

Analytic State Transition and Sensitivity Matrices for Real-Time Orbit Determination of LEO Satellites

Quinten Eduardus van Woerkom

Delft University of Technology



Analytic State Transition and Sensitivity Matrices for Real-Time Orbit Determination of LEO Satellites

by

Quinten Eduardus van Woerkom

in partial fulfilment of the requirements to obtain the degree of Master of Science
at the Delft University of Technology,
to be presented and defended on Friday, September 29, 2023 at 09:30 AM.

Student number: 4676998
Project duration: March 15, 2023 – September 29, 2023
Thesis committee: Dr. ir. E.J.O. Schrama, TU Delft, supervisor
Dr. ir. E. Mooij, TU Delft, committee chair
Dr. ir. E. van Kampen, TU Delft, external committee member

Cover: “Satellite Earth Orbiting” by SpaceX under CC0

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

In front of you, you find the end product of my time at TU Delft. This thesis, as well as the education leading up to it, have been the product of uncountable hours of work: hard work, but work I enjoyed doing.

Still, it is far from my work alone that made this possible. This thesis was made possible in no insignificant part by my supervisor, dr.ir. Ernst Schrama: I would like to thank him for guiding me through this tough process, and for providing me with invaluable feedback. Without him, this work would be nowhere close to what it is now.

Then, I must thank all the friends that have accompanied me during my time here in Delft. From childhood friends to fellow students, from the friends I made at the VSV “Leonardo da Vinci” to the office mates in NB 2.56, I have made friends for life. I could not and would not have done it without you.

Lastly, I am grateful for the unconditional support of my family. Many thanks go to my parents, Eddy and Aaltje, without whom I would not have been where I am today. You have always supported me and my choices, in spite of what those choices would bring, and for that I am thankful. I am also indebted to my brother, Quirijn, who has always been my mirror and my conscience. Let us keep pushing each other. *Dit is ook jullie werk.*

Quinten E. van Woerkom
Delft, September 2023

Summary

For Earth observation missions, awareness of the orbit of a spacecraft is more than just an operational necessity: observation of specifically the mass properties of our planet requires intimate knowledge of the trajectory of a spacecraft. For other types of measurements like altimetry, Radio Occultation (RO), or Synthetic Aperture Radar (SAR), knowledge of the spacecraft position is a fundamental necessity in correct interpretation of the observables. The quality or even the feasibility of these observables is limited in a large part by the accuracy of the spacecraft position used to derive them: it follows that improved real-time knowledge of the spacecraft position leads directly to improved quality of the near-real-time interpretation of these observables. Beyond direct scientific application, improved real-time onboard positional knowledge of spacecraft is also useful for improved operational performance in aspects such as formation flying and rendezvous, with the accompanying applications.

In recent years, the quality of Global Positioning System (GPS) and Galileo broadcast ephemerides, used to estimate their real-time antenna positions and internal clock offsets, have improved to decimetre-level accuracies [41, 42]. Combining measurements from both systems has been shown to allow real-time orbit accuracies down to 3.5 cm in the radial direction [14, 38]. Simultaneously, the Doppler Orbitography and Radiopositioning Integrated by Satellite (DORIS) system has made similar leaps in ground system quality to support radial orbit determination performance to the level of 3 cm [31]. To make optimal use of these increasingly accurate observables, the dynamic models underlying the orbit determination filters must also become more and more accurate, with an associated increase in computational effort. Hence, reduction of the computational effort needed to model spacecraft dynamics to this degree of accuracy becomes an effective way to expand the set of scenarios in which centimetre-level orbit determination is possible. We link this in particular to the prospect of centimetre-level onboard orbits, which must be computed under the even tighter computational constraints of an embedded system.

An important component of the dynamical model inside an orbit determination filter is its state transition matrix, which describes the response of the system to changes in its initial state. This matrix, also known as the fundamental matrix or matrizant, is used inside an orbit determination filter for two distinct though deeply connected steps: propagation of the error covariance estimate and updating the state using new observations. Naive computation of this matrix using the same trajectory model as used to propagate the spacecraft trajectory estimate is expensive, due to the quadratically higher dimension of the matrizant as compared to the orbital state. We contrast this high computational effort with the fact that the state transition matrix need not be known to the same accuracy as the state estimate itself. After all, it is used only to model *perturbations* of the trajectory: errors in the modelling of perturbations will naturally be smaller in absolute magnitude than equivalent errors in the trajectory model itself. It follows that efficient approximations of the state transition matrix are a prime candidate for computational savings inside the dynamical model.

Earlier work in the field presented analytic or semi-analytic approximations of the *orbital* state transition matrix for elliptical orbits [8, 9] or even when including higher-order zonal harmonics [15, 16, 57]. None of these formulations, however, include the derivatives with respect to dynamic parameters that might be part of the estimated state. For the purposes of centimetre-level orbit determination in low Earth orbit, examples of such parameters that need to be estimated due to imperfect a priori models are the radiation pressure and aerodynamic drag coefficients. State-of-the-art real-time filters make no mention of the approximations used in their state transition matrices to compute this sensitivity matrix. For DLR's Real-Time Navigation Filter (RTNAV), precursor work suggests that a numerically-integrated simplified force model was used [39]; for the DORIS Immediate on-board Orbit Determination (DIODE) and University of Padua filters, the used approximations are undocumented [14, 31], although it is certainly expected that computational simplifications were applied to make real-time application feasible.

In this research contribution, we develop and present an analytic approximation of the state transition matrix based on the Hill-Clohessy-Wiltshire equations of relative motion for circular orbits. This approximation improves over other state transition matrix approximations in the published state-of-the-art in its ability to include sensitivities to perturbing accelerations. Additionally, its fully-analytic nature

makes it computationally much cheaper to evaluate than the numerically-integrated simplifications known to be applied in other real-time filters. This combination of both computational efficiency and the ability to model perturbing accelerations makes it a good candidate for efficient real-time precise orbit determination.

The Hill-Clohessy-Wiltshire approximation of the state transition matrix is evaluated by means of a case study of filtering orbit determination performance of Sentinel-6A on 1 January 2022. Sentinel-6A “Michael Freilich” is an altimetric satellite [19], whose highly-circular low Earth orbit is representative of the class of orbits that one would encounter in the science missions that would benefit most from real-time centimetre-level orbits. In this context, the Hill-Clohessy-Wiltshire approximation is compared with two intermediate approximations that serve to show the effects of neglecting the eccentricity of the orbit and the oblateness of the Earth. This comparison is done in three metrics, representing the quality of the overall orbit estimate, of the error covariance propagation, and of the state update for data inclusion. In this manner, a holistic picture is sketched of not just the overall quality of the Hill-Clohessy-Wiltshire approximation, but also the specific errors introduced by each assumption as well as their effects on each specific step in the orbit determination filter.

Results from this case study show that, using the reference state transition matrix, the implemented orbit determination filter supports positioning accuracies of 0.5 ± 10 , -1.3 ± 16 , and -1.6 ± 7 cm in the radial, along-track, and cross-track direction, respectively. Although not strictly matching it, this makes the filter performance sufficiently representative of the 3 cm achieved in the state-of-the-art. The difference in performance is attributable to the neglect of ocean tides in the dynamical model and of carrier-phase wind-up in the observation model. The former neglect is based on the current lack of such a model in Tudat, the astrodynamics toolkit used for the dynamical model. The latter was needed to keep the scope of the present observation model implementation feasible.

Application of the Hill-Clohessy-Wiltshire state transition matrix results only in small deterioration of the filter performance. It results in deviations from the reference solution by 0.1 ± 0.8 , 0.1 ± 0.9 , and 0.0 ± 0.2 mm in radial, along-track, and cross-track direction, respectively. During filter convergence slightly larger deviations are found, but they never exceed 5, 5, and 1 mm in each of the aforementioned directions. This confirms that the Hill-Clohessy-Wiltshire approximation is indeed compatible with the 3 cm radial orbital accuracies achieved in the current state-of-the-art, and that it can support future reductions down to a centimetre. At a decrease in total filter runtime of more than 10% as compared to other approximations and of more than 50% compared to the reference model, this poses significant computational savings.

Comparison with an intermediary state transition matrix based on point-mass Earth gravity shows that a minor but observable part of the deviation from the reference solution is caused by the neglect of orbital eccentricity. As the assumption of negligible orbital eccentricity is fundamental to the derivation of the presented Hill-Clohessy-Wiltshire approximation, this provides an indication of the lowest accuracies that could be achieved with this method. The deviation found in this comparison is up to several millimetres, but converges over time to only fractions of a millimetre. This indicates that the Hill-Clohessy-Wiltshire approximation could in principle support much lower accuracies than indicated by the presented implementation.

A more significant part of the deviation from the reference solution in the present implementation of the Hill-Clohessy-Wiltshire approximation can be attributed to the neglect of the second order zonal term J_2 in computation of Earth gravity. This is not a fundamental limitation of the Hill-Clohessy-Wiltshire approximation: instead, it would be a rather natural future extension of the presented method. Indeed, the largest part of the J_2 effect could be captured as additional forcing term in the Hill-Clohessy-Wiltshire equations of relative motion used to derive the state transition matrix. The results of this work suggest that this would allow further reduction of the magnitude of deviations due to state transition matrix simplifications by more than a factor of two at no significant additional computational cost. This would extend the applicability of this approximation even further, down to far-future real-time orbital accuracies of mere millimetres.

Contents

Preface	i
Summary	ii
Acronyms	vi
Symbols	vii
1 Introduction	1
1.1 Technical background	1
1.2 Societal relevance	2
1.3 Research contribution	3
1.4 Report outline	3
2 Research Objectives	5
2.1 Scope	5
2.2 Knowledge gaps	6
2.3 Research question and subquestions	6
3 Methodology	9
3.1 Orbital context	9
3.2 Choice of variational equations	10
3.3 Analyses	12
3.4 Software architecture	14
4 Observation Model	16
4.1 Observation equations	16
4.2 GNSS ephemerides	20
4.3 Transformation to the GCRS	20
4.4 Code noise analysis	22
5 Trajectory Model	24
5.1 Space-time reference system	24
5.2 Benchmark model	24
5.3 Integrator	25
5.4 Modelled accelerations	27
5.5 Final model	29
6 Variational Equations	33
6.1 Solution through numerical integration	33
6.2 Analytical solution via Hill-Clohessy-Wiltshire equations	36
6.3 Numerical solution using the difference quotient	39
7 State Estimation	41
7.1 Filter mechanization	41
7.2 State design	41
7.3 Data processing	44
7.4 Outlier detection	46
7.5 Parameter initialization	47
7.6 Parameter removal	48
7.7 Covariance propagation	48
7.8 Filter parameter tuning	49

8	Results	51
8.1	Filter performance	51
8.2	Inter-solution differences	55
8.3	Error ellipsoids	56
8.4	Differential correction error	56
8.5	Runtimes	59
9	Discussion	61
9.1	Agreement between state transition matrix accuracy metrics	61
9.2	Effect of force model simplifications	62
9.3	Limitations of this case study	63
10	Conclusion	65
	References	67
A	Derivation of the Differential Correction Error	71
A.1	Differential correction	71
A.2	Application in GNSS-based orbit determination	72
B	State Transition Matrix Differences for the Precise GPS Filter	73

Acronyms

Acronym	Definition
ARF	Antenna Reference Frame
BDS	BeiDou Navigation Satellite System
CIO	Celestial Intermediate Origin
CNAV	GPS Civil Navigation Message
DIODE	DORIS Immediate on-board Orbit Determination
DLR	Deutsches Zentrum für Luft- und Raumfahrt
DORIS	Doppler Orbitography and Radiopositioning Integrated by Satellite
ECI	Earth-centered inertial frame
EKF	Extended Kalman Filter
EOP	Earth Orientation Parameters
FNAV	Galileo FNAV Navigation Message
GCRS	Geocentric Celestial Reference System
GGTO	Galileo-GPS Time Offset
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GTRF	Galileo Terrestrial Reference Frame
IAU	International Astronomical Union
IERS	International Earth Rotation and Reference Systems Service
ISC	Inter-Signal Correction
ITRF	International Terrestrial Reference Frame
ITRS	International Terrestrial Reference System
LEO	Low Earth Orbit
LNAV	GPS Legacy Navigation Message
MGEX	Multi-GNSS Pilot Project
NGA	National Geospatial-Intelligence Agency
POD	Precise Orbit Determination
PODRIX	Precise Orbit Determination Receiver
PPN	Parameterized Post-Newtonian
RINEX	Receiver Independent Exchange Format
RMS	Root Mean Square
RO	Radio Occultation
RTNAV	DLR's Real-Time Navigation Filter
SAR	Synthetic Aperture Radar
SISRE	Signal-In-Space Ranging Error
SLR	Satellite Laser Ranging
SRF	Spacecraft Reference Frame
SRIF	Square Root Information Filter
SRP	Solar Radiation Pressure
TDB	Barycentric Dynamical Time
TGD	Timing Group Delay
TIO	Terrestrial Intermediate Origin
TIRS	Terrestrial Intermediate Coordinate System
TT	Terrestrial Time
UKF	Unscented Kalman Filter
UTC	Universal Coordinated Time
WGS84	World Geodetic System 1984

Symbols

Symbol	Definition
J_2	Earth gravitational field expansion term due to its oblateness
P	Covariance matrix
R	Square root information matrix, such that $P = R^{-1}R^{-T}$
V	Volume of an error ellipsoid
Φ	State transition matrix
δt_{DCB}	Differential code bias
δt_{ISB}	Inter-system clock bias
δt_{TGD}	Time group delay
δt_{GGTO}	Inter-system bias of Galileo with respect to GPS time
δt_{clk}^{rel}	Relativistic clock offset
δt_{rcv}	Receiver clock offset
δt_{stc}^{rel}	Shapiro time delay
δt_{trm}	Transmitter clock offset
λ	Eigenvalue of a matrix
μ	Gravitational parameter of the Earth
ϕ	Carrier-phase observation
τ	GNSS signal travel time
θ	Relative orientation between two ellipsoids
a_e	Empirical accelerations
a	Acceleration
b	Carrier-phase ambiguities
p	Force model parameters estimated as part of the state
r	Position
s	SISRE-compensating state parameters
t	Time offset parameters estimated as part of the state
v	Eigenvector of a matrix
v	Velocity
x	Full system state
y	Orbital state, describing the current position and velocity
z	State observation inside the Square Root Information Filter
a_r	Radial empirical acceleration
a_s	Along-track empirical acceleration
a_w	Cross-track empirical acceleration
a	Semi-major axis
b	Carrier-phase ambiguity
c_d	Aerodynamic drag coefficient
c_r	Solar radiation pressure coefficient
c	Speed of light
i	Orbital inclination
n	Mean motion
p	Pseudorange
$r_{\oplus}$	Equatorial radius of the Earth
t	Time
u	Argument of latitude

1

Introduction

This report presents an analysis of the effectiveness of an analytic approximation for the state transition matrix of a typical altimetric low-Earth orbit, based on the Hill-Clohessy-Wiltshire equations of relative orbital motion. Before delving further into this subject in the remainder of this report, we shall first start with an introduction of this topic, and its place within the field of orbit determination. In Section 1.1, we start with a technical introduction to the topic, with a description of the relevance of the state transition matrix within real-time orbit determination. This is followed in Section 1.2 by a contextualization of such real-time orbits in terms of their greater societal applications. Afterwards, with this greater context in mind, we describe the contribution of this research towards the state-of-the-art in Section 1.3. Finally, we close the introduction with a description of the contents of the report. This can be found in Section 1.4.

1.1. Technical background

Over the years, different tracking systems have been developed and used to determine spacecraft orbits, whether it be for operational or scientific purposes. Early satellites such as Explorer I and Sputnik I were tracked using doppler tracking and radio angular measurements from specific ground tracking stations [26, 27]. Even more accurate approaches such as Satellite Laser Ranging (SLR) were also developed, though such optical techniques are weather-dependent [36]. Later, near-globe-covering networks would emerge, such as the Doppler Orbitography and Radiopositioning Integrated by Satellite (DORIS) system specifically developed for satellite orbit determination by CNES in the 1980s [52], and Global Positioning System (GPS), developed by the U.S. Department of Defense in the 1970s, with first application to satellite orbits onboard Landsat-4 [6, 56].

Since recent years, decimetre-level Root Mean Square (RMS) Signal-In-Space Ranging Error (SISRE) can be achieved with GPS and Galileo broadcast ephemerides [41, 42]. This results in the possibility of real-time onboard orbit determination to orbital accuracies of several centimetres: radial accuracies of 3.5 cm RMS were achieved using combined GPS and Galileo measurements with DLR's Real-Time Navigation Filter (RTNAV) for Sentinel-6A [14, 38]. With observations from BeiDou Navigation Satellite System (BDS)-3, a slightly reduced radial accuracy of 6.1 cm was demonstrated for Haiyang-2D [34]. Similarly, recent enhancements in the DORIS system and associated data products have permitted onboard orbit determination on the level of 3 cm RMS in the radial direction for Jason-2, Saral, CryoSat-2, and Sentinel-6A using the DORIS Immediate on-board Orbit Determination (DIODE) filter [31, 38].

To allow the Extended Kalman Filters (EKFs) used in these reduced-dynamic orbit determination methods to make optimal use of these advances in the quality of observables, higher-quality dynamic models are necessary. In computing such sophisticated trajectory models, the inherently limited processing power available onboard spacecraft becomes a bottleneck: a need follows for improved computational efficiency without reduction in orbit quality. Beyond direct modifications to the trajectory model itself, one avenue for such improvement is the approach to solution of the variational equations. These equations describe the variation over time of the system's state due to perturbations in an earlier state. Their solution, referred to as the state transition matrix, fundamental matrix, or matrizant $\Phi(t, t_0) = \partial \mathbf{x}(t) / \partial \mathbf{x}(t_0)$, describes the change in system state $\mathbf{x}(t)$ due to infinitesimal changes in the initial state $\mathbf{x}(t_0)$. Accurate yet computationally efficient approximations to this state transition matrix




Figure 1.1: Schematic representation of filter update steps making use of the state transition matrix $\Phi(t, t_0)$

are the principal topic of this thesis.

Based on the chosen state formulation, the state transition matrix can take many forms. Of particular interest in this work are two specific submatrices of this matrix that appear when the system state x contains some orbital state y and dynamic parameters p . Then, the matrix will contain an *orbital* state transition matrix $\Phi_y(t, t_0) = \partial y(t) / \partial y(t_0)$ and a sensitivity matrix $\Phi_p(t) = \partial y(t) / \partial p$. As opposed to most other components of the state transition matrix, computation of these submatrices is non-trivial due to their dependence on the dynamical model underlying the system. In this work, these submatrices will be referred to as the orbital state transition matrix and sensitivity matrix, respectively. In other works, the orbital state transition matrix is generally referred to as simply the state transition matrix, with the intended usage following from the context. For clarity, this will *not* be done in this work: any reference to the state transition matrix denotes the sensitivity of the total system state. Still, the reader is advised to be mindful of this distinction when consulting other sources.

The state transition matrix is needed for two separate steps within reduced-dynamic orbit determination using an EKF. In the first place, it is used during data inclusion to estimate the change in initial state needed to better fit those new observations. A second application, relevant specifically to the EKF, is propagation of the state error covariance estimate of the filter. This step is required to provide a suitable estimate of the uncertainty associated with a given a priori estimate after propagation in time. Both steps are drawn schematically in Figure 1.1.

Neither step requires an exceedingly accurate state transition matrix to be performed to satisfaction. The differential correction is applied repeatedly throughout the orbit determination process, meaning that small residuals after one correction are permissible. Similarly, the error covariance need not be known to the same accuracy as the orbital trajectory itself. After all, it is used only to weigh the importance of the a priori state estimate relative to new observations, the accuracy of which is also known only to rough approximation. This restricted benefit from a high-accuracy matrix is contrasted with the computational effort needed to solve the variational equations, which are of similar complexity and quadratically higher dimension than the equations of motion, to full accuracy. It follows that solution of a reduced set of variational equations is both justifiable and computationally attractive.

1.2. Societal relevance

In the first place, orbital positioning is a pragmatic requirement for spacecraft operation. Correct antenna and camera pointing, precise measurement localization, and orbital manoeuvring all necessitate awareness of the location and orientation of a spacecraft in its orbit. Such operational positioning generally does not require a high-accuracy solution: indeed, a relatively coarse and easy-to-compute “navigation solution” is used for these real-time operational scenarios [37, 53]. This navigation solution is generally purely kinematic: this has the advantage of being robust and computationally light-weight, but comes at the cost of reduced accuracy.

For a large class of missions, orbit determination holds more significance than being a mere operational necessity: it can be one of the core mission drivers, particularly in the context of Earth observation. Beyond simply encoding measurement positions, a spacecraft’s orbit is one of the primary geodetic tools available for high-accuracy gravimetric [48], altimetric [30], oceanographic [22, 47], and Radio Occultation (RO) [56] measurements of our Earth and nearby celestial bodies. These observations

allow the measurement of properties such as its gravity field, mass distribution, tropospheric activity, and sea level, or changes therein: these are all important tools for characterizing a planet. In the case of Earth, they are core tools for monitoring its climate. In general, such missions resort to ground-based post-processing of the data to achieve the highest level of accuracy, at a sub-centimetre 3D RMS position accuracy [56]. This is not solely due to the more limited computational power typically found onboard spacecraft, but also due to the lower accuracy of the Global Navigation Satellite System (GNSS) broadcast ephemerides and Earth Orientation Parameters (EOP) employed in generating (near-)real-time solutions [38]. For numerous applications, the longer latency incurred by this ground-based post-processing is a reasonable compromise, as they are not time-sensitive.

Nevertheless, some specific applications necessitate or find value in the lower latency of real-time orbit solutions permitted by usage of broadcast or ultra-rapid GNSS ephemerides and EOPs. Illustrative instances are the near-real-time data products from Synthetic Aperture Radar (SAR), RO, or altimetric measurements, which must be provided to end users with latencies on the order of three hours [17]. All three aforementioned data products are limited in quality in part by the accuracy of the orbit solutions used to generate them. There are use cases with even more exacting requirements on the timing and accuracy of near-real-time solutions, leading to loss of life or loss of mission if not met. Examples include sea level anomaly detection for oceanic eddy identification (which has implications on submarine detection) [55] or future tsunami warning systems [47]. Finally, certain approaches to formation flying, such as the GNSS-based PROBA-3 mission [35], also benefit from improved onboard orbit determination solutions. Additionally, formation-flying missions also benefit indirectly from further analysis of solutions to the variational equations: indeed, the mathematics behind formation flying bear strong resemblance to the process of orbit determination.

1.3. Research contribution

This research contributes to the state-of-the-art in three manners. Firstly and principally, an analytic state transition matrix approximation, based on solutions to the Hill-Clohessy-Wiltshire equations, is presented. This approximation is suitable for application in particular to altimetric low Earth orbits, one of the primary use cases for real-time orbits. The accuracy and computational efficiency of this approximation are characterized through comparison with a reference matrizant obtained through numerical solution of the full-fidelity variational equations. Additionally, it is compared with two other, purportedly more accurate, approximations. This comparison serves to establish the sources of the errors introduced by the approximation, and to determine the effectiveness relative to existing state transition matrix estimates from literature.

The second contribution is the manner of comparison. We introduce a novel comparison based on three key performance characteristics, which together sketch a holistic picture of the accuracy of a state transition matrix approximation in the context of an orbit determination filter. Besides the primary system performance indicator, being simply the difference in the resulting orbit estimates, two specific characterizations are introduced to interpret the accuracy of a matrizant in terms of its effect on the different filtering steps it is applied in. The first measure follows from a geometric interpretation of the error covariance matrix, previously applied already in analysis of state transition matrix approximations by Der and Danchick and by Der [15, 16]. The second measure is a generalization of the differential correction error measure as introduced by Montenbruck and Gill, extended in this work to remove the need for normalization of the state transition matrix prior to application [36].

Finally, this work is, to the author's knowledge, the first to introduce a representative real-time orbit determination filter into the public domain, based only on algorithms that could be suitable for real-time and onboard orbit determination. Similar filters exist in literature, among which RTNAV [38], DIODE [31], the Haiyang-2D filter [34], and a filter by the University of Padua [14]. However, all these filters provide only limited information on their implementation details, limiting the possibility of follow-up research. It is hoped that, by publishing the details behind the filter design used in this work, others can extend the implementation to eventually meet, and potentially surpass, the orbit accuracies achieved by these closed-source filters.

1.4. Report outline

This report is structured as follows. In Chapter 2, the scope of this work is sketched, combined with the associated research gaps and the subsequent research question and objective. To achieve this

research objective, a case study was performed of orbit determination of Sentinel-6A. The methodology applied in this case study is described in Chapter 3: this includes the choice of orbital platform, variational equations, and performed analyses.

For this case study, four different orbit determination filters were implemented. These depend in their implementation on an underlying observation model, trajectory model, variational model, and state estimator, the designs of which are described in Chapters 4 to 7. In Chapters 4 and 5, the reader can find a description and justification of the chosen observation and trajectory models, respectively, used to model the observables and their temporal correlations. This is followed in Chapter 6 by a description of the different approaches used to solve for the matizant. The design of the Square Root Information Filter, used in this work to reconcile observations with the observation and trajectory models, is given in Chapter 7.

Afterwards, the results obtained from applying the described filter models in orbit determination of Sentinel-6A are presented in Chapter 8. These results imply certain answers to the research questions posed earlier; discussion of these implications can be found in Chapter 9. Finally, this report is closed off in Chapter 10 with a summary of the conclusions drawn from this work, as well as recommendations for the directions to take when building upon this research.

2

Research Objectives

In this chapter, the objectives of the presented research are described and justified. Firstly, in Section 2.1, the reader can find a delineation of the scope of applications and methods considered in this work. Within this scope, a research gap is identified in terms of analytic state transition matrix approximations when the estimated state includes not just the orbital state, but also dynamic parameters. A further description of this gap is found in Section 2.2. Indeed, this knowledge gap motivates the principal research objective of this work, namely to establish the effectiveness of the Hill-Clohessy-Wiltshire equations in approximating the state transition matrices of altimetric, near-circular, low Earth orbits. To fulfil this objective, a research question is formulated and split into four subquestions: by answering each one of these subquestions, this work aims to characterize the impact of each assumption on the quality of the resulting orbit estimates. More detail on this research question and its constituent subquestions can be found in Section 2.3.

2.1. Scope

This research contribution focuses on the application of analytic approximations to the state transition and sensitivity matrices to the real-time orbit determination of satellites. The computational simplification afforded by such solutions to simplified sets of variational equations is useful in a wide variety of mission scenarios: naturally, a reduction in computational effort is always desirable. In this work, we will focus on the niche of altimetric satellites in Low Earth Orbit (LEO). This choice is twofold: in the first place, such altimetric orbits have convenient properties, most importantly their circularity, that make them particularly attractive for analytic solution. In the second place, it is specifically these altimetric, or similarly SAR and RO, missions, that benefit most from lower-latency, high-accuracy products.

Keeping in mind that these points of application would, for the lowest latency, be situated onboard the observing spacecraft, the analysis in this work is limited to only data and algorithms that would be suitable for an onboard environment. This is also necessary from a pragmatic point of view: implementation of the more complex analyses of, for example, the spacecraft antenna phase patterns, could not feasibly be executed within the timeframe allotted for this thesis. Indeed, this also means that we do not consider hard real-time constraints in the software presented in this work, and that the software implementation is not restricted to embedded computational resources. Such considerations are outside the author's area of expertise, and are not necessary for the derived conclusions to be valid. It is emphasized that this is in line with the primary use case for higher-accuracy real-time orbit solutions, which is faster availability of near-real-time products to end users rather than platform support. Here, "near-real-time" refers to within three hours after data acquisition, the definition used for Sentinel-6A science data [11].

Also driven by the necessity to maintain a manageable scope, this research will be confined to encompass solely GNSS pseudorange and carrier-phase observations. The preference for GNSS observations over alternative options is rooted in several factors: its widespread integration in contemporary spacecraft [56], low latency, high availability, and cost-effectiveness compared to similar systems such as SLR and DORIS. Combined with the recent improvement in SISRE qualities for multiple constellations, these attributes position GNSS as a suitable choice for this analysis. This decision does not preclude the applicability of the results to other measurement techniques. Particularly, the outcomes

hold direct relevance to DORIS, which has also demonstrated the capacity to attain high-accuracy real-time orbits [29].

Lastly, in terms of state estimators, the analyses in this work are limited to the effectiveness of state transition matrices in the context of filters, specifically formulations of the EKF. Where batch-based solvers, which can also benefit from improved transition matrix approximations, tend to yield improved accuracy, that comes at the cost of higher computational effort and the inherent latency associated with batch-based processing. This makes such state estimators unsuitable for the use cases under consideration. In the same vein, alternative filters like the Unscented Kalman Filter (UKF), despite being superior in the sense of a higher-accuracy posterior mean and covariance estimate for nonlinear dynamical systems [58], are excluded due to their higher computational demands. On the timescales considered in this work, with tracking data available every 10 s to a minute, the orbital dynamics are sufficiently linear to be estimable with an EKF [5, 56].

2.2. Knowledge gaps

The process of solving the variational equations can be computationally demanding due to that fact that they are of quadratically higher dimension than the equations of motion. Additionally, it is worth noting that the significance of its solution lies not directly in the solved-for values but rather, indirectly, in their use in obtaining a suitable approximation of the error covariance matrix and in establishing a proper differential correction. Consequently, the required accuracy for the state transition matrix is lower than for the trajectory estimate itself. This suggests that it is reasonable to apply simplifications to the variational equations in search of approximations to the state transition matrix.

In fact, early work on orbit determination by Rice [51] already showed that such approximations can be used to great effect in reducing the computational effort associated with orbit determination. Soon after, it was shown by Broucke [8] and Broucke and Cefola [9] that a fully-analytic orbital state transition matrix, without singularities, exists for unperturbed Keplerian orbits. Later work by Der and Danchick [16] and Der [15] built on top of this to establish the effectiveness of different analytical and semi-analytical solutions for the state transition matrix. They showed that such simplifications could greatly reduce computational effort without compromising orbit determination performance. Still, these works are all characterized by a common limitation: they only approximate the state transition matrix for the *orbital* state — state transition matrices for more intricate state definitions, such as those including a radiation pressure coefficient, an aerodynamic drag coefficient, or empirical accelerations, were not derived.

It follows that recent works on precise onboard orbit determination cannot directly apply these prior findings: the degree of accuracy achievable with modern GNSS and DORIS observables necessitates application of trajectory models of greater precision, which do require estimation of empirical accelerations and radiation pressure and aerodynamic drag coefficients. Indeed, recent works on precise onboard orbit determination adopt an alternative approach in simplifying their variational equations to approximate the state transition matrix. In the precursor work to RTNAV, Montenbruck and Ramos-Bosch simplify the variational equations by neglecting Lunisolar perturbations and all Earth spherical harmonics above the second order zonal term. The resulting simplified set of variational equations is solved through numerical integration [39]. In RTNAV, a simplified model is mentioned to be applied, but no note is made of the specific simplifications introduced [38]; presumably, the approach is similar. Similarly, no documentation could be found on the choices made with respect to these simplifications, if any, in DIODE and the Haiyang-2D filter.

This kind of simplification of the variational equations is flexible, supporting practically any choice of state with any choice simplifications. Yet, it does not remove the need for numerical integration, and the corresponding overhead, before one can obtain a solution to the variational equations. A natural research gap arises in analytical approximations to the state transition matrix of satellite orbits, which can be significantly cheaper to compute at near-equivalent accuracy. In this research, we shall investigate this family of approximations by a case study of one such analytical solution, based on the Hill-Clohessy-Wiltshire equations of relative orbital motion in near-circular orbits.

2.3. Research question and subquestions

To fill the knowledge gap sketched above, we aim to establish the sufficiency of the Hill-Clohessy-Wiltshire solutions in modelling the variational equations of the class of altimetric, near-circular orbits

such as the one followed by Sentinel-6A. Additionally, to the degree that this approximation is insufficient, the goal is to assess the root assumptions that introduce the relevant inaccuracies. Ultimately, the objective is to apply this knowledge to allow reduction of the computational effort needed to estimate the resulting orbit. In particular, this would be applicable within the context of real-time orbit determination, as needed for the generation of near-real-time geodetic products.

From this objective, combined with restrictions on the research scope described earlier, a suitable research question follows to be:

“To what extent can simplified, approximate solutions to the variational equations be used to improve the computational efficiency of real-time filtering low-Earth orbit determination with GNSS observations for the purposes of near-real-time altimetry?”

It is noted that this research question explicitly does not cover the Hill-Clohessy-Wiltshire solutions: in the first place, the goal of this research is to assess what the impact is of each considered simplification of the variational equations on the effectiveness of the resulting matrix. Only from this does the conclusion flow as to what degree the additional assumptions imposed by the Hill-Clohessy-Wiltshire solutions are valid, and whether that is acceptable in the context of real-time altimetric orbits such as that of Sentinel-6A.

This overarching research question can be split into four subquestions:

1. What is the required accuracy of state transition matrix estimates to achieve the desired accuracy in the trajectory solution?
2. What is the contribution of small non-circularities in an orbit to its state transition matrix?
3. How significant is the contribution of the Earth spherical harmonic C_{20} term to the variational equations?
4. To what extent do the remaining perturbing force terms contribute to the variational equations?

Each of these four questions describes one core aspect of the full research question that must be answered to sufficiently establish the adequacy of the Hill-Clohessy-Wiltshire solution or related approximations. A brief justification of each follows below.

The first subquestion considered is the required accuracy in solutions to the variational equations. This is important to answer, because it delineates the errors that are permissible in each component of the state transition matrix. It establishes which types of simplifications and in which state components are permissible. This is crucial because it helps create measures by which we can contextualize the answers to the other three questions. After all, the aim of this research is not to determine the influence of the given approximations on the accuracy of the state transition matrix itself, but on the accuracy of the resulting orbit determination results.

The subsequent three questions construct a framework that distinguishes between the validity of different approximations of the matrix. The first of these subquestions concerns the contribution of small non-circularities to the validity of the Hill-Clohessy-Wiltshire approximations. This can be assessed by comparison with more general methods that do support arbitrary elliptical orbits, albeit at computational expense. Because the Hill-Clohessy-Wiltshire method depends in its derivation inherently on near-circularity of the orbit, this is an important error to quantify. This is necessary in the first place to ensure that the method can be validly applied. Secondly, quantification of this error is useful to indicate whether it is or will become relevant at near-future orbital accuracies: this could indicate a good area for future research.

The contribution of the C_{20} Earth spherical harmonic term is relevant when assessing the difference in quality of the Keplerian solutions versus an oblate Earth model. The C_{20} spherical harmonic term is identified as topic of interest based on its inclusion in the state transition matrix approximation used by Montenbruck and Ramos-Bosch [39]; this is the only work on real-time filtering orbit determination that makes mention of the specific simplification applied in estimation of the matrix. The results of similar work by Der and Danchick on the *orbital* state transition matrix suggest that a transition matrix with inclusion of the second zonal harmonic term is indeed significantly more accurate [16]. However, present state transition matrix solutions taking into account C_{20} , like Vinti's method or numerical integration using oblate Earth gravity, are computationally significantly more expensive. Establishing the magnitude of this difference would help identify whether the Hill-Clohessy-Wiltshire model would indeed benefit from extension to oblate Earth gravity.

Ultimately, the last subquestion must be answered to determine the necessity for numerical integration. Given that numerical integration of the variational equations is nearly always viable, it is important to ascertain at which point perturbing accelerations beyond the effect of C_{20} become significant. Addressing this subquestion would provide bounds on the cut-off point where an end user must switch to the computationally expensive “last resort” of numerical integration when solving the variational equations.

3

Methodology

In this chapter, we will lay out the high-level approach used to answer the research question and subquestions that were formulated in the previous chapter. We start with a description of the orbital platform and associated dataset used in this evaluation, given in Section 3.1. Afterwards, in Section 3.3, we provide an overview of the processing needed and measures used to qualitatively and quantitatively compare the accuracy of the different solutions to the variational equations. In Section 3.2, we describe and justify the four approaches to solution of the variational equations that we compare in this work. Finally, in Section 3.4, we describe the architecture of the program used to generate our results.

3.1. Orbital context

We shall study the four different variational models through a case study of their performance in onboard orbit determination of Sentinel-6A “Michael Freilich”. This choice of platform is based primarily on the necessity of a representative, near-circular low-Earth orbit, as is the case for Sentinel-6A. Additionally, a sufficient GNSS measurements is necessary for reliable and accurate orbit determination. With its two Precise Orbit Determination Receiver (PODRIX) receivers, which provide dual-frequency pseudorange and carrier-phase observations of up to 18 Galileo or GPS satellites [43], Sentinel-6A more than meets this requirement.

The PODRIX receiver code and carrier-phase observations are available in Receiver Independent Exchange Format (RINEX) format. These observations consist of post-processed data rather than the raw ranging signal transit time and oscillator phase measurements obtained from the PODRIX receiver. Instead, the data is provided as uncombined pseudorange and carrier-phase measurements; the relevant phase shifts and temperature-dependent code biases have already been applied. Beyond the GNSS observations themselves, the satellite’s operational attitude quaternions, manoeuvre history, mass history and centre-of-mass history are also available to TU Delft. Pre-operational estimates of the satellite’s physical characteristics are publicly available in the Sentinel-6A Precise Orbit Determination (POD) context [25]; we apply known in-flight deviations from those characteristics based upon the results presented in [43].

The orbit determination filter will in the first place make use of broadcast GPS and Galileo ephemerides from the navigation messages of each constellation to estimate the GNSS satellite positions. The reasoning to use broadcast GPS and Galileo ephemerides instead of a posteriori precise ephemerides is twofold: in the first place, their quality is more representative of the data available for real-time orbit determination purposes. Secondly, from a more pragmatic point of view, it is noted that their implementation is more straightforward than the application of precise ephemerides, which are generally provided in terms of satellite centre of mass rather than antenna phase centre positions. For GPS satellites, alternative precise ephemerides in terms of the antenna phase centres *are* available, provided by the National Geospatial-Intelligence Agency (NGA) [45]. For Galileo, to the author’s knowledge, no such post-processed source is readily available. As verification step, the orbit determination process will be repeated using these precise GPS antenna ephemerides to provide an indication of the deterioration in filter performance attributable to imperfect real-time GNSS position knowledge: this filter will be referred to as the “precise GPS filter”, in contrast with the “broadcast GNSS filter”, which is the primary

investigation subject.

While the PODRIX receiver onboard Sentinel-6A does in principle support output of the received GNSS navigation messages, this is not fully enabled on the satellite [38]. As a result, the exact set of navigation messages as received by the spacecraft are not available for processing. Instead, we use the collection of all daily broadcast GNSS messages provided by Deutsches Zentrum für Luft- und Raumfahrt (DLR) as part of its Multi-GNSS Pilot Project (MGEX) BRDX product [40]. Barring reception errors and missing transmissions, this is a representative substitute of the navigation messages that are available onboard the satellite.

This analysis covers a period of one day from 2022-01-01 00:00:00 to 2022-01-01 23:59:59 GPS time. This is the first day on which GPS broadcast ephemerides are publicly available in version 4.00 of the RINEX format [40], which is required for the Timing Group Delay (TGD) and Inter-Signal Correction (ISC) parameters broadcast in the GPS GPS Civil Navigation Message (CNAV) message utilized in our analysis. During this timeframe, the Galileo Open Service performance satisfied all minimum performance level targets [21], as did GPS [2]. Additionally, no orbital manoeuvres were performed by Sentinel-6A. This time span of one day is chosen to be long enough to contain multiple once-per-orbit events, such as Earth eclipses, while still being computationally tractable.

As point of truth on Sentinel-6A's orbit, we use the TU Delft precise orbit solution [10, 54]. This is an ambiguity-fixed solution generated with GipsyX, with root-mean-square observation residuals of 5.10 mm and 0.14 m for the carrier-phase and code measurements, respectively. The solution provides position and velocity estimates at a ten-second interval, and clock offset estimates every five minutes. In-between those times, we interpolate using a seventh-order Lagrangian interpolant. To prevent Runge's phenomenon, the formation of oscillations at the edges of interpolation, use is made of three hours of orbit solution before and after the start of the day.

3.2. Choice of variational equations

We compare four distinct approaches for solving the variational equations, each based in its own set of assumptions on the orbital dynamics of the system. Three of these methods obtain the resulting matrix through numerical integration; the other permits analytical solution. We shall indicate these approaches, and their accompanying set of assumptions, as the reference, Hill-Clohessy-Wiltshire, oblate Earth, and spherical Earth models. These assumptions are given below; for details on the resulting procedures for solution, the reader is referred to Chapter 6. We stress that each of these models differ only in the dynamical model used to approximate the *state transition matrix*; for propagation of the state, the same force model is used in all cases.

Reference model

As baseline, we utilize a trajectory model that incorporates all accelerations that are relevant for propagation of the orbital state over a period of two hours to centimetre-level accuracy. The full derivation and description of the relevant accelerations in the case of Sentinel-6A is found in Chapter 5. It can safely be assumed that any unmodelled acceleration is also irrelevant for solutions to the variational equations: inherently, any acceleration term that is not relevant in propagation of the orbit itself will be even less relevant when propagating *changes* to that orbit.

For this reference model, we only assume:

1. All modelled accelerations are continuously differentiable with respect to the orbital state and dynamic parameters.

Under purely this assumption, it is possible to express the state transition matrix as the solution to an ordinary differential equation that depends only on time and the partial derivatives of the acceleration with respect to the system state. These equations are called the variational equations: their formulation is derived in Chapter 6. This set of variational equations is quite amenable to solution through numerical integration over time. Where this model is the most accurate, it is also computationally the most expensive of the considered approaches.

Hill-Clohessy-Wiltshire model

The principal simplification of the variational equations to be assessed is based on the following set of assumptions on a reference orbit and its perturbed equivalent:

1. The reference orbit is sufficiently circular.
2. The relative change in position is of sufficiently small magnitude with respect to the orbital radius of the reference orbit.
3. Solar radiation pressure, aerodynamic drag, and empirical accelerations are of constant magnitude and direction in the Hill frame over one propagation period.
4. Perturbations due to relativistic, tidal, and third-body effects, and spherical harmonic Earth gravity can be neglected.

Under these restrictions, it is possible to derive closed-form solutions to the equations of relative motion of the two orbits. These solutions are described in detail in Section 6.2. We note that the latter two assumptions are not inherent to the Hill-Clohessy-Wiltshire method itself, but only to the specific implementation presented in this work, as needed to keep the scope feasible. In future work, it would certainly be possible to relax them; in essence, the true underlying assumption in the Hill-Clohessy-Wiltshire derivation is only that these perturbing accelerations are sufficiently small and independent of the orbital state.

As a consequence of these solutions, fully analytic approximations can be obtained for the state transition matrix of a satellite. This approach requires significantly less computational effort than the reference model. On the other hand, due to the additional assumptions on the dynamical model, this approximate solution will also necessarily be less accurate. By comparison of these two methods, this work aims to establish that this penalty in accuracy is acceptable for the purposes of (near-)real-time orbit determination.

Oblate Earth model

The second simplification with respect to the baseline model that we investigate is an oblate Earth model, which assumes:

1. Earth spherical harmonics above degree 2 and order 0 are neglected.
2. Relativistic effects are negligible.
3. Tidal variations in Earth's gravity can be ignored.
4. Third-body perturbations are insignificant.

Not neglected are solar radiation pressure, Earth gravitational harmonics up to degree 2 and order 0, and empirical accelerations. While this set of assumptions again requires numerical solution of the variational equations, the computational effort associated with evaluation of the resulting accelerations is significantly reduced in comparison with the reference model.

This approach follows work by Montenbruck and Ramos-Bosch [39] on precise onboard orbit determination: using the same model, they achieved decimetre-level orbits. Follow-up work by, in part, the same authors in [38] used a simplified, although unspecified, set of variational equations to establish the system's state transition matrix. Presumably, the same or a similar simplification was used.

Inclusion of this oblate Earth model in our comparison allows the obtained results to be validated against the expected orbit determination performance achieved in literature. Additionally, it is more widely applicable than the Hill-Clohessy-Wiltshire model due to its independence of the eccentricity of the orbit. Comparing its physical accuracy and computational performance with that of the Hill-Clohessy-Wiltshire model provides valuable information on the applicability of both models to different classes of orbits, computational environments, and accuracy requirements. Finally, inclusion of this model is useful because comparison with the spherical Earth and reference models allow us to quantify the relative effect of the C_{20} gravitational harmonic, as needed to answer research subquestions 3 and 4, respectively.

Spherical Earth model

In a further simplification of the variational equations, we will also consider a spherical Earth model. This expands on the assumptions introduced in the oblate Earth model:

1. Earth gravity is a point-mass attraction.
2. Relativistic effects are negligible.
3. Tidal variations in Earth's gravity can be ignored.

4. Third-body perturbations are insignificant.

In other words, only point-mass Earth gravity, Solar radiation pressure, aerodynamic drag, and empirical accelerations are modelled.

In terms of comparative accuracy or performance, this model is itself not strictly interesting: it is only marginally more efficient than the oblate Earth model in terms of computational resources, while introducing a relatively significant error through its neglect of C_{20} . This model is included primarily as intermediary between the Hill-Clohessy-Wiltshire and oblate Earth models. It allows separation of the reduction in accuracy caused by neglect of C_{20} from the pessimization induced by the assumption of negligible eccentricity introduced in the Hill-Clohessy-Wiltshire model.

3.3. Analyses

We assess the impact of simplifications in the variational equations on orbit determination quality from five different points of view. The primary measures, following the research objective of improvement of the computational efficiency for precise real-time orbits, are naturally the accuracy of and difference between the resulting orbit solutions, as well as the runtime necessary to compute them. As secondary measures, to better quantify the precise sources of the resulting filter performance differences, we also assess the filter's error covariance estimates as well as a measure that we will refer to as the "differential correction error". These measures correspond to the two places in filtering orbit determination where the state transition matrix is used: propagation of the state error covariance and estimation of the observation Jacobian, for the error covariance ellipsoids and the differential correction error, respectively.

Runtimes

Referring back to the fundamental objective of this thesis, the reduction in computational effort is a key goal of the simplified variational equations. To quantify this reduction in computational resources, we assess the runtimes that are necessary for orbit determination of Sentinel-6A for 1 January 2022 using each of the approaches to estimation of the state transition matrix. Necessarily, these runtime measurements are not meant to be indicative of actual runtime on a representative microcontroller as would be used onboard a spacecraft. Rather, the resulting distribution of runtimes serves as indication of the relative decrease in computational effort that can be achieved by solving a simplified set of variational equations.

This runtime is measured for single-threaded execution on a machine with two Intel Xeon CPU E5-2683 v3 processors, each clocked at 2.00 GHz. The benchmarked executables were compiled using Clang 13.0.1 in a release configuration, and run under OpenSUSE Leap 15.4. Finally, to ensure that none of the computations are removed by the compiler during optimization, all results are stored on the filesystem *after* the measurements have been finished.

Filter performance

In the context of orbit determination, the deciding measure on the acceptability of an approximate solution to the variational equations will naturally be its effect on the resulting orbit solution. In the first place, we will therefore compare the accuracy of all four orbit solutions. This will serve to show what the practical impact of each approximation is. Secondly, it serves to show the degree of accuracy that is applicable in the later analyses.

Beyond the accuracy of the orbit solution itself, it is also necessary to establish the nature of the remaining errors, by means of comparison of the filters' a priori and a posteriori residuals. This is useful in verification of the filter implementation, by checking that these residuals have the expected distributions, but more importantly, to determine whether the filters are more kinematic or dynamic in nature. Such analysis is necessary to establish whether the root cause in differences between filter solutions is to be sought primarily in the error covariance propagation or in the observation inclusion steps of each filter.

Inter-solution differences

Beyond comparison of the orbit solutions with respect to the precise orbit, to determine their accuracy, we also contrast the difference of all three simplifications with respect to the reference orbit as generated using the reference model. This complements the filter performance analysis by allowing a clearer

view of the concrete differences in dynamics, separate from the error in the solution that is introduced by shared sources such as the observation or trajectory propagation model. Where the overall filter performance establishes to what degree simplifications in the variational equations affect the final estimate, this analysis serves to indicate the fundamental cause of these differences.

Error covariance ellipsoids

The first of two places where solutions to the variational equations are of importance, is the propagation of the error covariance ellipsoids. Here, an orbit determination filter approximates, based on its error covariance for a current state estimate, what the uncertainty in its state estimate will be at a later point along its orbit. The filter uses this to quantify the importance to place on its a priori state estimate when incorporating a new batch of observations.

We can characterize the error in the position and velocity covariance with respect to the reference model in three dimensions through the eigenvalues $\lambda_1, \lambda_2, \lambda_3$ and eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ of their 3×3 covariance matrices. Using these definitions, we can speak of the error ellipsoid volume as

$$V = \frac{4}{3} \pi \lambda_1 \lambda_2 \lambda_3, \quad (3.1)$$

following the definition used in [15, 16]. This is a measure of the total estimated uncertainty in the position or velocity estimate. As such, the difference in error ellipsoid volume can be used to quantify the relative over- or underestimation of the uncertainty of a given state estimate.

Similarly, we will define the “orientation” of a position or velocity error ellipsoid as the orientation of the eigenvector corresponding to the eigenvalue with the largest magnitude, which we shall denote with index 1. This allows us to speak of the relative orientation θ of two error ellipsoids with eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ and $\mathbf{v}'_1, \mathbf{v}'_2, \mathbf{v}'_3$:

$$\theta = \arccos(\mathbf{v}_1 \cdot \mathbf{v}'_1). \quad (3.2)$$

By comparing the orientation of the reference model ellipsoids with those obtained from the simplified variational models, we can summarize the ability of an approximate solution to the variational equations to properly capture changes in orientation of the uncertainty.

Differential correction error

The second aspect of filtering orbit determination making use of solutions to the variational equations is the inclusion of new observations to obtain an a posteriori state estimate. Here, the state transition and sensitivity matrices are used by a batch orbit determination filter to determine the change in observation residuals that is achieved by a change in the initial state estimate. To quantify the error introduced by approximations $\hat{\Phi}$ to the state transition matrix Φ in this step of the filtering process, we compute the “differential correction error” of each approximation along the precise Sentinel-6A orbit.

With “differential correction error”, we refer to the norm of the matrix

$$M(\Phi, \hat{\Phi}, t, t_0) = I - \Phi(t, t_0) \hat{\Phi}(t, t_0)^{-1}. \quad (3.3)$$

Conceptually, this matrix describes the residual state error after correction of an initial state error $\Delta \mathbf{y}$ using a differential correction $\Delta \mathbf{y}_0 = \hat{\Phi}^{-1} \Delta \mathbf{y}$ to an initial state [36]. We justify this choice of metric by noting that the measurement update step of a Kalman filter is analogous to application of a differential correction that minimizes the residuals among the a priori estimates and the newly-introduced measurements.

For this analysis, only the orbital state $\mathbf{y}$ and dynamic parameters $\mathbf{p}$ are directly affected by simplifications of the system’s variational equations, so we choose

$$\Phi(t, t_0) = \begin{bmatrix} \Phi_{\mathbf{y}}(t, t_0) & \Phi_{\mathbf{p}}(t, t_0) \\ 0 & I \end{bmatrix}$$

as state transition matrix. As matrix norm, Montenbruck and Gill suggest the maximum-magnitude element of M . However, this norm requires that all elements of the transition matrix are of the same magnitude, requiring explicit normalization of the matrix elements to achieve physically meaningful results. Instead, we propose the largest singular value of M , which is independent of the state transition matrix units, and is physically more meaningful: it describes the maximum relative magnitude, as

a fraction of the magnitude of $[\Delta y, \Delta p]^T$, that remains after application of the differential correction $[\Delta y_0, \Delta p_0]^T$. For a full derivation of the differential correction error measure, the reader is referred to Appendix A.

3.4. Software architecture

For the processing of Sentinel-6A GNSS observation data as necessary for this comparison, an orbit determination filter was written to estimate a trajectory that is representative of real-time performance. This software makes use of Tudat to implement the underlying dynamical models [33]. Additionally, use is made directly of the SPICE [1] and SOFA [3] interfaces provided by Tudat: as these are only exposed in the C++ implementation of the toolkit, the orbit determination filter subroutines were also written in C++. This choice of language has the additional advantage that more fine-grained control could be exercised upon the memory allocation behaviour of the program to significantly reduce its runtime.

The source code is available in a Git repository, to which access can be provided by the author on request¹. An overview of the most important source files for the filter implementation is given in Figure 3.1, along with an indication of the filter component they implement. This figure references the C++ header (.h) containing the code interfaces; the full implementation for each file can, per C++ convention, be found in the identically-named source .cpp file. We stress that this figure is not exhaustive: for more information on the implementation details, the reader is encouraged to investigate the source code itself. Note that this is not necessary to understand the rest of this work: all algorithmic information and full descriptions of the model semantics are included in this report itself.

Three primary executables exist to run the different analyses presented in this work: for the force model comparison detailed in Chapter 5, `trajectoryModelComparison` is run; estimation of the Sentinel-6A orbit using broadcast GNSS ephemerides for each of the four state transition matrices is done using `broadcastGnssFilterComparison`; and `preciseGpsFilterComparison` does the same, but applies precise GPS ephemerides instead. A secondary set of executables is described in the .cpp source files found in the `verification/` folder. As the name suggests, these executables are applied in verification of certain aspects of the implementation. The relevant results from these verification steps will be highlighted throughout this report.

¹Reachable after graduation at quintenvanwoerkom@gmail.com




Figure 3.1: Schematic representation of the filter software architecture

4

Observation Model

In this chapter, the reader can find a description of the observation model used to incorporate the measurements from the PODRIX receiver into the filter's orbit estimate. We start with a description of the used observation equations in Section 4.1. The procedure used to estimate the GNSS satellite ephemerides is given in Section 4.2. The resulting GNSS satellite positions are given in the World Geodetic System 1984 (WGS84) and Galileo Terrestrial Reference Frame (GTRF) frames. To express them in the inertial frame used for propagation, we apply a set of rotations, given in Section 4.3. Finally, we analyse the noise characteristics of the resulting measurements in Section 4.4. These properties are used to model the observations in the Square Root Information Filter (SRIF).

4.1. Observation equations

The Sentinel-6A PODRIX receiver, after post-processing, provides dual-frequency GNSS pseudorange and carrier-phase measurements for both Galileo and GPS satellites. Each is received on the E1 and E5a frequencies for Galileo, and for L1 C/A and L2C or L1 P(Y) and L2 P(Y) for GPS satellites. To compensate for the first-order ionospheric delay, we combine these measurements into the ionosphere-free pseudorange and carrier-phase [56], for which we use the observation models described below. In principle, we follow a similar accuracy requirement as for the trajectory model, including all terms that are significant on the level of more than a centimetre.

Pseudorange

Pseudorange observations, also called code observations, can be derived from the PODRIX output every 10 seconds. They are a measurement of the apparent signal travel time from the transmitting GNSS navigation satellite to the receiver. We model the ionosphere-free pseudorange combination as

$$p = |\mathbf{r}_{rcv}(t) - \mathbf{r}_{trm}(t - \tau)| + c (\delta t_{rcv}^{tot} - \delta t_{trm}^{tot} + \delta t_{stc}^{rel} + \delta t_{ISB} + \delta t_{DCB}) + s + \epsilon_p. \quad (4.1)$$

Here, we distinguish the following terms:

- $\mathbf{r}_{rcv}(t)$ is the receiver antenna position at signal reception in metres, in an inertial reference frame, computed based on the estimated Sentinel-6A position.
- $\mathbf{r}_{trm}(t - \tau)$ is the transmitting antenna position in metres, in the same inertial reference frame, computed based on broadcast or precise GNSS ephemerides.
- c is the speed of light, in metres per second.
- τ is the signal travel time in seconds, computed based on the receiver and transmitter positions.
- δt_{rcv}^{tot} is the total receiver clock offset in seconds, estimated as part of the filter state.
- δt_{trm}^{tot} is the total transmitter clock offset in seconds, computed based on the broadcast GNSS navigation message.
- δt_{stc}^{rel} is the relativistic signal delay, in seconds, due to the Shapiro effect, computed based on the receiver and transmitter positions.

- δt_{ISB} is the inter-system clock bias, in seconds, set to 0 s for GPS observations and part of the estimated state for Galileo observations.
- δt_{DCB} is the signal's differential code bias, in seconds, computed based on the broadcast GNSS navigation messages.
- s is the GNSS-satellite-specific SISRE, in metres, caused by errors in the ephemeris used for the GNSS transmitter positions and clocks, estimated as part of the filter state.
- ϵ_p is the observation error, in metres, due to unmodelled factors such as multipath, receiver noise, and transmitter noise, an unmodelled random variable.

Temperature-dependent receiver code biases have already been applied to the pseudoranges as used in the filter. Concerning the differential code biases δt_{DCB} , those vanish for GPS L1 P(Y) and L2 P(Y) signals with the GPS Legacy Navigation Message (LNAV). In the case of GPS L1 C/A with L2C measurements, this is not the case, and we compute the differential code bias in seconds as [24]

$$\delta t_{DCB} = \frac{1}{1 - \gamma_{12}} (\delta t_{ISC,L2C} - \gamma_{12} \delta t_{ISC,L1C/A}) - \delta t_{TGD}. \quad (4.2)$$

Here, δt_{ISC} are the inter-signal corrections and δt_{TGD} is the time group delay, all distributed as part of the GPS CNAV message. The factor γ_{12} is a constant dependent on the L1 and L2 frequencies $\gamma_{12} = (77/60)^2$. For the Galileo ionosphere-free E1 and E5a signal in combination with the Galileo FNAV Navigation Message (FNAV), the differential code bias also vanishes, although it is necessary to account for the time offset between the GPS and Galileo constellations. We do this through δt_{ISB} , the inter-system bias, which we set to 0 s for GPS signals. For Galileo signals, we include this as an estimated parameter in our filter.

Carrier-phase

A second type of ranging measurement available onboard Sentinel-6A is carrier-phase. Similarly to the pseudorange, the carrier-phase represents the range to a given transmitter. It differs from the pseudorange in the way that this range is tracked: where the pseudorange measurements represent the apparent signal travel time, the carrier-phase is obtained by determining the fractional phase shift of the carrier signal since acquisition. These measurements are provided as a number of cycles; for convenience, we shall multiply each signal by the corresponding wavelength in metres to obtain the ionosphere-free carrier-phase measurement in units of metres:

$$\phi = |\mathbf{r}_{rcv}(t) - \mathbf{r}_{trm}(t - \tau)| + c (\delta t_{rcv}^{tot} - \delta t_{trm}^{tot} + \delta t_{stc}^{rel} + \delta t_{ISB}) + s + b + \epsilon_\phi. \quad (4.3)$$

Most terms coincide with those provided in Equation (4.1), due to the similar nature of both measurement types. We do explicitly note that the error term ϵ_ϕ , while similar in nature, is distinct from the pseudorange error term ϵ_p . Another new term is the carrier-phase ambiguity b , expressed in metres. Apart from cycle slips, which we shall disregard, this is a conceptually constant parameter representing the geometric distance between the receiver and transmitter antennae at signal acquisition. This ambiguity parameter is estimated as part of the filter state during orbit determination.

It should be noted that we ignore the effect of carrier-phase wind-up in our measurements. In principle, this effect is significant, as it constitutes a near-constant drift on the order of several centimetres over a single observation arc. However, modelling this term would require a detailed model of each of the GNSS satellite orientations, which can differ significantly between types. Hence, this would require considerable effort to implement. Additionally, as its main impact is a slow drift, the resulting error is largely absorbed by the ambiguity estimate. As a pragmatic trade-off, we therefore disregard this effect in our observation model, even though it does introduce a non-negligible error.

Antenna phase centre

To compute the Sentinel-6A antenna phase centre position, we utilize the relative position of the Sentinel-6A antenna reference point $\delta \mathbf{r}_{rcv}$ with respect to the Spacecraft Reference Frame (SRF), as well as the rotation $R_{ARF \rightarrow SRF}$ from the antenna to the spacecraft reference frame. These are found in the Sentinel-6A POD context [25]. The spacecraft centre of mass position $\mathbf{r}_{com}$ in the SRF, and the SRF orientation with respect to the Geocentric Celestial Reference System (GCRS), $R_{SRF \rightarrow GCRS}$, are extracted from the operational data. This allows us to obtain an estimate of the antenna phase centre as

$$\mathbf{r}_{rcv} = \mathbf{r}_{com} R_{SRF \rightarrow GCRS} (\delta \mathbf{r}_{rcv} - \delta \mathbf{r}_{com} + R_{ARF \rightarrow SRF} \delta \mathbf{r}_{pco}). \quad (4.4)$$

Here, δr_{pco} is the phase centre offset, expressed in the antenna reference frame. Its values are obtained from the in-flight calibration results, which indicate values of 75 mm and 93 mm for GPS and Galileo, respectively [38].

Clock offsets

In estimation of both the transmitter and receiver clock offsets, we split each into a slowly-varying bias term and periodic offsets caused by relativistic effects. For the transmitter clock bias δt_{trm} , we estimate this offset based on the GNSS clock offset polynomials distributed in the GPS LNAV and Galileo FNAV messages. The receiver clock bias δt_{rcv} is determined as part of the filter's state estimation.

Relativistic effects

On top of the clock bias terms for the transmitter and receiver clocks, we also add the relativistic clock offset, in seconds, caused by orbital eccentricity [56]:

$$\delta t_{clk}^{rel} = -\frac{2}{c^2} \mathbf{r} \cdot \mathbf{v}. \quad (4.5)$$

The position $\mathbf{r}$ and velocity $\mathbf{v}$ as used here must be expressed in metres in the GCRS or an equivalent inertial, Earth-centered frame. This correction is largest, in the order of metres, for the GNSS satellites. For the Sentinel-6A clock it is smaller, yet still comprises a consistent bias in the order of a centimetre, so we also include it.

In addition, we offset each clock with the relativistic effect caused by the varying gravity potential induced by Earth's oblateness [56]:

$$\delta t_{clk, J_2}^{rel} = -J_2 \frac{3r_\oplus^2}{2c^2} \sqrt{\frac{\mu}{a^3}} \sin^2 i \sin 2u. \quad (4.6)$$

In this expression: J_2 is the Earth gravitational field expansion term due to its oblateness; $r_\oplus$ is the mean equatorial radius of the Earth, in metres; i is the orbital inclination, in radians; and u is the argument of latitude, also in radians. This effect is on the order of several centimetres for both the Sentinel-6A clock and the GNSS clocks.

The final relativistic effect incorporated in both the pseudorange and carrier-phase models is the Shapiro effect due to space-time curvature. This effect manifests itself as a propagation delay δt_{stc}^{rel} of the signal. This delay, in seconds, is estimated as

$$\delta t_{stc}^{rel} = \frac{2\mu}{c^3} \ln \frac{|\mathbf{r}_{rcv}| + |\mathbf{r}_{trm}| + |\mathbf{r}_{rcv}(t) - \mathbf{r}_{trm}(t - \tau)|}{|\mathbf{r}_{rcv}| + |\mathbf{r}_{trm}| - |\mathbf{r}_{rcv}(t) - \mathbf{r}_{trm}(t - \tau)|}. \quad (4.7)$$

For ranging between Sentinel-6A and GPS or Galileo satellites, this effect results in a delay of 40 ps, or one centimetre ranging error.

Light-time correction

In our application to decimetre-level orbit determination, the distance travelled by a GNSS satellite in the time between signal transmission and reception is significant. Thus, it is necessary to model and correct for the signal travel time, τ . For this work, we use a fixed-point iteration to find the light-time:

$$\tau_{n+1} = \frac{1}{c} |\mathbf{r}_r(t) - \mathbf{r}_t(t - \tau_n)|, \quad \tau_0 = 0. \quad (4.8)$$

This iteration converges in all cases for our dataset to within nanosecond precision in three iterations.

Consistency with observations

To verify the correctness of this observation model, we compare the actual Sentinel-6A GPS observations with the expected values computed using the precise TU Delft Sentinel-6A orbit and receiver clock estimates and the precise NGA GPS antenna phase centre ephemerides [45]. For the pseudoranges, we use the direct estimates; for the carrier-phase observations, we subtract the carrier-phase measurement at acquisition-of-signal from all subsequent measurements in the observation arc to reduce the impact of the ambiguity. By computing the mean of the resulting observation errors at each observation epoch, we obtain a (noisy) estimate of any remaining ranging biases.




Figure 4.1: Mean residual observation error per observation batch and signal type with precise Sentinel-6A and GPS clocks and orbits

The resulting residuals are found in Figure 4.1. For completeness, the L1 and L2 P(Y) signal results are also included, although we do note the presence of relatively more outliers in the L1 and L2 P(Y) combination, particularly for the pseudorange. This is caused by the lower number of measurements in each observation set, rather than an actually higher noise level. Indeed, inspection shows that the outliers, as particularly visible for the pseudorange, all consist of single measurements. As such, we do not take these outliers to be indicative of any biases.

Overall, no significant constant biases are evident. We do note a periodic offset, as particularly visible in the carrier-phase measurements for the L1 C/A and L2C signal, in Figure 4.1b. It recurs with a period of just under two hours, corresponding to one orbital period of Sentinel-6A. This suggests that it is likely an error in the modelling of the receiver clock offset of the signal. Because this offset is shared identically across all observed signals, it is highly observable, and as a result, easily correctable by the filter.

Beyond the biases that are common across all observations, we also consider the spread of carrier-phase errors within all observations at the same epoch. This spread is relatively much smaller than the amplitude of the observed periodic bias: for the L1 C/A with L2C signal, the standard deviation around the mean per observation epoch is 3 cm on average. This is relatively high for carrier-phase measurements, yet in line with expectations: we refer back to the choice to neglect carrier-phase wind-up. At one wavelength per revolution, it is entirely reasonable to expect that this standard deviation is caused primarily by the effects of carrier-phase wind-up over each observation arc. Otherwise, the observation model does indeed seem to be a good approximation of reality.

4.2. GNSS ephemerides

To estimate the orbits and clock offsets of the GNSS satellites, their broadcast ephemerides are used, based on the navigation messages recorded in the DLR MGEX dataset [40]. For both GPS combinations, the LNAV message is used — for the Galileo E1 and E5a combination, FNAV ephemerides are used. In simulation of the availability of messages in an actual real-time or onboard scenario, the contents of each message are used in the orbit determination filter only after the reception timestamp indicated in the MGEX dataset.

To verify the correctness of the procedure used to compute the broadcast ephemerides, a comparison was performed with the NGA precise antenna ephemerides. The results of this comparison, in terms of the distance between the estimated and precise antenna position, as well as the ranging error resulting from the clock error, are given in Figure 4.2. The difference between the broadcast and precise ephemerides is on the order of one metre, which is in agreement with expectations. This also serves as indication of the expected magnitude of the SISRE to design for when tuning the SRIF.

One navigation message was excluded from this analysis: between 2022-01-01:13:00 and 2022-01-01:14:00, the dataset contained an erroneous semi-major axis in the navigation message for satellite G07. It is unclear whether this is an error in the dataset or in the actual broadcast navigation message on this day. For the purposes of this work, it sufficed to remove the outlier manually: of course, in an actual real-time or onboard scenario, this would have to be done automatically.




Figure 4.2: Mean and standard deviation over time of the clock and orbit errors for GPS

4.3. Transformation to the GCRS

Because the GNSS ephemerides are given in each constellation's realization of the International Terrestrial Reference System (ITRS), they are not directly compatible with the observation equations as currently formulated. Indeed, it is necessary to account for the rotation of these reference frames during signal propagation. This can be done either by application of the Sagnac correction or by means of expressing the observation equations, and thus the geometric range, in terms of a non-rotating reference frame. Considering that the position estimate following from the observation equations must also be compatible with the propagated equations of motion, the choice is made to express the observation equation in the same inertial reference frame also used for propagation, the J2000-oriented Earth-centered inertial frame (ECI) frame. This comes down primarily to a rotation from the ITRS to the GCRS: the transformation from the GCRS to the J2000-oriented ECI frame constitutes only a constant rotation. This choice for an inertial formulation is made in spite of the computational overhead introduced by the explicit transformations required to do so: it ensures that Coriolis forces need not be accounted for, which would significantly complicate the trajectory modelling.

To achieve this transformation, the International Earth Rotation and Reference Systems Service (IERS) 2010 Conventions are adopted for rotation to the GCRS. Rapid Earth rotation parameters from




Figure 4.3: Comparison of the Earth orientation rotation matrices with those computed by l'Observatoire de Paris [20]

ITRS Bulletin A are used, representing the accuracy available for real-time operation. All GNSS positions shall be converted to the GCRS via the transformation [49]:

$$\mathbf{r}_{GCRS} = CRW\mathbf{r}_{ITRS}. \quad (4.9)$$

Where velocity conversions are needed, the following *simplified* relation is applied, per [7]:

$$\mathbf{v}_{GCRS} = CR(W\mathbf{v}_{ITRS} + \boldsymbol{\omega}_{\oplus} \times W\mathbf{r}_{ITRS}). \quad (4.10)$$

Both rotations consist of the matrices C , for bias-precession-nutation, R , for proper motion, and W , for polar motion. These are computed using the approximations recommended in [7]. The vector $\boldsymbol{\omega}_{\oplus}$ represents the Earth rotation in the Terrestrial Intermediate Coordinate System (TIRS), and is estimated using the approximation

$$\boldsymbol{\omega}_{\oplus} = \begin{bmatrix} 0 & 0 & \omega_{\oplus} \end{bmatrix}, \text{ where } \omega_{\oplus} = 7.2921151467064 \cdot 10^{-5} \left(1 - \frac{LOD}{86400}\right) \text{ rad/s}. \quad (4.11)$$

Here, LOD is the length-of-day in seconds, as also provided in IERS Bulletin A. This approximation introduces an error into the computed velocity in the inertial frame that could be quite significant if it were used directly to estimate the velocity. Fortunately, the orbit determination filter in this work does not directly make use of this transformation to compute the inertial velocity, so is not affected by this error. Instead, this transformation is used in this work only to transform the precise Sentinel-6A velocity between the inertial and Earth-fixed reference frames when comparing the estimated and precise orbits: at this point, the resulting error, expected to be on the order of micrometers per second, is acceptable.

The EOPs in IERS Bulletin A are tabulated at 24-hour intervals at midnight Universal Coordinated Time (UTC). Between those times, a seventh-order Lagrangian interpolant is applied to estimate these parameters. Based on the resulting values, the Earth rotation angle, Terrestrial Intermediate Origin (TIO) locator, and Celestial Intermediate Origin (CIO) locator are computed using the International Astronomical Union (IAU) SOFA subroutines, following the IAU 2006 precession and 2000A nutation models [3]. It is noted that ocean tides are explicitly not included in computation of the polar motion; as demonstrated in the verification below, this was not necessary.

The resulting Earth rotation matrices were verified against a set of rotation matrices between 1 January 2022 00:00:00 and 23:59:59, computed using a publicly available online tool provided by l'Observatoire de Paris [20]. The reference matrices are computed from the IERS combined series C04 EOPs, and include variations caused by ocean tides. Intermediate matrices for the separate polar motion, proper motion, and bias-precession-nutation rotations are also available and compared against. The angular error between the resulting transformation matrices is given in Figure 4.3.

As evident, the total error never exceeds 2 mas, which corresponds to a positional difference of about 8 cm at the orbital altitude of Sentinel-6A. While this would be a significant error source if the ITRS-to-GCRS rotation were one-way, that is not the case in this application to orbit determination in (an approximation of) the ITRS. Indeed, the roundtrip transformation is consistent and accurate to far below the centimetre threshold. Any residual errors introduced by this mismodelling thus come from any differences in acceleration that arise from this decimetre-level difference in position: as the relevant accelerations in the case of Sentinel-6A are relatively insensitive to such small changes in position, that is not significant.

4.4. Code noise analysis

To determine a reasonable set of observation weights for inclusion of the pseudorange measurements in the SRIF, a code noise analysis was performed. Of particular interest here was the relation between the code noise and the carrier-to-noise density of a measurement, as the latter is provided by the Sentinel-6A PODRIX for every pseudorange observation.

To establish this relation, the multipath combination of each pseudorange was analysed [56]. For a given pseudorange measurement p_A , in metres, on frequency A, and carrier-phase measurements ϕ_A, ϕ_B on frequencies A and B, also in metres, this is

$$o_A = p_A - \phi_A - 2 \frac{f_B^2}{f_A^2 - f_B^2} (\phi_A - \phi_B). \quad (4.12)$$

This combination can be interpreted as a range, in metres. It is both geometry-free and ionosphere-free; the remaining terms encompass primarily nuisance parameters, such as multipath, receiver noise, and signal biases of all three observations. Assuming that the carrier-phase noises and multipath effects are significantly smaller than those of the code observations, the variation in time of the resulting variable provides an estimate of the variation in time of the pseudorange noise and multipath effects for the code measurement on frequency A. In Figure 4.4, the resulting distribution of multipath combinations versus carrier-to-noise density is given for each of the dual-frequency combinations. To achieve approximately zero-mean estimates, the initial value of this multipath combination at the start of each observation arc has been subtracted.

It is evident from the figures that the magnitude of the multipath combination, representing the expected observation uncertainty for each, varies greatly between signals and carrier-to-noise densities. It follows that state estimation would benefit from a tailored estimate of the standard deviation of each observation rather than an overall approximation. To this end, a fit of the empirically-observed standard deviations of the multipath is applied. This is achieved by collecting all multipath observations by carrier-to-noise density in 21 equally-sized bins of 5 dB-Hz width between the minimum and maximum carrier-to-noise densities of 10 and 60 dB-Hz. This allowed the standard deviation in the multipath combination for each carrier-to-noise density range to be computed. Subsequently, this per-bin estimate is fit using a third-degree polynomial. The resulting polynomial fits can be found in Figure 4.4. They represent the approximate relation between the carrier-to-noise density of pseudorange measurements and their standard deviation used to estimate the standard deviation of pseudorange measurements in the SRIF.

It is evident from these results that the noise standard deviation does increase considerably for smaller carrier-to-noise densities. Hence, it could also be considered useful to simply discard all measurements below a given noise density ratio. This is, consciously, not done in the present filter implementation, as these figures show that the resulting noise is still approximately zero-mean, even for lower signal-to-noise ratios. As a consequence, albeit not much, there is still information to be extracted from these measurements. This means that an orbit determination filter should still benefit from the inclusion of such data, as we appropriately weigh such measurements down with an increased observation standard deviation.




Figure 4.4: Multipath combinations of the utilized GNSS signals, and their associated standard deviation fit

5

Trajectory Model

In this chapter, the reader can find a derivation and justification for the trajectory model used to approximate the Sentinel-6A orbit. The chapter starts with a brief description of the reference frame and time system used in trajectory propagation, in Section 5.1. In Section 5.2, the benchmark model on which the trajectory model design is based is described. By comparison with this benchmark model, we shall first choose an appropriate integrator design in Section 5.3. Afterwards, in Section 5.4, we determine the set of accelerations that are relevant for the accuracy required by our filter implementation. Finally, in Section 5.5, we assess the quality of the resulting trajectory model, in terms of the magnitude of unmodelled accelerations and the ability of the model to compute the precise TU Delft orbit.

5.1. Space-time reference system

For the coordinate system, we follow the default choice of Tudat for a J2000-oriented, ECI reference frame [18]. As origin, we choose the instantaneous Earth centre of mass. Throughout the rest of this report, we shall indicate this reference frame as the ECI J2000 frame. It coincides with the GCRS, minus a constant frame bias of $[-0''.016617, -0''.0068192]$ [3]; this frame bias, while minimal, is applied before propagation to ensure consistency.

For time, used as independent variable in propagation, we also follow the Tudat conventions in using the Barycentric Dynamical Time (TDB) timescale. TDB varies with respect to Terrestrial Time (TT) only with periodic variations, the largest being less than two milliseconds in magnitude [32] — as a result, for the purposes of centimetre-level propagation, we can safely neglect this difference. This allows for a computationally convenient transformation from GPS time to TT, which becomes a constant time offset.

5.2. Benchmark model

As benchmark orbit, we propagate the TU Delft Sentinel-6A precise orbital state forward from 2022-01-01 00:00:00 GPS time forward for 2 hours. The benchmark force model used for this propagation is summarized in Table 5.1. It includes accelerations down to a minuscule significance level and utilizes an integrator step size that was purposefully chosen to be exceedingly low, such that round-off error is dominant over truncation error. This incurs a quite significant computational overhead, making this filter unsuitable for actual application in the orbit determination filter. Indeed, this is *not* the actual trajectory model used inside the final orbit determination filter: rather, this benchmark trajectory model is used to determine an appropriate integrator and appropriate set of accelerations in the final trajectory model.

We design the final trajectory model to be accurate to the level of a centimetre over a period of two hours, or approximately one orbit. This accuracy requirement is based on two considerations. In the first place, the trajectory model must be sufficiently accurate to allow correlation of measurement over longer periods of time, in the order of minutes. A trajectory model of the given accuracy would be more than capable of allowing the measurements, which are accurate to the order of a centimetre, to be correlated. In addition, this accuracy is needed to propagate through longer periods without observations without deterioration of the estimated orbit. This is necessary in particular for DORIS-based filters which, due to the ground-based locations of DORIS beacons, can experience black-outs that can last for large

Table 5.1: Benchmark dynamical model for trajectory propagation

<i>Celestial ephemerides and orientations</i>
Celestial ephemerides from SPICE, tabulated every 10 s with Lagrangian interpolator [1]
<i>Propagation</i>
Runge-Kutta 4 with step size $10 \cdot 2^{-9}$ s Cowell formulation
<i>Gravitational forces</i>
GOCO05c Earth gravity (100×100) [23] k_2 and k_3 Lunisolar dynamic solid Earth tides Point-mass Sun, Moon, Venus, Mars, Jupiter, and Saturn gravity
<i>Surface forces</i>
Cannon-ball aerodynamic drag with NRLMSISE-00 density [50] Cannon-ball Solar radiation pressure with conical shadow model [36]
<i>Relativistic accelerations</i>
De Sitter acceleration Lense-Thirring acceleration Schwarzschild acceleration

portions of an orbit. To allow our conclusions to also be applicable to such filters, a sufficiently accurate trajectory model needs to be analysed that would last through such black-outs without requiring a cold start of the filter.

A notable exclusion in this benchmark acceleration model is the Earth radiation pressure, while it is known to (relatively) significantly affect the Sentinel-6A orbit [43]. However, Tudat does not include the means to model this acceleration to sufficient accuracy: instead, we model it as a constant radial acceleration of 30 nm/s^2 , representing the empirically-observed time-averaged influence [38]. Similarly, the ocean tides are excluded due to their unavailability in Tudat. However, unlike the Earth radiation pressure, there is no simple substitute to compensate for this neglect. Indeed, this is the most significant simplification in the trajectory model. At around one-tenth the effect of solid Earth tides [36], this means that the model will *not* strictly be accurate to the 1 cm mark given earlier. Still, because this perturbation is slowly-varying, at a period on the order of one orbit, it can be compensated quite well using empirical accelerations: hence, it is not expected to limit the filter performance too significantly.

We approximate the physical properties of Sentinel-6A, as needed for the surface forces, based on the physical dimensions provided in the POD context [25]. The frontal area, as used to estimate the aerodynamic drag, is estimated at 2.99 m^2 , the area of its front-facing body surface. This is certainly an under-approximation, as it neglects the sides of the solar panels and any protrusions or slightly-incident surfaces. Similarly, we approximate the Solar radiation pressure reference surface area as 10 m^2 , just below the sum of the solar panel areas. For purposes of this analysis, we take both the aerodynamic drag coefficient c_d and radiation pressure coefficient c_r to be unity: in the final orbit determination filter, these parameters are adjusted to fit the observed dynamics.

5.3. Integrator

For the choice of integrator, we compare four fixed-step and four variable-step integrators by substituting them for the integrator in the benchmark model. The maximum position error encountered over the two-hour propagation period is given in Figure 5.1, plotted against the propagation-time-normalized number of function evaluations used to compute each estimate. This propagation-time-normalized number of function evaluations can be interpreted as the number of function evaluations, on average, that is needed to propagate the state by one second; lower is better. The accuracy requirement of 1




Figure 5.1: Comparison of integrator performance for a representative Sentinel-6A orbit

cm over the two-hour period is indicated by the dashed grey line: any integrator set-up below this line satisfies the accuracy requirement.

It should be noted that all four fixed-step, and the two Runge-Kutta-Fehlberg variants among the variable-step methods, are of relatively low orders of four or lower. This is a result of our application [37]: in the context of real-time orbit determination, we update the state relatively frequently using new observations — the exact interval depends on the batch size, but will be not much more than one minute. In our specific filter design, this is ten seconds: after each inclusion of new measurements, the state is updated and must be propagated again. As a consequence, we are restricted to an upper bound of ten seconds in our step size: this eliminates the computational advantage of higher-order methods.

This constraint is reflected in our integrator comparison: for the variable-step integrators, a maximum step size of 100 seconds was imposed. The chosen tolerances differ per variable-step integrator: they were chosen to cover as large a part of the function evaluations and function error space as possible. For the fixed-step integrators, step sizes were tested from $10 \cdot 2^{-9}$ to 10 s: below this, the number of function evaluations became too large, and above this, the position errors for all fixed-step integrators exceeded a centimetre.

Perhaps counterintuitively, it is evident that both higher-order methods, Adams-Bashforth-Moulton and Dormand-Prince 8(7) are unable to outperform the three fourth-order fixed-step methods. This is in part due to aforementioned time step restrictions, which mean that these methods cannot make optimal use of their variable-step nature to enlarge their step size even further. However, it is also observed that their error behaviour is quite erratic, meaning that they are evidently quite sensitive to the chosen tolerances. Whether this is an issue with the configuration or implementation is unclear: as both perform worse in terms of computational effort for a given accuracy, this is not further investigated. The same goes for both lower-order Runge-Kutta-Fehlberg methods.

All fixed-step methods, as evident from their exponential decrease, can be seen to be dominated by truncation error near the 1 cm accuracy requirement. Only near an error of 10^{-5} m do the fourth-order methods transition to round-off-dominated error, as visible from the transition to a roughly constant position error: for the third-order Runge-Kutta 3 method, this transition point is not reached, and could be located at an even smaller error. This truncation-error-dominant behaviour is desirable, because it means that the error can be estimated quite well for any given step size.

Among the fourth-order fixed-step methods, Runge-Kutta 4 requires the fewest function evaluations




Figure 5.2: Significance of third-body gravitational perturbations

near the accuracy requirement of 1 cm. As such, this is integrator chosen for the final trajectory model. The corresponding step size is chosen to be 1 s, resulting in a positional accuracy well under one millimetre over a two-hour period. This specific step size is convenient not just from a computational point of view, but also because it matches the carrier-phase observation interval of Sentinel-6A's PODRIX receivers exactly. This means that no interpolation is needed to obtain the state estimates needed for each observation. It should be noted that this introduces a slight performance pessimization into the orbit determination process: a time step of 5 s would also be compatible with the given centimetre-level accuracy. Despite introducing interpolation overhead, which is relatively small, the reduction in number of function evaluations would more than likely still mean a reduced runtime. Still, the choice is made to forego interpolation: this keeps the orbit propagation implementation simple, and removes the possibility of interpolation error.

5.4. Modelled accelerations

With a known bound on the integrator-related error, it becomes possible to assess the influence of each dynamic model component on the total model accuracy. We do this by comparison of the benchmark orbit with separately-computed orbits, each with one model factor removed, to assess the magnitude of the removed acceleration. Assuming that the accelerations do not interact non-linearly to create greater perturbations, this gives an indication of the accuracy that is lost by removing a model factor.

To save on computational effort, all orbits for this analysis were computed using a Runge-Kutta 4 integrator with a time step of 0.125 seconds. This corresponds to a maximum position error of about 10^{-6} m. As such, any differences below this threshold are not visible: as our position error requirement is one centimetre, that is not a problem.

Third-body gravitational perturbations

To assess the relevance of gravitational perturbations by third bodies, we consider each of the closest and most massive bodies: the Sun, Moon, Venus, Mars, Jupiter, and Saturn. By virtue of their distance, we can model them as point masses. In Figure 5.2, one can find the error that is introduced by removal of each body's gravitational influence from propagation of Sentinel-6A.

It is evident that only the Sun and Moon's gravitational influences are significant to the centimetre level. Because of the small magnitude of the influence of Venus, we can also conclude that Mercury, which is smaller and farther away, is not of influence. Similarly, it can be deduced from Jupiter's insignificance that Uranus and Neptune, which are both smaller and farther away, need not be accounted for either.

This choice of model, considering only Lunisolar perturbations, matches the trajectory models used in [14, 38], as well as DIODE [29]. All three implementations even go one step further, and use simplified




Figure 5.3: Significance of removal of non-third-body accelerations on simulated Sentinel-6A position

ephemeris models to reduce the computational effort. The choice was made not to do this in our software, considering that our primary goal is not in the first place to produce an implementation that is compatible with embedded computational resources. This allows us to stay with the well-tested SPICE ephemerides implemented by default in Tudat.

Relativistic effects

For our modelling of relativistic effects, we follow the recommendations posited in [49], which indicates that only Schwarzschild acceleration is relevant for our propagation. This is in accordance with the supporting material of the conventions [28], which indicates that the Schwarzschild acceleration is 12 nm/s^2 for the Jason satellite, whereas the Lense-Thirring and De Sitter terms only account for accelerations of 0.1 nm/s^2 and below. As Sentinel-6A, by design, follows a similar orbit to Jason, we can conclude that similar magnitudes apply here. Indeed, as visible in Figure 5.3, we verify that the Lense-Thirring and De Sitter accelerations are not significant.

Surface forces

In terms of surface forces, we verify only the effect of aerodynamic drag and Solar radiation pressure. For Earth radiation pressure, we follow the recommendation of [38] to model it as a constant radial acceleration of 30 nm/s^2 . In Figure 5.3, the effects of aerodynamic drag and Solar radiation pressure are visible, modelled as cannon-ball drag and radiation pressure, respectively. Evidently, both accelerations are significant.

We note that the relative magnitude of both accelerations also indicates that mismodelling of the coefficients in their cannon-ball models can introduce relatively significant errors. Indeed, it is likely that these surface forces, particularly Solar radiation pressure, form the largest source of errors in our dynamical model, necessitating adjustment of both coefficients as state parameters. This modelling is similar to RTNAV, which also applies a cannon-ball model with adjusted coefficients to model Sentinel-6A [38], as does the derived filter presented in [14].

To keep the trajectory model implementation simple, the default high-accuracy NRLMSISE-00 atmospheric model available in Tudat is used also in our final trajectory model [50]. This model has been applied before by Darugna et al. to centimetre-level accuracy orbit determination of Sentinel-6A [14], indicating its applicability to this orbital environment. It is certainly possible that a simplified atmospheric model, such as the static Harris-Priester density model applied by Montenbruck et al. in RTNAV, would achieve similar results. Still, for the purposes of this work, the computational effort needed for NRLMSISE-00 is acceptable, so we do not examine further simplifications.




Figure 5.4: Significance of truncated Earth spherical harmonics on simulated Sentinel-6A position

Lunisolar dynamic solid tides

We evaluate the effect of geopotential variations due to Lunisolar solid Earth tides for both the k_2 and k_3 Love numbers. In this, we follow the IERS Conventions 2010 elastic Earth model [49]. It is clear that the k_2 tides are significant, affecting the orbit by more than a centimetre within five minutes. For the k_3 tides, this is not the case: although it is close to the centimetre threshold, it never exceeds it. In addition, the resulting orbit perturbation is quite slowly-varying, at an apparent frequency of about once per orbit. This kind of slow variation is well-compensated for by empirical accelerations. As such, we choose to neglect this acceleration due to k_3 tides in the final trajectory model.

Earth spherical harmonics

To assess the required degree and order of the modelled Earth spherical harmonics, we compare the performance of pairs of degree and order from 50 down to 0 with the benchmark degree 100, order 100 GOCO05c model. The results of this comparison are found in Figure 5.4. We see here that a degree and order 50 model comfortably fulfils the centimetre accuracy requirement, so that is the model chosen for modelling of the Sentinel-6A orbit. This is in agreement with the model used in RTNAV, which also applies a set of degree and order 50 harmonics. We contrast this with the model used in DIODE, where degree and order 78 harmonics are applied [38].

5.5. Final model

The selected set of accelerations and choice of integrator lead to the final trajectory model summarized in Table 5.2. This combination is, based on the previous analysis, accurate to 1 cm over one orbital period: this should be sufficient to sustain centimetre-level orbital accuracy over observation gaps or longer periods of low-quality observations. In addition, it is sufficiently accurate to allow correlation of measurements that are accumulated over longer periods of time, as needed for reduced-dynamic orbit determination.


To establish a lower bound on the expected magnitude of empirical accelerations to be modelled in our filter, the benchmark orbit is also compared with the derived trajectory model in terms of accelerations. Additionally, this serves as verification of the quality of the trajectory model, to ensure that removal of the total set of accelerations indeed does not result in errors exceeding the centimetre mark over two hours. Next to this, the trajectory model is validated through comparison with the precise TU Delft orbit.

Empirical accelerations

To compensate for the chosen model simplifications, as well as for modelling errors and stochastic perturbations, a set of empirical accelerations are introduced. These are modelled in the filter as

Table 5.2: Final dynamical model used for trajectory propagation in the orbit determination filter

<i>Celestial ephemerides and orientations</i>
Celestial ephemerides from SPICE, tabulated every 10 s with Lagrangian interpolator [1]
<i>Propagation</i>
Runge-Kutta 4 with step size 1 s Cowell formulation
<i>Gravitational forces</i>
GOCO05c Earth gravity (50×50) [23] k_2 Lunisolar dynamic solid Earth tides Point-mass Sun and Moon gravity
<i>Surface forces</i>
Cannon-ball aerodynamic drag with NRLMSISE-00 density [50] Cannon-ball Solar radiation pressure with conical shadow model [36]
<i>Relativistic accelerations</i>
Schwarzschild acceleration
<i>Empirical accelerations</i>
Exponentially-correlated constant radial, along-track, and cross-track accelerations

**Figure 5.5:** Differences in accelerations between the Sentinel-6A trajectory models of Table 5.1 and of Table 5.2




Figure 5.6: Propagation error of precise ephemerides over a propagation time of 1 s, time-averaged over 100 s to reduce noise




Figure 5.7: Propagation error of precise ephemerides over a propagation time of 10 s, time-averaged over 100 s to reduce noise

exponentially-correlated stochastic variables [53] in the radial, along-track, and cross-track orbit directions. To determine the magnitude of these accelerations, and to verify that their behaviour is properly modelled in terms of exponentially-correlated random variables, the relative accelerations due to *unmodelled* accelerations are analysed for the precise Sentinel-6A TU Delft orbit during the first two hours of 1 January 2022, GPS time. Naturally, *mismodelled* accelerations are not captured by this analysis: hence, the resulting relative accelerations should be seen as a lower bound.

The resulting relative accelerations can be found in Figure 5.5. In the first place, we observe that the magnitude of these accelerations is on the order of nm/s^2 . This is in line with the expected relative acceleration for an orbital accuracy of a centimetre over two hours. A second observation is to be made regarding the timescale of the dynamics: as expected, the unmodelled accelerations are clearly correlated in time, varying only over a timescale of minutes. This suggests that it is indeed reasonable to model these accelerations as exponentially-correlated random variables with correlation times of ten minutes, as previously done in RTNAV [38]. For more detail on the resulting empirical acceleration model, the reader is referred to Section 7.2.

Model quality

Finally, to validate the trajectory model, it is compared against the TU Delft precise Sentinel-6A orbit [10, 54]. For this purpose, the precise TU Delft orbit, given in the WGS84 realization of the ITRS, is transformed to the ECI J2000 frame used for propagation. For every second over the first two hours of 1 January 2022 GPS time, we propagate the precise initial state over an arc of one second forward in time. This is done instead of a single propagation over the full two-hour period because we are not

interested in the full model error integrated over time, but in the instantaneous error at each given point along the orbit. Then, the resulting errors provide an indication of the magnitude of both mismodelled and unmodelled accelerations at each instant in time. These position errors are given in Figure 5.6.

In the first place, this shows that the entire trajectory model is accurate, over propagation arcs of one second, to $100 \mu\text{m}$. This is sufficient for usage with the current observation model: as shown in Chapter 4, the neglect of carrier-phase wind-up in the observation model already introduces errors that are larger than this, in the order of millimetres per second. However, it *is* significantly larger than the error that would be expected from the neglected accelerations: shown in the previous subsection to be on the order of nm/s^2 , the associated position error after one second would also be on the order of nanometres. Similarly, it is also larger than the error that could reasonably be caused by mismodelling of the Solar radiation pressure or aerodynamic drag, each being under $1 \mu\text{m/s}^2$ in magnitude for Sentinel-6A.

Indeed, it follows that the cause for this position error is to be sought elsewhere. The root cause can be identified by repeating the comparison with the precise TU Delft orbit using longer propagation intervals of ten seconds. This results in the position errors graphed in Figure 5.7. Where the error is qualitatively similar to the error for the one-second propagation arc, its magnitude also grows ten-fold. If this error were caused by an unmodelled or mismodelled acceleration, approximately quadratic growth would be expected for such small errors. The fact that this growth is linear with an increased propagation time suggests that it can be brought back to an error in the initial orbital state rather than in the trajectory model. This is not unexpected: the initial orbital velocity is computed through a transformation of the precise TU Delft orbit to inertial space, described in Section 4.3, that is known to be simplified. Indeed, the dynamical model at this point is limited by this rotation rather than the resulting force model.

6

Variational Equations

Among the four approximations to solve the variational equations compared in this work, we distinguish two overall approaches. These will be described in this chapter. The first approach, used in the reference as well as the oblate and spherical Earth models, solves the variational equations through numerical integration; this is described in Section 6.1. The second approach, described in Section 6.2, is based on analytical solutions to the Hill-Clohessy-Wiltshire equations of relative motion for a circular target orbit. At the end of this chapter, in Section 6.3, a third approach is introduced for approximation of the matrizant, based on a finite difference quotient. This method is used to verify the other two approaches.

6.1. Solution through numerical integration

For the reference, oblate, and spherical Earth models of the orbital variations, it is not tractable to solve for an analytic expression of the state transition matrix. Instead, their transition matrices are obtained via numerical integration of the variational equations, which describe the evolution of the orbital state transition and sensitivity matrices over time.

The variational equations can be derived directly from the equations of motion of the system in combination with the definitions of the state transition matrix and sensitivity matrix:

$$\dot{\mathbf{y}}(t) = f(t, \mathbf{y}(t), \mathbf{p}) = \begin{bmatrix} \mathbf{v}(t) \\ \mathbf{a}(t, \mathbf{y}(t), \mathbf{p}) \end{bmatrix}, \quad \Phi_{\mathbf{y}}(t, t_0) = \frac{\partial \mathbf{y}(t)}{\partial \mathbf{y}(t_0)}, \quad \Phi_{\mathbf{p}}(t) = \frac{\partial \mathbf{y}(t)}{\partial \mathbf{p}}.$$

The variational equations for the orbital state transition matrix can be obtained by differentiating the equations of motion with respect to the initial state $\mathbf{y}(t_0)$:

$$\frac{\partial}{\partial \mathbf{y}(t_0)} \dot{\mathbf{y}}(t) = \frac{\partial}{\partial \mathbf{y}(t_0)} f(t, \mathbf{y}(t), \mathbf{p}) = \frac{\partial f(t, \mathbf{y}(t), \mathbf{p})}{\partial \mathbf{y}(t)} \frac{\partial \mathbf{y}(t)}{\partial \mathbf{y}(t_0)} = \frac{\partial f(t, \mathbf{y}(t), \mathbf{p})}{\partial \mathbf{y}(t)} \Phi_{\mathbf{y}}(t, t_0) = \begin{bmatrix} 0_{3 \times 3} & I_{3 \times 3} \\ \frac{\partial \mathbf{a}}{\partial \mathbf{r}} & \frac{\partial \mathbf{a}}{\partial \mathbf{v}} \end{bmatrix} \Phi_{\mathbf{y}}.$$

Similarly, the variational equations for the sensitivity matrix are obtained by differentiating with respect to the dynamic parameters $\mathbf{p}$:

$$\frac{\partial}{\partial \mathbf{p}} \dot{\mathbf{y}}(t) = \frac{\partial}{\partial \mathbf{p}} f(t, \mathbf{y}(t), \mathbf{p}) = \frac{\partial f(t, \mathbf{y}(t), \mathbf{p})}{\partial \mathbf{y}(t)} \frac{\partial \mathbf{y}(t)}{\partial \mathbf{p}} + \frac{\partial f(t, \mathbf{y}(t), \mathbf{p})}{\partial \mathbf{p}} = \begin{bmatrix} 0_{3 \times 3} & I_{3 \times 3} \\ \frac{\partial \mathbf{a}}{\partial \mathbf{r}} & \frac{\partial \mathbf{a}}{\partial \mathbf{v}} \end{bmatrix} \Phi_{\mathbf{p}}(t) + \begin{bmatrix} 0_{3 \times n_p} \\ \frac{\partial \mathbf{a}}{\partial \mathbf{p}} \end{bmatrix}.$$

As evident, these differential equations fully describe the evolution of the orbital state transition and sensitivity matrices over time in terms of the acceleration partials $\frac{\partial \mathbf{a}}{\partial \mathbf{r}}$, $\frac{\partial \mathbf{a}}{\partial \mathbf{v}}$, and $\frac{\partial \mathbf{a}}{\partial \mathbf{p}}$. Both initial value problems are completed by the initial transition matrix values $\Phi_{\mathbf{y}}(t_0, t_0) = I_{6 \times 6}$ and $\Phi_{\mathbf{p}}(t_0) = 0_{6 \times n_p}$, which follow from their definitions.

Consequently, the combined variational equations take the following form:

$$\begin{bmatrix} \dot{\Phi}_{\mathbf{y}} & \dot{\Phi}_{\mathbf{p}} \end{bmatrix} = \begin{bmatrix} 0_{3 \times 3} & I_{3 \times 3} \\ \frac{\partial \mathbf{a}}{\partial \mathbf{r}} & \frac{\partial \mathbf{a}}{\partial \mathbf{v}} \end{bmatrix} \begin{bmatrix} \Phi_{\mathbf{y}} & \Phi_{\mathbf{p}} \end{bmatrix} + \begin{bmatrix} 0_{3 \times 6} & 0_{3 \times n_p} \\ 0_{3 \times 6} & \frac{\partial \mathbf{a}}{\partial \mathbf{p}} \end{bmatrix}. \quad (6.1)$$

This method allows for accurate estimation of the transition matrices. The only limiting factors are the knowledge of the acceleration partials with respect to the orbital state and dynamical parameters, and

the numerical precision available for the integration. Because the accuracy of these choices affects the confidence with which conclusions can be drawn about the difference in performance of each state transition matrix approximation, those choices are clarified in detail in the remainder of this section.

Acceleration partials

In computation of the acceleration partials $\frac{\partial \mathbf{a}}{\partial \mathbf{r}}$, $\frac{\partial \mathbf{a}}{\partial \mathbf{v}}$, and $\frac{\partial \mathbf{a}}{\partial \mathbf{p}}$, the default Tudat implementations are used, as described in [18]. These will be summarized in the following. Overall, similar methods are used to compute these partials for the reference model as well as for the oblate and spherical Earth models: we shall indicate so explicitly where this is not the case.

Earth gravity

For the full reference model, the partials of acceleration with respect to Earth gravity are computed according to its full spherical harmonic gravity. Analytic expressions exist for this partial, with examples found in [36, 44]. In the Earth-centered, Earth-fixed International Terrestrial Reference Frame (ITRF), this expression is

$$\left(\frac{\partial \mathbf{a}}{\partial \mathbf{r}}\right)_{\text{grav}, \text{ITRF}} = \sum_{n,m} \left(\frac{\partial \mathbf{a}_{nm}}{\partial \mathbf{r}}\right)_{\text{ITRF}}$$

for a given set of degrees n and orders m , where the partial derivative $(\partial \mathbf{a}_{nm} / \partial \mathbf{r})_{\text{ITRF}}$ can be expressed analytically by differentiation of the regular expressions for the acceleration $\mathbf{a}_{nm}$ due to Earth's spherical harmonic gravity of degree n and order m . Afterwards, the same acceleration partial is expressed in the ECI propagation frame by means of the transformation

$$\left(\frac{\partial \mathbf{a}}{\partial \mathbf{r}}\right)_{\text{grav}, \text{ECI}} = R_{\text{ITRF} \rightarrow \text{ECI}} \left(\frac{\partial \mathbf{a}}{\partial \mathbf{r}}\right)_{\text{grav}, \text{ITRF}} R_{\text{ECI} \rightarrow \text{ITRF}}.$$

Tudat contains an implementation for such analytic computation up to arbitrary degree and order, which is used in the orbit determination filter to determine the resulting partials. For the reference model, this formulation is used to compute the acceleration partial to full degree and order. Necessarily, this calculation is computationally costly, and contributes the most significant computational effort in the computation of the acceleration partials in the reference model. Accordingly, for the spherical and oblate Earth models, this computation is truncated to point-mass respectively J_2 -perturbed Earth gravity models; hence, the names of these models.

Lunisolar dynamic solid tides

Tudat implements the effect of Lunisolar dynamic solid tides via variations in the spherical harmonic coefficients of the Earth model used for computation of the gravity acceleration [18]. As a result, the effect of dynamic solid tides on the acceleration partials is included directly in the formulation of the Earth gravitational harmonic acceleration partials, without the necessity for any further adjustment. In the reference model, therefore, the effect of this variation is included by means of inclusion of the Lunisolar dynamic solid tides in its force model. For the two simplified numerical models, this effect is not taken into account. This is partly by necessity, as all or most of the spherical harmonic coefficients that would be adjusted are not modelled in the first place in either model.

Third-body perturbations

As the third-body perturbations for the used trajectory models are all point masses, their acceleration partials are computable analytically as the derivative of the central gravity term. For the partials with respect to the position components:

$$\left(\frac{\partial \mathbf{a}}{\partial \mathbf{r}}\right)_{n\text{-body}} = -\mu \left(\frac{1}{|\mathbf{r}|^3} I_{3 \times 3} - 3 \frac{\mathbf{r} \mathbf{r}^T}{|\mathbf{r}|^5} \right). \quad (6.2)$$

Note that the position $\mathbf{r}$ as used here must be expressed relative to the third-body under consideration, with orientations that match those of the central reference frame. Additionally, μ refers here to the gravitational parameter of the third body, *not* of the Earth. Naturally, the partial with respect to velocity is zero. It can be deduced directly from Equation (6.2) that this sensitivity is small, in the order of nms^{-2}/m , for orbital radii like those of Sentinel-6A. Hence, this partial is included only in the reference model: in the oblate and spherical Earth models, it is neglected.

Schwarzschild acceleration

The Schwarzschild acceleration affects the variation of the orbital state due to changes in the initial orbital state. Therefore, it manifests itself as an additional term to be computed for the acceleration partials with respect to position and velocity. Tudat incorporates this effect into the equations of motion according to a Parameterized Post-Newtonian (PPN) formalism [18]. The resulting acceleration can be computed analytically:

$$\left(\frac{\partial \mathbf{a}}{\partial \mathbf{r}}\right)_{rel} = \left(2(\gamma + \beta) \frac{\mu}{|\mathbf{r}|} \left(I - \frac{\mathbf{r}\mathbf{r}^T}{|\mathbf{r}|^2}\right) - \gamma \mathbf{v}\mathbf{v}^T I + 2(1 + \gamma) \mathbf{v}\mathbf{v}^T\right) \frac{\mu}{c^2 |\mathbf{r}|^3} - 3 \frac{\mathbf{a}\mathbf{r}^T}{|\mathbf{r}|^2}, \quad (6.3)$$

$$\left(\frac{\partial \mathbf{a}}{\partial \mathbf{v}}\right)_{rel} = \frac{\mu}{c^2 |\mathbf{r}|^3} (-2\gamma \mathbf{r}\mathbf{v}^T + 2(1 + \gamma) (\mathbf{r} \cdot \mathbf{v} I + \mathbf{v}\mathbf{r}^T)). \quad (6.4)$$

Here, γ and β are the PPN parameters, which are taken to have a value of 1 in accordance with general relativity. It is evident that the magnitude of these sensitivities is very low. Indeed, these terms are considered only for the reference model, and are neglected in the oblate and spherical Earth models.

Aerodynamic drag

Because NRLMSISE-00, an empirical atmospheric model, is used, it is not possible to compute the partial of aerodynamic drag with respect to position or velocity analytically. Instead, Tudat allows for estimation of this acceleration partial by means of a central difference quotient approximation, based on perturbations of 10 m in position and 1 cm/s in velocity in each direction. This approach is applied in the oblate and spherical Earth models, as well as in the reference model. This finite difference estimate, while necessary, does fundamentally limit the accuracy of even the reference model with respect to the aerodynamic drag partials. It is necessary to take note of this inexactness when comparing the simplified models with respect to the reference model.

The acceleration partial with respect to the drag coefficient, fortunately, can be computed analytically from the expression

$$\frac{\partial \mathbf{a}}{\partial c_d} = R_{\mathbf{v}_{atm} \rightarrow ECI} \begin{bmatrix} -0.5\rho |\mathbf{v}_{atm}|^2 S m^{-1} \\ 0 \\ 0 \end{bmatrix}. \quad (6.5)$$

This expression can be recognized as the derivative of classic lift formula, with a rotation applied from the aerodynamic frame to the ECI, J2000-oriented reference frame used for propagation: ρ denotes the atmospheric density, S the aerodynamic reference surface area, m the spacecraft mass, and $\mathbf{v}_{atm}$ the spacecraft velocity with respect to the Earth atmosphere. For the aerodynamic frame, the Tudat convention is followed of a positive x-axis in the direction of the spacecraft velocity with respect to the atmosphere, $\mathbf{v}_{atm}$. Naturally, this representation of aerodynamic drag, and thus also its acceleration partials, assumes the absence of wind.

Solar radiation pressure

The Solar Radiation Pressure (SRP) partials can be computed analytically again. Tudat computes the partial with respect to position as

$$\left(\frac{\partial \mathbf{a}}{\partial \mathbf{r}}\right)_{SRP} = \frac{c_r A P_\odot}{m} \left(I \frac{1}{|\mathbf{r} - \mathbf{r}_\odot|} - 3(\mathbf{r} - \mathbf{r}_\odot) \frac{(\mathbf{r} - \mathbf{r}_\odot)^T}{|\mathbf{r} - \mathbf{r}_\odot|^3} \right). \quad (6.6)$$

Here, $P_\odot$ is the Solar radiation pressure magnitude, A the Solar radiation pressure reference area, m the spacecraft mass, and $\mathbf{r}_\odot$ the position of the Sun in the inertial reference frame of the propagation. Since velocity does not influence the radiation pressure, the SRP contribution to the acceleration-velocity partial is zero.

Similarly to the partial with respect to the drag coefficient, the acceleration partial with respect to the radiation pressure coefficient is also computed analytically:

$$\frac{\partial \mathbf{a}}{\partial c_r} = \frac{1}{c_r} \mathbf{a}_{SRP} = -\frac{P_\odot A}{m} (\mathbf{r}_\odot - \mathbf{r}). \quad (6.7)$$

This partial, as well as the SRP partial of acceleration with respect to position, is included in all three numerically-integrated variational models. It would be justifiable, due to the large distance to the Sun,




Figure 6.1: Differential correction error introduced by numerical integration error at different integration time steps

to ignore the contribution of the SRP to the acceleration-position partial. This is not done, mostly for pragmatic reasons: it is not very expensive to compute this contribution, and it is not possible by default in Tudat to disable its computation.

Empirical accelerations

Finally, the impact of empirical accelerations is considered. Because of their nature, it is quite straightforward to compute this contribution analytically:

$$\frac{\partial \mathbf{a}}{\partial \mathbf{a}_e} = R_{\text{HCW} \rightarrow \text{ECI}} I_{3 \times 3}. \quad (6.8)$$

Indeed, only a rotation from the local Hill-Clohessy-Wiltshire frame to the propagation frame is necessary. As the empirical accelerations are, conceptually, independent of position and velocity, their to the corresponding partials is negligible. In theory, a small effect arises due to the change in orientation of the Hill-Clohessy-Wiltshire frame due to changes in position or velocity, but this effect is ignored in the Tudat implementation. This same implementation is used for the reference model as well as the spherical and oblate Earth models.

Numerical integration

For integration of the variational equations over time, the same integrator and step size are used as derived in Chapter 5 for the trajectory propagation itself. This choice is justified by the fact that the dynamics of the variational equations evolve at the same or lower pace as the dynamics of the trajectory. Also, when an acceleration is modelled to sufficient accuracy by a given integrator, it is to be expected that perturbations of this acceleration, which are generally smaller in magnitude, are also modelled to at least the same accuracy.

To verify this assertion, the orbital state transition and sensitivity matrices were computed over the entire precise TU Delft Sentinel-6A orbit for 1 January 2022, for all power-of-two integrator step sizes from 2^0 to 2^{-7} seconds. These step sizes were chosen to cover a regime that covers errors which are dominated by truncation error (near 2^0 s) down to a step size that is almost assuredly dominated by round-off errors (near 2^{-7} s). This gives the most complete picture of the integration error that is to be expected near the actual used step size of 1 second, as it allows us to show that it is indeed in the truncation error regime. The average differential correction that follows from comparison of these results with those obtained with a 2^{-8} -second step size are given in Figure 6.1. It follows that an inherent truncation error of about 10^{-12} is to be expected when computing approximations of the transition matrices via numerical integration at a 1-second time step. This is far more accurate than necessary for the analyses carried out in this work: as will be shown in Chapter 8, the approximate solutions to the variational equations result in significantly greater errors.

6.2. Analytical solution via Hill-Clohessy-Wiltshire equations

The Hill-Clohessy-Wiltshire approximation of the state transition matrix does not make use of numerical integration. Rather, it follows from differentiation of the analytical trajectory solution described by the Hill-Clohessy-Wiltshire equations. Below, we derive this trajectory solution under the assumption of

constant perturbing accelerations. From this, the orbital state transition and sensitivity matrices are obtained directly through differentiation. Finally, the resulting expressions are transformed from the Hill-Clohessey-Wiltshire frame, in which they are derived, to the inertial frame in which the orbital state is expressed.

Trajectory solution

For approximation of the state transition and sensitivity matrices of a near-circular orbit, we can use analytic solutions to the Hill-Clohessey-Wiltshire equations of relative motion [12]:

$$\begin{aligned}\ddot{x} - 3n^2x - 2n\dot{y} &= a_x \\ \ddot{y} + 2n\dot{x} &= a_y, \\ \ddot{z} + n^2z &= a_z,\end{aligned}\tag{6.9}$$

where $n = \sqrt{\mu/a^3}$ is the mean motion, and a_x, a_y , and a_z are time-invariant accelerations. The equations are defined in the Hill-Clohessey-Wiltshire, or HCW, frame, such that the x -axis is in positive radial direction, y -axis along-track (positive in the velocity direction), and z is the cross-track axis that completes the right-hand coordinate frame. The origin coincides with the unperturbed orbit, and x, y , and z describe deviations from this trajectory in the co-rotating frame. We note that the equations as stated above already incorporate the assumptions that the deviation from the target orbit is small compared to the orbital radius, and that the orbital eccentricity is negligible.

Applying $a_x = a_y = a_z = 0 \text{ ms}^{-2}$, these equations can be shown to have the homogeneous solutions

$$\begin{aligned}x_h(t) &= x(t_0)(4 - 3\cos n\Delta t) + \dot{x}(t_0)\frac{\sin n\Delta t}{n} - \dot{y}(t_0)\frac{2}{n}(\cos n\Delta t - 1), \\ y_h(t) &= 6x(t_0)(\sin n\Delta t - n\Delta t) + y(t_0) + \dot{x}(t_0)\frac{2\cos n\Delta t - 2}{n} + \dot{y}(t_0)\left(-3\Delta t + \frac{4\sin n\Delta t}{n}\right), \\ z_h(t) &= \frac{1}{n}\dot{z}(t_0)\sin n\Delta t + z(t_0)\cos n\Delta t.\end{aligned}\tag{6.10}$$

Here, we define $\Delta t = t - t_0$ for brevity of notation. Similarly, we can consider the general case where a_x, a_y, a_z are allowed to be non-zero constants. In that case, the homogeneous solutions are extended with particular components to obtain the general solutions

$$\begin{aligned}x(t) &= x_h(t) + a_x\frac{1}{n^2}(1 - \cos n\Delta t) + a_y\frac{2}{n^2}(n\Delta t - \sin n\Delta t), \\ y(t) &= y_h(t) + a_x\frac{-2n\Delta t + 2\sin n\Delta t}{n^2} + a_y\left(-\frac{3\Delta t^2}{2} - \frac{4\cos n\Delta t}{n^2} + \frac{4}{n^2}\right), \\ z(t) &= z_h(t) + \frac{1}{n^2}a_z(1 - \cos n\Delta t).\end{aligned}\tag{6.11}$$

Expressions for the relative velocities can be obtained directly from these solutions by differentiating them with respect to time:

$$\begin{aligned}\dot{x}(t) &= 3x(t_0)n\sin n\Delta t + \dot{x}(t_0)\cos n\Delta t + 2\dot{y}(t_0)\sin n\Delta t + a_x\frac{\sin n\Delta t}{n} + a_y\frac{2(1 - \cos n\Delta t)}{n}, \\ \dot{y}(t) &= 6x(t_0)n(\cos n\Delta t - 1) - 2\dot{x}(t_0)\sin n\Delta t + \dot{y}(t_0)(4\cos n\Delta t - 3) \\ &\quad + a_x\frac{2(\cos n\Delta t - 1)}{n} + a_y\left(-3\Delta t + \frac{4\sin n\Delta t}{n}\right), \\ \dot{z}(t) &= -z(t_0)n\sin n\Delta t + \dot{z}(t_0)\cos n\Delta t + a_z\frac{\sin n\Delta t}{n}.\end{aligned}\tag{6.12}$$

State transition matrix

With the analytical solutions known, the state transition matrix can be readily derived through differentiation of the position and velocity solutions with respect to the initial positions and velocities:

$$\Phi_{\mathbf{y}}(t, t_0) = \frac{\partial \mathbf{y}(t)}{\partial \mathbf{y}(t_0)} = \begin{bmatrix} 4 - 3 \cos n\Delta t & 0 & 0 & \frac{1}{n} \sin n\Delta t & \frac{2}{n} (1 - \cos n\Delta t) & 0 \\ -6nt + 6 \sin n\Delta t & 1 & 0 & \frac{2}{n} (\cos n\Delta t - 1) & -3\Delta t + \frac{4}{n} \sin n\Delta t & 0 \\ 0 & 0 & \cos n\Delta t & 0 & 0 & \frac{1}{n} \sin n\Delta t \\ 3n \sin n\Delta t & 0 & 0 & \cos n\Delta t & 2 \sin n\Delta t & 0 \\ 6n (\cos n\Delta t - 1) & 0 & 0 & -2 \sin n\Delta t & 4 \cos n\Delta t - 3 & 0 \\ 0 & 0 & -n \sin n\Delta t & 0 & 0 & \cos n\Delta t \end{bmatrix}$$

Sensitivity matrix

For the sensitivity matrix, we must similarly obtain

$$\Phi_{\mathbf{p}}(t, t_0) = \frac{\partial \mathbf{y}(t)}{\partial \mathbf{p}} = \begin{bmatrix} \frac{\partial x}{\partial c_r} & \frac{\partial x}{\partial c_d} & \frac{\partial x}{\partial a_{e,x}} & \frac{\partial x}{\partial a_{e,y}} & \frac{\partial x}{\partial a_{e,z}} \\ \frac{\partial y}{\partial c_r} & \frac{\partial y}{\partial c_d} & \frac{\partial y}{\partial a_{e,x}} & \frac{\partial y}{\partial a_{e,y}} & \frac{\partial y}{\partial a_{e,z}} \\ \frac{\partial z}{\partial c_r} & \frac{\partial z}{\partial c_d} & \frac{\partial z}{\partial a_{e,x}} & \frac{\partial z}{\partial a_{e,y}} & \frac{\partial z}{\partial a_{e,z}} \\ \frac{\partial \dot{x}}{\partial c_r} & \frac{\partial \dot{x}}{\partial c_d} & \frac{\partial \dot{x}}{\partial a_{e,x}} & \frac{\partial \dot{x}}{\partial a_{e,y}} & \frac{\partial \dot{x}}{\partial a_{e,z}} \\ \frac{\partial \dot{y}}{\partial c_r} & \frac{\partial \dot{y}}{\partial c_d} & \frac{\partial \dot{y}}{\partial a_{e,x}} & \frac{\partial \dot{y}}{\partial a_{e,y}} & \frac{\partial \dot{y}}{\partial a_{e,z}} \\ \frac{\partial \dot{z}}{\partial c_r} & \frac{\partial \dot{z}}{\partial c_d} & \frac{\partial \dot{z}}{\partial a_{e,x}} & \frac{\partial \dot{z}}{\partial a_{e,y}} & \frac{\partial \dot{z}}{\partial a_{e,z}} \end{bmatrix}.$$

However, these terms cannot be obtained directly from the previously derived Hill-Clohessy-Wiltshire solutions. Instead, we rewrite the acceleration terms as follows:

$$\begin{aligned} \ddot{x} - 3n^2 x - 2n\dot{y} &= a_x = a_{e,x} + \hat{x} \cdot (a_r + a_d), \\ \ddot{y} + 2n\dot{x} &= a_y = a_{e,y} + \hat{y} \cdot (a_r + a_d), \\ \ddot{z} + n^2 z &= a_z = a_{e,z} + \hat{z} \cdot (a_r + a_d). \end{aligned}$$

Here, $\hat{x}$, $\hat{y}$, and $\hat{z}$ are the unit vectors in direction of the radial, along-track, and cross-track axes of the Hill-Clohessy-Wiltshire frame, respectively, expressed in the ECI, J2000-oriented frame. These unit vectors, representing the orientation of the Hill-Clohessy-Wiltshire frame in the inertial frame, are also assumed to be constant over the propagation time. We also approximate the radiation pressure and drag terms: we take them to be constant in magnitude and direction in the Hill-Clohessy-Wiltshire frame over a single propagation arc. In the context of application to propagation between Sentinel-6A measurements, which are spaced one second apart, these assumptions introduce only negligible error.

Then, we can express the partial derivatives of the orbital state with respect to the dynamic parameters using the intermediate derivative of the total acceleration $\mathbf{a}$ in the Hill-Clohessy-Wiltshire frame with respect to the dynamic parameters $\mathbf{p}$:

$$\Phi_{\mathbf{p}}(t, t_0) = \frac{\partial \mathbf{y}}{\partial \mathbf{a}} \frac{\partial \mathbf{a}}{\partial \mathbf{p}} = \begin{bmatrix} \frac{1 - \cos n\Delta t}{n^2} & \frac{2n\Delta t - 2 \sin n\Delta t}{n^2} & 0 \\ -\frac{2n\Delta t + 2 \sin n\Delta t}{n^2} & -\frac{3\Delta t^2}{2} - \frac{4 \cos n\Delta t - 4}{n^2} & 0 \\ 0 & 0 & \frac{1 - \cos n\Delta t}{n^2} \\ \frac{\sin n\Delta t}{n} & \frac{2 - 2 \cos n\Delta t}{n} & 0 \\ -\frac{2 \cos n\Delta t - 2}{n} & -\frac{3n\Delta t + 4 \sin n\Delta t}{n} & 0 \\ 0 & 0 & \frac{\sin n\Delta t}{n} \end{bmatrix} \begin{bmatrix} \hat{x} \cdot \frac{\partial a_r}{\partial c_r} & \hat{x} \cdot \frac{\partial a_d}{\partial c_d} & 1 & 0 & 0 \\ \hat{y} \cdot \frac{\partial a_r}{\partial c_r} & \hat{y} \cdot \frac{\partial a_d}{\partial c_d} & 0 & 1 & 0 \\ \hat{z} \cdot \frac{\partial a_r}{\partial c_r} & \hat{z} \cdot \frac{\partial a_d}{\partial c_d} & 0 & 0 & 1 \end{bmatrix}$$

Here, we can express the partial derivatives of the radiation pressure and drag accelerations with respect to their derivatives as

$$\frac{\partial a_r}{\partial c_r} = \frac{1}{c_r} a_r, \quad \frac{\partial a_d}{\partial c_d} = \frac{1}{c_d} a_d,$$

based on the assumption of constant mass for the spacecraft.

Transformation to the propagation frame

Note that both the state transition and sensitivity matrix are expressed in Hill-Clohessy-Wiltshire coordinates. For application in orbit determination, we need to be able to express both in the J2000-oriented

propagation frame. Deviations in position and velocity can be transformed to and from the Hill-Clohesy-Wiltshire and J2000-oriented ECI frame using the following rotation matrices [13]:

$$M_{\text{HCW} \rightarrow \text{ECI}} = \begin{bmatrix} A_{\text{HCW} \rightarrow \text{ECI}} & 0 \\ \tilde{\omega} A_{\text{HCW} \rightarrow \text{ECI}} & A_{\text{HCW} \rightarrow \text{ECI}} \end{bmatrix},$$

$$M_{\text{ECI} \rightarrow \text{HCW}} = \begin{bmatrix} A_{\text{ECI} \rightarrow \text{HCW}} & 0 \\ -A_{\text{ECI} \rightarrow \text{HCW}} \tilde{\omega} & A_{\text{ECI} \rightarrow \text{HCW}} \end{bmatrix},$$

where

$$A_{\text{ECI} \rightarrow \text{HCW}} = \begin{bmatrix} \hat{x}^T \\ \hat{y}^T \\ \hat{z}^T \end{bmatrix}, \quad A_{\text{HCW} \rightarrow \text{ECI}} = [\hat{x} \quad \hat{y} \quad \hat{z}],$$

and

$$\hat{x} = \frac{\mathbf{r}_{\text{ECI}}}{\|\mathbf{r}_{\text{ECI}}\|}, \quad \hat{y} = \hat{z} \times \hat{x}, \quad \hat{z} = \frac{\mathbf{r}_{\text{ECI}} \times \mathbf{v}_{\text{ECI}}}{\|\mathbf{r}_{\text{ECI}} \times \mathbf{v}_{\text{ECI}}\|}, \quad \tilde{\omega} = \begin{bmatrix} 0 & -\omega_1 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}.$$

These expressions can be used to transform the state transition and sensitivity matrices as follows:

$$\Phi_{\mathbf{y}, \text{ECI}}(t, t_0) = M_{\text{HCW} \rightarrow \text{ECI}}(t) \Phi_{\mathbf{y}, \text{HCW}}(t, t_0) M_{\text{ECI} \rightarrow \text{HCW}}(t_0),$$

$$\Phi_{\mathbf{p}, \text{ECI}}(t, t_0) = M_{\text{HCW} \rightarrow \text{ECI}}(t) \Phi_{\mathbf{p}, \text{HCW}}(t, t_0).$$

6.3. Numerical solution using the difference quotient

To verify that the solution of the variational equations as implemented in Tudat is consistent with its regular trajectory propagation model, the resulting state transition matrix estimates are compared with a finite difference approximation of its constituent derivatives. We follow [16] in applying a central difference approximation to approximate column i of the orbital state transition matrix

$$\Phi_{\mathbf{y}, i} = \frac{\partial \mathbf{y}(t)}{\partial y_i(t_0)} = \frac{(\mathbf{f}_1 - \mathbf{f}_2) - \rho^3(\mathbf{f}_3 - \mathbf{f}_4)}{\rho h_i(1 - \rho^2)} + \mathcal{O}(h_i^4), \quad (6.13)$$

where, writing the perturbed orbital states $\mathbf{y}(t)$ as a function $\mathbf{f}(\mathbf{y}_0 + \mathbf{h}_i)$ of some perturbed initial orbital state $\mathbf{y}_0 + \mathbf{h}_i$:

$$\begin{aligned} \mathbf{f}_1 &= \mathbf{f}(y_1, y_2, \dots, y_i + \frac{h_i}{2}, \dots, y_6), \\ \mathbf{f}_2 &= \mathbf{f}(y_1, y_2, \dots, y_i - \frac{h_i}{2}, \dots, y_6), \\ \mathbf{f}_3 &= \mathbf{f}(y_1, y_2, \dots, y_i + \frac{\rho h_i}{2}, \dots, y_6), \\ \mathbf{f}_4 &= \mathbf{f}(y_1, y_2, \dots, y_i - \frac{\rho h_i}{2}, \dots, y_6). \end{aligned}$$

Similarly, we approximate column j of the sensitivity matrix as

$$\Phi_{\mathbf{p}, j} = \frac{\partial \mathbf{y}(t)}{\partial p_j} = \frac{(\mathbf{f}_1 - \mathbf{f}_2) - \rho^3(\mathbf{f}_3 - \mathbf{f}_4)}{\rho h_j(1 - \rho^2)} + \mathcal{O}(h_j^4), \quad (6.14)$$

where we now instead use $\mathbf{f}(\mathbf{p} + \mathbf{h}_i)$ with

$$\begin{aligned} \mathbf{f}_1 &= \mathbf{f}(p_1, p_2, \dots, p_i + \frac{h_i}{2}, \dots, p_5), \\ \mathbf{f}_2 &= \mathbf{f}(p_1, p_2, \dots, p_i - \frac{h_i}{2}, \dots, p_5), \\ \mathbf{f}_3 &= \mathbf{f}(p_1, p_2, \dots, p_i + \frac{\rho h_i}{2}, \dots, p_5), \\ \mathbf{f}_4 &= \mathbf{f}(p_1, p_2, \dots, p_i - \frac{\rho h_i}{2}, \dots, p_5). \end{aligned}$$

These approximations result in an order $\mathcal{O}(h^4)$ approximation of the orbital state transition and sensitivity matrices. The value ρ can be chosen relatively freely: in this application, $\rho = 1/4$ is used.

This difference quotient is relatively expensive to compute, requiring four additional state propagations for each column of the state transition matrix to compute. Moreover, it is quite sensitive to numerical round-off errors because of the small sensitivities of the trajectory to changes in the aerodynamic drag and radiation pressure coefficients, which limits the achievable accuracy. Consequently, it is not suitable for direct application as state transition matrix approximator.

This approximation is, on the other hand, useful to verify the consistency of the Tudat variational equations with their corresponding equations of motion: by construction, it requires only the trajectory model itself to compute a reasonable estimate of solutions to the variational equations. To this end, we compare the state transition matrix approximation obtained through numerical integration of the reference model described in Section 3.2 with the difference-quotient approximation described above. Both are computed for each ten-second arc along the precise TU Delft orbit for all of 1 January 2022.

The similarity of both matrices is measured in terms of the differential correction error of the difference-quotient state transition matrix in estimating the reference state transition matrix: these are plotted in Figure 6.2. It is visible that the differential correction error is approximately 10^{-3} at its maximum. Referring to the interpretation of the differential correction error in Appendix A as the maximum relative magnitude of the residual state perturbation after a differential correction, it becomes evident that this is quite small: it corresponds to a mismodelling of state perturbations by at most 0.1%. This means that, indeed, it can be concluded that the variational equations and equations of motion implemented in Tudat are consistent.




Figure 6.2: Differential correction error obtained when comparing the difference quotient approximation with the reference state transition matrix

7

State Estimation

In this chapter, the reader can find an overview of the approach to state estimation used in all compared orbit determination filters. While this approach is neither novel nor the principal research topic of this work, it is the framework on which the comparison between the different state transition and sensitivity matrices rests. Hence, to draw a complete picture of the orbit determination filters, their workings in terms of state estimation are presented here. This chapter starts with a brief justification of the used filter mechanization in Section 7.1. This is followed by a brief description of the chosen set of state parameters in Section 7.2, including the assumptions made in modelling each of them. Afterwards, the reader will find a description of the exact algorithms used for state estimation: in Sections 7.3 to 7.6 one can find a description of the processes applied to reconcile the a priori state estimate with a batch of Sentinel-6A GNSS observations. In Section 7.7, this is complemented by the algorithms applied to propagate the covariance estimates forward in time. Finally, in Section 7.8, the design of the free filter parameters is given.

7.1. Filter mechanization

For state estimation inside our orbit determination algorithm, the Square Root Information Filter (SRIF) mechanization of the Kalman filter is used, based on the description given by Bierman in [5]. This approach has four advantages over other mechanizations: in the first place, its propagation of the square root information matrix rather than the full information matrix permits improved numerical stability. Secondly, its inclusion of a priori data in the shape of additional observation equations lends itself quite conveniently to implicit initialization of the carrier-phase ambiguities and SISRE estimates in the filter. Thirdly, the SRIF lends itself quite naturally to outlier detection, as the a posteriori residuals are a direct by-product of the data processing step. The fourth and most significant point is that the SRIF can split the state into noisy and deterministic parameters to reduce the computational effort associated with a state update; as nearly half of the state parameters, the carrier-phase ambiguities, are deterministic, this can make a significant difference.

This choice of mechanization is in line with similar approaches found in the state-of-the-art: both the Haiyang-2D filter and the forward filter used in GipsyX/RTGx apply the SRIF [4, 34]. Indeed, in both cases extended numerical stability and implicit outlier detection are cited as reasons. In DIODE, similarly, the UD formulation of the Kalman filter is used [29]: in that mechanization, the UD factors of the covariance matrix are propagated to achieve extended numerical stability. Only RTNAV seems to make use of a regular extended Kalman filter, showing that an extended-precision mechanization is not strictly necessary for orbital accuracies of 3 cm [38]. Still, the other advantages of the SRIF mechanization mean that it was the mechanization of choice for the presented filter implementation.

7.2. State design

In the treatment of filter state estimates, five categories of filter parameters are distinguished, roughly following the approach to partitioning used by Bierman in his formulation of the SRIF [5]. We extend his partitioning to include four criteria: whether a parameter is deterministically time-varying, whether it is stochastically time-varying, whether it affects the system dynamics, and whether it directly affects the

GNSS observations. As will be evident in Sections 7.3 and 7.7, this allows several computational and notational simplifications.

This results in the following split:

$$x = \begin{bmatrix} y \\ p \\ t \\ s \\ b \end{bmatrix}, \quad \text{where} \quad \begin{cases} y & \text{is the orbital state.} \\ p & \text{is the set of dynamic parameters.} \\ t & \text{are the estimated clock offsets.} \\ s & \text{are the SISRE-compensating parameters.} \\ b & \text{are the carrier-phase ambiguities.} \end{cases}$$

Notably, the correlated process noise from the state formulation used by Bierman is split by means of the dynamic or static nature and type of stochastic process into dynamic parameters p , clock offsets t , and SISRE-compensating parameters s . In the subsections below, we delve further into the properties of and assumptions underlying each set of parameters.

This state formulation is similar to other formulations from the state-of-the-art. It coincides exactly with the set of state parameters used in the real-time Haiyang-2D filter [34]. The state formulations used in the works of Montenbruck et al. and of Darugna et al. differ slightly due to their choice to ignore the SISRE-compensating parameters, instead opting to absorb those errors in the estimated carrier-phase ambiguities [14, 38]. While these works show that this choice, which introduces a slight inconsistency in the observation equations, does not preclude orbit accuracies down to 3 cm, Li et al. show that it does reduce the resulting orbit accuracy by up to 30% [34]; hence, the choice is made to separately model the SISRE-compensating parameters and carrier-phase ambiguities, in spite of the higher computational effort required to do so. In DIODE [29], the filter in its entirety is split into orbit and clock filters, to achieve a purportedly greater robustness: this was not investigated in the present work, but could be of interest in later contributions.

Orbital state

The principal system state to be estimated, naturally, is the orbital state of Sentinel-6A. To simplify the computation of the state transition and sensitivity matrices, these parameters are formulated as the Cartesian state, to match the formulation used in the trajectory model:

$$y = \begin{bmatrix} r \\ v \end{bmatrix}, \quad \text{where} \quad \begin{cases} r & \text{is the Cartesian position.} \\ v & \text{is the Cartesian velocity.} \end{cases}$$

Note, additionally, that the orbital states are expressed in the J2000-oriented ECI frame also used for propagation, to prevent repeated reference frame transformations when propagating the trajectory for a single observation batch. This is not the state estimate used in comparison with the precise orbits, which are given in one of several instantiations of the ITRS. Indeed, the real output state estimate of the orbit determination filter would be in terms of an Earth-fixed state estimate. Consequently, unless specified otherwise, all precise-versus-estimated orbital state comparisons in this work will thus be in terms of the ITRF.

In principle, the used filter implementation supports stochastic orbital state parameters, based on the successful application of white process noise in the work of Li et al. [34]. Preliminary results, however, showed this noise to be detrimental to filter performance in the scenario considered in the present work; hence, it was lowered to a negligible 10^{-20} m and m/s in each position and velocity component, respectively. For all intents and purposes in further analysis, the orbital state is therefore considered deterministic.

It is noted that state estimation is performed using an epoch state filter, i.e., we accumulate observations over multiple points in time. As such, the orbital state y as estimated in the filter is the estimated orbital state at some chosen start epoch of the observation batch, *not* at the time(s) of observation. For this implementation, this is chosen to be the time of the last observation from the previous epoch, to maximize the observability of the dynamic parameters in the observations.

Dynamic parameters

Several parameters are used to compensate deficiencies in the trajectory model. These parameters all affect the system dynamics, but have no direct impact on GNSS observations; hence, they are modelled

separately from the orbital state. Five parameters are accounted for:

$$\mathbf{p} = \begin{bmatrix} c_r \\ c_d \\ a_r \\ a_s \\ a_w \end{bmatrix}, \quad \text{where} \begin{cases} c_r & \text{is the Solar radiation pressure coefficient, random walk process noise.} \\ c_d & \text{is the aerodynamic drag coefficient, random walk process noise.} \\ a_r & \text{is the radial empirical acceleration, correlated process noise.} \\ a_s & \text{is the along-track empirical acceleration, correlated process noise.} \\ a_w & \text{is the cross-track empirical acceleration, correlated process noise.} \end{cases}$$

The first two parameters correspond to the surface forces modelled for Sentinel-6A. They are considered as random walk process noise, to compensate for the deficiencies in the radiation and atmosphere models. Additionally, this compensates for the inherent inability of a cannonball spacecraft model to incorporate the changes in incident surface geometry due to rotation.

The latter three parameters are used to compensate both known and unknown deficiencies in the trajectory model. These deficiencies are modelled as exponentially-correlated random variables, also referred to as a Gauss-Markov process [53]. Such random variables have the convenient property that their autocorrelation fades over time. This is desirable because the largest force model deficiencies, as shown in Section 5.5, are known to have this property. In this work, a correlation time τ of 600 s, conform the value used in similar application to Sentinel-6A [38], was found to be suitable.

For convenience and flexibility in implementation, all five dynamic parameters are modelled as exponentially-correlated process noise, resulting in a dynamic model

$$\mathbf{p}(t) = \mathbf{M}_p \mathbf{p}(t_0),$$

where $\mathbf{M}_p$ is a diagonal matrix containing the exponential scaling factors. These are of value $\exp(-\frac{t-t_0}{\tau})$ for the empirical accelerations, resulting in exponentially-correlated process noise. For the Solar radiation pressure and aerodynamic drag coefficients, these are unity, which is equivalent to random walk process noise.

Clock offsets

In this work, two clock offsets are estimated explicitly during orbit determination: the receiver clock offset δt_{rcv} , and the Galileo-GPS inter-system clock bias δt_{GGTO} . Similarly to the aerodynamic drag and Solar radiation pressure coefficients, the two clock offsets are modelled as random walk variables to account for the inherently stochastic behaviour of these clock offsets:

$$\mathbf{t} = \begin{bmatrix} \delta t_{rcv} \\ \delta t_{GGTO} \end{bmatrix}, \quad \text{where} \begin{cases} \delta t_{rcv} & \text{is the receiver clock offset, random walk process noise.} \\ \delta t_{GGTO} & \text{is the GGTO, random walk process noise.} \end{cases}$$

The primary difference between the force coefficients and the clock offsets is that the clock offsets affect the system observations rather than its dynamics. These offsets are estimated inside the filter as ct rather than merely t ; this brings their magnitude nearer to unity, which is desirable in terms of numerical stability.

Signal-in-space ranging errors

To compensate the ranging errors arising from mismodelling of the transmitter clocks and GNSS antenna positions, a SISRE-compensating parameter is introduced for every tracked GNSS satellite. These parameters are estimated as follows:

$$\mathbf{s} = \begin{bmatrix} s_1 \\ \vdots \\ s_N \end{bmatrix}, \quad \text{where } s_n \text{ is the SISRE for a GNSS satellite } n, \text{ exponentially-correlated process noise.}$$

Because these errors coincide between pseudorange and carrier-phase measurements, and over different times, they are estimable.

The choice to include process noise in these parameters is based on their function of compensating, among others, errors in the GNSS onboard clocks, which are known to be stochastic. The reason for exponentially-correlated process instead of random walk parameters is based on the desire to keep their

value near zero. If no restrictions are placed on the magnitude of the SISRE parameters, the receiver clock offset and GGTO would no longer be estimable [46]. Indeed, any offset in the receiver clock or GGTO could equivalently be applied in all SISRE estimates, rendering the resulting estimates wholly meaningless. A relatively large correlation time τ of 6000 s suffices to bound the SISRE estimates near zero without resulting in too significant underestimation of the pseudorange. The result is the following dynamical model for the SISRE-absorbing parameters:

$$\mathbf{s}(t) = \mathbf{M}_s \mathbf{s}(t_0).$$

This is a pragmatic solution, justified by the fact that an accurate receiver clock offset is not of primary interest for purposes of altimetry. In this work, it is used only to validate the resulting system state estimates by comparison with the precise TU Delft solution. If a more accurate and consistent estimate of the SISRE-absorbing parameters or receiver clock offset δt_{rcv} were needed, a formal choice of constraints would be needed to make the system consistent.

Carrier-phase ambiguities

Finally, for each active carrier-phase lock, a float carrier-phase ambiguity is estimated:

$$\mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_M \end{bmatrix}, \quad \text{where } b_m \text{ is the carrier-phase ambiguity for an observation arc } m, \text{ deterministic.}$$

The ambiguities are modelled as deterministic parameters, considering that their stochastic variation is, largely, below the level of observability. Cycle slips can introduce stochastic and discontinuous changes in the ambiguity at the level of one wavelength, but those are neglected: instead, they are filtered out through outlier detection. A rudimentary form of loss-of-lock detection is present: a new ambiguity parameter is introduced whenever a signal is lost for one or more observation epochs inside the same batch before being reacquired.

It is noted that, contrary to Montenbruck et al. [38], who combined the SISRE-absorbing parameters and carrier-phase ambiguities into one pseudo-ambiguity, the choice was made here to keep them separate. This choice is based on the analysis performed in [34], where it was shown that introduction of a separate carrier-phase ambiguity improved 3D RMS orbit determination performance by up to 30% for the Haiyang-2D satellite. The computational simplification permitted by combining both types of parameters was not deemed significant enough to justify such a reduction in performance.

7.3. Data processing

The inclusion of observations to obtain an a posteriori state estimate is done using batched observations, that is, using observations from multiple subsequent observation times. This improves the observability of the dynamic parameters, as those require some time before they start significantly affecting the orbital state. Additionally, this can be advantageous from a computational standpoint; for the SRIF, however, this effect is not as significant as for other Kalman filter mechanizations. After collection of all observations for one batch, they are reconciled with the a priori state estimate.

To this end, we transform our a priori covariance-estimate information $\tilde{P}_x$ and $\delta \tilde{\mathbf{x}}$ into the equivalent a priori data equation

$$\tilde{R}_x \delta \tilde{\mathbf{x}}(t_0) = \tilde{\mathbf{z}}_x - \tilde{\boldsymbol{\nu}}_x, \quad (7.1)$$

where $\delta \tilde{\mathbf{x}}(t_0) = \tilde{R}_x^{-1} \tilde{\mathbf{z}}_x$, $\tilde{P}_x = \tilde{R}_x^{-1} \tilde{R}_x^{-T}$, and $\boldsymbol{\nu}_x$ is the observation noise, assumed to have zero mean and unity covariance. Splitting the state and covariance estimate into their constituent components, the a priori data equation becomes

$$\begin{bmatrix} \tilde{R}_y & \tilde{R}_{yp} & \tilde{R}_{yt} & \tilde{R}_{ys} & \tilde{R}_{yb} \\ \tilde{R}_{py} & \tilde{R}_p & \tilde{R}_{pt} & \tilde{R}_{ps} & \tilde{R}_{pb} \\ \tilde{R}_{ty} & \tilde{R}_{tp} & \tilde{R}_t & \tilde{R}_{ts} & \tilde{R}_{tb} \\ \tilde{R}_{sy} & \tilde{R}_{sp} & \tilde{R}_{st} & \tilde{R}_s & \tilde{R}_{sb} \\ \tilde{R}_{by} & \tilde{R}_{bp} & \tilde{R}_{bt} & \tilde{R}_{bs} & \tilde{R}_b \end{bmatrix} \begin{bmatrix} \delta \tilde{\mathbf{y}}(t_0) \\ \delta \tilde{\mathbf{p}}(t_0) \\ \delta \tilde{\mathbf{t}}(t_0) \\ \delta \tilde{\mathbf{s}}(t_0) \\ \delta \tilde{\mathbf{b}}(t_0) \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{z}}_y \\ \tilde{\mathbf{z}}_p \\ \tilde{\mathbf{z}}_t \\ \tilde{\mathbf{z}}_s \\ \tilde{\mathbf{z}}_b \end{bmatrix} - \begin{bmatrix} \tilde{\boldsymbol{\nu}}_y \\ \tilde{\boldsymbol{\nu}}_p \\ \tilde{\boldsymbol{\nu}}_t \\ \tilde{\boldsymbol{\nu}}_s \\ \tilde{\boldsymbol{\nu}}_b \end{bmatrix}. \quad (7.2)$$

Note that, because we are linearizing around the mean, $\delta\tilde{x} = 0$ and, therefore, $\tilde{z}_x = 0$. Notably, contrary to the filter developed in Bierman's derivation, we do not make use of the fact that the carrier-phase ambiguities are deterministic and constant in time to set the bottom-left entries of the square root information matrix to zero. While this would largely be possible, our approach to initialization, described in Section 7.5, means that they are not always fully decoupled. A more sophisticated filter implementation could apply this relatively significant computational simplification: that is left for future work. This would further reduce the computational effort needed in performing the full filter update.

This equivalent representation of the covariance-estimate information in terms of a data equation is convenient because it has the same representation as the linearized observation equations. Indeed, reconciliation of the a priori information with a given set of observations becomes equivalent to solving the linearized augmented data equation

$$\begin{bmatrix} \tilde{R}_x \\ A_x \end{bmatrix} \delta\hat{x} = \begin{bmatrix} \tilde{z}_x \\ z \end{bmatrix} - \begin{bmatrix} \tilde{\nu}_x \\ \nu \end{bmatrix}. \quad (7.3)$$

Here, z is the set of a priori observation residuals, and A are the Jacobians of these residuals with respect to the state parameters at the *initial* time, or $\partial z / \partial x(t_0)$. The observation residuals are normalized to have unity covariance through division by their expected standard deviations σ_z : details on estimation of this standard deviation can be found in Section 4.4.

The measurement Jacobian can be computed relatively straightforwardly by splitting its components into each of the different state parameter categories:

$$A_y = \frac{\partial z}{\partial y(t)} \Phi_y, \quad A_p = \frac{\partial z}{\partial y(t)} \Phi_p, \quad A_t = \delta_t, \quad A_s = \delta_s, \quad A_b = \delta_b.$$

Here, δ_t , δ_s , and δ_b are matrices whose components are one if the observation corresponding to the component's row contains the corresponding time offset, SISRE, or ambiguity, and zero otherwise. This computation of the measurement Jacobian is the first place where the components of the matrizant come into play: to relate changes in dynamic parameters and the initial state to differences in the expected observations. It is this part of the filter that is analogous to the differential correction, justifying the use of the differential correction error as performance indicator for the accuracy of a state transition matrix.

It turns out to be computationally convenient and numerically effective to solve the augmented data equation using triangularization via Householder transformations:

$$T \begin{bmatrix} \tilde{R}_x \\ A \end{bmatrix} \delta\hat{x} = T \begin{bmatrix} \tilde{z}_x \\ z \end{bmatrix} - T \begin{bmatrix} \tilde{\nu}_x \\ \nu \end{bmatrix} \Rightarrow \begin{bmatrix} \hat{R} \\ 0 \end{bmatrix} \delta\hat{x} = \begin{bmatrix} \hat{z}_x \\ \epsilon \end{bmatrix} - \begin{bmatrix} \hat{\nu}_x \\ \nu_\epsilon \end{bmatrix}. \quad (7.4)$$

This triangularized form of the augmented data equation permits solution for $\delta\hat{x}$ via the triangular inverse of the a posteriori information square root $\hat{R}_x$. This is the crux of the SRIF: instead of propagation of the state error covariance P_x , the square root R_x of its inverse is propagated [5].

Triangularization of the resulting augmented data equation as in Equation (7.4) is done in two steps, making use of the special properties of the augmented information matrix obtained through the choice of state formulation [5]. This starts with a Householder transformation to triangularize the information matrix components corresponding to the orbital state y , dynamic parameters p , time offsets t , and SISRE-absorbing parameters s :

$$\hat{T}_x \begin{bmatrix} \tilde{R}_y & \tilde{R}_{yp} & \tilde{R}_{yt} & \tilde{R}_{ys} & \tilde{R}_{yb} & 0 \\ \tilde{R}_{py} & \tilde{R}_p & \tilde{R}_{pt} & \tilde{R}_{ps} & \tilde{R}_{pb} & 0 \\ \tilde{R}_{ty} & \tilde{R}_{tp} & \tilde{R}_t & \tilde{R}_{ts} & \tilde{R}_{tb} & 0 \\ \tilde{R}_{sy} & \tilde{R}_{sp} & \tilde{R}_{st} & \tilde{R}_s & \tilde{R}_{sb} & 0 \\ A_y & A_p & A_t & A_s & A_b & z \\ \tilde{R}_{by} & \tilde{R}_{bp} & \tilde{R}_{bt} & \tilde{R}_{bs} & \tilde{R}_b & 0 \end{bmatrix} = \begin{bmatrix} \hat{R}_y & \hat{R}_{yp} & \hat{R}_{yt} & \hat{R}_{ys} & \hat{R}_{yb} & \hat{z}_y \\ 0 & \hat{R}_p & \hat{R}_{pt} & \hat{R}_{ps} & \hat{R}_{pb} & \hat{z}_p \\ 0 & 0 & \hat{R}_t & \hat{R}_{ts} & \hat{R}_{tb} & \hat{z}_t \\ 0 & 0 & 0 & \hat{R}_s & \hat{R}_{sb} & \hat{z}_s \\ 0 & 0 & 0 & 0 & \hat{R}_b & \hat{z}_b \\ 0 & 0 & 0 & 0 & 0 & \hat{\epsilon} \end{bmatrix}. \quad (7.5)$$

With this triangularization, the augmented data equation is solved. Note is made of the remaining quantity ϵ on the right-hand side of this equation. It consists of the residuals of all observations: these are used in the outlier detection step.

After this triangularization, the augmented data equation is reduced into an equivalent a posteriori data equation of the following form:

$$\begin{bmatrix} \hat{R}_y & \hat{R}_{yp} & \hat{R}_{yt} & \hat{R}_{ys} & \hat{R}_{yb} \\ 0 & \hat{R}_p & \hat{R}_{pt} & \hat{R}_{ps} & \hat{R}_{pb} \\ 0 & 0 & \hat{R}_t & \hat{R}_{ts} & \hat{R}_{tb} \\ 0 & 0 & 0 & \hat{R}_s & \hat{R}_{sb} \\ 0 & 0 & 0 & 0 & \hat{R}_b \end{bmatrix} \begin{bmatrix} \delta \hat{y}(t_0) \\ \delta \hat{p}(t_0) \\ \delta \hat{t}(t_0) \\ \delta \hat{s}(t_0) \\ \delta \hat{b}(t_0) \end{bmatrix} = \begin{bmatrix} \hat{z}_y \\ \hat{z}_p \\ \hat{z}_t \\ \hat{z}_s \\ \hat{z}_b \end{bmatrix} - \begin{bmatrix} \hat{\nu}_y \\ \hat{\nu}_p \\ \hat{\nu}_t \\ \hat{\nu}_s \\ \hat{\nu}_b \end{bmatrix}. \quad (7.6)$$

From this a posteriori data equation, the least-squares adjustment to the initial system state can be obtained as

$$\begin{bmatrix} \delta \hat{y}(t_0) \\ \delta \hat{p}(t_0) \\ \delta \hat{t}(t_0) \\ \delta \hat{s}(t_0) \\ \delta \hat{b}(t_0) \end{bmatrix} = \begin{bmatrix} \hat{R}_y & \hat{R}_{yp} & \hat{R}_{yt} & \hat{R}_{ys} & \hat{R}_{yb} \\ 0 & \hat{R}_p & \hat{R}_{pt} & \hat{R}_{ps} & \hat{R}_{pb} \\ 0 & 0 & \hat{R}_t & \hat{R}_{ts} & \hat{R}_{tb} \\ 0 & 0 & 0 & \hat{R}_s & \hat{R}_{sb} \\ 0 & 0 & 0 & 0 & \hat{R}_b \end{bmatrix}^{-1} \begin{bmatrix} \hat{z}_y \\ \hat{z}_p \\ \hat{z}_t \\ \hat{z}_s \\ \hat{z}_b \end{bmatrix}. \quad (7.7)$$

To transform this a posteriori *initial* state estimate to the equivalent adjustment of the *final* state, the variational model of the state is applied:

$$\begin{bmatrix} \delta \hat{y}(t) \\ \delta \hat{p}(t) \\ \delta \hat{t}(t) \\ \delta \hat{s}(t) \\ \delta \hat{b}(t) \end{bmatrix} = \begin{bmatrix} \Phi_y & \Phi_p & 0 & 0 & 0 \\ 0 & M_p & 0 & 0 & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & 0 & 0 & M_s & 0 \\ 0 & 0 & 0 & 0 & I \end{bmatrix} \begin{bmatrix} \delta \hat{y}(t_0) \\ \delta \hat{p}(t_0) \\ \delta \hat{t}(t_0) \\ \delta \hat{s}(t_0) \\ \delta \hat{b}(t_0) \end{bmatrix} \quad (7.8)$$

Substituting these relations, the a posteriori final state estimate becomes:

$$\begin{bmatrix} \hat{y}(t) \\ \hat{p}(t) \\ \hat{t}(t) \\ \hat{s}(t) \\ \hat{b}(t) \end{bmatrix} = \begin{bmatrix} \tilde{y}(t) \\ \tilde{p}(t) \\ \tilde{t}(t) \\ \tilde{s}(t) \\ \tilde{b}(t) \end{bmatrix} + \begin{bmatrix} \Phi_y & \Phi_p & 0 & 0 & 0 \\ 0 & M_p & 0 & 0 & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & 0 & 0 & M_s & 0 \\ 0 & 0 & 0 & 0 & I \end{bmatrix} \begin{bmatrix} \delta \hat{y}(t_0) \\ \delta \hat{p}(t_0) \\ \delta \hat{t}(t_0) \\ \delta \hat{s}(t_0) \\ \delta \hat{b}(t_0) \end{bmatrix}. \quad (7.9)$$

This is the second place where the state transition matrix is useful: it removes the need for costly repropagation of the orbital state after adjusting for the observations. Effectively, this is the second part of the differential correction discussed before.

7.4. Outlier detection

Before the established least-squares adjustment is actually applied to the state estimate, outlier detection is performed based on the posterior residuals ϵ that remain after combination of the a priori information with the observations. Here, use is made of the fact that these are conveniently computed in the SRIF data processing step at no additional cost: this allows them to be inspected for quality control. To establish the presence of outliers, the maximum-magnitude element $\epsilon_{\max} = \max(|\epsilon|)$ is inspected, as this is most likely to be an outlier, if any exist.

If $\epsilon_{\max}$ exceeds an empirically-determined threshold value k , it is expected to belong to an outlier, and the corresponding measurement is removed from consideration. Explicitly, all other residuals exceeding this threshold are *not* removed, as it is possible that they exceeded the threshold only to compensate for the largest outlier. Then, the augmented data equation is solved again by repeating Equation (7.5), and the maximum-magnitude posterior residual is evaluated again to determine the presence of new outliers. If none remain, or none were found in the first place, the state update is performed according to Equation (7.9), and the remainder of the estimate-covariance propagation is performed.

A value of 20 was found to be a suitable threshold, corresponding to outliers above 20 standard deviations for the observations. This corresponds to about 10-20 cm for the carrier-phase observations. This threshold is, purposefully, relatively large: during poorly-predictable scenarios such as entering or exiting eclipse, it is entirely expected that the estimated system dynamics differ from the measurements

Table 7.1: Initial state-covariance estimates for persistent parameters

Parameter	Estimate	Variance
Position, r	+ 1 m	1000 m ²
Velocity, v	+ 0.1 m/s	200 m ² /s ²
Solar radiation pressure coefficient, c_r	1	0.001
Aerodynamic drag coefficient, c_d	1	0.1
Empirical accelerations, a_e	0 m/s ²	1 pm ² /s ⁴
Receiver clock offset, δt_{rcv}	0 m	500 m ²
Galileo-GPS Time Offset, δt_{GTO}	0 m	10 m ²

by several centimetres. A lower threshold would, in such scenarios, lead the filter to erroneously reject correct measurements because they do not match the inaccurate estimated state. Clearly, more refined outlier detection strategies are possible, but this approach suffices to detect the most important and largest sources of error: erroneous GNSS broadcast ephemerides and loss-of-lock in the phase-lock-loop.

The choice is made to perform this outlier detection *after* the prediction step, using the a posteriori residuals, rather than *before*, using the a priori observation residuals. This can incur additional overhead: for every detected outlier, the entire data processing step has to be repeated. Still, this approach is better at correcting for errors in the a priori state estimates than a filter with a priori outlier detection based on observation residuals. This is useful in correcting for cycle slips or for errors in the dynamical model, which result in errors that are correlated between observations. By applying a posteriori outlier detection, such observations can still be included in state estimation, where a priori outlier detection would likely result in such measurements being discarded. Experimentation showed that the repeated execution of the data processing step, at several milliseconds, does not significantly affect the total runtime of the filter, which is on the order of 50 milliseconds.

7.5. Parameter initialization

In terms of parameter initialization, two kinds of parameters can be distinguished: persistent parameters, which are initialized once at filter start-up, and temporary parameters, which can be created and removed throughout filter operation. Initialization for the first type is of little consequence: it is useful to give the filter a good starting point for convergence, but preliminary results showed that an orbit within several kilometres and decametres per second of the actual estimate sufficed. For the results given in this work, the precise initial orbit is used, offset by 1 m and 0.1 m/s in each direction. The same insensitivity is true for the force coefficients, which are initialized at a value of unity; for the empirical accelerations, which are set to zero; and the time offsets, which are estimated to be zero. The exact values used are given in Table 7.1.

Proper initial variance estimates were chosen through trial-and-error. For the orbital state and time offsets, relatively large values sufficed. The force coefficients, on the other hand, are more sensitive to the initial covariance estimate. The filter tends to overadjust these in the early convergence phase to compensate for observation mismatches, in particular near eclipse entry or exit. As a result, their initial variances are set to relatively low values, to ensure slower convergence. Similarly, the empirical accelerations are initialized at a low variance, even lower than their process noises, to ensure slow adaptation to early observation mismatches. This allows these early mismatches to be captured by changes in the orbital state and time offsets rather than the dynamic parameters.

There are two sets of parameters whose elements are added and removed throughout the process of orbit determination, namely the carrier-phase ambiguity and SISRE-absorbing parameters. Because those correspond to specific observation arcs and tracked satellites, they exist only for the associated durations. In initialization of those parameters, use is made of the SRIF property that a priori covariance-estimate information is equivalent to observations: instead of providing an arbitrary a priori estimate for any new ambiguity or SISRE, the first measurement in which they are encountered is used to initialize the parameter. To prevent the linear dependency of those parameters from creating singularities in solution of the augmented data equation, this is done step-by-step: before carrier-phase measurements, which need estimates of the ambiguity, are included, the equivalent SISRE must already be initialized using a

pseudorange measurement. Any carrier-phase measurements observed before this first pseudorange observation are discarded.

7.6. Parameter removal

To prevent unbounded growth of the filter memory footprint over time, it is necessary to remove filter parameters that will no longer be used. This goes in the first place for the carrier-phase ambiguities, which are fundamentally valid only for a single observation arc. Similarly, the SISRE-absorbing parameters are also removed when they are no longer encountered in an observation batch. While it could be justified that the SISRE for a given satellite conceptually persists over time, its estimate is also invalidated every time a new navigation message is distributed for a GNSS satellite. The timing of this navigation message update differs per constellation, but is on the order of several hours: accordingly, the SISRE estimate for any given GNSS satellite is invalidated nearly every orbit of Sentinel-6A, such that there is little advantage to storing an old estimate as new a priori for the next observation arc.

Due to the usage of a square root filter formulation, removal of a parameter is not as straightforward as removal of the corresponding rows and columns in the covariance or information matrix. Instead, a separate approach is derived to achieve this same effect through manipulation of the square root information matrix itself. To achieve this, the a posteriori information matrix $\hat{C}_x$ and its square root $\hat{R}_x$, which is known to be upper-diagonal after data processing, are partitioned into 3×3 block matrices. These blocks must be manipulated such that the row and column of $\hat{C}_x$ corresponding to the to-be-removed parameter become zero:

$$\hat{C}_x = \begin{bmatrix} C_{11} & C_{12} & C_{13} \\ C_{12}^T & C_{22} & C_{23} \\ C_{13}^T & C_{23}^T & C_{33} \end{bmatrix} = \begin{bmatrix} R_{11} & R_{12} & R_{13} \\ 0 & R_{22} & R_{23} \\ 0 & 0 & R_{33} \end{bmatrix}^T \begin{bmatrix} R_{11} & R_{12} & R_{13} \\ 0 & R_{22} & R_{23} \\ 0 & 0 & R_{33} \end{bmatrix}, \quad (7.10)$$

$$\hat{C}'_x = \begin{bmatrix} C_{11} & 0 & C_{13} \\ 0 & 0 & 0 \\ C_{13}^T & 0 & C_{33} \end{bmatrix} = \begin{bmatrix} R'_{11} & R'_{12} & R'_{13} \\ 0 & R'_{22} & R'_{23} \\ 0 & 0 & R'_{33} \end{bmatrix}^T \begin{bmatrix} R'_{11} & R'_{12} & R'_{13} \\ 0 & R'_{22} & R'_{23} \\ 0 & 0 & R'_{33} \end{bmatrix}. \quad (7.11)$$

By equating the four corner submatrices C_{11} , C_{13} , C_{13}^T , and C_{33} of the information matrices, we find $R'_{11} = R_{11}$, $R'_{13} = R_{13}$, $R'_{12} = 0$, $R'_{22} = 0$, $R'_{23} = 0$, and

$$R_{33}^T R'_{33} = R_{33}^T R_{33} + R_{23}^T R_{23}.$$

The last expression can be computed without explicit construction of $R_{33}^T R'_{33}$, $R_{33}^T R_{33}$, or $R_{23}^T R_{23}$ by means of a rank-one update of the Cholesky factorization. In this way it is possible to manipulate R_{33} in-place to achieve the same effect without explicitly needing to construct $\hat{C}_x$, which is computationally intensive and introduces numerical error. After this, the square root information matrix $\hat{R}'_x$ can be shrunk to remove the now-zero components R'_{12} , R'_{22} , R'_{23} entirely. Finally, the parameter estimate itself can be removed directly from the state estimate δx , completing the parameter removal.

7.7. Covariance propagation

The steps described in Section 7.3 provide us with an a posteriori estimate of the final *state*, but not yet of the *error covariance* at that point. Propagation of this error covariance estimate is done in two steps: first, its deterministic evolution is computed. Afterwards, the effect of process noise is included to obtain the final covariance estimate, which functions as new a priori estimate for the next filter update.

The deterministic propagation follows from the dynamical model of the state variations, given already in Equation (7.8). Using the same approach as Bierman, the deterministically-propagated square root information matrix is obtained by substituting Equation (7.8) [5], to obtain:

$$\begin{bmatrix} \hat{R}_y \Phi_y^{-1} & \hat{R}_{yp} M_p^{-1} - \hat{R}_y \Phi_y^{-1} \Phi_p M_p^{-1} & \hat{R}_{yt} & \hat{R}_{ys} M_s^{-1} & \hat{R}_{yb} \\ 0 & \hat{R}_p M_p^{-1} & \hat{R}_{pt} & \hat{R}_{ps} M_s^{-1} & \hat{R}_{pb} \\ 0 & 0 & \hat{R}_t & \hat{R}_{ts} M_s^{-1} & \hat{R}_{tb} \\ 0 & 0 & 0 & \hat{R}_s M_s^{-1} & \hat{R}_{sb} \\ 0 & 0 & 0 & 0 & \hat{R}_b \end{bmatrix}. \quad (7.12)$$

It is noted here that the matrices M_p and M_s , the scale factors for the corresponding exponentially-correlated process noise parameters, are diagonal and therefore trivially invertible. This deterministic propagation of the error ellipsoids is the third place in the filter where the orbital state transition and sensitivity matrices are used. Indeed, the performance of an approximation to the matrixant in specifically this part of the filter is assessed by means of comparison of the error ellipsoids.

After the deterministic covariance propagation, the process noises w_y, w_p, w_t , and w_s are included as a separate step. These process noises describe the differences between the deterministically-determined state estimate and the stochastically perturbed equivalent. For this, the square root noise matrices $\hat{R}_{wy}, \hat{R}_{wp}, \hat{R}_{wt}$, and $\hat{R}_{ws}$ are introduced, being the Cholesky factorizations of the inverse covariance matrices of the corresponding process noises. For the purposes of this work, all these covariance matrices are diagonal. Consequently, their square root noise matrices are computed trivially as the element-wise inverse square roots of the diagonal components.

The effect of these process noises is incorporated by considering the a priori data equations of the noises [5]:

$$\begin{bmatrix} \hat{R}_{wy} w_y \\ \hat{R}_{wp} w_p \\ \hat{R}_{wt} w_t \\ \hat{R}_{ws} w_s \end{bmatrix} = \begin{bmatrix} \hat{z}_{wy} \\ \hat{z}_{wp} \\ \hat{z}_{wt} \\ \hat{z}_{ws} \end{bmatrix} - \begin{bmatrix} \nu_{wy} \\ \nu_{wp} \\ \nu_{wt} \\ \nu_{ws} \end{bmatrix}. \quad (7.13)$$

Because, a priori, we expect the impact of the noise to be zero-mean, $\hat{z}_{wx} = 0$. Similarly, ν_{wx} is a random variable with zero mean and unity variance. Note the exclusion of the ambiguities in this system of equations, since $\hat{b}(t) = \hat{b}^d(t)$.

Combining Equations (7.12) and (7.13) leads to an augmented system, which we solve through triangularization via Householder transformations:

$$\tilde{T}_{ypts} \begin{bmatrix} -\hat{R}_{wy} & 0 & 0 & 0 & \hat{R}_{wy} & 0 & 0 & 0 & 0 \\ 0 & -\hat{R}_{wp} & 0 & 0 & 0 & \hat{R}_{wp} & 0 & 0 & 0 \\ 0 & 0 & -\hat{R}_{wt} & 0 & 0 & 0 & \hat{R}_{wt} & 0 & 0 \\ 0 & 0 & 0 & -\hat{R}_{ws} & 0 & 0 & 0 & \hat{R}_{ws} & 0 \\ \hat{R}_y \Phi_y^{-1} & \hat{R}_{yp} M_p^{-1} - \hat{R}_y \Phi_y^{-1} \Phi_p M_p^{-1} & \hat{R}_{yt} & \hat{R}_{ys} M_s^{-1} & 0 & 0 & 0 & 0 & \hat{R}_{yb} \\ 0 & \hat{R}_p M_p^{-1} & \hat{R}_{pt} & \hat{R}_{ps} M_s^{-1} & 0 & 0 & 0 & 0 & \hat{R}_{pb} \\ 0 & 0 & \hat{R}_t & \hat{R}_{ts} M_s^{-1} & 0 & 0 & 0 & 0 & \hat{R}_{tb} \\ 0 & 0 & 0 & \hat{R}_s M_s^{-1} & 0 & 0 & 0 & 0 & \hat{R}_{sb} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \hat{R}_b \end{bmatrix} \quad (7.14)$$

$$= \begin{bmatrix} * & * & * & * & * & * & * & * & * \\ 0 & * & * & * & * & * & * & * & * \\ 0 & 0 & * & * & * & * & * & * & * \\ 0 & 0 & 0 & * & * & * & * & * & * \\ 0 & 0 & 0 & 0 & \tilde{R}_y & \tilde{R}_{yp} & \tilde{R}_{yt} & \tilde{R}_{ys} & \tilde{R}_{yb} \\ 0 & 0 & 0 & 0 & \tilde{R}_{py} & \tilde{R}_p & \tilde{R}_{pt} & \tilde{R}_{ps} & \tilde{R}_{pb} \\ 0 & 0 & 0 & 0 & \tilde{R}_{ty} & \tilde{R}_{tp} & \tilde{R}_t & \tilde{R}_{ts} & \tilde{R}_{tb} \\ 0 & 0 & 0 & 0 & \tilde{R}_{sy} & \tilde{R}_{sp} & \tilde{R}_{st} & \tilde{R}_s & \tilde{R}_{sb} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \tilde{R}_b \end{bmatrix}.$$

The bottom-right submatrix of the result becomes the updated square root information matrix, which can be used as a priori for the next data processing step. All elements indicated by an asterisk contain unused but non-zero values: those could be used for smoothing, but that is not done in the present application.

7.8. Filter parameter tuning

There are three types of free parameters that must be tuned in the filter design: the process noises, exponentially-correlated process noise correlation times, and observation batch sizes. This parameter tuning was done twice, separately for the filters using broadcast GPS and Galileo ephemerides and for those using precise GPS antenna ephemerides. The resulting values are summarized in Table 7.2.

Table 7.2: SRIF parameters used for orbit determination

Parameter	Configuration	
<i>Observations</i>		
GNSS ephemerides	Broadcast GPS/Galileo	Precise GPS
Batch size	10	10
<i>Correlation times</i>		
Empirical accelerations	600 s	600 s
SISRE	6000 s	6000 s
<i>Process noise standard deviations</i>		
Radial empirical acceleration	0.75 nm/s ²	0.75 nm/s ²
Along-track empirical acceleration	0.5 nm/s ²	0.5 nm/s ²
Cross-track empirical acceleration	1 nm/s ²	1 nm/s ²
SISRE	0.8 mm	0.6 mm
Receiver clock offset	5 m	5 m
GGTO	5 cm	5 cm


The justification for the correlation times was given before, in Section 7.2; a justification for the chosen process noises and batch sizes follows below.

For the batch size between data processing updates, intervals ranging from 5 to 120 s were tried. This covers the range of batch sizes used by DIODE, RTNAV, and the Haiyang-2D filters, at 10, 60, and 30 s, respectively [34, 38]. For the broadcast ephemerides configuration, the best performance follows for an observation interval of 5 s; for the precise GPS ephemerides configuration, this is 10 s. This difference makes sense: with the reduced set of observations that is considered in the precise GPS filter, a longer observation interval is needed to obtain the same number of observations in one batch update.

Suitable values for the process noises are determined by iterative reduction until filter divergence or significant degradation of the filter performance. This results in the process noise standard deviations as given in Table 7.2. It is noted that not all stochastic parameters are listed: for the unlisted parameters, it was found that the presence of process noise did not significantly improve filter performance. Hence, it was set to 10^{-20} for those values, practically resulting in deterministic behaviour of those variables.

In the observation model, we apply three kinds of process noises: in the SISRE-absorbing parameters, the receiver clock offset, and the GGTO. The standard deviations for the GGTO follow to be 5 cm, corresponding to around 0.2 ns in units of time. This is slightly larger than the expected stability of the GPS-Galileo system time offset: over 4 January 2021, Montenbruck et al. find this offset to have a standard deviation of 2 cm [43]. For the receiver clock offset, significantly larger values of 5 m are obtained: this makes sense, as the Sentinel-6A receiver clock is expected to be significantly less stable in time than the drift between two collectives of stable GNSS clocks. In essence, this magnitude of process noise standard deviation makes the receiver clock offset a free variable. This is fine, as it is effectively constant for all observations in a single observation batch: this makes it highly observable. Finally, the SISRE-absorbing parameters obtain standard deviations of 0.8 and 0.6 mm. This is in line with the Haiyang-2D filter: with a similar configuration, it applies standard deviations in the SISRE process noise of 1 mm and 3 mm for BDS-3 and GPS observations, respectively [34].

Among the dynamic parameters, process noise is applied only for the empirical accelerations. As visible from Table 7.2, these have standard deviations of 0.75, 0.5, and 1 nm/s² in the radial, along-track, and cross-track directions, respectively. This is slightly larger than values found among state-of-the-art filters: Li et al. report values of 0.05 nm/s² every 30 seconds for the Haiyang-2D filter [34], and Darugna et al. applied values of 0.01-0.05 nm/s² in the radial and along-track directions, and 0.12-0.40 nm/s² in the cross-track direction every 10 seconds, depending on the configuration [14]. This can probably be attributed at least in part to the neglect of ocean tides in our dynamic model, which slightly reduces the quality of our dynamic model with respect to those filters.



Results

To provide a real-world context within which to assess the quality of the Hill-Clohessy-Wiltshire approximation, a case study was performed of orbit determination of Sentinel-6A on 1 January 2022. The results of this case study can be found in this chapter. First, in Section 8.1, we present the overall filter performance. This serves to place the differences presented afterwards in context relative to the overall filter quality. Secondly, the differences in orbit solutions are compared for each of the state transition matrices in Section 8.2. This quantifies the impact of each simplification in the estimated state transition matrices, separate from the shared behaviour that dominates the total orbit error. Afterwards, Section 8.3 reports the difference in estimated error ellipsoids between each of the state transition matrix approximations. This is complemented in Section 8.4 by the found transition matrix qualities in terms of differential correction errors. The last characterization of the state transition matrix for this case study is the runtime associated with each: these are given in Section 8.5.

8.1. Filter performance

The overall filter performance is established from three points of view. In the first place, we assess the accuracy achieved by the different filter configurations through comparison with a precise orbit. As a secondary assessment, the a priori and a posteriori residuals are inspected to determine their consistency with the expected behaviour. Thirdly, the resulting dynamical behaviour is presented, to assess their consistency with the expected trajectory model behaviour. These two measures serve primarily to verify the correctness of the filter implementation and to contextualize the accuracy with which later results can be interpreted. Explicitly, we do *not* present differences in results between different state transition matrix approximations here: those can be found in later sections.

Error with respect to precise orbit

The primary viewpoint from which to assess the efficacies of the different filter implementations, naturally, is their overall performance in estimating the Sentinel-6A orbit. We shall assess this in the first place based on their actual accuracy with respect to the precise TU Delft orbit. This solution, produced using GipsyX, covers a period of three hours before and after 1 January 2022. It has root-mean-square observation residuals of 5.10 mm for carrier-phase measurements and 0.14 m for code measurements, suggesting notably higher accuracy than the centimetre-level orbits possible with real-time filters. This makes it a suitable precise orbit to compare against.

As will be evident from the inter-filter differences presented in Section 8.2, the differences between different state transition matrix approximations are marginal with respect to the total filter error and, consequently, do not lead to qualitatively different error behaviour. Hence, only the results for the filter with the reference state transition matrix are given in Figure 8.1; in principle, all observations on the overall filter behaviour hold equally well for the simplified variational equations. Hence, the remainder of this section does *not* serve to compare the difference in orbit quality between different state transition matrix estimates. Instead, it serves to characterize the overall orbit quality, as needed to contextualize the certainty with which conclusions can be drawn on the significance of inter-filter differences. To this end, we also present results from a filter that uses precise GPS ephemerides, because they are




Figure 8.1: Filter positional errors and their normalized frequency distribution when using broadcast GNSS (blue) and precise GPS ephemerides (red)

qualitatively different. For all data presented later in this chapter, this is not the case: unless specified otherwise, the presented results refer to a filter with broadcast GNSS ephemerides, which is more representative of a real-time onboard scenario.

Where the choice of matrizant does not significantly impact the overall filter performance, the choice of ephemerides does. As expected, the usage of precise GPS ephemerides, in spite of fewer observations compared to the combined GPS and Galileo broadcast ephemerides, does result in improved filter performance, albeit marginal. Overall, both sets of ephemerides exhibit similar error characteristics: generally, errors have the same sign, and certain error peaks are clearly related, for example at 4 h in the along-track direction. Still, distinct differences in performance are evident. The means and standard deviations of each error component are given above each subfigure of Figure 8.1. Note that these values exclude the first hour of data, to reduce transient effects. Generally, the standard deviation of the error using precise GPS ephemerides can be seen to be lower; only in the cross-track component is a slight degradation in performance visible. Additionally, we observe that the precise GPS filter has relatively significant biases in both radial and along-track direction, at around 5 cm: these seem to be the result of transient behaviour, as their magnitudes reduce when a longer convergence time is used before errors are included in the statistics.

When considering the Fourier transforms, given on the right-hand side of Figure 8.1, several additional differences in behaviour become clear. In particular, we see that usage of broadcast GNSS ephemerides results in relatively large error peaks near an orbital frequency of one, especially in the radial and cross-track direction. This implies that the filter is more dynamic than kinematic in nature, as it is errors in the dynamic model specifically that appear at a once-per-orbit frequency. Usage of precise GPS ephemerides, on the other hand, does not have such once-per-orbit components. Indeed, this suggests that the precise GPS filter is more kinematic in nature. This makes sense, considering that the higher quality of the resulting observations allows the precise GPS filter to extract more information from each measurement instead of following the a priori state estimate.

Overall, it must be noted that the filter does not yet meet the performance previously demonstrated by other filters, such as RTNAV and DIODE: these achieve radial accuracies of as low as 3 cm for Sentinel-6A, with similar observation qualities. It is likely that this is caused primarily by the absence of an ocean tide model in the trajectory model, in combination with the neglect of carrier-phase




Figure 8.2: Mean magnitude of pseudorange (left) and carrier-phase (right) residuals per observation batch, normalized as a multiple of the estimated standard deviation

wind-up in the observation model. Consequently, this slightly larger error does not preclude the filter from being representative of real-time orbit determination performance of state-of-the-art filters. Indeed, it is not expected that any of the two aforementioned limitations of the filter lead to qualitatively different dependence on the state transition matrix.

Consistency of residuals

To ascertain that the filter is properly tuned, and to verify the expected dynamic and kinematic nature of all filters, the a priori and a posteriori observation residuals are also compared. As with the orbit error, only the reference filters with full variational equations are presented here: the qualitative behaviour is not affected by approximate state transition matrices. The results are presented in Figures 8.2a and 8.2b, each split into pseudorange and carrier-phase observations due to their qualitatively different behaviour. To reduce noise, the mean magnitude per observation batch is graphed, and not individual observations: the trend is of interest here, not individual residuals. Additionally, it is noted that the residual of each observation is expressed in a normalized manner, as a multiple of its estimated standard deviation: this is also the scale with which it is represented inside the SRIF itself. Finally, the figures refer to the residuals after the outlier detection step of the filter, such that those do not affect the presented data: still, at an average of 0.4% of observations, the fraction of rejected observations is negligible.

In Figure 8.2a, the residuals are shown for the reference filter using broadcast GNSS ephemerides. As expected, its posterior residuals are near one standard deviation in magnitude, indicating that the estimated observation standard deviations match reality. It is visible that the a priori residuals of the carrier-phase observations are significantly larger than those of the pseudorange observations: this is expected due to their smaller standard deviation, which means that differences in the a priori estimate due to errors in the dynamical model are relatively of larger impact on the carrier-phase estimates. Additionally, in part due to the same reason, it is visible that the reduction in residuals after the measurement update is more significant for the carrier-phase measurements: more information is extracted from them.




Figure 8.3: Distribution of the estimated empirical accelerations when using the full state transition matrix approximation

The residuals for the reference filter with precise GPS antenna ephemerides is shown in Figure 8.2b. Overall, it shows similar behaviour to the broadcast GNSS filter, with posterior residuals on the order of one standard deviation in magnitude. One stark contrast with the broadcast GNSS filter, though, is the significantly larger reduction in carrier-phase residuals achieved in each measurement update. This indicates that the filter is able to extract significantly more information from its carrier-phase measurements when using precise GPS ephemerides: this makes sense, considering the predominantly kinematic behaviour found in the analysis of Figure 8.1, compared with the more dynamic behaviour of the filter when using broadcast GNSS ephemerides.

Overall, it can be concluded that the posterior residuals for both choices of ephemerides, all near one standard deviation in magnitude, are consistent with a converged and properly tuned filter. Additionally, the observed reduction in residuals is in line with the suspected dynamic and kinematic nature of each filter deduced from the error behaviour. This suggests that the broadcast GNSS filter depends more on its trajectory model, while the precise GPS filter is able to rely to a higher degree on its observation model.

Consistency of dynamical model

As final measure of filter quality, the behaviour of the dynamical model is considered. To this end, we present the estimated dynamic parameters, to characterize the ability of the dynamical model to capture the physical reality. For the empirical accelerations, we are interested in their distribution and whether it matches the known trajectory model deficiencies. For the aerodynamic drag and radiation pressure coefficients, their convergence is of primary interest, as they indicate whether the dynamical model can adequately discern these accelerations from other, unmodelled accelerations.

We start with the distribution of empirical accelerations, given in Figure 8.3. The first hour of estimates is excluded, because the estimated empirical accelerations are much larger during filter convergence. We contrast this distribution with the analysis presented earlier in Section 5.5: there, unmodelled force model terms were found to account for relative accelerations with means and standard deviations of $0.1 \pm 1.5 \text{ nm/s}^2$, $-0.4 \pm 0.9 \text{ nm/s}^2$, and $0.0 \pm 0.8 \text{ nm/s}^2$ in radial, along-track, and cross-track direction, respectively.

These expectations are matched quite well by the standard deviations of 4.6 and 5.1 nm/s^2 of the estimated empirical radial and along-track accelerations. They are a factor 4-5 higher than the 0.9 and 1.5 nm/s^2 that can be attributed solely to known simplifications in the force model. Yet, the extra accelerations fall within the range that is attributable to force mismodelling and the known lack of ocean tides. Their biases of -2.2 and 0.9 nm/s^2 fall entirely within the range that can be expected based on mismodelling of the Earth radiation pressure and aerodynamic drag. The cross-track empirical accelerations tend to be a factor of 3 larger still, with a standard deviation of 16.7 nm/s^2 . Part of this is explainable by misestimation of the radiation pressure; the remainder seems to be caused by other unidentified modelling errors.

The evolution of the aerodynamic drag and radiation pressure coefficient estimates can be found in Figure 8.4. Both eventually settle around a value of 1.5. This is quite high: in comparison, the TU Delft precise orbit estimate has values of 1.00 ± 0.00 for the aerodynamic drag coefficient and 1.04 ± 0.03 for the radiation pressure coefficient. This can be brought back in large part to uncertainty in our reference




Figure 8.4: Estimated values of the radiation pressure and aerodynamic drag coefficient

area estimates, which are relatively arbitrary. Hence, this is not necessarily a deficiency in the force model.

Of larger interest is how well the filter can actually identify these expected accelerations in the observed system dynamics, i.e., how well it can converge to single coefficient values. For the Solar radiation pressure coefficient, it is clear that a converged value is eventually achieved, although the convergence period is exceedingly long, at multiple hours. The aerodynamic drag coefficient approaches its equilibrium value faster, but has trouble converging past that. This is not necessarily unexpected: with a magnitude in the order of 1 nm/s^2 , the aerodynamic acceleration is quite difficult to observe. Overall, it can at least be seen that both estimates approach reasonable values, indicating that the filter can indeed discern the aerodynamic drag and radiation pressure from other accelerations.

8.2. Inter-solution differences

As noted when describing the overall filter errors, the differences between different state transition matrix formulations are comparatively small. This is illustrated in Figure 8.5, where the reader can find the differences with respect to the full-fidelity solution resulting from application of an approximate state transition matrix. The mean and standard deviation of each component is given in the legend, excluding the first hour of data to remove any bias due to differences in convergence behaviour. The values given here were obtained using broadcast GNSS ephemerides, which are most representative of a real-time scenario. Results from the precise GPS filter are not meaningfully different. Hence, they will not be presented here: the interested reader can find those results in Appendix B.

As visible, these differences reduce to less than a millimetre in all three directions for all state transition matrix approximations. Both the Hill-Clohessey-Wiltshire and spherical Earth simplifications do result in a small but non-negligible offset of several millimetres in both along-track and radial direction. For both of these approximations, the differences in radial and along-track direction are flipped in sign and clearly correlated in time; this suggests that they are manifestations of the same dynamical difference, simply related through the known linked behaviour of perturbations in either direction. The periodic component of this radial-along-track error behaviour is consistent with the difference expected from neglecting either one or both of the J_2 effect and ellipticity. Indeed, it seems that both are relevant to approximately one millimetre for the orbit of Sentinel-6A orbit.

With deviations of at most one millimetre in magnitude, it becomes clear that all three approximations can permit orbit determination accuracies down to several centimetres. Still, it is evident that the oblate Earth model is significantly more accurate. At differences of mere micrometres in magnitude, this approximation to the transition matrix could support orbit determination down to the level of millimetres. The presented implementation of the Hill-Clohessey-Wiltshire approximation, on the other hand, would not support this: only accuracies down to a centimetre are to be expected.

It is also notable that the differences between solutions are only slightly bigger during the convergence phase of the filter. Even during this phase, the differences between solutions are insignificant with respect to the filter error. Particularly during the first 30 minutes, where it is evident from Figure 8.1 that the filter error is still in the order of metres, the inter-solution differences are at most 5 mm in each component. This suggests that the filter's convergence behaviour is not affected by any of the approximations.




Figure 8.5: Differences in final position estimate as a result of different approximations to the variational equations

8.3. Error ellipsoids

The effect of different approximations to the matizant on the estimated uncertainty in position and velocity is shown in Figures 8.6a and 8.6b, respectively. As with the inter-solution differences, the choice of GNSS ephemerides does not meaningfully change behaviour, so only the broadcast GNSS ephemerides results are shown here. The results from the precise GPS filter can be found in Appendix B. Overall, the results are in line with the behaviour observed in terms of the inter-solution difference. In terms of orientation for both position and velocity, the Hill-Clohesy-Wiltshire and spherical Earth approximations result in similar errors, at a difference of under 0.1° : the oblate Earth approximation improves on this even further, by two orders of magnitude.


In terms of error covariance volume, a significant departure is visible from the nearly identical behaviour shown by the inter-solution differences of the Hill-Clohesy-Wiltshire and spherical Earth approximations. Indeed, the effect of neglecting orbital ellipticity shows to be significantly larger here than the effect of the J_2 zonal harmonic: the Hill-Clohesy-Wiltshire approximation results in oscillations of up to 0.2% around the correct uncertainties in position and velocity. For the velocity, it even results in a consistent overestimation of the uncertainty. This makes sense: using a Hill-Clohesy-Wiltshire state transition matrix, the filter is unable to capture the periodic contraction and expansion of the orbit. It seems like this “error” is absorbed by an overall increase of the uncertainty in velocity.

This supposed error source is also in line with the frequency of this difference in volume and orientation, which is around once every hour, or twice per orbital revolution. Indeed, this contraction and expansion happens twice every orbit: once when transitioning from apoapsis to periapsis, and once vice versa. Similar error behaviour can be seen at a smaller amplitude in the spherical Earth model of the variational equations: indeed, a similar contraction-expansion behaviour is expected here due to the periodic nature of the J_2 acceleration.


8.4. Differential correction error

As final accuracy metric, the differential correction error of each of the different state transition matrix approximations is considered. For the full state transition matrix, consisting of both the orbital state transition and sensitivity matrices, this measure is illustrated in Figure 8.7a. Here, significantly larger differences become evident between the Hill-Clohesy-Wiltshire and spherical Earth approximations.

In the first place, the difference in magnitude is striking. Where the inter-solution differences and error



(a) For the position error



(b) For the velocity error

Figure 8.6: Difference in volume and orientation of the error covariance ellipsoids with respect to the full variational equations




Figure 8.7: Differential correction error of the full matrizant, and the contributions of its orbital state transition and sensitivity matrices

ellipsoid geometries as well as the actual orbit solution errors of the Hill-Clohessy-Wiltshire and spherical Earth approximations are relatively similar, their differential correction errors are of entirely different magnitudes. Evidently, this has relatively little impact on the filter performance. This makes sense: the differential correction error does not incorporate the effects of error sources such as non-linearity and unmodelled or mismodelled dynamics on the quality of a state update. It seems that the effect of these factors already starts to dominate the error of the state update at the 10% differential correction error achieved by the Hill-Clohessy-Wiltshire approximation. Consequently, below this threshold, the differential correction error is apparently no longer representative of the actual state update quality.

Splitting the state transition matrix into its constituent submatrices, it becomes evident that the primary cause for the differential correction error of the Hill-Clohessy-Wiltshire approximation is to be found in its sensitivity matrix. From Figure 8.7c, it is visible that errors in this submatrix' components are responsible for the largest part of the full matrix error. In contrast, it is evident from Figure 8.7b that the error contribution of the orbital state transition matrix is several orders of magnitude lower. This serves as indication of the “weakest link” in the Hill-Clohessy-Wiltshire approximation: on improvement of the trajectory estimate quality, it is likely that it is the simplified approximation of the sensitivity matrix that will be the largest source of error in the differential correction. Similarly, the sensitivity matrix of the other two approximations is also one order of magnitude worse than their orbital state transition matrix; at differential correction errors under 10^{-5} , this is less significant.

8.5. Runtimes

Lastly, the distribution of runtimes needed for computation of each approximate state transition matrix is presented. It is reiterated that these measurements are indicative only very roughly of the comparative computational effort: this benchmark was performed on a shared machine, and is based on code that was written with performance only as secondary consideration. Moreover, the machine on which it was run is not representative of the computational machinery available in an onboard scenario. Considering all aforementioned points, the overall distribution of runtimes does still provide meaningful insight into the *relative* computational effort.

In Figure 8.8, the distribution of runtimes is given for a full filter update step over a ten-second interval. These measurements were obtained during the orbit determination run of Sentinel-6A using its GNSS observations from 1 January 2022 of which the results were presented in this chapter; hence, they are representative of the relative difference in runtime during actual filter operation. It is directly evident that computation of the state transition matrix through numerical integration of the full variational equations, requiring between 80 and 90 ms, is exceedingly expensive at double the computational effort of the other approximations. The runtime of the spherical and oblate Earth models, both just under 60 ms, are significantly smaller. The Hill-Clohessy-Wiltshire approximation is faster to compute still, requiring just under 50 ms. These values are in line with expectations: the full variational model requires evaluation of all derivatives in the force model; the spherical and oblate Earth models, in turn, only require evaluation of one or a few spherical harmonics, the difference being negligible; the Hill-Clohessy-Wiltshire approximation, finally, circumvents integration of these derivatives entirely.

To contextualize what fraction of the runtime can be attributed to the filter update step and what




Figure 8.8: Distribution of runtimes for ten-second propagation and filter update steps over 24 hours of data processing

fraction to trajectory propagation and solution of the variational equations, Figure 8.9 is given. This runtime fraction determines how much room for improvement there still is in more efficient computation of the state transition matrix. Figure 8.9 contains the distribution of runtimes for pure orbit propagation, without filter processing such as inclusion of observations or covariance propagation. To indicate what fraction of runtime is due to computation of the state transition matrices, numbers are also included for state propagation only without estimation of the state transition matrix.

From these measurements, it is evident that the filter processing itself requires around 25 ms of runtime, and the state propagation another 25 ms; state transition matrix computation is responsible for the remainder. This provides an important result: the computational runtime expended on the Hill-Clohessey-Wiltshire solution itself is negligible, at mere milliseconds on top of the state propagation itself. It follows that further reduction of the computational effort beyond the Hill-Clohessey-Wiltshire approximation, while certainly feasible, is relatively meaningless: the effort is better spent reducing the computational effort of the trajectory propagation itself, or of the data inclusion step in the filter.




Figure 8.9: Distribution of runtimes for only ten-second propagations over 24 hours of data processing

9

Discussion

From the results of the case study, specifically in terms of differences between different models of the variational equations, several observations can be made. In this chapter, we compare and contrast the results obtained using the different metrics and discuss the implication that these results have on the research questions posed in Chapter 2. In Section 9.1, we start with the match and mismatch between the performance indicators used to characterize the state transition matrix quality and the actual observed orbit quality. Afterwards, Section 9.2 follows with a discussion on the effects of certain force model simplifications on the state transition matrix accuracy. Finally, in Section 9.3, we reflect on the limitations of this case study, and in particular on the effect that they have on the conclusions that can be derived from this data.

9.1. Agreement between state transition matrix accuracy metrics

In the presented work, both the geometry of the error covariance ellipsoids and the differential correction error were applied as indicators of the orbit determination quality to expect for a given state transition matrix approximation. Comparison of the quality indicated by these metrics with the actual observed trajectory estimate quality is useful from two points of view. In the first place, it serves as verification of the internal consistency of results: the relative performance of each matrix approximation is expected to match the relative trajectory quality. As a secondary purpose, these indicators are metrics by which we can assess the quality of a state transition matrix in isolation from the full filter performance. Consequently, these metrics can be used to answer research subquestion 1. This is of little value in assessing the specific quality of the Hill-Clohessy-Wiltshire approximation, but *is* useful in helping quantify the quality of state transition matrix needed for other orbital accuracies in future work.

First, we consider the error ellipsoid geometries. By comparing the inter-solution differences found in Section 8.2 with the differences in error ellipsoids presented in Section 8.3, we can establish, overall, that the results are in good agreement with each other. In terms of both error ellipsoid orientation and volume, the quality indicated by the ellipsoid geometry is about two orders of magnitude worse for the Hill-Clohessy-Wiltshire and spherical Earth models than for the oblate Earth model. The quality of the Hill-Clohessy-Wiltshire and spherical Earth models as suggested by the error ellipsoid geometries, on the other hand, are similar. This is in perfect agreement with the actually-observed inter-solution differences, suggesting that the resulting error ellipsoid orientation and volume are indeed good indicators for the quality of a state transition matrix approximation within a filter.

The differential correction error metric also correctly discerns the fact that the oblate Earth model is more accurate than the spherical Earth model, which is in turn more accurate than the Hill-Clohessy-Wiltshire approximation. In terms of actual relative orbit determination performance, however, the differential correction error is less reflective of actual performance. Where the spherical Earth approximation, for example, has a differential correction error that is around two to three orders of magnitude smaller than that of the Hill-Clohessy-Wiltshire approximation, their actual orbit determination performances are similar. This is not unexpected: the influence of non-linearities and unmodelled orbital mechanics, such as ocean tides, cannot be accounted for in this metric. Below the differential correction error of 0.1 of the Hill-Clohessy-Wiltshire approximation, these error sources apparently start to dominate

the actual error in the differential correction. This is in line with the observation from Section 8.2 that the convergence behaviour of the filter is left relatively unaffected by the simplified state transition matrices. Apparently, convergence of the filter is limited more by *unmodelled* errors in the system dynamics than by the modelled simplifications inside the state transition matrix.

9.2. Effect of force model simplifications

Referring back to the subquestions that were formulated for this research, the effect of certain force model simplifications is of interest. We start with a discussion of the effects of neglecting orbital eccentricity. Afterwards, we detail the validity of the assumption that variations in gravity due to Earth oblateness are negligible for purposes of approximating the state transition matrix. Finally, we consider the relative effect of ignoring the remaining perturbations.

Impact of orbital eccentricity

The most important assumption within the Hill-Clohessy-Wiltshire solutions is the neglect of orbital eccentricity. Where the neglect of perturbations such as Earth oblateness or variations in radiation pressure is only necessary to simplify the Hill-Clohessy-Wiltshire solutions, the assumption of negligible eccentricity is fundamental to their derivation: without it, the equations of relative motion are time dependent and not straightforwardly analytically solvable. To verify that this assumption indeed holds for the highly-circular orbit of Sentinel-6A, we compare the results obtained using the Hill-Clohessy-Wiltshire with those obtained using the spherical Earth model.

As these models differ primarily in terms of the assumption on the shape of the orbit, the differences between either result can be used to bound the effect of neglecting ellipticity. It is to be noted, however, that the models also differ in the fact that the spherical Earth model applies a numerical integrator, and that the Hill-Clohessy-Wiltshire model approximates the aerodynamic drag and radiation pressure accelerations as arc-wise constant accelerations. Hence, differences in performance cannot be attributed *solely* to neglecting ellipticity, but they do provide a useful upper bound.

In general, the difference between both models can be said to be negligible. As evident from Figure 8.5, the inter-solution difference between both converges to fractions of a millimetre. Similarly, as reported in Section 8.3, the error covariance estimates are in close agreement, although qualitative differences in behaviour can be seen in the velocity error covariance volume in particular. Finally, Figure 8.7a shows that the differential correction error does exhibit a relatively large difference in behaviour: as discussed in section 9.1, that difference is seemingly less significant in light of the overall state transition matrix error. Indeed, it can be concluded that the effect of ignoring eccentricity is small with respect to the overall filter error of multiple centimetres. This is unsurprising, but an important step: it shows that the fundamental assumption permitting the computational simplification behind the Hill-Clohessy-Wiltshire model is valid.

Effect of Earth oblateness

The relative impact of the neglect of Earth oblateness in the variational model becomes visible through comparison of the results from the oblate and spherical Earth models. Here, the overall inter-solution differences are significantly larger than between the spherical Earth and Hill-Clohessy-Wiltshire models: where the difference between the spherical Earth model and the full variational model, previously shown in Figure 8.5, is on the order of a millimetre, the oblate Earth model approximates the full model to within micrometres. Similar orders-of-magnitude reductions in model differences were found in Sections 8.3 and 8.4 for the other two metrics. This is expected: considering the low-Earth orbit of Sentinel-6A, the effect of Earth oblateness is quite pronounced with respect to other accelerations.

This indicates a promising avenue for future research: the effect of Earth oblateness can be described perfectly well as an extra acceleration in the Hill-Clohessy-Wiltshire equations. Because of its slowly-varying nature, it could even be described quite well using the piecewise-constant acceleration model already applied to model the effects of radiation pressure and aerodynamic drag. All acceleration terms needed to compute the acceleration due to the second order zonal harmonic are already computed within the trajectory propagator itself, meaning that this extension would incur no significant additional computational costs. This would directly remove the largest remaining error source in the variational equations.

Magnitude of remaining perturbations

We can quantify the relative impact of the larger remaining accelerations by comparing the quality of the oblate Earth model with the reference force model. This illuminates the effects of the higher-order spherical harmonics up to degree and order 50, Sun and Moon gravity, k_2 solid Earth tides, and Schwarzschild acceleration. These accelerations are included in the full trajectory propagation itself but not in any of the simplified variational models.

As expected, exclusion of these accelerations in the variational model leads to relatively insignificant changes in the final trajectory estimate: the inter-trajectory comparison in Figure 8.5 shows differences of mere micrometres. This is supported by the observed error ellipsoid geometries: matching to within 0.001° in orientation and 0.01% in volume, these differences are negligible. Similarly, the differential correction error associated with these perturbations was found to be on the order of 10^{-6} , far below the level of other error sources such as non-linearity of the variational equations. In contrast with the stark increase in runtime of over 200% with respect to the simplified variational models, it is clear that these neglects are justifiable.

9.3. Limitations of this case study

Naturally, this case study is characterized by several limitations. Here, we will briefly go over the three limitations that we deem most significant, to establish the degree to which they restrict the applicability or certainty of the conclusions that can be drawn from this research. In no particular order, these limitations are: the limited accuracy of the presented orbit determination filter, the duration of the analysed time period, and the fact that only Sentinel-6A has been considered.

Orbit determination accuracy

In the first place, we consider the filter accuracy: the radial accuracy, at 10 cm RMS, does not yet fully match the 3 cm radial accuracies that have been achieved in the state-of-the-art with RTNAV [38] and DIODE [31]. It is reiterated that matching these performances was not the goal of this research; still, it begs the question as to what degree a more accurate filter might react differently to simplifications in estimation of the state transition matrix. This relative reduction in accuracy is in line with what is to be expected based on the known neglect of ocean tides in the trajectory model and of carrier-phase wind-up in the observation model. Overall, inclusion of these factors would be expected to improve the quality of *all* estimates, suggesting that the absolute differences would also decrease similarly. However, their inclusion could also lead to slightly different qualitative behaviour: hence, we shall specifically assess what the expected impact is of each of both factors.

Firstly, we consider the effect of ocean tide inclusion, which would lead to an improved trajectory model. Probably the largest consequence of this would be that it permits an increase in the observation batch interval from 10 s to a larger value. In the state-of-the-art, values from 30 s to 60 s are observed. This would mean that the transition matrix approximations would have to be valid for longer, which could lead to slightly larger differences between approximation methods. Still, batch sizes of even 60 s span only around 1% of an orbital period, meaning that this difference is unlikely to be significant with respect to other errors. In addition, this improved dynamical model would lead to an improved estimate for the reference state transition matrix. This would likely lead to a slightly larger difference in behaviour between the reference matrix and the three approximations presented in this work; however, this difference is unlikely to result in estimation differences larger than a millimetre, as the effect of ocean tides over 10 seconds of propagation on the trajectory itself is already much smaller than that. Indeed, it can therefore be expected that the inclusion of ocean tides in the trajectory model will not change the effectiveness of the presented state transition matrices.

Secondly, we assess the expected impact of carrier-phase wind-up. Here, we refer back to the comparison between broadcast GNSS and precise GPS ephemerides in Section 8.1. As indicated already by the resulting error frequency spectrum, this does lead to a meaningfully different dependence on the kinematic model. Hence, this provides some insight on the effect of an improved observation model. Yet, as mentioned in the results, this does not lead to a meaningful change in the inter-solution differences. Because the inclusion of carrier-phase wind-up also does not affect the computation of the sensitivity and orbital state transition matrices, there is no reason to believe that it would affect the results differently. Indeed, a preliminary conclusion that can be drawn from this is that the effect of simplifications in the state transition matrix will not change when the observation model improves in

quality.

Duration of the analysed time period

To keep the computation time needed for the presented analysis limited, the time period to be analysed was also necessarily restricted to 24 h. This could affect our analysis in two ways. In the first place, longer-term filter behaviour that only arises after one day could be hidden. Secondly, this time period explicitly does not cover events such as orbit-raising manoeuvres or observation black-outs, to keep the filter implementation as simple as possible.

Regarding the hidden long-term filter behaviour, it becomes clear from the results that this is unlikely. As shown in the Section 8.1, the orbital state and dynamic parameters converge mostly within an hour, and fully within several hours. Only for the aerodynamic drag coefficient it can be said that it is not fully converged; however, due to the unobservability of this parameter with respect to the empirical accelerations, further convergence is unlikely. Indeed, all other parameters are reset every few hours, meaning that no new longer-term behaviour is expected to arise there, either.

As for the effect of orbit-raising manoeuvres or observation black-outs, we can say that their effects are likely best described by the filter behaviour after a cold start. In other words, for differences in state transition matrix performance, we can look at the convergence behaviour of each transition matrix approximation. As observed in Section 8.2, the transition matrix differences are only slightly larger during filter convergence, and do not differ meaningfully in behaviour: the relative differences between approximations still hold. Hence, it is not expected that the filter properties during orbital manoeuvres or observation black-outs are significantly affected by state transition matrix approximations.

Applicability to other spacecraft

Finally, it is noted that this case study was performed specifically for Sentinel-6A. This is based fundamentally on the fact that it is a well-studied spacecraft with an orbit that near-perfectly fits the properties needed for the Hill-Clohessy-Wiltshire approximation: its orbit is highly circular and, while in low Earth orbit, still sufficiently high for the effects of higher-order Earth spherical harmonic terms to be reduced. These properties are quite significant for the applicability of this analysis, as they are used to derive the presented Hill-Clohessy-Wiltshire approximation.

As discussed previously in Section 9.2, the most fundamental assumption in this derivation is that of circularity of the orbit. This property is satisfied by most LEO spacecraft to some degree, due to the circularizing effect of aerodynamic drag. Still, the relatively high near-circularity achieved by Sentinel-6A, based in part on its application to altimetry, is relatively rare. Hence, care should be taken when extrapolating the results presented in this work to other spacecraft: it has only been shown that the Hill-Clohessy-Wiltshire approximation holds for *sufficiently* circular orbits, but the exact meaning of *sufficient* has not been investigated.

Another property of the presented case study that restricts its applicability is the fact that only GNSS observations were considered. Strictly, the choice of observations does not affect the computation of the state transition matrix, so the core results can be applied to other types of observables. However, it is noted that other observables can have different observation accuracies and observation intervals. DORIS, for example, is characterized by relatively long and frequent observation black-outs due to the absence of ground stations at certain points on Earth. Similarly, SLR observations are dependent on a clear atmosphere to make observations, so can be even more irregular in timing. In such scenarios, the time over which covariance propagation and observation correlation is to be performed becomes even longer than using regularly-spaced GNSS measurements. Indeed, it is certainly reasonable to believe that this does change the effectiveness of certain approximations to the state transition matrix. Consequently, care must also be taken when extrapolating the presented results to other measurement types.

10

Conclusion

In recent years, improvements in navigation systems like the Global Positioning System (GPS) and Galileo and in satellite orbit determination systems such as Doppler Orbitography and Radiopositioning Integrated by Satellite (DORIS), have resulted in higher-quality observations. This introduced the potential for real-time orbit determination with centimetre-level accuracies. Notably, radial Root Mean Square (RMS) accuracies have been demonstrated of 3.5 cm for Sentinel-6A using combined GPS and Galileo measurements [14, 38], and of 3 cm for Jason-2, Saral, CryoSat-2, and Sentinel-6A using DORIS observations [31, 38].

To make most optimal use of these improved observables, reduced-dynamic orbit determination filters must apply increasingly refined dynamic models. Enhancing computational efficiency is vital in supporting further improvement of real-time orbits, in particular with an eye on potential future application to improved onboard orbits. This thesis focuses on accurate and efficient approximations of the state transition matrix, a key component of the orbit determination process which describes the effect of state perturbations over time. Specifically, this work addresses approximations for both the orbital state transition and sensitivity matrices; these state transition submatrices are crucial for observation residual correction and state error covariance propagation in an Extended Kalman Filter (EKF).

In this work, four such approximations are compared in terms of accuracy and computational intensity, in the context of a case study of Global Navigation Satellite System (GNSS)-based orbit determination of Sentinel-6A on 1 January 2022. These approximations include one reference matrix, computed through numerical integration of the full variational equations; one novel analytic approximation, based on the Hill-Clohessy-Wiltshire equations; one matrix neglecting all but oblate Earth gravity, known to be applied in the current state-of-the-art; and one intermediate matrix, based on a spherical Earth assumption, used to separate the effect of distinct assumptions. All four state transition matrix approximations reach near-identical radial RMS accuracies of 14 cm. While this accuracy is not yet on par with the 3 to 4 cm achieved by state-of-the-art filters, likely in a large part because of known deficiencies in the trajectory and observation models, the resulting *differences* between orbit solutions are expected to be representative of actual performance of each matrix approximation.

The difference in accuracy among these matrices is contrasted through a comparison of three key performance characteristics. Collectively, these characteristics offer a holistic overview of the accuracy of a state transition matrix approximation within the framework of filtering orbit determination. Apart from the principal performance indicator, which is the difference in orbit determination accuracy, two more specific characterizations are introduced to elucidate the accuracy of a matrix in terms of its effect on the different filtering steps it is applied in. The first measure is based in a geometric interpretation of the error covariance matrix; this interpretation follows from previous work by Der and Danchick and by Der on approximations to the state transition matrix [15, 16]. The second measure entails a generalization of the differential correction error metric as initially proposed by Montenbruck and Gill [36]. In this work, it is extended to remove the need for normalization of the state transition matrix prior to application.

These metrics are used to differentiate the effect of each of the assumptions made in application of the Hill-Clohessy-Wiltshire state transition matrix approximation. For the orbit of Sentinel-6A, neglect of orbital ellipticity in computation of the state transition matrix affects the trajectory estimate quality by less than a millimetre on average. The effect of Earth oblateness is slightly more significant, but also

does not exceed the level of a millimetre. Neglection of all other orbital perturbations in approximation of the matrizant is even less significant, at mere micrometres in magnitude. Finally, our results show that none of these simplifications appreciably affect filter convergence. It is thus concluded that these assumptions are more than valid in the context of centimetre-level orbit determination for altimetric low Earth orbits like that of Sentinel-6A.

It follows that an analytic matrizant approximation based on the Hill-Clohessy-Wiltshire equations is exceedingly suitable for the niche of real-time onboard orbit determination of altimetric low Earth orbits. Compared to existing approximations known to be applied in the state-of-the-art, a reduction is achieved of up to 10% in the *total* runtime needed for a single filter propagation and update step. The orbit solution computed using this Hill-Clohessy-Wiltshire matrizant never deviates more than 5 mm from the reference orbit computed with a full force model matrizant; the standard deviation is under a millimetre in all position components. This makes the simplification more than suitable for real-time onboard orbit determination scenarios where computational power is at a premium, with currently representative accuracies of 3 to 4 cm [38].

This does not mean that the presented Hill-Clohessy-Wiltshire approximation is perfectly suitable for all applications: applications with orbit accuracies nearer to 5 mm, as is the case when real-time corrections are available, require more sophisticated approximations of the state transition matrix. Applications where computational power is more readily available would benefit from opting for more accurate approximations, including effects such as Earth oblateness and ellipticity of orbits. Our results show that, at 10% computational overhead, a numerically-integrated oblate Earth model can improve the accuracy across all metrics by up to three orders of magnitude, which would permit application down to millimetre-level positioning accuracies. This is in line with its usage in earlier research [39]. Additionally, it relaxes the assumptions on orbital eccentricity, permitting even highly-elliptical orbits.

This suggests a natural recommendation for future research: while state-of-the-art real-time onboard orbital accuracies of 3-4 cm in near-circular orbits are achievable with application of the presented Hill-Clohessy-Wiltshire matrizant, further enhancement of this accuracy would require more sophisticated approximations. Our results show that the presented Hill-Clohessy-Wiltshire approximation is limited in the first place by its neglection of Earth oblateness. Extension of this Hill-Clohessy-Wiltshire implementation with an additional acceleration term to account for the J_2 gravity perturbation would significantly improve its accuracy. As this term is already computed as part of the trajectory propagator, this would incur no appreciable additional computational overhead. Aforementioned results with a numerically-integrated oblate Earth model suggest that such a state transition matrix approximation would be compatible with subcentimetre orbit determination of altimetric low Earth orbits. This would allow application of the presented Hill-Clohessy-Wiltshire approximation at the even lower real-time orbit accuracies that are to be expected in future years.


References

- [1] CH Acton. "Ancillary Data Services of NASA's Navigation and Ancillary Information Facility". In: *Planetary and Space Science* 44.1 (Jan. 1996), pp. 65–70. ISSN: 00320633. DOI: 10.1016/0032-0633(95)00107-7. URL: <https://linkinghub.elsevier.com/retrieve/pii/0032063395001077> (visited on 07/31/2023).
- [2] L Banfield and M Pattinson. *Global Positioning System (GPS) Performance*. Quarterly report 1 GMVNSL-CAA-GPS-SPS-Q1-22. GMV NSL, Apr. 28, 2022. URL: <https://www.caa.co.uk/Documents/Download/1851/6c8c7359-e7fa-4904-8b81-60a20aa1fdbd/76> (visited on 09/01/2023).
- [3] J Bangert et al. *SOFA Tools for Earth Attitude*. International Astronomical Union, Mar. 21, 2014. URL: https://www.iausofa.org/sofa_pn_f.pdf (visited on 09/01/2023).
- [4] W Bertiger et al. "GipsyX/RTGx, a New Tool Set for Space Geodetic Operations and Research". In: *Advances in Space Research* 66.3 (Aug. 1, 2020), pp. 469–489. DOI: 10.1016/j.asr.2020.04.015.
- [5] GJ Bierman. *Factorization Methods for Discrete Sequential Estimation*. Courier Corporation, May 26, 2006. 260 pp. ISBN: 978-0-486-44981-4.
- [6] WP Birmingham, BL Miller, and WL Stein. "Experimental Results of Using the GPS for Landsat 4 Onboard Navigation". In: *NAVIGATION* 30.3 (1983), pp. 244–251. DOI: 10.1002/j.2161-4296.1983.tb00843.x.
- [7] BK Bradley, A Sibois, and P Axelrad. "Influence of ITRS/GCRS Implementation for Astrodynamics: Coordinate Transformations". In: *Advances in Space Research* 57.3 (Feb. 2016), pp. 850–866. ISSN: 02731177. DOI: 10.1016/j.asr.2015.11.006. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0273117715007929> (visited on 05/31/2023).
- [8] RA Broucke. "On the Matrizant of the Two-Body Problem". In: *Astronomy and Astrophysics* 6 (June 1970), p. 173. URL: <https://ui.adsabs.harvard.edu/abs/1970A&A.....6..173B> (visited on 09/01/2023).
- [9] RA Broucke and PJ Cefola. "On the Equinoctial Orbit Elements". In: *Celestial Mechanics* 5.3 (May 1972), pp. 303–310. DOI: 10.1007/BF01228432.
- [10] *Copernicus POD Regular Service Review Jan - Mar 2022*. Regular service review deliverable GMV-CPOD-RSR-0024. Copernicus Sentinel-1, -2, -3 and -6 Precise Orbit Determination Service (CPOD), GMV, Apr. 27, 2022. URL: https://sentinel.esa.int/documents/247904/4599719/GMV-CPOD-RSR-0024_v1.0_Copernicus_POD_Regular_Service_Review_Jan_Mar_2022.pdf (visited on 09/04/2023).
- [11] R Cullen et al. "Sentinel-6 Michael Freilich Poseidon-4 Altimeter In-Orbit Performance". In: *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*. IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium. July 2022, pp. 6791–6794. DOI: 10.1109/IGARSS46834.2022.9883552.
- [12] HD Curtis. *Orbital Mechanics for Engineering Students*. Google-Books-ID: 2U9Z8k0TITYC. Butterworth-Heinemann, Oct. 5, 2013. 931 pp. ISBN: 978-0-08-097748-5.
- [13] Z Dang. "New State Transition Matrix for Relative Motion on an Arbitrary Keplerian Orbit". In: *Journal of Guidance, Control, and Dynamics* 40.11 (Nov. 2017), pp. 2917–2927. ISSN: 0731-5090, 1533-3884. DOI: 10.2514/1.G002723. URL: <https://arc.aiaa.org/doi/10.2514/1.G002723> (visited on 06/21/2023).
- [14] Francesco Darugna et al. "Sentinel-6A GPS and Galileo Dual-Frequency Real-Time Reduced-Dynamics P2OD". In: *2022 10th Workshop on Satellite Navigation Technology (NAVITEC)*. 2022 10th Workshop on Satellite Navigation Technology (NAVITEC). Apr. 2022, pp. 1–12. DOI: 10.1109/NAVITEC53682.2022.9847561.

- [15] GJ Der. "An Elegant State Transition Matrix". In: *The Journal of the Astronautical Sciences* 45.4 (Dec. 1, 1997), pp. 371–390. DOI: 10.1007/BF03546398.
- [16] GJ Der and R Danchick. "Analytic and Numerical Error Covariance Matrix Propagation". In: *Astrodynamics Conference*. Astrodynamics Conference. San Diego, CA, U.S.A.: American Institute of Aeronautics and Astronautics, July 29, 1996. DOI: 10.2514/6.1996-3661.
- [17] S Dinardo et al. "Sentinel-6 MF Poseidon-4 Radar Altimeter In-Flight Calibration and Performances Monitoring". In: *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022). Conference Name: IEEE Transactions on Geoscience and Remote Sensing, pp. 1–16. DOI: 10.1109/TGRS.2022.3216595.
- [18] D Dirx. *Tudat Mathematical Model Definition*. Unpublished notes. Delft University of Technology, Jan. 24, 2022. URL: https://raw.githubusercontent.com/tudat-team/tudat-space/master/Tudat_mathematical_model_definition.pdf (visited on 08/01/2023).
- [19] CJ Donlon et al. "The Copernicus Sentinel-6 Mission: Enhanced Continuity of Satellite Sea Level Measurements from Space". In: *Remote Sensing of Environment* 258 (June 2021). DOI: 10.1016/j.rse.2021.112395.
- [20] *Earth Orientation Matrix/Quaternion and Rotation Vector*. Online software. l'Observatoire de Paris, June 2023. URL: <https://hpiers.obspm.fr/eop-pc/index.php?index=rotation&lang=en> (visited on 07/29/2023).
- [21] European GNSS Service Centre. *European GNSS (Galileo) Services Open Service Quarterly Performance Report January - March 2022*. 1.0. 2022. URL: https://www.gsc-europa.eu/sites/default/files/Galileo-OS-Quarterly-Performance_Report-Q1-2022.pdf (visited on 09/01/2023).
- [22] JH Faghmous et al. "A Daily Global Mesoscale Ocean Eddy Dataset from Satellite Altimetry". In: *Scientific Data* 2.1 (June 9, 2015). Number: 1 Publisher: Nature Publishing Group, p. 150028. DOI: 10.1038/sdata.2015.28.
- [23] T. Fecher et al. "GOCO05c: A New Combined Gravity Field Model Based on Full Normal Equations and Regionally Varying Weighting". In: *Surveys in Geophysics* 38.3 (May 1, 2017), pp. 571–590. ISSN: 1573-0956. DOI: 10.1007/s10712-016-9406-y. URL: <https://doi.org/10.1007/s10712-016-9406-y> (visited on 08/01/2023).
- [24] A Flores and D Nicholas. *NAVSTAR GPS Space Segment/Navigation User Interfaces*. Interface specification IS-GPS-200. SAIC, May 21, 2021.
- [25] L Giulicchi and R Cullen. *Sentinel-6 Michael Freilich POD Context*. Noordwijk: European Space Research and Technology Centre, Jan. 22, 2022. URL: https://ids-doris.org/documents/BC/satellites/Sentinel6A_PODcontext.pdf (visited on 09/01/2023).
- [26] WH Guier and GC Weiffenbach. "A Satellite Doppler Navigation System". In: *Proceedings of the IRE* 48.4 (Apr. 1960). Conference Name: Proceedings of the IRE, pp. 507–516. DOI: 10.1109/JRPROC.1960.287399.
- [27] WH Guier and GC Weiffenbach. "The Doppler Determination of Orbits". In: *Proceedings of the National Aeronautics and Space Administration Conference on Orbit and Space Trajectory Determination*. National Aeronautics and Space Administration Conference on Orbit and Space Trajectory Determination. U.S. Naval Research Lab, Washington, D.C., Mar. 12, 1959.
- [28] U Hugentobler. *Orbit Perturbations due to Relativistic Corrections*. Supporting material to the IERS conventions. Aug. 18, 2008. URL: https://impacts.obspm.fr/content/supporting_material/chapter10/relativity_hugentobler.pdf.
- [29] C Jayles, JP Chauveau, and A Auriol. "DORIS/DIODE: Real-Time Orbit Determination Performance on Board SARAL/AltiKa". In: *Marine Geodesy* 38 (sup1 Sept. 10, 2015), pp. 233–248. DOI: 10.1080/01490419.2015.1015695.
- [30] C Jayles, JP Chauveau, and F Rozo. "DORIS/Jason-2: Better than 10cm On-Board Orbits Available for Near-Real-Time Altimetry". In: *Advances in Space Research*. DORIS: Precise Orbit Determination and Applications to Earth Sciences 46.12 (Dec. 15, 2010), pp. 1497–1512. DOI: 10.1016/j.asr.2010.04.030.

- [31] C Jayles et al. "DORIS System and Integrity Survey". In: *Advances in Space Research*. Scientific Applications of DORIS in Space Geodesy 58.12 (Dec. 15, 2016), pp. 2691–2706. DOI: 10.1016/j.asr.2016.05.032.
- [32] GH Kaplan. *The IAU Resolutions on Astronomical Reference Systems, Time Scales, and Earth Rotation Models*. Circular No. 179. Washington, D.C.: U.S. Naval Observatory, Oct. 20, 2005. URL: https://aa.usno.navy.mil/downloads/Circular_179.pdf (visited on 08/01/2023).
- [33] K Kumar et al. "TUDAT: a Modular and Robust Astrodynamics Toolbox". In: *Conference proceeding of the 5th ICATT Conference* (2012). Ed. by A Benoit. Place: Noordwijk Publisher: ESA, pp. 1–8.
- [34] M Li et al. "Performance of Multi-GNSS in the Asia-Pacific Region: Signal Quality, Broadcast Ephemeris and Precise Point Positioning (PPP)". In: *Remote Sensing* 14.13 (Jan. 2022). Number: 13 Publisher: Multidisciplinary Digital Publishing Institute, p. 3028. DOI: 10.3390/rs14133028.
- [35] JS Llorente et al. "PROBA-3: Precise Formation Flying Demonstration Mission". In: *Acta Astronautica*. 6th International Workshop on Satellite Constellation and Formation Flying 82.1 (Jan. 1, 2013), pp. 38–46. DOI: 10.1016/j.actaastro.2012.05.029.
- [36] O Montenbruck and E Gill. "Satellite Orbits, Models, Methods and Applications". In: Springer-Verlag, 2000, pp. 1–369. ISBN: 978-3-540-67280-7.
- [37] O Montenbruck and E Gill. "State Interpolation for On-Board Navigation Systems". In: *Aerospace Science and Technology* 5.3 (Mar. 2001), pp. 209–220. DOI: 10.1016/S1270-9638(01)01096-3.
- [38] O Montenbruck, F Kunzi, and A Hauschild. "Performance Assessment of GNSS-Based Real-Time Navigation for the Sentinel-6 Spacecraft". In: *GPS Solutions* 26.1 (Jan. 2022), p. 12. DOI: 10.1007/s10291-021-01198-9.
- [39] O Montenbruck and P Ramos-Bosch. "Precision Real-Time Navigation of LEO Satellites using Global Positioning System Measurements". In: *GPS Solutions* 12.3 (July 2008), pp. 187–198. DOI: 10.1007/s10291-007-0080-x.
- [40] O Montenbruck and P Steigenberger. "BRD400DLR: DLR's Merged Multi-GNSS Broadcast Ephemeris Product in RINEX 4.00 Format DLR/GSOC". In: (2022).
- [41] O Montenbruck, P Steigenberger, and A Hauschild. "Broadcast versus Precise Ephemerides: a Multi-GNSS Perspective". In: *GPS Solutions* 19.2 (Apr. 1, 2015), pp. 321–333. DOI: 10.1007/s10291-014-0390-8.
- [42] O Montenbruck, P Steigenberger, and A Hauschild. "Multi-GNSS Signal-in-Space Range Error Assessment - Methodology and Results". In: *Advances in Space Research* 61.12 (June 15, 2018), pp. 3020–3038. DOI: 10.1016/j.asr.2018.03.041.
- [43] O Montenbruck et al. "Sentinel-6A Precise Orbit Determination using a Combined GPS/Galileo Receiver". In: *Journal of Geodesy* 95.9 (Sept. 5, 2021), p. 109. DOI: 10.1007/s00190-021-01563-z.
- [44] TD Moyer. *Mathematical Formulation of the Double Precision Orbit Determination Program (DPODP)*. JPL-TR-32-1527. NTRS Author Affiliations: Jet Propulsion Lab., California Inst. of Tech. NTRS Document ID: 19710017134 NTRS Research Center: Legacy CDMS (CDMS). Jet Propulsion Laboratory, May 15, 1971. URL: <https://ntrs.nasa.gov/citations/19710017134> (visited on 08/03/2023).
- [45] *NGA GNSS Division Precise Ephemeris Parameters*. Product documentation. National Geospatial-Intelligence Agency, Feb. 20, 2020. URL: <https://earth-info.nga.mil/php/download.php?file=gnss-precise> (visited on 09/04/2023).
- [46] D Odijk et al. "On the Estimability of Parameters in Undifferenced, Uncombined GNSS Network and PPP-RTK User Models by Means of S-System Theory". In: *Journal of Geodesy* 90.1 (Jan. 1, 2016), pp. 15–44. DOI: 10.1007/s00190-015-0854-9.
- [47] EA Okal, A Piatanesi, and P Heinrich. "Tsunami Detection by Satellite Altimetry". In: *Journal of Geophysical Research: Solid Earth* 104 (B1 1999), pp. 599–615. DOI: 10.1029/1998JB000018.
- [48] JA O'Keefe, A Ekeis, and RK Squires. "Vanguard Measurements Give Pear-Shaped Component of Earth's Figure". In: *Science* 129.3348 (Feb. 27, 1959), pp. 565–566. DOI: 10.1126/science.129.3348.565.

- [49] G Petit and B Luzum. *IERS Conventions (2010)*. Technical note No. 36. Frankfurt am Main: IERS Conventions Centre, 2010. URL: <http://www.iers.org/TN36/> (visited on 05/31/2023).
- [50] JM Picone et al. "NRLMSISE-00 Empirical Model of the Atmosphere: Statistical Comparisons and Scientific Issues". In: *Journal of Geophysical Research: Space Physics* 107 (A12 Dec. 2002), SIA 15–1–SIA 15–16. ISSN: 01480227. DOI: 10.1029/2002JA009430. URL: <http://doi.wiley.com/10.1029/2002JA009430> (visited on 08/01/2023).
- [51] D Rice. "An Investigation into the Effects of Using Simplified Force Models in the Computation of State Transition Matrices". In: *5th Aerospace Sciences Meeting*. 5th Aerospace Sciences Meeting. New York, NY, U.S.A.: American Institute of Aeronautics and Astronautics, Jan. 23, 1967. DOI: 10.2514/6.1967-123.
- [52] J Saunier. "The DORIS Network: Advances Achieved in the Last Fifteen Years". In: *Advances in Space Research* (July 19, 2022). DOI: 10.1016/j.asr.2022.07.016.
- [53] B Schutz, B Tapley, and GH Born. *Statistical Orbit Determination*. Elsevier, June 26, 2004. 564 pp. ISBN: 978-0-08-054173-0.
- [54] W Simons, PNAM Visser, and M Naeije. "Precise Orbit Determination of Sentinel Satellites with GIPSY-X and GPS/Galileo Data". In: 44 (July 1, 2022). Conference Name: 44th COSPAR Scientific Assembly. Held 16-24 July ADS Bibcode: 2022cosp...44.3408S, p. 3408. URL: <https://ui.adsabs.harvard.edu/abs/2022cosp...44.3408S> (visited on 09/13/2023).
- [55] D Song et al. "Guidance and Control of Autonomous Surface Underwater Vehicles for Target Tracking in Ocean Environment by Deep Reinforcement Learning". In: *Ocean Engineering* 250 (Apr. 2022), p. 110947. DOI: 10.1016/j.oceaneng.2022.110947.
- [56] PJG Teunissen and O Montenbruck. *Springer Handbook of Global Navigation Satellite Systems*. 2017. ISBN: 978-3-319-42926-7. DOI: 10.1007/978-3-319-42928-1.
- [57] JP Vinti. "Inclusion of the Third Zonal Harmonic in an Accurate Reference Orbit of an Artificial Satellite". In: *Journal of Research of the National Bureau of Standards Section B Mathematics and Mathematical Physics* 70B.1 (Nov. 17, 1965), p. 17. DOI: 10.6028/jres.070B.003.
- [58] EA Wan and R Van Der Merwe. "The Unscented Kalman Filter for Nonlinear Estimation". In: *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*. Symposium on Adaptive Systems for Signal Processing Communications and Control. Lake Louise, Alta., Canada: IEEE, 2000, pp. 153–158. ISBN: 978-0-7803-5800-3. DOI: 10.1109/ASSPCC.2000.882463.



Derivation of the Differential Correction Error

When introducing approximations to the orbital state transition and sensitivity matrices $\Phi_{\mathbf{y}(t,t_0)}$ and $\Phi_{\mathbf{p}}$, it is useful to have a meaningful measure of their accuracy. While it is possible to compare the relative or absolute errors in the elements of some approximations $\bar{\Phi}_{\mathbf{y}}(t, t_0)$ and $\bar{\Phi}_{\mathbf{p}}(t, t_0)$, that comparison carries no semantic meaning. More specifically, it does not account for the interactions between elements, and for the relative magnitude of each element's effect on the final state approximation.

Instead, it is possible to derive a more meaningful, objective measure of state transition matrix accuracy based on its application in a differential correction. Conceptually, this can be connected to its use in a Kalman filter: the state update applied to reconcile the a priori state estimate with incoming observations is essentially a differential correction. It approximates the state perturbation needed to minimize the sum-of-squares of the observation residuals and the deviation from the a priori state. Similarly, the propagation of the covariance matrix is equivalent to propagation of a set of (infinitesimal) state perturbations. As such, one measure to quantify the error introduced into an orbit determination filter by a state transition matrix approximation is the residual state perturbation $\Delta\bar{\mathbf{x}}$ that remains after correcting for an observed perturbation $\Delta\mathbf{x}$. We shall derive this error measure in two steps: first, in Appendix A.1, we shall derive it in the context of a general differential correction step. Afterwards, in Appendix A.2, we shall formulate it in the context of our GNSS-based orbit determination filter.

A.1. Differential correction

With “differential correction”, we refer to the process of refining a state estimate $\bar{\mathbf{x}}(t)$, known to deviate from some other, “perfect” state estimate $\mathbf{x}(t)$ by a difference $\Delta\mathbf{x} = \bar{\mathbf{x}}(t) - \mathbf{x}(t)$, by adjustment of an earlier state estimate $\bar{\mathbf{x}}(t_0)$ ¹. This is the technique that is also applied in a single iteration of non-linear least squares:

$$\Delta\mathbf{x}(t_0) = \Phi^{-1}(t, t_0)\Delta\mathbf{x}(t).$$

This is the ideal case, when we have perfect knowledge of $\Phi(t, t_0)$. When we consider an approximation $\bar{\Phi}(t, t_0)$, we instead obtain an approximate correction $\Delta\bar{\mathbf{x}}(t_0)$:

$$\Delta\bar{\mathbf{x}}(t_0) = \bar{\Phi}^{-1}(t, t_0)\Delta\mathbf{x}(t).$$

Note that we have still assumed perfect knowledge of $\Delta\mathbf{x}$: after all, in the differential correction error measure, we are not interested in the error introduced by errors in the observation model.

Due to the error introduced by this imperfect approximation $\Delta\bar{\mathbf{x}}(t_0)$, we also find an imperfect final state update $\Delta\bar{\mathbf{x}}(t) = \Phi(t, t_0)\Delta\bar{\mathbf{x}}(t_0)$. The residual error after this differential correction can be expressed as the difference between $\Delta\mathbf{x}(t)$ and its approximation $\Delta\bar{\mathbf{x}}(t)$:

$$\Delta\mathbf{x}(t) - \Delta\bar{\mathbf{x}}(t) = \Phi(t, t_0)\Delta\mathbf{x}(t_0) - \Phi(t, t_0)\Delta\bar{\mathbf{x}}(t_0) = \left(I - \Phi(t, t_0)\bar{\Phi}^{-1}(t, t_0)\right)\Delta\mathbf{x}(t).$$

¹N.B.: The term “differential correction” is also used to refer to the technique of observation differencing to eliminate shared error sources. This is a distinct concept.

While this expression is insightful, its value depends on the specific perturbation $\Delta x(t)$ used to assess it. This makes it an imperfect measure to assess the quality of an approximation $\bar{\Phi}(t, t_0)$, which should conceptually be independent of $\Delta x(t)$. Instead, we can only assess the matrix

$$M(t, t_0) = I - \Phi(t, t_0)\bar{\Phi}^{-1}(t, t_0),$$

which represents the growth of the residual state error after a single differential correction. This same expression is also given, without derivation, by Montenbruck and Gill in [36].

Still, this matrix operator alone is not sufficient to summarize the error growth over time due to a simplified state transition matrix in one performance measure. To achieve this, we introduce the largest singular value, or operator norm, $\sigma_1(M) \in \mathbb{R}$, which has the property that

$$\|Mx\| \leq \sigma_1(M)\|x\|, \quad \text{for all } x \in \mathbb{R}^n,$$

where n is the number of state parameters. The largest singular value can be straightforwardly obtained as the largest-magnitude diagonal element of Σ in the singular value decomposition $M = U\Sigma V^*$ of M . This operator norm represents the maximum growth in the norm $\|x\|$ that can occur by application of M to any arbitrary choice of $x \in \mathbb{R}^n$. Indeed, in the context of a differential correction, it follows that $\sigma_1(M(t, t_0))$ represents the relative maximum residual error that remains after a correction with an approximation $\bar{\Phi}(t, t_0)$ of the state transition matrix for *any* choice of $\Delta x(t)$.

A.2. Application in GNSS-based orbit determination

In the above derivation, we have explicitly left out the chosen formulation of state parameters x and the precise form of $\Phi(t, t_0)$. In principle, that is of no consequence: the derivation of the differential correction does not rely on those facts. To compute the actual value of the differential correction error, we do need an explicit formulation. For the GNSS-based state formulation used in the orbit determination filter presented in this work, the full state and state transition matrix are such that

$$\begin{bmatrix} \Delta y(t) \\ \Delta p(t) \\ \Delta t(t) \\ \Delta s(t) \\ \Delta b(t) \end{bmatrix} = \begin{bmatrix} \Phi_y(t, t_0) & \Phi_p(t, t_0) & 0 & 0 & 0 \\ 0 & M_p & 0 & 0 & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & 0 & 0 & M_s & 0 \\ 0 & 0 & 0 & 0 & I \end{bmatrix} \begin{bmatrix} \Delta y(t_0) \\ \Delta p(t_0) \\ \Delta t(t_0) \\ \Delta s(t_0) \\ \Delta b(t_0) \end{bmatrix}$$

for a time update from t_0 to t .

For the purposes of this analysis, the choice of exponential damping factors M_p and M_s is of little consequence on the qualitative behaviour of the error growth, so we can safely assume them to be identity. Additionally, we see that only Δp and Δy affect subsequent estimates of the orbital state y , which is the error we are most interested in for an orbit determination filter. This means that, for purposes of transition matrix error analysis, we can consider the reduced system

$$\begin{bmatrix} \Delta y(t) \\ \Delta p(t) \end{bmatrix} = \begin{bmatrix} \Phi_y(t, t_0) & \Phi_p(t, t_0) \\ 0 & I \end{bmatrix} \begin{bmatrix} \Delta y(t_0) \\ \Delta p(t_0) \end{bmatrix}.$$

The associated differential correction error growth matrix, as used to compare the reference model of the variational equations with the analytic Hill-Clohessy-Wiltshire solutions becomes

$$M(t, t_0) = I - \begin{bmatrix} \Phi_y(t, t_0) & \Phi_p(t, t_0) \\ 0 & I \end{bmatrix} \begin{bmatrix} \bar{\Phi}_y(t, t_0) & \bar{\Phi}_p(t, t_0) \\ 0 & I \end{bmatrix}^{-1}.$$

Here, $\Phi_y(t, t_0)$ and $\Phi_p(t, t_0)$ are the reference orbital state transition matrix and sensitivity matrix, respectively. Their approximations are denoted as $\bar{\Phi}_y(t, t_0)$ and $\bar{\Phi}_p(t, t_0)$.

B

State Transition Matrix Differences for the Precise GPS Filter

In Chapter 8, the differences in filter estimates are given for the broadcast GNSS ephemerides filter. These inter-estimate differences can, naturally, also be derived for the precise GPS filter; as expected, the results are not meaningfully different. Hence, these differences are not treated in detail in the results presented there. Still, for completeness, the inter-solution differences as well as the error ellipsoid geometries are included in this appendix in Figures B.1 and B.2, for the interested reader.






Figure B.1: Differences in final position estimate as a result of different approximations to the variational equations



(a) For the position error



(b) For the velocity error

Figure B.2: Difference in volume and orientation of the error covariance ellipsoids with respect to the full variational equations