17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023.

SpaceOps-2023, ID # 223

Fixed Covariance Estimation in GPS-based Operational Orbit Determination: A Realistic Input for Collision Probability Monitoring

Petr Kuchynka

GMV for the Earth Observation Section and Flight Dynamics Division at the European Space Operations Centre (ESOC), Darmstadt, Germany, pkuchynka@gmv.com

Abstract

The European Space Operations Centre currently operates 10 spacecraft in low Earth orbit. For each mission, the Flight Dynamics team runs a daily on-ground orbit determination process that estimates the spacecraft state vector. For 9 out of the 10 spacecraft the process relies exclusively on the navigation solution generated on-board from GPS measurements. The estimation of a realistic covariance matrix on the state vector is not straightforward, because of the non-white distribution of errors in the observations. Yet the covariance is required for realistic collision probability monitoring. This paper presents the derivation of a single fixed covariance matrix for the Sentinel missions, which make up half of the fleet of the 10 spacecraft operated at low orbits. The derivation consists in comparing the history of estimated state vectors with their counterparts obtained with precise orbit determination. The comparison allows to define a conservative ellipsoidal confidence region, which is then translated into a covariance. The approach is simple to implement and provides a realistic input for collision probability monitoring in routine operational situations.

Keywords: Flight Dynamics, Covariance, Orbit Determination.

Acronyms/Abbreviations

Flight Dynamics (FD), European Space Operations Centre (ESOC), Sentinel-1 (S1), Sentinel-2 (S2), Sentinel-5P (S5P), Orbit Determination (OD), Consultative Committee for Space Data Systems (CCSDS), Navigation Data Message (NDM), Space Debris Office (SDO), precise orbit determination (POD), multivariate (mv), Right Ascension of Ascending Node (RAAN).

1. Introduction

The FD team within the Earth Observation section at ESOC currently operates 10 spacecraft in low Earth orbit. The spacecraft are the 5 Sentinels (S1A, S1B, S2A, S2B, S5P) and the 5 Earth Explorers (Aeolus, Cryosat-2, Swarm-A, Swarm-B, Swarm-C). The highest flying is S5P with a reference orbit at an altitude of about 800 km. The lowest is Aeolus at an altitude of about 300 km. For each of these missions an on-ground OD process is run daily, where a propagated orbit is adjusted to available observations and used as a basis for product generation. For some missions the OD runs several times per day. A single product, in the CCSDS NDM format, is delivered to the SDO at ESOC and used for collision monitoring. The file contains the operational orbit, together with a state vector and covariance matrix at an epoch close to the last available observation. Except for Cryosat-2, the OD process relies exclusively on GPS data, more precisely on the GPS-based navigation solution generated on-board. As discussed in [1], the noise contained in the data is not white and, in consequence, the derivation of a realistic covariance for the NDM product is not straightforward. In particular, the covariance provided by the usual least-squares formula, after weighting the observations with the rms of the residuals, is incorrect. It predicts uncertainties inconsistent with day-to-day jumps observed in the operational orbit. The manufacturer of the GPS unit may provide estimates of uncertainties in the navigation solution, but the documented information is not necessarily sufficient to whiten the noise. For example, for S1, the documentation does not quantify the amount of autocorrelation between measurements. In [1], the proposed solution to the problem is to use over-estimated weights calibrated against differences observed in the past between the FD operational orbit and the orbit determined by POD. However, the calibration is sensitive to gaps in the GPS data and to additional estimated parameters such as maneuver calibrations. In practice the calibration can be derived only for a typical scenario well represented in the past operational history, that is for data without gaps and without maneuvers. In the end, the approach in [1] does not provide a significant advantage over the use of a fixed covariance matrix.

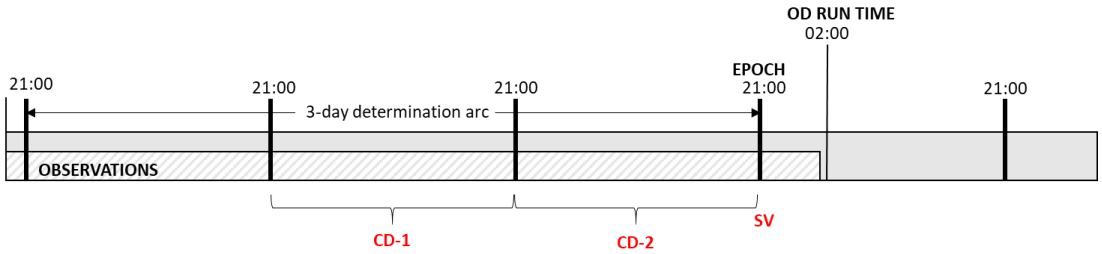

Figure 1. The setup of the S1 OD process. Estimated parameters include the drag coefficients (CD-1, CD-2) and the state vector (SV)

The purpose of this paper is to derive a realistic fixed covariance matrix for the 5 Sentinel spacecraft operated from ESOC. The approach is to compare FD operational orbits with POD over an extended period. The differences allow to define a realistic confidence region which is then translated into a covariance. For S5P, POD data is not available. The covariance obtained for S1 is reused and validated with cross-comparison experiments. In a separate work [2], the same approach was applied to estimate fixed covariances for the 5 Earth Explorers. Both for the Sentinels and the Earth Explorers missions, the fixed covariances are currently implemented in the NDM interface between FD and SDO. They are undergoing a review process within ESOC for operational use in collision probability monitoring.

The next paragraph describes the Sentinel OD setup and some related concepts. The methodology section recalls properties of mv normal distributions and presents the approach to translate a confidence region into a covariance. Systematic comparisons with POD and the actual derivation of the fixed covariance are presented in the results section. Finally, the last section discusses the operational circumstances in which the covariance is valid.

The S1 OD setup is illustrated in Figure 1. The setup in place for S2 and S5P is similar. The OD process runs daily with a determination arc spanning 3 days. Its start and end are fixed with respect to current day. Although fixed in advance, the end time is set shortly before the last available observation nominally expected at run time. The least-squares process estimates the state vector at the end of the arc and 2 daily drag coefficients. The time for which the state vector is estimated, that is the end of the arc, will be referred to as the epoch. The drag coefficient and the adjusted orbit for the middle day are archived. For the first day, drag is treated as a fixed parameter and taken from the archived drag history. The operational archived orbit consists of the archived patches of orbits spanning each the middle day of the determination arc. The last day of the OD arc is also archived. In the following, the orbit made from archived patches of the last day will be called pseudo-operational archived orbit.

S1A and S1B execute orbit control maneuvers every 1 to 2 weeks inside a fixed slot between Wednesday 21:15Z and Thursday 01:45Z. Calibration takes place every hour in the 24 hours following the maneuver. It is carried out independently of the OD process in which the maneuver calibrations are considered as fixed parameters. S2A, S2B and S5P execute only 4 to 5 orbit control maneuver sequences per year. As mentioned already, for the 5 spacecraft, the OD process relies on Earth-fixed positions provided by the navigation solution. Velocity estimates are also available, but they are not used.

2. Methodology

The mathematical concepts necessary to interpret and manipulate covariance matrices are well known. A covariance matrix is the generalization of the concept of variance. The components of the matrix partially characterize the probability density function of a multidimensional random variable. In our case, the random variable is the 6-dimensional state vector derived in the OD process. If the random variable is normally distributed, the covariance matrix, together with the expectancy, fully define its probability density function. There are 3 properties of the mv normal distribution that are relevant for this paper. First, if a random variable $\mathbf{X}$ follows a mv normal distribution, any linearly dependent random variable defined as $\mathbf{Y} = A\,\mathbf{X}$ will also follow a normal distribution with:

$$\mathrm{cov}(\mathbf{X}) = A\ \mathrm{cov}(\mathbf{Y})\ A^T. \qquad (1)$$

Second, the probability density function of a normally distributed variable has a geometric interpretation. Any surface of constant density is defined by:

$$(\mathbf{X}-E[\mathbf{X}])^T\ \mathrm{cov}^{-1}(\mathbf{X})\ (\mathbf{X}-E[\mathbf{X}]) = d^2,$$

where d is a constant. For a given d, the equation defines an ellipsoid centred at E(**X**). Its principal axes are the eigenvectors of cov(**X**), its semi-major axes are the square roots of the eigenvalues of cov(**X**) multiplied by d. In the following, d will be called scaling factor and the ellipsoid defined by d=1 will be called unitary. Third, the probability for the value of **X** to fall within the unitary ellipsoid of dimension n scaled by a factor d is straightforward to compute. The probability is given by the so-called chi-squared distribution with n degrees of freedom. Probabilities for various values of n and d are provided in Table 1 below. The table represents the generalization of the 68-95-99.7 rule known for the 1-dimensional normal distribution. With higher dimensions, the chance of having a significant fraction of the observations within the unitary '1-sigma' ellipsoid drops dramatically.

Table 1. Probability of the outcome of a mv normal distribution of dimension n to be within a scaling factor d of the unitary ellipsoid.

	d = 1	d = 2	d = 3	d = 4	d = 5	d = 6
n = 1	68.269	95.450	99.730	99.994	100.000	100.000
n = 3	19.875	73.854	97.071	99.887	99.998	100.000
n = 6	1.439	32.332	82.642	98.625	99.966	100.000

The simplicity of propagating a mv normal distribution through linear operators makes it the distribution of choice in practical implementations. A typical OD process relies on least-squares for the estimation of the state vector at epoch. The assumption of normally distributed errors in measurements will translate into a normal distribution of errors in the estimated state vector. Similarly, the linear dependence between state vector at epoch and state vector at a later time guarantees that normally distributed errors at epoch will also translate into normally distributed errors at a later time. However, the errors in a particular problem may not be normally distributed, as it's the case with the errors in the navigation solution. The consequence is a non-normal distribution of errors on the estimated parameters. Although the agreed interface with SDO is a covariance matrix, it has to be kept in mind that the covariance matrix will be interpreted as defining a normal distribution with negligible expectancy. In other words, the actual interface is a mv normal distribution of errors centred on zero. It is then necessary to model the state vector errors as normal distributions, even if they are not. This imperfection is understood by both FD and SDO. It is considered, in operational practice, as unavoidable.

In this paper, we use the concept of confidence region to translate an arbitrary distribution of errors into a mv normal distribution. For an ellipsoidal volume encompassing some percentage of the errors, the translation is operated by matching the surface of constant density of a mv normal distribution to the ellipsoidal volume. The unitary ellipsoid is then determined by scaling the ellipsoidal volume with the scaling factor of Table 1 that corresponds to the portion of errors in the confidence region. For example, in 6 dimensions, for a confidence region encompassing about 99% of the errors, the unitary ellipsoid is obtained by scaling down the confidence region by a factor of 4. For the given scaling factor, the percentage of errors within the confidence region matches the percentage in the normal distribution. The approach is necessarily imperfect. For example, if the modelled distribution is fat-tailed, the unitary ellipsoid scaled by a factor smaller than 4 will encompass more points than expected from a normal distribution. For scaling factors larger than 4, the situation is reversed and less points then predicted will be observed inside the iso-density surfaces.

3. Results

3.1 Sentinel-1

Figure 2 shows differences in the orbital frame, in both position and velocity, between the S1A orbit and the orbit derived from POD. Only the last orbit before epoch is plotted. The differences are smaller than 1 m in position and smaller than 1 mm/s in velocity. The comparison spans 2.5 years, from January 2019 to July 2021. A small portion of the data was considered as degraded and was discarded from the comparison. This includes either data within +/- 2 orbits around maneuvers or full days where the difference with POD appears unusually large. The latter case is limited to 6 specific days in the 2.5-year interval, likely due to unexpected operational events, such as gaps in GPS data or a failed OD.

The components of the differences in the orbital frame are strongly correlated. The left side of Figure 3 represents the correlation matrix estimated from sample covariance. Recall this is equal to the covariance matrix with a change of units on each dimension. Radial position correlates with along-track velocity on one hand, and, on the other hand, along-track position correlates with radial velocity.

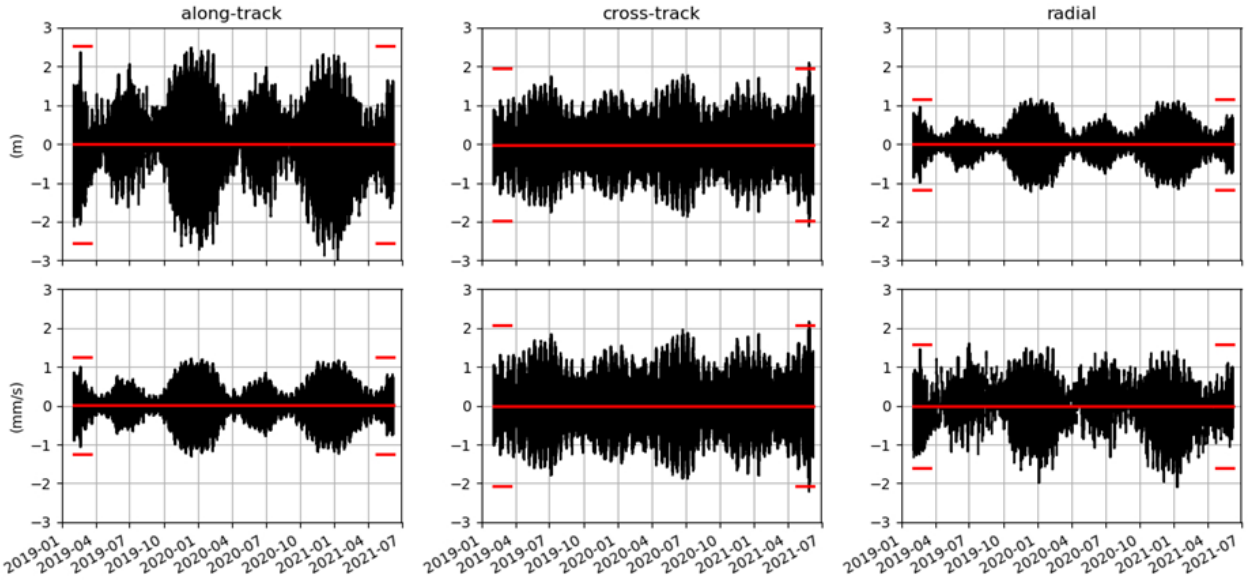

Figure 2. Difference between the S1A operational orbit and POD in the orbital frame. The red lines indicate mean and 3 standard deviations.

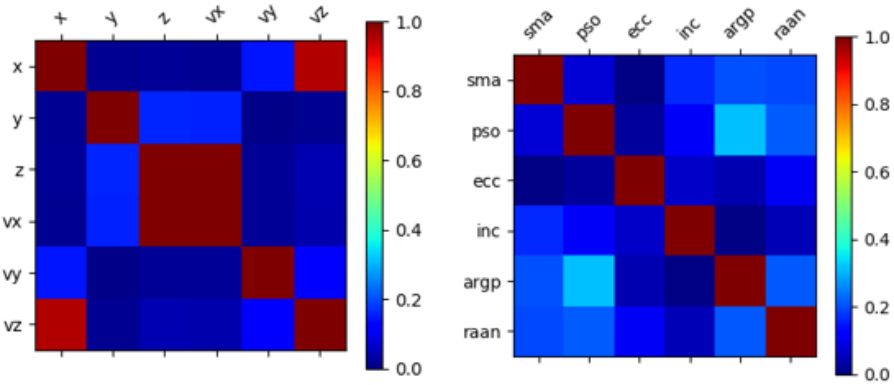

Figure 3. Correlation matrix for differences with POD in orbital frame and Keplerian elements.

Finding an ellipsoidal confidence region thus requires defining both the ellipsoid size and orientation. The right side of Figure 3 shows the correlation matrix computed for differences in Keplerian orbital elements. The matrix appears approximately diagonal making it easier to derive a confidence region. The evolution of the differences in orbital elements are shown on Figure 4. On some components, in particular the semi-major axis, the differences contain a systematic bias. The differences exhibit a seasonal trend, being for some components significantly larger around the winter and summer solstices. We define a conservative ellipsoidal confidence region, aligned with the Keplerian coordinate system, with the following semi-major axes:

- 0.17 m in orbital semi-major axis (sma)
- 0.17 arcsec in argument of latitude (pso)
- 3E-7 in eccentricity (ecc),
- 0.08 arcsec in inclination (inc),
- 90 arcsec in argument of perigee (argp)
- 0.12 arcsec in right ascension of ascending node (raan)

The confidence region contains 100% of all differences over the considered time span. Although the comparison includes 160000 state vectors, the number of independent samples is in fact given by the number of OD runs or, equivalently, the number of days in the 2.5-year interval, that is about 900. The 6 outlier days, removed from the comparison, represent about 1% of the independent data.

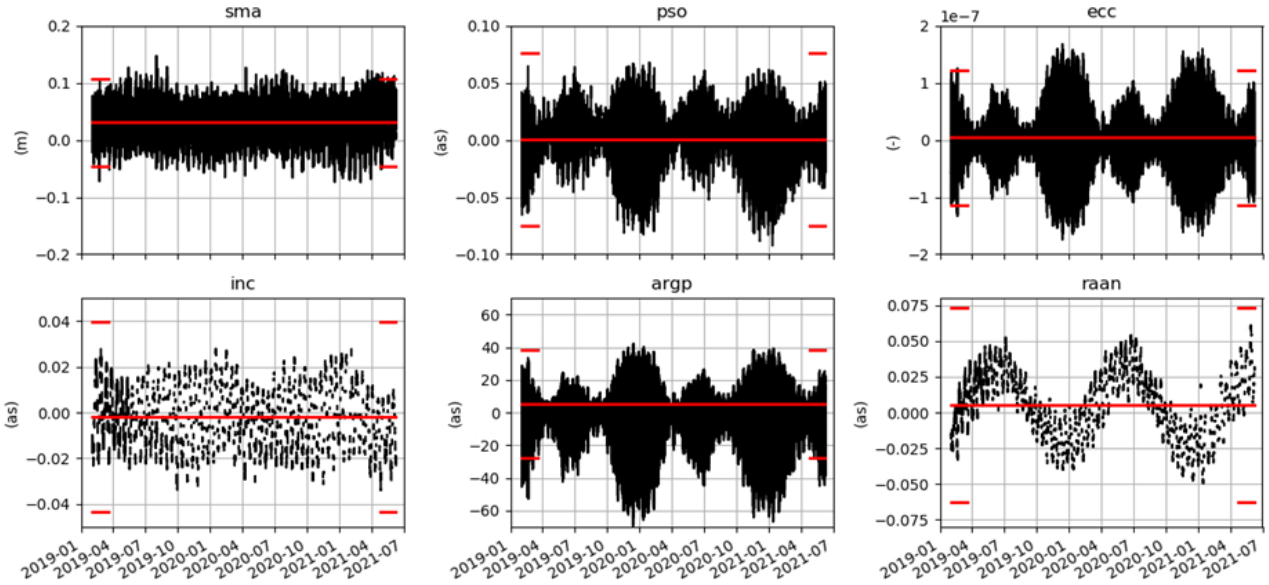

Figure 4. Difference between the S1A operational orbit and POD in Keplerian elements. The red lines indicate mean and 3 standard deviations.

The confidence region must be translated into a mv. normal distribution. In Table 1, for a 6-dimensional distribution, a fraction of approximately 99% of samples inside the scaled unitary ellipsoid is reached for a scaling factor of 4. The confidence region defined previously is thus scaled down by a factor of 4 and interpreted as the unitary ellipsoid of a mv normal distribution. Table 2 below provides the corresponding diagonal covariance matrix. This is the default covariance matrix representing the uncertainty of the state vector at epoch. Table 3 compares, for the default covariance, the fraction of samples within various scalings of the unitary ellipsoid to the expected fraction from a mv normal distribution. The comparison is provided, both for the total 6-dimensional sample population as well as for a 1-dimensional projection on the orbital semi-major axis coordinate. The fraction of differences within the scaled unitary ellipsoid is in general larger than expected, indicating that the mv normal distribution is conservative in modelling the error distribution.

Table 2. Fixed constant covariance matrix, non-diagonal elements are zero.

sma (km^2)	pso (rad^2)	ecc (-)	inc (rad^2)	argp (rad^2)	raan (rad^2)
1.806e-09	4.245e-14	5.625e-15	9.402e-15	1.190e-08	2.115e-14

Table 3. Fraction of differences, for one orbit preceding epoch, within the unitary ellipsoid of the fixed constant covariance scaled by a factor d. In parentheses the corresponding values expected from a mv normal distribution.

	d = 1	d = 2	d = 3	d = 4	d = 5	d = 6
(6D, total)	(1.439)	(32.332)	(82.642)	(98.625)	(99.966)	(100.000)
S1A, total	9.736	78.745	98.886	100.000	100.000	100.000
S1B, total	9.436	79.743	99.002	100.000	100.000	100.000
S2A, total	7.192	87.951	99.982	100.000	100.000	100.000
S2B, total	4.320	90.469	99.881	99.999	100.000	100.000
S5P, total	57.535	99.995	100.000	100.000	100.000	100.000
(1D)	(68.269)	(95.450)	(99.730)	(99.994)	(100.000)	(100.000)
S1A, sma	67.672	98.538	99.991	100.000	100.000	100.000
S1B, sma	68.178	98.589	99.995	100.000	100.000	100.000
S2A, sma	65.534	99.067	100.000	100.000	100.000	100.000
S2B, sma	65.502	98.676	99.958	99.999	100.000	100.000
S5P, sma	99.973	100.000	100.000	100.000	100.000	100.000

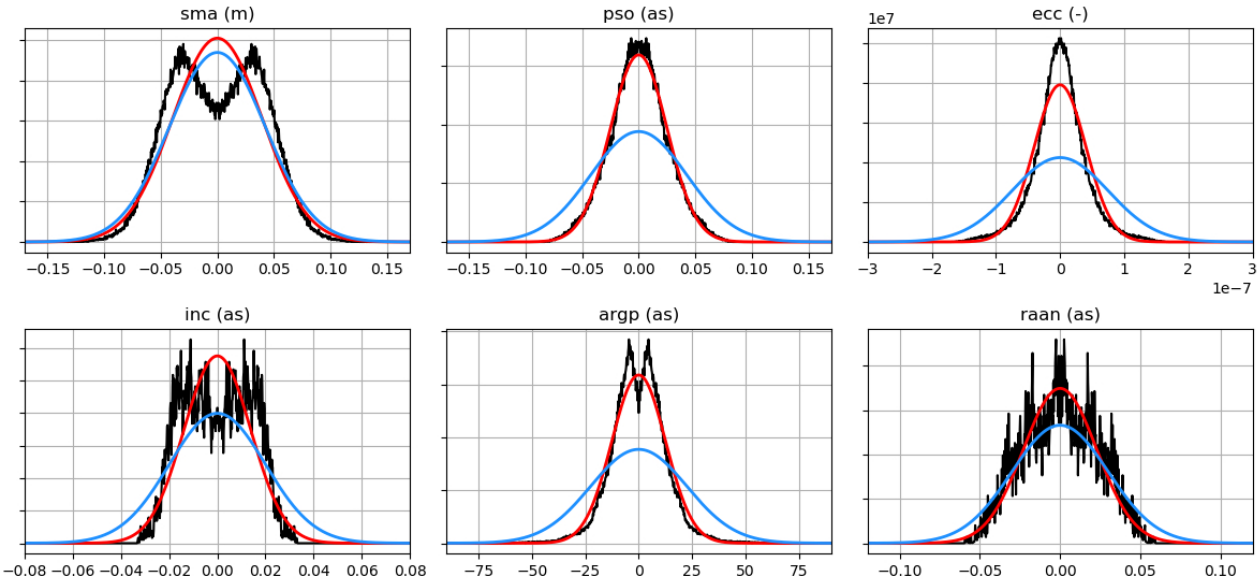

Figure 5. Symmetrized 1-D distributions of S1A differences with POD in Keplerian coordinates. In blue, projection of the mv normal distribution derived from the confidence region. In red, normal distribution derived from sample mean and standard deviation.

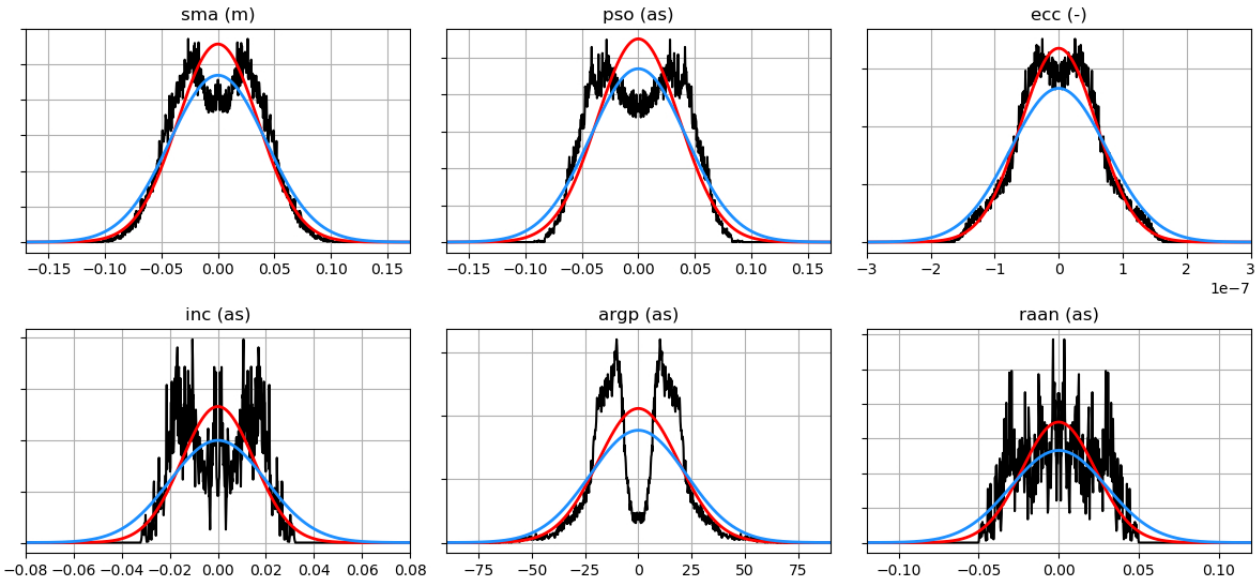

Figure 6. Analogue of Figure 5 obtained using only data between November and March.

Figure 5 compares different projections of the error distribution with its mv normal distribution model. For convenience, the 1-D distributions have been mirrored around the origin by adding for each error an error of the opposite sign. This allows to compare the unbiased model distribution with the biased data. The information content of the data is diluted by the mirror transformation, but the data becomes easier to model. The figure confirms that the model is conservative in all coordinates, except in semi-major axis, where the standard deviation of the model follows the standard deviation of the mirrored distribution. Figure 6 is analogous to Figure 5, except it has been derived using only data between November and March, both for 2019 and 2020, where the comparison with POD is particularly unfavourable. The exact numbers used earlier to define the confidence region were chosen to achieve the match observed in Figure 6 between the model and projected error distributions. During the unfavourable comparison period, around the winter solstice, the model distribution does not appear particularly conservative.

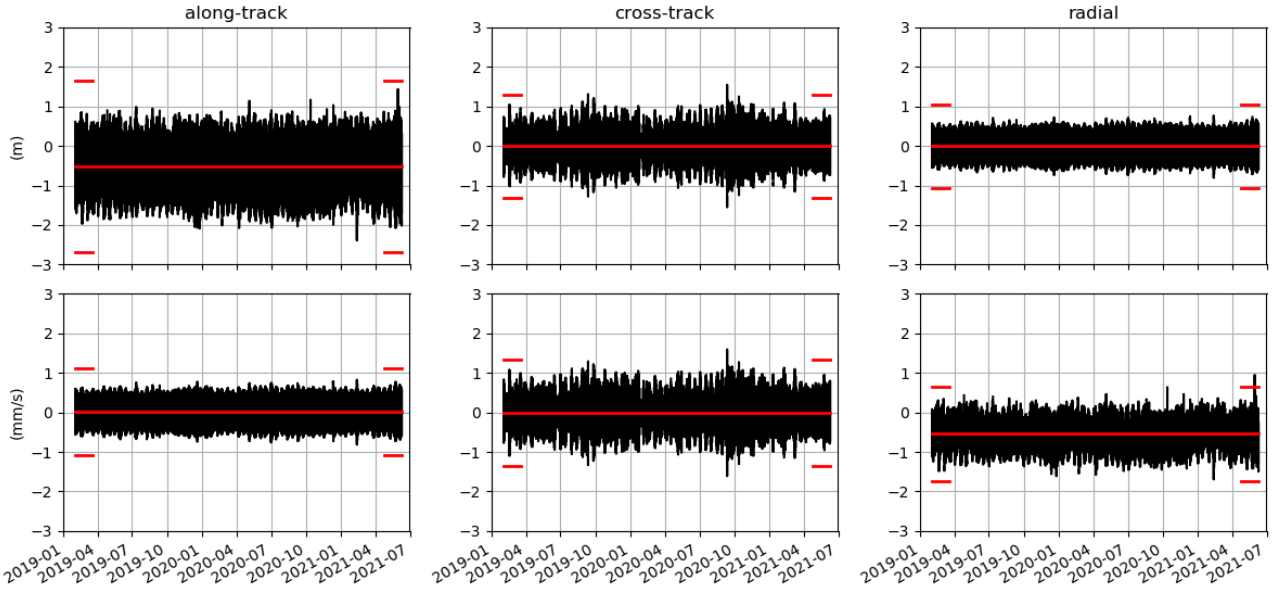

Figure 7. Difference between the S2A operational orbit and POD in the orbital frame. The red lines indicate mean and 3 standard deviations.

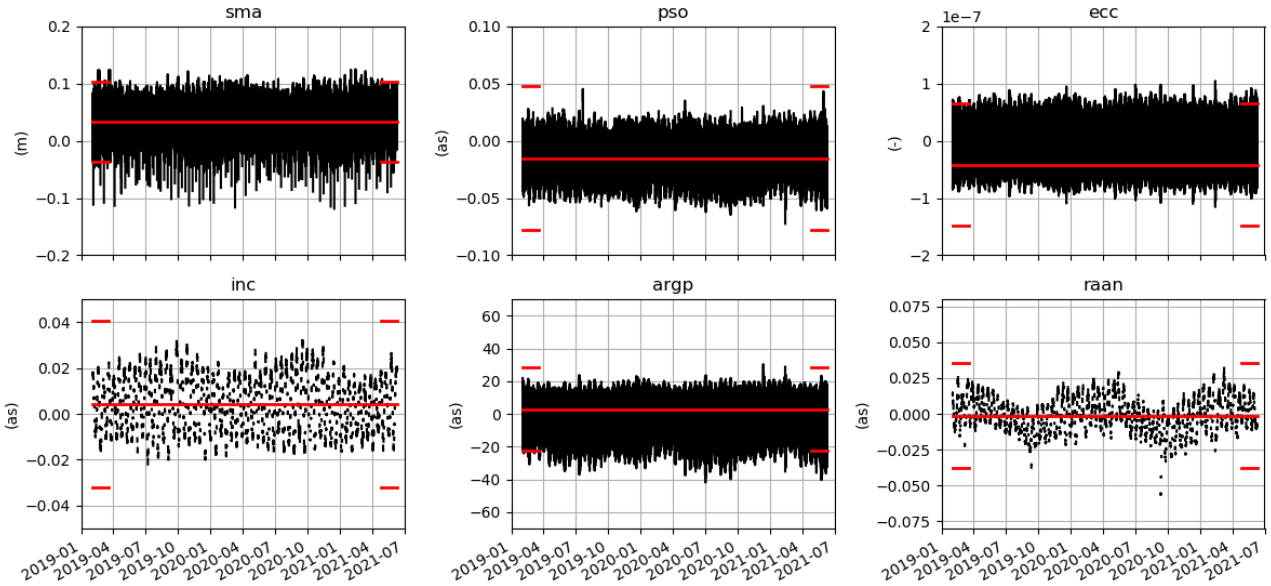

Figure 8. Difference between the S2A operational orbit and POD in Keplerian elements. The red lines indicate mean and 3 standard deviations.

The default covariance derived for S1A is also valid for S1B. Table 3 indicates that the covariance is indeed also conservative in representing the S1B error during the one orbit preceding epoch. The iso-density surfaces of the model mv normal distribution encompass in general more sample points than theoretically predicted. The S1B comparisons with POD are not shown, because the figures are almost identical to those derived for S1A. The same applies for the 1-D projections of the distributions.

3.2 Sentinel-2

The approach to define a constant default covariance matrix is to re-use the covariance derived for S1 and to show that it is conservative in representing the errors on the S2A and S2B state vectors at epoch. Figure 7 shows differences in the orbital frame, in both position and velocity, between the S2A orbit and the orbit derived from POD. As for S1A, only the last orbit before epoch is plotted.

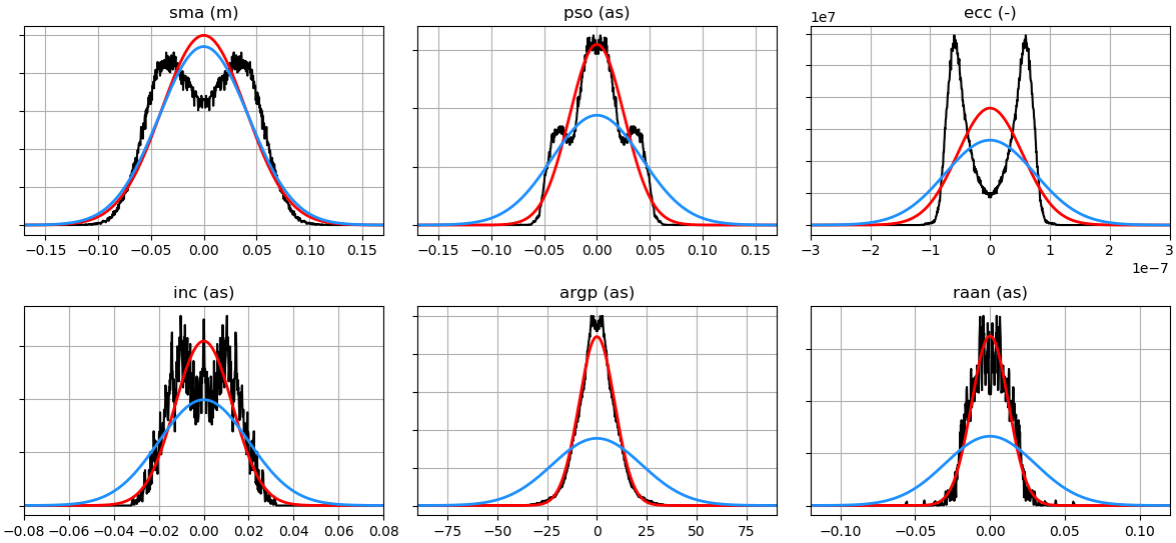

Figure 9. Symmetrized 1-D distributions of S2A differences with POD in Keplerian coordinates. In blue, projection of the mv normal distribution derived from the confidence region. In red, normal distribution derived from sample mean and standard deviation.

The differences are smaller than 1 m in position and smaller than 1 mm/s in velocity. As for S1, the comparison spans 2.5 years. A small portion of the data was discarded. This includes either data within +/- 2 orbits around maneuvers or full days where the difference with POD appears unusually large. The latter case is limited to 6 specific days out of which 2 days contain large out-of-plane maneuvers. This suggests that such maneuvers may significantly degrade the OD solution beyond +/- 2 orbits around execution time. The differences are not affected by a seasonal variation as for S1. On the other hand, the along-track position and radial velocity components are suffering from a bias of 0.5 m and 0.5 mm/s respectively. This corresponds to a constant shift in the PSO of the orbit.

Figure 8 shows the differences in orbital elements. Besides the bias in PSO, a bias appears also on eccentricity and semi-major axis. The latter was also observed on S1. A small seasonal trend is present within the RAAN component, but its magnitude is significantly smaller than for S1. On the figure, the average value for eccentricity is shifted with respect to the middle of the band representing the difference. This is because, over one orbital period, the difference is nearly constant except for a short spike. The spike contributes to most of the width of the band on the figure. A similar effect is observable on the argument of the perigee.

Table 3 compares for the fixed covariance in Table 2, the fraction of S2A differences within various scalings of the unitary ellipsoid to the expected fraction from a mv normal distribution. As for S1A, the fraction is in general larger than expected. Figure 9 compares different projections of the error distribution with the model mv normal distribution based on the default covariance. As for S1, the 1-D distributions have been mirrored around the origin by adding for each error an error of the opposite sign. The figure confirms that the model is conservative in all coordinates, except in semi-major axis, where the standard deviation of the model follows the standard deviation of the mirrored distribution. The default covariance is also valid for S2B. This is confirmed by Table 3. S2B comparisons with POD in Cartesian and Keplerian coordinates are not shown. The figures are almost identical to those obtained for S2A, except the along-track bias amounts only to 0.2 m instead of 0.5 m.

3.3 Sentinel-5P

The approach to define a constant default covariance matrix is to re-use the covariance derived for S1 and S2. The validation process cannot rely on POD, because the GPS pseudo-range measurements, required in the POD process, are currently not transmitted by the spacecraft. As a replacement for POD, we use the archived operational orbit. Recall that this orbit is made from the middle day of the determination arc. Thus, it includes the segment before epoch, but estimated one day later. In general, this estimation is more accurate than the estimation from the 3rd day of the determination arc, which, at a given day, contains the state vector at epoch transmitted in the NDM.

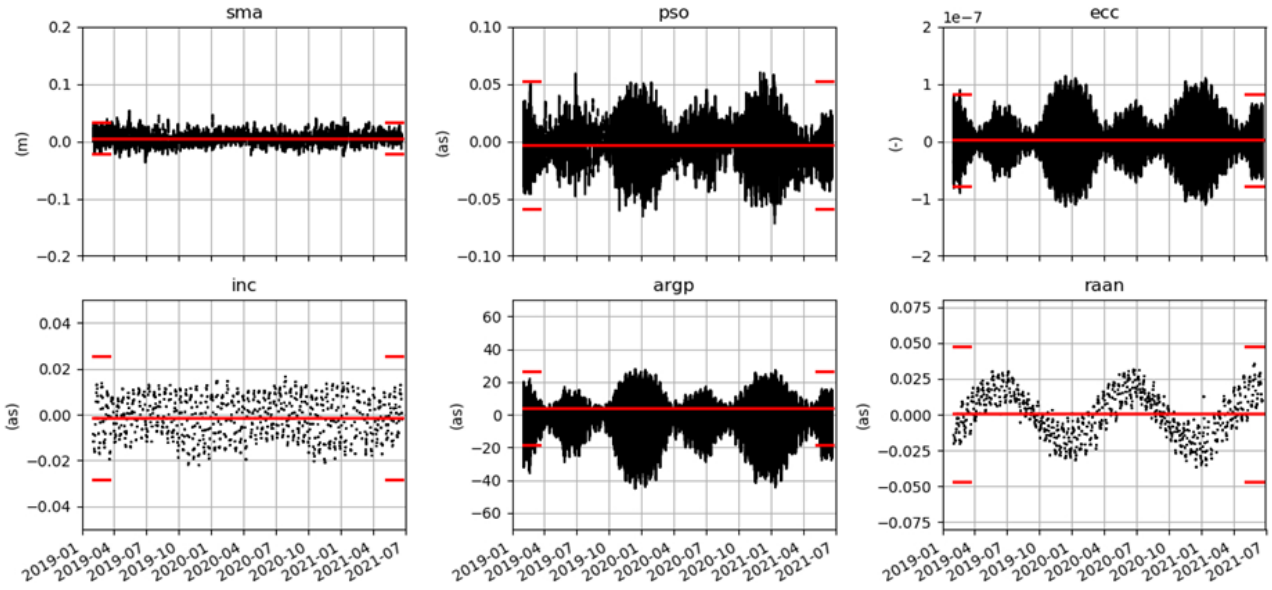

Figure 10. Difference between the S1A operational and pseudo-operational orbits in Keplerian elements. The red lines indicate mean and 3 standard deviations.

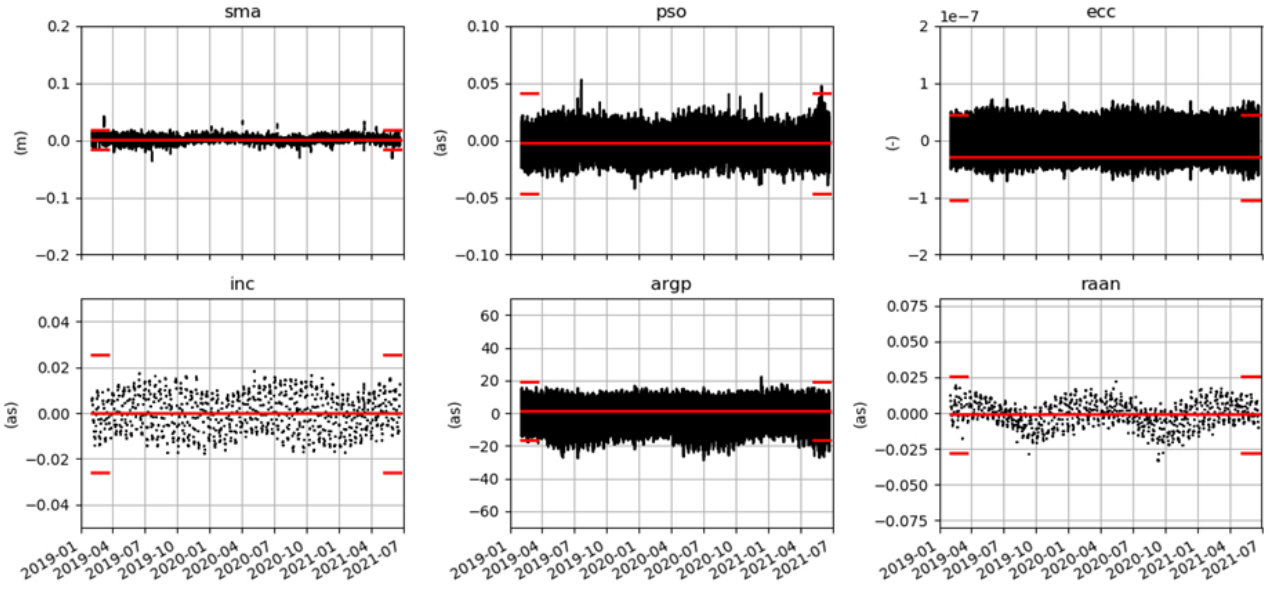

Figure 11. Difference between the S2A operational and pseudo-operational orbits in Keplerian elements. The red lines indicate mean and 3 standard deviations.

As discussed in the introduction, the 3rd day of the determination arc is stored in the archived pseudo-operational orbit. In the end, the proposed validation process consists in a comparing the archived operational and pseudo-operational orbits on a 2-year interval. This cross-comparison experiment is carried out for S1A and S2A as well. We can then evaluate to which extent the cross-comparison reproduces the actual errors as derived from POD.

Figure 10 and Figure 11 compare the pseudo-operational and operational orbits for S1A and S2A. The differences in the cross-comparisons are about two times smaller than the differences with POD (see Figure 4 and Figure 8). An exception is the semi-major axis, where the cross-comparison underestimates at least by an order of magnitude the actual error. The cross-comparison for S5P is provided in Figure 12. The estimated errors are very similar to those obtained with the same approach on S1A and S2A.

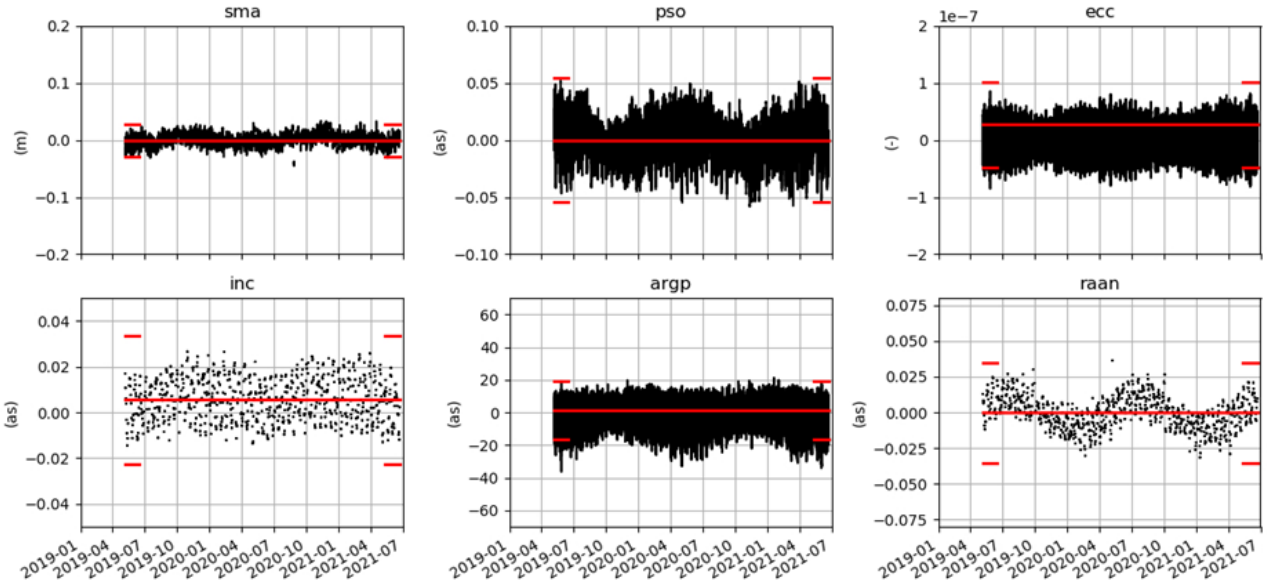

Figure 12. Difference between the S5P operational and pseudo-operational orbits in Keplerian elements. The red lines indicate mean and 3 standard deviations.

This suggests similar errors in the OD process and the use, on S5P, of the same default covariance as the one validated for S1 and S2 through comparisons with POD. The fraction of the S5P cross-comparison differences in Table 3 is systematically larger than the expected theoretical values. This is of course due to the likely under-estimation of the errors by the cross-comparison approach.

5. Discussion

In Figure 4, the RAAN is subject to a conspicuous oscillation that reaches its maximum around the summer and winter solstices. The origin of the oscillation has been tracked down to the effect of ocean tides, which are not included in the operational FD dynamical model. The effect is much less prominent on S2. The impact of ocean tides on the propagated orbit is rather small, at its peak, around the solstices, the effect causes a cross-track oscillation with an amplitude growing by about 0.6 m per day.

In the previous section, the fixed covariance has been derived and validated for the 5 Sentinel spacecraft. Although the covariance in Table 2 is in Keplerian elements, the NDM file includes a state vector and covariance in ICRF Cartesian elements. The necessary conversion is achieved with Eq. 1 where the A matrix is set equal to the Jacobian of the Cartesian elements with respect to the Keplerian ones. The Jacobian is time-dependent and must be calculated at epoch.

The fixed covariance in Table 2 it is to be used only in typical operational circumstances well represented in the 2.5-year comparison interval. In case of large gaps in the GPS data, solar storms, the use of radiometric tracking, etc., the covariance must be marked as invalid. A priori, the covariance is also invalid in situations where a maneuver is inside the determination arc. For S2A/B and S5P, this concerns a few days per year. For S1A/B, maneuvers are executed every week. Fortunately, if they are executed in the weekly slot, from 21:15Z to 01:45Z, the maneuvers do not invalidate the covariance. As shown in Figure 1, the OD epoch is set shortly before the slot. Any data after epoch is not considered in the OD and maneuvers executed inside the slot will appear in the determination arc at least 24 hours after run time. Any such maneuver will be well calibrated and with minimum effect on the state vector uncertainty.

In case the NDM is marked as invalid, collision monitoring must rely on the last valid product, typically the product automatically distributed on the previous day. The propagation of the covariance over a longer timespan will naturally result in a larger uncertainty on the spacecraft position. Nevertheless, this inflation seems preferable to the use of an incorrect input. It is also possible that the FD operator manually adapts the OD setup and generates an NDM file with the epoch set at the latest possible time where the validity is still guaranteed. Even if the latest

automatically distributed covariance is valid, a manual intervention may shift the determination arc to include the latest available observations, thus shortening the time between epoch and the closest approach.

The fixed covariance has been estimated on the most unfavourable parts of the S1 vs POD comparisons (around the winter equinox). It is thus rather conservative. Though conservative, the contribution to the propagated position is typically not much larger than the uncertainty introduced by imperfect modelling of atmospheric drag. The fixed covariance was derived for a scaling factor of 4 on the initially defined confidence region. An error on the state vector falling with this confidence region propagates and grows with time. However, spot check propagations indicate that the effect is very limited in the cross-track and radial directions. For propagations of several days, the errors remain on the order of a few meters. This is small compared with the uncertainty on the diameter of the spacecraft sphere used in the computation of the collision probability. Indeed, the body of S1 fits in a diameter of 4 m, on the other hand solar panels require a diameter of 20 m leaving most of the defined sphere empty. In the along-track direction, the error will increase linearly with time at a rate determined by the initial error in semi-major axis. Within the confidence region, the maximum error in semi-major axis is 0.17 m. This induces an along-track drift of 25 m per day. Table 1 indicates that in 70% of the cases, the drift will be 4 times smaller, that is about 6 m per day. Typical day-to-day variations in the estimated drag coefficients are about 15% of the nominal value. Numerical propagations, carried out near solar minimum, show that an error of this magnitude in the predicted drag coefficient induces an along-track shift of 3 m in the 1st day and then a drift of about 7 m per day. The effect is similar to a small maneuver spread over 1 day. These effects randomly cumulate or compensate at later days, resulting in a random-walk in along-track drift with a step of 7 m per day. On the short term, say within 3 days of epoch, the drag contribution to the propagated uncertainty is about the same as the contribution from the uncertainty on the state vector at epoch. On the long term, drag effects become dominant. Similar considerations apply also for S2 and S5P. Tightening the fixed covariance would not bring any benefit in terms of collision monitoring unless day-to-day variations on the drag coefficients is reduced well below the 15%. This would require a substantial improvement in the atmospheric drag model.

6. Conclusions

In the present paper we derived from systematic comparison with POD a single fixed covariance for the Sentinel family of spacecraft currently operated at ESOC. The covariance represents the state vector uncertainty at epoch in the FD operational OD process and is intended to support collision probability monitoring. It is to be considered valid only in typical operational circumstances. For S1A/B, these include the presence of orbit control maneuvers inside the determination arc. Although conservative, the matrix is also optimal in the sense that its further reduction would not bring any benefit in terms of spacecraft propagated uncertainty.

Acknowledgements

The Navigation Support Office at ESOC has kindly provided access to the Sentinel POD orbits. The author thanks GMV for the organisational and financial support necessary for attending the SpaceOps conference. Thanks to Z.K. for proofreading the paper.

References

[1] P. Kuchynka, M. A. Serrano, K. Merz, J. Siminski. Uncertainties in GPS-based operational orbit determination: A case study of the Sentinel-1 and 2 satellites. The Aeronautical Journal, 124, 2020.

[2] OPS-GFE ORBMAN Team. Approach to the generation of covariance matrices for Earth orbiting missions. ESA Technical Note (ESA-DOPS-FD-TN-0004), 2023.