RARE EVENT SAMPLING SCHEMES FOR THE EFFICIENT COMPUTATION OF IN-ORBIT COLLISION PROBABILITIES

Yannick Sztamfater-García, Guillermo Escribano, Manuel Sanjurjo-Rivo, and Joaquín Míguez

Universidad Carlos III de Madrid, Av. De La Universidad, 30, Madrid, Spain, Email: ysztamfa@pa.uc3m.es

ABSTRACT

The near-Earth environment is rapidly changing. The associated requirements to preserve its safety and achieve automation of collision avoidance procedures require a search for alternative approaches to improve the current state-of-the-art. In-orbit collisions are considered rare events, as they are seldom observed. The accurate computation of rare event probabilities has been a matter of extensive research in statistics and fields as varied as climatology, road-traffic management and radio communications, as these events are often those with the direst consequences.

We propose a novel method for the computation of the in-orbit probability of collision (PoC), based on the combination of rare event importance sampling and an approximate dynamical model which exploits the ease of computation of collision scenarios using Keplerian orbital dynamics. This approach enables us to maximise the accuracy of the resulting estimator of the PoC for a given computational budget.

Keywords: SSA, SST, Space Debris, Collision Probability, Rare Event Sampling, Conjunction Analysis.

1. INTRODUCTION

Rare events, as the name implies, are infrequent. However, oftentimes, it is these events which have the highest impact on their surroundings, so accurate ways to model and predict them is crucial in many fields. A prime example of a high impact event is an in-orbit collision between space objects. The near-Earth space environment is experiencing a rapid increase in resident space objects (RSOs), both operational and inoperative. This change is due to the improvement of launch capabilities, the reduction in cost of space missions themselves, the expansion of internet provision infrastructure carried out by Starlink and OneWeb and the advent of nano-satellite technologies, allowing for the simultaneous deployment of multiple, inexpensive satellites. As of 2022, an estimated 36,000 objects currently orbit the Earth, out of which just over 7000 are operational spacecraft, the remainder including long de-serviced satellites and space debris [1]. To worsen the panorama, events such as the 2009 Iridium 33-Cosmos 2251 collision, as well as several guided anti-satellite missile tests, carried out by some of the world's superpowers, have expelled thousands of new fragments into Earth orbit, each of which represents a deadly threat in itself. In addition, due to their size, most are almost impossible to accurately track.

Space situational awareness (SSA) is the framework which houses the combination of reconnaissance, environmental monitoring, and space surveillance activities [2]. Its aim is to ease space traffic management (STM), or the process of coordinated response to newly identified threats to safety, as well as maintaining safety guidelines to prevent catastrophe. Space surveillance and tracking (SST) and collision avoidance (CA) procedures have been shown to suffice for current RSO population safety needs but may soon become insufficient [2]. To prevent the NEO environment from becoming inoperable, improved collision risk computation techniques are required.

Contributions. In this work, we expand on the work started in [3], in which a new orbit conjunction metric is proposed. We begin by further validating the approximate model in LEO and GEO scenarios. Then we delve into the application of rare event sampling, by introducing an adaptive importance sampling approach, achieving a more accurate PoC calculation and a fair representation of the domain of attraction of the extreme event. We propose and apply an adaptive framework, which can be iterated based on the predictability of the method. Finally the the algorithmic efficiency of the entire procedure is enhanced by parallelizing the dynamical propagation of the method, bringing computation run-time down considerably.

Organization of the paper. The rest of the paper is organized as follows: Section 2 reviews the state of the art regarding probability of collision computation and extreme value theory applications. Section 3 provides a description of the problem state space, as well as defining the cited metric, and the extreme event bounds. Section 4 discusses the different rare-event methodologies implemented and the proposed framework. Section 5.1 and Section 5.2 show the results for a LEO and a GEO collision respectively, ending with the conclusions in Section 6.

2. STATE OF THE ART

2.1. Probability of collision computation

Current collision probability methods can be somewhat restrictive, and their assumptions may or may not hold true in general. They focus on the objects' covariance matrices and modelling their states as Gaussian distributions. Methods like Foster's, Chan's, Patera's and Alfano's are all analytical or numerical methods employing transformations of the relevant pdfs and assumptions such as rectilinear motion, spherical shapes, Gaussian positional errors, and a relative velocity that is large enough that its covariance can be deemed static. See [4, 5] for an in-depth review of these methods. These assumptions, although currently necessary for an estimation of in-orbit probabilities, limit the applicability of such methods on a wider range of cases such as those with non-Gaussian uncertainty.

A standard method to compute the probability of collision between space objects is the crude Monte Carlo (CMC) simulation. This method entails simulating N independent and identically distributed (i.i.d.) samples from a starting distribution and applying some transformation, over which the extreme event set is defined based on a threshold of exceedance [6]. The CMC provides an unbiased estimate of the probability, and makes no prior assumptions, making it a suitable benchmark method with which to compare results. The drawback is that if the event in question is rare, the number of samples required to obtain even just one successful event is extraordinarily large, which makes this kind of method infeasible for rare event computations. However, it is in principle, possible to represent this target distribution in ways that allow for the study of rare events with more ease. This is possible with the help of extreme value theory (EVT) and rare-event sampling techniques (REST).

2.2. Extreme Value Theory

One way to study the behaviour of the tails of a probability density function (pdf) is to extrapolate beyond the available data and determine the characteristics of the extreme event, making use of methods such as block maxima [7], peak-over-threshold or large deviation theory. A thorough introductory review of EVT can be found in [8], and an extensive explanation of different rare-event simulation approaches can be found in [9].

2.3. Rare Event Sampling Techniques

Rare event sampling relies on a numerical model that is enticed to generate more realizations of a rare event than it would normally, provided the conditions for the rare event are known to an extent. The prime examples of this methodology include importance splitting and importance sampling (both static and dynamical) algorithms.

Importance splitting. Importance splitting is a method by which one generates more realizations of a rare event by resampling promising trajectories [10]. Splitting methods compute small tail probabilities p' as a product of not-so-small conditional probabilities [11]. However these methods are more suitable for sequential processes, instead of static ones which will be the focus of this work, as will be explained shortly.

Importance sampling. Importance sampling involves sampling regions of the state space which are more likely to lead to the event in question, with the aim of reducing estimate variance and computational cost [12, 6, 13, 9], which, with CMC, it is rather large for even just one successful realization. There are various flavours of IS, which depend on the problem at hand. Static IS refers to using a fixed change of measure throughout the simulation (generating one proposal), whilst adaptive IS involves updating and learning an improved change of measure based on the samples which have been simulated [6]. A change of measure is introduced by means of an alternative probability measure, such that the likelihood ratio, the Radon-Nykodyn derivative, between the nominal probability measure (the one associated to the rare event we are trying to estimate) and our proposal is well defined on the rare event of interest [13]. This reduces the computational cost of producing more realizations of the rare event.

Works which apply the EVT and REST to safety and collision avoidance endeavours in other fields include [15, 14]. Splitting techniques are applied to the satellite conjunction assessment problem in [11], where they use an adaptive splitting technique to analyse the Iridium-Cosmos collision. In 2006, [16] used a sequential Monte Carlo (SMC) method to estimate the collision risk of aircraft in a given Air Traffic Area, whilst [17, 18] employed an SMC^2 on the satellite conjunction problem, by nesting an SMC inside another, for parameter estimation at time of closest approach.

All of these approaches deal with distance as the target metric on which to impose the threshold of exceedance. While this approach may be suitable for problems in which distance is a smooth function, in orbital dynamics, velocities and distances at key time-intervals in which collisions typically occur (which are extremely short), may fluctuate enormously, making it less suitable for application.

3. APPROXIMATE DYNAMICAL MODEL

In this section, a dimensionality reduction and conjunction mapping method introduced in [3] is described. The mapping represents two space objects' states in terms of a simpler metric in $\mathbb{R}^2$ with which we study collision events.

3.1. State Space Definition

The state of each of these two RSOs is represented by a 6-dimensional vector composed of the position and velocity vectors in earth-centred-inertial (ECI) reference frame coordinates. Let $\mathbf{x}_{i,t=0} = [\mathbf{r}_{i,t=0}^\top \ \mathbf{v}_{i,t=0}^\top]^\top$, the initial state vector of object $i = 1, 2$, each of which may be represented by probability distribution $\pi_{i,t=0}(\mathbf{x}_{i,t=0})$. For simplicity, take $\mathbf{x}_{i,0}$ to be equivalent to $\mathbf{x}_{i,t=0}$. The two $\mathbb{R}^6$ vectors are then concatenated into a single $\mathbb{R}^{12}$ vector composed of both objects' initial states. This way, $\mathbf{X}_0 = [\mathbf{x}_{1,0}^\top \ \mathbf{x}_{2,0}^\top]^\top \in \mathbb{R}^{12}$ is their joint state vector, and $\pi_0(\mathbf{X}_0)$, the joint initial pdf. Both objects are subject to their corresponding orbital regimes' dynamics, so care must be taken when propagating their orbits, possibly separately. A deterministic dynamical model for the propagation of both objects' orbits, which may be high fidelity (HF) or low-fidelity (LF), can be denoted by the map

$$\begin{aligned} \Phi : \mathbb{R}^{12} &\rightarrow \mathbb{R}^{12} \\ \mathbf{X}_0 &\rightarrow \mathbf{X}_t \end{aligned} \tag{1}$$

3.2. 2D Mapping

During propagation, at some finite timestamp t_c after t_0, objects 1 and 2 may come into close range of one another. A metric of distance, d_{min} is usually used to quantify this proximity, and an associated threshold, R_{th} is defined to classify a possible collision. A collision occurs when R_{2,t_c}–$R_{1,t_c} \leq R_{th}$ at t_c. Complex models for accurate determination of R_{th} may be employed which account for the geometry of the object, though in our case, R_{th} is fixed and equal to the sum of the radii of the spheres which contain each object, for simplicity.
Determining the probability of collision between space objects, with distance as a lone metric may be suboptimal, since velocities of space objects in, say, a standard LEO orbit exceed the $7\ km/s$ mark, so collisions occur at very small timescales over which the distance between the objects fluctuates enormously. Numerical methods may fail to capture collisions due to the lack of smoothness in the distance function. A collision metric is therefore used to define an alternative geometrical space, less sensitive to the inherent magnitudes of orbital dynamics and can in principle capture events in a smoother fashion. The central idea exploits the fact that a collision must occur at the intersection node of the two orbital planes. By studying the angular momentum, given by $\mathbf{h}_i = \mathbf{r}_{i,0} \times \mathbf{v}_{i,0}$ for one object, this collision node is given by

$$\mathbf{n}_c = \pm \frac{\mathbf{h}_1 \times \mathbf{h}_2}{|\mathbf{h}_1 \times \mathbf{h}_2|} \tag{2}$$

The mentioned 2D collision vector is defined as $\xi = [\Gamma \ \Delta]^\top$ and composed of two variables defined below (see Fig. 1 for an illustration). The quantities expressed in the metric arise due to two conditions which must be fulfilled for a collision to occur, in the case where one assumes Keplerian dynamics alone (no perturbations):

- The radii of the orbits along n_c should be equal. This leads to the ratio of radii Γ at collision node, given by
$$\Gamma = 1 - \frac{r_{2,c}}{r_{1,c}} \tag{3}$$
where $r_{i,c} = (h_i^2/\mu)/(1 + e_i \cos \nu_{i,c})$, μ is the gravitational constant of the Earth, e_i, the eccentricity of the orbit of object $i = 1, 2$ ($\mathbf{e}_i$, the eccentricity vector of the orbit of object $i = 1, 2$) and $\cos(\nu_{i,c}) = \mathbf{n}_c^\top \cdot \mathbf{e}_i / |\mathbf{e}_i|$.

- The angular position of both objects within their orbits should be along $\mathbf{n}_c$ at the same time. Their angular distance Δ is given by
$$\Delta = \nu_2(t_{c,m}) - \nu_{2,c} \tag{4}$$
where $t_{c,m}$ corresponds to the time at which object 1 is along the collision direction, and the angular velocity can be computed from Kepler's equations.

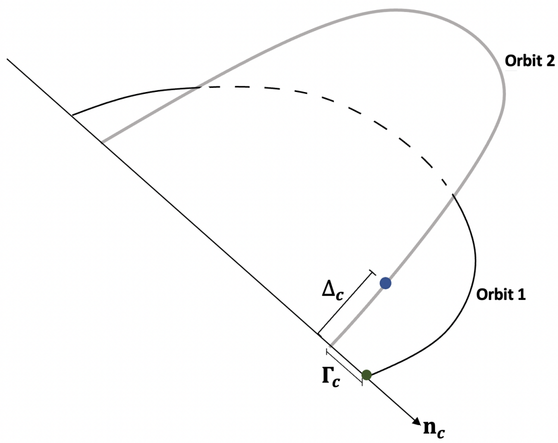

Figure 1. Illustration of the two variables, Γ and Δ, for the collision metric used to define the dimensionality reduction mapping. The line of nodes of satellites 1 and 2 is also shown. It is easy to see that if Γ and Δ are both zero, the objects are at the same point in space.

3.3. Approximate Model

Keplerian dynamics alone, however, rarely result in acceptable levels of accuracy, especially in LEO, where drag deviates the line of nodes significantly from the Keplerian estimate. This implies that perturbations must be considered, increasing the computational cost significantly. By generating several samples from the initial distribution of states and propagating them according to the HF model in eq. 1, the nodes between the two objects can be recorded at finite time-steps by using a stop function in the propagation. The $\mathbf{X}_t$ states at the crossings are then mapped into the 2D conjunction space.

With infinite computational power, we would compute the 2D mapping in HF for every case. With realistic computational capabilities, however, the process is expensive. Reasonable results, however can be obtained by employing a first-order linear approximation of the HF estimate, for which we must evaluate the following expansion around the initial condition or that with the minimum ξ_c, which reads

$$\Xi_{c,i} = \Xi_{c,m} + \mathbf{J}_c \left(\mathbf{X}_{i,0} - \mathbf{X}_{c,m,0}\right). \tag{5}$$

The full linear approximation can then be denoted by

$$\mathbf{\Psi}_c = \xi_i^{LF} + \Xi_{c,m} + \mathbf{J}_c \left(\mathbf{X}_{i,0} - \mathbf{X}_{c,m,0}\right). \tag{6}$$

We find $\Xi_{c,m} = \xi_c^{HF} - \xi_c^{LF}$ by simply saving the 2D HF value of the sample closest to [0,0], obtaining ξ_c^{HF}, whilst at the same time, saving the lowest values in the Keplerian propagation, obtaining ξ_c^{LF}. The caveat is this approximation probably does not apply in LEO past 7 days or so. Also, to be able to perform the correction, we must necessarily first propagate using the HF model to obtain the required correction terms.

Next, the full approximate model draws a very large number of samples from the initial distribution, and maps them onto the 2D plane via Keplerian dynamics, obtaining ξ^{LF} in eq. 6. $\mathbf{J}_c$ is defined as $\mathbf{J}_c = \partial \Xi_c / \partial \mathbf{X}_0$, i.e., the partial derivative of the Ξ values over the partial derivative of the initial state $\mathbf{X}_0$. Finite differences may be applied, giving $\mathbf{J}_c = \Delta \mathbf{\Xi_c} / \Delta \mathbf{X}_0$.

One can now linearly correct the estimate of the LF model for the large pool of samples to enhance accuracy, via the mapping above in eq. 6, defined as

$$\begin{aligned} \Upsilon(\mathbf{X}_{c,m,0}) : \mathbb{R}^{12} &\to \mathbb{R}^2 \\ \mathbf{X}_0 &\to \mathbf{\Psi}_c \end{aligned} \tag{7}$$

3.4. Extreme event definition

At this point we have obtained possible collision event samples given an initial state. Now we must compute their probability of occurrence. Traditionally, methods to compute this collision probability (PoC) study the likelihood of two objects coming into a distance lower than a certain distance threshold. By transforming this threshold into a collision condition in the conjunction plane, we can represent the extreme domain as a circle of radius $||\xi|| \leq R_{th}/r_{1,c}$ around [0,0]. The collision circle in a typical LEO scenario is seen in Fig. 2.

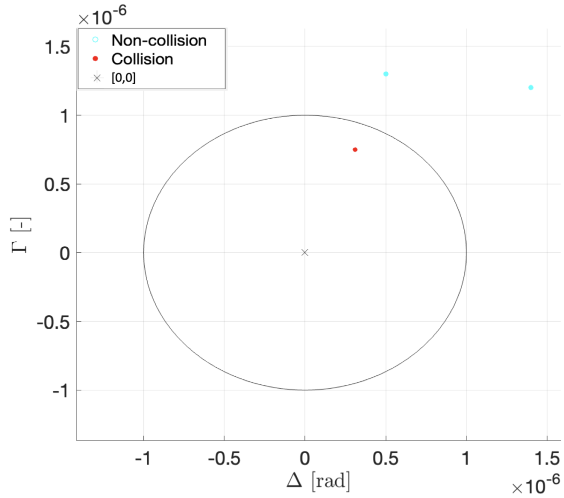

Figure 2. Illustration of the extreme event threshold in the conjunction space for a typical LEO scenario. The central cross represents certain collision (0,0), the red dot is inside the collision circle bounded by the threshold, and the cyan samples are points further away, and therefore non-colliding samples.

According to theory, the extreme event set $E(f_e)$ is the set of samples which fulfil the extreme event condition above. The proposed method aims to compute the PoC, and determine a well-defined domain of attraction (DoA). The DoA is described as the set of all initial states which lead to an observable transformation below the threshold defined in section 3.2. Regardless of the method employed, PoC will be the probability that the joint state vector defined in sec. 3.1 belongs to the DoA, provided the latter has been adequately defined (which is a challenge in itself).

3.5. Probability of Collision

There are several ways of computing the PoC. With the approximate model described above, using a very large (but inexpensive) number of samples centered on the point the linearization was performed, the PoC may be computed as the Monte Carlo sum of the generated samples which exist in the domain of attraction set, i.e., those whose L_2 norm in $\mathbb{R}^2$ is equal or lower than a specified threshold. More formally, we define the domain of attraction as $D_\chi(R_{th}/r_{1,c}) = \left\{\mathbf{X} \in \mathbb{R}^{12} \cap \mathbf{\Upsilon}(\mathbf{X}_0) \leq R_{th}/r_{1,c}\right\}$, implying that the theoretical probability of collision translates to

$$\mathbb{P}_\pi(D_\chi) = \int I_{D_\chi}(\mathbf{X}_0)\pi(\mathbf{X}_0)\mathrm{d}\mathbf{X}_0 \,, \tag{8}$$

where $I_{D_\chi}(\mathbf{X}) = 1$ if $\mathbf{X} \in D_\chi$ and $I_{D_\chi}(\mathbf{X}) = 0$ if $\mathbf{X} \notin D_\chi$. Hence by using the N particles generated, the probability is approximated as

$$\mathbb{P}_\pi(D_\chi) \approx \frac{1}{n}\Sigma_i^n I_{D_\chi}(\mathbf{X}_0^i) \qquad \mathbf{X}_0^i \sim \pi(\mathbf{X}_0). \tag{9}$$

Applying the above to the computation of a rare event, estimate variance will be high and efficiency, low.

4. PROPOSED METHODOLOGY

4.1. Importance Sampling

The rare-event sampling method of choice in this communication is importance sampling, which will be described in more depth in these pages. In order to study the occurrence of extreme events, we turn to these strategies to ease and optimize computations. More specifically, we begin by explaining the methodology behind importance sampling (IS), as applied to our problem. In the extreme case where we are able to generate enough particles to “fill” the collision area in $\mathbb{R}^2$ (see Fig. 2) down to infinitesimal levels, we should in principle be able to completely characterize the domain of attraction.

The first step of the general approach is to build the approximate model following the previous sections, to obtain suitable values for the Jacobian and $\mathbf{\Psi}_c$. The following are thorough descriptions of the fundamentals of the approaches followed:

4.1.1. Approach 1: standard IS

One must construct a proposal function $\mu(\mathbf{X}_0)$ which will generate more rare event instances, thereby reducing estimate variance, compared to the crude MC approach. If the proposal used was zero variance, it would generate exactly and only collision events. Clearly this requires prior information on the collision event, which is impossible a-priori, but by performing a cheap simulation of the afore-mentioned approximate process, one can close in on said knowledge. Our proposal function is chosen to be a Gaussian distribution, based on a mean value and an associated covariance. The more information one has, the more varied and complex the distribution type can be. In our case it was defined on the initial samples in $\mathbb{R}^{12}$.

- **Obtaining the mean value**
  A secant root-finding method is used on the linear approximation function to determine the $\mathbf{X}_0$ which leads to the lowest $\mathbf{\Psi}$. The algorithm requires two guesses, out of which at least one must be sufficiently good. Once the tolerance criterion is adjusted to output the desired accuracy, the resulting $\mathbf{X}_0$ is selected as the proposal distribution mean $\mathbf{X}_{0,min}$.
- **Obtaining the covariance $\mathbf{C}_0$**
  A trial run of a very large number of inexpensive approximate model samples is performed with a distribution built around $\mathbf{X}_{0,min}$, with a sufficiently large covariance. Samples on the outer edge of the collision area are identified, and a covariance matrix is built based on the difference between them and $\mathbf{X}_{0,min}$, in $\mathbb{R}^{12}$. This is the proposal distribution covariance $\mathbf{C}_0$.

In the standard IS implementation, the algorithm is run once, and the number of samples drawn from the proposal is considerably smaller compared to the approximate model implementation. These are then mapped onto the conjunction plane as before. In this case, assuming the proposal $\mu(\mathbf{X}_0)$ now exists, the probability of collision can be computed as

$$\mathbb{P}_\pi(D_\chi) \approx \sum_i^n I_{D_\chi}(\mathbf{X}_0^i) w^i \qquad \mathbf{X}_0^i \sim \mu(\mathbf{X}_0)\,, \quad (10)$$

where

$$w^i = \pi(\mathbf{X}_0^i)/\mu(\mathbf{X}_0^i) \quad (11)$$

Eq. 11 implies that the weights of the samples are a measure of how well the proposal distribution ”matches” the initial distribution.

4.1.2. Approach 2: adaptive IS

Alternatively, an iterative version of the above, AIS, can be applied, which employs a static form of the sequential Monte Carlo algorithm. The initialization follows the same lines as described in the standard version, though the probability of collision computation employing the quotient of the two distributions (see eq. 11) does not occur until the end of the algorithm, once sufficient iterations have occurred. Fictitious measurements $\mathbf{Z}_j \in \mathbb{R}^2$ are then generated, with a covariance spanning several orders of magnitude less than the extreme event threshold and centred around $\mathbf{\Psi}_{Z,c} = [0\ 0]^\top$. The $\mathbb{R}^2$ particles are given temporary weights according to the likelihood of the $\mathbf{\Psi}_{i=1:N}$ samples relative to the measurements $\mathbf{Z}_j$, i.e., these weights reflect how well samples match the fictitious collision measurements, and form the basis of the resampling step. Resampling means those with higher weights are duplicated, and those with lower weights removed, so the following iteration draws samples from a reduced set and concentrates more samples around the collision area. At the end of each iteration, a new proposal function is created with a new mean and covariance, for its use in the next iteration. After a certain number of iterations, the probability of collision is computed as shown in equation 10 above.

4.2. Proposed Method

Algorithm 1 describes the procedure for implementing the proposed framework based on iterative AIS. For optimality, we present an all-encompassing framework which draws on the convenience of the approximate model's computation of the conjunction metric and rare event sampling's ability to reduce estimate variance and achieve a more efficient computation of the PoC. We begin by processing the initial state and covariance of the two satellites and obtain the ξ_0 and ξ_0^{kep} at collision, as well as the Jacobian, enabling us to build the linearization in equation 6. The method then applies AIS, by obtaining a suitable proposal with a mean value centred on the collision domain in space, and a covariance that contains the DoA a.s. Once a satisfactory coverage of the extreme event area is achieved, we measure predictability, or our ability to predict collisions with accuracy, by HF-propagating the state samples with the highest value in $\mathbb{R}^2$ (those at the frontiers of the collision zone) to determine how many are correctly identified.

Following a successful run of this algorithm, the collision zone may be adapted to provide a more realistic threshold area for the rare-event zone by selecting the samples which lie closest to the threshold boundary (but lie outside), and propagating them in HF. We can then study whether we have miss-classified potential collision events as safe, which is less admissible than identifying risk-free events as collision events. The procedure can be concisely seen in Algorithm 1 below.

Algorithm 1 Proposed Method

1. Obtain initial state vector and covariance, giving prior distribution $\pi(\mathbf{X}_0)$.
2. Propagate with full-fidelity dynamics, saving the nodes of intersection. Compute the Jacobian and save the $\xi_{c,m}^{LF}$ and $\xi_{c,m}^{HF}$ to construct the approximate model using eq. 6.
3. Construct and sample from the initial proposal $\mu(\mathbf{X}_0)$ by finding the mean $\mathbf{X}_{0,min}$ and reduced covariance $\mathbf{C}_0$, as seen in Sec. 4.1.1.
4. Begin AIS in Sec. 4.1.2 by approximating the collision zone iteratively based on likelihood weights. These are associated to similarity between our samples in $\mathbb{R}^2$ and fictitious samples $\mathbf{Z}_j$ around (0,0).
5. Compute final probability of collision following eq. 10.
6. Propagate outer boundary collision samples in HF, and compute predictability.
7. If predictability is low, perform a second iteration by setting $\mathbf{X}_0 = \mathbf{X}_{0,min}$ and perform steps 2) to 6).

5. TEST CASES: RESULTS

In this section, technical information about the configurations of the two case studies investigated here is included. The two cases correspond to a LEO encounter and a GEO encounter, and the results obtained for the algorithms outlined above are shown after a brief description of each.

5.1. LEO Case

This case study constitutes the famous collision which occurred on February 10, 2009 between operational communications satellite Iridium-33 and the defunct USSR satellite Cosmos-2251. The time of closest approach (TCA) of the two satellites was t_c = Feb 10, 2009, 16:55:59 UTC, obtained from the observed conjunction geometry [2]. The states of both spacecraft at TCA containing the position and velocity expressed in ECI coordinates (km and km/s respectively), were the following:

$$\mathbf{x}_1 = [\text{-1457.77 1589.34 6814.11 -7.00 -2.44 -0.93}]^\top$$
$$\mathbf{x}_2 = [\text{-1457.77 1589.36 6814.19 3.58 -6.17 2.20}]^\top$$

The position and velocity uncertainties for the covariance matrix of the initial state were extracted from the Consultative Committee for Space Data Systems historical archives. It is assumed equal for both objects and expressed in the cross-track, along-track, and out-of-plane (RTN) coordinate system, which must be converted to ECI coordinates to match the coordinate system of both states.

$$\mathbf{P}_0^{RTN} = \text{diag}[41.42,\ 2533,\ 70.98,\ 5.7e\text{-}3,\ 1.1e\text{-}5,\ 1e\text{-}6]^\top$$

To obtain the initial states, these conjunction conditions were HF-propagated backwards in time to find suitable initial conditions for the collision probability methods. This simulation is expected to be very computationally costly, due to its high-fidelity propagation in LEO, where calls to atmospheric models for drag computation are numerous and expensive. Note that in a typical LEO scenario, the collision threshold is assumed to be $\mathbf{R}_{\xi_{coll}} = 1 \times 10^{-6}$.

5.1.1. Performance Comparison

In this section, the performance of the algorithms is compared for the LEO case, based on probability of collision and run-time. Initially, the computational cost of the algorithm for this LEO scenario was relatively high, reaching past the 10h mark. Though not an unacceptable run-time for a method which predicts collisions in the distant future, it is a method which can be parallelized in the time-consuming stages of the code, namely, the HF Jacobian computation. By observing the data structure presented throughout the code, it was parallelized to achieve a lower run-time, reaching the 3h mark. The results for this orbital configuration are summarized in table 1. The proposed algorithm benefits from this reduction in run-time as it involves a potential double computation of the Jacobian. In light of the extremely low probability value inferred from the HF Monte Carlo simulation, we can assume the lowest of the predicted probabilities is the most accurate, given the initial state and associated uncertainty. In this case, the proposed method achieves the desired accuracy, with a sufficiently low simulation time.

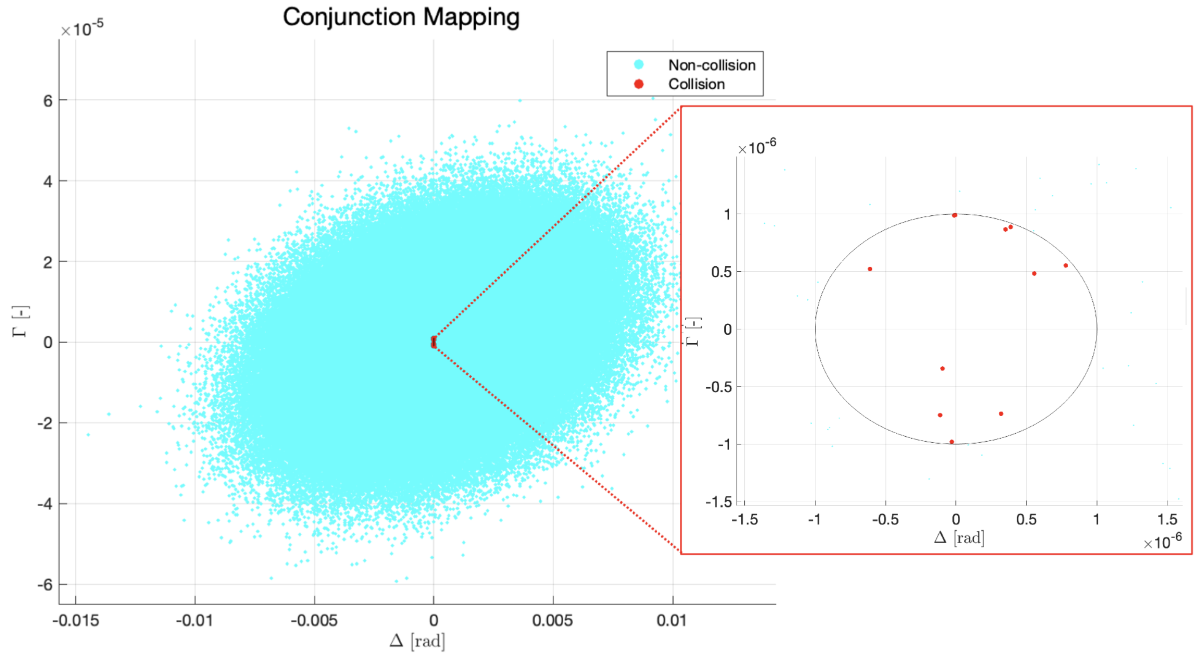

Figure 3. Conjunction mapping showing the defined 2D metric samples for the approximate model in LEO.

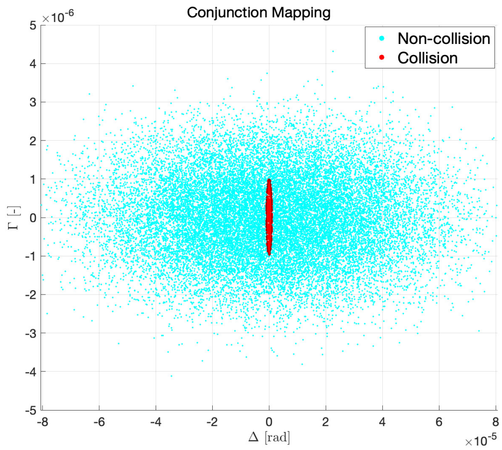

Figure 4. Conjunction mapping showing the defined 2D metric for the proposed model samples in LEO. Note the difference in scales of the axis, compared to the previous Fig.

	Probability of Collision	Run-time (s)*
MC. High-Fidelity	$< 1 \cdot 10^{-5}$	1814400
MC. Approximate Model	$1.1 \cdot 10^{-5}$	68.17
Importance Sampling	$9.2 \cdot 10^{-6}$	64.21
Proposed Method	$6.1 \cdot 10^{-6}$	74.12/150

Table 1: LEO Performance comparison showing results of probability of collision and run-time for each algorithm.

* The computational run-time is quoted for each algorithm (except for the MC High Fidelity), non-inclusive of the computation of the Jacobian (both in the standard run and in the parallelized run). The standard computation of the Jacobian for this LEO scenario takes 42,065 seconds, and 14,110 seconds in the parallelized case, and these must be added to the quoted run-times above in each case.

5.1.2. The collision zone

Figures 3 and 4 show the collision zone in $\mathbb{R}^2$ for the LEO case, for the approximate model and the proposed method, respectively. The red dots represent collision samples predicted by the respective method, while the cyan dots are non-collision samples in this space.

From the first plot in Fig. 3, we can readily observe that the variance is significant, showing a spread of several orders of magnitude around the initial state. By zooming into the centre (second plot of the same Fig.), we see that there are very few particles inside the collision zone, more precisely, 11. By crude Monte Carlo, the approximate model's probability of collision with 1 million samples is low, as seen in table 1 above. This has the general explanation of relatively poor knowledge of the state at time $t = 0$ resulting in a low probability of collision, due to the magnitude of the covariance matrix. Either way, it shows sufficient accuracy given the level of uncertainty. When plotting the proposed method samples in Fig. 4, a clear reduction in variance is observed, with a substantially larger proportion of particles residing inside the event zone: 972. The simulation was run with 30,000 samples. This illustrates the point that in an ideal sce-

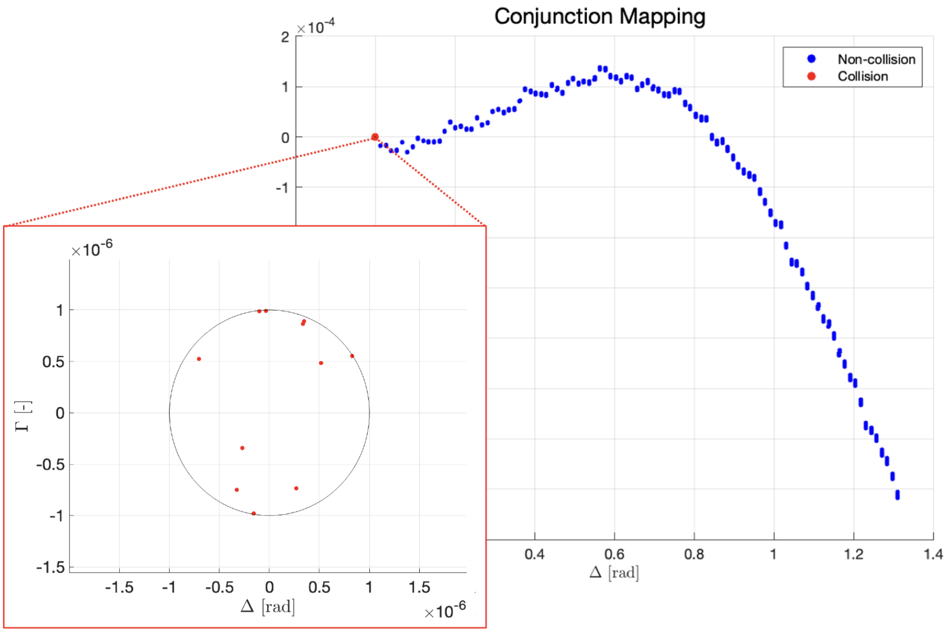

Figure 5. HF mapping showing the approximate model collision samples in the conjunction space as they evolve through each node up to conjunction. Seeing how many are correctly identified as collision samples gives an estimate of predictability.

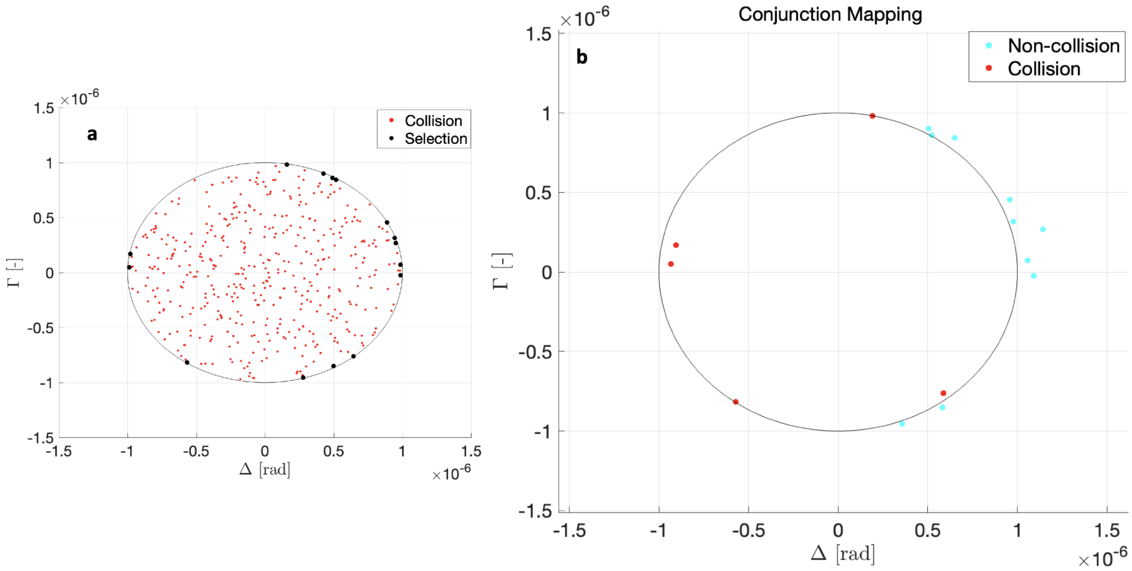

Figure 6. HF mapping showing the outer collision samples of the proposed model in the conjunction space and seeing how many are correctly identified as collision samples.

nario, a zero-variance IS estimator would generate only points within the collision zone, i.e., one would be able to fully characterize the extreme event set, as well as the domain of attraction, as N approached infinity. In any case, by employing equation 10 in section 4.1.1 above, the probability of collision can be seen as a more accurate approximation of the probability of collision given our initial degree of uncertainty.

5.1.3. Predictability graphs

A necessary test for our methods is whether the particles which lead to collision with our proposals based on the linear approximation actually lead to collision when propagated in high fidelity. Hence, we can test for predictability and validity of our method (a test inherently

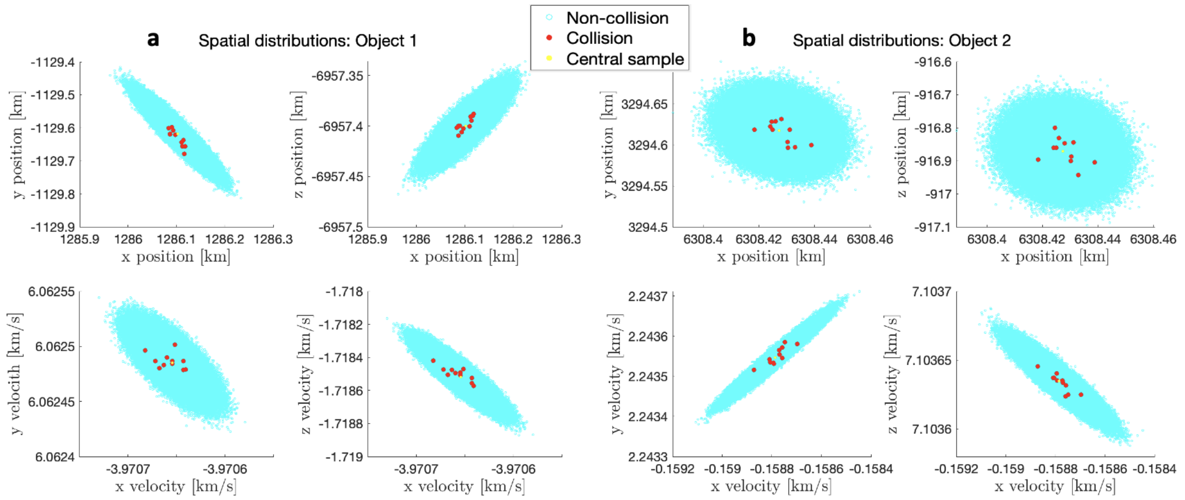

Figure 7. a) shows an illustration of the distribution of object 1 samples generated by the approximate model both in position and velocity, where each subfigure represents a pair of dimensions. b) shows the same for object 2.

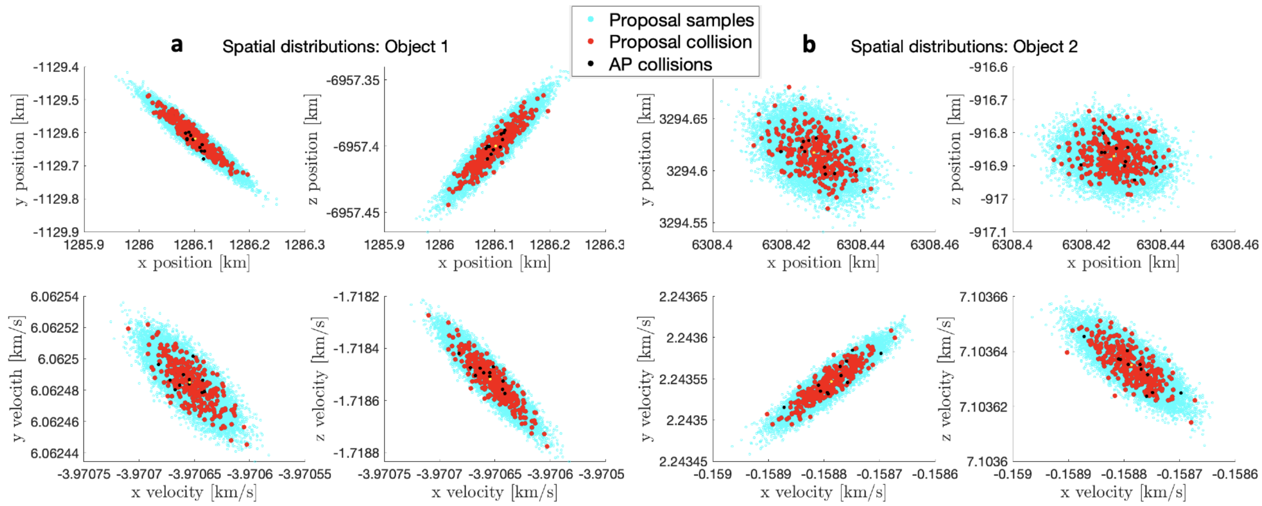

Figure 8. a) shows an illustration of the distribution of object 1 samples generated by the proposed method both in position and velocity, where each subfigure represents a pair of dimensions. b) shows the same for object 2.

present in the proposed method).

Approximate Model. The HF mapping for the approximate model samples is shown in Fig. 5, spanning all orbital nodes (in blue) up until collision node (red dot on the larger plot). When zooming into the collision area for the relevant $\mathbf{X}_0$ samples drawn, we see that all the collision samples identified in Fig. 3 are correctly characterized as collision samples in HF. This may be due to the low number of collision samples, i.e., they may just happen to be within the collision zone in HF, so predictability of 100% cannot be claimed yet. A larger subset of samples should be compared in the HF to determine how well we are predicting the risk of collision. This subset may be defined by the circle corresponding to a radius of 2×10^{-6}. This is left for a future study.

Proposed Method. The HF comparison for the proposed method samples is shown on Fig. 6. The entire collision sample set (essentially Fig. 4) is depicted on the left of the figure, with the outermost 15 samples selected for HF propagation shown in black. The same samples shown after HF propagation are shown on the on the right of said figure. We see successful prediction, taking into consideration that the location is roughly the same in the HF map, indicating continuity in the mapping. However half of them lie just outside the boundaries, though these are indeed the outermost samples, so it may be assumed that samples which are not at the limits on the left will be correctly classified. A study of the predicted non-collision particles in the surrounding space as suggested above would be ideal, to see whether they lead to collision in HF, in the hopes of determining possible under-estimation of risk. This leads to the conclusion that for the vast majority of samples, our methods are reliable, as seen by their high predictability.

5.1.4. Space Domain graphs

In this section, we plot all the proposal initial state $\mathbf{X}_0$ samples (in cyan), as well as the collision samples (in red) for the algorithms above, with the aim of determining a possible domain of attraction in $\mathbb{R}^{12}$ (i.e., the set which will end up colliding in the future). For each method, 4 subplots are shown for each of the two objects, showing the space distribution in the $X - Y$ and $X - Z$ axis, and the velocity distribution in the $v_X - v_Y$ and $v_X - v_Z$ axis, respectively. As seen in Figs. 7 and 8, generally samples lie in the center of the sample cloud in space and velocity for both objects. By plotting the approximate model samples which we know lead to collision (in black) as verified in the previous section, we can add to the answer to the question of whether we are underpredicting risk by observing where these points lie relative to the proposed algorithm's collision points. Generally good agreement is observed, with the black dots being mostly contained in the collision cloud, though outlier detection analysis can be performed here. In the limit of N approaching infinity, these dots would be completely covered inside a red homogeneous collision cloud. Upon closer inspection, one may question the continuity properties of the mapping, since there seem to be non-collision samples in the middle of an apparently exclusive collision area. However, by observing these points on a case by case basis to determine which may be outliers, it can be shown that this is an illusion produced by the high-dimensionality of the problem, and the inability to fully represent the domain of attraction of a 12 dimensional problem in just two dimensions. See next section 5.1.5 for a spatial depth analysis of these potential outliers.

5.1.5. Spatial Depth

Spatial depth analysis is performed to study whether there is a continuity issue in the domain of attraction computation. We select a few samples which seem to lie within the DoA for the approximate model and for the proposed method. The Mahalanobis distance (MD) is computed from the centre of the distribution to determine their distance and find possible "outlyingness" [20]. Note that for the plots, we decided to represent the 12D state samples by way of a variable K which holds no physical meaning, it is devised merely to represent the state samples in 1D clearly and the MD associated to them. It is defined as $K = |r_0| + |v_0|\Delta t$ where both $|r_0|$ and $|v_0|$ represent the euclidean norm of all position vectors and all velocity vectors of both objects, respectively, and $\Delta t = 1$. Indeed, Figs. 9 and 10 below show that this subset of points seems to lie at MD values further than the cloud of collision points, from the centre of the distribution, meaning that there should apparently be no issues regarding estimate continuity. Note that while the MD values may be potentially very high, they need not be; most non-collision samples should be outside or bordering a threshold. For example, the lowest non-collision sample in Fig. 10 is roughly a distance of 4 MDs from the red cloud.

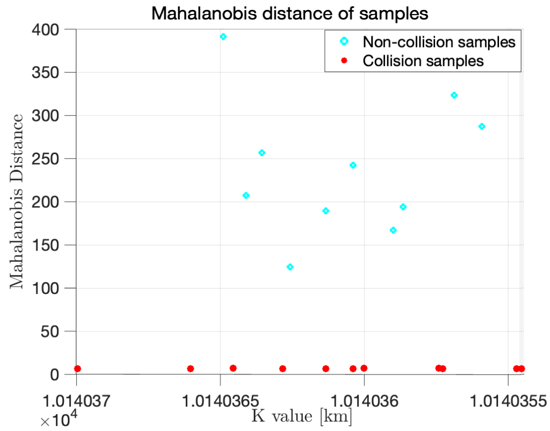

Figure 9. Figure showing the difference in Mahalanobis distance between points which lie within the domain of attraction of the approximate model when observing the space domain figures in two dimensions. Due to the scales, collision samples appear close to zero, but they are in fact in the MD = 3 - 5 vicinity.

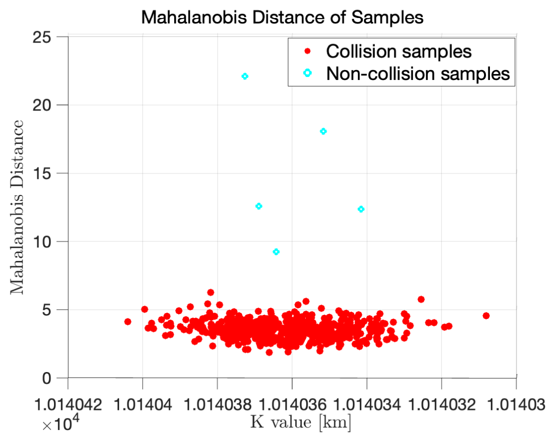

Figure 10. Figure showing the difference in Mahalanobis distance between points which lie within the domain of attraction of the proposed method when observing the space domain figures in two dimensions.

5.2. GEO Case

This case study constitutes a simulated collision in geostationary orbit between two satellites. The time of closest approach (TCA) of the two satellites was t_c = Nov 17, 2016, 08:25:43 UTC, obtained from the observed conjunction geometry. The states of both spacecraft at TCA containing the position and velocity expressed in ECI coordinates (km and km/s respectively), were the following:

$$\mathbf{x}_1 = [42094.92\ \text{-}1113.11\ 2170.41\ 0.08\ 3.07\ \text{-}0.02]^\top$$
$$\mathbf{x}_2 = [42094.92\ \text{-}1113.10\ 2170.41\ 0.15\ 3.04\ \text{-}0.39]^\top$$

The position and velocity uncertainties for the covariance matrix of the initial state have the same source as the LEO case, but are of lower magnitude to ensure collision is appreciable with the presented methods. Its RTN coordinates must be converted to ECI coordinates to match the coordinate system of both states.

$$\mathbf{P}_0^{RTN} = \text{diag}[41,\ 2533,\ 0.71,\ 5.7e\text{-}5,\ 1.1e\text{-}7,\ 1e\text{-}8]^\top$$

These conjunction conditions were propagated backwards in time 15 days by means of the high-fidelity propagation to find suitable initial conditions for the collision probability methods. Note that in a typical GEO scenario, the collision threshold is assumed to be $\mathbf{R}_{\xi_{coll}} = 1 \times 10^{-7}$.

5.2.1. Performance Comparison

The computational cost for the high fidelity propagation of a GEO configuration does not pose a problem as it involves no drag calculations, (as opposed to LEO). The algorithm runs in a fraction of the time of the LEO case, because the calls to the atmospheric model do not involve atmospheric density computations. In addition, due to a larger orbital semi-major axis, the number of passes is lower, so the calculation of the Jacobian is significantly more efficient. So much so, that we forgo the need for parallelization schemes, as in practice, they were found to add no significant speedup improvements. The results are summarised in table 2 below. In light of the probability value inferred from the HF Monte Carlo simulation, we can assume the highest of the predicted probabilities is the most accurate, given the initial state and associated uncertainty. In this case, the proposed method achieves the desired accuracy, with sufficiently low simulation time.

	Probability of Collision	Run-time (s)*
MC. High-Fidelity	$3.14 \cdot 10^{-4}$	604800
MC. Approximate Model	$7.6 \cdot 10^{-5}$	64.45
Importance Sampling	$8.1 \cdot 10^{-5}$	66.28
Proposed Method	$1.19 \cdot 10^{-4}$	67.46/-

Table 2: GEO Performance comparison showing results of probability of collision and run-time for each algorithm

* The computational run-time is quoted for each algorithm (except for the MC High Fidelity), non-inclusive of the computation of the Jacobian. The standard computation of the Jacobian for this GEO scenario takes 134.6 seconds and this must be added to the quoted run-times above in each case.

5.2.2. The collision zone: a comparison

Figures 11 and 12 show the collision zone in $\mathbb{R}^2$ for the approximate model and the proposed method, respectively, for the GEO case. The red dots represent the collision samples predicted by the respective method, while the cyan dots are non-collision samples in this space.

From the first plot in Fig. 11, it can be observed that the variance is significant, showing a spread of several orders of magnitude around the initial state. By zooming into the centre (second plot of the same Fig, we see a non-negligible number of particles inside the collision zone, 76. By crude Monte Carlo, the approximate model's probability of collision with 1 million samples is lower than desired, as seen in table 2 above. This has the same explanation as in the LEO case, and it is due to the magnitude of the covariance matrix. Either way, it is an acceptable value, with this level of uncertainty. When plotting the proposed method samples (Fig. 12) a clear reduction in variance is observed, with a significantly larger proportion of particles residing inside the event zone: 4258. The simulation was run with 20,000 samples, and a more accurate probability of collision is achieved (see table 2).

5.2.3. Predictability graphs

The same test as performed for LEO is repeated here to see whether the particles which are predicted to lead to collision with the proposals based on the linear approximation actually lead to collision when propagated in high fidelity. This way, one can test for predictability and validity of the method.

Fig. 13 shows the approximate model's HF mapping of all nodes into the conjunction space, with a zoom on the collision circle for the relevant $\mathbf{X}_0$ samples drawn. It can be seen that all samples are predicted to be in the vicinity of the collision circle, though not all of them are *inside*. A larger subset of samples should also be compared in the HF to determine if we are indeed overpredicting the risk of collision. This subset may be defined by the circle corresponding to a radius of 2×10^{-7}.

In Fig. 14, bearing in mind that over 4000 samples are predicted to collide, due to computational costs, even in GEO, not all are propagated in HF. The outermost 200 samples are propagated, and we see a very large proportion of these are correctly classified. Again, not all lie strictly within the collision zone, but at a small distance. indicating continuity in the mapping. We can study surrounding non-collision particles to see whether they lead

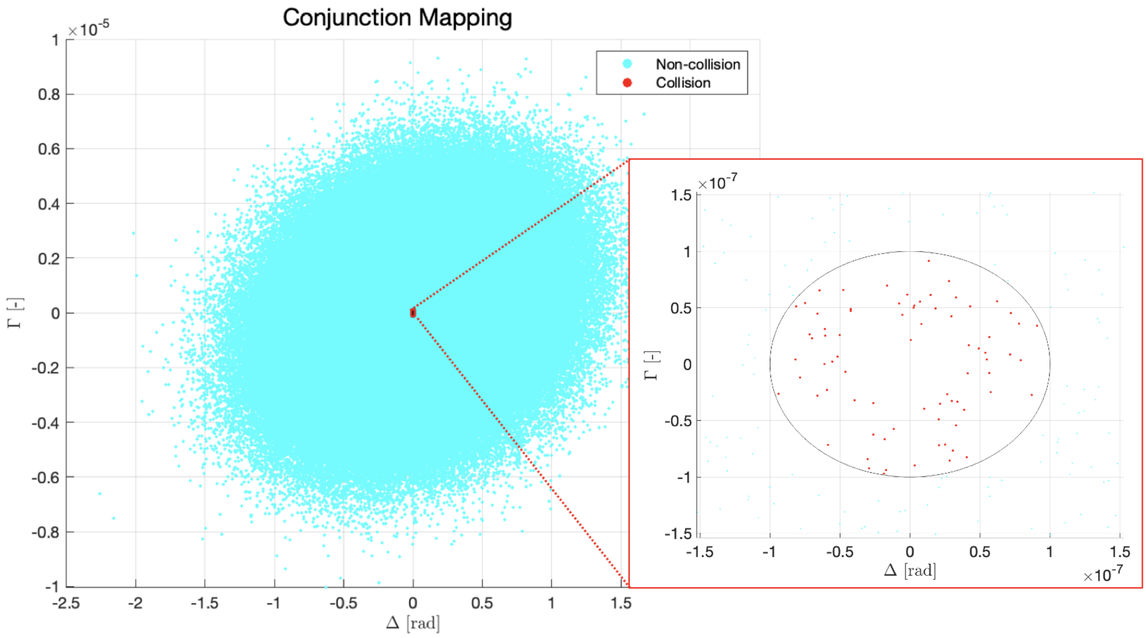

Figure 11. Conjunction mapping showing the defined 2D metric for the approximate model samples in GEO.

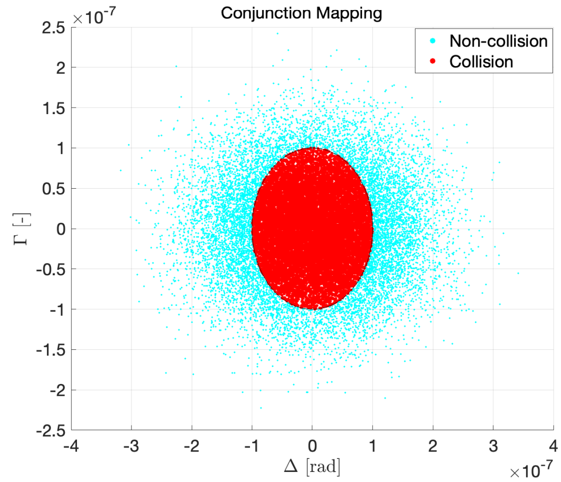

Figure 12. Conjunction mapping showing the defined 2D metric for the proposed model samples in GEO.

to collision in HF, in the hopes of determining possible underestimation of risk. These results lead to the conclusion that for the vast majority of samples, our methods are reliable, as seen by their high predictability.

5.2.4. Space Domain graphs

In Figs. 15 and 16, we plot all the proposal initial state $\mathbf{X}_0$ samples (cyan), as well as the collision samples (in red) for the algorithms above, with the aim of determining a possible domain of attraction (i.e., the set which will end up colliding in the future). For each algorithm the position and velocity axis are plotted. In Fig. 16, black dots represent the collision samples predicted by the approximate model (though bear in mind that this time, not all of them represent a 100% certain collision). As seen in Figs. 15 and 16, generally samples lie in the center of the sample cloud in space and velocity for both objects, though a more distinguishable collision area can be observed when compared to these figures' LEO counterparts. This fact, along with the conjunction plots, lead us to believe that the proposal methods work better in GEO. By observing the black benchmark points, we can add to the answer to the question of whether we are underpredicting risk by observing where these points lie relative to the proposed algorithm's collision points. Good agreement is observed, with the black dots being entirely contained in the collision clouds.

Continuity properties of the mapping can be examined again, since there seem to be non-collision samples in the middle of an apparently exclusive collision area (not entirely obvious in the IS methods but easy to see in the first approximate model figures). We observe these points on a case by case basis to determine which may be outliers. Again this seems to be an illusion produced by the high-dimensionality of the problem, and the inability to fully represent the domain of attraction of a 12 dimensional problem in 2D.

5.2.5. Spatial Depth

As mentioned, spatial depth analysis is performed to study whether there is a continuity issue in the domain of

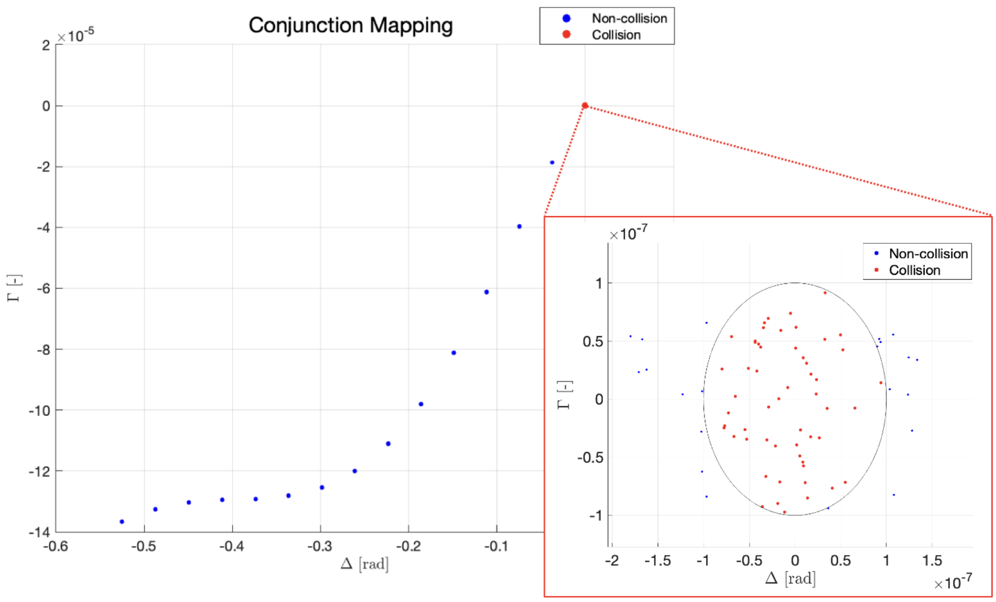

Figure 13. HF mapping showing the approximate model collision samples in the conjunction space as they evolve through each node up to conjunction. Seeing how many are correctly identified as collision samples gives an estimate of predictability.

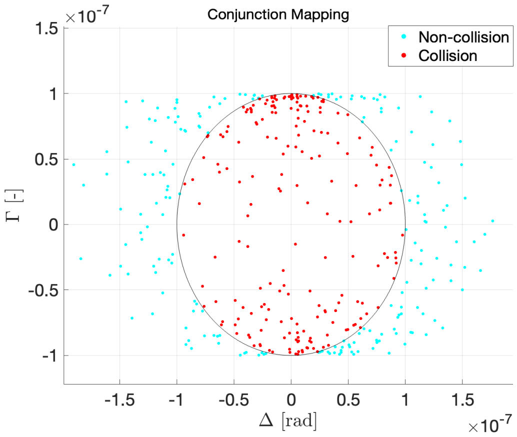

Figure 14. Conjunction mapping showing the defined 2D metric evaluated in high fidelity to examine which of the outermost 200 proposed model samples really lie within the collision zone, i.e., estimating predictability.

attraction computation. We select a small subset of samples which seem to lie within the DoA for the approximate model and for the proposed method. The Mahalanobis distance (MD) is computed from the centre of the distribution to determine their distance and find "outlyingness". Here, again we represent the 12D state samples by way of variable K defined in 5.1.5 to clearly represent the state samples in 1D and the MD. The approximate model and the proposed method spatial depth analysis can be seen in Figs. 17 and 18. Indeed, these figures show that these points lie at a Mahalanobis distance further than the cloud of collision points, from the centre of the distribution, meaning that there should be apparently no issues regarding estimate continuity. Note that due to the magnitude ranges in the MD of the non-collision samples, the collision set MD appears close to zero, but it lies in the range MD = 1 - 4.

6. CONCLUSION

The method proposed in this communication builds upon a recently developed collision risk metric for orbital objects, with the aim of reliably computing their probability of collision. This could satisfy the need to provide reliable safety guidelines to satellite operators to prevent collisions, as well as unnecessary collision avoidance maneuvers. The method proposed draws on the principles of rare event sampling to devise a way to characterize as accurately as possible, the possible initial states which may end up in collision at TCA. Two cases are studied in this communication, the Iridium-Cosmos collision of 2009 in LEO, and a simulated collision in GEO, with synthetic data. The proposed method is based heavily on adaptive importance sampling. It is validated in terms of accuracy, reaching probabilities which are limited only by our knowledge of the state of the satellite, but also in terms of computational cost, which, in stark contrast to the costly CMC, takes only a few hours to run. The predictability of our method was shown to remain consistently high in both cases which ensures reliability of the principles employed, and a domain of attraction can be somewhat characterized in all dimensions. Further work may involve a more complete characterization of the domain of attraction by studying the particles surrounding the circle of

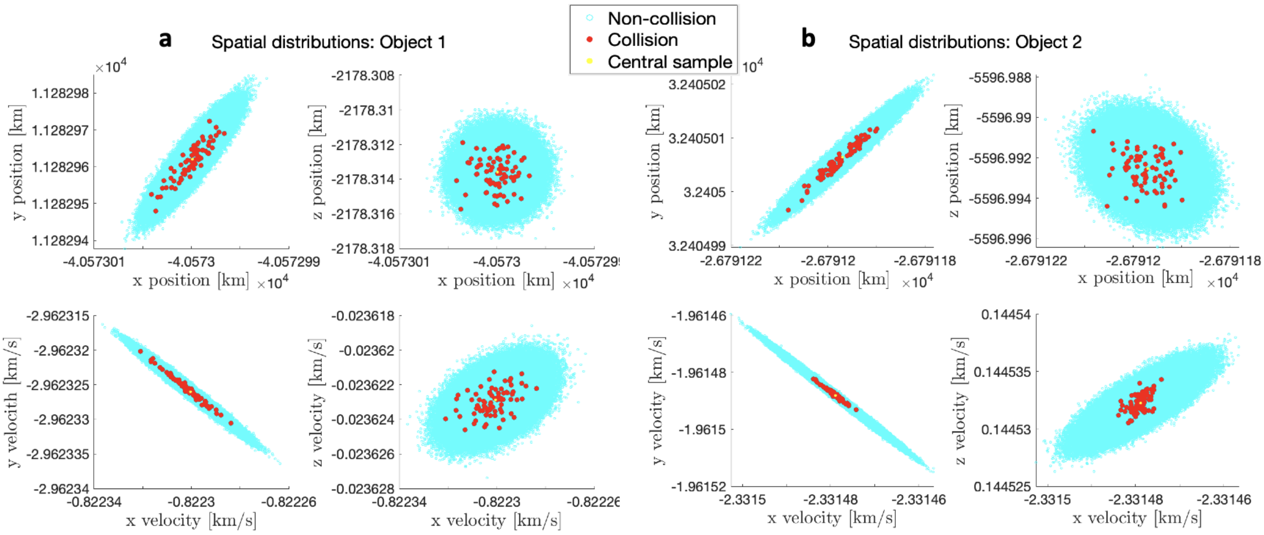

Figure 15. a) shows an illustration of the distribution of object 1 samples generated by the approximate model both in position and velocity, where each subfigure represents a pair of dimensions. b) shows the same for object 2.

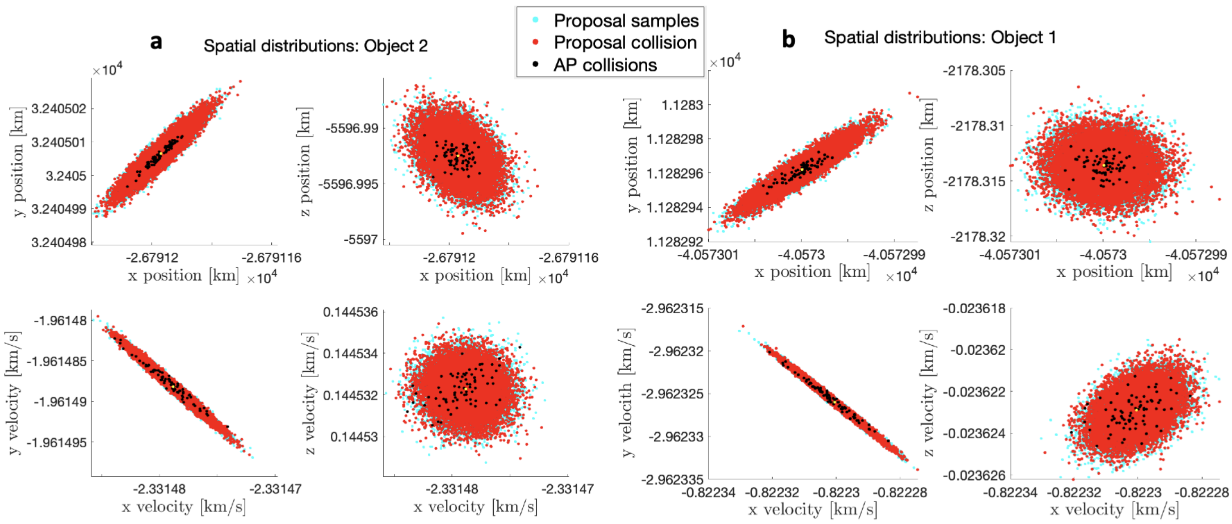

Figure 16. a) shows an illustration of the distribution of object 1 samples generated by the proposed method both in position and velocity, where each subfigure represents a pair of dimensions. b) shows the same for object 2.

collision and their HF mapping, to determine, as N goes to ∞, the correct shape of the extreme event set area in $\mathbb{R}^2$, assumed here to be circular and constant. Higher order approximations can be implemented, such as a second order approximation instead of first order, which would involve the calculation of the Hessian matrix to further increase the accuracy and predictability of the proposed method. The limits of the method can be studied in different cases where the initial state estimate is worse, or the uncertainty is significantly higher, and cases where a collision has definitely **not** occurred, to see if we can determine an extremely low probability of collision accurately.

In addition, the algorithm could be re-formulated as a sequential method, which deals with measurements as they are received, and adapts the proposal function based on each new measurement. In any case, further work on the proper development of extreme value theory applied to the satellite conjunction case is a promising line of work, and one which has not been exploited thus far. This work is a first step in said direction.

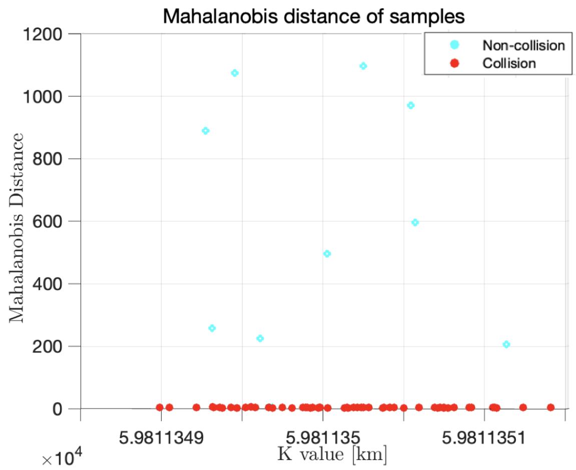

Figure 17. Figure showing the difference in MD between points which seem to lie within the domain of attraction when observing the approximate model space domain figures in two or three dimensions.

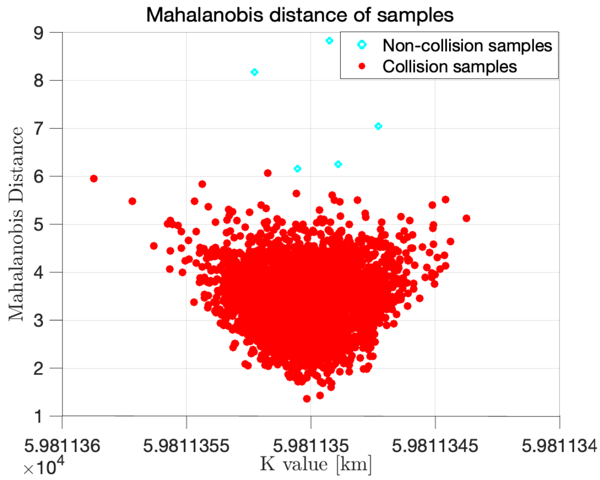

Figure 18. MD for points which lie within the domain of attraction of the proposed model when observing the space domain figures in two dimensions.

REFERENCES

1. Space Debris Environment Report 6th edition (04/2022), Space Environment Statistics, ESA Space Debris Office.
2. T.S. Kelso (2009). Analysis of the Iridium 33-Cosmos 2251 Collision. Advances in the Astronautical Sciences, p1099–1112.
3. G. Escribano, R. M. Motiño, Y. Sztamfater García, M. Sanjurjo Rivo, J. Lopez Santiago, M. Vazquez, and J. Miguez Arenas (2022). A Novel Method for Computation of Satellite Collision Probability Based on EVT. IAC, Paris, France.
4. S. Alfano (2007). Review of conjunction probability methods for short-term encounters. American Astronautical Society.
5. J-S. Li, Z. Yang and Y-Z. Luo (2022). A review of space-object collision probability computation methods. Astrodynamics, Springer 6, pages 95–120
6. S. Juneja, P. Shahabuddin (2006). Rare-event simulation techniques: An introduction and recent advances. Elseiver. Handbook in OR, Vol. 13. Chapter 11.
7. A Ferreira, L. De Haans (2015). On the block maxima method in extreme value theory: PWM estimators. The Annals of Statistics Vol. 43, No. 1. pp. 276-298 (23 pages)
8. L. De Haan, A. Ferreira (2006). Extreme value theory: an introduction, volume 21. Springer.
9. J. Morio, M. Balesdent, D. Jacquemart, and C. Vergé (2014). A survey of rare event simulation methods for static input–output models. Simulation Modelling Practice and Theory, 49:287–304.
10. F Bouchet, J Rolland, and J Wouters (2019). Rare event sampling methods. Chaos, 29(8):080402–080402,
11. R. Pastel (2011). Estimating satellite versus debris collision probabilities via the adaptive splitting technique. In Proc. 3rd Int. Conf. Comp. Modeling Simulation, p.1–6.
12. P. L'Ecuyer, G. Rubino and S. Saggadi (2011). Approximate Zero-Variance Importance Sampling for Static Network Reliability Estimation. IEEE Transactions on Reliability 60(3):590 – 604.
13. J. Blanchet and H. Lam (2012). State-dependent importance sampling for rare-event simulation: An overview and recent advances. Surveys in Operations Research and Management Science. Vol. 17, Issue 1, P.38-59.
14. F. Orsini, G. Gecchele, M. Gastaldi, and R. Rossi (2019). Collision prediction in roundabouts: a comparative study of extreme value theory approaches. Transport science, 15(2):556–572.
15. H. Farah and C. Lima Azevedo (2017). Safety analysis of passing maneuvers using extreme value theory. IATSS research, 41(1):12–21.
16. H.A.P. Blom, J. Krystul, G.J. Bakker, M.B. Klompstra, and B. K. Obbink (2007). Free flight collision risk estimation by sequential mc simulation. Stochastic hybrid systems, p249–281
17. C. Vergé, J. Morio, P. Del Moral, and J. C. D. Pérez (2016). Probabilistic safety analysis of the collision between space debris and satellite with an island particle algorithm. In Space Engineering, p443–457. Springer.
18. C. Vergé, J. Morio, P. Del Moral (2015). An Island particle algorithm for rare event analysis. Journal of Reliability Engineering and System Safety, V149; p63-75
19. M. Farazmand and T. P. Sapsis (2019). Extreme events: Mechanisms and prediction. Applied Mechanics Reviews, 71(5)
20. Y. Chen, X. Dang, H. Peng, (2009). Outlier detection with the kernelized spatial depth function. In IEEE Transactions on machine intelligence, VOL. 31 (2).