IAA-AAS-DyCoSS2 14-05-04

LONG TERM NONLINEAR PROPAGATION OF UNCERTAINTIES IN PERTURBED GEOCENTRIC DYNAMICS USING AUTOMATIC DOMAIN SPLITTING

Alexander Wittig[*], Pierluigi Di Lizia[†], Roberto Armellin[‡], Franco Bernelli Zazzera [§], Kyoko Makino [¶], and Martin Berz [‖]

Current approaches to uncertainty propagation in astrodynamics mainly refer to linearized models or Monte Carlo simulations. Naive linear methods fail in nonlinear dynamics, whereas Monte Carlo simulations tend to be computationally intensive. Differential algebra has already proven to be an efficient compromise by replacing thousands of pointwise integrations of Monte Carlo runs with the fast evaluation of the arbitrary order Taylor expansion of the flow of the dynamics. However, the current implementation of the DA-based high-order uncertainty propagator fails in highly nonlinear dynamics or long term propagation. We solve this issue by introducing automatic domain splitting. During propagation, the polynomial of the current state is split in two polynomials when its accuracy reaches a given threshold. The resulting polynomials accurately track uncertainties, even in highly nonlinear dynamics and long term propagations. Furthermore, valuable additional information about the dynamical system is available from the pattern in which those automatic splits occur. From this pattern it is immediately visible where the system behaves chaotically and where its evolution is smooth. Furthermore, it is possible to deduce the behavior of the system for each region, yielding further insight into the dynamics. In this work, the method is applied to the analysis of an end-of-life disposal trajectory of the INTEGRAL spacecraft.

INTRODUCTION

Nonlinear propagation of uncertainties plays a key role in astrodynamics. Orbit determination is affected by measurement errors; consequently, the knowledge of the state of any spacecraft or celestial body is characterized by an estimable level of uncertainty. Even the dynamical models

[*] AstroNet II experienced researcher and post-doctoral research fellow, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milano, Italy, alexander.wittig@polimi.it

[†] Post-doctoral research fellow, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milano, Italy, pierluigi.dilizia@polimi.it

[‡] Lecturer in Astronautics, School of Engineering Sciences, University of Southampton, Southampton, SO17 1BJ, UK, roberto.armellin@soton.ac.uk

[§] Full professor, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milano, Italy, franco.bernelli@polimi.it

[¶] Full professor, Department of Physics and Astronomy, Michigan State University, 567 Wilson Road, East Lansing, MI 48824, USA, makino@msu.edu

[‖] Full professor, Department of Physics and Astronomy, Michigan State University, 567 Wilson Road, East Lansing, MI 48824, USA, berz@msu.edu

used to propagate the motion are synthesized by adopting approximations that affect the accuracy of the orbital predictions. In addition, the size of the uncertainty set tends to quickly increase along the trajectory due to nonlinearities. Nonlinearities are not confined to object dynamics: even simple conversions between different coordinate systems (e.g. the conversion from polar to Cartesian coordinates that forms the foundation for the observation models of many sensors) introduce significant nonlinearities and, thus, affect the accuracy of classical uncertainty propagation techniques.

Uncertainty propagation and estimation in nonlinear systems is extremely difficult. Present-day approaches mainly refer to linearized propagation models[1,2,3] or full nonlinear Monte Carlo simulations.[4] The linear assumption significantly simplifies the problem, but the accuracy of the solution drops off in case of highly nonlinear systems and/or long time propagations. On the other hand, Monte Carlo simulations provide true trajectory statistics, but are computationally intensive and therefore, in many cases, unmanageable. A significant step towards the solution of such issues was taken by Park and Scheeres ,[5] who derived an alternate way to map the statistics by approximating the flow of the dynamics in Taylor series. This approach attains a good agreement with Monte Carlo simulations. However, it suffers of high computational complexity to derive the dynamics of the high order tensor, especially for high fidelity dynamics.

Alternative methods based on the use of differential algebra were already presented by the authors in past works. Differential algebra (DA) supplies the tools to compute the derivatives of functions within a computer environment.[6] More specifically, by substituting the classical implementation of real algebra with the implementation of a new algebra of Taylor polynomials, any function f of n variables is expanded into its Taylor polynomial up to an arbitrary order k with limited computational effort. As a main consequence, the Taylor expansion of the flow of any ordinary differential equation can be readily obtained by carrying out all the operations of any explicit integration scheme in the DA framework. Nonlinear uncertainty propagation can take advantage of the resulting high order expansions. For example, as the accuracy of the Taylor expansion can be tuned by adjusting the expansion order, the approach of classical Monte Carlo simulations can be enhanced by replacing thousands of integrations with evaluations of the Taylor expansion of the flow. Moreover, since the DA approach can provide the higher order derivatives of any function, also higher order propagation schemes can be enhanced and speeded up.

Differential algebra has already proven its efficiency in the non-linear propagation of uncertainties within different dynamical models, including two-body dynamics,[7] (n+1)- body dynamics,[8] and geocentric models (including Earth's gravitational harmonics, solar radiation pressure, shadows, and third body perturbations).[9] Nonetheless, the accuracy of the method tends to decrease drastically in the long term propagation of even relatively small sets of initial conditions under nonlinear dynamics. The long-term propagation of a disposal orbit for an Earth satellite is such an example: initially rather small uncertainties on the spacecraft velocity propagate into very different final positions over the typical time frames of hundreds of revolutions. While performing better than classical linearized methods, the current implementation of the DA uncertainty propagator tends to be inaccurate and impractical in such cases, due to the extremely high order required to describe the resulting uncertainty sets by a single polynomial.

In the field of verified numerics, domain decomposition is a common tool used in interval arithmetic as well as Taylor Model integration.[10] This paper introduces the concept of automatic domain splitting to the non-verified DA framework and applies it to the problem of uncertainty propagation in order to overcome the previously described issues. In the course of the integration of the initial condition, the polynomial representing the current state is constantly monitored. When the nonlin-

earities of the system cause the high order terms of the polynomial to grow too large, integration is paused and the domain of the polynomial is split into two halves along one of the variables. This yields two new polynomials, one covering each half of the initial conditions. Since the split of the domain of one variable into half causes the n-th order terms of that variable to shrink by a factor of 2^n, this method efficiently reduces the size of the highest order terms. Integration is then resumed on each one of the two new polynomials until either further splits are required or the final state is reached. The final result of this procedure is a list of final state polynomials, each describing the evolution of some automatically determined subset of the initial condition.

The paper is organized as follows. In the next section, a brief introduction is given about DA techniques. The use of DA to compute high order expansion of the flow of ODEs is then described. A simple application to the propagation of uncertainties in Kepler's dynamics is presented to show the advantages of high order propagation with respect to linear methods and to illustrate its limits for large uncertainty sets and nonlinear dynamics. Automatic domain splitting is then introduced and the uncertainty propagation in Kepler's dynamics is resumed to show the advantages of domain splitting over standard high order propagation. Lastly, the performances of the resulting splitting DA-based integrator are assessed on the propagation of an end-of-life disposal trajectory of the INTEGRAL spacecraft.

NOTES ON DIFFERENTIAL ALGEBRA

DA techniques find their origin in the attempt to solve analytical problem by an algebraic approach.[11] Historically, the treatment of functions in numerics has been based on the treatment of numbers, and the classical numerical algorithms are based on the mere evaluation of functions at specific points. DA techniques are based on the observation that it is possible to extract more information on a function rather than its mere values. The basic idea is to bring the treatment of functions and the operations on them to the computer environment in a similar way as the treatment of real numbers. Referring to Figure 1, consider two real numbers a and b. Their transformation into the floating point representation, $\overline{a}$ and $\overline{b}$ respectively, is performed to operate on them in a computer environment. Then, given any operation $\times$ in the set of real numbers, an adjoint operation $\otimes$ is defined in the set of FP numbers such that the diagram in figure commutes*. Consequently, transforming the real numbers a and b in their FP representation and operating on them in the set of FP numbers returns the same result as carrying out the operation in the set of real numbers and then transforming the achieved result in its FP representation. In a similar way, suppose two sufficiently regular functions f and g are given. In the framework of differential algebra, the computer operates on them using their Taylor series expansions, F and G respectively. Therefore, the transformation of real numbers in their FP representation is now substituted by the extraction of the Taylor expansions of f and g. For each operation in the function space, an adjoint operation in the space of Taylor polynomials is defined such that the corresponding diagram commutes; i.e., extracting the Taylor expansions of f and g and operating on them in the function space returns the same result as operating on f and g in the original space and then extracting the Taylor expansion of the resulting function. The implementation of differential algebra in a computer allows the computation of the Taylor coefficients of a function up to a specified order n, along with the function evaluation, with a fixed amount of effort. The Taylor coefficients of order n for sums and product of functions, as well as scalar products with reals, can be computed from those of summands and factors; therefore, the set of equivalence classes of functions can be endowed with well-defined operations, leading to

*The diagram commutes approximately in practice, due to truncation errors.

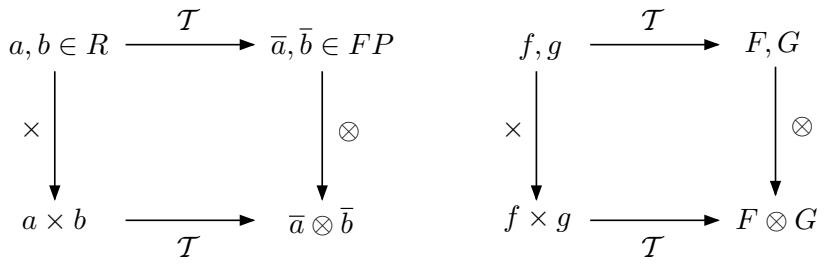

Figure 1. Analogy between the floating point representation of real numbers in a computer environment (left figure) and the introduction of the algebra of Taylor polynomials in the differential algebraic framework (right figure).

the so-called truncated power series algebra.[12,13]

Similarly to the algorithms for floating point arithmetic, the algorithm for functions followed, including methods to perform composition of functions, to invert them, to solve nonlinear systems explicitly, and to treat common elementary functions.[14] In addition to these algebraic operations, also the analytic operations of differentiation and integration are introduced, so finalizing the definition of the DA structure. The differential algebra sketched in this section was implemented in the software COSY INFINITY.[15]

HIGH ORDER EXPANSION OF THE FLOW OF ODES

Differential algebra allows the derivatives of any function f of n variables to be computed up to an arbitrary order k, along with the function evaluation. This has an important consequence when the numerical integration of an ODE is performed by means of an arbitrary integration scheme. Any integration scheme is based on algebraic operations, involving the evaluation of the ODE right hand side at several integration points. Therefore, carrying out all the evaluations in the DA framework allows differential algebra to compute the arbitrary order expansion of the flow of a general ODE with respect to the initial condition.

Without loss of generality, consider the scalar initial value problem

$$\begin{cases} \dot{x} = f(x,t) \\ x(t_0) = x_0 \end{cases} \tag{1}$$

and the associated phase flow $\varphi(t; x_0)$. We now want to show that, starting from the DA representation of the initial condition x_0, differential algebra allows us to propagate the Taylor expansion of the flow in x_0 forward in time, up to the final time t_f.

To this aim, replace the point initial condition x_0 by the DA representative of its identity function up to order k, which is a $(k+1)$-tuple of Taylor coefficients. (Note that x_0 is the flow evaluated at the initial time; i.e, $x_0 = \varphi(t_0; x_0)$.) As for the identity function only the first two coefficients, corresponding to the constant part and the first derivative respectively, are non zeros, we can write $[x_0]$ as $x_0 + \delta x_0$, where x_0 is the reference point for the expansion. If all the operations of the numerical integration scheme are carried out in the DA framework, the phase flow $\varphi(t; x_0)$ is approximated, at each fixed time step t_i, as a Taylor expansion in x_0.

For the sake of clarity, consider the forward Euler's scheme

$$x_i = x_{i-1} + f(x_{i-1})\Delta t \tag{2}$$

and substitute the initial value with the DA identity $[x_0] = x_0 + \delta x_0$. At the first time step we have

$$[x_1] = [x_0] + f([x_0]) \cdot \Delta t. \tag{3}$$

If the function f is evaluated in the DA framework, the output of the first step, $[x_1]$, is the k-th order Taylor expansion of the flow $\varphi(t; x_0)$ in x_0 for $t = t_1$. Note that, as a result of the DA evaluation of $f([x_0])$, the $(k+1)$-tuple $[x_1]$ may include several non zero coefficients corresponding to high order terms in δx_0. The previous procedure can be inferred through the subsequent steps. The result of the final step, $[x_f]$, is the k-th order Taylor expansion of $\varphi(t; x_0)$ in x_0 at the final time t_f. Thus, the flow of a dynamical system can be approximated, at each time step t_i, as a k-th order Taylor expansion in x_0 in a fixed amount of effort.

The conversion of standard integration schemes to their DA counterparts is straightforward both for explicit and implicit solvers. This is essentially based on the substitution of the operations on real numbers with those on DA numbers. In addition, whenever the integration scheme involves iterations (e.g. iterations required in implicit and predictor-corrector methods), step size control, and order selection, a measure of the accuracy of the Taylor expansion of the flow needs to be included. For the numerical integrations presented in this paper, a DA version of the Dormand-Prince (8-th order solution for propagation, 7-th order solution for step size control) implementation of the Runge-Kutta scheme is used.

The main advantage of the DA-based approach is that there is no need to write and integrate variational equations in order to obtain high order expansions of the flow. This result is basically obtained by the substitution of operations between real numbers with those on DA numbers, and therefore the method is ODE independent. Furthermore, the efficient implementation of the differential algebra in COSY INFINITY allows us to obtain high order expansions with limited computational time.

Test case: Uncertainty Propagation in Kepler's Dynamics

The previous DA-based numerical integrators pave the way to numerous practical applications. An example is given hereafter pertaining the propagation of errors on initial conditions. The Taylor polynomials resulting from the use of DA-based numerical integrators expand the solution of the initial value problem of Eq. (1) with respect to the initial condition. Thus, at each step i, the dependence of the solution on the initial condition is available in terms of an high order polynomial map $[x_i] = \mathcal{M}_{x_i}(\delta x_0)$, where δx_0 is the displacement of the initial condition from its reference value x_0. Suppose now the reference value x_0 represents a nominal initial condition for a dynamical system and assume an error δx_0 occurs between the actual initial condition and the nominal one. The mere evaluation of the Taylor polynomial $\mathcal{M}_{x_i}(\delta x_0)$ supplies the new solution x_i at time t_i corresponding to the displaced initial condition. More precisely, the Taylor polynomial $\mathcal{M}_{x_i}(\delta x_0)$ delivers a Taylor approximation of the new solution x_i, whose accuracy depends on the expansion order n and the size of the displacement δx_0. The main advantage of the DA-based integrator is that the new solution is obtained by means of the evaluation of a polynomial, so avoiding a new numerical integration corresponding to the displaced initial condition. Moreover, the same Taylor polynomial can be used to compute the solution corresponding to any error δx_0. Consequently, if many values of δx_0 are to be processed, the polynomial evaluations efficiently replace the multiple necessary numerical integrations.

The results of the application of the previous procedure are illustrated in the following example. The dynamics of an object moving in the solar system is integrated in the framework of the two

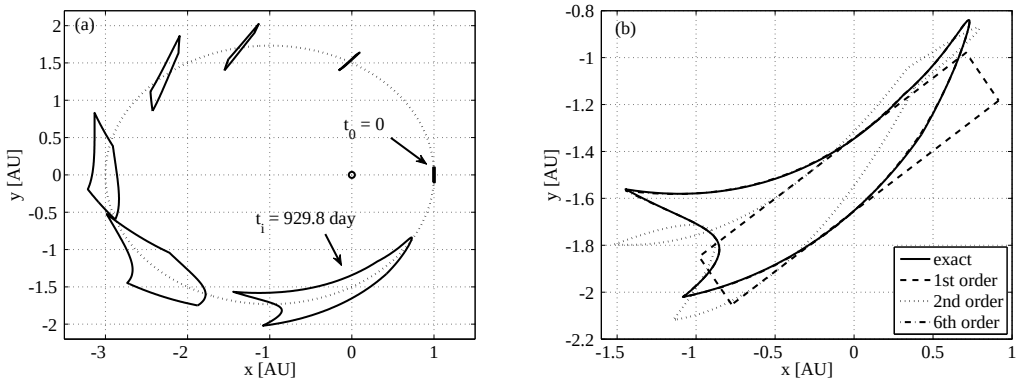

Figure 2. (a) Propagation of a box of initial positions in the two-body problem using the 6th order Taylor expansion of the flow of the associated ODE; (b) accuracy analysis on one of the resulting boxes ($t_i = 929.8$ day).

body problem:

$$\begin{cases} \dot{\boldsymbol{r}} &= \boldsymbol{v} \\ \dot{\boldsymbol{v}} &= -\frac{\mu}{r^3}\,\boldsymbol{r}, \end{cases} \tag{4}$$

where $\boldsymbol{r}$ and $\boldsymbol{v}$ are the object position and velocity vectors respectively, and μ is the Sun gravitational parameter. The nominal initial conditions are set such that the object starts moving from the pericenter of an elliptic orbit, lying on the ecliptic plane (see the dotted line in Figure 2a). The pericenter radius is 1 AU, whereas the magnitude of the initial velocity is selected to have a resulting orbit of eccentricity 0.5. The DA-based integrator is used to compute a 6th order expansion of the ODEs flow along the orbit.

The x and y components of the initial position are then supposed to lie in an uncertainty box of size 0.008 AU and 0.08 AU in the x and y direction respectively. This is a rather unrealistic uncertainty set, which has been exaggerated for illustrative purposes. The evolution of the resulting initial box is now investigated by propagating its boundary. More specifically, a uniform sampling of the boundary is performed. Then, for each sample, the displacement with respect to the nominal initial conditions is computed and the 6th order polynomial maps obtained with the DA-based integrator are evaluated. In this way, for each integration time, the evolved box can be plotted by means of mere polynomial evaluations. The evolved box is reported in Figure 2a corresponding to 9 integration times uniformly distributed over the orbital period. The time required by COSY-Infinity for the computation of the 6th order map is about 0.15 s on a 2.4 GHz Intel Core i5 MacBook Pro running Mac OS X 10.9.1.

The accuracy of the Taylor expansion of the flow is better highlighted in Figure 2b. Focusing on the integration time $t_i = 929.8$ day, the figure reports the box obtained with a multiple point-wise integration of the samples (solid line). The propagated boxes obtained by the evaluation of the polynomial maps representing the flow of the ODE in Eq. (4) are then plotted for comparison, corresponding to different expansion orders. The figure shows that a 6th order expansion of the flow is necessary to achieve a visually accurate representation of the exact box.

Unfortunately, the accuracy of the 6th order Taylor expansion drastically decreases for longer integration times. Figure 3a focuses on the integration time $t_i = 2025.9$ day, which corresponds to only about 1.96 revolutions. The figure compares the box obtained by a multiple point-wise

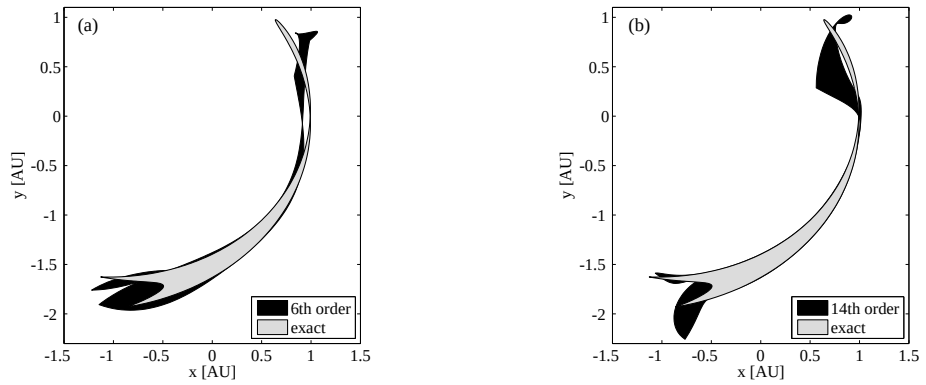

Figure 3. Uncertainty set at $t_i = 2025.9$ day: 6th order (a) and 14th order (b) Taylor expansions.

integration of the samples with that resulting from the evaluation of the 6th order polynomial map. The 6th order expansion is not able to accurately describe the exact box. Even increasing the order of the Taylor expansion does not improve the accuracy. This is confirmed in Figure 3b, where the results of a 14th order expansion of the ODEs flow are compared with the exact box.

Figures 3a and 3b demonstrate that a single Taylor expansion of reasonable order is not always able to catch the typical nonlinearities of orbital mechanics problems. Consequently, while performing better than classical linearized methods, the high order integrator described above may fail to accurately track uncertainties depending on the nonlinearity of the dynamics, the size of the uncertainty set to be propagated, and the propagation time.

Automatic domain splitting can play a crucial role to solve the previously described issues. In the course of the integration of the initial conditions, the uncertainty set is split along its variables when the nonlinearities of the system cause the Taylor expansion to lose accuracy. This yields the final set to be described with new polynomials, each one covering a subset of the initial conditions. The technique and its advantages with respect to the current implementation of the DA-based integrator are described in the next section.

AUTOMATIC DOMAIN SPLITTING

The approximation error between an $n+1$ times differentiable function $f \in C^{n+1}$ and its Taylor expansion P_f of order n, without loss of generality taken around the origin, is given by Taylor's theorem:[16]

$$|f(\delta x) - P_f(\delta x)| \leqslant C \cdot \delta x^{n+1}, \tag{5}$$

for some constant $C > 0$. We remark in passing that Taylor's theorem does not require f to be analytic, it is sufficient that $f \in C^{n+1}$.

Consider now the maximum error e_r of P_f on a domain B_r of radius $r > 0$ around the expansion point. By equation 5 we have that

$$|f(\delta x) - P_f(\delta x)| \leqslant C \cdot \delta x^{n+1} \leqslant C \cdot r^{n+1} = e_r.$$

If the domain of P_f is reduced from B_r to a ball $B_{r/2}$ of radius $r/2$, the maximum error of P_f

over $B_{r/2}$ will decrease by a factor of $1/2^{n+1}$, i.e.

$$|f(\delta x) - P_f(\delta x)| \leqslant C \cdot \delta x^{n+1} \leqslant C \cdot \left(\frac{r}{2}\right)^{n+1} = \frac{e_r}{2^{n+1}}.$$

We observe that for sufficiently large expansion orders, such as e.g. $n = 9$, the effect of reducing the size of the domain by half is thus greatly amplified and the maximum error is reduced by a factor of $\frac{1}{2^{10}} \approx 10^{-4}$. One solution to the previously described problem of non-convergence of the polynomial expansion over its initial domain is therefore to subdivide the initial box into smaller boxes and compute the Taylor expansion around the center point of each of the new boxes. Then the error of the new polynomial expansions in each sub box is greatly reduced, while taken in their entirety, the expansions still cover the entire initial set.

This process is often referred to as a divide and conquer strategy, and is very common in the field of numerical analysis.[17] However, the method described before suffers from an important drawback. By manually subdividing the initial box into smaller subsets of a predefined size, it is necessary to know a priori the required size of the subdivided boxes to obtain the desired error. If the initial boxes are chosen too small, precious computational time is wasted computing expansions over several small boxes where one large box would have sufficed.

Furthermore, for practical reasons such subdivisions are typically performed in a uniform manner, producing a uniform grid of boxes. This adds to the computational cost as often times the dynamical behavior of the function f being expanded differs significantly over the various parts of its domain. In some regions larger boxes will yield the required accuracy, while other regions may be more critical and require a more finely spaced box cover.

Lastly, in the case of the expansion of the flow of an ODE the a priori splitting of the initial box into sub boxes causes computational yet additional unnecessary overhead due to the fact that the flow at the beginning of the integration ($t = t_0$) is just the identity, i.e. $\Phi_{t_0}(x) = x$.

At this initial time, the entire flow over the initial condition box can be accurately represented by the identity polynomial. As the integration of the dynamics progresses, the flow Φ is distorted away from the identity until such a time t^* at which the polynomial approximation P_Φ surpasses some pre-specified maximum error. Up until that time t^*, however the flow is described well by just one polynomial expansion over the entire initial condition, there is no need to perform the integration between t_0 and t^* twice using two separate expansions.

Automatic Domain Splitting, originally introduced by Berz and Makino in the verified Taylor Model integrator COSY VI,[18] builds on this observation to circumvent these problems by automatically detecting at which time the flow expansion over a given box of initial conditions is becoming too inaccurate. Once this case has been detected, the domain of the original polynomial expansion is divided into two domains of half their original size each, with two separate polynomial expansions on each. Each of the new polynomials represents the exact same function over its domain as the original polynomial, but expanded around the center point of the new domain.

More specifically let $P(x)$ be the polynomial representation of the flow $\Phi_{t^*}(x)$ at some time t^*, with the domain $x \in [-1, 1]$. Then the split of P into P_1 and P_2 is defined as

$$\begin{aligned} P_1(x) &= P\left(\frac{1}{2} \cdot x - \frac{1}{2}\right) \\ P_2(x) &= P\left(\frac{1}{2} \cdot x + \frac{1}{2}\right), \end{aligned}$$

with the domains of P_1 and P_2 again being taken as $x \in [-1, 1]$.

From this definition it is evident that P_1 is covering the left half ($[-1, 0]$) of the domain of P and P_2 covering the right half ($[0, 1]$). Since both P_1 and P_2 are again polynomials of the same degree as P, this splitting operation can be performed exactly in DA arithmetic without adding any truncation errors. The new polynomials P_1 and P_2 represent exactly the same graph as that of P, just expanded around a different expansion point. However, in accordance with equation 5, the terms of any order n in P_1 and P_2 will be smaller by a factor of 2^n than the corresponding terms in P.

After such a split occurs, the integration can continue on each one of P_1 and P_2 in the same manner as described in the previous section until further splits are required or the final integration time is reached. The result is a list of polynomial expansions, each covering a specific part of the domain of initial conditions.

The decision when exactly a polynomial needs to be split, and in the case of multivariate polynomials the direction of the split, is in general difficult to answer. We use a heuristic method which estimates the size of the $n + 1$ order of the polynomial based on an exponential fit of the size of the known coefficients up to order n. If the size of this truncated order becomes too large, we decide to split the polynomial.

In the case of multivariate polynomials $P(x_1, x_2, \ldots, x_v)$, the split is only performed in one variable x_j. The splitting direction x_j is determined by factoring the known coefficients of order up to n with respect to each x_i, i.e. writing

$$P(x_1, x_2, \ldots, x_v) = \sum_{k=0}^{n} x_i^k \cdot q_{i,k}(x_1, x_2, \ldots, x_v)$$

where the polynomials $q_{i,k}$ do not depend on x_i. Then the size of the coefficients in $q_{i,n+1}$ are estimated again by an exponential fit, and the direction x_i with the largest $q_{i,n+1}$ is chosen as the splitting direction.

In this way, all splits are performed in the direction of the variable that has the largest estimated contribution to the total truncation error of the polynomial P, and thus the splits have the maximal impact on reducing the approximation error.

Test case: Uncertainty Propagation in Kepler's Dynamics

To demonstrate the domain splitting technique described in the previous section, we apply it to the problem of propagating Kepler's dynamics as in the previous section. The computations are performed at order 14 with the same initial condition box as in the previous example. The splitting precision is set to $\varepsilon = 3 \cdot 10^{-4}$, meaning that when the estimated truncation error of an expansion exceeds this limit a split is triggered. We remark that for actual applications this limit is typically chosen much lower, it was chosen this high to allow for a better visualization of the splitting process.

Integrating the dynamics from time $t_0 = 0$ to time $t_f = 2906.6$ day (2.81 revolutions), the entire computation takes about 22 seconds on the same machine used for the example in the previous section, and produces 23 final polynomial expansions covering the initial condition.

Figure 4 shows the resulting boxes at various times during the integration. Up until time $t_a = 930.1$ day (0.90 revolutions), the entire set is well described by a single DA expansion. At time $t_b = 988.3$ day (0.96 revolutions), just before completing the first revolution 2 splits have occurred,

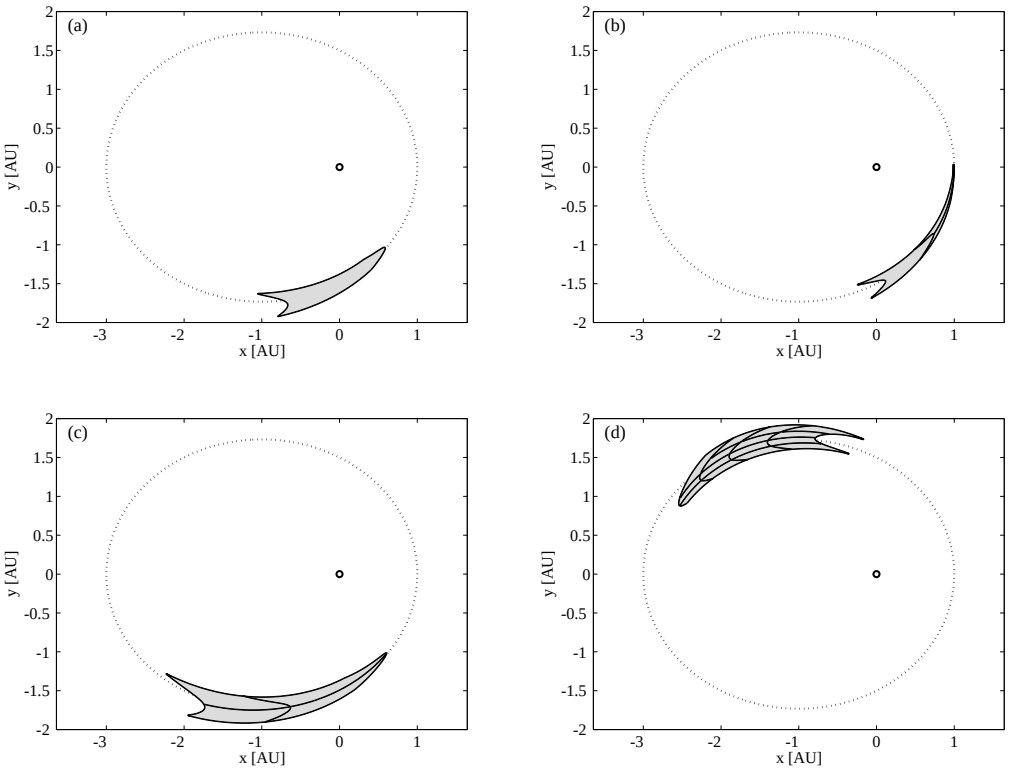

Figure 4. Propagation of the initial uncertainty set in the two-body dynamics using 14th order Taylor expansions of the flow and automatic domain splitting: (a) $t_a = 930.1$ **day (0.90 revolutions); (b)** $t_b = 988.3$ **day (0.96 revolutions); (c)** $t_c = 1918.4$ **day (1.86 revolutions); (d)** $t_d = 2325.3$ **day (2.25 revolutions)**

leading to three polynomial patches. Another split is performed at time $t_c = 1918.4$ day (1.86 revolutions). Figure 4d shows the 15 DA patches that are necessary to accurately track the uncertainty set at time $t_d = 2325.3$ day (2.25 revolutions). Then, the number of patches increases to 23 at the final integration time.

For the sake of completeness, similarly to the previous section, Figure 5 focuses on the integration time $t_i = 2025.9$ day. More specifically, Figure 5a illustrates how automatic domain splitting subdivides the initial domain in 9 boxes during the 14th order integration. On the other hand, Figure 5b reports all 9 propagated boxes at $t_i = 2025.9$ day and the box resulting from a single 14th order Taylor expansion of the flow on the entire uncertainty set. The comparison between Figure 5b and Figure 3b shows that automatic domain splitting allows the exact propagated set to be accurately described with the 9 Taylor polynomials.

END-OF-LIFE DISPOSAL OF INTEGRAL

The improvements that automated domain splitting brings into the DA-based integration are now investigated in the long-term propagation of an end-of-life disposal trajectory of the INTEGRAL spacecraft. The awareness of the risk of uncontrolled accumulation of man-made objects in orbit around the Earth became significant only in the late 70s. In 1978 Donald J. Kessler tackled for the first time the problem of the collision between orbiting objects.[19] He observed that, since many ar-

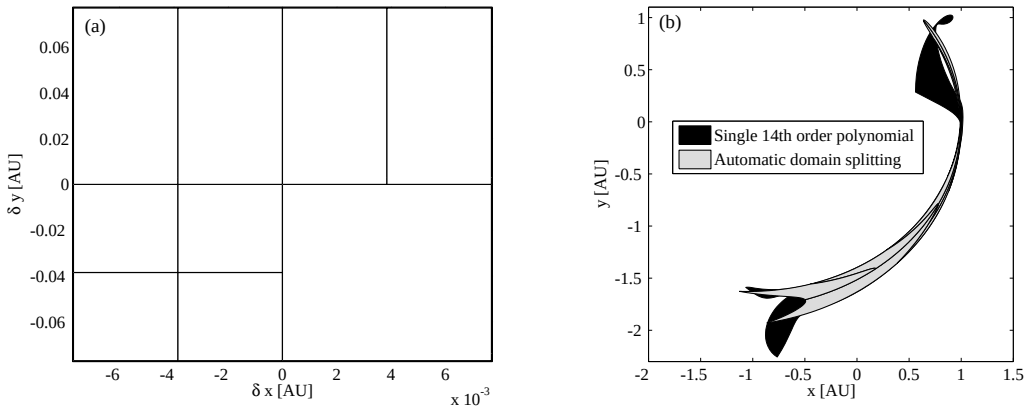

Figure 5. Propagation of the initial uncertainty set in the two-body dynamics to $t = 2025.9$ day with automatic domain splitting: (a) final subdivision of the initial domain in 9 subdomains; (b) comparison between a single 14th order Taylor expansion of the flow and a 14th order Taylor expansion including automatic domain splitting.

tificial satellites are in orbits that cross one another, there is a finite probability of collision between them. Using a statistical method developed to simulate the fragmentation within the asteroid belt, he discovered that, following a possible collision cascade process, there were the conditions necessary for the formation of a debris belt around the Earth. The term "Kessler Syndrome" was introduced to describe this phenomenon. From that year on, National Governments and Space Organizations began to keep track of all orbiting objects, including space debris and rocket bodies, while previously mainly satellites were catalogued.

In parallel to the effort towards accurate debris tracking, there is international consensus that space activities must be managed to minimize debris generation and risk. This consensus is embodied in space debris mitigation guidelines published by various organizations such as the International Telecommunication Union (ITU), the Inter-Agency Space Debris Coordination Committee (IADC), and the United Nations (UN). The general aim of space debris mitigation is to reduce the growth of space debris by ensuring that space systems are designed, operated, and disposed of in a manner that prevents them from generating debris throughout their orbital lifetime.

Within this context, this paper investigates the end-of-life disposal of the INTEGRAL spacecraft. INTEGRAL is a space observatory to simultaneously observe objects in gamma ray, X-ray, and visible light. Its principal targets are violent explosions known as gamma-ray bursts, powerful phenomena such as supernova explosions, and regions in the Universe thought to contain black holes. The spacecraft was successfully launched by a Proton rocket on October 17, 2002 from the Baikonur Cosmodrome in the Republic of Kazakhstan. The INTEGRAL operational orbit is a 72-hour orbit with an initial perigee altitude of 9000 km and the apogee altitude of 153600 km.[20] A mid-term review of the mission is expected to occur in 2014 in order to postpone the mission end to December 16, 2016. However, for the following analysis, we assume the nominal mission end occurs on December 1, 2014 and we are looking for a disposal maneuver in 2015. At this epoch, a single impulsive ΔV of 60 m/s is applied to the spacecraft to inject it into a disposal trajectory to reenter Earth's atmosphere.

The motion of INTEGRAL is modeled using Kepler's equation in Gauss' form including

- luni-solar gravitational perturbations, where Moon and Sun's ephemerides are computed using a simplified analytical model obtained by fitting JPL DE405 ephemerides with third order polynomials;
- non-spherical harmonics, including J_2 and J_3 effects;
- atmospheric drag, where a simple exponential model is used to compute the atmosphere density.

The spacecraft is supposed to have an area-to-mass ratio of 0.01 m^2/kg and a drag coefficient of 2.2. The Gauss' equations are modified to have the true anomaly as integration variable, moving time into the set of parameters defining the spacecraft state as $(a, e, i, \Omega, \omega, t)$.

The design of the optimal disposal maneuver for INTEGRAL has already been addressed in the work by Di Mauro et al.[21] Based on this work, a disposal maneuver has been designed for the epoch $t_0 = 5755.074$ MJD2000 (October 4, 2015). The resulting nominal spacecraft orbital parameters after the maneuver are

$$\begin{cases} a_0 &= 87550.4340 \text{ km} \\ e_0 &= 0.812471 \\ i_0 &= 0.953952 \text{ rad} \\ \Omega_0 &= 3.825817 \text{ rad} \\ \omega_0 &= 4.480676 \text{ rad} \\ f_0 &= 2.846481 \text{ rad.} \end{cases} \tag{6}$$

The spacecraft state in Eq. (6) is used as reference initial condition for the disposal phase. We assume the reentry condition occurs when the altitude of the spacecraft is less than 120 km. Figure 6 illustrates the evolution of the reference INTEGRAL disposal trajectory in the above dynamical model. More specifically, for the sake of clarity, the figure reports the profile of the perigee altitude of the osculating orbit. The reference disposal trajectory turns out to reenter Earth's atmosphere on March 9, 2028.

The reference initial velocity is now supposed to be affected by a given uncertainty deriving from possible errors of the disposal maneuver. More specifically, an uncertainty of ± 1.5 m/s is assigned to each component of the velocity vector in the local orbital reference frame at t_0, where $\boldsymbol{i}_{LOF}$ is the direction parallel to the velocity, $\boldsymbol{k}_{LOF}$ identifies the out-of-plane direction parallel to the instantaneous direction of the orbital angular momentum, and $\boldsymbol{j}_{LOF}$ is the in-plane direction that completes the reference frame. The DA integrator is used to propagate the uncertain initial conditions with domain splitting enabled. The robustness of the disposal maneuver is then investigated by verifying the occurrence of INTEGRAL reentry over the entire uncertainty domain.

A numerical integration is performed using 8th order expansions and the integrator settings have been tuned to split the initial domain so to meet the requirement of tracking the uncertainties with an accuracy on the order of about 6 m in the semimajor axis, 10^{-6} in the eccentricity, 10^{-5} rad in i, Ω, ω, and 14 min in time. In order to limit the number of generated subdomains and polynomials, domain splitting is disabled on any box whose volume is less than 2^{-15} times that of the initial domain. This means no more than 15 splits are performed on any box, limiting a priori the maximum number of possible final boxes generated to $2^{15} = 32768$. This limit was introduced to reduce the computational time for this exploratory work, so to enable a relatively quick parametric analysis

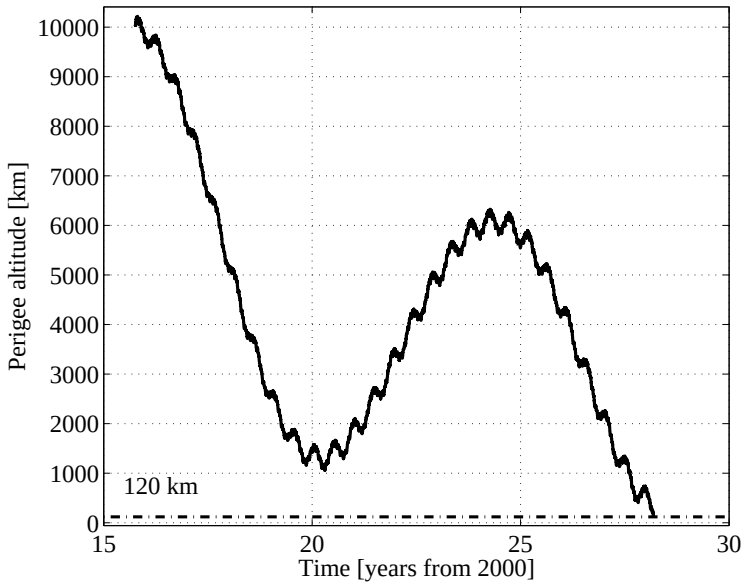

Figure 6. Perigee altitude of the INTEGRAL osculating orbit for the reference initial conditions of Eq. (6)

of the integrator performances. Furthermore, as will be shown later, it is not necessary to propagate boxes below a certain threshold as the dynamic behavior exhibited over these boxes is more non-linear than is desirable from a mission design point of view. Nonetheless, the limit could be relaxed or removed in future work, especially considering that automatic domain splitting can take significant advantage of parallelization. Thus, the computational time can be drastically reduced.

In addition to the size limitation, the integration of any box is stopped when the altitude of the center of the box is lower than 120 km, as it is assumed to have entered the atmosphere at this point. The check on only the center point is a small simplification, as this does not automatically imply that all points in the box have reentered. However, since the automatic splitting algorithm ensures the size of each box to be such that the polynomial expansion converges well, in practice that means the deviation from the center point is small enough that a reentry event of the center point also implies a reentry event of the entire box.

The integrator is able to propagate the uncertainty using only one 8th order Taylor polynomial until January 15, 2020 (7319.665 MJD2000). Then, the nonlinearities over the now relatively large uncertainty set prevent the integrator from meeting the accuracy requirements set for the semi major axis a. The largest contribution to the polynomial truncation error of a is determined to be the uncertainty along the velocity direction $\boldsymbol{i}_{LOF}$. Thus, a first split of the initial uncertainty domain occurs along the velocity direction $\boldsymbol{i}_{LOF}$, which causes the initial domain to split in two boxes. The resulting two polynomials for each state variable $(a, e, i, \Omega, \omega, t)$ then continue to be propagated forward in true anomaly. Two additional splits occur on December 12, 2024 (9112.996 MJD2000) and January 7, 2025 (9138.644 MJD2000), respectively. Figure 7 spotlights the second and third splits on the altitude profile of the osculating perigee. The lines represent the trajectories followed by the center points of each box. Note that since time itself is one of the phase space variables, the

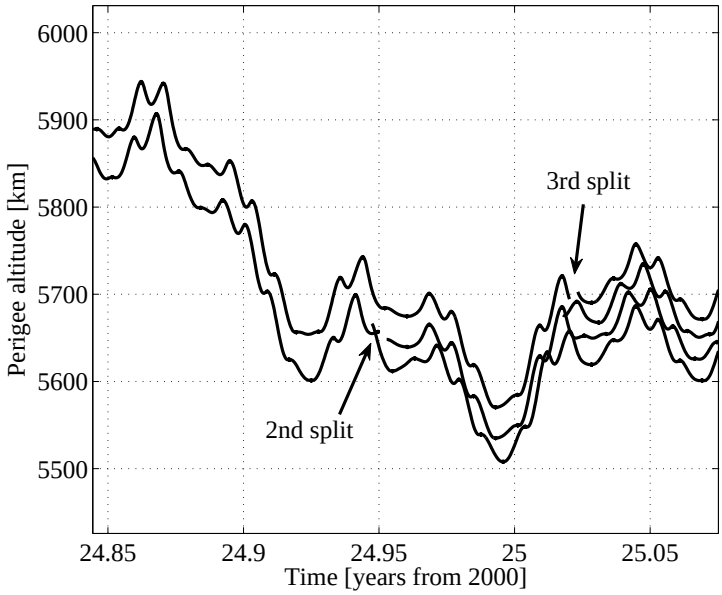

Figure 7. Perigee altitude of INTEGRAL osculating orbit resulting from the propagation of the initial uncertainty set: detail on the second and third split. Lines represent the trajectories followed by the center points of each box.

two center points of the split boxes do not continue the integration from the same time. They do, of course, restart from the same true anomaly, which is the independent variable in our dynamical system.

The number of splits drastically increases towards the end of the integration, when the spacecraft reaches lower altitudes and the atmospheric drag starts to manifest itself. The perturbation induced by the drag greatly intensifies the nonlinearities and causes a significant number of splits compared to the previous part of the integration. This is clearly illustrated in Figure 8, which zooms on the last part of the disposal phase. The final number of generated boxes is 7224, whereas the associated computational time is 15.4 hr on a 2.4 GHz Intel Core i5 iMac running Mac OS X 10.9.2.

The propagation of most generated boxes stops because the altitude threshold of 120 km is reached and the box is declared to reenter. However, some of the generated boxes get close but do not reach the altitude threshold. Thus, their propagation and interaction with Earth's atmosphere continues, which causes additional splits to occur until the minimum box size is reached. At this stage, even if the reentry conditions are not achieved, the resulting boxes can no longer be propagated without failing to meet the accuracy requirements. Thus, their propagation is stopped. The fact that the atmospheric drag is introducing most of the nonlinearity is confirmed in Figure 9, where the same uncertain initial conditions are propagated by removing the drag from the dynamical model. As can be seen, the number of splits drastically decreases and the number of final boxes is limited to four. In addition, all boxes reenter Earth's atmosphere, which shows the importance of including the atmospheric drag to improve the fidelity of the dynamical model.

During the integration, most splits occur in the velocity direction, $\boldsymbol{i}_{LOF}$, and in the in-plane direction $\boldsymbol{j}_{LOF}$. Consequently, the section of the three-dimensional initial uncertainty set in the

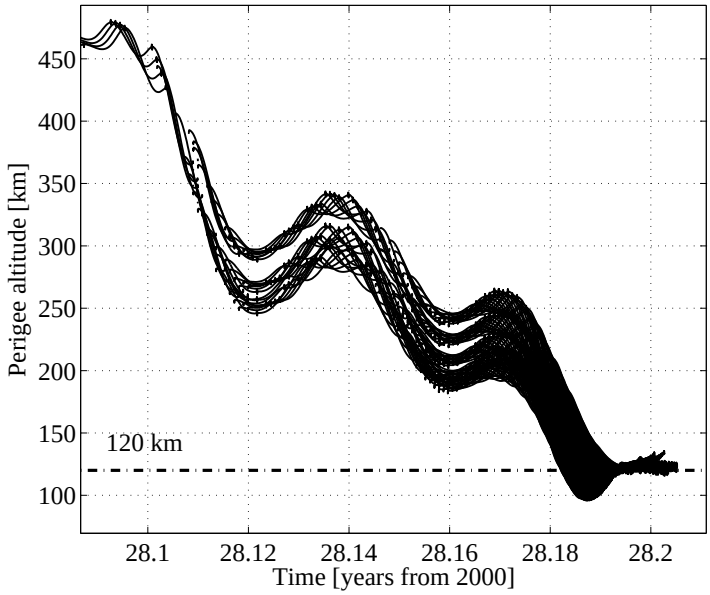

Figure 8. Perigee altitude of INTEGRAL osculating orbit resulting from the propagation of the initial uncertainty set: detail on the last part of the disposal phase.

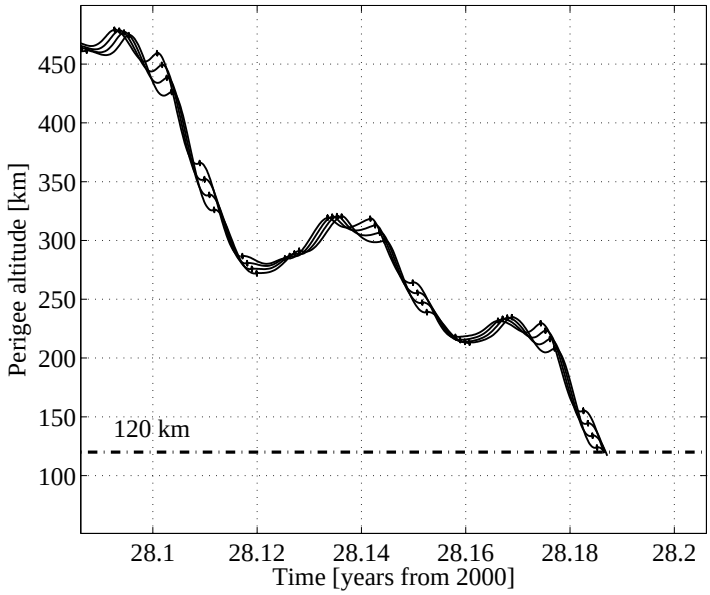

Figure 9. Perigee altitude of the INTEGRAL osculating orbit resulting from the propagation of the initial uncertainty set by excluding atmospheric drag from the dynamical model.

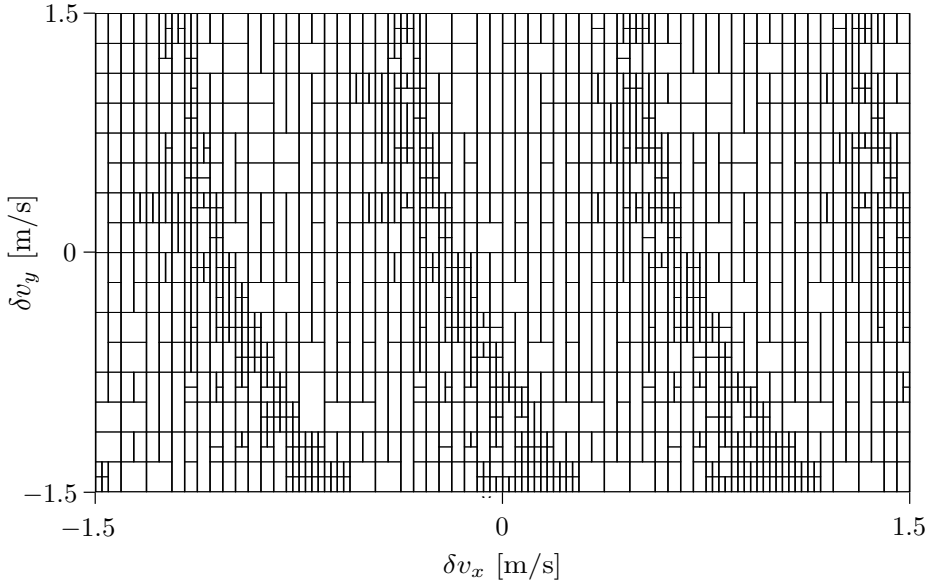

Figure 10. Final boxes on the initial domain in the v_x, v_y-plane in the local orbital reference frame.

v_x, v_y-plane of the local orbital reference frame corresponding to the reference out-of-plane initial velocity (i.e., $\delta v_z = 0$ m/s) is investigated in Figure 10. More specifically, the figure reports all the subdomains generated by the automatic domain splitting during the DA-based integration. Regions of larger final boxes can be easily distinguished from areas where most splits occur. This figure represents a precious source of information. As explained in the previous sections, the splits occur when the nonlinearities are too strong to be managed with a single Taylor polynomial. Consequently, the areas in Figure 10 where most splits concentrate coincide with regions of strong nonlinearity. These regions include boxes of late reentry, which are then integrated for longer time in a dynamical environment including the highly nonlinear effect of atmospheric drag.

This observation is confirmed by Figure 11, which superimpose a color map to Figure 10. More specifically, the color map of Figure 11 illustrates the flight path angle corresponding to the center point of each box at its final integration time. As can be seen, the regions of larger boxes of Figure 10 match the red areas in Figure 11. This means that larger boxes tend to show higher flight path angles at their final time, which allow them to steeply and robustly reach the altitude threshold earlier than the other boxes. The white areas correspond to boxes whose altitude is always above 120 km during the entire integration. Consequently, they correspond to boxes that does not reach the altitude threshold and spend more time in the atmosphere above 120 km. This is confirmed by the fact that they match the region of smaller boxes in Figure 10. The blue areas are regions of transition between the previous conditions. More specifically, the initial conditions lying in the blue areas eventually reenter Earth's atmosphere. However, their flight path angle tend to be small when they reach the atmosphere. Thus, they are close to skip reentry even for relatively small deviations from the nominal initial state.

For the sake of completeness, Figures 12 and 13 report the same color map of Figure 11 on the sections of the three-dimensional initial uncertainty set corresponding to $\delta v_z = -1.5$ m/s and $\delta v_z = 1.5$ m/s, respectively. As can be seen, the white boxes near the top in Figure 11 turn blue

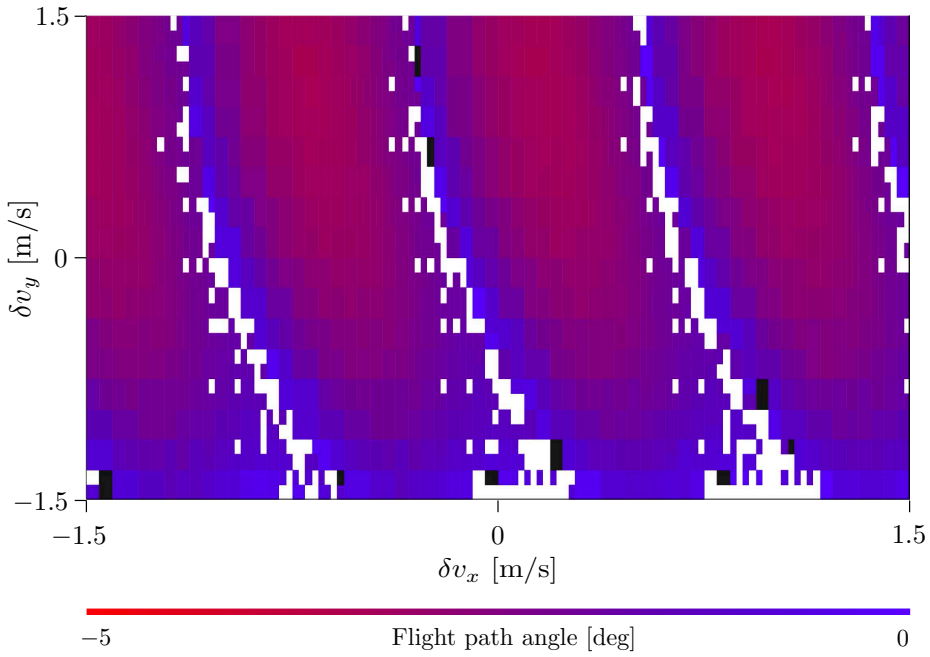

Figure 11. Color map superimposed over the final boxes on the initial domain in the v_x, v_y-plane for $\delta v_z = 0\ m/s$: final flight path angle.

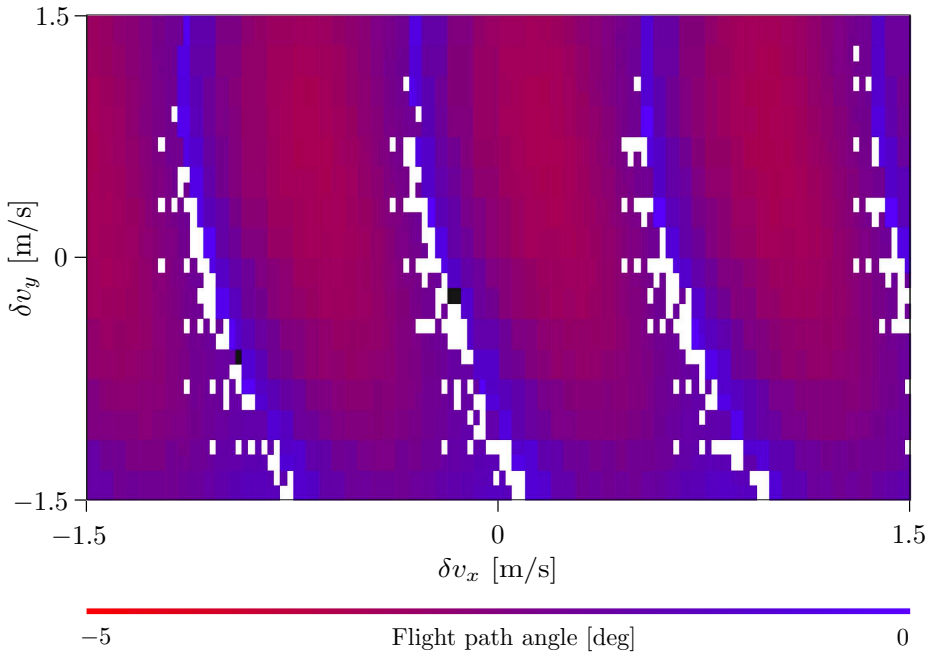

Figure 12. Color map superimposed over the final boxes on the initial domain in the v_x, v_y-plane for $\delta v_z = -1.5\ m/s$: final flight path angle.

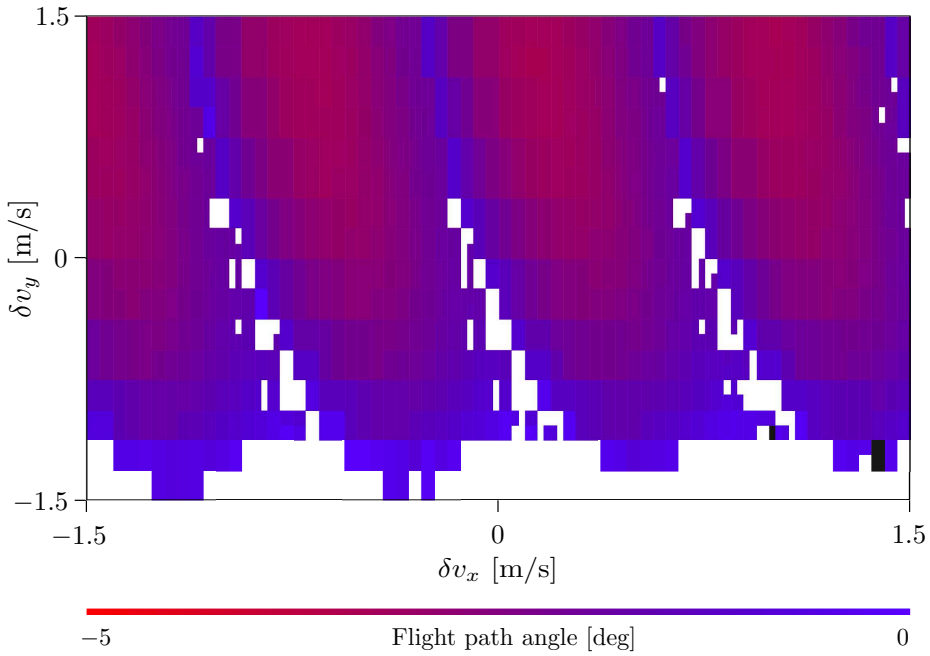

Figure 13. Color map superimposed over the final boxes on the initial domain in the v_x, v_y-plane for $\delta v_z = 1.5\ m/s$: final flight path angle.

in Figure 13. Thus, the corresponding initial conditions now reenter Earth's atmosphere at the end of the integration. However, the associated flight path angle is small, which confirms that the blue boxes correspond to the continuous transition from steep reentry to skip reentry also in v_z; i.e., in the third dimension of the initial domain. However, those initial conditions that barely reenter the atmosphere exhibit in their further orbital evolution an undesirable near-circularization of the orbit along with highly chaotic motion before slowly spiraling towards Earth. For a disposal maneuver such as the one sought after in this paper those orbits are not suitable and must be avoided. This is because circularization leads to longer interactions with the atmosphere and hence large uncertainties in the reentry point.

These considerations are confirmed in Figures 14a and 14b. Sample initial conditions are taken from the three main areas of Figure 11 as illustrated in Figure 14a. Sample 4 corresponds to the nominal initial conditions of Eq. (6). The resulting perigee altitude profiles are plotted in Figure 14b. Even if taken from different red areas, samples 1 to 3 propagate close one to each other and safely reenter Earth's atmosphere. The remaining samples tend to reenter later and show a higher sensitivity to initial errors and a tendency to form circularized orbits that remain around 120 km altitude for several weeks.

One last observation deserves to be mentioned. The previous analysis shows that the dynamics is highly sensitive to uncertainties on the initial velocity. In fact, besides the large red areas of safe and robust reentry, the ± 1.5 m/s initial uncertainty domain includes risky initial conditions that may skip reentry even in case of relatively small errors on the initial conditions. In particular, as illustrated in Figure 14, the reference initial condition of Eq. (6) lie in the region of transition between a safe reentry and a skip reentry. Thus, the maneuver as designed here is not robust with respect to the possible errors associated with the disposal maneuver. If robustness is a requirement, the disposal

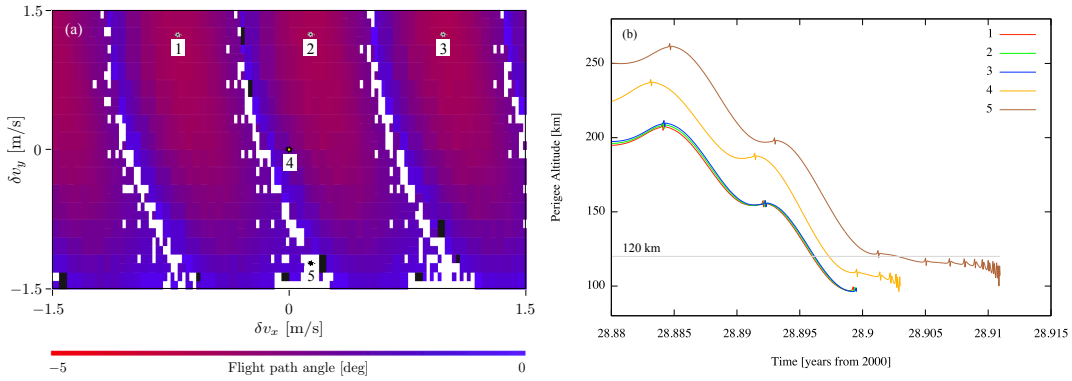

Figure 14. Sample initial conditions from the three areas of Figure 11 (a) and associated perigee altitude profiles (b).

maneuver shall be designed to have a post-maneuver spacecraft state lying close to the center of the red areas in Figure 11. Even considering such a favorable situation, the admissible uncertainty to ensure a robust reentry is of the order of ± 0.4 m/s, which highlights the stringent requirements that must be imposed to design end-of-life disposal maneuvers for spacecraft in highly elliptical orbits.

CONCLUSION

This paper introduced automatic domain splitting into the high order DA-based integration to accurately propagate large sets of uncertainties in highly nonlinear dynamics and long term integrations. The resulting integrator splits the initial uncertainty domain in subdomains along the integration when the polynomials representing the current state do not meet the accuracy requirements. The final result is a list of final state polynomials, each describing the evolution of some automatically determined subset of the initial condition. Thus, altogether, the Taylor polynomials accurately map the entire initial domain in the final set. Consequently, given any generated sample in the initial domain, its integration to the final time can be replaced with the fast evaluation of the polynomial map that propagates the subdomain it belongs to. Monte Carlo simulations can then be enhanced by replacing thousands of integrations with evaluations of the Taylor expansions of the flow. Furthermore, the dynamical behavior of the system in different regions of phase space can be understood directly from the resulting splitting pictures showing the final box sizes.

The performances of the integrator have been assessed on the end-of-life disposal of the INTEGRAL spacecraft. A single impulse is applied to the spacecraft at the end of its mission. The maneuver is designed to put INTEGRAL in a disposal trajectory to ballistically reenter Earth's atmosphere. In order to limit the number of generated subdomains and polynomials, a minimum box size was fixed. The relation between the size of the final boxes and the robustness of the reentry condition has been identified. For the case of INTEGRAL, the admissible error on the post-maneuver velocity turns out to be of the order of ± 0.4 m/s, which highlights the stringent requirements that must be posed to design end-of-life disposal maneuvers for spacecraft in highly elliptical orbits. Future work will be devoted to verify that less stringent requirements can be admitted if additional corrective impulses are allowed along the disposal trajectory.

ACKNOWLEDGMENTS

A. Wittig gratefully acknowledges the support received by the EU Marie Curie fellowship PITN-GA 2011-289240 (AstroNet-II).

REFERENCES

[1] H. R. Battin, *An introduction to the Mathematics and Methods of Astrodynamics*. Reston, VA: AIAA Education Series, 1999.

[2] O. Montenbruck and G. Eberhard, *Satellite orbits*. New York: Springer-Verlag, 2nd ed., 2001.

[3] J. L. Crassidis and J. L. Junkins, *Optimal estimation of dynamic systems*. Boca Raton, FL: CRC press, 2004.

[4] P. S. Maybeck, *Stochastic Models, Estimation and Control*. New York: Academic press, 1982.

[5] R. S. Park and D. J. Scheeres, "Nonlinear Mapping of Gaussian Statistics: theory and applications to spacecraft trajectory design," *Journal of guidance, Control, and Dynamics*, Vol. 29, No. 6, 2006, pp. 1367–1375.

[6] M. Berz, *Modern Map Methods in Particle Beam Physics*. New York: Academic press, 1999.

[7] M. Valli, R. Armellin, P. Di Lizia, and M. Lavagna, "Nonlinear Mapping of Uncertainties in Celestial Mechanics," *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 1, 2012, pp. 48–63.

[8] R. Armellin, P. Di Lizia, F. Bernelli-Zazzera, and M. Berz, "Asteroid close encounters characterization using differential algebra: the case of Apophis," *Celestial Mechanics and Dynamical Astronomy*, Vol. 107, No. 4, 2010, pp. 451–470.

[9] A. Morselli, R. Armellin, P. Di Lizia, and F. Bernelli-Zazzera, "Computing Collision Probability Using Differential Algebra and Advanced Monte Carlo Methods," *Proceedings of the 63rd International Astronautical Congress (IAC), Napoli, Italy, 1-5 October*, 2010, pp. 2311–2324.

[10] M. Berz and K. Makino, *Lecture Notes MSUHEP-080609*. East Lansing, MI: Michigan State University, 2008.

[11] M. Berz, *Modern Map Methods in Particle Beam Physics*. New York: Academic Press, 1999.

[12] M. Berz, *The new method of TPSA algebra for the description of beam dynamics to high orders*. Los Alamos National Laboratory, 1986. Technical Report AT-6:ATN-86-16.

[13] M. Berz, "The method of power series tracking for the mathematical description of beam dynamics," *Nucl. Instrum. Methods*, Vol. A258, No. 3, 1987, pp. 431–436.

[14] M. Berz, *Differential Algebraic Techniques, Entry in Handbook of Accelerator Physics and Engineering*. New York: World Scientific, 1999.

[15] M. Berz and K. Makino, *COSY INFINITY Version 9 Reference Manual*. East Lansing, MI: Michigan State University, 2006.

[16] R. Walter, *Principles of Mathematical Analysis, Third Edition*. New York: McGraw-Hill, Inc., 1976.

[17] R. E. Moore, *Interval analysis*, Vol. 2. Englewood Cliffs, NJ: Prentice-Hall, 1966.

[18] K. Makino and M. Berz, "Suppression of the wrapping effect by Taylor model-based verified integrators: Long-term stabilization by preconditioning," *International Journal of Differential Equations and Applications*, Vol. 10, No. 4, 2006, pp. 353–384.

[19] D. J. Kessler and B. G. Cour-Palais, "Collision frequency of artificial satellites: The creation of a debris belt," *Journal of Geophysical Research: Space Physics (1978–2012)*, Vol. 83, No. A6, 1978, pp. 2637–2646.

[20] P. Jensen, K. Clausen, C. Cassi, F. Ravera, G. Janin, C. Winkler, and R. Much, "The INTEGRAL spacecraft-in-orbit performance," *Astronomy and Astrophysics-Les Ulis*, Vol. 411, No. 1, 2003, p. L7.

[21] G. Di Mauro, M. Rasotto, R. Armellin, and P. Di Lizia, *End-of-Life Disposal Concepts for Lagrange-Points and HEO Missions*. ESA Technical report, 2013.