Available online at www.sciencedirect.com

ScienceDirect

Advances in Space Research xx (2022) xxx-xxx

ADVANCES IN SPACE RESEARCH (a COSPAR publication)

www.elsevier.com/locate/asr

Manoeuvre detection in Low Earth Orbit with Radar Data

Jose M. Montilla[a], Julio C. Sanchez[a], Rafael Vazquez[a], Jorge Galan-Vioque[a], Javier Rey Benayas[b], Jan Siminski[c]

[a] *Universidad de Sevilla, Escuela Técnica Superior de Ingeniería, Camino de los Descubrimientos s.n 41092 Sevilla, Spain*
[b] *Indra Sistemas, Crta Loeches 9, 28850 Torrejón de Ardoz, Madrid, Spain*
[c] *Space Debris Office, ESA/ESOC, Darmstadt, Germany*

Abstract

This work outlines and assesses several methods for the detection of manoeuvres in Low Earth Orbit (LEO) from surveillance radar data. To be able to detect manoeuvres, the main starting assumption is that the object under analysis has an orbit known with a sufficient degree of precision. Based on the precise (a posteriori) orbit and radar data, several manoeuvre detection methods are presented; one is based on unscented Kalman filtering, whereas two others algorithms are based on reachability analysis of the state, which correlates its prediction set with the next track from the radar. The filtering algorithm can be extended for several radar tracks, whereas the reachability-based methods are more precise in detecting manoeuvres. Then, to inherit the best properties of both classes of algorithms, a manoeuvre detection filter that combines both concepts is finally presented. Manoeuvre detection results are analysed first for simulated scenarios—for validation and calibration purposes—and later for real data. Radar information comes from the Spanish Space Surveillance Radar (S3TSR), with real manoeuvre information and high-quality ephemerides. The results show promise, taking into account that a single surveillance radar is the only source of data, obtaining manoeuvre detection rates of more than 50% and false positive rates of less than 10%.

1. Introduction

In the field of Space Surveillance and Tracking (SST), accurate orbital determination and manoeuvre detection is of upmost importance to infer object's orbital information and their future behaviour, as well as to be able to carry out tasks such as prediction of potential conjunctions with operating satellites, taking avoidance orbital corrections, predicting re-entries, identifying fragmentations or updating orbital elements of known satellites, among others.

Satellites performing unknown manoeuvres pose a challenge when trying to associate the new collected observations (obtained by laser, radar, or by any other means from the SST infrastructure) with the previously known reference orbits (which are stored in SST catalogues). Indeed, one of the main motivations of manoeuvre detection is that it can significantly reduce the number of uncorrelated targets detected by the SST sensors infrastructure. Most of these uncorrelated objects are known satellites, which have performed unpublished manoeuvres, in such a way that their new orbits do not match with the predictions.

Email addresses: jmontillag@us.es (Jose M. Montilla), jsanchezm@us.es (Julio C. Sanchez), rvazquez1@us.es (Rafael Vazquez), jgv@us.es (Jorge Galan-Vioque), jreyb@indra.es (Javier Rey Benayas), Jan.Siminski@esa.int (Jan Siminski)

https://dx.doi.org/10.1016/j.jasr.xxxx.xx.xxx
0273-1177/

This work develops several methods for the detection of manoeuvres in Low Earth Orbit (LEO) from radar data, providing first some preliminary numerical initial results obtained from simulated orbits and radar data. Since the final aim is to integrate these algorithms in the S3T (Spanish Space Surveillance Tracking) Cataloguing System in order to provide routine automatic manoeuvre detection capabilities to the system in the future, a validation of all the algorithms is carried out with real tracks from S3TSR (Gomez et al., 2019), the Spanish surveillance radar developed, installed and validated by Indra with the funding of the Spanish Government under the technical and contractual management of ESA on behalf of Centro de Desarrollo Tecnológico e Industrial (CDTI). Manoeuvre information and ephemerides are obtained from ESA/ESOC and DLR/GSOC to assess the results, for several scenarios. The methods were implemented by making use of the space-dynamics library OREKIT (see Mainsonobe et al. (2012)).

Maneuver detection has received increased attention in recent times and, correspondingly, diverse results have already appeared in the literature. For instance, Jaunzemis et al. (2016) consider Gaussian-mixture models and binary hypothesis testing; manoeuvre detection is developed using Mahalanobis distance and control distance metrics, and tested for geostationary satellites with optical measurements. In Wang et al. (2021), using an in-house simulated space catalog environment, the Gaussian Binary Classification method is used to detect several types of maneuvers; results are validated only by simulation. In Escribano et al. (2022), a Markov Chain Monte Carlo sampling scheme is used to test the post-maneuver observation in terms of control distance metric, the result is tested for geostationary satellites with optical measurements. Both Yu et al. (2021) and Clark & Lee (2020) use TLE analysis; the first work employs a dynamical model of the manoeuvre (considering intense manoeuvres), whereas the second uses parallel processing to analyze TLE historical data. Finally, in Serra et al. (2021), an approach combining optimal control theory and admissible regions is developed and tested for geostationary satellites with optical measurements.

The contribution of this paper, compared with other existing works in the literature, is the development of algorithms and metrics that can work with a single surveillance radar, a situation rarely considered in the previous literature; having a single radar, one has a rather small number of measurements (tracks are usually 5-20 plots) and large gaps in between tracks (1 to 3 days). In addition, the considered set of satellites are in Low Earth Orbit (thus subject to considerable perturbations), and perform moderately small manoeuvres, compared with most analysis. Based on previous works, we use two manoeuvre detection metrics and formulate two orbit determination filters with manoeuvre detection capabilities; additionally, we develop a novel method to interpret the control distance metric based on percentiles. The algorithms are tested not only in simulation, but also against real data for several LEO satellites. Part of these results were already presented in conference form in Vazquez et al. (2021a,b).

The structure of this paper is as follows. After this brief introduction, a literature review is performed for the two main families of methods that can be used to detect manoeuvres, namely: Kalman filters (based on orbit determination approaches) in Section 2, and reachability analysis-based methods (which compare reachable predicted sets with obtained measurements) in Section 3. The particular implementations selected for this work are presented in Sections 4 and 5, respectively for each family, together with some preliminary proof-of-concept results. Then, in Section 6 a manoeuvre detection filter based on a combination of filtering and reachability is developed. Next validation results are presented for simulated (Section 7) and real scenarios (Section 8). The paper is concluded in Section 9 with some final remarks and future work.

2. Manoeuvre detection filters

Manoeuvre detection filters (MDFs) employ orbit determination in the process of detecting if some manoeuvre has been performed; they are quite useful, since they are able to correlate new (post-manoeuvre) orbits with previous (pre-manoeuvre) known

orbits, thus paving the way to perform orbit determination using quite fewer measures (as compared to a conventional orbit determination problem). This fact is shown, for instance, in Goff (2015) where the author compares the accuracy and cost of orbit determination using a manoeuvre detection filter on known flying objects, versus a conventional (cold-started) orbit determination procedure. At its simplest, a manoeuvre detection algorithm relies on Statistical Orbit Determination (SOD) methods (see Schutz et al. (2004) for a general overview). Much has been written about the estimation and tracking of spacecraft using radar (or laser) measures, and classical methods like the Batch Least Squares (BLS) method, Extended Kalman Filter (EKF), and the Unscented Kalman Filter (UKF)—and even non-gaussian techniques like the Particle Filter—are well known in the literature. See Vetter (2007) for a brief revision on the historical and contemporary methods for orbit determination.

The problem with traditional methods arises when the target performs unknown manoeuvres in-between the measurements windows. Then, the propagated orbit in which the estimation methods are based may become too inaccurate (since they do not take into account the manoeuvres) as they have become "overconfident" due to their covariance becoming too small (a filter exhibiting such behaviour is known as a "smug" filter). Thus, a manoeuvring scenario may lead to severe outliers and convergence problems in conventional filtering techniques; there exist methods to handle these issues and enhance the robustness of these classical algorithms, to avoid divergence problems. For instance, some possible techniques are covariance inflation or fading memory (Goff, 2015), among others (Jiang et al., 2019). These mechanisms would be activated if a manoeuvre is detected; using the filter residuals, a decision logic can be put in place (Goff, 2015; Guang et al., 2018) to estimate when a manoeuvre has been performed.

Although not considered for this work there exist solutions with extra layers of complexity for problems with highly unstructured uncertainty. These are the Multiple Model filters, suited for tracking problems and based on a family of elemental filters which can be designed to model different aspects of the system behaviour, together with a probabilistic mixing logic to select the best estimation combining all the outputs of the elemental filters (Li & Jilkov, 2005).

2.1. State of the art on manoeuvre detection filters

Next, a representative sample of MDFs from the literature are analysed.

Woodburn et al. (2003) presents a fixed interval smoother for manoeuvre reconstruction; this algorithm gives a simple approach that consists of a sequential filter used to move forward across the manoeuvre and a fixed interval smoother to move backwards across the manoeuvre. The sequential filter serves to process all the tracking data prior to the manoeuvre to provide an optimal pre-manoeuvre state estimate and covariance. Radar data is processed after the time of the manoeuvre until the uncertainty in the state estimate returns to a normal non-manoeuvre condition. The difference between the post-manoeuvre and pre-manoeuvre smoothed states may now be extracted as the estimate of the impulse (with covariance information used to characterize manoeuvre uncertainty). A drawback of this method is that it is assumed, without guarantees, that the filter will quickly converge on a good state estimate after a manoeuvre has passed.

A variable structure estimator is proposed in Guang et al. (2018), where a manoeuvre detection metric is used to design an estimator with an additional manoeuvre observer module. In this scheme an EKF is used together with a manoeuvre observer (which is in turn triggered when the manoeuvre detection metric reaches a certain threshold). Then, the manoeuvre observer estimates the manoeuvre acceleration, and sends that information to the EKF, which takes into account the estimated acceleration to improve the orbit propagation in its algorithm. The "manoeuvre observer" is based on a simple first-order observer, which produces an estimation of the acceleration to be fed back to the EKF. The method is simple, but assumes that radar measurements are always available (with a frequency of 5 Hz) so that manoeuvres are always observed; that would require having data from a very large radar network.

A joint kinematic/dynamic filter is proposed in Ye et al. (2021) (with simulations of three radars and a satellite in a Medium Earth Orbit) and consists of two filters running in parallel. The first and main filter is a traditional orbit determination Kalman filter whereas the second filter is a kinematic filter and utilizes some representative random processes (with design parameters) to describe the orbital motion. While the detailed motion is not captured at all, the changes caused by orbital manoeuvres can be captured by those flexible random processes, although it requires a complex tuning.

In Goff et al. (2015b) (see also Goff et al. (2015a)) the results show that a filter-through Interacting Multiple Model orbit determination filter (EKF or UKF) can converge on a post-manoeuvre orbit with similar performance to Initial Orbit Determination (IOD) approaches, based on multiple filters running with different levels of covariance inflation. Once the post-manoeuvre orbit is known with a certain degree of accuracy, to reconstruct a single manoeuvre, one determines the time when the orbits intersected or became very close.

3. Reachability analysis

In this section the idea of Reachable Sets (RS) and their analysis (Reachability Analysis, RA) is introduced, as well as their relationship with control and estimation of systems. The technique has been used, for instance, in the context of rendezvous of spacecraft (Sanchez et al., 2019), and has many applications in the area of safety for trajectories of vehicles (Xu et al., 2019). This background material constitutes the foundational framework for Section 5.

Citing Holzinger & Scheeres (2009), "the concept of reachability is central to Space Situational Awareness (SSA)", which underscores the interest of this concept for the present work. Reachability Analysis deals with the study and applications of Reachable Sets, which are defined as follows: given a system that evolves from an initial condition (or set of initial conditions), and possibly has some control inputs, the reachable set is the set of states at which the system can arrive (i.e., the states that can be reached) at a given time.

To more formally define a RS, let us consider the system governed by the differential equation $\dot{x} = f(t, x, u)$ (where x is the state, t the time and u a possible control input), which for a given initial condition and control spawns a trajectory $x(t)$. This solution, if the dependencies are explicit, defines the state trajectories flow $x(t) = \varphi^{t,t_0}(x_0, u)$. Consider an initial set (instead of a point) of initial conditions at t_0, and denote it by Ω_0. Consider the set of all possible actuations U. Then, the RS from Ω_0 at time t, denoted as $\Omega(t)$, is defined (assuming there are no collisions or singularities for the flow) as

$$\Omega(t) = \{x \in \mathbb{R}^n : x = \varphi^{t,t_0}(x_0, u), x_0 \in \Omega_0, u \in U\}. \tag{1}$$

Even in the linear case, the dependence on the control and initial conditions can make the computation of these sets quite difficult. If the dynamics are non-linear (as in orbital mechanics), a state transition matrix is not available, and therefore the computation of reachable sets becomes highly intensive (Kurzhanski & Varaiya, 2000). Our approach interprets the differential equation in a stochastic sense (Jain et al., 2019), with initial conditions given as a certain initial probability distribution, so that one can consider the starting set Ω_0 a confidence region of that probability distribution, and the RS its evolution, $\Omega(t)$. In principle, with six states (three pertaining to position and three to velocity) that may have some degree of uncertainty, one would require to propagate the boundary of a six-dimensional closed manifold, as well as the probability distribution function inside of it.

Thus, in this project a particle-based approach is applied instead (very much in the spirit of the Montecarlo method), in which one samples the initial confidence region, to then propagate those sample points. Since a large number of particles (trajectories) need to be propagated, the use of differential algebra techniques such as Taylor expansion over an initial condition can be employed

in order to obtain $\Omega(t)$ in a reasonable amount of time (see Armellin et al. (2010) and Pérez et al. (2013)). Notice that this in fact represents a higher-order approach than the classical propagation of covariances (linear approach), that relies on Jacobians and the assumption that Gaussian distributions keep being Gaussian, which does not hold true here since the non-linearity of the propagation "distorts" the distribution, making it lose its Gaussian shape (Holzinger & Scheeres, 2009).

In addition, one of the most interesting applications to SST of reachability analysis is the problem of object correlation. Looking at the literature, this problem has indeed received considerable attention in the last years. There are a number of metrics that can be used such as the Mahalanobis distance (Hall & Singla, 2019) and techniques that can help when several measurements are present, such as the use of attributables (Vananti et al., 2017; Reihs et al., 2021), but these do not explicitly take into account the possibility of manoeuvring objects, which is critical since small orbital corrections can produce outsized state discrepancies at the long term. This problem is tackled in Singh et al. (2012); Holzinger et al. (2012) computing (by means of optimal control) the minimum possible manoeuvre that connects the previous orbit with the new measurements. In Siminski et al. (2017) this optimal control approach is compared with the use of historical data, which is found more accurate when available (at least for the GEO example considered in that paper) and if the manoeuvres are predictable. These ideas are used in this work to develop manoeuvre detectors.

4. An orbit determination filter with basic manoeuvre detection capabilities

To decide which filter to develop for this work, it is important to take into account that the the Spanish survey radar S3TSR (Gomez et al., 2019) is the only source of measurement data considered for this project. Being a single radar, this implies that objects will have long windows without observation in-between, from about half a day up to 3 days, and then a radar track, typically with 5-20 individual plots, will become available. Therefore, designs relying on a large number of measurements and/or frequent measurements are not implementable. The scheme of Guang et al. (2018) is adapted, with manoeuvre detection based on residue analysis. As for the choice of the filter type itself, the UKF seems to be the superior choice. The rationale of this choice is as follows. Since an EKF relies on linearisation to obtain the evolution of the state error covariance, scenarios with long propagation times such as the ones considered in this work may degrade its accuracy, depending on the starting covariance. To overcome this drawback, an UKF is considered instead, since it provides a higher-order approximation for covariance evolution which can withstand longer propagations. The UKF is based on the "unscented transformation" first proposed by Julier & Uhlmann (1997) and later improved by Wan & Van Der Merwe (2000) to compute the first two moments of the probability density distribution of a random variable given by some transformation $y = h(x)$, assuming that the mean and the covariance of the variable x are known. The idea behind the unscented transformation is to use a set of points x^i (sigma-points) in such a way that their mapping $y^i = h(x^i)$ can be used to accurately approximate the exact mean and covariance of y (by using a predefined set of weights).

4.1. UKF algorithm

Considering a system with n states, given by the following process and observation equations

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}, t), \tag{2}$$

$$\boldsymbol{y} = h(\boldsymbol{x}, t), \tag{3}$$

and a set of weights to estimate the mean and the covariance (denoted by $w_m^{(j)}$ and $w_c^{(j)}$ respectively, for $j = 1, \ldots, 2n+1$), together with a tuning parameter κ (see Goff (2015) for a description of the weights and the tuning parameter values), the UKF algorithm is (obtained from Goff (2015)):

1. Start from the previous estimate of the state and the covariance of its error ($\hat{\boldsymbol{x}}_0$ and $\hat{\mathbf{P}}_0$ at first).
2. Read the next observation and its covariance: $t_i, \boldsymbol{y}_i, \mathbf{R}_i$.
3. Perform the decomposition $\mathbf{P}_{i-1} = \mathbf{A}^T\mathbf{A}$, and denote $\boldsymbol{a}^{(j)}$ as the j-th column of $\mathbf{A}$.
4. Calculate the sigma points:

$$\tilde{\boldsymbol{x}}_{i-1}^{(j)} = \hat{\boldsymbol{x}}_{i-1} + \check{\boldsymbol{x}}^{(j)}, \text{ for } j = 0, \dots, 2n, \text{ and } \check{\boldsymbol{x}}^{(0)} = 0,$$

$$\check{\boldsymbol{x}}^{(j)} = \boldsymbol{a}^{(j)}\sqrt{n+\kappa}, \text{ for } j = 1, \dots, n,$$

$$\check{\boldsymbol{x}}^{(n+j)} = -\boldsymbol{a}^{(j)}\sqrt{n+\kappa}, \text{ for } j = 1, \dots, n.$$

5. Propagate all the sigma points using numerical integration: Initial conditions $\tilde{\boldsymbol{x}}_{i-1}^{(j)}$, differential equation $\dot{\boldsymbol{x}} = f(\boldsymbol{x}, t)$, integration results $\tilde{\boldsymbol{x}}_i^{(j)}$.
6. Calculate the propagated a priori state and covariance (adding the process noise as in Section 4.3):

$$\bar{\boldsymbol{x}}_i = \Sigma_{j=0}^{2n} w_m^{(j)} \tilde{\boldsymbol{x}}_i^{(j)}, \; \bar{\mathbf{P}}_i = \Sigma_{j=0}^{2n} w_c^{(j)} \left(\tilde{\boldsymbol{x}}_i^{(j)} - \bar{\boldsymbol{x}}_i\right)\left(\tilde{\boldsymbol{x}}_i^{(j)} - \bar{\boldsymbol{x}}_i\right)^T + \mathbf{Q}_i.$$

7. Transform the sigma-points and calculate the predicted observation:

$$\tilde{\boldsymbol{y}}_i^{(j)} = h(\tilde{\boldsymbol{x}}_i^{(j)}, t_i), \; \bar{\boldsymbol{y}}_i = \Sigma_{j=0}^{2n} w_m^{(j)} \tilde{\boldsymbol{y}}_i^{(j)}.$$

8. Calculate the predicted observation covariance and results:

$$\mathbf{S}_i = \Sigma_{j=0}^{2n} w_c^{(j)} \left(\tilde{\boldsymbol{y}}_i^{(j)} - \bar{\boldsymbol{y}}_i\right)\left(\tilde{\boldsymbol{y}}_i^{(j)} - \bar{\boldsymbol{y}}_i\right)^T + \mathbf{R}_i, \boldsymbol{\nu}_i = \boldsymbol{y}_i - \bar{\boldsymbol{y}}_i.$$

9. Compute the Kalman gain and update the estimate of the state and its covariance:

$$\mathbf{V}_i = \Sigma_{j=0}^{2n} w_c^{(j)} \left(\tilde{\boldsymbol{x}}_i^{(j)} - \bar{\boldsymbol{x}}_i\right)\left(\tilde{\boldsymbol{y}}_i^{(j)} - \bar{\boldsymbol{y}}_i\right)^T, \; \mathbf{K}_i = \mathbf{V}_i(\mathbf{S}_i)^{-1},$$

$$\hat{\boldsymbol{x}}_i = \bar{\boldsymbol{x}}_i + \mathbf{K}_i\boldsymbol{\nu}_i, \; \hat{\mathbf{P}}_i = \bar{\mathbf{P}}_i - \mathbf{K}_i\mathbf{S}_i\mathbf{K}_i^T.$$

10. Return to step 1 and process the next observation.

4.2. Smoothing

In the simulated scenario considered in this work, measurements from the radar come in tracks of 5-20 plots, with long intervals in-between them (hours). While the BLS approach is simultaneous in nature, the KF approaches (EKF/UKF) process the measurements sequentially, in the order they were obtained; thus, the output of the filter can be improved via a backwards smoother. This additional algorithm propagates the filter backwards in time, starting from the last plot in a track, up to the first one (or even to previous tracks), modifying estimates accordingly (Goff, 2015). It is well-known that smoothers provide considerable improvement for orbit determination.

4.3. UKF process noise estimation

The UKF algorithm requires the process noise covariance as an input. This quantifies mismatches with respect to the real process. Consequently, it is a key factor in the filter as it will balance the credibility of the process with respect to the measurements. In any case, the process covariance is unknown, as its exact knowledge would imply perfect modelling, and has to be tuned. Initial covariance needs also to be estimated to be as realistic as possible (Poore et al., 2016). In Carpenter & D'Souza (2018), the state noise compensation technique described next is recommended as a good practice for navigation filters and has been adopted.

Denote by LVLH a Local-Vertical, Local-Horizontal frame. Assuming LVLH acceleration error as Gaussian white noise with covariance

$$\boldsymbol{\mathcal{S}}_{LVLH} = \begin{bmatrix} q_x & 0 & 0 \\ 0 & q_y & 0 \\ 0 & 0 & q_z \end{bmatrix}, \tag{4}$$

then, the transformation to inertial coordinates can be made using the rotation matrix $\mathbf{R}_{LVLH}$ which transforms LVLH coordinates to the inertial frame

$$\boldsymbol{\mathcal{S}}_k = \mathbf{R}_{LVLH}\boldsymbol{\mathcal{S}}_{LVLH}\mathbf{R}_{LVLH}^T. \tag{5}$$

Dividing the elapsed time between radar tracks in increments Δt, where the inertial orientation of the LVLH frame is assumed constant, the full state inertial covariance grows during an interval k as the equivalent discrete-time model of the corresponding Wiener process (Särkkä & Solin (2019), p.82):

$$\mathbf{M}_k = \begin{bmatrix} \boldsymbol{\mathcal{S}}_k\Delta t^3/3 & \boldsymbol{\mathcal{S}}_k\Delta t^2/2 \\ \boldsymbol{\mathcal{S}}_k\Delta t^2/2 & \boldsymbol{\mathcal{S}}_k\Delta t \end{bmatrix}, \tag{6}$$

and then process noise covariance estimation at the time instant i of the UKF algorithm is

$$\mathbf{Q}_i = \begin{cases} \sum_{k=1}^{N} \mathbf{M}_k, & \text{for filter calls between tracks,} \\ 0_{6\times 6}, & \text{for filter calls within a track.} \end{cases}$$

Where N is the number of time increments in-between filter calls. Within a track and between plots, where measurements are obtained every few seconds, the process mismatch is negligible.

4.4. UKF measurements and radar characterization

As a first step, the UKF must be tuned to work correctly in the absence of manoeuvres. The radar measurements (range, range rate, azimuth and elevation) are simulated as a noisy process following a gaussian distribution, with the mean on the true state, and power given by a fixed covariance. The complete characterization of the S3TSR is confidential, as it has numerous modes of operation with different capabilities, but for our testing purposes a single diagonal covariance matrix has been used in the (simulated) results and examples that are presented here and onward:

$$\sigma_\rho = 7 \text{ m},\ \sigma_{El,Az} = 0.3^\circ - 0.5^\circ,\ \sigma_{\dot\rho} = 0.4 \text{ m/s} \tag{7}$$

The radar accuracy on range is in the order of metres, whereas one gets errors under a metre per second for range rate, but the angular error translates into position errors of kilometres for a LEO satellite, as expected for a surveillance radar of this size. An intense tangential manoeuvre (0.1 m/s for our scenarios) at 600Km altitude would cause (under assumption of quasi-circular orbit) an angular phase difference of about 0.014° per revolution, or 150 m, way under the possible mismatch that angular measurements can suggest given it's inaccuracy. Thus, only range and range rate measurements are considered as the filter's inputs.

4.5. UKF preliminary testing results

Numerical results are shown in order to justify the chosen implementation. The considered scenario is the LEO satellite Sentinel-1A (with initial orbital elements taken from public TLEs and assumed precise, propagated with second-order gravity harmonics and drag using OREKIT) between 16:00:00 08/07/2015 to 16:00:00 12/07/2015. The following results assume a model mismatch in drag, with $C_D = 2.2$, $S = 10$ m^2 the "real" drag coefficient and exposed surface, and $C_D = 2$, $S = 9.5$ m^2 the assumed ones. The LVLH acceleration errors in Eq. (4) are

$$q_x = 10^{-9} \text{ m}^2 s^{-3},\ q_y = q_z = 5 \cdot 10^{-10} \text{ m}^2 s^{-3},$$

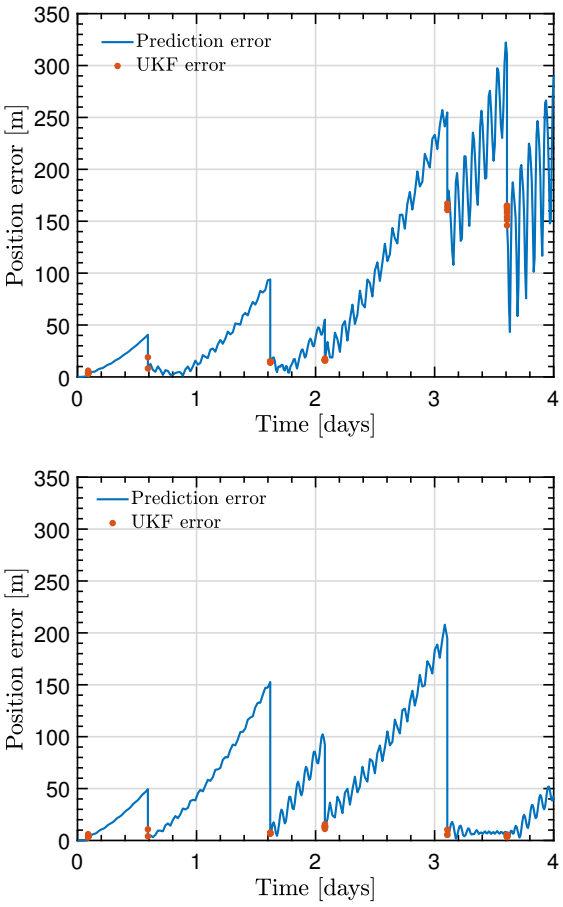

Fig. 1. Position error without process noise (top) and estimating process noise (bottom). Red dots indicate the UKF error at each plot (measurement points).

where more process noise has been assumed in the tangential direction (due to the drag modelling uncertainty being dominant in LEO). The discretization time period is taken as $\Delta t = 10$ min. A comparison with a simulation assuming a null process covariance noise is shown in Figure 1 (red dots indicate the mismatch between measurements and the predicted state after the filter update). It can be seen how the inclusion of some process noise greatly improves the filter's convergence.

In the results, the initial covariance is assumed small and with a realistic shape (obtained from running the filter for a few days); using a diagonal shape resulted in a much poorer performance of the filter.

4.6. Manoeuvre detection

The filter is adapted to estimate the presence of manoeuvres. Thus, in the 8th step of the UKF's algorithm, a manoeuvre prediction metric can be included, which reads:

$$\Psi_i = \sqrt{\boldsymbol{\nu}_i^T \mathbf{S}_i^{-1} \boldsymbol{\nu}_i}, \tag{8}$$

with $\boldsymbol{\nu}_i$ being the residuals and $\mathbf{S}_i$ the observation covariance, for $i \in \{1, \ldots, n\}$, n being the number of plots of a given track. This term can be used to estimate model mismatches (due to manoeuvres), and then trigger other manoeuvre detection algorithms, every time a radar track is processed.

For manoeuvre detection, three metrics based on Eq. (8) were proposed; either Ψ_1 (with the logic that the first plot may show

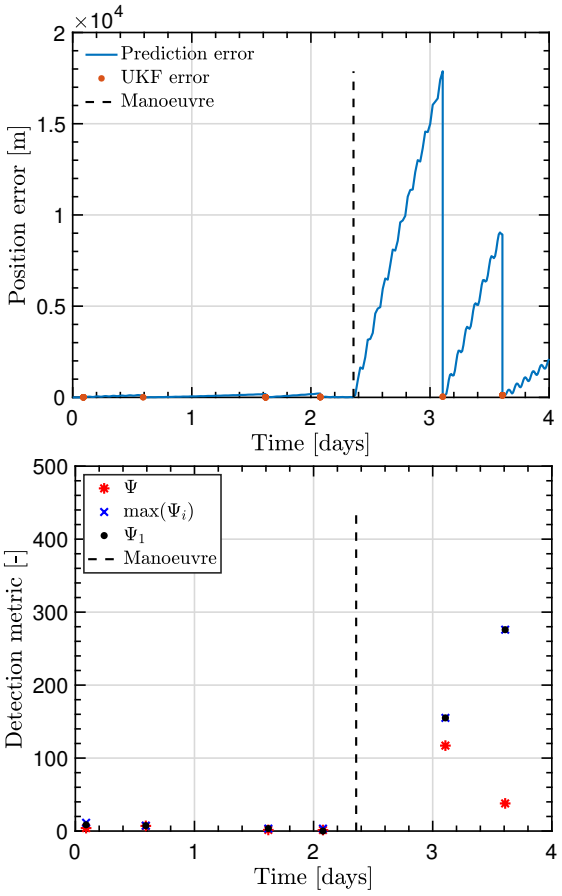

Fig. 2. UKF position error, on order of 10 metres, with manoeuvre (top) and detection metrics comparison (bottom). The vertical dashed line represents the manoeuvre

the largest impact from a previous manoeuvre), $\max_i\{\Psi_i\}$ or

$$\Psi = \frac{1}{n}\sqrt{\sum_{i=1}^{n}\Psi_i}, \tag{9}$$

As the detection method is based on finding significant discrepancies between the predicted orbit and the actual one, the residuals are computed prior to smoothing. A proof of concept is shown for the same scenario with a manoeuvre. The manoeuvre start is 00:34:58 11/07/2015 and ends at 00:35:24 11/07/2015 with a constant acceleration of $u = \left[0.31 \cdot 10^{-2}, -0.35 \cdot 10^{-3}, 0.37 \cdot 10^{-5}\right] m/s^2$ in the LVLH frame. In Figure 2, the UKF demonstrates its capability to recover the orbit after the manoeuvre is applied, even without any particular mechanism for covariance inflation. Moreover, a comparison of the possible detection metrics is also shown, with Ψ seemingly being the metric with a better trade-off between detection and false positives.

5. Reachability-based manoeuvre detection algorithms

In this section, RA as outlined in Section 3 is applied to the specific problem of manoeuvre detection. Thus, the starting inputs are the precise orbit of the objects (pre-manoeuvre) and a radar track (possible, post-manoeuvre), and the output is a metric that establishes a confidence in the detection of a manoeuvre.

As a first step, the theory of attributables (Vananti et al., 2017; Reihs et al., 2021) is introduced; it allows to “compress” several plots into a single, higher-quality measurement, fitting a full track into a single polynomial expression whose order needs to be determined.

Next, two algorithms are explained; the first is based on comparing the range and range rate attributables $(\rho, \dot{\rho})$ obtained from measurements with the one obtained from the initial uncertain orbit, by means of confidence regions and the Mahalanobis distance, which is a measure of the distance between a point and a distribution (Hall & Singla, 2019).

The second algorithm is based on the use of optimal control theory. Following the ideas of Singh et al. (2012) and Holzinger et al. (2012) one can compute by means of stochastic optimal control a distribution of the ΔV that connects the uncertain orbit with the measurement. This distribution can then be used to obtain the likelihood of a manoeuvre having been performed.

5.1. Attributables

We use attributables to condense the information of all plots in each track (Reihs et al., 2021). A radar provides range ρ, range rate $\dot{\rho}$, elevation El and azimuth Az, that, coupled with the chosen reference epoch t_{ref}, form the attributable

$$\mathcal{A} = \left\{ t_{ref}, \rho, \mathrm{El}, \mathrm{Az}, \dot{\rho} \right\}. \tag{10}$$

Elevation and Azimuth, as mentioned before, are not used here (save when explicitly said), but are included nonetheless for completeness. Fitting the information of the observables independently is one option, but it is possible to reduce the uncertainty of the resulting virtual measurement if one incorporates the definition of range-rate into the modelling, so that it shares the parameters with the range, as follows:

$$\rho\,(t) = \rho_0 + \rho_1 \Delta t + \rho_2 \frac{\Delta t^2}{2!} + \cdots + \rho_n \frac{\Delta t^n}{n!}\,, \tag{11}$$

$$\mathrm{El}\,(t) = \mathrm{El}_0 + \mathrm{El}_1 \Delta t + \mathrm{El}_2 \frac{\Delta t^2}{2!} + \cdots + \mathrm{El}_n \frac{\Delta t^n}{n!}\,, \tag{12}$$

$$\mathrm{Az}\,(t) = \mathrm{Az}_0 + \mathrm{Az}_1 \Delta t + \mathrm{Az}_2 \frac{\Delta t^2}{2!} + \cdots + \mathrm{Az}_n \frac{\Delta t^n}{n!}\,, \tag{13}$$

$$\dot{\rho}\,(t) = \frac{d\rho\,(t)}{dt} = \rho_1 + \rho_2 \frac{2\Delta t}{2!} + \cdots + \rho_n \frac{n\Delta t^{n-1}}{n!}. \tag{14}$$

In the expressions above $\Delta t = t - t_{ref}$. This method manages to average out noise and reduce the standard deviation of the virtual measurement. Following the nomenclature in Reihs et al. (2021), the set of equations that allows to solve the parameters in the sense of least-squares is:

$$\boldsymbol{m} = \begin{bmatrix} \boldsymbol{\rho} \\ \mathbf{El} \\ \mathbf{Az} \\ \dot{\boldsymbol{\rho}} \end{bmatrix} = A_{SYS}\,\boldsymbol{p} + \boldsymbol{\upsilon} = \begin{bmatrix} A & 0 & 0 \\ 0 & A & 0 \\ 0 & 0 & A \\ A_{\dot{\rho}} & 0 & 0 \end{bmatrix} \boldsymbol{p} + \boldsymbol{\upsilon}, \tag{15}$$

where $\boldsymbol{m}$ contains the measurements of all observables in the track, $\boldsymbol{p}$ collects the parameters ρ_i, Az_i, and El_i that one wants to calculate, and the matrices A and $A_{\dot{\rho}}$ have the coefficients in the formulas (11)–(13) and (14), respectively, evaluated at the times of the corresponding plots. The error $\boldsymbol{\upsilon}$ is assumed to follow a Gaussian distribution. Then the problem to solve is posed using weighted least-squares:

$$\min_{p}\, \boldsymbol{\upsilon}^T W \boldsymbol{\upsilon} = \min_{p}\, (\boldsymbol{m} - A_{SYS}\,\boldsymbol{p})^T W (\boldsymbol{m} - A_{SYS}\,\boldsymbol{p}), \tag{16}$$

whose solution is well-known:

$$\boldsymbol{p} = (A_{SYS}^T W A_{SYS})^{-1} A_{SYS}^T W \boldsymbol{m}. \tag{17}$$

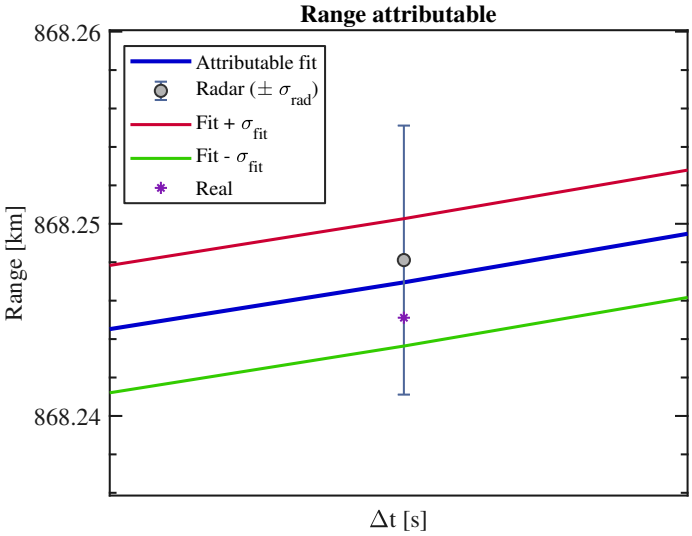

Fig. 3. Range attributable and reduction of error in a simulated radar track of 15 plots. Radar characterization is as in all simulated scenarios (see defined standard deviations in Eq. (7)), and "Real" value is that of the simulated low earth orbit at that instant.

The weighting matrix W is chosen to be the inverse of the covariance matrix of the measurements, Σ_υ, thus the attributable errors covariance matrix is:

$$\Sigma_p = \left(A_{SYS}^T \Sigma_\upsilon^{-1} A_{SYS}\right)^{-1}. \tag{18}$$

This allows to estimate how good a virtual measurement $\alpha(t)$ is expected to be at any point of the fit, by computing

$$Var[\alpha(t)] = \sum_{i,j} \sigma_{ij} \frac{\Delta t^i \Delta t^j}{i!\, j!}, \tag{19}$$

where σ_{ij} are the coefficients of Σ_p corresponding to the sub-matrix of each observable; these are sufficient to provide the covariance matrix of the complete attributable ($\Sigma_{\mathcal{A}}$) at the epoch without further processing.

A test track shown in Figure 3, with standard deviation that of our simulated radar (See Eq. (7)), has been used as an example of range attributable. The uncertainty is mitigated with a noticeable reduction of 50% in the standard deviation (in the case of the range rate, the reduction is of 60%), a consequence of the methodology used for the definition of the range-rate attributable.

In addition, one could even try to estimate the azimuth and elevation rates and use it for Initial Orbit Determination (IOD), but as one cannot expect to perform well as a good IOD with such a short observation arc, this method is not pursed in this work.

5.2. Algorithm 1: Comparison of real and projected attributable through Mahalanobis distance

Applying the nomenclature of Section 5.1, from all the plots of a track one can obtain the virtual values of range, range-rate, azimuth and elevation at the middle of a track (t_{ref}), namely $\mathcal{A} = (t_{ref}, \rho_0, \mathrm{Az}_0, \mathrm{El}_0, \rho_1)$, as well as the associated uncertainty in the form of a covariance matrix $\Sigma_{\mathcal{A}}$.

The following algorithm is used to obtain a "projected" (or predicted) measurement from the initial value of the reference orbit, which is assumed to follow a certain known distribution:

1. Sample the PDE of the initial condition obtaining m sample points. Denote these as x_{0j} for $j = 1, \ldots, m$. The set of initial conditions Ω_0 is then approximated by these points.

2. Propagate the sampled points using an OREKIT propagator up to time t_i (next radar attributable epoch). Taylor differential algebra methods can be used to greatly speed up this computation, at the price of a lengthy initial calculation (Andrea & Maisonobe, 2016). Thus, one obtains m trajectories $x_j(t)$.
3. Projected values at the attributable time t_i (middle of the track) are obtained as a cloud of points $x_j(t_i)$, with the density of points giving an approximate measure of the probability associated to the real trajectory.
4. Now for each sampled orbit, one can compute the corresponding radar measurements at t_i, obtaining a "cloud" of measurements, from which one can obtain its mean $(\widehat{\rho_0}, \widehat{\mathrm{Az}_0}, \widehat{\mathrm{El}_0}, \widehat{\rho_1})$ as well as the associated uncertainty (assumed Gaussian) in the form of a covariance matrix $\hat{\Sigma}$. This is denoted as the projected measurement (in the sense that it is the attributable value one would expect given the distribution of the initial condition).
5. Finally, the attributables and projected measurements can be compared. If no manoeuvre has been performed, one would expect that both values should somewhat agree. To formulate this more precisely, define:

$$\begin{bmatrix} \Delta\rho_0 \\ \Delta\mathrm{El}_0 \\ \Delta\mathrm{Az}_0 \\ \Delta\rho_1 \end{bmatrix} = \begin{bmatrix} \rho_0 \\ \mathrm{El}_0 \\ \mathrm{Az}_0 \\ \rho_1 \end{bmatrix} - \begin{bmatrix} \widehat{\rho_0} \\ \widehat{\mathrm{El}_0} \\ \widehat{\mathrm{Az}_0} \\ \widehat{\rho_1} \end{bmatrix}, \Delta\Sigma = \Sigma_{\mathcal{A}} + \hat{\Sigma}. \tag{20}$$

6. Then, if there is no manoeuvre, one would expect that, under an assumption of normality, $(\Delta\rho_0, \Delta Az_0, \Delta\mathrm{El}_0, \Delta\rho_1)$ should belong to a normal distribution of zero mean and covariance $\Delta\Sigma$. This can be checked either by computing confidence regions or equivalently through the Mahalanobis distance, as briefly explained next.

5.2.1. Use of confidence regions and Mahalanobis distance

For a n-dimensional multivariate normal distribution with mean m and covariance matrix Σ, the $p-$ level confidence ellipsoid (this is, the ellipsoid containing with probability p samples from the distribution) is given by

$$\left\{x \in \mathbb{R}^n : (x-m)^T \Sigma^{-1} (x-m) \leq \chi_n^2(p)\right\}, \tag{21}$$

where $\chi_n^2(p)$ is the inverse cumulative distribution function of the Chi-square distribution with n degrees of freedom (the dimension of the vector x), evaluated at the probability value p. Similarly, the Mahalanobis distance is a measure of the distance of a point x from a distribution. It is unitless, scale-invariant and takes into account the correlations of the distribution. Concretely, if the distribution has mean m and covariance matrix Σ the Mahalanobis distance (MD) of a point x is computed as

$$\mathrm{MD}(x) = \sqrt{(x-m)^T \Sigma^{-1} (x-m)}\,. \tag{22}$$

In particular if the distribution is a multivariate normal, then the MD^2 has a chi-square distribution with n degrees of freedom; thus, it is equivalent to the use of confidence ellipsoids. This property can be used to compute probabilities of manoeuvre. Next, an example is shown where the comparison of real and projected attributables is carried out for two cases: one example with manoeuvre and one without. Figure 4 shows that the confidence intervals and MD are able to discriminate the manoeuvred case from the non-manoeuvred one, at least for a simple basic simulation, using range and range-rate.

A manoeuvre detection metric has been computed based on the MD being distributed as a χ^2 distribution function with as many degrees of freedom (n) as variables, by fixing a threshold of being outside the 50% ellipsoid. Thus, a number $\mathrm{PR_{MD}}$ (which is a manoeuvre detection confidence) is computed as follows:

$$\mathrm{PR_{MD}} = max\left\{0, 2(\chi^2\left(\mathrm{MD}^2; n\right) - 0.5)\right\}. \tag{23}$$

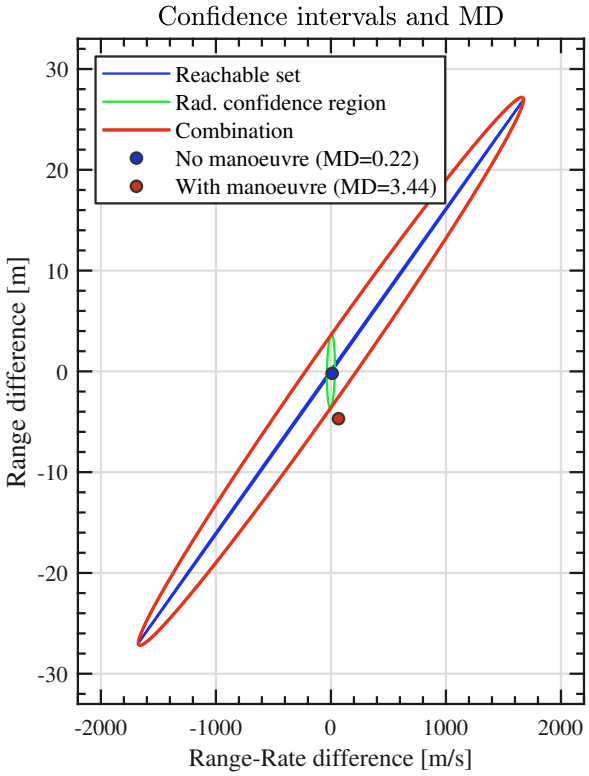

Fig. 4. Confidence intervals (90%) and Mahalanobis distance, considering only range and range-rate, without manoeuvre (blue circle) and with manoeuvre (red circle). The green ellipse represents the measurement uncertainty (radar, as characterized by Eq. (7)), the blue one the orbit uncertainty, and the red one the total uncertainty.

This way, if the MD has a probability of 50% or less of occurring, it is assumed that there is no manoeuvre, to reduce false positives. If the MD has a probability of more than 50% of happening, then one subtracts 50% and multiplies it by two; if one gets, e.g., a probability of a certain MD of 80%, then $\mathrm{PR}_{\mathrm{MD}} = 60\%$.

5.3. Algorithm 2: use of optimal control to compute a ΔV bounded measurement of distance

As a more sophisticated alternative to the Mahalanobis distance (Singh et al., 2012; Holzinger et al., 2012) one can compute by means of stochastic optimal control a distribution of the minimum ΔV that connects the uncertain orbit around it. This distribution can then be used as a metric to obtain the likelihood of a manoeuvre having been performed. The optimal control problem is posed as follows:

$$\begin{aligned} \mathcal{J} &= \min_{\Delta V_i} \sum_{i=1}^{N} \Delta V_i^2, \\ s.t. \quad x'(t) &= f(x(t), u(t), t), \\ x(t_0) &= x_0, \\ h\left(x\left(t_f\right)\right) &= [\rho\, \dot{\rho}]^T. \end{aligned} \tag{24}$$

In the above optimal control problem, the initial point is known from the precise orbit whereas the function h at the final point represents the function relating position and velocity with range and range-rate (the most precise measurements) which should take the value obtained with attributables as explained in Section 5.1. The function f represents the orbital dynamics, including any desired perturbation. The selected functional would represent the specific energy of the manoeuvre.

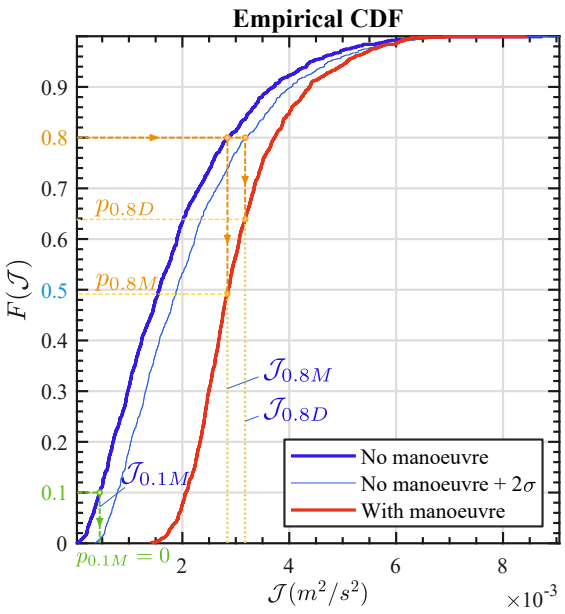

Fig. 5. Empirical distribution of manoeuvre energy $\mathcal{J}$ computed from the stochastic optimal control problem (1000 samples). The dark blue line represents the mean CDF of the non-manoeuvred case, and the light blue is the mean plus 2-sigma distribution, whereas the red plot represents a manoeuvred cases.

The problem is solved with CasADi (Andersson et al., 2019), an open-source solver for MATLAB, with a multiple shooting method discretizing the orbital dynamics in N time intervals; for each of these, since manoeuvres are small, the orbital dynamics is replaced with a linearized model obtained from OREKIT (computing the State Transition Matrix, or STM), with discrete ΔV's applied at the beginning.

As a first step, the problem has been solved in a deterministic way. Since, once the STM is computed, the solution is fast (seconds or less), to incorporate the stochasticity of the problem (both in initial orbit and measurements), a Monte Carlo algorithm has been implemented as a simple solution, albeit rather time-consuming. Figure 5 shows the obtained cumulative empirical distribution of $\mathcal{J}$ (from 1000 samples) for two cases (with and without manoeuvre).

In addition, a novel method to discriminate potential manoeuvres is now described. Qualitatively, it is clear that the distribution without manoeuvre is "smaller" than the one with manoeuvre. In the case without manoevre, we can derive a "mean distribution" as well as a distribution at a 2-sigma distance from the mean, which is helpful to avoid false positives. From these distributions some metrics have been defined, by using its 10%, 50% and 80% percentiles.

The metrics are exemplified in Fig. 5. As shown, take the 80% percentile of the estimated mean non-manoeuvred distribution, $\mathcal{J}_{0.8M}$, and compare with the probability $p_{0.8M} = \Pr(\mathcal{J} \leq \mathcal{J}_{0.8M})$ of the (potentially manoeuvred) distribution to be below that energy value, as graphically shown in Fig. 5. The higher that probability, the less likely of a manoeuvre to have happened according to that particular percentile (as there would be more cases that require the same or less energy than the non-manoeuvred case to connect the initial orbit and the measurement). In the figure it can be seen how using the 2-σ non-manoeuvred distribution (thus using $\mathcal{J}_{0.8D}$ instead, leading to $p_{0.8D}$) is more conservative. To use this idea to establish a metric for manoeuvre detection from a given percentile d the following scaling is used:

$$\mathrm{P}_d = max\left\{0, \frac{(d - \Pr(\mathcal{J} \leq \mathcal{J}_d))}{d}\right\}. \tag{25}$$

For instance, for $d = 0.1$ and calling $p_{0.1} = \Pr(\mathcal{J} \leq \mathcal{J}_{0.1})$, if $p_{0.1}$ is above 10% the metric becomes zero, and if not, the difference is multiplied by 10, which means that $p_{0.1} = 0$ would indicate total confidence of manoeuvre, for that metric (as shown in Fig. 5).

When using the mean distribution we will refer to this metric as P1M, with P1D reserved for the 2-sigma metric. Other metrics we use employ the 50% and 80% percentiles and are denoted as P5M, P5D, P8M and P8D.

6. A manoeuvre detection filter using reachability analysis

The results of sections (8.4)-(8.6) suggest the need of a combined manoeuvre detection tool-filter, that is able both to predict an orbit and detect and take into account manoeuvres. This section presents our work on such a filter, which we call the Manoeuvre Detection Filter (MDF). The idea is to follow the scheme of the UKF of Section (4), combined with the Algorithm 1 of Section (5.2), to detect manoeuvres. Once a manoeuvre is detected, the idea of covariance inflation is followed (see Section 3.10.7 of (Goff, 2015) and (Chul Ko & Scheeres, 2016)); thus, the state covariance is increased up to the point where a manoeuvre is no longer detected, which would imply that the uncertainty of the state is able to include the possibility of such a manoeuvre having been performed. The reasons to choose Algorithm 1 instead of 2 are that it performs better for scenarios 1 to 13, which are the ones best suited to the filter. It is hopeless to expect the filter to perform well in scenarios with too many manoeuvres and few radar tracks. In addition, Algorithm 1 fits quite well with the philosophy of the filter: the unscented transform can be used to estimate the state covariance by using the attributable as a "virtual measurement" used only for purposes of manoeuvre detection, but not for updating the state. This considerably reduces the computational burden. Also Algorithm 1 gave less false positives in the real testing.

To be more precise, using the UKF notation of Section 4.1 and skipping the unchanged steps, the MDF algorithm is:

1. Start from the previous estimate of the state and the covariance of its error ($\hat{\boldsymbol{x}}_0$ and $\hat{\boldsymbol{P}}_0$ at first).
2. Compute the attributable time $t_i^{\mathcal{A}}$, the values of the next track $\boldsymbol{y}_i^{\mathcal{A}}$ and its covariance $\boldsymbol{\Sigma}_i^{\mathcal{A}}$.
3. Compute the sigma-points of the unscented transform.
4. Propagate all the sigma points using numerical integration until the attributable time $\tilde{\boldsymbol{x}}_i^{(j)}$.
5. Compute the weighted mean and the covariance matrix of the transformed sigma-points: $\bar{\boldsymbol{x}}_i$ and $\bar{\mathbf{P}}_i$.
6. Compute the manoeuvre detection metric p_i by transforming the sigma-points (using the observation equation) to get the predicted observation $\hat{\boldsymbol{y}}_i$, the residuals $\boldsymbol{\nu}_i = \boldsymbol{y}_i^{\mathcal{A}} - \hat{\boldsymbol{y}}_i$, and the observation covariance $\mathbf{S}_i$ to compute the MD_i and the associated metric p_i (using Eq. (23)):
$$\mathbf{S}_i = \Sigma_{j=0}^{2n} w_c^{(j)} \left(\tilde{\boldsymbol{y}}_i^{(j)} - \hat{\boldsymbol{y}}_i\right)\left(\tilde{\boldsymbol{y}}_i^{(j)} - \hat{\boldsymbol{y}}_i\right)^T + \boldsymbol{\Sigma}_i^{\mathcal{A}}$$
$$\mathrm{MD}_i = \sqrt{\boldsymbol{\nu}_i^T \mathbf{S}_i^{-1} \boldsymbol{\nu}_i}.$$
7. If $p_i \geq 0,5$, multiply the covariance of the predicted state $\bar{\mathbf{P}}_i$ by 2 and return to Point 6; otherwise, continue.
8. Transform the sigma-points (using the observation equation) and calculate the predicted observation, the residuals, and the observation covariance.
9. Calculate the predicted observation covariance and the residuals.
10. Compute the Kalman gain and update the state estimate.
11. Return to step 1 and continue propagating.

In addition, a "long smoothing" is implemented: if no manoeuvre is detected, a smoothing is performed backwards until the previous radar track, and again forwards.

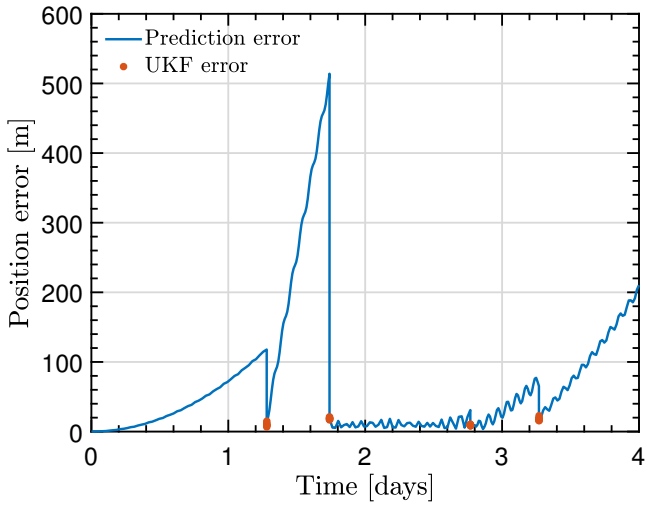

Fig. 6. Position error for Sentinel-1A with respect to reference orbit without manoeuvre. Red dots indicate the mismatch (which becomes of the order of 10 metres) between measurements and the predicted state after the filter update.

7. Results for simulated scenarios

An OREKIT-based simulator, both for the manoeuvres and for the radar observations, has been developed. They provide realistic (though not accurate) testing examples. They are very useful to tune and validate the different algorithms and filters. Starting points are generated from public TLEs which are used to define reference orbits with propagators including J_2 and aerodynamic drag, as explained in Section 4.5.

The algorithms, besides the model mismatches explained in that section, start from initial conditions within the expected limits of error of the real precise orbits (metres). Two main scenarios, respectively based on the satellites Sentinel-1A and Swarm-C, are considered. A manoeuvre either tangential (T), out-of-plane (OOP) or hybrid (with components both tangential and out of plane) is simulated, maintaining a constant acceleration of $10^{-3}\ m/s^2$ and characterized by the following fields:

1. Manoeuvre intensity (regulated through the duration): low (5 seg $\rightarrow 5 \cdot 10^{-3}\ m/s$) / medium (30 seg $\rightarrow 3 \cdot 10^{-2}\ m/s$) / high (120 seg $\rightarrow 1.2 \cdot 10^{-1}\ m/s$).
2. Manoeuvre location with respect to a radar track: 2 h, 6 h or 12 h before radar.
3. The Sentinel-1A scenario spans from 00:00:00 18/08/2020 to 00:00:00 22/08/2020. The manoeuvre starts at 18:25:00 20/08/2020.
4. The Swarm C scenario spans from 00:00:00 14/07/2020 to 00:00:00 20/07/2020. The manoeuvre starts at 12:30:00 17/07/2020.

Thus, combining all these factors, one gets 18 simulation scenarios per satellite to analyse the influence of these factors for the algorithms. Due to space limitations, only selected results are shown, with general conclusions drawn from the complete set.

7.1. UKF simulated results

The result without manoeuvre for Sentinel-1A is presented in Figure 6, whereas the manoeuvred case (tangential) is shown in Figure 7. The value of Ψ, which should be useful for detecting manoeuvres, is given for some Sentinel-1A cases in Table 1.

It can be observed e.g. in Fig. 6 that the filter takes some time to stabilize. This is probably due to the incorrect initial covariance. Since in real scenarios the covariance will not be perfectly known this can be expected. On the other hand, it is clear that the filter is

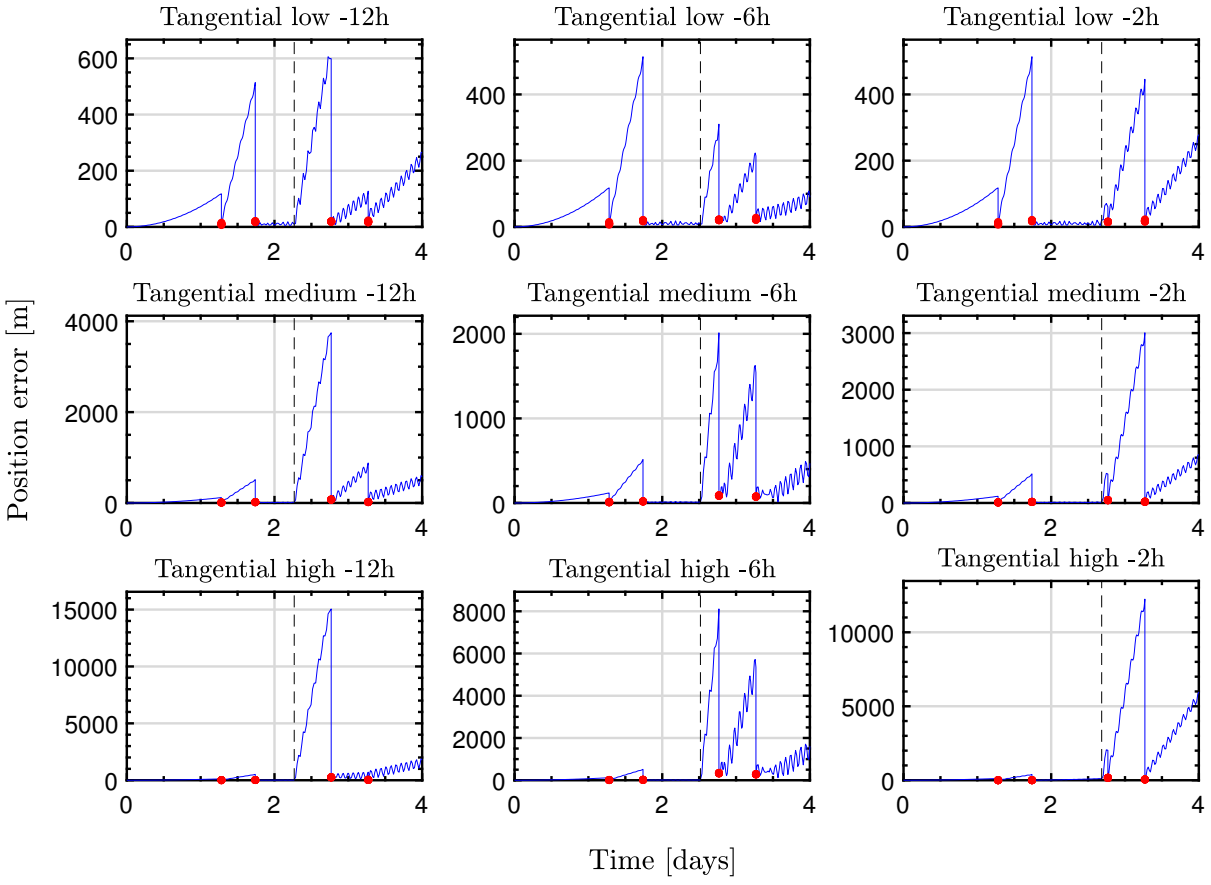

Fig. 7. Position error with respect to reference orbit for Sentinel-1A manoeuvred scenarios. Red dots indicate the mismatch (which becomes of the order of 10 metres) between measurements and the predicted state after the filter update.

***Case*/Ψ**	***Pre − man.Max.***	***Post − man.Max.***
No manoeuvre	4.435	2.341*
low − 2h	4.435	2.386*
low − 6h	4.435	2.411*
low − 12h	4.435	2.335*
medium − 2h	4.435	3.478*
medium − 6h	4.435	10.33*
medium − 12h	4.435	6.260
high − 2h	4.435	34.39*
high − 6h	4.435	132.0*
high − 12h	4.435	75.89

Table 1. Maximum value of filter detection metric before and after Sentinel-1A tangential manoeuvre, simulated results. The asterisk indicates that the maximum arises after the first post-manoeuvre track (i.e., at a later track). This metric only uses range and range rate.

working correctly in all cases; since the measurements are scarce it is unavoidable that the position errors grow, however, they are clearly mitigated at each measurement. It can be seen in Fig. 7 that manoeuvres induce large errors after they happen, since they are unaccounted for in the process covariance. The largest the manoeuvre the larger the error and the more it takes to recover from it. From Table 1, one can observe that the value of Ψ is indicative of the presence of a manoeuvre only in medium and specially in high-intensity cases. Low-intensity manoeuvres are indistinguishable from process noise. In addition, the distance to the radar measurement does not seem to have much influence in the value of Ψ

In the Swarm C case (not shown), the value of Ψ was indicative of the presence of a manoeuvre only in high-intensity cases. For low- and medium-intensity manoeuvres, they were, in principle, indistinguishable from process noise, unless the manoeuvre happened at a long enough distance from the first radar measurement. The main cause of this was, besides the long gap without measurement, having less radar measurements; in the case of Sentinel-1A, nine values were obtained at the pass after the manoeuvre, whereas in the case of Swarm-C, only five values are obtained.

7.2. Algorithm 1 simulated results

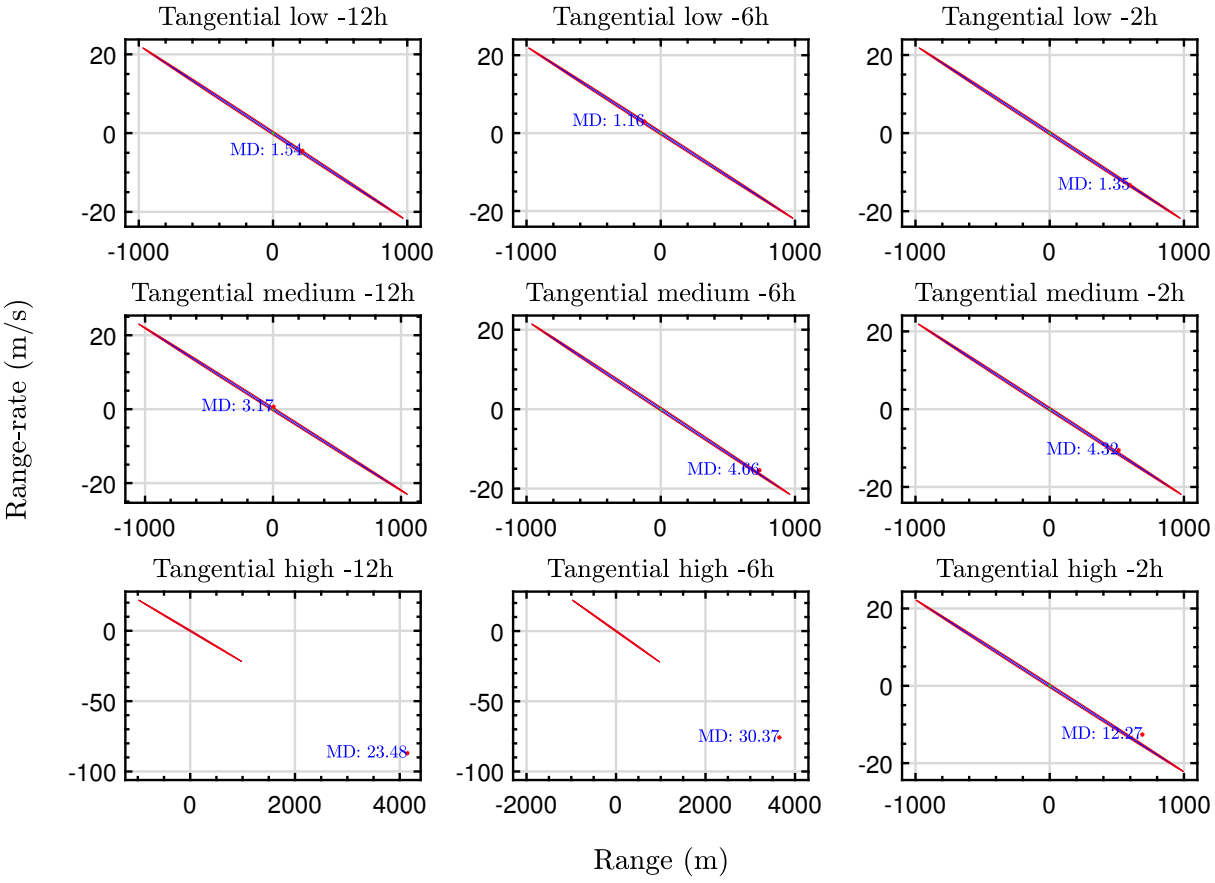

Fig. 8. Reachability results (Algorithm 1) for Sentinel-1A manoeuvre, Range vs Range-rate.

In Table 2, the Mahalanobis distance (MD) has been computed considering only range-range rate (column 2), El-Az (column 4) and all four measurements (column 6). As seen in the table, in general, using only range and range-rate is more sensitive in more cases; using elevation and azimuth can induce false positives. All high- and medium-intensity manoeuvres are detected, but low-intensity manoeuvres are usually not detected. In addition, the distance to the radar track does not seem to affect these results.

The results can be inspected visually in Figure 8. Note that due to the propagation "stretching" the orbit uncertainty in the range-range rate plane, it is hard to verify if measurements belong to the confidence region, except in high-intensity cases. In the Swarm C case (not shown), only high-intensity manoeuvres are detected, with varying success for medium-intensity manoeuvres.

7.3. Algorithm 2 simulated results

The results are presented in Table 3. The metrics computed from Algorithm 2 detect all high- and (except P8M and P8D) medium-intensity manoeuvres. P1M is the most sensitive algorithm being able to detect even some low-intensity manoeuvres.

Case / Metrics	MD ($\rho, \dot{\rho}$)	PR_{MD} (%)	MD (El, Az)	PR_{MD} (%)	MD (All)	PR_{MD} (%)
WoM	0.29	0	2.18	33	2.60	0
$L-12h$	1.54	8	1.66	13	2.22	0
$L-6h$	1.16	0	1.52	7	1.96	0
$L-2h$	1.35	0	2.10	30	2.16	0
$M-12h$	3.17	59	2.15	32	3.97	18
$M-6h$	4.66	81	0.97	0	4.81	39
$M-2h$	4.32	77	2.29	36	4.80	38
$H-12h$	23.48	100	3.61	67	23.53	100
$H-6h$	30.37	100	4.44	78	30.38	100
$H-2h$	12.27	100	1.23	0	12.44	97

Table 2. Sentinel-1A reachability analysis with Algorithm 1 and detection metric from MD. WoM=without manoeuvre, L=low, M=medium, H=high. Tangential case.

Case / Metrics	P1M (%)	P5M (%)	P8M (%)	P1D (%)	P5D (%)	P8D (%)
FP	10	1	0	3	0	0
$L-12h$	100	72	0	70	40	0
$L-6h$	100	42	0	50	0	0
$L-2h$	0	0	0	0	0	0
$M-12h$	100	96	0	100	84	0
$M-6h$	100	100	95	100	100	95
$M-2h$	100	98	0	100	96	0
$H-12h$	100	100	100	100	100	100
$H-6h$	100	100	100	100	100	100
$H-2h$	100	100	100	100	100	100

Table 3. Sentinel-1A reachability analysis with Algorithm 2 and detection metric from optimal control approach. FP=false positives, L=low, M=medium, H=high. Tangential case. This metric only uses range and range rate.

However, it has a non-negligible rate of false positives (a false positive is defined as a non-manoeuvred case from the Monte Carlo simulation being detected with $p \geq 50\%$). P1D is only slightly less sensitive and reduces the number of false positives. Other metrics seen to perform worse than P1M and P1D. In the Swarm-C case (not shown), only high-intensity manoeuvres are detected, with varying success for medium-intensity ones. As in Section 7.1, main causes are long gaps without measurement and having less radar measurements right after the manoeuvre.

8. Results for real scenarios

Given the algorithms already presented and validated through simulated scenarios on previous sections, this section presents the results obtained when they were tested on real-world data, for satellites of the Sentinel and Swarm family, and also for TerraSAR-X and TanDEM-X satellites. The data used and the sources were, for the satellites' orbits, OEM data, this is, accurate position and velocity information of the satellites under study (with precisions one order of magnitude better than the radar data, i.e., with position error of about 1 metre), provided by ESA/ESOC and DLR/GSOC. For radar data, real tracks from the Spanish survey radar S3TSR were used, with the necessary uncertainty information for the algorithms. Finally, for testing purposes, manoeuvre data, providing accelerations in a local reference frame as well as the duration, were provided by ESA/ESOC and DLR/GSOC.

First, the selected scenarios are presented and briefly described in Section 8.1. In Section 8.2 the particularities of the dynamical modelling are detailed, followed by a brief comment on the data consistency check (Section 8.3). The last subsections (8.4)–(8.5) contain the numerical results of UKF and reachability analysis using the real data.

Scenario	Epoch of most intense man.	Initial epoch	Final epoch	# of segments	# seg. with man.
1	14-Aug-2019 23:11:03	06-Aug-2019 06:52:43	16-Aug-2019 18:09:54	12	2
2	29-Jan-2020 23:11:08	23-Jan-2020 06:36:14	31-Jan-2020 07:09:05	8	1
3	21-May-2020 01:06:08	17-May-2020 06:27:59	25-May-2020 18:01:43	10	1
4	19-Aug-2020 22:29:28	19-Aug-2020 17:45:21	27-Aug-2020 18:18:18	9	2
5	14-Aug-2019 23:59:53	05-Aug-2019 07:00:17	17-Aug-2019 07:00:15	13	2
6	19-Sep-2019 00:57:21	10-Sep-2019 18:00:59	22-Sep-2019 07:00:19	12	1
7	20-May-2020 22:37:36	12-May-2020 18:09:12	28-May-2020 17:36:30	17	1
8	17-Jun-2020 22:05:17	09-Jun-2020 17:36:29	24-Jun-2020 18:00:59	18	1
9	05-Feb-2020 16:09:00	03-Feb-2020 22:25:25	13-Feb-2020 22:25:22	8	1
10	10-Sep-2020 17:10:27	03-Sep-2020 11:22:01	10-Sep-2020 22:25:35	6	1
11	18-Sep-2019 16:59:55	11-Sep-2019 11:01:44	21-Sep-2019 22:25:30	5	1
12	17-Sep-2020 16:10:04	08-Sep-2020 11:22:02	19-Sep-2020 10:51:42	9	1
13	15-Jul-2020 17:02:15	06-Jul-2020 01:47:22	20-Jul-2020 00:05:35	12	1
14	22-Aug-2020 00:24:06	17-Aug-2020 17:49:25	24-Aug-2020 06:48:33	4	4
15	22-Aug-2020 00:24:07	16-Aug-2020 18:06:50	22-Aug-2020 17:58:07	4	4

Table 4. List of real testing scenarios. A segment is defined as the elapsed time between a radar track and the next.

8.1. Real testing scenarios

The list of scenarios is in Table 4 with the corresponding epochs. The satellites used to create these scenarios are: Sentinel-1A, Sentinel-1B, Sentinel-2A, Sentinel-2B, Swarm-C, TanDEM-X and TerraSAR-X, but the relation between satellite and scenario has been omitted here. Scenarios are divided in a number of segments, which start and end at consecutive radar tracks. For the purposes of testing the RA algorithms, these segments are considered and processed individually (using the precise orbits to determine the starting point for each segment), whereas the filter runs for a full scenario, processing each segment consecutively.

A Gantt-like representation was produced to exemplify how the scenarios are generally distributed, see Figure 9 for Scenario 1, with the radar observation (red) and manoeuvres (blue). For completeness, the plots include the simulated radar observation (black circles), which in some cases reveal missing radar tracks from the real data.

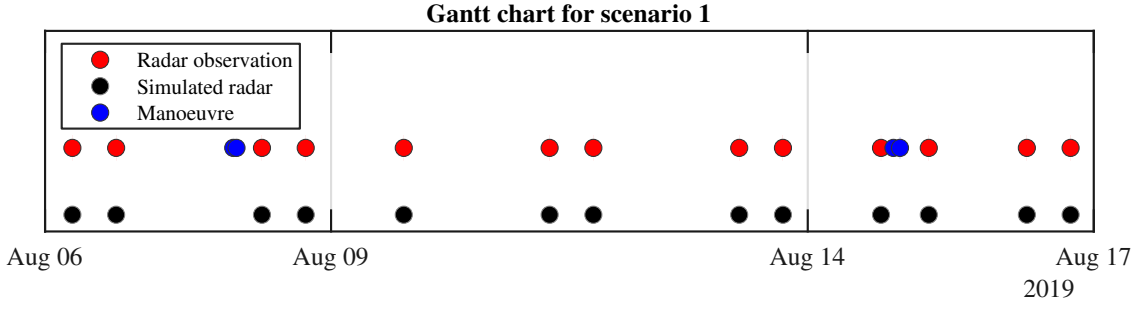

Fig. 9. Gantt chart for one real scenario. This is representative of the frequency of radar tracks for the orbits under study.

The simulations studied in Section 7 were carried out for a single type manoeuvres only, namely a uniform acceleration in a single time segment, either IP, OOP or hybrid. However, real operation of satellites show that orbit corrections are generally a combination of more than one type, in consecutive and close time segments (usually low impulse IP and medium impulse OOP). These consecutive manoeuvres may have a bigger impact on the orbit by increasing its detectability.

8.2. Dynamical modelling

Since real orbits are subject to multiple complex perturbations, an important initial step is to determine the dynamical model for each satellite. In order to use the dataset provided for the real testing scenarios several improvements had to be made to the modelling of the dynamics with respect to the one used in Section 7. The most important ones are the changes of the Earth gravity and atmosphere models. For one, the degree and order of the harmonics for the earth gravity field has been considerably increased. Testing has been done to discern the relation between the computational costs of increasing this parameter against the changes in the simulation error (measured with respect OEM data).

Given the uncertainty of the data, values of the harmonic's degree/order above 40 have a negligible effect and can be discarded, as it would significantly slow down the computation without any relevant benefit. To justify this, a comparison has been made with increasing degree/order of the harmonics, see Figure 10 for a representation of the position error evolution (against OEM data points) along a 24-hour simulation for Sentinel-1A. The time required to simulate is in the legend, where up to degree/order 40 it is affected very little when considering the great reduction in prediction error. Going above this value has a measurable effect in the time required, but with almost no impact on the error (there is even some small random increase possibly due to other perturbations and misfits). This result supports the decision to keep the harmonics only up to degree/order 40. Just to make sure that these results hold for lower orbits, a similar test has been done for a 1-day simulation interval with the OEM data of Swarm-C, which despite not being shown here points to the same conclusions.

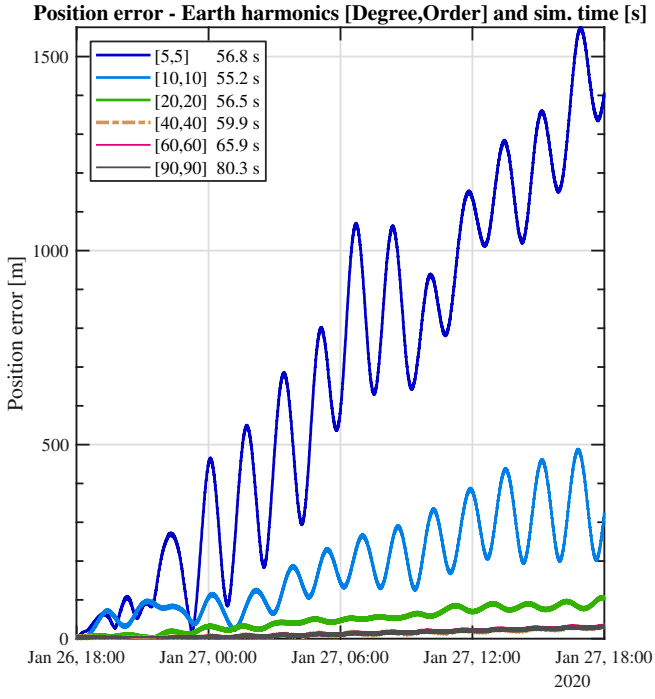

Fig. 10. Position error (m) for Sentinel-1A in a 24-hour segment, comparing the effect of increasing the Earth harmonics.

8.2.1. Atmosphere model and other orbital perturbations

The atmosphere model has been changed from a Modified Harris Priester model (static atmosphere) on the simulated scenarios to the 2001 Naval Research Laboratory Mass Spectrometer and Incoherent Scatter Radar Exosphere of the MSIS model (Picone et al., 2002), or NRLMSISE-00 model (also used by ESA for prediction and orbit determination). This last model is empirical and needs real weather data to compute the neutral atmosphere from the surface to the lower exosphere. The model feeds from the

MSAFE (Marshal Solar Activity Future Estimation) data implemented inside OREKIT, which provides the mean and daily solar flux and geomagnetic indices. From the carried out testing, see Figure 11, the MSAFE atmosphere provides a much better prediction (against OEM data), but is considerably more expensive to compute, as it requires to perform interpolations from the empirical data to compute density values. These tests show us that, once the Earth shape has been modelled accurately (degree/order of 40), the gain from a more advanced atmosphere model is also very significant, even close to one order of magnitude. Another comparison has been made to measure the relative importance of the solar radiation pressure in simulations of these periods (the order of days) when the other forces are modelled with as much precision as possible, to conclude that this perturbation's relative significance is minimal, see Figure 12 (the effect is masked by other modelling errors for the length of these simulations).

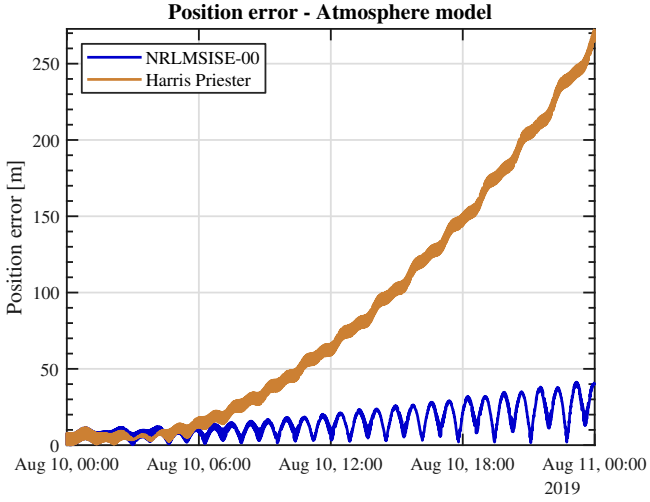

Fig. 11. Position error (m) for Sentinel-1A in a 24-hour segment, comparing the effect of two atmosphere models. The Harris Priester is only used in the simulated scenarios, while the NRLMSISE-00 model is used for the real scenarios to maximize fidelity. No other perturbations have been added for this plot (aside from earth shape with order/degree 40).

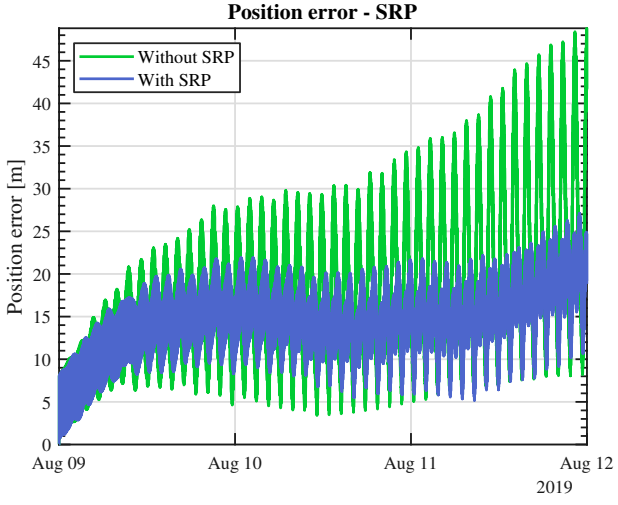

Fig. 12. Position error (m) for Sentinel-1A in a 72-hour segment, comparing the effect of the solar radiation pressure. This simulation includes luni-solar perturbations, NRLMSISE-00 atmosphere and earth shape with order/degree 40.

8.2.2. Satellite parameters

The last consideration has been the model of the satellite itself, which is defined by the drag coefficient (C_D), the frontal area for the drag force (S) and the area affected by the solar radiation pressure (S_{SRP}). The mass of the satellite is considered different for

each case and has been chosen to be the mean value between the wet and dry mass of each satellite. Simple as the model is, the 3 parameters are adjusted to get a good fit with the real data provided. In order to do this, an iterative optimization process has been performed, with enough iterations so that the changes at the end are sufficiently small to consider that it has reached a minimum (normally at around 4-5 iterations). The algorithm is quite expensive computationally speaking, as for each of the optimization steps the orbit must be simulated several times for it to find a solution. This method has another drawback, and it is that depending on the chosen segment the final values may vary slightly. Although marginal, this effect can be palliated if the final values are averaged between different segments.

8.3. Data consistency check

Several sanity checks were performed to ensure the consistency of the different data sources and with the propagators, namely, verifying: that the precise orbit replicate, approximately, the radar measurements (to measurement and orbit error); that the propagators do not have much error with respect to the precise orbits and the measurement in the absence of manoeuvres; that the manoeuvre file was consistent with the OEMs, which can be verified by the error of the propagators growing rapidly in the presence of manoeuvres.

These checks are essential to ensure that false positives or false negatives are not in fact detecting inconsistencies in the data sets, and although very extensive, only general results will be mentioned here.

The first check (consistency between OEM files and radar data) shows that the range differences are in the order of the combined error of the radar measurements and the OEMs themselves. The second check, consistency between our propagator and the OEM data, shows that in general the propagation performs well (errors about 60 metres maximum for a 24-hour period of propagation, as found when adjusting the values of the dynamics model), but as expected in the presence of manoeuvres errors grow rapidly, so that the last check is confirmed as well (see Figure 13 for an example).

8.4. UKF real results

This section tests the developed UKF against real radar data of manoeuvring satellites. The filter was improved, including also the ballistic coefficient in the estimation. In addition, since the metric Ψ has a close relation to the Mahalanobis distance, Equation (23) was directly used to derive a manoeuvre detection metric.

The smoothed prediction errors with respect to OEMs, the smoothed range residuals as well as the detection metric derived from Ψ are shown for one scenario of Sentinel-1A in Figure 14. The filter error increases slightly at the beginning and then the filter converges; later, after the manoeuvre, errors start to increase considerably. In any case, the steep increment in the residuals allows the detection of this manoeuvre by the filter metric.

Space restrictions do not allow for a more comprehensive revision of the results, but the following conclusions were derived. The estimation of the ballistic coefficient does not vary much and does not seem to have significant impact. It was observed that the filter was, in general, well-behaved even in the presence of manoeuvres. When there are no manoeuvres, it tends to converge, albeit sometimes slowly, to errors of the order of just a few hundred metres compatible with the errors of the propagator (for long propagations). However, there were some instances of rapid increase of errors due to the presence of manoeuvres. The manoeuvre detection metric Ψ^* allowed for the computation of a confidence in detection but was not very sensitive. Occasionally it detected a manoeuvre after one or two radar tracks have passed. The need of a combined manoeuvre detection tool integrated in the filter became clear: then, the filter can react to a manoeuvre by increasing the process noise (covariance inflation) and thus take into account the presence of the unknown manoeuvre. In addition, it would allow for longer smoothings reducing the risk of using a segment with a manoeuvre for a long smoothing. This filter is shown in Section 6.

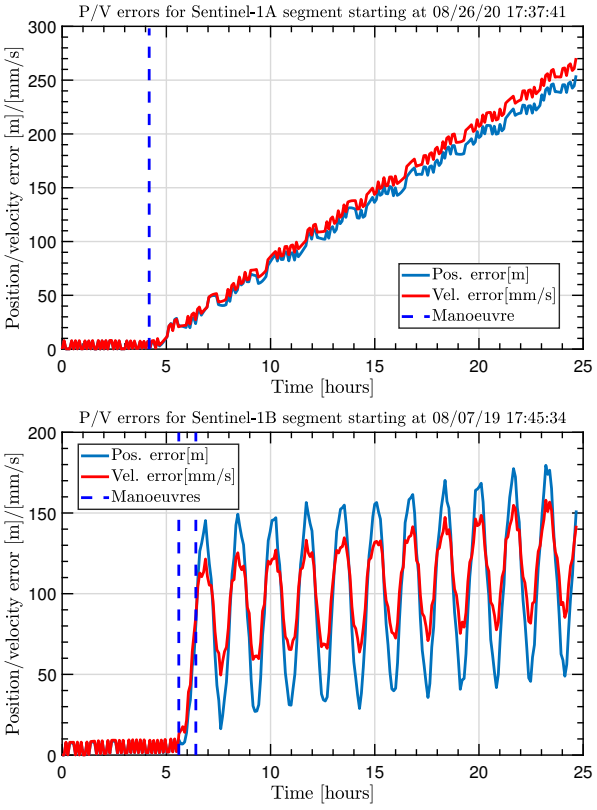

Fig. 13. Examples of how manoeuvres make prediction errors (with respect to OEMs) grow large.

8.5. Algorithm 1 real results

For each segment of the real scenarios, Algorithm 1 of Section 5.1 was implemented. Initialization of the predicted orbit is done using the precise orbit information at the closest OEM point right after a radar track, and it is stopped at the attributable epoch of the next track (the middle of it). Each of these simulations is done using Taylor differential algebra and an assumed covariance matrix Σ_{OEM} for initialization (see Equation 26, obtained by scaling the output of the filter of Section 4), where the coefficient block matrices in Eq. 26 are in Eq. 27. This, combined with the confidence of the radar attributable can be used to compute a cloud of points and from that, the $\mathrm{PR_{MD}}$ metric defined in Equation (23); here MD is calculated from range and range-rate only.

$$\Sigma_{OEM} = \begin{bmatrix} \Sigma_{P/P} & \Sigma_{P/V} \\ \Sigma_{P/V}^T & \Sigma_{V/V} \end{bmatrix} \tag{26}$$

$$\begin{aligned} \Sigma_{P/P} &= \begin{bmatrix} 2.094 & -1.684 & -2.281 \\ -1.684 & 2.686 & 0.673 \\ -2.281 & 0.673 & 6.733 \end{bmatrix} m^2 \\ \Sigma_{V/V} &= 10^{-4} \begin{bmatrix} 4.096 & 0.1959 & 4.097 \\ 0.1959 & 3.710 & 0.9932 \\ 4.097 & 0.9932 & 5.27 \end{bmatrix} m^2/s^2 \\ \Sigma_{P/V} &= 10^{-2} \begin{bmatrix} -2.025 & -0.7215 & -1.646 \\ 1.110 & -0.1791 & 0.7416 \\ 4.250 & 1.982 & 5.580 \end{bmatrix} m^2/s \end{aligned} \tag{27}$$

The results are summarized and analysed, with each satellite being grouped in Table 5. The results of group *1-4*, are quite positive, as there is only one false negative and no false positives, out of 33 cases. Thus, out of 6 manoeuvres, 5 were detected

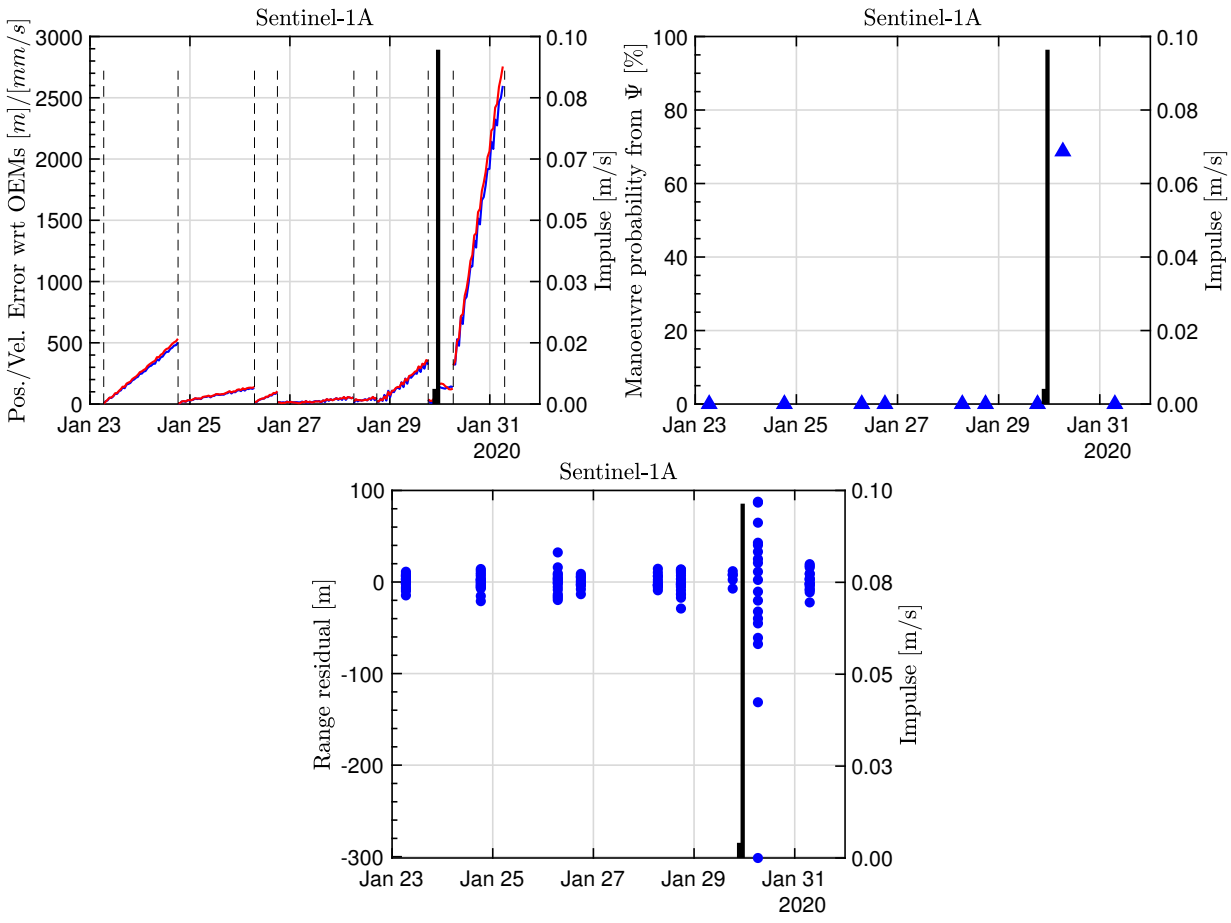

Fig. 14. UKF results for one scenario of Sentinel-1A

and with high confidence value in general. This is a rather satisfactory result, as the only manoeuvre that was not detected is also the smallest. Group *5-8* cases contain 5 manoeuvres, but only 1 is detected (another one results in a non-zero value but with low confidence). There are also 3 false negatives, from 55 total cases. Group *9-10* results are not very reassuring as no manoeuvres are detected; and with one false positive (from 22 cases). The 2 manoeuvres present in the *11-12* scenario are detected, with no false negatives (from 13 cases). Group *13* presents no detected manoeuvres (out of one) but no false negatives. For group *14-15* (which is a high-manoeuvring case) only 25% of manoeuvres are detected.

In global, the results are in need of improvement, as only about 40% of manoeuvres are detected, even with some of them being rather intense. The rate of false positives is quite good on the other hand. Analysing the results, the causes of errors were as follows; from a total of 5 false positives, all except one present less than 10 plots in the track following the manoeuvre. Thus, the main cause of false positives is tracks with fewer plots. From a total of 14 false negatives, all except one were segments of length equal to or larger than one day. Thus, the main cause of false negatives is longer propagations accumulating additional propagation error. Sometimes these longer propagation periods are due to missed radar observations right after the manoeuvre. The most challenging scenarios were those with TanDEM-X/TerraSAR-X due to the abundance of manoeuvres and the scarcity of data. This algorithm performed poorly in those scenarios compared with the others.

8.6. Algorithm 2 real results

The summarized results can be seen in Table 6 (results have to consider there are 134 segments without manoeuvres).

Scenario	# man.	% man. detected	% false positives	% false negatives
1	2	100	0	0
2	1	100	0	0
3	1	100	0	0
4	2	50	0	50
1-4	6	83.3	0	16.7
5	2	50	0	50
6	1	0	0	100
7	1	0	6.25	100
8	1	0	11.76	100
5-8	5	20	5.45	80
9	1	0	9.09	100
10	1	0	0	100
9-10	2	0	4.55	100
11	1	100	0	0
12	1	100	0	0
11-12	2	100	0	0
13	1	0	0	100
14	4	25	N/A	75
15	4	25	N/A	75
14-15	8	25	N/A	75
All	**16**	**41.66**	**2.98**	**58.34**

Table 5. Algorithm 1 summarized results.

Scenario	# man.	P1M % man. detected	P1M % false positives	P1M % false negatives	P1D % man. detected	P1D % false positives	P1D % false negatives
1	2	100	0	0	100	0	0
2	1	100	14.29	0	100	14.29	0
3	1	100	0	0	100	0	0
4	2	50	0	50	50	0	50
1-4	6	83.3	3.03	16.7	83.3	3.03	16.7
5	2	50	0	50	50	27.3	50
6	1	0	0	100	0	9.1	100
7	1	0	12.50	100	0	12.50	100
8	1	0	11.76	100	0	11.76	100
5-8	5	20	7.27	80 20	5.45	80	
9	1	0	20	100	0	20	100
10	1	0	25	100	0	25	100
9-10	2	0	22.73	100	0	22.73	100
11	1	100	0	0	100	0	0
12	1	0	10	100	0	10	100
11-12	2	50	7.69	50	50	7.69	50
13	1	0	0	100	0	0	100
14	4	50	N/A	50	25	N/A	75
15	4	100	N/A	0	100	N/A	0
14-15	8	75	N/A	25	62.5	N/A	37.5
All	24	54.16	45.84	8.21	50	7.46	50

Table 6. Algorithm 2 summarized results.

The results of group *1-4* are quite positive for both P1M and P1D, as there is only one false negative and few false positives, out of 33 cases. Thus, out of 6 manoeuvres, 5 were detected. Only the smallest manoeuvre was not detected (similarly to the Algorithm 1 results). Group *5-8* cases contain 5 manoeuvres, but only 1 is detected for both P1M and P1D (another one is felt but with low confidence). There are also some false positives, from 55 total cases, with P1M obtaining a slightly higher rate of false positives. Group *9-10* results are quite bad for both metrics as no manoeuvres are detected; and with a rather high rate of false positive (from 22 cases). One of the 2 manoeuvres present in group *11-12* scenarios is detected, with some false negatives (from 13 cases). Results are the same for both metrics. Group *13* presents no detected manoeuvres (out of one) but no false negatives, for both metrics. Finally, group *14-15* gives 8 segments, all of them with manoeuvres; out of these, 5 are detected for P1D and 6 for P1M. In global, the results can be considered positive, as more than half the manoeuvres are detected with a low rate of false negatives for both metrics, but in need of improvement. P1M seemed to perform better than P1D, with a minimal increase in false negatives. Analysing the results, the causes of errors are similar as for Algorithm 1: From a total of 11 (10) false positives for P1M (resp., P1D), all except two (resp., one) present less than 10 plots in the track following the manoeuvre. Thus, the main cause of false positives is tracks with fewer plots. From a total of 11 (12) false negatives for P1M (resp. P1D), all except one were segments of length equal to or larger than one day. Thus, the main cause of false negatives is longer propagations accumulating additional propagation error. Sometimes these longer propagation periods are due to missed radar observations right after the manoeuvre. The most challenging scenarios are those with TanDEM-X/TerraSAR-X due to the abundance of manoeuvres and the scarcity of data. As opposed to Algorithm 1, this algorithm performs excellently in those scenarios compared with the other.

8.7. MDF real results

Figure 15 has one example of the MDF results. In this case, the filter performs initially quite well thanks to the long smoothing, and the manoeuvre is detected. Even though the initial spike after the manoeuvre is quite large, the filter recovers quickly thanks to the inflation mechanism; comparing with the UKF of Figure 14 the behaviour is much improved. However, a false positive also happens before the end, but it does not impact the filter's performance. Note that the UKF metric did not detect manoeuvres in this segment.

The following conclusions can be derived from the results obtained from the MDF, which cannot be shown here due to lack of space. It can be observed that the filter is, in general, well-behaved even in the presence of manoeuvres and detects many of them. When there are no manoeuvres, it tends to converge, quicker than the UKF without long smoothing, to errors of the order of just a few hundred metres compatible with the errors of the propagator. As in the UKF, there are some instances of rapid increase of errors due to the presence of manoeuvres, particularly when undetected. The MD metric allows for detection of many manoeuvres but also produces a considerable number of false positives. It is not as straightforward as for the RA algorithms to obtain fair statistics, since it is unclear if a detection after one or two tracks should be considered a true or false positive; this is due to the sequential nature of this algorithm, which considers scenarios as a whole, instead of processing segments separately. Thus, the history of each scenario influences the results in several ways. Covariance inflation works well in most cases but in some instances, it might be excessive, inducing large errors in the state.

9. Conclusions and future work

Several methods for the detection of manoeuvres in LEO from radar data have been presented, based on UKF, attributable theory and reachability analysis. Initial simulation results showed that the filter did not detect manoeuvres unless they are rather intense,

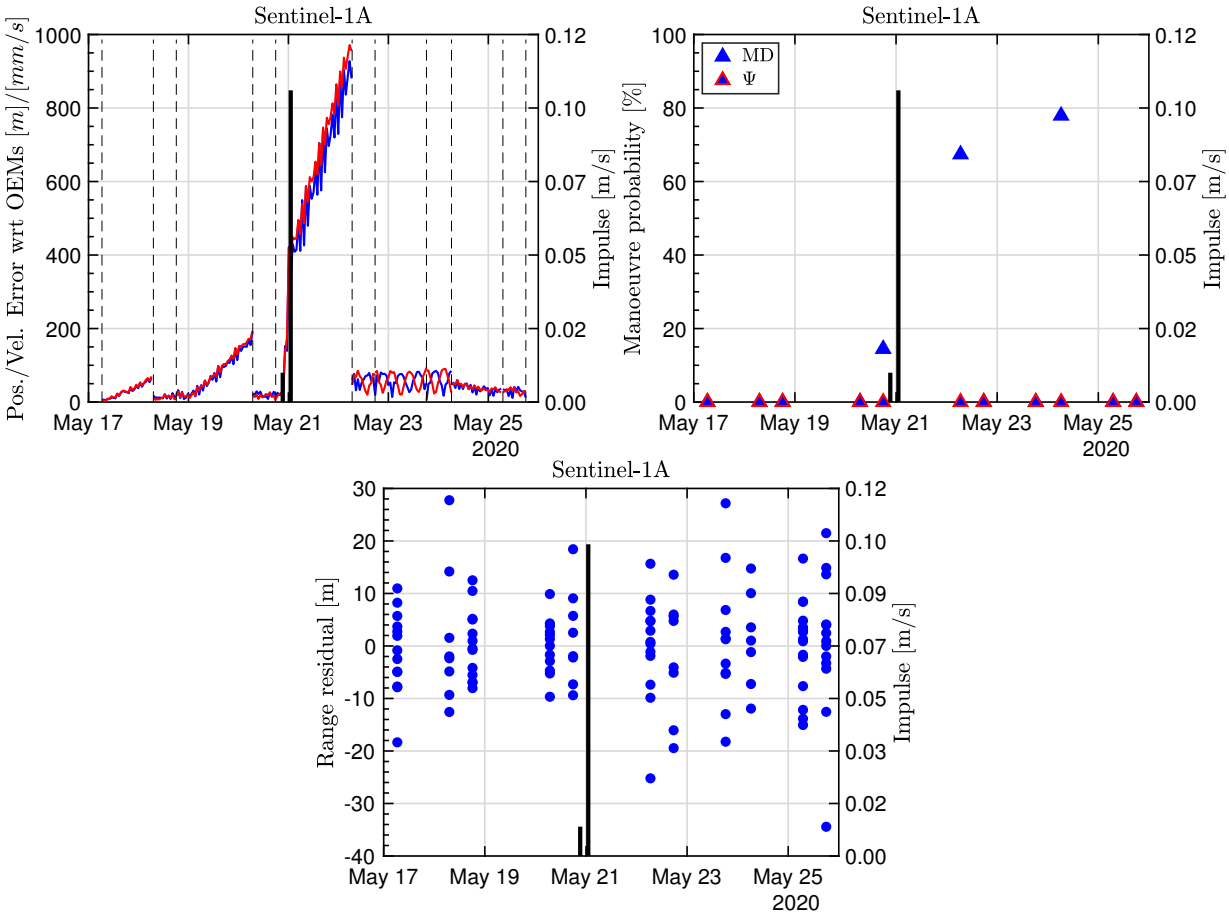

Fig. 15. MDF results for one scenario of Sentinel-1A.

whereas the reachability approach was more sensitive at the price of longer computation times. When tested with real data, the results, while not bad, are in need of improvement. Since the quality of the data was verified, the main identified difficulty was the scarcity of measurements (low number of tracks resulting in long propagation times without information and/or low number of plots in some cases), due to the fact of having a single surveillance radar, the Spanish survey radar S3TSR, as the source of data. Sometimes, more than 24 hours or more went without a measurement and the manoeuvres were not very intense. Future ideas to address this challenge include improvements in the propagator, in the description of uncertainty (e.g. the use of Gaussian mixtures to better describe the error distribution after long propagations), and the use of the intensive surveillance mode of the radar, which can provide hundreds of plots for a single track.

From an innovation point of view, the results can be considered of interest, since most of the results in the literature depend on having numerous measurements (oftentimes, almost continuous data is assumed, which is only realistic in GEO with optical sensors), which was not the case here. The final aim is to have these algorithms integrated in the S3T Cataloguing System in order to provide routine automatic manoeuvre detection capabilities, but all methods presented here can be refined and extended in many directions. Beyond obvious extensions or refinement of the algorithms, there is much to be gained from additional measurements (additional stations, more regular radar tracks, SLR measurements, etc.) as it was shown that RA algorithms under-perform when the number of radar plots are lower. It would be of great interest to draw a set of minimum measurement requirements for the successful application of the manoeuvre detection algorithms.

Acknowledgements

The authors would like to acknowledge Centro de Desarrollo Tecnológico e Industrial (CDTI) for providing funding for this work under the S3T program and in particular Cristina Pérez Hernández for her help and support. The authors are also grateful to Ralph Kahle from DLR for providing the orbital and manoeuvre data for TerraSAR-X and TanDEM-X satellites. Finally, Universidad de Sevilla is acknowledged for partial funding under V PPIT-US and VI PPIT-US.

References

Andersson, J. A., Gillis, J., Horn, G. et al. (2019). CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, *11*(1), 1–36.

Andrea, A., & Maisonobe, L. (2016). Automatic differentiation for propagation of orbit uncertainties on Orekit. In *Stardust Conference 2016*. URL: `https://www.orekit.org/doc/Antolino-2016-automatic-diff-for-prop-of-orbit-uncertainties.pdf`.

Armellin, R., Di Lizia, P., Bernelli-Zazzera, F. et al. (2010). Asteroid close encounters characterization using differential algebra: the case of Apophis. *Celestial Mechanics and Dynamical Astronomy*, *107*(4), 451–470.

Carpenter, J. R., & D'Souza, C. N. (2018). *Navigation Filter Best Practices*. Technical Report TP-2018-219822 NASA. URL: `https://ntrs.nasa.gov/api/citations/20180003657/downloads/20180003657.pdf`.

Chul Ko, H., & Scheeres, D. J. (2016). Tracking maneuvering satellite using thrust-fourier-coefficient event representation. *Journal of Guidance, Control, and Dynamics*, *39*(11), 2554–2562.

Clark, R., & Lee, R. (2020). Parallel processing for orbital maneuver detection. *Advances in Space Research*, *66*(2), 444–449.

Escribano, G., Sanjurjo-Rivo, M., Siminski, J., Pastor, A., & Escobar, D. (2022). Automatic maneuver detection and tracking of space objects in optical survey scenarios based on stochastic hybrid systems formulation. *Advances in Space Research*, *69*(9), 3460–3477.

Goff, G. M. (2015). *Orbit Estimation of Non-Cooperative Maneuvering Spacecraft*. Ph.D. thesis Air Force Institute of Technology. URL: `https://scholar.afit.edu/etd/202/`.

Goff, G. M., Black, J. T., & Beck, J. A. (2015a). Tracking maneuvering spacecraft with filter-through approaches using interacting multiple models. *Acta Astronautica*, *114*, 152–163.

Goff, G. M., Showalter, D., Black, J. T. et al. (2015b). Parameter requirements for noncooperative satellite maneuver reconstruction using adaptive filters. *Journal of Guidance, Control, and Dynamics*, *38*(3), 361–374.

Gomez, R., Salmerón, J. M.-V., Besso, P. et al. (2019). Initial operations of the breakthrough Spanish Space Surveillance and Tracking Radar (S3TSR) in the European context. In *1st ESA NEO and Debris Detection Conference. Presented paper. Darmstadt, Germany*. URL: `https://conference.sdo.esoc.esa.int/proceedings/neosst1/paper/479`.

Guang, Z., Xingzi, B., Hanyu, Z. et al. (2018). Non-cooperative maneuvering spacecraft tracking via a variable structure estimator. *Aerospace Science and Technology*, *79*, 352–363.

Hall, Z., & Singla, P. (2019). A probabilistic approach for reachability set computation for efficient space situational awareness. In *29th AAS/AIAA Space Flight Mechanics Meeting, 2019* (pp. 3001–3020). Univelt Inc.

Holzinger, M., & Scheeres, D. (2009). *Reachability Analysis Applied to Space Situational Awareness*. Technical Report Colorado University at Boulder, departament of aerospace engineering sciences. URL: `https://apps.dtic.mil/sti/pdfs/ADA531761.pdf`.

Holzinger, M. J., Scheeres, D. J., & Alfriend, K. T. (2012). Object correlation, maneuver detection, and characterization using control distance metrics. *Journal of Guidance, Control, and Dynamics*, *35*(4), 1312–1325.

Jain, A., Guého, D., Singla, P. et al. (2019). Stochastic reachability analysis for the hypersonic re-entry problem. In *29th AAS/AIAA Space Flight Mechanics Meeting, 2019* (pp. 2455–2476). Univelt Inc.

Jaunzemis, A. D., Mathew, M. V., & Holzinger, M. J. (2016). Control cost and Mahalanobis distance binary hypothesis testing for spacecraft maneuver detection. *Journal of Guidance, Control, and Dynamics*, *39*(9), 2058–2072.

Jiang, Y., Yang, H., Baoyin, H., & Ma, P. (2019). Extended Kalman filter with input detection and estimation for tracking manoeuvring satellites. *The Journal of Navigation*, *72*(3), 628–648.

Julier, S. J., & Uhlmann, J. K. (1997). New extension of the Kalman filter to nonlinear systems. In *Signal processing, sensor fusion, and target recognition VI* (pp. 182–193). International Society for Optics and Photonics volume 3068.

Kurzhanski, A. B., & Varaiya, P. (2000). Ellipsoidal techniques for reachability analysis: internal approximation. *Systems & control letters*, *41*(3), 201–211.

Li, X. R., & Jilkov, V. P. (2005). Survey of maneuvering target tracking. part v. multiple-model methods. *IEEE Transactions on Aerospace and Electronic Systems*, *41*(4), 1255–1321.

Mainsonobe, L., Cefola, P., Frouvelle, N. et al. (2012). Open governance of the Orekit space flight dynamics library. In *Proceedings of the International Conference on Astrodynamic Tools and Techniques (ICATT)* (pp. 327–343). ESA/ESTEC volume 29.

Pérez, D., Masdemont Soler, J., & Gómez Muntané, G. (2013). Jet Transport propagation of uncertainties for orbits around the earth. In *64rd International Astronautical Congress* (pp. 1–8). URL: `https://upcommons.upc.edu/handle/2117/21808`.

Picone, J., Hedin, A., Drob, D. P. et al. (2002). NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues. *Journal of Geophysical Research: Space Physics*, *107*(A12, 1468), SIA–15.

Poore, A. B., Aristoff, J. M., Horwood, J. T. et al. (2016). *Covariance and Uncertainty Realism in Space Surveillance and Tracking*. Technical Report Numerica Corporation. URL: `https://apps.dtic.mil/sti/pdfs/AD1020892.pdf`.

Reihs, B., Vananti, A., Schildknecht, T. et al. (2021). Application of attributables to the correlation of surveillance radar measurements. *Acta Astronautica*, *182*, 399–415.

Sanchez, J. C., Louembet, C., Gavilan, F. et al. (2019). An Event-Triggered predictive controller for spacecraft rendezvous hovering phases. *IFAC-PapersOnLine*, *52*(12), 97–102.

Särkkä, S., & Solin, A. (2019). *Applied stochastic differential equations*. New York NY: Cambridge University Press.

Schutz, B., Tapley, B., & Born, G. H. (2004). *Statistical orbit determination*. Burlington MA: Elsevier Academic Press.

Serra, R., Yanez, C., & Frueh, C. (2021). Tracklet-to-orbit association for maneuvering space objects using optimal control theory. *Acta Astronautica*, *181*, 271–281.

Siminski, J., Flohrer, T., & Schildknecht, T. (2017). Assessment of post-maneuver observation correlation using short-arc tracklets. *Journal of the British Interplanetary Society*, *70*, 63–68.

Singh, N., Horwood, J. T., & Poore, A. B. (2012). 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* (pp. 2012–159). Univelt San Diego, CA volume 143.

Vananti, A., Schildknecht, T., Siminski, J. et al. (2017). Tracklet-tracklet correlation method for radar and angle observations. In *Proc. 7th European Conference on Space Debris, Darmstadt, Germany* (pp. 18–21). URL: `https://conference.sdo.esoc.esa.int/proceedings/sdc7/paper/925`.

Vazquez, R., Sanchez, J., Montilla, J. M. et al. (2021a). Manoeuvre detection for near-orbiting objects. In *Proc. 8th European Conference on Space Debris, Darmstadt, Germany*. URL: `https://conference.sdo.esoc.esa.int/proceedings/sdc8/paper/238/SDC8-paper238.pdf`.

Vazquez, R., Sanchez, J., Montilla, J. M. et al. (2021b). Two manoeuvre detection probability metrics based on radar measurements and validated with S3TSR data. In *Stardust-R 2nd Global Virtual Workshop (GVW-II) Book of Abstracts* (pp. 153–156).

Vetter, J. R. (2007). Fifty years of orbit determination. *Johns Hopkins APL technical digest*, *27*(3), 239.

Wan, E. A., & Van Der Merwe, R. (2000). The unscented Kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)* (pp. 153–158). IEEE.

Wang, Y., Bai, X., Peng, H., Chen, G., Shen, D., Blasch, E., & Sheaff, C. B. (2021). Gaussian-binary classification for resident space object maneuver detection. *Acta Astronautica*, *187*, 438–446.

Woodburn, J., Carrico, J., & Wright, J. R. (2003). Estimation of instantaneous maneuvers using a fixed interval smoother. *Advances in the Astronautical Sciences*, *116*, 243–260.

Xu, Z., Chen, X., Huang, Y., Bai, Y. et al. (2019). Collision prediction and avoidance for satellite ultra-close relative motion with zonotope-based reachable sets. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, *233*(11), 3920–3937.

Ye, L., Hua, Z., Chuankai, L. et al. (2021). Maneuver detection and tracking of a space target based on a joint filter model. *Asian Journal of Control*, *23*(3), 1441–1453.

Yu, S., Wang, X., & Zhu, T. (2021). Maneuver detection methods for space objects based on dynamical model. *Advances in Space Research*, *68*(1), 71–84.