CrossMark

ORIGINAL ARTICLE

Propagation of large uncertainty sets in orbital dynamics by automatic domain splitting

Alexander Wittig[1] · Pierluigi Di Lizia[1] · Roberto Armellin[2] · Kyoko Makino[3] · Franco Bernelli-Zazzera[1] · Martin Berz[3]

Abstract 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 when the non-linearities of the dynamics prohibit good convergence of the Taylor expansion in one or more directions. We solve this issue by introducing automatic domain splitting. During propagation, the polynomial expansion of the current state is split into two polynomials whenever its truncation error reaches a predefined threshold. The resulting set of polynomials accurately tracks uncertainties, even

✉ Alexander Wittig
  alexander.wittig@polimi.it

  Pierluigi Di Lizia
  pierluigi.dilizia@polimi.it

  Roberto Armellin
  roberto.armellin@soton.ac.uk

  Kyoko Makino
  makino@msu.edu

  Franco Bernelli-Zazzera
  franco.bernelli@polimi.it

  Martin Berz
  berz@msu.edu

[1] Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milan, Italy

[2] Aeronautics, Astronautics and Computational Engineering Unit, University of Southampton, Highfield campus, Southampton SO17 1BJ, UK

[3] Department of Physics and Astronomy, Michigan State University, 567 Wilson Road, East Lansing, MI 48824, USA

in highly nonlinear dynamics. The method is tested on the propagation of (99942) Apophis post-encounter motion.

1 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. Typically these uncertainties need to be propagated forward in time, for example for spacecraft navigation or to estimate the collision risk between artificial satellites or the threat from near Earth objects. As orbital dynamics is highly nonlinear the size of the uncertainty set tends to quickly increase along the trajectory. 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 in nonlinear systems is extremely difficult. Present-day approaches mainly refer to linearized propagation models (Battin 1999; Montenbruck and Eberhard 2001; Crassidis and Junkins 2004) or full nonlinear Monte Carlo simulations (Maybeck 1982). 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. Thus, the main challenge in uncertainty propagation is to develop methods that can accurately map uncertainties in nonlinear system with limited computational effort.

The unscented transformation (UT) was proposed to address the deficiencies of linearization: a limited number of samples are deterministically chosen to match the mean and covariance of a (non necessarily Gaussian-distributed) probability distribution (Julier and Uhlmann 2004; Julier 2002). The main drawback of the approach is that it delivers a second order approximation of the first two moments of the mapped statistical distributions. This can turn out to be inadequate for applications when higher order moments are needed to accurately describe the propagated uncertainty set. Gaussian mixtures (Giza et al. 2009; Horwood and Poore 2011; Jah and Kelecy 2009) have been introduced as a possible way to overcome this limitation. The method is based on the fact that any probability density function can be represented by a weighted sum of Gaussian distributions. As each of these distributions is defined on a reduced domain, a linearized approach or UT can be employed to individually map each of them without a significant loss in accuracy.

An alternative way to map the statistics is based on the approximation of the flow of the dynamics in Taylor series and the use of the resulting polynomials as dynamical substitutes. This approach was presented by Park and Scheeres (2006), and later investigated by other researchers (Majji et al. 2008; Vittaldev et al. 2012). It was shown that a good agreement with Monte Carlo simulations can be achieved, however the derivation of the dynamics of the high order tensor can be complex from the computational standpoint, especially for high fidelity dynamics.

These difficulties were addressed by the authors in past works, in which a technique called differential algebra (DA) was used to automatically expand the flow of the dynamics up to an arbitrary order (Di Lizia et al. 2008; Armellin et al. 2010). Nonlinear uncertainty propagation can take advantage of the high order expansions of the flow. 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. The most established DA computation tool is COSY INFINITY by Berz and Makino (Berz and Makino 2006) which is used in this work. Other implementations of DA techniques have recently been implemented in academia, e.g. Jet Transport (Alessi et al. 2009), or in industry such as the DA code JACK developed by Thales Alenia Space (Bignon et al. 2014).

Differential algebra has already proven its efficiency in the nonlinear propagation of uncertainties within different dynamical models, including two-body dynamics (Valli et al. 2012), $(n + 1)$- body dynamics (Armellin et al. 2010), and geocentric models (including Earth's gravitational harmonics, solar radiation pressure, shadows, and third body perturbations) (Morselli et al. 2010). Nonetheless, the accuracy of the method tends to decrease drastically when the uncertainty domain becomes too stretched in one or more directions. This can be due to one or a combination of the following causes: high nonlinearity of the dynamics, large initial uncertainty sets, and long term propagations.

The propagation of asteroids motion after a close encounter with a major body is a typical example. As reported by Valsecchi et al. (2006), the asteroid (99942) Apophis will have an extremely close approach to the Earth on 13 April 2029. The asteroid orbit will suffer a very large perturbation, opening the door to the possibility of a resonant return in 2036. The nonlinearities of the close encounter and of the post-encounter motion will make any uncertainty in the direction anti-parallel to the Earth heliocentric motion in 2029 diverge by a factor of 40000 in 2036. While performing better than classical linearized methods, the current implementation of the DA uncertainty propagator still is inaccurate and impractical in such cases, due to the prohibitively high order required to describe the resulting uncertainty sets by a single polynomial. Furthermore, in the case of impacts for some part of the initial conditions, it is mathematically impossible to represent the entire resulting uncertainty set by a single polynomial expansion (Di Lizia et al. 2009; Alessi et al. 2009).

To overcome these problems, this work introduces a novel method, referred to as automatic domain splitting, into the DA uncertainty propagator. It handles those situations in which a single Taylor expansion of the flow is not enough to accurately map the entire initial uncertainty set. The underlying idea is to split the initial domain into manageable subdomains over which the Taylor expansion shows good convergence properties. This "divide-and-conquer"style approach is common in the field of verified numerics. The DA algorithm developed in this work extends the traditional method in various ways, most importantly by performing the subdivision of the initial domain adaptively as needed during the integration instead of a priori. This yields not only a large gain in computational performance as the initial set is automatically subdivided into optimally sized subsets, but also reveals valuable information about the dynamical behavior in different regions of the initial conditions.

In past works, the authors have already studied the same asteroid (99942) Apophis (Di Lizia et al. 2009) using the self-validating Taylor Model integrator COSY VI (Makino and Berz 2009). Employing its own conceptually similar splitting technique, COSY VI was able to manage the propagation only until shortly after the first close encounter. This is because

the purpose of COSY VI is the rigorous verified integration of initial conditions, bringing along with it different, more stringent requirements.

In this work, this problem is overcome by introducing a dedicated automatic procedure to identify the necessity of splitting during the integration, determine the direction in which to split, and to manage the resulting sets. During the integration of the initial condition, the polynomial representing the current state is constantly monitored. When the nonlinearities of the system cause the estimated truncation error of the polynomial to grow too large, integration is paused and the domain of the polynomial is split into two halves along the variable with the largest contribution to the truncation error. This yields two new polynomials, one covering each half of the initial condition set. Since the split of the domain of one variable into half causes the nth 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 application of the resulting tool to the propagation of asteroid (99942) Apophis motion is addressed in the second part of this paper. Automatic domain splitting is shown to overcome the previously described issues of the DA uncertainty propagator and to pave the way to a innovative approach to study the challenging problem of resonant returns. As an additional contribution, the paper proposes the use of the resulting splitting structure at the end of the propagation as a novel method to infer the dynamical behavior of the system over the initial uncertainty domain.

The paper is organized as follows. In the next section, a brief introduction is given about DA techniques and the high order expansion of the flow of an ODE is 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 asteroid (99942) Apophis.

2 Differential algebra and high order flow expansion

DA supplies the tools to compute the derivatives of functions within a computer environment (Berz 1999a). More specifically, by substituting the classical implementation of real algebra with the implementation of a new algebra of Taylor polynomials, any function f of v variables is expanded into its Taylor polynomial up to an arbitrary order n with limited computational effort. In addition to basic algebraic operations, operations for differentiation and integration can be easily introduced in the algebra, thus finalizing the definition of the differential algebra structure of DA (Berz 1986, 1987). Similarly to algorithms for floating point arithmetic, also in DA various algorithms were introduced, including methods to perform composition of functions, to invert them, to solve nonlinear systems explicitly, and to treat common elementary functions (Berz 1999b). The differential algebra used for the computations in this work was implemented in the software COSY INFINITY (Berz and Makino 2006).

An important application of DA is the automatic high order expansion of the solution of an ODE in terms of the initial conditions (Berz 1999a; Armellin et al. 2010). This can be achieved by replacing the operations in a classical numerical integration scheme, including evaluation of the right hand side, by the corresponding DA operations. This way, starting from the DA representation of an initial condition x_0, DA ODE integration allows the propagation of the Taylor expansion of the flow in x_0 forward in time, up to any final time t_f. Any explicit ODE integration scheme can be rewritten as a DA integration scheme in a straightforward way. For the numerical integrations presented in this paper, a DA version of a 7/8 Dormand–Prince (8th order solution for propagation, 7th order solution for step size control) 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. It is therefore independent of the particular right hand side of the ODE and the method is quite efficient in terms of computational cost.

2.1 Kepler's dynamics example

To illustrate the method, we consider the dynamics of a celestial body moving in the framework of the two-body problem

$$\begin{cases} \dot{\boldsymbol{r}} = \boldsymbol{v} \\ \dot{\boldsymbol{v}} = -\frac{\mu}{r^3}\,\boldsymbol{r}, \end{cases} \tag{1}$$

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 orbit with eccentricity of 0.5, lying on the ecliptic plane (see the dotted line in Fig. 1a). The units are normalized in such a way that both the pericenter radius and μ are equal to 1, thus leading to the following initial conditions:

$$\begin{cases} x = 1 \\ y = 0 \\ \dot{x} = 0 \\ \dot{y} = \sqrt{1.5} \end{cases} \tag{2}$$

In these units the nominal orbital period is $2\pi\sqrt{8}$.
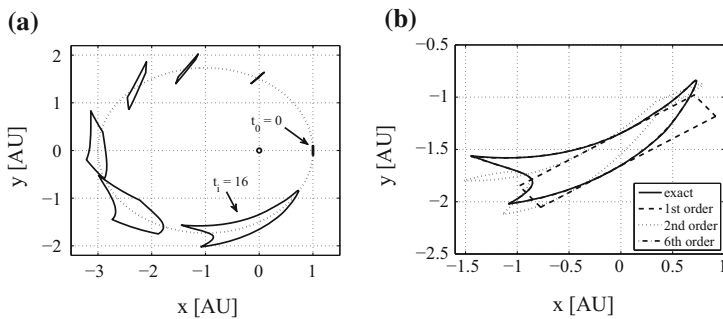
Fig. 1 Propagation of a set of initial positions in the two-body problem using the 6th order Taylor expansion of the flow of the associated ODE. **a** Evolution of the initial set over an orbital period. **b** Accuracy analysis on the set at $t_i = 16$ (after 0.9 nominal orbital revolutions)

Fig. 2 Uncertainty set at $t_i = 34.85$ (after 1.96 nominal revolutions). **a** 6th order Taylor expansion. **b** 14th order Taylor expansion
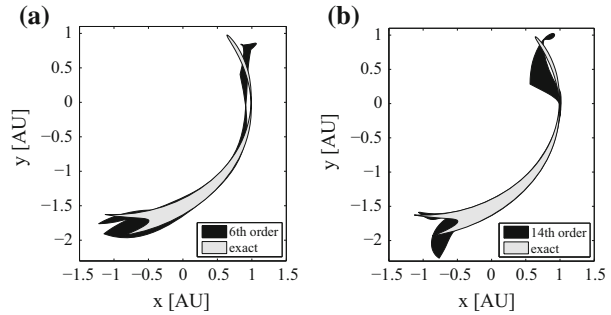
The DA-based integrator is used to compute a 6th order expansion of the ODEs flow along the orbit. The x and y components of the initial position are then supposed to lie in an uncertainty box of size 0.008 and 0.08 in the x and y direction respectively. As can be seen in Fig. 1, this is a rather unrealistic 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 set can be plotted by means of mere polynomial evaluations. The evolved set is reported in Fig. 1a corresponding to nine 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 Fig. 1b. Focusing on the integration time $t_i = 16$, the figure reports the set obtained with a multiple pointwise integration of the samples (solid line). The propagated sets obtained by the evaluation of the polynomial maps representing the flow of the ODE in Eq. (1) 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 result (i.e. a relative error of the order of 10^{-4}) for the large initial uncertainty set considered in this example. Note that depending on the actual application, stricter requirements may be appropriate.

Unfortunately, the accuracy of the 6th order Taylor expansion drastically decreases for longer integration times. Figure 2a focuses on the integration time $t_i = 34.85$, which corresponds to only about 1.96 revolutions of the nominal initial condition. The figure compares the set obtained by a multiple pointwise 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 set. Even increasing the order of the Taylor expansion does not improve the accuracy. This is confirmed in Fig. 2b, where the results of a 14th order expansion of the ODEs flow are compared with the exact set.

Figure 2a, b demonstrate the aforementioned fact that a single Taylor expansion of reasonable order is not always able to accurately describe the evolution of an initial uncertainty set. 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 a list of final polynomials, each one expanded around a different point in the initial condition box and 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.

3 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 (Rudin 1976):

$$\left| f(\delta x) - P_f(\delta x) \right| \leqslant C \cdot \delta x^{n+1} \tag{3}$$

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 Eq. 3 we have that

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

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.

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

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 domain into smaller domains and compute the Taylor expansion around the center point of each of the new domains. Then the error of the new polynomial expansions in each sub domain 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 (Moore et al. 2009). However, the method described before suffers from an important drawback. To manually subdivide the initial domain into smaller subsets of a predefined size, it is necessary to know a priori the required size of the subdivided domains to obtain the desired error. If the initial domains are chosen too small, precious computational time is wasted computing expansions over several small domains where one large domain would have sufficed.

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

Lastly, in the case of the expansion of the flow of an ODE, the a priori splitting of the initial domain into sub domains 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. $\varphi(t_0; x) = x$. At this initial time, the entire flow over the initial condition 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 step $t = t_{i+1}$ at which the polynomial approximation P_φ surpasses some pre-specified maximum error ε. Up until the previous time step t_i, 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_i several times using a fine cover of polynomial expansions.

Building on these observations, *Automatic Domain Splitting* employs an automatic algorithm to determine at which time the flow expansion over a given set of initial conditions does not describe the dynamics with sufficiently high accuracy any more. Once this case has been detected, the domain of the original polynomial expansion is divided along one of the expansion variables into two domains of half their original size. By re-expanding the polynomials around the new center points, two separate polynomial expansions are obtained. As the re-expansion of the polynomials does not change their order, each of the new polynomials is identical to the original polynomial on its respective domain. This process is illustrated in Fig. 3.

More specifically, let $P(x)$ be the polynomial representation of the flow $\varphi(t_i; x)$ at some time t_i, where $x \in [-1, 1]^k$. The choice of the domain to be $[-1, 1]$ for each component of x is an arbitrary one, yet this particular choice simplifies many of the following calculations and has various numerical advantages as well. Without loss of generality, the domain $[-1, 1]$ of any component x_i can be transformed linearly into any other desired interval $[a, b]$.

The split of P into P_1 and P_2 along component x_j is defined as

$$P_1(x) = P\left(x_1, \ldots, x_{j-1}, \frac{1}{2} \cdot x_j - \frac{1}{2}, x_{j+1}, \ldots, x_k\right)$$

$$P_2(x) = P\left(x_1, \ldots, x_{j-1}, \frac{1}{2} \cdot x_j + \frac{1}{2}, x_{j+1}, \ldots, x_k\right)$$

and the domains of P_1 and P_2 are again assumed to be $x \in [-1, 1]^k$.

From this definition it is evident that P_1 is covering the left half ($x_j \in [-1, 0]$) of the original domain of x_j and P_2 covering the right half ($x_j \in [0, 1]$). 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 have exactly the same graph as that of P, just expanded around a different expansion point. However, in accordance with Eq. 3, the terms of order n in x_i present in P_1 and P_2 will be smaller by a factor of 2^n than the corresponding terms in P.

Fig. 3 Illustration of the propagation process with domain splitting
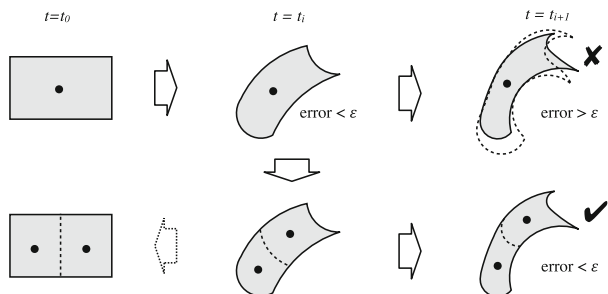
After such a split occurs, the integration can continue on each one of P_1 and P_2 in the same manner as described in the previous section until further splits are required or the final integration time is reached. The result is a list of polynomial expansions, each covering a specific part of the domain 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 method that estimates the size of the $n + 1$ order terms of the polynomial based on an exponential fit of the size of all the known non-zero coefficients up to order n. If the size of this truncated order becomes too large, we decide to split the polynomial.

This method allows us to take into account all the information available to us in the polynomial expansion in order to obtain an accurate estimate of the size of the $n + 1$ order. Compared to more trivial splitting criteria, such as only considering the size of the terms of order n, this method improves in particular the estimates for sparse functions such as functions with symmetries like sin and cos where many terms vanish. The exponential fit is chosen because after reducing the domain by a sufficient number of splits, the coefficients of the resulting polynomial expansion will in fact decay exponentially independently of the function being expanded. This is a direct consequence of Taylor's Theorem.

To be more precise, given a polynomial P of order n of the form

$$P(\pmb{x}) = \sum_\alpha a_\alpha \pmb{x}^\alpha$$

written using multi-index notation, the size S_i of the terms of order i is computed as the sum of the absolute values of all coefficients of exact order i:

$$S_i = \sum_{|\alpha|=i} |a_\alpha|.$$

We denote by I the set of indices i for which S_i is non-zero. A least squares fit of the exponential function

$$f(i) = A \cdot \exp(B \cdot i)$$

is used to determine the coefficients A, B such that $f(i) = S_i$, $i \in I$ is approximated optimally in the least squares sense. Then the value of $f(n + 1)$ is used to estimate the size S_{n+1} of the truncated order $n + 1$ of P. This method is illustrated in Fig. 4, where the polynomial under consideration is the Taylor expansion of $\sqrt{1 + x/2}$ up to 9th order. The size S_i of each order used for the least squares fit is shown as hatched bars, while the resulting fitted function f is shown as a line. As can be seen from this simple example, the size of the 10th order term, which was not used in the fitting, is approximated reasonably well by this method.

As described above, in the case of multivariate polynomials $P(\pmb{x}) = P(x_1, x_2, \ldots, x_k)$, the split is only performed in one component x_i. The determination of the splitting direction x_i is once again a non-trivial problem. We use a similar method to determine this direction i as we used previously for the decision to split. For each $j = 1, \ldots, n$ we begin by factoring the known coefficients of P of order up to n with respect to x_j, i.e. we write

$$P(x_1, x_2, \ldots, x_k) = \sum_{m=0}^n x_j^m \cdot q_{j,m}(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_k)$$

Fig. 4 Illustration of the estimation of the truncation error of the Taylor expansion of $\sqrt{1+x/2}$ via exponential fitting. Terms of order up to 9 are used for the fitting (*hatched*), while the 10th order term (*white*) is shown for reference
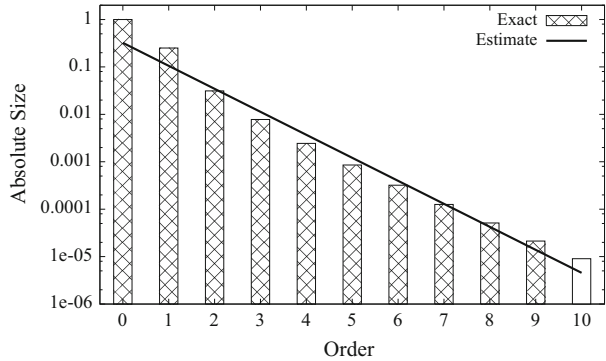
where the polynomials $q_{j,m}$ do not depend on x_j. Then the size $S_{j,m}$ of the polynomials $q_{j,m}$ is estimated by the sum of the absolute values of their coefficients and the same exponential fitting routine as described above is applied to obtain an estimate of the size $S_{j,n+1}$ of the truncated terms of order $n+1$ in x_j. Finally the splitting direction i is chosen to be the direction corresponding to the component x_j with largest truncation error $S_{j,n+1}$.

In this way, all splits are performed in the direction of the variable that currently 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. The splitting process described here in general, and the selection of the splitting direction in particular, are strongly dependent on the parametrization of the initial condition. The direction of maximum expansion in general is not aligned with a single direction of the parametrization, in which case several variables will contribute to the truncation error. In this case, splits occur automatically along all variables involved. However, the initial condition can often be parametrized such that expansion happens mainly along only a few or even just one of the directions.

3.1 Kepler's dynamics example: domain splitting illustration

Before we present a full analysis of the effect of the splitting precision on the accuracy, efficiency and number of final sets in the next subsection, we first demonstrate the domain splitting technique described in the previous section. We apply it to the same problem of propagating Kepler's dynamics as presented in the Sect. 2.1. Computations are performed at order 14 with the same initial condition box. The splitting precision is set to $\varepsilon = 3 \times 10^{-4}$, meaning that when the estimated truncation error of an expansion exceeds this limit a split is triggered. The limit was chosen this high to allow for a better visualization of the splitting process, in actual applications the limit is typically chosen much lower.

Integrating the dynamics from time $t_0 = 0$ to time $t_f = 50$ (2.81 nominal revolutions), the entire computation takes about 22 s on the same machine used for the example in the previous section, and produces 23 final polynomial expansions covering the initial condition.

Figure 5 shows the resulting sets at various times during the integration. Up until time $t_a = 16$ day (0.90 nominal revolutions), the entire set is well described by a single DA expansion. At time $t_b = 17$ (0.96 nominal revolutions), just before completing the first revolution 2 splits have occurred, leading to three polynomial patches. Another split is performed at time $t_c = 33$ (1.86 nominal revolutions). Figure 5d shows the 15 DA patches that are necessary

Fig. 5 Propagation of the initial uncertainty set in the two-body dynamics using 14th order Taylor expansions of the flow and automatic domain splitting **a** $t_a = 16$ (0.90 nominal revolutions). **b** $t_b = 17$ (0.96 nominal revolutions). **c** $t_c = 33$ (1.86 nominal revolutions). **d** $t_d = 40$ (2.25 nominal revolutions)
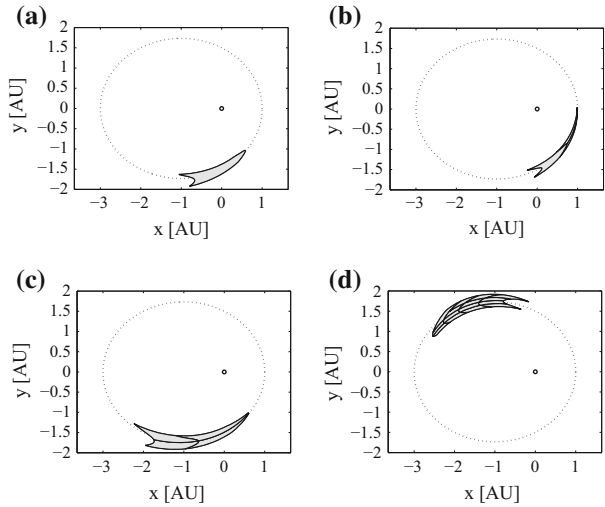

Fig. 6 Time versus the number of splits required by the automatic domain splitting algorithm. The times of perigee passages are highlighted with *grey bands*
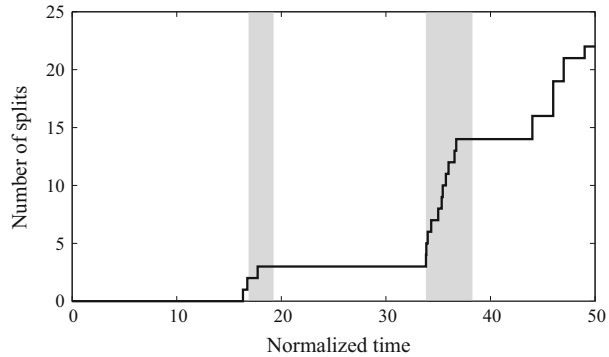
to accurately track the uncertainty set at time $t_d = 40$ (2.25 nominal revolutions). Then, the number of patches increases to 23 at the final integration time.

Figure 6 shows the number of splits as a function of the integration time. As is clear from this graph, the splits do not just happen at random times equally distributed over the integration period. Instead, splitting decisions are tightly correlated with the time of perigee passages (highlighted with grey band in Fig. 6). Clearly, splits cluster around the time of perigee passage, when the dynamics exhibit the strongest non-linearities and the propagated set is the most distorted. This is in accordance with the previous observations on the poor convergence of a single polynomial expansion close to the time or perigee passage in Fig. 2. It also indicates that the automatic splitting algorithm as introduced above is very efficiently determining the regions of phase space in which the dynamics require splitting due to high non-linearities without introducing unneeded splits. As the propagation continues, the time window in which splitting occurs become wider. This is in part due to the different initial conditions in the set having different periods, and the overall growth of the set over time. The splitting algorithm automatically takes these effects into account for each subset when deciding at which points along the trajectory splits occur.

Fig. 7 Propagation of the initial uncertainty set in the two-body dynamics to $t_i = 34.85$ with automatic domain splitting. **a** 14th order Taylor representation of the final set with automatic domain splitting. **b** Final subdivision of the initial domain in 9 subdomains
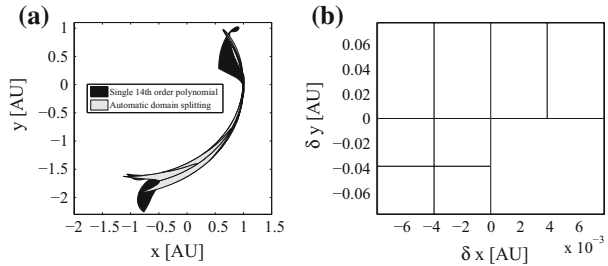
To further highlight the power of this method compared to the single polynomial expansion, Fig. 7 focuses on the same integration time $t_i = 34.85$ used in the example presented in the previous section. More specifically, the figure reports the 9 final resulting polynomials at $t_i = 34.85$ and the polynomial resulting from a single 14th order Taylor expansion of the flow on the entire uncertainty set. The comparison between Fig. 7a, b shows that automatic domain splitting allows the exact propagated set to be described to visual accuracy by the 9 Taylor polynomials. Figure 7b, on the other hand, illustrates how the automatic domain splitting subdivided the entire initial uncertainty set into 9 smaller sets of varying size during the 14th order integration. It is apparent that the dynamics in the bottom left quarter of the initial domain, starting closer to the primary and a little before perigee passage, requires the most splits.

While rather trivial in this example, *splitting pictures* such as the one shown in Fig. 7b in general can carry a significant amount of information about the dynamical behavior of the system under consideration in different phase space regions. In fact, the size of the boxes in a given region can give a strong indication of the non-linearity and hence the chaoticity of the motion there. Especially in higher dimensional cases, splitting pictures allow the identification of sensitivity to initial conditions by analyzing the number of splits in each direction of the final sets. Furthermore, different regions of dynamical regimes can be identified by analyzing the size of the resulting sets over the initial domain.

3.2 Kepler's dynamics: effect of splitting precision

Having illustrated the method using the previous example, we now proceed to add a quantitative analysis of the automatic domain splitting algorithm applied to Keplerian dynamics. In particular, we focus on the effect of the parameter to the splitting algorithm, the splitting precision ε, on the accuracy of the resulting polynomials, the computation time, and the number of splits required.

To this end, we consider the same problem described in Sect. 2.1, but propagated for 35 non-dimensional time units and with the following settings for the splitting precision ε:

$$\varepsilon \in \{10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}, 10^{-9}\}$$

These values are now chosen as they are more representative of a real world application. The integrations are performed using a Dormand–Prince 7/8 order Runge–Kutta integration scheme with automatic step size control set to a one step integration error of 10^{-12}.

In Table 1, for each such setting the computational time on a 2.9 GHz Intel Core i5 iMac with 8 GB DDR3 RAM is reported along with the final number of polynomials at the end of the integration. For each final set of the integration, the evaluation of the polynomial at the center, the corner points, and 100 randomly chosen points within the set is compared

ε	Comp. Time (s)	# of polynomials	Max. Error
10^{-3}	17.9	12	3.31×10^{-2}
10^{-4}	22.3	14	3.70×10^{-3}
10^{-5}	28.4	17	2.20×10^{-4}
10^{-6}	35.3	21	5.71×10^{-5}
10^{-7}	49.2	31	6.05×10^{-6}
10^{-8}	63.0	40	9.61×10^{-7}
10^{-9}	84.2	55	4.52×10^{-8}

Table 1 Dependence of computational time, maximum error of the Taylor expansion, and number of polynomials on the splitting precision ε

to pointwise integrations of each corresponding initial condition and the maximum error encountered is reported.

This analysis shows that the maximum error over the resulting set of polynomial expansions decreases linearly with the selected splitting precision as expected. However, the maximum error is always larger than the selected integration precision. This too is expected, as a close analysis of the description of the splitting algorithm reveals. The splitting precision plays a similar role as the one-step error set in the automatic step size control of the integration scheme. It is the maximum error that can accumulate at any time before the integrator takes action to reduce further error accumulation. However, the accumulated error at the time of the splitting cannot be undone as the splitting solely re-expands the (possibly already inaccurate) polynomial to prevent exponential error growth in future integration steps. The ideal tolerance for the splitting precision depends both on the dynamics, as more splits lead to more accumulation of approximation errors, and the integration time. It has to be chosen heuristically to ensure the final result satisfies the accuracy requirements of the application, exactly as the one-step tolerance of the automatic step size control of the integrator.

In the above analysis, we fixed the integrator precision to a low value of 10^{-12} in order to reduce the number of factors affecting the result and to isolate the true effect of the splitting precision setting on the result. In a practical application, the choice of the two tolerances is related, as the final error of the solution is affected by both values. In general it is advisable to choose the integration tolerance several orders of magnitude lower than the splitting tolerance. This is because in a typical integration the number of integrations steps, and hence the number of single-step errors, is much larger than the number of splits, and hence the approximation error. However, choosing the integration precision too low with respect to the splitting precision leads to unnecessary waste of computational time as the final accuracy in such a case will be almost solely due to the splitting error. In the above analysis, for the larger splitting precisions the integration tolerance could have been relaxed, reducing the time of the integration significantly without a noticeable loss of accuracy.

4 Long term propagation of (99942) Apophis

The improvements that automated domain splitting brings into the DA-based integration are now investigated in the practical application of long-term propagation of uncertainties during Apophis post-encounter motion. This test case lends itself to the automatic domain decomposition very well, as it involves a mix of highly non-linear dynamics during the close encounters with Earth, benign dynamics in the interim motion, and long overall integration time scales. It thus is a hard problem to treat for uncertainty propagation techniques.

The motion of Apophis in the Solar System is modeled including relativistic corrections to the well-known Newtonian forces. Specifically, the full equation of motion of our model is given by

$$\ddot{\boldsymbol{r}} = G \sum_i \frac{m_i(\boldsymbol{r}_i - \boldsymbol{r})}{r_i^3} \left\{ 1 - \frac{2(\beta + \gamma)}{c^2} G \sum_j \frac{m_j}{r_j} - \frac{2\beta - 1}{c^2} G \sum_{j \neq i} \frac{m_j}{r_{ij}} + \frac{\gamma |\dot{\boldsymbol{r}}|^2}{c^2} \right.$$

$$+ \frac{(1 + \gamma)|\dot{\boldsymbol{r}}_i|^2}{c^2} - \frac{2(1 + \gamma)}{c^2} \dot{\boldsymbol{r}} \cdot \dot{\boldsymbol{r}}_i - \frac{3}{2c^2} \left[\frac{(\boldsymbol{r} - \boldsymbol{r}_i) \cdot \dot{\boldsymbol{r}}_i}{r_i} \right]^2 + \left. \frac{1}{2c^2}(\boldsymbol{r}_i - \boldsymbol{r}) \cdot \ddot{\boldsymbol{r}}_i \right\}$$

$$+ G \sum_i \frac{m_i}{c^2 r_i} \left\{ \frac{3 + 4\gamma}{2} \ddot{\boldsymbol{r}}_i + \frac{\{[\boldsymbol{r} - \boldsymbol{r}_i] \cdot [(2 + 2\gamma)\dot{\boldsymbol{r}} - (1 + 2\gamma)\dot{\boldsymbol{r}}_i]\}(\dot{\boldsymbol{r}} - \dot{\boldsymbol{r}}_i)}{r_i^2} \right\} \quad (4)$$

where $\boldsymbol{r}$ is the position of Apophis in Solar System barycentric coordinates, G is the gravitational constant; m_i and $\boldsymbol{r}_i$ are the mass and the Solar System barycentric position of Solar System body i; $r_i = |\boldsymbol{r}_i - \boldsymbol{r}|$; c is the speed of light in vacuum; and β and γ are the parametrized post-Newtonian parameters measuring the nonlinearity in superposition of gravity and space curvature produced by unit rest mass (Seidelmann 1992).

In Eq. 4 it is assumed that the object we are integrating is affected by the gravitational attraction of n bodies, but has no gravitational effect on them; i.e., we are adopting the restricted $(n+1)$-body problem approximation. The positions, velocities, and accelerations of the n bodies are considered as given values, computed from the JPL DE405 ephemeris model. In our model, the n bodies include the Sun, the planets, the Moon, as well as the asteroids Ceres, Pallas, and Vesta. For planets with moons, with the exception of the Earth, the center of mass of the system is considered. The dynamical model is written in the J2000.0 ecliptic reference frame and is commonly referred to as the Standard Dynamical Model (Giorgini et al. 2008). This is the dynamical model used by NASA/JPL for the close encounters prediction in the frame of the Near Earth Object Program (http://neo.jpl.nasa.gov/index.html).

As illustrated in Fig. 8a, the asteroid Apophis will have an extremely close approach to the Earth on 13 April 2029 with a nominal closest distance of about 3.8×10^4 km. The asteroid orbit will then suffer a large perturbation on its orbital parameters, which will mainly affect its semimajor axis, inclination, and argument of the periapsis. No appreciable effect are expected on eccentricity and right ascension of the ascending node. The orbital period increases from 323.60 to 422.33 days (see Fig. 8b for a plot of Apophis trajectory before and after 2029 close encounter). This opens the door to the possibility of a resonant return to Earth in 2036.

The nominal initial state and the associated σ of Apophis at 3456 MJD2000 (June 18, 2009), expressed in equinoctial variables $\boldsymbol{p} = (a, P_1, P_2, Q_1, Q_2, l)$, are used as our test


Fig. 8 Apophis nominal motion before and after the close encounter on April 13th, 2029. **a** Distance from earth. **b** Trajectory
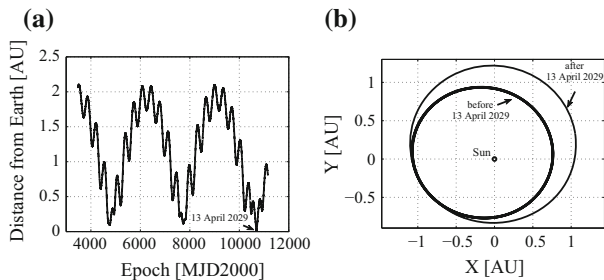
Table 2 Apophis' equinoctial variables at 3456 MJD2000 (June 18, 2009) and associated σ values obtained from the Near Earth Object Dynamic Site in September 2009

	Nom. Value	σ	
a	0.922438242375914	2.29775×10^{-8}	AU
P_1	-0.093144699837425	3.26033×10^{-8}	–
P_2	0.166982492089134	7.05132×10^{-8}	–
Q_1	-0.012032857685451	5.39528×10^{-8}	–
Q_2	-0.026474053361345	1.83533×10^{-8}	–
l	88.3150906433494	6.39035×10^{-5}	Deg

Fig. 9 Propagation of Apophis' uncertain initial conditions. **a** November 19th, 2028 (10550 MJD2000). **b** September 15th, 2029 (10850 MJD2000). **c** April 3rd, 2030 (11050 MJD2000). **d** November 8th, 2032 (12000 MJD2000)
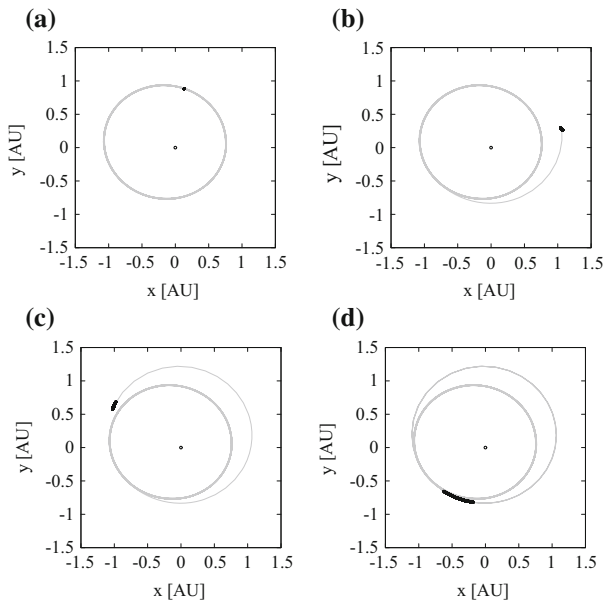
case in the following study. More specifically, Table 2 reports the Apophis' ephemerides derived from the observations available until late 2009, excluding the recent optical and radar observations performed from late 2011 onward. These data were obtained by accessing the Near Earth Object Dynamic Site (http://newton.dm.unipi.it/neodys) in September 2009.

The nonlinearities of the close encounter and of the subsequent post-encounter motion will make the uncertainty in the direction anti-parallel to the Earth heliocentric motion in 2029 drastically diverge in subsequent epochs. This is clearly illustrated in Fig. 9: the uncertain conditions of Table 2 are first propagated to epoch 10550 MJD2000 (before the close encounter) by sampling the edges of the 3σ uncertainty set with 2000 points and carrying out the associated pointwise integrations (see Fig. 9a). The integration of Apophis's motion is then continued to include the post-encounter motion. The resulting sets of final positions are illustrated in Fig. 9b–d. As can be seen, the uncertainty set tends to quickly spread along the orbit due to the perturbations induced by the close encounter and reaches a size on the order of 0.5 AU.

The performances of the standard DA-based integrator are assessed first. Apophis' initial conditions in Table 2 are initialized as DA variables, converted into Cartesian coordinates

Table 3 Maximum position error of a single 9th order Taylor representation of the flow of Apophis' dynamics over the initial uncertainty set of Table 2 for different epochs

Epoch (MJD2000)	Epoch (date)	Max. error (AU)
10550	19 Nov 2028	0.5×10^{-11}
12800	17 Jan 2035	0.5×10^{-6}
12825	11 Feb 2035	0.1×10^{-5}
12850	8 Mar 2035	0.8×10^{-3}
12875	2 Apr 2035	0.5×10^{-2}
12900	27 Apr 2035	0.2×10^{-1}

using the relations given in (Battin 1999), and then numerically propagated. Table 3 reports the maximum position error of a 9th order Taylor representation of the flow at the corners of the initial set, with respect to the pointwise integration of the same points. The errors are computed for increasing epochs. The table shows that a single Taylor polynomial of relatively high order cannot track uncertainties with sufficient accuracy for practical applications such as impact probability computation. In fact, already long before the resonant return in 2036, the accuracy of the single DA propagation becomes so low as to be practically useless for any sort of analysis or estimate.

Following this quick test of the dynamical behavior of the system, *automatic domain splitting* is enabled to treat the non-linearities in order to improve the accuracy of the standard DA-based integration. The initial conditions of Table 2 are propagated until May 1st, 2038 (14000 MJD2000), which includes the resonant return of 2036. A 9th order integration is performed and the integrator settings have been tuned to split the initial domain so to meet the requirement of tracking the uncertainties with an accuracy of the order of 10^{-9} AU and 10^{-9} AU/day for the asteroid position and velocity, respectively.

In order to limit the number of generated polynomials and associated subsets, domain splitting is disabled on any set whose volume is less than 2^{-12} times that of the initial domain. That is, any set is split at most 12 times. Instead of splitting a set further, integration is stopped at the attempt to perform a 13th split and the resulting polynomial expansion is saved as "incomplete". These incomplete polynomials are later treated separately in the analysis of the results. This limit was introduced to reduce the computational time for this exploratory work, enabling us to perform a relatively quick parametric analysis of the integrator performances. The limit can be relaxed or entirely removed in future in-depth work, especially considering that automatic domain splitting can take advantage of parallelization.

Figure 10a plots the results obtained until April 27th, 2035 (12900 MJD2000) in terms of Apophis' distance from Earth. The solid lines represent the trajectories followed by the center points of each set, whereas the lower and upper bounds of Earth's distance over the entire uncertainty set are given by the grey band. As can be seen, the integrator is able to propagate uncertainties using only one 9th order Taylor polynomial until March 3rd, 2034 (12490 MJD2000), far past the first close encounter of 2029. At this point the nonlinearities over the now relatively large uncertainty set prevent the integrator from meeting the accuracy requirements. Thus, a first split occurs in the semimajor-axis direction, which causes the initial domain to split into two sets with distinct Earth distances of their center points. The resulting two polynomials are propagated forward in time and only one additional split occurs before epoch 12700 MJD2000. The perturbation induced by the close encounter intensifies the nonlinearities and causes a significant number of splits to cascade from the three sets.

The number of splits drastically increases in subsequent epochs. This is clearly illustrated in Fig. 10b, which reports the results obtained until the final epoch 14000 MJD2000. The

Fig. 10 Earth distance profiles resulting from the propagation of the initial uncertainty set of Table 2. **a** Distance range over the entire uncertainty set (*grey band*) and trajectories of the center points (*black lines*) before the resonant return. **b** Trajectories of the center points of all sets until the final time
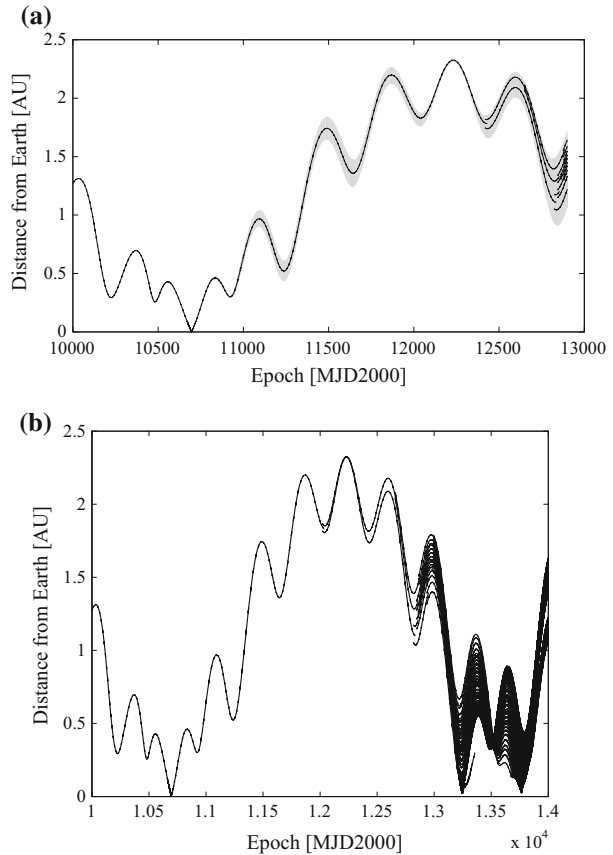

figure shows that most splits tend to occur when the trajectories get close to Earth. This is expected as Earth's gravitational perturbation is maximized there, resulting in strongly non-linear dynamics. The final number of generated domains is 2497, whereas the associated computational time is 28.13 h on a single 3.7 GHz Intel Core i7-4820K CPU. It is worth highlighting that not all domains reach the final epoch: this is due to the minimum allowed domain size we chose as described above. Once a domain of minimum size can no longer be propagated without failing to meet the accuracy requirements, its integration is stopped.

It is worth comparing the computational time required by the proposed method with the one of a classical pointwise integration. The pointwise propagation of Apophis motion from 3456 MJD2000 to 14000 MJD2000 takes an average of 50 s on the same machine. The ratio between the two computational times shows that the breakeven point between the proposed approach and a classical pointwise Monte Carlo method occurs for the propagation of a set of two thousand initial conditions, which are not sufficient to accurately represent the statistics of typical resonant returns. In addition, DA propagation with automatic domain splitting can be implemented to make use of parallelization techniques as, after splitting, each split box is propagated independently of all the others. Propagating several boxes in parallel can provide large savings in computational time. In comparison, to achieve similar results, classical pointwise Monte Carlo simulations would have to be run on massively parallel clusters with tens of thousands of cores.

Table 4 The average number of splits performed in each direction computed over all final sets at the end of the integration

Direction	a	P_1	P_2	Q_1	Q_2	l
# of splits	5.9	0	1.8	0.9	0	2.9

**Fig. 11** Final subdomains of the initial domain in the a, l-plane
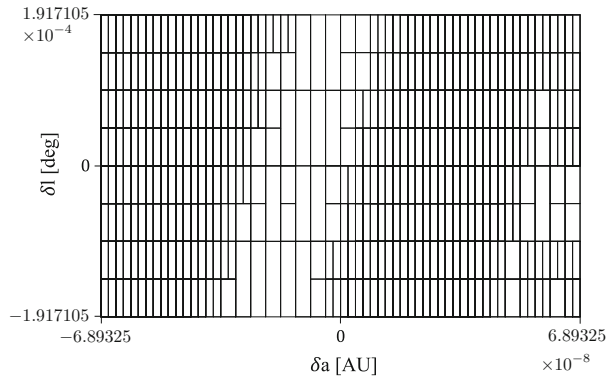
Most domains do indeed reach the minimum domain size, the average number of total splits for the final domains is 11.7. To analyze the splits more in depth, we report the average number of splits along each direction over all of the final domains in Table 4. Evidently, most splits during the integration occur in the semimajor axis and the true longitude directions, while P_1 and Q_2 are, in fact, never split. Thus without further analysis of the actual numerical results, the splitting information alone already reveals that the maximum sensitivity in the initial condition is in the a and l direction.

Consequently, even though the entire integration was performed using a full six dimensional initial condition box, we focus our analysis on the dynamics in the projection onto the a, l-plane of initial condition. Figure 11 shows this projection of the initial uncertainty set onto the a, l-plane along with all the resulting subdomains created during the propagation. This figure represents a precious source of information on the dynamics of the system. Regions of larger final domains can be easily distinguished from those areas where most splits occur causing the resulting sets to be smaller.

As explained in the previous sections, splits occur when the nonlinearities are too strong to be managed with a single Taylor polynomial. Therefore, the areas in Fig. 11 where most splits concentrate coincide with regions of strong nonlinearity, a strong indicator of chaotic behavior. The size of a box of initial conditions can be inversely related to the local expansion rate of the phase space of the system. The smaller the initial condition box must be split, the more expansive the dynamics dynamics are acting on this set of initial conditions. Intuitively, these regions are recognized to include initial conditions that attain the closest distances from Earth during their propagation. This conjecture is confirmed by Fig. 12, which superimposes a color map on Fig. 11.

The color map shown in Fig. 12 illustrates the final epoch reached by each set at the end of its propagation either because it reached the final epoch or the minimum size. As can be seen, the regions of larger domains in Fig. 11 match the black areas in Fig. 12. This means that larger sets smoothly propagate until the final time. Consequently, all the initial

Fig. 12 Color map superimposed to the final subdomains in the initial domain in the a, l-plane: final integration time
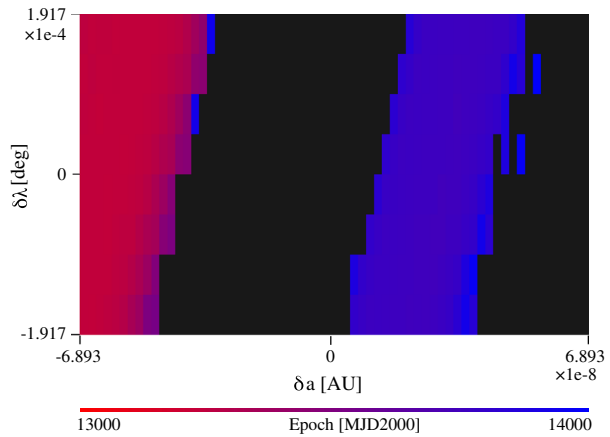
conditions lying within have no risk to impact Earth until the final epoch. On the other hand, the propagation of smaller sets tends to stop before 14000 MJD2000. This means that smaller domains and the associated colored areas might contain risky initial conditions and deserve additional analysis. Thus, Fig. 12 allows astrodynamicists to rule out the possibility of impacts from the black areas and can be used to estimate at least an upper bound for impact probability.

Additionally, the two distinct regions of different colors visible in Fig. 12 provide valuable information on the resonances. More specifically, moving from the left to the right side of the figure, the prevalent color of the first colored area indicates that the integration of most sets there stops around 2036. Similarly, the colors of the second area indicate that the integration of the associated sets are truncated around the year 2037. Consequently, the colored areas clearly mark the regions of possible resonances between Earth and Apophis: the initial conditions lying in the first and second colored areas can result in possible impacts in 2036 and 2037 respectively.

All three areas, black, blue, and red, each exhibit a distinct dynamical behavior; furthermore the blue band serves as a separatrix between the two black regions. This observation is confirmed in Fig. 13a, b. Sample initial conditions are taken from each of the four main areas of Fig. 11 as illustrated in Fig. 13a. The resulting distance profiles are plotted in Fig.13b: initial conditions (2) and (4) lying in the black areas show safe distances from Earth throughout the entire integration, whereas initial conditions (1) and (3) exhibit close encounters with Earth around the time of the associated resonance epoch.

One last result deserves to be mentioned. The initial conditions of Table 2 are derived from the observations available until late 2009. They allowed us to illustrate the advantages and potential of DA-based automatic domain splitting. In the meanwhile, additional optical and radar observations have been made available (see http://neo.jpl.nasa.gov/apophis/). Thanks to the new observations, more accurate initial conditions are available to astrodynamicists. Table 5 reports Apophis' ephemerides on July 5th, 2009 (3473.5 MJD2000) including all recent optical and radar observations. These data were obtained by accessing the Near Earth Object Dynamic Site in October 2013.

As illustrated in Table 5, recent observations allowed the standard deviation to be considerably reduced with respect to the previously available data (see Table 2). A 9th order DA-based integration of the new initial conditions reported in Table 5 shows that these new initial conditions can be smoothly propagated until the final epoch 14000 MJD2000 without requiring any split of the initial domain. Figure 14 reports the distance profile associated to

Fig. 13 Analysis for four
dynamical areas of Fig. 12.
a Selected initial conditions.
**b** Earth distance profiles
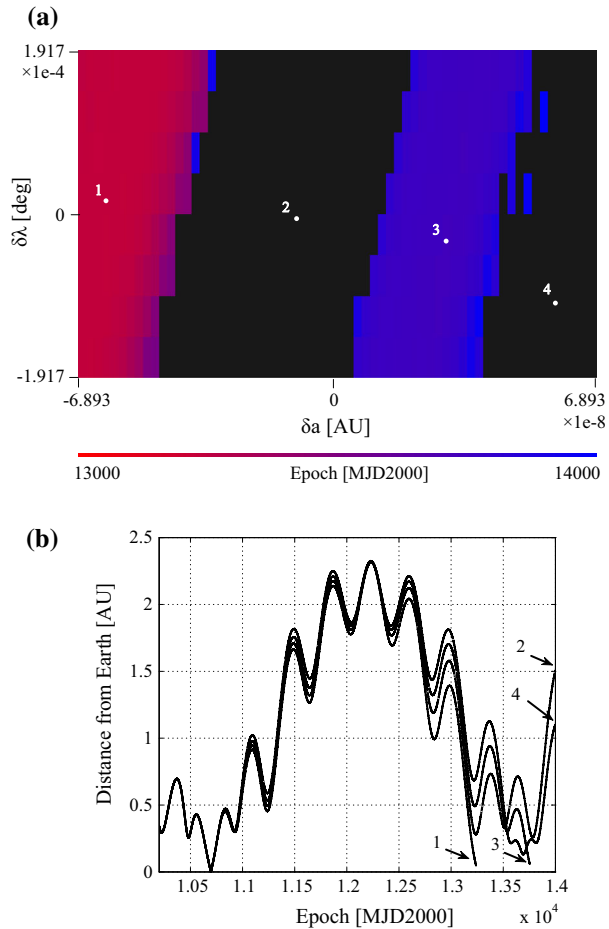
Table 5 Apophis' equinoctial
variables at 3473.5 MJD2000
(July 5, 2009) and associated σ
values obtained from the Near
Earth Object Dynamic Site in
September 2013

	Nom. Value	σ	
a	0.922443731280282	9.58934×10^{-11}	AU
P_1	-0.093137787707699	4.21204×10^{-9}	–
P_2	0.166984258496493	5.01301×10^{-9}	–
Q_1	-0.012032702063741	6.35114×10^{-9}	–
Q_2	-0.026474187976460	4.69853×10^{-9}	–
l	107.7856397515106	7.35568×10^{-7}	Deg

the new initial conditions. This matches the distances of closest approach reported in the
Near Earth Object Dynamic Site (http://newton.dm.unipi.it/neodys) and NASA's Jet Propulsion Laboratory website (http://neo.jpl.nasa.gov/apophis/). In order to compare where these
new initial conditions lie with respect to the old ones, we propagated the initial conditions
of Table 5 backward to epoch 3456 MJD2000. Figure 14a shows the result superimposed
on Fig. 12. As expected, the new initial conditions and their associated uncertainties form

Fig. 14 Apophis motion with new initial condition. **a** Uncertainty set of Table 5 projected onto the color map in Fig. 12. **b** Distance from Earth
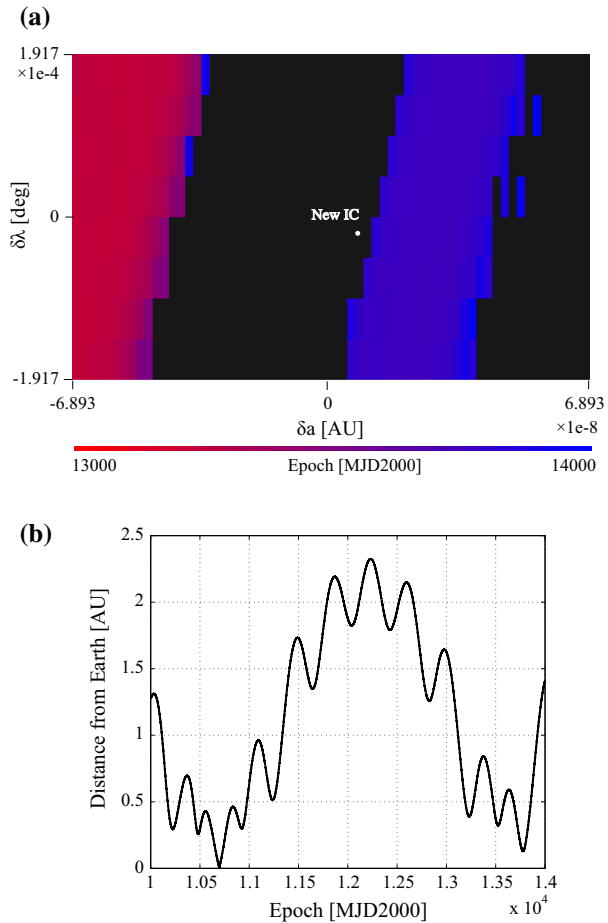
essentially a single point within the figure, located entirely within one of the black areas. This result further confirms that Apophis' impact in 2036 can be ruled out.

5 Conclusion

This paper introduced the automatic domain splitting technique into the classical high order Differential Algebra based integration with the goal of accurately propagating large sets of uncertainties in highly nonlinear dynamics and long term integrations. The resulting propagation algorithm automatically splits the initial uncertainty domain into subdomains during the integration when the polynomial expansions representing the current state do not meet a predefined accuracy requirement. 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 into the final set.

The performance of the splitting integrator has been assessed by applying it to the propagation of asteroid (99942) Apophis post-encounter motion. In order to limit the number of

generated subdomains and polynomials, a minimum domain size was fixed. Consequently, not all subdomain achieved the final integration time. It has been shown that these sets correspond to regions of strong nonlinearity in the dynamics, which are automatically identified by the integration algorithm. Each of the identified regions, in the case of Apophis, has been matched to regions where risky close encounters can either be ruled out or are possible.

References

Alessi, E.M., Farres, A., Vieiro, A., Jorba, A., and Simo, C.: Jet Transport and Applications to NEOs. In: Proceedings of the 1st IAA Planetary Defense Conference, Granada, Spain (2009)

Armellin, R., Di Lizia, P., Bernelli-Zazzera, F., Berz, M.: Asteroid close encounters characterization using differential algebra: the case of Apophis. Celest. Mech. Dyn. Astron. **107**, 451–470 (2010)

Battin, H.R.: An Introduction to the Mathematics and Methods of Astrodynamics. AIAA Education Series, Reston (1999)

Berz, M.: The new method of TPSA algebra for the description of beam dynamics to high orders. Technical report AT-6:ATN-86-16, Los Alamos National Laboratory (1986)

Berz, M.: The method of power series tracking for the mathematical description of beam dynamics. Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers Detectors Assoc Equip. **A258**, 431–436 (1987)

Berz, M.: Modern Map Methods in Particle Beam Physics. Academic press, New York (1999a)

Berz, M.: Differential Algebraic Techniques. Entry in Handbook of Accelerator Physics and Engineering. World Scientific, New York (1999b)

Berz, M., Makino, K.: COSY INFINITY Version 9 Reference Manual. MSU report MSUHEP-060803, Michigan State University, East Lansing, MI 48824, 1–84 (2006)

Bignon, E., Pinède, R., Azzopardi, V., Mercier, P.: JACK: an accurate Numerical Orbit Propagator using Taylor Differential Algebra Emmanuel. KePASSA workshop, Logroño, Spain, April 23–25, (2014)

Crassidis, J.L., Junkins, J.L.: Optimal Estimation of Dynamic Systems. CRC Press, Boca Raton (2004)

Di Lizia, P., Armellin, R., Lavagna, M.: Application of high order expansions of two-point boundary value problems to astrodynamics. Celest. Mech. Dyn. Astron. **102**, 355–375 (2008)

Di Lizia, P., Armellin, R., Bernelli Zazzera, F., Jagasia, R., Makino, K., Berz, M.: Validated Integration of Solar System Dynamics. In: Proceedings of the 1st IAA Planetary Defense Conference, Granada, Spain (2009)

Giorgini, J.D., Benner, L.A., Ostro, S.J., Nolan, M.C., Busch, M.W.: Predicting the Earth encounters of (99942) Apophis. Icarus **193**, 1–19 (2008)

Giza, D.R., Singla, P., Jah, M.: An approach for nonlinear uncertainty propagation: Application to orbital mechanics. In: Proceedings of the AIAA Guidance, Navigation, and Control Conference, Chicago, Illinois (2009)

Horwood, J.T., Poore, A.B.: Adaptive Gaussian sum filters for space surveillance. IEEE Trans. Automatic Control **56**, 1777–1790 (2011)

Jah, M., Kelecy, T. M.: Orbit Determination Performance Improvements for High Area-to-Mass Ratio Space Object Tracking Using an Adaptive Gaussian Mixtures Estimation Algorithm. In: Proceedings of the 60th International Astronautical Congress, Daejeon, Republic of Korea (2009)

Julier, S.J., Uhlmann, J.K.: Unscented filtering and nonlinear estimation. Proc IEEE **92**, 401–422 (2004)

Julier, S.J.: The scaled unscented transformation. In: Proceedings of the American Control Conference, Piscataway, New Jersey (2002)

Makino, K., Berz, M.: Rigorous Integration of Flows and ODEs using Taylor Models. In: Proceedings of the 2009 conference on Symbolic Numeric Computation, pp. 79–84 (2009)

Majji, M., Junkins, J.L., Turner, J.D.: A high order method for estimation of dynamic systems. J. Astronaut. Sci. **56**, 1–32 (2008)

Maybeck, P.S.: Stochastic Models, Estimation and Control. Academic press, New York (1982)

Montenbruck, O., Eberhard, G.: Satellite orbits. Springer, New York (2001)

Moore, R.E., Kearfott, R.B., Cloud, M.J.: Introduction to Interval Analysis. Society for Industrial and Applied Mathematics, Philadelphia (2009)

Morselli, A., Armellin, R., Di Lizia, P., and Bernelli-Zazzera, F.: Computing Collision Probability Using Differential Algebra and Advanced Monte Carlo Methods. In: Proceedings of the 63rd International Astronautical Congress, Napoli, Italy (2010)

Park, R.S., Scheeres, D.J.: Nonlinear mapping of Gaussian statistics: theory and applications to spacecraft trajectory design. J. Guidance Control Dyn. **29**, 1367–1375 (2006)

Rudin, W.: Principles of Mathematical Analysis. McGraw-Hill Inc., New York (1976)

Seidelmann, P.K.: Explanatory Supplement to the Astronomical Almanac. University Science Books, Sausalito, CA (1992)

Valli, M., Armellin, R., Di Lizia, P., Lavagna, M.: Nonlinear Mapping of Uncertainties in Celestial Mechanics. J. Guidance Control Dyn. **36**, 48–63 (2012)

Valsecchi, G., Milani, A., Rossi, A., Tommei, G.: The SRT, Near-Earth objects, and space debris. Memorie della Societa Astronomica Italiana Supplementi **10**, 186 (2006)

Vittaldev, V., Russell, R.P., Arora, N., Gaylor, D.: Second-order Kalman filters using multi-complex step derivatives. In: Proceedings of the AAS/AIAA Space Flight Mechanics Meeting, Kauai, Hawaii (2012)