A differential algebra-based importance sampling method for impact probability computation on Earth resonant returns of near-Earth objects

Matteo Losacco,[1]⋆ Pierluigi Di Lizia,[1]⋆ Roberto Armellin[2]⋆ and Alexander Wittig[3]

[1]*Department of Aerospace Science and Technology, Politecnico di Milano, Via G. La Masa 34, I-20156 Milano, Italy*
[2]*Surrey Space Centre, University of Surrey, Guildford GU2 7XH, UK*
[3]*Aeronautics, Astronautics and Computational Engineering Unit, University of Southampton, Southampton SO17 1BJ, UK*

ABSTRACT

A differential algebra-based importance sampling method for uncertainty propagation and impact probability computation on the first resonant returns of near-Earth objects is presented in this paper. Starting from the results of an orbit determination process, we use a differential algebra-based automatic domain pruning to estimate resonances and automatically propagate in time the regions of the initial uncertainty set that include the resonant return of interest. The result is a list of polynomial state vectors, each mapping specific regions of the uncertainty set from the observation epoch to the resonant return. Then, we employ a Monte Carlo importance sampling technique on the generated subsets for impact probability computation. We assess the performance of the proposed approach on the case of asteroid (99942) Apophis. A sensitivity analysis on the main parameters of the technique is carried out, providing guidelines for their selection. We finally compare the results of the proposed method to standard and advanced orbital sampling techniques.

Key words: methods: statistical – celestial mechanics – minor planets, asteroids: individual: (99942) Apophis.

1 INTRODUCTION

Over the last 30 yr, significant efforts have been devoted to develop new tools for detection and prediction of planetary encounters and potential impacts by near-Earth objects (NEOs). The task introduces relevant challenges due to the imperative of early detection and accurate estimation and propagation of their state and associated uncertainty set (Chesley 2005). The problem is made more complicated by the fact that the dynamics describing the motion of these objects is highly non-linear, especially during close encounters with major bodies. Non-linearities of the orbital dynamics tend to significantly stretch the initial uncertainty sets during the time propagation. Non-linearities are not confined to object dynamics only: even simple conversions between coordinate systems introduce non-linearities, thus affecting the accuracy of classical propagation techniques (Wittig et al. 2015). Present day approaches for robust detection and prediction of planetary encounters and potential impacts by NEO mainly refer to linearized models or full non-linear orbital sampling (Farnocchia et al. 2015). The impact probability computation by means of linear methods in the impact plane

was introduced by Chodas (1993), whereas the introduction of the Monte Carlo technique to this problem was developed by Yeomans & Chodas (1994) and Chodas & Yeomans (1999), who suggested to apply the method to sample the linear 6D confidence region at the observation epoch and then numerically integrate over the time interval of investigation using fully non-linear equations (Milani et al. 2002). Milani, Chesley & Valsecchi (1999), Milani (1999), Milani, Chesley & Valsecchi (2000a) and Milani et al. (2000b) applied the multiple solutions approach to sample the central line of variations (LOV) of the non-linear confidence region at the initial epoch and then numerically integrate over the time span of interest in a similar way. Within the framework of the impact probability computation of resonant returns, a well-known approach relies on the concept of keyholes, small regions of the impact plane of a specific close encounter such that, if an asteroid passes through one of them, it will hit the Earth on subsequent return (Gronchi & Milani 2001; Milani et al. 2002; Valsecchi et al. 2003).

The preferred approach to detecting potential impacts depends on the uncertainty in the estimated orbit, the investigated time window and the dynamics between the observation epoch and the epoch of the expected impact (Farnocchia et al. 2015). Linear methods are preferred when linear approximations are reliable for both the orbit determination and uncertainty propagation. When these assumptions are not valid, one must resort to more computationally

⋆ E-mail: matteo.losacco@polimi.it (ML); pierluigi.dilizia@polimi.it (PDL); r.armellin@surrey.ac.uk (RA)

intensive techniques: among these, Monte Carlo methods are the most accurate but also the most computationally intensive, whereas the LOV method guarantees compute times 3-4 orders of magnitude lower than those required in MC simulations, though the LOV analysis may grow quite complex after it has been stretched and folded by multiple close planetary encounters, leaving open the possibility of missing some pathological cases (Farnocchia et al. 2015).

Alternative approaches rely on the use of differential algebra (DA). DA supplies the tools to compute the derivatives of functions within a computer environment, i.e. it provides the Taylor expansion of the flow of ordinary differential equations (ODEs) by carrying out all the operations of any explicit integration scheme in the DA framework (Berz 1999; Wittig et al. 2015). DA has already proven its efficiency in the non-linear propagation of uncertainties (Armellin et al. 2010b; Valli et al. 2013). Nonetheless, the accuracy of the method drastically decreases in highly non-linear dynamics. The propagation of asteroids motion after a close encounter with a major body is a typical case.

A DA-based automatic domain splitting algorithm was presented by the authors in the past to overcome the limitations of simple DA propagation (Wittig et al. 2015). The method can accurately propagate large sets of uncertainties in highly non-linear dynamics and long-term time spans. The propagation algorithm automatically splits the initial uncertainty domain into subsets when the polynomial expansions representing the current state do not meet predefined accuracy requirements. The performance of the algorithm was assessed on the case of asteroid (99942) Apophis, providing a description of the evolution of the uncertainty set to the epoch of predicted close encounters with Earth in 2036 and 2037 (Wittig et al. 2015). Though representing a significant improvement with respect to simple DA propagation, the approach required a not negligible computational effort in propagating the whole set of generated subdomains. Moreover, no information about the impact probability for asteroid Apophis was provided, as the propagation of the uncertainty set was stopped before the close encounters.

We present in this paper an evolution of the automatic domain splitting algorithm. The method, referred to as automatic domain pruning, automatically identifies possible resonances after a close encounter with a major body. Then, assuming no intervening close approaches with other celestial bodies in between, it optimizes the propagation to the first resonant returns, by limiting the propagation of the uncertainty set to the regions that generate a close encounter with that celestial body at the investigated epoch. The result is a list of polynomial state vectors, each mapping only specific subsets of the initial domain to the resonant return epoch. Taking advantage of the availability of the polynomial maps, a DA-based Monte Carlo importance sampling technique is then used to generate samples in the propagated subsets and provide an estimate for the impact probability at the epoch of the selected resonant return. The proposed approach does not apply any simplification step on the uncertainty domain associated with the orbit determination process. Thus, the method is proposed as an alternative approach with respect to equivalent techniques, such as a full Monte Carlo simulation or other 6D-based orbital sampling techniques, which will represent the main term of comparison for our analysis.

The paper is organized as follows. First, we present a description of the automatic domain pruning and importance sampling techniques, showing the application to the case of the first resonant return. Then, we apply the method to the critical case of asteroid (99942) Apophis, providing an estimate of the impact probability for the resonant return in 2036. Finally, we carry out a sensitivity
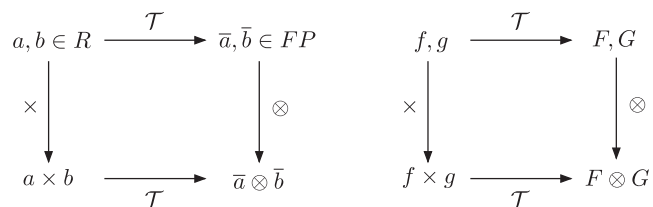
Figure 1. Analogy between the FP representation of real numbers in computer environment (left) and the algebra of Taylor polynomials in DA framework (right) (Di Lizia 2008).

analysis on the main parameters of the method, presenting a comparison with standard and advanced orbital sampling techniques.

2 DIFFERENTIAL ALGEBRA AND AUTOMATIC DOMAIN SPLITTING

Differential algebra provides the tools to compute the derivatives of functions within a computer environment (Ritt 1932, 1948; Risch 1969, 1970; Kolchin 1973; Berz 1999). Historically, the treatment of functions in numerics has been based on the treatment of numbers, and the classical numerical algorithms are based on the evaluation of functions at specific points. The basic idea of DA 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 (Berz 1999). Real numbers, indeed, are approximated by floating point (FP) numbers with a finite number of digits. With reference to Fig. 1, let us consider two real numbers a and b, and their FP counterpart $\bar{a}$ and $\bar{b}$, respectively: given any operation $\times$ in the set of real numbers, an adjoint operation $\otimes$ is defined in the set of FP numbers so 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 DA, these functions are converted into their Taylor series expansions, F and G, respectively. In this way, the transformation of real numbers in their FP representation is now substituted by the extraction of the Taylor expansions of f and g (see Fig. 1, right). For each operation in the function space, an adjoint operation in the space of Taylor polynomials is defined such that the corresponding diagram commutes.

The implementation of DA in a computer environment provides the Taylor coefficients of any function of v variables up to a specific order n. More specifically, by substituting classical real algebra with the implementation of a new algebra of Taylor polynomials, any function f of v variables can be expanded into its Taylor expansion up to an arbitrary order n, along with the function evaluation, with a limited amount of effort. The Taylor coefficients of order n for sum and product of functions, as well as scalar products with real numbers, can be directly computed from those of summands and factors. As a consequence, the set of equivalence classes of functions can be endowed with well-defined operations, leading to the so-called truncated power series algebra. In addition to basic algebraic operations, differentiation and integration can be easily introduced in the algebra, thus finalizing the definition of the differential algebra structure of DA (Berz 1986, 1987). The DA used in this work is implemented in the DACE software (Massari, Di Lizia & Rasotto 2017).

A relevant application of DA is the automatic high-order expansion of the solution of an ODE with respect to the initial conditions (Berz 1999; Di Lizia, Armellin & Lavagna 2008; Massari et al. 2017). This expansion can be achieved by considering that any integration scheme, explicit or implicit, is characterized by a finite number of algebraic operations, involving the evaluation of the ODE right-hand side (RHS) at several integration points. Therefore, replacing the operations between real numbers with those on DA numbers, it yields to the nth order Taylor expansion of the flow of the ODE, $\phi(t; \delta x_0, t_0) = \mathcal{M}_\phi(\delta x_0)$, at each integration time, assuming a perturbed initial condition $x_0 + \delta x_0$. Without loss of generality, consider the scalar initial value problem:

$$\dot{x}(t) = f(t, x), \quad x(t_0) = x_0 \qquad (1)$$

and the associated flow $\phi(t; \delta x_0, t_0)$. For simplicity, consider uncertain initial conditions only. Starting from the nth order DA representation of the initial condition, $[x_0] = x_0 + \delta x_0$, which is an $(n + 1)$-tuple of Taylor coefficients, and performing all the operations in the DA framework, we can propagate the Taylor expansion of the flow in x_0 forward in time, up to the final time t_f. Consider, for example, the forward Euler's scheme:

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

and replace the initial value with the DA expression $[x_0] = x_0 + \delta x_0$. The first time step yields

$$[x_1] = [x_0] + f([x_0])\Delta t. \qquad (3)$$

If the function f is evaluated in the DA framework, the output of the first step, $[x_1]$ is the nth order Taylor expansion of the flow $\phi(t; \delta x_0, t_0)$ in x_0 for $t = t_1$. Note that, as a result of the DA evaluation of $f([x_0])$, the $(n + 1)$-tuple $[x_1]$ may include several non-zero coefficients corresponding to high-order terms in δx_0. The previous procedure can be repeated for the subsequent steps. The result at the final step is the nth order Taylor expansion of $\phi(t; \delta x_0, t_0)$ in x_0 at the final time t_f. Thus, the flow of a dynamical system is approximated, at each time step t_i, with its nth order Taylor expansion in a fixed amount of effort. Any explicit ODE integration scheme can be rewritten as a DA scheme. 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 to obtain high-order expansions of the flow. It is therefore independent on the RHS of the ODE and it is computationally efficient. Unfortunately, DA fails to accurately describe, with a single polynomial map, the evolution in time of an uncertainty set in case of highly non-linear dynamics or long term propagation. The approximation error is strictly related to the size of the domain the polynomial is defined in (Wittig et al. 2015). 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:

$$|f(\delta x) - P_f(\delta x)| \leqslant C \cdot \delta x^{n+1} \qquad (4)$$

for some constant $C > 0$. Consider now the maximum error e_r of P_f on a domain B_r of radius $r > 0$ around the expansion point. Considering equation (4), we obtain:

$$|f(\delta x) - P_f(\delta x)| \leqslant C \cdot \delta x^{n+1} \leqslant C \cdot r^{n+1} = e_r. \qquad (5)$$
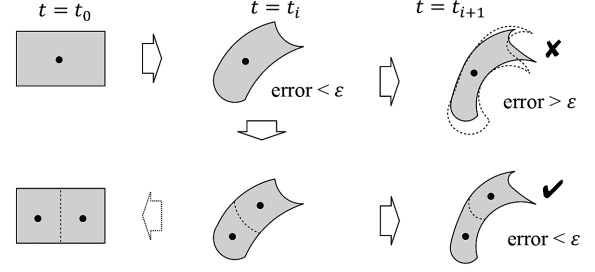
Figure 2. ADS algorithm schematic illustration (Wittig et al. 2015).

If the domain of P_f is reduced from B_r to $B_{r/2}$ of radius $r/2$, the maximum error of P_f over $B_{r/2}$ will decrease by a factor $1/2^{n+1}$:

$$|f(\delta x) - P_f(\delta x)| \leqslant C \cdot \delta x^{n+1} \leqslant C \cdot \left(\frac{r}{2}\right)^{n+1} = \frac{e_r}{2^{n+1}}. \qquad (6)$$

By subdividing the initial domain into smaller domains and computing the Taylor expansion around the centre points of the new domains, the error greatly reduces, whereas the expansions still cover the entire initial set. Starting from these considerations, automatic domain splitting (ADS) employs an automatic algorithm to determine at which time t_i the flow expansion over the set of initial conditions is no longer able to describe the dynamics with enough accuracy (Wittig et al. 2015). 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 centre points, two separate polynomial expansions are obtained. By defining with x_i the splitting direction, both generated polynomial expansions P_1 and P_2 have terms of order n in x_i smaller by a factor of 2^n with respect to the original polynomial expansion P. Thus, the splitting procedure guarantees a more accurate description of the whole uncertainty set at the current time epoch t_i. After such a split occurs, the integration process is resumed on both generated subsets, until new splits are required. A representation of the ADS procedure is shown in Fig. 2.

The decision on the splitting epoch and, in case of multivariate polynomials, the splitting direction relies on estimating the size of the $(n + 1)$th order terms of the polynomial using an exponential fit of the size of all the known non-zero terms up to order n. If the size of the truncated order becomes too large, we decide to split the polynomial. This method allows us to consider all the information available in the polynomial expansion and to obtain an accurate estimate of the size of the $(n + 1)$th order term, the first discarded order. The exponential fit is chosen because, after reducing the domain with a sufficient number of splits, the coefficients of the resulting polynomial expansion decay exponentially as a direct consequence of Taylor's theorem. A mathematical description is offered hereafter and follows the scheme presented in Wittig et al. (2015). Consider a polynomial P of order n of the form

$$P(\boldsymbol{x}) = \sum_\alpha a_\alpha \boldsymbol{x}^\alpha \qquad (7)$$

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:

$$S_i = \sum_{|\alpha|=i} |a_\alpha|. \qquad (8)$$
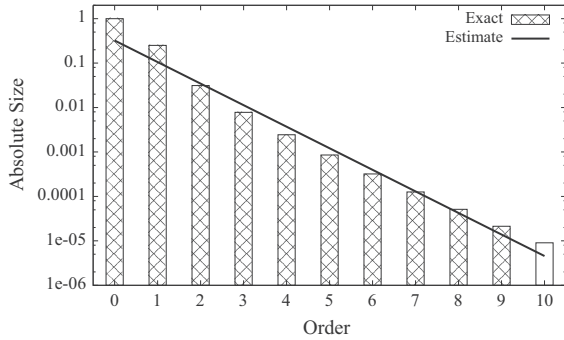
Figure 3. Truncation error estimation for the Taylor expansion of $\sqrt{1 + x/2}$ via exponential fitting. Terms of order up to 9 are used for the fitting (Wittig et al. 2015).

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

$$f(i) = Ae^{Bi} \tag{9}$$

is used to determine the coefficients A and B such that $f(i) = S_i$, $i \in I$, is approximated optimally in 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. An example of the application of the method is shown in Fig. 3, where the polynomial is the Taylor expansion of $\sqrt{1 + x/2}$ up to order 9. The size S_i of each order is shown as bars, whereas the resulting fitted function f is shown as a line.

In the case of multivariate polynomials $P(\boldsymbol{x}) = P(x_1, x_2, \ldots, x_v)$, the split is performed in one component x_i. We determine the splitting direction using a method similar to the one adopted for the splitting decision. For each $j = 1, \ldots, v$ we begin by factoring the known coefficients of P of order up to n with respect to x_j, i.e.

$$P(x_1, x_2, \ldots, x_v) = \sum_{m=0}^{n} x_j^m \cdot q_{j,m}(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_v), \tag{10}$$

where the polynomials $q_{j,m}$ do not depend on x_j. The size $S_{j,m}$ of the polynomials $q_{j,m}$ is estimated by the sum of the absolute values of their coefficients. Then, the exponential fitting routine is applied to estimate the size $S_{j,n+1}$ of the truncated terms of order $n + 1$ in x_j. Finally, the splitting direction i is chosen as the component x_j with the 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.

The main parameters of the algorithm are the tolerance for the splitting procedure and the maximum number of allowed splits $N_{\max}$. The first parameter is selected according to the required precision of the polynomial expansions and determines the splitting epochs: when the estimated truncation error exceeds the imposed tolerance, the current domain is split. As a direct consequence of the ADS procedure, the maximum error over the obtained set of polynomials decreases with the selected splitting precision. However, the maximum error is always larger than the selected integration precision. This difference is actually expected, as the splitting tolerance plays a similar role as the one-step error set in the automatic step size control of the integration scheme (Wittig et al. 2015). 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 only re-expands the polynomial to prevent

exponential growth in future integration steps. The ideal tolerance depends on both the dynamics and the integration time, and it has to be chosen heuristically to ensure that the final result satisfies the accuracy requirements of the application. A numerical example is shown in Section 7.1.

The second parameter plays the role of limiting the number of generated subdomains by imposing a minimum size for the generated subsets: domain splitting is disabled on any set whose volume is less than $2^{-N_{\max}}$ times that of the initial domain. That is, any set is split at most $N_{\max}$ times. Then, instead of splitting a set further, integration is stopped at the attempt to perform the $(N_{\max} + 1)$th split and the resulting polynomial expansion is saved as incomplete. Incomplete polynomials are later treated separately in the analysis of the results (Wittig et al. 2015).

When each generated subset reaches either the final simulation time or the minimum box size, the ADS propagation terminates, and the result is a list of polynomial expansions, each covering a specific subset of the domain of initial conditions. A more detailed description of the ADS algorithm can be found in Wittig et al. (2015).

3 AUTOMATIC DOMAIN PRUNING

As described in Section 2, automatic domain splitting provides an accurate description of the evolution in time of a given uncertainty set by splitting the domain in subsets when required. Unfortunately, this method may entail a not negligible computation effort, as all generated subsets are propagated to the final simulation time or till the minimum box size is reached. While this approach is unavoidable when the behaviour of the whole uncertainty set is analysed, it becomes a strong limitation when only a portion of the initial set is to be investigated. This is the case when the first resonant return of an NEO is studied. Resonant returns occur when, during a close encounter, an asteroid is perturbed into an orbit with a period $T' \sim k/h$ yr. Thus, after h revolutions of the asteroid and k revolutions of the Earth, both celestial bodies are in the same region of the first close encounter and a second one may occur. Given the initial uncertainty set, only a portion of it may lead to the resonant return. It would be therefore interesting to limit the propagation to this region only.

Starting from these considerations, the automatic domain pruning (ADP) we present in this paper combines the ADS algorithm with a pruning technique with the aim of limiting the number of propagated subsets. We make here the assumption of no close approaches with other celestial bodies between the first close encounter and the selected resonant return. This assumption is easily checked right before the ADP propagation, as later explained in Section 7.2.

The first phase of the algorithm consists in propagating the whole uncertainty set by means of ADS propagation up to the epoch of the first close encounter. The availability of the polynomial expansion of the state vector of the object with respect to the initial uncertainty provides the polynomial expansion of the orbital period of the object after the close encounter. By using a polynomial bounder, we can estimate the range of all possible values of the orbital period after the close encounter and, thus, retrieve all possible resonances with the planet, i.e. all orbital periods included in the computed orbital period range leading to a resonant return with the planet.

Once all resonances are computed, the analysis focuses on a single resonance, and the propagation is resumed. Every time a new subset is generated, the method automatically identifies if the set may lead to the investigated resonant return or not. By exploiting the knowledge of the DA state vector at the epoch of the first close
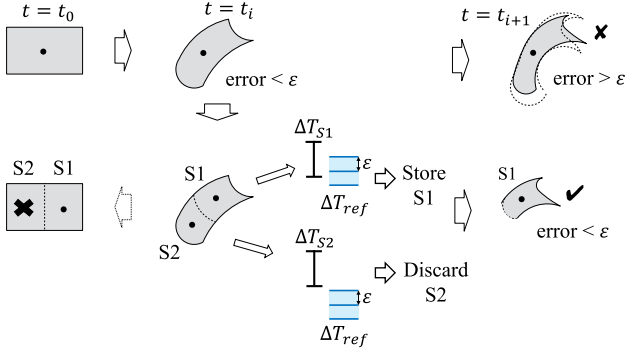
Figure 4. ADP algorithm illustration. Pruning is performed by comparing the estimated subset orbital period range ΔT_{sub} with the reference range ΔT_{ref}.

encounter, indeed, we can assign a given orbital period range to each generated subset. This range, defined as ΔT_{sub}, is compared to a reference range ΔT_{ref}, centred in the resonance period T' with a semi-amplitude ε, $\Delta T_{\text{ref}} = [(1 - \varepsilon)T', (1 + \varepsilon)T']$. We select the reference range in order to consider small dynamical perturbations between the first close encounter and the resonant return. If ΔT_{sub} is at least partially included in the reference range, then the current subset is retained, and its propagation is continued. If ΔT_{sub} is not included in the reference range, then the initial conditions included in the current subset do not lead to a resonant return at the investigated epoch, and so the subset is discarded. This way, subsets are dynamically pruned during the ADS propagation. An illustration of the ADP algorithm is shown in Fig. 4.

The ADP algorithm, therefore, does not alter the sequence of generated subdomains, but limits the propagation in time to those subsets that are involved in the investigated resonant return. This pruning action has a positive impact on the overall computational burden, since the computational effort required by the propagation of all the discarded subsets is saved. As only subsets with close approaches to the Earth at the epoch of the investigated resonant return are maintained, the result at the end is a set of subdomains whose propagation stops slightly before the epoch of the investigated resonant return for having reached their minimum box size.

4 IMPORTANCE SAMPLING METHOD

The output of the ADP propagation is a list of subsets at epochs close to the investigated resonant return. Still, no value for the impact probability is available. We obtain an estimate for the impact probability by sampling the generated subsets and propagating the samples till they reach their minimum geocentric distance. Among all possible sampling technique, we employ the Importance Sampling (IS) method (Zio 2013).

The IS method amounts to replacing the original probability density function (pdf) $q_x(x)$ with an importance sampling distribution (ISD) $\tilde{q}_x(x)$ arbitrarily chosen by the analyst so as to generate a large number of samples in the importance region of the phase space F, the region of initial conditions leading to an impact with Earth at the epoch of the resonant return. In the case under study, we select the auxiliary distribution in order to limit as much as possible the generation of the samples to the subsets that get through the dynamic pruning. The IS algorithm is the following:

(i) Identify a proper $\tilde{q}_x(x)$.

(ii) Express the impact probability $p(F)$ as a function of $\tilde{q}_x(x)$.

$$p(F) = \int I_{\text{F}}(x)q_x(x)\mathrm{d}x = \int \frac{I_{\text{F}}(x)q_x(x)}{\tilde{q}_x(x)}\tilde{q}_x(x)\mathrm{d}x, \qquad (11)$$

where $I_{\text{F}}(x) : \mathbb{R}^v \to \{0, 1\}$ is an indicator function such that $I_{\text{F}}(x) = 1$ if $x \in F$, 0 otherwise.

(iii) Draw N_{T} samples $\{x^k : k = 1, 2, \ldots, N_{\text{T}}\}$ from the importance sampling distribution $\tilde{q}_x(x)$. If a good choice for the auxiliary pdf is made, the generated samples concentrate in the region F.

(iv) Compute the estimate $\hat{p}(F)$ for the impact probability $p(F)$ by resorting to equation (11):

$$\hat{p}(F) = \frac{1}{N_{\text{T}}} \sum_{k=1}^{N_{\text{T}}} \frac{I_{\text{F}}(x^k)q_x(x^k)}{\tilde{q}_x(x^k)}. \qquad (12)$$

(v) Compute the variance of the estimator $\hat{p}(F)$ as:

$$\sigma^2(\hat{p}) = \frac{1}{N_{\text{T}}} \left(\int \frac{I_{\text{F}}^2(x)q_x^2(x)}{\tilde{q}_x^2(x)}\tilde{q}_x(x)\mathrm{d}x - p^2(F) \right) \qquad (13)$$

$$\approx \frac{1}{N_{\text{T}}} \left(\widehat{p^2(F)} - \hat{p}^2(F) \right).$$

The selection of the ISD represents the most critical point for the method. Several techniques have been developed in order to find the one giving small variance for the estimator (Zio 2013). In this paper, we shape the ISD according to the result of the ADP propagation. As described in Section 3, the ADP propagation provides a list of subsets whose propagation is stopped slightly before the resonant return. All subsets are identified as potentially hazardous subdomains (PHSs), but no probability ranking is provided by the ADP propagation. Starting from these considerations, we define the ISD as a uniform probability density function including all the generated subsets over the whole domain. This selection allows us to increase the number of samples drawn in the PHSs and, eventually, in the impact-leading region.

5 AUTOMATIC DOMAIN PRUNING IMPORTANCE SAMPLING METHOD

The combination of the methods presented in Sections 3 and 4 yields the ADP importance sampling method (ADP–IS) for uncertainty propagation and impact probability computation of the first resonant returns of NEO. The starting point is represented by the output of an orbit determination process of a given NEO at the observation epoch t_0. This output can be expressed in terms of estimated state vector and related covariance matrix. Then, the steps of the ADP propagation phase are the following:

(i) Consider the initial state vector and related pdf and perform an analysis to identify possible epochs of close encounters and resonant returns. The analysis is carried out by propagating the uncertainty set using ADS up to the first close encounter, computing the semi-major axis dispersion over the set with a polynomial bounder and identifying the resonant frequencies. The validity of the resonances is then checked as explained in Section 7.2.

(ii) Select a resonance and identify its epoch t_{res}.

(iii) Perform an ADP propagation till the epoch t_{res}. Every time a split is required, compare the orbital period range of the current subset ΔT_{sub} with the reference range ΔT_{ref}:

$$\Delta T_{\text{sub}} \cap \Delta T_{\text{ref}} \begin{cases} = 0 & \text{discard the current subset} \\ \neq 0 & \text{include the current subset} \end{cases} \qquad (14)$$

The method provides a set of n_{PHS} PHSs and related DA state vectors $[x_f^i]$ at the truncation time t_f^i, with $i = 1, \ldots, n_{\text{PHS}}$. Vector

Table 1. Apophis equinoctial parameters and related uncertainties on 2009 June 18 00:00:00 (TDB).

		Nominal value	
a	0.922 438 242 375 914	$2.297\,75 \times 10^{-8}$	au
P_1	−0.093 144 699 837 425	$3.260\,33 \times 10^{-8}$	–
P_2	0.166 982 492 089 134	$7.051\,32 \times 10^{-8}$	–
Q_1	−0.012 032 857 685 451	$5.395\,28 \times 10^{-8}$	–
Q_2	−0.026 474 053 361 345	$1.835\,33 \times 10^{-8}$	–
l	88.315 090 643 349 4	$6.390\,35 \times 10^{-5}$	°

$[\boldsymbol{x}_f^i]$ is a polynomial state vector, each component being a function of the initial conditions $\boldsymbol{x}_0^i$.

The IS phase is initialized by setting the value of the estimated impact probability $\hat{p}_{\text{old}}$ and the number of iterations n_{it} equal to zero. Then, the steps of the algorithm are the following:

(i) Define the ISD function $\tilde{q}_x(\boldsymbol{x})$ as a uniform pdf including all the generated PHSs.

(ii) Set $n_{it} = n_{it} + 1$ and draw one sample $\boldsymbol{x}_0^{it}$ from $\tilde{q}_x(\boldsymbol{x})$.

(iii) Check if the sample belongs to one of the PHSs: if it is out of the PHSs, go back to step (ii), otherwise identify the correct PHS i the sample belongs to.

(iv) Compute the algebraic state vector $\boldsymbol{x}_f^{it}$ corresponding to the drawn sample $\boldsymbol{x}_0^{it}$ at the truncation epoch t_f^i by performing a polynomial evaluation of the DA state vector $[\boldsymbol{x}_f^i]$ at $\boldsymbol{x}_0^{it}$. That is, $\boldsymbol{x}_f^{it} = [\boldsymbol{x}_f^i](\boldsymbol{x}_0^{it})$.

(v) Propagate the state vector $\boldsymbol{x}_f^{it}$ from t_f^i to the epoch of the selected resonant return.

(vi) Compute the minimum geocentric distance $d|_{\text{res}}^{it}$ and evaluate the indicator I_{F}^{it}

$$I_{\text{F}}^{it} = \begin{cases} 0 & \text{if } d|_{\text{res}}^{it} > R_{\oplus} \\ 1 & \text{if } d|_{\text{res}}^{it} < R_{\oplus} \end{cases} \tag{15}$$

(vii) If $I_{\text{F}}^{it} = 0$, go back to step (ii), otherwise evaluate the new impact probability $\hat{p}_{\text{new}}$. By reformulating equation (12), we obtain:

$$\hat{p}_{\text{new}} = \frac{1}{N_{\text{T}}} \sum_{k=1}^{N_{\text{T}}} \frac{I_{\text{F}}(\boldsymbol{x}^k) q_x(\boldsymbol{x}^k)}{\tilde{q}_x(\boldsymbol{x}^k)} = \frac{1}{n_{it}} (\hat{I} + q_x(\boldsymbol{x}_0^{it})) \frac{1}{\tilde{q}_x(\boldsymbol{x}_0^{it})} \tag{16}$$

where $q_x(\boldsymbol{x}_0^{it})$ is the value of the original pdf in $\boldsymbol{x}_0^{it}$, $\tilde{q}_x(\boldsymbol{x}_0^{it})$ is the value of the auxiliary pdf in $\boldsymbol{x}_0^{it}$, whereas the term $\hat{I}$ represents the summation of all terms $I_{\text{F}}(\boldsymbol{x}_0^k) q_x(\boldsymbol{x}_0^k)$ of the previous iterations. The total number of samples considered for the estimation is n_{it}, i.e. the number of drawn samples when the estimate is computed. Note that, since the ISD is uniform over the whole set of PHSs, it can be extracted from the summation.

(viii) Compare $\hat{p}_{\text{old}}$ and $\hat{p}_{\text{new}}$: if the relative difference is larger than an imposed tolerance, go back to step (ii), otherwise stop.

6 NUMERICAL SIMULATIONS: THE CASE OF ASTEROID (99942) APOPHIS

In this section, we assess the performance of the ADP–IS method on the evaluation of the impact probability for the test case of asteroid (99942) Apophis. Table 1 shows the nominal initial state and associated uncertainties σ for Apophis on 2009 June 18 expressed in terms of equinoctial parameters $\boldsymbol{p} = (a, P_1, P_2, Q_1, Q_2, l)$, considering a diagonal covariance matrix. Data were obtained from the

Near Earth Objects Dynamic Site[1] (Chesley & Milani 1999; Milani et al. 2005) in September 2009.

We selected a diagonal covariance matrix in order to help distinguish the contribution of the six orbital parameters and test our method in a scenario in which the uncertainty volume is maximized. In general, however, this selection may lead to quite inaccurate results as uncertainties may be highly correlated. Nevertheless, the method can be applied in the most general case of full covariance matrix exactly in the same way, with the only difference that the DA variables would be placed along the directions of the covariance eigenvectors to avoid artificially adding extra-volume in the initial domain definition.

As previously stated, the starting point, not including recent optical and radar observations performed from late 2011 onwards, was selected in order to test the algorithm against the most critical scenario. Asteroid Apophis will have a close encounter with Earth on 2029 April 13 with a nominal distance of 3.8×10^4 km (Chesley 2005). According to the selected initial conditions, though an impact in 2029 can be ruled out, the perturbations induced by the encounter open the door to resonant returns in 2036 and 2037. The aim is therefore to apply the presented method to provide an estimate for the impact probability at the epoch of the first resonant return, in 2036.

The motion of Apophis in the Solar system is modelled according to the $(N + 1)$ body problem, including relativistic corrections to the Newtonian forces (Seidelmann 1992; Wittig et al. 2015). Specifically, the full equation is

$$\ddot{\boldsymbol{r}} = G \sum_i \frac{m_i(\boldsymbol{r}_i - \boldsymbol{r})}{r_i^3} \left\{ 1 - \frac{2(\beta + \gamma)}{c^2} G \sum_j \frac{m_j}{r_j} - \frac{2\beta - 1}{c^2} \right.$$
$$\times G \sum_{j \neq i} \frac{m_j}{r_{ij}} + \frac{\gamma |\dot{\boldsymbol{r}}|^2}{c^2} + \frac{(1 + \gamma)|\dot{\boldsymbol{r}}_i|^2}{c^2} - \frac{2(1 + \gamma)}{c^2} \dot{\boldsymbol{r}} \cdot \dot{\boldsymbol{r}}_i$$
$$- \frac{3}{2c^2} \left[\frac{(\boldsymbol{r} - \boldsymbol{r}_i) \cdot \dot{\boldsymbol{r}}_i}{r_i} \right]^2 + \frac{1}{2c^2} (\boldsymbol{r}_i - \boldsymbol{r}) \cdot \ddot{\boldsymbol{r}}_i \right\}$$
$$+ G \sum_i \frac{m_i}{c^2 r_i} \left\{ \frac{3 + 4\gamma}{2} \ddot{\boldsymbol{r}}_i + \frac{\{[\boldsymbol{r} - \boldsymbol{r}_i] \cdot [(2 + 2\gamma)\dot{\boldsymbol{r}}}{r_i^2} \right.$$
$$\left. - \frac{(1 + 2\gamma)\dot{\boldsymbol{r}}_i]\}(\dot{\boldsymbol{r}} - \dot{\boldsymbol{r}}_i)}{r_i^2} \right\}, \tag{17}$$

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). The position and velocity vectors of all celestial bodies are computed with NASA's SPICE library[2] (Acton et al. 2018). We used the planetary and lunar ephemeris DE432s. The N bodies include the Sun, the planets, and the Moon. For planets with moons, with the exception of the Earth, the centre of mass of the system is considered. The dynamical model is written in the J2000 ecliptic reference frame.
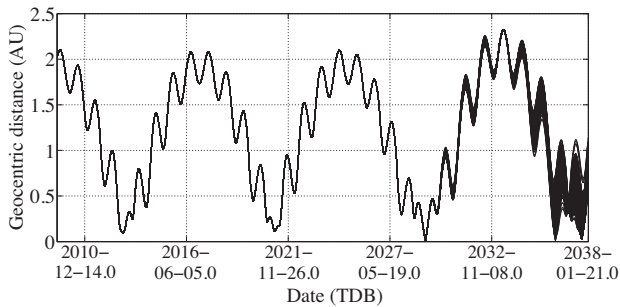
Fig. 6 shows the geocentric distance profile in time for 1000 samples from the initial Gaussian distribution. As expected, the uncertainties significantly increase after 2029 and pave the way to resonant returns in 2036 and 2037.

---

[1] http://newton.dm.unipi.it/neodys/
[2] http://naif.jpl.nasa.gov/naif/toolkit.html
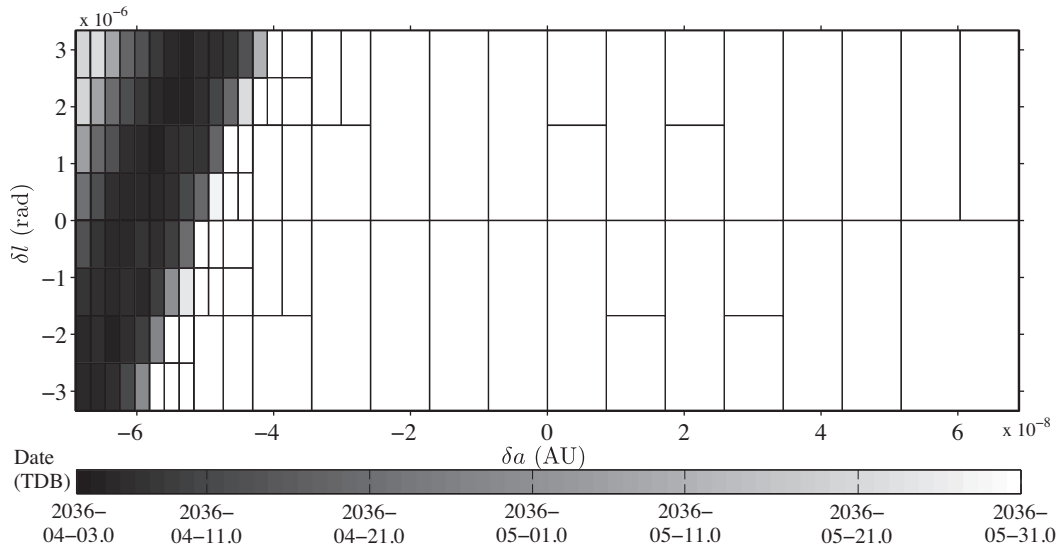
Figure 5. Projection of the generated subsets onto the $a - l$ plane of the initial conditions (ADS propagation, order 8, tolerance 10^{-10}, N_{max} 12, 3σ domain).

Figure 6. Geocentric distance profile up to 2038 January 21 for 1000 samples from the initial uncertainty set.

The authors showed an analysis of the performance of the ADS algorithm for the propagation of the whole uncertainty set up to the second resonant return in Wittig et al. (2015). The results are now limited to the first resonant return, and they will be used as a reference for the assessment of the performance of the ADP.

All the results presented in this section are obtained considering an expansion order equal to 8, a tolerance for the splitting procedure equal to 10^{-10}, a value of N_{max} equal to 12 and an initial uncertainty set with 3σ boundaries, i.e. a 6D rectangle with 3σ boundaries.

The initial uncertainty set should be properly selected, as the neglected part of the probability mass, i.e. the integral of the pdf over the domain outside the considered box, could significantly alter the estimated impact probability. For the case under study, in which we are considering a 6D problem with uncorrelated variables, the selection of a 6D rectangular domain with 3σ boundaries corresponds to considering the 98.4 per cent of the probability mass, and so the estimated impact probability may result underestimated. The accuracy of the estimate improves for larger initial uncertainty sets. A detailed sensitivity analysis on the uncertainty set size and all the other available parameters is offered in Section 7. All computations are performed on a single core Intel i7-3770 CPU @3.4 GHz, 16 GB RAM processor.

The number of subdomains obtained with ADS propagation without pruning is 653, while the computational time is 10 h 6 min. An analysis of the average number of splits per direction shows that most splits occur in the semi major axis (a) and true longitude (l) directions (Wittig et al. 2015). Thus, though the problem is 6D, the analysis on the dynamics can be focused on the projection onto the $a - l$ plane of the initial conditions.

Fig. 5 shows the projection of the initial uncertainty box onto the $a - l$ plane, along with the subdomains generated during the ADS propagation. Colours refer to the truncation epoch of the related subset: white regions represent subsets that were able to reach the final simulation time (2036 May 31, after the expected resonant return), coloured regions represent subsets whose propagation was stopped earlier because they reached their minimum box size. Fig. 5 can be exploited to easily identify the regions of the initial set that are involved in the resonant return in 2036. While all initial conditions lying within white regions have no risk to impact the Earth, coloured subdomains represent sets of initial conditions that might lead to close encounters with Earth at that epoch. That is, coloured regions represent PHSs. This behaviour is expected, as splits occur when the non-linearities increase, which happens when trajectories get closer to Earth. It is evident, however, that a significant portion of the computational effort required by the ADS propagation is spent on regions of the initial set that are not involved in the first resonant return. Thus, the application of a selective pruning technique as the ADP aims at alleviating this inefficiency.

We now investigate the performance of the ADP method. The first part of the analysis is represented by the propagation of the uncertainty set up to the epoch of the first close encounter in 2029. The DA propagation of the whole uncertainty set up to the close encounter in 2029 is performed with no splits. Therefore, the whole set can be described with a single polynomial map at the epoch of the first close encounter. The availability of the DA state vector of the asteroid, then, provides the polynomial expansion of its perturbed orbital period immediately after the close encounter with the Earth. This polynomial expansion allows us to estimate the asteroid orbital period range after the close encounter by means of a polynomial bounder: for the case under study, this range is equal to [415.02, 428.91] d. By looking at this range, we can identify the first resonances: $T_{res1} = 7/6\, T_\oplus = 426.12$ d (where $T_\oplus$ is the Earth orbital period) is the first resonant orbital period included in the computed range, and it represents a resonant return in 2036. This
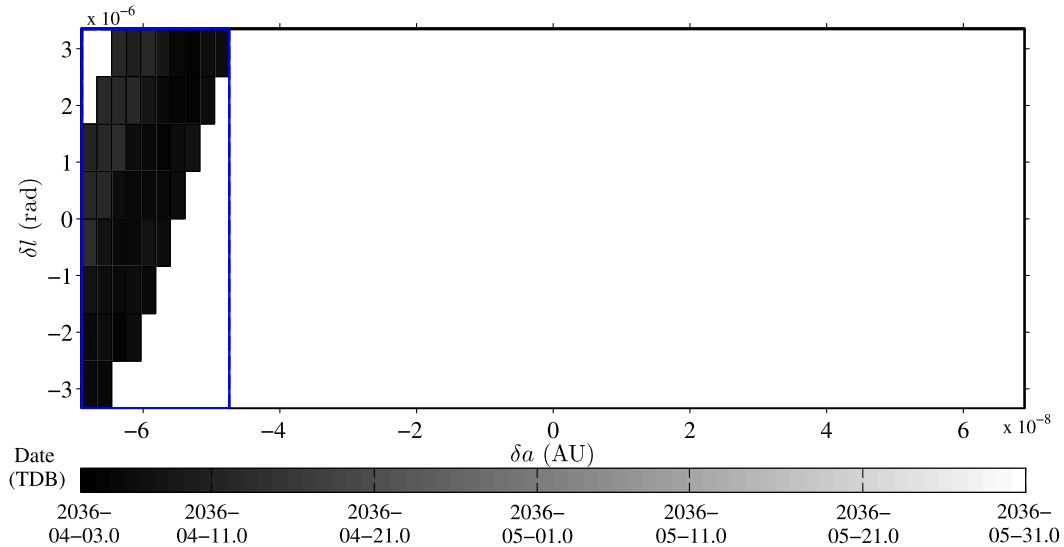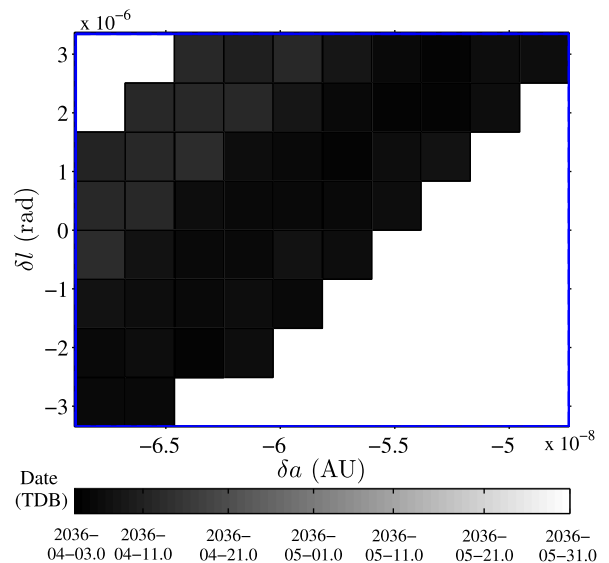
Figure 7. Projection of the generated subsets onto the $a - l$ plane of the initial conditions (ADP propagation, order 8, tolerance 10^{-10}, N_{max} 12, 3σ domain). In blue, boundaries of the ISD.

value is expected, as shown in Fig. 6. We can also notice that the expected second resonant return (in 2037, resonance 8:7), is also included ($T_{res2} = 417.43$ d).

The a priori identification of the resonances and the application of the ADP–IS method is strictly related to the assumption of no intervening close encounter with other major bodies in between. This assumption is checked immediately after the resonances computation, as explained in Section 7.2. For the case under study, the assumptions are verified. Therefore, we can now concentrate the analysis on the first resonant return, in 2036. Given a nominal value $T' = 7/6\, T_\oplus$, ΔT_{ref} is determined by setting a value of ε equal to 10^{-3}. The value of ε is selected in order to take into account small perturbations between the close encounter in 2029 and the resonant return in 2036. An analysis of the impact of ε on the results is carried out in Section 7.2. The propagation is then resumed as described in Section 5. Fig. 7 shows the results of the ADP propagation in terms of subdomains distribution on the $a - l$ plane. A comparison with Fig. 5 clearly shows how the ADP restricts the propagation of the generated subdomains to a limited portion of the PHSs. That is, only subsets that are actually involved in the resonant return in 2036 are propagated till the end of the simulation.

The pattern of subdomains is not altered by the introduction of the pruning. Simply, a large portion of the initial set is no longer investigated. This action has a strong impact on the number of propagated subdomains, that is now significantly lower (267). Consequently, the computational time required by the propagation reduces significantly (4 h 6 min).

The pattern of generated subdomains represents the starting point for the second phase, the application of the IS method for the computation of the impact probability in 2036. We initialize the method by defining a uniform pdf including all the generated PHSs as ISD. The boundaries of the ISD on the $a - l$ plane are represented in blue in Fig. 7. Then, samples are drawn from the ISD and each sample is associated with a PHS if possible. For samples belonging to the PHSs, the state vector corresponding to the drawn sample at the truncation epoch of the related PHS is reconstructed, and a point-wise propagation up to the epoch of minimum geocentric distance is performed. Fig. 8 shows a detail of the resulting subsets, whereas Fig. 9 shows the pattern of generated samples projected onto the
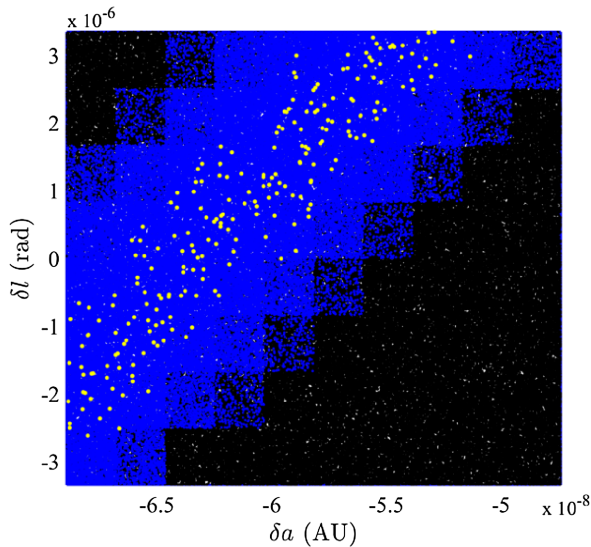

**Figure 8.** Projection of the generated subdomains onto the $a - l$ plane (detail of Fig. 7).

$a - l$ plane. Samples belonging to the PHSs are represented in blue, whereas impacting samples are represented in yellow. Black dots represent discarded samples. Not all samples belong to the PHSs, due to the shape of the selected ISD. A uniform ISD over a domain of regular shape enclosing all PHSs represents the easiest choice and can be applied regardless the complexity of the PHSs pattern. On the other side, this selection leads to the black dots shown in Fig. 9. These samples, however, have a minimal impact on the computational effort required by the method, as they are discarded as soon as they are identified.

The selection of the IS method as sampling technique allows us to increase significantly the number of samples lying within the PHSs with respect to a standard Monte Carlo approach, and this advantage is made possible by the pruning action of the ADP propagation. The analysis of the distribution of the impacting samples on the $a - l$ plane, however, shows that these are confined to a limited region

Figure 9. Projection of the generated samples onto the $a - l$ plane. In black, discarded samples. In blue, samples belonging to the PHSs. In yellow, impacting samples.
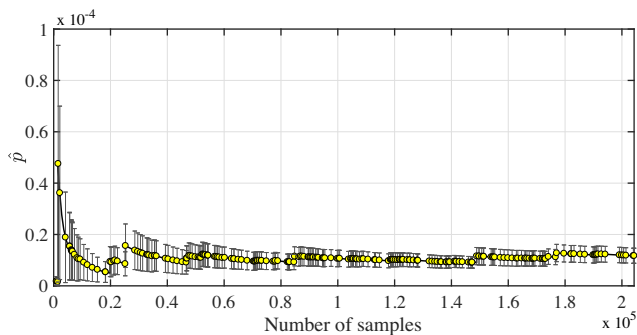
Figure 10. Estimated impact probability for Apophis resonant return 2036 as a function of the number of samples. In yellow, impacting samples. In grey, estimated Poisson statistics uncertainty (1σ).

inside the PHSs. That is, not all PHSs actually give a contribution to the impact probability in 2036. This result is related to the selection of the amplitude of ΔT_{ref}: the value was set in order to grant a conservative pruning action on the subsets. A more detailed analysis is offered in Section 7.2.

The trend of the estimated impact probability with the number of drawn samples is represented in Fig. 10. Impacting samples are represented with yellow circles. The tolerance for the stopping criterion was set equal to 0.01 per cent. After some initial significant oscillations, the impact probability asymptotically converges to the value of 1.17×10^{-5}. The estimate is of the same order of magnitude of the reference value (2.2×10^{-5}) obtained with a standard Monte Carlo analysis, though slightly lower (see Section 8). This difference can be explained considering the size of the propagated uncertainty set, as later explained in Section 7.3.

An overview of the main results of the simulation is shown in Table 2. Results are expressed in terms of number of PHSs n_{PHS}, computational time required by the ADP propagation t_{ADP}, number of generated samples for the IS method at convergence n_{samples}, computational time required by the IS method t_{IS}, overall computational time t_{all}, estimated impact probability value $\hat{p}$, and related Poisson statistics uncertainty $\sigma(\hat{p})$.

7 SENSITIVITY ANALYSIS

The analysis presented in Section 6 was carried out starting from predefined values of expansion order, tolerance for the splitting routine, maximum number of splits and size of the uncertainty set. In this section, we investigate the role of the different parameters and provide some guidelines for the selection of the most appropriate set of parameters. The discussion is carried out dividing parameters mostly affecting the ADP propagation (order, tolerance, minimum box size, and reference orbital period range) and parameters affecting the estimated impact probability (uncertainty box size).

7.1 Selection of splitting tolerance, expansion order and N_{max}

As described in Section 2, the main parameters for the ADS propagation are the tolerance for the splitting procedure, the expansion order and the maximum number of splits. The selection of the tolerance is strictly related to the accuracy required in the description of the subsets at the end of the simulation. This concept is valid in both ADS and ADP propagation. Due to error accumulation during the integration process, indeed, the actual accuracy of the ADP result tends to decrease with respect to the imposed accuracy. This effect becomes more significant as the non-linearities of the dynamics increase, so that, in order to grant a specific accuracy, the imposed tolerance must be in some cases some orders of magnitude lower.

Table 3 shows the average accuracy in position for the subsets at the epoch of the first resonant return considering an expansion order equal to 8 and decreasing values of tolerance. We estimated the accuracy by comparing the results of pointwise propagations and polynomial evaluations for random samples drawn in the generated subsets. The error in position shown in Table 3 represents an average of the computed errors. As expected, there is a difference of around 3 orders of magnitude with respect to the imposed tolerance.

For the case under study, the error is strictly related to the intervening close encounter in 2029. As an example, if we perform an ADS propagation with order 8 and tolerance 10^{-10}, and we stop the propagation 3 months before the close encounter in 2029, we obtain a position error of about 10^{-11} au. This error expands to 10^{-8} au 6 months later, i.e. 3 months after the close encounter. That is, the close encounter yields an increase of about 3 orders of magnitude in the position error. As the propagation continues, the error accumulates and reaches 5.95×10^{-7} au at the epoch of the first resonant return. Therefore, the splitting tolerance is a critical parameter and its selection must account for all the above aspects. For our analysis, we selected a tolerance capable of granting a maximum error in position of 100 km. This requirement results into a splitting tolerance of at least 10^{-10}.

Expansion order and minimum box size, instead, play quite different roles in ADS and ADP propagation. During a DA propagation, a reduction of the expansion order causes a decrease in the accuracy of the results at a specific integration epoch. This decreased accuracy yields an increase in the required number of splits during the ADS propagation and, overall, a larger number of generated subsets. The role of the minimum box size, instead, is to limit the number of splits, so that, overall, both parameters have a strong influence on the number of generated subdomains and, as a consequence, on the required computational effort. The role of the expansion order is twofold, since a decrease in the order causes the number of subdomains to increase, but reduces the computational time required to perform a single integration step. Thus, it is reasonable to imagine that there exists a specific expansion order capable of minimizing the computational effort required by the ADS propagation.

Table 2. Overall results for the ADP–IS method (order 8, tolerance 10^{-10}, N_{max} 12, 3σ domain).

n_{PHS}	t_{ADP}	$n_{samples}$	t_{IS}	t_{all}	$\hat{p}$	$\sigma(\hat{p})$
267	4 h 6 min	204 293	26 min	4 h 32 min	1.17×10^{-5}	2.93×10^{-6}

Table 3. Average error in position as a function of the imposed tolerance for subsets at the epoch of the first resonant return (ADP propagation, order 8, N_{max} 12, 3σ initial uncertainty set).

Tolerance	Error
10^{-8}	4.83×10^{-6} au
10^{-9}	1.24×10^{-6} au
10^{-10}	5.95×10^{-7} au

This value, obviously, changes according to the specific case under study. The role of the minimum box size, instead, is univocal: by increasing the value of N_{max}, the computational effort required by the ADS propagation increases.

In the case of ADP propagation, the analysis is quite different. The role of the two parameters for the two phases is reported in Table 6. More specifically, the ADP propagation aims to select only subsets whose integration stops before the resonant return of interest having reached their minimum box size. A change in the expansion order modifies the splitting history, which could, but not necessarily would, modify the overall number of splits. This behaviour has a direct impact on the required computational time, though the description of the role of the expansion order is not immediate. A decrease in the expansion order, indeed, may cause just earlier splits performed with the same splitting sequence, or a complete change in the splitting history. In the first case, the role of the expansion order becomes univocal: a reduction in the expansion order causes a decrease in the computational effort. In the second case, the changes in the splitting history and the number of generated subsets may be so relevant that what is gained in performing single integration steps may be lost in the longer propagation of the generated subsets. Overall, the role of the order in not univocal, and it exists an order that minimizes the computational time required by the ADP propagation.

The role of the minimum box size, instead, is the same as in the ADS propagation: a decrease in the value of N_{max} causes an earlier stop of the propagation of the subsets, and a reduction of the computational effort.

The whole procedure, however, includes both an ADP propagation and a sampling phase, and the role played by the two parameters during the sampling phase is different from the ADP phase. The role of the expansion order is, again, twofold: a reduction of the order causes longer pointwise propagations, but faster polynomial evaluations. The relative weight of the two effects essentially depends on the number of required samples. The role of the minimum box size, instead, is univocal and opposite with respect to the ADP propagation: a reduction of the value of N_{max} implies longer pointwise propagations.

The selection of the best combination of order and minimum box size, therefore, relies on all these aspects. Starting from these considerations, we performed a sensitivity analysis in order to quantify the impact of the two parameters on the performance of the ADP–IS method for the case of the first resonant return of asteroid Apophis.

The results of a sensitivity analysis on the expansion order are shown in Table 4, considering six different expansion orders. The comparison is performed by considering the same parameters of the analysis presented in Section 6. The second column shows the number of generated subdomains. The value is not affected by the expansion order till order 4, while for order 3 the value is more than doubled. This trend can be explained looking at the splitting history. For orders from 8 to 4, no split occurs before the close encounter in 2029, and the sequence of splits is exactly the same, though single splits are performed at different epochs. Things completely change with order 3, with 4 splits occurring before the 2029 close encounter. This change has a direct impact on the number of generated subsets. A difference can be detected also by looking at the required number of samples or at the estimated impact probability and related Poisson statistics uncertainty. Assuming not to alter the sequence of generated samples, values obtained with orders 8 to 4 are identical, whereas values obtained with order 3 are slightly different.

The expansion order has a significant impact on the computational effort required by the ADP and sampling phases (columns 3 and 5). As expected, the trends are not monotonic. Thus, it is possible to identify order 5 as the expansion order capable of minimizing the overall computational time. Once again, it is interesting to see what happens with order 3: the early splits at the epoch of the 2029 close encounter completely change the splitting history, causing subsets to stop much earlier than what happens with larger orders. Unlike orders from 8 to 4, where the subsets are stopped few days before the expected resonant return in 2036, with order 3 the subsets are stopped around 2030, 6 yr earlier. This earlier stop grants computational time saving for the ADP propagation, but has a tremendous backlash for the sampling phase, with each sample propagated for years instead of days. For this reason, the computational time required for the sampling phase is much larger (column 5). The trend of the computational time for the different orders is shown in Fig. 11.

We limited the analysis in Table 4 to order 3 as the minimum order, as early splits that appear with this order magnify with order 2, leading to 53 subsets generated before the 2029 close encounter. This behaviour exacerbates the limitations previously pointed out for low orders. Moreover, the error estimation procedure described in Section 2 does not work with linear approximation and tends to provide inaccurate estimates with order 2.

We performed a similar analysis by considering the effect of the minimum box size on the required computational effort and estimated impact probability. Table 5 shows the results of the analysis considering the optimal expansion order identified in Table 4 and the same values of tolerance and uncertainty box size of the previous simulations. As described before, the role of the minimum box size is univocal in the two phases, though opposite, and this trend is confirmed by the analysis: a decrease in the value of N_{max} causes a reduced computational effort required by the ADP propagation but longer pointwise propagations for all samples. For the case under study, N_{max} equal to 10 allows us to minimize the required computational effort.

As in the previous case, a change in the value of N_{max} does not alter the estimated impact probability, though the pattern of generated subsets is now modified. This trend can be explained considering the fact that, a reduction of the value of N_{max} generates larger

Table 4. Performance of the ADP–IS method for different expansion orders (tolerance 10^{-10}, N_{max} 12, 3σ initial uncertainty set).

Order	n_{PHS}	t_{ADP}	$n_{samples}$	t_{IS}	t_{all}	$\hat{p}$	$\sigma(\hat{p})$
8	267	4 h 6 min	204 293	26 min	4 h 32 min	1.17×10^{-5}	2.93×10^{-6}
7	267	1 h 53 min	204 293	24 min	2 h 17 min	1.17×10^{-5}	2.93×10^{-6}
6	267	49 min	204 293	23 min	1 h 12 min	1.17×10^{-5}	2.93×10^{-6}
5	267	24 min	204 293	23 min	47 min	1.17×10^{-5}	2.93×10^{-6}
4	267	28 min	204 293	25 min	53 min	1.17×10^{-5}	2.93×10^{-6}
3	589	8 min	273 035	9 h 47 min	9 h 55 min	8.60×10^{-6}	2.20×10^{-6}

Table 5. Performance of the ADP–IS method for different values of N_{max} (order 5, tolerance 10^{-10}, 3σ initial uncertainty set).

N_{max}	n_{PHS}	t_{ADP}	$n_{samples}$	t_{IS}	t_{all}	$\hat{p}$	$\sigma(\hat{p})$
12	267	24 min	204 293	23 min	47 min	1.17×10^{-5}	2.93×10^{-6}
11	148	17 min	204 293	26 min	43 min	1.17×10^{-5}	2.93×10^{-6}
10	84	12 min	204 293	29 min	41 min	1.17×10^{-5}	2.93×10^{-6}
9	47	9 min	204 293	34 min	43 min	1.17×10^{-5}	2.93×10^{-6}

Table 6. Trend of the computational times t_{ADP} and t_{IS} for increasing ($\uparrow$) values of the expansion order and minimum box size.

	ADP	IS
Order $\uparrow$	$t_{ADP}\uparrow\downarrow$	$t_{IS}\uparrow\downarrow$
$N_{max} \uparrow$	$t_{ADP}\uparrow$	$t_{IS}\downarrow$
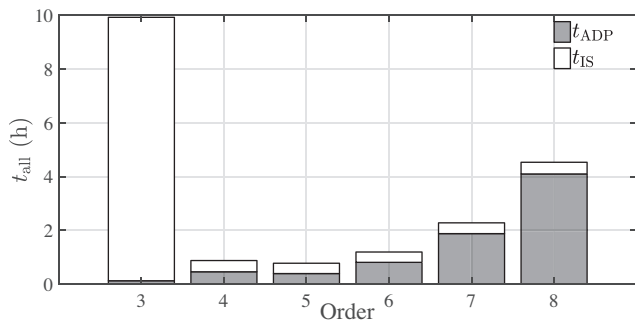
Figure 11. Computational time versus expansion order (ADP–IS method, tolerance 10^{-10}, N_{max} 12, 3σ uncertainty domain). Grey bars and white bars represent computational times required by the ADP propagation and IS phase, respectively.

subsets at earlier truncation epochs, but with the same accuracy. Thus, if the drawn samples are fixed, their mapping to the epoch of the first resonant return is essentially the same. That is, the pattern of impacting samples is not altered.

As described in the presented analysis, both expansion order and minimum box size influence the performance of the method. In particular, the expansion order plays a key role in the definition of the computational effort required by the method, while the minimum box size has a lower influence. As the method is composed by two phases, we can say that the selection of the order must be done in order to minimize the computational effort required by the heaviest one. In our method, the ADP propagation plays this role, so that, in order to decrease its impact on the overall computational time, the most effective way is to reduce the expansion order, still limiting as much as possible the number of generated subsets before the first close encounter.

7.2 Definition of ΔT_{ref}

As described in Section 3, the reference orbital period range ΔT_{ref} represents the key parameter for the ADP propagation. The range, centred in the selected resonant return period T', is defined to account both inaccuracies in the estimation of the orbital period range of the subdomains and small dynamical perturbations between the first close encounter and the predicted resonant return. The semi-amplitude of this range, ε, plays therefore a key role as it influences both the accuracy of the probability estimates and the required computational time.

Table 7 shows the performance of the ADP–IS method for different values of the reference range semi-amplitude ε. With respect to the previous analyses, two additional parameters are shown: the number of generated subdomains n_D and the number of subdomains that include impacting samples n_{imp}. These two parameters provide, along with the number of PHSs n_{PHS}, a clear picture of the pruning action performed during the ADP propagation. We performed the analysis considering four values of ε. In particular, it is interesting to analyse what happens considering the two limiting cases: $\varepsilon = 0$ and $\varepsilon = 1$.

In the first case, the reference orbital period range collapses to the value of T', i.e. subsets are maintained throughout the simulation only if their estimated orbital period range includes T'. In this case, the number of PHSs is lower than the one obtained with $\varepsilon = 10^{-3}$, but the value of impact probability is exactly the same. This result can be explained considering that the pattern of impacting samples is not altered, as confirmed by the parameter n_{imp}. That is, for the case under study, a less conservative selection of the parameter ε would allow us to obtain the same results, though no evident savings in computational time would be obtained. The selection of $\varepsilon = 0$ corresponds to considering a Keplerian motion between the two encounters. As described in Valsecchi et al. (2003), this assumption may provide quite accurate results for the timing, and for the case under study, where no significant perturbations between the two encounters exist, it can be considered acceptable.

Let us now analyse the second limit case. If ε is set to 1, we are essentially selecting a very large reference range for the orbital period. That is, the ADP propagation becomes an ADS propagation, i.e. no pruning is performed and all subsets are propagated until the final simulation time or the maximum number of splits is reached. The sampling phase, instead, is not significantly altered: the ISD is defined including subsets whose propagation stops before

Table 7. Performance of the ADP–IS method for different values of ε (order 5, tolerance 10^{-10}, $N_{\max}$ 10, 3σ initial uncertainty set).

n_D	n_{PHS}	n_{imp}	t_{ADP}	$n_{samples}$	t_{IS}	t_{all}	$\hat{p}$	$\sigma(\hat{p})$
0	71	71	43	11 min	204 293	29 min	40 min	1.17×10^{-5}	2.93×10^{-6}
10^{-5}	71	71	43	11 min	204 293	29 min	40 min	1.17×10^{-5}	2.93×10^{-6}
10^{-3}	84	84	43	12 min	204 293	29 min	41 min	1.17×10^{-5}	2.93×10^{-6}
1 (ADS)	470	121	41	1 h 31 min	216 299	31 min	2 h 02 min	1.11×10^{-5}	2.25×10^{-6}

the expected resonant return. Results are reported in the last row of Table 7. The value of impact probability is similar to the one obtained with the pruning action, but the number of generated subsets is much larger, which affects in turn the required computational time. That is, a very conservative selection of ε would yield a factor 3 increase in computational time.

The selection of the parameter ε is therefore crucial. Within the assumptions of our method, i.e. small perturbations between the two encounters, we can define an upper threshold for the ε value as

$$h\varepsilon T' v_{\oplus} < 0.05 \,\text{au}, \tag{18}$$

where h is the number of revolutions of the asteroid between the encounters, whereas $v_{\oplus}$ is the Earth heliocentric velocity. The expression on the left-hand side of the inequality represents the heliocentric arc covered by the Earth in the time range $h\varepsilon T'$. When we define the reference range ΔT_{ref} and we compare it with ΔT_{sub} of a given subset, therefore, we are verifying that the uncertainty in the position of the current subset with respect to the Earth position is lower than 0.05 au, i.e. the current subset can be labelled as potentially hazardous. For the case under study, this value is about 10^{-3}, the value selected for the analysis presented in the previous sections.

As previously stated in the paper, the application of the ADP–IS method is strictly related to the assumption of no intervening close approaches with other major bodies in the investigated time window. In case of expected close approaches, indeed, the situation drastically changes as the resonances estimated at the epoch of the first close encounter may lose their validity. In such cases, one must rely on the more conservative approach of ADS propagation for investigating a selected propagation window, obtaining accurate results with unavoidable drawbacks in efficiency.

The decision whether to perform an ADP propagation or disable pruning is made based on a preliminary analysis of the possible trajectories of the asteroid between the two encounters. For the case under study, the availability of the dispersion of Apophis' orbital parameters after the first close encounter allows us to estimate the minimum orbit intersection distance (MOID) dispersion between the asteroid and the other main bodies of the Solar System (see Armellin et al. 2010a). This fast survey provides us with an overview of possible close approaches between the two encounters and drives our decision on the propagation method. For the case under study, the analysis required less than 1 min and allowed us to exclude any significant encounter with other planets in between.

7.3 Effect of the size of the uncertainty domain

The analysis presented in the previous sections was done considering different values of expansion order, tolerance for the splitting procedure and minimum box size, whereas we considered only one size for the initial uncertainty set, that is a 3σ 6D rectangle. This selection, in a 6D problem with uncorrelated variables, consists in considering the 98.4 per cent of the probability mass. In the following paragraphs, we present the impact of the size of the uncertainty

set on the results of the ADP–IS method. It is worth noting that, as previously mentioned in Section 6, in case of full covariance matrix, the initial uncertainty box would be defined in the eigenvector space in order to avoid wrapping effect and including very low probability solutions, so all the analyses and values presented in this section hold for the more general case of correlated variables.

The ADP propagation and the IS phase are strongly influenced by the selection of the size of the initial uncertainty set. The ADP propagation, indeed, limits the generation of the subsets within the boundaries of the considered uncertainty set, and this aspect influences also the shape of the ISD for the impact probability computation phase. Samples, indeed, are confined within the initial uncertainty set, impacting samples are found only within these limits and possible impacting samples, that lie out of the initial uncertainty set, are discarded. This aspect distinguishes our sampling approach from a standard Monte Carlo method, where samples are drawn directly from the original probability density function, and the probability of drawing samples is determined by the pdf itself. In principle, samples could lie anywhere in the uncertainty region.

Starting from these considerations, it is therefore interesting to study how the estimate for the impact probability changes with an increasing size of the initial uncertainty set. We initially performed the analysis by considering a size for the initial uncertainty set of 3σ (i.e. what was previously presented), 4σ, and 5σ.

Fig. 12(a) shows the results for the ADP propagation considering order 5, tolerance 10^{-10}, $N_{\max}$ equal to 10 and an initial uncertainty set size of 4σ. The ISD boundaries on the $a - l$ plane are represented in light blue. On the same plot, the 3σ boundaries are represented with black dashed lines, whereas the ISD boundaries for the 3σ case are represented with dashed blue lines. The plot allows us to compare the sample regions in the two cases. In particular, a portion of the PHSs generated during the 4σ ADP propagation is not considered during the 3σ case. The 5σ case is shown in Fig. 12(b), representing the 5σ ISD in cyan.

Fig. 13 shows a comparison of the results of the sampling phase for the 3σ, 4σ, and 5σ cases. Blue points represent drawn samples belonging to the 3σ case, with impacting samples represented as yellow dots. Light blue points represent drawn samples belonging to the 4σ case, with impacting samples represented as black dots. Finally, cyan points represent accepted samples for the 5σ case, with impacting samples represented with red dots. The analysis of the plot offers a clear picture of how the sampling region changes in the three cases. Moreover, it is possible to see how the increasing size of the initial uncertainty set allows us to include impacting samples out of the 3σ domain. While a 3σ domain appears as a too narrow selection, the 4σ and 5σ domains offer a better description of the impact region. Figs 14(a) and (b) show the distribution of the impacting samples for the three different simulations, with colours showing the contribution to the overall impact probability.

Table 8 shows the results of the analysis, including the 6σ case. With reference to the previous analyses, we added the parameter p_{out}, which represents the probability mass outside the selected uncertainty set, i.e. the complementary to 1 of the integral of the pdf
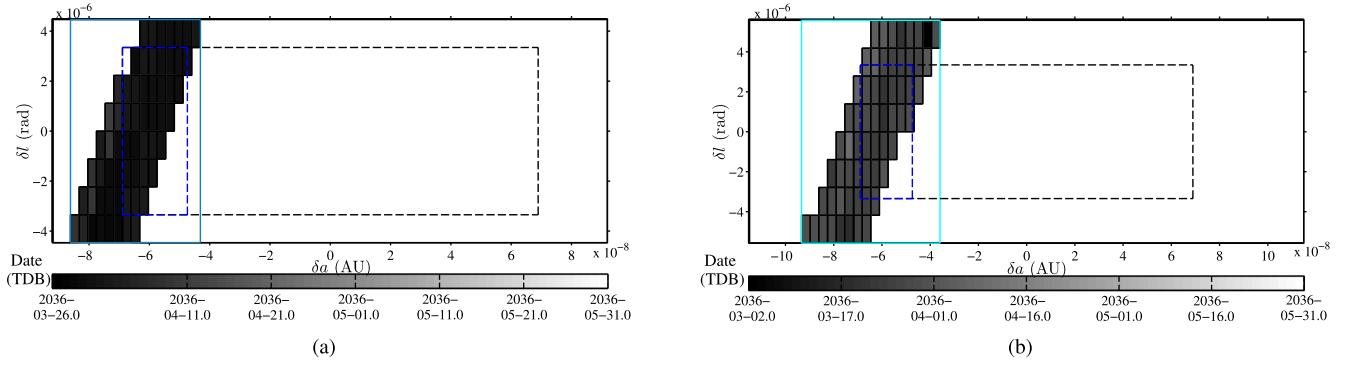
Figure 12. Comparison between 4σ (a) and 5σ (b) cases. The figure box and the dashed black lines represent the selected boundaries [4σ in (a) and 5σ in (b)] and the 3σ boundaries, respectively. In light blue and cyan, the ISD boundaries for the 4σ and 5σ, respectively. In dashed blue, the 3σ ISD boundaries.
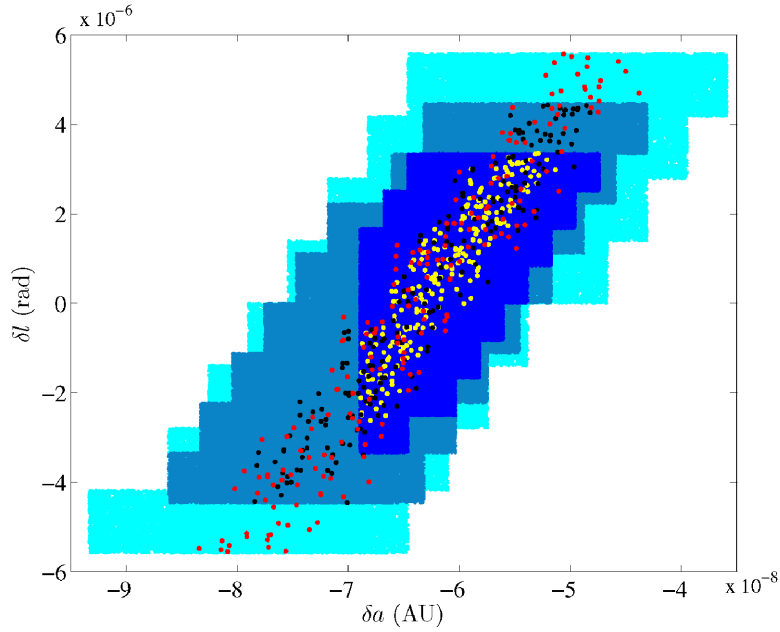
Figure 13. Comparison of the samples distribution for the 3σ, 4σ, and 5σ cases. In blue, light blue, and cyan, drawn samples for the 3σ, 4σ, and 5σ case, respectively. In yellow, black, and red, impacting samples for the 3σ, 4σ, and 5σ case, respectively.
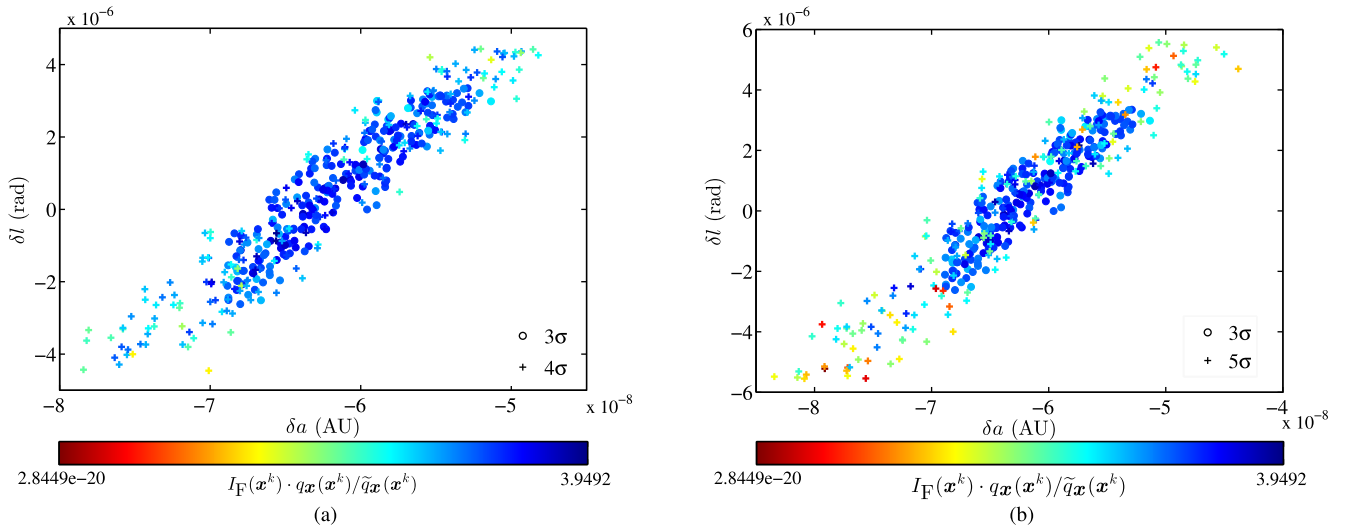
Figure 14. Impacting samples distribution for 4σ (a) and 5σ (b) cases. Colours are referred to the associated contribution to the impact probability.

Table 8. Performance of the ADP–IS method for different values of initial uncertainty domain size (order 5, tolerance 10^{-10}, N_{max} equal to 10).

Domain	p_{out}	n_{PHS}	t_{ADP}	$n_{samples}$	t_{IS}	t_{all}	$\hat{p}$	$\sigma(\hat{p})$
3σ	1.61×10^{-2}	84	12 min	204 293	29 min	41 min	1.17×10^{-5}	2.93×10^{-6}
4σ	3.80×10^{-4}	107	21 min	300 643	39 min	1 h	1.61×10^{-5}	5.70×10^{-6}
5σ	3.44×10^{-6}	104	29 min	341 804	41 min	1 h 10 min	1.90×10^{-5}	6.81×10^{-6}
6σ	1.18×10^{-8}	107	46 min	353 056	42 min	1 h 28 min	2.09×10^{-5}	7.00×10^{-6}

over the considered domain. We remark that, because we use rectangular uncertainty sets, the values of p_{out} shown in Table 8 are significantly smaller than those corresponding to the more commonly used ellipsoidal uncertainty regions, for which p_{out} is equal to 0.17, 1.38×10^{-2}, 3.41×10^{-4}, and 2.76×10^{-6} for the 3σ, 4σ, 5σ, and 6σ cases, respectively.

By increasing the size of the initial uncertainty set, the computational time required by the ADP propagation increases. This trend is essentially due to the larger computational effort required by a single integration step and the larger number of subsets. An increase in the number of generated PHSs can be detected passing from 3σ to 4σ, whereas this value remains essentially the same in the 5σ and 6σ cases.

The analysis of the sampling phase shows some interesting results. As expected, an increase in the initial uncertainty size causes an increase in the estimated impact probability value. Essentially, regions of the uncertainty set that were not studied during the ADP propagation for the 3σ case are now considered, and impacting samples can be found also in these regions. As a result, the estimated impact probability values for the 5σ and 6σ cases become very close to the reference value. The enlargement of the investigated region causes also an increase in the Poisson statistics uncertainty of the estimate. This result is expected too, as the variance is proportional to the sample region volume (see equation (13)). The analysis of the required number of samples at convergence shows that this value increases for larger initial uncertainty sets, and this trend reflects back on the computational time required by the sampling phase.

The size of the uncertainty set should be selected to achieve the desired resolution on the impact probability, which is directly expressed by the parameter p_{out}. For the case under study, with an estimated impact probability of the order of 10^{-5}, we selected a 6σ domain, which excludes only 1.18×10^{-8} of the probability mass.

8 COMPARISON WITH STANDARD AND ADVANCED ORBITAL SAMPLING TECHNIQUES

The analysis presented in the previous sections showed how the ADP–IS method represents a valuable tool for uncertainty propagation and impact probability computation for the first resonant return of a NEO. In order to assess the efficiency of the method with respect to other impact probability computation tools, we present in this section a comparison with standard and advanced orbital sampling techniques. In the first part, we compare our approach with Monte Carlo sampling techniques based on sample generation on the whole uncertainty set. Finally, we present a general comparison with the most used technique for impact probability calculation, the LOV method.

A first comparison can be done considering a standard Monte Carlo approach, where samples are drawn from the covariance matrix directly at the initial epoch (2009 June 18). This method is probably the most straightforward approach but also the most expensive one, as the sampling is performed on the whole domain,

and the propagation of each sample starts from the observation epoch. By performing the propagation of one million samples, the estimated impact probability results into 2.2×10^{-5}, whereas the Poisson statistics uncertainty is equal to 4.71×10^{-6}. We selected the number of samples in order to detect a non null value of impact probability (Farnocchia et al. 2015).

Unfortunately, if the Monte Carlo simulation is performed considering the same conditions of our method (i.e same dynamics, single core), the required computational time is much larger. The average computational time required to perform a single pointwise propagation from the initial epoch to the epoch of the first resonant return is ≈1.2 s. As a result, within the computation time required by the ADP–IS method for the 6σ case (see Table 8), about 4400 samples could be propagated, which is not enough to estimate the expected impact probability. All this would lead to an estimated computational time of around 2 weeks for propagating one million samples on a single core. This value is of course not realistic, as typically Monte Carlo analyses can be easily set up in a multithread environment, thus granting significant savings in computational time. It is interesting, however, to highlight the significant savings that our approach grants with respect to standard MC approach in the same conditions. The ADP–IS method, indeed, employs a lower number of samples, as samples are drawn just in a subset of the uncertainty set. Moreover, the propagation of all samples starts immediately before the resonant return, while in a standard Monte Carlo approach each sample is propagated starting from 2009 June 18. Therefore, the computational effort required by the ADP propagation is largely repaid later by shorter pointwise propagations and a reduced number of samples.

We show now a comparison with an advanced Monte Carlo technique called subset simulation (SS). The basic idea of SS is to compute small failure probabilities as the product of larger conditional probabilities (Au & Beck 2001; Cadini et al. 2012; Zuev et al. 2012). Given a target failure event F, let $F_1 \supset F_2 \supset ... \supset F_n = F$ be a sequence of intermediate failure events, so that $F_k = \bigcap_{i=1}^{k} F_i$, $k = 1, 2, ..., n$. Considering a sequence of conditional probabilities, then the failure probability becomes:

$$p(F) = p(F_1) \prod_{i=1}^{n-1} p(F_{i+1}|F_i), \tag{19}$$

where $P(F_{i+1}|F_i)$ represents the probability of F_{i+1} conditional to F_i. A detailed description of the algorithm can be found in Au & Beck (2001). In the problem under study, the failure F represents an impact with Earth, i.e. a geocentric distance smaller than the Earth radius. The method is initialized using standard MC to generate samples at the so-called conditional level (CL) 0 starting from the available nominal state vector and related uncertainty of the investigated object at the observation epoch. The number of samples generated at this level is maintained for each generated conditional level and it is referred to as N. Once the failure region F_1 is identified, a Monte Carlo Markov Chain (MCMC) Metropolis Hastings algorithm is used to generate conditional samples in the identified

intermediate failure region. Another intermediate failure region is then located, and other samples are generated by means of MCMC. The procedure is repeated until the target failure region is identified. An illustration of the method is shown in Fig. 15.

The approach was originally developed for the identification of structural failures, but it was also used in different research areas in reliability such as the definition of failure probabilities of thermo-hydraulic passive systems. The method was recently applied to the computation of space debris collisional probabilities by Morselli et al. (2014).

In the presented approach, the intermediate failure regions are identified by assuming a fixed value of conditional probability $p_0 = p(F_{i+1}|F_i)$. The identification of each conditional level, therefore, is strictly related to this value, and changes accordingly step by step, as explained in the followings. The resulting SS algorithm follows the general description presented in Morselli et al. (2014) and goes through the following steps:

(i) Set $i = 0$ and generate N samples $\boldsymbol{x}_0^{0,k}$, $k = 1, \ldots, N$ at conditional level 0 by standard MC starting from the available state estimate of the investigated object at the initial epoch t_0.

(ii) Propagate each sample up to the epoch of the first resonant return and compute its minimum geocentric distance. Note that, as in the ADP–IS method, the resonances can be easily determined by propagating the uncertainty set up to the epoch of the first close encounter by means of DA and evaluating the orbital period range.

(iii) Sort the N samples in descending order according to the associated geocentric distance at the epoch of the first resonant return.

(iv) Identify an intermediate threshold value D_{i+1} as the geocentric distance corresponding to the $(1 - p_0)N$th element of the sample list. Define the $(i + 1)$th conditional level as $F_{i+1} = \{d < D_{i+1}\}$, where d represents the geocentric distance. According to the definition of D_{i+1}, the associated conditional probability $p(F_{i+1}|F_i) = p_0$.

(v) If $D_{i+1} < R_\oplus$, i.e. the geocentric threshold distance is lower than the Earth radius, go the the last step, otherwise select the last samples of the list $\boldsymbol{x}_0^{i,j}$, $j = 1, \ldots, p_0N$. By definition, these samples belong to the $(i + 1)$th conditional level.

(vi) Using MCMC, generate $(1 - p_0)N$ additional conditional samples starting from the previously selected seeds belonging to F_{i+1}. A sample is set to belong to F_{i+1} according to the following performance function:

$$g_{\boldsymbol{x}}^{i+1}(\boldsymbol{x}_0) = d(\boldsymbol{x}_0) - D_{i+1} \begin{cases} > 0 & \boldsymbol{x}_0 \text{ is out of the } (i+1)\text{th CL} \\ = 0 & \boldsymbol{x}_0 \text{ is at the limit of the CL} \\ < 0 & \boldsymbol{x}_0 \text{ belongs to the } (i+1)\text{th CL} \end{cases} \quad (20)$$

(vii) Set $i = i + 1$ and return to step (ii).
(viii) Stop the algorithm.

The total number of generated samples is

$$N_{\mathrm{T}} = N + (n - 1)(1 - p_0)N, \quad (21)$$

where n is the overall number of conditional levels required to reach the impact region. Since the conditional probability is equal to p_0 for each level, the impact probability expressed by equation (19) becomes:

$$p(F) = p(F_n) = p(F_n|F_{n-1})p_0^{n-1} = p_0^{n-1}N_n/N \quad (22)$$

where N_n is the number of samples belonging to the last conditional level whose geocentric distance is lower than the Earth radius.

The main degrees of freedom of the method are the selected fixed conditional probability p_0, the number of samples per conditional

level N and the proposal auxiliary distribution for the MCMC phase, and they govern the accuracy and efficiency of the method (Zuev et al. 2012). We used for our analysis 1000 samples per conditional level and a value of conditional probability equal to 0.1. A normal distribution with spread equal to the original pdf was selected as proposal pdf for the MCMC algorithm.

A comparison between the SS technique and the ADP–IS method is shown in Table 9. The required number of samples, the overall computational time, the estimated impact probability and related Poisson statistics uncertainty are shown. Results for the ADP–IS method are the ones referring to the 6σ case.

SS and ADP–IS have a similar computational burden, though the required number of samples is very different. This result is expected, as the propagation windows for the two cases are different. Fig. 16 shows the distribution on the $a - l$ plane of the generated conditional samples obtained with SS, along with the thresholds per conditional level and related colours. Impacting samples at the last conditional level are represented in black. Conditional samples progressively move to the left, until impacting samples at conditional level 4 are identified. If compared to Figs 8 and 9, this region is practically coincident with the PHSs identified during the ADP propagation. That is, SS and ADP–IS allow us to identify the same region in two completely independent ways.

The advantage of the ADP–IS method is that, by identifying the PHSs, the propagation of the samples is drastically reduced in time, which yields a similar computational burden though the number of generated samples is significantly larger. Potentially, the ADP–IS method could take advantage of parallelization both during ADP propagation and the sampling phase, while the advantages for SS would be lower, as parallelization could be introduced only for specific phases of the algorithm. This approach would heighten the difference in efficiency between the two methods. However, the great savings granted by the SS could be included in the ADP–IS method during the sampling phase, by replacing the standard MC performed in the ISD with an SS limited to the unpruned subsets. This aspect may represent a future development of the method. Overall, the combination of ADP propagation and importance sampling allows us to achieve a computational burden that is competitive with both standard and advanced Monte Carlo techniques.

Finally, it is worth comparing the performance of the presented approach with the reference technique in the field of impact probability computation, the LOV method. The LOV method takes advantage of the fact that the orbital uncertainty grows with time by stretching into a long slender ellipsoid in Cartesian space (Farnocchia et al. 2015). The tendency of uncertainty to stretch during propagation suggests the possibility of a 1D parametrization of the uncertainty region, i.e. the sampling and the generation of the so-called virtual asteroids (VAs) is performed along the line of weakness of the orbit determination, and if all orbits are sufficiently close to the LOV, then significant savings in computational time with respect to a standard Monte Carlo approach are obtained without sacrificing reliability.

The analysis presented in Milani et al. (2005) offers a first term of comparison: the generic completion level of 10^{-7} can be obtained with the propagation of only $\sim 10^4$ VAs. If compared to a standard MC approach, it would lead to compute times 3–4 orders of magnitude below those required for similar completeness with MC simulations (Farnocchia et al. 2015). The analysis presented in the previous section showed that the ADP–IS method grants a reduction in computation burden of around two orders of magnitude with respect to standard MC. Therefore, the LOV shows better performance than the ADS–IS method in the current implementation.
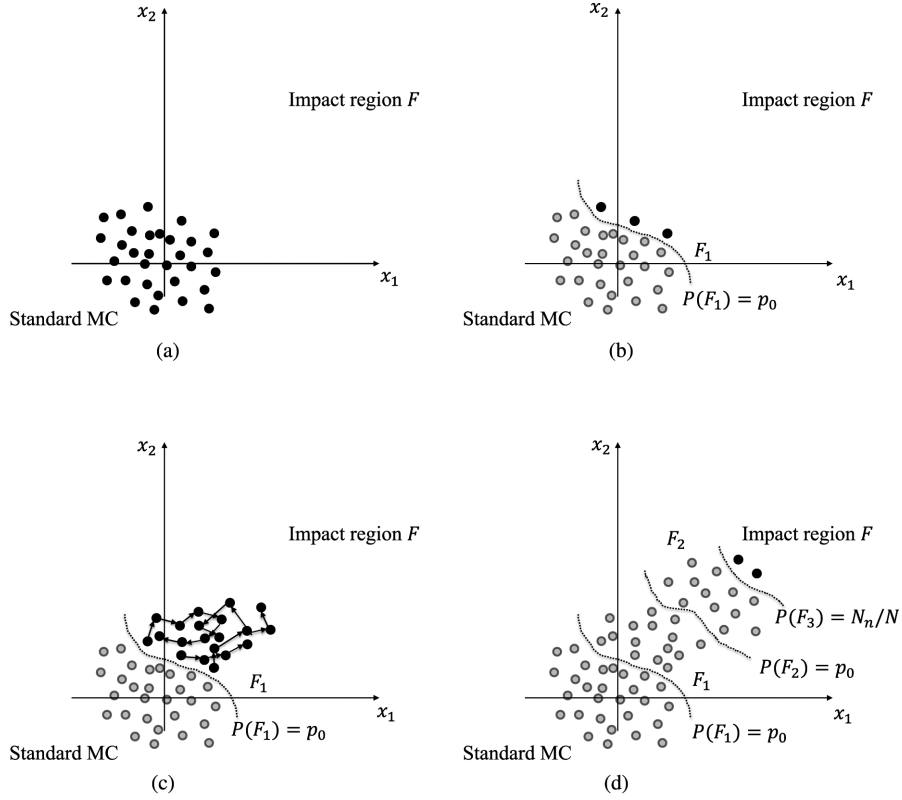
Figure 15. Subset Simulation process: (a) initialization by standard MC, (b) CL 1 identification, (c) samples generation by means of MCMC, and (d) new iterations and impact region identification.

Table 9. Comparison of the performance of ADP–IS method (order 5, tolerance 10^{-10}, N_{max} equal to 10, 6σ uncertainty set) and SS method (N equal to 1000, p_0 equal to 0.1).

	$n_{samples}$	t_{all}	$\hat{p}$	$\sigma(\hat{p})$
ADP–IS	353 056	1 h 28 min	2.09×10^{-5}	7.00×10^{-6}
SS	4 600	1 h 30 min	2.46×10^{-5}	5.04×10^{-6}

Nevertheless, there are some cases in which the LOV method does not guarantee the same level of accuracy of a standard MC approach. A first case occurs when the observed arc of the investigated object is very short, i.e. 1 or 2 d (Milani et al. 2005). In this case, the confidence region is wide in two directions and the unidimensional sampling may not be suitable. What happens is that different LOVs, computed with different coordinates, provide independent sampling and may provide different results. That is, if some impacting samples lie well of the LOV and are separated
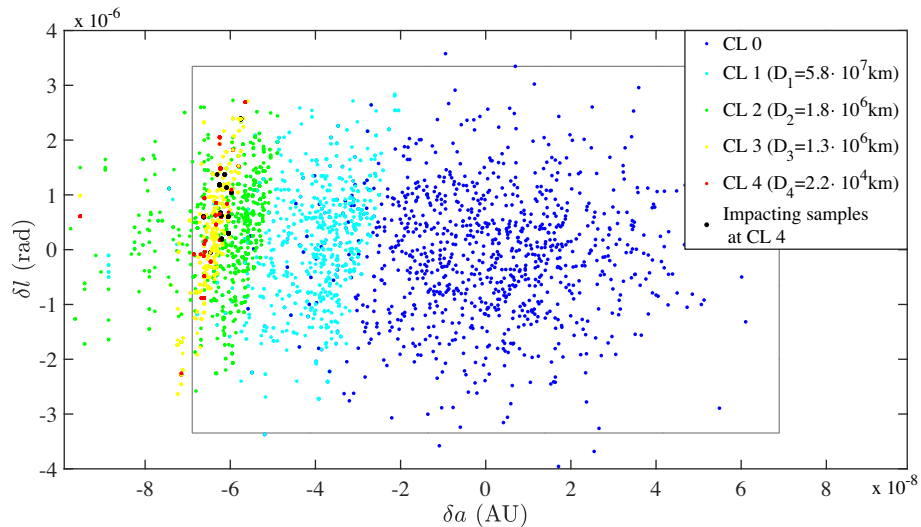
Figure 16. SS conditional samples projected onto the $a - l$ plane (N equal to 1000, p_0 equal to 0.1). In grey, boundaries of the 3σ domain.

from it by some strong non-linearity, then the VAs selected along the LOV may fail to indicate some potential threatening encounters (Milani et al. 2002). In such cases, a standard MC approach would result more reliable, with unavoidable drawbacks in terms of computational time. As presented in this paper, the ADP–IS method, though maintaining a 6D sampling, allows us to drastically reduce the computational effort by limiting the sampling to just specific regions. For these reasons, the method may be considered as a valuable trade-off between the efficiency of the LOV method and the reliability of standard MC in all those cases in which the former may result inaccurate. The possibility of improving the efficiency of the method by means of parallelization in both ADP propagation and sampling phases represents another step in this direction, as well as an optimized coding of the dynamics.

9 CONCLUSIONS

This paper introduced the combination of automatic domain pruning and importance sampling for uncertainty propagation and impact probability computation for Earth resonant returns of NEOs. The automatic domain pruning represents an evolution of the DA-based automatic domain splitting technique, it allows us to estimate possible resonances after a planetary close encounter and limit the propagation of an uncertainty set to those subsets that may be involved in the resonant return of interest. During the propagation, the uncertainty domain is divided into subsets (PHSs) whose propagation stops just before the epoch of the resonant return. The identification of PHSs represents the starting point for the sampling phase. An importance sampling pdf is defined over these subdomains and samples are drawn directly from this auxiliary pdf. We tested the ADP–IS method on the case of asteroid (99942) Apophis, providing an estimate for the impact probability in 2036. We carried out a sensitivity analysis on the main parameters of the method, providing general guidelines for their selection. The comparison with a standard Monte Carlo approach showed how the ADP–IS method can reduce the computation effort by more than 2 orders of magnitude, still granting the same accuracy level for the impact probability estimate. In addition, the current algorithm can be implemented to make use of parallelization techniques in both the ADP and the IS phase, thus significantly reduce the required computational time. All these considerations suggest that the method may be used as a valuable alternative to standard MC in all those cases in which the LOV method does not guarantee the required level of accuracy. Future developments include a more rigorous formulation of the reference orbital period for subsets pruning allowing us to extend the pruning algorithm to the more critical case of intervening close encounters with other celestial bodies between the two encounters, and the testing to a wider set of cases.

REFERENCES

Acton C., Bachman N., Semenov B., Wright E., 2018, Planet. Space Sci., 150, 9

Armellin R., Di Lizia P., Bernelli Zazzera F., Berz M., 2010b, Celest. Mech. Dyn. Astron., 107, 451
Armellin R., Di Lizia P., Berz M., Makino K., 2010a, Celest. Mech. Dyn. Astron., 107, 377
Au S.-K., Beck J. L., 2001, Probabilistic Eng. Mech., 16, 263
Berz M., 1986, Techreport AT-6:ATN-86-16, The new method of TPSD algebra for the description of beam dynamics to high orders. Los Alamos National Laboratory, New Mexico, USA
Berz M., 1987, Nucl. Instrum. Methods Phys. Res. A, A258, 431
Berz M., 1999, Modern Map Methods in Particle Beam Physics, 1st edn., Academic Press, New York
Cadini F., Avram D., Pedroni N., Zio E., 2012, Reliab. Eng. Syst. Safe., 100, 75
Chesley S. R., ed., 2005, Proc. IAU Symp. 229, Asteroids, Comets, Meteors, Rio de Janeiro, Brazil, p. 215
Chesley S. R., Milani A., 1999, Bull. Astron. Soc., 31, 1117
Chodas P. W., 1993, BAAS, 25, 1236
Chodas P. W., Yeomans D. K., 1999, AAS/Division of Dynamical Astronomy Meeting, 1227
Di Lizia P., 2008, PhD Thesis, Polytechnic Univ. Milan
Di Lizia P., Armellin R., Lavagna M., 2008, Celest. Mech. Dyn. Astron., 102, 355
Farnocchia D., Chesley S. R., Milani A., Gronchi G. F., Chodas P. W., 2015, in Michel P., DeMeo F. E., Bottke W. F., eds, Asteroids IV, 1st edn. Univ. Arizona Press, Tucson, USA, p. 815
Gronchi G. F., Milani A., 2001, Icarus, 152, 58
Kolchin E. R., 1973, Differential Algebra and Algebraic Groups. Academic Press, New York
Massari M., Di Lizia P., Rasotto M., 2017, J. Aerosp. Inf. Syst., 14, 493
Milani A., 1999, Icarus, 137, 269
Milani A., Chesley S. R., Boattini A., Valsecchi G. B., 2000b, Icarus, 145, 12
Milani A., Chesley S. R., Chodas P. W., Valsecchi G. B., 2002, in Bottke W. F., Jr, Cellino A., Paolicchi P., Binzel R. P., eds, Asteroids III, 1 edn. Univ. Arizona Press, Tucson, USA, p. 55
Milani A., Chesley S. R., Sansaturio M. E., Tommei G., Valsecchi G. B., 2005, Icarus, 173, 362
Milani A., Chesley S. R., Valsecchi G. B., 1999, A&A, 346, L65
Milani A., Chesley S. R., Valsecchi G. B., 2000a, Planet. Space Sci., 48, 945
Morselli A., Armellin R., Di Lizia P., Bernelli Zazzera F., 2014, Adv. Space Res., 55, 311
Risch R. H., 1969, Trans. Am. Math. Soc., 139, 167
Risch R. H., 1970, Bull. Am. Math. Soc., 76, 605
Ritt J. F., 1932, Differential Equations From the Algebraic Standpoint. American Mathematical Society, Rhode Island, USA
Ritt J. F., 1948, Integration in finite terms: Liouville's theory of elementary methods. Columbia Univ. Press, New York, USA
Seidelmann P. K., 1992, Explanatory Supplement to the Astronomical Almanac. Univ. Science Books, California, USA
Valli M., Armellin R., Di Lizia P., Lavagna M., 2013, J. Guid. Control Dyn., 36, 48
Valsecchi G. B., Milani A., Gronchi G. F., Chesley S. R., 2003, A&A, 408, 1179
Wittig A., Di Lizia P., Armellin R., Makino K., Bernelli Zazzera F., Berz M., 2015, Celest. Mech. Dyn. Astron., 122, 239
Yeomans D. K., Chodas P. W., 1994, in Gehrels T., Matthews M. S., Schumann A. M., eds, Hazards Due to Comets and Asteroids, Univ. Arizona Press, Tucson, USA, p. 241
Zio E., 2013, The Monte Carlo Simulation Method for System Reliability and Risk Analysis, 1 edn. Springer Series in Reliability Engineering, Springer-Verlag, Berlin
Zuev K. M., Beck J. L., Au S.-K., Katafygiotis L. S., 2012, Comput. Struct., 92-92, 283

This paper has been typeset from a TeX/LaTeX file prepared by the author.