Maneuvering Target Tracking Methods for Space Surveillance

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy

PhD Program in Aerospace Engineering
Universidad Carlos III de Madrid

by

Guillermo Escribano Blázquez

To my family and girlfriend, for their utmost support throughout these difficult times.

ACKNOWLEDGEMENTS

A special thank you to my supervisors, Manuel Sanjurjo and Jan Siminski for their huge effort in understanding me and making me understand. Not to forget Alejandro Pastor and Diego Escobar, whose operations oriented approach and mathematical and academical rigor were crucial throughout the development and spread of the research. I would like to thank Klaus Merz, Helene Ma and Beatriz Jilete for their exceptional support and guidance especially during the early stages of the thesis, that comes from countless years of experience and contact with cutting-edge space operations.

I will always be grateful to Dr. Brandon Jones first for giving me the opportunity to collaborate with him but foremost for guiding me through the intricacies of advanced statistical methods, as well as spreading his passion for teaching and academics. I would also like to thank Enrico Zuchelli for being such a good critical thinker and helping me unravel conceptual errors in the methodology, and also to Nick Ravago and Ben Reifler for their invaluable technical assistance in the particularities of multi-Bernoulli filters.

I will never forget the long discussions with Soumaya Azzi and Emiliano Cordelli, which significantly helped me better understand and formulate what later became a key aspect of the methodology.

There are no words that can express my gratitude to my family, especially my parents Esperanza and Emilio, for shaping my character and ensuring one day I could turn my passion into my career. It goes without saying this achievement belongs to a great extent to my girlfriend Bea, who has done what was in her power to help me succeed in this venture.

ABSTRACT

Earth orbits are a valuable natural resource that shall be preserved. A myriad of services currently relies on orbital stations, most of which support modern human society. To name a few, global navigation satellite services and weather forecasting, crucial for the vast majority of the population, require sensors and stations in Earth orbit. Recently, the space industry has seen a considerable decrease in launch and development costs, so access to space is more affordable than ever before. The latter has opened the gate for new actors in the space domain, such as startups and universities, which now operate small-sized spacecraft (up to 500 kg) at low orbital altitudes. This breed of new satellites adds up to an already numerous background population, composed not only of operative satellites but also derelict rockets and spacecraft as well as fragments originated from break-up events or explosions and in-orbit collisions. It is this last subset of the Earth orbital population that alarms the space community: for a sufficiently high congestion level, a single collision can trigger a cascade of additional events to the point that every object placed in orbit is guaranteed to collide against a neighboring fragment in a relatively short period of time. To prevent this catastrophic situation, the Earth orbital space is continuously monitored by surveillance networks composed of ground and space-based sensors of different types. Data acquired by these sensors is processed and fused with information coming from satellite operators to elaborate and maintain a comprehensive list, or catalog, of objects in Earth orbit. Generating a labelled map of objects can help to prevent future in-orbit collisions: by running one-to-one conjunction analysis over the entire population, it is possible to identify potentially hazardous conjunctions and issue warnings to spacecraft operators so they can take remedial action. Accurately tracking space objects with surveillance sensors is a complex task, especially if it is to be done automatically. In particular, the integrity of space object catalogs is compromised when spacecraft perform maneuvers and their operators do not report them in a timely manner, mainly because

automated space surveillance and tracking systems rely on natural satellite motion or very simple maneuver models at best. This thesis is aimed at advancing the tools and techniques required for automated space surveillance and tracking in the presence of unknown maneuvers, emphasizing on the definition of suitable maneuver models for space objects. One of the main problems associated with orbital maneuvers is data association, this is, being able to assess whether certain observation corresponds to a given target following a maneuver. In general, it is not easy to provide evidence supporting measurement to object correlation after or during a maneuvering period, especially if no knowledge regarding the maneuvering characteristics of the target can be assumed a priori. In the context of this thesis, orbital maneuvers are classified according to the different types of space propulsion technologies, namely, chemical and electrical. The former present high-thrust magnitudes and a poor propellant utilization, whereas the latter are significantly more fuel efficient at the cost of lower thrusting forces. Nonetheless, one can derive educated guesses for the expected control magnitudes of orbital maneuvers by analyzing missions equipped with either propulsive type. These control magnitudes can be used to bound the space that is accessible (or reachable) for a target given an initial state and some temporal bounds, for instance representing the last known state and the time elapsed since then. Still, application of these concepts requires the development of computationally efficient methods to compute orbital distances in terms of control, also known as control distance metrics, based either on chemical or electrical propulsion. In other words, a key enabler of the methodology proposed in this thesis is the development of inexpensive methods to evaluate control distance metrics, which shall capture the dynamical features of Earth's orbital environment and the different space propulsive technologies. To this end, two different surrogate models for orbital maneuvers have been developed, under low and high-thrust assumptions, and exploited to compute not only the control distance between two orbital states but also to approximate reachability bounds conditioned on an initial state and a time of flight. Statistically consistent maneuver models are then constructed from the developed surrogates, which can provide reasonable maneuver evidence for data association in space surveillance

applications. These models are embedded in an advanced multiple maneuvering target tracking filter, capable of managing ambiguity arising from target maneuvers, death, birth, missed sensor detections and false observations, thus being suitable for robust automated operations. Results obtained from synthetic datasets indicate the developed methods help the filter to maintain custody of maneuvering targets and resolve ambiguity in a wide range of scenarios. However, the strong dynamical impact of maneuvers leads to significant uncertainties in the state of space objects, to the point that the quality of conjunction analyses is severely degraded. Therefore, even if the developed methods are capable of maintaining custody of maneuverable space objects, it is advisable to further characterize maneuvers according to, for instance, historical data. Including heuristics in the definition of maneuver modes can help to reduce the effective uncertainty at early post-maneuver epochs, thereby increasing the overall quality of the space object catalog and allowing for accurate conjunction analyses in the presence of maneuvering targets.

RESUMEN

El espacio orbital terrestre es un recurso natural de gran valor que ha de ser preservado. Tanto es así que la sociedad moderna se desarrolla alrededor de multitud de servicios, la mayoría de los cuales requieren satélites en órbita terrestre. Entre ellos, cabe destacar los servicios globales de navegación por satélite y la predicción meteorológica, cruciales para la amplia mayoría de la población. Recientemente, los costes de lanzamiento y desarrollo relacionados con la industria espacial han disminuido de forma considerable, haciendo del espacio orbital terrestre un recurso natural relativamente accesible. Aprovechando la situación, numerosas startups y universidades ahora operan satélites pequeños, de hasta media tonelada, en órbita baja. Esta clase de nuevos satélites ha de convivir con una población de objetos ya de por sí numerosa, compuesta no sólo por satélites operativos sino también por cohetes y satélites abandonados, así como fragmentos originados por desintegraciones, explosiones y colisiones. De hecho, son estos últimos los que ponen a la comunidad espacial en alerta ya que, llegado cierto nivel de congestión en el entorno orbital, un único choque podría desencadenar una cascada de colisiones. Para evitar esta situación catastrófica, existen diversas redes de vigilancia espacial que monitorizan el espacio mediante sensores en tierra y en órbita. Los datos obtenidos a través de dichos sensores se procesan y combinan con información procedente de los operadores de satélites para elaborar y actualizar un catálogo detallado de objetos en órbita, garantizando así cierto nivel de conocimiento acerca de la población orbital terrestre. La generación de dicho catálogo puede ayudar a prevenir futuras colisiones orbitales: mediante un análisis individualizado de las posibles conjunciones entre objetos, se pueden identificar pares de objetos potencialmente peligrosos y así informar a los operadores de dichos objetos para que tomen las medidas pertinentes, es decir, maniobren para evitar la colisión. Realizar un seguimiento preciso de la población orbital con sensores de vigilancia es una tarea compleja, y más aún si ha de realizarse de forma autónoma. En particular, la integridad de los catálogos

de objetos se ve comprometida cuando éstos realizan maniobras sin que sus operadores las reporten, ya que los sistemas de catalogación automatizados se basan generalmente en el movimiento natural de los satélites o, a lo sumo, modelos de maniobra relativamente sencillos. Esta tesis tiene como principal objetivo avanzar en el desarrollo de herramientas y técnicas para la vigilancia y seguimiento del espacio considerando maniobras desconocidas a priori, centrándose en la definición de modelos de maniobra para objetos espaciales. Uno de los principales problemas asociados a las maniobras orbitales es la correlación de medidas, es decir, ser capaces de discernir si una observación se corresponde con un objeto que ha maniobrado recientemente. En términos generales, es complicado cuantificar la probabilidad de correlación entre una observación y un satélite que que acaba de maniobrar, más aún si se carece de información relativa a las características de maniobrabilidad del objeto en cuestión. En el contexto de esta tesis, se clasifican las maniobras orbitales en función de los distintos tipos de tecnologías de propulsión espacial (química y eléctrica). Los primeros se caracterizan por unos niveles de empuje altos y un uso ineficiente del combustible, mientras que los segundos poseen un consumo de combustible mucho más bajo que repercute en niveles de empuje relativamente pobres. No obstante, debido a la madurez tecnológica de ambos tipos de propulsión, se pueden obtener estimaciones fundamentadas de la capacidad o magnitud de control para ambos casos, por ejemplo, analizando misiones espaciales con uno u otro tipo de motor. Dichas estimaciones pueden ser de gran utilidad para limitar el espacio accesible (o conjunto alcanzable) para un objeto dado un estado u órbita inicial y una ventana temporal, que pueden corresponderse con el último estado orbital conocido y el tiempo transcurrido desde entonces. La aplicabilidad operativa de estos conceptos requiere de métodos computacionales (algoritmos) eficientes para calcular distancias orbitales en términos de control, también llamadas métricas de control, basadas tanto en propulsión química como eléctrica. En otras palabras, la viabilidad de la metodología propuesta en esta tesis pasa por el desarrollo de métricas de control eficientes que capturen las características dinámicas del entorno orbital terrestre y la tecnología propulsiva utilizada en una hipotética maniobra. Con este fin, se han desarrollado dos modelos

subrogados para maniobras orbitales, atendiendo a los tipos de propulsión espacial mencionados anteriormente, cuya aplicación no se limita al cálculo de distancias orbitales en términos de control si no que además pueden utilizarse para aproximar los límites del conjunto alcanzable. Dichos modelos subrogados permiten, a su vez, elaborar modelos de maniobra estadísticamente consistentes, con los que cuantificar la probabilidad de que un objeto maniobre en el proceso de correlación de medidas. Finalmente, se puede definir un sistema robusto y automatizado de vigilancia espacial mediante la incorporación de los modelos de maniobra propuestos en un filtro avanzado de seguimiento de múltiples blancos con capacidad de maniobra, capaz de gestionar no sólo la ambigüedad relativa a las maniobras per se, sino también la desaparición y nacimiento de objetos, así como falsas observaciones y detecciones omitidas por los sensores. Los resultados obtenidos en escenarios simulados indican que los métodos propuestos incrementan la capacidad actual de los sistemas de seguimiento, principalmente en lo referente al mantenimiento de custodia de objetos que maniobran. Sin embargo, el impacto dinámico de las maniobras conlleva altos niveles de incertidumbre en el conocimiento del estado orbital de los objetos, hasta el punto de degradar la calidad de los análisis de conjunción. Por lo tanto, aunque los métodos desarrollados son capaces de realizar un seguimiento ininterrumpido de objetos con capacidad de maniobra, se recomienda caracterizar las maniobras basándose en información adicional, tal como datos históricos. La inclusión de heurísticos de este tipo en la definición de los modos de maniobra puede ayudar a reducir de forma significativa la incertidumbre del estado en tiempos cercanos a la maniobra, incrementando de este modo la calidad del catálogo de objetos espaciales y permitiendo análisis de conjunción precisos en presencia de objetos maniobrables.

PUBLISHED AND SUBMITTED CONTENT

- Guillermo Escribano, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor and Diego Escobar, “Maneuver detection via combined heuristical and statistical methodologies", in *8th European Conference on Space Debris (ECDS)*, 2021. The contents from this conference paper can be considered an early version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.2, 3.3.2 and 4.3.1.

- Alejandro Pastor, Guillermo Escribano, Alejandro Cano, Manuel Sanjurjo Rivo, Cristina Pérez, Igone Urdampilleta and Diego Escobar. “Manoeuvre detection and estimation based on sensor and orbital data", in *8th European Conference on Space Debris (ECDS)*, 2021. The methodology proposed in this co-authored conference paper has been employed in validation tests within Section 4.1.1.

- Guillermo Escribano, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor and Diego Escobar, “Custody maintenance of low-thrust orbit raising satellite with surveillance radar data", in *3rd IAA Conference on Space Situational Awareness (ICSSA)*, 2022. The contents from this conference paper can be considered an early version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.3 , 3.1.4 and the results presented in Section 4.3.2.

- Alejandro Pastor, Guillermo Escribano, Manuel Sanjurjo-Rivo and Diego Escobar, “Satellite maneuver detection and estimation with optical survey observations,” *The Journal of the Astronautical Sciences*, vol. 69, no. 3, pp. 879–917, 2022. The methodology proposed in this co-authored journal paper has been employed in validation tests within Section 4.1.1.

- Lorenzo Porcelli, Alejandro Pastor, Alejandro Cano, Guillermo Es-

cribano, Manuel Sanjurjo-Rivo, Diego Escobar and Pierluigi Di Lizia. “Satellite maneuver detection and estimation with radar survey observations", *Acta Astronautica*, vol. 201, pp 274-287, 2022. DOI: https://doi.org/10.1016/j.actaastro.2022.08.021. The methodology proposed in this co-authored journal paper has been employed in validation tests within Section 4.1.1.

- Guillermo Escribano, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor and Diego Escobar, “Automatic maneuver detection and tracking of space objects in optical survey scenarios based on stochastic hybrid systems formulation”, *Advances in Space Research*, vol. 69, no. 9, pp. 3460–3477, 2022. DOI: https://doi.org/10.1016/j.asr.2022.02.034. The contents from this journal paper can be considered an early version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.2, 3.3.2 and 4.3.1. Results presented therein are included and discussed in Section 4.3.1.

- Guillermo Escribano, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor and Diego Escobar, “Data association for maneuvering space objects considering different control distance metrics", in *2022 AAS/AIAA Astrodynamics Specialist Conference*. The contents from this conference paper can be considered an early version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.2, 3.1.3, 3.1.4, 3.3.1 and the results presented in Section 4.4.

- Guillermo Escribano, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor, and Diego Escobar, “A filter for tracking non-cooperative low-thrust satellites using surveillance radar data,” *Acta Astronautica*, vol. 213, pp. 694–707, 2023. doi: https://doi.org/10.1016/j.actaastro.2023.09.026. The contents from this journal paper can be considered an early version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.3 and 3.1.4. Results presented therein are included and dis-

cussed in Section 4.3.2.

- Guillermo Escribano, Brandon Jones, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor and Diego Escobar. “A glmb filter for space objects with control metric based maneuver detection", in *33rd AAS/AIAA Spaceflight Mechanics Meeting*, 2023. The contents from this conference paper can be considered an early version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.2 and 3.2.3. Results presented therein are included and discussed in Section 4.4.

- Guillermo Escribano, Brandon Jones, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor and Diego Escobar. “A multiple target tracking filter for non-cooperative space objects", in *2nd NEO and Debris Detection Conference*, 2023. The contents from this conference paper can be considered an early version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.2 and 3.2.3.

- Guillermo Escribano, Brandon Jones, Manuel Sanjurjo-Rivo, Jan Siminski, Alejandro Pastor and Diego Escobar. “An efficient thrust-limited control metric: towards tractable reachability computations in the two-body problem", in *AIAA SciTech 2024 Forum*, 2024. The contents from this conference paper can be considered a preliminary version of the methodology developed in the context of this thesis, in particular the concepts discussed throughout Sections 3.1.3 and 3.2.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

IADC inter-agency space debris coordination committee

ASAT anti-satellite test

RSO resident space object

LEO low Earth orbit

ADR active debris removal

SST space surveillance and tracking

SSN Space Surveillance Network

TLE two-line element

SSA space situational awareness

COLA collision avoidance

GEO geosynchronous Earth orbit

HEO highly elliptical orbit

GNSS global navigation satellite services

MHT multiple hypothesis tracker

JPDA joint probabilistic data association

NN nearest-neighbour

PDE partial differential equation

MCMC Markov chain Monte Carlo

MCMCDA Markov chain Monte Carlo data association

RFS random finite set

FISST finite set statistics

PHD probability hypothesis density

SCR signal to clutter (ratio)

SNR signal to noise (ratio)

UCT uncorrelated tracklet

δ-GLMB delta generalized labeled multi-Bernoulli

GLMB generalized labeled multi-Bernoulli

JMS jump Markov systems

IMM interacting multiple model

MeMBer multi-target multi-Bernoulli

GPBI first order generalized pseudo-Bayesian estimator

GPBII second order generalized pseudo-Bayesian estimator

PDF probability density function

FPK Fokker Planck Kolmogorov

CK Chapman Kolmogorov

EKF extended Kalman filter

UKF unscented Kalman filter

UT unscented transform

CKF cubature Kalman filter

MHE moving horizon estimator

BLS batch least squares

MAP maximum a posteriori

PF particle filter

SMC sequential Monte Carlo

CDF cumulative distribution function

ESS effective sample size

SISR sequential importance sampling and resampling

BPF bootstrap particle filter

RPF regularized particle filter

GSF Gaussian sum filter

MH Metropolis Hastings

DREAM differential evolution adaptive metropolis

HMC Hamiltonian Monte Carlo

NUTS no U-turn sampler

IE input estimation

VSD variable state dimension

GARCH generalized autoregressive conditional heteroskedastic

STM state transition matrix

SSM state sensitivity matrix

GMM Gaussian mixture model

HCW Hill Clohessy Wiltshire

SGP simplified general perturbations

PoL patterns of life

TBP two-body problem

SRP solar radiation pressure

TFC thrust Fourier coefficients

TPBVP two-point boundary value problem

DSST Draper semi-analytical satellite theory

MEE modified equinoctial elements

HOSM higher order sensitivity matrix

AABB axis-aligned minimum bounding box

ITP interpolate, truncate and project

qMC quasi-Monte Carlo

KDE kernel density estimator

NLP non-linear programming

RMSE root mean squared error

LEOP launch and early operations phase

IOD initial orbit determination

AOS attributable optimized coordinate system

GTDS Goddard trajectory determination system

RAAN right ascension of the ascending node

AEGIS adaptive entropy-based Gaussian mixture information synthesis

OSPA optimal sub-pattern assignment

SNEES scaled normalized estimation error squared

SYMBOLS

$\mathbf{X}$	labeled multi-object state
$\mathbf{x}$	labeled individual object state
x	physical state
m	dynamical mode
ℓ	label or identifier
t	time
$\mathbb{R}^n$	n-dimensional subset of real-valued variables
$n_{\mathbf{X}}$	cardinality of multi-object state
n_x	dimensionality of single object state
$\mathbb{X}$	set of multi-object states
$\mathbb{M}$	finite set of dynamical modes
$\mathbb{L}$	finite set of unique labels
$\lvert\cdot\rvert$	cardinality operator
$\mathcal{F}(\cdot)$	enumerator for all finite subsets
$\Delta(\cdot)$	unique label indicator
$\mathcal{L}(\cdot)$	label indicator
$\delta_x[y]$	Kronecker's delta
$1_{\mathbb{X}}(x)$	set indicator function
$[\cdot]^{\mathbf{X}}$	multi-object exponential
$\langle\cdot,\cdot\rangle$	standard inner product
Z	multi-object observation
z	individual observation
n_Z	cardinality of multi-object observation
n_z	dimensionality of single object observation
θ	observation to target association map
π	multi-object state density
p	single object state density
ω	hypothesis or sample weight
ξ	hypothesis of observation to target associations

Ξ	set of observation to target association hypotheses
P_S	probability of survival
r_B	probability of birth
p_B	newborn target density
$\mathbb{B}$	label set of newborn targets
ψ	uncertain measurement to object likelihood
ϑ	mode transition probability
p_T	state transition density
p_+	prior state density
P_D	probability of detection
g	perfect measurement to object likelihood
κ	clutter density or clutter number
f	state space model
w	process noise
h	measurement mapping
v	measurement noise
$\mathcal{N}$	Gaussian distribution
$E[\cdot]$	expectation operator
μ	expected or mean value
$Cov[\cdot]$	covariance operator
Σ	covariance matrix
S	square-root covariance matrix
Σ_Q	process noise covariance
Σ_R	measurement noise covariance
K	Kalman gain
q	cubature point generating function
$\mathcal{X}$	set of cubature points
$\mathcal{Z}$	cubature points mapped to measurement space
$\delta\mathcal{X}$	weighted deviation of cubature points
$\delta\mathcal{Z}$	weighted deviation of cubature points mapped to measurement space
$\alpha', \beta', \gamma', \eta'$	unscented Kalman filter parameters

A	real-valued matrix
b	real-valued vector
F	gradient of state space model with respect to the state
J	scalar cost function
p_s	sampled density
$\delta(\cdot)$	Dirac's delta function
$\mathcal{U}$	uniform distribution
$P(\cdot)$	cumulative distribution of $p(\cdot)$
N_s	number of samples of a discrete distribution
ESS	effective sample size
$\mathcal{W}$	non-negative kernel function
ϱ	kernel bandwidth parameter
p_J	jumping distribution
α_R	acceptance ratio
$\mathcal{H}$	Hamiltonian scalar function
$\dot{x}$	momentum variables of Hamiltonian system
M	momentum mass matrix of Hamiltonian system
τ	fictitious time
∇_x	gradient operator
ε	stepsize of leapfrog scheme
N_ε	number of steps of leapfrog scheme
Δv	impulsive manoeuvre
ΔV	impulsive manoeuvre magnitude or total orbital transfer cost
u	control acceleration
d_u	control distance metric by Holzinger [1]
$\|\| \cdot \|\|_p$	p-norm of a vector
$\mathcal{R}$	reachable set
v	velocity
r	position
n	orbital mean motion
I_{sp}	specific impulse

m	mass
N_r	number of full revolutions for Lambert's problem
d_r	direction of Lambert's problem solution (pro- or retrograde)
$\mu_\oplus$	Earth standard gravitational parameter
a	semi-major axis
e	eccentricity
I	inclination
Ω	right ascension of the ascending node
ϖ	argument of periapsis
M	mean anomaly
E	eccentric anomaly
T	orbital period
d_I	impulsive control metric
a	acceleration
œ	state represented in the space of orbital elements
p	semi-latus rectum
h_1, k_1	non-singular components of eccentricity vector
L	true longitude
l	mean longitude
h_2, k_2	non-singular components of inclination vector
α, β	individual thrust Fourier coefficients
f_T	set of thrust Fourier coefficients
d_{LT}	low-thrust control distance metric
ϵ_1	mean anomaly phase change with respect to averaged mean motion
$\mathbb{Z}$	set of positive integer numbers
y, η	auxiliary variables used in the numerical reachability computation method
V_R	volume of the reachable set

J	Jacobian of the state transformation
ϕ	reachability inclusion function
ζ	feature vector for heuristic maneuver model
$\|\| \cdot \|\|_1$	scalar absolute value
r_p	periapsis radius of an orbit
V_{nE}	volume of an n-dimensional ellipsoid
V_{nR}	volume of an n-dimensional hyperrectangle
V_{AABB}	volume of the AABB approximation to the reachable set
c_R	reflectivity coefficient
λ	parameter of the Poisson distribution
d_M	Mahalanobis distance
$RMSE$	root mean squared error
c_D	atmospheric drag coefficient
B	ballistic coefficient
$B^\star$	B-star coefficient of SGP4 model
A	exposed area
ρ	atmospheric density
p_α	statistical significance threshold
$R_\oplus$	Earth radius
N_H	maximum number of parallel hypotheses
$D(\cdot)$	optimal sub-pattern assignment metric (OSPA)
c	cut-off distance for OSPA
d_b	base distance for OSPA
$SNEES$	scaled normalized estimation error squared

1. INTRODUCTION

Outer space is a valuable asset that has been exploited since the 1960s. Its unique characteristics have intrigued the most eloquent minds, perhaps because the motion of celestial bodies seems to follow beautifully simple physical laws and appears almost imperturbable - at least for human timescales. Arguably the most appealing feature of outer space is that objects can travel around massive bodies describing cyclic, or elliptical, orbits, driven by gravitational forces. Given the proper relative position and velocity, an object (or satellite) will remain in orbit indefinitely, as is the case of the Earth relative to the Sun: note this is extremely convenient for telecommunications and remote sensing. Once rocket technology enabled space access by the late 1950s, its unprecedented advantages for Earth observation, surveillance and telecommunications triggered a race for conquest between the United States and the Soviet Union. Soon, private enterprises realized the value of orbital services, with the first commercial satellite being launched in 1965. The number of objects in orbit has increased ever since, and by 1976, it was not lower than $3,866$. Concerned about this situation, Kessler and Cour-Palais performed an analysis on the evolution of the orbital population, based on previous knowledge on asteroid belt modelling [2]. Their study concluded that the Earth orbital population would see an exponential growth within a few decades, caused by cascaded in-orbit collisions, a situation that was later coined Kessler's syndrome.

During the following decades, these findings drew little interest among space operators and Earth's orbital congestion continued increasing without further consequences. There were, however, several groups of practitioners that identified the risk related to an uncontrolled growth in the debris population, pertaining to ESA [3], NASA [4], CNES [5] and the Russian federation [6], among others [7]. These resulted in the inter-agency space debris coordination committee (IADC) debris mitigation guidelines [8], thus formalizing the need to preserve the Earth orbital space. Unfortunately, these relatively mild suggestions saw poor acceptance except for the

geosynchronous Earth orbit (GEO) protected region. It was not until 2009 that Kessler's syndrome was broadly recognized from an operator's perspective: a derelict Russian satellite (Kosmos 2251) collided with an operational telecommunications spacecraft (Iridium 33) on February that same year. In fact, the latter took place two years after a Chinese anti-satellite test (ASAT) that created no less than 3,000 trackable debris fragments [9], and at a time where the cumulative number of satellite fragmentations (or break-ups) already exceeded 200 [10]. Most of these fragments remain in orbit to date, since only those with a sufficiently high area to mass ratio could be dragged to Earth in the lower altitude range ($h < 1,000$ km). As of June 2023, the number of resident space objects (RSOs) greater than 10 cm in size registered by the US Space Surveillance Network (SSN) climbs up to $26,648$, the 64 % of which are not related to active neither ceased operations [11]. Similarly, the ESA Space Debris Office estimates the number of catalogued and asserted objects to be greater than $32,000$, from which a total of objects not exceeding $10,000$ are currently in use [12]. Additional threats are envisaged due to the deployment of constellations aimed at providing global internet coverage. Starlink and OneWeb are two examples, currently accounting for almost 5,000 satellites in low Earth orbit (LEO) and expected to exceed 13,000 in case both constellations adhere to their plan[1].

Perhaps later than desired, the Kessler syndrome has been recognized, following the interest of space agencies and also numerous space activists. The Earth orbital space has yet to be formally (or institutionally) considered a finite natural resource, though the current trends in environmental consciousness have helped to showcase the problem of space debris to the general public. However, in this particular case, there is little the individual can do to mitigate the problem, rather than advocating for government and industry driven solutions. The most direct, yet costly, approach is active debris removal (ADR), by which dedicated missions aim at de-orbiting defunct objects, i.e. those with no maneuver capabilities. Note this is feasible for object sizes on the order of meters, e.g. rocket bodies, whose poten-

[1]The reader is referred to the table in newspace.im for a comprehensive list of satellite constellations in Earth orbit and their current status.

tial fragmentations would entail a significant debris growth. Still, there is a huge population of smaller fragments that can disable an active satellite or trigger a break-up event. Thus, it is of paramount importance not only to accurately track all the hazardous objects, this is, those that can cause a satellite failure (with a minimum size on the order of 10 cm), but also to assess the collision risk against the active object population.

To this end, there exist vast sensor networks distributed across the globe which continuously monitor the Earth orbital space. Object detections materialize in observations or measurements, which contain information regarding the physical state of the sensed subject. These data are then processed and used to elaborate a labeled map of RSOs, which is the main task of space surveillance and tracking (SST) systems. Historically, the most relevant SST system has been the SSN, built by the US in a military context and operated by the 18^{th} Space Defense Squadron. Unclassified data is made available at space-track.org in the two-line element (TLE) format, which provides information regarding the orbital state of catalogued targets and remains to date as one of the reference sources for orbital population status. In an effort to own a separate and independent surveillance framework, the EU has recently developed EUSST, much in line with the former Galileo approach to global navigation satellite services (GNSS) and seeking technological autonomy. They currently operate 38 ground-based sensors of different types [13], mainly surveillance and tracking radars, lasers and telescopes, allowing coverage of the full near-Earth environment. Their telescope network is distributed across the globe, yet radars are only located within the European territory in part due to their high cost and defence implications [14]. In parallel, private actors are showing interest in the deployment of sensor networks and SST systems, the most representative example of which may well be LeoLabs[2]. Based in California, they currently operate six radars spanning several longitudes and latitudes (Alaska, Texas, New Zealand, Costa Rica and Azores) and provide numerous space situational awareness (SSA) services such as satellite tracking and collision avoidance.

[2]Detailed information can be found at their website leolabs.space.

Satellite operators subscribed to such type of SSA services are relieved from conflict detection, being warned and asked for resolution if one of their satellites presents a potential close approach with another object, be it maneuverable or not. This framework is convenient to ensure a safe and orderly growth of space operations in the era of space tourism and space-based global internet coverage, though this relies on sensor networks acquiring vast amounts of data. To enable real time services, these measurements need to be autonomously processed by SST systems, which basically decide whether an incoming observation: 1) belongs to an already catalogued object, 2) can be associated to a previous uncorrelated observation to instantiate or confirm a new object or 3) shall remain uncorrelated as there is no sufficient evidence to prove its existence. This process, known as *observation-correlation*, mainly relies on orbit determination, orbit propagation and data association methods, and coherence between them is critical for real time SST.

1.1. Motivation

The dynamical characteristics of orbital motion are highly non-linear, and so are the transformations between the physical state and typically sensed quantities. To complicate things further, usual revisit times for objects in a surveillance framework are on the order of a few hours at best, sufficient for LEO satellites to perform several orbital revolutions. Fortunately, there is a vast knowledge regarding the physics that govern this environment, allowing for reasonably accurate state predictions in the context of SST. These predictions are used as prior evidence to statistically evaluate object detections and update the state of tracked individuals accordingly. Problems arise when observations cannot be elucidated from prior knowledge on the background population, yet from an expert's view they could hardly be caused by sensor malfunction. Small fragmentations or spawning are a potential cause, though the most frequent source is (unacknowledged) maneuvers.

Corrective actions are usually applied to maintain a satellite within pre-

scribed orbital bounds, commonly referred to as station-keeping maneuvers. The second most relevant type of maneuvers are aimed at avoiding potential collisions with other objects, known as collision avoidance (COLA) in the SSA jargon. Besides, spacecraft are often required to follow certain orientation or attitude laws, for they have to regularly perform slew maneuvers when attitude actuators such as reaction (or momentum) wheels saturate. The latter usually imply small deviations in translational motion, though appreciable for high precision applications or well-known dynamics, e.g. GEO. Other orbital maneuver types are associated to orbit re-positioning or relocation, disposal (move to graveyard orbit or atmospheric re-entry) and orbit acquisition, usually to raise the altitude or circularize the orbit after launch. These are only performed once or twice throughout the mission lifetime and present demanding control actions.

Orbital maneuvers threaten automated SST systems when they are not given proper notice by operators since objects are generally assumed to follow their natural ballistic motion in space. The two main problems caused by maneuvers in space are related to: 1) accuracy, undetected maneuvers necessarily introduce modelling bias that can degrade the system performance; and 2) custody loss, since the system may fail to associate a measurement after a maneuver as it lacks association evidence, thus not being able to keep track of the target. These issues are exacerbated by nonlinearities and the characteristic re-observation times of space surveillance since a relatively mild maneuver can lead to positional changes on the order of 100 km after several hours, well beyond the usual variance of capable SST systems, on the order of 10 m (see Figure 8-1 in [15] and [16]).

1.2. Gap Identification

Tracking maneuvering targets is by no means a novel research topic, as it has been studied since the dawn of radar technology [17]. The problem is that for considering dynamics other than natural (or ballistic) motion, there needs to exist a workable alternative model definition. A common approach in the tracking literature is to inflate the process noise to capture the lack of

dynamical knowledge stemming from unmodeled maneuvers, thus converting epistemic (i.e. modeling) uncertainty into aleatoric uncertainty. While this is a proven solution for low latency measurements or small control inputs [18], orbital corrections performed by operational satellites generally imply significant deviations from ballistic predictions. Thus, process noise levels need to be extremely high, to the point maneuver-conditioned probability dilutes and the true state falls to the tail of the distribution, this is, a region with very low probability density.

To overcome these limitations, some authors have resorted to ad-hoc maneuver models. As any other business-oriented activity, spaceborne missions aim to minimize costs and one of the limiting factors of satellites is mass. Historically, propellant efficiency for spacecraft has been low, as they mostly used chemical engines. Therefore, one of the cornerstones of mission design is trajectory optimization: maneuvers are designed to minimize mass consumption throughout the mission lifetime. The development of electric, or plasma, propulsion has enabled commercial missions with higher control requirements due to its lower propellant consumption. Launch costs in terms of price per kilogram are significant and hence resource (i.e. mass) optimization remains crucial. Inspired by this domain-specific knowledge, some maneuver detection and estimation methods applied to space surveillance make use of optimal control techniques. The orbital transfer between a past known state and a hypothesized observation is assumed optimal in a control sense and thus obtained as part of the solution to a (non-linear) optimal control problem.

The main concern regarding the above formulation is that state space filtering applies to stochastic problems, whereas this approach provides deterministic solutions. As a first approximation, one can use derivative information or cubature rules to perform rough uncertainty estimates, thereby assuming a distribution for the post-maneuver state. An additional caveat concerns the information content of observations themselves, from which accurate and / or complete state estimations cannot be retrieved in general. This implies that a subset of the state space is barely or to no degree observable, so the optimal control assumption leads to biased estimates by

definition. While the latter can be mitigated by the use of heuristics, it confronts with the foundations of data association and multi-object filtering. The lack of a proper prior density, i.e. a statistically consistent definition for the distribution of the state following a maneuver, conflicts with the observation-correlation process as hypotheses cannot be ranked based on measurement likelihoods. Therefore, SST system implementations that rely on these premises would be necessarily biased since 1) maneuver dynamics are defined on an ad-hoc basis and 2) data association hypotheses are not compared in a consistent manner.

Other approaches [19]–[21] have focused on the use of patterns of life (PoL) and historical data to elaborate consistent maneuver models. Under the assumption that spacecraft orbital corrections belong to one of a fixed number of types, it is natural to embrace a multiple-model formulation for state space filtering. In fact, this is consistent with the previous discussion of orbital maneuvers, provided their end goal can be used to infer their dynamical characteristics. If there is sufficient statistical knowledge on the maneuver pattern of certain target or target class, a multi-model approach of this kind would outperform process noise or optimal control based estimates. Note the latter only applies to cases when the target dynamics are governed by one of such proposed models, as otherwise the filter would produce biased estimates or even fail to maintain custody.

1.3. Objectives

In the interest of improving the capacity of current SST systems, it is convenient to devise a statistically consistent maneuvering target tracking framework. An adequate representation of the prior knowledge regarding the target dynamics can allow for optimal decision making in what concerns information usage, thereby enhancing data association performance. To this end, it is paramount to define a general maneuver model for space targets whose parametrization shall be reduced to the bare minimum. One of the main concerns regarding the definition of such a general model lies in the existence of different space propulsion systems, whose control char-

acteristics widely differ. Chemical engines boast high thrust levels to the detriment of propellant efficiency; electric propulsion on the other hand is significantly more frugal, yet it fails to deliver strong net forces. Accordingly, the former are fired for much shorter time spans and used occasionally, whereas the latter can be active for entire weeks, as required to attain the desired net orbital change.

As highlighted earlier, historical information can be made an integral part of the process. A framework that succeeds in tracking maneuvering targets without specific knowledge can still benefit from the inclusion of an a priori characterization of target patterns. The use of such patterns or heuristics is expected to yield state distributions with lower variances, leading to better informed data association decisions. This can potentially aid in resolving ambiguities arising from nearby objects and diffuse or diluted state distributions. In addition, reducing uncertainty for post-maneuver states improves the accuracy of conjunction analyses, and thus helps to prevent future in-orbit collisions.

This work is thus aimed at developing statistically consistent maneuver models for space surveillance applications, attending to the characteristics of both chemical and electric propulsion. Furthermore, it is of interest to elaborate additional heuristic maneuver modes that fully benefit from prior knowledge regarding orbital patterns. These models are to be implemented in advanced multiple maneuvering target tracking filters, rendering scalability and computational efficiency highly desirable features.

1.4. Contribution

The research conducted within this doctoral thesis has focused on identifying the needs of SST systems with regard to maneuvering target tracking. Defining a general model for orbital maneuvers, as a step forward from process noise formulations, has been the core subject of study. The orbital environment is characterized by non-linearity, affecting both natural forces and sensitivity to control actions. Finding efficient formulations that can be applied to a general case is hindered as simple linear, or higher order

polynomial, models fail at accurately capturing the periodic characteristics of orbital motion. In an effort to reduce the complexity of such model, the proposal presented herein advocates for characterizing the orbital space in terms of control, still retaining the most relevant features of the underlying dynamics. Control distance metrics, as they are commonly referred to, are not novel to SST, though the computational cost of former approaches poses tractability issues for complex statistical frameworks.

The work developed in the context of this doctoral thesis comprises a set of control distance metrics and surrogate dynamical models whose computational complexity allows for sampling-based and large scale applications. Each of the main space propulsion types, chemical and electric, deserves to be treated separately, so they have an associated control metric that best captures their propulsive characteristics. The computational efficiency of the aforementioned metrics is enhanced via the definition of simplified surrogate models that enable fast reachability analyses, which in turn provide a natural definition for maneuver models. The transition density for a general maneuver mode is then defined as a vacuous (least informative) distribution throughout the set of reachable (or accessible) states. The use of such an uninformative prior is intended to minimize estimation bias in the absence of additional maneuver knowledge, which can otherwise be included in alternative maneuver modes. Such a multiple model formulation responds to the special features of certain orbital maneuvers, in particular those related to station keeping.

Direct implementation of the proposed maneuver models in a conventional filtering framework is not straightforward. Therefore, possible filter implementations are elaborated first for single target, and then for multi-target environments. The former can be regarded as a proof of concept for maximal accuracy, and within this context, filters for typical space surveillance scenarios have been developed. The higher complexity of multi-object filtering demands resorting to simpler, though less accurate, filtering algorithms. Moreover, data association requires measurement likelihoods to be comparable so maneuver transition densities shall be adequately normalized. Reachability computations, and methods to calculate the volume

of the reachable set, are thus deemed necessary. To this end, an approximative approach is devised taking advantage of the developed surrogate models, which has potential applications to mission analysis and defense operations.

In summary, the main contributions of this thesis comprise the development of surrogate models for efficiently evaluating control distance metrics, tailored to both high- and low-thrust transfers; the definition of tractable least-informative (or vacuous) maneuver models that minimize estimation bias for uncooperative target tracking, enabled by efficient reachability computations; and the implementation of an advanced Bayesian multiple maneuvering target tracking filter aimed at space surveillance operations.

1.5. Organization of the thesis

The main components of the thesis and their interrelations are summarized in Figure 1.1. Recall the main aim is to develop a capable multiple maneuvering target tracking algorithm for space surveillance applications, so the literature review starts with a discussion on multi-object tracking methods in Section 2.1. Regardless of the type of multi-object filter used, state space filtering (or inference) is required for tracking activities, so different methods and techniques are presented and discussed throughout Section 2.2. In addition, a thorough literature review can help in identifying the current methodology gaps for tracking maneuverable spacecraft in a surveillance environment. To this end, Section 2.3 delves into the intricacies of maneuvering target tracking both from a general perspective and with application to space surveillance. The latter shows a lack of a proper maneuver model definition, one which considers non-linear dynamical features conditioned on the control performances of the target. The main contribution of the thesis is thus to provide dynamically consistent maneuver models for multi-object tracking applications, as those described in Section 3.3. These models make use of computationally efficient control distance metrics (see Section 3.1), which rely on surrogate dynamical approximations to reduce the computational complexity. Through the definition of such

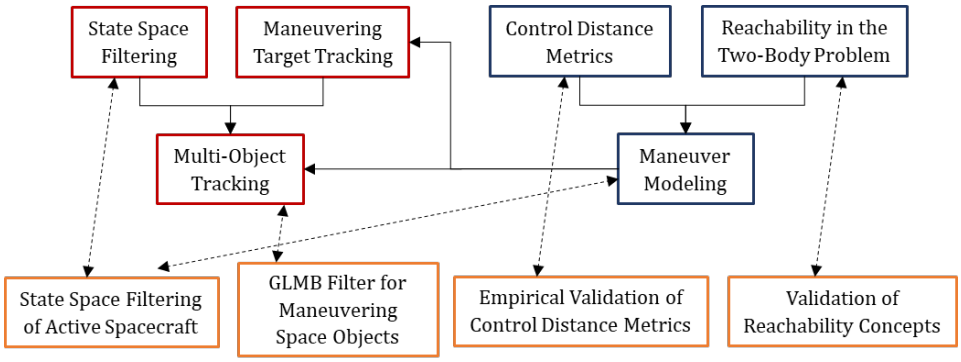

Figure 1.1: Breakdown of the thesis building blocks. Red items correspond to sections included in the literature review chapter, blue sections belong to the methodology chapter and orange elements are the different sections conforming the results chapter. Solid lines indicate knowledge dependence or contribution flow and dashed lines represent the relationship between results and methodology or literature review sections.

surrogate models, it is possible to perform efficient reachability computations in the two-body problem, as discussed in Section 3.2. The proposed methodology is validated through a series of tests of increasing complexity. First, the range of validity of the efficient control metrics is assessed in Section 4.1, followed by the validation of the above reachability concepts later in Section 4.2. The performance of the proposed maneuver models is evaluated in two different single target tracking scenarios in Section 4.3. Finally, the complete methodology is tested in Section 4.4 for an advanced multi-object filter implementation, enhanced with the developed maneuver models based on reachability.

2. LITERATURE REVIEW

The Earth orbital space is home to mostly quiescent objects, whose position highly depends upon environmental factors such as Earth's mass distribution, the presence of (massive) third bodies, solar radiation pressure and atmospheric drag. Models of varying complexity exist for determining the effects of these force fields on RSOs, which are essential for tracking purposes in cluttered regions. These models are usually the only source of evidence for data association, yet they are deemed to fail during maneuvering intervals. It is thus desirable to automatically infer when an object has maneuvered to make informed decisions on whether it has been observed afterwards. Maneuver detection and estimation methods, typically conceived as an extension to tracking algorithms, are significantly affected by the available information and the quality and timeliness of such information. In this regard, space surveillance typically employs two types of ground-based sensors: radars and optical telescopes. While the former are capable of measuring the target position up to certain accuracy, the latter only provide line-of-sight data. Measurements in the SST context are prolonged observations, referred to as *tracklets*, that render information about local temporal dependency. These can be processed as raw sequences or pre-conditioned to yield some compressed estimate referred to a single epoch [22]–[25]. However, the accuracy of sensed quantities is limited and information conveyed by a single tracklet is insufficient to derive a complete state estimate. The latter, coupled with the fact that objects can remain unobserved for extended periods, on the order of tens of hours, renders the problem of space object cataloguing rather challenging.

The remainder of the chapter comprises a thorough review on multiple maneuvering target tracking. Building upon the basics of Bayesian multi-object (or multi target) filtering discussed in Section 2.1, methods to perform Bayesian inference are presented in Section 2.2. Specific literature covering maneuvering target tracking is discussed later in Section 2.3, wherein Section 2.3.1 further elaborates on the state of the art in maneuvering target tracking applied to SST.

2.1. Multi-Object Tracking

Catalog build-up and maintenance in the context of SST can, and perhaps should, be cast as a multi-object filtering problem. Under the general formulation, a time varying number of objects are detected by some surveillance sensors, which partially observe the state of the targets. In fact, certain objects may not be detected and it is not uncommon for sensors to produce false alarms, i.e. measurements not originating from any trackable source. To complicate things further, observations are assumed untagged, so they are indistinguishable from each other; and targets could in principle spontaneously emerge or disappear, yet in general they are not expected to reappear as that can be attributed to a missed detection. Multi-object filters [26] thus handle uncertainty not only regarding the physical state of the system, i.e. single target state, but rather its configuration in terms of object cardinality. Statistical models that enable such flexibility in their description are seldom found, e.g. point process theory, and their influence in the tracking literature was almost non-existent by the time the first multiple target tracking filters were under development. To illustrate this fact, the reader is referred to the notion of expectation: no matter how complex a standard probability distribution function $p(x)$, the expected state is simply defined by an integral $E[x] = \int xp(x)dx$. In a multi-object scenario, expectation not only applies to the state of each individual target, information regarding its existence need also be conveyed, which necessarily yields a non-trivial definition of the expectation operator.

Former approaches to multiple target tracking separate the problem of measurement update, for which they follow the standard single object formulation, and data association, where most adopt hypothesis enumeration techniques. A good example is given by the multiple hypothesis tracker (MHT), which represents a direct and apparently simple approach to solve the data association problem. Applications in the context of SSA are ubiquitous as it is arguably the most adopted technique: several examples can be found in the literature [27], [28]. This method suggests elaborating parallel measurement to target association hypotheses based on a scoring function, usually the multi-object measurement likelihood. The expected population

state is conditioned on the maximum *a posteriori* hypothesis, i.e. the one with the highest score, though alternative hypotheses are kept alive to allow for ambiguity. The hypothesis ranking is updated as new observations are gathered and, in practice, the lowest-scoring ones are pruned to bound memory requirements. Note that as the number of tracks and objects increase, the combinatorial nature of this method poses tractability issues. In this regard, S.M. Gadaleta et al. [29] suggest utilizing a preliminary gating procedure to avoid the computation of associations that seem a priori unfeasible, commonly applied to other multi-target frameworks. MHT shows good performance in cluttered environments, supporting its application to LEO orbits. Of special note on the application of this method is that, in principle, one should elaborate and evaluate all the configuration hypotheses for a given measurement set, which not only considers individual measurement to object associations but also false alarms, missed detections, object birth, death and maneuvers, yielding tractability issues for complex applications.

A simplistic alternative to MHT can be found in joint probabilistic data association (JPDA) [30], in which one-to-one associations are replaced by joint updates. This way, the set of observations that fall within certain association gate are used to update the state of a target. Thus, JPDA considers association of multiple observations to multiple targets simultaneously, the latter being updated according to the individual measurement likelihoods. From a theoretical perspective, a statistical characterization of this kind is preferable to nearest-neighbour (NN) association, analogous to a MHT implementation with a single surviving hypothesis, since by assigning multiple weighted measurements to a target, erroneous associations can be resolved by future observations that may have fallen out of the association gate in the NN case. JPDA assumes the number of objects to be known a priori due to its decision-free approach and implies a sequential, also known as single-scan, treatment of measurement sets. Moreover, association of multiple observations to multiple targets can lead to state coalescence for closely separated objects.

Knowledgeable of the technical and theoretical pitfalls of the above

data association methods, S. Oh et al. [31] propose an alternative approach based on Markov chain Monte Carlo (MCMC). MCMC is known to be efficient in obtaining an approximation to unknown distributions with a modest computational effort, thus being well-suited for high dimensional applications such as multi-target tracking. Markov chain Monte Carlo data association (MCMCDA) aims at finding the optimal partition of the complete, or multi-scan, measurement set, formulated as a combinatorial optimization problem. Its randomized approach to finding a solution avoids enumeration of the entire hypothesis space, as opposed to MHT, thus improving computational efficiency while retaining the convergence properties of MCMC methods.

The above represent association-based approaches to Bayesian multi-object filtering, whose consistency has been objected [32] due to the lack of a proper definition for the probability of an association hypothesis. The main objection is that such probability, upon which decision-making is entirely based, is necessarily conditioned on the realization of the measurement set so it is not a valid prior for Bayesian filtering in a multi-object context. To overcome these limitations, Vo [32] suggests embracing the concept of random finite sets (RFSs), proposed by Mahler [33] under the framework of finite set statistics (FISST), which itself builds upon the broader point process theory [34]–[36]. An RFS is formally defined as a random variable with a finite random cardinality, this is, randomness lies within both its number of elements and values. The concept of RFS thus captures the characteristics of multi-object tracking models as both the target and measurement sets present time varying cardinality. In fact, the goal of FISST was to provide a principled solution to the multi-target tracking problem, with a first (tractable) proposal given by the probability hypothesis density (PHD) filter [37]. Therein, the state space is characterized by an intensity function (target intensity) that indicates the probability that a given number of targets exist at a certain location. Note this proposal does not explicitly generate tracks, i.e. it does not identify individual objects, but rather estimates the distribution of the population state. Therefore, it is best suited for high signal to clutter (SCR) and signal to noise (SNR) ratios, where clutter represents undesired detections arising from natural phenom-

ena and SNR refers to the absolute accuracy of the sensor, i.e. its ability to distinguish a target from the background.

When the SCR is sufficiently low, it is reasonable to assign a unique label or identifier to each trackable target, thereby representing the multi-object state space via a labeled RFS. Jones et al. [38] elaborate on the needs of advanced data association methods for space object cataloguing activities, which could potentially reduce the number of uncorrelated tracklets (UCTs), i.e. observations not attributed to any target due to absence of sufficient evidence. Therein, the authors provide a qualitative comparison of former methods with a special emphasis on their adequacy to the characteristics of space surveillance, among which they identify the delta generalized labeled multi-Bernoulli (δ-GLMB) [39] as the best candidate. The latter has seen a broad acceptance within the tracking community as it provides a full history of (labeled) object tracks via a closed form solution (conjugate prior) in the presence of undetected objects, false observations, and object birth. Thereafter, under any possible data association scenario, the multi-object state will always be backward traceable and take the form of a generalized labeled multi-Bernoulli (GLMB) distribution. In light of these findings, it is advisable to make use of the δ-GLMB, hereafter GLMB for ease of notation, formulation to perform multiple-maneuvering target tracking for SST. A detailed explanation of a valid GLMB formulation for jump Markov systems (JMS) is provided in the following section. The latter can be regarded as a general framework for multiple maneuvering target tracking since JMS are systems with multiple dynamical models from which only one is active between subsequent intervals and whose probability is only dependant on the immediately previous state, i.e. it fulfills the Markov property.

2.1.1. Generalized Labeled Multi-Bernoulli filter for JMS

The GLMB is an RFS filter that expresses the multi-object density in the form of a (generalized) labeled multi-Bernoulli distribution. It encompasses the concept of data association hypotheses by assigning certain individual probability to the existence of a joint population distribution, which

is itself conditioned on the history of individual target to observation correlations. The current section discusses the mathematical description of the formulation for JMS proposed by Punchihewa et al. [40], but note this is not the only possible means to accommodate for maneuvering targets. Ravago et al. [41] propose to hypothesize over the event of a maneuver, so that each multi-object density realization is conditioned on both a set of target to measurement correlations and a set of maneuvers. Explicitly hypothesizing over the active dynamical mode helps to eliminate the otherwise multi-modal nature of the JMS formulation resulting from [40], where single target dynamics can be filtered according to some sort of interacting multiple model (IMM) approach. Note this comes at the cost of a higher complexity when solving the assignment problem, since, despite retaining linearity, constraints on individual measurement to target associations need to be explicitly enforced [42]. Moreover, the hypothesis space is necessarily larger due to conditioning on the active mode so tractable implementations of the GLMB filter, i.e. featuring hypothesis truncation, may fail to consider the true event sequence at all times. Thereby, the filter adopts multi-modal single target distributions featuring different dynamical transition sequences and truncates them based on probability, effectively discarding unlikely states at a lower (single-target as opposed to hypothesis) level with the aim of reducing the size of the hypothesis space.

Preliminaries on Finite Set Statistics

Let the joint multi-target density of a set of jump Markov targets be expressed in terms of

$$\mathbf{X} = \left\{ \mathbf{x}_1 \quad \mathbf{x}_2 \quad \dots \quad \mathbf{x}_{n_\mathbf{X}} \right\}, \qquad \mathbf{x} = \left\{ x \quad m \quad \ell \right\}, \tag{2.1}$$

where $x \in \mathbb{R}^{n_x}$ represents an individual object state, $m \in \mathbb{M}$ indicates the dynamical model active between two subsequent epochs $[t_{k-1}, t_k]$ and $\ell \in \mathbb{L}_{0:k} \equiv \mathbb{L}_{0:k-1} \cup \mathbb{L}_k$ is a unique label devised for object identification. Thus, the multi-object state $\mathbf{X}$ is modeled as a random finite set according to FISST theory [33], i.e. a collection of random state vectors whose cardinality $|\mathbf{X}| = n_x$ is itself random. Let the set of all finite subsets of certain RFS

$\mathcal{S}$ be denoted as $\mathcal{F}(\mathcal{S})$ in what follows. Let also the unique label indicator be defined as

$$\Delta(\mathbf{X}) = \delta_{|\mathbf{X}|}\left[|\mathcal{L}(\mathbf{X})|\right], \tag{2.2}$$

being $\delta_x[y]$ Kronecker's delta, which takes the value of one if and only if $x = y$ and zero otherwise. Therefore, $\Delta(\mathbf{X})$ is a shorthand notation for enforcing unique labels, which must satisfy the condition that the cardinalities of the state $|\mathbf{X}|$ and the set of unique labels $|\mathcal{L}(\mathbf{X})|$ coincide. Similarly, it is convenient to introduce the indicator function,

$$1_{\mathbb{X}}(x) = \begin{cases} 1 & \text{if } x \in \mathbb{X}, \\ 0 & \text{otherwise}, \end{cases} \tag{2.3}$$

to express set containment. Provided certain single target state density f, a common practice in FISST is to make use of the multi-object exponential

$$[f]^{\mathbf{X}} = \prod_{\mathbf{x}\in\mathbf{X}} f(\mathbf{x}), \tag{2.4}$$

where $[f]^{\varnothing} = 1$ by convention. For two real-valued functions f and g, the standard inner product can be expressed as

$$\langle f, g\rangle = \int f(x)g(x)dx. \tag{2.5}$$

Given a set of observations belonging to difference sources

$$Z = \left\{ z_1 \quad z_2 \quad \ldots \quad z_{n_Z} \right\}, \tag{2.6}$$

with $z \in \mathbb{R}^{n_z}$, let an association map from the label space $\mathbb{L}$ to the space of individual measurements be defined as $\theta : \mathbb{L} \to \{0, 1, ..., |Z|\}$, which indicates a unique one-to-one correspondence between measurements and objects (or a missed detection if $\theta(\ell) = 0$). In the context of FISST, the Bernoulli distribution for an RFS is defined as

$$\pi(X) = \begin{cases} 1 - r & \text{if } X \equiv \varnothing, \\ r \cdot p(x) & \text{if } X \equiv \{x\}, \end{cases} \tag{2.7}$$

with r indicating the probability that state x exists and is distributed according to $p(\cdot)$. Note in this case bold letters are not used since so far all the referenced sets are unlabeled. Accordingly, a combination of n_x independent Bernoulli RFS distributions can be expressed as

$$\pi\left(\left\{x_{i_1}\ x_{i_2}\ ...\ x_{i_{n'_X}}\right\}\right) = \prod_{i=1}^{n_X}\left(1-r^{(i)}\right)\sum_{1\leqslant i_1\neq ...\neq i_{n'_X}\leqslant n_X}\prod_{k=1}^{n'_x}\frac{r^{(i_k)}p^{(i_k)}(x_i)}{1-r^{(i_k)}}, \quad (2.8)$$

where n_X is the number of Bernoulli components and $n'_X \in \{0, 1, ..., n_X\}$ is the cardinality of the RFS under evaluation [43]. Note that for some applications, it may be more convenient to parametrize such a multi-Bernoulli distribution at a population level, so that single target probabilities of existence (and state densities) are replaced by the probability of existence of certain multi-object configuration with a given multi-object state density. The latter is commonly referred to as generalized multi-Bernoulli distribution and takes the form

$$\pi(X) = \sum_{c\in\mathbb{C}}\omega^{(c)}\left[p^{(c)}\right]^X, \quad (2.9)$$

where $X = \left\{x_{i_1} \quad x_{i_2} \quad ... \quad x_{i_{n'_X}}\right\}$ is some multi-object realization, and $\omega^{(c)}$ and $\left[p^{(c)}\right]^X$ are the (non-negative) weight and collection of single target densities associated to the multi-object configuration c. Note that for $\pi(X)$ to be a proper density the sum of the weights over all possible configurations should be equal to one, i.e.

$$\sum_{c\in\mathbb{C}}\omega^{(c)} = 1. \quad (2.10)$$

Vo et al. [39] elaborate on the adequacy of this type of distribution to represent multi-object state densities. In particular, they propose the use of the generalized multi-Bernoulli distribution: a special case of multi-Bernoulli in which individual targets (and their density) only pertain to a certain label, thus allowing for keeping track of single target state trajectories. The GLMB distribution takes the form

$$\pi_k(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} \omega_k^{(I,\xi)} \delta_I(\mathcal{L}(\mathbf{X})) \left[p_k^{(\xi)}\right]^{\mathbf{X}}, \tag{2.11}$$

where $\xi \in \Xi$ represents a hypothesis for the history of associations $\xi = \theta_{0:k}$, $\omega_k^{(I,\xi)}$ are the weights associated to each hypothesis and every $p_k^{(\xi)}(\mathbf{x})$ is a proper probability density. Advantages of the GLMB filter over former, non FISST-based approaches mainly concern the provision of a closed form solution to the multi-target filtering problem, where the family of the filtered posterior distribution remains equal to that of the prior in the event of target birth, death, missed detection and false alarms. Note the latter is not generally the case for MHT [44] and JPDA [45] filters under the aforementioned events since their state representation does not allow for uncertainty in the number of targets. As compared to filters that consider RFS, such as the PHD [37] and the multi-target multi-Bernoulli (MeMBer) [43], the GLMB naturally allows for fast single target identification via labeling, so that recovering individual trajectories is straightforward.

GLMB Filter Equations

As previously discussed, it is convenient to express the state of a jump Markov target $\mathbf{x}$ as a combination of 1) a dynamical state x, e.g. Cartesian position-velocity coordinates, 2) the dynamical mode in effect m and 3) a unique label ℓ. This enables not only to recover individual target trajectories but also to identify maneuver intervals, without the need to explicitly hypothesize on them, i.e. there is no need to condition data association on the dynamical mode in effect as object's dynamics are inherently assumed hybrid at a multi-target filtering level. Herebelow, the joint prediction and update equations for the GLMB filter formulation of Punchihewa [40] are presented and discussed, emphasizing on relevant aspects regarding multiple maneuvering target tracking and maneuver detection. The posterior multi-object density for the filter in question can be expressed as

$$\pi(\mathbf{X}_+|Z_+) \propto \Delta(\mathbf{X}_+) \sum_{I,\xi,I_+,\theta_+} \omega^{(I,\xi)} \omega_{Z_+}^{(I,\xi,I_+,\theta_+)} \delta_{I_+}\left[\mathcal{L}(\mathbf{X}_+)\right] \left[p_{Z_+}^{(\xi,\theta_+)}\right]^{\mathbf{X}_+}, \tag{2.12}$$

with $I \in \mathcal{F}(\mathbb{L})$ and $\xi \in \Xi$ the prior label set and data association history, and $I_+ \in \mathcal{F}(\mathbb{L}_+)$ and $\theta_+ \in \Theta_+(I_+)$ the posterior label set and data association hypotheses for the current measurement frame. Thus, $\omega^{(I,\xi)}$ indicates the prior hypothesis weight and it is updated via the term

$$\omega_{Z_+}^{(I,\xi,I_+,\theta_+)} \propto \left[1 - \bar{P}_S^{(\xi)}\right]^{I \setminus I_+} \left[\bar{P}_S^{(\xi)}\right]^{I \cap I_+} \left[1 - r_{B,+}\right]^{\mathbb{B}_+ \setminus I_+} \left[r_{B,+}\right]^{\mathbb{B}_+ \cap I_+} \left[\bar{\psi}_{Z_+}^{(\xi,\theta_+)}\right]^{I_+}, \tag{2.13}$$

where the operator $\setminus$ indicates set difference, $r_{B,+}$ is the probability of target birth and $\mathbb{B}_+$ represents the label set corresponding to newborn targets. The (total) probability of survival for a given target

$$\bar{P}_S^{(\xi)}(\ell) = \sum_{m \in \mathbb{M}} \bar{P}_S^{(\xi)}(m, \ell) \tag{2.14}$$

is integrated over the past transition modes and also over all possible *dynamical* states at the previous step

$$\bar{P}_S^{(\xi)}(m, \ell) = \left\langle p^{(\xi)}(\cdot, m, \ell), P_S^{(m)}(\cdot, \ell) \right\rangle, \tag{2.15}$$

being $p^{(\xi)}(x, m, \ell)$ the filtered or posterior single target density and $P_S^{(m)}(x, \ell)$ the survival probability of target ℓ with state x conditioned on mode m. The likelihood corresponding to each individual target to measurement association is again marginalized, first over the current transition mode

$$\bar{\psi}_{Z_+}^{(\xi,\theta_+)}(\ell) = \sum_{m_+ \in \mathbb{M}} \bar{\psi}_{Z_+}^{(\xi,\theta_+)}(m_+, \ell) \tag{2.16}$$

and then the dynamical state at the current step

$$\bar{\psi}_{Z_+}^{(\xi,\theta_+)}(m_+, \ell) = \left\langle \bar{p}_+^{(\xi)}(\cdot, m_+, \ell), \psi_{Z_+}^{(\theta_+)}(\cdot, m_+, \ell) \right\rangle, \tag{2.17}$$

where the prior or predicted single target density

$$\begin{aligned} \bar{p}_+^{(\xi)}(x_+, m_+, \ell) = {} & 1_{\mathbb{B}_+}(\ell) p_B^{(m_+)}(x, \ell) + \\ & 1_{\mathbb{L}}(\ell) \frac{\sum_{m \in \mathbb{M}} \vartheta(m_+|m) \left\langle P_S^{(m)}(\cdot, \ell) p_{T,+}^{(m_+)}(x_+|\cdot), p^{(\xi)}(\cdot, m, \ell) \right\rangle}{\bar{P}_S^{(\xi)}(\ell)} \end{aligned} \tag{2.18}$$

is expressed in terms of the newborn target density $p_B^{(m_+)}(x,\ell)$, the mode transition probability $\vartheta(m_+|m)$ and the mode conditioned transition density $p_{T,+}^{(m_+)}(x_+|x)$ that is applied to the posterior at the previous step $p^{(\xi)}(x,m,\ell)$. Note marginalization is performed over the prior mode m to prevent the otherwise geometric growth $|\mathbb{M}|^k$, with k the number of filter steps, thereby following a second order generalized pseudo-Bayesian estimator (GPBII) for the single target multiple model filtering problem (see [46, Chapter 11]). As any proper transition density, the different $p_{T,+}^{(m_+)}$ satisfy Chapman Kolmogorov Equation (2.26) and represent the different dynamical models accessible to the target. The measurement to object likelihood

$$\psi_{Z_{\{1:|Z|\}}}^{(j)}(x_+,m_+,\ell) = \begin{cases} \dfrac{P_D^{(m_+)}(x_+,\ell)g^{(m_+)}(z_j|x_+,\ell)}{\kappa(z_j)+\delta_0\left[\kappa(z_j)\right]} & \text{if } j \in 1,...,|Z|, \\ 1-P_D^{(m_+)}(x_+,\ell) & \text{if } j=0, \end{cases} \tag{2.19}$$

equals to one minus the probability of detection $P_D^{(m_+)}(x_+,\ell)$ if the target is not assigned any observation and is otherwise proportional to $P_D^{(m_+)}(x_+,\ell)$ times the measurement likelihood function $g^{(m_+)}(z_j|x_+,\ell)$, which may or may not depend on the active dynamical mode depending on the application. Note the inclusion of the term $\kappa(z_j)+\delta_0\left[\kappa(z_j)\right]$, which is used to model false detections (clutter) that can be caused sensor malfunction or other anomalous phenomena. Finally, the posterior (or filtered) single target density is given by

$$p_{Z_+}^{(\xi,\theta_+)}(x_+,m_+,\ell) = \frac{\bar{p}_+^{(\xi)}(x_+,m_+,\ell)\psi_{Z_+}^{(\theta_+(\ell))}(x_+,m_+,\ell)}{\bar{\psi}_{Z_+}^{(\xi,\theta_+)}(m_+,\ell)} \tag{2.20}$$

under usual Bayesian assumptions. This density will then serve as prior density for the surviving targets at the following measurement step. An important caveat on the use of RFSs filters such as the GLMB is the relative complexity of the state extraction step, something that was already discussed at the beginning of this section. Usually, expectations on the multitarget state density are first computed over the hypothesis space, this resulting in the I,ξ,I^+ and θ^+ with the highest associated weight $\omega^{(I,\xi)}\omega_{Z_+}^{(I,\xi,I_+,\theta^+)}$.

Then, the expected mode at each filtering step can be readily obtained from

$$\hat{m}_+ = \underset{m_+ \in \mathbb{M}}{\operatorname{argmax}} \int p_{Z_+}^{(\xi,\theta_+)}(x_+, m_+|\ell)dx, \tag{2.21}$$

for each individual target by conditioning on the label of interest. This ultimately leads to the expected physical state of each object, expressed as

$$\hat{x}_+ = \int p_{Z_+}^{(\xi,\theta_+)}(x_+|\hat{m}_+, \ell)dx. \tag{2.22}$$

Nonetheless, the GLMB distribution is a rather complex information source and it is up to the practitioner to define the best procedure to extract the required knowledge, e.g. it might be of interest to recover individual state distributions without conditioning on the mode to perform conjunction analysis in an SSA scenario.

2.2. State Space Filtering

Multi-object filters focus on modelling the tracked entities (target intensity in PHD filter, labeled multi-object state in GLMB formulation, etc) and the mechanisms by which uncertain or imperfect observations are mapped to them. The final aim is to fuse information derived from prior knowledge on the state of the (multi-object) system with observations thereof. Thus, for a strong measurement mapping definition, as is the case of labeled filters, the multi-object measurement update step reduces to solving a collection of single target inference problems, this is, Equation (2.20). Under these circumstances, there are no sensible differences with respect to common state space filtering, or single target filtering, methods. Therein, the target (or system) is assumed to be subject to a dynamical model, which is usually some mathematical approximation of the complete system dynamics. This dynamical model is typically referred to as state-space model provided the system is observable, i.e. defined by some internal state dynamics and a set of observed or measured variables. The general (and formal) definition reads

$$x_k = f(x_{k-1}, t_k) + w(t_k), \tag{2.23}$$

$$z_k = h(x_k, t_k) + \nu(t_k), \tag{2.24}$$

where x_k represents the state of the system at time t_k and z_k is the measurement (observed variables) at time t_k. Equation (2.23) is the so-called *state equation* and describes the temporal evolution of the state in a Markov sense, i.e. the state at a any future time instant is only dependant on the last state (and not on the sequence of states). The term $w(t_k)$ in Equation (2.23) represents a Wiener process, introduced to account for a stochastic behavior of the system dynamics. The latter description also applies to $\nu(t_k)$ in Equation (2.24), which in this case is used to model the measurement noise. State space filtering refers to the process of estimating the density of the state $x_{0:N}$ according to a set of system observations $z_{1:N}$ at times $T = \{t_1, t_2, ..., t_N\}$. In general, one needs to solve for the evolution of the probability density function (PDF) $p(x,t)$ with time, which is given by Fokker Planck Kolmogorov (FPK) equation

$$\frac{\partial p(x,t)}{\partial t} = -\sum_{i=1}^{n_x} \frac{\partial}{\partial x_i}\Big[f(x,t)p(x,t)\Big] + \sum_{i=1}^{n_x}\sum_{j=1}^{n_x} \frac{\partial^2}{\partial x_i \partial x_j}\Big[w(t)p(x,t)\Big] \tag{2.25}$$

for continuous time systems. It is common practice to employ the generalization to discrete time systems, known as Chapman Kolmogorov (CK) equation

$$p_+(x_k) = \langle p_{T,+}(x_k|\cdot), p(\cdot)\rangle = \int p_{T,+}(x_k|x_{k-1})p(x_{k-1})dx_{k-1}. \tag{2.26}$$

Note in the above $p_{T,+}(x_k|x_{k-1})$ represents a state transition density compliant with the stochastic dynamical model defined in Equation (2.23). Similarly, one needs to find a solution to the measurement update or Bayesian inference (Equation (2.20)) which under certain simplifications, namely no clutter or missed detections, can be rewritten as

$$p(x_k) = \frac{p_+(x_k)g(z_k|x_k)}{\bar{g}(z_k)} \tag{2.27}$$

for the single target single model, i.e. non-maneuvering, case. Under Markovian assumptions, i.e. provided the current state x_k is only dependant

on the immediately previous realization x_{k-1}, sequentially solving Equations (2.26) and (2.27) yields an updated representation of the system uncertainty conditioned on the latest known information, this is $p(x_N|z_{1:N})$.

The remainder of this section discusses different single target filtering methods applied in the context of the thesis, highlighting their advantages and limitations. Special emphasis is given to non-linear filters since linearity assumptions generally fall short for space surveillance applications.

2.2.1. Kalman Filter

By the time operational filters were required to operate in on-board computers, the state of the art in processing power was very limited. Thereby simple and efficient methods to solve complex dynamical problems had to be devised, especially if one was eager to send a manned spacecraft to the moon. In this regard, R. Kalman came up with a solution to the linear filtering problem for discrete data sources, dubbed the *Kalman filter*. Assuming Gaussian distributions for both the measurements and state of the system, and always under linear transformations, Kalman devised the optimal recursion to combine the predicted and observed state of the system, with the ultimate goal of maximizing accuracy while minimizing the covariance of the results. Soon this approach was extended to approximate non-linear transformations, giving birth to many different implementations considering various uncertainty propagation methods. The simplest approach, based on a linearization around the expected or mean state, is the extended Kalman filter (EKF), valid for applications with *sufficiently* linear dynamics, e.g. short re-observation times or low degree of non-linearity. More advanced non-linear implementations make use of cubature (integration) rules to approximate FPK or CK equation, and most of them share the common scheme summarized in Algorithm 1. Therein, the predicted state density is approximated as $p_+(x_k) \sim \mathcal{N}(x_k; \hat{x}_k^+, \Sigma_k^+)$, where $\hat{x} = E[x]$ and $\Sigma = Cov[x]$. Note the adoption of the square-root version, in which the covariance matrix Σ is replaced by a lower triangular matrix resulting from the QR decomposition so that $\Sigma = SS^T$, or equivalently S =Tria(Σ). The latter helps to

ensure positive semi-definiteness of the co-variance matrices for computer applications, which can be lost due to round-off errors. The first $\hat{x}_k^+$ and second S_k^+ order moments of $p_+(x_k)$ are approximated by propagating a set of cubature points, defined according to the generating function $q(\cdot)$ that depends upon the Kalman filter scheme. Likewise, weights ω_i^m and ω_i^c are determined by the cubature rule used to approximate the transformation. In practice, the process noise term $w(t)$ in Equation (2.23) is approximated by a zero-mean Gaussian with square-root covariance S_Q. Similarly, the measurement noise $v(t)$ is substituted by a zero-mean Gaussian with covariance $\Sigma_R = S_R S_R^T$. The contribution of Kalman lies in the derivation of the gain matrix K_k, that corresponds to the optimal measurement update, i.e. solution to Equation (2.27), under linear Gaussian assumptions.

One of the simplest non-linear Kalman schemes is the unscented Kalman filter (UKF) [47], an evolution of the EKF in which the probability distribution is matched up to third degree. Based on the unscented transform (UT), the UKF makes use of a set of weighted points (sigma-points). These points are propagated according to the full non-linear mapping in Equation (2.23) and their resulting scatter is used to obtain the distribution at the subsequent step without linearization assumptions, thus greatly improving the accuracy of the a priori covariance estimate. The UT is essentially a means to approximate the statistics of a random variable after undergoing a non-linear transformation. For this purpose, a set $\mathcal{X}_i$ of $N_{UT} = 2n_x + 1$ points is defined according to the prior mean and covariance of such a variable or variable set, where n_x is the number of state variables of the system.

$$\begin{aligned} \mathcal{X}_1 &= \hat{x} \\ \mathcal{X}_i &= \hat{x} + \sqrt{(n_x + \gamma')\Sigma}\eta'_{i-1} \qquad i = 2, ..., n_x + 1 \\ \mathcal{X}_i &= \hat{x} - \sqrt{(n_x + \gamma')\Sigma}\eta'_{i-n_x-1} \quad i = n_x + 2, ..., 2n_x + 1 \end{aligned} \tag{2.28}$$

$$\begin{aligned} \omega_1^m &= \frac{\gamma'}{n_x+\gamma'} \\ \omega_1^c &= \frac{\gamma'}{n_x+\gamma'} + (1 - \alpha'^2 + \beta') \\ \omega_i^m &= \omega_i^c = \frac{1}{2(n_x+\gamma')} \qquad i = 2, ..., 2n_x + 1 \end{aligned} \tag{2.29}$$

Algorithm 1: Non-Linear Kalman Filter

INITIALIZATION: determine $\hat{x}_0^+ = x_0$ and $S_0^+ = S_0$, the state and square-root covariance factor, respectively

for $k > 1$ **do**

TIME UPDATE

1) Evaluate the uncertainty propagation points $i = 1, ..., m$:

$$\mathcal{X}_{i,k-1} = q(\hat{x}_{k-1}, S_{k-1})$$

2) Propagate the points $i = 1, .., m$:

$$\mathcal{X}_{i,k}^+ = f(\mathcal{X}_{i,k-1})$$

3) Estimate the predicted mean state:

$$\hat{x}_k^+ = \sum_{i=1}^{m} \omega_i^m \mathcal{X}_{i,k}^+$$

4) Obtain the predicted (square-root) error covariance $i = 1, ..., m$:

$$\delta\mathcal{X}_{i,k}^+ = \sqrt{\omega_i^c}(\mathcal{X}_{i,k}^+ - \hat{x}_k^+)$$

$$S_k^+ = \text{Tria}([\delta\mathcal{X}_{i,k}^+ \;\; S_Q])$$

MEASUREMENT UPDATE

1) Evaluate the uncertainty propagation points $i = 1, ..., m$:

$$\mathcal{X}_{i,k}^* = q(\hat{x}_k^+, S_k^+)$$

2) Obtain the estimated measurement set $i = 1, ..., m$:

$$\mathcal{Z}_{i,k}^* = h(\mathcal{X}_{i,k}^*)$$

3) Estimate the predicted measurement:

$$\hat{z}_k^* = \sum_{i=1}^{m} \omega_i^m \mathcal{Z}_{i,k}^*$$

4) Determine the (square-root) innovation covariance:

$$\delta\mathcal{Z}_{i,k}^* = \sqrt{\omega_i^c}(\mathcal{Z}_{i,k}^* - \hat{z}_k^*)$$

$$S_{zz,k}^* = \text{Tria}([\delta\mathcal{Z}_{i,k}^* \;\; S_R])$$

5) Estimate the (square-root) cross covariance matrix $i = 1, ..., m$:

$$\delta\mathcal{X}_{i,k}^* = \sqrt{\omega_i^c}(\mathcal{X}_{i,k}^* - \hat{x}_k^+)$$

$$\Sigma_{xz,k}^* = X_k^* Z_k^{*T}$$

6) Compute the Kalman Gain:

$$K_k = \Sigma_{xz,k}^*(S_{zz,k}^* S_{zz,k}^{*T})^{-1}$$

7) Update the state and (square-root)covariance estimates:

$$\hat{x}_k = \hat{x}_k^+ + K_k(z_k - \hat{z}_k^*)$$

$$S_k = \text{Tria}([\delta\mathcal{X}_k^* - K_k\delta\mathcal{Z}_k^* \;\; K_k S_R])$$

The definition of $\mathcal{X}_i$ and $\omega_i^{m,c}$ is given in Equations (2.28) and (2.29) for the scaled version of the UT, where η_k' is a vector of dimension n_x in which the only non-zero entry is located at element k. The weights with superscript ^{m} are used to estimate the mean and ^{c} are applicable to covariance calculations. γ' is a scaling parameter which determines the distance from the sigma points to the mean, and is defined according to $\gamma' = \alpha'^2(n+\lambda) - n_x$, where λ is usually set to 0 and α' may be tuned to render the sum $n+\gamma' \sim 3$. Finally the parameter β' can be used to introduce information regarding the expected distribution, where a value of $\beta' = 2$ corresponds to a Gaussian. There exist different alternatives to the UKF, such as the cubature Kalman filter (CKF) [48], in which the parameters γ', α' and β' are set to satisfy the spherical radial cubature rule for integration. Higher order filters [49], [50] have been derived in the literature, though their increase in accuracy is limited to the similarity of the approximated posterior to a (uni-modal) Gaussian.

Non-linear Kalman filters present a computationally efficient solution to the Bayesian inference problem in relatively benign settings, i.e. cases with mild non-linearity. There are applications, among which one can find space surveillance, where the predicted density is not well represented by a Gaussian, for instance showing some type of banana shape [51]. In these cases, approximating the state density as normally distributed usually leads to overestimating the support (defined as the non-zero probability region) of $p_+(x)$ since the mode no longer coincides with the expected value and higher order moments are not negligible. Moreover, for diffuse or partial observations, i.e. measurements with low information content, non-linearities in the measurement mapping can result in posterior densities that significantly deviate from Gaussianity, even for a perfectly Gaussian prior. The latter is exemplified by range-only measurements for an observer close to the target, where the posterior would present a banana shape for a sufficiently high prior uncertainty.

2.2.2. Moving Horizon Estimator

The moving horizon estimator (MHE) or batch least squares (BLS) is aimed at fitting a finite set of observations to the dynamical model in a least squares sense, under the assumption that the state follows a Gaussian probability distribution. Note in doing so the Markovian assumption, and thus its computational benefits, is relaxed in favor of a higher accuracy in the estimation of the expected state. The least-squares method seeks the solution to a curve fitting problem. It can be considered the first and simplest optimization method as it is devised for minimizing the L_2-norm (or Euclidean norm) of an overdetermined linear system of equations. Considering a system of the form

$$Ax = b, \tag{2.30}$$

the least squares solution x_{LS} to this equation is the one that minimizes the norm of the residuals, i.e.

$$\frac{\partial}{\partial x}\Big((Ax_{LS} - b)^T(Ax_{LS} - b)\Big) = 0, \tag{2.31}$$

thus leading to

$$x_{LS} = (A^TA)^{-1}A^Tb. \tag{2.32}$$

The basic linear least-squares problem can then be refined in order to account for differences in the scales (or uncertainties) of b_i, leading to the so-called generalized (or weighted) least-squares. When applied to estimation and curve-fitting problems considering uncertainty in the measurements, the weighted least squares approach provides a robust solution by including the information on the measurement covariance Σ_Q, so that

$$x_{GLS} = (A^T\Sigma_Q^{-1}A)^{-1}A^T\Sigma_Q^{-1}b. \tag{2.33}$$

Note the above solution is only applicable to linear systems, so certain modifications are required to target non-linear problems. Under the assumption of a Gaussian state probability density, the usual approach to non-linear least-squares consists in linearizing the system dynamics around a reference state, iteratively solving an approximate linear system centered at the esti-

mated state until convergence is found. The non-linear BLS or MHE [52] is aimed at determining an estimate of the state, usually referred to the estimation epoch x_N, that minimizes the residuals with respect to a batch of measurements $z_{1:N} = z(T)$, obtained at $T = \{t_1, t_2, ..., t_N\}$, provided certain knowledge on the initial estimate Σ. Usually, the measurement uncertainty at each time instant is known and given by the covariance Σ_Q. Accordingly, the output of the estimator conveys not only the aforementioned maximum a posteriori but its associated uncertainty in the form of a covariance. The measurement residuals take the form

$$\varepsilon(x_N^+) = \begin{bmatrix} z_1 - h(f(x_N^+, t_1)) \\ z_2 - h(f(x_N^+, t_2)) \\ \vdots \\ z_N - h(f(x_N^+, t_N)) \end{bmatrix} \approx \begin{bmatrix} z_1 - h(f(x_N, t_1)) \\ z_2 - h(f(x_N, t_2)) \\ \vdots \\ z_N - h(f(x_N, t_N)) \end{bmatrix} + F \cdot \left[x_N^+ - x_N \right], \tag{2.34}$$

where the initially non-linear system is approximated by a constant term plus a linear variation (assumed small) in the vicinity of a reference solution x_N. Therein, the linear mapping F is defined as

$$F = \begin{bmatrix} \frac{\partial h(f(x_N, t_1))}{\partial x_{N,1}} & \frac{\partial h(f(x_N, t_1))}{\partial x_{N,2}} & \cdots & \frac{\partial h(f(x_N, t_1))}{\partial x_{N,n}} \\ \frac{\partial h(f(x_N, t_2))}{\partial x_{N,1}} & \frac{\partial h(f(x_N, t_2))}{\partial x_{N,2}} & \cdots & \frac{\partial h(f(x_N, t_2))}{\partial x_{N,n}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial h(f(x_N, t_N))}{\partial x_{N,1}} & \frac{\partial h(f(x_N, t_N))}{\partial x_{N,2}} & \cdots & \frac{\partial h(f(x_N, t_N))}{\partial x_{N,n}} \end{bmatrix}. \tag{2.35}$$

Since the system is linearized, it is then possible to determine the optimal local update x_N^+ that minimizes the following loss (or cost) function

$$J = \varepsilon(x_N^+)^T \Sigma_Q^{-1} \varepsilon(x_N^+) + \left[x_N^+ - x_N \right]^T \Sigma^{-1} \left[x_N^+ - x_N \right], \tag{2.36}$$

which is given by the normal equations

$$x_N^+ = x_N + (\Sigma^{-1} + F^T \Sigma_Q^{-1} F)^{-1} \begin{bmatrix} z_1 - h(f(x_N, t_1)) \\ \vdots \\ z_N - h(f(x_N, t_N)) \end{bmatrix}. \tag{2.37}$$

The above leads to an iterative algorithm on the best estimate of the state, expressed as

$$x_N^{k+1} = x_N^k + (\Sigma^{-1} + F^T \Sigma_Q^{-1} F)^{-1} \left(\begin{bmatrix} z_1 - h(f(x_N, t_1)) \\ \vdots \\ z_N - h(f(x_N, t_N)) \end{bmatrix} \Sigma^{-1} \left[x_N - x_N^k \right] \right). \tag{2.38}$$

Note the difference in notation between the reference solution around which the dynamics are linearized, x_N, and the best estimate at the last iteration x_N^k since they need not necessarily coincide. In practice, the reference solution is updated according to the solver scheme used, as it has a strong influence on the computational complexity of the approach. The original method, dubbed Gauss-Newton algorithm, was conceived as an extension of Newton's method to find the minimum of a function. An evolution over Gauss-Newton, and currently an industry standard, is the Levenberg-Marquardt algorithm, which combines weighted and raw gradient information to improve robustness, though there has been more recent proposals [53].

An important caveat on the use of the MHE is that the derived Gaussian statistics do not necessarily represent a good estimate but rather a linear approximation of the scaled observation residuals centered at the maximum a posteriori (MAP) estimate. This is a well-known problem and has been assessed in the literature under the subject of covariance (or uncertainty) realism [54], [55]. Nevertheless, the MHE has long been and is still used for orbit determination, i.e. tracking space objects, due to long observation gaps where uncertainty growth leads to divergence for non-linear Kalman filter implementations.

2.2.3. Sequential Monte Carlo

Advances in computational technology within the last few decades have triggered the development of new approaches, especially those whose associated computational burden rendered them infeasible in the past: this is the case of Monte Carlo methods. Former Bayesian inference approaches assume the posterior distribution belongs to a given family, which is arguably a good approximation. Inconsistencies between the parameter-fitted distribution and the actual (analytical) one are critical for systems exhibiting high non-linearities. The latter fact triggered the adoption of Monte Carlo methods, whose initial aim was to numerically approximate complex high dimensional integrals. When applied to sequential state space filtering, Monte Carlo approaches are typically referred to as particle filters (PFs) or sequential Monte Carlo (SMC). PFs rely on sampled densities, as opposed to predefined distribution families, so their PDF representation consists in a set of sample realizations (particles) and their associated weights. Each particle is then propagated according to the system dynamics, thus exactly following the underlying transformation at the sampled states. Therefore, the sampled population is a discrete approximation of the posterior distribution that converges to the true solution as the number of particles increases. Conceptually, PFs are appealing since they allow for an arbitrarily accurate non-linear propagation of the state statistics [56].

The basic concept underlying SMC methods is the use of a sampled density, i.e. a collection of particles or states aimed at approximating a smooth statistical distribution. In doing so, each of the different particles need to have an associated weight, which indicates its relative importance with respect to the rest of particles. A sampled distribution (or sampled density) $p_s(x) \approx p(x)$ of N_s particles can be defined as follows

$$p(x) \approx p_s(x) = \sum_{i=1}^{N_s} \omega_i \delta(x - x_i), \tag{2.39}$$

where ω_i is the weight of the i^{th} particle and $\delta(\cdot)$ is the Dirac's delta function, which takes the value $\delta(0) > 1$ and zero otherwise, and satisfies

$\int \delta(x)dx = 1$. Similarly, one can define the cumulative distribution function (CDF) $P_s(x)$ of a sampled density $p_s(x)$ as follows

$$P_s(x) = \sum_{i \in \mathcal{I}(x)} \omega_i, \qquad \mathcal{I}(x) = \{i | x_i \leqslant x\}, \tag{2.40}$$

which is simply the cumulative sum of the weights ω_i whose associated components x_i are lower or equal than x. It can be demonstrated that as the number of samples $N_s \to \infty$, the sampled distribution converges to the smooth analytical PDF and, so that, for sufficiently large N_s, the sampled density can be assumed unbiased. Numerically approximating smooth distributions by sampled densities can be beneficial in a number of applications; unfortunately, sampling from arbitrary distributions is not trivial, especially in high dimensional spaces.

In order to obtain samples from a general distribution $p(x)$, it is necessary to compute first its cumulative distribution $P(x) = \int p(x)dx$ and then the inverse of such CDF, $P^{-1}(x)$. Then, one can sample u_i from a uniform distribution $\mathcal{U}(0,1)$ bounded between 0 and 1. By assuming that $u_i = P(x_i)$, one can readily obtain samples as $x_i = P^{-1}(u_i)$. There are, however, a few caveats on the applicability of this approach: 1) obtaining the cumulative distribution may require some computational effort, 2) finding the inverse CDF may be even more costly or even impossible (in case $P(x)$ is a non-invertible function). The latter is inherent to multivariate distributions since multiple state realizations have the same associated probability, rendering the transformation for $P(x)$ not bijective by definition. Efficiently sampling from arbitrary distributions is thus still an issue as of today, and the recommended approach is in fact one of the pillars of Monte Carlo methods: *importance sampling*.

Given a distribution $p'(x)$ has full support over $p(x)$, i.e.

$$p'(x) > 0 \quad \forall\, x \mid p(x) > 0, \tag{2.41}$$

it can be a good candidate importance function of $p(x)$ provided if can be efficiently sampled from. Depending on the target distribution to sample,

it is convenient to select a different importance function: for bounded intervals one may use the uniform distribution, while for semi-infinite and infinite intervals a valid choice may be the log-normal and normal distributions, respectively. The definition of importance distribution is particularly useful for Monte Carlo integration, to which SMC owes its nomenclature. Assume one is eager to compute a higher-dimensional integral given by

$$\mathrm{I} = \int f(x)dx. \tag{2.42}$$

In principle, it is possible to factorize $f(x) = g(x)p(x)$, where $p(x)$ is an arbitrary probability density function satisfying $p(x) \geqslant 0$ and $\int p(x)dx = 1$, without altering the resulting value of the integral I. Recalling the concept of sampled densities, one can approximate the integral I by means of a sampled density $p_s(x) \approx p(x)$, yielding

$$\mathrm{I} \approx \mathrm{I}_s = \sum_{i=1}^{N_s} \omega_i g(x_i). \tag{2.43}$$

Integrals of the form given by $\mathrm{I} = \int g(x)p(x)dx$ may be known to be better solved by means of a certain probability density $p(x)$, though it can occur that directly sampling from such density is computationally expensive. Under these circumstances, one can make use of an importance function $p'(x)$, from which a sampled distribution $p'_s(x)$ can be readily obtained, so that

$$\begin{aligned}
\mathrm{I} &= \int g(x)p(x)dx \\
&= \int g(x)p'(x)\frac{p(x)}{p'(x)}dx \\
&\approx \sum_{i=1}^{N_s} g(x_i)\omega_i \frac{p(x_i)}{p'(x_i)}.
\end{aligned}$$

The above result indicates one can sample from a distribution $p(x)$ given

an importance density $p'(x)$ as follows

$$p_s(x) = \sum_{i=1}^{N_s} \tilde{\omega}_i \delta(x - x_i), \qquad \tilde{\omega}_i = \omega_i \frac{p(x_i)}{p'(x_i)}. \tag{2.44}$$

These concepts become particularly useful when applied to the Bayesian inference problem, this is, solving Equations (2.26) and (2.27). Assume the distributions of $w(t)$ and $v(t)$ are known at any time instant and the initial state density, denoted as $p(x_0)$, is also known. The goal is to infer the distribution $p_+(x_k|z_{k-1})$, which can be derived by sampling from $p(x_{k-1}, w(t_k)|z_{k-1})$. Note that since $w(t)$, and also $v(t)$, are independent from both x_{k-1} and z_{k-1}, they can be sampled independently at any time instant. Then, the discrete probability $p_{s,+}(x_k|z_{k-1})$ can be defined as

$$p_s(x_{k-1}|z_{k-1}) = \sum_{i=1}^{N_s} \omega_{k-1,i} \delta(x_{k-1} - x_{k-1,i}), \tag{2.45}$$

$$p_{s,+}(x_k|z_{k-1}) = \sum_{i=1}^{N_s} \omega_{k-1,i} \delta(x_k - f(x_{k-1,i}, t_k) - w_i), \tag{2.46}$$

so that $\omega_{k,i}^+ = \omega_{k-1,i}$ and $x_{k,i} = f(x_{k-1,i}, t_k) + w_i$, i.e. the weight associated to each particle is kept the same while its corresponding state is propagated according to the (stochastic) dynamical model. After approximating the prior distribution for x_k, the measurement update can be performed following Bayes' theorem, such that

$$p(x_k|z_k) = \frac{g(z_k|x_k, z_{k-1}) p_+(x_k|z_{k-1})}{g(z_k|z_{k-1})}. \tag{2.47}$$

A reasonable assumption is to consider the dynamical model only dependant on the previous state, hence the process is Markov and the density $g(z_k|x_k, z_{k-1}) = g(z_k|x_k)$. Recalling the concept of importance sampling, the distribution of Equation (2.47) can be expressed solely as a function of $g(z_k|x_k)p_+(x_k|z_{k-1})$, thus

$$p(x_k|z_k) = \frac{g(z_k|x_k) p_+(x_k|z_{k-1})}{\int g(z_k|x_k) p_+(x_k|z_{k-1}) dx}. \tag{2.48}$$

The latter can be readily obtained provided a sampled approximation for the prior since evaluating the likelihood $g(z_k|x_k)$ is inexpensive in general. Finally, the posterior density $p_s(x_k|z_k)$ takes the form

$$p_s(x_k|z_k) = \frac{\sum_{i=1}^{N_s} \omega_{k,i}^{+} g(z_k|x_{k,i}) \delta(x_k - x_{k,i})}{\sum_{i=1}^{N_s} \omega_{k,i}^{+} g(z_k|x_{k,i})}, \tag{2.49}$$

so that the location of the particles is not modified with respect to Equation (2.47) but the weights are updated according to the information conveyed by the observation z_k.

This sequential filtering algorithm presents a key implementation issue: *sample degeneracy*. Note that the same particle set that was drawn initially is subsequently propagated according to the model dynamics and only the weights are updated to account for incoming observations. Thereby, as the propagation advances, higher weights will concentrate in regions of high probability leading to homogenization of the sample density, i.e. very few particles define the sampled distribution. In fact, if the model used to propagate the sampled distribution does not properly capture the stochasticity in the dynamics, all the weights may drop to zero given they are sufficiently far from the observations. In this regard, a very simple algorithm can be used to rejuvenate the particles and avoid filter divergence: *resampling*. An illustrative example of a resampling scheme can be consulted in Algorithm 2, which makes use of the CDF of the particle weights and exploits the fact that samples with a higher weight present a higher share of the cumulative distribution.

This way, it is possible to avoid propagating particles that add little to no information regarding its sampled density. Nonetheless, resampling at each time step is not recommended as the space covered by the samples is reduced, potentially failing to approximate multi-modal densities or simply the tails of the distribution. For this particular purpose, the effective sample

Algorithm 2: Resampling procedure

Given $p_s(x) = \sum_{i=1}^{N_s} \omega_i \delta(x - x_i)$:

Compute $P_s(x) = \sum_{i \in \mathcal{I}(x)} \omega_i \quad \text{where} \quad \mathcal{I}(x) = \{i | x_i \leqslant x\}$

Draw N_s random samples u_i from $\mathcal{U}(0,1)$

Assign the new $x_i = x_j$ according to $P(x_{j-1}) \leqslant u_i \leqslant P(x_j)$

Update the weights as $\omega_i = \frac{1}{N_s}$

size (ESS)

$$ESS = \frac{1}{\sum \omega_i^2}, \tag{2.50}$$

provides a simple metric [57] to monitor the homogeneity of the particle population. In general, in case the ESS falls below certain threshold, commonly $ESS_{min} = N_s/2$, the sampled density is assumed to have a significant number of non-informative particles so performing a resampling step is advisable.

Bootstrap Particle Filter

The combination of sequential importance sampling and resampling (SISR) conforms the bootstrap particle filter (BPF). Though not the most advanced PF method, the BPF is a good starting point due to its simplicity and intuitiveness, as indicated by the reduced number of steps in Algorithm 3. The number of samples required for PFs to perform vary depending on the application, but they are typically on the order of $N_s \sim O(10^3 - 10^4)$. There is a big caveat on the use of this type of filters [58], especially for high-dimensional systems, as the sample size should increase exponentially with the number of state variables in order to avoid collapse, this is, converging to a sampled population with extremely low state variability. There are applications, such as space surveillance, in which the dynamical model $f(\cdot)$ can be expensive to evaluate, e.g. simulating multi-day orbital trajec-

tories with high fidelity atmospheric density models, potentially hindering the applicability of this type of filters. It is, however, important to recall that each particle can be propagated independently, thus benefiting from the use of parallel and distributed computing. In fact, opposed to methods based on optimization or differential algebra, the computational time of a particle filter can a priori be estimated so an online SMC implementation is feasible given the computational burden can be afforded.

Algorithm 3: Bootstrap Particle Filter

INITIALIZATION: sample N_s particles $x_{0,i}$ from $p(x_0)$ with weights

$\omega_i = \frac{1}{N_s}$

for $k > 1$ **do**

IMPORTANCE SAMPLING

1) Approximate $p_+(x_k|z_{k-1})$, assuming a distribution for $w(t_k)$:

$$\omega_{i,k}^- = \omega_{i,k-1} \qquad x_{k,i}^- = f(x_{k-1,i}, t_k) + w_i$$

2) Apply the measurement update to obtain $p(x_k|z_k)$:

$$\omega_{i,k} = \frac{\omega_{i,k}^- g(z_k|x_{k,i})}{\sum_{i=1}^{N_s} \omega_{i,k}^- g(z_k|x_{k,i})} \qquad x_{k,i} = x_{k,i}^-$$

RESAMPLING

3) Compute the Effective Sample Size:

$$ESS = \frac{1}{\sum_{i=1}^{N_s} \omega_{i,k}^2}$$

if $ESS \leqslant ESS_{min}$ **then**

4) Compute the cumulative distribution of $p_s(x_k|z_k)$:

$$P_s(x_k|z_k) = \sum_{i \in \mathcal{I}(x)} \omega_{i,k} \qquad \mathcal{I}(x) = \{i | x_k \leqslant x\}$$

5) Draw u_i from $\mathcal{U}(0,1)$ and update the particles according to:

$$\omega_{i,k} = \frac{1}{N} \qquad x_{k,i} = x_{k,j} \qquad P_s(x_{k,j-1}|z_k) \leqslant u_i \leqslant P_s(x_{k,j}|z_k)$$

Regularized Particle Filter

Ideally, a sampled distribution should present similarly weighted samples to better represent the support of the approximated density. Note this is not a sufficient condition for the latter since, especially in high-dimensional spaces, it needs to be coupled with heterogeneous sample realizations. This renders naive resampling techniques as the one used in the BPF rather inefficient, ultimately leading to particle depletion, i.e. homogeneous sample populations. Various methods have been proposed to circumvent this issue, among which one can find the regularized particle filter (RPF). Relying on smoothing (or regularization), the RPF aims at approximating the sampled density to a continuous form by assigning a kernel to each sample realization as

$$p_c(x) \approx \frac{1}{N_s \varrho} \sum_{i=1}^{N_s} \mathcal{W}\left(\frac{x - x_i}{\varrho}\right), \tag{2.51}$$

where $\mathcal{W}(\cdot)$ is a non-negative *window* or *tapering* function and ϱ its bandwidth. A thorough discussion on this type of filters can be consulted in [59], where guidelines on filter design are supported by analyses on window function and bandwidth selection. They suggest a good starting point is to adopt a Gaussian kernel $K \sim \mathcal{N}(x; \hat{x}_i, \Sigma_i)$ with a co-variance set according to Silverman's rule of thumb [60], i.e.

$$\Sigma_i = \frac{\left(\frac{4}{n_x+2}\right)^{\frac{1}{n_x+4}} N_s^{\frac{-1}{n_x+4}}}{2} \mathrm{Cov}[x_i]. \tag{2.52}$$

The latter allows for an adaptive tuning of the kernel bandwidth based on the overall population co-variance, enhancing particle diversity and delaying particle depletion. Besides the resampling scheme, the RPF essentially follows the scheme in Algorithm 3 and particle rejuvenation is typically performed only when the ESS falls below certain threshold.

PFs or SMC methods are recommended for applications with significant non-linearities in the state or measurement spaces and relatively high

process and / or measurement noise levels as they can accommodate for any type of colored noise ($w(t)$ and $\nu(t)$) and dynamical and measurement models. There are, however, certain applications for which PFs may fail to converge: systems with relatively high process noise and accurate observations. Under these circumstances, very few samples from the predicted distribution are expected to fall within the support of the posterior, demanding an additional procedure that ensures a sufficiently accurate representation of the filtered density.

2.2.4. Markov Chain Monte Carlo

The aforementioned importance sampling approach presents good convergence provided the support of the importance function fully covers that of the prior, i.e. the one it is difficult to sample from. Problems arise when the support of the posterior, approximated by updating the population weights proportional to the observation likelihood in SMC, is limited to a relatively small region within the support of the prior. This directly incurs in particle degeneracy since most or even all the samples will present negligible weights, thereby leading to premature filter divergence. Under these circumstances, posterior exploration methods such as MCMC can be extremely useful as they directly focus on approximating the posterior given some surrogate prior or importance function. Bayes' rule,

$$p(x|z) \propto g(z|x)p_+(x), \tag{2.53}$$

states that the probability density of the state x conditioned on the observation z is directly proportional to the likelihood of the state realization $g(z|x)$ times the probability of the realization itself $p_+(x)$. Opposed to the particle filtering approach, in this case the problem does not consider solving FPK Equation (2.25) (or, for discrete systems, the CK Equation (2.26)) since a sufficiently accurate representation of $p_+(x)$ is assumed to be available. The aim is to efficiently sample from the posterior $g(z|x)p_+(x)$ making use of a Markov Chain. There are two general approaches to explore the posterior distribution using MCMC: 1) assuming a proposal distribution and

a starting point or 2) generating samples based on a single starting point. In the former case, subsequent samples are drawn from an evolving proposal distribution and its acceptance is based on a probability ratio. In the latter, however, a common approach is to model the posterior density as a Hamiltonian scalar function, so that the initial state follows some fictitious dynamics whose trajectory directly serves as a sampling source. In either case, the result is a set (or *chain*) of samples that eventually converge to the target posterior distribution. Examples of these two approaches are provided in the following sections, which also discuss their applicability to inference in space surveillance applications.

Metropolis-Hastings

The Metropolis Hastings (MH) algorithm [61] can be regarded as the precursor of all MCMC methods. It requires a prior *jumping* distribution $p_J(x^{(k+1)}|x^{(k)})$ and an initial point $x^{(0)}$ to generate candidate samples $x^{(k+1)\prime}$, distributed as

$$x^{(k+1)\prime} \sim p_J(x^{(k+1)}|x^{(k)}). \tag{2.54}$$

Acceptance of these candidate samples is based on the ratio

$$\alpha'_R = \frac{p\left(x^{(k+1)\prime}|z\right)}{p\left(x^{(k)}|z\right)}, \tag{2.55}$$

so that if $\alpha'_R \geqslant u$, where $u \sim \mathcal{U}(0,1)$, the proposed sample $x^{(k+1)} = x^{(k+1)\prime}$ is accepted (otherwise $x^{(k+1)} = x^{(k)}$). Since samples are randomly proposed and accepted, the method is expected to statistically converge to the desired posterior distribution $p(x|z)$ for a sufficiently high number of samples. This is attributed to the fact that the sampling process simulates a random walk whose moves are governed by the ratio α'_R, so that the chain will be captured by high probability regions and will remain there. Conveniently, design choices for this algorithm are limited to the jumping distribution, which can be modeled as a Gaussian centered on $x^{(k)}$ with a certain covariance matrix. Note that the magnitude of the covariance dictates the distance travelled by successive candidates, so a proper tuning of these parameters

is crucial to avoid frequent travels to low probability regions. The MH algorithm is not particularly suited to high dimensional problems since each candidate sample comprises an entirely new state realization. This renders tuning of the covariance parameters a crucial step for efficiently drawing samples in the support of $p(x|z)$, i.e. samples whose posterior density is not negligible. Different MH-based method have been proposed to deal with this issue, for instance, the Gibbs sampler [62] only considers varying a single state dimension at each subsequent candidate sample, thereby enhancing the probability of accepting subsequent candidates. Modern approaches, such as the differential evolution adaptive metropolis (DREAM) algorithm [63], seek to improve convergence by simulating multiple parallel chains that benefit from differential evolution. DREAM automatically adapts the subspace that is updated (much like the Gibbs sampler) but also the scale and orientation of the jumping (or proposal) distribution, trying to maximize the distance travelled by each chain. Successful applications of this algorithm for high-dimensional Bayesian inference can be consulted in [64], [65], with a focus on parameter estimation for complex environmental models. MH-like methods are well suited for situations in which there is some knowledge on the expected variance of the posterior, so the jumping distribution can accommodate for this, e.g. through some scaling. MH methods can be regarded as a high-performance version of classic importance sampling since they can tractable approximate complex posterior distributions in an efficient manner and are less affected by the selection of the importance function, via the use of a jumping distribution.

Hybrid Monte Carlo

Hamiltonian Monte Carlo (HMC), first proposed by Duane et al. [66], is a type of MH algorithm with a concrete definition of the *proposal* distribution. The HMC algorithm is initialized with a single starting point $x^{(0)}$, whose dynamics are assumed to be subject to the Hamiltonian defined in Equation (2.56). Therein, the state is augmented with some momentum vector $\dot{x}$, whose impact in the *kinetic energy* component is influenced by some mass matrix M. The (fictitious) temporal evolution of the state vari-

ables and their momenta is then governed by Equation (2.57).

$$\mathcal{H}(x,\dot{x}) = \frac{1}{2}\dot{x}^T \mathrm{M}\dot{x} + log\left[p(x|z)\right] \tag{2.56}$$

$$\begin{gathered} \frac{dx}{d\tau} = \nabla_{\dot{x}}\mathcal{H}(x,\dot{x}) \quad \frac{d\dot{x}}{d\tau} = -\nabla_x \mathcal{H}(x,\dot{x}) \\ x(0) = x^{(0)} \qquad \dot{x}(0) \sim \mathcal{N}(0,I) \end{gathered} \tag{2.57}$$

Following HMC, these Hamiltonian dynamics are propagated by means of a time-reversible and volume-preserving propagator, usually the fixed-step leapfrog scheme detailed in Equation (2.58).

$$\begin{aligned} \dot{x}^{(\tau+\varepsilon/2)} &= \dot{x}^{(\tau)} - \frac{\varepsilon}{2}\nabla_x \mathcal{H}(x^{(\tau)}, \dot{x}^{(\tau)}) \\ x^{(\tau+\varepsilon)} &= x^{(\tau)} + \varepsilon \nabla_{\dot{x}} \mathcal{H}(x^{(\tau)}, \dot{x}^{(\tau+\varepsilon/2)}) \\ \dot{x}^{(\tau+\varepsilon)} &= \dot{x}^{(\tau+\varepsilon/2)} - \frac{\varepsilon}{2}\nabla_x \mathcal{H}(x^{(\tau+\varepsilon)}, \dot{x}^{(\tau+\varepsilon/2)}) \end{aligned} \tag{2.58}$$

The total propagation time for each candidate sample is $\tau_f = \varepsilon N_\varepsilon$, where N_ε represents the number of steps of the leapfrog scheme. Acceptance or rejection of the propagated particle $x^{(\tau_f)}$ is based on the Metropolis rule, i.e. if the ratio

$$\alpha'_R = \frac{p(x^{(\tau_f)}|z)}{p(x^{(\tau_0)}|z)} \tag{2.59}$$

exceeds some number $u \sim \mathcal{U}(0,1)$, then the sample $x^{(\tau_f)}$ is accepted and will replace $x^{(0)}$ in the following iteration. Though there exist methods to tune the mass matrix M, the number of steps N_ε and step size ε present more difficulties, and in fact they have a great impact on the performance of the algorithm. A small step size can lead to inefficient sampling as each candidate would be very similar to the previous. High ε values trigger a broader exploration of the distribution but can have a detrimental effect on the probability of acceptance. Moreover, selecting a proper N_ε is influenced by the local behavior of the flow field, this is, the gradient of the posterior density

function. Perhaps motivated by the broad applications of HMC methods, Hoffman et al. [67] have recently developed a workaround to defining the number of steps. Their method, dubbed no U-turn sampler (NUTS), suggests propagating the state dynamics until they are found to make a U-turn. Results indicate a huge increase in computational efficiency as the only parameter to be optimized is the step size. An advantage of HMC methods over Metropolis algorithms is that one need not define any jumping distribution, not even an initial one as per the requirements of DREAM. This comes at a higher computational cost since Hamiltonian dynamics need to be simulated for subsequent sampling, and at least one chain simulation should be performed to tune the mass matrix and the step size. Therefore, HMC methods are recommended for applications where sampling from the posterior is an arduous task and derivative information for the likelihood and / or posterior distributions is readily available, preferably in an analytical form.

2.3. Maneuvering Target Tracking Methods

Single target tracking or state space filtering methods work under the assumption of a known transition density. Maneuvering target tracking on the other hand requires inference over the active dynamical mode between sequential observations, resembling to some extent the concept of data association in a multi-target setting. Tracking of maneuvering targets has been a topic of active research for several decades, see for instance [68]–[72]. The different approaches to maneuvering target tracking can be segmented in two broad categories: 1) adjustable level process noise, input estimation (IE) and Variable State Dimension (VSD) filters; and 2) multiple model formulations. The former characterize maneuvers as unknown accelerations applied to the usual non-maneuvering (ballistic or constant velocity) models. These are either approximated as an increase in process noise, an input estimation problem or an augmentation of the state dimension, thus requiring means to detect maneuver onset. Multiple model formulation, usually employed in the context of stochastic hybrid systems, considers the target

a state machine whose dynamics are defined according to one of a finite set of models. This formulation presents significant similarities with JMS, for which there is a vast literature [73]. Filtering methods for JMS aim at both estimating the state and active mode of the system, thus not requiring maneuver detection methods as opposed to adjustable process noise, IE or variable state dimension (VSD) filters. Moreover, the latter find little to no application in clutter or multiple maneuvering target tracking as they are to some extent conditioned by observations, thereby biasing data association.

Besides the characterization of the system dynamics, maneuvering target tracking methods can be classified according to their intended filtering framework, mainly Kalman or SMC based. Typical implementations of these methods require on-line state estimation so their associated filters present little to no recursion: i.e. they are sequential. Implementations considering Kalman filters for maneuvering targets can be consulted in [46, Chapter 11], where adjustable process noise, IE and VSD approaches directly use the Kalman recursion, and multiple-model formulations consist in a bank of Kalman filters embedded in a (pseudo) Bayesian framework. Among the methods therein presented, the IMM [74] formulation seems to be the most capable approach judging by its high acceptance [75]–[79]. Adjustable process noise approaches are deemed to employ PF methods in order to properly capture the details of the propagated prior density, as suggested in [80]. E. Hajiramezanali et al. [80] propose modeling maneuvers as generalized autoregressive conditional heteroskedastic (GARCH) processes in which the conditional variance evolves as a function of past innovations. Thereby, the state of the system is found to have periods of high and low variance that should, in principle, correlate with maneuvering and non-maneuvering intervals, respectively. The latter fact relieves the method from maneuver detection procedures while preventing the definition of multiple dynamical models, i.e. hybrid systems. A similar approach is presented in [81], where the process noise is adjusted whenever the sum of innovations accrued in a given sliding window overcomes some predefined threshold. More interestingly, the method therein presented features an adaptive number of particles, which is directly proportional to the aforementioned sum of innovations. These novel methods arise in juxtaposition

to the various PF approaches to JMS, relying on multiple model formulations for approximating system dynamics during maneuver intervals. C. Andrieu et al. [82] suggest utilizing a particle filtering framework for estimating the state and active mode of a jump Markov system. The latter implementation can be regarded as a first order generalized pseudo-Bayesian estimator (GPBI) applied to each individual particle (see [46, Chapter 11]). Therein, every particle is propagated according to each system mode via the UKF with the aim of assessing the posterior distribution associated to each mode. This way, it is possible to resample the prior distribution in order to maximize the posterior sample diversity. H. Blom and E. Bloem [83] propose a different formulation that extends to the general case of stochastic hybrid systems (not necessarily Markov) for which the transitional boundaries between modes are arbitrary functions of the state and active mode. Their approach does not consider the realization of each mode in particular but rather its probability of occurrence, resulting in a Rao-Blackwellization [84] of the estimation problem in which only the state is estimated via SMC whereas the mode is inferred following an IMM scheme. There are two additional approaches related to IE and multiple model formulations: optimization techniques and heuristics, respectively. Nonetheless, their implementation is significantly dependant on the particular application as the use of surrogate models for optimization or patterns of life for heuristic characterization require an in-depth analysis of the underlying system. These are discussed in the following section, which gathers relevant multiple maneuvering target tracking approaches in the context of space surveillance.

2.3.1. Maneuver Detection and Estimation Methods applied to SSA

Perhaps the simplest way of detecting, and implicitly estimating, a maneuver follows from one of the most used estimation methods in orbit determination: the moving horizon estimator (see Section 2.2.2). As mentioned earlier in the text, in non-linear BLS the model dynamics are linearized at each iteration of the algorithm, effectively resulting in the calculation of state transition matrix (STM) and state sensitivity matrix (SSM). An STM is defined as the mapping of the state from one epoch to another, obtained

via linearization of the dynamics. The effects of a maneuver can then be (linearly) approximated thanks to the STM, and its accuracy will be subject to the magnitude of the maneuver itself. J. Huang et al. [85] present an application of this approach, only considering single along track impulsive maneuvers and Keplerian dynamics. In fact, observations comprise the entire state of the target as they are obtained from radar sensors. The proposed method consist in, given an initial guess for the maneuver magnitude Δv and epoch t_M, obtaining the state transition matrices that map the initial state to t_M and the state at t_M to the last observation epoch. The maneuver Δv can then be applied as a perturbation, resulting in the final state estimate. Maneuver parameters and initial state are updated until a local minimum of the observation residuals is found. Alternative methodologies, initially aimed at uncertainty propagation considering known impulsive maneuvers, can be extended to account for maneuver detection and estimation. A representative example can be found in the procedure developed in [86], which employs state transition tensors for state propagation and Gaussian mixture models (GMMs) for uncertainty characterization. Thus, an iterative nonlinear BLS scheme can be employed in the context of the previous work in order to estimate an expected maneuver sequence. Note, however, that it is subject to a predefined number of maneuvers and an initial guess for the maneuver impulses and epochs needs to be provided. As opposed to [85], the work in [86] is not only generalized to account for higher order tensors, but conveys a proper computation of the transition from the post-maneuver state to the last observation epoch. Recent works have been devoted to the application of these methods to SST [87], [88], wherein single and double burn approaches are used for data association and maneuver characterization, respectively. The perturbed dynamics (IE) approach seems to be the most usual procedure to account for maneuvering targets as it relies on the main methods used for orbit determination. The assumption of a fixed number of maneuvers does not seem to hinder the applicability of the method since two maneuvers are sufficient to transfer between two arbitrary orbits. Nevertheless, the accuracy of the method may be significantly affected by the approximations used to infer the effects of maneuvers in the object dynamics, especially for relatively long propagation times, high maneuver

magnitudes and continuous thrust, since these methods approximate maneuvers as sudden velocity variations.

Typically embedded in sequential filters, multiple model formulations consider additional models for the system dynamics, aimed at capturing the target motion during maneuvering intervals. Most of the multiple model implementations applied to the space environment follow the IMM approach, in part due to a lower presence of particle filter implementations. The latter may stem from the relatively high dimensionality of the system, typically ranging from $\mathbb{R}^6$ when only considering position and velocity to $\mathbb{R}^{10}$ if including force model parameters for atmospheric drag and solar radiation pressure; but also the computational cost of evaluating high fidelity dynamical models[3]. Moreover, the definition of maneuvering modes is complicated by the high non-linearity of orbital motion and the reduced timeliness of observation data. Accordingly, most IMM implementations for maneuver detection of space objects rely on multiple process noise level formulations. In this regard, B. Jia et al. [89] propose an interacting multiple model approach based on CKFs with varying process noise levels. A low noise model is intended to characterize ballistic motion, tuned for consistency between the dynamical model and the observations. An additional Kalman filter assumes a similar dynamical model featuring an increased process noise, in this case tuned to cover an accessibility region for the maneuvering object. The IMM method then is used to marginalize over both models, resulting in an increased tracking performance when compared to a single model formulation. The work in [90] can be considered an extension of [89] but applied to space based observations only using a simplified relative motion model, i.e. the Hill Clohessy Wiltshire (HCW) equations. Nonetheless, extensive testing is performed on the method, considering different measurement update schemes and a wide variety of process noise levels. Q. M. Lam [91] presents a different approach for the definition of a maneuver model. Therein, two models are updated according to an EKF, which requires the computation of linearized dynamics. In

[3]Models with the highest fidelity consider the roto-translational dynamics that stem from the interaction between external forces and the attitude of the spacecraft, resulting in a system with 12 or more dimensions.

the former case, the STM is computed for the position and velocity components of the object, while in the latter, sensitivity to a control acceleration is also introduced. The state estimate of both models is then fed to the IMM to determine the most feasible dynamics. By carefully analyzing these formulations, one may realize they are intended to work in a rather benign environment. Their primary use may be limited to operator surveillance, close formations and, in general, low re-observation times scenarios. Higher revisit times, on the order of several hours, present unfavourable conditions on 1) the level of additional process noise to be considered and 2) the validity of linearized dynamics. Nonetheless, IMM methods are appealing provided some prior knowledge on the maneuvering modes, e.g. orbit raising modes for electric propelled spacecraft and east-west / north-south station keeping modes for GEO operational satellites.

Space objects are typically subject to constant, daily (or monthly) periodic and seasonal perturbations, hence inference based on historical data is expected to deliver reasonable performances at a modest computational effort. The most frequent maneuver sources are due to orbit maintenance, implying that most configurations triggering control actions may well be characterized in terms of orbital regions or in a case-by-case basis using yearlong data. One of the most complete publicly available catalogs, maintained by the 18^{th} Space Defense Squadron and available at space-track.org, reports the state of objects in the TLE format, which condenses one or more observations into a set of initial conditions to be propagated via simplified general perturbations (SGP) model solvers. T. Kelecy et al. [92] suggest analyzing historic TLE data to detect and characterize past maneuvers. This information can then be used within online applications to detect maneuvers in an automated fashion. Their results are promising, yet lacking an appropriate uncertainty characterization that stems from the use of TLE data. Moreover, there seems to be a lag in the maneuver onset time, potentially intrinsic to the method, the use of TLE data, or the higher observability of position with respect to velocity, which delays the observed difference between the expected and actual object state. Siminski et al. [93] elaborate on the use of heuristics for maneuver detection and apply the concept to optical tracklets. Historic maneuver characterization is based on a kernel density

estimator over a reduced set of pre- and post-maneuver orbital elements plus the time of occurrence. The latter is used to infer a typical revisit time, which is expected to coincide for station keeping maneuvers of the same type. This method is then compared against an optimal control based approach, showing a clear advantage in terms of accuracy and efficiency. The simplicity of the kernel density estimator for maneuver characterization has been revised by other authors, who advocate for more advanced machine learning methodologies. In this regard, C. Shabarekh et al. [19] suggest utilizing a combination of an unsupervised and a supervised machine learning algorithm to predict and approximate future maneuvers based on historical data, therein referred to as PoL. Their main claim is being capable of predicting the maneuver epochs with a 3-day confidence interval of 70% seven days in advance. Moreover, this approach can be used to detect deviations from PoL or the usual station keeping strategies, supporting surveillance and collision avoidance related tasks. Heuristics, and in particular the use of historical data, has shown to bring in many benefits to a wide range of applications. Estimation of data sparse systems, as is the case of space object tracking with long re-observation times, can be eased by incorporating prior information regarding the expected behavior of the target. It is obvious that in the absence of previous data, the methodologies presented above are hardly applicable, in which case they might be augmented with some sort of general maneuver model.

Stemming from the belief that operators may be willing to perform fuel-optimal maneuvers, the maneuvering target tracking problem may be posed in the form of an optimal control problem (under the IE class). Subject to a set of constraints, i.e. an initial state estimate, a (yet) uncorrelated observation and a dynamical model, a cost function based on the required control input is minimized to determine the distance from an initial state to an observation in terms of control. The work in [1], [94] considers the use of optimal control methods for determining the association between two different orbits (uncertain two-point boundary-value problem) and an orbit and an UCT (measurement residual boundary-value problem). While both are aimed at determining the minimum control effort required to perform the association, the latter poses an observation as final state constraint.

Holzinger et al. [1] consider (Gaussian) uncertainty in both the initial and final conditions so the association probability is effectively given in the form of a (unimodal) control metric distribution. Means to analyze the resulting distribution are also provided, both to accept or discard the association and to detect whether a maneuver had occurred. Therein, the proposed control distance metric is defined as $d_u = \int_{t_0}^{t_f} \frac{1}{2} u^T u \, d\tau$. No actuator performances are therefore assumed for the control, which is consistent with the intended application as no knowledge regarding the target maneuver characteristics is assumed available. As an extension to this work, N. Singh et al. [95] proposed embedding the optimal control method in an operational framework. Such framework is composed of a MHT used to infer associations without considering maneuvers, a high process noise filter for assessing the possibility of a maneuver and an optimal control method to estimate and characterize the maneuver, if feasible. The optimal control problem is posed between a pair of expected initial and final states with the velocity increment ΔV as cost function, opposed to the energy shape defined in [1]. The solution to the maneuver characterization problem is provided in the form of a probability distribution function for the required control effort ΔV, defined as $\Delta V = \int_{t_0}^{t_f} ||u||_2 \, d\tau$. Various actuator models are considered and analyzed, ranging from multiple impulsive burns to continuous thrust with limited control authority. Significant differences between the various actuator models are found, which can be used as indicative of the propulsive technology equipped by the target given the estimations are sufficiently representative of the ground truth. Finally, acceptance of the maneuver hypothesis is solely based on a threshold, ΔV_{th}, beyond which the possibility of a maneuver is discarded. D.P. Lubey presents a thorough work on this topic in [96]. Therein, an estimation filter is developed based on an optimal control framework. Note the approach is to use optimal control policies to derive 1) mismodeled dynamics and 2) unknown maneuvers. To this end, an exhaustive discussion on various metrics is provided, together with statistical tests to validate maneuver hypotheses. Linear and non-linear versions of the algorithm are developed and tested in LEO and GEO environments, highlighting the limits of applicability of the linearized estimator. The GEO scenario, however, does not consider purely optical measure-

ments as the distance to the target is known and an observation is assumed to extend for two consecutive hours every night. Maneuver detection and characterization via optimal control is a promising solution in the absence of prior information about the target, or when heuristics-based approaches have shown to fail. The latter may well be the case for outlier maneuvers such as collision avoidance, failure conditions or re-positioning. It is important to emphasize on the fact that operators may not always perform fuel optimal maneuvers and even if they did, whenever there is not full knowledge on the post-maneuver state, optimality assumptions inevitably lead to biased estimates. Moreover, the ultimate decision on maneuver acceptance must be defined by an assumed threshold on some control metric, which may significantly vary depending on the orbital regime and / or propulsion system and can hinder data association. There exist additional tractability issues associated to the application of optimal control as convergence of non-linear programming problems cannot be guaranteed for the general (non-convex) case.

3. METHODOLOGY

One of the main pitfalls identified in Section 2 regarding tracking maneuvering space objects is the lack of proper maneuver models, or transition densities for non-ballistic trajectories. A wide range of methods are aimed at solving the multiple maneuvering target tracking problem, but the high non-linearities of space surveillance severely complicate a successful implementation. As identified in Section 2.1, there exist modern capable methods to perform multi-object filtering under uncertainty, specifically tailored to cataloguing activities. In fact, they have been generalized to accommodate hybrid systems, i.e. targets with multiple accessible dynamical modes, though the main difficulty lies in finding a proper definition for such modes. Proposals based on variable process noise levels fail at characterizing moderate control efforts, as these would require unaffordable noise magnitudes; whereas optimal control based solutions cannot provide statistically consistent estimates as they are inherently conditioned on a certain data association decision. One of the main aims of this work is to provide proper maneuver modes for space object tracking that can be readily applied to multi-object filtering. In doing so, the use of efficient control distance metrics (see Section 3.1) is explored, in line with proposals based on optimal control. A dynamics-aware metric can help in developing reasonable transition densities for orbital maneuvers, for instance expressed in terms of the integral acceleration (or net work) required to travel from certain state x^- to a different state x^+. By adapting such dynamics-aware distances to represent current space propulsion technologies, one can derive realistic measures of the control effort required to perform certain orbital transfer. Inversely, technological and operational limits exist on such control efforts, and thus providing upper bounds for orbital transfers can be done with relative ease. Furthermore, based on these control bounds, it is possible to define a set within the state space that is assumed accessible to certain target: the *reachable set*. Section 3.2 deals with the formal definition of such set and introduces different methods to perform approximative

reachability computations. Finally, Section 3.3 provides feasible maneuver mode definitions based on the concept of reachability, which benefit from the computational efficiency of the control metrics developed in Section 3.1. Note these maneuver modes comprehend general design proposals based on the requirements of SST systems and current limitations on computational tractability, yet there is a myriad of possible alternative formulations that may improve performance under particular circumstances or further assumptions.

3.1. Control Distance Metrics

The term *metric* or *distance metric* is mathematically described as a (positive, real-valued, scalar) distance $d(x, y)$ over a set $\{x,\ y\} \in X$, satisfying the following conditions [97]:

- $d(x, y) = 0$ if and only if $x = y$, known as identity axiom;
- $d(x, y) + d(y, z) \geqslant d(x, z)$, the triangle axiom;
- $d(x, y) = d(y, x)$, referred to as the symmetry axiom.

Thus, distance metrics are used to express how far two states x and y are provided certain topological characteristics, which are themselves dictated by the distance function $d(\cdot)$.

Orbital distances in the two-body problem (TBP) have long been expressed in terms of position, velocity, and primarily orbital energy, yet these fail to provide a deep dynamical insight. Some applications, such as SSA, require not only to constrain such distance in terms of a time of flight but also to consider orbital perturbations: luni-solar gravitational forces, solar radiation pressure (SRP) and atmospheric drag, among others. Note that these dynamics cannot be expressed in terms of a Hamiltonian system, since atmospheric drag is a non-conservative force, and therefore the symmetry axiom cannot be fulfilled by definition. Therefore, any time-constrained dynamics-based *metric* that considers drag shall be strictly referred to as *quasi-metric*. However, for the sake of readability, and with

an abuse of notation, the author will refer to such quasi-metrics as distance metrics throughout the remainder of this dissertation. Following the requirements of SSA, Holzinger et al. [98] and later Lubey [96] proposed a control distance metric based on the optimal solution to a boundary value problem, namely the time integral of the required control acceleration (squared) which serves as the cost function itself. The latter has been broadly adopted in the SST community (see Serra et al. [99], [100] and Pirovano et al. [101]) due to its dynamical consistency and solid theoretical foundations. However, one can identify two issues concerning the applicability of such control distance metric: 1) it requires solving a non-linear optimal control problem, which (in general) does not present an analytical solution and whose convergence is sensitive to some initial guess, and 2) the orbital distance is provided in terms of an integral change in velocity (in fact an upper bound as given by the Cauchy-Schwarz inequality) which does not directly relate to the control capabilities of the spacecraft. To illustrate this fact, take a satellite equipped with chemical engines, whose thrust is sufficiently high so as to safely assume impulsive maneuvers. The optimal solution subject to such control performance would correspond to singular solutions as viewed from the former optimal control problem perspective, in which the control (i.e. the perturbing acceleration) is expected to behave smoothly. Moreover, direct applicability to low-thrust transfers is hindered by the mere characteristics of electric propulsion: they are not power (and thus energy) limited but rather constrained by thrust (and hence acceleration). Therefore, a better posed low-thrust metric would be given in terms of the maximum acceleration requirements associated to certain (optimal) orbital transfer. In the following, two approximate and efficient metrics are proposed for each of the main space propulsion types. The impulsive approach takes advantage of local minima in the phasing (or targeting) problem to improve the convergence of a Lambert problem iteration over the epoch of two impulses. The low-thrust metric is based on the thrust Fourier coefficients (TFC) formulation, only retaining the secular effects of the applied thrust profile. Phasing between the initial and final orbits is solved by targeting an intermediate mean motion that ensures the final phase is consistent with the prescribed time of flight, still

considering orbital-averaged dynamics. Note both proposals aim at finding the optimal solution to a surrogate problem, since optimality ensures compliance with the triangle axiom [98]. The general formulation described in Sections 3.1.1-3.1.3 follows simple Keplerian motion, but note this can yield significant estimation errors when applied to realistic scenarios. Section 3.1.4 is thus aimed at providing a series of guidelines and additional techniques to better approximate perturbed two-body motion.

3.1.1. High-Thrust Orbital Transfers

Orbital maneuvers have long been characterized in terms of their impact in the orbital velocity, or incremental velocity change. Former spacecraft were equipped with chemical reaction engines, whose thrust is high enough to produce significant orbital changes in short time scales, barely affecting position. Therefore, the effects of an orbital maneuver can be assumed instantaneous and thus be well characterized by a sudden change in orbital velocity delta-V, $\Delta V = ||v^+ - v^-||$. This approximation is convenient from the maneuver design perspective as it provides an analytical means to infer the effect of a maneuver, this is, one simply needs to determine the required (vectorial) change in velocity that injects the spacecraft into the desired transfer trajectory. Note that precisely due to these dynamical characteristics, one single impulsive maneuver is insufficient to accomplish an arbitrary orbital transfer as both the initial and final trajectories necessarily intersect at the firing point. Thus, a minimum of two impulsive burns are required to perform an arbitrary orbital maneuver, yet an even higher number of impulses may lead to better cost-effectiveness or operational robustness. The main intention is to efficiently characterize any orbital transfer, given a pair of pre- (x_0) and post-maneuver (x_f) states referred to t_0 and t_f respectively. This drives the decision of opting for the least required modelling complexity and hence impulsive transfers will hereafter be approximated by two sequential burns.

Among the most influential theorems in orbital dynamics one can find Lambert's, which states that:

> The transfer time of a body moving between two points and on a conic trajectory is a function only of the sum of the distances of the two points from the origin of the force, the linear distance between the points, and the semi-major axis of the conic section. [102]

Note this only applies to restricted two-body dynamics, and thus no orbital perturbations are considered. The above statement gave rise to the definition of Lambert's problem, which aims at determining the characteristics of such a conic trajectory, ultimately finding the orbital velocities at each point. This two-point boundary value problem (TPBVP) has been a recurrent subject of study throughout the past 200 years [103]–[109] (see [110] for a thorough review of Lambert's problem history), mainly due to its relevance in trajectory design. Certain mission analysis applications do not present stringent time of flight requirements, so that this constraint can be relaxed to derive delta-V optimal transfers. It has been shown that solving the time-free minimum delta-V Lambert problem requires finding the roots of an octic polynomial [111], though via a modification of the cost function (ΔV^2 vs ΔV_{tot}) it is possible to reduce it to a quartic [112]. The latter has been recently revisited [113] as it has applications to station keeping and transfer orbits, yet there is some relevant aspect of the optimal transfer that has not been covered: the global optimal bi-impulsive transfer can (and perhaps should) be posed as an optimization over the location of each impulsive burn within the initial and final orbits, thereby allowing for a lower propellant consumption and thus lower ΔV_{tot} values. There is early research on this regard, as McCue et al. [114] provided a numerical analysis on this topic, yet still not considering it as a full transfer problem in the sense implied by SST, i.e. by ensuring the spacecraft state matches the boundary state conditions at the prescribed epochs. This complicated problem has been explored with application to maneuver detection and reconstruction applied to orbit-to-orbit correlation [115], [116], where there seems to be strong local minima for certain initial and final orbit configurations that follow certain periodicity stemming from the relative phase between the initial and final orbits.

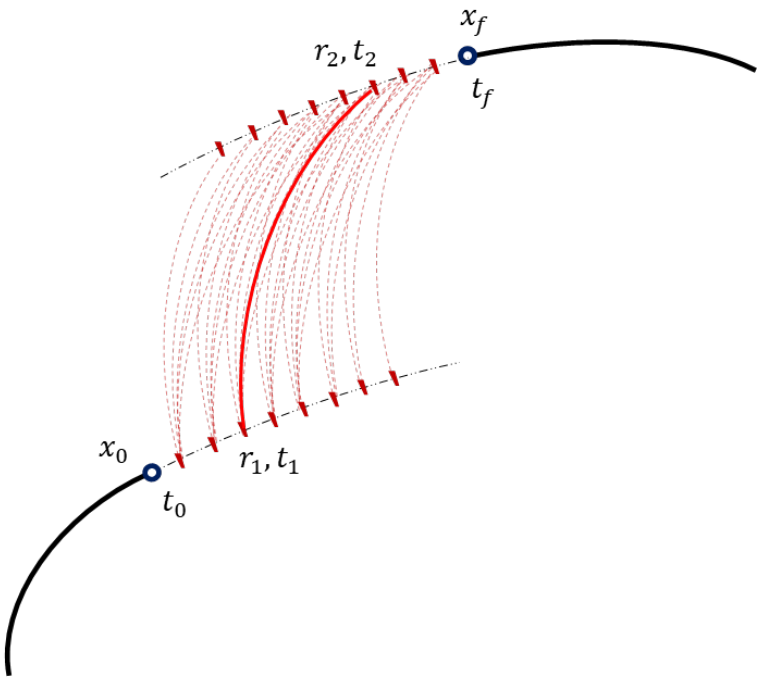

Figure 3.1: Sketch of a double impulse transfer from an initial x_0 to a final orbit x_f provided certain time bounds $[t_0, t_f]$. Boundary conditions for Lambert's problem are given by $[r_1, t_1]$ and $[r_2, t_2]$, where the direction and number of revolutions during the intermediate orbit are those that minimize the total delta-V of the transfer.

Figure 3.1 illustrates an orbital transfer featuring two impulsive maneuvers. Therein, dashed lines emphasize on the fact that there are infinitely many possible trajectories that join the initial and final states in the prescribed time frame, whereas red segments highlight the combinatorial nature of the problem and its temporal bounds, i.e. $t_1 \in [t_0, t_f]$ and $t_2 \in [t_1, t_f]$. The goal is to determine a pair of epochs (where to perform each impulse) that are optimal in some sense. At the time of writing, the cost of injecting 1 kg in orbit is significant (67 M$ for 22,000 kg in LEO[4]), so orbital maneuvers are usually not driven by robustness or safety but rather fuel efficiency. The latter is equivalent to minimizing the required delta-V cost under impulsive (high-thrust) assumptions since, according to Tsiolkovsy's equation,

$$\Delta V = I_{sp} \ln \left(\frac{\mathrm{m}_0}{\mathrm{m}_f} \right). \tag{3.1}$$

Therein, the specific impulse or effective exhaust velocity I_{sp} indicates the efficiency of the propulsive system and the mass ratio $\mathrm{m}_0/\mathrm{m}_f$ refers to the mass fuel expenditure due to the maneuver. The problem of finding the optimal two-impulse transfer can be expressed as

$$\underset{t_1,t_2,N_r,d_r}{\textit{minimize}} \quad ||v_1^+ - v_1^-|| + ||v_2^+ - v_2^-|| \tag{3.2}$$

$$\textit{subject to} \quad [r_1, v_1^-] = f(x_0, t_0, , t_1), \quad [r_2, v_2^+] = f(x_f, t_f, t_2), \tag{3.3}$$

$$[v_1^+, v_2^-] = \textit{Lambert}(r_1, r_2, t_2 - t_1, N_r, d_r), \tag{3.4}$$

where $f(x_i, t_i, t_j)$ represents the dynamical mapping of x_i from t_i to t_j based on Keplerian motion, N_r represents the number of full revolutions elapsed at the intermediate or transfer orbit and d_r indicates the direction of the orbital transfer, i.e. prograde or retrograde. Optimization over the last two parameters is an integral process of applying Lambert's problem solution to trajectory design, and their values are usually close to educated initial guesses, yet as discussed earlier there are specific methods aimed at providing efficient and robust solutions. On the other hand, the dependency of the cost function on the epoch of both impulses is highly non-linear, as

[4]Source: SpaceX Falcon 9 published capabilities and services, updated as of March 2022.

there are characteristics of the orbit that can only be modified at certain orbital positions, e.g. inclination change maneuvers are only effective near the line of nodes. Therefore, it is convenient to express the time of flight in terms of the (mean) angle ΔM covered throughout each leg, i.e.

$$t_f - t_0 = \frac{\Delta M_0}{n_0} + \frac{\Delta M_i}{n_i} + \frac{\Delta M_f}{n_f}, \tag{3.5}$$

where the mean motion can be written in terms of the semi-major axis as $n = \sqrt{\mu_\oplus / a^3}$. The latter provides a convenient scaling since local minima are strongly related to the burn location within each revolution. In order to deal with the complex topology of the problem, one may define the following set of initial conditions

$$t_{1,j} = t_0 + \left\{ 0, \frac{t_f - t_0}{3}, \frac{t_f - t_0}{2}, t_f - t_0 \right\} \pm \left\{ \frac{T_0}{3}, \frac{2T_0}{3} \right\}, \tag{3.6}$$

$$t_{2,j} = t_f - \left\{ 0, \frac{t_f - t_0}{3}, \frac{t_f - t_0}{2}, t_f - t_0 \right\} \pm \left\{ \frac{T_f}{3}, \frac{2T_f}{3} \right\}, \tag{3.7}$$

$$t_f \geqslant t_{2,j} \geqslant t_{1,j} \geqslant t_0, \quad j = 1, ..., N_j, \tag{3.8}$$

where $T = 2\pi/n$ is the orbital period. An example of a multi-revolution transfer in LEO can be consulted in Figure 3.2, wherein the total delta-V cost, $\Delta V = ||v_1^+ - v_1^-|| + ||v_2^+ - v_2^-||$, is given as a function of the burn epochs. Note therein, the initial conditions defined in Eqs. (3.6-3.8) are included. The number of local minima, defined as the number of unconnected regions with $\Delta V \leqslant 1.1\Delta V_{min}$, is on the order of 1,200 for this example, thus posing a challenge for non-linear solvers. The proposed approach consists in the use of a broad set of initial guesses (up to $N_j = 66$) fed to a Gauss-Newton algorithm with a fixed number of maximum iterations. In general, this approach leads to a suboptimal solution as it usually converges to a local minimum. Depending on the required solution accuracy, one can provide a higher number of initial conditions that span different regions of the solution space, but also tune the algorithm to perform a thorough (yet less

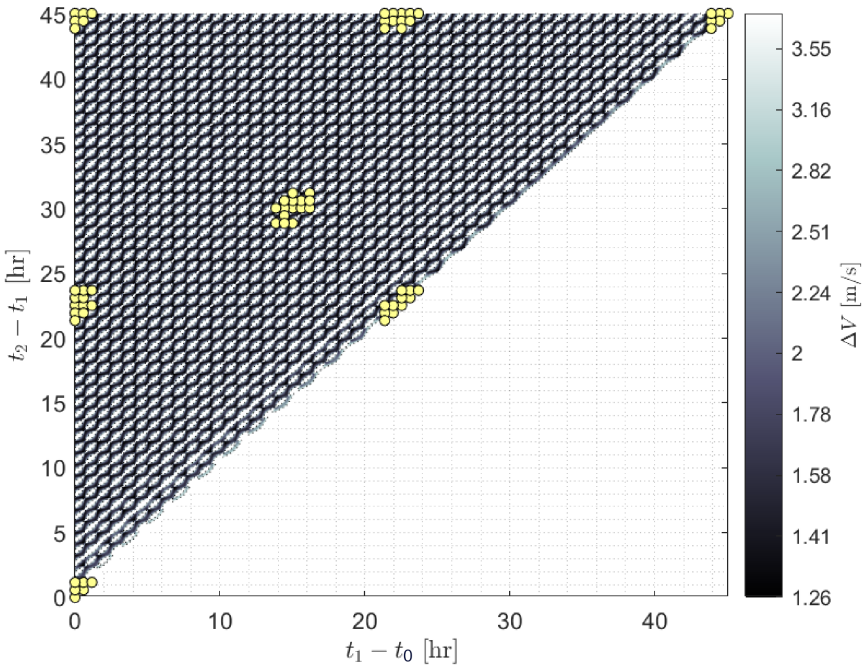

Figure 3.2: Transfer ΔV cost as a function of the first t_1 and second t_2 burn epochs. Yellow dots indicate the proposed set of initial guesses for the Gauss-Newton algorithm.

efficient) search[5]. Accordingly, the impulsive control metric is defined as

$$d_I(x_0, x_f, t_0, t_f) : \underset{j \in N_j}{argmin} \; ||v^+_{1,j} - v^-_{1,j}|| + ||v^+_{2,j} - v^-_{2,j}||, \tag{3.9}$$

where the index j refers to the converged (or iterated) solution of its corresponding initial guess upon the application of Gauss-Newton algorithm.

3.1.2. Linearized High-Thrust Approach

Within the context of this thesis, an alternative approach for characterizing impulsive transfers in terms of control has also been developed. The method presented in Section 3.1.1 is still preferred as the formulation be-

[5]Results included later in Section 4.1.1 were obtained after performing such tuning, based on a battery of randomized test cases.

low presents difficulties when accounting for phase changes, though it is better suited for other type of computations (see Section 3.2.1). A convenient way of expressing the acceleration-perturbed two-body dynamics is given by Gauss planetary equations, which take the following form when expressed in modified equinoctial elements (MEE)

$$\frac{dp}{dt} = \frac{2p}{w}\sqrt{\frac{p}{\mu_\oplus}}\mathrm{a}_\theta, \tag{3.10}$$

$$\frac{dh_1}{dt} = \sqrt{\frac{p}{\mu_\oplus}}\left[\sin(L)\mathrm{a}_r + ((w+1)\cos(L)+h_1)\frac{\mathrm{a}_\theta}{w} - (h_2\sin(L)-k_2\cos(L))\frac{k_1}{w}\mathrm{a}_h\right], \tag{3.11}$$

$$\frac{dk_1}{dt} = \sqrt{\frac{p}{\mu_\oplus}}\left[-\cos(L)\mathrm{a}_r + ((w+1)\sin(L)+k_1)\frac{\mathrm{a}_\theta}{w} + (h_2\sin(L)-k_2\cos(L))\frac{h_1}{w}\mathrm{a}_h\right], \tag{3.12}$$

$$\frac{dh_2}{dt} = \sqrt{\frac{p}{\mu_\oplus}}\frac{s^2\cos(L)}{2w}\mathrm{a}_h, \tag{3.13}$$

$$\frac{dk_2}{dt} = \sqrt{\frac{p}{\mu_\oplus}}\frac{s^2\sin(L)}{2w}\mathrm{a}_h, \tag{3.14}$$

$$\frac{dL}{dt} = \sqrt{\mu_\oplus p}\left(\frac{w}{p}\right)^2 + \frac{1}{w}\sqrt{\frac{p}{\mu_\oplus}}(h_2\sin(L)-k_2\cos(L))\mathrm{a}_h, \tag{3.15}$$

where $s^2 = 1+h_2^2+k_2^2$, $w = 1+h_1\cos(L)+k_1\sin(L)$ and $\mathrm{a}_r, \mathrm{a}_\theta, \mathrm{a}_h$ represent the radial, azimuthal and normal Cartesian components of a perturbing acceleration. Under the assumption of impulsive maneuvers, $\{\mathrm{a}_r, \mathrm{a}_\theta, \mathrm{a}_h\} \propto \delta(t-t_M)$ with t_M the maneuver epoch, and provided Equations (3.10-3.15) behave smoothly during the thrust arc, a somewhat accurate linear approximation for the effect of an impulsive burn on the set of orbital elements is given by

$$\begin{bmatrix}\Delta p\\ \Delta h_1\\ \Delta k_1\\ \Delta h_2\\ \Delta k_2\\ \Delta L\end{bmatrix} \approx \begin{bmatrix} 0 & \frac{2p}{w}\sqrt{\frac{p}{\mu_\oplus}} & 0\\ \sqrt{\frac{p}{\mu_\oplus}}\sin(L) & \sqrt{\frac{p}{\mu_\oplus}}[(w+1)\cos(L)+h_1]\frac{1}{w} & -(h_2\sin(L)-k_2\cos(L))\frac{k_1}{w}\\ -\sqrt{\frac{p}{\mu_\oplus}}\cos(L) & \sqrt{\frac{p}{\mu_\oplus}}[(w+1)\sin(L)+k_1]\frac{1}{w} & (h_2\sin(L)-k_2\cos(L))\frac{h_1}{w}\\ 0 & 0 & \sqrt{\frac{p}{\mu_\oplus}}\frac{s^2\cos(L)}{2w}\\ 0 & 0 & \sqrt{\frac{p}{\mu_\oplus}}\frac{s^2\sin(L)}{2w}\\ 0 & 0 & \frac{1}{w}\sqrt{\frac{p}{\mu_\oplus}}(h_2\sin(L)-k_2\cos(L)) \end{bmatrix}\begin{bmatrix}\Delta v_r\\ \Delta v_\theta\\ \Delta v_h\end{bmatrix}. \tag{3.16}$$

The above will hereafter be written in the compact form

$$\Delta\text{œ} \approx A(\text{œ}, L)\Delta v. \tag{3.17}$$

where œ is the set of orbital parameters that define an orbit. Hereunder, the pertubing acceleration is assumed to be negligible except for thrust arcs, and hence only the *running variable L*, representing the angular location of the object within its orbit, shows a dependency on time. Following the approximation in Equation (3.17), one can express the change in orbital elements caused by a sequence of n_M burns as

$$\Delta\text{œ} \approx \sum_{j=1}^{n_M} A(\text{œ}_j, L_j)\Delta v_j. \tag{3.18}$$

Note that time dependency is implicit in the true longitude L_j of each impulse. The aim is to solve the inverse problem, this is, determine the set of maneuvers that results in a prescribed change in orbital elements. According to the dimensionality of the system, a minimum of two impulses is required to acquire an arbitrary orbit, so one can further approximate a bi-impulsive transfer as

$$\Delta\text{œ}_e = \begin{bmatrix} A(\text{œ}_0, L_1) & A(\text{œ}_f, L_2) \end{bmatrix} \begin{bmatrix} \Delta v_1 \\ \Delta v_2 \end{bmatrix}. \tag{3.19}$$

Therein, the sensitivity matrix for the first burn is computed with respect to the initial orbit, whereas that of the second impulse is referred to the final state. Conveniently, the problem reduces to finding a pair of true longitudes $[L_1, L_2]$ such that the augmented matrix is full rank and the cost function

$$J = \Delta v_{tot}^T \Delta v_{tot} + c_1(\Delta\text{œ}_t - \Delta\text{œ}_e)^T(\Delta\text{œ}_t - \Delta\text{œ}_e), \tag{3.20}$$

is minimized. Therein, the burn couple is rearranged as $\Delta v_{tot} = \begin{bmatrix}\Delta v_1, \Delta v_2\end{bmatrix}^T$ and c_1 is a relative cost index that captures the error committed when solving the linear system, i.e. target $\Delta\text{œ}_t$ vs estimated $\Delta\text{œ}_e$ change. The goal is thus to find the bi-impulsive maneuver that minimizes the above cost function, which reduces to

$$\Delta v_{tot}^* = \left(2\left(\mathbb{I} + c_1 A'^T A'\right)\right)^{-1} \cdot 2c_1 A'^T \Delta\text{œ}_t, \tag{3.21}$$

being $\mathbb{I}$ the identity matrix and A' the augmented sensitivity matrix in Equation (3.19), i.e.

$$A' = \begin{bmatrix} A(\text{œ}_0, L_1) & A(\text{œ}_f, L_2) \end{bmatrix}. \tag{3.22}$$

Note this is an extremely convenient way of approximating the bi-impulsive transfer problem as one need only find the optimal pair of true longitudes (L_1, L_2) that solve the following constrained optimization problem

$$\begin{aligned} &\text{Minimize:} \quad J = \Delta v_{tot}^T \Delta v_{tot} + c_1 (\Delta\text{œ}_t - \Delta\text{œ}_e)^T (\Delta\text{œ}_t - \Delta\text{œ}_e) \\ &\text{subject to:} \quad L_0 \leqslant L_1 \leqslant L_2 \leqslant L_f. \end{aligned} \tag{3.23}$$

Relatively simple methods can be applied to solve this, such as a gradient descent seeded with multiple initial guesses, similar to the one proposed in Section 3.1.1.

Accordingly, an alternative impulsive control metric can be defined as

$$d_I'(x_0, x_f, t_0, t_f) = ||\Delta v_1^*||_2 + ||\Delta v_2^*||_2, \tag{3.24}$$

where $[\Delta v_1^*, \Delta v_2^*]$ are the first and second burns that stem from evaluating Equation (3.21) at the optimum set of true longitudes (L_1^*, L_2^*). A final caveat on the application of this method is that directly targeting phasing via the sensitivity matrix is highly inefficient, as this only considers shifts in true longitude due to plane changes. However, this does not impose limitations on the intended usage of the above metric, as discussed later in Section 3.2.1. In fact, it is still possible to perform efficient phasing maneuvers by introducing an initial change in semi-major axis compliant with the overall shift in anomaly, thereby replacing the last row in Equation (3.16) by an additional change in semi-latus rectum that is only affected by Δv_1.

3.1.3. Low-Thrust Orbital Transfers

Electric propulsion has matured over the last decades to the point of being the preferred solution for constellations in LEO [117]. Plasma thrusters, either gridded ion or Hall effect, can deliver high total impulses with a limited

amount of onboard mass, rendering them a cost effective choice for orbit acquisition and station keeping under certain circumstances. As opposed to chemical thrusters, these type of engines are not energy limited, but rather power limited, so that the amount of thrusting force they are able to produce is highly dictated by the electrical power available to them. Thus, control capabilities are subject to the *normal operation mode* requirements, sufficiently high to accommodate for this type of engines in the case of commercial telecommunication buses.

There is a vast literature regarding optimal low-thrust orbital transfers. Of particular interest is to develop surrogate models that provide analytical, yet sub-optimal, solutions, as these can be valuable starting points for non-linear optimal control problem solvers. A compelling approach is to employ acceleration profiles compliant with first integrals of motion, as described in Petropoulos et al. [118] and Roa et al. [119]. These only consider in-plane dynamics and, though there exist a formulation for non co-planar transfers [120], it employs a decoupled out-of-plane model and the resulting control profile can be far from optimal in certain cases. Hudson and Scheeres [121] propose an alternative characterization of the thrust profile that can be used to solve a general orbital transfer. By expanding any type of acceleration as a Fourier series in the eccentric anomaly, i.e.

$$\mathrm{a}_r = \sum_{k=0}^{\infty} \left[\alpha_k^{r,E} \cos(kE) + \beta_k^{r,E} \sin(kE) \right], \tag{3.25}$$

$$\mathrm{a}_\theta = \sum_{k=0}^{\infty} \left[\alpha_k^{\theta,E} \cos(kE) + \beta_k^{\theta,E} \sin(kE) \right], \tag{3.26}$$

$$\mathrm{a}_h = \sum_{k=0}^{\infty} \left[\alpha_k^{h,E} \cos(kE) + \beta_k^{h,E} \sin(kE) \right], \tag{3.27}$$

along the radial, azimuthal, and normal directions, one can accurately approximate control profiles with periodic components up to one orbital revolution. In practice, thrust-limited transfers require multiple revolutions and optimal solutions to the primal (unperturbed) problem require either constant or orbital periodic thrust profiles (see Kluever and Oleson [122]),

providing theoretical support to this parametrization.

To make the problem more tractable without the need to explicitly truncate the infinite series, Hudson [123] propose to average the *control perturbed* two-body dynamics over the mean anomaly M, effectively resulting in the following set of *essential* coefficients

$$f_T = \Big[\alpha_0^{r,E} \quad \alpha_1^{r,E} \quad \alpha_2^{r,E} \quad \beta_1^{r,E} \quad \alpha_0^{\theta,E} \quad \alpha_1^{\theta,E} \quad \alpha_2^{\theta,E} \quad \beta_1^{\theta,E} \quad \beta_2^{\theta,E} \quad \alpha_0^{h,E} \quad \alpha_1^{h,E} \quad \alpha_2^{h,E} \quad \beta_1^{h,E} \quad \beta_2^{h,E}\Big]^T. \tag{3.28}$$

In fact, this approximation is most appealing due to the linear dependency of the dynamics $d\bar{\mathfrak{e}}/dt = G(\bar{\mathfrak{e}})f_T$ on the reduced set of TFC, which take the form of Equations (2.66-2.72) in [123] for a set of non-singular (equinoctial-like) elements. As an additional correction, one should account for the initial conditions of the transfer, which present a set of averaged elements that depart from the prior orbit due to the action of the control profile (refer to Section 2.5 in [123] for a deeper explanation). Thus, the secular orbital elements evolution can be expressed as

$$\bar{\mathfrak{e}}_f = \bar{\mathfrak{e}}_0 + \int_{t_0}^{t_f} G(\bar{\mathfrak{e}})f_T dt + H(\bar{\mathfrak{e}})f_T. \tag{3.29}$$

Note the intention is to find the control required to perform certain orbital transfer so $\bar{\mathfrak{e}}_0$, $\bar{\mathfrak{e}}_f$, as well as t_0, t_f are inputs to the problem. However, this problem is closed if one assumes a shape for the transfer in terms of the orbital elements, following a smooth evolution from the initial to the final state. Though general transfers might adopt a sequencing strategy for the change in orbital elements, e.g. first increase the orbital energy and then change the inclination, as a first approximation one can presume a linear behavior, so that

$$\bar{\mathfrak{e}}(t) = \bar{\mathfrak{e}}_0 + \frac{\bar{\mathfrak{e}}_f - \bar{\mathfrak{e}}_0}{t_f - t_0}(t - t_0), \tag{3.30}$$

thus allowing to solve the integral in Equation (3.29) via a simple colloca-

tion method, e.g.

$$\hat{G} = \int_{t_0}^{t_f} G(\bar{\mathfrak{e}})dt \approx \sum_{i=0}^{N_i} G(\bar{\mathfrak{e}}, t_i)(t_i - t_{i-1}), \tag{3.31}$$

where $t_i = t_0 + i(t_f - t_0)/N_i$. Conveniently, this leads the following linear system relating the change in orbital elements with the set of TFCs

$$\Delta\bar{\mathfrak{e}} = \left[\hat{G}(\bar{\mathfrak{e}})\Delta t + H(\bar{\mathfrak{e}})\right] f_T, \tag{3.32}$$

and thus the minimum norm solution of this linear system, hereafter denoted as f_T^*, can be used as a proxy for the control law that minimizes the thrust magnitude. The latter can be obtained by finding the minimum of the cost function

$$\begin{aligned} J = \Big\{ &\left[\alpha_0^{r,E} + \alpha_1^{r,E}\cos(E) + \beta_1^{r,E}\sin(E) + \alpha_2^{r,E}\cos(2E)\right]^2 + \\ &\left[\alpha_0^{\theta,E} + \alpha_1^{\theta,E}\cos(E) + \beta_1^{\theta,E}\sin(E) + \alpha_2^{\theta,E}\cos(2E) + \beta_2^{\theta,E}\sin(2E)\right]^2 + \\ &\left[\alpha_0^{h,E} + \alpha_1^{h,E}\cos(E) + \beta_1^{h,E}\sin(E) + \alpha_2^{h,E}\cos(2E) + \beta_2^{h,E}\sin(2E)\right]^2 \Big\}^{\frac{1}{2}} \end{aligned} \tag{3.33}$$

with respect to the eccentric anomaly, which means obtaining the E^* satisfying

$$\left.\frac{\partial J}{\partial E}\right|_{E^*} = 0, \qquad \left.\frac{\partial^2 J}{\partial E^2}\right|_{E^*} < 0. \tag{3.34}$$

Finally, the proposed thrust-limited control metric is defined as the maximum instantaneous acceleration compliant with the optimal set of TFC, this is

$$d_{LT}(x_0, x_f, t_0, t_f) = J(f_T^*, E^*). \tag{3.35}$$

A note on efficient phasing

There is a caveat with regard to the use of the above formulation: the running variable ϵ_1 used in [121], [123] represents the additional change in mean anomaly due to the control perturbation when compared to the average mean motion. Therefore, its sensibility to the TFC is rather limited

and any effort towards performing phasing maneuvers via targeting this variable is highly inefficient. It is thus reasonable to use a reduced set of (non-singular) orbital elements $\text{œ} = [a, h_1, k_1, h_2, k_2]$ (c.f. [123]) that drops the incremental anomaly variation ϵ_1. In practice, phasing maneuvers are accomplished by adjusting the relative or synodic period to reach the target phase at the desired epoch. Therefore, one can divide the transfer in two stages, targeting an intermediate orbital period that ensures the overall phase shift is compliant with the final conditions.

According to the definition of mean motion n, the change in mean anomaly can be expressed as

$$M_f = M_0 + \int_{t_0}^{t_f} n dt. \tag{3.36}$$

The goal then is to derive an optimal law for $n(t)$ subject to the boundary conditions imposed by n_0 and n_f that is compliant with the required phasing $\Delta M = M_f - M_0$. An obvious choice would be to define a cost function of the form

$$J'_n = \frac{1}{2} \int_{t_0}^{t_f} \left(\frac{dn}{dt} \right)^2 dt, \tag{3.37}$$

thus minimizing the integral change in mean motion. This boundary value problem has in fact an analytical solution with a quadratic form for $n(t)$. However, and again due to the assumption of thrust limited control, a better definition for the cost function would be

$$J_n = \underset{t}{max} \left(\frac{dn}{dt} \right), \tag{3.38}$$

as it directly relates to the maximum instantaneous acceleration. Note that a quadratic law for $n(t)$ would not be optimal with respect to the new cost function and, in fact, it turns out that the optimal law in this case corresponds to a bang-bang solution of the type

$$\frac{dn}{dt} = \begin{cases} \dot{n} & \text{if } t_0 \leqslant t \leqslant t_1, \\ -\dot{n} & \text{if } t_1 < t \leqslant t_f. \end{cases} \tag{3.39}$$

Conveniently, such a linear law for $n(t)$ is compliant with the assumption in Equation (3.30) and simply requires to solve two separate transfers targeting an intermediate mean motion $n_1 = n_0 + \dot{n}(t_1 - t_0)$. Accordingly, one can re-define the low-thrust control metric in Equation (3.40) to account for the phasing between the initial and final orbits, which can be expressed as

$$d_{LT}(x_0, x_f, t_0, t_f) = min\left[d_{LT}(x_0, x_1, t_0, t_1), d_{LT}(x_1, x_f, t_1, t_f)\right]. \tag{3.40}$$

Note in the above, following the assumption of a smooth transition from the initial to the final state, the intermediate orbital state is approximated as

$$\text{œ}_1 = \text{œ}_0 + \frac{t_1 - t_0}{t_f - t_0}(\text{œ}_f - \text{œ}_0), \tag{3.41}$$

modified to account for the desired intermediate mean motion, i.e. $a_1 = \sqrt[3]{\mu/n_1^2}$.

For the two-phase control policy to be completely characterized, one needs to find 1) the optimal intermediate time $t_1 \in [t_0, t_f]$ and 2) the required phase change ΔM. In fact, the latter can be attained after an arbitrary number of revolutions, i.e. $\Delta M_k = \Delta M + 2\pi k$ for $k \in \mathbb{Z}$ are all valid solutions to the problem. In general, it is sufficient to retain the value of k yielding n_1 that is closest to n_0 or n_f. With respect to the policy change epoch t_1, Figure 3.3 illustrates its effect in the control metric d_{LT} for two different scenarios, both with a strong in-plane phasing component. The LEO case (on the right) shows an asymmetric behavior, attributed to the secular effects of atmospheric drag and, to a lower extent, Earth's oblateness. In both cases, a linearized version of SGP4 is used as dynamical model (see Section 3.1.4). These effects are much less noticeable in the geosynchronous scenario, which is almost purely symmetrical. For Keplerian dynamics, the minimum of J_n has an analytical solution that corresponds to

$$t_1 = \frac{t_0 + t_f}{2}. \tag{3.42}$$

Orbital perturbations shift this minimum to either direction depending on the relative geometry of the boundary conditions, but in general it remains

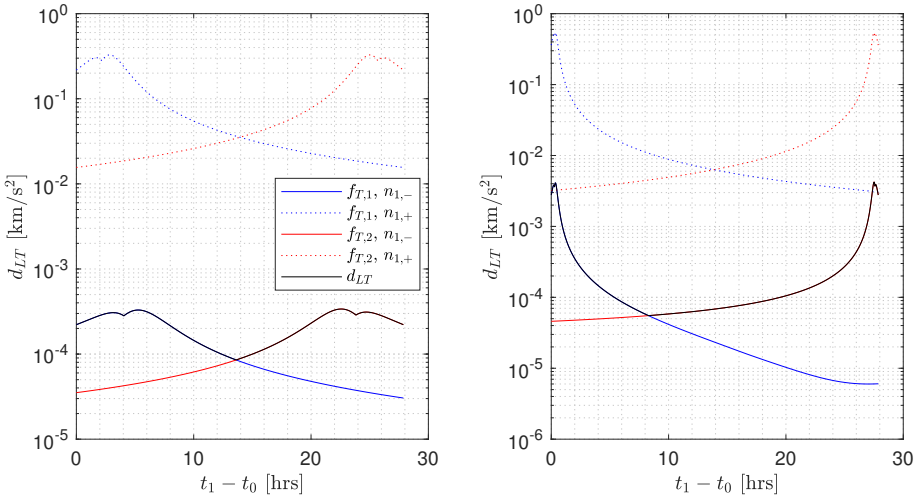

Figure 3.3: Low-thrust metric d_{LT} as a function of intermediate (policy change) time t_1 for a phasing maneuver in GEO (left) and LEO (right). $n_{1,+}$ and $n_{1,-}$ represent the two closest solutions of the intermediate mean motion n_1 to the average transfer value $(n_0+n_f)/2$. The blue color indicates the maximum acceleration for the first segment ($t_0 \leqslant t \leqslant t_1$), whereas the red line is associated to the second segment ($t_1 < t \leqslant t_f$).

in the vicinity of the above solution and features the value of n_1 that is closest to the average mean motion$(n_0 + n_f)/2$. Similar to the approach followed in Section 3.1.1, one can use a Gauss-Newton iteration to refine the estimate of t_1.

3.1.4. Implications of perturbed two-body dynamics

The control metric formulations discussed in previous sections rely on an unperturbed two-body problem as dynamical model, which may incur significant approximation errors, e.g. for long times of flight in LEO. As formulated in previous sections, the initial state x_0 is referred to the initial time t_0, though due to SST system requirements, a high fidelity estimate can be assumed available at t_f, hereafter referred to as x_0^+. Provided the post maneuver state x_f is sufficiently close to said ballistic estimate, one can expect the perturbative forces to be of similar magnitude. Therefore,

back-propagating both states following the surrogate dynamics assumed in Sections 3.1.1 and 3.1.3 would lead to control requirements that resemble those of a full-fidelity simulation, in spite of the differences in the estimates themselves with respect to those obtained via high-fidelity modeling.

However, the above does not circumvent the main concern regarding perturbations: orbital elements, as a convenient parametrization of the state, no longer remain time invariant but show long and short term periodic oscillations as well as secular drifts. Thus, it is convenient to derive analytical transformations that smooth these effects, for instance following orbital averaging methods. In this regard, Brouwer [124] and Kozai [125] devised analytical methods to account for secular and short-periodic effects caused by (low-order) spherical harmonics (see Draper semi-analytical satellite theory (DSST) [126] for an alternative semi-analytical approach). In fact, Brouwer's method was extended to include an exponential atmospheric density model, which served as baseline for the semi-analytical SGP4 model that is still used nowadays to propagate TLEs [127]. The effect of secular terms can be mitigated by referring both the pre- and post-maneuver orbits to a common epoch, as discussed in the previous paragraph, yet short periodic oscillations can still cause inconsistencies. In lower orbital regimes, they present significant amplitudes ($O(10\,km)$ in semi-major axis) and one possible solution to limit their impact is to *regularize* the orbital space adopting orbital averaging techniques. Throughout this work, the SGP4 state representation will be used for all computations concerning the control distance metrics discussed in previous sections, thus transforming osculating states to TLE variables. The selection of SGP4 over other methods such as DSST is mainly driven by computational efficiency, yet a key factor to take into account is the availability of an open source MATLAB implementation. Though the developed impulsive metrics internally rely on Keplerian dynamics (efficient solvers for the perturbed Lambert problem are seldom found in the literature), one can improve the fidelity of the low-thrust transfer with a modest computational effort. These kind of transfers are assumed to be smooth, so the approximation in Equation (3.31) can be modified to account for the average state evolution as given by semi-

analytical models, and thus Equation (3.30) now becomes

$$\bar{\mathfrak{E}}(t) = \frac{\bar{\mathfrak{E}}_0(t) \cdot (t_f - t) + \bar{\mathfrak{E}}_f(t) \cdot (t - t_0)}{t_f - t_0}. \tag{3.43}$$

3.2. Reachability in the two-body problem

The concept of reachability is not novel to SSA as there are numerous texts devoted to its formal definition [128], computation [129], [130] and applications [131], [132]. In the former, the authors elaborate on the mathematical description of the reachable set as a particularization of an optimal control problem. Therein, strong theoretical background is given on the concept of reachability and the authors derive analytical solutions to linearized (circular) relative dynamics. Special emphasis is given to the problem of *ellipsoidal* reachability, dealing with the temporal evolution of an initial ellipsoidal region in the state space. In a follow-up paper, Holzinger at al. [129] studied reachability applied to impulsive maneuvers, where the optimal control problem is reformulated setting the ΔV cost as independent parameter. The resulting time-free optimal control policies derived from this approach are aligned with common literature in orbital maneuvering, serving as proofs for the optimality of Hohmann and bi-elliptic transfers. However, the proposed method is rather costly as it involves solving a partial differential equation (PDE) and, due to the time-free assumption, does not consider reachability in terms of the direction of the eccentricity and inclination vectors. Another relevant contribution, in this case targeting low-thrust spacecraft, can be consulted in [130]. Therein, the authors propose an indirect multi-stage formulation for the optimal control problem, ultimately leading to a step-wise linearized model. This allows for computationally efficient reachability calculations that depend upon the control discretization, this is, the number of thrust segments or stages. Their proposal relies on Cartesian coordinates, arguably due to its focus on cislunar applications, and they identify convergence and accuracy issues for concave reachable sets and multi-revolution transfers. Aguilar et al. [131] make use of reachability approximations to maintain custody of maneuvering space

objects via sensor (telescope) tasking, under the assumption of low-thrust control. Similarly, Hall et al. [132] developed a method to regain custody of maneuvering targets based on control assumptions, where the search space is restricted to lie within the set of reachable states. Therein, an approximation of the reachable set is sought via the higher order sensitivity matrix (HOSM) method, based on polynomial approximations of the (non-linear) dynamics. Note in this case the definition of the reachable set is stochastic, rather than deterministic, and is linked to the uncertainty inherent to the maneuver execution process. Accordingly, the performance of the method is severely affected by prior knowledge regarding the PDF of the control input: vacuous maneuver priors may be poorly characterized by the proposed cubature rules for control sensitive systems since certain levels of optimality are expected. The latter is to say that reachable bounds are strictly defined by optimal control laws, which fall in regions of very low probability density for vacuous control input PDFs. Under these circumstances it is more convenient to directly characterize transfers in the state space, rather than the control sequence, assigning certain control cost to a given pair of pre- and post-maneuver states. Therefore, the remainder of this section discusses an alternative (deterministic) formulation that builds upon the control metrics defined Section 3.1.

Irrespective of the control measure $d(\cdot)$ used, one can express the reachable set, i.e. the set assumed accessible to a target conditioned on its control capabilities, as

$$\mathcal{R}(x; x_0, t_0, t_f) = \{x : d(x_0, x, t_0, t_f) \leqslant \lceil d \rceil\}, \tag{3.44}$$

where $\lceil d \rceil$ represents the maximum control performance. Provided the set $\mathcal{R}(\cdot)$ is connected, which can be shown to hold for the reachable frontier [128], there exist different approximative approaches. This is because reachability computations are complex tasks, considered intractable in most cases. Note for general non-linear systems the reachable set may well be non-convex and its mere representation can be a problem on its own right. Approximations are thus deemed necessary to find and express $\mathcal{R}(\cdot)$: in general one is eager to reduce the complexity of the dynamics or find a

canonical transformation (for the state space) wherein quasi-linear assumptions hold to a certain degree. Tractability also spans its representation, and the most usual approach is to make use of a bounding box or a zonotope [133]. The following sections elaborate on efficient methods for reachability computations in the two-body problem. Section 3.2.1 describes an axis-aligned minimum bounding box (AABB) solution to a linear metric-bounded system, which stems from the definition of the control metrics in Sections 3.1.2 and 3.1.3. An efficient search algorithm is proposed in Section 3.2.2, which can be applied to alternative state space representations and, in general, applications where set containment evaluation is relatively inexpensive. Finally, Section 3.2.3 is devoted to volumetric set computations, emphasizing on non-linear state transformations and high-dimensional systems.

3.2.1. Linear Reachability

Reachability computations are inexpensive tasks in linear problems as the principal directions are invariant. To illustrate this concept, consider a linear system of the form

$$Ax = b, \tag{3.45}$$

and suppose the goal is to maximize b_j subject to the condition $||x||_2 \leqslant c$. The solution to this linear optimization problem is given by the direction of maximum variation for b_j, multiplied by the maximum norm of x, i.e.

$$x^* = \frac{\nabla_x(b_j)}{||\nabla_x(b_j)||_2}c \equiv \frac{A^T_{j,:}}{||A_{j,:}||_2}c, \tag{3.46}$$

where $A_{j,:}$ indicates the j^{th} row of matrix A. In the following, this extremely simple principle is applied to the linear impulsive and low-thrust surrogate models developed in Sections 3.1.2 and 3.1.3, thereby exposing simple and elegant methods to approximate the reachable bounds for chemical and electric propulsion spacecraft, respectively.

Linear impulsive metric

According to the definition of the linear impulsive metric $d_I'(x_0, x_f, t_0, t_f)$ in Section 3.1.2, it is possible to approximate the effect of a two-impulse maneuver in the space of MEEs as

$$\Delta\text{œ} \approx A(\text{œ}_0, L_1)\Delta v_1 + A(\text{œ}_f, L_2)\Delta v_2. \tag{3.47}$$

Therefore, by extrapolation of the linear reachability concepts discussed previously, one can express the maximum variation along the j^{th} state dimension as

$$\Delta\text{œ}_{j,max} = A_{j,:}(\text{œ}_0, L_1)\frac{A_{j,:}^T(\text{œ}_0, L_1)}{||A_{j,:}(\text{œ}_0, L_1)||_2}c_1 + A_{j,:}(\text{œ}_f, L_2)\frac{A_{j,:}^T(\text{œ}_f, L_2)}{||A_{j,:}(\text{œ}_f, L_2)||_2}c_2, \tag{3.48}$$

where c_1 and c_2 represent the relative control exerted at the first and second burn, respectively, and satisfy $c_1 + c_2 = \lceil d_I' \rceil$, i.e. the total control effort coincides with the maximum allowable. Note again, the above results in a linear reachability problem, this time in terms of the relative control share via c_1 and c_2, which now take the values

$$c_1 = \frac{A_{j,:}(\text{œ}_0, L_1)\dfrac{A_{j,:}^T(\text{œ}_0, L_1)}{||A_{j,:}(\text{œ}_0, L_1)||_2}}{A_{j,:}(\text{œ}_0, L_1)\dfrac{A_{j,:}^T(\text{œ}_0, L_1)}{||A_{j,:}(\text{œ}_0, L_1)||_2} + A_{j,:}(\text{œ}_f, L_2)\dfrac{A_{j,:}^T(\text{œ}_f, L_2)}{||A_{j,:}(\text{œ}_f, L_2)||_2}}\lceil d_I' \rceil \tag{3.49}$$

and

$$c_2 = \frac{A_{j,:}(\text{œ}_f, L_2)\dfrac{A_{j,:}^T(\text{œ}_f, L_2)}{||A_{j,:}(\text{œ}_f, L_2)||_2}}{A_{j,:}(\text{œ}_0, L_1)\dfrac{A_{j,:}^T(\text{œ}_0, L_1)}{||A_{j,:}(\text{œ}_0, L_1)||_2} + A_{j,:}(\text{œ}_f, L_2)\dfrac{A_{j,:}^T(\text{œ}_f, L_2)}{||A_{j,:}(\text{œ}_f, L_2)||_2}}\lceil d_I' \rceil. \tag{3.50}$$

Recall that there is no knowledge on the final state after the maneuver, œ_f, which might seem obvious as the goal is to derive it. Therefore, one shall

first solve the single burn reachability problem to find

$$\Delta v_1' = \frac{A_{j,:}^T(\text{œ}_0, L_1)}{||A_{j,:}(\text{œ}_0, L_1)||_2}\lceil d_I' \rceil, \tag{3.51}$$

from where to obtain an initial guess for the final state $\text{œ}_f' = \text{œ}_0 + \Delta\text{œ}'$, with

$$\Delta\text{œ}' = A(\text{œ}_0, L_1)\Delta v_1'. \tag{3.52}$$

Several iterations can be performed over Equation (3.49) to better approximate œ_f, though convergence to acceptable tolerances is typically reached after one or two iterations. Note throughout this derivation, references to L_1 and L_2 have been kept. That is because the problem is highly sensitive to the epoch (or true longitude) of each maneuver, and thus it is customary to find the optimal L_1^* and L_2^*. Recalling each $\Delta v' \propto A_{j,:}^T(\text{œ}, L)/||A_{j,:}(\text{œ}, L)||_2$, the above reduces to finding

$$\underset{L}{\text{argmax}} \;\; A_{j,:}(\text{œ}, L)A_{j,:}^T(\text{œ}, L). \tag{3.53}$$

Matrix $A(\cdot)$ was earlier defined in Equation (3.16), from which it is possible to obtain analytical expressions for the first and second order derivatives of $A_{j,:}(\text{œ}, L)A_{j,:}^T(\text{œ}, L)$ with respect to L. Unfortunately, there are no analytical solutions for the critical points as they imply complex trigonometric functions (note w appears in the denominator in several occasions). In any case, efficient derivative based root finding algorithms can be applied to solve Equation (3.53), such as Newton's, Halley's or higher order Householder's. This work implements a Newton's iteration with an additional check to ensure solutions are properly bounded, i.e. $L_0 \geqslant L_1 \geqslant L_2 \geqslant L_0 + n_0(t_f - t_0)$: a choice driven by its ease of implementation. In addition, multiple initial guesses for different quadrants are also used, with the purpose of providing robustness to local minima.

Note that the linear impulsive metric is not well suited to estimate phasing maneuvers since these are best performed via changes in the orbital period, as noted in the discussion of the low-thrust metric in Section 3.1.3. Therefore, the maximum reachable anomaly necessarily coincides with the

maximum negative variation in the orbital period, and thus, semi-major axis. Similarly, the upper bound on the reachable semi-major axis gives the minimum reachable anomaly. The latter is very convenient from a computational perspective since only $2(n_x - 1)$ problems need to be solved, thus reducing the computational complexity of an already inexpensive method.

Low-thrust metric

Direct application of the above reachability concepts to the low-thrust metric (see Section 3.1.3) may seem straightforward, since the dynamics are simply given by a linear system. Furthermore, since one is eager to obtain maximal variations along each state dimension, the double control policy approach aimed at satisfying phasing constraints is no longer required. This, in principle, leads to the following definition for the set of TFC that maximizes the change in the j^{th} orbital element

$$f^*_{T,j} = \frac{\hat{G}_{j,:}(\bar{\mathfrak{e}})\Delta t + H_{j,:}(\bar{\mathfrak{e}})}{||\hat{G}_{j,:}(\bar{\mathfrak{e}})\Delta t + H_{j,:}(\bar{\mathfrak{e}})||_2} c_T. \tag{3.54}$$

Again, since a good estimate of the final state after the maneuver is in general not available, one can initially approximate $\hat{G}(\bar{\mathfrak{e}})$ as $G(\bar{\mathfrak{e}}_0)$ and perform successive iterations making use of Equation (3.31) until some tolerance is met, for which two to three iterations might suffice. Note in Equation (3.54) f^*_T is known up to a proportionality constant, as in this case the control metric is not defined by its magnitude but rather a different nonlinear function (see Equations (3.33-3.35)). In this regard, one can also apply a simple iterative process over c_T with $c'_T = \lceil d_{LT} \rceil$ as initial guess, yielding relative errors on the order of 10^{-5} after one to two iterations. Recall, however, that the thrust law is parameterized as a Fourier series, so that there is not a linear dependency between the maximum acceleration and sine-cosine combinations, i.e.

$$\max_t \left[\cos(t) + \sin(t)\right] < 2. \tag{3.55}$$

Therefore, it is of interest to find

$$c : \max_{t} \left[c \cdot \cos(t) + \sin(t) \right] = 2, \tag{3.56}$$

or equivalently

$$c : \max_{t} \left[\cos(t) + c \cdot \sin(t) \right] = 2, \tag{3.57}$$

which turns out to be $c = \sqrt{3}$ in both cases, and apply that factor to the smallest component of each combination of $\sin(kE)$-$\cos(kE)$. This way, the utilization of each basis function of the Fourier series is maximized so that the instantaneous acceleration magnitude is close to the allowable control threshold $\lceil d_{LT} \rceil$ at all times. Finally, as discussed for the impulsive metric, bounds on the angular position (anomaly) are readily obtained from semi-major axis limits, so it is not necessary to explicitly derive them.

3.2.2. N-dimensional search

The previous linear dynamical approximations are certainly appealing from a computational perspective, yet their applicability is limited to a particular state space representation, i.e. modified equinoctial elements. Note it might be of interest to further constrain the reachable set making use of some type of knowledge, e.g. observational, thus utilizing a different state representation. Under these circumstances, the resulting non-linear reachability problem could, for instance, be solved via the method proposed in [133]. However, the computational complexity can be further reduced by exploiting the characteristics of the problem. As defined in Equation (3.44), the reachability problem can be recast as finding the (closed) hyper-surface parameterized by $d(x_0, x, t_0, t_f) = \lceil d \rceil$. In the following, $d(x_0, x, t_0, t_f)$ is assumed to be a black-box or opaque function. Furthermore, since $d(\cdot)$ requires finding the solution to an optimization problem, no assumptions can be made regarding its smoothness, thus discouraging the use of derivative-based algorithms. This severely restricts the set of available tools, especially considering the dimensionality of the problem, i.e. $n_x = 6$.

Within this work, a two-step algorithm is proposed to find an axis aligned

bounding box approximation AABB approximation for $\mathcal{R}(\cdot)$ and, as discussed, only robust, derivative-free algorithms were considered for this purpose. A first step is used to infer the relative scale of the problem: a bisection method is run along each basis direction within the search space. The used scheme can be consulted in Algorithm 4, which requires some initial guess as well as relative and absolute tolerances. To simplify notation, the control distance metric function $d(x_0, x, t_0, t_f)$ is rewritten as $d(y)$, which is an abbreviated form of $d(x_0, y + x^*, t_0, t_f)$. Therein, the solution is computed with respect to some reference state x^*, for instance one can use the ballistic state $x^* \equiv x_0 \equiv \{x : d(x_0, x, t_0, t_f) = 0\}$, i.e. the state at t_f that follows from a null control law. More advanced techniques exist for solving this particular problem, such as the interpolate, truncate and project (ITP) method [134], that exhibits super-linear convergence under any circumstances. Nonetheless, for illustrative and software maintenance purposes, this work has been limited to a simpler, yet less efficient, alternative. The outcome of this first step are the intersections between the reachable set and each basis direction centered at x^*.

Note in general one would expect some inter-dimensional correlations and the maximum along certain dimension may well have some cross-directional component. Accordingly, an additional step can help in approximating the bounds of the reachable set. In this second step, the goal is to find the location of the maximum and minimum distances along each direction, and hence the previous root finding now becomes a set of constrained optimization problems. Again, it is desirable to find a gradient-free method that can be easily tailored to the application at hand. Examples of derivative-free optimization schemes are Powell's methods [135], Nelder-Mead [136], simulated annealing [137] and pattern search [138]. The latter has been a subject of study for several decades and it is still under research [139], arguably due to its flexibility and ease of implementation. It consists in defining a grid or pattern referred to some reference point: if there exist a point within the grid whose cost is lower than that of the reference, such point is selected as the new reference; otherwise the pattern is squeezed to refine the current estimate. There are studies on the convergence properties of pattern search algorithms [140], proven for certain

Algorithm 4: Finding the AABB for the reachable set I. Bisection method step.

REQUIRE: $\lceil d \rceil$ admissible control effort, y_0 initial guess, ε_d absolute control metric tolerance, ε_y relative state tolerance.

for $j' \leqslant 2n$ **do**

  if $j' \leqslant n$ **then**

    $j = j', \quad y_j = y_{0,j}$

  else

    $j = j' - n, \quad y_j = -y_{0,j}$

  $y_j^+ = \infty, \quad y_j^- = -\infty$

  if $d(y_j) > \lceil d \rceil$ **then**

    while $d(y_j) > \lceil d \rceil$ **do**

      $y_j^+ = y_j, \quad y_j = \dfrac{y_j}{2}$

    $y_j^- = y_j$

  else

    while $d(y_j) < \lceil d \rceil$ **do**

      $y_j^- = y_j, \quad y_j = 2y_j$

    $y_j^+ = y_j$

  while $||d(y_j) - \lceil d \rceil||_1 > \varepsilon_d$ **and** $\left|\left|\dfrac{y_j^+ - y_j^-}{y_j}\right|\right|_1 > \varepsilon_y$ **do**

    $y_j = \dfrac{y_j^+ + y_j^-}{2}$

    if $d(y_j) > \lceil d \rceil$ **then**

      $y_j^+ = y_j$

    else

      $y_j^- = y_j$

  if $j' \leqslant n$ **then**

    $y_{max,j} = y_j$

  else

    $y_{min,j} = y_j$

algorithm classes and problem topologies.

The particular characteristics of the problem allow to express it as

$$\begin{array}{ll} \text{Minimize / Maximize:} & y_j \\ \text{subject to:} & d(y) \leqslant \lceil d \rceil, \end{array} \tag{3.58}$$

with $j = \{1, ..., n_x\}$, resulting in a total of $2n_x$ optimization problems as both the maximum and minimum y_j are required. The proposed pattern search, described in Algorithm 5, makes use of the problem scale identified in the first step, i.e. y_{max}, y_{min}, to define the pattern. For each state dimension j it generates an axis aligned grid with all the possible three-dimensional combinations for $\eta_j = 0$ and two-dimensional combinations for $\eta_j = \pm 1$, where η are the individual directions used in the pattern search method. Note only steps with $\Delta y_j \geqslant 0$ are allowed so the pattern only considers the η_j value that advances towards the required direction. These design choices reduce the number of grid points to 141 for $n_x = 6$, which is already a considerable amount. If a grid point with a better cost is found, then it becomes the new reference; otherwise the algorithm looks for the point that is furthest from the control metric constraint. If no grid points satisfy the constraint or their distance to the barrier is lower that the current reference, then the pattern scale is reduced to refine the search. As described, this type of algorithm is not robust to non-convex topologies, where back-steps may be required to find a global minimum. However, it provides a good compromise between accuracy, efficiency and robustness. Moreover, note this scheme can easily accommodate for additional constraints on y, and thus represents a good trade-off between flexibility and computational cost when compared to the linear approach discussed in Section 3.2.1.

3.2.3. Volumetric computations

When defining a least informative transition density for maneuverable objects, as suggested later in Section 3.3.1, it can be of interest to determine the size (volume) of the space assumed accessible to a given target. In this regard, there are two relevant aspects to be considered: 1) the state space

Algorithm 5: Finding the AABB for the reachable set II. Pattern search step.

$\Delta y_0 = (y_{max} - y_{min})/3$
for $j' \leqslant 2n$ **do**
  $u = 0, \quad y = 0$
  if $j' \leqslant n$ **then**
    $j = j', \quad y_j = y_{max,j}, \quad \eta_j = 1$
  else
    $j = j' - n, \quad y_j = y_{min,j}, \quad \eta_j = -1$
  $\Delta y = \Delta y_0, \quad \eta_\perp = [], \quad \eta_\parallel = []$
  for $l \leqslant n$ **do**
    $\eta' = 0$
    if $l \neq j$ **then**
      $\eta'_l = 1$
      for $k \leqslant n, k \neq l$ **do**
        $\eta'' = 0, \eta''_k = 1$
        $\eta_\perp = \begin{bmatrix} \eta_\perp & \eta'' + \eta' & -\eta'' + \eta' & \eta'' - \eta' & -\eta'' - \eta' \end{bmatrix}$
    else
      for $k \leqslant n, k \neq l$ **do**
        $\eta'_k = 1$
        $\eta_\perp = \begin{bmatrix} \eta_\perp & \eta' & -\eta' \end{bmatrix}$
        $\eta_\parallel = \begin{bmatrix} \eta_\parallel & \eta' + \eta & -\eta' + \eta \end{bmatrix}$
  while $\left|\left|\frac{\Delta y \cdot \eta}{y \cdot \eta}\right|\right|_1 > \varepsilon_y$ **do**
    $Y_\parallel = y + \eta_\parallel \cdot \Delta y$
    $y_\parallel = \text{argmin}(d(Y_\parallel))$
    if $d(y_\parallel) \leqslant \lceil d \rceil$ **then**
      $y = y_\parallel$
    else
      $Y_\perp = y + \eta_\perp \cdot \Delta y$
      $y_\perp = \text{argmin}(d(Y_\perp))$
      if $d(y_\perp) \leqslant \lceil d \rceil$ **then**
        $y = y_\perp$
      else
        $\Delta y = \frac{\Delta y}{2}$
  if $j' \leqslant n$ **then**
    $y_{max,j} = y_j$
  else
    $y_{min,j} = y_j$

representation in which the bounds of $\mathcal{R}(\cdot)$ are obtained and 2) the space in which the volumetric integral is computed. Their relevance can be easily appreciated for the case of an Earth orbiting object equipped with chemical propulsion. Assuming the considered time of flight is sufficiently long, the reachable set of such a target is expected to resemble a torus in the Cartesian position space, whose main bounds are related to the energy of the orbit (semi-major axis) and the direction of its angular momentum (inclination and right ascension of the ascending node). Such scenario can hinder the search for an outer approximation of the reachable set, e.g. the numerical AABB method proposed in the previous section would need to be robust to complex geometries and even when finding a proper outer bound, its ratio with respect to the true volume of $\mathcal{R}(\cdot)$ would be high. To mitigate this problem, one can seek a state transformation under which the volume of the reachable set can be well approximated by a bounding box (or a fraction thereof). However, the volume in the more linear (or better behaved) state space representation is prone to lack an intuitive physical interpretation: if used as a normalization constant for Bayesian inference, angular variables would require directional statistics that wrap around a complete revolution. Consider again the case of an Earth orbiting object, and assume its reachable set is well defined by some bounds given in the Keplerian orbital element space. One can readily compute the volume in such a space, yet the information that is conveyed by that integral does not have a clear meaning: the resulting volume would have units of positional distance times angular distance to the power of four. Moreover, if one were to use such volume to normalize a uniform probability density, lower semi-major axis regions would be assigned a higher probability due to the PDF being equally distributed along the different angular elements. Therefore, such a vacuous density may be better expressed in the Cartesian space, as it represents a set of coordinates that do not convey information regarding the dynamical environment, thereby being more compliant with a least-informative (or vacuous) assumption.

Still, one can take advantage of both state space representations since the *dynamically-informed* coordinates are rather convenient to approximate the limits of integration, as detailed in Section 3.2.1. The volumetric inte-

gral can then be computed in any desired state space representation (e.g. Cartesian) making use of the relation

$$\mathcal{V}_{\mathcal{R}} = \int_{\mathcal{R}(x)} dx \equiv \int_{\mathcal{R}(\text{œ})} \det\left(\mathcal{J}|_{\text{œ}}\right) d\text{œ}, \tag{3.59}$$

where

$$\mathcal{J} = \begin{bmatrix} \dfrac{\partial x_1}{\partial \text{œ}_1} & \cdots & \dfrac{\partial x_1}{\partial \text{œ}_{n_x}} \\ \vdots & \ddots & \vdots \\ \dfrac{\partial x_{n_x}}{\partial \text{œ}_1} & \cdots & \dfrac{\partial x_{n_x}}{\partial \text{œ}_{n_x}} \end{bmatrix} \tag{3.60}$$

is the Jacobian of the transformation from the orbital element (œ) to, in this case, the Cartesian (x) space. The problem thus reduces to solving a high-dimensional integral ($n_x = 6$) with approximate (enclosing) limits of integration. In practice, analytical integrals for the determinant of the Jacobian are seldom found and the present case is not an exception, so it is necessary to resort to numerical integration techniques. One of the most capable integration methods for high-dimensional problems relies on the use of randomized lattice rules, which combine deterministic low-discrepancy, also called quasi-Monte Carlo (qMC), sequences with some small aleatoric noise that improves convergence for arbitrary integrands [141]. Note there may be numerical resonances between the deterministic (low-discrepancy) sequence and the function to be integrated, thus adding a small perturbation can improve convergence for general applications. An instance of a randomized lattice rule method is the scrambled Sobol approach $\mathcal{S}(\cdot)$, which generates a quasi-deterministic sequence in the $\mathbb{R}^{n_x}$ unitary hypercube with a sampling efficiency that outperforms other alternatives, e.g. latin hypercube sampling, in relatively high dimensions. Following this approach, the integral in Equation (3.59) can be approximated as

$$\mathcal{V}_{\mathcal{R}} \approx \frac{1}{N_s} \prod_{k=1}^{n} (\text{œ}_{max,k} - \text{œ}_{min,k}) \sum_{i \in N_{\mathcal{R}}(\text{œ})} \det\left(\mathcal{J}|_{\mathbb{E}_i}\right), \tag{3.61}$$

where $\text{Œ}_i \sim \mathcal{S}(\text{œ}_{min}, \text{œ}_{max})$ and

$$N_{\mathcal{R}}(\text{œ}) = \{\text{œ} : d(x_0, \text{œ}, t_0, t_f) \leqslant \lceil d \rceil\} \tag{3.62}$$

is the subset of samples that belong to the reachable set. Hence, if the limits $[\text{œ}_{min}, \text{œ}_{max}]$ serve as a proper bounding box for the reachable set, then the approximation in Equation (3.61) converges to the true value of $\mathcal{V}_{\mathcal{R}}$ as the number of samples $N_s \to \infty$.

3.3. Maneuver modeling

One of the main problems associated to tracking maneuvering spacecraft is the lack of a proper dynamical characterization. The high level of non-linearity, coupled with sparse observations, render common methods in the tracking literature unusable for space surveillance. Some of these include inflated process noise levels or constant acceleration models, which do not provide reasonable physical bounds for maneuvering capabilities in space: note for instance, that plane change maneuvers require symmetric control policies with respect to the nodes or apices of the orbit. The control metrics developed in Sections 3.1.1 and 3.1.3 allow for a better characterization of orbital maneuvers, while ensuring a reduced computational cost. Therefore, these can be used to elaborate maneuver models that are consistent with the special dynamical and current technological characteristics of two-body problem transfers. Based on these metrics, and the ad-hoc reachability computation methods proposed in Section 3.2, one can readily elaborate control bounded maneuver transition densities to be used in advanced multi-object filtering frameworks such as the GLMB filter discussed in Section 2.1.1. The remainder of this section presents a series of maneuver models that can be used in a multi-target SST context.

3.3.1. Uncooperative tracking

SST applications mostly reduce to uncooperative tracking since surveillance sensors lack target specific knowledge and, though SST entities de-

mand spacecraft operators to share their maneuver plan, the latter may not be willing to provide such operational knowledge or may not deliver it in a timely manner. Therefore, it is convenient to derive agnostic maneuver models that can accommodate for general orbital transfers, let them be performed by commercial or military targets. As thoroughly discussed in the literature, and also earlier in this document, the two main types of spacecraft propulsion present diametrically opposed dynamical characteristics, and thus should be treated differently. Nonetheless, bounds on the maneuvering capabilities of a certain target can be estimated with relative ease, be it equipped with chemical or electric thrusters. Let the reachability inclusion function be defined as

$$\phi(x, x_+) = \begin{cases} 1 & \text{if } d(x, x_+) \leqslant \lceil d \rceil, \\ \\ 0 & \text{if } d(x_+, x) > \lceil d \rceil, \end{cases} \tag{3.63}$$

indicating whether the state x_+ is accessible given the prior ballistic realization x. Making use of the above definition, one can express the least-informative (or vacuous) maneuver transition density as

$$p_T^{(m_+)}(x_+|x) = \frac{\phi(x, x_+)}{\bar{\phi}(x)}. \tag{3.64}$$

Note as per the requirements of the CK Equation (2.26), the equality $\int p_T^{(m_+)}(x_+|x)dx_+ = 1$ shall be enforced for $p_T^{(m_+)}(x_+|x)$ to be a proper density. Hence, it is necessary to compute the integral

$$\bar{\phi}(x) = \int \phi(x, x_+)dx_+, \tag{3.65}$$

as it gives the required normalization constant. It is perhaps not surprising that the above expression exactly matches the volume of the reachable set defined in Equation (3.59), since the support of the reachability inclusion function used above precisely coincides with the bounds of $\mathcal{R}(x)$. The maneuvering capabilities of the target are thus characterized by the control distance metric: they are assumed to be isotropic in the Cartesian space

bounded by an admissible control effort $\lceil d \rceil$, itself based on some expected propulsion system capabilities. A possible target agnostic filter scheme could rely on a blend of dynamical models: natural or ballistic, chemical thruster propelled based on the control metric defined in Section 3.1.1, and electric propulsion powered making use of the metric introduced in Section 3.1.3. In principle, such a scheme could maintain custody of a general uncooperative target orbiting the Earth regardless of its maneuver characteristics, provided none of its exerted maneuvers exceed the bounds prescribed by $\lceil d \rceil$ in the metric space.

3.3.2. Heuristics

The vast majority of spacecraft maneuvers are intended to maintain the satellite within pre-assigned orbital slots, commonly referred to as station keeping. Therefore, maneuver detection and estimation methods are sometimes built around the so-called patterns of life, i.e. the circumstances surrounding previous maneuvers, see [19], [21], [142]. Knowledge from the series of past events is analyzed to derive heuristic models, commonly following statistical [21] or machine learning [19] approaches. Despite the different techniques used therein, their proposal basically reduces to finding patterns within the set of past orbital maneuvers that can be used to predict and approximate future events, preferably in a statistical sense. Station keeping is mostly concerned about the overall trend of the orbital shape, rather than specific point to point features. Thus, patterns of life only apply to certain orbital elements or reduced combinations thereof. This implies there is some subset of the state space for which these models are not defined, where the concept of reachability can provide reasonable bounds. In fact, there is evidence regarding the expected cost of previous maneuvers of certain type (say east-west or north-south station keeping in GEO), so one can apply this knowledge to further constrain the support of the heuristic model compared to the vacuous reachable set formulation described in Section 3.3.1. Note however, that absolute changes in orbital elements already present an associated control cost, hence the magnitude of the maneuver itself is somewhat implicit to them.

The current work can be considered an incremental evolution over the proposal of Siminski et al. [21], thereby characterizing maneuver modes in terms of the relative variation in certain orbital elements

$$\zeta = \left\{ \Delta a \quad \Delta e \quad \Delta I \right\}. \tag{3.66}$$

The selection of these three elements is largely influenced by expectations on operator maneuver policies: phasing (east-west for GEO orbits) and Ω-change maneuvers in LEO are typically executed through variations in the semi-major axis and eccentricity [143], [144], whereas inclination needs to be targeted on its own [145]. Accordingly, each detected maneuver has an associated feature vector distribution that results from comparing the pre- and post-maneuver state densities. Patterns can then be inferred from the maneuver history, for instance via clustering methods such as the expectation-maximization algorithm [146], resulting in a kernel density estimator (KDE) characterization of the heuristic modes

$$p_T^{(m_+)}(x_+|x) = \frac{\phi(x, x_+)}{C(x)} \sum_{i=1}^{n_M} \omega_i \mathcal{N}(\zeta; \hat{\zeta}_i . \Sigma_{\zeta,i}), \tag{3.67}$$

Therein, n_M represents the number of independent (clustered) maneuver modes and the feature vector ζ corresponds to the change in a, e, and I from x to x_+. In principle, each individual maneuver mode could be represented by a single transition density, to the detriment of tractability. The use of a mixture distribution for the heuristics-based mode reduces the complexity of the filtering process to a minimum, since alternative heuristic mode hypotheses are combined. An additional complication of this approach concerns determining the normalization constant for the transition density in Equation (3.67). It is not trivial to approximate the support of these maneuver modes, but one can take advantage of the qMC method presented in Section 3.2.3 and re-weight each individual sample according to every Gaussian component. The normalizing constant for the heuristic transition density will then be the weighted sum of each individual normalization

constant, i.e.

$$C(x) \approx \frac{1}{N_s} \prod_{k=1}^{n} (\text{œ}_{max,k} - \text{œ}_{min,k}) \sum_{j=1}^{N_s} |J| \phi(\text{Œ}_j) \sum_{i=1}^{n_M} \omega_i \mathcal{N}(\text{Œ}_j; \hat{\zeta}_i, \Sigma_{\zeta,i}), \quad (3.68)$$

making use of Equation (3.61). Though it might seem computationally expensive, note these samples are already available from evaluation of the general (vacuous) maneuver mode, and thus require little additional computational effort.

4. RESULTS AND DISCUSSION

Space surveillance and tracking systems demand methods and techniques to maintain custody of maneuvering spacecraft, with the ultimate goal of providing accurate and updated collision risk assessments. Tracking maneuvering spacecraft can be framed under the general stochastic hybrid systems framework, in which the system dynamics are given by one of a finite number of a priori characterized modes. Advanced multi-target maneuvering target tracking methods embrace this type of framework for state space filtering as it provides a natural approach under Bayesian assumptions. However, the high non-linearities present in general SST operations render common approaches inaccurate for tracking maneuvering spacecraft. Therefore, one needs to resort to surrogate models for maneuver characterization, which can help in reducing the computational burden associated to orbital maneuver estimation while ensuring a certain level of dynamical fidelity. Under the absence of a priori knowledge on maneuver characteristics, it is reasonable to assume vacuous maneuver transition densities within some reachable domain, this is, the region of the state space assumed accessible to a certain target. Bounds for such reachable set can be derived based on control distance metrics: acceleration (or thrust) magnitude for electrically propelled spacecraft and ΔV cost for chemical thrusters. Still, there does not seem to exist efficient methods to infer the above metrics for a general time-constrained two-point boundary value problem. To advance further in the orbital maneuver modelling literature, different surrogate models that enable an efficient estimation of such control metrics have been developed, tailored to the different types of space propulsion (see Sections 3.1.1 and 3.1.3). Results obtained from the use of such surrogate models are analyzed and compared to higher fidelity approaches later in Section 4.1, thereby assessing their accuracy and range of validity. Furthermore, approximating the bounds of the reachable set, defined in terms of the aforementioned control distance metrics, can be an arduous task if numerical methods such as the one proposed in Section 3.2.2 are used. Fortunately, one can make use of linear surrogates

to allow for tractable and robust reachability computations, as discussed throughout Section 3.2.1. It is nonetheless necessary to numerically evaluate the accuracy of such linear approaches when compared to direct evaluation of the metrics themselves, and calibrate them if necessary. These tests, together with a convergence analysis on the volumetric computation techniques described in Section 3.2.3, are presented and discussed in Section 4.2.2, thereby supporting a tractable implementation of the maneuver models proposed in Section 3.3. These transition densities are tested in two different single target scenarios in Section 4.3: a chemically powered GEO target detected by electro-optical sensors and a low-thrust orbit raising LEO target sensed by surveillance radars. Results from these tests are used to define a sensible implementation of the GLMB filter (see Section 2.1.1), described and evaluated in Section 4.4. Despite being simplified test cases with pre-processed data, results obtained for the implemented GLMB are quite promising for SST as it allows to maintain custody of maneuvering spacecraft with surveillance data in a multi-target environment, even providing consistent state estimates with a rather limited amount of post-maneuver observations.

4.1. Empirical validation of proposed control distance metrics

Analytical proofs of the convergence and accuracy of the developed methods is, to the author knowledge, non-existent. Note these require finding optimal solutions to surrogate problems, which may or may not adequately represent the higher fidelity model they approximate. Complex non-linear transformations, that include but are not limited to orbital averaging, are ubiquitous. Therefore, to assess the validity of the proposed control metrics within the expected application context, several test scenarios have been devised. Tailored to the characteristics of both the high- and low-thrust control metrics, these tests are aimed at capturing the combined inaccuracies derived from the dynamical approximation and the proposed solution approach. They are arranged as a series of randomized conditions, comprising the pre- and post-maneuver states as well as the time of flight. Ballistic state sampling is concentrated around highly populated areas of Earth's orbital

space, whereas the (randomized) orbital changes caused by maneuvers are designed to yield reasonable control distance magnitudes. Time spans for maneuver execution correspond to typical revisit periods in space surveillance, ranging from a few hours to a couple of days.

These analyses are intended to shed light on the applicability of the developed metrics, both from an accuracy and computational perspective. Provided the estimated control effort correlates well with higher fidelity maneuver estimation models, they may provide a good trade-off for large scale applications. Furthermore, if the computational complexity is reduced to a minimum, the developed metrics could be used in iterative or even Monte Carlo-like procedures. Besides these highly desirable features, it is also paramount for them to show robustness in the expected application domain, i.e. elliptical near-Earth orbits. Ideally, one would hope the proposed approaches to yield estimation errors on the order of 10%, with a clear bias towards higher control magnitudes. This would stem from the inability of the proposed numerical approaches to locate the global minimum of the cost function, thus yielding sub-optimal solutions. Moreover, dynamical approximations may also play a relevant role, especially if the assumptions made when developing the surrogate model do not hold. These may require to set biases and scaling factors to the output of the proposed computationally efficient methods, so that the analyses included in the following are effectively used for model calibration.

4.1.1. High-thrust Lambert-based metric

As developed in Section 3.1.1, the impulsive control metric requires finding the epochs that minimize the total delta-V of a two-burn transfer between some boundary conditions. The expected topology of this optimization problem is quite complex, with a significant number of local minima. Thus, it is of interest to first evaluate the ability to find a sufficiently good solution, whose associated cost is not far from the global minimum. To this end, Lambert's problem is solved in a grid of 500 by 500 points in t_1 and t_2, similar to the one used in Figure 3.2. The coordinates with minimum control cost within such grid are then evolved through a Gauss Newton

method, namely the same used to refine the initial conditions described in Equations (3.6-3.8). In addition, it is of interest to evaluate the accuracy of these estimates in a realistic environment, thereby inferring the dynamical mis-modelling addressed in Section 3.1.4. To this end, the *orbit-to-orbit* method developed by Pastor et al. [115] has been employed, where the effect of the first and second impulsive maneuvers is approximated via the STMs of the pre- and post-maneuver orbits, respectively. These matrices are computed for a temporal grid of certain size, set to 20 points per orbit for the present test case, wherein the minimum total delta-V is sought. Note the latter would always correspond to a suboptimal solution as the optimal burn epochs are not solved for, though it provides a computationally feasible approximation to a more realistic dynamical environment.

These three solution types are compared in randomized test scenarios in LEO, highly elliptical orbit (HEO) and GEO, whose characteristics are summarized in Tables 4.1, 4.2 and 4.3, respectively. Differences in the distributions for relative orbital element variations are driven by the sensitivities of each orbital regime: phasing maneuvers (performed via the semi-major axis) are less costly for lower orbits since the time of flight is larger compared to their orbital period, and absolute changes in perigee radius require less energy the higher the semi-major axis of the baseline orbit. An additional set of parameters needs to be defined as per the requirements of the testing procedure, whose statistics are described in Table 4.4. These comprise environmental parameters used in high-fidelity modelling, such as the drag coefficient c_D, the reflectivity coefficient c_r and the area to mass ratio A/m. Furthermore, realistic simulations require defining a proper initial epoch besides the time of flight, in this case spanning two full years to cover seasonal and annual variations, for instance in the solar flux. The high-fidelity dynamical model used for testing presents the following features:

- The software is implemented in MATLAB.
- Dynamics are propagated using a variable-step variable-order Adams-Bashford-Moulton method (*ode113* in MATLAB).

$r_{p,0} \sim R_\oplus + \mathcal{U}(500, 1000)$ km	$\Delta r_{p,0} \sim \mathcal{U}(-1, 1)$ km
$e_0 \sim 10^{[\mathcal{U}(-5,-0.8)]}$	$\Delta e_0 \sim \text{sgn}\left[\mathcal{U}(-1, 1)\right] 10^{[\mathcal{U}(-5,-3.5)]}$
$I_0 \sim \mathcal{U}(40, 100)$ [deg]	$\Delta I_0 \sim \mathcal{U}(-0.01, 0.01)$ [deg]
$\Omega_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\Omega_0 \sim \mathcal{U}(-0.01, 0.01)$ [deg]
$\varpi_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\varpi_0 \sim \mathcal{U}(-0.1, 0.1)$ [deg]
$M_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta M_0 \sim \mathcal{U}(-1, 1)$ [deg]

Table 4.1: Distributions of the randomized LEO test case for the high-thrust control metric. Each individual orbital element change yields a control cost around 2 m/s.

- Solar radiation pressure is considered, using a cannonball model and conical eclipses.
- Atmospheric density follows Jacchia-Bowmann 2008 [147] model, mainly due to its trade off between accuracy and computational efficiency when compared to, for instance, NRLMSISE00 [148].
- Third body perturbations of the Sun, Moon and main planets are also accounted for, using NASA's JPL DE403 ephemerides model.
- The Earth gravity field is modelled with sectorial and tesseral harmonics up to degree and order 10.
- Solid and ocean Earth tides are also included for completeness.

The test configuration detailed in Tables 4.1, 4.2, 4.3 and 4.4 leads to the ΔV cost distributions depicted in Figure 4.1, whose values correspond to the solution given by the STM-based higher fidelity approach. The three test scenarios present a mode around 1.5 m/s, though LEO and HEO show a thicker right tail due to higher ΔV cost samples. Typical orbital maneuvers of operational spacecraft are of a lower magnitude, on the decimeter per second level for east-west station keeping in GEO, though the broader support of the test definition can help to better understand the accuracy of the

$r_{p,0} \sim R_\oplus + \mathcal{U}(500, 10^4)$ km	$\Delta r_{p,0} \sim \mathcal{U}(-5, 5)$ km
$e_0 \sim \mathcal{U}(0.2, 0.9)$	$\Delta e_0 \sim \text{sgn}\left[\mathcal{U}(-1, 1)\right] 10^{[\mathcal{U}(-5,-3.5)]}$
$I_0 \sim \mathcal{U}(10, 80)$ [deg]	$\Delta I_0 \sim \mathcal{U}(-0.02, 0.02)$ [deg]
$\Omega_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\Omega_0 \sim \mathcal{U}(-0.02, 0.02)$ [deg]
$\varpi_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\varpi_0 \sim \mathcal{U}(-0.04, 0.04)$ [deg]
$M_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta M_0 \sim \mathcal{U}(-0.2, 0.2)$ [deg]

Table 4.2: Distributions of the randomized HEO test case for the high-thrust control metric. Each individual orbital element change yields a control cost around 2 m/s.

$r_{p,0} \sim 42157 + \mathcal{U}(-10, 10)$ km	$\Delta r_{p,0} \sim \mathcal{U}(-10, 10)$ km
$e_0 \sim 10^{[\mathcal{U}(-5,-3.5)]}$	$\Delta e_0 \sim \text{sgn}\left[\mathcal{U}(-1, 1)\right] 10^{[\mathcal{U}(-5,-3.5)]}$
$I_0 \sim \mathcal{U}(-5, 5)$ [deg]	$\Delta I_0 \sim \mathcal{U}(-0.02, 0.02)$ [deg]
$\Omega_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\Omega_0 \sim \mathcal{U}(-0.06, 0.06)$ [deg]
$\varpi_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\varpi_0 \sim \mathcal{U}(-0.06, 0.06)$ [deg]
$M_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta M_0 \sim \mathcal{U}(-0.06, 0.06)$ [deg]

Table 4.3: Distributions of the randomized GEO test case for the high-thrust control metric. Each individual orbital element change yields a control cost around 2 m/s.

$c_D \frac{A}{\text{m}} \sim \mathcal{U}(8.46 \cdot 10^{-3}, 0.254)$ m^2/kg	$c_r \frac{A}{\text{m}} \sim \mathcal{U}(4.81 \cdot 10^{-3}, 144)$ m^2/kg
$t_0 \sim \mathcal{U}(01/01/2018, 31/12/2019)$	$t_f - t_0 \sim \mathcal{U}(3, 48)$ [hr]

Table 4.4: Distributions for the parameters of the randomized high-thrust metric tests.

proposed approach. In fact, for sufficiently small maneuvers, e.g. 1 cm/s, first or second order Taylor approximations for the dynamics can yield accurate results and they might even be difficult to distinguish from process noise in certain environments, e.g. LEO. The proposed Lambert-based solution approach for high-thrust optimal two-burn transfers was designed to maximize computational efficiency, which has succeeded up to certain level judging by the results displayed in Table 4.5. Therein, the high-fidelity approach considers back-propagation of both the pre- and post-maneuver orbits to find the state transition matrices throughout a dense temporal grid. Note that, in principle, these STMs might be available for the ballistic state but not for the hypothesized post-maneuver one, so that a fair comparison for SST applications would reduce the high-fidelity runtimes to half of the presented values. Regardless, the proposed method shows a two-order of magnitude decrease in computational cost, thereby enabling a broad set of sampling based applications. Note that the cost of GEO and LEO / HEO tests differ almost by an order of magnitude for both approaches. In the high-fidelity case this is due to the lack of atmospheric drag and its corresponding *easier* dynamics from a variable step integrator perspective. For the Lambert-based approach there is a lower number of initial conditions since some of those specified in Equations (3.6) and (3.7) do not meet the requirement in Equation (3.8) and, as such, are automatically discarded. In fact, these numbers correspond to a MATLAB implementation, that can be sped up by one order of magnitude if implemented in a compiled language, e.g. C or C++.

Comparative results for the accuracy of the proposed metric are illustrated in Figures 4.2, 4.3 and 4.4 for the LEO, HEO and GEO test scenarios, respectively. The histogram corresponds to the relative error of the approach described in Section 3.1.1 with respect to the STM-based method proposed in [115]. Note two different cumulative mass functions are included in these figures, where the one labeled *Minimum Lambert* corresponds to an iterated global minimum obtained from a 500 by 500 temporal grid. As expected, the latter is shifted towards lower ΔV cost estimates as it seeks to provide the global optimal solution, or at least a better approximation. In this regard, Figure 4.5 depicts the distribution of the number

	Lambert Approach	High-Fidelity STM Grid
LEO	0.26 s	140 s
HEO	0.26 s	110 s
GEO	0.02 s	10 s

Table 4.5: Average runtimes for estimating the ΔV cost of an impulsive transfer using the Lambert-based method described in Section 3.1.1 and the STM-based approach proposed in [115]. These correspond to MATLAB implementations executed on an Intel Core i7-8750H CPU.

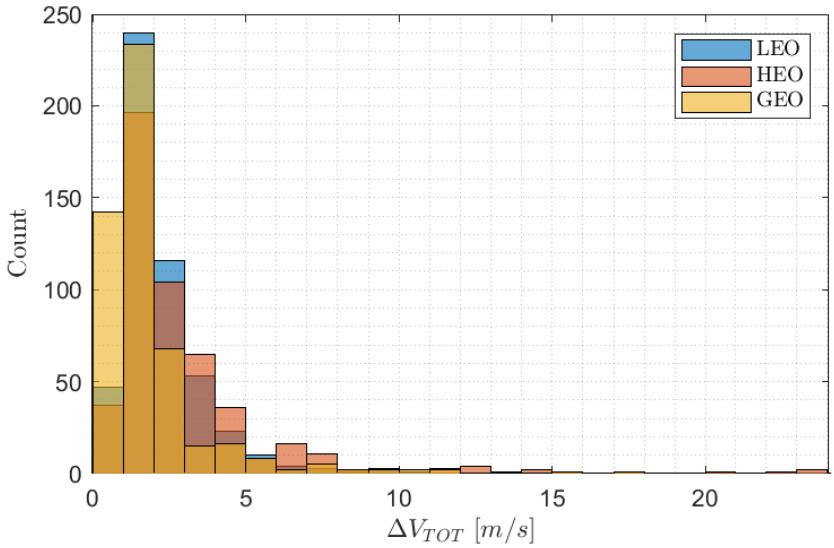

Figure 4.1: Distribution of ΔV cost for each test scenario.

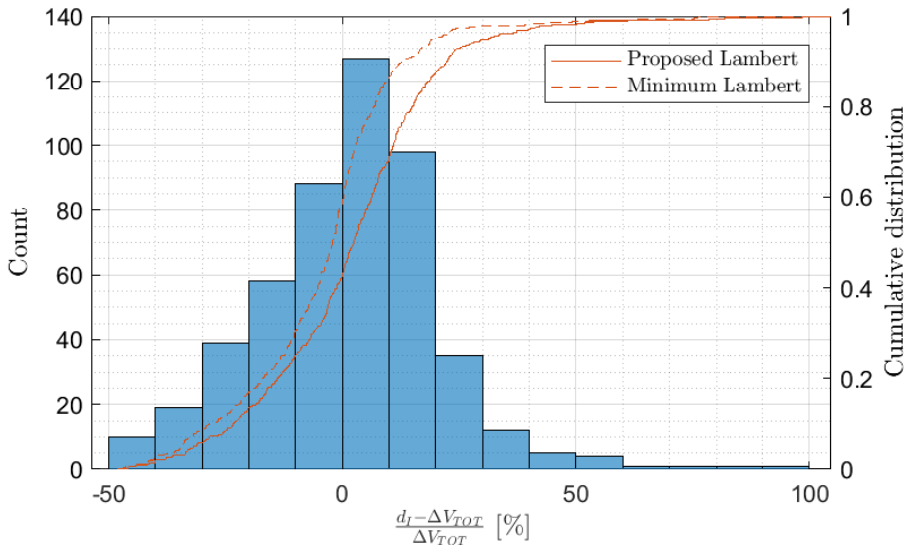

Figure 4.2: ΔV estimation error statistics for the LEO tests.

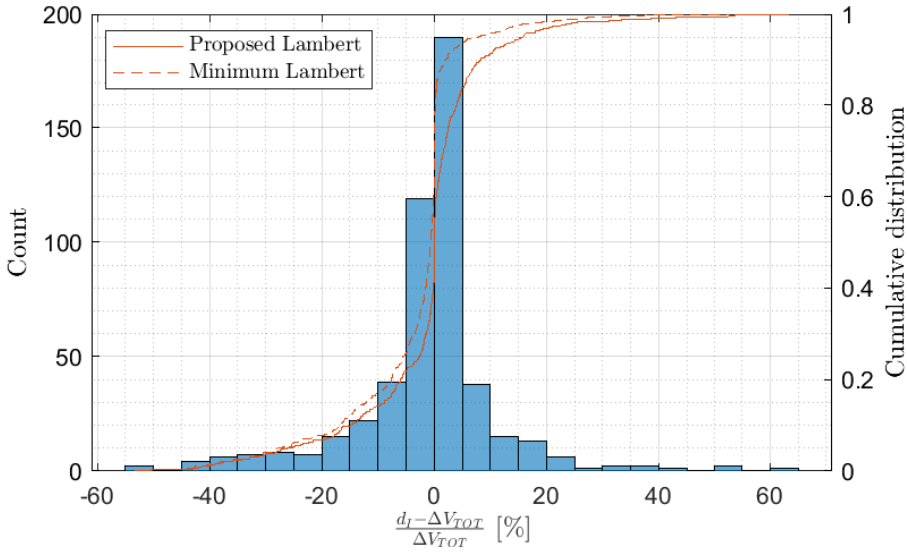

Figure 4.3: ΔV estimation error statistics for the HEO tests.

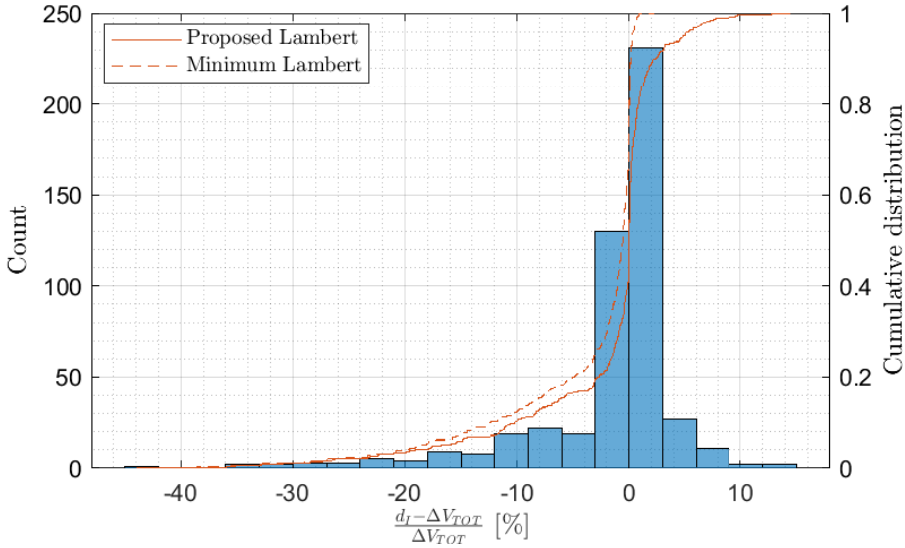

Figure 4.4: ΔV estimation error statistics for the GEO tests.

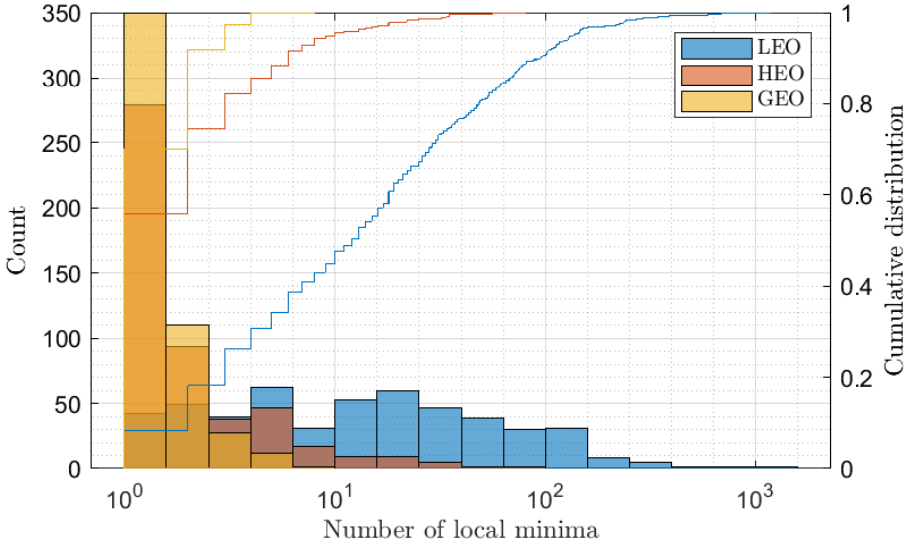

Figure 4.5: Distribution of number of local minima for each test scenario. A local minima is identified as each separable region with a ΔV cost lower than 110% of the global minimum cost of the Lambert grid.

of local minima, i.e. connected regions with an estimated control effort below 110% the reference magnitude. Note this is highly correlated with the number of orbital revolutions elapsed during the time of flight, and thus LEO frequently presents a number of local minima in excess of 100. In principle, this would negatively affect the comparison with respect to the iterated global minimum approach, to be decoupled from the STM-based one [115] as the latter implements a higher fidelity dynamical model. Errors associated to not finding the global optimum within the search space reduce to a bias of -5% in LEO and less than -1% in HEO and GEO, well below the variance when compared to the reference solution. In fact, recall the reference high-fidelity method does not estimate the optimal burn epoch beyond a temporal grid of 20 points per orbit, so that this reference estimate is guaranteed to be suboptimal. Therefore, and provided the solution to the Lambert-based methods can be considered locally optimal, the high number of underestimations are somewhat justified. Moreover, the ΔV estimation error statistics for GEO (see Figure 4.4) are rather promising: the Keplerian dynamics used to solve Lambert's problem do not seem to negatively affect the estimation, apart from certain samples with highly underestimated costs. Overall, the 95% confidence interval for GEO lies within the $[-25, 5]$ % range, which is something affordable given the decrease in computational requirements. The HEO scenario presents a higher error variance, as in this case the confidence interval is approximately given by $[-35, 10]$ %. Note there is a high portion of HEO orbits whose perigee is affected by atmospheric drag, so that the assumed Keplerian dynamics present inconsistencies when compared to a higher fidelity approach. The latter fact is magnified for lower Earth orbits, where the dominant effects of J_2 and drag render the unperturbed approximation rather inaccurate. While the expected value coincides for the proposed method and the high-fidelity approach, which holds true for all the scenarios, the error variance in this case is higher: the 95% confidence interval now reads $[-40, 40]$ %. Moreover, the error distribution is no longer skewed towards lower ΔV estimates, suggesting that the proposed approach underestimates the required control. However, the estimates are shown to remain within ± 50 % relative error in the worst case, which may be absorbed to some extent by uncertainty from

modelling drag forces or measurement noise.

4.1.2. Low-thrust metric based on thrust Fourier coefficients

To avoid solving a non-linear optimal control problem, the metric developed in Section 3.1.3 adopts a periodic parametrization of the control acceleration in the form of a Fourier series. Conveniently, the resulting averaged dynamics only present a linear dependency on a reduced set of fourteen coefficients. Therefore, one can pose a surrogate linear problem to derive the (optimal) set of TFC compliant with certain orbital boundary conditions, given by the pre- and post-maneuver states. As a benchmark, this control distance metric is compared to the solution of the complete non-linear optimal control problem. This problem is usually expressed as

$$\begin{aligned} &\text{Minimize:} \quad J = \int_{t_0}^{t_f} u^T u dt, \\ &\text{subject to:} \quad x(t_0) = x_0, \quad x(t_f) = x_f, \quad \frac{dx}{dt} = f(x,u,t), \end{aligned} \tag{4.1}$$

which specifies the control input sequence required to minimize the integral control magnitude (squared). As already discussed throughout Section 3.1, a better characterization of the cost associated to low-thrust transfers is sought, preferably one which decouples the influence of the time of flight. A naive approximation would consider the use of a time-averaged control effort, e.g. $d = \frac{\sqrt{J}}{t_f - t_0}$, though this definition may not adequately represent the propulsive requirements of the underlying transfer. Therefore, the maximum instantaneous control input is proposed as a measure of the orbital distance for low-thrust equipped spacecraft. In fact, one can reformulate the optimization problem in Equation (4.1) as

$$\begin{aligned} &\text{Minimize:} \quad J = \max_t \quad u^T u, \\ &\text{subject to:} \quad x(t_0) = x_0, \quad x(t_f) = x_f, \quad \frac{dx}{dt} = f(x,u,t), \end{aligned} \tag{4.2}$$

to find the minimum thrust requirements to travel from the pre-maneuver state x_0 to the post-maneuver state x_f. These non-linear programming

(NLP) problems are solved using a collocation method with third order Legendre polynomials and 201 equidistant points in the time domain. This configuration is selected to ensure sufficient smoothness in the solution, while preventing singularities that may arise from higher order polynomial approximations. CasADi software [149] is used to pose the optimization problem, which is later solved via an interior point (ipopt [150]) solver. The solution to Equation (4.2) shows rather poor convergence for arbitrary initial conditions, mainly since the transcription method used requires to reformulate it as

$$\begin{aligned} &\text{Minimize:} \quad J = u_{max}, \\ &\text{subject to:} \quad x(t_0) = x_0, \quad x(t_f) = x_f, \quad \frac{dx}{dt} = f(x,u,t), \quad u_{max}^2 \geqslant u^T u. \end{aligned} \tag{4.3}$$

To improve convergence, the problem in Equation (4.1) is solved and then its optimal solution is used as initial guess the for non-linear problem in Equation (4.3). This way it is possible recover a local optimum in the vicinity of the solution to the general optimal low-thrust transfer problem. Furthermore, algorithmic differentiation frameworks such as CasADi require differentiable functions to operate, colliding with common ephemeris-based high fidelity dynamical models. Thus, only J_2-perturbed Keplerian dynamics are used for the optimal control problem, so the proposed control metric implementation only considers such perturbative force (see Section 3.1.4).

The efficient low-thrust metric is compared against the solution to the optimal control problem defined in Equations (4.1) and (4.3) for a series of randomized conditions. Three different scenarios were designed to cover the three main types of Earth orbits: LEO, HEO and GEO, whose conditions are given in Tables 4.6, 4.7 and 4.8, respectively. For each scenario, 500 samples are drawn from the distributions specified in these tables, resulting in a total of 1,500 individual test cases. Note all the scenarios present the same time of flight characteristics, and these corresponds to expected re-visit times for SST operations. The latter has severe implications in the performance and validity of the proposed approach: LEO orbits easily perform more than 10 full orbital revolutions, whereas geosynchronous orbits barely reach two in the limit case. This is exemplified in Figures 4.6,

$r_{p,0} \sim R_\oplus + \mathcal{U}(500, 1000)$ km	$\Delta r_{p,0} \sim \mathcal{U}(-2.5, 2.5)$ km
$e_0 \sim 10^{[\mathcal{U}(-5,-0.8)]}$	$\Delta e_0 \sim \text{sgn}\left[\mathcal{U}(-1,1)\right] 10^{[\mathcal{U}(-5,-4)]}$
$I_0 \sim \mathcal{U}(40, 100)$ [deg]	$\Delta I_0 \sim \mathcal{U}(-0.01, 0.01)$ [deg]
$\Omega_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\Omega_0 \sim \mathcal{U}(-0.01, 0.01)$ [deg]
$\varpi_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\varpi_0 \sim \mathcal{U}(-0.5, 0.5)$ [deg]
$M_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta M_0 \sim \mathcal{U}(-0.5, 0.5)$ [deg]
$t_f - t_0 \sim \mathcal{U}(3, 48)$ [hr]	

Table 4.6: Distributions of the randomized LEO test case for the low-thrust control metric. Each individual orbital element change yields a control cost around 10^{-4} m/s^2.

$r_{p,0} \sim R_\oplus + \mathcal{U}(500, 10^4)$ km	$\Delta r_{p,0} \sim \mathcal{U}(-2.5, 2.5)$ km
$e_0 \sim \mathcal{U}(0.2, 0.9)$	$\Delta e_0 \sim \text{sgn}\left[\mathcal{U}(-1,1)\right] 10^{[\mathcal{U}(-5,-3.5)]}$
$I_0 \sim \mathcal{U}(10, 80)$ [deg]	$\Delta I_0 \sim \mathcal{U}(-0.01, 0.01)$ [deg]
$\Omega_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\Omega_0 \sim \mathcal{U}(-0.01, 0.01)$ [deg]
$\varpi_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\varpi_0 \sim \mathcal{U}(-0.05, 0.05)$ [deg]
$M_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta M_0 \sim \mathcal{U}(-0.15, 0.15)$ [deg]
$t_f - t_0 \sim \mathcal{U}(3, 48)$ [hr]	

Table 4.7: Distributions of the randomized HEO test case for the low-thrust control metric. Each individual orbital element change yields a control cost around 10^{-4} m/s^2.

$r_{p,0} \sim 42157 + \mathcal{U}(-10, 10)$ km	$\Delta r_{p,0} \sim \mathcal{U}(-10, 10)$ km
$e_0 \sim 10^{[\mathcal{U}(-5,-3.5)]}$	$\Delta e_0 \sim \text{sgn}\left[\mathcal{U}(-1, 1)\right] 10^{[\mathcal{U}(-5,-3.5)]}$
$I_0 \sim \mathcal{U}(-5, 5)$ [deg]	$\Delta I_0 \sim \mathcal{U}(-0.005, 0.005)$ [deg]
$\Omega_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\Omega_0 \sim \mathcal{U}(-0.05, 0.05)$ [deg]
$\varpi_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta\varpi_0 \sim \mathcal{U}(-0.05, 0.05)$ [deg]
$M_0 \sim \mathcal{U}(-\pi, \pi)$ [rad]	$\Delta M_0 \sim \mathcal{U}(-0.05, 0.05)$ [deg]
$t_f - t_0 \sim \mathcal{U}(3, 48)$ [hr]	

Table 4.8: Distributions of the randomized GEO test case for the low-thrust control metric. Each individual orbital element change yields a control cost around 10^{-4} m/s^2.

4.7 and 4.8, which depict the thrust control laws required to perform certain orbital transfer for low Earth, highly elliptical and GEO orbits, respectively. Therein, one can appreciate the difference in control policies that results from minimizing the integral cost (blue), the maximum required thrust (red) and solving the surrogate (linear) TFC-parameterized problem (dark yellow). Perhaps unsurprisingly, the surrogate solution matches quite well for the three cases presented in terms of overall trend and magnitude. The benefits of splitting the TFC in two segments to target the final orbital phase can be visualized in these examples, as the optimal solution tends to evolve from t_0 to t_f in the along-track u_T direction. Throughout the testing process, the NLP problem solver is initialized both with a null control input and the solution to the aforementioned TFC surrogate, and the one featuring the lowest cost function from these two is selected as a reference. The latter improves convergence since, especially for multi-revolution (i.e. LEO) transfers, a naive null initial guess may not be sufficiently informative for interior point NLP solvers to find an optimal solution within reasonable time. As expected, the minimum thrust control law (red) yields lower maximum instantaneous acceleration requirements, yet the proposed surrogate estimates lower maximum thrust demands than the baseline opti-

TFC Surrogate	Optimal Control Problem
0.06 s	31.2 s

Table 4.9: Average runtimes for estimating the thrust requirements for the developed surrogate approach (TFC) and the optimal control problem defined in Equation (4.1). These correspond to MATLAB implementations executed on an Intel Core i7-8750H CPU.

mal control problem. By looking at the complete evolution, the overall (or integral) thrust demands of the TFC approach seem higher but the optimal control solution presents certain peaks in acceleration magnitude. These peaks are mostly predominant in the along-track component of the velocity and, in general, associated to orbital phasing constraints. Therefore, despite responding to the needs of a time-constrained two-point boundary value problem, one would expect operational thrust control laws to be significantly smoother, more in line with the TFC solution, and thereby yielding lower maximum instantaneous thrust requirements. Accordingly, the optimal thrust law that results from solving Equation (4.3) may be considered a better proxy for estimating the minimum thrust magnitude required to estimate a given uncertain orbital transfer since 1) it yields a proper lower bound for the required thrust level and 2) can devote *spare* thrust during the transfer to accommodate for uncertainty in the boundary conditions, and possibly dynamical characteristics.

Figure 4.9 depicts the statistics of the maximum acceleration requirements for the designed test scenarios. LEO and HEO sample transfers present similar magnitudes, with a single mode around $3 \cdot 10^{-5}\ m/s^2$ that would correspond to 30 mN for a $1,000$ kg satellite. GEO transfers, on the other hand, present lower average thrust requirements, yet still within the same order of magnitude. In fact, both GEO and HEO feature thicker right tails, and thus a higher number of cases with relatively high thrust requirements (on the order of 1 N for a 1-tonne spacecraft). The latter is attributed to the relatively short times of flight (lower bound of 3 hours) when compared to their orbital period, which severely affects the thrust de-

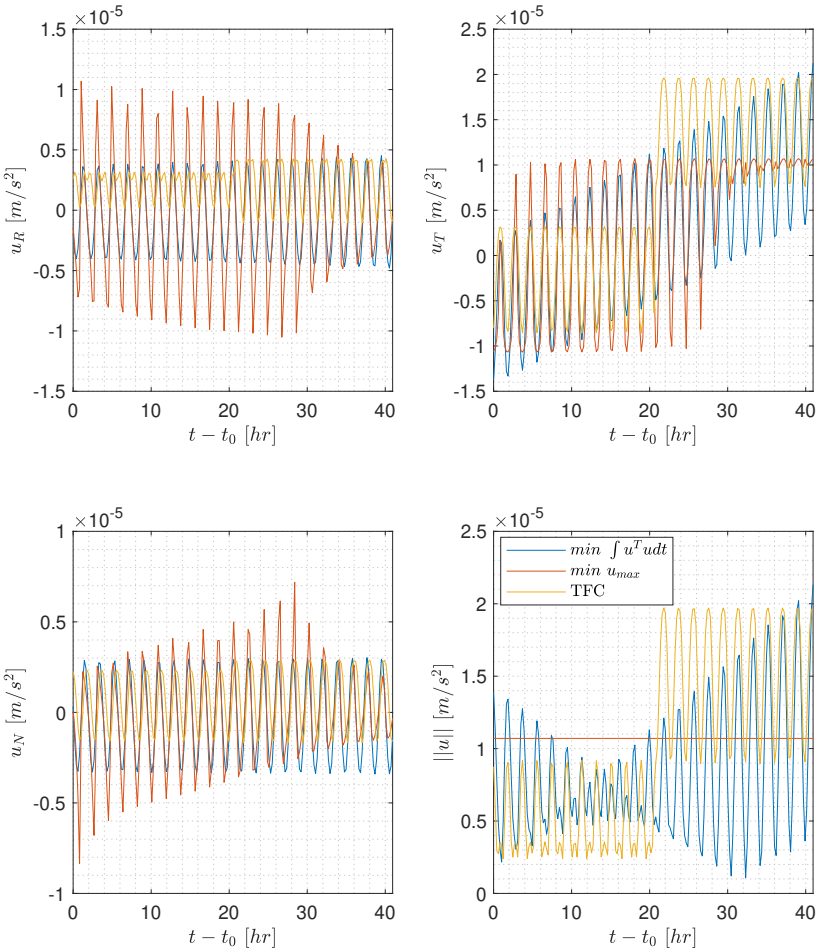

Figure 4.6: Example of derived control laws for a LEO transfer. The blue line corresponds to solving Equation (4.1), the red one to Equation (4.3) and the yellow line represents the TFC-parameterized law.

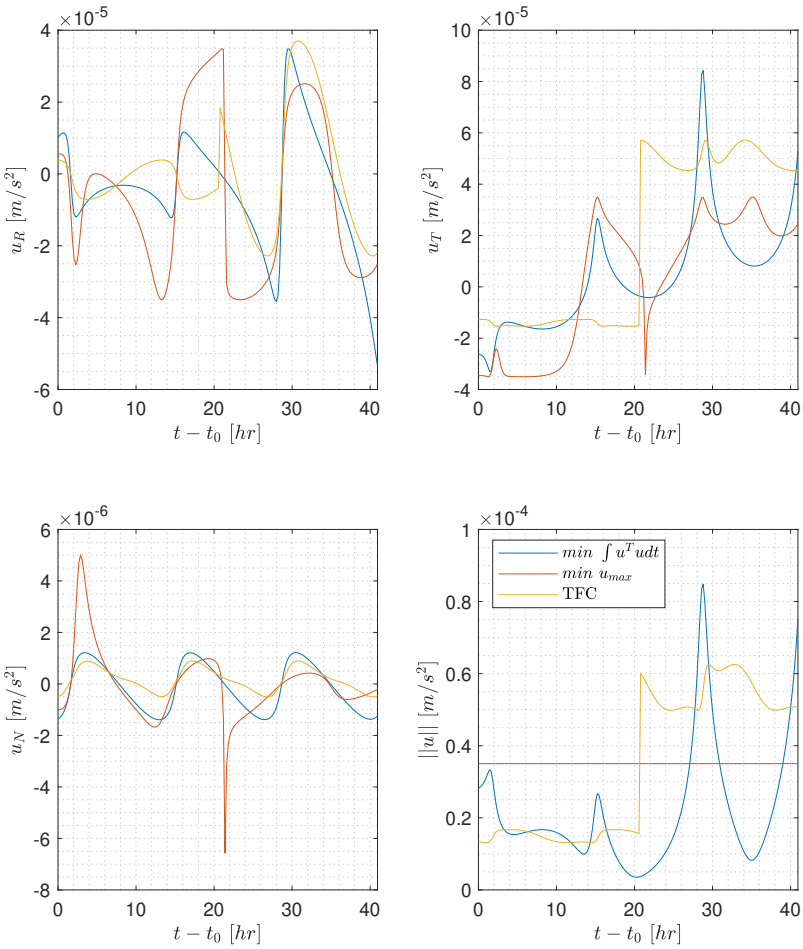

Figure 4.7: Example of derived control laws for a HEO transfer.

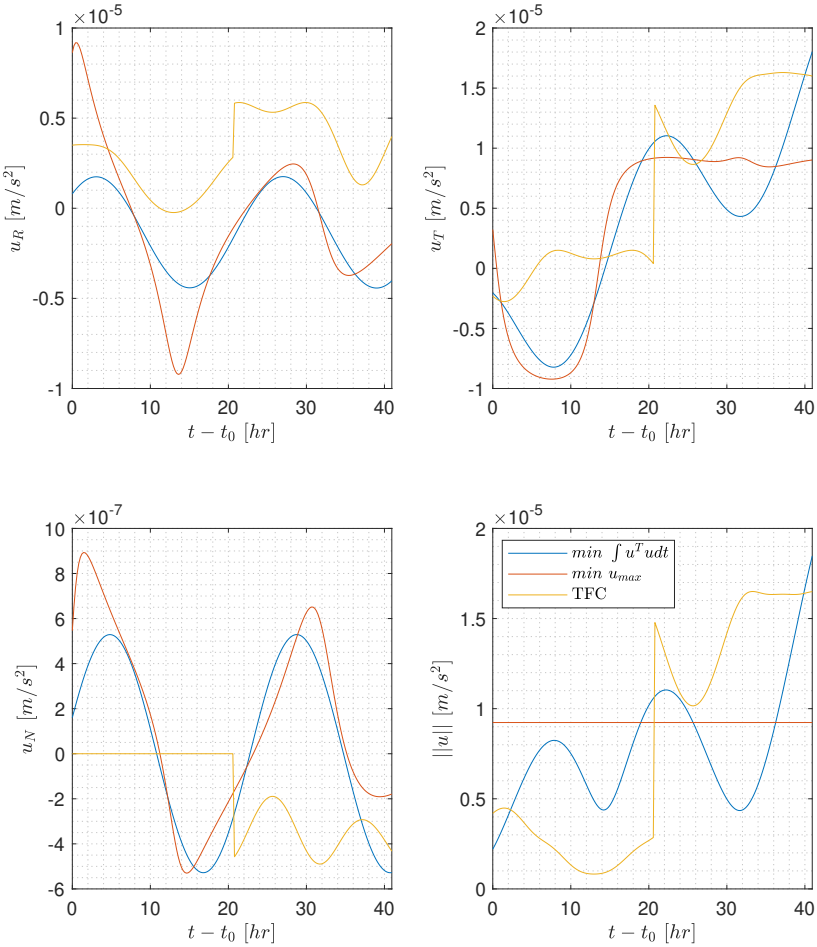

Figure 4.8: Example of derived control laws for a GEO transfer.

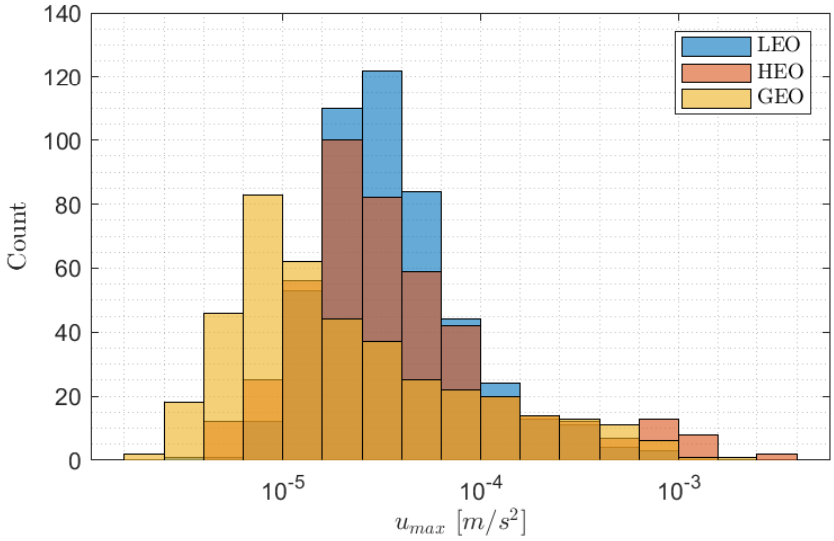

Figure 4.9: Distribution of minimum thrust requirements for each sample transfer.

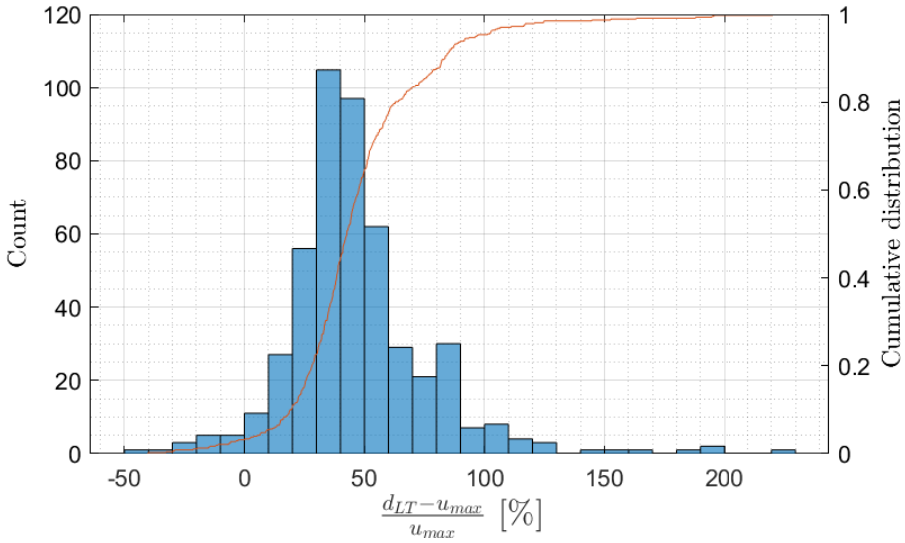

Figure 4.10: Thrust estimation error statistics for the LEO tests.

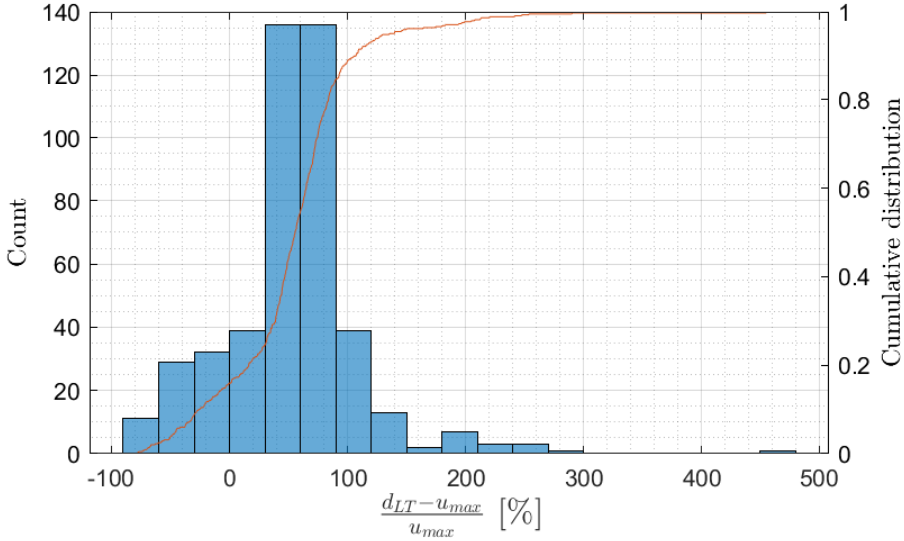

Figure 4.11: Thrust estimation error statistics for the HEO tests.

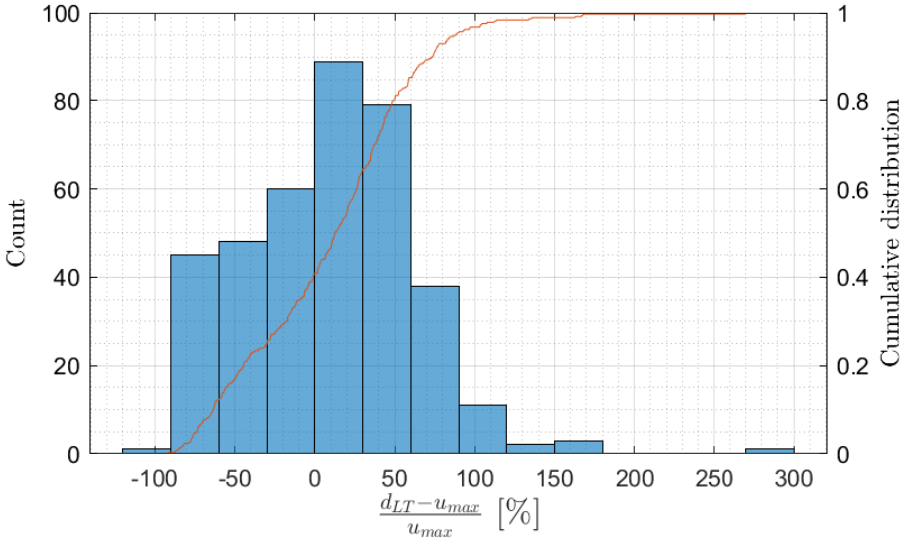

Figure 4.12: Thrust estimation error statistics for the GEO tests.

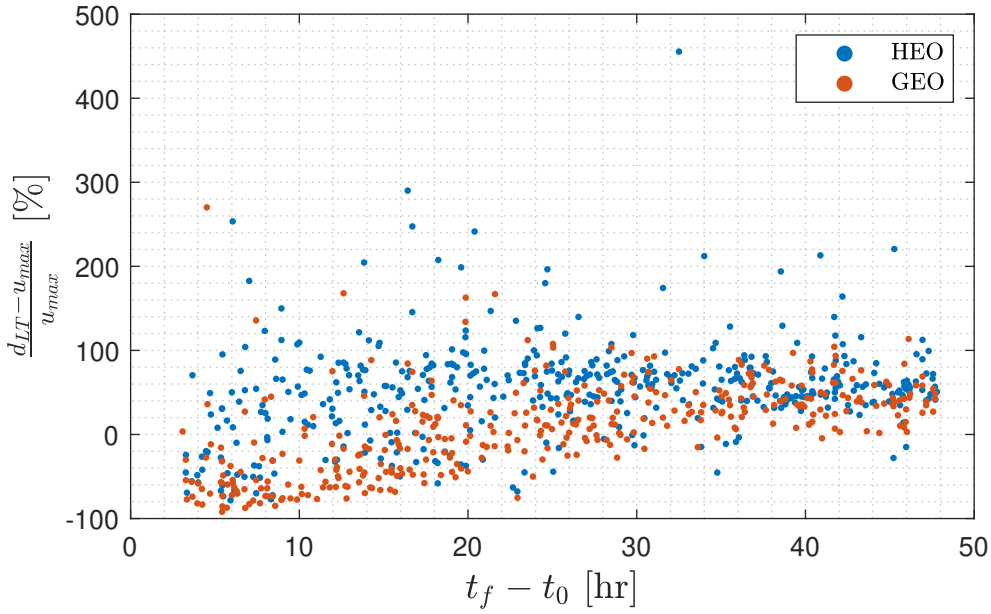

Figure 4.13: Thrust estimation error as a function of time of flight for HEO and GEO samples.

mands of plane-change maneuvers. Note in practice the observed orbital transfers may present long coasting arcs, yet ballistic dynamics are bound to fail in accurately approximating them. These cases would correspond to mild thrust requirements as derived by the proposed methods, in line with the majority of sampled transfers. Station keeping and collision avoidance maneuvers would ascribe to the previous characteristics, only extending for a fraction of the unobserved arc.

The estimation error statistics for the proposed control metric are summarized in Figure 4.10 for the LEO scenario. Therein, the TFC surrogate approach is compared against the solution to the NLP problem defined in Equation (4.3). As shown, the 95% confidence interval is lower bounded by an underestimation of the 10% and upper bounded by an overestimation around 100%, i.e. twice the reference value. Though these results might seem inaccurate, it should be noted the proposed approach is able to estimate the required thrust within a factor of two: $\Delta d_{LT} = 45 \pm 55\%$.

And it does so in a fraction of the computational time required to solve the complete non-linear problem, as indicated by Table 4.9. For the sake of comparison, these results are obtained from a rather inefficient implementation in MATLAB, but the *production*-ready C++ code presents a reduction of one order of magnitude in the computational runtimes. Results for HEO and GEO orbits, depicted in Figures 4.11 and 4.12, are less promising. Underestimations are more frequent for these test scenarios, and the variance of the control metric estimation error is much higher than for the former LEO samples. This erratic behavior is attributed to the assumptions made in the development of the surrogate approach: a periodic thrust profile and control input-averaged orbital dynamics. LEO orbits perform a significant number of full revolutions within the prescribed time spans[6], as opposed to HEO and GEO, and thus the above assumptions barely hold for these higher orbits. In fact, Figure 4.13 shows evidence of the strong correlation between the time of flight and the error committed in estimating the minimum thrust requirements.

In light of the above results, it can be concluded that the developed metric definition and proposed solution approach provide a convenient trade-off between accuracy and computational complexity. While it is true that its precision is somehow limited, in the worst case providing estimates within a factor of 2 (or $1/2$), the reduction in computational demand is unparalleled, on average 500 times faster than the high fidelity alternative. Moreover, the proposed approach can be safely computed in a fixed amount of time, which is barely affected by the time of flight and boundary conditions. NLP problems, on the other hand, cannot be solved in finite time as they require convergence to, at least, a local minimum. In fact, the number of orbital revolutions severely affects the required grid size, since otherwise a direct collocation method as the one employed may lack sufficient dynamical accuracy. Thus, the method developed throughout Section 3.1.3 has been shown to provide robust and efficient solutions for the low-thrust time-constrained boundary value problem, which can also be highly beneficial in providing a sound initial guess for the complete non-linear problem.

[6]Note these are obtained from expected SST system performances, this is, the target scenarios where the developed methods are to be applied.

4.2. Performance analysis and validation of reachability concepts

Section 3.2 comments on the increased tractability of reachability computations as enabled by the proposed efficient control metrics. Note these are approximative approaches to their respective complete problems so their performance and accuracy needs to be somehow tested. Results of the calibration and validation processes for the methods described in Section 3.2 are presented in the following, giving special emphasis to their robustness and computational efficiency, as opposed to accuracy, since they are to be implemented in a large-scale SST system. A positive evaluation of the developed methods is considered a strong indicator of tractability for reachability-based filters applied to the space domain, enabling the use of the maneuver transition densities proposed in Section 3.3.

4.2.1. Outer approximations of the reachable set

As discussed in Section 3.2.1, the outer boundary of the reachable set, conditioned either of the proposed metrics, can be efficiently approximated in the space of SGP4-averaged orbital elements (see Section 3.1.4). However, recall an alternative algorithm has been developed to find an outer bound approximation of the reachable set (see Section 3.2.2), providing a general solution for any kind of state space representation. In fact, the latter algorithm can be used as a benchmark to test the accuracy of the former linear reachability approach. In the following, a series of randomized tests are used to evaluate the performance of both approaches. The definition of these tests is borrowed from Section 4.1 for the sake of compactness, so that the ballistic state from which the reachable set is derived matches the statistics described previously and the control metric bound for each test sample is given by the control distance metric associated to the pre- and post-maneuver sample pair.

Table 4.10 collects the expected computational cost for each method and propulsion type, showing a two-order of magnitude increase in efficiency for the linear-based reachability computation methods. In fact, the linear method presents little variance with respect to its expected runtime as

	Linear approach	Derivative-free search
High-thrust	0.26 s	140 s
Low-thrust	0.05 s	8 s

Table 4.10: Average runtimes for approximating the axis aligned bounds of the reachable set in the orbital elements space following the linear surrogate method described in Section 3.2.1 and the derivative-free algorithm detailed in Section 3.2.2. These correspond to MATLAB implementations executed on an Intel Core i7-8750H CPU.

it only involves a naive iteration, this is, until the solution no longer varies for increasingly accurate reference states. On the other hand, the derivative-free approach potentially shows a much greater variance in practice, as each iteration of the n-dimensional pattern search is rather costly and the number of required iterations significantly depends on the correlations within the state space. Therefore, provided both approaches yield similar results, it is advisable to apply the former linearization-based method for reachability computations in the complete state space.

Comparison between the bounds derived using both methods are based on the following error metric

$$\Delta_L^N \text{œ}_i = \left(\frac{\text{œ}_{i,max}^{(N)} - \text{œ}_{i,min}^{(N)}}{\text{œ}_{i,max}^{(L)} - \text{œ}_{i,min}^{(L)}} - 1 \right) /2, \tag{4.4}$$

where the superscripts $^{(N)}$ and $^{(L)}$ represent the numerical (derivative-free) and linear approaches, respectively. This error metric thus compares the bounds estimated by each method, with negatives values indicating the linear approach overestimates the bounds or otherwise the numerical pattern search algorithm underestimates them. Results for the 1,500 individual test cases are depicted in Figure 4.14 under low-thrust assumptions and Figure 4.15 for impulsive maneuvers. In the former case, both methods largely agree for LEO and HEO orbits, with an error distribution centered at 0 and variances on the order of 10 % at most. However, GEO test cases show a lower level of compliance between the two approaches. The semi-major

axis bounds appear to be overestimated by the linear method in almost half of the test cases, which might be caused by the inability of the numerical method to unravel the strong correlation between the semi-major axis and true longitude. Similarly, differences in estimating the true longitude bounds are unexpectedly high, with the numerical method suggesting bounds that are more than three times greater that the linear method in some cases. This is attributed to the singularity presented by GEO orbits, whose eccentricity is negligible in practice, and the inability of the surrogate method to naturally deal with a fast rotating eccentricity vector (see Figure 4.23 and its discussion in Section 4.2.2). This suggests the bounds used should be those corresponding to the linear approach, as in fact it is more consistent with the formulation of the metric and does not rely on the convergence of a complex algorithm.

Results for high-thrust transfers are much more consistent, as indicated by Figure 4.15. Recall the compared methods use different formulations: linear reachability computations are based on the surrogate model discussed in Section 3.1.2, whereas the approach used everywhere else throughout the thesis follows an unperturbed central gravity model (see 3.1.1). Variances in the estimation differences are relatively small in the majority of cases, showing a considerable degree of compliance between both approaches. A higher inconsistency is found for highly eccentric orbits, where the linear model underestimates the reachable bounds in eccentricity and inclination and overestimates the semi-major axis for a number of cases ($\sim$ 40 %). These errors are below 15 % in relative magnitude except for the inclination vector components, h_2 and k_2, for which a 10% of test cases deviate by more than a 25%. While this is not a negligible amount, further analysis indicates some kind of over-estimation by the numerical approach. In fact, Figure 4.22 supports this statement, as all samples belonging to the reachable set are contained within the linear bounds with no signs of constriction at the inclination and right ascension of the ascending node (RAAN) limits. The results for GEO in terms of true longitude are much in line with those obtained for the low-thrust case and are, again, attributed to the singularity in the eccentricity. However, in this case the results given by the numerical approach can be trusted to a higher extent since these adopt a non-singular

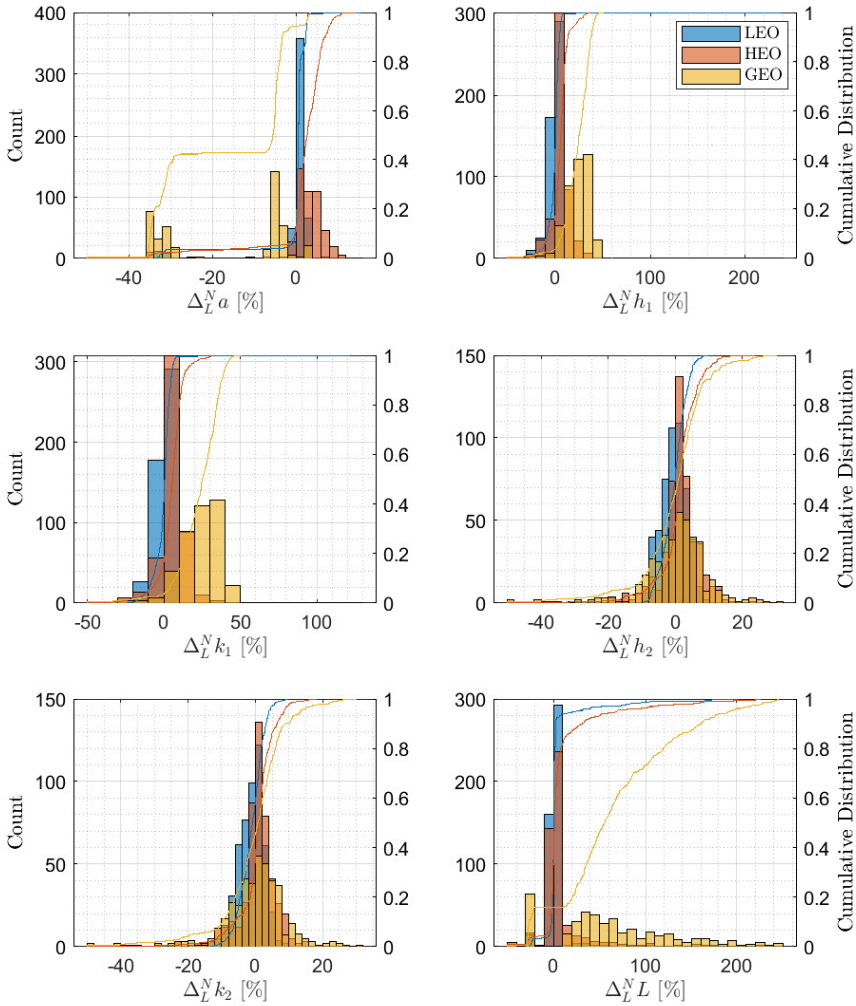

Figure 4.14: Comparison of reachable bounds for the low-thrust metric, showing the difference between the results from the linear method described in Section 3.2.1 and the numerical one in Section 3.2.2 in terms of the metric defined in Equation (4.4).

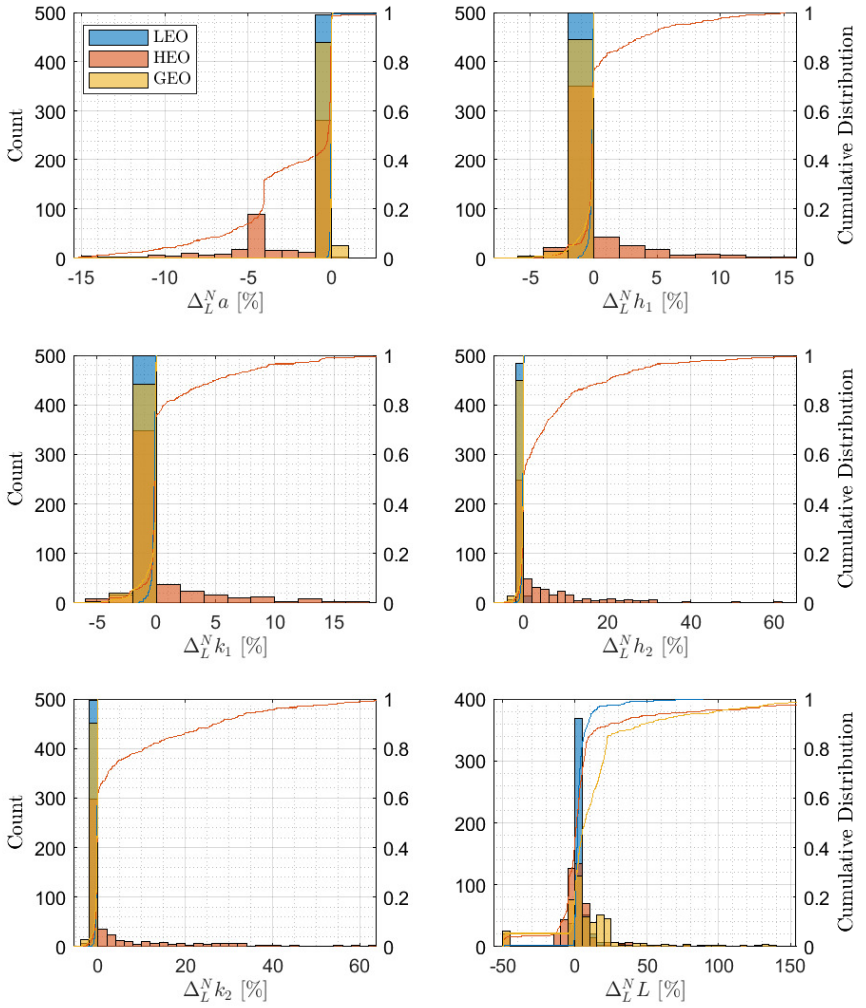

Figure 4.15: Comparison of reachable bounds for the impulsive metric, showing the difference between the results from the linear method described in Section 3.2.1 and the numerical one in Section 3.2.2 in terms of the metric defined in Equation (4.4).

formulation based on Cartesian elements, though convergence to the global maximum can still not be guaranteed in all the cases. The small number of cases where the linear approach underestimates the bounds may due to the numerical scheme not performing adequately, e.g. being stuck on a local minimum. On the other hand, positive $\Delta_L^N L$ values can be safely attributed to the linear approach. In fact, for cases with a relatively low number of orbital revolutions, i.e. GEO and HEO, upper and lower bounds on the semi-major axis do not necessarily coincide with those in the phase angle. Limit values in a require perigee burns, whereas maximal variations in the angular variable are directly proportional to the time elapsed since a change in orbital period. Therefore, for time-constrained transfers, limits on a and L do not necessarily coincide as the latter require a trade-off between the time of flight and the change in orbital period. Though these results indicate further research needs to be carried out to identify the proper optimal burn policy for approximating the bounds in true longitude and also isolate the source of inconsistencies for low-thrust transfers, it is still reasonable to employ these linear methods for efficient reachability computations as their accuracy is deemed sufficient for target tracking purposes.

4.2.2. Volumetric reachability computations

Coherent maneuver modes for multi-object filtering necessitate proper transition densities, as this is paramount to provide comparable measurement likelihoods among different targets. Though one can elaborate arbitrarily complex maneuver models based on the proposed metrics, obtaining the normalization constant of their corresponding transition densities may be challenging. Section 3.2.3 presents a method to approximate the volume of the reachable set in an arbitrary state space representation. In fact, it can be easily tuned to perform other types of volumetric integrals, such as those required to obtain the normalization constants of the heuristics-based maneuver model described in Section 3.3.2, though convergence for each particular application may need to be studied on a case by case basis. In the following, a collection of individual test cases are selected at random from the reachability scenarios discussed in Section 4.2. For each indi-

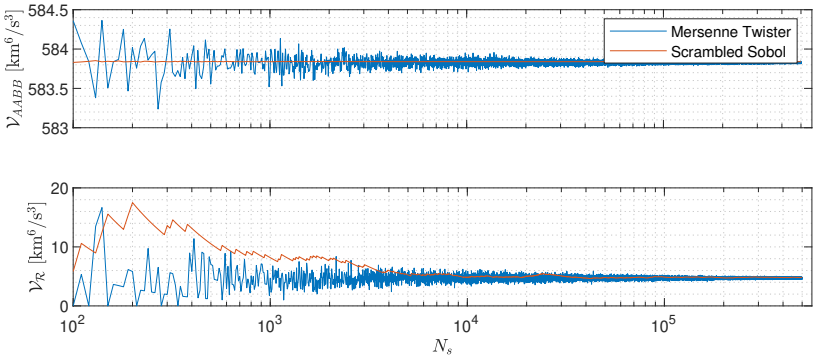

Figure 4.16: Convergence in volumetric reachability computations for a GEO target using the low-thrust metric. The upper plot shows the volume of the AABB approximation in Cartesian coordinates, whereas the bottom one corresponds to the volume of the reachable set.

vidual test case, a convergence analysis is run on the number of samples used in the approximation given by Equation (3.61). These samples, drawn from the SGP4-averaged modified equinoctial element bounds obtained via the linear reachability method described in Section 3.2.1, are evaluated for containment within the reachable set via the reachability inclusion function (see Equation 3.63). Two different sampling methods are used: a Mersenne Twister [151] pseudo-random number generator, the default algorithm implemented by MATLAB software, and the qMC scrambled Sobol sequence proposed in [141]. As discussed in Section 3.2.3, the latter method is expected to present improved convergence properties for high-dimensional and relatively complex integrals like the one at hand.

Results of the convergence analysis for the different test cases can be consulted in Figures 4.16 to 4.21. Under all circumstances, the qMC method based on scrambled Sobol sampling outperforms the default Mersenne Twister sampling algorithm, showing much greater numerical stability for volumetric integral calculations. In fact, the volume of the AABB in Cartesian coordinates is estimated to significant accuracy with less than 300 samples, mainly due to its lattice characteristics. This feature renders qMC smoother for sequential approximations, especially compared to the highly oscillating estimations given by its alternative. However, such smooth behavior is

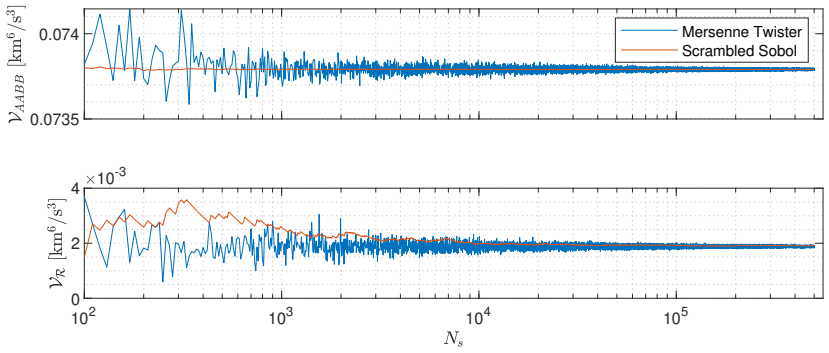

Figure 4.17: Convergence in volumetric reachability computations for a HEO target using the low-thrust metric. The upper plot shows the volume of the AABB approximation in Cartesian coordinates, whereas the bottom one corresponds to the volume of the reachable set.

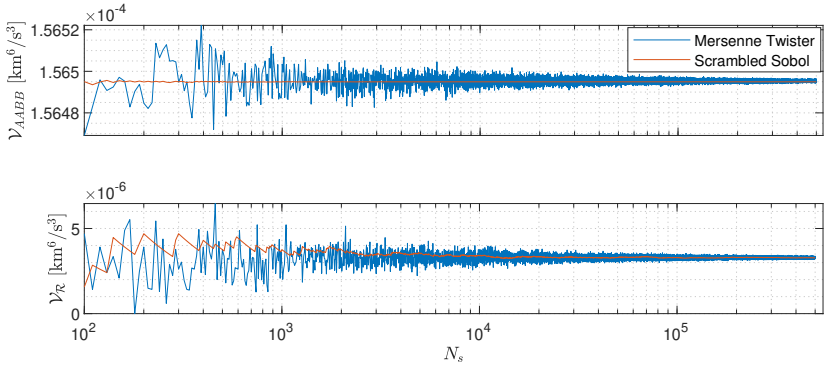

Figure 4.18: Convergence in volumetric reachability computations for a LEO target using the low-thrust metric. The upper plot shows the volume of the AABB approximation in Cartesian coordinates, whereas the bottom one corresponds to the volume of the reachable set.

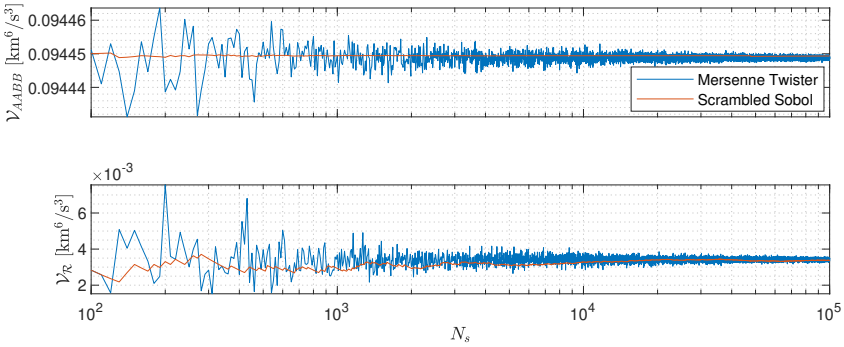

Figure 4.19: Convergence in volumetric reachability computations for a GEO target using the impulsive metric. The upper plot shows the volume of the AABB approximation in Cartesian coordinates, whereas the bottom one corresponds to the volume of the reachable set.

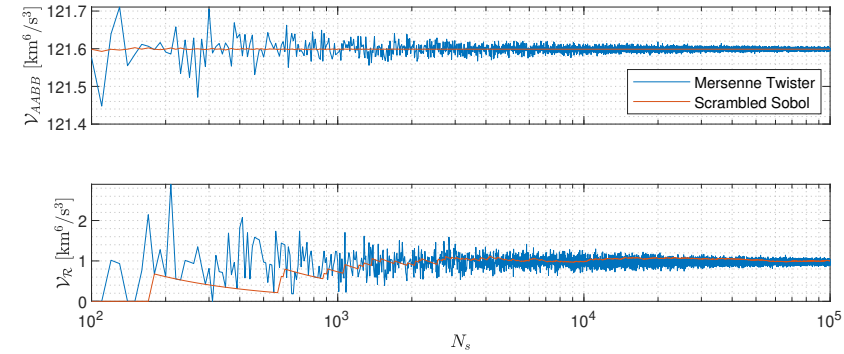

Figure 4.20: Convergence in volumetric reachability computations for a HEO target using the impulsive metric. The upper plot shows the volume of the AABB approximation in Cartesian coordinates, whereas the bottom one corresponds to the volume of the reachable set.

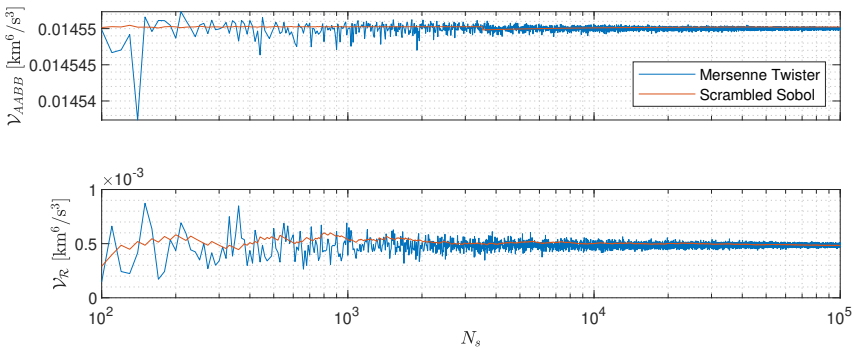

Figure 4.21: Convergence in volumetric reachability computations for a LEO target using the impulsive metric. The upper plot shows the volume of the AABB approximation in Cartesian coordinates, whereas the bottom one corresponds to the volume of the reachable set.

not necessarily desirable, as shown in Figures 4.16, 4.17 and 4.20, where the qMC estimates initially present a considerable bias when approximating $\mathcal{V}_{\mathcal{R}}$ that slowly decreases towards reasonable values before stabilizing around $3,000$ samples. While it is true that Mersenne Twister approximations are erratic, their average shall in principle converge to the true value, so it remains to be tested whether there exists a combination of tests (i.e. 3 runs of 1,000 samples or 30 runs of 500 samples) that presents a higher convergence rate than the quasi-deterministic sequence. Alternatively, it might be of interest to discover some sort of mathematical relation between the volume of the bounding box and the reachable set, since this would reduce the computational complexity even further. In principle, one would expect $\mathcal{V}_{\mathcal{R}}$ to be a consistent fraction of $\mathcal{V}_{AABB}$ and, due to correlations among different dimensions, to present some kind of ellipsoidal shape. Under this hypothesis, the volumetric ratio could be vaguely approximated as

$$\frac{\mathcal{V}_{nE}}{\mathcal{V}_{nR}} = \frac{2}{n}\frac{\pi^{n/2}}{\Gamma(n/2)}\prod_{i=1}^{n}\frac{1}{2}, \tag{4.5}$$

where $\Gamma(\cdot)$ is the Gamma function and $\mathcal{V}_{nE}/\mathcal{V}_{nR}$ is the fraction between an n-dimensional axis aligned ellipsoid and and n-dimensional hyperrectangle. For the case of $n = 6$, the above expression takes the value $\mathcal{V}_{nE}/\mathcal{V}_{nR} =$

0.0807. For the sake of comparison, Table 4.11 gathers the computed volumetric ratio for the different test cases considered, which are for the most part on the same order of magnitude, except for the low-thrust GEO and impulsive HEO cases. Note in all cases the estimated ratio is lower than the hypothesized geometric one, suggesting that the reachable set may resemble a tilted, and thus not axis aligned, ellipse. This is illustrated in Figures 4.22 to 4.25, which depict samples belonging to the reachable set projected in the a-e, a-I, a-L and a-Ω planes. Figure 4.22 partially explains the relatively low volumetric ratio for the HEO impulsive case, as the bounds in semi-major axis, eccentricity and right ascension of the ascending node seem to be overestimated. Still, there is a strong correlation between eccentricity and semi-major axis, mostly caused by the relatively high eccentricity magnitude of the orbit as this is not observed in the LEO case depicted in Figure 4.24. Note both present strong correlations between the semi-major axis and true longitude, since, as discussed in Sections 3.1.3 and 3.2.1, a is the most efficient driver for changes in the orbital phase. Figure 4.23 depicts the projected bounds of the reachable set for a GEO orbit under low-thrust control, whose volumetric ratio with respect to the sampling bounds used was rather low compared to the rest of cases. Though not appreciable from the plot, the latter is attributed to an overestimation of the eccentricity and Ω bounds, which in turn yield a rather inefficient sampling. However, note true longitude bounds seem to be artificially capped by the linear method when compared to the numerical one. In fact, this is due to the relatively low eccentricity of the baseline orbit, which renders the linear surrogate model inaccurate under certain circumstances (see Hudson [123]), as opposed to the more reliable phase change approach followed in the linear approach. The expected shape of the projected bounds of the reachable set under low-thrust assumptions resemble those illustrated in Figure 4.25, with strong linear correlations between the semi-major axis, eccentricity and true longitude. Therein, one can also identify a quadratic relationship between in-plane and out-of-plane orbital parameters, as these are related through the norm of the applied control in spite of the dynamics. Recall from the results included in Section 4.2.1 that, for a reduced number of cases, bounds estimated by the numerical derivative-free approach were

	GEO	HEO	LEO
High-thrust	0.03581	0.00826	0.03322
Low-thrust	0.00823	0.0257	0.0205

Table 4.11: Ratio between the volume of the reachable set and that of the AABB for different test cases, considering converged values.

	GEO	HEO	HEO
High-thrust	111 s	155 s	173 s
Low-thrust	112 s	151 s	99 s

Table 4.12: Average runtimes for approximating the volume of the reachable set in the orbital elements space following the approach described in Section 3.2.3 using 3,000 samples. The code is implemented in MATLAB and executed on an Intel Core i7-8750H CPU.

higher than those estimated by the linear method, though in principle the former should be more accurate. Figure 4.22 in fact corresponds to one of those cases, suggesting that the linear approach may be more precise when estimating the reachability bounds and perhaps shall be used as basis for the sampling since samples belonging to the reachable set do not seem to concentrate around the linear bounds.

Table 4.12 gathers the expected computational cost of these volumetric computations for a sufficiently accurate approximation using the randomized lattice (qMC) approach, this is, using 3,000 samples. Note these runtimes can be reduced by around one order of magnitude if implemented in pre-compiled code such as C++. In fact, since the scrambled Sobol sequence is almost deterministic, it is reasonable to monitor convergence past a certain minimum (representative) number of samples, seeking to further reduce the computational requirements while targeting some accuracy in the estimate.

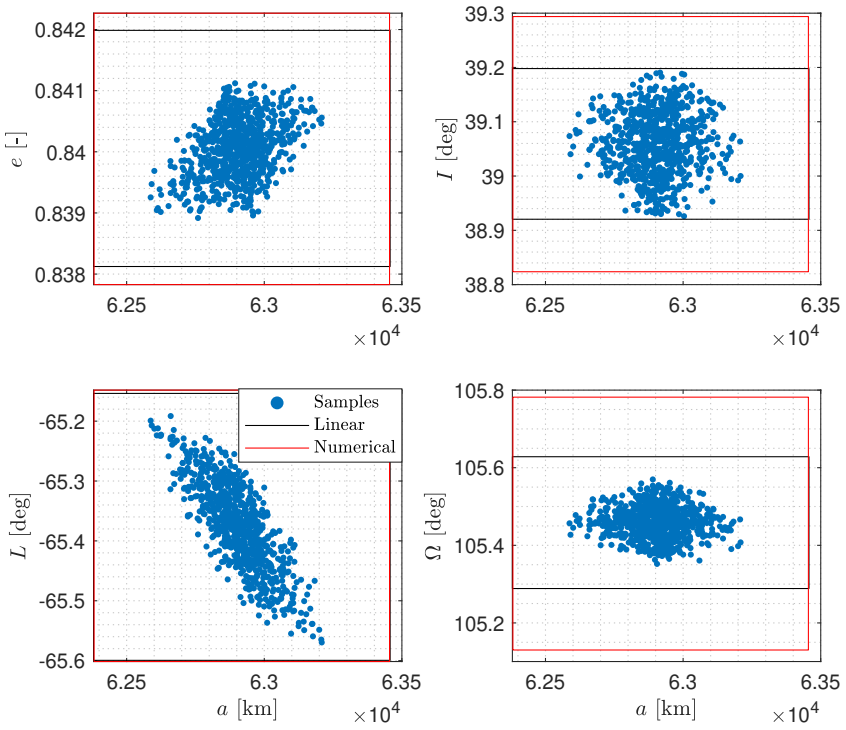

Figure 4.22: Two-dimentional projections of samples belonging to the reachable set for a high-thrust HEO target. Boxes indicate the maximum and minimum magnitudes along each dimension as given by the linear (red) and numerical (black) approaches.

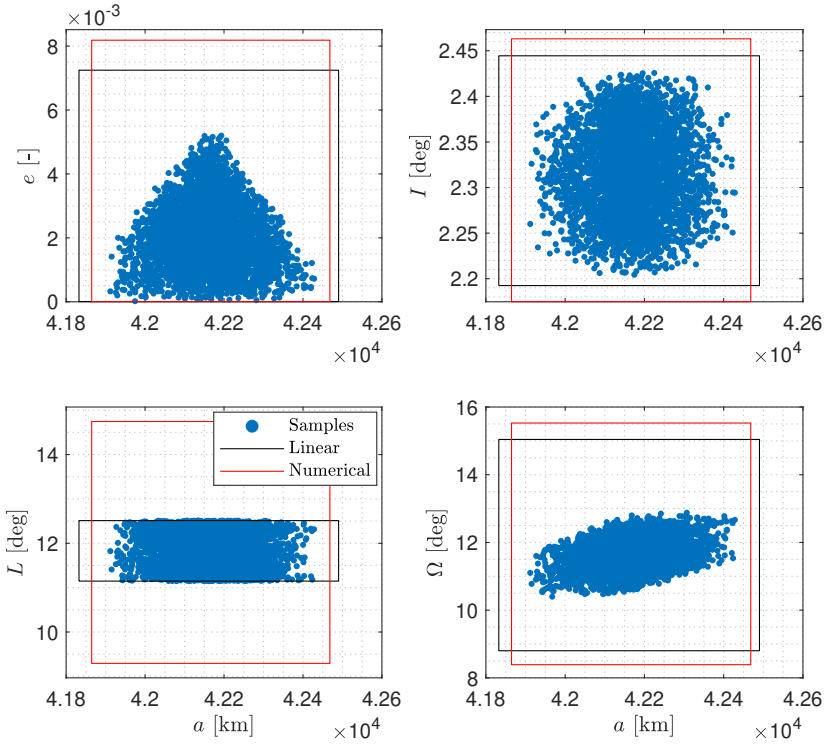

Figure 4.23: Two-dimentional projections of samples belonging to the reachable set for a low-thrust GEO target. Boxes indicate the maximum and minimum magnitudes along each dimension as given by the linear (red) and numerical (black) approaches.

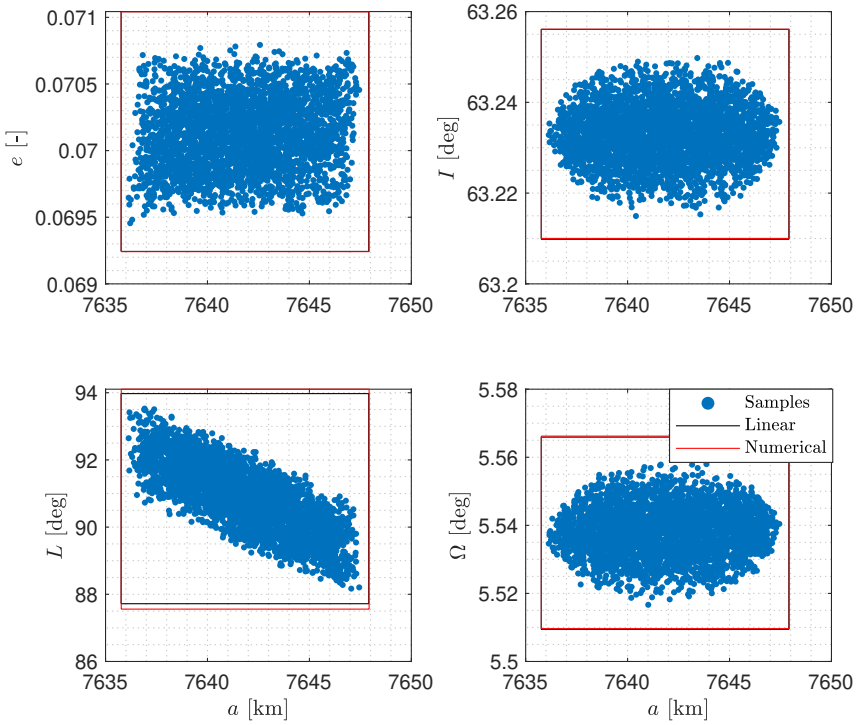

Figure 4.24: Two-dimentional projections of samples belonging to the reachable set for a high-thrust LEO target. Boxes indicate the maximum and minimum magnitudes along each dimension as given by the linear (red) and numerical (black) approaches.

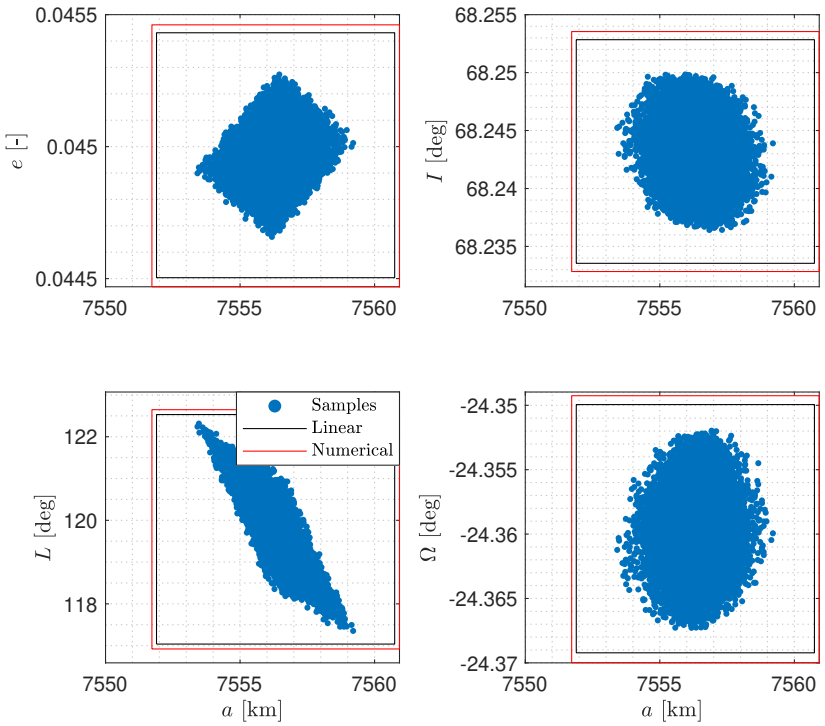

Figure 4.25: Two-dimentional projections of samples belonging to the reachable set for a low-thrust LEO target. Boxes indicate the maximum and minimum magnitudes along each dimension as given by the linear (red) and numerical (black) approaches.

4.3. State space filtering of active spacecraft

Prior to evaluating the performance of the maneuver modes developed in Section 3.3 in a multi-object framework, it is convenient to analyze their implications with regard to filtering. This can well be performed in a single-target scenario, where measurement to object correlations may be simply based on certain statistical threshold within the measurement space. This simplified problem can help to identify the requirements of a capable algorithm for sequential state space filtering applied to space surveillance, with special emphasis given to the ability of the proposed maneuver priors to approximate the true posterior distribution conditioned on a measurement sequence. Moreover, it simplifies the benchmarking process of different filter types, thus easing the evaluation of alternatives to provide a trade-off between accuracy, robustness and computational cost. In the remainder of this section, two representative test scenarios for space surveillance are presented, each with different observability and orbital characteristics. The former, described in Section 4.3.1, deals with a geosynchronous satellite performing high-thrust station keeping maneuvers. Observations are gathered from ground-based electro-optical sensors and the proposed maneuver modes (see Section 3.3) are evaluated in different filtering schemes: moving horizon estimators and particle filters. Later, Section 4.3.2 presents a surveillance scenario wherein a LEO spacecraft equipped with electric propulsion performs orbit raising and orbit acquisition maneuvers. In this case, only the vacuous model described in Section 3.3.1 was employed since the target is continuously maneuvering, thus hindering the application of smoothing recursions, which are needed for maneuver characterization. Moreover, only a single filter implementation based on sequential Monte Carlo is evaluated, mainly because batch processing was complicated by the long duration and relatively high magnitude of the orbital maneuvers.

4.3.1. Optical surveillance data for a high-thrust target in GEO[7]

Spacecraft orbiting within the geosynchronous belt are subject to several perturbations that force them away from their operational orbital slots. There-

[7] The results and test description gathered in this section have been previously published in the journal Advances in Space Research [152].

fore, they are required to perform frequent orbital correction or station-keeping maneuvers. Though these usually present periodic patterns, in some cases it remains a challenge to automatically detect and model them within the observation-correlation process. The following paragraphs describe a surveillance scenario where measurements are line of sight observations gathered from a ground-based electro-optical sensor:

- The target is a GEO satellite whose operational orbital slot is restricted to the mean longitude band $l = -4.8 \pm 0.2$ deg and the inclination band $I = 2 \pm 0.05$ deg. Equipped with chemical propulsion, the object is simulated for 401 days spanning from March 2009 to May 2010 using the following dynamical model:
  - Non-spherical Earth model of degree and order 70.
  - Sun, moon and planetary third-body perturbations.
  - SRP model with conical solar and lunar eclipses, where the target is assumed a perfect sphere.
  - To account for dynamical mismodelling, variations in the reflectivity coefficient c_R are modeled as a Poisson random variable with parameter $\lambda = 7$ days and magnitude $\Delta c_R \sim \mathcal{N}(0, 10^{-2})$.
  - Solid Earth and ocean tides.
  - General Relativity.
- Simulated observations are gathered from two electro-optical sensors located at Zimmerwald (AIUB Zimmerwald's Observatory) and Tenerife (ESA Optical Ground Station), according to the following characteristics:
  - Minimum elevation above the local horizon of 20 deg.
  - Solar phase angle in the range $[0, 90]$ deg.
  - Angular distance to Earth shadow greater than 0 deg.
  - Observables are gathered as a series of right ascension and declination pairs, commonly referred to as optical tracklet, with

zero-mean Gaussian noise and a standard deviation of 1 arcsecond.

- On average, each sensor detects the target every two days.
- The length of acquired tracklets is distributed as $\mathcal{U}(2, 10)$ min.
- To ease computations, tracklets are pre-processed to obtain their *attributables* following the approach in [23]. Accordingly, the estimated observation covariance is derived from a second order fit on the right ascension-declination pair sequence, retaining the angles and their rates as observations and discarding the second order components of the fit.

362 tracklets are gathered from the simulation timeframe, wherein the target performs 20 detectable maneuvers that consist of 29 detectable impulsive burns. From these, 4 correspond to changes in inclination or north-south station-keeping maneuvers (single-burn), 15 to longitude corrections or east-west station-keeping maneuvers, and the remaining one is a combined east-west nort-south correction. Note for a maneuver to be detectable on its own, it needs to be bounded between subsequent tracklets, and hence it is possible to find combined maneuvers or observe individual impulses from a common east-west maneuver.

This data set is processed via different filter implementations, aimed at analyzing the performance and capabilities of the maneuver models and control metrics developed in the context of this thesis. In the light of a fair comparison, all the methods present the same dynamical modelling and data association characteristics. To allow for certain model mismatch between the simulated and predicted dynamics, the state prediction is computed using the following perturbations:

- Non-spherical Earth of degree and order 10.
- Third-body perturbations of Sun and Moon.
- Cannonball model for the SRP with a conical solar and lunar eclipse model, using fraction of illumination for penumbra regions.

- Estimation is performed in the space of Cartesian position and velocity coordinates, augmented with the reflectivity coefficient.

Despite not being representative of a multi-target scenario, it is convenient to allow for miscorrelation, which may stem from poor post-maneuver state estimates. In this case, measurement association is based on the Mahalanobis distance

$$d^2_{M,k} = (z_k - h(x_k))^T \Sigma^{-1}_{R,k} (z_k - h(x_k)), \tag{4.6}$$

assumed to follow a χ^2 distribution with 4 degrees of freedom. Under the hypothesis that the observation corresponds to the target following ballistic motion, the above distance is expected to fall within the 3-σ bound of its corresponding distribution. The maximum expected control effort for this particular application has been set to $\lceil d_I \rceil = 10$ m/s to account for the relatively high cost of the north-south station keeping maneuvers, on the order of 5 m/s. The baseline filter scheme, considered an operational standard, is a moving horizon estimator (see Section 2.2.2), hereafter referred to as MHE, that sequentially processes observations in batches up to 6 consecutive tracklets. Whenever the ballistic model fails to meet the above criteria for data association, a maneuver hypothesis is evaluated. If (and only if) the reachable set and the support of the measurement distribution intersect in the measurement space, the filter estimates a post-maneuver state. Such estimation is conditioned on the sequence of observations after the detected maneuver (no-association), and the loss function of the least-squares problem is augmented to include the expected control distance $\hat{d}_I$ to the previous orbit. As a benchmark, the same filtering framework is applied to the true maneuver sequence, relieving the filter from maneuver detection. The latter method is labeled as MHE II in the following, and provides the maximum a posteriori estimate in a Bayesian sense when no heuristic or prior maneuver information is available. Note these implementations assume a Gaussian state distribution and the estimated covariance is only a function of the local measurement partials and the observation residuals [55].

As an alternative to the previous deterministic, linearization based ap-

proach, this work implements a regularized particle filter (see Section 2.2.3) with a variable population size, termed SHF, wherein different maneuver hypotheses can co-exist. Each individual hypothesis is represented by a subpopulation consisting of $N_s = 1000$ particles, which are individually tested for association. Post-maneuver state estimation for the general maneuver model proposed in Section 3.3.1 is based on MCMC sampling via the DREAM algorithm (see Section 2.2.4), where the prior distribution is defined as the intersection between the support of the tracklet attributable and the reachable set transformed to the measurement (or attributable) space. An alternative implementation, SHF II, makes use of the history of detected maneuvers to elaborate a heuristic maneuver model, as detailed in Section 3.3.2. Therefore, heuristics are automatically generated based on already detected and consolidated maneuvers, and again estimation is based on a MCMC move that is further conditioned on said heuristic models. Both SHF implementations feature backwards simulation particle smoothing [153] not only to improve sequential state estimation accuracy but also to properly isolate the maneuver epochs, as this is paramount to ensure a proper definition of the maneuver heuristics.

Maneuver detection results for the MHE and SHF II filters can be consulted in Table 4.13. The latter shows a higher sensibility as it was able to properly locate the different orbital maneuvers. Nonetheless, the MHE filter also identified all the different maneuvers, yet in four occasions the detection was delayed by one observation. It shall be noted that the impact of individual east-west station keeping burns in the orbital state is so small that neither of the methods are capable of identifying isolated east-west impulses but rather their combined effect. Due to its higher sensitivity, the SHF II filter triggered two false detections. These are attributed to an overestimation of the eccentricity, which implies inaccuracies in the mean longitude for an east-west maneuver and an underestimated inclination magnitude in the case of a north-south maneuver. Note since the MHE directly performs orbital updates as long as the estimates are sufficiently close within the measurement space, no false detections have been reported.

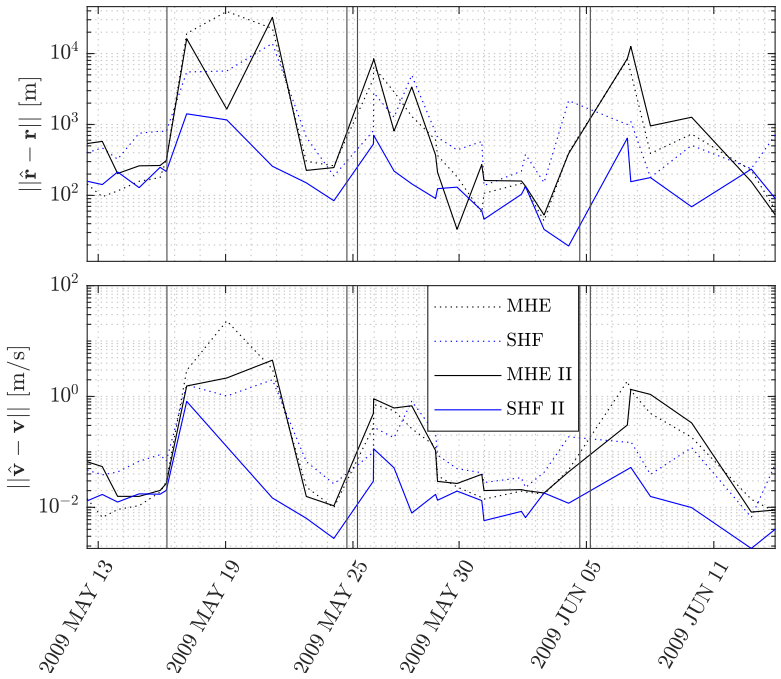

Figure 4.26: Position (top) and velocity (bottom) estimation errors for a sequence of north-south, east-west, east-west station keeping maneuvers.

This comes at the cost of state estimation performance, since even when knowing the true maneuver sequence, MHE II fails to improve on the state estimates given by the particle filter implementations. The latter is exemplified in Figure 4.26, where MHE, MHE II and the (uninformed) SHF show difficulties when recovering the post-maneuver orbits. Not only their estimates are relatively inaccurate, with errors on the order of 10 km and 1 m/s, but they require two to three additional observations to decrease them to within usual ballistic magnitudes. On the contrary, the heuristics-informed filter, SHF II, presents a faster convergence to the reference orbit, which is attributed to the lower magnitude of the initial post-maneuver state errors, in this case on the order of 1 km and 0.5 m/s.

A comprehensive analysis of the estimation performance of the different filters can be drawn from Figure 4.27, which shows the root mean squared error (RMSE) as a function of the number of tracks elapsed after every

Maneuvers - Burns	Correct	Delayed	False
MHE	16 - 16	4 - 4	0 - 0
SHF II	20 - 21	0 - 0	2 - 2

Table 4.13: Maneuver detection accuracy shown by the moving horizon estimator (MHE) and the heuristics-based stochastic hybrid filter (SHF II). The total number of detectable maneuvers is 20, whereas the number of detectable burns is 29.

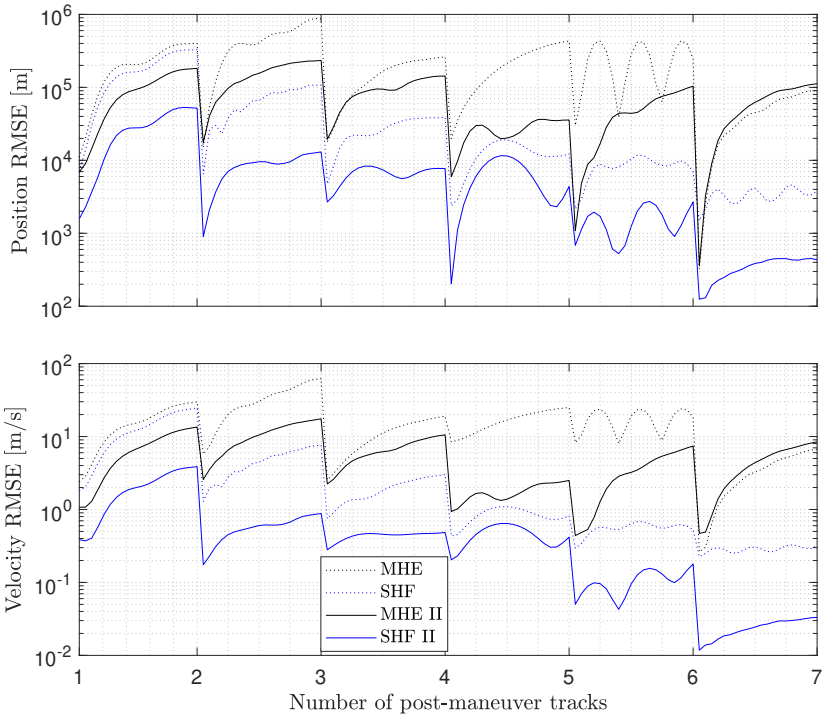

Figure 4.27: Position (top) and velocity (bottom) root mean squared error (RMSE) as a function of elapsed tracks after each maneuver. The true (and not detected) maneuver sequence is used as a reference. Note the results are averaged over a total of 20 maneuvers

Method	MHE	MHE II	SHF	SHF II
Comp. Time [hrs.]	2.7	3.5	8.7	10.7

Table 4.14: Runtime comparison. These values are obtained for a MATLAB implementation on an Intel Core i7-8750H laptop CPU.

maneuver, i.e.

$$RMSE_x(n_T) = \sqrt{\frac{\sum_{j=1}^{N_{s,i}}(x_{i,j} - x_i^*)^2}{N_{s,i}}}, \quad \forall i : m_{i-n_T:i} \neq 0, \tag{4.7}$$

where x_i^* represents the reference or true state at the corresponding epoch. These results are in line with those already shown in Figure 4.26, with SHF being on par or slightly more accurate than MHE II and MHE showing the poorest performance. This might be caused by the presence of multiple local minima in the least squares minimization problem solved in the MHE filters, or it may occur that the global minima happen to be relatively far from the true post-maneuver state. In any case, as more tracklets become available, these batch filters are able to converge towards the true state. The regularized particle filters mitigate this problem by allowing for multi-modal state distributions, potentially capturing different local minima and hopefully providing support over the true state. This seems to provide a more accurate state characterization within the first two to five post-maneuver tracks, since local minima that do not comply with the observation sequence are not evolved but pruned during the filtering process. Moreover, the inclusion of maneuver heuristics has proved to be beneficial in early post-maneuver estimation as the prior density is less vacuous and concentrated around regions closer to the true solution.

The statistics of the characterized maneuvers are depicted in Figure 4.28 as blue ellipses, whereas the true maneuver features are indicated by black circles. Note every true maneuver is enclosed by a blue ellipse, indicating that the estimated maneuvers correlate well with the true ones at least in terms of the orbital elements depicted therein. Characterized maneuvers present different levels of confidence or variance, which is indicated by the

size of the ellipses themselves. Two particular north-south station keeping characterizations present a high variance in terms of eccentricity and, in fact, one of these corresponds to the false maneuver detection in the top right corner of the l^- vs I^- plot. This is attributed to the relatively low cost of eccentricity variations compared to changes in inclination and, in particular, to the fact that under certain observability conditions (if not observed at the antinodes or apices of the orbit) orbits with higher eccentricity and lower inclination leave traces similar to more inclined circular orbits. Note eccentricity variances across east-west maneuvers are more consistent, yet semi-major axis variations are more spread. Most Δa estimates, especially those greater than zero, present a significant standard deviation, on the order of 1 km. Due to the relatively high distance between the sensors and the target, information regarding the semi-major axis is conveyed in the time stamp of the observations, i.e. information regarding the period of an orbit is encoded in the time elapsed between subsequent observations. Therefore, uncertainty in the semi-major axis and eccentricity remains high until there is a sufficient temporal separation between tracklets. It is thus crucial to perform some kind of smoothing recursion to improve subsequent state estimates and provide accurate maneuver characterizations.

Table 4.14 gathers the computational cost of the different filter implementations, where SHF filters present up to four times the computational complexity of MHE. Part of this is attributed to the sampled nature of the former filters, as they require to propagate a considerable amount of state realizations. However, certain improvements could be attained in this regard by resorting to more efficient sample propagation methods, such as multi-fidelity propagation [154]. Other alternative might consider the use of Gaussian sum filter implementations [155], for which the required number of individual samples or components is expected to be considerably lower. Nevertheless, this test has compared the solutions expected from purely deterministic and operational algorithms, such as the MHE, and advanced non-linear filtering methods such as the RPF with MCMC post-maneuver sampling. The latter has shown to be more robust and accurate at a higher computational cost, and it is left for future work to study alternative schemes that provide a better trade-off in terms of computational cost

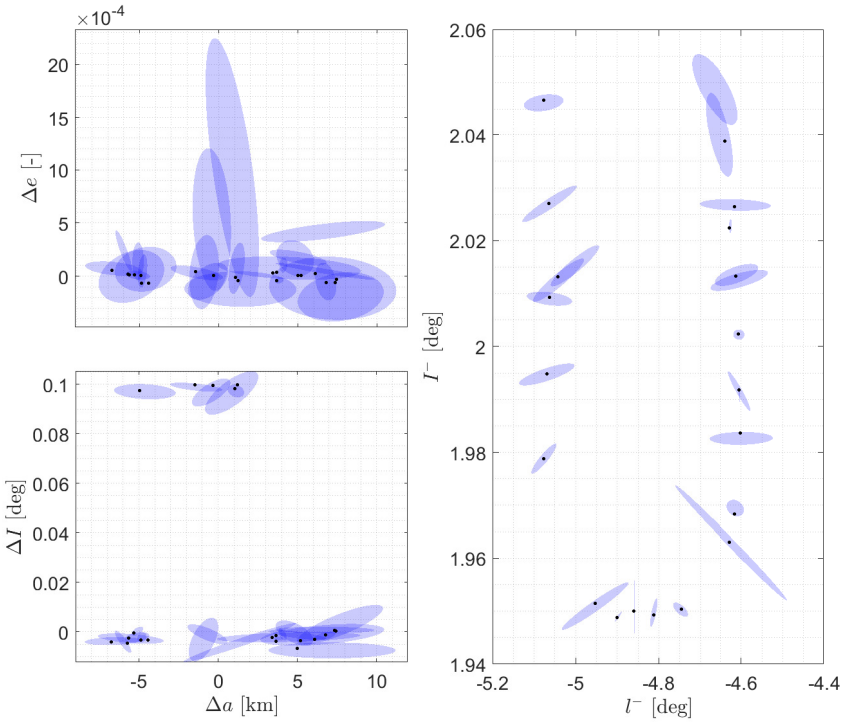

Figure 4.28: Maneuver characterization in terms of relative change in orbital elements (left) and pre-maneuver inclination and mean longitude (right). Blue ellipses are 3-σ co-variance contours for the characterized maneuvers, while black dots correspond to true maneuvers.

and state estimation accuracy.

4.3.2. Radar surveillance data for a low-thrust target in LEO[8]

The efficiency of electric propulsion allows for relatively inexpensive high ΔV cost maneuvers, as propellant and thus mass consumption is significantly reduced when compared to chemical engines. However, their low thrust magnitude require longer maneuver arcs that can span from days to weeks and even months. That is the case of orbit raising or orbit acquisition maneuvers, performed by the satellites themselves upon launcher

[8]The results and test description gathered in this section have been previously published in the journal Acta Astronautica [156].

separation in order to reach their operational orbit. These type of controlled transfers are a challenge for SST systems in the absence of maneuver knowledge as dynamics deviate from the expected ballistic motion. Therefore, a simulated surveillance scenario has been devised to shed light on the applicability of the proposed maneuver models to this situation of particular interest. The aim is to maintain custody of an electric-propelled target during its launch and early operations phase (LEOP), whose physical characteristics are summarized in Table 4.15. The reference trajectory is based a TLE sequence obtained from space-track.org for Starlink 1511 (NORAD ID 45767). The sequence is first smoothed, to minimize the impact of possible outliers or inaccurate estimates, and then used as target orbit for the Q-Law path controller developed by Petropoulos [157]. The outcome of this process is a high-fidelity simulation of a realistic orbital transfer, which is compared against the reference TLE sequence in Figure 4.29. The dynamical model used in such simulation considers the Sun, Moon and planets, Earth non-spherical model up to degree and order 30, solid and ocean Earth tides, Jachia-Bowman 2008 model [147] atmospheric density model and a cannonball model for the solar radiation pressure with conical eclipses. Note no thrusting occurs during eclipses as the propulsion system is only assumed to be operative under direct sunlight. One can appreciate five main segments in the reference trajectory: 1) a ballistic segment until day 3, 2) an orbit raising and circularization maneuver spanning 14 days, 3) a coasting or ballistic phase from day 17 to day 65, aimed at targeting the required operational orbital plane while avoiding potentially hazardous conjunctions with other space objects, 4) a final ascent to the operational orbit lasting 25 days, and 5) a nominal ballistic segment with a duration of 10 days. In line with the expectations [122], the executed maneuvers only present in-plane thrust components.

The target is sensed by a network composed of seven space surveillance radars distributed accross the globe, distributed as shown in Figure 4.30. Their performances, included in Table 4.16, are set according to the state of the art in space surveillance radar sensor technology [14], [158]. Observables are given as batches of azimuth, elevation, range and Doppler readings, hereafter referred to as radar tracklets, that last for as long as the

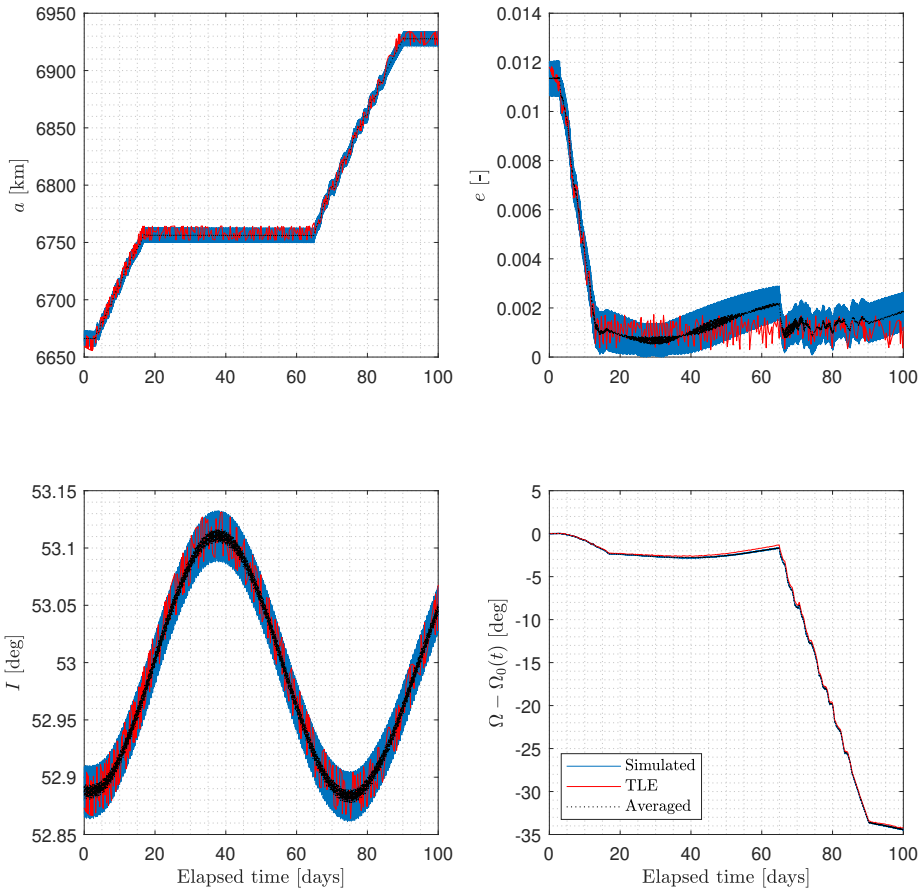

Figure 4.29: Starlink 1511 TLE sequence (red) and reference orbit (blue). Dashed black line corresponds to averaged orbital elements from the reference orbit. $\Omega_0(t)$ refers to the J_2 secular drift in right ascension of the ascending node of the initial orbit.

Initial Mass	260 kg
Area Drag	6 m^2
Area SRP	6 m^2
Drag Coefficient	2.1
Reflectivity Coefficient	1.3
Thrust Force	60 mN
Specific Impulse	1200 s

Table 4.15: Spacecraft physical properties for LEO single target test case.

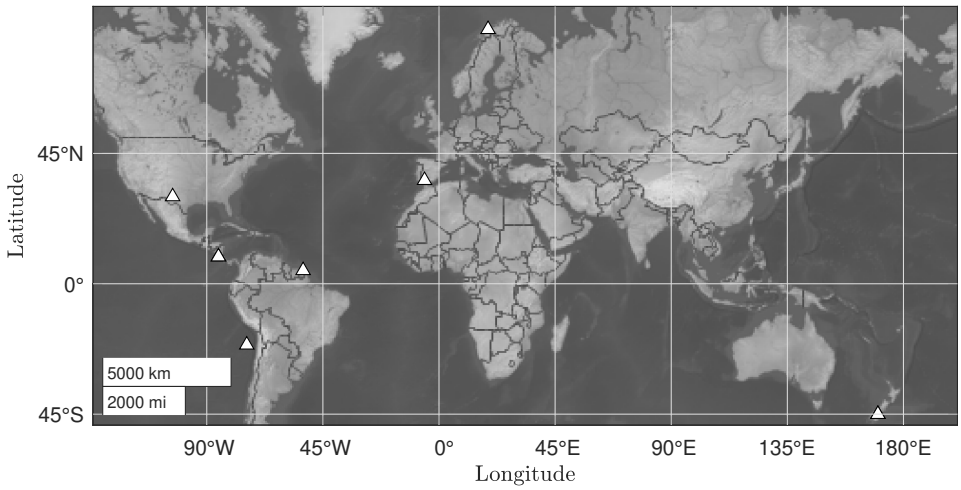

Figure 4.30: Sensor network used for the single target surveillance radar scenario.

satellite remains within the field of view (or field of regard) of the sensor. A total of 138 tracklets were obtained throughout the simulation time frame, showing the duration statistics gathered in Figure 4.31. On average, the target is detected by the sensor network every 17 hours (around 10 orbital revolutions), with a maximum blackout time of 38 hours. The particular detection sequence and mass utilization can be consulted in Figure 4.32, where a lower detection cadence is observed during the first 10 days of simulation. Note due to the relatively low observation cadence, variable process noise filters may not yield accurate results as the expected noise levels are of significant magnitude.

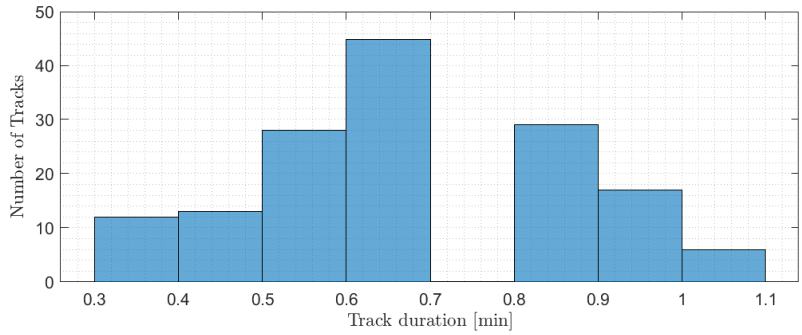

Figure 4.31: Tracklet duration statistics for the surveillance radar scenario.

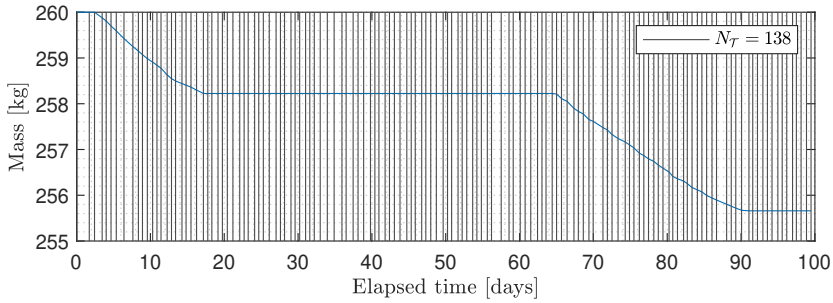

Figure 4.32: Tracklet sequence and mass profile throughout the simulated LEOP phase of Starlink 1511.

Field of Regard	180 ± 43 deg (azimuth) $\times$ 60 ± 15 deg (elevation)
Range detectability threshold	2,500 km
Range standard deviation	6.5 m
Angular standard deviation	0.1 deg
Doppler standard deviation	0.275 m/s
Scanning Frequency	7 seconds

Table 4.16: Surveillance network characteristics for the LEO single target test case.

State space filtering is performed via sequential Monte Carlo, in particular, a slight modification over the SHF MATLAB implementation used for the GEO optical survey test case (see Section 4.3.1). The main characteristics and configuration of the filtering scheme are enumerated in the following:

- To improve computational efficiency while still ensuring a certain level of dynamical accuracy, uncertainty propagation follows a bi-fidelity approach [154].
- The high-fidelity model is the same used to simulate the trajectory, except for a lower fidelity non-spherical Earth, up to degree and order 10.
- The low-fidelity model used is the semi-analytical SGP4 [159], where the B-star has been set to comply with the high-fidelity mode, this is, $B^\star = \rho_0 \frac{B}{2}$ where $\rho_0 = 0.1570$ kg/m$^2/R_\oplus$ being $R_\oplus$ the Earth radius and $B = \frac{c_D \mathrm{A}}{\mathrm{m}}$ is the ballistic coefficient.
- The maneuver mode transition density follows the definition given in Section 3.3.1, thus employing a least informative prior defined over the reachable set. Note this may yield pessimistic post-maneuver estimates though a more optimistic filter could present bias levels that hinder custody maintenance.
- The maximum expected control effort subject to a given maneuver is adaptively set as

$$\lceil d_{LT} \rceil = \min \left[d_{LT,max},\ \max \left(d_{LT,min},\ k_{d,LT} d(x^\star) \right) \right], \tag{4.8}$$

  where $d_{LT,max} = 10^{-3}$ m/s^2, $d_{LT,min} = 10^{-5}$ m/s^2 and $k_{d,LT} = 3$. Therein, the centroid of the admissible control region $x^\star$ is defined as

$$x_k^\star \approx \underset{x_k}{\operatorname{argmin}} \left[d_{LT}(x_k; \hat{x}_k^-) \right],\ \chi^2_{CDF}(d_M^2; n_z) \leqslant 1 - p_\alpha, \tag{4.9}$$

  where d_M is the Mahalanobis distance (see Equation (4.6)), $p_\alpha \ll 1$ is a statistical significance threshold and $\hat{x}_k^-$ is the expected state es-

timate under ballistic dynamics. Such an adaptive definition for the allowable control effort results in an improved filter convergence at the end of the different thrust legs. Note the larger the accessible space, the less likely is to obtain samples sufficiently close to the true state estimate, thereby leading to a delay in convergence. The latter fact could be mitigated by implementing other types of measurement updates, e.g. based on the Kalman recursion [155] or MCMC, improving the radius of convergence of the filter by preventing premature particle depletion.

- The regularized particle filter uses $N_s = 10^4$ samples to approximate the state distribution for each individual maneuver hypothesis, and the resampling threshold is set to $ESS_{min} = \frac{1}{4}N_s$.

- The filter is initialized with the true mean and an uncertainty $\sigma_{r,0} = 100$ m, $\sigma_{v,0} = 0.01$ m/s, $\sigma_{c_R,0} \equiv \sigma_{c_D,0} = 10^{-7}$ [-]. Process noise levels are set as $\sigma_{\mathrm{a}} = 10^{-7}$ m/s^2 and $\sigma_{c_R} \equiv \sigma_{c_D} = 10^{-14}$ [-]/s for the state and physical parameters, respectively.

- As in the test case presented in Section 4.3.1, measurement association under natural dynamics is based on the Mahalanobis distance and a 3-σ confidence interval is required for measurement to tracklet correlation.

- Whenever a measurement fails to correlate against the ballistic mode, samples from post-maneuver state density are obtained from the DREAM algorithm by making use of the linear reachability computation methods and derived low-thrust control metric.

- To improve convergence after maneuver segments, the filter implements a backwards simulation particle smoothing recursion [153].

Note the information encoded in a radar tracklet is much richer than that of an optical one since the latter lacks positional data. In fact, the analogous to an optical attributable [23] for a radar tracklet is directly the outcome of an initial orbit determination (IOD) method. A wide variety

of IOD methods have been proposed in the literature, among which one can find the vector approach by Gibbs (1839–1903) and the later modification proposed by Herrick (1911 – 1974) [160]. The attributable optimized coordinate system (AOS) [24] is another approach that relies on spherical coordinate transformations, as opposed to the usual Cartesian space. Uncertainty quantification has also been explored and applied to former methods, e.g. for the Goddard trajectory determination system (GTDS) range and angles method [161], among others [162], [163]. This effectively provides a fair approximation to the inverse of the measurement mapping, $h^{-1} : \mathbb{R}^{n_z} \rightarrow \mathbb{R}^{n_x}$, that can be used to characterize the uncertainty of IOD methods when applied to radar tracklets of varying measurement accuracy. Such an analysis has been performed on an operational environment by Lossaco et al. [164], indicating median state estimation errors of the order of 600 m in position and 50 m/s in velocity for the BIRALES radar under a general surveillance configuration. This type of analysis can help to set search bounds for the centroid defined in Equation (3.19), which is itself a representative state realization for the post-maneuver state density. Figure 4.33 depicts the error distribution of the IOD estimates, computed via the GTDS method, in the local RTN frame. These were applied to 1,000 realizations for each of the tracklets obtained within the current simulated radar scenario. In line with the expectations, uncertainty along the radial component is much lower due to the accuracy of the range measurements $\mathcal{O}(10 \text{ m})$ compared to azimuth and elevation $\mathcal{O}(1 \text{ km})$. Precision regarding the Cartesian position is relatively high, with a 95% confidence interval around ± 10 km. Velocity estimates, on the other hand, are much less accurate, showing bounds of ± 500 m/s for the 95% confidence interval. The latter fact is attributed to the relatively short time span of radar tracklets, on the order of one to two minutes. Knowledge regarding the maneuver characteristics or simply the thrust levels of a target can then be used to further reduce uncertainty in the Cartesian velocity space. Thus, the GTDS IOD estimates and uncertainty bounds from the numerical simulation test discussed earlier were used to solve Equation (4.9), both for providing a basis for the initial guess and a simplified alternative to the Mahalanobis distance.

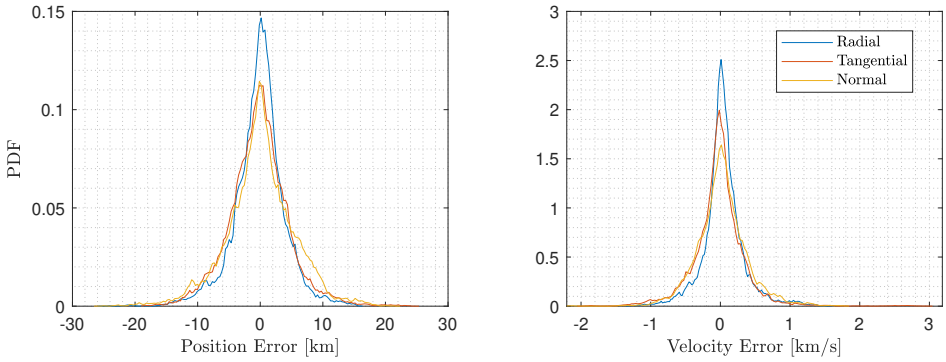

Figure 4.33: Initial orbit determination error in the local radial, along-track and cross-track Cartesian reference frame.

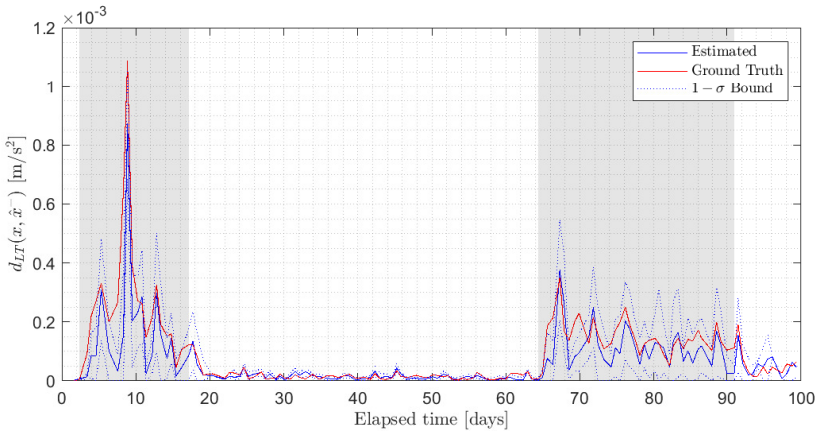

Figure 4.34: Low-thrust control distance metric sequence. Pre-maneuver state $\hat{\mathbf{x}}^{-}$ corresponds to the filter estimate at the previous observation epoch. Ground truth refers to the control distance between the estimated pre-maneuver state and the reference state at the estimation epoch. Shaded areas represent maneuver periods.

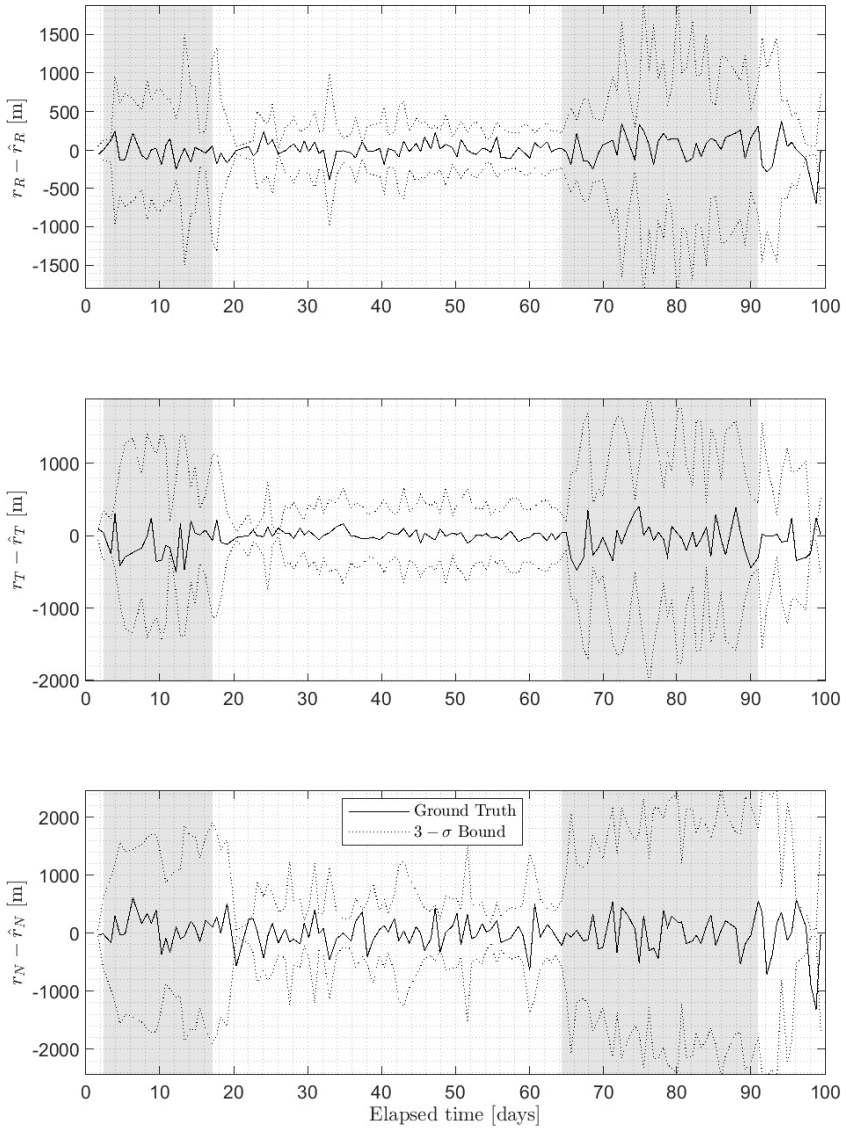

Figure 4.35: Position estimation error sequence in the radial-tangential-normal (RTN) reference frame. Shaded areas correspond to maneuver periods.

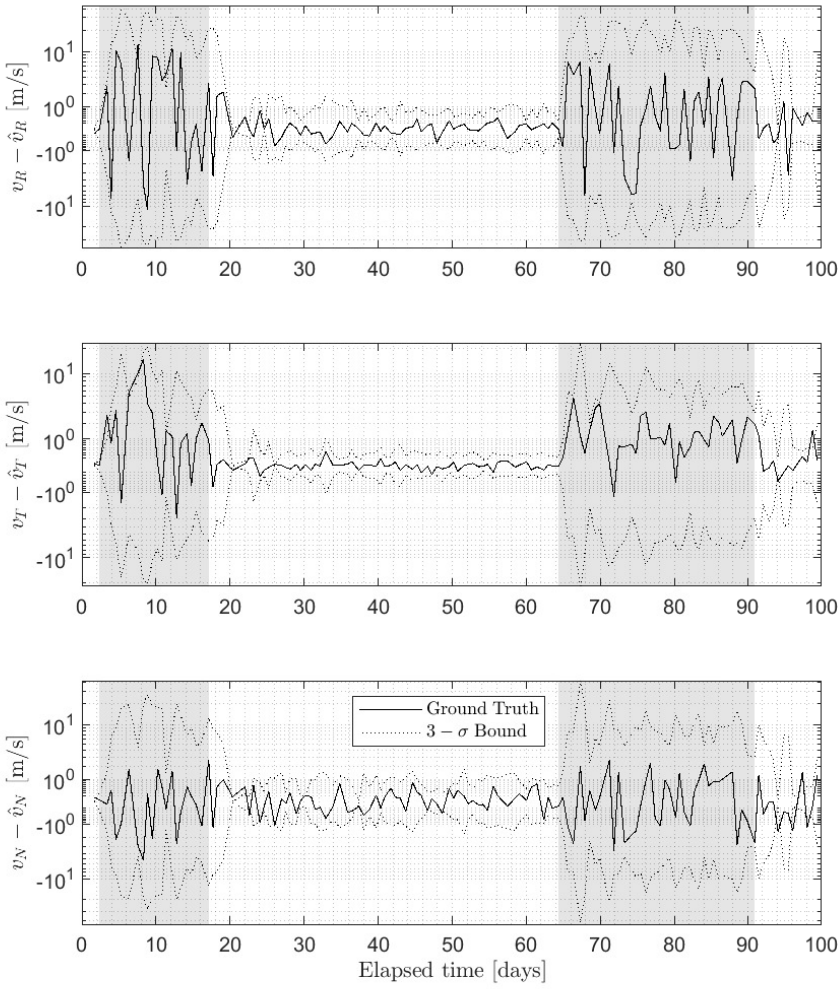

Figure 4.36: Velocity estimation error sequence in the radial-tangential-normal (RTN) reference frame. Shaded areas correspond to maneuver periods. Note the y-axis is represented in logarithmic scale.

The implemented RPF processes the observation sequence in 9.5 hours in a laptop computer equipped with an Intel Core i7-8750H CPU. Note this code was not explicitly optimized and no parallelization was performed, other than vectorization to improve handling of sampled populations. Accordingly, each tracklet requires 4 minutes to be processed on average, which in theory would allow to simultaneously track 200 maneuvering satellites on a single laptop. The filter failed at supporting measurement correlation in 3 occasions, i.e. it could not find a tracklet-compatible state estimate with a control distance lower that the prescribed threshold $d_{LT,max} = 10^{-3}$ m/s^2. Even if these three tracklets could not be associated to the target, custody was regained afterwards and so the filter was able to support subsequent measurement associations. The allowable control is a key driver for the accuracy of the filter: if set too low the filter could miss associations and fail to recover from biased estimates, but if it is too high then it may oscillate between biased estimates and yield overly pessimistic results even during ballistic sections. Figure 4.34 shows the sequence of expected (blue) and real (black) low-thrust control distance metric values, together with the estimated variance. The reference orbit from with the control effort is derived corresponds to the expected state at the previous associated observation, i.e. the maximum likelihood hypothesis prior to the current tracklet. Note the expected d_{LT} is lower than the reference value throughout the maneuvering segments, so that more frugal maneuvers are favoured according to the definition of the admissible control given in Equation 4.8. In general, $\lceil d_{LT} \rceil$ should have a fixed value corresponding to the expected thrust onboard the target, though this would negatively affect the filter convergence during the ballistic sections that follow thrusting segments. In fact, the position and velocity estimation errors, shown in Figures 4.35 and 4.36, respectively, show a delay in convergence. The high uncertainty during thrusting arcs prevents ballistic particles to fall within the support of the measurement distribution or, at least, to be sufficiently close to the true state in terms of variance. A potential solution to this maneuver termination biased might be the inclusion of some sort of maneuver smoother, that yields coherence between maneuvered states. However, this type of processing confronts with the proposed vacuous maneuver transition density

and, to the author knowledge, none of the existing methods (c.f. [101]) provide uncertainty estimates beyond first order sensitivities. Moreover, such smoothing recursions would not exempt the filter from estimating maneuver termination, which is a crucial input to such class of smoothers.

The estimated thrust levels, mostly between 10^{-4} and $3 \cdot 10^{-4}$ m/s^2, are aligned with the simulated value of $2.3 \cdot 10^{-4}$ m/s^2, so the proposed control metric is able to capture the dynamical characteristics of the transfer. Notwithstanding, there is a significant peak near the ninth day of simulation, where the estimated thrust is close to the prescribed threshold, almost reaching $9 \cdot 10^{-4}$ m/s^2. Prior to such estimate, the filter experienced a sequence of modest control effort estimates, resulting in higher tangential velocity estimates and thus lower orbital altitudes. Accordingly, and possibly due to an increase in the information content within the radar tracklet, the filter estimates a corrective maneuver of significant cost to reduce the mean orbital altitude bias. These findings support the use of a vacuous maneuver prior, or at least to avoid employing control effort-weighted maneuver transition densities. Note the latter may further accumulate bias along the tangential velocity component as transfers with a lower control cost are assumed more likely, thus trying to place the apoapsis of the orbit at each subsequent observed state during orbit raising maneuvers. Nonetheless, provided it is contained within the reachable set, the proposed methodology allows for an arbitrary definition of the maneuver mode, as exemplified by the heuristics-based transition proposed in Section 3.3.2 and later tested in Section 4.3.1. Moreover, the true state evolution depicted in Figure 4.29 suggests the control policy follows some kind of periodic pattern that might be inferred by looking at the sequence of state estimates. Thus, similar to the proposal in [152], it might be possible to learn maneuver patterns in terms of the rate of change of the semi-major axis for instance, potentially leading to an increase in state estimation accuracy. Note the latter would not be exempt from estimation bias as these heuristic models are highly dependant on maneuver characterization and the level of similarity between past and present maneuvers.

4.4. Generalized Labeled Multi-Bernoulli filter for maneuvering space objects[9]

This section builds on the results obtained in Section 4.3. In doing so, a computationally efficient, yet sufficiently robust, filtering scheme is selected to be used within a multi-object filtering framework. Note the added complexity of filters such as the GLMB demands simplified approaches for measurement update, discouraging the use of SMC-like methods in favor of Kalman filter implementations. Moreover, the are some concerns regarding data processing in what regards measurements. Multi-object filters significantly benefit from measurement pools, i.e. the fact that a single observation corresponds to at most one target. This is intrinsic to most (ground focused) surveillance sensors, which provide low-latency coverage for a certain subset of the state space, this is, their field of view. Unfortunately, space surveillance sensors present rather narrow fields of view, especially electro-optical ones, so measurement data is mostly collected as an information stream of a reduced set of targets rather than isolated frames containing a global snapshot of a sub-population of considerable size. Moreover, as discussed in Section 4.3, a space object observation is effectively retrieved as a single pass through the sensor's field of view, thus being a short time sequence of observables commonly referred to as tracklet. Nonetheless, it is possible to efficiently evaluate tracklet pairs for correlation, for instance employing one of the various IOD methods that exist for both optical and radar measurements (see Pastor et al. [25] for a comprehensive discussion). Therefore, tracklets can be grouped into the same observation pool as long as there is no single tracklet pair combination whose IOD estimate yields an Earth orbit within the detectability characteristics of the surveillance network. To simplify the analysis, in the remainder of this section detectable targets are assumed to be simultaneously observed.

The GLMB filter discussed in Section 2.1.1 is an exact filter, though

[9]The results and test description gathered in this section have been previously presented at the 33rd AAS/AIAA Spaceflight Mechanics Meeting[165]

a direct implementation is deemed intractable for most applications as the number of hypotheses to be considered grow geometrically. Therefore, a common approximation in the literature is the δ-GLMB, whose only modification consists in sequentially pruning hypotheses to comply with certain computational requirements. Typically, hypothesis truncation is based on thresholding, so that those featuring a sufficiently low likelihood are discarded, and truncation, so that at each step only the most likely N_H hypotheses are retained. Hypothesis scoring and ranking, commonly referred to as the linear assignment problem, is usually performed via Murty's algorithm [166], especially for problems of moderate size. For large scale problems, the required matrix operations lose tractability fairly rapidly and randomized algorithms, such as Gibbs sampling (a kind of MCMC method), are preferred [167]. Under non-linear transformations, single object density representations may also pose tractability issues. For general non-linear inference problems, particle filters can provide arbitrarily accurate approximations that depend upon the size of the sampled population (see Sections 4.3.1 and 4.3.2). However, as discussed earlier in Section 2.2.3, these type of methods suffer from the curse of dimensionality, thus requiring a number of particles that grows exponentially with the number of state variables. Alternatively, closed-form solutions to non-linear state space filtering problems typically assume a Gaussian state density representation, e.g. the MHE discussed in Section 2.2.2 or the class of non-linear Kalman filters described in Section 2.2.1. Hence, the accuracy of these approximate filters is dictated by the divergence of the filtered PDF from the Gaussian family. Some schemes, such as the Gaussian sum filter (GSF) [168], seek to bridge the gap between sampling and Gaussian based approaches. In this particular case, the state density is approximated by a GMM and the measurement update is performed over each individual component following certain Kalman scheme, usually an EKF. Therefore, provided the Gaussian sum approximates the true prior density, a linear measurement update on each component results in a fair approximation to the true posterior. For the latter assumption to hold, it is customary for the filter to behave similar to a particle filter, thus, capable automatic component splitting is deemed crucial. To this end, the adaptive entropy-based Gaussian mixture infor-

mation synthesis (AEGIS) method [169] monitors the differential entropy during uncertainty propagation and splits individual components whenever the accumulated entropy exceeds a prescribed threshold, being the latter an indicative of the linearization error.

Recall the intention is to devise a tractable GLMB filter implementation (see Section 2.1.1) that can manage the characteristics of space surveillance, i.e. non-linear transformations and data sparsity. In fact, since each individual target is modeled as jump Markov, it is mandatory for the filter to allow for mixture or sample based state representations. Therefore, and in light of a more efficient implementation, the proposed GLMB scheme implements a GSF to perform state space filtering at a single target scale, augmented with AEGIS to allow for automatic component splitting during uncertainty propagation. In Section 3.3.1, the transition density for a general maneuver was defined as a uniform distribution within the set of reachable states. Such a space can be instantiated via a Gaussian mixture, for instance following the approach proposed by DeMars & Moriba for the admissible region [170]. This mixture would then be updated subject to certain measurement correlation, resulting in a filtered post-maneuver state density. However, the accuracy of a Kalman update significantly depends upon the validity of linear assumptions within the support of the prior Gaussian distribution, which essentially means the size of each individual component shall be intractably small.

Hereafter, this limitation is circumvented by simply instantiating components that belong to the admissible control region, i.e. the intersection between the reachable set and the support of a hypothesized measurement. Therefore, the filter directly approximates the posterior distribution under the event of a maneuver, represented by a Gaussian mixture of a fixed size within the admissible control region. The latter is performed via the following procedure:

1. The single target density is approximated as

$$p(x) \approx \sum_{i=1}^{n_C} \omega_i \mathcal{N}(x; \mu_i, \Sigma_i), \tag{4.10}$$

where ω_i is i-th component weight and μ_i and Σ_i their expected state and covariance, respectively. Reachability computations are then obtained with respect to each component mean, so that in practice the covariance is assumed to be negligible when evaluating the control distance metric.

2. Thus, the reachable set is defined as

$$\mathcal{R}(x_+; x) = \bigcup_{i=1}^{n_C} \mathcal{R}(x_+; \mu_i) \tag{4.11}$$

where each individual $\mathcal{R}(x_+; \mu_i)$, computed using the linear approach defined in Section 3.2.1, is assigned the corresponding component weight ω_i for filtering purposes.

3. Recall that the vacuous transition density defined in Section 3.3.1 requires computing the volume of the reachable set, for which the techniques discussed in Sections 3.2.3 and 4.2.2 are employed. Therefore, 10,000 samples are generated from the AABB approximation of the reachable set and test them for containment, checking whether their control distance is lower than the prescribed $\lceil d \rceil$.

4. Samples that fall within the reachable set are then transformed to the measurement space, from which the admissible control region can be extracted via containment analysis. Note as in Sections 4.3.1 and 4.3.2, a statistical significance threshold p_α is used to numerically approximate the support of the measurement likelihood.

5. Finally, such an admissible region is instantiated following the approach proposed in [170], given proper bounds for the uncertainty along each dimension. As discussed by DeMars & Moriba, each component weight is set proportional to the volume it represents within the uniform distribution. Therefore, the (pre-conditioned) prior density following a maneuver transition is given by

$$p_+(x_+) = \sum_{j=1}^{n_{C+}} \omega_j^+ \mathcal{N}(x_+; \mu_j^+, \Sigma_j^+), \tag{4.12}$$

where each individual component weight is computed as

$$\omega_j^+ = \sum_{i=1}^{n_C} \frac{\sqrt{2\pi^{n_x}\det(\Sigma_j^+)}}{\mathcal{V}_{\mathcal{R}}(\mu_i)} \omega_i \phi(\mu_j^+, \mu_i). \tag{4.13}$$

Thus, ω_j^+ is computed as the weighted sum of the ratio between its representative volume and the volume of the reachable set for each pre-maneuver state density component, multiplied by the reachability inclusion function $\phi(\cdot)$ defined in Equation (3.63).

The benefits of advanced multiple maneuvering target tracking filters are best assessed in cluttered regions wherein targets maneuver in close proximity. An optical surveillance scenario in the Geosynchronous belt is designed to stress the proposed filter scheme. Therein, a set of five spacecraft (chasers) perform impulsive maneuvers to rendez-vous a second set of quiescent objects (targets), with the same multiplicity as the former. The so-called chasers perform a pair of impulsive burns within two temporal windows, $t_{m,1} \in [11, 14]$ and $t_{m,2} \in [17, 20]$ elapsed hours, both randomly directed and with magnitude $\Delta v_{1,2} \sim \mathcal{U}(0.5, 1.0)$ m/s. Upon application of the second maneuver, each chaser is located at a distance $\Delta r \sim \mathcal{U}(200, 1000)$ m (again in a random direction) from a target spacecraft, thus being grouped in pairs. A ground-based electro-optical is used to survey these targets for a total of 5 days, resulting in a collection of attributables gathered every 3 hours, yet only between dawn and dusk. The observational accuracy of the sensor is set to 1 arcsec for angular observables and 0.1 arcsec/s for their rates. The simulated dynamical environment considers a non-spherical Earth model of degree and order 20, the Sun and Moon as third body perturbations and a cannonball model for the solar radiation pressure.

The proposed GLMB filter, implemented in C++, takes the following configuration:

- Probability of survival $P_S^{(m)}(x, \ell) = 1 - 10^{-7}$.
- Probability of maneuver $\vartheta(m_+|m) = 0.01$.

- Probability of detection $P_D^{(m)}(x,\ell) = 1 - 10^{-3}$.
- To emphasize on the impact of maneuvers on the performance of the filter, no birth, clutter or missed detections are simulated.
- Maximum number of parallel hypotheses $n_H = 50$.
- Minimum relative hypothesis weight $\omega_{min} = 10^{-7}$.
- Measurement gate for data association $d_{M,th} = 6.0$ in the Mahalanobis distance.
- Process noise for AEGIS propagation $\sigma_{\mathrm{a}} = 10^{-8}$ m/s^2. This noise directly affects automatic splitting by inflating the covariance of each Gaussian component.
- Maximum expected control effort $\lceil d_I \rceil = 5$ m/s measured via the impulsive control distance described in Section 3.1.1.
- Single target filtering is performed via a Gaussian sum filter with a UKF measurement update step.
- The proposed JMS GLMB filter is compared against an equivalent process noise approach, wherein the required noise is estimated via the low-thrust control distance metric developed in Section 3.1.3. This implementation is essentially a non-linear alternative to the proposal in [171], which proposed a linear model for the unknown control acceleration. The maximum expected noise level is set to $\lceil d_{LT} \rceil = 10^{-5}$ m/s^2, which emulates a thrust of 50 mN exerted by a spacecraft with a mass of 5 metric tons.

The filter is tested for a total of 10 randomized runs, each requiring an average computational time of 9.5 hours in a laptop computer equipped with an Intel Core i7-8750H. The main share of the computational burden is devoted to control distance metric evaluations, expected to significantly decrease if these are performed on a GPU. The equivalent process noise implementation, on the contrary, only requires 9 minutes per test on average: a reduction of more than one order of magnitude. This is attributed to

the relatively simple computations required by the low-thrust metric, which reduce to a few matrix inversions. Table 4.17 summarizes the measurement to object correlation errors, defined as any type of erroneous data association. The majority of these correspond to cross-tagging, i.e. assigning the detection of a chaser to a target and vice versa, though missed detections are also found. Note that test case V is missing from the result because it failed in maintaining custody of one of the chasers, for which yielded subsequent measurement correlations with a position bias around 70 km in alternating directions. In fact, the proposed filter failed to recover the correct data association sequence for all the simulation runs. This is attributed to the fact that at the first post-maneuver observation epoch, the chasers are less than 1 km away from their corresponding targets, which is considerably close given the angular observation variance itself leads to 200 m 1-σ uncertainty at Geosynchronous altitudes. Moreover, the true association hypothesis did not qualify among the best 50 alternative for those scenarios with less than 5 surviving hypotheses, emphasizing on the level of complexity of a priori scenarios like this one. The latter fact suggests that the number of maximum parallel hypotheses needs to be increased to enhance data association performance, though another alternative might be to employ multi-scan algorithms [172] at the cost of computational complexity and tractability. Cross correlations not only occur between chasers and targets, but also among maneuvering targets themselves. Provided their reachable sets intersect, there is no a priori information that can be used to resolve the ambiguity apart from, perhaps, the required control cost of the transfer. However, as discussed earlier in Section 4.3.2, a maneuver transition density proportional to the control effort is deemed to yield biased state estimates in general, mainly because the complete state space is not observable. In any case, cross-tagging between chasers is the main source of miscorrelations for test runs number IV, VI and X, where a pair of chasers is swapped from observation frame 5 onwards, resulting in 32 individual association errors. It shall be noted that embedding maneuver detection within the data association problem, i.e. hypothesizing over the event of individual maneuvers, would result in an even greater size of the hypothesis space, negatively affecting the filter performance for a fixed n_H.

Estimation performance of multi-object filters is usually given in terms of the optimal sub-pattern assignment (OSPA) metric [173], which is defined for a set of estimated $\mathbf{X}$ and reference $\mathbf{X}^*$ targets and as

$$D(\mathbf{X},\mathbf{X}^*) = \left[\frac{1}{|\mathbf{X}|}\left(\min_{\gamma\in\Gamma_{|\mathbf{X}^*|}}\sum_{i=1}^{|\mathbf{X}|} d_c(\mathbf{x}_i,\mathbf{x}^*_{\gamma(i)}) + (|\mathbf{X}^*|-|\mathbf{X}|)\cdot c^p\right)\right]^{\frac{1}{p}}, \quad (4.14)$$

being Γ_n the set of all permutations $\{1,...,n\}$ for $n \in \mathbb{N}$, and $\gamma \in \Gamma_n$ a sequence $\{\gamma(1),...,\gamma(n)\}$. The distance $d_c(\mathbf{x}_i,\mathbf{x}^*_{\gamma(i)}) = min\left(c, d_b(\mathbf{x}_i,\mathbf{x}^*_{\gamma(i)})\right)$ is defined as the minimum between a cut-off c and a base distance $d_b(\cdot,\cdot)$. Within this work, the cutoff parameter is set to $c = 100\ km$ and the order to $p = 2$ so that a missing target is translated to an estimation error of 100 km and $d_b(\cdot,\cdot)$ is the Euclidean position distance. Figure 4.37 gathers the results in terms of OSPA for the proposed filter (top) and the equivalent process noise alternative (bottom). Note in both cases, the implemented GLMB filter can maintain custody of all the targets within the simulation window, and hence OSPA represents the population-averaged position estimation error with respect to the ground truth. As the chaser satellites perform the impulsive maneuvers, between 12 and 24 hours of simulation, the OSPA metric raises from $\sim$ 200 m to around 20 km. Between the first and second days of simulation, estimation error levels are maintained, as the amount of available information is still insufficient to further reduce the positional uncertainty without affecting the filter consistency. As a measure of the latter, one can use the scaled normalized estimation error squared (SNEES)

$$SNEES_i = \frac{1}{n_x}(E[x_i]-x_i^*)^T Cov[x_i]^{-1}(E[x_i]-x_i^*), \quad (4.15)$$

which shall remain close to 1 for consistent filters. Figure 4.38 shows the results for the impulsive metric based filter (top) and variable process noise implementation (bottom). Therein, the proposed filter is shown to outperform the latter approach in terms of consistency, arguably due to a better characterization of the maneuvering characteristics of the chasers. However, there are occasions in which the SNEES increases beyond reasonable

Test Case	I	II	III	IV	VI
Sequential	14 (75)	8 (50)	15 (54)	29 (70)	28 (60)
Hypothesis I	6 (83)	4 (51)	10 (49)	40 (67)	34 (64)
Hypothesis II	- (83)	6 (51)	8 (50)	39 (66)	34 (64)
Hypothesis III	- (68)	35 (49)	12 (48)	41 (67)	36
Hypothesis IV	- (68)	-	12 (49)	41 (68)	-
Hypothesis V	- (69)	-	10 (51)	40 (68)	-

Test Case	VII	VIII	IX	X
Sequential	10 (44)	14 (55)	14 (17)	32 (51)
Hypothesis I	6 (36)	6 (51)	10 (12)	36 (50)
Hypothesis II	6 (37)	8 (50)	12 (10)	35 (53)
Hypothesis III	6 (78)	38 (52)	10 (12)	37 (67)
Hypothesis IV	-	8 (53)	12 (12)	36 (51)
Hypothesis V	-	40 (53)	8 (14)	36 (52)

Table 4.17: Data association errors for a total of 200 observations. Results for the equivalent process noise implementation are shown in parentheses. Dashed line indicates no additional hypotheses were promoted.

values, i.e. $SNEES > 3$. These correspond to miscorrelations for the target satellites, whose uncertainty was believed to be small and thus present moderately biased estimates. Despite the complexity of the proposed test scenario, with objects orbiting in close proximity, the tested algorithms show convergence in terms of consistency and accuracy, suggesting advanced multi-object filtering methods such as the GLMB can be a valuable asset for modern SST systems.

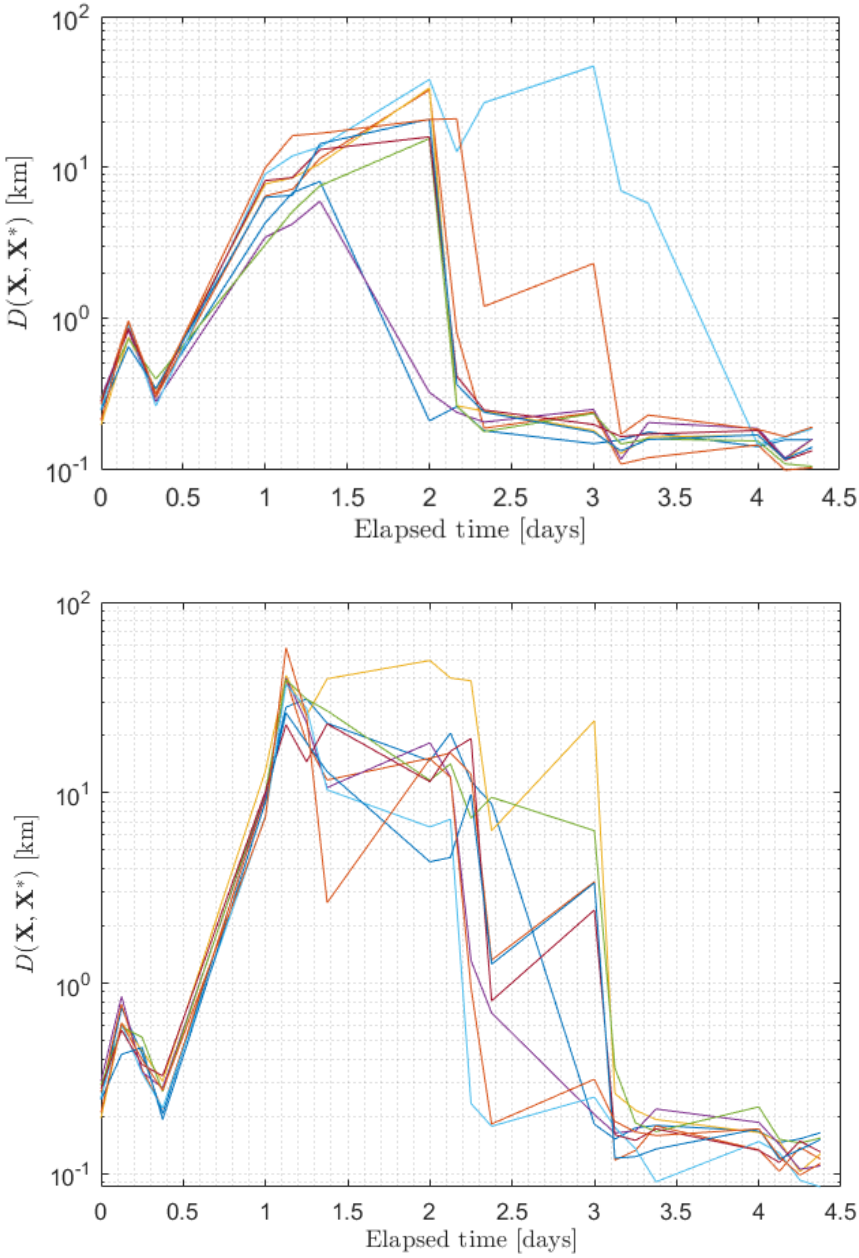

Figure 4.37: OSPA metric for the test cases considered. The single-object expected state is computed as the weighted sum of each Gaussian mixture mean. The top figure corresponds to the proposed approach and the bottom one shows the results for the equivalent process noise alternative.

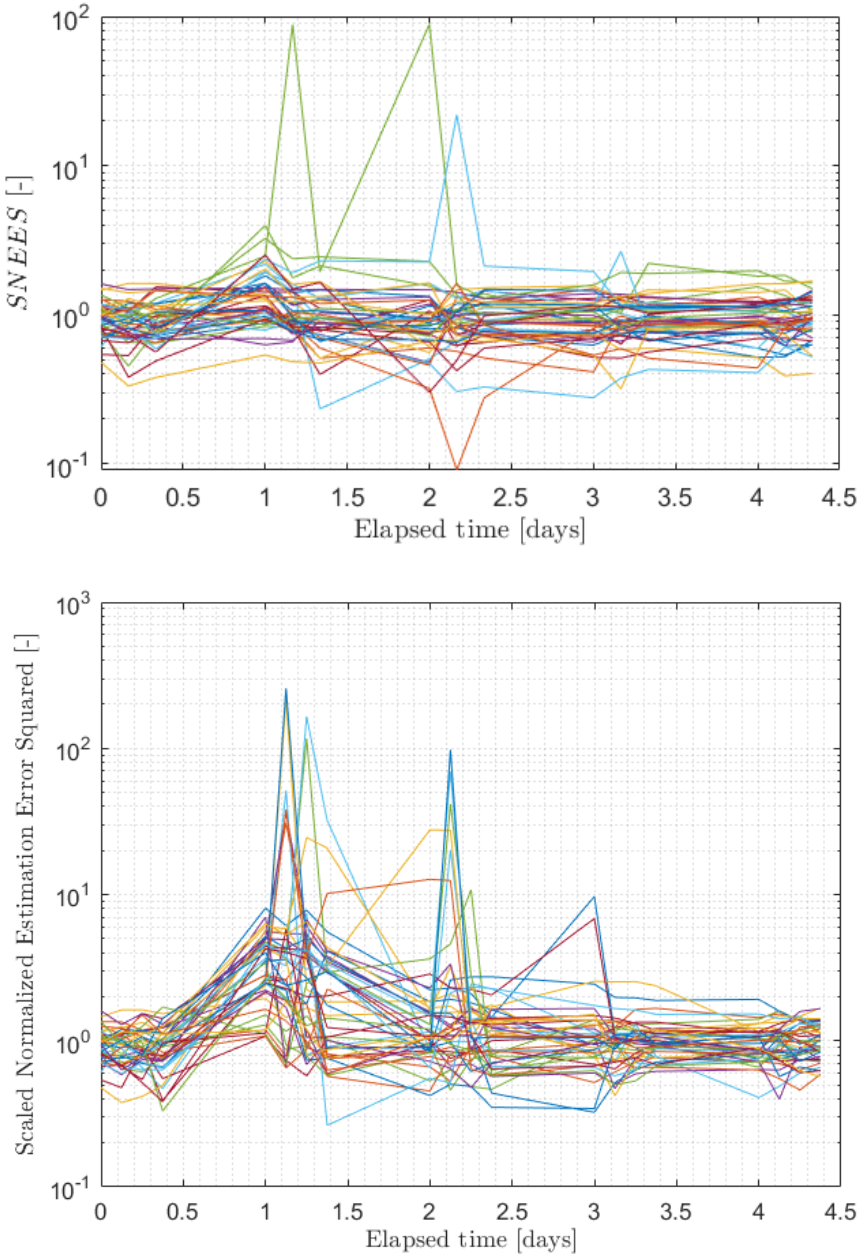

Figure 4.38: SNEES evolution for the five maneuvering objects computed as the weighted sum for each individual Gaussian component. The top figure corresponds to the proposed approach and the bottom one shows the results for the equivalent process noise alternative.

5. CONCLUSIONS

The current situation of Earth's orbital population demands better SST system performances and capabilities, especially in what regards automated operations. Unacknowledged satellite maneuvers are among the primary sources of SST system malfunctions, resulting in custody losses and unacceptable estimation biases. Advanced maneuvering target tracking methods for multi-object scenarios are hence demanded by the industry, mainly because the current state of the art in spacecraft maneuver modelling for tracking scenarios does not meet the requirements of space surveillance data. The latter is dominated by sparsity, whose detrimental effect is exacerbated by the highly non-linear characteristics of orbital motion and observations. Therefore, simplistic models based on inflated process noise levels fail to characterize orbital maneuvers of moderate control effort. Moreover, surveillance observations convey limited information regarding the orbital state of the sensed object, not allowing for a complete state estimation. The latter fact renders optimal control based solutions necessarily biased, since uncertainty within the unobservable space is artificially reduced by the optimality assumption. In addition, there exist two different types of propulsive systems for space operations, which are indistinctively used onboard spacecraft. Systems based on chemical processes benefit from high thrust levels at the cost of propellant efficiency. On the other end, electric thrusters, relying on electromagnetic forces, present a frugal alternative to the detriment of power demands and thrust magnitude. Accordingly, maneuvers executed with different propulsive systems show distinct characteristics.

Based on these facts, it appears natural to develop novel models for general orbital maneuvers, treating chemical and electric thrusters as separate problems. Such models shall be easily implementable, so that their input and tunable parameters must be reduced to a bare minimum. They should also capture the dynamical characteristics of orbital motion while maintaining the computational complexity low, as demanded by the large scale of SST applications. Note the main characteristic of a maneuver may

well be its control effort, which ultimately translates into cost for the spacecraft operator. Therefore, it is natural to define maneuver models based on the expected control capabilities of a target, for instance in terms of fuel consumption or required thrust level. The former is highly correlated with the ΔV, a magnitude that is commonly used in astrodynamics to express the control cost of an orbital transfer. Since chemical thrusters are not limited by thrust magnitude but propellant utilization, it is reasonable to characterize their maneuvers in terms of ΔV, being directly related to the net orbital velocity change. Electric propulsion, on the other hand, is largely limited by the exerted thrust rather than the accrued fuel consumption, supporting the definition of an effort metric based on the thrust magnitude. In light of this analysis, two efficient control distance metrics are proposed in the context of this thesis. The former, based on the solution to Lambert's problem, provides a reasonably accurate estimate of the minimum ΔV required to transfer from one orbital state to another within some temporal bounds. For low-thrust transfers, the proposed metric corresponds to the minimum control acceleration magnitude required to perform certain orbital transfer within a given time window. These metrics, though based on surrogate dynamical models, have shown exceptional computational efficiency and reasonable levels of accuracy. Moreover, the surrogate models used are quasi-linear functions of the change in orbital elements, thus allowing for extremely efficient reachability computations subject to prescribed control effort bounds. Again, there exist some limitations on the accuracy of said reachability computations, though these can be mitigated in stochastic environments by the noise inherent to the estimation process. The concept of reachable set, that is, the set of accessible states provided a reference initial state, a maximum control capacity and a time horizon, has multiple applications in spacecraft operations and SST. In fact, the reachable set is a very good candidate to define the support of a general orbital maneuver model, as it is completely defined by the maximum allowable control for tracking purposes. Also note that in the absence of prior information, there is no evidence for the spacecraft to be located at a given orbital state other than its propulsive system being capable of such an orbital transfer. Therefore, a reasonable and potentially unbiased maneuver model may well be defined

as a vacuous (or least informative) distribution over the set of reachable states. In fact, such model can be used as a proper transition density when normalized, which simply requires finding the volume of the reachable set. Efficient volumetric computations are enabled by the frugal computational demands of the developed metrics, though they require the use of sampling techniques. In addition, it is possible elaborate further maneuver modes based on historical data, after identifying patterns in the net orbital effect and control cost of maneuvers. An illustrative example are station keeping maneuvers performed by satellites in the geosynchronous belt, which need to perform periodic maneuvers to maintain their assigned orbital slot. A heuristic characterization of such maneuvers can then be combined with the vacuous reachability based model to define alternative transition densities or maneuver modes. Note these would result in accurate estimation of the orbital state provided the maneuver in question resembles previous patterns, potentially resolving ambiguities.

These novel maneuver models have been tested in different surveillance scenarios and filter implementations. A simulated test case consisting of a single target in GEO, observed by two ground-based electro-optical sensors, is used to analyze the validity of the proposed vacuous and heuristic models. This test compares two filter implementations, the MHE as an industry standard and a SMC as state of the art in non-linear filtering. The developed vacuous maneuver model allows to keep track of the target throughout the simulation window, which extends for more than a year. This is true for the Gaussianity-based MHE as well as the stochastic particle filter, both showing similar state estimation errors yet the former providing overly confident uncertainty levels. Maximal performance is obtained for a self-reinforcing implementation, which makes use of previously detected maneuvers to propose additional maneuver modes. This is shown to considerably decrease estimation errors within the first day following a maneuver, something that might be critical to perform accurate conjunction screening at orbital population levels. The capabilities of the low-thrust variant of the general maneuver model are tested in a rather challenging scenario: the LEOP phase of a Starlink satellite, observed by a network of space surveillance radars. This test presents two different maneuver pe-

riods, each with a duration in excess of 10 days, separated by a loitering segment of 45 days. Despite maneuvers are quite demanding in terms of control, the implemented particle filter is capable of maintaining custody of the target throughout the complete simulation window. In fact, estimates are consistent in general, though velocity estimation errors are quite significant during maneuver intervals. Interestingly, midway along the first orbit raising segment, the filter consistently underestimates the required control effort, and thus the mean orbital altitude, modifying the orbital eccentricity and anomaly in a way such that the state matches the observations. Immediately after, it estimates a maneuver of significant magnitude to drive the tangential velocity estimate closer to the true value, therefore simulating a strong orbit raising. This suggests that the assumption of control optimal orbital transfers can lead to unbounded bias accumulation, and ultimately a loss of custody. Finally, the proposed maneuver models are employed in an advanced multi-object filter, the JMS GLMB. An implementation using a GSF single target PDF representation and UKF measurement update is tested in a simulated optical surveillance scenario in the GEO belt. Therein, five objects maneuver to rendezvous another set of (five) targets within the first day of simulation. After the maneuver period, the trackable population is grouped in pairs with a positional distance not greater than one kilometer, and thus barely distinguishable from the sensor perspective. The GLMB filter is evaluated for 10 Monte Carlo runs, succeeding in accurately tracking all the targets in all but one case, where it lost custody of one object. Despite being incapable to support the true data association sequence, mainly due to the relatively close distance between object pairs and the ambiguity arising from intersecting reachable sets, the filter maintains custody of all the targets and is capable of providing consistent state estimates. Even if this is a largely simplified surveillance scenario, the developed test case is rather challenging from data association and state estimation perspectives. Nonetheless, the JMS GLMB filter implementation is able to find reasonable measurement to object association maps that lead to filter convergence after three to four days of simulation. This is in part attributed to the fact that the JMS GLMB does not hypothesize over the event of an individual maneuver during data association but rather whether an observation belongs

to an object regardless of its motion model.

This thesis has advanced in the development of maneuver models for space surveillance and tracking applications, providing a set of efficient methods to perform control effort and reachability calculations in Earth orbit. It has shown the applicability of the proposed methods to realistic surveillance scenarios, highlighting the benefits of a vacuous maneuver model and the use of maneuver heuristics. The developed concepts and techniques have helped to solve complex tracking problems, such as maintaining custody of a Starlink during LEOP making use of space surveillance radars. In addition, the applicability and tractability of the proposed methods has been evaluated in a simplified multi-object tracking environment, making use of advanced multi-target filters such as the GLMB. The developed maneuver models have proved consistent in terms of state estimation, supporting data association decision in the presence of ambiguity.

BIBLIOGRAPHY

[1] M. J. Holzinger, D. J. Scheeres, and K. T. Alfriend, "Object correlation, maneuver detection, and characterization using control distance metrics," *Journal of Guidance, Control, and Dynamics*, vol. 35, no. 4, pp. 1312–1325, 2012.

[2] D. J. Kessler and B. G. Cour-Palais, "Collision frequency of artificial satellites: The creation of a debris belt," *Journal of Geophysical Research: Space Physics*, vol. 83, no. A6, pp. 2637–2646, 1978.

[3] "Esa space debris mitigation handbook," European space agency ESA, Apr. 7, 1999.

[4] "Guidelines and assessment procedures for limiting orbital debris," National aeronautics and space administration NASA, Aug. 1995.

[5] "Cnes standards collection, method and procedure space debris – safety requirements," Centre national d'études spatiales CNES, 1999.

[6] "Space technology items. general requirements mitigation of space debris population," Russian aviation and space safety standards, 2000.

[7] "Technical report on space debris," United nations office for outer space affairs UNOOSA, 1999.

[8] "Iadc space debris mitigation guidelines," Inter-agency space debris coordination committee, version 0, Oct. 15, 2002.

[9] C. Pardini and L. Anselmo, "Assessment of the consequences of the fengyun-1c breakup in low earth orbit," *Advances in Space Research*, vol. 44, no. 5, pp. 545–557, 2009.

[10] P. Anz-Meador, J. Opiela, and J.-C. Liou, "History of on-orbit satellite fragmentations," Tech. Rep., 2023.

[11] U. S. S. Network, "Satellite box score," *NASA Orbital Debris Quarterly News*, vol. 27, Jun. 2023.

[12] E. S. D. Office, "Esa's annual space environment report," Jun. 2023.

[13] EUSST. "About the european space surveillance and tracking (EUSST) consortium." (2021), [Online]. Available: `https://www.eusst.eu/about-us/` (visited on 10/09/2023).

[14] G. Muntoni, G. Montisci, T. Pisanu, P. Andronico, and G. Valente, "Crowded space: A review on radar measurements for space debris monitoring and tracking," *Applied Sciences*, vol. 11, no. 4, 2021.

[15] J. M. Hermoso *et al.*, "System approach to analyze the performance of the eu space surveillance and tracking system," in *AMOS conference proceedings*, 2021.

[16] S. Lopez-Jimenez, A. Pastor, and D. Escobar, "Improving orbital uncertainty realism through covariance determination," *Acta Astronautica*, vol. 181, pp. 679–693, 2021.

[17] L. Pode, "An approach to the study of automatic target tracking," *Transactions of the American Institute of Electrical Engineers, Part I: Communication and Electronics*, vol. 78, no. 6, pp. 806–811, 1960.

[18] G. M. Goff, J. T. Black, and J. A. Beck, "Tracking maneuvering spacecraft with filter-through approaches using interacting multiple models," *Acta Astronautica*, vol. 114, pp. 152–163, 2015.

[19] C. Shabarekh, J. Kent-Bryant, G. Keselman, and A. Mitidis, "A novel method for satellite maneuver prediction," in *Advanced Maui Optical and Space Surveillance Technologies Conference.(Maui, Hawaii, USA)*, 2016.

[20] R. Linares and R. Furfaro, "Space objects maneuvering detection and prediction via inverse reinforcement learning," in *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, 2017, p. 46.

[21] J. Siminski, H. Fiedler, and T. Flohrer, "Correlation of observations and orbit recovery considering maneuvers," *AAS/AIAA Space Flight Mechanics*, 2017.

[22] A. Milani, G. F. Gronchi, M. De' Michieli Vitturi, and Z. Knezevic, "Orbit determination with very short arcs. i - admissible regions," *Celestial Mechanics and Dynamical Astronomy*, vol. 90, pp. 57–85, Jul. 2004.

[23] J. Maruskin, D. Scheeres, and K. Alfriend, "Correlation of optical observations of objects in earth orbit," *Journal of Guidance, Control, and Dynamics*, vol. 32, Jan. 2009.

[24] B. Reihs, A. Vananti, T. Schildknecht, J. A. Siminski, and T. Flohrer, "Application of attributables to the correlation of surveillance radar measurements," *Acta Astronautica*, vol. 182, pp. 399–415, 2021.

[25] A. Pastor, M. Sanjurjo-Rivo, and D. Escobar, "Initial orbit determination methods for track-to-track association," *Advances in Space Research*, vol. 68, no. 7, pp. 2677–2694, 2021.

[26] B.-n. Vo *et al.*, "Multitarget tracking," *Wiley encyclopedia of electrical and electronics engineering*, no. 2015, 2015.

[27] J. Aristoff, J. Horwood, N. Singh, A. Poore, C. Sheaff, and M. Jah, "Multiple hypothesis tracking (mht) for space surveillance: Theoretical framework," *Advances in the Astronautical Sciences*, vol. 150, pp. 55–74, Jan. 2014.

[28] N. Singh, J. T. Horwood, J. M. Aristoff, A. B. Poore, C. Sheaff, and M. K. Jah, "Multiple hypothesis tracking (mht) for space surveillance: Results and simulation studies," in *Advanced Maui Optical and Space Surveillance Tech. Conf.(AMOS)*, 2013.

[29] S. M. Gadaleta, J. T. Horwood, and A. B. Poore, "Short arc gating in multiple hypothesis tracking for space surveillance," in *Sensors and Systems for Space Applications V*, K. D. Pham, J. L. Cox, R. T. Howard, and H. Zmuda, Eds., International Society for Optics and Photonics, vol. 8385, SPIE, 2012, pp. 248–262.

[30] J. Stauch, T. Bessell, M. Rutten, J. Baldwin, M. Jah, and K. Hill, "Joint probabilistic data association and smoothing applied to multiple space object tracking," *Journal of Guidance, Control, and Dynamics*, vol. 41, no. 1, pp. 19–33, 2018.

[31] S. Oh, S. Russell, and S. Sastry, "Markov chain monte carlo data association for multi-target tracking," *IEEE Transactions on Automatic Control*, vol. 54, no. 3, pp. 481–497, 2009.

[32] B. T. Vo, "Random finite sets in multi-object filtering," 2008.

[33] R. P. S. Mahler, "Random-set approach to data fusion," in *Automatic Object Recognition IV*, F. A. Sadjadi, Ed., International Society for Optics and Photonics, vol. 2234, SPIE, 1994, pp. 287–295.

[34] D. J. Daley, D. Vere-Jones, *et al.*, *An introduction to the theory of point processes: volume I: elementary theory and methods*. Springer, 2003.

[35] D. J. Daley and D. Vere-Jones, *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media, 2007.

[36] M. Jacobsen and J. Gani, "Point process theory and applications: Marked point and piecewise deterministic processes," 2006.

[37] R. Mahler, "Multitarget bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, 2003.

[38] B. A. Jones, D. S. Bryant, B.-T. Vo, and B.-N. Vo, "Challenges of multi-target tracking for space situational awareness," in *2015 18th International Conference on Information Fusion (Fusion)*, 2015, pp. 1278–1285.

[39] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, 2013.

[40] Y. Punchihewa, "Efficient generalized labeled multi-bernoulli filter for jump markov system," in *2017 International Conference on Control, Automation and Information Sciences (ICCAIS)*, 2017, pp. 221–226.

[41] N. Ravago and B. A. Jones, "Tracking multiple maneuvering satellites using a generalized labeled multi-bernoulli filter," in *AIAA/AAS Astrodynamics Specialist Conference*, 2019.

[42] M. Yuan, W. Shen, J. Li, Y. Pavlidis, and S. Li, "Auction/belief propagation algorithms for constrained assignment problem," in *Algorithms and Discrete Applied Mathematics*, S. Ganguly and R. Krishnamurti, Eds., Cham: Springer International Publishing, 2015, pp. 238–249.

[43] J. L. Williams, "Marginal multi-bernoulli filters: Rfs derivation of mht, jipda, and association-based member," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 3, pp. 1664–1687, 2015.

[44] S. Blackman, "Multiple hypothesis tracking for multiple target tracking," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 5–18, 2004.

[45] T. E. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Multi-target tracking using joint probabilistic data association," in *1980 19th IEEE Conference on Decision and Control including the Symposium on Adaptive Processes*, 1980, pp. 807–812.

[46] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.

[47] S. J. Julier and J. K. Uhlmann, "New extension of the kalman filter to nonlinear systems," in *Signal processing, sensor fusion, and target recognition VI*, International Society for Optics and Photonics, vol. 3068, 1997, pp. 182–193.

[48] I. Arasaratnam and S. Haykin, "Cubature kalman filters," *IEEE Transactions on automatic control*, vol. 54, no. 6, pp. 1254–1269, 2009.

[49] H. Jiang and Y. Cai, "Adaptive fifth-degree cubature information filter for multi-sensor bearings-only tracking," *Sensors*, vol. 18, no. 10, p. 3241, 2018.

[50] D. Meng, L. Miao, H. Shao, and J. Shen, "A seventh-degree cubature kalman filter," *Asian Journal of Control*, vol. 20, no. 1, pp. 250–262, 2018.

[51] Y.-z. Luo and Z. Yang, "A review of uncertainty propagation in orbital mechanics," *Progress in Aerospace Sciences*, vol. 89, pp. 23–39, 2017.

[52] D. G. Robertson, J. H. Lee, and J. B. Rawlings, "A moving horizon-based approach for least-squares estimation," *AIChE Journal*, vol. 42, no. 8, pp. 2209–2224, 1996.

[53] M. K. Transtrum and J. P. Sethna, *Geodesic acceleration and the small-curvature approximation for nonlinear least squares*, 2012.

[54] A. B. Poore *et al.*, "Covariance and uncertainty realism in space surveillance and tracking," *Report of the Air Force Space Command Astrodynamics Innovation Committee*, vol. 27, 2016.

[55] A. Cano, A. Pastor, D. Escobar, J. Míguez, and M. Sanjurjo-Rivo, "Covariance determination for improving uncertainty realism in orbit determination and propagation," *Advances in Space Research*, vol. 72, no. 7, pp. 2759–2777, 2023, Space Environment Management and Space Sustainability.

[56] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman filter: Particle filters for tracking applications*. Artech house, 2003.

[57] L. Martino, V. Elvira, and F. Louzada, "Effective sample size for importance sampling based on discrepancy measures," *Signal Processing*, vol. 131, pp. 386–401, 2017.

[58] T. Bengtsson, P. Bickel, and B. Li, "Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems," *Probability and Statistics: Essays in Honor of David A. Freedman*, pp. 316–334, 2008.

[59] C. Musso, N. Oudjane, and F. Gland, "Improving regularized particle filters," in *Sequential Monte Carlo Methods in Practice*, Springer, New York, Jan. 2001.

[60] B. W. Silverman, *Density estimation for statistics and data analysis*. Routledge, 2018.

[61] W. K. Hastings, "Monte carlo sampling methods using markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.

[62] A. E. Gelfand and A. F. M. Smith, "Sampling-based approaches to calculating marginal densities," *Journal of the American Statistical Association*, vol. 85, no. 410, pp. 398–409, 1990.

[63] J. A. Vrugt, C. Ter Braak, C. Diks, B. A. Robinson, J. M. Hyman, and D. Higdon, "Accelerating markov chain monte carlo simulation by differential evolution with self-adaptive randomized subspace sampling," *International journal of nonlinear sciences and numerical simulation*, vol. 10, no. 3, pp. 273–290, 2009.

[64] B. Minasny, J. A. Vrugt, and A. B. McBratney, "Confronting uncertainty in model-based geostatistics using markov chain monte carlo simulation," *Geoderma*, vol. 163, no. 3, pp. 150–162, 2011.

[65] D. Lu, D. Ricciuto, A. Walker, C. Safta, and W. Munger, "Bayesian calibration of terrestrial ecosystem models: A study of advanced markov chain monte carlo methods," *Biogeosciences*, vol. 14, no. 18, pp. 4295–4314, 2017.

[66] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, "Hybrid monte carlo," *Physics Letters B*, vol. 195, no. 2, pp. 216–222, 1987.

[67] M. D. Hoffman and A. Gelman, *The no-u-turn sampler: Adaptively setting path lengths in hamiltonian monte carlo*, 2011.

[68] G. Pulford, "A survey of manoeuvring target tracking methods," Dec. 1998.

[69] X. Rong Li and V. P. Jilkov, "Survey of maneuvering target tracking. part i. dynamic models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333–1364, 2003.

[70] X. Rong Li and V. P. Jilkov, "Survey of maneuvering target tracking. part v. multiple-model methods," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1255–1321, 2005.

[71] X. Rong Li and V. P. Jilkov, "A survey of maneuvering target tracking-part iv: Decision-based methods," in *SPIE proceedings series*, 2002, pp. 511–534.

[72] X. Wang, T. Li, S. Sun, and J. M. Corchado, "A survey of recent advances in particle filters and remaining challenges for multitarget tracking," *Sensors*, vol. 17, no. 12, p. 2707, 2017.

[73] P. Shi and F. Li, "A survey on markovian jump systems: Modeling and design," *International Journal of Control, Automation and Systems*, vol. 13, no. 1, pp. 1–16, 2015.

[74] Y. bar-shalom, S. Challa, and H. Blom, "Imm estimator versus optimal estimator for hybrid systems," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 41, pp. 986–991, Aug. 2005.

[75] W. D. Blair, G. A. Watson, and T. R. Rice, "Tracking maneuvering targets with an interacting multiple model filter containing exponentially correlated acceleration models," in *[1991 Proceedings] The Twenty-Third Southeastern Symposium on System Theory*, 1991, pp. 224–228.

[76] R. Torelli, A. Graziano, and A. Farina, "Im3ht algorithm: A joint formulation of imm and mht for multitarget tracking," in *1997 European Control Conference (ECC)*, 1997, pp. 2979–2984.

[77] H. A. P. Blom and E. A. Bloem, "Combining imm and jpda for tracking multiple maneuvering targets in clutter," in *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002. (IEEE Cat.No.02EX5997)*, vol. 1, 2002, 705–712 vol.1.

[78] K. Yoo, D. H. Won, S. Sung, and Y.-J. Lee, "Hybrid tracking of maneuvering multiple-aircraft in 3d space," *ICCAS-SICE 2009 - ICROS-SICE International Joint Conference 2009, Proceedings*, Jan. 2009.

[79] G. Yuan, W. Zhu, W. Wang, and B. Yin, "Maneuvering target tracking algorithm based on interacting multiple models," *Mathematical Problems in Engineering*, vol. 2015, pp. 1–7, May 2015.

[80] E. Hajiramezanali, S. Fouladi, and H. Amindavar, "A novel maneuvering target tracking approach by stochastic volatility garch model," Feb. 2019.

[81] B. Li, Z. Shi, and J. Chen, "Residual-feedback particle filter for maneuvering target tracking," in *2010 6th International Conference on Wireless Communications Networking and Mobile Computing (WiCOM)*, 2010, pp. 1–4.

[82] C. Andrieu, M. Davy, and A. Doucet, "Efficient particle filtering for jump markov systems. application to time-varying autoregressions," *IEEE Transactions on Signal Processing*, vol. 51, no. 7, pp. 1762–1770, 2003.

[83] H. A. P. Blom and E. A. Bloem, "Exact bayesian and particle filtering of stochastic hybrid systems," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 1, pp. 55–70, 2007.

[84] D. Blackwell, "Conditional expectation and unbiased sequential estimation," *Ann. Math. Statist.*, vol. 18, no. 1, pp. 105–110, Mar. 1947.

[85] J. Huang, W.-D. Hu, Q. Xin, and X.-Y. Du, "An object correlation and maneuver detection approach for space surveillance," *Research in Astronomy and Astrophysics*, vol. 12, no. 10, pp. 1402–1416, Sep. 2012.

[86] Z. Yang, Y. Luo, and J. Zhang, "Nonlinear semi-analytical uncertainty propagation of trajectory under impulsive maneuvers," *Astrodynamics*, vol. 3, pp. 61–77, Mar. 2019.

[87] A. Pastor, G. Escribano, M. Sanjurjo-Rivo, and D. Escobar, "Satellite maneuver detection and estimation with optical survey observations," *The Journal of the Astronautical Sciences*, vol. 69, no. 3, pp. 879–917, 2022.

[88] L. Porcelli *et al.*, "Satellite maneuver detection and estimation with radar survey observations," *Acta Astronautica*, vol. 201, pp. 274–287, 2022.

[89] B. Jia *et al.*, "Space object tracking and maneuver detection via interacting multiple model cubature kalman filters," *IEEE Aerospace Conference Proceedings*, vol. 2015, Jun. 2015.

[90] J. D. Katzovitz, "Space-based maneuver detection and characterization using multiple model adaptive estimation," M.S. thesis, 2018.

[91] Q. M. Lam, "Resident space object tracking using an interacting multiple model mixing scheme," in *Sensors and Systems for Space Applications VII*, K. D. Pham and J. L. Cox, Eds., International Society for Optics and Photonics, vol. 9085, SPIE, 2014, pp. 218–232.

[92] T. Kelecy, D. Hall, K. Hamada, and D. Stocker, "Satellite maneuver detection using two-line element (tle) data," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui Economic Development Board (MEDB) Maui, HA, 2007.

[93] J. Siminski, T. Flohrer, and T. Schildknecht, "Assessment of post-maneuver observation correlation using short-arc tracklets," *Journal of the British Interplanetary Society*, vol. 70, pp. 63–68, 2017.

[94] M. Holzinger and D. Scheeres, "Object correlation and maneuver detection using optimal control performance metrics," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference E*, vol. 26, 2010.

[95] N. Singh, J. T. Horwood, and A. B. Poore, "Space object maneuver detection via a joint optimal control and multiple hypothesis tracking approach," in *Proceedings of the 22nd AAS/AIAA Space Flight Mechanics Meeting*, Univelt San Diego, CA, vol. 143, 2012, pp. 2012–159.

[96] D. P. Lubey, "Maneuver detection and reconstruction in data sparse systems with an optimal control based estimator," Ph.D. dissertation, 2015.

[97] M. I. Voitsekhovskii, "Metric," *Encyclopedia of mathematics*,

[98] M. J. Holzinger, D. J. Scheeres, and K. T. Alfriend, "Object correlation, maneuver detection, and characterization using control distance metrics," *Journal of Guidance, Control, and Dynamics*, vol. 35, no. 4, pp. 1312–1325, 2012.

[99] R. Serra, C. Yanez, and C. Frueh, "Tracklet-to-orbit association for maneuvering space objects using optimal control theory," *Acta Astronautica*, vol. 181, pp. 271–281, 2021.

[100] R. Serra, C. Yanez, and E. Delande, "Tracklet-to-orbit association under uncertainty applied to maneuvering space objects," *Acta Astronautica*, 2022.

[101] L. Pirovano and R. Armellin, "Detection and estimation of spacecraft maneuvers for catalog maintenance," 2022.

[102] J. F. Jordan, *Application of Lambert's Theorem to the Solution of Interplanetary Transfer Problems*. Jet Propulsion Laboratory, California Institute of Technology, 1964.

[103] C. F. Gauss, *Theory of the motion of the heavenly bodies moving about the sun in conic sections, a translation of Gauss's Theoria motus*. Little Brown and Company, Boston, 1857.

[104] E. Lancaster, R. Blanchard, and R. Devaney, "A note on lambert's theorem.," *Journal of Spacecraft and Rockets*, vol. 3, no. 9, pp. 1436–1438, 1966.

[105] R. H. Battin, "Lambert's problem revisited," *AIAA Journal*, vol. 15, no. 5, pp. 707–713, 1977.

[106] R. H. Gooding, "A procedure for the solution of lambert's orbital boundary-value problem," *Celestial Mechanics and Dynamical Astronomy*, vol. 48, no. 2, pp. 145–165, 1990.

[107] G. Avanzini, "A simple lambert algorithm," *Journal of guidance, control, and dynamics*, vol. 31, no. 6, pp. 1587–1594, 2008.

[108] D. Izzo, "Revisiting lambert's problem," *Celestial Mechanics and Dynamical Astronomy*, vol. 121, pp. 1–15, 2015.

[109] R. P. Russell, "Complete lambert solver including second-order sensitivities," *Journal of Guidance, Control, and Dynamics*, vol. 45, no. 2, pp. 196–212, 2022.

[110] A. Albouy, "Lambert's theorem: Geometry or dynamics?" *Celestial Mechanics and Dynamical Astronomy*, vol. 131, no. 9, Aug. 2019.

[111] S. P. Altman and J. S. Pistiner, "Minimum velocity increment solution for two-impulse coplanar orbital transfer," *AIAA Journal*, vol. 1, no. 2, pp. 435–442, 1963.

[112] M. Avendano and D. Mortari, "A closed-form solution to the minimum lambert's problem," *Celestial Mechanics and Dynamical Astronomy*, vol. 106, no. 1, p. 25, 2010.

[113] J. McElreath, M. Majji, S. Kumar Singh, and J. L. Junkins, "Optimal bi-impulse orbital transfers: Station keeping applications," *Journal of Guidance, Control, and Dynamics*, vol. 44, no. 11, pp. 2057–2066, 2021.

[114] G. A. McCue and D. F. Bender, "Numerical investigation of minimum impulse orbital transfer.," *AIAA Journal*, vol. 3, no. 12, pp. 2328–2334, 1965.

[115] A. Pastor, G. Escribano, M. Sanjurjo-Rivo, and D. Escobar, "Satellite maneuver detection and estimation with optical survey observations," *The Journal of the Astronautical Sciences*, vol. 69, no. 3, pp. 879–917, 2022.

[116] L. Porcelli *et al.*, "Satellite maneuver detection and estimation with radar survey observations," *Acta Astronautica*, vol. 201, pp. 274–287, 2022.

[117] R. Bhatia *et al.*, "The operational realities of electric propulsion," 2022.

[118] A. E. Petropoulos and J. M. Longuski, "Shape-based algorithm for the automated design of low-thrust, gravity assist trajectories," *Journal of Spacecraft and Rockets*, vol. 41, no. 5, pp. 787–796, 2004.

[119] J. Roa, J. Peláez, and J. Senent, "New analytic solution with continuous thrust: Generalized logarithmic spirals," *Journal of Guidance, Control, and Dynamics*, vol. 39, no. 10, pp. 2336–2351, 2016.

[120] J. Roa and J. Pelaez, "Three-dimensional generalized logarithmic spirals," in *Proceedings of the 26th AAS/AIAA Space Flight Mechanics Meeting, Napa, CA, USA*, 2016, pp. 14–18.

[121] J. S. Hudson and D. J. Scheeres, "Reduction of low-thrust continuous controls for trajectory dynamics," *Journal of Guidance, Control, and Dynamics*, vol. 32, no. 3, pp. 780–787, 2009.

[122] C. A. Kluever and S. R. Oleson, "Direct approach for computing near-optimal low-thrust earth-orbit transfers," *Journal of Spacecraft and Rockets*, vol. 35, no. 4, pp. 509–515, 1998.

[123] J. S. Hudson, "Reduction of low-thrust continuous controls for trajectory dynamics and orbital targeting," Ph.D. dissertation, University of Michigan, 2010.

[124] D. Brouwer, "Solution of the problem of artificial satellite theory without drag," Yale University, New Haven, United States, Tech. Rep., 1959.

[125] Y. Kozai, "The motion of a close earth satellite," *The Astronomical Journal*, vol. 64, p. 367, 1959.

[126] S. J. Setty, P. J. Cefola, O. Montenbruck, and H. Fiedler, "Application of semi-analytical satellite theory orbit propagator to orbit determination for space object catalog maintenance," *Advances in Space Research*, vol. 57, no. 10, pp. 2218–2233, 2016.

[127] F. R. Hoots and R. L. Roehrich, "Models for propagation of norad element sets," Aerospace Defense Command Peterson AFB CO Office of Astrodynamics, Tech. Rep., 1980.

[128] M. J. Holzinger and D. J. Scheeres, "Reachability results for nonlinear systems with ellipsoidal initial sets," in 2. 2012, vol. 48, pp. 1583–1600.

[129] M. J. Holzinger, D. J. Scheeres, and R. S. Erwin, "On-orbit operational range computation using gauss's variational equations with j2 perturbations," *Journal of Guidance, Control, and Dynamics*, vol. 37, no. 2, pp. 608–622, 2014.

[130] P. R. Patel and D. J. Scheeres, "Rapid and automatic reachability estimation of electric propulsion spacecraft," *The Journal of the Astronautical Sciences*, vol. 70, no. 6, p. 45, 2023.

[131] D. Aguilar Marsillach and M. J. Holzinger, "Spacecraft custody maintenance and maneuver detection using robotic telescopes and reachable sets," *Journal of Guidance, Control, and Dynamics*, vol. 44, no. 4, pp. 667–683, 2021.

[132] Z. Hall, P. Singla, and K. Johnson, "Reachability-based search for tracking of noncooperative maneuvering satellites in data sparse environment," *The Journal of the Astronautical Sciences*, vol. 70, no. 2, p. 9, 2023.

[133] M. Althoff, O. Stursberg, and M. Buss, "Reachability analysis of nonlinear systems with uncertain parameters using conservative linearization," in *2008 47th IEEE Conference on Decision and Control*, 2008, pp. 4042–4048.

[134] I. F. D. Oliveira and R. H. C. Takahashi, "An enhancement of the bisection method average performance preserving minmax optimality," *ACM Trans. Math. Softw.*, vol. 47, no. 1, Dec. 2020.

[135] R. Fletcher and M. J. Powell, "A rapidly convergent descent method for minimization," *The computer journal*, vol. 6, no. 2, pp. 163–168, 1963.

[136] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The computer journal*, vol. 7, no. 4, pp. 308–313, 1965.

[137] P. J. Van Laarhoven, E. H. Aarts, P. J. van Laarhoven, and E. H. Aarts, *Simulated annealing*. Springer, 1987.

[138] N. V. Findler, C. Lo, and R. Lo, "Pattern search for optimization," *Mathematics and computers in simulation*, vol. 29, no. 1, pp. 41–50, 1987.

[139] F. Neri and S. Rostami, "Generalised pattern search based on covariance matrix diagonalisation," *SN Computer Science*, vol. 2, no. 3, p. 171, 2021.

[140] V. Torczon, "On the convergence of pattern search algorithms," *SIAM Journal on Optimization*, vol. 7, no. 1, pp. 1–25, 1997.

[141] H. Hong and F. Hickernell, "Algorithm 823: Implementing scrambled digital sequences.," *ACM Trans. Math. Softw.*, vol. 29, pp. 95–109, Jan. 2003.

[142] R. Abay, S. Gehly, S. Balage, M. Brown, and R. Boyce, "Maneuver detection of space objects using generative adversarial networks," in *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2018.

[143] P. Romero and J. M. Gambi, "Optimal control in the east/west station-keeping manoeuvres for geostationary satellites," *Aerospace Science and Technology*, vol. 8, no. 8, pp. 729–734, 2004.

[144] S. G. Team, "Starlink conjunction avoidance with crewed space stations," SpaceX, Technical Memo SAT-51385, Apr. 2022.

[145] B.-S. Lee, J.-S. Lee, and K.-H. Choi, "Analysis of a station-keeping maneuver strategy for collocation of three geostationary satellites," *Control Engineering Practice*, vol. 7, no. 9, pp. 1153–1161, 1999.

[146] T. K. Moon, "The expectation-maximization algorithm," *IEEE Signal Processing Magazine*, vol. 13, no. 6, pp. 47–60, 1996.

[147] B. Bowman, W. K. Tobiska, F. Marcos, C. Huang, C. Lin, and W. Burke, "A new empirical thermospheric density model jb2008 using new solar and geomagnetic indices," in *AIAA/AAS astrodynamics specialist conference and exhibit*, 2008, p. 6438.

[148] J. M. Picone, A. E. Hedin, D. P. Drob, and A. C. Aikin, "Nrlmsise-00 empirical model of the atmosphere: Statistical comparisons and scientific issues," *Journal of Geophysical Research: Space Physics*, vol. 107, no. A12, SIA 15-1-SIA 15–16, 2002.

[149] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.

[150] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, pp. 25–57, 2006.

[151] M. Matsumoto and T. Nishimura, "Mersenne twister: A 623 dimensionally equidistributed uniform pseudo-random number generator," *ACM Trans. Model. Comput. Simul.*, vol. 8, no. 1, pp. 3–30, Jan. 1998.

[152] G. Escribano, M. Sanjurjo-Rivo, J. Siminski, A. Pastor, and D. Escobar, "Automatic maneuver detection and tracking of space objects in optical survey scenarios based on stochastic hybrid systems formulation," *Advances in Space Research*, vol. 69, no. 9, pp. 3460–3477, 2022.

[153] A. Doucet, A. M. Johansen, *et al.*, "A tutorial on particle filtering and smoothing: Fifteen years later," *Handbook of nonlinear filtering*, vol. 12, no. 656-704, p. 3, 2009.

[154] B. A. Jones and R. Weisman, "Multi-fidelity orbit uncertainty propagation," *Acta Astronautica*, vol. 155, pp. 406–417, 2019.

[155] S. Yun, R. Zanetti, and B. A. Jones, "Kernel-based ensemble gaussian mixture filtering for orbit determination with sparse data," *Advances in Space Research*, vol. 69, no. 12, pp. 4179–4197, 2022.

[156] G. Escribano, M. Sanjurjo-Rivo, J. Siminski, A. Pastor, and D. Escobar, "A filter for tracking non-cooperative low-thrust satellites using surveillance radar data," *Acta Astronautica*, vol. 213, pp. 694–707, 2023.

[157] A. E. Petropoulos, "Refinements to the q-law for the low-thrust orbit transfers," 2005.

[158] C. Perez Hernandez, M. A. Ramos Prada, M. Alessandrini, D. Cano Mañanes, and J. Siminski, "Operation review and analysis of the s3t surveillance radar," in *8th European Conference on Space Debris*, 2021.

[159] N. Z. Miura, "Comparison and design of simplified general perturbation models (sgp4) and code for nasa johnson space center, orbital debris program office," 2009.

[160] P. Escobal, *Methods of Orbit Determination*. J. Wiley, 1965.

[161] J. Siminski, "Techniques for assessing space object cataloguing performance during design of surveillance systems," in *6th International Conference on Astrodynamics Tools and Techniques (ICATT)*, 2016, pp. 14–17.

[162] A. Pastor, M. Sanjurjo-Rivo, and D. Escobar, "Initial orbit determination methods for track-to-track association," *Advances in Space Research*, vol. 68, no. 7, pp. 2677–2694, 2021.

[163] S. M. Lenz, H. G. Bock, J. P. Schlöder, E. A. Kostina, G. Gienger, and G. Ziegler, "Multiple shooting method for initial satellite orbit determination," *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 5, pp. 1334–1346, 2010.

[164] M. Losacco *et al.*, "Initial orbit determination with the multibeam radar sensor birales," *Acta Astronautica*, vol. 167, pp. 374–390, 2020.

[165] G. Escribano, B. Jones, M. Sanjurjo-Rivo, J. Siminski, A. Pastor, and D. Escobar, "A glmb filter for space objects with control metric based maneuver detection," in *Proceedings of AAS/AIAA Spaceflight Mechanics Meeting*, 2023.

[166] K. Murty, "An algorithm for ranking all assignments in order of increasing cost," *Operations Research*, vol. 16, pp. 682–687, Jun. 1968.

[167] C. Shim, B.-T. Vo, B.-N. Vo, J. Ong, and D. Moratuwage, "Linear complexity gibbs sampling for generalized labeled multi-bernoulli filtering," *IEEE Transactions on Signal Processing*, vol. 71, pp. 1981–1994, 2023.

[168] J. T. Horwood, N. D. Aragon, and A. B. Poore, "Gaussian sum filters for space surveillance: Theory and simulations," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1839–1851, 2011.

[169] K. J. DeMars, R. H. Bishop, and M. K. Jah, "Entropy-based approach for uncertainty propagation of nonlinear dynamical systems," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 4, pp. 1047–1057, 2013.

[170] K. J. DeMars and M. K. Jah, "Probabilistic initial orbit determination using gaussian mixture models," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 5, pp. 1324–1335, 2013.

[171] G. M. Goff, J. T. Black, and J. A. Beck, "Orbit estimation of a continuously thrusting spacecraft using variable dimension filters," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 12, pp. 2407–2420, 2015.

[172] B.-N. Vo and B.-T. Vo, "A multi-scan labeled random finite set model for multi-object state estimation," *IEEE Transactions on Signal Processing*, vol. 67, no. 19, pp. 4948–4963, 2019.

[173] B. Ristic, B.-N. Vo, D. Clark, and B.-T. Vo, "A metric for performance evaluation of multi-target tracking algorithms," *IEEE Transactions on Signal Processing*, vol. 59, no. 7, pp. 3452–3457, 2011.