

Advanced Collision Avoidance NorthStar's Solution for Safe and Efficient Satellite Operations

Matteo Budoni^{a*}, Brendan Thorn^b, Priyatharsan Rajasekar^c, Manuel Sanjurjo-Rivo^d, Srinivas J. Setty^e

^a NorthStar Earth and Space, 24-26 Rue Goethe, Luxembourg, matteo.budoni@northstar-data.com

^b NorthStar Earth and Space, 24-26 Rue Goethe, Luxembourg, brendan.thorn@northstar-data.com

^c NorthStar Earth and Space, 24-26 Rue Goethe, Luxembourg, priyatharsan.rajasekar@northstar-data.com

^d NorthStar Earth and Space, 24-26 Rue Goethe, Luxembourg, manuel.sanjurjo@northstar-data.com

^e NorthStar Earth and Space, 24-26 Rue Goethe, Luxembourg, srinivas.setty@northstar-data.com

* Corresponding Author

Abstract

With the rapid increase in satellite deployments and the consequent rise in space objects, the need for robust collision avoidance systems has never been more urgent. NorthStar Earth and Space (NorthStar) has developed an advanced system that integrates conjunction analysis, conjunction assessment improvement, constellation risk assessment, and maneuver recommendations to ensure the safety and efficiency of satellite operations.

This paper will detail the architecture and capabilities of NorthStar's conjunction analysis system. We will demonstrate how the system synthesizes maneuver recommendations by considering collision risk alongside the specific orbit maintenance and station-keeping requirements of each satellite. Through continuous monitoring and data analysis, our system provides timely, accurate assessments, helping operators to mitigate potential conjunctions while optimizing satellite longevity and operational performance. Key areas of focus in the presentation will include generic maneuver optimization strategy and orbit improvement with new observations.

In addition to these points, the presentation will provide an outlook of the Probability of Collision methodologies employed by our system, including high-velocity, low-velocity, and Monte Carlo methods. We will explain the decision logic behind applying the appropriate method in different scenarios.

We will emphasize NorthStar's commitment to advancing space safety through continuous innovation, and how our system positions itself as a critical tool in the ongoing effort to manage and protect valuable orbital assets in an increasingly crowded space environment.

Keywords: Conjunction analysis, Collision avoidance, Space sustainability

Acronyms/Abbreviations

BFMC = Brute Force Monte Carlo

CA = Conjunction Assessment

CAM = Collision Avoidance Maneuver

CARA = Conjunction Assessment Risk Analysis

CAS = Conjunction Analysis System

CATS = Camera Target Sun

CCSDS = Consultative Committee for Space Data Systems

CDM = Conjunction Data Message

CE = Conjunction Event

CORAM = Collision Risk Assessment and Avoidance Maneuvers

CORCOS = Collision Risk Computation Software

DCA = Distance of Closest Approach

ESA = European Space Agency

GEO = Geostationary Orbit

HVO = High Value Object

IaC = Infrastructure as Code

ID = Identifier

KKT = Karush–Kuhn–Tucker

LEO = Low Earth Orbit

MEO = Medium Earth Orbit

NASA = National Aeronautics and Space Administration

NoC = Number of Collisions
NORAD = North American Aerospace Defense Command
NS = NorthStar Earth and Space
OEM = Orbit Ephemeris Message
PoC = Probability of Collision
RIC = Radial, In-track, Cross-track
RSO = Resident Space Object
SDMC = Simplified Dynamics Monte Carlo
SP = Special Perturbation
SSA = Space Situational Awareness
SST = Space Surveillance and Tracking
STM = Space Traffic Management
TCA = Time of Closest Approach
TLE = Two-Line Elements set
TRIC = Time, Radial, In-track, Cross-track
UR = User Requirements

1. Introduction

The rapid growth of satellite launches and the proliferation of space activities have created significant challenges in managing orbital congestion and mitigating the risks posed by space debris. According to the European Space Agency (ESA) Space Debris Office [1], there are more than 39000 objects regularly tracked and maintained in the Space Surveillance Network catalog, whereas the number of space debris larger than 1 cm is estimated to have exceeded 1 million.

Additionally, there are about 411 commercial satellite constellations tracked by the NewSpace Index [2,3], some of which are mega-constellations consisting of (or planned to consist of) tens of thousands of satellites, e.g., Starlink and Qianfan (or Thousand Sails). The risk posed by this type of constellation has been widely analyzed in [4].

All these factors are contributing to making the near-Earth environment increasingly crowded and hazardous. This congestion raises the likelihood of collisions, which could trigger a cascade effect, known as the Kessler Syndrome, where debris from one collision generates further collisions, potentially rendering certain orbits unusable for decades or centuries [5]. In such an environment, governments, international organizations, and private companies are actively working to mitigate the risks in various ways. De-orbit mitigation policies have been set and reinforced, technological innovation is driving towards the development of autonomous collision avoidance systems onboard the satellites. The greatest effort, however, is towards an improvement of Space Traffic Management (STM). STM frameworks have the goal to establish guidelines for tracking space objects, sharing data and coordinating maneuvers to avoid collisions. In this respect, the United States Space Force is leading the way with the 18th and 19th Space Defense Squadrons very actively contributing to the topic. In the ‘Spaceflight Safety Handbook for Satellite Operators’ [6], the two squadrons document the processes for on-orbit conjunction assessment and collision avoidance. Additionally, these entities contribute to sharing Conjunction Data Messages (CDMs), via the Space-Track.org website [7], to all the satellite owner/operators who subscribe to their services.

As described in [6] and [8], the process of Conjunction Analysis consists fundamentally of three key phases: the screening, in which satellite trajectories are compared against the space object catalog to identify close approaches within certain protective volumes; the risk assessment, in which the collision probability associated with each close approach identified in the previous step is computed; and the risk mitigation, in which avoidance maneuvers (or minor attitude adjustments) are performed when the risk exceeds pre-determined thresholds. The National Aeronautics and Space Administration (NASA) Conjunction Assessment Risk Analysis (CARA) program exemplifies modern best practices and released a handbook [9] that represents a solid reference in the field.

There are several pressure points in the current STM framework. The first one concerns data limitations, as not all the cataloged objects have reliable tracking data. Another major point is that the current methods used for assessing conjunctions generate millions of warnings annually, many of which are false positives [10], i.e., non-critical conjunctions that do not require any action. That is why some research [11] is focusing on defining completely new metrics to estimate the likelihood of satellite collisions.

Within this environment, NorthStar Earth and Space is positioning itself as a transformative player with a twofold contribution. Its space-based Space Situational Awareness (SSA) satellite constellation, using optical sensors, would allow in the long run to compensate for the data scarcity since it will provide coverage and tracking for space objects in all orbital regimes (Low Earth Orbits (LEO), Medium Earth Orbits (MEO), Geostationary Orbits, (GEO)) with

higher revisit rates than existing ground-based systems. Indeed, a space-based asset dedicated to SSA does not suffer from weather or geographical constraints and offers more geometrical observation opportunities compared to ground sensors (requiring a lower number of these latter at the cost of a more complex data processing scheme).

Additionally, NorthStar is integrating NorthStar Conjunction Analysis System (NS CAS), a software system with the overarching objective to produce accurate and actionable estimates of risk for upcoming conjunction events. This paper will focus on the description of this latter system.

Specifically, Section 2 will analyze the algorithms and methods used for assessing different types of conjunction events as well as the Collision Avoidance Maneuver (CAM) strategy, Section 3 will present NS CAS, introducing its architecture and key functionalities. Section 4 will present some results in terms of CAS implementation, showcasing actual snapshots from the current tool and presenting one selected conjunction case. Lastly, we will finish the paper with some outlooks and conclusions, grouped into Section 5.

2. Algorithms and methods

CAS constitutes a full framework for conjunction assessment and collision avoidance and, therefore, implements all the required astrodynamics and statistical methods for that purpose. In what follows, we present a review of the methods implemented for two specific functionalities: the computation of the probability of collision and the strategy for planning collision avoidance maneuvers. Other functionalities included in CAS have been presented and discussed elsewhere, for example the generation of synthetic covariance matrices from state vector data in [12].

2.1 Computation of the Probability of Collision

Assessment of the risk posed by a conjunction event is key within Space Surveillance and Tracking (SST) and SSA services. Establishing a metric that condenses in a concise manner the relevant information pertaining to a conjunction is not straightforward. However, the space industry has been using the probability of collision (PoC) as the ‘foundational element of their collision likelihood assessment’ [9]. The problem of computing the PoC of two space resident objects (RSOs) can be posed as follows: obtain the probability that the distance between two independent and identically distributed samples from the probability density function of the two RSOs at the initial time, falls below a certain threshold within a given time window. Usually, the combined hard body radius of the two objects is used as a conservative value for the collision threshold distance.

In this regard, the so-called Monte Carlo approach is a direct approach to computing the probability by performing the analysis numerous times and counting the proportion of these which result in a collision. However, the number of required samples for estimating the PoC is prohibitive when its value is small (for example, for a probability of the order of 10^{-4} with 10% relative error, 10^6 samples are needed according to [13]).

Leveraging from the specific conditions of in-orbit conjunctions, several alternative methods have been developed to tackle the PoC computation. For a recent literature review, the reader is referred to [14]. In this latter reference, there is a classification of the different methods for computing the PoC, a schematic of which is shown in Fig. 1.

NASA CARA handbook [9] describes two different approaches to compute the PoC:

- i) the ‘two-dimensional’ PoC (2D-PoC) calculation approach, originally introduced by Foster [15] in 1992. This method is a durable simplification of the calculation and is valid in most instances;
- ii) the ‘three-dimensional’ PoC (3D-PoC) calculation method, originally formulated by Coppola [16] in 2012. This method is accurate for nearly every situation, and it is computationally efficient.

The dimension here refers to the number of integrals to compute: in the 2D there are two integrations over a surface (the analysis is performed in a geometrical plane, discarding time and velocity variables); in the 3D there is a third integration over the time variable. Basically speaking, the 2D-PoC applies to ‘short term’ encounters, whereas the 3D-PoC to ‘long-term’ encounters. In the first case, the two RSOs involved have a high relative velocity at the close approach and are thus moving rapidly relative to one another. According to Chan [17], a sufficient condition for short-term encounters is that the relative path is a straight line over a distance of 8-25 km with a deviation of 2 m at most. If this condition is fulfilled, then it can be assumed that: i) the relative position of the two objects is linear; ii) the relative velocity is constant; iii) there is no uncertainty in the velocity; iv) the position uncertainty is constant and equal to the value at the closest approach point. On the other hand, long-term encounters occur when the relative velocity is low and, therefore, the RSOs involved are moving slowly relative to one another, whether intentionally or naturally. In this case, the assumptions of the short-term encounters are no longer valid.

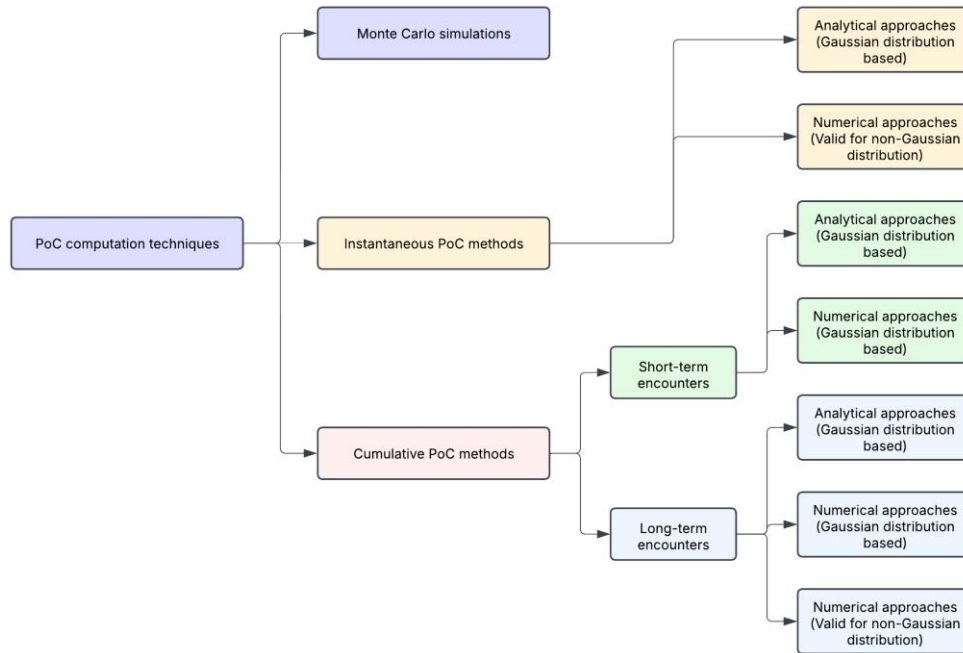


Fig. 1. Collision probability methods, extracted from [14]

2D-PoC methods, analytical or numerical (semi-analytical), are efficient in terms of computational cost and accurate whenever the underlying assumptions are fulfilled. Nevertheless, there exist known situations in which these approaches do not meet operational accuracy requirements. These 2D failure-case geometries usually occur among GEO objects, where relative velocities are typically in the region of m/s to hundreds of m/s and the approaches can be repeating in quick succession every half orbital period. While it is challenging to construct an exact delineation between low and high velocity cases, [18] performed an analysis of 80,453 conjunctions and found that the 2D analytical method produced results that strongly differed from the 3D in only 5.6% of the cases; in another study of 1.73 million conjunctions this figure was found to be 5.3%. This indicates that 2D-PoC methods can inadequately describe some significant events: these are the cases to be described via a 3D or alternative approaches.

As broadly discussed, many algorithms for computing PoC have been published and are widely used in the space industry. Space agencies have developed standard methods and tools to deal with both types of conjunction events. The next subsections review the tools implemented by ESA and NASA and finally present CAS' approach.

2.1.1. ESA CORAM

The Collision Risk Assessment and Avoidance Maneuvers (CORAM) software is one of ESA's tools used for assessing collision risks and computing avoidance maneuvers. Its Collision Risk Computation Software (CORCOS) module addresses the evaluation of the collision risks between two objects. CORCOS resorts to several algorithms to compute the PoC according to the collision speed and the object geometries.

In the case of simple spherical geometries and high-speed conjunctions, the algorithms used are Alfriend & Akella [19], Patera [20], maximum probability assuming spherical covariance using the maximum likelihood approach [21], maximum probability according to Klinkrad's algorithm [22] and covariance scaling (useful when the covariance is not well known).

The use of real geometry instead of the spherical assumption is required to properly estimate the encounters only in the presence of small miss distances and very accurate orbital data (covariance values about 1 m). For those cases, the objects are discretized into multiple small boxes, the PoC is determined for each single box individually (by using [19] or [20]) and then the sum of all the contributions is reconstructed to obtain the total PoC [23].

For low-velocity conjunctions, an interval-slicing method based on Patera's work [24] is applied. The collision interval is divided into several slices, small enough to allow to fall back into the high-speed conjunction solutions. This slicing method holds for both spherical and complex geometries.

Lastly, CORCOS features the Monte Carlo approach as well, used to check and validate the other algorithms.

2.1.2. NASA CARA

For 2D calculations, NASA CARA uses Foster [15] (established standard) and Elrod [25] (very fast technique that builds upon Foster) methods. Additionally, the PoC computed according to Frisbee [26] is used in the case where one object comes without any associated covariance information.

When one of the four assumptions necessary to fall under the 2D approximation does not hold anymore, a different method has to be used. NASA CARA resorts to the Coppola-Hall method [26], also called NoC method because it counts the statistically expected Number of Collisions (NoC). NoC is a 3D method, but a simplified 2D version (NoC-2D) also exists: this latter can be used when one or more of the 2D assumptions fail, but the event can still be treated as a temporally isolated one.

Lastly, NASA CARA includes the Monte Carlo approach as well, but as for the ESA software, it is mainly used for validation purposes.

2.1.3. CAS

As discussed, there are a variety of methods to evaluate the PoC and each of these entails assumptions about the conjunction dynamics. It is clear that there are varying degrees of complexity associated with these various algorithms, with 2D-PoC methods being the least computationally intensive and Monte Carlo the most. As highlighted by Hall [28], it is advisable to first invoke the least computationally intensive method which can produce an accurate PoC estimate for a given conjunction, invoking 3D NoC or PoC methods only in cases where the 2D approach is likely to fail. To this extent, we put in place preliminary steps to check for violations of any assumptions of a given method, prior to the computation of PoC. This is the logic embodied in CAS and shown in Fig. 2.

When it comes to the algorithms integrated in CAS, for the ‘short term’ encounters, we implemented Foster’s method [15], as well as Alfano’s [29] and two versions of Alfriend’s [30]. One of these latter computes the maximum PoC and is the technique used when the covariance information is not available for one of the involved RSOs.

For the ‘long term’ encounters (2D computation not applicable), the need for a semi-analytical approach has been addressed by integrating an operational implementation of the Hall’s method [27]. To date, this is the algorithm that has been proved to yield the most accurate results at a reasonable computational cost.

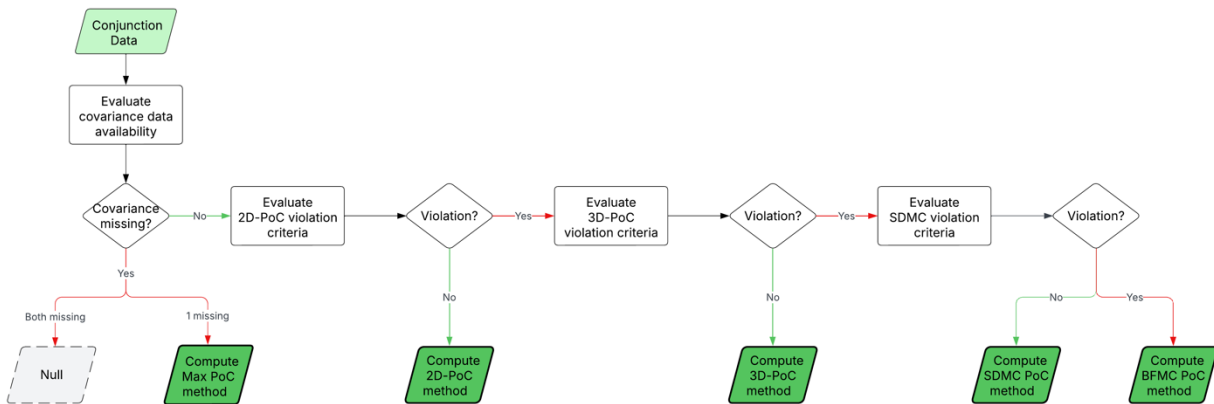


Fig. 2. Workflow showing the choice of the algorithm for the PoC computation

Furthermore, and initially intended for validation purposes, Monte Carlo implementation was integrated as well in an agnostic manner. This means that the user can specify as input the conjunction data at the Time of Closest Approach (TCA) and use a simplified two-body (Keplerian) dynamics to mimic the Simplified Dynamics Monte Carlo (SDMC) method shown in [28], but also provide initial conditions at estimation epoch for the primary and secondary to run a high-fidelity simulation analogous to the Brute Force Monte Carlo (BFMC) approach [28].

2.2 Collision avoidance maneuver service

According to the usual procedures of satellite operators, a service for collision avoidance can be designed for providing two outputs: a trade-off analysis of the possible available maneuvers and a detailed description of a maneuver to be implemented onboard, compatible with all the platform and ground segment constraints. Following this approach, CAS can assist in the decision-making process, as well as providing a collision avoidance maneuver solution

compatible with all the operating constraints of a satellite. These are two quite different objectives that are usually treated with different mathematical methods. The decision-making support methods are based on heuristic methods that explore the solution space using surrogate models, so that the availability of the necessary information for decision-making is balanced with the associated computational cost. Once the decision to carry out a collision avoidance maneuver has been made, the mathematical methods used to obtain the optimal maneuver according to the stated objective and in accordance with all the existing constraints, are based on optimal control theory. Although there are different options to solve the resulting optimal control problem, a direct transcription method has been chosen given its flexibility when adapting to different constraints and maturity in terms of the available tools. In exchange, the computational cost can be potentially high. Having an adequate initial guess can be decisive in reducing it.

Considering the different methods involved in the problem resolution, a two-step approach has been devised. The User will have the information about the possible CAMs at the end of the first step, helping in the decision-making process. The second step provides the User with the propulsion control inputs to fly a safe trajectory compliant with all the constraints. The detailed process flow is described in Subsection 3.2.3.

The first step is, thus, posed as a multi-objective optimization problem. Following a literature review, a genetic algorithm has been identified as the best solution for this problem. Specifically, the algorithm Non-dominated Sorting Genetic Algorithm-II (NSGA-II) [31] has been selected. NSGA-II is a non-sorted algorithm that provides the Pareto front of all non-dominated optimal solutions. The performance of the individual solutions is used in the selection process. The evaluation of the performance is done computing the objective functions of each individual. For this purpose, surrogate models are used to reduce computational time. The Pareto front is the output of the first step, providing a global view of the solution space to the operator. The mathematical description of the first step can be found in [32].

The second step is posed as an optimal control problem. A direct transcription method is used to solve it. On a case-by-case basis, the provided User Requirements (UR) specify the constraints to be considered in the design of the CAM. All these constraints are incorporated into the optimal control problem and the direct transcription method allows us to select or unselect them in a straightforward way. Requirements can be added or removed from the problem without modification of the solution process. This is a great advantage with respect to indirect methods, which require a reformulation of the problem. The direct transcription method translates the optimal control problem into a non-linear programming problem. In a general non-linear program, under certain conditions, the Karush–Kuhn–Tucker (KKT) conditions [33] provide necessary conditions for a solution to be optimal. The mathematical description of the second step, and the corresponding mapping with the UR can be found in [32].

The algorithm process flow is presented in Fig. 3, where it is highlighted that the CAM algorithm is implemented as a two-step process, where the first step uses a genetic algorithm to obtain the Pareto front of a multi-objective optimization, and the second step is a gradient-based optimization. Additional information on the process flow can be found in [32].

3. CAS

As mentioned above, the NorthStar Conjunction Analysis System (NS CAS) is envisioned as a software system that integrates essential building blocks related to conjunction analysis and collision avoidance. The overarching objective of CAS is to produce accurate and actionable estimates of risk for upcoming conjunction events (days before the events will be realized) and propose avoidance maneuvers in case the risk is above the operator's threshold. By generating meaningful early warnings, additional input data can be requested and advanced methods can be applied to improve the assessment quality. When NorthStar, its customers, or its data consumers (e.g., satellite constellation owner/operators) are confident in the conclusion of a high-risk event, the parties can plan maneuvers and act towards collision avoidance or risk mitigation.

The building blocks are software services that implement all the required functionalities: catalog management, orbit and covariance propagation, conjunction screening, conjunction assessment with the computation of the relevant risk metrics (PoC, miss distance, Mahalanobis distance, etc.) and a smooth integration with the collision avoidance module to provide suggested avoidance paths if needed. Other advanced functionalities include the CDM refinement to exploit

all the available data for the conjunction assessment or orbit improvement using additional observations of the secondary RSO in a conjunction. All these functionalities are briefly reviewed in section 3.2.

CAS integrates these components, along with their external upstream and downstream dependencies, into a coherent and optimized solution that answers to the conjunction analysis needs of NorthStar customers.

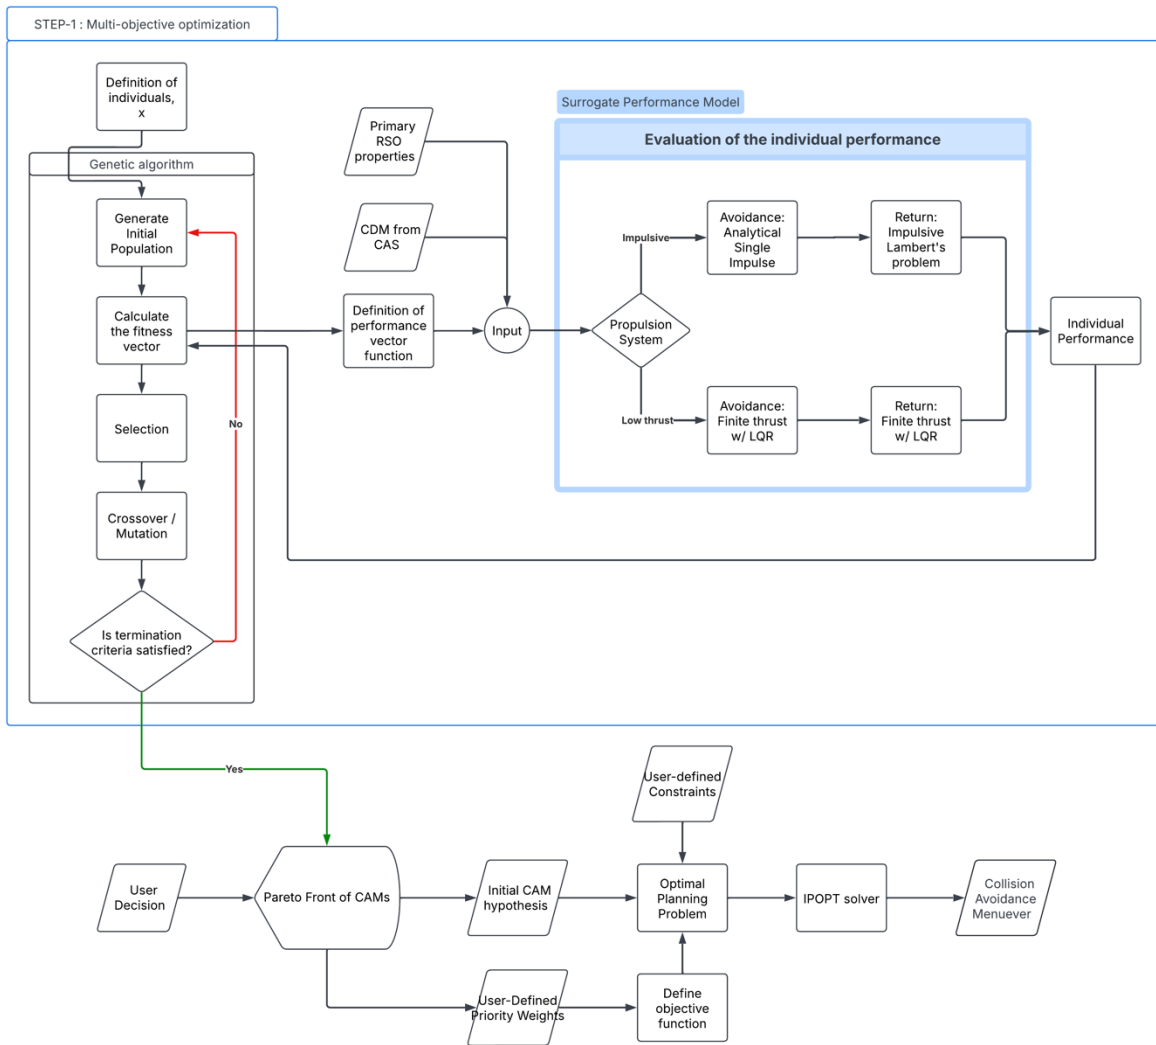


Fig. 3. Scheme of the flow of information in the CAM service algorithm

3.1 CAS Architecture

The NS CAS Platform has been designed as a cloud-based system with an Infrastructure as Code (IaC) philosophy. The core business functionalities are achieved via a Service Oriented Architecture orchestrated with a central workflow engine. The system deployment is comprised of four key entities:

1. Apps of Apps. A definition for an App of Apps is a ‘pattern that allows to define a root ArgoCD application which, rather than pointing to application manifest files, points to a directory containing application spec files for all micro-services’
2. Services (sourced as Helm charts from our services repository and deployed via apps of apps). CAS services provide the business logic for the conjunction analysis system and are orchestrated together by the workflow engine
3. Platform Dependencies (sourced as Helm charts from dedicated repositories). Platform dependencies are tools maintained by the NS team which do not provide direct business functionality but are key for supporting the system. They provide orchestration for enabling CAS workflows, as well as maintenance support via observability (e.g. monitoring, alerting)

4. Infrastructure dependencies (pure Terraform deployment). Finally, there is an underlying infrastructure to provide safe data persistency to services via databases and object storage for Temporal (workflow) history and visibility

All these types of deployments are managed through a dedicated repository, which provisions infrastructure and platform components using Terraform, and handles services via an ArgoCD App of Apps.

3.2. CAS key functionalities

In the following subsections, the key functionalities of CAS are described. Specifically, we will introduce the concept of Screening and Conjunction assessment (Subsection 3.2.1), CDM refinement (Subsection 3.2.2), Collision avoidance maneuver recommendation (Subsection 3.2.3) and lastly, Orbit improvement (Subsection 3.2.4).

3.2.1 Screening and Conjunction assessment

The core objective of CAS is to identify for a list of RSOs of interest (the High Value Objects (HVOs)) all the close approaches under a certain threshold distance in a given period of time and then to compute their PoC. This process consists in running a so-called N vs. All analysis, in which we look for close encounters of each object in the HVO list against all the objects in the RSO catalog, resorting to the basic tools of CAS, like the orbit propagator and interpolator, the TCA finder, and the PoC algorithms (described in Section 2).

The computation of the TCA and the corresponding Distance of Closest Approach (DCA) for a given pair of RSOs is not a computationally intensive process but considering the size of the space catalog (including tens of thousands of objects), the computational load escalates quite quickly. To this extent, the screening functionality helps in removing all the object pairs for which a conjunction is not possible before running any heavy calculation. Fig. 4 illustrates some of the key steps of the screening and conjunction assessment functionalities.




Fig. 4. Visualization of the core concepts of CAS screening and conjunction assessment

Specifically, Panel A) shows CAS starting point: from the orbit determination process, we know the estimated state of each RSO in our catalog at a specific epoch, where the epoch is generally different, but recent, for each RSO. The drawing shows only three RSOs, but it could be extended to show an entire catalog. The state of each RSO could be described by a position and velocity vector (i.e., a state vector) or by other commonly used representations, like the Keplerian orbital elements or the Two-Line Elements sets (TLEs). In the case of conjunction screening, we often use the TLE representation, even though this one is known for not being highly accurate, because we can afford a small loss in precision as a trade-off for the computationally efficient data manipulation it offers.

Panel B) depicts the concept behind our first screening algorithm: the Apogee-Perigee filter. For each of RSOs 1, 2, and 3, an elliptical orbit is drawn to show potential points of intersection. At this stage, all we need to know are the approximate apogee and perigee for each object, that can be quickly computed from the state vectors or TLEs. We then generate all possible RSO pairs based on the catalog and retain all pairs that satisfy the following criteria: the difference between the maximum perigee and the minimum apogee is less than a configurable threshold.

Panel C) zooms in on the specific case of applying the Screening Volume filter to the conjunction pair RSO1-RSO3. We have constructed a screening volume (of configurable dimensions) and propagated the initial orbital state of each RSO throughout a specified screening period. Since we do not have to find an exact TCA or compute an exact probability of collision at this stage, we can use a coarse-grained ephemeris and ‘fill the gaps’ with an interpolator. In the figure, we observe the case where the secondary RSO will soon intersect with the primary RSO’s screening volume. This intersection means that we have to retain the conjunction pair for a detailed assessment, whereas we would have filtered out the pair if there was no intersection. The times at which the secondary RSO enters and then subsequently leaves the screening volume provide a rough window within which we should find the TCA. This will be helpful in the next step, i.e., the conjunction assessment.

Apogee-Perigee and Screening Volume are the main filters we apply before assessing the conjunctions in detail. In the literature, there are additional filters that could further reduce the number of pairs to analyze (e.g., the orbital path filter for which we have an implementation as well), but due to the high parallelization in place at software level, we already satisfy our timing constraints and we do not need to resort to additional filters. The schematic of the Screening Service software design is shown in Fig. 5.

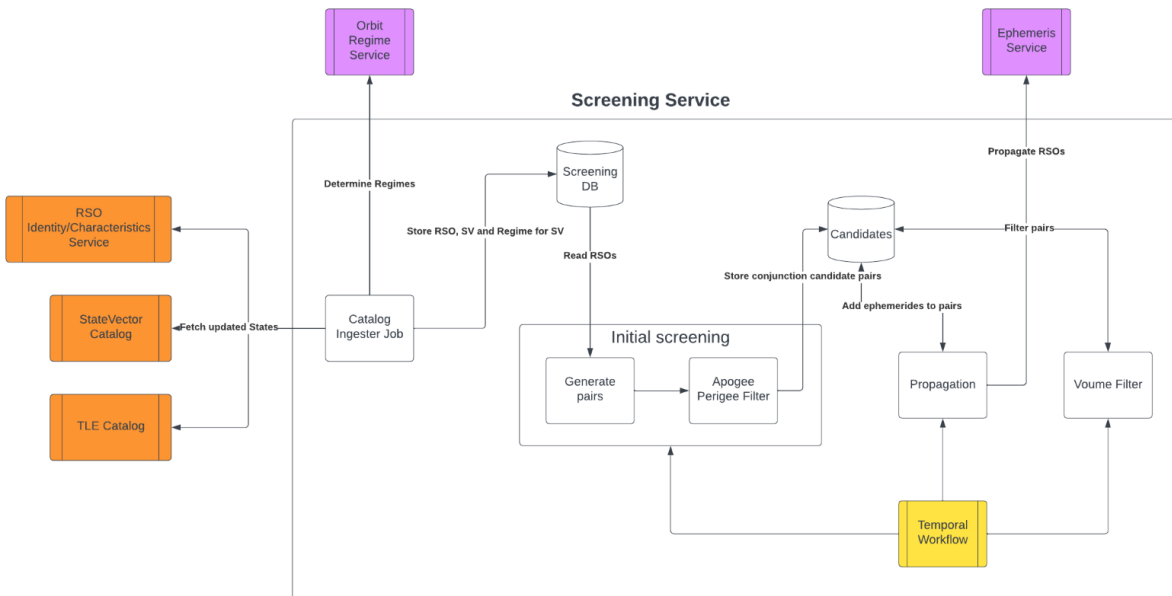


Fig. 5. Schematic of the Screening Service software design

Following the screening functionality, the conjunction assessment (also called 1 vs. 1 analysis) takes place between all the retained pairs for which we now know that there is an upcoming interesting conjunction event. Panel D) in Fig. 4 shows the next step to perform, i.e., generate an accurate ephemeris by propagating the RSO state via a numerical propagator and no longer resorting to the TLE representation. Since the screening filter already provided the TCA window, we know that the fine-grained propagation must only be performed until the end of the window, and that we do not need to output each ephemeris step before the start of the window (even if the propagator will compute them internally). For this stage of the analysis, the propagator must also evolve the state uncertainty of each RSO so that we can use the covariance matrices at TCA as a key input for the PoC algorithms. As shown in the image, the uncertainty (covariance ellipsoid) grows as we propagate the state further forward in time. It is for this covariance minimization reason that we are interested in having as many recent observations of each object in our catalog as possible.

At this point, the first goal of the conjunction assessment functionality is to refine the estimate of TCA between the two RSOs in the conjunction pair, and their corresponding DCA. This task is executed through an algorithm that is based on the Orekit library [34] and takes the previously computed ephemerides as input. The algorithm can be

data (as introduced in [12]). In [12], these matrices were computed starting from the TLE history, but the same is applicable using the more accurate Special Perturbation (SP) catalog from the 18th Space Defense Squadron as data source (which is the actual source used by CAS at conjunction assessment/refinement level).

Fig. 7 shows the details of the CDM refinement flow. In a first step, the latest known states for both objects are loaded, as well as any operator ephemerides (if provided) to refine the primary and secondary state vectors using a Batch Least Squares filter. Based on these inputs, refined ephemerides are propagated to a time window around the ‘presumed TCA’ (the TCA of the initial CDM). Since a conjunction is assumed in this case, no screening filters are applied, and the next step is to calculate the exact TCA that matches the input CDM. After this step, the covariances are propagated to the new TCA and the probabilities of collision are calculated, resulting in a refined CDM, which is saved in the Conjunction Database. Note that since no re-screening is performed before evaluating the conjunction, our assessment might result in very low PoC values if the latest information available on one of the RSO is showing a significant change with respect to the original CDM (meaning that the conjunction would no longer take place).

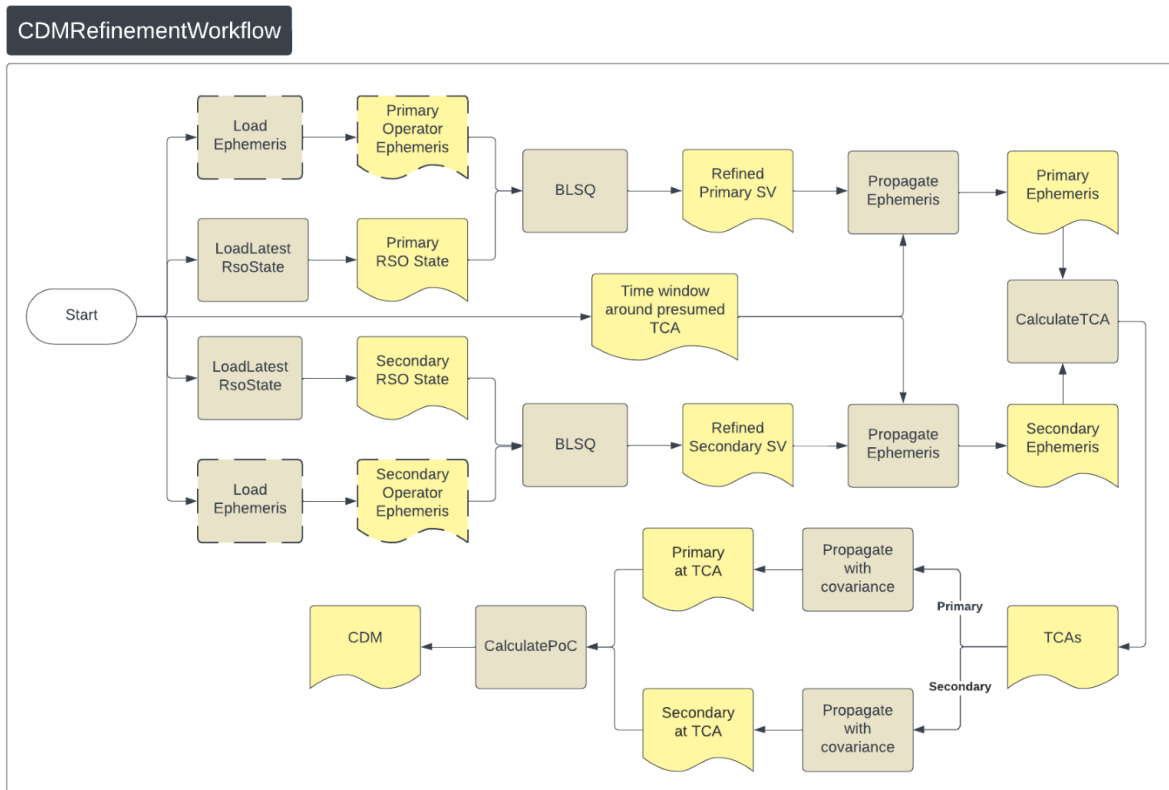


Fig. 7. CDM Refinement workflow

3.2.3 Collision avoidance maneuver recommendation

Fig. 8 describes the user flow for CAM. There are two envisioned ways to use CAM: i) creating a CAM scenario via a button inside CAS, thus creating a scenario directly linked to that specific conjunction event; ii) manually creating the scenario and uploading a CDM. After the scenario has been created, its settings will automatically be prefilled with the available information from the conjunction event. After this, the user will be asked to configure the relevant spacecraft settings and constraints. CAM will offer the possibility to define default settings for each spacecraft so that these settings can also be prefilled and just require confirmation from the user. In the next step, a new CAM assessment will be automatically created, linked to the relevant conjunction assessment in CAS (if the scenario was created via CAS). When CAS creates new conjunction assessments for an event, CAM will automatically create new corresponding CAM assessments. After the CAM assessment creation, the first step of the CAM optimization will be automatically started. After completing the first optimization step, the user will be notified, and the Pareto front will be visualized. The user now chooses one or multiple sample points on the Pareto front, which will automatically create

a Proposed Maneuver for each selected sample point and run the second CAM optimization step to give a maneuver solution. Once this has finished, the user will be notified again and presented with the optimized maneuver.

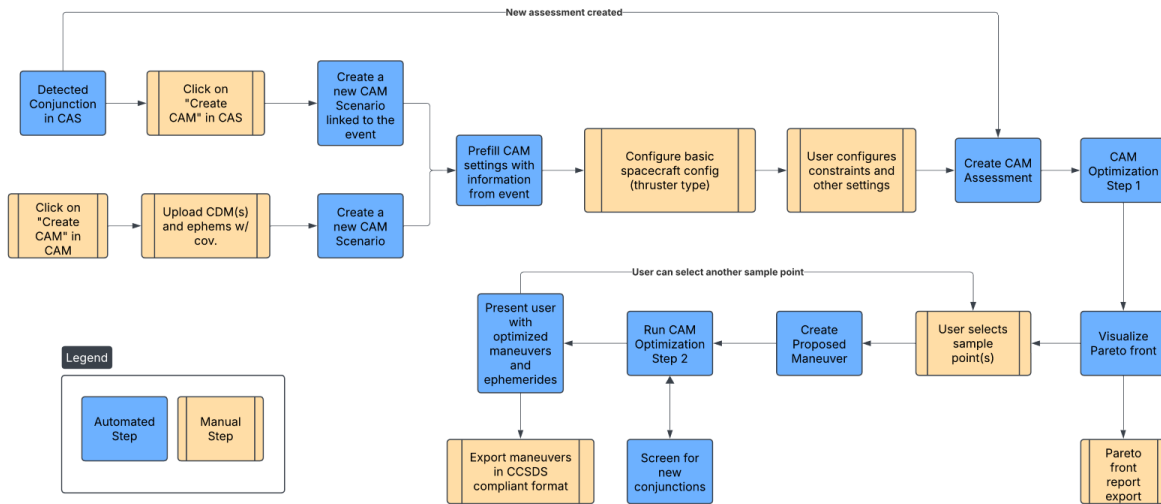


Fig. 8. General User flow for CAM

As introduced above, the CAM application will employ the concept of scenarios (i.e., the CAM scenarios). A CAM scenario is linked to a conjunction event, either one detected by CAS or an event uploaded by the user. As such, the lifetime of a CAM scenario is until the conjunction event has been resolved, after which it is no longer needed. The CAM scenario contains settings about the spacecraft, like thrusting capability, as well as the operational constraints that apply. The CAM scenario holds multiple CAM assessments which are in turn linked to conjunction assessments. The CAM assessment represents the maneuver planning related to a specific assessment of the conjunction event. For each assessment, the Pareto front is generated and stored. Finally, each CAM assessment holds multiple Proposed Maneuvers, which represent an optimized maneuver based on the user choice of the solution in the Pareto front.

3.2.4 Orbit improvement

PoC is the preferred metric for assessing risk in conjunctions, as mentioned earlier in the paper. However, researchers have identified a critical flaw called probability dilution, where reduced data quality paradoxically decreases calculated collision probabilities [35]. This phenomenon stems from false confidence in probabilistic models, where certain propositions maintain high belief regardless of the truth. In cases in which the uncertainty on the secondary in a conjunction is too high, the way the conjunction assessment can be safely done is by obtaining additional observations for that object. CAS has implemented the capability to ingest on-request observations from SST data providers and produce a new OD solution combining the catalogue information and the newly acquired data in a more accurate solution allowing for a safer conjunction assessment. This process is what we call within CAS ‘Orbit improvement’.

4. Results

This section will present the current implementation of CAS via several snapshots coming from the latest graphical user interface. At the same time, one selected event is analyzed in some detail to briefly explain the logic behind the conjunction analysis operations.

Similar to the CAM application, CAS is built upon the concept of scenarios. The scenario settings allow for definition of the HVOs, i.e., the list of RSOs of interest that we want to monitor, as well as for the setup of the analysis variables. These latter include the time schedule and the screening options and are fully customizable to the client’s needs. Two functionalities can be enabled or disabled at the scenario level: the screening and the CDM refinement ones. If the screening is disabled, the scenario will create conjunction events and assess their risk only for the RSO pairs (with one object belonging to the HVO list) analyzed in the CDMs retrieved from third parties (mostly from Space-Track.org website).

Once a scenario is created, run for the first time and selected, CAS will present the user with the scenario home page, shown in Fig. 9. In this interface, useful information is displayed both in a graphical and a tabular format. On the top graph, all the Conjunction Events (CEs) detected for the selected HVOs are displayed plotting their PoC against

the TCA. The PoC is computed using the method that is indicated right on top of the plot (Foster1992 in the case of Fig. 9). According to the computed PoC, the events are classified in different categories of risk level. In the plot, we can see blue squares representing the ‘Info’ risk level, yellow pointing-up triangles for the ‘Warning’ risk level and finally orange pointing-down triangles for the ‘Alert’ risk level. In addition to these and not visible in Fig. 9, there is the ‘Critical’ risk level as well, represented by a red diamond. Clearly, different levels of risk correspond to different actions that the operator responsible for the platform will have or not to take to ensure a good CE resolution. In the bottom table, the same kind of information is represented with some additional details. Various sorting and filtering mechanisms can be applied to the data listed here. On the right side of Fig. 9, three boxes can be observed. These boxes appear only after selecting one of the table rows. In here, information about the CE is displayed together with details on the Primary and Secondary RSOs (including amongst others the RSO NORAD ID, the object class and the hard body radius). Additionally, a timer to the TCA (constantly updating) is also displayed.

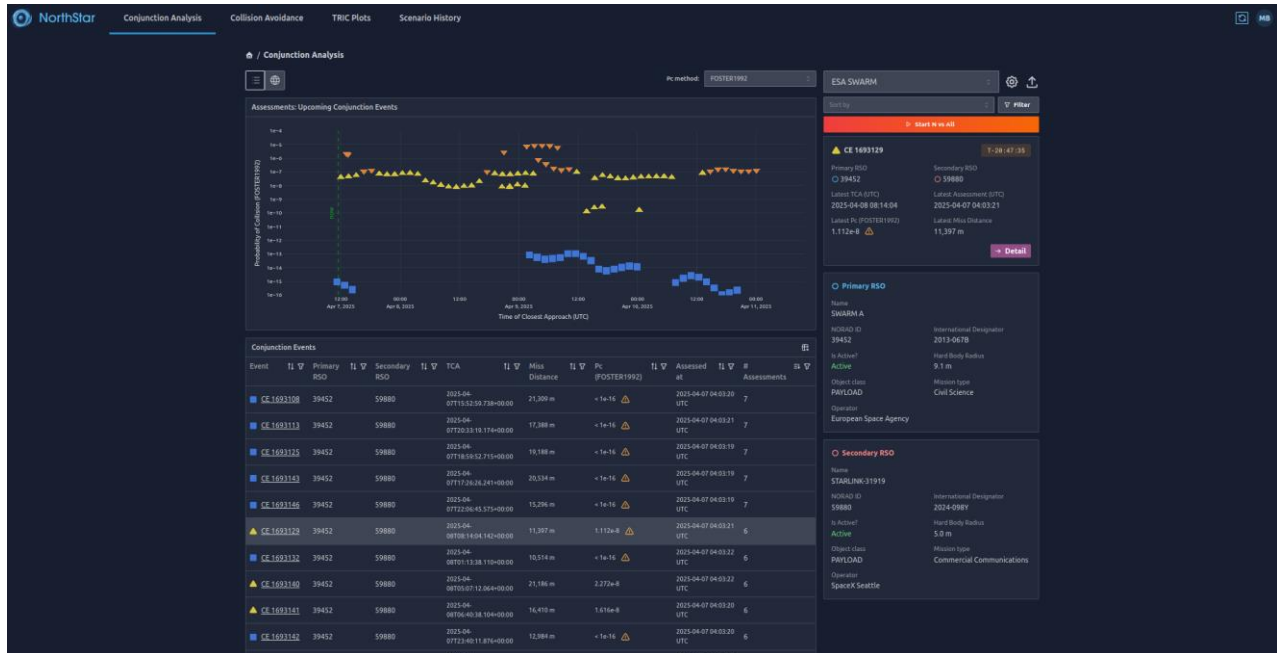


Fig. 9. CAS scenario home page

In what follows, the tool will be presented by analyzing one selected CE. Let us assume that of all the events listed in the table visible in Fig. 9 the user/operator is interested in a specific one. He/she clicks on it, thus opening the specific CE page. This latter presents again the CE and RSO information as in Fig. 9, with some main additions that allow the user to perform several important functions such as uploading an OEM file for the primary RSO via the ‘Upload Flight Plan’ button, requesting additional measurements on the secondary RSO via the ‘Improve Orbit’ button or contacting the operator for escalation processes via the ‘Contact Operator’ button.

Additionally, this page presents the user with several plots that make it a very important page for the analysis of the conjunction. Here we can indeed find the Evolution of the PoC (Fig. 10), DCA (Fig. 11) and TCA (Fig. 12) with respect to the Time of the assessment and finally a graph representing the relative distance between the involved RSOs against the time both in norm and separated on the three components in the Radial, In-track and Cross-track (RIC) reference frame.

The evolution plots can include a varying number of points (starting from one) according to how many Conjunction Assessments (CA) of the same event have been analyzed up until the moment the plot is consulted. To have a clearer view of this information, the list of all the CAs corresponding to the same CE is presented as well in the ‘Assessments’ tab. In the selected case, eleven different assessments of the same event were performed in a time period spanning a bit more than three days.

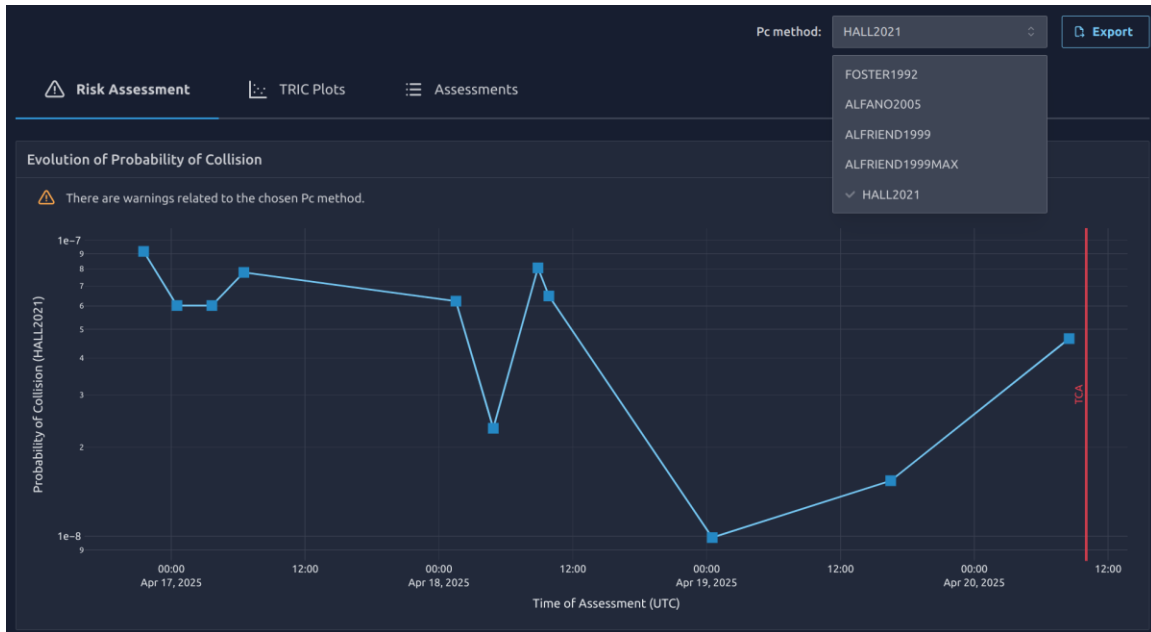


Fig. 10. Evolution of Probability of Collision

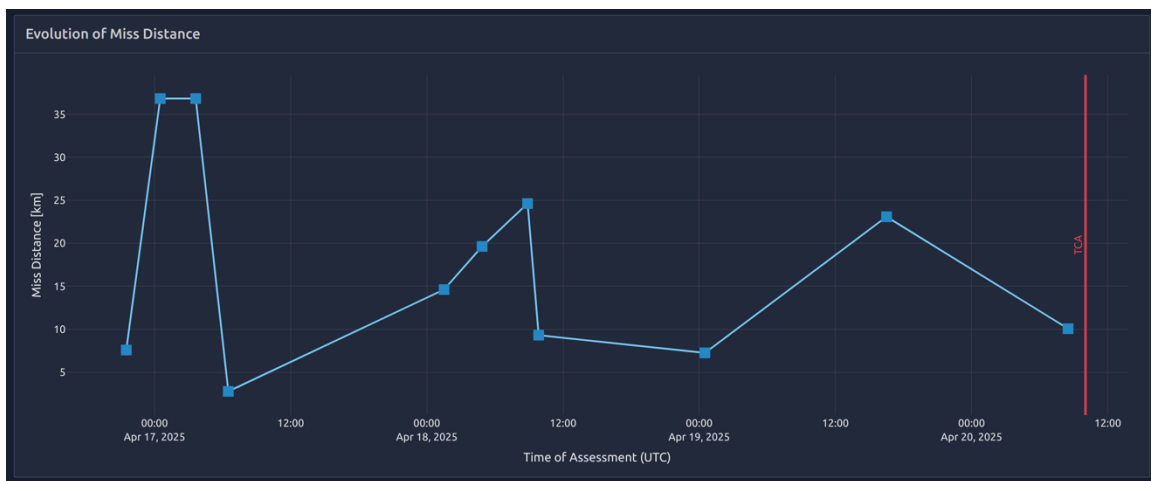


Fig. 11. Evolution of Miss Distance (or Distance of Closest Approach)

Fig. 10, Fig. 11 and Fig. 12 show how the PoC, DCA and TCA change over time as new information on either the primary or the secondary RSO becomes available. These plots reveal that the schedule chosen for this specific scenario implied computing an assessment every three hours. This is not something that happens blindly though: CAS is only running a new assessment if new information on either of the objects with respect to the previous estimate is accessible (this is also the case between the second and third point of the graphs, even though the difference looks minimal from the figures).

Since the selected event represented a CE between two GEO objects, the chosen PoC method is Hall2021, but Fig. 10 illustrates how all the different PoC methods introduced in Subsection 2.1.3 are available for selection as well. Along the different assessments, the PoC oscillated between $1e-8$ and $1e-7$, classifying this event at a 'Warning' level never exceeding the threshold to further escalate it. As depicted in Fig. 11, the DCA instead varied between 2 and 37 km: while these values might not seem very small for a conjunction assessment to a LEO operator, for a GEO conjunction they instill already some concerns. Fig. 12 shows how the TCA remained more or less constant over all the assessments oscillating within a range of 15 s only.

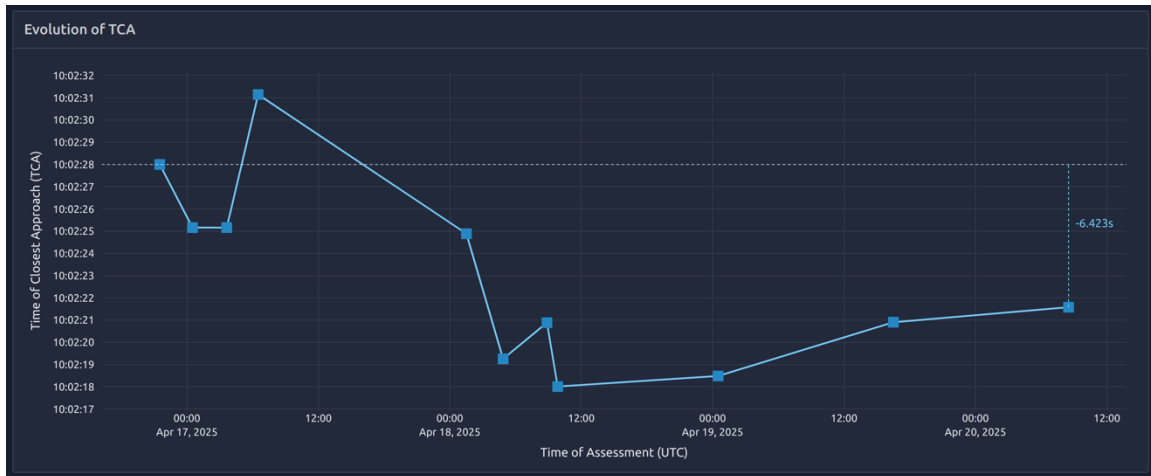


Fig. 12. Evolution of Time of Closest Approach

By clicking on any of the single assessment points in the evolution plots, the operator is redirected to the overview of the assessment itself. This is another important page for the conjunction analysis process. Here indeed, one can find a representation of the B-plane that depicts the relative encounter geometry for the CA. In this plot, visible in Fig. 13, the covariance matrix is represented using three different scale factors (from one to three standard deviations) and the operator can see where the secondary object falls with respect to the uncertainty of the assessment. In the selected example, the secondary RSO is located quite close to the primary (within the one standard deviation scaled covariance) making it an interesting event to follow. This plot is useful as well to control the size of the covariance matrix (i.e., the uncertainty on the objects). As discussed in the paper, when this is too large, the assessment results lose significance, making it necessary to request new measurements to reduce the uncertainty. This is the reason why this plot presents the button to request additional measurements on the secondary object via the ‘Orbit improvement’ functionality.

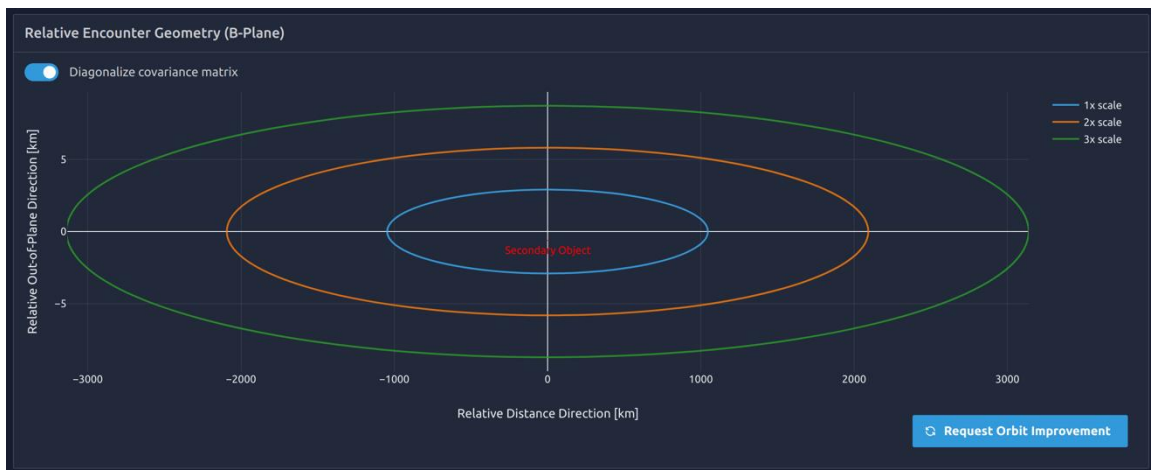


Fig. 13. Representation of the B-plane, indicative of the relative encounter geometry

Following up, in the single assessment page, there is the so-called Time, Radial, In-track, Cross-track (TRIC) plot which shows the relative position between the two RSOs involved in norm and in the three components of the RIC reference frame (in these subplots the green and red diamonds represent the start and end time, respectively). This plot, presented in Fig. 14, is particularly meaningful to understand the relative orbital geometry between the two objects. In this specific example, for instance, this plot clarifies that the objects have a very strong cross-track separation over time indicating that the inclination of the two orbits is quite different.

Finally, there are two additional graphs in this tab which are not shown in the paper. The first one is a simplified 3D relative position plot, whereas the last one depicts the Camera Target Sun (CATS) angle over time. This last

quantity helps understanding the visibility conditions for the primary RSO over the secondary one with respect to the Sun position.

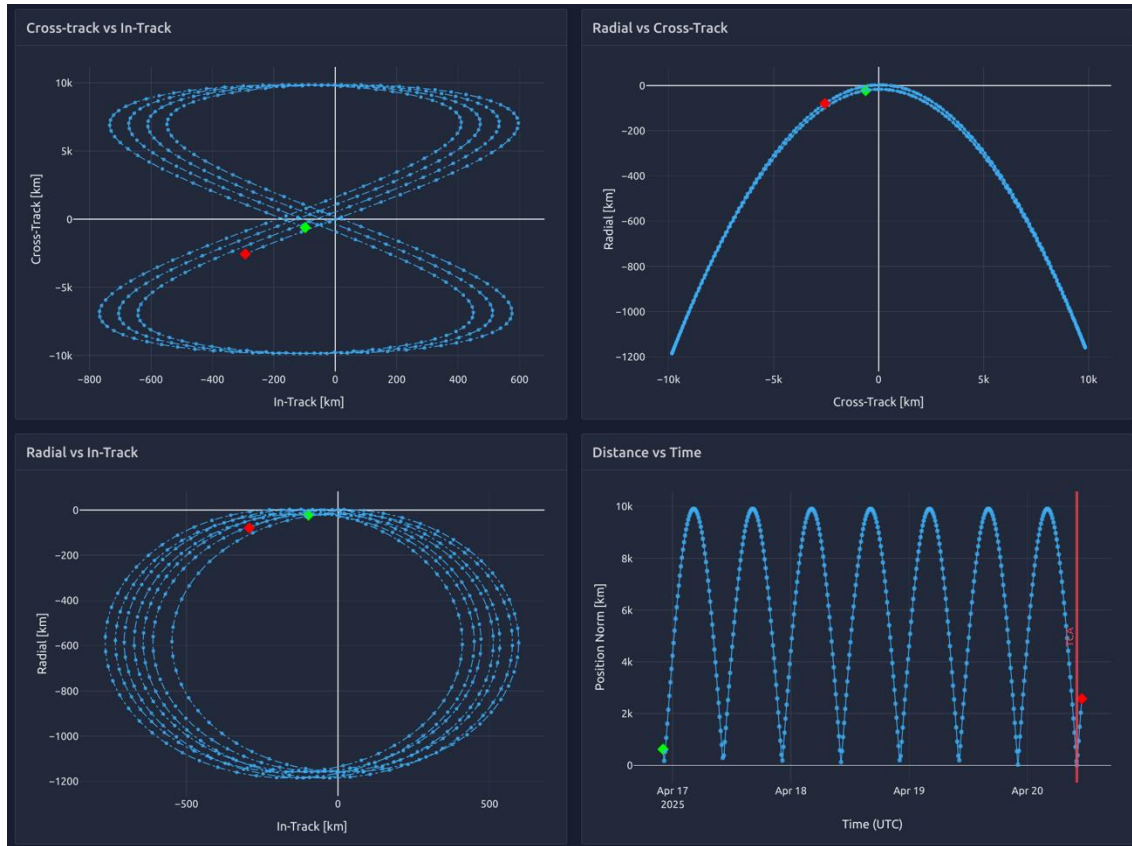


Fig. 14. TRIC plot, indicative of the relative orbital geometry

In the operational context, the thresholds used to decide when to escalate an event or to plan a CAM are customizable for each customer and the conjunction analysis operator will adhere to them. Let us briefly describe here a possible course for the event introduced above. The two GEO objects under evaluation have a predicted close approach on April 20th at 10:02. CAS platform detected this event (based on the screening settings used in this scenario) on April 16th at 21:30, i.e., more than three days in advance of the TCA. Since the PoC is decently high (at ‘Warning’ level), when three days are left to the TCA, the operator informs the customer of the upcoming event, detailing that there is no action to take as of now, but that CAS will follow and monitor the event in the upcoming days. As time goes by and we approach the TCA, in this specific case, the new assessments show that nothing has significantly changed and as a consequence, there is no need for any sort of escalation for the event. As a rule of thumb (but again customizable by the client), the moment to start escalating an event into CAM planning comes around one day ahead of the TCA: if the PoC is above the threshold of $1e-4$ by that time, it is necessary to take action. The operator would have immediately started a CAM analysis and called a meeting with the customer to explain the situation in detail. In fact, ultimately, the maneuver–no maneuver decision is always the responsibility of the satellite owner/operator.

5. Conclusions, Outlooks

As analyzed throughout this work, NS CAS represents a critical advancement in ensuring the safety and sustainability of satellite operations in the increasingly congested orbital environment. By integrating real-time risk assessment, optimized avoidance maneuver planning, and on-request orbit refinement, CAS is capable of delivering actionable insights that help protect valuable space assets and extend mission lifespans.

The assessment of the collision risk relies on the correct computation of the PoC. To this end, CAS employs state-of-the-art algorithms whose selection was motivated following a similar logic as in ESA and NASA in-house dedicated software. In future, a statistical comparison of all the different methods will be carried out.

In this paper, our innovative CAM strategy has been introduced; CAS architecture and key functionalities (ranging from screening to orbit improvement) have been thoroughly dissected and finally a comprehensive presentation of the current tool with its latest graphical user interface has been provided.

Undoubtedly, the tool is still evolving and not yet in its final version: as a clear example, the possibility of invoking the planning of a CAM from the interface is not yet available, even though the strategy is already defined and has been presented in [32]. Furthermore, the algorithm lying behind the proposed CAM approach has been developed and the infrastructure has been already prepared for the upcoming integration into the tool.

In future, when NS space-based satellite constellation will be capable of providing enough data to allow us to maintain our own catalog (complementing the publicly available ones with more accurate and frequent updates), this data will be included in CAS, helping it stand out from the solutions currently available on the market. In the meantime, however, we need to employ different data sources in the conjunction assessment process. The best practice guidelines [9] state that using the publicly available Two-Line Elements sets (TLEs) is not a recommended practice despite some entities still resorting to them. The problem with this approach is that the TLE accuracy is not sufficient to perform the necessary conjunction assessment calculations. As described in this paper, in CAS, we utilize the TLE catalog only to cold start our computation in the initial steps of our screening process (applying larger thresholds to make up for the lost accuracy). This type of data indeed still provides some advantages such as the extremely low computational cost required for the propagation. For the actual conjunction assessment process though, we switch to the more accurate SP catalog from the 18th Space Defense Squadron (trusted industry-standard data source).

To date, NS is performing a validation campaign against existing systems providing Conjunction Analysis services to prove that CAS is ready to be used at an operational level. In this respect, we are collaborating with international partners to compare our system's performance against the existing warnings and CDMs from the Space-Track.org website as well as from our partner's solutions, with the aim to demonstrate improvements in accuracy and reliability.

In the meantime, while some commercial satellite operators are already in the onboarding phase with CAS, we are setting in place foundational agreements with selected ground-based optical sensor providers, ensuring downstream data fusion and enhanced system capabilities (especially required for the orbit improvement functionality).

Acknowledgements

The authors would like to acknowledge that the project is supported by LuxIMPULSE funding, managed by the Luxembourg Space Agency and implemented by the European Space Agency.

References

- [1] ESA, Space debris by the numbers, 21 March 2025, https://www.esa.int/Space_Safety/Space_Debris/Space_debris_by_the_numbers (accessed 30.03.25)
- [2] E. Kulu, NewSpace Index, <https://www.newspace.im/> (accessed 30.03.25)
- [3] E. Kulu, Satellite Constellations - 2024 Survey, Trends and Economic Sustainability, IAC-24.E6.1.13, 75th International Astronautical Congress (IAC 2024), Milan, Italy, 14-18 October 2024
- [4] B. Bastida Virgili, J.C. Dolado, H.G. Lewis, J. Radtke, H. Krag, B. Revelin, C. Cazaux, C. Colombo, R. Crowther and M. Metz, Risk to space sustainability from large constellations of satellites, *Acta Astronautica* 126 (2016) 154-162
- [5] D. J. Kessler and B. G. Cour-Palais, Collision frequency of artificial satellites: the creation of a debris belt, *Journal of Geophysical research* 83-6 (1978) 2637-2646
- [6] 18th & 19th Space Defense Squadron, Spaceflight safety handbook for satellite operators. Version 1.7, April 2023
- [7] SAIC, Space-track, <https://www.space-track.org/> (accessed 30.03.25)
- [8] National Aeronautics and Space Administration, Conjunction assessment, <https://www.nasa.gov/cara/> (accessed 30.03.25)
- [9] F. J. Krage, NASA Spacecraft Conjunction Assessment and Collision Avoidance Best Practices Handbook, February 2023
- [10] B. Lal, A. Balakrishnan, B. M. Caldwell, R. S. Buenconsejo, S. A. Carioscia, Global Trends in Space Situational Awareness (SSA) and Space Traffic Management (STM), April 2018
- [11] M. Hejduk, D. E. Snow, Setting the Bar for the Replacement of the Probability of Collision Metric, 29th International Symposium on Space Flight Dynamics (ISSFD 2024), Darmstadt, Germany, 2024, 22-26 April

- [12] G. Escribano, P. Rajasekar, Y. Khatri, N. Gollu, S. Setty and M. Sanjurjo-Rivo, Statistically Consistent Synthetic Covariance to Quantify Operational Uncertainty, 9th European Conference on Space Debris (ECSD 2025), Bonn, Germany, 2025, 1-4 April
- [13] J. Morio, M. Balesdent, D. Jacquemart and C. Vergé, A survey of rare event simulation methods for static input-output models, *Simulation Modelling Practice and Theory* 49 (2014), 287–304
- [14] J. Li, Y. Zhen and L. Ya-Zhong, A review of space-object collision probability computation methods, *Astrodynamics* 6 (2 2022), 95–120
- [15] J.L. Foster and H.S. Estes, A parametric analysis of orbital debris collision probability and maneuver rate for space vehicles, 1992 NASA/JSC-25898
- [16] V.T. Coppola, Including velocity uncertainty in the probability of collision between space objects, AAS 12-247, 22nd AAS/AIAA Space Flight Mechanics Meeting, Charleston, South Carolina, USA, 2012, 29 January - 2 February
- [17] K. Chan, Short-term vs. long-term spacecraft encounters, AIAA 2004-5460, AIAA/AAS astrodynamics specialist conference and exhibit, Providence, Rhode Island, USA, 2004, 16-19 August
- [18] D. T. Hall, M. D. Hejduk, and L. C. Johnson, Time Dependence of Collision Probabilities During Satellite Conjunctions, AAS 17-271, 27th AAS/AIAA Space Flight Mechanics Meeting, San Antonio, Texas, USA, 2017, 5-9 February
- [19] M. R. Akella and K. T. Alfriend, Probability of collision between space objects, *Journal of Guidance, Control and Dynamics* 23-5 (2000) 769-772
- [20] R. P. Patera, General method for calculating satellite collision probability, *Journal of Guidance, Control and Dynamics* 24-4 (2001) 716-722
- [21] F. K. Chan, *Spacecraft collision probability*, Aerospace Press, 2008
- [22] H. Klinkrad, *Space debris: models and risk analysis*, Springer Praxis Books / Astronautical Engineering, 2006
- [23] N. Sanchez-Ortiz, *Methodologies for collision risk computation*, Doctoral Dissertation, Polytechnic University of Madrid, 2015
- [24] R. Patera, Satellite collision probability for nonlinear relative motion, *Journal for Guidance, Control and Dynamics* 26-5 (2003) 728-733
- [25] C. Elrod, *Computational bayesian methods applied to complex problems in bio and astro statistics*, Doctoral Dissertation, Baylor University, 2019
- [26] J. Frisbee, An upper bound on high-speed satellite collision probability when only one object has position uncertainty information, AAS 15-717, AAS/AIAA Astrodynamics Specialist Conference, Vail, Colorado, USA, 2015, 9-13 August
- [27] D.T. Hall, Expected collision rates for tracked satellites, *Journal of Spacecraft and Rockets* 58-3 (2021), 715-728
- [28] D. T. Hall, L. G. Baars and S. J. Casali, A Multistep Probability of Collision Computation Algorithm, AAS/AIAA Astrodynamics Specialist Conference, Big Sky, Montana, USA, 2023, 13-17 August
- [29] S. Alfano, A numerical implementation of spherical object collision probability, *Journal of the Astronautical Sciences* 53-1 (2005), 103-109
- [30] K. T. Alfriend, M. R. Akella, J. Frisbee and J. L. Foster, Probability of collision error analysis, *Space debris* 1-1 (1999) 21-35
- [31] K. Deb, A. Pratap, S. Agarwal and T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE transactions on evolutionary computation* 6-2 (2002) 182-197
- [32] S. Dutta, M. Budoni, G. Escribano, L. Pirovano, M. Sanjurjo-Rivo, N. Gollu, and S. Setty, A holistic approach for deciding on and designing collision avoidance maneuvers, 9th European Conference on Space Debris (ECSD 2025), Bonn, Germany, 2025, 1-4 April
- [33] H. W. Kuhn and A.W. Tucker, Nonlinear programming, MR 0047303, 2nd Berkeley Symposium on mathematical statistics and probability, Berkeley, California, USA, 1950, 31 July – 12 August
- [34] CS Group, <https://www.orekit.org/> (accessed 30.03.25)
- [35] M. S. Balch, R. Martin and S. Ferson, Satellite conjunction analysis and the false confidence theorem, *Proceedings of the royal society A: mathematical, physical and engineering sciences* 475-2227 (2019)