, could oscillate between being correctly and incorrectly pointed towards the target. Nevertheless, considering lower batch sizes than 5 could certainly be studied, and might even help the batch filters perform better than what is shown in this thesis.

11.3.4 Scenario

When considering the variants of the scenario being analysed, there were also some avenues which were left unconsidered. In the baseline analysis, a significant number of simulations were run with each filter, meaning that the results can statistically be considered robust. However, this was only done for the baseline scenario. Despite this, some of its main aspects were set and not variable, most importantly, the fly-by relative speed. The magnitude of this parameter is so crucial that it can make the whole encounter be seven times longer in one extreme than in the one being considered. Thus, it would be interesting to study other fly-by speeds.

Once the outbursts are introduced, the number of variables skyrockets again, with a whole 3D space of possible source positions, and varying intensity, duration, and starting time. Additionally, a threshold for the measurement outlier rejection must be studied. Assessing all the impacts of these parameters for different filters and randomly generated scenario noises would be an insane task, and that is why simplifications had to be made regarding some of these aspects. For example, considering the outbursts intensities, initially 2 and 15 times the radius of the comet were studied, only to find that the former barely impacted the navigation, whereas the latter had the opposite effect, and thus the considered intensities were changed to focus more on more intermediate values. For the case of the outburst origin, the chosen rationale was to only consider those pointing outward from the nucleus' horizon as seen by the spacecraft, as these should be the ones with the greatest impact. But, for outbursts taking place right before the closest approach, perhaps these orientations are not the most consequential, because, as the spacecraft approaches the nucleus, the outbursts increasingly point not outward from the horizon but towards or away from the vehicle, in case they are in the comet-Sun line. Thus, it would be wise to broaden the orientations considered for these outbursts in future studies, and follow a similar approach with some of the other outburst-related variables.

11.3.5 Software framework

Finally, this discussion would not be complete without mentioning the FCS-ATOMIC framework. Although incredibly versatile when handling the creation of scenarios due to its modularity, this also implies that there are certain aspects of the code which are not relevant to the specific mission being simulated and thus complicate and burden the execution. When performing a thorough analysis, which requires enormous amounts of data, these small burdens start to add up, both in terms of computational time and storage space of variables saved after the simulation. As such, and especially considering that the Comet Interceptor is such a simple mission, with easy to compute dynamics, only navigation and no guidance or control, a single sensor taking measurements, and an almost open loop simulation for the first 18-19 hours, it would be extremely beneficial to perform these analyses within a simplified framework.

All in all, this thesis study did not spend enough time performing these optimisations – despite those mentioned in Section 7.3, which were almost accidental – and such an effort would have greatly facilitated the task of gathering data by running multiple simulations. That could mean that more data could have been gathered in the same time, thus strengthening the confidence in the results.

11.4 Applicability in future missions

Having discussed the conclusions drawn from the study performed throughout this document, it is now essential to establish applicability criteria for other missions. Ideally, these results would be useful for every future space endeavour that shared some characteristic with the Comet Interceptor, mainly being aimed at a comet, or at a body of similar dimensions and with the possibility of presenting outburst-like phenomena. Another requirement needed to compare future missions with the Comet Interceptor is navigation-related. Since this whole analysis relies on the usage of optical images to estimate the comet's position, then such a navigation strategy would be required to establish parallels with the results gathered in this thesis.

However, the very nature of this mission, especially that of the encounter phase, when the spacecraft actually flies by the target's nucleus, is quite particular. The very high relative high speed at which the vehicle travels with respect to the comet makes it so that the mission must rely on autonomous navigation during this phase. Additionally, the small time period that it spends close enough to the nucleus to actually capture it for what it really is, and not just a point of light in the camera images as seen from far away, results in a profile where there is little to no information to update the spacecraft's own estimated state for the majority of the time, which then quickly turns into a few minutes when most of the information must be gathered in this short period as the nucleus is close enough to be seen. Because of this concentration of the estimation effort in the few minutes preceding the closest approach, there is a heavy weight on the navigation procedure, which is particularly influenced, in the case of the Comet Interceptor, by the measurement-related uncertainties, and by the deviation of the trajectory from the established nominal.

Thus, summarising all the points made above:

- High speed fly-by mission.
- Target body with a small mass – preferably a comet.
- Autonomous navigation.
- Optical driven navigation.
- Presence of outburst-like jet phenomena.

Missions obeying all or some of the laid out conditions can take advantage of the performed analysis as a source of information regarding the expected navigation results from different filters, the impact of the usage of outlier rejection techniques, as well as the effect of including outbursts into the captured images. Some of the missions described in Chapter 2, such as Giotto, have followed this recipe in the past. In fact, this mission was the source of much of the theoretical assumptions used to start building this Comet Interceptor endeavour. Others, on the other hand, like Rosetta, apply a more thorough methodology, deciding to spend more time inside the comet's vicinity and studying it from close, instead of just performing a quick fly-by.

This does not mean, however, that these other kinds of missions might not find use in the subjects discussed here. A mission similar to Rosetta might not fall within the scope of "high-speed fly-by mission". Nevertheless, the type of outbursts described and analysed in this thesis are one and the same phenomenon as experienced in other comets, regardless of the spacecraft's orbit. As such, the conclusions to be drawn from their impact on the overall performance of the mission might be dissimilar, but the way that these events are described and even modelled can be the same.

Not only that, but there are also some other lessons to be learned that are not mission dependent. For example, the trade-off between the UKFs' reliability versus performance and the impact of rejecting outliers on their performance. As was seen in the past few chapters, the batch UKF, especially, really stood out as the best performed for all the scenarios being analysed, despite being imperfect in terms of robustness. This analysis on the dependability of the UKFs is likely less reliant on the mission's characteristics than the other conclusions discussed above, and can show a balance between a better performance and a higher probability of failure that is inherent to these filters when used in similar conditions and with a similar formulation as done in this thesis.

Bibliography

Bhaskaran, S., Riedel, J. E., and Synnott, S. P. "Autonomous target tracking of small bodies during flybys". *AAS/AIAA Spaceflight Mechanics Meeting, Maui, HI*, Feb. 2004.

Bierman, G. J. *Factorization Methods for Discrete Sequential Estimation*, volume 128. Elsevier, 1977. doi: 10.1016/S0076-5392(08)60888-7.

Bitzer, S. *The UKF exposed: How it works, when it works and when it's better to sample*. 2016. doi: 10.5281/zenodo.44386.

Carpenter, J. R., D'Souza, C. N., Markley, F. L., and Zanetti, R. *Navigation Filter Best Practices*. NASA Engineering and Safety Center, Hampton, VA, 2018.

Di Gangi, M., Cantarella, G., and Vitetta, A. Solving stochastic frequency-based assignment to transit networks with pre-trip/en-route path choice. *EURO Journal on Transportation and Logistics*, 8, 05 2019. doi: 10.1007/s13676-019-00142-9.

Edenhofer, P., Bird, M. K., Brenkle, J. P., Buschert, H., Kursinki, E. R., Mottinger, N. A., Porsche, H., Stelzried, C. T., and Volland, H. "Dust Distribution of Comet p/ Halley's Inner Coma Determined from the Giotto Radio Science Experiment". *Astronomy and Astrophysics*, 187:712, Nov. 1987.

Fung, W. K. Outlier diagnostics in several multivariate samples. *"Journal of the Royal Statistical Society: Series D (The Statistician)"*, 48(1):73–84, 1999.

Golish, D. R., d'Aubigny, C. D., Rizk, B., DellaGiustina, D. N., Smith, P. H., Becker, K., Shultz, N., Stone, T., Barker, M. K., Mazarico, E., et al. "Ground and In-Flight Calibration of the OSIRIS-REx Camera Suite". *Space Science Reviews*, 216(12):259–264, 2020. doi: 10.1007/s11214-019-0626-6.

Groves, P. D. *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems (2nd Edition)*. Artech House, 2013. ISBN 978-1-60807-005-3.

Hashimoto, T., Kubota, T., Kawaguchi, J., Uo, M., Shirakawa, K., Kominato, T., and Morita, H. "Vision-based guidance, navigation, and control of Hayabusa spacecraft - Lessons learned from real operation". *IFAC Proceedings Volumes*, 43(15):259–264, 2010. doi: 10.3182/20100906-5-JP-2022.00045.

Julier, S. J. and Uhlmann, J. K. "Unscented filtering and nonlinear estimation". *Proceedings of the IEEE*, 92(3): 401–422, 2004. doi: 10.1109/JPROC.2003.823141.

Keller, H. U. and Kührt, E. "Cometary Nuclei—From Giotto to Rosetta". *Space Science Reviews*, 216, 2020. doi: 10.1007/s11214-020-0634-6.

Królikowska, M. Long-period comets with non-gravitational effects. *American Astronomical Society*, 427: 1117–1126, Dec. 2004. doi: 10.1051/0004-6361:20041339.

Leonard, J. M., Antreasian, P. G., Jackman, C. D., Page, B., Wibben, D. R., and Moreau, M. C. "Orbit Determination Strategy And Simulation Performance For Osiris-rex Proximity Operations". *10th International ESA Conference on Guidance, Navigation & Control Systems*, 2017.

Lissauer, J. J. and de Pater, I. *Fundamental Planetary Science: Physics, Chemistry and Habitability*. Cambridge University Press, Cambridge, 2013.

Lorenz, D. A., Olds, R., May, A., Mario, C., Perry, M. E., Palmer, E. E., and Daly, M. "Lessons learned from OSIRIS-REx autonomous navigation using natural feature tracking". *2017 IEEE Aerospace Conference*, pages 1–12, 2017. doi: 10.1109/AERO.2017.7943684.

Mastrodemos, N., Kubitschek, D. G., and Synnott, S. P. "Autonomous Navigation for the Deep Impact Mission Encounter with Comet Tempel 1". *Space Science Reviews*, 117:95–121, 2005. doi: 10.1007/s11214-005-3394-4.

Park, E., Park, S., Roh, K., and Choi, K. Satellite orbit determination using a batch filter based on the unscented transformation. *Aerospace Science and Technology*, 14(6):387–396, 2010. ISSN 1270-9638. doi: https://doi.org/10.1016/j.ast.2010.03.007.

Pellacani, A., Graziano, M., Fittock, M., Gil, J., and Carnelli, I. "HERA vision based GNC and autonomy". *8th European Conference for Aeruonautics and Space Sciences*, 2019. doi: 10.13009/EUCASS2019-39.

Ridolfi, G. and Mooij, E. Regression-based sensitivity analysis and robust design. In *Space Engineering: Modelling and Optimization with Case Studies*, pages 303–336. Springer International Publishing, 2016. doi: 10.1007/978-3-319-41508-6_12.

Riedel, J. E., Bhaskaran, S., Desai, S., Han, D., Kennedy, B., et al. "Using Autonomous Navigation for Interplanetary Missions: The Validation of Deep Space 1 AutoNav". *International Conference on Low-Cost Planetary Missions*, 2000.

Space Telescope Science Institute. *Wide Field and Planetary Camera 2 Instrument Handbook for Cycle 14*. Maryland, US, 2004.

Süli, E. and Mayers, D. *An Introduction to Numerical Analysis*. Cambridge University Press, Cambridge, UK, 2003.

Tapley, B. D., Schutz, B. E., and Born, G. H. *Statistical Orbit Determination*. Elsevier Academic Press, California, USA, 2004.

Thorton, C. L. and Bierman, G. J. *UDUT Covariance Factorization for Kalman Filtering*, volume 16 of *Control and Dynamic Systems*. Academic Press, 1980. doi: 10.1016/B978-0-12-012716-0.50011-X.

Vincent, J. B., A'Hearn, M. F., Lin, Z. Y., El-Maarry, M. R., Pajola, M., et al. "Summer fireworks on comet 67P". *Monthly Notices of the Royal Astronomical Society*, 462:184–194, Sept. 2016. doi: 10.1093/mnras/stw2409.

Wakker, K. F. *Fundamentals of Astrodynamics*. Delft University of Technology, 2015.

Weissman, P. R. and Lowry, S. C. "Structure and density of cometary nuclei". *Meteoritics & Planetary Science*, 43(6):1033–1047, 2008. doi: 10.1111/j.1945-5100.2008.tb00691.x.

Wesołowski, M., Gronkowski, P., and Tralle, I. "Selected mechanisms of matter ejection out of the cometary nuclei". *Icarus*, 338, 2020. doi: 10.1016/j.icarus.2019.113546.

Williams, B., Antreasian, P., Carranza, E., Jackman, C., Leonard, J., Nelson, D., Page, B., Stanbridge, D., Wibben, D., Williams, K., et al. "OSIRIS-REx Flight Dynamics and Navigation Design". *Space Science Reviews*, 214(4):69, June 2018. doi: 10.1007/s11214-018-0501-x.

Zanetti, R. and D'Souza, C. Recursive Implementations of the Schmidt-Kalman 'Consider' Filter. *The Journal of the Astronautical Sciences*, 60, 12 2013. doi: 10.1007/s40295-015-0068-7.

Internal GMV and ESA documents

ESA - Concurrent Design Facility. *CDF Study Report - Comet Interceptor - Assessment of Mission to Intercept a Long Period Comet or Interplanetary Object*. Dec. 2019.

GMV. *GMV'S Proposal To TAS UK For Comet Interceptor PHASE A/B/C/D/E1 - GNC Subsystem*. 2020a.

GMV. *D1- Hypervelocity Fly-by GNC Detailed Design and Justification*. 2020b.

GMV. *Comet Interceptor: Full Mission Characterisation Document*. 2021a.

GMV. *GNC DESIGN JUSTIFICATION FILE - COMET-INTERCEPTOR*. 2021b.