

National Aeronautics and Space Administration

A large, colorful galaxy with a central yellowish-white core, surrounded by swirling arms of blue, purple, and red, set against a dark space background with scattered stars.

**NASA Spacecraft Conjunction
Assessment and Collision Avoidance
Best Practices Handbook**

This document is made available through the declassification efforts
and research of John Greenewald, Jr., creator of:

The Black Vault



The Black Vault is the largest online Freedom of Information Act (FOIA)
document clearinghouse in the world. The research efforts here are
responsible for the declassification of hundreds of thousands of pages
released by the U.S. Government & Military.

Discover the Truth at: <http://www.theblackvault.com>



December 2020

NASA/SP-20205011276

The cover photo is an image from the Hubble Space Telescope showing galaxies NGC 4038 and NGC 4039, also known as the Antennae Galaxies, locked in a deadly embrace. Once normal spiral galaxies, the pair have spent the past few hundred million years sparring with one another.

Image can be found at:

https://images.nasa.gov/details-GSFC_20171208_Archive_e001327



DOCUMENT HISTORY LOG

Status (Baseline/Revision/ Canceled)	Document Version	Effective Date	Description
Baseline	V0		



THIS PAGE INTENTIONALLY LEFT BLANK



Table of Contents

Preface.....	1
1. Introduction.....	3
2. Roles and Responsibilities.....	5
3. History.....	7
3.1. USSPACECOM CA Process	8
3.1.1. SSN Tracking Functionality Description	8
3.1.2. USSPACECOM CA Screening Timeline and Process.....	10
3.1.3. Data Sharing with Other O/Os.....	10
3.1.4. SSA Sharing Agreements.....	11
3.2. NASA Partnership with USSPACECOM.....	11
4. Spacecraft and Constellation Design	13
4.1. Ascent to / Disposal from the Constellation’s Operational Orbit.....	13
4.2. General Orbit Selection: Debris Object Density	14
4.3. Vehicle and Constellation Specific Orbit Selection: Spacecraft Colocation.....	15
4.4. Launch-Related Conjunction Assessment	16
4.5. Spacecraft Trackability	17
4.6. Spacecraft Reliability	18
4.7. Development of Capabilities for Ephemeris Generation and CA Risk Assessment and Mitigation.....	18
5. Pre-Launch Preparation and Early Launch Activities.....	21
5.1. CONOPS Discussions and Arrangements with USSPACECOM Pertaining to Launch Phase.....	21
5.2. CONOPS Discussions and Arrangements with USSPACECOM and NASA Pertaining to On-Orbit Mission Phase.....	23
5.3. CONOPS Discussions and Arrangements with USSPACECOM and NASA Pertaining to Conjunction Assessment.....	24
5.4. <i>In situ</i> Launch Products and Processes.....	26



6. On-Orbit Collision Avoidance.....	27
6.1. Spacecraft Information and Orbital Data Needed for CA	27
6.2. CA Screenings.....	30
6.3. CA Risk Assessment	31
6.4. CA Risk Mitigation	34
6.5. Automated Trajectory Guidance and Maneuvering.....	37
6.6. Other Considerations.....	38
7. Contact Information	40
Appendix A. Acronyms	41
Appendix B. Glossary.....	44
Appendix C. Best Practices List	46
Appendix D. Best Practices for NASA Missions	55
Appendix E. Use of Analytic Theory Orbital Data in Conjunction Assessment..	69
Appendix F. Expected CA Event Rates	73
Appendix G. Orbital Debris Density.....	75
Appendix H. Satellite Colocation Analysis.....	79
Appendix I. Satellite Covariance Realism Assessment Procedures	84
Appendix J. CARA CA Risk Assessment Tools Validation	95
Appendix K. R-15 Message.....	96
Appendix L. Commercial Data in NASA Conjunction Assessment	99
Appendix M. Use of the Probability of Collision (Pc) as the Risk Assessment Metric for Conjunction Assessment	101
Appendix N. Pc Calculation Approaches	108
Appendix O: Collision Consequence	134
Appendix P: Event Actionability	140
Appendix Q. List of Works Cited.....	154



List of Figures

Figure 3-1 Locations and Types of SSN Sensors and Nodes	9
Figure E-1 MD Distributions for Critical ($P_c > 1E-04$) and Non-critical ($P_c < 1E-04$) Events	70
Figure E-2 eGP and Regular SGP4 Errors: LEO > 500 km (log scale)	71
Figure E-3 Expansion of Figure E-1 to Include CDFs for P_c Thresholds of 1E-03 and 1E-02	72
Figure F-1 CA Event Frequencies as a Function of Orbit Regime and Event Severity	74
Figure G-1 Debris Flux as a Function of Orbit Inclination, Orbit Altitude, and Object Size.....	76
Figure H-1 Cartesian, Inertial, Earth-centered Position Vectors.....	79
Figure H-2 Numerical vs. Analytic (Linear) α for Highly Eccentric Orbit	81
Figure H-3 Co-planar Orbits with One Circular and One Elliptical.....	81
Figure H-4 Primary Object Tool Inputs	82
Figure H-5 Tool Output for Synodic Situation	83
Figure H-6 Tool Output for Elliptical Case	83
Figure N-1 Relative Positions and Uncertainties of Two Objects	109
Figure N-2 Combining the Sizes of the Primary and Secondary Objects	110
Figure N-3 Geometry at TCA	110
Figure N-4 Representation of the Conjunction Plane	111
Figure N-5 Comparison of Monte Carlo Collision Probabilities with the 2D- P_c Method (left) and the 3D- N_c Method Introduced Here (Right), for a Large Set of CARA Conjunctions	114
Figure N-6 Cumulative 3D- N_c Risk Over Four Repeating Conjunctions.....	115
Figure N-7 Behavior of Relative Density Error by Perigee Height and Solar Activity.....	117
Figure N-8 Two-phase Application of Consider Parameters	119
Figure N-9 Profiles of 250 Conjunctions with Primary and Secondary Satellites of Non-negligible Drag	122
Figure N-10 3D- N_c Temporal Risk Plot with Monte Carlo Result Overlay (in Pink)	126



Figure N-11 Mismatch between Elongated In-track Covariances and Forced Cartesian Ellipsoidal Representation	127
Figure N-12 Comparative Results of 373 High-Pc Conjunctions	129
Figure N-13 Comparative Results of 44,000 Conjunction Events	130
Figure O-1 Range of Debris Production Possibilities	136
Figure O-2 Confidence in NanoSat Mass Ratio	138
Figure O-3 Catastrophic/Non-Catastrophic Composition as a Function of Mass Quantile	139
Figure P-1 SWTS Plot for Invariant Situation	149
Figure P-2 SWTS Plot for “Ridge” Situation	150
Figure P-3 SWTS Plot for Variation Situation	151
Figure P-4 LUPI Upper Bound for Certain EDR Values	153



List of Tables

Table C-1 Best Practices Summary	46
Table D-1 Best Practices for NASA Projects	56
Table P-1 Geopotential, Atmospheric Drag, and Solar Radiation Pressure	152
Table P-2 Ballistic Coefficient and Solar Radiation Pressure Coefficient Reasonability	152
Table P-3 LUPI Minimum and Maximum Values	153



Preface

Since Explorer 1 was launched on January 31, 1958, the United States (U.S.) has reaped the benefits of space exploration. New markets and new technologies have spurred the economy and changed lives in many ways across the national security, civil, and commercial sectors. Space technologies and space-based capabilities now provide global communications, navigation and timing, weather forecasting, and more.

Space exploration also presents challenges that impact not only the U.S. but also its allies and other partners. A significant increase in the volume and diversity of activity in space means that it is becoming increasingly congested. Emerging commercial ventures such as satellite servicing, in-space manufacturing, and tourism as well as new technologies enabling small satellites and large constellations of satellites present serious challenges for safely and responsibly using space in a stable, sustainable manner.

To meet these challenges, the U.S. seeks to improve global awareness of activity in space by publicly sharing flight safety-related information and by coordinating its own on-orbit activity in a safer, more responsible manner. It seeks to bolster stability and reduce current and future operational on-orbit risks so that space is sustained for future generations. To this end, new and better Space Situational Awareness (SSA) capabilities are needed to keep pace with the increased congestion, and the U.S. seeks to create a dynamic environment that encourages and rewards commercial providers who improve these capabilities.

The National Aeronautics and Space Administration (NASA) Spacecraft Conjunction Assessment and Collision Avoidance Best Practices Handbook:

- Reflects the goal of Space Policy Directive-3, the National Space Traffic Management Policy (SPD-3), to develop safety standards and best practices that consider “maneuverability, tracking, reliability, and disposal.”
- Reflects how NASA currently operates in space, which evolves over time. Consideration is given to important topics such as spacecraft and constellation design; spacecraft “trackability;” pre-launch preparation and early launch activities; on-orbit collision avoidance; and automated trajectory guidance and maneuvering.
- Encourages the use of commercially available SSA data and information, and for large constellations, it encourages mitigating light pollution to ground-based astronomy.
- Focuses on current NASA best practices. It will be updated as new approaches for Conjunction Assessment (CA) and on-orbit operations are developed.

A parallel document may be developed to highlight additional practices as the Department of Commerce (DOC) develops the Open Architecture Data Repository (OADR) mandated by SPD-3. Other commercial services may be offered to supplement or replace the government services currently offered.

The approaches outlined in this document are offered to spacecraft Owners/Operators (O/Os) as an example of responsible practices to consider for lowering collision risks and operating safely



in space in a stable and sustainable manner. Entities offering, or intending to offer, SSA or CA services should consider the information in this handbook from the perspective of augmenting or improving upon existing capabilities as the entire space industry benefits from advancing these capabilities. In the near term, raw observation data can be used to improve close approach predictions. Longer-term improvements might be found in enhancing notifications and data sharing, developing new models, and enabling increased automation.

NASA continuously examines and actively updates its best practices for CA as the industry undergoes rapid evolution. Large constellations of satellites, for example, comprise a new and evolving paradigm for which NASA is developing in-house expertise. NASA seeks input from the community to improve the content presented in this document. Comments or suggestions may be submitted to ca-handbook-feedback@nasa.onmicrosoft.com.

For space operations regulated by other U.S. agencies such as DOC, the Federal Aviation Administration (FAA), and the Federal Communications Commission (FCC), NASA defers to those agencies. As part of interagency consultations and to contribute to safe and sustainable space operations, NASA partners such as the FAA and FCC request NASA review of license, payload, and/or policy applications made by commercial space operators to U.S. Government (USG) regulatory agencies. In addition to the information required by those regulatory agencies, NASA has prepared examples of information for various types of missions that is valuable in expediting NASA's review. Current examples can be found at <https://www.nasa.gov/recommendations-commercial-space-operators>. Commercial space operators are welcome to contact NASA with questions about these examples.

This document was developed in close collaboration with the U.S. Space Command (USSPACECOM), one of NASA's closest interagency partners in ensuring safe operations in space, as part of the NASA/Department of Defense (DOD) Interagency Working Group for Large Constellation Conjunction Assessment Best Practices. Special thanks are due to working group members Mr. Jeff Braxton of USSPACECOM's Strategy, Plans, and Policy Directorate, and Ms. Diana McKissock and Ms. Cynthia Wilson of the U.S. Space Force's 18th Space Control Squadron (18 SPCS).



1. Introduction

This document is intended to provide NASA spacecraft Owner/Operators (O/Os) with more details concerning the implementation of requirements in NASA policies and to provide non-NASA O/Os with a reference document describing existing CA practices.

Within NASA, space flight missions are typically implemented through a program/project office structure. For example, a human exploration program might involve multiple missions (i.e., discrete launches) to a common goal. In this document, the term “mission” means the end-to-end activity that results in the creation, launch, operation, and eventual disposal of a spacecraft.

This document provides detailed information on spacecraft CA topics. Best practices for each topic are described and justified throughout this document by mission phase. A summary list of these best practices without the supporting explanatory text is provided in appendices C and D. Appendix C is a complete listing of all the best practices so that non-NASA mission managers can review them to determine which are useful for their mission. The mission manager would then need to identify the person or group responsible for performing the best practice and arrange for that effort. For NASA missions, Appendix D lists the best practices and the party responsible for performing them.

Different organizations use the term “conjunction assessment” in different ways, but NASA defines a 3-step process:

1. **Conjunction assessment** – The process of comparing trajectory data from the asset to be protected against the trajectories of the objects in the space object catalog to predict when a close approach will occur within a chosen protective volume placed about the asset.
2. **CA risk assessment** – The process of determining the likelihood of two space objects colliding and the expected consequence if they collide in terms of lost spacecraft and expected debris production.
3. **Conjunction mitigation** – An action taken to remediate conjunction risk including a propulsive maneuver, an attitude adjustment (e.g., for differential drag or to minimize frontal area), or providing ephemeris data to the secondary O/O to enable an avoidance maneuver.

This document compiles best practices based on the U.S. Government (USG) process that exists today where CA screening is performed by USSPACECOM/18 SPCS.

The risk assessment process is necessary because the orbit solutions of the catalog objects have varying accuracies depending on factors described in this document. Currently, O/Os need to perform this function for themselves (or hire a third party) because 18 SPCS is not tasked to perform this necessary risk assessment function. In the future, DOC may offer risk assessment services as it takes on a space traffic management role, and broader commercial services will be available.



Some entities use Two-Line Elements (TLEs) to perform CA. This practice is not recommended because the TLE accuracy is not sufficient to perform the necessary CA calculations. (See Appendix E for more information on TLEs.)

The terms “satellite” and “spacecraft” are interchangeable in this document. The word “object” means any discretely identifiable debris or other cataloged item in addition to satellites and spacecraft.

The term “large constellation” is defined loosely in this document. The USG Orbital Debris Mitigation Standard Practices (ODMSP) defines “large constellation” as containing 100 or more spacecraft. However, because constellations having as few as 10-20 spacecraft can experience greater conjunction risk, the O/O of any constellation of spacecraft is asked to consider the intent of the best practices in this document and implement all of them to the degree possible.



2. Roles and Responsibilities

This section provides a brief introduction to the organizations referenced in this document.

1. **The U.S. Space Command (USSPACECOM)**, one of the combatant commands within DOD, is responsible for U.S. military space operations. For the purposes of this document, it is the entity responsible for establishing SSA sharing agreements with domestic and international entities, both governmental and private. It establishes guidance and direction governing execution of the congressionally mandated SSA sharing program on behalf of the Secretary of Defense. It also oversees execution of the SSA sharing program as well as day-to-day operations of the USSPACECOM SSA and space flight safety support.
2. **The U.S. Space Force's 18th Space Control Squadron (18 SPCS)** maintains the U.S. catalog of space objects and provides 24/7 space flight safety support on behalf of USSPACECOM. 18 SPCS also conducts advanced analysis, sensor optimization, CA, human space flight support, reentry/break-up assessment, and launch analysis on behalf of USSPACECOM.
3. **The NASA Human Space Flight Operations Directorate (FOD)** of the Johnson Space Center (JSC), through the console positions Trajectory Operations Officer (TOPO) and Flight Dynamics Officer (FDO), provides conjunction risk analysis support to the space flight missions that fall under NASA human space flight. The International Space Station (ISS) and vehicles visiting the ISS receive conjunction risk assessment support from the TOPO,¹ while Artemis space flight missions receive conjunction risk assessment support from their assigned FDO. The TOPO group is the sole liaison to USSPACECOM and the U.S. Space Force for matters related to trajectory maintenance and the orbital safety of human space flight assets.
4. **The NASA Conjunction Assessment Risk Analysis (CARA) Program** located at the Goddard Space Flight Center (GSFC) provides conjunction analysis and risk assessment services for all NASA spacecraft not affiliated with human space flight. CARA is responsible for protecting the orbital environment from collision between NASA non-human space flight missions and other tracked and cataloged on-orbit objects. CARA is responsible for routinely collecting predicted orbital information from NASA spacecraft operators, passing it to NASA Orbital Safety Analysts (OSAs) for screening against the space object catalog, analyzing the screening results to determine the risk posed by predicted close approaches, and working with NASA spacecraft operators to determine an appropriate mitigation strategy for the risks posed by close approaches. CARA is the sole entity with authority to submit Orbital Data Requests (ODRs) to DOD on behalf of NASA non-human space flight entities in accordance with NASA Procedural

¹ However, for missions deployed from ISS or visiting vehicles, FOD certifies that the mission has a CA process as outlined in jettison policies but does not provide direct conjunction risk assessment support.



Requirements (NPR) 8715.6, *NASA Procedural Requirements for Limiting Orbital Debris and Evaluating the Meteoroid and Orbital Debris Environments*. CARA is the designated point of contact between NASA and USSPACECOM and the U.S. Space Force for matters related to trajectory maintenance and orbital safety of non-human space flight assets.



3. History

This section explains why the CA and related safety-of-flight services currently available exist in their present forms. It provides a brief history of the CA function, including an overview of who the actors are and why the activities are arranged as they are.

The present CA process, which predicts close approaches between space objects, began in support of the NASA Space Shuttle Program return-to-flight effort after the Challenger accident in 1986. Drawing on the DOD's ability to track space objects for missile defense and SSA, the two USG entities partnered to expand the capability to offer a methodology that would protect humans in space.

What most people consider the "space catalog" is a database of TLE sets that permits a medium-fidelity propagation and orbit prediction of all tracked objects larger than approximately 10 cm in Low Earth Orbit (LEO) using Simplified General Perturbation Theory #4 (SGP4).² The satellite catalog was originally designed not for CA but for generalized space safety, predicting the location of objects with sufficient accuracy for sensors in the network to reacquire those objects, thereby maintaining custody for purposes such as ensuring that intercontinental ballistic missile threats could be differentiated from satellites in orbit. The catalog was then, and is today, maintained using the radars, telescopes, and other sensors comprising the Space Surveillance Network (SSN). In recent years, however, USSPACECOM has been working to augment the existing network by incorporating observations from international and commercial entities.

For ISS CA to better protect humans in space, NASA collaborated with the Air Force to develop an improvement to the TLE catalog that would be accurate enough to compute a Probability of Collision (P_c) between two objects. To compute P_c , orbit determination covariance data were needed. Since the general perturbations theory used to maintain the TLE catalog (SGP4) did not produce a usable covariance, the current Special Perturbations (SP) Satellite Catalog was developed. 18 SPCS, the U.S. Space Force unit that currently maintains the catalog at Vandenberg Air Force Base (VAFB) on behalf of USSPACECOM, has greatly expanded the process of screening protected assets since that time (especially since the Iridium 33/COSMOS 2251 collision in 2009) to include all operational assets currently on orbit.

However, since the process was developed 50 years ago and was grown piecemeal to meet existing needs, it contains certain oddities of evolution that may not be anticipated by operators who are accustomed to using current technology and methods. Understanding the way the 18 SPCS process works is key to using it properly to protect on-orbit assets and the space environment.

² Hoots and Roehrich 1980.



3.1. USSPACECOM CA Process

The full CA screening process used by USSPACECOM with all accompanying details is documented in the *Spaceflight Safety Handbook for Satellite Operators* (the Spaceflight Safety Handbook), which can be obtained along with other helpful information from the USSPACECOM website Space-Track.org. This website is the principal way USSPACECOM communicates and exchanges CA-related information with O/Os.

- All non-NASA spacecraft operators should familiarize themselves with the Spaceflight Safety Handbook as well as other website contents and offerings.
- NASA spacecraft operators do not need to be familiar with the Spaceflight Safety Handbook since CARA or FOD will perform all actions on their behalf.

(See Section 7 in this document for USSPACECOM contact information including Uniform Resource Locators (URLs) to obtain products. See Sections 5.1 through 5.3 for details of some topics from the Spaceflight Safety Handbook to provide end-to-end understanding of the CA process.)

Commercial CA providers, often using space catalogs assembled from private sector space tracking assets, are now an available resource for O/Os who wish to obtain CA services. NASA encourages the use of validated commercial CA services when they can augment the existing USSPACECOM capabilities. NASA recommends that all O/Os avail themselves of the USSPACECOM CA service at a minimum as an adjunct to a commercially procured service. (See Section 6.1 for additional information about orbital data for CA, and see Appendix I for a discussion of data validation.)

To address this topic, the following specific practices are recommended:

- A. Obtain a Space-Track.org account.
- B. Become familiar with the Spaceflight Safety Handbook contents.
- C. Use USSPACECOM CA service as a minimum for screenings, even if additional commercial data or services are used.

3.1.1. SSN Tracking Functionality Description

The Space Surveillance Network (SSN), a global sensor network made up of optical telescopes and radars, is used to collect the tracking observations that are processed to maintain the High Accuracy Catalog (HAC) of on-orbit space objects. Figure 3-1 shows the current locations of the member sensors.

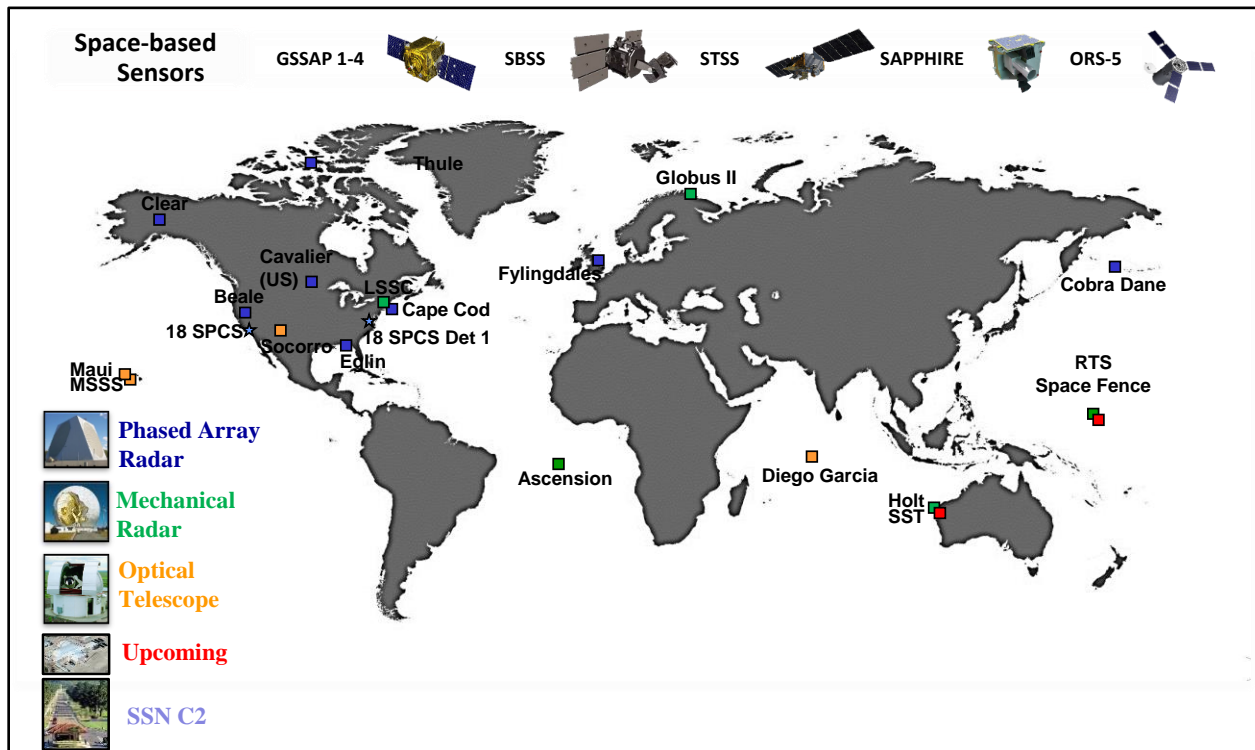


Figure 3-1 Locations and Types of SSN Sensors and Nodes³

Because most sensors have other purposes in addition to collecting data for SSA, tracking data cannot be simply requested and obtained on demand. Objects are placed in tasking categories, and a list of objects to attempt to track by category is sent to each sensor, which then sorts the list to determine which requests are feasible given the sensor’s calculation of the probability of detection of the requested satellite and competing activities.

Priority is given to military operational needs and special needs within CA such as human space flight. If a USSPACECOM CA operator determines that there is insufficient tracking data available to create a good orbit determination solution for an object, the operator may attempt to obtain additional data by increasing the priority of the object or by tasking additional sensors to track it. Sensors may not be able to track an object due to geometric and power constraints, weather, equipment outages, and other exigencies.

The USSPACECOM process is different from a commercial process in which a vendor is compensated directly for the tracking of an object with the flexibility of being able to obtain additional data if more funds are available to support the request. Because of the time required to collect additional data on an object of concern using the USSPACECOM process, it is recommended that O/Os provide a significant length of ephemeris prediction (including planned maneuvers) to USSPACECOM for screening. This screening duration allows time to request and

³ Figure courtesy of USSPACECOM.



obtain increased tasking for identifying secondary objects with insufficient orbit determination solutions, and thus to enable the best Pc to be computed for use in CA risk assessment decisions. (See Section 6.1 for a more detailed discussion of these issues.)

To address this topic, the following specific practice is recommended:

- A. Provide seven (7) days of predicted ephemeris (including maneuvers) to USSPACECOM for screening for LEO spacecraft, and provide 14 days for other orbits; e.g., High Earth Orbit (HEO)/Geosynchronous Equatorial Orbit (GEO).

3.1.2. USSPACECOM CA Screening Timeline and Process

Personnel at the USSPACECOM facility at VAFB currently perform CA screenings of protected assets against the entire HAC at minimum once per day and up to three times per day every day. Screenings are performed using both the ephemeris data that O/Os share (which include planned maneuvers for the asset) as well as the SSN-derived orbit determination solution for the asset (which does not consider predicted maneuvers).

Results from the screenings, in the form of Conjunction Data Messages (CDMs), are posted to Space-Track.org for customers with accounts. Entities without accounts who have provided an email address to USSPACECOM may receive basic prediction information by email if an emergency is detected.

Special/off-cycle screenings can be requested per the guidance in the Spaceflight Safety Handbook. It is important for safety of flight for all spacecraft operators to ensure that there are no large latencies or gaps in the process as uncertainties in the propagated orbit continue to grow over time. For instance, submitting an ephemeris for screening today at noon that has an epoch of today at 0 hours means that the data is already 12 hours old. If it then takes 8 hours to screen that data and another 2 hours to analyze the results, they will be very old indeed. Ensure that the data screened is for a future time and is analyzed and acted upon before it becomes stale. (See Section 6.2 for more information.)

To address this topic, the following specific practices are recommended:

- A. Provide at least one ephemeris per day to be screened and three ephemerides per day in the lower-drag LEO regime (perigee height less than 500 km).
- B. Determine whether the O/O's process for obtaining screening results and performing CA risk assessment aligns with the timeline of the USSPACECOM process. If the timelines do not align in such a way as to enable timely and efficient screening support, pursue a rearrangement of the O/O's process to minimize data latency and optimize screening efficiency.

3.1.3. Data Sharing with Other O/Os

To facilitate communication between operators to mitigate a close approach, USSPACECOM maintains a list of contact for each asset. Maintaining this list of contacts is critical so that someone can always be reached. Data sharing is further facilitated through the use of standard formats and coordinate frames. (See Section 6.4 for a more thorough treatment.)



To address this topic, the following specific practices are recommended:

- A. Populate and maintain the point of contact section on Space-Track.org with your operations contact data. Be sure that the operations contact can be reached 24/7 due to time zone differences between operators and the immediate nature of certain CA emergencies.
- B. Use standard ephemeris, CDM, and maneuver notification formats defined by the Consultative Committee for Space Data Systems (CCSDS).

3.1.4. SSA Sharing Agreements

USSPACECOM/J535, SSA Data Sharing Branch, negotiates SSA sharing agreements, which establish the parameters within which data will be exchanged by the signing parties to facilitate ongoing cooperation and advance space flight safety. These agreements are useful because they substantially expand the types of data products that can be received over the default products provided to entities without agreements.

Any member of the space community, including satellite operators, launching agencies, commercial service providers, and research/academic institutions, that wishes more and more frequent SSA products than are provided through a generic Space-Track.org account should contact USSPACECOM. (See Section 7 of this document for contact information.)

USG organizations and their contractors have implied agreements and do not need to pursue a separate formal arrangement with USSPACECOM.

3.2. NASA Partnership with USSPACECOM

Since the development of the HAC and its use in routine CA screenings, NASA has partnered with USSPACECOM and its predecessor organizations. That partnership currently includes special, dedicated CA screening support. The human space flight program uses Air Force civilians for CA screening support, while NASA's non-human space flight missions are supported by NASA contractors called Orbital Safety Analysts (OSAs) who work on the operations floor in the VAFB operations center. While both groups do essentially the same work with the same input data, NASA OSAs focus specifically on NASA needs and can write scripts that are used to tailor the output data for use by CARA.

CA is a 3-step process:

1. The first step is CA "screening," which involves computing the predicted close approaches between the protected asset and the catalog of space objects. This is the step performed at the USSPACECOM facility for mission customers, and screening results are provided that describe predicted close approaches. Both the USSPACECOM CA cell and NASA OSAs perform this function.
2. The second step is CA "risk assessment," in which the screening results are analyzed to determine the level of risk posed by each predicted close approach and to determine whether the predicted close approaches warrant additional investigation and, ultimately, mitigation. This step is critical because not all predicted close approaches require



mitigation, and often close approaches require analysis to determine what action is warranted. This step is not performed by USSPACECOM as it is a responsibility allocated to O/Os or organizations that perform this activity on the O/O's behalf. NASA JSC FOD performs this function for human space flight program assets, and NASA CARA performs this function for all other NASA spacecraft.

3. The third step, if required, is CA “mitigation,” in which the O/O plans, and perhaps executes, a collision avoidance maneuver or other mitigation solution to reduce the collision risk to an acceptable level. This activity is allocated to the individual O/Os although NASA CARA does provide some basic support tools to aid in the initial choice of mitigation actions for its mission customers.

Because risk assessment requires different products and yields different analyses than screening, it is important to consider the tools and processes required to assess risk properly and adequately. For large constellations, the cadence of conjunctions and the consequence of a collision for the proper operation of the constellation are both heightened. (See Section 6 for more detail on risk assessment processes.)

To address this topic, the following specific practices are recommended:

- A. Develop a robust safety-of-flight process that includes both CA screening and risk assessment to inform close approach mitigation decisions.
- B. Large constellation operators should work with USSPACECOM pre-launch to determine if variations from the standard approach are necessary and, if so, to define a tailored screening process.
- C. Large constellation operators should consider working with NASA to define a risk assessment process. NASA is an experienced representative operator. Including NASA in discussions regarding establishing a CA process will ensure that the process will work with most operators for risk assessment purposes.



4. Spacecraft and Constellation Design

Safety of flight is an integral aspect of satellite space operations and should be considered as part of design decisions to enable cost- and mission-effective solutions so that potential impacts to other operators in the space environment can be avoided. Specific design areas for which the consideration of on-orbit safety issues is appropriate include:

- orbit selection,
- spacecraft ascent and disposal activities,
- sensor trackability of the spacecraft,
- spacecraft reliability,
- capabilities for ephemeris generation, and
- capabilities for risk assessment and mitigation.

Each of these items is treated in more detail in this section.

4.1. Ascent to / Disposal from the Constellation's Operational Orbit

Safety-of-flight issues may be considerable when related to the ascent from the launch injection point to the satellite's on-station position or to the descent to an orbit from which disposal/re-entry can be accomplished directly. Typical ascent and descent trajectory design using now common electric propulsion or low-thrust chemical propulsion can take months to accomplish, potentially passing through highly populated regions of space along the way. For large constellations, the amount of nearly continuous transiting satellite traffic could be especially large as old satellites pass out of service and replacement satellites are added.

All missions should plan for orbit disposal to reduce long-term orbital debris. The USG has established Orbital Debris Mitigation Standard Practices (ODMSP) including for post-mission disposal.⁴ NASA missions are required by NPR 8715.6 to comply with NASA Standard 8719-14B, *Process for Limiting Orbital Debris*, which defines how NASA implements the ODMSP. Reducing the number of inactive space objects minimizes the probability of generating debris through orbital collisions, and the fastest disposal option should be pursued; for example, existing large constellations are promptly deorbiting end-of life spacecraft.

During ascent and descent, a spacecraft will pass by other assets that are already on orbit. Some but not all of those may be maneuverable. The ascending/descending spacecraft that is equipped to maneuver needs to yield the right-of-way to existing on-orbit assets by performing risk mitigation maneuvers or ascent/descent trajectory alterations.

Special caution is needed to protect humans on orbit. If the ascent or descent trajectory will pass through the ISS altitude, operators should coordinate with the NASA JSC FOD to avoid perigee-

⁴ U.S. Government Orbital Debris Mitigation Standard Practices. 2019 version
https://orbitaldebris.jsc.nasa.gov/library/usg_orbital_debris_mitigation_standard_practices_november_2019.pdf.



lowering approaches that pose persistent and problematic orbital crossings with ISS and other HSF assets.

To address this topic, the following specific practices are recommended:

- A. Perform a study to compute the number of expected close approaches anticipated during ascent and descent as well as the imputed additional satellite reliability that will be required to meet satellite disposal requirements at the chosen operational orbit.
- B. If the results of the study show a large burden, consider choosing a different mission orbit with a lower burden.
- C. All missions should review and follow the ODMSP guidance standards.
- D. When practicable, pursue active disposal using the fastest disposal option available.
- E. Recognize that transiting spacecraft should yield way to on-station spacecraft and thus take responsibility for any CA risk mitigation maneuvers or transient trajectory alterations that may be required.
- F. When planning descent using thrusters, if not planning an approach of circular-orbit altitude reduction, coordinate with the NASA JSC TOPO to ensure that perigee-lowering approaches do not present persistent and problematic orbital crossings with the ISS and other human space flight assets.

4.2. General Orbit Selection: Debris Object Density

Space debris is not uniformly distributed about the Earth. Because most of the present space debris is generated from a relatively small number of satellite collisions or explosions, the orbital parameters of these collided or exploded satellites determine where the large debris fields reside.

For example, the debris object density is much greater in the 750-900 km altitude band than in other parts of LEO. If possible, this region is one to avoid as a destination orbit since a larger number of serious satellite conjunctions with debris objects can be expected. An orbit that results in a high number of close approaches can lead to the need to perform many maneuvers that may not fully mitigate the risk of close approach. Spacecraft that cannot maneuver will operate at higher risk of debris creation, effectively presenting additional risk to other operators.

Orbit selection should be informed by a study to determine expected conjunction rates; that is, the expected number of lifetime conjunction risk mitigation maneuvers and the amount of satellite fuel needed to permit the number of desired years on orbit. Published results by NASA CARA are available and can be consulted to obtain first-order high-interest event density information. (See Appendix F for more information and examples.) Use this information to assess the mission impact for CA risk mitigation maneuvers, given the expected lifetime of each satellite (e.g., additional propellant required, impacts on data collection).



To address this topic, the following specific practices are recommended:

- A. Include considerations of orbital debris density when selecting candidate mission orbits. Debris density diagrams, indexed by object size, are published by the NASA Orbital Debris Program Office (ODPO). (See Appendix G for more information and examples.)
- B. When choosing among the final set of candidate orbits, perform a conjunction frequency analysis to determine the number of CA high-interest events per spacecraft per year that can be expected in each of the proposed orbital areas.

4.3. Vehicle and Constellation Specific Orbit Selection: Spacecraft Colocation

While there is a process used by the International Telecommunication Union (ITU) to ensure that spacecraft locations are deconflicted from a radio frequency perspective, there is no similar centralized process for ensuring that spacecraft locations are deconflicted from a location perspective. This gap has resulted in spacecraft injecting into their planned mission orbit only to find another spacecraft already in or very near that location (colocated), which causes close approaches that then have to be mitigated, potentially very frequently, over the whole mission lifetime.

Another similar problem is that of systematic or repeating conjunctions in which a spacecraft has multiple repeated close approaches on successive orbits with another object over a long period of time; for example, when using electric propulsion to move slowly through a congested area. Two actively maintained spacecraft with very similar orbits can present a permanent recurring collision hazard because both spacecraft can be planning future maneuvers that result in a collision if not coordinated. Therefore, active coordination for every maneuver is required to prevent the execution of simultaneous maneuvers that could create a collision.

Due to the lack of centralized coordination, it is incumbent upon individual operators to protect themselves and the space environment by performing a colocation analysis on their own. The goal would be to prevent choosing a location already occupied by another spacecraft, or, if that is not possible, to work with the other operator before launch to devise a strategy to share the space. Colocation situations can often be ameliorated through relatively small changes to the planned orbits during the design phase. For NASA missions, CARA can perform this analysis using the current USSPACECOM catalog data. (See Appendix H for more information and examples.) For non-NASA missions, CARA will make their software tool available via public release.

In performing the analysis, it is not necessary to consider proximity to objects that are not maneuverable (debris, rocket body, dead spacecraft) since the non-maneuverable object will pass through the area under the control of non-conservative forces such as atmospheric drag and not present a future issue for the (largely) static orbital parameters of the asset that is maintaining its orbit using maneuvers.



To address this topic, the following specific practices are recommended:

- A. During orbit selection, perform an orbit colocation analysis to determine whether any of the proposed orbits is likely to create systematic conjunctions with existing actively maintained satellites.
- B. If the orbit colocation analysis identifies systematic conjunctions, consider modifying the proposed orbit(s) slightly to eliminate this possibility. Optimizing orbit placement may be appropriate.
- C. If the selected orbit is likely to present systematic conjunctions with a pre-existing spacecraft, then coordinate with the other operator(s) to arrange a process to coordinate maneuver plans routinely during the life of the mission.

4.4. Launch-Related Conjunction Assessment

Launch Collision Avoidance (Launch COLA or LCOLA) is the evaluation of a particular launch trajectory to determine whether the launch assembly, following this trajectory, will produce any close approaches with other space objects. LCOLA is sometimes expanded to include more indirect analyses to ensure that launch offscourings (rocket bodies, fairings, etc.) do not produce close approaches with other objects during the first three days after launch, after which they can be presumed to be cataloged and thus handled by the usual on-orbit CA methods.

LCOLA requirements are governed by the individual launch range or launch licensing agency. Requirements are in the form of stand-off distances P_c depending on whether the on-orbit conjuncting asset is a piece of debris, operational payload, or human space flight mission. Examples of the current requirements are 200 km stand-off distance or $1E-6 P_c$ vs human space flight assets and 25 km or $1E-5 P_c$ for operational payloads. In practice, LCOLA covers launched objects from the time they reach 150 km through three hours. Launch times found in violation of these requirements are enforced as launch hold periods by the respective wing commanders.

A large study conducted by NASA (Hejduk et al. 2014) questioned the added value of LCOLA, given the uncertainties typical of most predicted launch trajectories. However, recent improvements in the fidelity of these predicted trajectories has made the LCOLA enterprise potentially more meaningful, suggesting that the situation may need to be reevaluated.

The period between when traditional LCOLA ends and an on-orbit asset can mitigate a risk using standard on-orbit CA methods is called the COLA gap. The O/O should ensure that newly launched objects and their detritus will not come into conjunction with protected assets during the COLA gap (i.e., until the 18 SPCS can catalog the new object(s) and is available for CA screening). Such analyses should consider not just the predicted nominal launch injection but the expected (up to three-sigma) launch dispersions.

Specific criteria defining COLA gap launch cutouts do not exist since they are dependent on both what is being launched and the capabilities of the protected on-orbit asset. For example, ISS needs 36 hours for the COLA gap, which allows 24 hours for 18 SPCS to track the launch objects and then 12 hours for ISS to assess the risk and mitigate if required. Other assets might



require more time to close this risk gap. The methodology for this risk assessment can be a stand-off distance phasing analysis, nodal separation and in-track screening, or a probability density model. Note that this analysis is typically performed by Aerospace⁵ for DOD missions and the NASA Launch Services Program for NASA missions. These organizations can be consulted for details on the analysis techniques and criteria, which are also listed in Appendix Q (Hametz and Beaver 2013; Jenkin et al. 2020). NASA JSC FOD can also be consulted for analysis methods and to get details on the assets that could be at risk for each specific launch.

To address this topic, the following specific practices are recommended:

- A. Protect human space flight assets from close approaches during the COLA gap using stand-off distance or statistical measures.
- B. Conform to additional LCOLA requirements that the launch range may impose.

4.5. Spacecraft Trackability

Effective CA requires an accurate orbital state. Accurate orbital state data requires non-cooperative tracking data to maintain a comprehensive space catalog available to all space operators. The USSPACECOM catalog serves this purpose and uses the data collected by the SSN. Because even active spacecraft become debris after their end-of-life and therefore need non-cooperative tracking to obtain an orbit solution, all launched satellites need to be acquirable and trackable by the SSN so that they can be cataloged and maintained using SSN capabilities alone. Launching an untrackable spacecraft increases risk to all operators.

Objects are trackable if they have a large enough radar or optical cross section to be tracked by at least two SSN sensor assets. Analytical evaluations of trackability consider both object size and material properties, so there is no absolute size threshold that is determinative. But as a rule of thumb, satellites need to have characteristic dimensions of 10 cm in each major dimension for spacecraft with perigee less than 2000 km and greater than 50 cm in each major dimension for spacecraft with perigee greater than 2000 km.

To address this topic, the following specific practices are recommended:

- A. Through selection of physical design and materials, ensure that the satellite is trackable by SSN.
 - a. For spacecraft with perigee heights less than 2000 km, the spacecraft should have characteristic dimensions of at least 10 cm in each major dimension.
 - b. For spacecraft with perigee heights greater than 2000 km, the spacecraft should have characteristic dimensions of at least 50 cm in each major dimension. (*continued*)

⁵ Aerospace is an independent, nonprofit corporation operating the only Federally Funded Research and Development Center (FFRDC) for the space enterprise.



- c. If the spacecraft cannot meet these dimension constraints, use a proven detectability enhancement.

4.6. Spacecraft Reliability

Studies related to active debris removal strategies have shown that the greatest contributor to long-term space debris growth is derelict spacecraft left on orbit. They produce opportunities for collision with smaller debris objects that cannot be mitigated, which will then generally produce large amounts of debris. So it is important to ensure that spacecraft survive until they can be disposed of by using appropriate disposal as described in the OMDSP and by ensuring that the probability of successful post-mission disposal meets or exceeds 99%.

To address this topic, the following specific practices are recommended:

- A. Ensure that spacecraft reliability is high enough that the likelihood of each spacecraft remaining fully functional until it can be disposed of meets or exceeds 99%.
- B. Reassess the 99% analysis whenever the underlying assumptions change; for example, extending operations beyond design life or failures in key systems.

4.7. Development of Capabilities for Ephemeris Generation and CA Risk Assessment and Mitigation

CA capabilities are needed immediately after launch. The software and process components of the CA risk assessment and mitigation capabilities should be fully developed and tested before launch including elements of both the flight hardware and the ground system.

Three main capabilities should be developed:

1. **Predicted ephemeris capability.** Predicted ephemeris data that can be shared with other O/Os is needed to assess the potential for collisions between space objects. The current USG CA process uses two predicted ephemeris solutions for the asset: one generated using non-cooperative SSN tracking data and one shared with 18 SPCS by the O/O. For active, maneuverable spacecraft, the O/O-generated ephemeris is often the best representation of the future position because the O/O usually has access to more cooperative tracking data to feed the orbit-determination process (for example, from on-board Global Positioning System (GPS) receivers); has an accurate model of the spacecraft's drag coefficient and frontal area; and, most importantly, can model planned future maneuvers in the predicted trajectory. To enable the CA enterprise, predicted ephemerides must be furnished frequently, span an appropriate period of predictive time, employ point spacing close enough to enable interpolation, provide a full state (position and velocity) for each ephemeris point, and provide a realistic 6 x 6 covariance matrix (with both variance and covariance terms) for each ephemeris point. (See Appendix I for more details about covariance realism characterization and evaluation.)



2. **Risk assessment capability.** The results of satellite conjunction screening analyses are sent from USSPACECOM to O/Os in the form of CDMs that contain the states, covariances, and amplifying orbit determination information for the primary and secondary object at the Time of Closest Approach (TCA). These messages, however, are only a notification of a predicted potential close approach. For a robust safety-of-flight process, additional risk assessment analysis is needed to determine whether the close approach warrants mitigation. CA risk assessment tools are needed to perform this risk assessment, which includes calculation of P_c and other relevant information such as expected collision consequence (addressed in Appendix O) and whether the state and covariance information is sufficiently accurate to subtend the risk assessment process (addressed in Appendix P). For conjunction events in which the P_c exceeds a severity threshold, mitigation action planning is necessary to plan a mitigation option that will lower the P_c to below an acceptable threshold, usually conducted with the aid of trade-space plots that give conjunction P_c reduction as a function of spacecraft trajectory modification size and execution time. Validating these tools is important well before flight. (See Appendix J.) Some commercial vendors offer risk assessment services. Spacecraft operators who choose to purchase risk assessment services should make these arrangements well before launch to ensure adequate time for testing and development of a corresponding process for choosing a mitigation option given the risk assessment results received from the service.
3. **Mitigation capability.** A spacecraft can mitigate a close approach in several ways:
 - Propulsive maneuvers (e.g., performed using thrusters).
 - Attitude changes to take advantage of altered drag coefficients from altered frontal areas to effect a relative velocity change between the two objects.
 - Changing the spacecraft attitude to present a minimal frontal area to the relative velocity vector to minimize the likelihood of collision.
 - Sharing ephemeris data with the operator of the other spacecraft involved in the close approach so that they can perform an avoidance maneuver.

Because the design of the spacecraft and/or ground system needs to accommodate the ability to perform these actions, spacecraft operators should choose a mitigation option when they select their mission orbit and ensure that the capability is built into the system appropriately in time for testing and use.

To address this topic, the following specific practices are recommended:

- A. Develop and implement a capability to generate and share accurate predicted spacecraft ephemerides including any planned maneuvers. (See Section 3.1 for a discussion of ephemeris sharing for screening. See Section 6.2 for specific guidance on ephemeris generation frequency, length, point spacing, formatting, and contents.) (*continued*)



- B. Determine during spacecraft design what risk mitigation approaches are possible for the spacecraft. While trajectory modification via thrusting maneuver is the most common approach, other approaches such as differential drag orbit modification are possible.
- C. Develop and implement or arrange to acquire a capability to process CDMs and to compute conjunction risk assessment parameters such as the P_c .
- D. Develop and implement or arrange to acquire a risk analysis capability to select mitigation actions that will lower the P_c for dangerous conjunctions to a user-selected value.
- E. Validate CA tools well before flight, typically 6-12 months prior.



5. Pre-Launch Preparation and Early Launch Activities

The pre-launch and launch/early orbit phase of a satellite's lifetime is the period during which safety-related policy decisions are rendered and associated data interfaces established. Safety-related policy decisions carry forward as precedent and actual operational arrangement to form the O/O's safety posture over the satellite's orbital lifetime.

Many pre-launch activities such as establishing interfaces, data sharing agreements, and a Concept of Operations (CONOPS) require coordination. Other activities pertaining to the launch and early orbit period such as outlining particular practices and reporting are needed to conduct CA activities during this phase of orbital life and to facilitate a smooth transition into on-orbit CA. Each activity is treated individually in this section.

5.1. CONOPS Discussions and Arrangements with USSPACECOM Pertaining to Launch Phase

A spacecraft's journey from the launch facility to its on-station positioning in its service orbit is complicated and open to mishaps that can damage both the launching spacecraft and other spacecraft in proximity. The USG performs required activities in support of the launch enterprise including cataloging the spacecraft and any other space objects, debris and otherwise, related to the launch. To meet these responsibilities and provide a safe exercise of the launch process, the spacecraft operator needs to execute several pre-launch activities and points of coordination. This is best done by coordinating with USSPACECOM to discuss all of the related issues, exchange information, and generate a comprehensive launch and deployment plan so that all affected agencies can exercise their proper roles.

First, early cataloging is extremely important to enable the CA process to protect all other on-orbit assets. To assist the cataloging function, all aspects of the launch delivery and deployment methodology need to be fully documented and communicated to USSPACECOM so that those performing the launch cataloging process know what to expect at each point and can respond appropriately.

Aspects of the launch delivery that may affect cataloging and should be documented would include the launch vehicle configuration (rocket versus air launch) and the launch phases (parking orbits, transfer orbits, injection, and post-injection transiting to on-station locations).

The deployment methodology includes timing post-separation and number of objects deployed. Examples of unusual deployment features that might create cataloging issues include high-velocity deployments, non-immediate deployments (i.e., subsequent deployment of a satellite from one of the deployed satellites), and tethered satellites. Multiple nearly simultaneous deployments are a challenge for tracking and will delay cataloging which in turn poses a safety-of-flight risk. Deployments should be spaced to allow acquisition of each object. For multiple nearly simultaneous deployments, the expected range of deployment speeds and directions should be provided.



Input for the R-15 launch form (launch plan and orbital parameters) should also be discussed with USSPACECOM so that they may render any necessary clarification or aid. The launch provider is responsible for completing and submitting the R-15 form; however, the O/O should discuss the content with USSPACECOM well in advance of the launch. (See Appendix K for more information.)

Second, detailed launch trajectory information for all launch objects needs to be submitted to USSPACECOM well before the launch date so that a colocation and object avoidance analysis for the launch sequence can be conducted and, if necessary, adjustments made to ensure safety. To avoid close approaches with the ISS and other active spacecraft, it may be necessary to modify the launch sequence slightly and eliminate certain launch times from the launch window. Launch providers frequently take responsibility for these activities, but it is important for O/Os to be aware of the availability and benefits of this service.

Third, to avoid any data transfer failures while the launch is in progress, the required exchange of data between O/O and USSPACECOM during the launch process and the associated formats and timelines need to be established, understood, and exercised prior to launch.

To address this topic, the following specific practices are recommended:

- A. Establish contact with USSPACECOM to describe and discuss all aspects of the launch including delivery and deployment methodologies.
- B. Space the deployment of multiple spacecraft in a way that enhances the ability of USSPACECOM to quickly identify and differentiate the spacecraft using the SSN.
- C. Discuss the preparation of the R-15 launch form with USSPACECOM so that they understand its contents and any implications for safety of flight.
- D. Ensure that the launch provider submits launch trajectory information to USSPACECOM, including ephemerides for the powered flight portions and orbital elements for the parking orbits, so that potential colocation and deconfliction potentialities can be discovered.
- E. Provide to USSPACECOM as soon as possible launch-related information (e.g., injection vectors and initial ephemerides for deployed spacecraft) that can be used to assist with the cataloging process, especially to confirm the identity of launch-related objects. Coordinate a satellite numbering scheme (potentially including temporary satellite numbers) appropriate to the launch type and expected degree of cataloging difficulty.
- F. Coordinate with USSPACECOM any potential launch anomaly diagnostic products that can be provided if issues arise during the launch and early orbit sequence.



5.2. CONOPS Discussions and Arrangements with USSPACECOM and NASA Pertaining to On-Orbit Mission Phase

USSPACECOM is required to maintain orbit determination solutions on all objects arising from the launch as part of its military space object custody responsibilities. USSPACECOM also provides CA information to O/Os for safety of flight. In order for USSPACECOM to maintain the catalog most easily for safety of flight for all, the O/O should provide relevant details about its assets. The pre-launch period is the appropriate time to conduct exchanges of information that facilitate these activities.

1. **Provide operational contact information.** It is important for the satellite operator to provide operational point-of-contact information so that USSPACECOM can coordinate spacecraft tracking, cataloging, identification, and provision of space flight safety data for the satellite. Submitting this information allows expeditious contact if issues arise during the launch and on-orbit periods of the satellite's lifespan and provides the information necessary for USSPACECOM to arrange for the delivery of basic space flight safety services.
2. **Provide information about satellite construction and operation.** Basic information about the satellite to be deployed, such as satellite dimensions, presence of deployable structures (e.g., solar panels, antennae), satellite material properties, and expected satellite regular attitude, needs to be communicated to USSPACECOM. These aspects of satellite construction and operation affect trackability, and an understanding of these construction features allows USSPACECOM to assemble a tracking approach, assign appropriate sensors, and predict the regularity of attempted tracking success.
3. **Provide satellite orbit maintenance strategy.** The orbit maintenance strategy for the satellite should be communicated, at least at a high level. Information such as the logic for determining orbit maintenance maneuvers, the maneuver thruster technology, and the frequency/size/duration of orbit maintenance burns will help USSPACECOM to set up the orbit maintenance parameters properly for each satellite in order to perform the most reliable orbit maintenance possible.
4. **Provide the satellite flight control and navigation paradigm.** Understanding the flight control and navigation paradigm of the satellite is needed, especially whether a traditional ground-control paradigm has been followed or whether autonomous navigation is used. If the latter, the amount of ground-knowledge of satellite activities, such as whether there is foreknowledge of maneuvers and an opportunity to override these maneuvers from the ground before they are executed, is important in order to understand the degree to which the satellite's navigation can be manually influenced if necessary. In addition to setting general expectations for the satellite's flight dynamics, this information is helpful in determining the degree of expected fidelity for regularly generated satellite O/O ephemerides.



To address this topic, the following specific practices are recommended:

- A. Register the spacecraft with USSPACECOM using the Satellite Registration form on Space-Track.org.
- B. Provide USSPACECOM with basic construction and mission information about the satellite such as stowed dimensions; deployable structures such as solar panels, antennae, booms, including all of their (rough) dimensions; satellite material properties and colors; regular satellite attitude; and registered operational radio frequencies.
- C. Provide USSPACECOM with a basic overview of the satellite's orbit maintenance strategy including the paradigm for determining when orbit maintenance maneuvers are required; the maneuver technology used (for this relates to burn duration and expected accuracy); and the frequency, duration, and magnitude of typical burns.
- D. Provide USSPACECOM with an understanding of the flight control and navigation paradigm, principally whether a ground-based control approach is followed or some degree of (or full) autonomous control is used. If the satellite control does include some autonomous flight dynamics or control features, indicate how much (if any) foreknowledge ground controllers have of autonomous maneuver actions, the amount of information that is communicated to the ground both before and after the maneuver (e.g., maneuver time, delta-V, direction), and whether ground-based overrides are possible.

5.3. CONOPS Discussions and Arrangements with USSPACECOM and NASA Pertaining to Conjunction Assessment

On-orbit CA is needed to protect the satellite and keep valuable orbital corridors free of debris pollution and sustainable for the indefinite future. The screening process is the first step in completing a risk analysis for a potential conjunction.

CA screenings typically use O/O ephemerides as a statement of the primary object's position. Ephemeris formats, delivery mechanisms, and screening timetables need to be coordinated with the screening provider. The screening provider should use this information to monitor the satellite's projected position over time and identify potential conjunctions. The O/O needs to choose which screening service provider will be used (i.e., the USSPACECOM free service, a validated commercial CA service, or both) and establish a robust interface prior to launch.

USSPACECOM currently provides a free CA service that performs CA screenings on behalf of an O/O and sends proximity warnings (CDMs) for each situation in which the distance between the O/O's satellite and another cataloged object is smaller than a set threshold. These messages give O/Os the data they need to assess the collision likelihood of a particular conjunction and, if necessary, plan and take mitigative action. As noted in Section 3.1, NASA recommends at a minimum using the USSPACECOM CA service, which can be augmented with validated



commercial CA services when desired. (See the USSPACECOM Spaceflight Safety Handbook for additional information regarding the USSPACECOM CA service.)

Commercial providers should provide equivalent support and data formats. Many commercial CA services work from the USSPACECOM CA data and not from separate or unique commercial SSA data.

Based on satellite orbit and mission characteristics, a screening volume size will need to be assigned. (See the USSPACECOM Spaceflight Safety Handbook for more details.) The screening volume is the physical volume (generally an ellipsoid) that is “flown” along the primary’s orbit during the screening process, with any objects found within this volume considered to be conjunctions and associated CDMs generated. If a commercial CA provider is selected, the O/O will need to work directly with the provider to discuss the type, timing, and format of the information needed. (See Appendix L for Information about how NASA validates and uses commercial data.)

USSPACECOM distributes CDMs through Space-Track.org. An O/O needs to register on Space-Track.org and provide contact information to USSPACECOM to receive the basic level of CDMs. If desired, a test instantiation of Space-Track.org is available to allow O/Os to practice generating, receiving, and processing the CA data products.

Finally, a formal Orbital Data Request (ODR) for any of the desired USSPACECOM-generated CA information beyond the most basic products will need to be submitted to USSPACECOM and adjudicated before information exchange can begin. An ODR form is a method to request data and services beyond the basic ones. The form and instructions for filling it out are posted on Space-Track.org. Non-USG entities are highly encouraged to sign SSA sharing agreements with USSPACECOM, which expedites the ODR process. For NASA missions, ODRs are sent to CARA and JSC FOD, who submit them on behalf of the mission. ODRs are used to request permission to redistribute USSPACECOM data to other entities, obtain specialized one-time analyses, and make other specific requests.

Large constellations may require special considerations such as quantity or timeframe of data used for screening and risk analysis. NASA has considerable operational experience with CA risk assessment and has assisted previous large constellation operators in designing a CA risk assessment process that scales appropriately.

To address this topic, the following specific practices are recommended:

- A. Decide whether the USSPACECOM free service, a validated commercial CA service, or both will be used by the mission.
- B. Establish a service with the selected service provider.
- C. Implement a SSA sharing agreement with USSPACECOM to receive advanced data support and services. (*continued*)



- D. Through the registration of the satellite with USSPACECOM, begin the process of arranging for CA data exchange including O/O ephemerides, maneuver notification reports, and CDMs. USSPACECOM uses the Space-Track.org account as the mechanism for product exchange.
- E. If needed, complete an Orbital Data Request (ODR) form to arrange for delivery of USSPACECOM advanced CA products.
- F. For large constellations, coordinate with NASA and the screening provider to identify and address any special considerations.

5.4. *In situ* Launch Products and Processes

Once liftoff is achieved, the foci become those of a safe journey to the final on-station destinations of the spacecraft and the efficient performance of the launch cataloging process. To determine whether the launch is nominal and has deposited its spacecraft as expected, the initial injection vector should be provided to USSPACECOM as soon as it is available. Additionally, once initial contact has been made with each spacecraft and initial position information has been downlinked, the generation and forwarding to USSPACECOM of associated predicted ephemerides is very helpful in properly identifying the new spacecraft and, in some cases, finding them in the first place. Finally, to render any desired anomaly support and to assume the appropriate CA posture for non-functional spacecraft, it is important to forward (or update on Space-Track.org) spacecraft status information, especially in the era of “disposable” satellites in which infant mortality is higher.

To address this topic, the following specific practices are recommended:

- A. To aid in satellite tracking and identification, provide injection vector(s) to USSPACECOM as soon as they are available.
- B. To assist in spacecraft identification for the cataloging process and provide general awareness among all O/Os, generate and forward predicted ephemerides for these spacecraft to USSPACECOM and publish these ephemerides (and all subsequent ephemeris updates) publicly as soon as contact is established with each deployed spacecraft.
- C. If USSPACECOM has issued TLEs for launched objects, notify USSPACECOM of the TLE and object number that are associated with your spacecraft.
- D. Provide early reporting to USSPACECOM of any spacecraft failures or other operational difficulties, both to obtain any available anomaly support and to assign the appropriate CA approach to the spacecraft (i.e., inactive and thus handled in a manner equivalent to a debris object).
- E. If using a commercial provider, make sure it has access to information from items A-D.



6. On-Orbit Collision Avoidance

For nearly all spacecraft, the on-orbit phase of their life cycle is the longest, meaning that the CA practices and CONOPS in place during this phase will have the greatest impact on the satellite's overall risk exposure.

As explained in Section 3, the CA process as presently structured comprises three phases:

1. **CA screenings** identify close approaches between a protected asset, called the primary satellite, and any other space objects, called (from the vantage point of the primary) "secondaries."
2. **CA risk assessment** examines each of the close approaches produced by the screening activity to determine which may represent dangerous situations and therefore require a mitigation action.
3. **CA mitigation** constructs a mitigation action, usually a trajectory change for the primary object, that will both reduce the collision risk of the close approach to an acceptable level and not create any fresh high-risk conjunction events.

Before any of these three activities can take place, certain satellite data and position information needs to be produced and made available to the CA enterprise.

These four aspects of CA (required input data and the three phases of the CA process) are each treated separately below.

6.1. Spacecraft Information and Orbital Data Needed for CA

Because CA identifies (and if necessary, mitigates) close approaches between spacecraft, it requires access to a comprehensive space catalog of orbital information. The USSPACECOM space catalog is the base catalog used by nearly all CA practitioners. Earlier sections of this document described the pre-launch registration and coordination process with USSPACECOM to receive basic launch and on-orbit CA services. Several commercial CA service providers offer CA products derived from alternative space catalogs. O/Os are encouraged to pursue commercial services, particularly when such services offer improvements and innovations above and beyond what is available from USSPACECOM. However, NASA recommends O/Os use the service offered by USSPACECOM as both a baseline and a supplement to commercially procured services. This is because the objects contained in the commercial catalog may not be the same as those in the DOD catalog and because the commercial vendor may provide CA services based on publicly available TLEs⁶ along with solutions for the objects in their own catalog. These TLE-based solutions are not sufficiently accurate to be used for CA. (See Appendix L for information about how NASA validates and uses commercial data.)

⁶ See Appendix E for more information about the use of TLEs.



The USSPACECOM space catalog used for CA contains position and uncertainty information for the primary object.

- If this object is not maneuverable, it is in principle possible to perform CA on its behalf entirely from catalog-based information. However, satellite O/Os generally have a better understanding of the spacecraft's construction and thus its non-conservative force parameters such as the ballistic coefficient and solar radiation pressure coefficient. So the O/Os' prediction of the satellite's future position, captured in a predictive ephemeris, is often more reliable or at the least an important adjoining datum to the future position information calculated from a space catalog entry.
- For satellites that do maneuver, future position calculations from catalog information alone will not capture any planned trajectory changes and thus will leave undiscovered any satellite conjunctions that could arise from the modified trajectory.

Therefore, for both non-maneuverable and maneuverable spacecraft, but especially for the latter, it is necessary that O/Os furnish predicted satellite state and uncertainty information, usually in the form of an ephemeris that includes state covariance at each ephemeris point. The details of the generation of such ephemerides that result in the most useful product for CA are described in the details portion that concludes this section.

While it is most important to submit such ephemerides to the screening entity, it is also helpful to place them on a public-facing website for any space operator to download and process. Mutual sharing of expected future positions is the best way for active satellites to avoid collisions with each other. Claims that predicted ephemerides contain proprietary information of any consequence are simply not compelling and, in any case, are outweighed by the safety benefit of exchanging such information. Such ephemerides should be recomputed and reissued/reposted as soon as a change to a spacecraft's intended trajectory is planned.

In determining how to react when a primary satellite is found to be in conjunction with a secondary, the logic path is governed heavily by whether or not the secondary object is an active, maneuverable spacecraft.

- If the object is non-maneuverable, then the object will follow a Keplerian orbit without unexpected perturbations. Future positions can be predicted in a straightforward way using a dynamical model.
- If the object is an active spacecraft that either has shown a history of maneuvering or is believed to be maneuverable, then it is quite possible that trajectory-changing maneuvers are planned and the assumption of a Keplerian orbit is not appropriate.

For this reason, it is important to know whether any given spacecraft is active, capable of maneuvering, and presently in a phase of satellite life in which maneuvering is possible. This status can be documented and communicated by the O/Os' setting of the spacecraft's active/dead and maneuverable/non-maneuverable flags in its Space-Track.org record. This allows other O/Os to determine whether the satellite can or will change its trajectory in a non-Keplerian manner and thus consider this possibility in CA.



O/Os of active, maneuverable spacecraft should send USSPACECOM a maneuver notification message outlining basic information about each planned maneuver. While in principle such information could be reconstructed from submitted ephemerides, it is simpler and more accurate to provide the information directly in this form. Providing maneuver notifications assists USSPACECOM in updating their own catalog as they know when to look for maneuver activity and, if discovered (and subsequently tracked), can expeditiously update the satellite's catalog entry to reflect the new trajectory. Of course, forwarding maneuver notifications is not a substitute for sending updated ephemerides that contain the intended maneuvers. Rather, it is an accompanying notification that allows better use of the received ephemerides and facilitates the USSPACECOM mission.

More O/Os are considering or are including autonomous flight control features, especially in large constellations in which the constellation management functions are complex. Autonomous flight control features may include orbit maintenance maneuvers and in some cases, even CA risk mitigation maneuvers that are developed, scheduled, and executed from the on-board control system without any active ground-based participation. While such approaches can offer improved efficiencies, they present their own challenges.

Even if satellite maneuvers are planned and executed autonomously, these planned maneuvers must be included in ephemerides and made available in near real time both to the screening provider and more broadly. This is because, recognizing the latencies of the download and distribution mechanisms, satellite maneuvers must not be executed without sufficient advance notice to allow the CA process to become aware of the intended maneuver and ensure its safety. The amount of advance notice required is governed by the latencies in the selected CA process, which includes the O/O infrastructure to receive and react to CA information from the screening provider. The principal difficulty is notifying other active, maneuverable satellites that may be contemplating their own maneuvers. Maneuver intentions must be shared with other such O/Os in a satellite's vicinity to ensure that intended maneuvers by either or both operators, if executed, do not place both satellites on a collision course.

To address this topic, the following specific practices are recommended:

- A. Actively maintain the Space-Track.org record for the satellite, updating the active/dead and maneuverable/non-maneuverable flags to reflect the satellite's current status.
- B. Furnish predicted ephemerides that include state covariances to USSPACECOM (and any additional commercial screening provider) and set the privileges to allow any interested party to access and download this information.
- C. Furnished ephemerides should possess the following characteristics:
 - a. Be of a 7-day predictive duration for LEO and 14 days for other orbits;
 - b. Be issued at least at the following frequencies:
 - i. Three times daily for spacecraft with perigee heights less than 500 km;
 - ii. Daily for other LEO orbits; and
 - iii. Twice weekly for other orbits. (*continued*)



- c. Include all known maneuvers within the ephemerides' prediction duration.
 - d. Provide ephemeris point spacing of approximately $1/100^{\text{th}}$ of an orbit, in either time or true anomaly. (Certain scientific missions with extremely long orbits or high eccentricities may require more specialized approaches.)
 - e. Contain a realistic covariance at each ephemeris point for at least the six estimated state parameters. (See Appendix I for a practical guide for assessing covariance realism and some general principles for tuning the covariance production process.)
 - f. Be formatted and distributed in the Consultative Committee for Space Data Systems (CCSDS) standard Orbital Ephemeris Message (OEM) format, preferably in the J2000 reference frame.
- D. Furnish maneuver reports to USSPACECOM for any trajectory-altering satellite maneuvers sufficiently in advance of maneuver execution to enable an O/O evaluation of the maneuver's safety. Employ the standard maneuver reporting message for this notification.
- E. When a maneuver becomes part of a satellite's trajectory plan, generate and submit to the screening provider an updated ephemeris that contains this new maneuver as early as is feasible but certainly with sufficient advance notice to enable an O/O evaluation of the maneuver's safety.

6.2. CA Screenings

The first step in the CA process is to find close approaches between protected (primary) objects and other objects, with the latter object set's positions represented either by ephemerides (if they are active spacecraft) or space catalog entries. Finding these close approaches is accomplished by CA screenings in which the predicted positions of the primary object and all other space objects that survive a pre-filtering process are compared.

At USSPACECOM, CA screenings are executed using a volumetric-based approach: a physical screening volume (usually an ellipsoid) is "flown" along the primary object's trajectory, and any secondary trajectories that penetrate this volume are considered conjuncting objects. This screening is accomplished using a tool called the Astrodynamics Support Workstation (ASW).

Volumetric screenings generate more close approaches than probabilistic or "covariance-based" methods but are preferable because they essentially give a snapshot of the satellite catalog in the vicinity of the primary's ephemeris. If the screening volume is large enough, the conjunction information can be used to determine the safety of not just the nominal trajectory but also any reasonably sized maneuvers that the primary may choose to make.

Once these close approach objects are identified, further processing is invoked to determine the precise time of closest approach between the two objects and the two objects' states and covariances at that time.

A protected asset should be screened for close approaches against a comprehensive satellite catalog at least daily with the results of this process obtained and processed by the O/O, also at least daily. For active, maneuverable satellites that submit ephemerides to the screening process



as a way of including their planned maneuvers into their predicted trajectories, these ephemerides should be screened against each other (subject to appropriate pre-filtering) in near real time whenever an O/O submits an updated ephemeris. This is called an “O/O vs O/O” screening. O/O vs O/O screenings are the best way to ensure that maneuver plans are communicated among O/Os to prevent simultaneous maneuvers from causing a collision.

Current USSPACECOM practice is to conduct three screenings per day. Each screening predicts close approaches between:

1. The ASW solution for each protected (active) asset against the full catalog;
2. O/O ephemerides submitted after the last screening against the full catalog; and
3. O/O ephemerides submitted after the last screening against all other unexpired O/O-submitted ephemerides.

In other words, O/O-submitted ephemerides are screened against the full catalog when initially submitted, then retained and screened against other O/O ephemerides until they expire.

Three screenings per day is a minimum frequency for higher-drag orbit regimes such as LEO orbits with perigee heights below 500 km; daily screenings for other orbit regimes yields sufficient accuracy.

(See the USSPACECOM Spaceflight Safety Handbook for details about the cadence and conduct of CA screenings.)

To address this topic, the following specific practices are recommended:

- A. Submit predicted ephemerides for the spacecraft to a screening provider to be screened for conjunctions at least daily with spacecraft in higher-drag orbit regimes screened at least three times per day.
- B. Ensure that an O/O ephemeris for an active, maneuverable spacecraft is screened against other ephemerides from active, maneuverable spacecraft in near real time after any such ephemeris is submitted to the screening provider.
- C. Obtain and process these screening results from the screening provider at the same frequency at which they are produced for both the full-catalog and O/O vs O/O screening cases described above.

6.3. CA Risk Assessment

Satellite conjunctions are approaches between two satellites closer than a specified set of distances, which is often chosen to be much larger than the set that would pose an actual collision threat. A key next step, therefore, is to evaluate each conjunction to determine if it poses, or is likely to pose, a substantial risk of collision. This evaluation activity, called CA risk assessment, comprises both the key calculations that feed the concept of risk (that is, both collision likelihood and collision consequence) and the evaluations of the input data that subtend these calculations to ensure that they constitute a basis for decision making. Only conjunctions that are determined to be high risk merit the consideration of mitigation actions.



Methods proposed in the scientific literature for calculating and assessing satellite collision likelihood have different merits and different risk tolerance orientations. Of all the possibilities, P_c is the oldest, most straightforward, and most widely embraced collision likelihood parameter. Because it is a present industry standard, NASA recommends that operators employ P_c as the foundational element of their collision likelihood assessment. (See Appendix M for an extended treatment of the recommendation for the use of P_c as the collision likelihood parameter.)

Two approaches to calculating the P_c are:

1. The “two-dimensional” P_c calculation approach, which was originally introduced by Foster and Estes in 1992 with theoretical clarifications by Alfano (2005b) and Chan (2008). This method is a durable simplification of the calculation that is valid in most instances, but there are situations in which it does not perform accurately.
2. The “three-dimensional” P_c calculation method, which was originally formulated by Coppola (2012) and “repaired” and extended by CARA. Once it has completed final NASA validation, it is expected to be the recommended analytic approach. This method is accurate for nearly every situation, tests exist to identify those very few cases in which it may not be fully adequate, and it is computationally efficient.

Additionally, the primary and secondary object covariances for some orbit regimes can contain significant correlated error, thus introducing inaccuracy in the calculated P_c . (See Appendix N for a more detailed discussion of P_c calculation methods and related issues.)

Software (including source code, test cases, and documentation) to calculate the P_c using both the traditional two-dimensional and (when validated) the presently recommended three-dimensional technique and to remove correlated error between the primary and secondary covariances can be obtained free of charge at the public-facing CARA software repository. (See Section 7 in this document for the specific URL.)

The value of the P_c threshold at which an O/O would choose to mitigate a conjunction is a function of both the O/O’s risk tolerance and the volume of conjunctions that the O/O is able to mitigate. In the CA mission area, however, there has been broad convergence on a per-event P_c mitigation threshold value of $1E-04$, meaning that remediation actions are recommended when the likelihood of collision is greater than 1 in 10,000. Missions that wish a more conservative risk posture can select a more demanding P_c threshold such as $1E-05$. This infused conservatism may allow more streamlined risk assessment techniques in other areas.

Finally, some practitioners like to include an additional imperative to mitigate when the predicted miss distance is smaller than the hard body radius or a distance close to the hard-body-radius value. This predicted miss distance approach would need to be invoked only rarely because a small miss distance typically produces a P_c that violates most commonly accepted P_c mitigation thresholds. (See Appendix N for more information about the hard-body-radius calculation.)

Other risk assessment methods exist (e.g., Alfano 2005b, Carpenter and Markey 2014, Balch et al. 2019) and are generally considered to be more conservative than the P_c -based methodology advocated above. O/Os are encouraged to employ a more conservative CA approach if it suits



their risk posture. However, the P_c method and the mitigation threshold of $1E-04$ enjoy wide acceptance in the CA industry and historically have provided a sustainable balance between safety and mission impact.

Risk is the product of the likelihood of an unfavorable event and the consequence of that event, should it occur. For some time, it has been presumed that any satellite collision is a catastrophic event that is to be avoided by any means possible. In most cases, such events will render an active satellite unusable, presenting a very serious risk to the O/O. Therefore, the likelihood (P_c) was computed, and the consequence was not factored into the computation.

However, as the orbit environment grows, there may come a point when an O/O is faced with too many conjunctions to avoid individually. In this case, one method to triage the conjunctions to determine which should be mitigated to best protect both the spacecraft and the orbital environment is computation and application of the consequence as part of the risk determination.

Satellite collisions, depending on satellite masses and relative velocity, can produce wildly different amounts of orbital debris from a handful of pieces to many thousands of pieces large enough to critically damage a spacecraft. To be sure, introducing any debris at all into the space environment is to be avoided. Conjunctions that, if they were to result in a collision, would produce only a small amount of debris (maybe fewer than 50 pieces) could be addressed using a more lenient mitigation threshold to prioritize them appropriately against those that would create more debris. A threshold that is an order of magnitude more lenient than what is used for high-debris conjunctions would align in magnitude with other situations in which relaxing the mitigation threshold is warranted (Hejduk et al. 2017). (See Appendix O for a more expansive treatment of the assessment of a conjunction's debris production potential. See the CARA software repository for software to calculate debris production potential; see Section 7 in this document for the specific URL.)

Finally, the data used to calculate the parameters that feed the risk assessment decision, namely the state estimates and accompanying uncertainty volumes (covariances) for the primary and secondary objects, must be examined to determine whether they manifest problematic elements that would prevent the calculated parameters from serving as a basis for CA mitigation actions. In such cases, it is possible that executing a mitigation action based on those data could make the CA situation worse rather than better.

A non-actionability situation with USSPACECOM CA products occurs occasionally. For example, an object whose state has been propagated longer than the orbit determination fit-span in order to reach TCA is considered under most conditions to be insufficiently tracked and over-propagated and thus not suitable as a basis for CA mitigation actions. (See Appendix P for a recommended procedure to evaluate USSPACECOM CA products for actionability.)

Occasionally one of the two objects in a conjunction lacks a covariance, making probabilistic CA impossible but still enabling other risk assessment methods. (See Appendix P for a discussion of such situations.)

To address this topic, the following specific practices are recommended:

- A. Use the Probability of Collision (P_c) as the principal collision likelihood assessment metric. (*continued*)



- B. Pursue mitigation if the P_c value exceeds $1E-04$ (1 in 10,000).
- C. Pursue mitigation if the estimated total miss distance is less than the hard-body-radius value.
- D. Employ the current operational NASA P_c calculation methodology for routine P_c calculation. Consider removing correlated error from the primary and secondary object joint covariance.
- E. As a prioritization method for situations in which the number of conjunctions meeting mitigation criteria exceeds the ability of the O/O to mitigate, estimate the amount of debris that a conjunction would produce if it were to result in a collision. A less stringent P_c an order of magnitude lower could be appropriate in such cases.
- F. If employing USSPACECOM data products for CA, use the procedure given in Appendix P to determine whether the data for a particular conjunction are actionable and thus constitute a basis for CA-related decisions.
- G. If a different CA product provider is chosen, develop and employ data actionability criteria for this provider's CA information to determine CA event actionability.

6.4. CA Risk Mitigation

O/Os with satellites that possess the ability to mitigate conjunctions have a responsibility to perform mitigation actions when required. In practical terms, this means that a mitigation action should be tendered when a conjunction's risk parameter, usually the P_c , exceeds the mitigation threshold at the "maneuver commitment point," which is that point in time before the TCA when a decision to mitigate is needed in order for the mitigation action to be in place before the TCA actually occurs. The most typical (and effective) mitigation action is changing the satellite's trajectory to avoid a possible collision with the secondary either by thrusting to effect a satellite maneuver or, in some cases, changing the satellite's attitude to alter its drag coefficient and thus its trajectory. An additional, although generally much less effective approach is to change the satellite's attitude simply to reduce the satellite's cross-sectional area in the direction of the oncoming secondary. This does not fully mitigate the close approach but reduces the likelihood of an actual collision between the two objects.

When planning a mitigation action, the general practice is to choose a trajectory alteration that reduces the P_c for the conjunction by 1 to 2 orders of magnitude. A recent study (Hall 2019a) showed that a P_c reduction by 1.5 orders of magnitude (that is, by a factor of ~ 0.03) marked the beginning of diminishing return with regard to lifetime satellite collision risk. Therefore NASA recommends a minimum of 1.5 orders of magnitude reduction in P_c as a post-mitigation P_c goal; i.e., if the recommended mitigation threshold of $1E-04$ is used, the post-mitigation goal would be $\sim 3.2E-06$ or lower.



In general, there is a trade-off between mitigation maneuver size and maneuver execution time. To mitigate a conjunction adequately, smaller maneuvers can be used if the maneuver is made earlier; that is, farther in advance of TCA. Waiting until closer to TCA will typically require a larger maneuver to achieve the same P_c reduction. However, because most CA events drop off to a low level of risk before TCA (due to additional tracking and improved state estimates and covariances), waiting until closer to TCA to perform a mitigation action increases the likelihood that it can be determined not to be necessary. Thus, O/Os must decide in each case whether they wish to act earlier with a less invasive action required or whether they wish to wait until closer to TCA in the hopes of observing the collision likelihood drop below the mitigation threshold and thus being able to waive off the maneuver, but knowing that, if an action is still required, it will be larger and more disruptive. There is no clear rubric for how to proceed in such cases, but external considerations, such as staffing availability for a later maneuver and amount of disruption that a large maneuver would require, often govern the decision.

When performing a trajectory-changing mitigation action, “new” conjunctions and/or elevated P_c values of existing conjunctions often occur. O/Os should ensure any mitigation action does not render the overall safety-of-flight evaluation worse by producing a more dangerous situation than would have existed without the mitigation action. One could examine the amalgamated risk of all of the conjunctions in the near future and conduct mitigation planning on this basis, but it is generally acceptable to pursue mitigation of the one conjunction that violates the threshold (and remediate it to a P_c 1.5 orders of magnitude below the threshold) while ensuring that no new conjunctions are produced that exceed the threshold. Using the thresholds recommended here, this means bringing the violating conjunction down to a P_c of $3.2E-06$ without introducing or raising any other conjunction P_c values above $1E-04$. The best way to ensure that these standards are met is to generate an ephemeris that includes the mitigation action and submit it to the screening provider for a special screening action. This will produce a fresh set of CDMs against the planned mitigation action and allow easy and reliable verification that the P_c values for both the principal conjunction and any ancillary conjunctions remain below the desired levels.

Using larger screening volumes and analyzing the maneuver against the resulting CDMs is another method of determining whether the maneuver creates close approaches. This method creates CDMs for conjunctions that are somewhat far away from the nominal trajectory but might constitute worrisome conjunctions after a maneuver. This method for ensuring the safety of the post-maneuver trajectory is serviceable for debris objects, which will not alter their orbits beyond Keplerian models. However, this method does not protect against potential, yet-to-be-disclosed trajectory changes by other maneuverable satellites.

Choosing a mitigation action for a conjunction against a non-maneuverable secondary is generally straightforward, but the situation is much more complicated when the secondary object is an active, maneuverable spacecraft. While submitted ephemerides from that secondary’s O/O is a statement of that spacecraft’s intended trajectory at the time of production, it is likely that the secondary’s O/O has also noticed the conjunction and may be planning a mitigation action of its own. In addition to the fact that two mitigation actions for the same close approach are



unnecessary, there is a real danger that the two mitigation actions could place the two satellites on even more of a collision course than taking no action at all.

Establishing contact with the O/O of an active secondary is critical to jointly establishing a way forward for the particular conjunction. The O/Os should jointly decide which satellite will pursue mitigation (should it still remain necessary at the maneuver commitment point) and what the mitigation action will be, and constrain the trajectory of the non-mitigating satellite to that given in its published ephemeris until after the TCA has passed.

Given the recent industry trend to extremely large satellite constellations, the individual reaching out to other O/Os is unlikely to scale adequately as such constellations continue to grow. A more automated mechanism for exchanging maneuver intentions among O/Os will likely be necessary. SPD-3 directs the movement of CA activities away from DOD to a Federal civil agency by 2024, so it makes sense for this new entity to develop an architecture and protocol for the exchange of this type of information. For the present, O/Os need to share their maneuver intentions with other O/Os through direct contact or a third-party organization, such as the Space Data Association (SDA),⁷ with some amount of manual interaction required.

To address this topic, the following specific practices are recommended:

- A. When a conjunction's P_c at the mitigation action commitment point exceeds the mitigation threshold (recommended to be $1E-04$), pursue a mitigation action that will reduce the P_c by at least 1.5 orders of magnitude from the remediation threshold.
- B. Ensure that an ephemeris containing the mitigation action is screened against the full catalog, not a large screening volume collection of CDMs.
- C. Ensure that the mitigation action does not create any additional conjunctions with a P_c value above the mitigation threshold (for which the recommended value is $1E-04$).
- D. When the secondary object is an active, maneuverable spacecraft, reach out to the secondary's O/O and jointly establish a way forward for the particular conjunction, including deciding which spacecraft will maneuver and freezing the other spacecraft's planned trajectory until the TCA has passed.
- E. Use Space-Track.org contact information to engage other O/Os.

⁷ An international organization that brings together satellite operators to support the controlled, reliable, and efficient sharing of data critical to the safety and integrity of the space environment. The SDA membership includes the world's major satellite communications companies.



6.5. Automated Trajectory Guidance and Maneuvering

Satellite design is increasingly taking advantage of the automation of trajectory guidance to simplify mission operations. Particularly with large constellations of satellites, automation can substantially increase efficiency over traditional flight control techniques. However, the automation of trajectory guidance and control in which flight dynamics computations and decision making are placed entirely onboard presents real concerns related to space traffic management and the prevention of collisions between satellites. With traditional ground-computed and executed trajectory adjustments, usually with humans in the loop, satellite operators can compute maneuvers, produce predicted ephemeris data to transmit to the rest of the community, and command the execution of those maneuvers to effect that advertised trajectory. Admittedly, such a technique does not scale easily for large constellations, but it has proven reliable and serviceable for securing safety of flight over years of experience.

However, highly automated systems now can be designed in which satellites, using onboard navigation and a programmed target, can compute their own maneuvers and execute them with no input from the ground. The knowledge of an intended trajectory and intended maneuvers might not even be communicated to ground control in advance of maneuver execution, making timely notification to USSPACECOM and thereby the rest of the space community difficult and in some cases impossible. If other space operators are not aware that a satellite is planning a maneuver, the other operator may also plan a maneuver and the two maneuvers taken together may cause a collision.

Operators should therefore ensure that when an autonomously controlled satellite computes any maneuver to alter its orbit for mission or CA purposes, it communicates the intended time, direction, and magnitude of that maneuver (as well as other event-related data for CA maneuvers) to the operating ground control a sufficient interval before the maneuver such that the maneuver plan can be transmitted to USSPACECOM for screening. This communication is most important in ensuring safety of flight between active maneuverable spacecraft, each of which may be generating maneuver plans that, if executed as intended, will create a collision.

Maneuver plans should be shared with USSPACECOM both as a maneuver notification report and a predicted ephemeris. Operators should ensure that the automated process includes an accurate on-ground satellite emulation capability that can, with reasonable accuracy, predict the behavior of the autonomously controlled satellite and generate related trajectory products such as maneuver reports and predicted ephemerides. This emulation may, in some cases, provide the best (and only) position predictions available to the space community.

Once a maneuver is autonomously planned and communicated to the ground, the maneuver should be executed as planned unless an alteration is required for safety of flight. Remaining with the communicated plan, even if more efficient possibilities arise in the interval leading up to maneuver execution, enables the rest of the space community to use the updated ephemeris as the basis for their own trajectory planning

There may be situations in which a satellite should not execute an intended maneuver that the automation computes, such as a surprise post-maneuver conjunction that has not been uploaded or a conjunction with an active satellite involving cooperative, human-in-the-loop CA event



management. It is thus imperative for safety reasons that an automated maneuvering system include the ability to pause or abort from the ground any maneuver planned by the on-board automated controller.

At some point in the future, an automated clearinghouse may be possible for such information to allow near real time receipt and circulation of plans/updates to other O/Os, but for the present, these data will need to be transmitted and shared using existing mechanisms. For this reason, systems that intend to incorporate autonomous satellite operations should follow the practices enumerated below.

To address this topic, the following specific practices are recommended

- A. When an autonomously controlled satellite computes any maneuver to alter its orbit for mission or CA purposes, communicate maneuver details to the operating ground control early enough to enable a USSPACECOM screening and appropriate action in response to the screening results.
- B. Share maneuver plans with USSPACECOM both as a maneuver notification report and a predicted ephemeris.
- C. Deploy an accurate on-ground satellite emulation capability to predict and share the behavior of the autonomously controlled satellite in predicted ephemerides.
- D. Execute the maneuver as planned unless an alteration is required for safety of flight once a maneuver is autonomously planned and communicated to the ground.
- E. Include the ability to pause or abort from the ground, for safety reasons, any maneuver planned by the on-board automated controller.

6.6. Other Considerations

The items in this section are best practices that affect orbit design and so, while not directly related to CA process, are appropriate to consider.

Light Pollution. After the launch of the initial portion of the SpaceX Starlink constellation, astronomers noticed that their observing campaigns were being affected by extremely bright, detector-saturating streaks due to the Starlink vehicles. To minimize interference with Earth-based astronomy, O/Os should ensure that spacecraft are built and flown in such a way as to minimize the creation of light pollution. Observatories are most challenged by spacecraft in higher orbits and those having inclinations close to the observatory latitude. Counterintuitively, the higher altitudes result in a lower relative velocity between satellite and telescope, which increases target dwell-time on each pixel, which facilitates saturation. Similarities between orbit inclinations and observatory latitudes result in regular crossings of the telescope's field of view.

At Starlink altitudes (around 550 km), objects having a brightness dimmer than the 7th stellar magnitude will not affect astronomical observations. Factors that affect the visual magnitude of an object include altitude, attitude, albedo, size, surface characteristics, degree of specular versus diffuse reflection, self-shadowing, and solar phase (and sometimes solar declination) angle,



although ground-based astronomy will be affected more by the sheer numbers of large objects in orbit than the individual object's photometric brightness as such.

In the case of large constellations in which the same satellite design will be duplicated a large number of times, O/Os should not rely on analyses with first-order assumptions about the satellite's expected brightness. Rather, a fully modeled brightness solution that includes a Bidirectional Reflectance Distribution Function (BRDF) characterization of the satellite's construction materials should be undertaken during the design phase. O/Os should consider all phases of the lifetime including launch, ascent, and descent in addition to mission orbit (Seitzer 2020).

To address this topic, the following specific practices are recommended:

- A. As part of spacecraft physical design and orbit selection, perform a spacecraft photometric brightness analysis to determine whether the spacecraft is likely to present an impediment to ground-based astronomy. Consider changes to the satellite's construction, materials, or operating attitudes to reduce expected photometric brightness to levels that will not impede ground-based astronomy.
- B. If a large constellation is being planned, construct a full photometric model of the spacecraft, including full BRDF information for the satellite surfaces, to obtain a durable estimate of the satellite's expected photometric brightness.
- C. If the constellation, given its population, orbit, and constituent satellites, is likely to affect ground-based astronomy, reassign the satellite orbits or modify the satellite construction to eliminate this effect.



7. Contact Information

For feedback on this handbook: ca-handbook-feedback@nasa.onmicrosoft.com

Organization	Contact Information and URLs of Interest
USSPACECOM / 18 SPCS	Diana McKissock; Cynthia Wilson 18SPCS.DOO.CustomerService@us.af.mil +1-805-606-2675 <i>Spaceflight Safety Handbook for Satellite Operators</i> https://www.space-track.org/documents/Spaceflight Safety Handbook for Operators.pdf <i>Launch Conjunction Assessment Handbook</i> https://www.Space-Track.org/documents/LCA Handbook.pdf <i>Space-Track Handbook for Operators</i> https://www.Space-Track.org/documents/Spacetrack Handbook for Operators.pdf Description of Space Situational Awareness Sharing basic and advanced services, forms, and data examples: https://www.Space-Track.org/documentation#/odr
USSPACECOM	Scott Van Sant Scott.m.vansant.civ@mail.mil
NASA Examples	NASA Examples of Information to Expedite Review of Commercial Operator Applications to Regulatory Agencies: https://www.nasa.gov/recommendations-commercial-space-operators
NASA JSC FOD	Human Space Flight Conjunction Assessment Operations jsc-dl-topo-iwg@mail.nasa.gov
NASA CARA	NASA CARA Operations cara-management@lists.nasa.gov
NASA CARA Software Repository	https://github.com/nasa/CARA_Analysis_Tools



Appendix A. Acronyms

18SPCS	18 th Space Control Squadron
ANOVA	Analysis of Variance
ASW	Astrodynamics Support Workstation (tool)
BRDF	Bidirectional Reflectance Distribution Function
CA	Conjunction Assessment
CAD	Computer-Aided Design
CARA	NASA Conjunction Assessment Risk Analysis (Program)
CCSDS	Consultative Committee for Space Data Systems
CDF	Cumulative Distribution Function
CDM	Conjunction Data Message
CME	Coronal Mass Ejection
COLA	Collision Avoidance
CONOPS	Concept of Operations
CPU	Central Processing Unit
DCP	Dynamic Consider Parameter
DLA	Dynamic LUPI Assignment
DOC	U.S. Department of Commerce
DOD	U.S. Department of Defense
ECI	Earth-Centered Inertial
EDF	Empirical Distribution Function
EDR	Energy Dissipation Rate
eGP	Extrapolated General Perturbation
FAA	U.S. Federal Aviation Administration
FCC	U.S. Federal Communications Commission
FDO	Flight Dynamics Officer (JSC FOD)
FFRDC	Federally Funded Research and Development Center
FOD	(NASA Human Space) Flight Operations Directorate (JSC)
GEO	Geosynchronous Equatorial Orbit
GMT	Greenwich Mean Time
GPS	Global Positioning System
GSFC	Goddard Space Flight Center
HAC	High Accuracy Catalog



HASDM	High Accuracy Satellite Drag Model
HEO	High Earth Orbit
ISS	International Space Station
ITU	International Telecommunication Union
JBH09	Jacchia-Bowman HASDM 2009
JSC	Johnson Space Center
L	Launch
LCOLA	Launch Collision Avoidance
LEO	Low Earth Orbit
LUPI	(Dynamic) Length of Update Interval
MC	Monte Carlo
MET	Mission Elapsed Time
NASA	National Aeronautics and Space Administration
NPR	NASA Procedural Requirements (document)
OADR	Open Architecture Data Repository
ODMSP	USG Orbital Debris Mitigation Standard
ODPO	NASA Orbital Debris Program Office
ODR	Orbital Data Request
OEM	Orbit Ephemeris Message (CCSDS)
O/O	Owner/Operator
OPM	Orbital Parameter Message (CCSDS)
OSA	Orbital Safety Analyst
P _c	Probability of Collision
PDF	Probability Density Function
R-15	Ready Minus 15
RAAN	Right Ascension of the Ascending Node
RIC	Radial – In-track – Cross-track
SDA	Space Data Association
SDK	Software Development Kit
SGP4	Simplified General Perturbation Theory #4
SMA	Semi-Major Axis
SP	Special Perturbations
SPCS	Space Control Squadron
SPD-3	Space Policy Directive-3



SSA	Space Situational Awareness
SSN	Space Surveillance Network
SWTS	Space Weather Trade Space
TCA	Time of Closest Approach
TLE	Two-Line Element
TOPO	Trajectory Operations Officer (JSC FOD)
URL	Uniform/Universal Resource Locator
U.S.	United States
USG	United States Government
USSPACECOM	United States Space Command
VAFB	Vandenberg Air Force Base
VCM	Vector Covariance Message
VIM	(R-15) Vehicle Information Message



Appendix B. Glossary

COLA Gap. The period between when traditional LCOLA ends and an on-orbit asset can mitigate a risk using standard on-orbit CA methods (i.e., until the object can be cataloged and is available for CA screening).

Collision Avoidance. The planning and execution of risk mitigation strategies.

Conjunction. A close approach between two objects that is predicted to occur because the secondary object passes within a chosen geometric or statistical safety volume about the primary (protected) asset.

Conjunction Assessment. The identification of close approaches using ephemeris screening against a catalog of identified resident space objects.

Conjunction Mitigation. An action taken to remediate conjunction risk, including a propulsive maneuver, an attitude adjustment (e.g., for differential drag or to minimize frontal area), or providing ephemeris data to the secondary operator so that they can perform an avoidance maneuver.

Conjunction Risk Assessment. The determination of the likelihood of two space objects colliding and the expected consequence if they collide in terms of lost spacecraft and expected debris production.

Ephemeris (plural: ephemerides). The trajectory (i.e., the position (and possibly velocity) over time) of objects in the sky.

Geosynchronous Equatorial Orbit (GEO). An orbit with period between 1300 and 1800 minutes, eccentricity (e) <0.25 , and inclination (i) <35 degrees. Note that this definition is chosen to align with the definition used by the 18 SPCS, who provides CA screenings.

High Earth Orbit (HEO). An orbit with period greater than 225 minutes and $e >0.25$. Note that this definition is chosen to align with the definition used by the 18 SPCS, who provides CA screenings.

Low Earth Orbit (LEO). An orbit with period less than 225 minutes and $e <0.25$. Note that this definition is chosen to align with the definition used by the 18 SPCS, who provides CA screenings.

Maneuver Plan. The specific parameters that represent a planned spacecraft maneuver, including execution time, burn duration, and delta- v . The industry standard for this message is the Consultative Committee for Space Data Standards (CCSDS) Orbital Parameter Message (OPM).

Maneuverable. Capability to carefully and skillfully guide or manipulate a spacecraft, such as to avoid an orbital conjunction.

Maneuverable Object. An object that can alter its trajectory substantially such that standard orbital dynamics models cannot predict its location.



Non-Human Space Flight Mission. A NASA space flight mission that is not related to human space flight. These space flight missions are supported by the CARA Program.

Orbit Regime. The term orbit regime refers to the general location of an orbit in space, such as the perigee height and whether the orbit is eccentric, in order to assign it to a general orbit category.

Space Flight Mission. NASA space flight programs, projects, and activities (including spacecraft, launch vehicles, instruments developed for space flight programs and projects, some research and technology developments funded by and to be incorporated into space flight programs and projects, critical technical facilities specifically developed or significantly modified for space flight systems, highly specialized information technology acquired as a part of space flight programs and projects (non-highly specialized IT is subject to NPR 7120.7, *NASA Information Technology Program and Project Management Requirements*), and ground systems that are in direct support of space flight operations). This also applies to reimbursable space flight programs and projects performed for non-NASA sponsors and to NASA contributions to space flight programs and projects performed with international and interagency partners. (See NPR 7120.5, *NASA Space Flight Program and Project Management Requirements*.)

Space Surveillance Network. A network of radar and optical sensors used by DOD to track space objects. Tracking data is used to perform orbit determination and maintain the space object catalog.

Systematic Conjunction. A situation in which two space objects repeatedly experience close approaches with each other due to their similar orbits.



Appendix C. Best Practices List

This appendix contains a summary list of the best practices found in this document. NASA missions should use Appendix D Best Practices for NASA Missions instead for a list that integrates internal NASA services and requirements.

Table C-1 Best Practices Summary

Section	Best Practice	Comment
3.0 History		
3.1 USSPACECOM CA Process		
3.1.A	Obtain a Space-Track.org account.	
3.1.B	Become familiar with the Spaceflight Safety Handbook contents.	
3.1.C	Use USSPACECOM CA service as a minimum for screenings, even if additional commercial data or services are used.	
3.1.1.A	Provide seven (7) days of predicted ephemeris (including maneuvers) to USSPACECOM for screening for LEO spacecraft, and 14 days for other orbits (e.g., HEO/GEO).	
3.1.2.A	Provide at least one (1) ephemeris per day to be screened; three (3) ephemerides per day in the lower-drag LEO regime (perigee height less than 500 km).	
3.1.2.B	Determine whether the O/O's process for obtaining screening results and performing CA risk assessment aligns with the timeline of the USSPACECOM process. If the timelines do not align in such a way as to enable timely and efficient screening support, pursue a rearrangement of the O/O's process to minimize data latency and optimize screening efficiency.	
3.1.3.A	Populate and maintain the point-of-contact section on Space-Track.org with your operations contact data. Be sure that the operations contact can be reached 24/7 due to time zone differences between operators and the immediate nature of certain CA emergencies.	
3.1.3.B	Use standard ephemeris, CDM, and maneuver notification formats defined by CCSDS.	
3.1.3.C	Large constellation operators should work with USSPACECOM pre-launch to determine if variations from the standard approach are necessary and, if so, to define a tailored screening process.	
3.2 NASA Partnership with USSPACECOM		
3.2.A	Develop a robust safety-of-flight process that includes both CA screening and risk assessment to inform close approach mitigation decisions.	



Section	Best Practice	Comment
3.2.B	Large constellation operators should work with USSPACECOM pre-launch to determine if variations from the standard approach are necessary and, if so, to define a tailored screening process.	
3.2.C	Large constellation operators should consider working with NASA to define a risk assessment process. NASA is an experienced representative operator. Including NASA in discussions regarding establishing a CA process will ensure that the process will work with most operators for risk assessment purposes.	
4.0 Spacecraft and Constellation Design		
4.1 General Orbit Selection: Debris Object Density		
4.1.A	Include considerations of orbital debris density when selecting candidate mission orbits. Debris density diagrams, indexed by object size, are published by the NASA ODPO; more information and examples are provided in Appendix G.	
4.1.B	When choosing among the final set of candidate orbits, perform a conjunction frequency analysis to determine the number of CA high interest events per spacecraft per year that can be expected in each of the proposed orbital areas.	
4.2 Vehicle and Constellation Specific Orbit Selection: Spacecraft Colocation		
4.2.A	During orbit selection, perform an orbit colocation analysis to determine whether any of the proposed orbits are likely to create systematic conjunctions with existing actively maintained satellites.	
4.2.B	If the orbit colocation analysis identifies systematic conjunctions, consider modifying the proposed orbit(s) slightly to eliminate this possibility. Optimizing orbit placement may be appropriate.	
4.2.C	If the selected orbit is likely to present systematic conjunctions with a pre-existing spacecraft, then coordinate with the other operator(s) to arrange a process to coordinate maneuver plans routinely during the life of the mission.	
4.3 Ascent to/Disposal from the Constellation's Operational Orbit		
4.3.A	Perform a study to compute the number of expected close approaches anticipated during ascent and descent, as well as the imputed additional satellite reliability that will be required, to meet satellite disposal requirements at the chosen operational orbit.	
4.3.B	If the results of the study show a large burden, consider choosing a different mission orbit with a lower burden.	



Section	Best Practice	Comment
4.3.C	All missions should review and follow the ODMSP guidance standards.	
4.3.D	When practicable, pursue active disposal using the fastest disposal option available.	
4.3.E	Recognize that transiting spacecraft should yield way to on-station spacecraft and thus take responsibility for any CA risk mitigation maneuvers or transient trajectory alterations that may be required.	
4.3.F	When planning descent using thrusters, if not planning an approach of circular-orbit altitude reduction, coordinate with NASA JSC TOPO to ensure that perigee-lowering approaches do not present persistent and problematic orbital crossings with the ISS and other human space flight assets.	
4.4 Spacecraft Trackability		
4.4.A	<p>Through selection of physical design and materials, ensure that the satellite is trackable by SSN.</p> <ol style="list-style-type: none"> For spacecraft with perigee heights less than 2000 km, the spacecraft should have characteristic dimensions of at least 10 cm in each major dimension. For spacecraft with perigee heights greater than 2000 km, the spacecraft should have characteristic dimensions of at least 50 cm in each major dimension. If the spacecraft cannot meet these dimension constraints, use a proven detectability enhancement. 	
4.5 Spacecraft Reliability		
4.5.A	Ensure that spacecraft reliability is high enough that the likelihood of each spacecraft to remain fully functional until it can be disposed meets or exceeds 99%.	
4.5.B	Reassess the 99% analysis whenever the underlying assumptions change, for example, extending operations beyond design life or failures in key systems.	
4.6 Development of Capabilities for Ephemeris Generation and CA Risk Assessment and Mitigation		
4.6.A	Develop and implement a capability to generate and share accurate predicted spacecraft ephemerides, including any planned maneuvers. (See Section 3.1 for a discussion of ephemeris sharing for screening. See Section 6.2 for specific guidance on ephemeris generation frequency, length, point spacing, formatting, and contents.)	



Section	Best Practice	Comment
4.6.B	Determine during spacecraft design what sorts of risk mitigation approaches are possible for the spacecraft. While trajectory modification via thrusting maneuver is the most common, there are other possibilities such as differential drag orbit modification.	
4.6.C	Develop and implement or arrange to acquire a capability to process CDMs and to compute conjunction risk assessment parameters such as the Pc.	
4.6.D	Develop and implement or arrange to acquire a risk analysis capability to select mitigation actions that will lower the Pc for dangerous conjunctions to a user-selected value.	
4.6.E	Validate CA tools well before flight, typically 6-12 months prior.	
5.0 Pre-Launch Preparation and Early Launch Activities		
5.1 CONOPS Discussions/Arrangements with USSPACECOM Pertaining to Launch Phase		
5.1.A	Establish contact with USSPACECOM to describe and discuss all the different aspects of the launch including delivery and deployment methodologies.	
5.1.B	Deployment of multiple spacecraft should be spaced in a way that enhances the ability of USSPACECOM to quickly identify and differentiate the spacecraft using the SSN.	
5.1.C	Discuss the preparation of the R-15 launch information message with USSPACECOM so that they understand its contents and any implications for safety of flight.	
5.1.D	Ensure that the launch provider submits launch trajectory information to USSPACECOM, including ephemerides for the powered-flight portions and orbital elements for the parking orbits so that potential colocation and deconfliction potentialities can be discovered.	
5.1.E	Provide launch-related information to USSPACECOM as soon as possible, such as injection vectors and initial ephemerides for deployed spacecraft, which can be used to assist with the cataloging process, especially to confirm the identity of launch-related objects. Coordinate a satellite numbering scheme (potentially including temporary satellite numbers) appropriate to the launch type and expected degree of cataloging difficulty.	
5.1.F	Coordinate with USSPACECOM any potential launch anomaly diagnostic products that can be provided should issues arise during the launch and early orbit sequence.	



Section	Best Practice	Comment
5.2 CONOPS Discussions and Arrangements with USSPACECOM Pertaining to Catalog Maintenance		
5.2.A	Register the spacecraft with USSPACECOM using the Satellite Registration form on Space-Track.org.	
5.2.B	Provide USSPACECOM with basic construction and mission information about the satellite, such as stowed dimensions; deployable structures such as solar panels, antennae, booms, including all of their (rough) dimensions; satellite material properties and colors; regular satellite attitude; and registered operational radio frequencies.	
5.2.C	Provide USSPACECOM with a basic overview of the satellite’s orbit maintenance strategy, which would include the paradigm for determining when orbit maintenance maneuvers are required; the maneuver technology used (as this relates to burn duration and expected accuracy); and the frequency, duration, and magnitude of typical burns.	
5.2.D	Provide USSPACECOM with an understanding of the flight control and navigation paradigm, principally whether a ground-based control approach is followed or some degree of (or full) autonomous control is used. If the satellite control does include some autonomous flight dynamics or control features, indicate how much (if any) foreknowledge ground controllers have of autonomous maneuver actions, the amount of information that is communicated to the ground both before and after the maneuver (e.g., maneuver time, delta-V, direction), and whether ground-based overrides are possible.	
5.3 CONOPS Development with Screening Provider Pertaining to CA Screening and Risk Assessment		
5.3.A	Decide whether the USSPACECOM free service or a validated commercial CA service will be used by the mission.	
5.3.B	Establish a service with the selected service provider.	
5.3.C	Implement a SSA sharing agreement with USSPACECOM to receive advanced data support and services.	
5.3.D	Through the registration of the satellite with USSPACECOM, begin the process of arranging for CA data exchange, including O/O ephemerides, maneuver notification reports, and CDMs. USSPACECOM uses the Space-Track.org account as the mechanism for product exchange.	
5.3.E	If needed, complete an ODR form to arrange for delivery of USSPACECOM advanced CA products.	



Section	Best Practice	Comment
5.3.F	For large constellations, coordinate with NASA and the screening provider to identify and address any special considerations.	
5.4 Launch-Related Conjunction Assessment		
5.4.A	Protect human space flight assets from having close approaches during the COLA gap using stand-off distance or statistical measures.	
5.4.B	Conform to any LCOLA requirements that the launch range may impose.	
5.5 <i>In situ</i> Launch Products and Processes		
5.5.A	Provide injection vector(s) to USSPACECOM as soon as they are available to aid in satellite tracking and identification.	
5.5.B	As soon as contact is established with each deployed spacecraft, generate and forward predicted ephemerides for the spacecraft to USSPACECOM, as well as publish these ephemerides (and all subsequent ephemeris updates) publicly, to assist in spacecraft identification for the cataloging process and provide general awareness among all O/Os.	
5.5.C	If USSPACECOM has issued TLEs for launched objects, notify USSPACECOM of the TLE and object number that is associated with your spacecraft.	
5.5.D	Provide early reporting to USSPACECOM of any spacecraft failures or other operational difficulties, both to obtain any available anomaly support and to assign the appropriate CA approach to the spacecraft (i.e., inactive and thus handled in a manner equivalent to a debris object).	
5.5.E	If using a commercial provider, make sure they have access to information from A-D.	
6.0 On-Orbit Collision Avoidance		
6.1 Spacecraft Information and Orbital Data Needed for CA		
6.1.A	Actively maintain the Space-Track.org record for your satellite, updating the active/dead and maneuverable/non-maneuverable flags to reflect the satellite's current status.	
6.1.B	Furnish predicted ephemerides that include state covariances to USSPACECOM (and any additional commercial screening provider) and set the privileges to allow any interested party to access and download this information.	
6.1.C	Furnished ephemerides should possess the following characteristics:	



Section	Best Practice	Comment
	<ul style="list-style-type: none"> a. Be of a 7-day predictive duration for LEO and 14 days for other orbits; b. Be issued at least at the following frequencies: <ul style="list-style-type: none"> i. Three times daily for spacecraft with perigee heights less than 500 km; ii. Daily for other LEO orbits; and iii. Twice weekly for other orbits. c. Include all known maneuvers within the ephemerides' prediction duration. d. Provide ephemeris point spacing of approximately 1/100th of an orbit, in either time or true anomaly (certain scientific missions with extremely long orbits or high eccentricities may require more specialized approaches). e. Contain a realistic covariance at each ephemeris point for at least the six estimated state parameters. (Appendix I provides a practical guide for assessing covariance realism, as well as some general principles for tuning the covariance production process.) f. Be formatted and distributed in the CCSDS standard OEM ephemeris format, preferably in the J2000 reference frame. 	
6.1.D	Furnish maneuver reports to USSPACECOM for any trajectory-altering satellite maneuvers, sufficiently in advance of maneuver execution to enable an O/O evaluation of the maneuver's safety. Employ the standard maneuver reporting message for this notification	
6.1.E	When a maneuver becomes part of a satellite's trajectory plan, generate and submit to the screening provider an updated ephemeris that contains this new maneuver as early as is feasible, but certainly with sufficient advance notice to enable an O/O evaluation of the maneuver's safety.	
6.2 CA Screenings		
6.2.A	Submit predicted ephemerides for your spacecraft to a screening provider to be screened for conjunctions at least daily, with spacecraft in higher-drag orbit regimes screened at least three (3) times per day.	
6.2.B	Ensure that an O/O ephemeris for an active, maneuverable spacecraft is screened against other ephemerides from active, maneuverable spacecraft in near real time after any such ephemeris is submitted to the screening provider.	



Section	Best Practice	Comment
6.2.C	Obtain and process these screening results from the screening provider at the same frequency at which they are produced, for both the full-catalog and O/O vs O/O screening cases described above.	
6.3 CA Risk Assessment		
6.3.A	Use the P_c as the principal collision likelihood assessment metric.	
6.3.B	Pursue mitigation if the P_c value exceeds $1E-04$ (1 in 10,000).	
6.3.C	Pursue mitigation if the estimated total miss distance is less than the hard-body-radius value.	
6.3.D	Employ the current operational NASA P_c calculation methodology for routine P_c calculation. Consider also removing correlated error from the primary and secondary object joint covariance.	
6.3.E	As a prioritization method for situations in which the number of conjunctions meeting mitigation criteria exceeds the ability of the O/O to mitigate, estimate the amount of debris that a conjunction would produce were it to result in a collision. A less stringent P_c an order of magnitude lower could be appropriate in such cases.	
6.3.F	If employing USSPACECOM data products for CA, use the procedure given in Appendix P to determine whether the data for a particular conjunction are actionable and thus constitute a basis for CA-related decisions.	
6.3.G	If a different CA product provider is chosen, develop and employ data actionability criteria for this provider's CA information to determine CA event actionability.	
6.4 CA Risk Mitigation		
6.4.A	When a conjunction's P_c at the mitigation action commitment point exceeds the mitigation threshold (recommended to be $1E-04$), pursue a mitigation action that will reduce the P_c by at least 1.5 orders of magnitude from the remediation threshold.	
6.4.B	Ensure that an ephemeris containing the mitigation action is screened against the full catalog, not a large screening volume collection of CDMs.	
6.4.C	Ensure that the mitigation action does not create any additional conjunctions with a P_c value above the mitigation threshold (for which the recommended value is $1E-04$).	



Section	Best Practice	Comment
6.4.D	When the secondary object is an active, maneuverable spacecraft, reach out to the secondary's O/O and jointly establish a way forward for the particular conjunction, including deciding which spacecraft will maneuver and freezing of the other spacecraft's planned trajectory until the TCA is passed.	
6.4.E	Use Space-Track.org contact information to engage other O/Os.	
6.5 Automated Trajectory Guidance and Maneuvering		
6.5.A	When an autonomously controlled satellite computes any maneuver to alter its orbit for mission or conjunction assessment purposes, communicate maneuver details to the operating ground control early enough to enable a USSPACECOM screening and appropriate action in response to the screening results.	
6.5.B	Share maneuver plans with USSPACECOM both as a maneuver notification report and a predicted ephemeris.	
6.5.C	Deploy an accurate on-ground satellite emulation capability to predict and share the behavior of the autonomously controlled satellite in predicted ephemerides.	
6.5.D	Execute the maneuver as planned unless an alteration is required for safety of flight once a maneuver is autonomously planned and communicated to the ground.	
6.5.E	Include the ability to pause or abort from the ground, for safety reasons, any maneuver planned by the on-board automated controller.	
6.6 Other Considerations		
6.6.A	As part of spacecraft physical design and orbit selection, perform a spacecraft photometric brightness analysis to determine whether the spacecraft is likely to present an impediment to ground-based astronomy. Consider changes to the satellite's construction, materials, or operating attitudes to reduce expected photometric brightness to levels that will not impose an impediment to ground-based astronomy.	
6.6.B	If a large constellation is being planned, construct a full photometric model of the spacecraft, including full BRDF information for the satellite surfaces, to obtain a durable estimate of the satellite's expected photometric brightness.	
6.6.C	If the constellation, given its population, orbit, and constituent satellites, is likely to affect ground-based astronomy, reassign the satellite orbits or modify the satellite construction to eliminate this effect.	



Appendix D. Best Practices for NASA Missions

This appendix contains a summary list of the best practices for NASA space flight missions (i.e., projects). It uses the NASA term “project” and “mission” interchangeably. NASA uses dedicated staff to perform CA risk assessment: NASA Johnson Space Center (JSC) Flight Operations Directorate (FOD) for human space flight assets and visiting vehicles and NASA Conjunction Assessment Risk Analysis (CARA) Program for all other spacecraft. The table below differentiates between actions recommended for the project and actions taken by FOD or CARA on the project’s behalf.

- The numbers of the best practices in the list below correspond to the section in the main document where they are fully described.
- A check mark in the “Project” column indicates that the project performs the best practice.
- A check mark in the “CARA/FOD Interfaces” column indicates that the project performs the best practice but will use CARA or FOD as their interface with USSPACECOM.
- A check mark in the “CARA/FOD” column indicates that CARA or FOD will perform the best practice on behalf of the project.
- Check marks in parenthesis indicate the best practice is optional.

Table D-1 Best Practices for NASA Projects

Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
3.0 History					
3.1 USSPACECOM CA Process					
3.1.A	Obtain a Space-Track.org account.	✓			Can optionally be done by project
3.1.B	Become familiar with the Spaceflight Safety Handbook contents.			✓	
3.1.C	Use USSPACECOM CA service as a minimum for screenings, even if additional commercial data or services are used.			✓	



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
3.1.1.A	Provide seven (7) days of predicted ephemeris (including maneuvers) to USSPACECOM for screening for LEO spacecraft, and 14 days for other orbits (e.g., HEO/GEO)		✓		
3.1.2.A	Provide at least one ephemeris per day to be screened; three ephemerides per day in the lower-drag LEO regime (perigee height less than 500 km).		✓		
3.1.2.B	Determine whether the O/O's process for obtaining screening results and performing CA risk assessment aligns with the timeline of the USSPACECOM process. If the timelines do not align in such a way as to enable timely and efficient screening support, pursue a re-arrangement of the O/O's process to minimize data latency and optimize screening efficiency.			✓	
3.1.3.A	Populate and maintain the POC section on Space-Track.org with your operations contact data. Be sure that the operations contact can be reached 24/7 due to time zone differences between operators and the immediate nature of certain CA emergencies.			✓	
3.1.3.B	Use standard ephemeris, CDM, and maneuver notification formats defined by CCSDS.	✓			
3.1.3.C	Large constellation operators should work with USSPACECOM pre-launch to determine if variations from the standard approach are necessary and, if so, to define a tailored screening process.			✓	
3.2 NASA Partnership with USSPACECOM					
3.2.A	Develop a robust safety-of-flight process that includes both CA screening and risk assessment to inform close approach mitigation decisions			✓	
3.2.B	Large constellation operators should work with USSPACECOM pre-launch to determine if variations from the standard approach are necessary and, if so, to define a tailored screening process.			✓	



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
3.2.C	Large constellation operators should consider working with NASA to define a risk assessment process. NASA is an experienced representative operator. Including NASA in discussions regarding establishing a CA process ensures that the process will work with most operators for risk assessment purposes.			✓	
4.0 Spacecraft and Constellation Design					
4.1 General Orbit Selection: Debris Object Density					
4.1.A	Include considerations of orbital debris density when selecting candidate mission orbits. Debris density diagrams, indexed by object size, are published by the NASA ODPO. (See Appendix G for more information and examples.)	✓			
4.1.B	When choosing among the final set of candidate orbits, perform a conjunction frequency analysis to determine the number of CA high interest events per spacecraft per year that can be expected in each of the proposed orbital areas.			✓	
4.2 Vehicle and Constellation Specific Orbit Selection: Spacecraft Colocation					
4.2.A	During orbit selection, perform an orbit colocation analysis to determine whether any of the proposed orbits are likely to create systematic conjunctions with existing actively maintained satellites.			✓	
4.2.B	If the orbit colocation analysis identifies systematic conjunctions, consider modifying the proposed orbit(s) slightly to eliminate this possibility. Optimizing orbit placement may be appropriate.	✓			
4.2.C	If the selected orbit is likely to present systematic conjunctions with a pre-existing spacecraft, then coordinate with the other operator(s) to arrange a process to coordinate maneuver plans routinely during the life of the mission.	✓			Include CARA in initial contact with other operator



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
4.3 Ascent to/Disposal from the Constellation's Operational Orbit					
4.3.A	Perform a study to compute the number of expected close approaches anticipated during ascent and descent, as well as the imputed additional satellite reliability that will be required to meet satellite disposal requirements at the chosen operational orbit.			✓	
4.3.B	If the results of the study show a large burden, consider choosing a different mission orbit with a lower burden.	✓			
4.3.C	All missions should review and follow the ODMSP guidance standards.	✓			NASA projects comply with NPR 8715.6.
4.3.D	When practicable, pursue active disposal using the fastest disposal option available.	✓			
4.3.E	Recognize that transiting spacecraft should yield way to on-station spacecraft and thus take responsibility for any CA risk mitigation maneuvers or transient trajectory alterations that may be required.	✓			
4.3.F	When planning descent using thrusters, if not planning an approach of circular-orbit altitude reduction, coordinate with NASA JSC TOPO to ensure that perigee-lowering approaches do not present persistent and problematic orbital crossings with the ISS and other human space flight assets.	✓			CARA to be included in these discussions
4.4 Spacecraft Trackability					
4.4.A	Through selection of physical design and materials, ensure that the satellite is trackable by SSN. a. For spacecraft with perigee heights less than 2000 km, the spacecraft should have characteristic dimensions of at least 10 cm in each major dimension.	✓			CARA can perform trackability analysis for NASA missions



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
	<p>b. For spacecraft with perigee heights greater than 2000 km, the spacecraft should have characteristic dimensions of at least 50 cm in each major dimension.</p> <p>c. If the spacecraft cannot meet these dimension constraints, use a proven detectability enhancement.</p>				
4.4.B	For spacecraft with perigee heights less than 5000 km and with dimensions less than 10 cm in each major dimension, or spacecraft with perigee heights greater than 5000 km with dimensions less than one meter, submit a schematic of the satellite, along with basic material properties, to USSPACECOM for a trackability analysis.		✓		
4.5 Spacecraft Reliability					
4.5.A	Ensure that spacecraft reliability is high enough that the likelihood of each spacecraft to remain fully functional until it can be disposed meets or exceeds 99%.	✓			
4.5.B	Reassess the 99% analysis whenever the underlying assumptions change, for example, extending operations beyond design life or failures in key systems.	✓			
4.6 Development of Capabilities for Ephemeris Generation and CA Risk Assessment and Mitigation					
4.6.A	Develop and implement a capability to generate and share accurate predicted spacecraft ephemerides, including any planned maneuvers. (See Section 3.1 for a discussion of ephemeris sharing for screening. See Section 6.1 for specific guidance on ephemeris generation frequency, length, point spacing, formatting, and contents.)	✓			



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
4.6.B	During spacecraft design, determine what risk mitigation approaches are possible for the spacecraft. While trajectory modification via thrusting maneuver is the most common, there are other possibilities such as differential drag orbit modification.	✓			
4.6.C	Develop and implement or arrange to acquire a capability to process CDMs and to compute conjunction risk assessment parameters, such as the Pc.			✓	
4.6.D	Develop and implement or arrange to acquire a risk analysis capability to select mitigation actions that will lower the Pc for dangerous conjunctions to a user-selected value.			✓	CARA to review project-selected mitigation actions
4.6.E	Validate CA tools well before flight, typically 6-12 months prior.	✓			
5.0 Pre-Launch Preparation and Early Launch Activities					
5.1 CONOPS Discussions and Arrangements with USSPACECOM Pertaining to the Launch Phase					
5.1.A	Establish contact with USSPACECOM to describe and discuss all the different aspects of the launch including delivery and deployment methodologies.		✓		
5.1.B	Deployment of multiple spacecraft should be spaced in a way that enhances the ability of USSPACECOM to quickly identify and differentiate the spacecraft using the SSN.		✓		
5.1.C	Discuss the preparation of the R-15 launch information message with USSPACECOM so that they understand its contents and any implications for safety of flight.		✓		
5.1.D	Ensure that the launch provider submits launch trajectory information to USSPACECOM, including ephemerides for the powered-flight portions and orbital elements for the parking orbits so that potential colocation and deconfliction potentialities can be discovered.	✓			



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
5.1.E	Provide launch-related information to USSPACECOM as soon as possible, such as injection vectors and initial ephemerides for deployed spacecraft, which can be used to assist with the cataloging process, especially to confirm the identity of launch-related objects. Coordinate a satellite numbering scheme (potentially including temporary satellite numbers) appropriate to the launch type and expected degree of cataloging difficulty.		✓		
5.1.F	Coordinate with USSPACECOM any potential launch anomaly diagnostic products that can be provided should issues arise during the launch and early orbit sequence		✓		
5.2 CONOPS Discussions and Arrangements with USSPACECOM Pertaining to Catalog Maintenance					
5.2.A	Register the spacecraft with USSPACECOM using the Satellite Registration form on Space-Track.org.			✓	
5.2.B	Provide USSPACECOM with basic construction and mission information about the satellite, such as stowed dimensions; deployable structures such as solar panels, antennae, booms, including all of their (rough) dimensions; satellite material properties and colors; regular satellite attitude; and registered operational radio frequencies.		✓		
5.2.C	Provide USSPACECOM with a basic overview of the satellite's orbit maintenance strategy, which would include the paradigm for determining when orbit maintenance maneuvers are required; the maneuver technology used (for this relates to burn duration and expected accuracy); and the frequency, duration, and magnitude of typical burns.		✓		



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
5.2.D	Provide USSPACECOM with an understanding of the flight control and navigation paradigm, principally whether a ground-based control approach is followed or some degree of (or full) autonomous control is used. If the satellite control does include some autonomous flight dynamics or control features, indicate how much (if any) foreknowledge ground controllers have of autonomous maneuver actions, the amount of information that is communicated to the ground both before and after the maneuver (e.g., maneuver time, delta-V, direction), and whether ground-based overrides are possible.		✓		
5.3 CONOPS Development with Screening Provider Pertaining to CA Screening and Risk Assessment					
5.3.A	Decide whether the USSPACECOM free service or a validated commercial CA service will be used by the mission.			✓	
5.3.B	Establish a service with the selected service provider.			✓	CARA coordination required for use of commercial data
5.3.C	Implement a SSA sharing agreement with USSPACECOM to receive advanced data support and services.			✓	
5.3.D	Through the registration of the satellite with USSPACECOM, begin the process of arranging for CA data exchange, including O/O ephemerides, maneuver notification reports, and CDMs. USSPACECOM uses the Space-Track.org account as the mechanism for product exchange.			✓	
5.3.E	If needed, complete an ODR form to arrange for delivery of USSPACECOM advanced CA products.		✓		
5.3.F	For large constellations, coordinate with NASA and the screening provider to identify and address any special considerations.			✓	



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
5.4 Launch-Related Conjunction Assessment					
5.4.A	Protect human space flight assets from having close approaches during the COLA gap using stand-off distance or statistical measures.			✓	
5.4.B	Conform to any LCOLA requirements that the launch range may impose.	✓			
5.5 <i>In situ</i> Launch Products and Processes					
5.5.A	Provide injection vector(s) to USSPACECOM as soon as they are available to aid in satellite tracking and identification.		✓		
5.5.B	As soon as contact is established with each deployed spacecraft, generate and forward predicted ephemerides for these spacecraft to USSPACECOM, as well as publish these ephemerides (and all subsequent ephemeris updates) publicly, to assist in spacecraft identification for the cataloging process and provide general awareness among all O/Os.		✓		
5.5.C	If USSPACECOM has issued TLEs for launched objects, notify USSPACECOM of the TLE and object number that associates to your spacecraft.		✓		
5.5.D	Provide early reporting to USSPACECOM of any spacecraft failures or other operational difficulties, both to obtain any available anomaly support and to assign the appropriate CA approach to the spacecraft (i.e., inactive and thus handled in a manner equivalent to a debris object).		✓		
5.5.E	If using a commercial provider make sure they have access to information from A-D.			✓	



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA/FOD	Comment
6.0 On-Orbit Collision Avoidance					
6.1 Spacecraft Information and Orbital Data Needed for CA					
6.1.A	Actively maintain the Space-Track.org record for your satellite, updating the active/dead and maneuverable/non-maneuverable flags to reflect the satellite's current status.			✓	
6.1.B	Furnish predicted ephemerides that include state covariances to USSPACECOM (and additionally the screening provider, if this is different from USSPACECOM) and set the privileges to allow any interested party to access and download this information.		✓		
6.1.C	<p>Furnished ephemerides should possess the following characteristics:</p> <ul style="list-style-type: none"> a. Be of a 7-day predictive duration for LEO and 14 days for other orbits; b. Be issued at least at the following frequencies: <ul style="list-style-type: none"> i. Three times daily for spacecraft with perigee heights less than 500 km; ii. Daily for other LEO orbits; and iii. Twice weekly for other orbits. c. Include all known maneuvers within the ephemerides' prediction duration. d. Provide ephemeris point spacing of approximately 1/100th of an orbit, in either time or true anomaly (certain scientific missions with extremely long orbits or high eccentricities may require more specialized approaches). 		✓		



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
	<p>e. Contain a realistic covariance at each ephemeris point for at least the six estimated state parameters. (See Appendix I for a practical guide for assessing covariance realism as well as some general principles for tuning the covariance production process.)</p> <p>f. Be formatted and distributed in the CCSDS standard OEM ephemeris format, preferably in the J2000 reference frame.</p>				
6.1.D	Furnish maneuver reports to USSPACECOM for any trajectory-altering satellite maneuvers, sufficiently in advance of maneuver execution to enable an O/O evaluation of the maneuver's safety. Employ the standard maneuver reporting message for this notification		✓		
6.1.E	When a maneuver becomes part of a satellite's trajectory plan, generate and submit to the screening provider an updated ephemeris that contains this new maneuver as early as is feasible but certainly with sufficient advance notice to enable an O/O evaluation of the maneuver's safety.		✓		
6.2 CA Screenings					
6.2.A	Submit predicted ephemerides for your spacecraft to a screening provider to be screened for conjunctions at least daily with spacecraft in higher-drag orbit regimes screened at least three times per day.		✓		
6.2.B	Ensure that an O/O ephemeris for an active, maneuverable spacecraft is screened against other ephemerides from active, maneuverable spacecraft in near real time after any such ephemeris is submitted to the screening provider.			✓	



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
6.2.C	Obtain and process these screening results from the screening provider at the same frequency at which they are produced, for both the full-catalog and O/O vs O/O screening cases described above.			✓	
6.3 CA Risk Assessment					
6.3.A	Use the Pc as the principal collision likelihood assessment metric.			✓	
6.3.B	Pursue mitigation if the Pc value exceeds 1E-04 (1 in 10,000).	✓			
6.3.C	Pursue mitigation if the estimated total miss distance is less than the hard-body-radius value.	✓			
6.3.D	Employ the current operational NASA Pc calculation methodology for routine Pc calculation. Consider also removing correlated error from the primary and secondary object joint covariance.			✓	
6.3.E	As a prioritization method for situations in which the number of conjunctions meeting mitigation criteria exceeds the ability of the O/O to mitigate, estimate the amount of debris that a conjunction would produce were it to result in a collision. A less stringent Pc an order of magnitude lower could be appropriate in such cases.			✓	
6.3.F	If employing USSPACECOM data products for CA, use the procedure given in Appendix P to determine whether the data for a particular conjunction are actionable and thus constitute a basis for CA-related decisions.			✓	
6.3.G	If a different CA product provider is chosen, develop and employ data actionability criteria for this provider's CA information to determine CA event actionability.			✓	



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
6.4 CA Risk Mitigation					
6.4.A	When a conjunction's Pc at the mitigation action commitment point exceeds the mitigation threshold (recommended to be 1E-04), pursue a mitigation action that will reduce the Pc by at least 1.5 orders of magnitude from the remediation threshold.	✓			
6.4.B	Ensure that an ephemeris containing the mitigation action is screened against the full catalog, not a large screening volume collection of CDMs.			✓	
6.4.C	Ensure that the mitigation action does not create any additional conjunctions with a Pc value above the mitigation threshold (for which the recommended value is 1E-04).		✓		
6.4.D	When the secondary object is an active, maneuverable spacecraft, reach out to the secondary's O/O and jointly establish a way forward for the particular conjunction, including deciding which spacecraft will maneuver and freezing of the other spacecraft's planned trajectory until the TCA is passed.			✓	
6.4.E	Use Space-Track.org contact information to engage other O/Os.			✓	
6.5 Automated Trajectory Guidance and Maneuvering					
6.5.A	When an autonomously controlled satellite computes any maneuver to alter its orbit for mission or conjunction assessment purposes, communicate maneuver details to the operating ground control early enough to enable a USSPACECOM screening and appropriate action in response to the screening results.	✓			
6.5.B	Share maneuver plans with USSPACECOM both as a maneuver notification report and a predicted ephemeris.		✓		
6.5.C	Deploy an accurate on-ground satellite emulation capability to predict and share the behavior of the autonomously controlled satellite in predicted ephemerides.	✓			



Section	Best Practice	Project	Project with CARA/FOD providing interface to USSPACECOM	CARA /FOD	Comment
6.5.D	Execute the maneuver as planned unless an alteration is required for safety of flight once a maneuver is autonomously planned and communicated to the ground.	✓			
6.5.E	For safety reasons, include the ability to pause or abort from the ground any maneuver planned by the on-board automated controller	✓			
6.6 Other Considerations					
6.6.A	As part of spacecraft physical design and orbit selection, perform a spacecraft photometric brightness analysis to determine whether the spacecraft is likely to present an impediment to ground-based astronomy. Consider changes to the satellite’s construction, materials, or operating attitudes to reduce expected photometric brightness to levels that will not impose an impediment to ground-based astronomy.			✓	
6.6.B	If a large constellation is being planned, construct a full photometric model of the spacecraft, including full BRDF information for the satellite surfaces, to obtain a durable estimate of the satellite’s expected photometric brightness.	✓			
6.6.C	If the constellation, given its population, orbit, and constituent satellites, is likely to affect ground-based astronomy, reassign the satellite orbits or modify the satellite construction to eliminate this effect.	✓			



Appendix E. Use of Analytic Theory Orbital Data in Conjunction Assessment

The following is an amplification of a statement made by the Conjunction Assessment Technical Advisory Committee in 2016 on the use of analytic theory, and in particular the Two-Line Element (TLE) set publicly accessible catalog produced by USSPACECOM.

Due to enormous increases in computational capacity over the last twenty years, the space surveillance industry has been able to transition from analytic to higher-order theory approaches to space catalog maintenance. This transition brings a number of significant advantages to modern space catalogs such as the modeling of additional perturbations that were not typically represented in analytic theories (e.g., solar radiation pressure and solid earth tides), an improved fidelity in modeling the major perturbations (e.g., geopotential and atmospheric drag), and a durable covariance matrix to accompany each state estimate. Some of these features existed with previous analytic and semi-analytic theories, but it is only in recent times that all the additional attributes are routinely produced and available for space surveillance applications. This development has prompted the question of how data from analytic models (such as the Simplified General Perturbations Theory #4 (SGP4), which is used to produce the publicly available TLE catalogs) and higher-order theory models (such as the USSPACECOM Special Perturbations (SP) Space Catalog, although higher-order theory data can also be furnished from other sources) should be employed in the CA risk assessment process.

Risk assessment, whether it is quotidian consumer risk, the extremely serious discipline of nuclear accident risk, or the present problem of satellite collision risk, must always contain within its calculus some determination of the likelihood of occurrence of the event in question. For CA, this determination is most commonly assessed by calculating the Probability of Collision (P_c) (see Appendix I for a development of the calculation theory), and this calculation requires both the state estimates for the primary and secondary objects and the accompanying measures of these estimates' uncertainties, typically with a covariance matrix.

Most implementations of analytic orbital theories lack the ability to produce a meaningful covariance (this is true of the SGP4 TLEs), so they cannot enable a probabilistic CA calculation and thus for CA applications will yield only a miss distance at the time of the two objects' closest approach (TCA). Because no uncertainty information on this miss distance is provided, no probabilistic conclusion can be drawn. One can (arbitrarily) select a miss-distance threshold to enable a binary "safe / not safe" conjunction evaluation, but there is no ability to determine, for example, how often the real miss distance, for which the nominal miss distance is only an estimated mean value, is actually likely to exceed or fall below such a threshold. No probabilistic assessment of whether a conjunction will be a dangerously close encounter is possible, so the likelihood of occurrence of a collision cannot be determined.

This dynamic is illustrated in Figure E-1, which gives a Cumulative Distribution Function (CDF) plot of miss distance distributions, both for critical ($P_c > 1E-04$) and non-critical ($P_c < 1E-04$) events. One observes immediately the lack of any strong correlation between P_c and miss distance: the span of miss distances for the critical events (less than 100 m to more than 10 km)



is extremely broad, and it overlaps significantly with the results for the non-critical events. If, for example, one wished to eliminate 95% of the high- P_c events using a miss-distance criterion alone, one would need to choose a 10 km miss-distance threshold—a value that would also include more than 50% of the non-worrisome events, creating a density of events to mitigate so great that an unsustainable number of such mitigation actions would be required.




Figure E-1 Miss-Distance Distributions for Critical ($P_c > 1E-04$) and Non-critical ($P_c < 1E-04$) Events

Additionally, the inherent theory error in general perturbation approaches makes them ill-suited for even a risk-tolerant, miss-distance-based CA. A published study (Hejduk et al. 2013) provided TLE accuracy information for the two methods presently used by USSPACECOM to generate them: traditional SGP4 orbit determination and Extrapolated General Perturbation (eGP) TLE generation, which performs the SGP4 orbit determination not against sensor measurements but against pseudo-observations generated from a higher-order theory trajectory for the object, which includes some forward prediction of the trajectory. Both methods are in use presently to generate the publicly released TLE catalog, and it is not always easy to determine which method was used to produce any given TLE. Performance data for the orbital regime corresponding to Figure E-1 above, namely LEO with perigee heights greater than 500 km, are given in Figure E-2. These results show that, while the eGP methodology outperforms traditional SGP4 as expected, both are noisy in the epoch to 72-hour propagation states, which is the time-frame during which most CA decisions are made; errors from 600 m to 3 km are observed.




Figure E-2 eGP and Regular SGP4 Errors: LEO > 500 km (log scale)⁸

Figure E-3 below is an expansion of Figure E-1 to include CDF plots of miss-distance distributions for events with both $P_c > 1E-03$ and, separately, $1E-02$. Here, one can see that it is possible that a miss-distance criterion of several hundred meters, or perhaps 1-2 km, could be used as a simplified method of performing a risk-tolerant CA: remediate any conjunctions with a miss distance less than 1-2 km. In addition to this being an ill-advised move because it would result in a huge number of unnecessary mitigations (as an example, for events with a miss distance less than 1 km, the number with a P_c less than $1E-03$ and thus not serious is six times greater than the number with a P_c greater than $1E-03$), the inherent theory error (600 m to 3 km, from above) is a substantial portion of the miss-distance threshold of 1-2 km that would be imposed—in some cases even greater than the threshold itself. So even under the very limited case described here, the use of general perturbation results for CA risk assessment is not tenable. To this end, analytic theory data, such as that contained in the public TLE catalogs, should not, taken alone, serve as a basis for risk assessment decisions and subsequent mitigation actions.

⁸ From Hejduk et al. 2015.

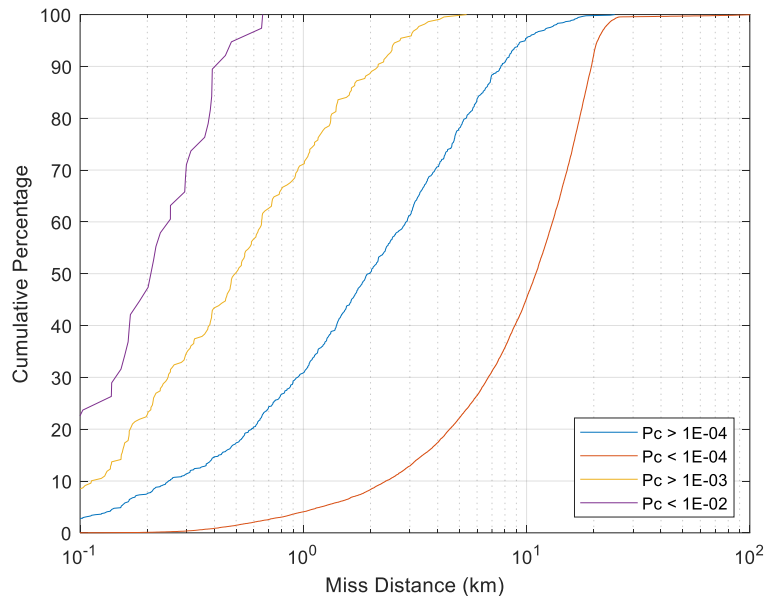


Figure E-3 Expansion of Figure E-1 to Include CDFs for P_c Thresholds of 1E-03 and 1E-02

Nonetheless, analytic theory orbital information is not without its uses in CA. It is frequently employed helpfully in the following ways:

- Mission planning activities such as the selection of satellite future orbits and final disposal orbits;
- General situational awareness information such as the positions of neighboring objects and the general effects of the basic perturbations (e.g., atmospheric drag);
- Pre-filtering of candidates for conjunction assessment screenings for which simple TLE-based filters can eliminate substantial numbers of candidate pairs from requiring explicit screening runs;
- Assessment of the status of a payload that appears as a conjuncting secondary as active secondaries require special analysis and alerting when performing conjunction remediation; and
- Consistency checks of the higher-order theory data. While the higher-order theory data are typically more accurate, they themselves are vulnerable to updates with bad or inadequate tracking data, errors in force model and integration control settings, and incorrect maneuver solutions. Comparisons with the analytic data can often identify such problems and, while not enabling their repair, can at least make clear that such cases constitute non-actionable CA data.

The above uses of analytic theory data, usually furnished as TLEs, can be helpful to the CA enterprise, but these uses do not themselves directly support conjunction risk assessment for which more accurate state estimate data accompanied by uncertainty information is required.



Appendix F. Expected CA Event Rates

The CA workload that a mission experiences is a function of the total number of CDMs that it receives, the number of associated events with risk levels high enough to require enhanced monitoring, and the number of events that have a risk level so high as to require mitigation planning and potential execution. During mission planning and development, it is helpful to have a general idea of the frequency of such situations in order to model spacecraft fuel consumption, identify staffing needs, and determine the appropriate degree of CA-related tool automation desired. The following profiling of CA event densities is intended to provide at least some first-order information to answer these questions. These data summarize three years' recent CA data history for the payloads that NASA protects in the LEO orbit regime (event densities in the non-LEO region are both low and highly dependent on the specific orbit, so summary information is not given here).

The graphs presented in Figure F-1 contain several component parts. The x-axis gives the hard-body radius for the conjunction, which is essentially the combined sizes of the primary and secondary objects. A number of methods are used to determine this size; perhaps the most common is to construct a circumscribing sphere about the primary object and add its radius to a standard radius size for the type of secondary object (i.e., payload, rocket body, or debris) in the conjunction (Mashiku and Hejduk 2019). HBR values of 20 m, 10 m, 5 m, 2 m, and 1 m are used, the last few values included to give some statement of the expected situation for small satellites (all the CARA events were reprocessed with these different hard-body-radius values to produce these varied results). The y-axis gives the number of expected events per payload per year, shown on a logarithmic scale, with the type of event indicated by the color in the stacked bar charts, which are to be read in the following way:

- The height of the red bar indicates the frequency of red events, which are defined as events containing at least one CDM with a P_c value $> 1E-04$. These are the most serious events; they require detailed monitoring, mitigation planning, and in a minority of cases, the actual execution of a mitigation action.
- The height of the yellow bar indicates the frequency of yellow events, which are defined as events containing at least one CDM with a P_c value $> 1E-07$. These events generally merit increased monitoring and attention, including in some situations, mitigation planning (if they are close to the red threshold). Based on the definition above, all red events are also yellow events.
- The height of the green bar indicates the frequency of events of any type, including those that never manifest a $P_c > 1E-07$.

The overall height of each bar thus indicates the CDM generation rate. It should be pointed out that this number, while in some ways a useful loading figure, is not all that meaningful in itself because it is a function of the size of the screening volume that happens to be employed. If one chooses a screening volume similar in size to those employed by NASA CARA, then the green event generation rates (which include the yellow and red events) are likely to be similar. Because they require relatively small miss distances, the yellow and red generation rates are expected to



be largely independent of the screening volume size selected as long as a reasonably sized volume is used.

Three different graphs are presented representing three different regions of LEO sorted by perigee height (p): a high-drag regime ($p < 500$ km), a low-drag regime (perigee height between 500 and 750 km) that includes a large concentration of satellites somewhat higher than 700 km, and a second low-drag regime from 750 to 1400 km. There are moderate differences in event frequency among the three different regimes, although some of that variation is due to the different levels of sampling among the three, given that NASA operates substantially different numbers of payloads in each of these three regions. So, one should take these results to be more notional than precise; nonetheless, they do give some level of orientation to the expected loading for each event type (green, yellow, and red).




Figure F-1 CA Event Frequencies as a Function of Orbit Regime and Event Severity



Appendix G. Orbital Debris Density

Orbital debris can be generated by a number of processes: detritus cast off as part of the launch and injection process; off-scourings from healthy satellites themselves, such as peeling of Mylar thermal insulation; collisions between satellites; explosions of satellites; and satellite debris-generating events in which (without any clear cause) a satellite is observed to break into more than one piece. All these methods have conspired to produce the present debris environment. Because of the orbital locations in which major debris-producing events have taken place, and due to the differences in efficiency of natural debris cleansing arising from varying drag at these different altitudes, orbital debris density is not uniform. It is thus helpful to examine the current debris density situation when choosing an operational orbit as the orbital debris density does influence the rate of close approach events as well as the likelihood of satellite damage due to an object too small to be tracked.

The graphs in Figure G-1 are produced from the NASA Orbital Debris Engineering Model (ORDEM) 3.1 release developed by the NASA Orbital Debris Program Office (ODPO). Generated for the 2020 debris environment, these graphs give debris flux (per satellite unit area and time on orbit) as a function of orbital altitude (a circular orbit is presumed) and inclination. One perhaps could argue that, given that satellites often have many years of orbital lifetime, the graphs should be projected forward in time and thus reflect debris generation events and natural debris depletion processes over as much as the next decade. However, such a prediction requires a stochastic execution of the model in order to try to guess at the number, severity, and location of future debris-producing events, and as the purpose here is to compare the relative debris densities of different regimes, the simplicity of a single rather than probabilistic presentation is preferable.

The graphs are arranged by inclination, and while it is observed that the highest inclination band represented here does have a slightly higher debris density than the two lower bands, the differences in both absolute debris flux values and the overall curve morphology among the three inclination bands are not great, indicating that inclination is not a major contributor to debris density. Variation with orbital altitude, however, is more marked: the density peaks at about 850 km and tails off in either direction from that peak at relatively equal rates. And the reduction is significant: about an order of magnitude in each direction within the bounds of the graphs (400 to 1200 km orbital altitude).

Three lines are shown on each graph. The amber line represents objects of sizes greater than 10 cm, which is the size of objects that the present space catalog comprises without any contributions from the new Space Fence SSN radar and is thus related to satellite conjunctions presently tracked and reported. The reddish line is for objects of sizes greater than 5 cm, which is the size of objects that can be expected to be included in a catalog that does include tracking data from the new Space Fence radar, which should be contributing data to the CA enterprise in the near future. The blue line is for objects of sizes greater than 1 cm, which is the size of objects large enough to be generally considered lethal to a satellite if there is a collision.

The gap between the reddish and blue lines thus represents the flux level of objects that are too small to be tracked yet large enough to potentially be lethally damaging, which represents a



collision risk that is simply accepted as the cost of operating a satellite in space as essentially, it cannot be mitigated.

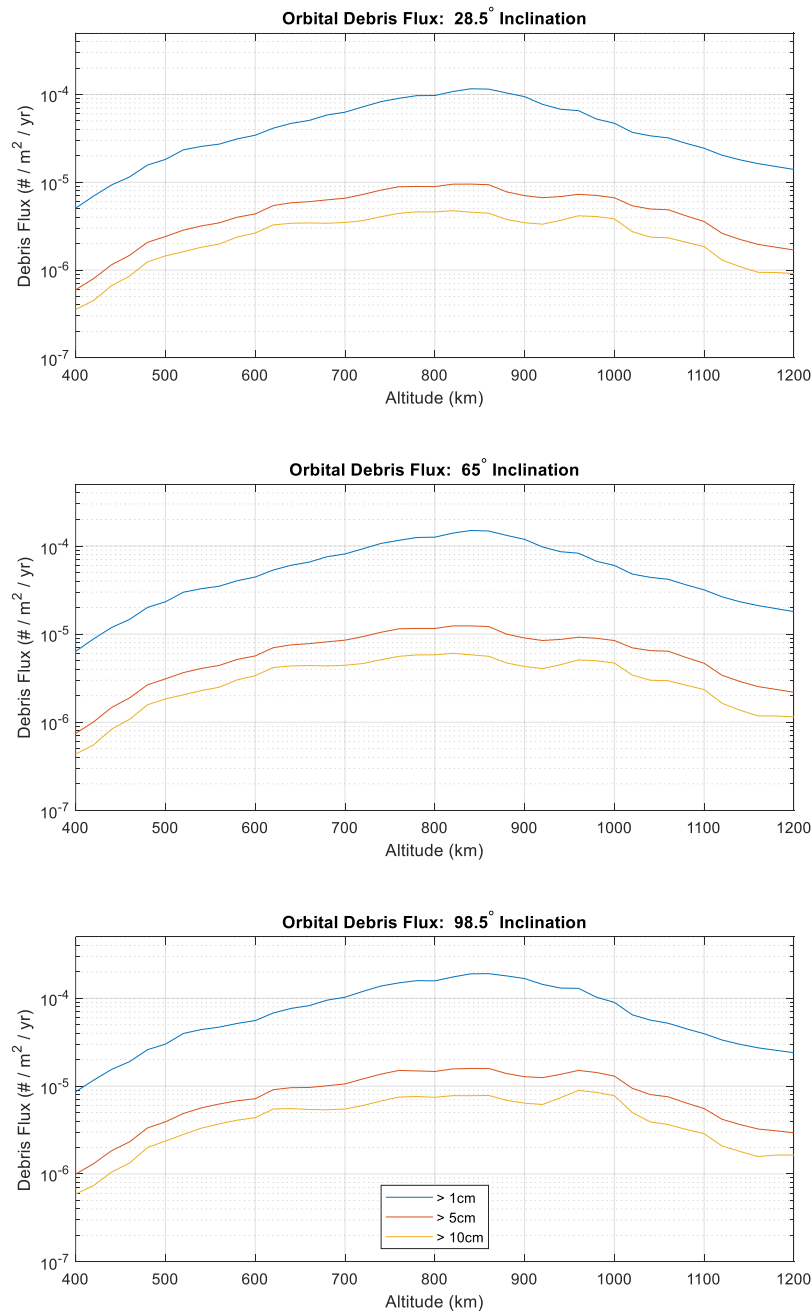


Figure G-1 Debris Flux as a Function of Orbit Inclination, Orbit Altitude, and Object Size

In some cases, the flux densities between the 1 cm and 5 cm bands represent almost an order of magnitude difference. Considering this difference, some might be led to question why CA as an



activity is being pursued at all, given that so much of the collision risk cannot be remediated and simply must be accepted. To this question are two distinct responses:

1. First, within reasonable bounds, it makes sense to protect against safety-of-flight hazards that are known and that can, if necessary, be mitigated. NASA CARA's statement of purpose is "to take prudent measures, at reasonable cost, to enhance safety of flight, without placing an undue burden on mission operations." One must guard against an unjustified over-investment in safety, and mission objectives should not be regularly and substantially impaired by CA considerations. But since it is not particularly expensive or difficult to identify high-risk conjunctions from the existing space catalog, and since such conjunctions are relatively infrequent, and furthermore, since the need for mitigation action for such conjunctions is less than an order of magnitude less frequent than that, it is good safety practice to conduct CA operations according to the CARA mission statement. Space actors thus employ the CA enterprise to "prevent the preventable."
2. Second, the satellite conjunctions that occur between objects large enough to be in the space catalog are the ones that have the potential to generate large amounts of space debris. As the NASA ODPO has shown (Johnson et al. 2001), the debris production potential of a collision is a function of the two objects' relative kinetic energy and masses: the ratio of their two masses and their relative velocity. Collisions in which the two objects' masses are quite different will generally result in only the lighter of the two objects being fragmented with the heavier object being merely cratered. Objects in which the two masses are closer to each other will generally result in complete fragmentation of both of the objects and thus very large amounts of resultant debris. Extremely small secondary objects (i.e., a few centimeters in size and thus below the level of trackability), even at the large relative velocities of LEO conjunctions (~10,000 m/s), are quite unlikely to cause fragmentation of larger active payloads, so the amounts of debris produced by collisions with untracked objects will in almost all cases be relatively small. As stated previously, collisions with untracked objects can still be lethal to a satellite and do constitute an accepted risk simply of operating a satellite in space, but CA against tracked and cataloged objects prevents the collisions that will produce large amounts of debris and thus makes a substantial contribution to space sustainability, even though it cannot protect payloads from debilitating conjunctions with very small objects.

Choosing a mission trajectory that occupies a region of space with a lower debris density is thus a favorable selection for two reasons: first, it lowers the risk of a lethal collision with an untracked object; and second, it in general reduces the rate at which serious conjunctions with cataloged objects will occur. As many O/Os have determined, choosing orbits with lower orbital altitudes both achieves the above two goals (as the curves in Figure G-1 demonstrate) and facilitates complying with the 25-year disposal guidelines

As a final remark, it is important to recognize that, while there may be a general correlation, there is not in fact a direct relationship between the debris flux for catalog-size objects and the actual rate at which conjunctions, even serious conjunctions, will be identified. The debris flux gives the expected frequency of penetration of the satellite's cross-sectional area by some object in space. According to the curves in Figure G-1, a sun-synchronous object at 700 km, placed at



the middle of a 20 m radius sphere (to employ an extremely conservative statement of its exposed area), would over a one-year period expect only 0.006 penetrations of this 20 m sphere. In other words, it should take ~160 years on orbit for a penetration to occur. However, predicted close approaches for such an object are identified daily, and a few times per year, mitigation actions would be warranted. What is the meaning of this lack of alignment?

When a P_c for a specific event is calculated, it is interpreted as representing the likelihood that the “true” miss distance (which ultimately cannot be actually known) is smaller than the hard-body radius defined about the object (e.g., the 20 m radius of the constructed sphere described in the above paragraph). The “true” P_c corresponding to the “true” miss distance is either 0 or 1; i.e., either the satellites will pass within 20 m of each other or they will not. However, the orbit determination process does not predict future positions without error, so a P_c value somewhere between 0 and 1 represents a calculation of the likelihood of a penetration given the uncertainties in the input measurement data to the orbit determination, lack of comprehensiveness of the dynamical model, and additional uncertainties expected in predicting the satellites’ epoch states to their TCA (such as atmospheric density forecast error). Generally, when the P_c is greater than $1E-04$ (1 in 10,000 likelihood of a penetration), a mitigation action is recommended and executed; this threshold has emerged in the industry as an acceptable balance between safety-of-flight considerations and additional mission burden and encumbrance. Because mitigation actions are recommended when the likelihood of penetration exceeds 1 in 10,000, there is little surprise that the frequency of CA serious events exceeds that projected by the debris flux value: in only 1 out of 10,000 cases would the penetration actually take place if the mitigation action were not executed. Fortunately, a conservative approach such as this can be implemented and is still consistent with the mandates of prudent action, reasonable costs, and lack of undue burden on mission operations.

(See Appendix F for information on actual expected rates of serious CA events.)



Appendix H. Satellite Colocation Analysis

During their on-orbit residence, most spacecraft with some frequency come into conjunction with other objects, which in CA terminology are called “secondary” objects. In the great majority of these cases, these objects are not active satellites, and thus even if a repeating conjunction situation arises in which a chain of conjunctions between any two objects is observed, because the secondary object is not maintained in a particular orbit, the situation naturally dissipates. However, if two objects are maintained in similar orbits or in orbits that have consistent intersections, synodic alignment between the objects can occur, producing an entire orbital lifetime of frequent conjunction. Because conjunction management between active payloads is complex, requiring the sharing and coordination of maneuver plans and joint decisions about near-term courses of action, it is desirable to minimize such situations. To this end, when an operational orbit is being selected during mission design, a satellite colocation analysis should be performed to determine whether there are any existing active spacecraft or planned missions that will occupy the same or similar orbits and thus potentially produce synodic alignment and therefore frequent recurring conjunctions. It is often possible to make modifications to an intended mission orbit to avoid this unfavorable conjunction situation yet still meet mission objectives.

This section describes the approach used by NASA CARA to perform colocation analyses for proposed orbits. The NASA CARA instantiation of this capability is presently being refactored to become more portable and better documented. When this process is finished, NASA intends to add it to the suite of CARA tools that can be freely downloaded and used by the public. (See Section 7, Contact Information in this document for the URL to the NASA public-facing software repository.)

H.1 Algorithmic Description

For two objects in orbit to be in conjunction, two conditions must be met: as seen in Figure H-1, the magnitude of each object’s position vectors must be equal ($R_1 = R_2$) and the angle between the position vectors, α , must go to zero.




Figure H-1 Cartesian, Inertial, Earth-centered Position Vectors



The difference in radius, $\Delta R = R_1 - R_2$, and α , are then the key metrics in determining satellite colocation. Both metrics are functions of the well-defined Keplerian orbital elements and their rates. The Keplerian orbital elements and rates for active spacecraft are derived and cataloged from TLE histories. The analyst then provides this same orbital information for a proposed mission, and the algorithm determines the frequency at which both conditions are met with the cataloged objects.

The radius metric for a single object is easily determined by analyzing the osculating radius magnitude. The osculating radius vector is chosen over other orbit size metrics, such as osculating or mean semi-major axis or apoapsis and periapsis axes, because short periodic perturbations can cause deviation of the radius vector from these metrics by tens of kilometers, even for low eccentricities. Likewise, the short periodic and secular perturbations will not be completely captured by TLEs with a sparse (e.g., daily) reporting frequency. Therefore, each TLE is propagated for at least an orbit to capture the true maximum and minimum radius magnitudes for an object. The radius metric will be evaluated as a simple comparison of the maximum and minimum observed radius for two objects. If the orbit radii overlap, the orbits will have met the first criterion for potential colocation.

The angle α is a function of the Keplerian orbital element given by the equation

$$\begin{aligned} \cos\alpha = & K_1 \cos(\Delta\nu - \Delta\Omega + \Delta\omega) + K_2 \cos(\Delta\nu + \Delta\Omega + \Delta\omega) \\ & + K_3 \text{other as through} \end{aligned} \quad (\text{H-1})$$

A full derivation of Equation H-1 is documented in a separate technical memorandum, which can be obtained from NASA CARA. The frequency of α is a function of the differences and sums of the Right Ascension of the Ascending Nodes (RAAN, Ω) and argument of latitude ($L = \nu + \omega$, with ν as true anomaly). The amplitudes are functions of the constants K_i in Equation H-1, which are determined by the objects' inclinations. Using the fact that inclination is, in general, slowly varying, and that RAAN and L are primarily linear in time for nearly circular orbits, α can be characterized by recording the maximum and minimum inclination and the initial value and rates of RAAN and L .

When secular (linear) time dependence in RAAN and L are included, Equation H-1 holds for small eccentricities. As eccentricity increases, the assumption of linear L rate begins to break down. However, it is observed that Equation H-1 still captures the primary frequency in minimum of alpha as shown in Figure H-2. Therefore, Equation H-1 still serves as a sufficient estimate of the frequency with which eccentric orbits have the potential for $\alpha \rightarrow 0$.

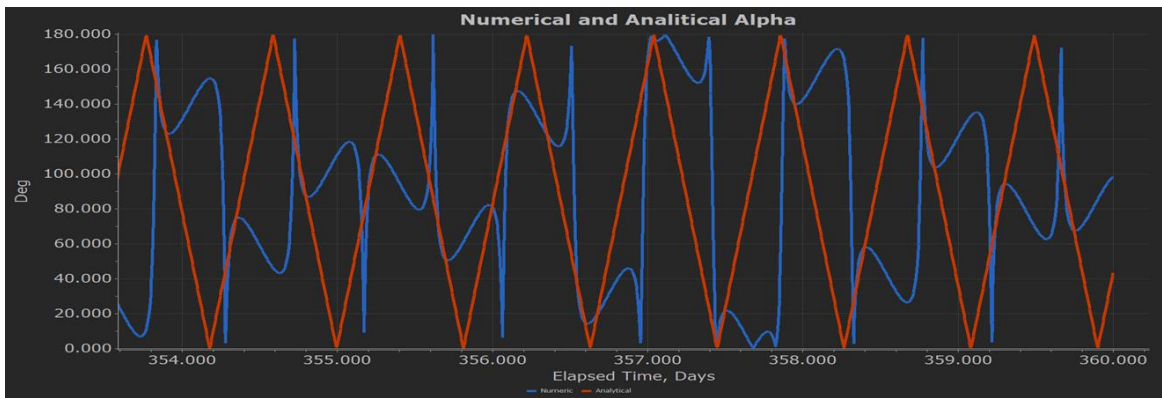


Figure H-2 Numerical vs. Analytical (Linear) α for Highly Eccentric Orbit

A second concern, however, exists for eccentric orbits. Consider the co-planar orbits in Figure H-3 below. At both times A and B, α is zero. However, due to the eccentricity of the inner orbit, they will not have the potential to conjunct at B. Since these orbits have the same radius at A, they would have passed the radius condition check as it is a simple time-independent overlap of radius maximum and minimum. A third method is needed for eccentric orbits to add time dependence to the radius metric.




Figure H-3 Co-planar Orbits with One Circular and One Elliptical

To do this, it is observed that the following constraint must hold

$$N_{\alpha}P_{\alpha} = N_oP_o \quad (H-2)$$

where N_{α} is the integer number of periods of α , P_{α} is the dominant period of α determined by Equation H-1, N_o is the integer number of eccentric object orbital periods, and P_o is the eccentric object's orbital period. $N_{\alpha}P_{\alpha} = N_oP_o$ will then be the period of expected conjunction with the eccentric objects. Note that solving Equation H-2 while enforcing integer N_{α} and N_o is not trivial. The implemented algorithm provides candidate resonate values for N_{α} and N_o for analyst investigation.



H.2 Tool Usage Methodology

There are two separate execution aspects to the colocation tool that implements the algorithm described in the previous section. First, the TLE histories of actively maintained objects need to be examined to understand the canonical values and rates of the orbital parameters and additional factors that the algorithm employs. One must first collect TLE histories for such objects and execute the tool in this mode for these histories to establish the database of parameters needed for the follow-on comparison analysis. Second, the parameters of interest of the orbit to be examined for colocation must be provided. The screen shot in Figure H-4 is a display confirmation of these inputs.

```
Test SC Co-Location.Rpt
Jul 07 2020 13:49:27

Primary Object Parameters

Object ID:      Test SC
Anchor Epoch:  Dec 30 2017 21:47:46.056
Radius, km:    7071 to 7090
Inclination, deg: 98.2 to 98.1
RAAN:         303.100 deg, 1.132477403e-05 deg/s
Arg. Of Latitude: 75.100 deg, 6.064034343e-02 deg/s

Analysis Parmaters

Miss Distance Limit, km: 25
Analysis Duration, days: 365
Time Step, sec:         10
```

Figure H-4 Primary Object Tool Inputs

Outputs are given in terms of maximum, minimum, and average frequencies in which a synodic situation is observed. Figure H-5 below gives sample output for a situation that does constitute a recurring conjunction:



```

SATCAT ID: 22491

Days to first violation:           0.842
Minimum Days Between Violations:  0.034
Maximum Days Between Violations:  57.481
Average Days Between Violations:  25.315

Alpha Constants and Frequencies
Ki          wi
0.027      1.568196447e-06
0.409      1.339622195e-06
0.209      1.142871259e-07
0.209      2.116861592e-03
0.544      7.819864102e-01
0.020      7.777526870e-01
    
```

Figure H-5 Tool Output for Synodic Situation

For elliptical cases (currently considered to be those with eccentricities greater than 0.1), straightforward identification of synodic cases is not possible; instead, the tool identifies candidate situations that will require more specialized follow-up, an example of which is given in Figure H-6:

```

Object is in eccentric orbit with E = 5.51014e-01
Alpha angle alone may not fully characterize co-location.
Investigate the resonance between Alpha Angle and eccentric orbit's period.
Possible resonance values between eccentric orbit and mean violation frequency are:

n1          n2          Period, days    Period Delta, days
1           170          37.1           0.0
2           340          74.2           0.1
3           510          111.3          0.1
4           680          148.3          0.2
5           850          185.4          0.2
6          1021          222.5          0.0
7          1191          259.6          0.0
8          1361          296.7          0.1
9          1531          333.8          0.1
10         1701          370.9          0.2

Note: n1 limited to total time spans (n1*Period1) less than 365 days.
    
```

Figure H-6 Tool Output for Elliptical Case

As can be seen, a certain specialization is required to interpret the results of this tool correctly. It is recommended that qualified personnel be engaged to set up and interpret runs of this utility.



Appendix I. Satellite Covariance Realism Assessment Procedures

I.1 Introduction

This section outlines standard methods for assessing the realism of a covariance produced as part of the state estimates for a spacecraft. Because probability-based methods for satellite conjunction risk assessment in almost all cases require a covariance for their calculation, realistic covariances for both the primary and secondary objects in a conjunction are important.

USSPACECOM has pursued algorithmic improvements over the last twenty years to improve the realism and representativeness of their catalog's covariance matrices, especially in LEO, so covariances that arise from this source have undergone efforts to assess and sustain their trustworthiness. For covariances produced for and contained in O/O ephemerides, the situation is much more uneven. While it is true that in the majority of conjunctions the covariance for the secondary object predominates, in 5% of the conjunctions the primary covariance is larger than the secondary, meaning that it is likely to play a substantial role in collision likelihood assessment in over 10% of conjunctions. This is a large enough subset of a satellite's conjunction history to make covariance realism for the primary object ephemeris an important consideration.

Because the focus of this section is more practical than theoretical, broader questions of covariance formation, methods of covariance propagation, or approaches within the orbit determination process to improve covariance realism are not addressed in detail. Modern orbit determination packages, such as Orbit Determination Tool Kit produced by Analytical Graphics Inc., include features for covariance formation, propagation, and tuning, and their accompanying technical documentation is extensive.

What this section does attempt to do is enumerate the data types needed for covariance realism investigations, introduce the appropriate test statistics, and give some guidance for the evaluation and interpretation of these tests. Most of the information here derives from NASA CARA's practical experience with this problem as there are few published studies on the subject and no practical guides. The particular focus is on determining the realism of predicted covariance generation for actively maintained spacecraft in the presence of precision position data (e.g., on-board GPS or telemetry tracking information) to use for a realism assessment because this is the particular situation presented to O/Os.

Covariance realism assessment consists of three parts: collection/calculation of position error data, calculation of covariance realism test statistics, and proper assessment of those test statistics. After some introductory discussion, each of the three parts of this process will be addressed in detail.

I.2 General Remarks

It is important to recall that a spacecraft state estimate generated through an orbit determination process is an estimate of a mean state (i.e., mean position and mean velocity), and a covariance is a stochastic characterization of the expected errors about that mean state. The governing assumption is that these errors in the individual components conform to a Gaussian distribution;



this is how a covariance matrix alone is able to represent the error distribution without requiring higher-order tensors. The activity of covariance realism is to determine how well actual state errors conform to the Gaussian distribution that the covariance specifies. This statement needs some remarks.

First, the present treatment focuses on the assessment of the realism of the position portion of the covariance, although in most cases the methods advanced here are applicable to a full 6 x 6 or even larger covariance matrix. There certainly are reasonable arguments for evaluating the covariance in element space (usually equinoctial elements) and considering all six elemental representations in any evaluation (Woodburn and Tanygin 2014). However, since the principal purpose of covariance generation at present is for CA applications, and for many CA risk assessment approaches only the position portion of the covariance is used (typically to calculate the probability of collision employing the two-dimensional simplification of the situation; see Chan 2008), it is acceptable to limit oneself to a Cartesian representation of the covariance and focus only on its position portion. Furthermore, a previous study has demonstrated that potential non-Gaussian covariance behavior brought about by working in a Cartesian rather than curvilinear framework rarely affects the outcome of higher-probability CA events, namely those with P_c values greater than $1E-04$ (Ghrist and Plakalovic 2012). Given all of this, limiting the analysis to the position portion of the covariance is reasonable.

Second, because a covariance represents an error distribution, its adequacy as a representation can be evaluated only by examining a set of actual error data and determining whether these data conform to the distribution specified by the covariance. In actuality, this is usually not possible in a straightforward way because a particular covariance is propagated to a particular point in time and relevant only at that moment, and at that moment there is only one state estimate and thus only one error point. It is not possible to determine in any definitive way whether a single point conforms to a distribution; one can determine a percentile level for that point (e.g., for a normal distribution with mean 0 and standard deviation 1, absolute values of 3 or greater should occur only 0.13% of the time) and discern that such points will not occur frequently in such a distribution, but they will occur occasionally, and the situation in the above example may well be one such instance. Typically, this problem is addressed by generating a test statistic that compares each error point to its associated covariance, determining what the distribution of such test statistics should be, and evaluating the conformity of the set of test statistics to the expected distribution. More will be said about this in the following sections that directly address the calculation of test statistics and evaluation of their statistical properties.

Third, as stated in the previous paragraph, a covariance is always a covariance propagated to a particular time. If it is not propagated at all, it is an epoch covariance and is purported to reflect the error of the fit;⁹ if it is propagated, it is intended to be a reflection of the state estimate error

⁹ It is questionable whether the *a priori* covariance emerging from the batch process does indeed do this as it is produced by a calculation that involves only the amount of tracking, times of the tracks, and *a priori* error variances for each observable. It does not give a statistical summary of the actual correction residuals. For this reason, some practitioners have argued for an alternative formulation of the covariance that considers residual error.



of that propagated state. Many different covariance propagation methods are available, each of which achieves a different level of fidelity, so one source of covariance irrealism can be the propagation method itself. It is thus important to group realism evaluations by propagation state; for example, all the calculated test statistics for covariances propagated from 1.5 to 2.5 days could be collected into a single pool and evaluated as a group, and the evaluation could be said to be applicable to covariances propagated about two days. It should also be pointed out that covariance irrealism, however it may be evaluated, is unlikely to correlate directly with propagation time in the sense that if the covariance is made realistic for a particular propagation state, it cannot necessarily be expected to be realistic for all propagation states. When pursuing covariance realism remediation, one frequently needs to decide whether some sort of omnibus improvement over all common propagation states is desired or whether it is preferable to focus on a particular propagation state. Many aspects of the USSPACECOM orbit determination process are tuned to optimize performance at 3 days' propagation; this is not a terrible propagation state to choose for O/O covariance realism optimization for CA applications, although one could also choose something a bit shorter if that were to align more closely with maneuver commitment timelines.

Fourth, in the context of a covariance error calculation, the individual state error values will need to be presumed to constitute independent samples, free of correlation between them. Such an assumption is common in most statistical processes, and it is not surprising to see it arise here. While true sampling independence is unlikely to be achieved, there are certain measures that can be taken to promote it. For example, if one has set up a covariance realism assessment scenario by comparing a predicted ephemeris (with covariance) to a definitive ephemeris, it would be best to take only one point of comparison for each propagation state bin that one wishes to evaluate. For example, in evaluating the two-day propagation state, there would be a temptation to take all of the ephemeris points from, say one hour before the two-day propagation point to one hour after that. If the ephemeris were generated with one-minute time steps, then this would produce 120 points instead of just one, which on the surface would seem to give a nice sample set for a statistical evaluation. However, because the propagated and definitive ephemeris points are highly correlated, in fact little additional information is brought to the evaluation by including the entire dataset. Additionally, all of the statistical evaluation techniques that will be deployed to generate a realism evaluation assume sample independence; in such a case as the one described above, they will miscarry and produce what is likely to be an overly optimistic result. In the present scenario, it is best to limit oneself to a single comparison point per propagation state from this pair of ephemerides and seek out a group of such ephemeris pairs rather than use closely grouped data points from them to try to broaden the sample set.

The procedure recommended above eliminates a particular strain of correlation, but it leads naturally to the consideration of a second type: that between successive ephemerides. Definitive ephemerides are generally formed by piecing together sections of ephemeris from essentially the middle of the fit-spans of moving-window batch orbit determinations, and filters process data sequentially but with a forgetting-rate matrix that reduces the influence of older data as a function of data age. But in each case, a fair amount of ephemeris time needs to pass before the generated points are statistically independent of a given predecessor set. The independence condition could be fulfilled by forcing the generation of products from entirely different datasets,



but this is not a practical procedure. In past projects, trying to reduce the data overlap to less than 50%, meaning that subsequent ephemerides to be used in covariance realism evaluations need to be spaced so that they share less than 50% of the generating data with the previous ephemeris, has been considered acceptable in practice, although a lower figure would be desirable. One need not become overly fastidious about this issue, but it is important to take whatever practical steps one can to reduce the influence of sample data correlation on the realism assessment.

I.3 Part I: Error Computation

State errors are computed by comparing a predicted state estimate to some type of state truth source. The predicted data typically come from a predicted ephemeris with a predicted covariance associated with each ephemeris point. Truth data can come from a variety of sources: on-board GPS precision position data, precision telemetry tracking data, or a precision ephemeris constructed from either of these. Ideally, both the predicted and truth data should align temporally (i.e., their time points should be exactly the same), and they should both possess covariance data at each time point. In actuality, it is rare that both of these conditions are met, and often neither is.

Covariance data are often not available for the truth data source. Precision tracking or GPS data can include covariance information as an output from a tracking or combinatorial filter; when precision ephemerides are constructed, sometimes covariances can be synthesized from the abutment differences observed between the joined “pieces” of ephemeris. However, it is frequently the case that even when this is theoretically possible, the construction software has failed to include this as a feature. If no covariance data are available for the truth source, one must then determine whether the errors in the truth data are so much smaller than the errors in the predicted ephemeris under evaluation that the errors in the truth data can safely be neglected. The degree of accuracy difference that would allow this assumption is a matter of opinion, but most commentators would probably agree that an order of magnitude difference would be an acceptable decrement; and in some cases, even smaller ratios could be tolerated. These differences, however, must be considered at the individual component level. If the error in one component tends to dominate the entire arrangement (as is often encountered with the in-track component due to inadequate drag modelling), then the overall normalized error of all three position points in the truth ephemeris can be much smaller than the overall error in each point in the predicted ephemeris, yet component errors for the cross-track and radial components could be of a similar magnitude in both truth and prediction. In such a case, the omnibus test statistic to be described in the next section would be distorted because it treats the normalized errors in each component essentially equally.

If the time-points between truth and predicted data do not align, then some sort of intermediate-value-determination scheme must be used for one of the two datasets. The highest-fidelity approach would be to use an actual numerical propagator to produce aligning values, but if the points for the source to be adjusted are spaced sufficiently closely, then a number of different interpolation approaches produce results quite adequate for a covariance realism assessment. While the interpolation of state values is relatively straightforward and can be accomplished satisfactorily with different approaches (e.g., Lagrange, Hermite), the interpolation of covariances is much more problematic. Different interpolation methods appear to work better in



different situations; there is no single accepted method to compare multiple covariances to evaluate their relative fidelity against each other, and there is a danger of certain interpolation methods producing non-positive-definite covariances, which do not make physical sense with the current problem. Given this situation, if both data sources (truth and prediction) can be interpolated, but only one of the two data sources possesses accompanying covariance data (presumably the predicted source), then it is usually preferable to interpolate the data source that lacks the covariance as this eliminates the question of how to handle covariance interpolation. If both sources possess covariance, then one would in general interpolate the source with the closer point spacing. Alfano (2004) outlines certain covariance interpolation techniques and provides test results. There is also an interest in interpolating not the covariance, but the state transition matrices associated with the straddling ephemeris points. While this approach can in some cases produce slightly less desirable interpolated covariances, it does essentially guarantee that the interpolated result will be positive definite.

Once position and covariance data for the predicted and truth datasets are aligned at the time-point of interest, calculation of the error is straightforward: it is simply the (subtracted) difference between states as long as they both are rendered in the same coordinate system. The subtraction convention does not actually matter given the way the test statistic is computed, but for consistency, one can follow the “observed – expected,” or truth – prediction, paradigm. The two covariances for each comparison point (one for the truth position and one for the predicted position) are simply added together, presuming, of course, that they are in the same coordinate system. Thus, for each point of comparison, one will have a position difference (with three components) and a corresponding combined covariance (or just the predicted source’s covariance if the truth source’s errors are presumed to be extremely small).

While one can conduct an entire covariance realism evaluation in the Earth-Centered Inertial (ECI) reference frame, it is often more meaningful to move it into the Radial – In-track – Cross-track (RIC, which is also called the UVW) frame—a reference frame that is centered on the object itself. If only one of the evaluation data sources has an associated covariance, then it probably makes sense to make that source the center for the RIC coordinate frame. If a covariance is provided for both, then it does not really matter which is selected to ground the RIC frame. The use of the RIC frame is helpful in moving from a finding of irrationalism to remediation suggestions.

I.4 Part II: Test Statistic Computation

As stated previously, a position covariance, which in this treatment is the portion of the matrix to be tested for proper error representation, describes a three-dimensional distribution of position errors about the object’s nominal estimated state, and the test procedure is to calculate a set of these state errors and determine whether their distribution matches that described for the position covariance matrix. To understand the test procedure, it is best to consider the problem first in one dimension, perhaps the in-track component of the state estimate error. Given a series of state estimates for a given trajectory and an accompanying truth trajectory, one can calculate a set of in-track error values, here arranged in vector form and given the designation ϵ , as the differences between the estimated states and the actual true positions. According to the assumptions previously discussed about error distributions, this group of error values should conform to a



Gaussian distribution. As such, one can proceed to make this a “standardized” normal distribution, as is taught in most introductory statistics classes, by subtracting the sample mean and dividing by the sample standard deviation:

$$\frac{\varepsilon - \mu}{\sigma} \tag{I-1}$$

This should transform the distribution into a Gaussian distribution with a mean of 0 and a standard deviation of 1, a so-called “z-variable.” Since it is presumed from the beginning that the mean of this error distribution is 0 (because the state estimate is an estimate of the mean and is presumed to be unbiased), the subtraction as indicated in the numerator of Equation I-1 should be unnecessary, simplifying the expression to:

$$\frac{\varepsilon}{\sigma} \tag{I-2}$$

It will be recalled that the sum of the squares of n standardized Gaussian variables constitutes a chi-squared distribution of n degrees of freedom. As such, the square of Equation I-2 should constitute a one-degree-of-freedom chi-squared distribution. This particular approach of testing for normality—evaluating the square of the sum of one or more z-variables—is a convenient approach for the present problem, as all three state components can be evaluated as part of one calculation (ε_u represents the vector of state errors in the radial direction, ε_v the in-track direction, and ε_w the cross-track direction):

$$\frac{\varepsilon_u^2}{\sigma_u^2} + \frac{\varepsilon_v^2}{\sigma_v^2} + \frac{\varepsilon_w^2}{\sigma_w^2} = \chi_{3\,dof}^2 \tag{I-3}$$

One could calculate the standard deviation of the set of errors in each component and use this value to standardize the variable, but it is the covariance matrix that is providing, for each sample, the expected standard deviation of the distribution. Since the intention here is to test whether this covariance-supplied statistical information is correct, the test statistic should use the variances from the covariance matrix rather than a variance calculated from the actual sample of state estimate errors. For the present moment, it is helpful to presume that the errors align themselves such that there is no correlation among the three error components (for any given example, it is always possible to find a coordinate alignment where this is true, so the presumption here is not far-fetched; it is merely allowing that that particular coordinate alignment happens to be the RIC or UVW coordinate frame). In such a situation, the covariance matrix would lack any off-diagonal terms and thus look like the following:

$$C = \begin{bmatrix} \sigma_u^2 & 0 & 0 \\ 0 & \sigma_v^2 & 0 \\ 0 & 0 & \sigma_w^2 \end{bmatrix} \tag{I-4}$$

and its inverse would be straightforward:



$$C^{-1} = \begin{bmatrix} 1/\sigma_u^2 & 0 & 0 \\ 0 & 1/\sigma_v^2 & 0 \\ 0 & 0 & 1/\sigma_w^2 \end{bmatrix}. \quad (\text{I-5})$$

If the errors for a single state are formulated as:

$$\boldsymbol{\varepsilon} = [\boldsymbol{\varepsilon}_u \quad \boldsymbol{\varepsilon}_v \quad \boldsymbol{\varepsilon}_w] \quad (\text{I-6})$$

then the pre-and post-multiplication of the covariance matrix inverse by the vector of errors (shown in Equations I-3 through I-6) will produce the expected chi-squared result:

$$\boldsymbol{\varepsilon} C^{-1} \boldsymbol{\varepsilon}^T = [\boldsymbol{\varepsilon}_u \quad \boldsymbol{\varepsilon}_v \quad \boldsymbol{\varepsilon}_w] \begin{bmatrix} 1/\sigma_u^2 & 0 & 0 \\ 0 & 1/\sigma_v^2 & 0 \\ 0 & 0 & 1/\sigma_w^2 \end{bmatrix} \begin{bmatrix} \boldsymbol{\varepsilon}_u \\ \boldsymbol{\varepsilon}_v \\ \boldsymbol{\varepsilon}_w \end{bmatrix} = \frac{\boldsymbol{\varepsilon}_u^2}{\sigma_u^2} + \frac{\boldsymbol{\varepsilon}_v^2}{\sigma_v^2} + \frac{\boldsymbol{\varepsilon}_w^2}{\sigma_w^2} = \chi_{3\text{ dof}}^2 \quad (\text{I-7})$$

What is appealing about this formulation is that, as the covariance becomes more complex and takes on correlation terms, the calculation procedure need not change: the matrix inverse will formulate these terms so as properly to apportion the variances among the U , V , and W directions, and the chi-squared variable can still be computed with the $\boldsymbol{\varepsilon} C^{-1} \boldsymbol{\varepsilon}^T$ formulary. For such a situation in two-dimensions (chosen here for illustrative purposes because the expression is less complex) in which the error quantities are $\boldsymbol{\varepsilon}_x$ and $\boldsymbol{\varepsilon}_y$ and the correlation coefficient is ρ , the entire equation, with correlation terms included, assumes the form:

$$\boldsymbol{\varepsilon} C^{-1} \boldsymbol{\varepsilon}^T = \frac{1}{(1-\rho)^2} \left(\frac{\boldsymbol{\varepsilon}_x^2}{\sigma_x^2} + \frac{\boldsymbol{\varepsilon}_y^2}{\sigma_y^2} - \frac{2\rho\boldsymbol{\varepsilon}_x\boldsymbol{\varepsilon}_y}{\sigma_x\sigma_y} \right) \quad (\text{I-8})$$

One can observe that if the correlation coefficient is zero, the equation reduces to the two-dimensional equivalent of the form shown in Equation I-7 above. As the correlation coefficient moves from zero to a more substantial value, the test statistic encounters a trade-off between the overall inflating effect of the $(1-\rho)$ multiplier and the subtracted correlation term.

The quantity $\boldsymbol{\varepsilon} C^{-1} \boldsymbol{\varepsilon}^T$ is called the Mahalanobis distance (technically, it is the square of the Mahalanobis distance), named after the mathematician P.C. Mahalanobis (who, by the way, was a friend of the 20th-century mathematical genius Ramanujan). This construct is a convenient and useful way to calculate a normalized distance.

I.5 Part III: Test Statistic Evaluation

It is very well that a test statistic can be derived that, if the covariance is realistic, should conform to a known statistical distribution, but of course, there needs to be some method for testing a group of these test statistics to determine if in fact they do conform to the expected distribution. Such a desire leads the investigation to the statistical subdiscipline of “goodness of fit.”



Every student of college statistics learns about Analysis of Variance (ANOVA), the particular procedure for determining whether two groups of data can essentially be considered the same or different. More particularly, it is a procedure for determining whether the experimental distribution, produced by the research hypothesis, can be considered to come from the parent distribution represented by the null hypothesis, and the operative statistic arising from the analysis is the p -value: the likelihood that the research dataset can be considered a sample drawn from the null hypothesis's parent distribution. If this value becomes small, such as only a few percent, it means that there are only about two or three chances in one hundred that the experimental dataset would have been generated from sampling from the null hypothesis dataset. In this case, the research and the null hypothesis outcomes can be considered to be palpably different, which would be likely to lead to the rejection of the null hypothesis and the embrace of the research hypothesis. This procedure is a specific example of statistical hypothesis testing.

A similar procedure can be applied to evaluate goodness of fit, namely, to evaluate how well a sample distribution corresponds to a hypothesized parent distribution. In this case, the general approach is the reverse of the typical ANOVA situation: it is to posit for the null hypothesis that the sample distribution does indeed conform to the hypothesized parent distribution with a low p -value result counseling the rejection of this hypothesis. This approach does favor the association of the sample and the hypothesized distribution, which is why it is often called "weak-hypothesis testing". It is not an unreasonable method: what is being sought is not necessarily the "true" parent distribution but rather an indication of whether it is reasonable to propose the hypothesized distribution as the parent distribution for the experimental data that one has generated or calculated. In the present case, the question to be posed is whether the squared Mahalanobis distance histories, calculated by the procedure in Part II, conform to a 3-degree-of-freedom chi-squared distribution, namely, whether they can be considered to have been drawn from a 3-degree-of-freedom chi-squared distribution as a parent distribution.

There are several different mainstream techniques for goodness-of-fit weak-hypothesis testing: moment-based approaches, chi-squared techniques (not in any way linked to the fact that the present application will be testing for conformity to a chi-squared distribution), regression approaches, and Empirical Distribution Function (EDF) methods. The easiest and most direct of these is simply a test of the first moment of the distribution (that is, the mean), which, if normalized by the degrees of freedom of the distribution, should be unity (or, in the present case, should take on an unnormalized value of 3). The square root of this mean, the so-called Mahalanobis distance, is a good estimate of a single-value scale factor describing the covariances departure from reality, and as such, it is a convenient way to compare the results from different covariance correction techniques as well as estimate a single-value scale factor that could potentially be used to remediate an irrealism situation.

While this test is easy to apply and expeditious for generating comparative results, in comparison to other goodness-of-fit tests it lacks power. Indeed, many different distributions could have the same mean yet be substantially different in the overall behavior or in the tails: a close match of the mean is a necessary but not sufficient condition for matching a distribution. To evaluate the match between entire distributions, the EDF methodology is generally considered to be both the most powerful and most fungible to different applications. For this reason, it does not make



sense to apply formal goodness-of-fit tests to first-moment results (i.e., matching of the mean); these should be used merely as a methodology to compare performance for the corrected versus uncorrected cases and intra-correction-methodology performance.

The general EDF approach is to calculate and tabulate the differences between the Cumulative Distribution Function (CDF) of the sample distribution and that of the hypothesized distribution, to calculate a goodness-of-fit statistic from these differences, and to consult a published table of p -values for the particular goodness-of-fit statistic to determine a significance level. Specifically, there are two goodness-of-fit statistics in use with EDF techniques: supremum statistics, which draw inferences from the greatest deviation between the empirical and idealized CDF (the Kolmogorov-Smirnov statistics are perhaps the best known of these); and quadratic statistics, which involve a summation of a function of the squares of these deviations (the Cramér – von Mises and Anderson-Darling statistics are the most commonly used). It is claimed that the quadratic statistics are the more powerful approach, especially for samples in which outliers are suspected, so it is this set of goodness-of-fit statistics that are recommended for covariance realism evaluations. The basic formulation for both the Cramér – von Mises and Anderson-Darling approaches is of the form:

$$Q = n \int_{-\infty}^{\infty} [F_n(x) - F(x)]^2 \psi(x) dx \quad (\text{I-9})$$

The two differ only in the weighting function ψ that is applied. The Cramér – von Mises statistic is the simpler:

$$\psi(x) = 1 \quad (\text{I-10})$$

setting ψ to unity. The Anderson-Darling is the more complex, prescribing a function that weights data in the tails of the distribution more heavily than those nearer the center:

$$\psi(x) = \{F(x)[1 - F(x)]\}^{-1} \quad (\text{I-11})$$

The Anderson-Darling construct is thus more sensitive to outliers. Given NASA CARA’s experience that outliers frequently creep into covariance realism evaluations and introduce statistical processing issues, it is recommended to choose the somewhat more permissive Cramér – von Mises statistic for covariance realism investigation purposes.

It is a straightforward exercise to calculate the statistic in Equation I-9, discretized for the actual individual points in the CDF for each trajectory (that is, changing the integral into a summation), and for convenience, this quantity is called the “Q-statistic.” The step after this is, for each Q-statistic result, to consult a published table of p -values (determined by Monte Carlo studies) for this test to determine the p -value associated with each Q-statistic (a good source for these tables as well as an excellent treatment of the overall subject is given in D’Agostino and Stephens 1986). The usual procedure is to set a p -value threshold (e.g., 5%, 2%, 1%) and then to determine whether the sample distribution produces a p -value greater than this threshold (counseling the retention of the null hypothesis: sample distribution conforms to hypothesized distribution) or less than this threshold (counseling rejection of the null hypothesis: sample



distribution cannot be said to derive from the hypothesized distribution as a parent). One can also interpolate to determine the precise p -value for each test situation. NASA CARA has source code and test cases that perform all of the above calculations. As soon as it is approved for public release, such code (along with a worked example) will be placed in the NASA CARA software repository for free download (See Section 7, Contact Information in this document for the specific URL.)

Finally, it should be noted that, even though there are normalization provisions within the EDF formulation, the results to some degree do depend on the size of the sample. In a way this is considered already by accommodation within the p -value tables for sample size, but because of the quadratic-sum nature of the test statistic, the procedure can still be overwhelmed by large sample sizes. One approach to mitigating this situation is to pick a standard sample size—perhaps somewhere in the neighborhood of 50 samples—and calculate the test statistic in a (with-replacement) resampled manner, producing a CDF of the p -values attained for each sample. As an example, suppose that for a particular evaluation there are 100 error vectors with associated covariance and therefore 100 test statistic ($\epsilon C^{-1} \epsilon^T$) results. When running the goodness-of-fit test, one might choose 1000 random, 50-point samples (with replacement) from this set of 100 values and test each, producing a CDF chart of the 1000 p -values obtained from the resampling investigation.

What is an acceptable level for a p -value result—one that would indicate that the error distribution matches that of the covariance? In goodness-of-fit practice, rarely is a significance level greater than 5% required, and levels of 2% or even 1% are often accepted. It probably can be said that values less than 1% cannot allow the conclusion that there is any real conformity to the hypothesized distribution. At the same time, it should be added that the calculation is rather sensitive to outliers and that this should be kept in mind when interpreting results.

If results from different remediation approaches match the hypothesized distribution closely enough, then comparison of different p -value levels can serve as a notional indication of the relative performance of these different approaches. However, if the performance is such that the hypothesized distribution is not approached all that closely by any of the correction mechanisms, then a situation can be encountered in which the results fall off of the published tables of p -values, and it becomes extremely difficult to compare the results of the different methods conclusively. In such a case, it may be possible to draw some broad comparative results from looking at the Q-values rather than the associated p -values, but in all likelihood, it will be necessary to revert to simply a comparative-results set such as first-moment tests.

I.6 Part IV: Data Outliers and Conditioning

As remarked in the above section, goodness-of-fit test results can be sensitive to outliers; this is true whether one interprets results visually or uses a formal technique. The deleterious impact on test results is mitigated somewhat using the resampling approach discussed above, by which the contribution of the outliers to each individual test is lessened. However, the fact remains that bad data do enter the orbit determination process, and the failure to take some account of this reality can produce situations in which the covariance realism assessment problem becomes intractable.



The entire covariance realism assessment process is grounded on the notion that individual component errors are normally distributed, and this situation allows for certain techniques to identify outliers. The usual “x-sigma” filter is a naïve and unscientific method for outlier identification and is especially difficult to justify for the Gaussian distribution, where there is some developed theory for outlier identification. A superior approach is the Grubbs outlier test, which provides a formal statistical test for outliers but works only in situations with a single outlier and cannot be applied recursively (Grubbs 1969). In the case of multiple outliers, the procedures of Rosner (1975, 1977) are applicable but must be applied with an *a priori* guess of the number of potential outliers. That is, one must first inspect the data to assemble a proposed set of outliers and then test this group as outliers for a particular significance level. Because only a single component can be assessed at a time, to use this approach it is probably best to design a tool that can examine a particular error point’s behavior in all three components (compared to the rest of the main distribution) in order to determine which set of points might constitute an outlier set and test that set. Several such attempts may be needed before a set of points can be identified as outliers to a given significance level.



Appendix J. CARA CA Risk Assessment Tools Validation

CARA and FOD are chartered to perform CA screenings, risk assessment, and mitigation action guidance on behalf of NASA missions, so typically there is no particular need for missions to obtain or develop CA risk assessment evaluation tools. However, cases can occasionally arise in which it is not practical or desirable for CARA/FOD to perform these calculations such as potential autonomously controlled missions that will perform on-board CA. In such cases, the needed calculations or tools will be validated by CARA/FOD before operational use. Methods used for validation include the following:

- **Inspection.** Used to verify that a particular datum or feature is present, that a display or graphic is properly constructed to convey a particular concept, etc.
- **Analysis/Documentation.** Used to ensure that a particular algorithm is theoretically or practically sound, that a needed parameter is properly converted to a different reference frame or units, etc.
- **Formal Test.** Used to ensure that critical calculations are performed correctly by executing a formal test with preconfigured input data and expected results, which are compared to the test article's results and any differences satisfactorily explained.

Formal test is the most frequently used methodology within CARA validation. CA risk assessment calculations for which formal test cases with expected results presently exist include the following:

- Two-dimensional Pc calculation with and without covariance correlation correction and non-positive-definite covariance correction
- Two-dimensional Pc calculation when one covariance is missing (Frisbee method)
- Three-dimensional Pc calculation (Coppola-Hall method) with and without covariance correlation correction and non-positive-definite covariance correction
- Pc calculation by Brute Force Monte Carlo (from epoch)
- Pc sensitivity to space weather mismodeling
- Conjunction relative state comparison
- Collision consequence evaluation (anticipated number of resultant debris pieces above a certain size)
- Maneuver trade space calculations for single and multiple conjunctions

The NASA CARA software repository includes test cases for software validation. (See Section 7, Contact Information in this document for the specific URL.)



Appendix K. R-15 Message

The Ready minus 15 (R-15) Vehicle Information Message (VIM) is a collection of launch-related information that is distributed by the launch vehicle operator 15 days before launch to inform concerned parties about upcoming launch activity. It contains key information to support the tracking and early acquisition of the injected payloads as well as other launch-related objects such as rocket bodies. The extract below gives preparation instructions for this message. An official copy of the R-15 form can be obtained from the Space-Track.org website maintained by USSPACECOM. A sample structure for the R-15 is provided below for reference only.

K.1 R-15 Preparation Instructions

How to fill out a R-15/VIM report:

ITEM ENTRY

- 1 Launch site (Site, Pad, Country)
- 2 Launch date (Greenwich Mean Time (GMT))
- 3 Earliest launch time. (GMT, HH:MM:SS)
- 3A Latest launch time. (GMT, HH:MM:SS)
- 4 List the total number and name of each object to achieve orbit.
- 4A Payload(s) to achieve orbit. Include the nominal (operational) lifetime and operating position for each.
- 4B Rocket bodies (booster segments) to achieve orbit. If none achieve orbit, enter "none."
- 4C All other objects achieving orbit, including debris, debris clusters, bolts, and so forth. If none will achieve orbit, enter "none."
- 5 Launch booster and sustainer description. If booster is augmented by strap-on motors, list the number and type.
- 6 Point of contact for the launch.
- 7 Mission brief of payload(s).
- 8 Transmitting frequency and power of all devices (including booster segments and continuous radio transmissions) and schedule and power of all lights (if any) throughout the operational life. Statement of whether emission is fixed by program, command, or transponder tracking signal.



R-15/VIM REPORT ATTACHMENT A: Keplerian orbital parameters to include sequence of events from liftoff (HH:MM:SS = 00:00:00) to final injection into operational orbit. Require times for each in HH:MM:SS from liftoff. Events include: separation of booster(s)/stage(s), motor ignition(s)/cutoff(s), jettison of pieces (fairings etc.), maneuvers and reorientation, deorbit and ejections) of packages/booms and so forth.

K.2 R-15/VIM Report Structure

ITEM 1	Launch Site:
ITEM 2	Launch Date:
ITEM 3	Earliest and latest possible launch time (GMT):
ITEM 3A	Latest possible launch time (GMT):
ITEM 4	List the total number and name of each object to achieve orbit.
ITEM 4A	Payload(s) to achieve orbit. Include the nominal (operational) lifetime and operating position for each:
ITEM 4B	Rocket bodies (booster segments) to achieve orbit. If none achieve orbit, enter "none."
ITEM 4C	All other objects achieving orbit, including debris, debris clusters, bolts, and so forth. If none will achieve orbit, enter "none."
ITEM 5	Launch booster and sustainer description. If booster is augmented by strap-on motors, list the number and type:
ITEM 6	Point of contact for the launch:
ITEM 7	Mission brief of payload(s):
ITEM 8	Transmitting frequency and power of all devices (including booster segments and continuous radio transmissions) and schedule and power of all lights (if any) throughout the operational life. Statement of whether emission is fixed by program, command, or transponder tracking signal.

K.3 R-15/VIM Report Attachment A Structure

MET (s)	Event timeline	Altitudes (km)***		Inclination (deg)	Eccentricity	SMA (km)	Arg. of perigee (deg)	RAAN ** (deg)	True Anomaly (deg)	Latitude (deg)	Longitude (deg)	Relative Velocity (m/s)	Current Altitude (km)	Period (min)
		Apogee	Perigee											

NOTES:



Appendix L. Commercial Data in NASA Conjunction Assessment

In this last decade there have been remarkable developments in the breadth and quality of commercial SSA data from essentially no commercial presence at all to a number of different vendors, each specializing individually in radar, optical, or passive radio frequency satellite tracking. Many of these vendors include among their offerings not just tracking measurement data but additional SSA products as well such as vectors, ephemerides, and even CA analysis outputs such as Conjunction Data Messages (CDMs).

In examining and facilitating the use of commercial data in NASA CA calculations, CARA and JSC FDO adhere to the following principles:

1. **Use raw observation data only, and combine them with SSN observations for a single solution.** Of all of the data types offered by commercial providers, only satellite tracking data themselves are fully and unambiguously of use to the NASA CA enterprise. Because NASA has access to both the USSPACECOM satellite tracking data from the Space Surveillance Network (SSN) and its space catalog maintenance operational system, NASA has the ability to include raw commercial observation data into a combined close approach prediction. This process involves combining all of the measurement data on objects of interest from both the SSN and commercial sources to perform a single orbit determination. The predicted close approaches are calculated from this combined source. The CARA and FOD personnel, embedded in the 18 SPCS work center at VAFB, presently have this capability in a “sandbox” area on the operational system that is not used by 18 SPCS or other operators. The sandbox area enables operators to combine the SSN and commercial data in a native and compatible way to produce a single risk assessment solution that can be acted upon unambiguously. Receiving higher-level products from commercial vendors, such as vectors or even CA outputs like CDMs, introduces multiple solutions into the enterprise and creates significant operational difficulties in trying to adjudicate among them as there is no demonstrated straightforward method for combining these higher-level products into a reliable single ensemble solution. Therefore, the user would be left with multiple solutions that cannot be adjudicated without truth data and are therefore not useful for time-critical mitigation decision making.
2. **All data must be validated.** All data used in orbit determination solutions to maintain the catalog are calibrated on a daily basis to ensure the removal of any data biases that might skew the CA screening results. Using raw observation data from commercial vendors (as opposed to post-processed solutions) enables straightforward characterization and numerical validation. Commercial vendors supporting NASA would be asked regularly to track and submit the tracking data taken on calibration satellites for which precision ephemerides are available; these satellites can be both dedicated calibration spheres (such as those used by the International Laser Ranging Service) and USG payloads with stable orbits. The commercial tracking data are compared to the precision ephemerides and the residuals characterized: for each observable, mean errors can be



subtracted from the data themselves and thus do not affect the data quality, error variances can be used as sensor data weighting factors in the orbit determination process, and non-Gaussianity and variation of the means and variances with time can indicate an unstable situation that may render a particular commercial provider's data unsuitable for use. For stable commercial data sources, quantification of mean error and error variance allows the data to be combined with SSN data (which also have quantified error means and variances) to produce an orbit determination result that has weighted the constituent data correctly and thus properly considers the relative accuracy of the inputs in forming the desired single solution.

- 3. Cost/benefit analysis must be undertaken before purchasing.** While following the recommendations of 1) and 2) above would allow the introduction and proper use of commercial SSA data in NASA CA, a third consideration focuses on whether the use of these data actually would result in a significant operational benefit. NASA, as a USG entity, follows the guidance of the National Space Policy to receive CA services from the 18 SPCS, who maintains the U.S. catalog of space objects. While it is true that adding properly calibrated commercial tracking data should result in a more reliable orbit determination solution with a smaller uncertainty, the effect may be such that the use of these additional data rarely, if ever, would actually result in a different operational decision regarding whether or not to require and execute a CA mitigation action. Most CA situations receive adequate SSN tracking data to the degree that adding some additional data does not change the risk assessment appreciably and therefore does not alter the operational decisions. A complex cost-benefit analysis would need to be performed in order to understand the financial obligation appropriate to purchase commercial data to augment the existing NASA CA enterprise based on the risk mitigation expected. If, for example, having commercial data changed a conjunction P_c value from $5E-05$ to $1E-04$, this would cause a mitigation action to be made when, without the data, it probably would not have been if the P_c threshold for maneuvering were set at $1E-04$, so the operational outcome was changed to lower the risk. However, the risk exposure if the mitigation action were not taken is still very low: an actual risk of 1 in 10,000 chance of a collision versus the 1 in 20,000 chance calculated with the SSN data alone. Benefit would depend on how many times per year these decisions would be altered and thus how much risk is mitigated. By most estimates, over 90% of the objects larger than 1cm and thus able to render a satellite inoperative in a collision are too small to be tracked and thus are not considered in the CA enterprise at all. Therefore, part of the trade space is to consider how much adding commercial data to the CA process reduces in collision compared to the background risk. This decision would likely vary by orbit altitude based on trackability of objects of various sizes and existing debris population.



Appendix M. Use of the Probability of Collision (P_c) as the Risk Assessment Metric for Conjunction Assessment

M.1 Introduction

Satellite conjunction assessment comprises two principal goals: to protect important assets from premature and unplanned mission failure due to preventable collisions and to assist in the regulation and minimization of space debris production. The first of these goals is in its details largely left up to the individual O/O to adjudicate. After all, the particular “value” of any given on-orbit asset, and therefore the level of risk of its possible loss that one is willing to bear, are governed by the specifics of the situation including the age (and perhaps importance) of the satellite and mission, the degree to which the mission objectives have already been met, and mission redundancy provided by additional spacecraft. It is best left to the individual mission project office to choose appropriate CA actions (or inactions) relating directly to mission preservation. However, the second goal of overall space debris production minimization is much broader in scope. Born of the desire to keep key orbital corridors free from debris pollution to allow their perpetual continued use by all space actors, this goal transcends the health and preservation of the individual mission. It is in response to this second goal that actual CA risk assessment and mitigation thresholds are articulated. Mission operators can always take a more conservative posture than what is appropriate for debris mitigation, but they must at a minimum embrace a minimum level of responsibility to preserve key orbital corridors for future use.

To this end, the purpose of satellite conjunction assessment and risk analysis is neither to maximize safety, nor fully to minimize risk of satellite conjunction with other objects, nor absolutely eliminate the likelihood of space debris production. Instead, the official CARA statement of purpose is more nuanced: to take prudent measures, at reasonable cost, to improve safety of flight, without placing an undue burden on mission performance. This statement includes a number of qualitative terms, such as “prudent,” “reasonable,” and “undue burden” for ultimately, it is not a scientific conclusion but a series of prudential judgments to assemble guidelines that work to prevent space debris pollution while accommodating the competing claims of mission execution and inherent levels of risk assumed simply by launching a satellite at all.

It is this last consideration, namely the risk necessarily assumed simply by launching a satellite, that requires attention, especially in selecting a collision likelihood metric and evaluation paradigm. A 2020 NASA ODPO space debris catalog generated for CARA contained over 300,000 objects larger than 1 cm in physical size, which is the level at which an object is typically considered able to penetrate a satellite’s shielding and render it inoperative. After the U.S. Space Force’s Space Fence radar becomes operational and objects down to 5 cm in physical size are routinely trackable, the CARA Team estimates that the size of the space catalog will eventually reach in the neighborhood of 50,000 objects. This means that about five-sixths of the objects large enough to leave a satellite in a failed and uncontrolled state if a collision should occur with them will be untracked. Presumably, that same proportion of the close-approach events that actually occur are not discoverable with current sensor technology and cannot be mitigated. Given then that only about one-sixth of the close-approach events are trackable and in



principle actionable, the overall CA risk analysis process (and the risk assessment parameters that drive it) should be commensurate with a situation in which greater than 85% of similarly likely collisions cannot be addressed at all and thus constitute a risk that simply has to be accepted. It is not reasonable to adopt an extremely conservative risk management position for known conjunctions when so substantial a portion of the actual collision risk is accepted without any possible mitigation as part of the cost of space operations.

However, despite the large background collision likelihood that is accepted simply by launching a satellite at all, there is still great value to on-orbit conjunction analysis and risk assessment / mitigation. First, because collisions between protected payloads and large secondary objects produce by orders of magnitude the most debris, there is substantial benefit to avoiding collisions of this type. These large secondary objects are resident and well maintained in the present satellite catalog, so they represent rather straightforward cases for conjunction assessment. Second, when serious conjunctions of any type are identified and are well determined, these also represent straightforward situations for which due diligence requires risk assessment and possible mitigation action. It is only when the situation is poorly determined—when some information about a conjunction is present but not of the sufficiency needed to conclude that a problematic situation is in fact at hand—that requiring remediation actions through an overweening conservatism is not appropriate, given the size of the accepted background risk.

The purpose of the this appendix is to explain the selected risk assessment metric and interpretive framework for that metric, chosen to enable proper balance between mitigation of truly problematic conjunctions and an unnecessary conservatism that reduces spacecraft lifetime and unduly imposes on science mission objectives.

M.2 Point of Reference: CA based on Miss Distance

In the early days of CA, satellite position uncertainty data were not regularly available, so the only orbital data product on which to base CA decisions was the miss distance at TCA between the two conjuncting satellites. If the miss distances were smaller than the combined sizes of the two objects (a quantity usually on the order of 10-20 meters), then it was clear what was to be done: a mitigation action would be necessary to modify the primary satellite's orbit to increase this miss distance to a safe level. But it was also known that there was often significant uncertainty in the two objects' predicted state estimates, so miss distances that were larger than the objects' combined size were also likely to be threatening—but how threatening, precisely? Similarly, if a mitigation action such as a maneuver were pursued, how large should such an action be to guarantee a desirable level of safety? Because the answers to these questions were not known, very conservative miss-distance thresholds were then embraced, often based on very little analysis, which led to a great deal of operational anxiety and unnecessary mitigation actions. For example, nearly all conjunctions with miss distances (at TCA) less than a value of 5 km present perfectly safe encounters, but all of these could be treated as worrisome and worthy of potential mitigation actions. Such a situation could not be sustained operationally: it went beyond prudence and it presented an undue burden to mission operations. Methods were thus developed that considered the characterized error in the state estimates to determine the likelihood that the two trajectories at TCA would have a close approach smaller than the combined sizes of the two satellites.



Approaches that view the problem in this way, namely those that give an actual likelihood that the miss distance will be small enough to cause a collision, are certainly an advance over the use of the miss distance alone. It is important to remember, however, that the probabilistic answer they produce is solely a function of the quality of the astrodynamics data that are used as input to the computation—the tracking information available and the errors encountered in predicting states from epoch to TCA. In truth, the probability of collision for any given conjunction is actually either 1 or 0: the conjunction is either going to result in a collision, or it is not. The probabilistic framing comes from the degree of predictive certainty of the outcome brought by the quality of the astrodynamics data in the particular case. There is no “definitive” collision likelihood value for a particular conjunction. Improvements in data and data quality will push the calculated answer further toward the values of 0 or 1 depending on whether the two satellites actually are on path to collide, and because satellite collisions are extremely rare events, in nearly every case the true likelihood of collision is 0. Sometimes interest is expressed in wishing to obtain sufficient tracking data to allow the “true” collision likelihood to be calculated, or after the event, to rework the solution using the tracking data collected *ex post facto* to establish what the “definitive” collision likelihood actually was. None of these considerations is in the end helpful, or for that matter even possible. In the first case, while more and better data will provide a superior answer, any value short of 0 or 1 is just an estimate constrained ultimately by the inadequacy of the tracking data. In the second case, the *ex post facto* solution is not definitive but is yet another estimate, this time with somewhat better input data. The probabilistic framework is helpful and desirable, but in the end, it reflects what is known about the conjunction situation rather than being a stand-alone, absolute assessment of collision likelihood.

M.3 Use of P_c for CA Risk Assessment Requirements

Many CA risk assessment metrics that have been proposed in the critical literature; Hejduk and Snow (2019) give a useful overview, description, and attempted categorization. In choosing a particular CA collision likelihood metric for requirements specification purposes, different considerations, technical and otherwise, are relevant:

- Suitability to NASA’s risk posture regarding CA risk assessment;
- Ability to represent the collision likelihood in a manner commensurate with this risk posture;
- Straightforwardness and tractability of calculation;
- Ease of conceptual grasp of the metric and its origins; and
- Acceptance within the CA industry.

From these considerations, CARA has chosen the Probability of Collision (P_c) as the metric for high-risk conjunction identification. Of the different options examined, this metric was the best candidate to evaluate collision risk given the desired context of balancing safety of flight with minimizing the disruption of mission objectives. It is straightforward to calculate and, with certain recent improvements, can be computed quite accurately in many astrodynamically demanding circumstances. It can be explained relatively easily to, and be interpreted by, decision makers. It can be turned into a statement about miss distances if this framework is desired. Finally, it is the emerging metric of choice among the military, civil, and commercial CA industries.



The P_c metric, as calculated by analytic methods, was first explained formally in 1992 (Foster and Estes). The most comprehensive treatment of the calculation and associated issues can be found in the monograph by Chan (2008). (See Appendix K for an explanation of the acceptable calculation methods for this parameter, along with the tests that have to be performed to determine which calculation approach to use in any particular case. It may be helpful to the reader new to CA to refer to Appendix K before proceeding as it gives a precise definition of terms used in the following treatment.)

M.4 P_c and Probability Dilution

Despite the previous section's account of selecting a metric, the P_c is not a parameter without controversy. The main objection to the metric is a phenomenon called "probability dilution" in which low values of the P_c do not necessarily guarantee safety although they can offer support for refraining from a mitigation action. This section discusses the overall phenomenon of probability dilution and explains why it is viewed as a manageable issue when using the P_c in CA threshold development.

As stated previously, probabilistic CA metrics take on their probabilistic nature from the uncertainties in the state estimation data used to characterize the miss distance at TCA. It is not difficult to understand the alignment of inputs that creates a high P_c value: the calculated miss distance between the two objects is small, and the primary and secondary object covariances (and thus the uncertainty on the miss distance) are reasonably small also, so one easily envisions that of the entire set of possible actual miss distances, some notable portion will be smaller than the hard-body radius. The P_c , which is the portion of the possible miss distances smaller than the hard-body radius, will be relatively large. Imagine a similar situation for a small P_c : the nominal miss distance is large, and the covariances and miss-distance uncertainty are small, so the range of possible actual miss distances are nearly all rather large. One can conclude quite satisfactorily in such a case that there is little chance that the actual miss distance will be less than the hard-body radius.

But what if, for either of the previous cases, one or both covariances are extremely large? This would happen if, for at least one of the satellites, the position at TCA, due either to a poor orbit determination update or difficulties in accurate propagation, were extremely badly known. The predicted miss distance is by definition the expected value, but it is a poor proximate prediction of the miss distance because there is so much uncertainty in one or both of the satellites' position at TCA. The range of possible miss distances now becomes extremely large. Therefore, a relatively smaller portion of the possible miss-distance values will be smaller than the hard-body radius, and the P_c will take on a small value. In the earlier case in which the nominal miss distance was large but the covariances small, one could with confidence conclude, based on the quality of the data, that a collision is unlikely. In this latter case, the likelihood of collision is also small, but only because one has very little idea where one or both satellites actually will be at TCA, so one cannot conclude that a small miss distance is likely when so large a range of miss distances is possible. This latter case is called "probability dilution" (a term introduced by the first treatment of this subject in the literature (Alfano 2005b)) because the small value of the P_c is not achieved by the certainty of well-established state estimates but through "dilution" by relatively poor state estimates.



An analogy that is sometimes helpful is that of determining whether two cars parked in a large parking lot are parked next to each other. If one knew that one car was parked on the left side of the parking lot and the other on the right side of the parking lot, then one could conclude that it is unlikely that the two cars are parked next to each other: only if both cars park on the boundary between the left and right sides, and then only if they actually choose adjacent spaces, will the two cars be parked next to each other. This is a conclusion enabled by the definiteness of the data. Now, suppose that one knows nothing at all about where the two cars are parked in the parking lot; what is the likelihood of their being parked next to each other? It is also low, but in this case not because one knows that the cars have been placed in different parts of the lot but rather because, in a large parking lot, it is simply unlikely that any two cars will happen to be next to each other. It is a conclusion that follows not from what is known but rather from what is not known; one cannot conclude from the available data that the cars are probably far apart; but at the same time there is no evidence to indicate that they are close. If, therefore, one is required to hazard a guess, the large size of the parking lot makes the supposition that they are not adjacent reasonable.

So, what is the risk assessor to do in such a situation? One could embrace a very conservative position and require that, unless it can be demonstrated conclusively that the two satellites are virtually certain not to collide, a mitigation action is warranted. Some commentators have suggested this position as the natural one for advancing safety of flight (e.g., Balch et al. 2019). There are two sets of difficulties with such a position, one of which is practical and the other conceptual.

The practical difficulty centers on the number of increased mitigation actions and the magnitude of those actions if an extremely conservative risk analysis metric/approach were to be embraced operationally. An analysis by Hejduk et al. (2019) indicated that, for a protected asset in a 700 km circular orbit, the adoption of the conservative “ellipse overlap” technique, for which the two covariances’ overlap has to be limited to a very small value to certify that a collision is extremely unlikely, would increase the number of mitigation actions by a factor of anywhere from 6 to 10. Most O/Os believe they are already assuming a relatively safe posture in that they have set the thresholds on the risk-assessment metrics they currently use to require mitigation actions close to the mitigation action rate that they believe they can sustain without unacceptable mission disruption. While perhaps a relatively bounded increase could be countenanced, any expansion beyond that would be extremely difficult to accept. Increases by a factor of as much as an order of magnitude simply cannot be borne without resulting in major disruption, if not nullification, of the mission activities that most satellites are launched to perform. Furthermore, remediation actions based on minimizing covariance ellipsoid overlap rather than reducing the P_c value to an acceptable level will need to be large—much larger than the typical satellite maneuvers that are presently performed. Large maneuvers expend additional fuel and, if pursued more frequently, will also require subsequent response maneuvers to restore the orbital parameters that were violated by the initial mitigation action. An extremely conservative risk analysis approach such as the one described here is simply not seen as operationally viable for most missions.

In addition to these practical impediments, a conceptual dissonance is introduced in embracing an extremely conservative risk-assessment strategy given the much larger collision background



risk from untracked/uncataloged objects. As previously explained, the dynamics of the dilution region result in a reduction in severity and therefore a reduction in perceived necessity to act as less is known about the conjunction. As the amount of tracking information (or, similarly, the amount of certainty in prediction) for a dilution-region conjunction is reduced, the satellite covariance(s) increase in size, and the P_c decreases until ultimately it results in a value of 0 (to machine precision). This asymptotic value of nullity aligns nicely with the zero-value P_c that is imputed to conjunctions with untracked/uncataloged objects and about which the risk-assessment analyst can do nothing at all. As one progresses from knowing a lot, to a little, to nothing at all about a conjunction, the P_c moves conceptually from a higher value, to a lower value, to a zero value. For an extremely conservative approach such as that of minimizing covariance ellipsoid overlap, an opposing and inconsistent dynamic is observed: the less one knows about the conjunction, the larger the covariance ellipsoids are, and therefore larger and more frequent maneuvers are required in order to achieve minimum ellipsoid overlap to ensure safety until one moves from knowing a small amount about a conjunction to knowing nothing at all, and then—there is a sudden disjunction from taking large and invasive actions to doing nothing. Given that about 85% of the conjunctions that a satellite actually faces cannot be addressed at all because nothing is known about them, it does not make sense that, for the small number of conjunctions for which only poor data are available, large and disruptive mitigations are required.

Instead, the use of the P_c , with the dilution region dynamic explicitly recognized, is a more natural fit for the situation that is actually encountered. When outside the dilution region, a satisfactory amount of orbital data and reasonably small covariances exist. In this situation, the P_c accurately reflects a situation that is either worrisome or safe, and one can be confident in the conclusion in either case. If in the dilution region the P_c is high despite the poor data quality and large covariances, a worrisome situation has very much been identified, so appropriate mitigation action should be taken. If in the dilution region the P_c is low, one cannot, as has been explained previously, necessarily infer that the situation is safe; the two satellites might actually be poised for a dangerous close approach with this fact obscured by poor data quality. It is appropriate to recognize that this situation (state estimate with large covariance) is not dissimilar to that of a conjunction with an untracked/uncataloged object about which nothing is known (random state estimate and infinite covariance). It is thus acceptable to treat low- P_c dilution region cases as a small extension of the untracked/uncataloged object conjunctions and not pursue a mitigation action. Requiring conjunction remediation for low- P_c cases with poor data quality when such cases are truly very similar to the ~85% of conjunctions that are believed to exist but for which no mitigation is possible is not considered to be a prudent measure and thus runs counter to the CARA statement of purpose.

While not appropriate for broad implementation, conservative approaches such as ellipsoid overlap minimization are not without merit. In situations in which the collision risk may be pushed higher by other factors, such as a high-consequence conjunction due to the potential to produce large amounts of debris, it may be desirable to employ a conservative method when in the dilution region. In such a case, one might first analyze the situation by manipulating the covariance and determining the highest P_c value that could be realized with the miss distance (following Alfano 2005b); if that P_c is still below the threshold, then the conservative approach would be difficult to justify. But if this “maximum P_c ” is above the threshold, then the



conjunction could in fact be worrisome, and a conservative approach could be properly applied. Such an approach is not known to be operationally employed in any of the major CA risk assessment groups, but one could well understand its adoption.

M.5 Conclusion

It is emphasized that the above guidance describes and explains why P_c was chosen by NASA CARA as the risk-assessment metric to be used to determine the circumstances under which a mitigation action is required of missions to promote preservation of orbital corridors from debris pollution. O/Os are always welcomed to pursue additional mitigation actions in a manner that increases the conservatism of their safety profile. For the case of low- P_c conjunctions in the dilution region, missions may elect to pursue mitigation actions. If they wish, they may even use extremely conservative methods such as ellipsoid overlap minimization to size these mitigation actions. This guidance indicates that mitigation actions are required only when the P_c exceeds the appropriate mitigation threshold. Missions are always at leisure to pursue additional such actions, either episodically or as part of a standing strategy, if they wish a more risk-adverse safety posture.



Appendix N. Pc Calculation Approaches

N.1 Introduction

This appendix discusses the selection of the Probability of Collision (Pc) as the risk assessment parameter to use for CA requirements compliance. While this parameter is widely used in the CA industry, issues related to its calculation exist that merit extended discussion. Analytical techniques to calculate the Pc are well established and computationally efficient, but they make assumptions that restrict their use and make the calculation vulnerable to error. Numerical techniques also exist. Predictably, these make fewer assumptions and are more widely applicable, but they are much more computationally demanding. There are also issues related to the regularization and interpretation of the input data to the Pc calculation, some of which are resolved by techniques that are now becoming standard practices. In calculating the Pc, it is necessary to examine and prepare the input datasets carefully and then to choose the calculation approach that is appropriate to the situation. The purpose of this appendix is to amplify the Pc calculation-related recommendations by providing an extended technical explanation of the data preparation and Pc calculation issues so that implementers and users of these calculations can proceed with a better understanding of the different options and resultant fidelities of Pc calculation.

To accomplish this, this appendix will treat the following technical areas:

- A step-by-step description of the two-dimensional Pc calculation, which is the most established and widely used analytical Pc computation technique, and its enabling simplifications and assumptions;
- Examination and repair/expansion of input data to the calculation, focusing mostly on the covariances for the two objects in conjunction;
- Discussion of the use of Monte Carlo techniques, which is the chief numerical method used for Pc computation;
- A test to determine whether an analytical or numerical technique should be used for a particular conjunction;
- Approaches to choosing for the Pc calculation the hard-body radius, which gives a statement of the combined sizes of the primary and secondary objects; and
- Discussion of alternative analytic Pc calculation methods, specifically the three-dimensional Pc and its promise as the most desirable of the analytical calculation possibilities.

N.2 Two-Dimensional Pc Calculation

The two-dimensional Pc calculation, which is by far the most widely used Pc calculation approach in the CA industry, was developed for the Space Shuttle Program and first appeared in the literature in 1992 (Foster and Estes). There have been a number of important treatments since that time (e.g., Akella and Alfriend (2000), Patera (2001), Alfano (2005a), Chan (2008)), but all rely on the same basic methodology: applying reasonable assumptions to reduce the

dimensionality of the problem to make it analytically tractable. While there are variations in the particulars, the assumptions and concept of dimensionality reduction are shared by all these approaches, so the procedure can be walked through generically.

One begins with the states and covariances for the primary and secondary objects' TCA. An important set of questions should be addressed concerning whether these data, especially the covariances, are truly representative of the expected states and uncertainties at TCA or whether the propagation process has distorted them. These questions will be addressed in the next section once a more focused context for them has been established through the present discussion of the general procedure.

The first step is to recognize that this calculation centers on the relative positions and uncertainties of the two objects, so moving to a framework that views the problem this way is helpful. Subtraction of the two objects' positions produces a relative position vector, the magnitude of which is the miss distance between the two objects (at TCA). Similarly, because interest is in the relative rather than absolute position uncertainty, it is possible to combine the two objects' covariances to form a joint (relative) covariance and allocate that, if desired, to one "end" of the relative position vector (by convention the end for the secondary object), as shown in Figure N-1. There are, of course, questions regarding whether the two covariances are independent and can be combined by simple addition; this question is addressed in a later section, but it is typically acceptable to presume independence and combine the covariances by addition.




Figure N-1 Relative Positions and Uncertainties of Two Objects

The second step relates to modeling the combined sizes of the primary and secondary objects. The general approach is to place a circumscribing sphere about the primary (whose size is well known by the O/O because it is their satellite) and then do the same thing for the secondary but via an estimation technique as the secondary is usually a debris object for which there is no definitive size information. If the satellite sizes are modeled by these two spheres as the two satellites pass through their TCA and if the two spheres touch each other, a potential collision has been identified. Again, keeping in mind that a relative framework is useful here, the two objects' size spheres can be combined into one supersphere (sometimes called the "collision sphere") and placed at one end, by convention the primary object end, of the relative miss vector, as shown in Figure N-2. If the miss vector should shrink to be smaller than the radius of the supersphere (called the hard-body radius), that would be the equivalent of the two original spheres encroaching on each other at TCA.




Figure N-2 Combining the Sizes of the Primary and Secondary Objects

The third step is to envision the situation at TCA in which all the uncertainty is assigned to the secondary object's end of the relative miss vector held in a fixed position in the mind. The primary object end of the relative miss vector is moving along through TCA and bringing with it the sphere representing both objects' combined size. Even though the satellites follow curved trajectories and the covariance evolves and changes at each instant, if the encounter is presumed to take place extremely quickly—and this is in most conjunctions a good assumption because the satellites' relative velocity usually exceeds 10 km/sec—then two assumptions can be made: the trajectories are essentially rectilinear during the encounter period, and the covariances (and thus the joint covariance) can be considered static. This means that the encounter can be envisioned as in Figure N-3:




Figure N-3 Geometry at TCA

The passing of the primary object by the secondary can be seen as following a “soda straw” straight trajectory whose cylindrical radius is the same as that for the hard-body radius and whose placement is one miss distance away from the position of the secondary at TCA. Since the joint covariance shown as the ellipsoid above represents the uncertainty of one “end” of the miss-distance vector (as shown by the central dot in the above diagram), this dot can be presumed to be potentially in any place within the ellipsoid, meaning that in some portion of those realizations, it will fall within the soda straw. When this is the case, a collision is presumed. The task is to determine the likelihood that the dot will, in the actual realization of this conjunction, fall within the cylindrical path (“soda straw”) swept out by the motion of the

primary object. This probability can be decomposed to be rendered as the product of the individual probabilities that each component of the secondary object position (the dot) will fall within the soda straw pathway, i.e., if an orthogonal x-y-z coordinate system is presumed, this overall probability can be generated as the product of the probability that the x-component of the dot’s position will fall within the straw, the y-component of the dot’s position will fall within the straw, and the z-component of the dot’s position will fall within the straw. If this coordinate system is aligned so that one of the axes (perhaps the x axis) aligns with the direction of the straw, because one is assuming rectilinear motion, the soda straw can be presumed to be unbending and infinite in length, and as such, it will contain all of the x-component probability of the dot’s falling within the straw. As such, the x-component probability in this arrangement will be unity and will be what is called “marginalized out,” meaning that the overall probability can be fully represented by the probability remaining with the y- and z-components. The entire situation can thus be reduced from three to two dimensions and can be analyzed as a phenomenon that occurs in the y-z plane, this plane being defined as orthogonal to the soda-straw direction, which is the direction of the relative velocity vector. This procedure defines the “conjunction plane,” a representation of which is given in Figure N-4:

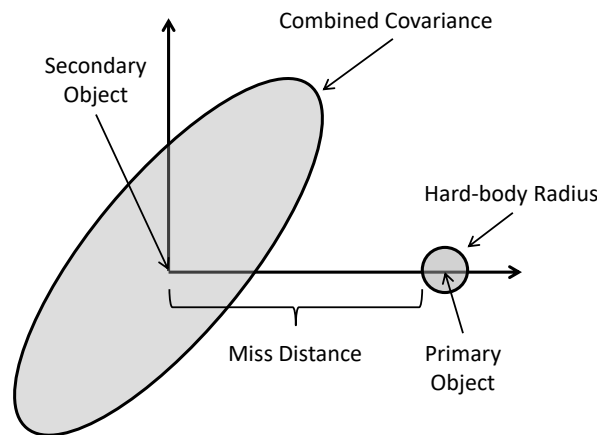


Figure N-4 Representation of the Conjunction Plane

By convention, in this plane the secondary object location is placed at the origin and the local coordinate system configured so that the miss vector lies along the horizontal axis. The “soda straw” is coming out of the page and represented as the hard-body-radius circle. With this planar reduction, the P_c is now the frequency with which the miss-distance vector will fall within the hard-body-radius circle; this is equivalent to the amount of joint covariance probability density (reduced here to an ellipse) falling within that circle. Since the projected joint covariance represents a bivariate Gaussian distribution, the amount of its probability density falling within the hard-body-radius circle is given by the following integral:

$$P_c = \frac{1}{\sqrt{(2\pi)^2 |C^*|}} \iint_A \exp\left(-\frac{1}{2} \vec{r}^T C^{*-1} \vec{r}\right) dYdZ \quad (N-1)$$



in which r is the miss vector, C is the covariance, and A is the hard-body-radius circular area. There are different approaches to evaluating this integral. Foster and Estes (1992) applied a quadrature technique; this works well, and MATLAB®'s¹⁰ *quad2d* integrator is quite equal to the task, although perhaps not the most computationally efficient of all possible choices. Chan (2008) uses equivalent-area transforms to produce a single-dimensional integral, which has a series solution; this approach is more computationally efficient but occasionally produces errors and is not commonly used in the industry. Alfano (2005a) formulated the computation in terms of error functions and employed a midpoint-rule integrator. Elrod (2019) expanded this approach to frame the calculation in terms of complementary error functions (to improve accuracy for conjunctions with small P_c values) and employed Chebyshev quadrature with nodal symmetry. The Foster and Elrod approaches (the former being the established standard and the latter being extremely fast, especially with a vectorized MATLAB implementation) are both included in the P_c Computation Software Development Kit (SDK) in the NASA CARA software repository. (See Section 7, Contact Information in this document for the specific URL.)

It is perhaps helpful at this point to review the assumptions that this two-dimensional P_c calculation technique employs as alternative techniques are needed when these assumptions do not inhere:

- The two objects' uncertainty distributions are statistically independent so that the joint covariance can be obtained by simple addition of the two covariances. This assumption is largely true but can break down for objects that share global atmospheric density forecast error in a manner that influences the conjunction. This issue will be discussed in the next section of this appendix.
- The conjunction circumstances are such that it is reasonable to presume rectilinear motion and static covariances during the encounter. These conditions inhere for most conjunctions. The objects' relative velocity at TCA is in some places used as an indication of the reasonability of these assumptions, but this parameter alone is not sufficient to identify situations in which the two-dimensional P_c will miscarry.
- The conjunction presents a well-defined single event so that the collision likelihood can be ascertained by examining that single instance. Objects that stay in close proximity for extended periods do not have a well-defined TCA and accumulate risk throughout these long encounters. A different approach is required for their evaluation, which is discussed in the section that addresses Monte Carlo techniques.
- The state errors at TCA follow Gaussian distributions, which allows evaluation techniques that work with known distributions to be used. While it can often be shown that satellite covariances represented in Cartesian coordinates do not represent the actual state uncertainties, there are both approaches to address this divergence and studies

¹⁰ MATLAB is a registered trademark of The MathWorks, Inc.



documenting its relative insignificance for CA operations. These are also discussed in the section that addresses Monte Carlo techniques.

N.3 Three-Dimensional Analytic P_c Calculations

Conjunctions that do not satisfy the assumptions discussed in the previous section must be addressed with a different methodology, and in response to this need, several authors have formulated semi-analytical approaches relaxing some or all of these assumptions. In 2012, Coppola (2012) proposed a method for single encounters designed to account for non-linear orbital motion and velocity uncertainties, resulting in an approximation for the probability rate, $\dot{P}_c(t)$, calculated using two-dimensional numerical integration over the surface of a unit sphere. When combined with a numerical one-dimensional time integration, this yields a “three-dimensional P_c ” approximation. Chan (2015) contested Coppola’s formulation, arguing that a proper approach must employ a set of random variables associated with a time-invariant Probability Density Function (PDF). NASA CARA implemented the three-dimensional P_c method in software (Hall et al. 2017a) and subsequently discovered that, for some conjunctions, it can produce P_c estimates that differ significantly from high-fidelity Monte Carlo P_c calculations, even though all of the required three-dimensional P_c assumptions are satisfied.

Shelton and Junkins (2019) provided a key insight into why the original three-dimensional P_c approximation fails in certain situations. Their analysis indicates that accurate P_c approximations require that the state uncertainty PDFs of the two satellites be estimated accurately *in the volume of space where they overlap the most*. The original three-dimensional P_c formulation did not incorporate this concept, but NASA CARA has recently reformulated the method to do so explicitly. Specifically, for single encounters, the reformulated approach approximates the curvilinear motion of each satellite using a first-order Taylor series expansion not centered on the mean orbital state but centered instead on a state that coincides with the maximum overlap of the PDFs of the two satellites. The analysis demonstrates that these expansion-center states can be determined using an iterative calculation. The resulting “three-dimensional N_c ” (for “number” of collisions) approximation can be expressed as a three-dimensional numerical integral with the same form as the original three-dimensional P_c method but with an integrand function that contains different specific terms.

Figure N-5 shows a comparison of collision probabilities for 63,603 temporally isolated conjunctions extracted from the NASA CARA database for the period 2017-05-01 and 2019-08-15 for events with $2D P_c > 10^{-7}$. The vertical axes on both panels plot high-fidelity Monte Carlo (MC) estimates for the collision probability (specifically, from-TCA TBMC- P_c estimates to be described in the next section), with error bars that show 95% confidence intervals estimated using the Clopper-Pearson method (1934). Several of the conjunctions had zero hits registered in the MC simulations, which are represented in Figure N-5 using downward-pointing triangles and a single-sided error bar. The horizontal axes plot the corresponding semi-analytical approximations: two-dimensional P_c on the left graph and three-dimensional N_c on the right graph. The colored points indicate the results of a binomial proportion statistical test that evaluates the agreement between the estimates. Specifically, black points in Figure N-5 indicate P_c values that agree reasonably well as they do not violate a null-hypothesis that the two are equal at a p -value $\leq 10^{-3}$ significance level. However, those highlighted in yellow do violate the



hypothesis at this significance level, and those in red at a more stringent level of p -value $\leq 10^{-6}$. Overall, the two-dimensional P_c comparison contains 254 yellow and 436 red points, which both significantly exceed the number of disagreements expected from purely statistical variations, even though together they represent a small fraction ($\sim 1\%$) of the original conjunctions. The three-dimensional N_c approximations clearly match the MC P_c estimates better, producing only 65 yellow and zero red points. The comparison indicates that, for temporally isolated conjunctions, the three-dimensional N_c method produces more accurate estimates on average than the two-dimensional P_c method (although both are much more accurate than the original implementation of the three-dimensional P_c method, which is not plotted here).

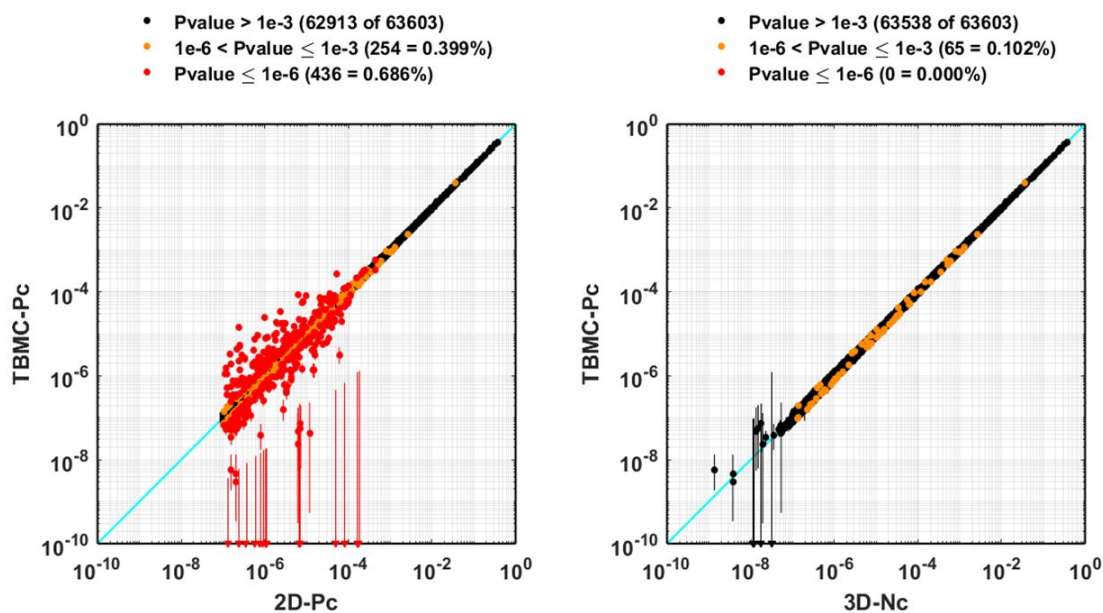


Figure N-5 Comparison of Monte Carlo Collision Probabilities with the Two-Dimensional P_c Method (left) and the Three-Dimensional N_c Method (right) for a Large Set of CARA Conjunctions

At the time of this writing, the three-dimensional N_c method is still sustaining extended check-out, but it is expected to become the principal method of P_c calculation for NASA CARA. As soon as it has sustained public release, the source code and accompanying journal article outlining the development of the subtending theory will be posted in the NASA CARA software repository for download. (See Section 7, Contact Information in this document for the specific URL.)

An additional feature of the three-dimensional N_c approach is its ability to address explicitly the amalgamated risk of repeating conjunctions. While most conjunctions are substantially temporally isolated, there are two conjunction types that exhibit different behaviors: conjunctors in orbits that produce synodic alignment and will generate a conjunction every revolution (or multiple of a revolution) and conjunctions that linger close to each other for long periods, generating multiple close approaches of a very similar miss distance. If each of these events / close approaches is treated as a separate event, situations can arise in which each event in the series falls below a mitigation threshold, but the combined risk of all of the events would exceed



that threshold. The three-dimensional N_c methodology takes account of this phenomenon and gives an upper-limit bound on the combined risk, an example of which is given in Figure N-6. Four identifiable conjunction events take place over about a five-hour period, each of which produces a P_c value in the upper yellow region (between $1E-05$ and close to $1E-04$); these data are shown in the bottom plot. The solid line in the top plot shows an upper limit on the combined collision likelihood from all of the individual events, taken serially; one observes that after the third event, the cumulative likelihood exceeds $1E-04$ —a value frequently selected as a mitigation threshold. So, while these events would not necessarily prompt a mitigation action if examined individually, when considered collectively they do appear to represent a situation of sufficiently high collision likelihood to warrant mitigation. In such a case, it is advisable to run a Monte Carlo investigation (discussed in a subsequent section of this appendix) to verify that the upper-limit value generated by the three-dimensional N_c is in fact representative of the actual cumulative risk. The utility of the analytic calculation, however, should be clear: if the upper-bound calculation for repeating events is below the mitigation threshold, then there is no need to marshal computationally intensive methods, such as Monte Carlo analyses, for evaluating such repeating events for it has been shown that even in the worst case, the likelihood does not merit a mitigation action.

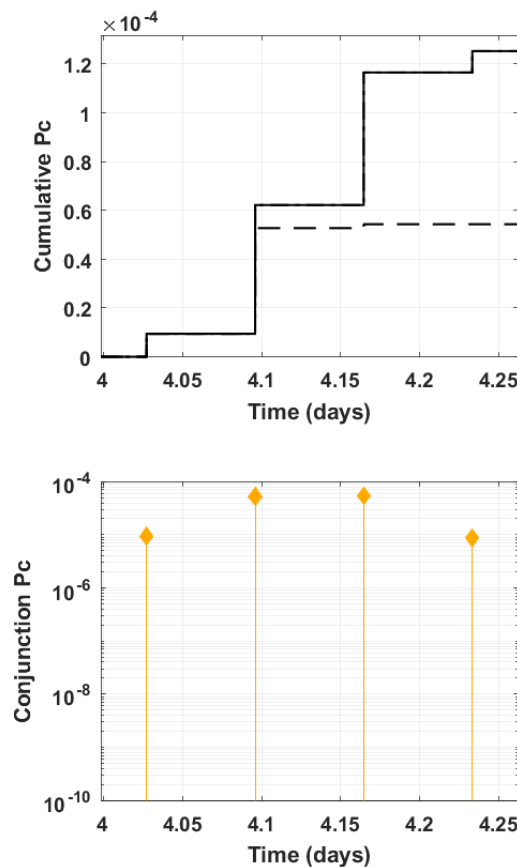


Figure N-6 Cumulative Three-Dimensional N_c Risk Over Four Repeating Conjunctions



N.4 Input Covariance Data Considerations

Most calculations are only as good as the input data that drive them, and the Pc calculation is no exception. Appendix J discusses orbit determination quality and propagation issues for individual objects and addresses the question of the circumstances under which the state estimate and covariance might be considered sufficiently poor so as not to constitute a basis for CA mitigation actions. The purpose of this section is to address the routine improvement and basic error checking extended to covariances as part of the Pc calculation. These activities fall into three basic groups: correction for known problems in propagation, covariance formation and stability issues, and correlation between primary and secondary covariances. Each of these topics will be treated in turn.

N.5 Propagation Error Compensation

Historically, accuracy analysis of state estimate and uncertainty estimates has focused on products that emerge directly from the orbit determination fit. These direct orbit-determination products include best-estimate states and covariances of the satellite at an epoch that usually coincides with the acquisition time of the last tracking metric observation incorporated into the analysis. While these direct orbit-determination products are of course foundational, it is important to remember that most SSA activities, and CA in particular, are not usually conducted with these direct products but rather with predictions propagated from the orbit-determination epoch solution, often over a non-trivial duration that spans many orbital revolutions into the future. The batch orbit-determination analysis method used by the DOD produces a formation covariance that represents the expected at-epoch state uncertainty based on the number, quantity, and temporal spacing of the incorporated metric observations; when the state is propagated forward, a parallel process can also be used to propagate the covariance forward in time. The same dynamical models used for the orbit-determination analysis as well as the state propagation itself are used to perform this covariance propagation, although in a linearized way. This means that the propagated covariance will be sized (mostly) appropriately for both the propagation duration and the final prediction point in the orbit.

Despite the use of appropriate models to propagate the covariance forward in time, a number of additional sources of error manifest themselves during the propagation interval yet are not part of the dynamical model used during the fit; these errors are therefore neither incorporated into the orbit-determination-produced covariance nor added as part of the regular propagation process. Because of the prevalence of such outside-of-model errors, techniques have been developed to account for them, the most familiar of which is the addition of process noise during propagation. Originally developed to account for acceleration errors that, due to model inadequacy, were to some degree known, this method propagates a noise matrix alongside the propagated covariance and combines both matrices as part of the overall process. The result is a covariance that is larger than it would have been otherwise to account for these (characterized) errors in the force model(s). A second approach, which is the one used by DOD in the propagation of their “e” products, applies parameters to the covariance before propagation to guide the propagation process in producing a more realistic result. Because this is the approach reflected in the CDM covariances that CA practitioners receive, it will be described here in some detail.



Orbit determination makes a distinction between “solved-for” parameters that are actually estimated during an orbit-determination activity, and “consider” parameters that are not “solved for” but represent *a priori* information that is “considered” as part of the orbit-determination process. In the present case, the use of the term “consider parameter” is somewhat non-nominal in that it is referring not to additions or alterations made during the fit but to modifications to the fit-produced covariance so that when it is propagated, it may be a more realistic representation of the expected state errors. For DOD covariances, two different consider parameters are applied to compensate for two distinct expected errors during propagation: atmospheric density forecast error and satellite frontal area uncertainty.

Atmospheric drag is a significant force, affecting satellite orbits with perigee heights less than 1000 km, and the calculation of the expected atmospheric drag at any particular moment requires an estimate of the neutral atmospheric density that that satellite will encounter at that moment. Because the atmospheric models that generate this estimate are driven by space weather indices, the ability to predict these indices accurately affects the fidelity of the predicted atmospheric density and thus the atmospheric drag. Unfortunately, it is difficult to predict future space weather indices well, primarily because they are affected by activities on the part of the sun’s surface that has not yet rotated into view from the Earth. This particular issue was studied with the Jacchia-Bowman High Accuracy Satellite Drag Model (HASDM) 2009, which is the atmospheric density model used by DOD, by comparing predicted densities (using predicted space weather indices) to actual densities and producing polynomial fits of the relative density error as a function of satellite perigee height and the A_p (major magnetic storms list) and Dst (disturbance storm-time) geomagnetic indices. Figure N-7 shows the behavior of these polynomial fits divided into four different classes of A_p/Dst activity; y-axis values are omitted here to allow full releasability of the figure:




Figure N-7 Behavior of Relative Density Error by Perigee Height and Solar Activity



These fits produce a variance term that can be added to the ballistic coefficient variance in the covariance: because in the drag equation the ballistic coefficient and the atmospheric density estimate are multiplicatively coupled, changing one of these parameters has the same effect as changing the other. When the covariance is propagated, this augmented term will appropriately expand the other covariance terms.

The amount of drag acceleration a satellite encounters is also governed by the frontal area that the satellite presents to the atmosphere; this makes intuitive sense (amount of resistance is a function of area presented to the resisting fluid) and is reflected in the ballistic coefficient (B) equation

$$B = C_D \frac{A}{M} \quad (\text{N-2})$$

in which C_D is the (dimensionless) drag coefficient, which indicates the degree to which the satellite surface is susceptible to drag resistance; M is the satellite mass; and A is the satellite frontal area. As A increases, B increases as well, increasing the overall drag acceleration value. Stabilized satellites should manifest a stable B term, but rotating satellites, because their frontal area term is continuously changing, can exhibit a range of B terms. Three outcomes are possible depending on the rapidity of the rotation: the effect can be washed out during the fit (because the rotation is so rapid that an average value is quite representative), not be relevant during the fit (because the rotation is so slow that it does not affect near-term look-back and look-forward), or be such that the rotation does affect results fit-to-fit; it is this last case for which compensation is helpful. A history of regularized B histories for individual satellites is examined and a relative error and variance calculated, and this variance is added to the drag variance in the covariance as a corrective term whose influence will be realized in propagation.

There is some additional subtlety regarding the exact way these consider parameters are applied. A typical propagation consists of two conceptual stages: the first stage is the propagation forward from the epoch time of the orbit determination to the present time, which can make use of measured and thus highly reliable space weather indices; and the second stage is from the present time to the desired final propagation time, which has to use predicted space weather indices and the errors that these introduce. The two consider parameters are thus applied at different times. Because satellite rotation and its resultant uncertainty will occur for the entire interval from epoch time to the propagation end point, that consider parameter is applied at epoch. Atmospheric density forecast error, however, is encountered only forward in time from the present moment, so it is added only for that portion of the propagation. Figure N-8 outlines the two-phase application of these consider parameters:




Figure N-8 Two-phase Application of Consider Parameters

If the CDMs generated by the DOD are used for CA applications, the good news is that all of the consider parameter activities described above are already performed—the propagated covariances that the CDM contains have had these two consider parameters applied during the covariance propagation executed at the DOD work center. If one is working with epoch state estimates, which is sometimes necessary when performing P_c calculations with Monte Carlo techniques, then manual application of the consider parameters may be necessary. This issue is discussed at greater length in the section that addresses Monte Carlo P_c calculations.

N.6 Defective Covariances

There are several ways in which a covariance contained in a CDM can be defective.

- **Null Covariances.** All-zero, or null, covariances are sometimes observed, usually arising from a CA screening for which the O/O-provided ephemeris does not contain covariance data. In such a case, it is possible to compute the P_c either using only the one covariance that the CDM message contains or by applying a special technique that determines the maximum P_c possible presuming that the null covariance could take on any possible value (developed and described in Frisbee 2015).
- **Default Covariances.** Default covariances are diagonal covariances that contain a value of ten earth radii squared for each of the three position variances. The presence of this covariance indicates that a precision, special-perturbation solution for the object was not possible; the state estimate provided arose from a general-perturbation solution, and an orbit-determination-produced covariance was not available. Such a result is not a precision solution and does not constitute a basis for CA mitigation actions.
- **Non-Positive-Semidefinite Covariances.** Another defective covariance type found in CDMs, now quite rare due to improvements to the DOD operational system, is a covariance that fails to be positive semidefinite. A positive semidefinite matrix is one that contains no negative eigenvalues. Because the covariance represents a hyperellipsoid of actual state error information, it must have a set of eigenvalues all greater than or equal to zero for error information to consist of real rather than imaginary quantities. The orbit-determination mechanism that generates the covariance should always produce at least a positive semidefinite matrix for the linear algebra function involves the product of a square matrix and its transpose, and one can prove that this procedure always produces a positive semidefinite result. Due to either numerical precision limits or interpolation, the



provided matrix is sometimes not positive semidefinite. If the 2×2 projection of the covariance matrix into the conjunction plane is not positive semidefinite, the two-dimensional Pc calculation is not possible. If the full 6×6 or 8×8 matrix is not positive semidefinite, then Monte Carlo sampling on the entire matrix is not possible either.

As such, some attention must be paid to this issue of positive semidefinite matrix conditioning. A recent paper on this subject (Hall 2017b) examined the issue in some detail and compared different matrix repair algorithms to minimally adjust the covariance to make it positive semidefinite compliant; it found that most repair approaches yield equivalent answers in terms of the resultant calculated Pc. An “eigenvalue clipping” procedure was developed in which any negative eigenvalues (which are almost always small) are set to a small positive or zero value, as required.

The CARA operational implementation of this method proceeds parsimoniously, namely by directing such repair only to the level needed to perform a calculation. For example, a covariance used for the two-dimensional Pc calculation would neither be tested for positive semidefinite compliance in its full 8×8 form nor its position-only 3×3 form; instead, the 2×2 projection of the joint covariance into the conjunction plane would be tested and repaired only when necessary to enable the two-dimensional Pc calculation. To do otherwise is to make repairs and potentially alter the result when this is not strictly necessary.

The Pc Computation CARA SDK includes the source code for identifying the need for and making the covariance matrix repairs described above; it is available in the NASA CARA software repository. (See Section 7, Contact Information in this document for the specific URL.)

N.7 Covariance Correlation

For nearly all the broader CA conduct during the past decade, practitioners operated with the presumption that the primary and secondary objects’ covariances could be considered uncorrelated. Not only was this the “right” answer in that it greatly simplified the Pc computation because the joint covariance could be formed by simple addition of the two covariances, but there was also a good intuitive justification for the presumption. Because the focus had been on the two objects’ orbit-determination fits, which are based on separate sets of observations, there was no expectation that there would exist any significant correlation between the two objects’ covariances. The principal source of potentially correlated error was presumed to be uncharacterized but correlated errors in space sensor observations used by both primary and secondary objects. Because most primaries receive many observations from many different sensors, it was seen as unlikely that this particular source would introduce much correlation. Correlation between covariances was thus expected to be small, and CA operators proceeded as though it were.

With the initiative several years ago to include outside-of-fit prediction error characterization into DOD satellite covariances (see the above section on Propagation Error Compensation), the issue of covariance cross-correlation began to be rethought. The principal outside-of-fit prediction error is global atmospheric density forecast error due to inadequate space weather index prediction. Because this is a global error, it is likely to be shared among large classes of objects, some of which might constitute both the primary and secondary objects in a conjunction.



As discussed previously, this global density forecast error has been parameterized by satellite perigee height and predicted geomagnetic index, so the degree of such error, both identified separately and injected into each DOD-furnished covariance by means of a consider parameter, is known for each satellite. It is possible to determine the degree of shared error from this source and account for it when forming the joint covariance.

A study by Casali et al. (2018) provides a full development of this theory and presents results for an evaluation set of conjunction data. Essentially, one has to determine the global density forecast error relevant to each satellite and the degree to which the associated drag error induced by this density forecast error will manifest itself in position error relevant to the particular conjunction. The governing equation is the following:

$$P_m = P_s + P_p - \sigma_{s/g} \sigma_{p/g} G_s G_p^T - \sigma_{s/g} \sigma_{p/g} G_p G_s^T \quad (N-3)$$

in which P_m is the decorrelated joint covariance at TCA, P_s is the secondary covariance at TCA, P_p is the primary covariance at TCA, $\sigma_{s/g}$ is the density forecast error germane to the secondary satellite, $\sigma_{p/g}$ is the density forecast error germane to the primary satellite, G_s is the sensitivity vector mapping drag acceleration error to secondary satellite position error at TCA, and G_p is the sensitivity vector mapping drag acceleration error to primary satellite position error at TCA. One could wonder how the CA practitioner would come upon all of the needed data to effect the proposed compensation. A recent enhancement to the DOD operational system has placed all of this information in the CDM itself, allowing the direct calculation of the decorrelated joint covariance. The CARA Pc Calculation SDK, available on the NASA CARA software repository, contains both a math specification outlining this calculation and source code to perform it. (See Section 7, Contact Information in this document for the URL.)

A heuristic probing of the situation reveals that, conjunction by conjunction, different levels of Pc changes are possible due to cross-correlation remediation. Orbits that are insensitive to atmospheric drag are little affected by density forecast error and will have Pc estimates that change little. Head-on conjunctions are expected to be left mostly unaffected as well, for the components of the error that govern the Pc are not subject to density forecast error perturbation. Crossing events are perhaps the most susceptible to cross-correlation effects, especially if the drag level of both satellites is similar.

The plot in Figure N-9 profiles 250 conjunctions in which the primary and secondary satellites are of non-negligible drag (i.e., Energy Dissipation Rate (EDR) values greater than 0.0006 watts/kg; see Hejduk 2008 for an explanation of energy dissipation rate) and plots the ratio of the Pc calculated with the decorrelated joint covariance to that of the Pc calculated with the unmodified joint covariance. One can see that for somewhat more than half of the events, the ratio hovers near unity, meaning that the conjunction is little affected by the compensation. For about one-third of the cases, the decrease in Pc is notable, in many instances more than an order of magnitude. For the remaining cases, there is an increase in Pc from a factor of 1.5 to 5.




Figure N-9 Profiles of 250 Conjunctions with Primary and Secondary Satellites of Non-negligible Drag¹¹

While the effects are not revolutionary, they are substantial enough to justify the integration of this additional consideration into the Pc computation, especially because it is a straightforward calculation from data provided directly in the CDM.

N.8 Monte Carlo Pc Calculation Techniques

Analytic approaches to Pc calculation, of which the two-dimensional Pc is the most widely used, are extremely computationally efficient, but they require certain enabling assumptions. Numerical techniques can often allow such assumptions to be relaxed but with a substantial computational penalty. Numerical approaches are not typically employed as the first line of computation methods but rather are marshalled for those cases that are suspected of violating the enabling assumptions of the analytic methods. Diagnostic techniques to identify cases for which the analytic assumptions that subtend analytic Pc calculations exist and will be addressed in a subsequent section. To discuss those most meaningfully, it is preferable first to present the chief numerical technique used to calculate Pc, which is a Monte Carlo method. There are two strains of this method that are regularly employed:

- “Monte Carlo from epoch,” in which the epoch mean states and covariances are used to generate sampled states, and potentially long (i.e., multiday) high-fidelity propagations are required for each Monte Carlo trial.
- “Monte Carlo from TCA,” in which the mean states and covariances predicted at TCA are used to generate sampled states, and only extremely short propagations to each trial’s new TCA are required.

Each of these two approaches will be discussed in turn.

¹¹ From Casali et al. 2018.



N.9 Monte Carlo from Epoch

The principal appeal of calculating the P_c using the Monte Carlo from epoch approach is that it requires almost no simplifying and restrictive assumptions, making it as close to a “gold standard” for P_c estimation as can be devised. The input information includes the states and covariances for the primary and secondary objects at their respective orbit-determination epoch times, the combined hard-body radius of the two objects, and an ensemble of environmental datasets required for the high-fidelity propagations (such as predicted space weather indices and atmospheric density debiasing data; see Hall et al. 2018). The most elaborate instantiation of this technique uses the full eight-dimensional orbital state vectors (each containing six coordinate or element variables, plus drag and solar radiation pressure variables) along with the associated 8×8 covariance matrices for the primary and secondary satellites.

For each Monte Carlo trial, a state perturbation is obtained by performing a random draw from the distribution for each of the eight variables using the covariance matrix to generate the associated Gaussian distribution for each variable. These perturbations are then used to alter the initial states by adding each to the appropriate epoch mean state estimates. To do this for the primary satellite, for example, one would generate random perturbations for the eight variables representing the primary satellite’s state based on their distribution as defined by the covariance, and then create a new, sampled state vector by adding the (signed) perturbations to the mean epoch state vector. This same procedure would also be performed for the secondary satellite. This sampling process generates epoch states for both the primary and secondary objects that represent statistically reasonable alternatives for their actual states. These two sampled epoch states are then propagated forward, a TCA identified, and a check performed to determine whether the miss distance at TCA is smaller than the hard-body radius; if it is, the trial resulted in a simulated collision and a “hit” is registered; if not, it is considered a “miss.” This sampling/propagation/ hit-registration procedure is then repeated for a large number of Monte Carlo trials, and the final number of hits divided by the total number of trials constitutes the Monte Carlo P_c estimate. There are algorithms that can be applied to determine the confidence interval on the Monte Carlo P_c after a given number of trials.

This procedure seems straightforward enough, and in many respects it is. But there are subtleties that require attention, especially if the technique is deployed for LEO conjunctions:

- For the result to be valid, the same force models and force model settings must be used for the Monte Carlo propagations as were used when generating the original orbit-determination solution. While it often is not difficult to apply the same general force model settings, there does need to be overall compatibility between the orbit-determination engine and the Monte Carlo propagator, such as the same set of geopotential Legendre polynomial coefficients (not just the same geopotential order) and, most critically for LEO, the same atmospheric density model. While it may be possible to deviate somewhat from this level of compatibility and still obtain reasonably accurate outcomes, the “gold standard” propriety of the result is lost.
- Correlation between the primary and secondary covariances, as described in the previous section, should be considered. Even though the covariances are not propagated because



there will be correlation in propagation error, this correlation can be modeled during the random draw process by forcing correlation in the primary and secondary objects' drag perturbations.

- Monte Carlo from epoch often requires extremely long computation run times. The run time is a function of the actual P_c since this will determine how often hits are likely to be observed and the number of trials required to obtain a result with a desired confidence level. Without marshalling extensive high-performance computing, the P_c levels that can be explored with this method have to remain relatively large. For example, to validate a $1E-07$ P_c event with the TCA five days from the epoch time with 95% confidence requires two years of execution time on a 20-CPU, reasonably modern server (Hall et al. 2018). It is typically necessary to reserve this calculation method for larger P_c events only, which means that one must trust analytic methods to identify a candidate subset of conjunctions for follow-up Monte Carlo analysis.
- High-fidelity orbital state propagations require the most processing during typical Monte Carlo from epoch computations, so when applying this approach, there is a temptation to “reuse” propagations to gain computational efficiency. Suppose that ten perturbations were performed for the primary satellite and ten propagated ephemerides were generated, and the same was done for the secondary object as well. Ephemeris #1 for the primary could be compared to ephemeris #1 for the secondary to determine whether a hit occurred, ephemeris #2 for the primary to ephemeris #2 for the secondary, ephemeris #3 for the primary to ephemeris #3 for the secondary, etc., and have as a result ten comparisons/trials. To go further, one could compare ephemeris #1 for the primary to ephemeris #2 for the secondary, ephemeris #1 for the primary to ephemeris #3 for the secondary, etc., and realize 100 comparisons/trials from merely 20 total propagations. Such a procedure certainly seems advantageous, given that processing time is the limiting factor to the deployment of Monte Carlo from epoch, and following such a procedure will produce a result that converges on the true P_c . The drawback is that the reuse of samples violates the conditions for the proper application of the formulae to calculate confidence intervals. Monte Carlo results without associated reliable confidence intervals are not operationally useful because it is never known how close one is to the true P_c value. Schilling et al. (2016) discuss this issue and confirm it to be a problem, and they recommend some estimation techniques that allow (large) confidence intervals to be constructed around Monte Carlo products that avail themselves of sample reuse. The CARA implementation has avoided any sample reuse to ensure that the “gold standard” status of the results not be in question and to produce more narrow confidence intervals.

Due to all the above considerations, but especially the run-time requirements, Monte Carlo from epoch is usually reserved only for those cases that require it for accurate P_c estimation.

A study effort discussed in greater detail in the next section determined that Monte Carlo from epoch appears to be needed only when the two objects stay in such close proximity that they experience a sequence of multiple, comparable-risk close approaches during a multiday risk assessment projection interval. For closely spaced co-orbiting objects, these conjunctions may also become effectively “blended” in time with one another such that collision probability



accumulates appreciably even during the mid-point times between the encounters rather than just during short bursts very near the close approaches (Baars et al. 2019). In such cases, two impediments arise to estimating accurate P_c values using methods formulated for single close approaches such as a two-dimensional P_c . First, there is no clear, single TCA at which to evaluate the collision likelihood. While one could in fact find the unique point of closest approach between the two nominal trajectories for the entire sequence and perform a two-dimensional P_c calculation for that encounter, there is no guarantee that another encounter in the sequence may actually possess a higher P_c due to different covariance orientations even though it has a larger nominal miss distance. Second, performing a single encounter P_c calculation at each of the close approaches in the sequence and then amalgamating these using the following formula (derived from DeMorgan's Law of Complements)

$$P_{Tot} = 1 - \prod_{i=1}^n (1 - P_{ci}) \quad (N-4)$$

is likely to overestimate the overall collision likelihood. This is because the above formula presumes that the individual events in the sequence are statistically independent, but in fact they are not, especially if blended in time. As explained previously, it is believed that the three-dimensional N_c functionality properly provides an upper bound for the amalgamated risk over the time interval of interest, so Monte Carlo from epoch would be run operationally in cases in which this upper bound is above a mitigation threshold and there was interest in determining whether the Monte Carlo higher-fidelity calculation would reduce this value to one much closer to or below the threshold. Monte Carlo from epoch can also be run if there is any question about the rectitude of the P_c calculation. As stated earlier, lower- P_c conjunctions may present intractable Monte Carlo execution times, but if one wishes only to ensure that the Monte Carlo P_c falls below the mitigation threshold (rather than establish a high-fidelity P_c value), this can usually be accomplished with far fewer Monte Carlo trials.

In applying the Monte Carlo from epoch P_c estimation methodology, what is envisioned here is not to execute the evaluation for a relatively short time interval bracketing a single conjunction's TCA but rather over a more extended interval that may span several close approach encounters. For example, the collision likelihood between two objects in conjunction would not be evaluated at the nominal TCA but, perhaps, over a risk assessment interval projecting forward seven days. This multiday interval would include the original nominal TCA but also a sequence of other encounters between the primary and secondary as well. In this case, each Monte Carlo trial would be propagated forward seven days and a hit registered at the earliest time that the hard-body radius is violated, if such a time exists (see Hall et al. 2018). Temporal risk plots can be produced using the sequence of hits registered during all the trials, an example of which is shown in Figure N-10 below. Related to Figure N-6 from the earlier section that discusses the three-dimensional N_c calculation method, the pink area shows the Monte Carlo confidence region and the pink line the actual result. As can be seen, the black upper line, which is the upper boundary estimate from the three-dimensional N_c function, is within the confidence interval of the Monte Carlo results and thus is a reasonable actual realization of the repeating conjunctions' cumulative risk.




Figure N-10 Three-Dimensional N_c Temporal Risk Plot with Monte Carlo Result Overlay (in Pink)

Because it is complicated to set up the execution environment for a computationally intensive calculation such as this, and because “gold standard” results require assembling environmental data and software settings identical to the original DOD orbit-determination solutions, it is envisioned that the ability to run this strain of Monte Carlo will remain with NASA CARA. However, a more computationally efficient mode of Monte Carlo, which is serviceable for several different applications and is easier to obtain and employ operationally, is described in the next section.

N.10 Monte Carlo from TCA

A much more computationally efficient variant on Monte Carlo from epoch, which has been used by CA practitioners for some time, is Monte Carlo conducted from TCA. As the definition implies, the Monte Carlo simulation begins with the primary and secondary objects’ equinoctial element states propagated to TCA. Perturbation and sampling of both states is conducted much as described earlier for Monte Carlo from epoch, and each sampled primary and secondary state is propagated both forward and backward in time to find the pair’s TCA and determine whether the corresponding miss distance is smaller than the hard-body radius (backward propagations are required to register hits that occur at times before TCA). The simplification arises from the fact that, since one is beginning from TCA, the propagations required will be short. A low-fidelity propagator (such as a two-body propagator) can be used, and this, combined with the very short propagation times, vastly improves the computational efficiency of the calculation—by a factor

of 10,000 to 100,000 according to the study by Hall et al. (2018). To use this simplified approach, the conjunction duration needs to be short so that one may safely presume a single, unblended event that does not require the Monte Carlo from epoch method. As a second condition, one must have confidence that both objects' states and covariances propagated to TCA are good representations of the states and state uncertainties at that point. Usually, there is reasonable confidence in the mean state estimates themselves, but the covariances are a different matter: a number of studies (e.g., DeMars et al. 2014) have indicated that propagated covariances represented in Cartesian space fail to represent the actual uncertainty distributions, due both to the potential failure of the linearized dynamics to remain representative over long propagation intervals and, more importantly, a mismatch between elongated in-track uncertainty volumes and the forced representation of these uncertainty volumes as Cartesian ellipsoids. The latter problem is illustrated in Figure N-11. The actual in-track error volume should follow the curved trajectory of the orbit, but the Cartesian covariance is limited to the rectilinear representation shown: as the elongation grows in the in-track direction (which occurs for longer propagations), the mismatch between the two representations also increases.




Figure N-11 Mismatch between Elongated In-track Covariances and Forced Cartesian Ellipsoidal Representation

To address the latter problem, the results from a study conducted by Sabol et al. (2010) are both important and extremely helpful. This study addressed directly the question of the optimal choice of orbital state representation for covariances, finding that it is not the specific state representation in which the propagation is executed but rather the one in which the propagated covariance is used that ultimately governs the realism of the uncertainty distribution and associated error volume. Specifically, if the covariance is rendered and used in a curvilinear state representation, such as equinoctial elements, then it tends to represent the error volume much more accurately than if it is transformed and used in Cartesian coordinates. The surprising result is that a non-representative Cartesian covariance transformed into an equinoctial covariance becomes a representative covariance. Furthermore, taking random samples using the equinoctial state representation and performing the non-linear conversion of each sample to Cartesian



coordinates generates a point cloud in the Cartesian frame that also approximates the true error volume much more accurately.

This latter procedure allows the Monte Carlo from TCA method to employ more realistic uncertainty volumes, at least with respect to orbital state representation-related mismatches. The detailed procedure is the following:

1. Convert both objects' states and covariances at TCA to equinoctial elements.
2. Generate a set of perturbations for each object based on the equinoctial covariances.
3. Combine these with the mean equinoctial states to generate sampled equinoctial states for the primary and secondary.
4. Convert these sampled states from equinoctial elements to Cartesian coordinates using the non-linear transformation.
5. Propagate the Cartesian states for both the primary and secondary using two-body equations of motion to find the new TCA, which may precede or follow the nominal TCA.
6. Determine whether the new miss distance is less than the hard-body radius, and if so, register a hit at the time that the hard-body radius sphere is violated.
7. Repeat steps 5-6 until the entire set of Monte Carlo sampling trials has been processed.
8. The P_c is the number of hits divided by the total number of trials, and the confidence interval can be calculated from an appropriate formula.

This approach seems reasonable enough; but it would be presumptuous to assert, without further study, that it is truly robust, especially since the question of the durability of the linearized dynamics typically used to propagate covariances was not directly addressed. As it is, additional study efforts have been performed to verify that it is indeed sufficiently representative for CA applications, and they are described below.

The first of these study efforts was performed as part of the previously cited study effort by Hall et al. (2018). A set of 373 high- P_c conjunctions was selected and evaluated with both Monte Carlo from epoch and Monte Carlo from TCA, and the comparative results are shown in Figure N-12. The top window is a scatter plot of the P_c calculated by Monte Carlo from epoch (called VCM mode here because it works from the Vector Covariance Message, which gives the orbit-determination results at epoch) versus that from Monte Carlo from TCA (called CDM mode here, because it can work directly from the data in the CDM). The intersection of each "plus" sign gives the scatter-plot point, and the length of the plus-symbol tails indicates the uncertainty of the calculation. One can see that the agreement is strong as all the points are close to the dashed y-x line that would indicate perfect equality. The bottom window provides the base-ten logarithm of the ratio of the VCM to CDM mode results, plotted against the CDM mode P_c . The largest deviations are about 0.2 of an order of magnitude in P_c , which is considered to be below operational significance. A separate statistical test for similarity of results produced p -values all less than $1E-03$, indicating that one should reject a hypothesis that these results arise from

different distributions. Good agreement is thus observed between the abbreviated Monte Carlo from TCA and Monte Carlo from epoch results.




Figure N-12 Comparative Results of 373 High-Pc Conjunctions¹²

The second study effort involved a comparison between Monte Carlo at TCA and the two-dimensional Pc calculation (Hall 2019b). About 44,000 conjunction events with two-dimensional Pc values greater than 1E-07 were subjected to Pc calculation by both the two-dimensional Pc algorithm and Monte Carlo at TCA, and the comparative results are shown in Figure N-13. This plot is similar to that shown above, except colored diamonds represent conjunctions in which the two-dimensional Pc calculation significantly underestimates the Monte Carlo from TCA Pc values, and their particulars are described in the plot legend. In the main, the agreement between the two is excellent: 99.48% of the cases showed no statistical difference between the Monte Carlo and two-dimensional Pc calculations. For those that did show large disparities, the subset that had true Pc values in the tractable range for Monte Carlo from epoch were validated with this methodology; and in each case the Monte Carlo from epoch reruns matched the output from the Monte Carlo from TCA. To the degree that non-representative covariances may be responsible for two-dimensional Pc failures (due to coordinate frame mismatches), Monte Carlo from TCA certainly appears to be able to recover the true Pc. As an aside, some of these differences between two-dimensional Pc and Monte Carlo are considerable—there are several

¹² From Hall et al. 2018.



cases in which the two-dimensional P_c understates the actual P_c by more than an order of magnitude.




Figure N-13 Comparative Results of 44,000 Conjunction Events¹³

Originally, Monte Carlo from TCA was advanced as a robust and computationally tractable way to ensure reliable P_c calculations in the face of occasional miscarriage of the two-dimensional P_c algorithm. However, the development and testing of the three-dimensional N_c analytic methodology have led NASA CARA to recommend that this analytic method supplant routine use of Monte Carlo from TCA. One may still resort to this more computationally efficient Monte Carlo method if desired, but testing has indicated that the analytical three-dimensional N_c calculation method outperforms Monte Carlo from TCA (in that it performs equally well for isolated conjunctions but also provides a cumulative risk upper bound for repeating conjunctions) and is probably at least an order of magnitude more efficient, even with an optimized Monte Carlo propagation engine. Monte Carlo from TCA is a capability available for download from the NASA CARA software repository,¹⁴ and O/Os who are currently using it or a

¹³ From Hall 2019b.

¹⁴ See Section 7, Contact Information in this document for the CARA software repository URL.



similar implementation are not urged to take it out of service, but as a P_c calculation approach, it does not seem to offer any enduring advantage over the three-dimensional N_c method.

Finally, a dedicated study on covariance Gaussianity was recently conducted (Lechtenberg 2019b). The same set of 44,000 conjunctions introduced above was analyzed for multivariate normality of the position covariance in Cartesian coordinates. The methodology was to convert the covariance to equinoctial elements, generate a set of random position samples from this covariance, convert the sample set to Cartesian coordinates, and apply a multivariate normality test (here the Henze-Zirkler test) to assess compliance to this distribution. At a 5% significance level, only 60% of the cases could be considered to conform to a multivariate Gaussian distribution in Cartesian coordinates. In principle, such conjunctions would be considered suspect cases whose two-dimensional P_c results would be considered doubtful. As it is, since 99.48% of the cases show good agreement between the Monte Carlo at TCA and the two-dimensional P_c , clearly the Gaussianity of the covariances in the Cartesian framework does not matter appreciably for the P_c calculation. This result corroborates that of an earlier study by Ghrist and Plakalovic (2012), which reported similar findings.

How can this be, given that the covariance is integral to the two-dimensional P_c calculation? It is important to remember that the analyses above are restricted to high- P_c events. For the probability of collision to be high, significant overlap needs to exist between the primary and secondary object covariances; this means that the central parts of the covariances, which are where most of the probability density lies, have to overlap substantially. Such a requirement makes the behavior of the tails of the covariance much less important, and it is in the tails that non-Gaussian behavior is most strongly manifested. Even though a good number of conjunctions fail tests for covariance Gaussianity in Cartesian coordinates, for high- P_c events, this result does not appear to affect the rectitude of the calculated P_c . A difference probably is observable for lower- P_c conjunctions, but because these conjunctions are not operationally significant, it is not operationally important to identify or characterize this phenomenon.

There is the lingering question of the eventual failure of the linearized dynamics in propagating covariances as DOD covariances are propagated through a pre- and post-multiplication by a linearized state transition matrix. It is agreed that, given sufficiently long propagations, such an eventuality should arise. However, the propagation duration required for this problem to manifest itself substantially is believed to be much longer than encountered for most conjunctions—on the order of weeks. This is why more attention is paid to this phenomenon in other areas of SSA such as uncorrelated track processing for which propagations of 30 days or more may be required, whereas it is rare for propagations longer than ten days to take place in CA. Of course, fresh tracking data are always appreciated as they shorten the propagation interval and lend additional confidence to the solution.

N.11 Choosing an Appropriate Hard-Body Radius

As discussed in the first section of this appendix, the hard-body radius represents the combined size of the primary and secondary objects; the word “radius” is used because this combined size is typically envisioned as a sphere, and the radius of a sphere is a convenient linear representation of its size. The hard-body radius is needed for the P_c calculation as it represents



the circle/sphere within which a simultaneous penetration by both the primary and secondary objects' trajectories will constitute a presumed collision. It is not just a required input to the Pc calculation, however, it is also one of the governing parameters of the calculation: the Pc value represents an integral over the area of the hard-body-radius circle (or surface area of the hard-body-radius sphere) and thus in many circumstances, varies roughly in proportion to the square of the hard-body radius, so an increase of the hard-body radius by a factor of three (3) increases the calculated Pc by nearly an order of magnitude. Because of this sensitivity, it is important not to overstate the hard-body radius simply for “conservatism” as the effect on the calculated Pc can be considerable. The best overall strategy for applying conservatism, should one wish that, is to apply it at the end of the process by lowering the Pc threshold at which mitigation actions should occur. Injecting conservatism into different places throughout the Pc calculation makes it difficult to determine how much conservatism has actually been introduced, whereas addressing this desire through a modification of the Pc threshold for mitigation allows it to be understood precisely.

Mashiku and Hejduk (2019) recently completed a study of different hard-body-radius calculation and setting strategies, and a streamlined summary of the possibilities examined is provided below:

1. **Large *a priori* value.** For some years it was standard practice simply to choose a hard-body-radius value that was notably larger than the expected actual size of the combined primary and secondary object; 20 meters is a value that was typically used. Perhaps in the early days of CA, this was an acceptable initial screening strategy to identify potential serious events, and the hard-body radius would then be reduced when analysis began in earnest. As the space catalog has grown in size, and especially with the expected near-term future growth through the deployment of the Space Fence radar, this particular strategy merely creates additional false alarms that unmeaningfully burden the process. Nearly all O/Os have moved away from this hard-body-radius strategy.
2. **Circumscribing sphere.** The use of a circumscribing sphere to set the primary object's hard-body radius is perhaps the most commonly used present operational technique, which admits of two typical variants: placing the sphere's center at the center of mass of the primary satellite and defining the radius by the length from this center to the satellite's most distant extremity; or allowing the center point to float freely and then defining the smallest circumscribing sphere. The overall sphere size then has to be increased by the expected size of the secondary object using either an averaged value or an estimate for the particular secondary encountered. This estimate could arise either from published dimensions (for intact spacecraft or rocket bodies) or attempts to estimate size from satellite signature data such as radar cross-section or visual magnitude.
3. **Maximum projected area into any plane.** Since the circumscribing sphere described above will end up being projected into the conjunction plane, it is instructive to examine in more detail the implications of such a projection. Clearly the sphere itself will project as a circle, but the projection of the three-dimensional spacecraft inside will necessarily be smaller in area than the projected circle, and for some spacecraft shapes and



orientations, it will be substantially smaller in area. In this latter case, the substantial “white space” within the projected circular area not occupied by the primary could justifiably be excluded in the P_c estimation process, especially for the most common debris-encounter scenario when the incoming secondary object is much smaller than the primary asset. A straightforward way to address this issue that does not require knowledge of the satellite’s actual orientation in relation to the conjunction plane is simply to determine in advance the maximum area that the satellite can possibly project into any plane (based on a three-dimensional CAD¹⁵ model of the satellite), and use a hard-body-radius circle of that equivalent area (which of course must then be enlarged to include the estimated size of the secondary) for P_c estimation. It is true that this is a conservative formulation in that it uses the maximum possible projected area, but this is often substantially smaller than the area of the projected circumscribing sphere. One could argue that uncertainty is introduced by using the equivalent circular area rather than a contour integral to perform the integration over the actual projected shape, but individual exploratory examples show that this difference is usually negligibly small, and in any case, the most conservative projection approach should compensate for any differences in the shape chosen to represent the hard-body-radius area for the P_c calculation.

4. **Projection into actual conjunction plane.** The most accurate, and at the same time the most difficult, approach is to perform a projection of the primary satellite’s shape into the actual conjunction plane. Specifically, this requires a three-dimensional CAD model of the satellite plus knowledge of its inertial attitude and the orientation of any major articulating components at TCA along with a calculation to project the resulting shape into the conjunction plane. Once this projection is obtained, its boundaries have to be augmented to account for the expected size of the secondary, and the integration of the joint covariance probability density can take place over this figure via contour integration or over a more convenient shape of equivalent area. Chan (2019) recently proposed a method to decompose complex shapes into individual triangles and use an equivalent-area method to evaluate the P_c for each triangle; the composite P_c is simply the sum of the P_c values for these individual trials of decomposition.

Each successive approach among the four presented brings greater precision to the hard-body-radius determination but at the same time, additional complexity. A reasonable *via media* would appear to be approach 3) above, which keeps the hard-body-radius value grounded in reality and free from excessive conservatism but avoids the difficulties of gathering and maintaining shape and attitude data to enable a detailed projection calculation for each conjunction.

¹⁵ Computer-aided Design



Appendix O. Collision Consequence

O.1 Introduction

It is common in CA circles to speak of “collision risk metrics” such as the Probability of Collision (P_c), but in fact, calling such parameters risk metrics is a misnomer. As was first laid out formally in the literature by Kaplan and Garrick (1981) has been adapted and reinforced by NASA’s own project risk assessment paradigms. Risk is actually the combination of the likelihood of a detrimental event and the consequence if such an event should come to pass. CA risk assessment metrics such as the P_c are simply measures of collision likelihood, which is only part of the overall risk assessment. Because CA practice arose originally from individual missions whose focus was to protect their own spacecraft, it seemed unnecessary to add any considerations beyond collision likelihood. The loss of the mission’s spacecraft was considered a catastrophic event to be avoided at all costs, so all that needed to be evaluated was the likelihood of a collision, which if it took place, was presumed to be fatal to the satellite’s operation and the mission.

When the conjunction assessment risk analysis for groups of satellites or an entire agency or nation is considered, however, the problem broadens. There is certainly the desire to see individual missions protected from catastrophic termination. The particular urgency of protecting any given spacecraft from premature end-of-mission life is a product of a number of factors such as the degree to which mission objectives have been met, the age of the satellite, and the functional redundancy within a constellation. Because the individual mission has the best sense of these considerations, they are in the best position to assess the consequences of potential spacecraft loss. However, there is a second parallel consideration, which is the preservation of orbital corridors from debris pollution to ensure their utility for future missions, both by the United States and all space-faring nations. The NASA Orbital Debris Program Office (ODPO) has for some years been the NASA project chartered with this concern. They have outlined spacecraft design and disposal recommendations to mitigate orbital debris production, but there has not yet been an effort to establish clear requirements for debris minimization through on-orbit CA activities during mission lifetimes. One of the purposes of the present document is to lay out best practices for CA operations that promote debris minimization. To reduce the production of debris, one must not only consider the likelihood of any given collision between a NASA primary satellite and other space objects but also evaluate the amount of debris that such a conjunction, if it results in a collision, may engender.

O.2 Debris Production Determination Methodology

It is very much the case that different types of satellite collisions can produce substantially different amounts of debris. The ODPO has studied this phenomenon in some depth through staged collisions of satellites with simulated debris objects in vacuum chambers and established relationships among conjunction parameters and the number and size of the resultant debris objects (Johnson et al. 2001). The basic distinction is between catastrophic collisions in which both the primary and secondary object are fully fragmented and thus generate large amounts of debris, and non-catastrophic collisions in which the smaller/lighter satellite is fully fragmented but the larger/heavier satellite only cratered, generating much smaller amounts of debris. The



catastrophic/non-catastrophic distinction is governed by the following relationship, based on the relative kinetic energy of the encounter:

$$\frac{mV_{rel}^2}{2M} > 40,000 \quad (O-1)$$

in which V_{rel} is the relative velocity between the two satellites, m is the mass of the lighter satellite, and M is the mass of the larger satellite. When the relative kinetic energy exceeds 40,000 Joules per kilogram, the collision is expected to be catastrophic and produce much larger amounts of debris. The ODPO has further developed a relationship that estimates the number of debris pieces greater than a specified size that the collision will generate:

$$N(L_c) = 0.1(P)^{0.75} L_c^{-1.71} \quad (O-2)$$

in which L_c is the characteristic length (size) above which one wishes to determine the number of debris pieces, and P is a momentum factor of sorts. If the collision is catastrophic, P is the sum of the two satellites' masses; if it is non-catastrophic, it is the product of the mass of the lighter satellite and the conjunction's relative velocity.

These relationships conform to first-order intuition regarding expected relative levels of debris production. Any schoolboy who has played with dirt clods can testify that a greater relative velocity between two conjuncting dirt clods will produce greater fragmentation. Furthermore, if the masses of the heavier and lighter object differ substantially, which probably also means that their relative sizes differ as well, one can imagine the (much) smaller object simply passing through (cratering) the larger one and, while itself fragmenting, leaving the larger object essentially intact. This full versus partial fragmentation division introduces a discontinuity in the debris production relationship (as a function of relative mass and relative velocity), and the entire dynamic, illustrated by Figure O-1, shows the extremely broad range of debris production possibilities—from very few debris pieces to thousands—all of which represent real possibilities for satellite conjunctions because the full range of secondary object mass values and relative velocities shown in the figure are frequently encountered in practice.

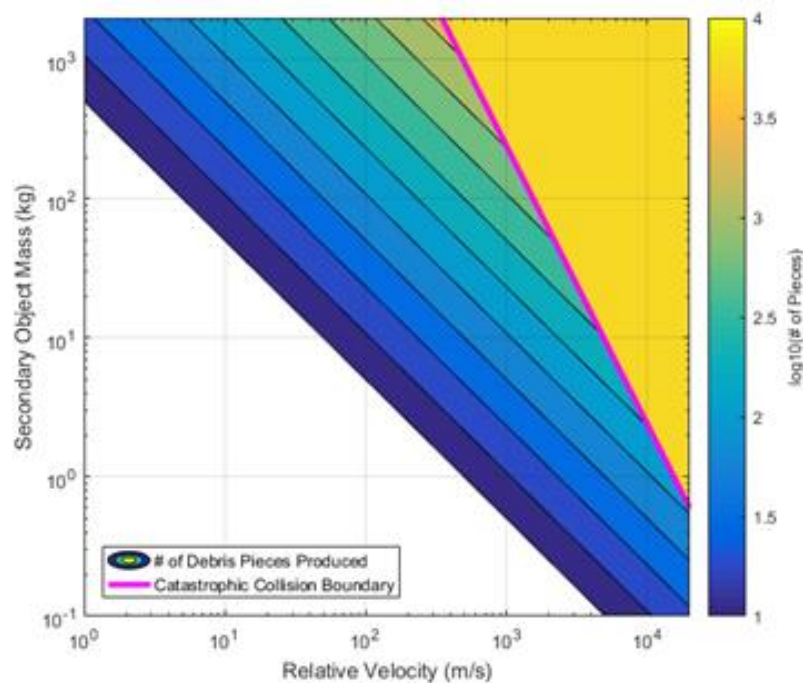


Figure O-1 Range of Debris Production Possibilities¹⁶

Since the debris production potential can vary so greatly among different conjunctions, it makes sense that conjunction remediation requirements should take cognizance of this parameter, presuming of course that it can be estimated from available conjunction information. The conjunction relative velocity is easy to calculate and is in fact provided directly as part of the CDM, and the mass of the primary object is known by the satellite O/O. The mass of the secondary object, however, is more elusive; in the great majority of conjunctions, this object is a debris object for which there is *no a priori* information. One must therefore look to estimation processes to estimate this object's mass.

In formulating a mass estimation methodology, an approach is desired that allows a predetermined amount of conservatism, namely a specified degree of belief that the amount of debris potential is not underestimated. If conjunction mitigation requirements are to be somewhat relaxed for low-debris-potential conjunctions, it is important to err on the side of debris quantity overestimation to guarantee that all situations in which this relaxed threshold may be applied are truly warranted. In nearly all practical cases, a larger secondary object mass value will result in larger produced debris counts. A conservative approach requires ensuring, to a desired degree, that the estimation process overestimate rather than underestimate the mass of the secondary object. Such an orientation can be tuned and verified by attempting to estimate the mass of small objects with known mass values such as microsats.

¹⁶ From Hejduk et al. 2019.



The most satisfactory approach to date has been to look to the orbit-determination solved-for parameter of the satellite ballistic coefficient. The formal definition of the ballistic coefficient (B) is given as:

$$B = C_D \frac{A}{M} \tag{O-3}$$

in which C_D is the drag coefficient, A is the satellite frontal area presented to the atmosphere, and M is the satellite mass. Because the parameters in the ballistic coefficient equation are difficult to represent as single values, the chosen approach is to represent each of them by a probability distribution, producing as an outcome a probability density function for the estimated mass. This approach will also be helpful in ensuring a conservative outcome as one can choose the percentile point of the mass distribution at which one wishes to operate to produce a desirably small frequency of underestimation.

The following estimation process was assembled initially in Hejduk et al. (2017) and refined by Lechtenberg and Hejduk (2019) and Lechtenberg (2019a):

- **B .** Each orbit-determination update for the secondary produces an estimate of the ballistic coefficient (B) and an estimation variance (σ_B). These values can be used to generate a distribution of possible B values for the secondary.
- **C_D .** Recent studies have established reasonable C_D values for different satellite shapes. Given expected debris object shape possibilities, a reasonable value and uncertainty can be chosen for the C_D parameter.
- **A .** Satellite frontal area is perhaps the most difficult of these parameters to estimate as it is not directly calculable or observable. However, radar cross section information is taken as part of satellite radar metric tracking, and these data can be used with ODPO-developed models to develop distributions of satellite characteristic dimensions. It is important to work with entire distributions of radar cross section data for a particular object rather than a single point estimate for this is the proper framework in which the ODPO size estimation models operate. From a distribution of characteristic dimensions, with certain assumptions, a distribution of satellite frontal areas can be assembled.

With distributions for the above three parameters, Monte Carlo techniques can be used in conjunction with Equation O-3 to generate a distribution of satellite masses.

This estimation approach was tested against two different classes of objects: sodium potassium (NaK) droplets formed from coolant leaking from certain rocket body types constituting a fairly benign satellite test set (due to their being spherical; see Rosenberg 2003), and about 400 small satellites that were deployed from the ISS for which definitive size and mass data were available. As expected, the latter test set proved more challenging. However, as Figure O-2 shows, choosing the 99.9 percentile level of the distribution of estimated masses as the mass estimate guarantees that none of the test objects' masses were underestimated.

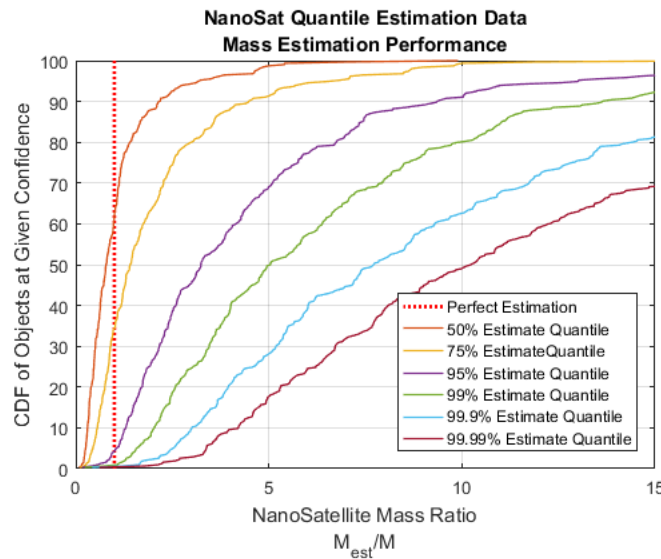


Figure O-2 Confidence in NanoSat Mass Ratio¹⁷

The obvious follow-on question is whether choosing so conservative an estimate level creates a situation in which virtually every conjunction is presumed to be a high-debris-producing event. If this high-debris threshold is set at the catastrophic collision level for three NASA primaries in approximately 700 km circular orbits, one can see in Figure O-3 that, following the 99.9% curve at the catastrophic threshold of 40,000 J/kg, about 70% of the conjunction events are non-catastrophic. If, for additional conservatism, the relative energy threshold is reduced to half its present value (that is, to 20,000 J/kg), the number of non-catastrophic events would still be a majority. As such, even an extremely conservative rendering of this criterion identifies a substantial number of events that can be legitimately and confidently treated as low-debris-producing situations if they result in collision.

¹⁷ From Lechtenberg 2019a.

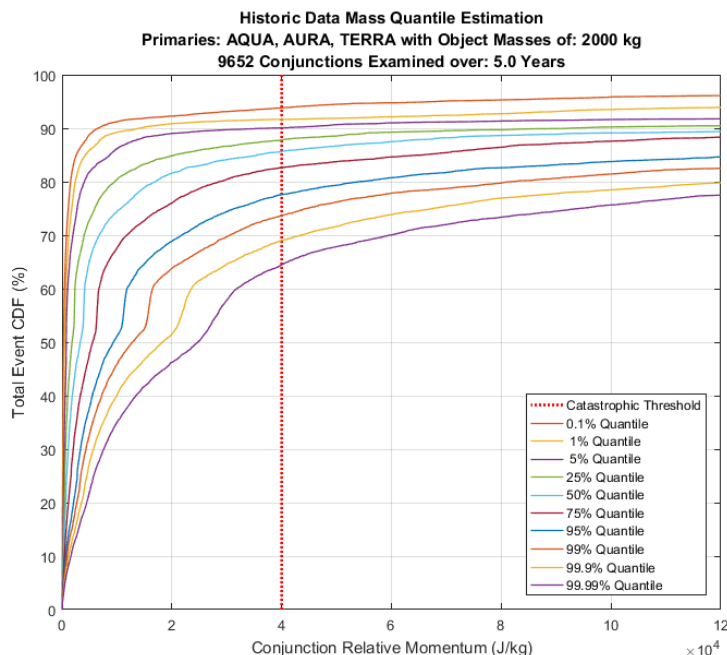


Figure O-3 Catastrophic/Non-Catastrophic Composition as a Function of Mass Quantile¹⁸

It must be emphasized that no O/O would be asked to take advantage of the low-debris-producing-event possibility for relaxed mitigation requirements. Whether to focus on the protection of their own payload or merely to take a more conservative CA posture generally, missions certainly may embrace the 1E-04 Pc threshold for all conjunctions. From an orbit regime protection point of view, however, it would not be unreasonable to relax this threshold mildly for situations that are non-catastrophic and only expected to create a small amount of debris. Such a consideration could also be desirable for triage situations in which a mission is experiencing more serious conjunctions than it is possible to remediate. In such cases, one could direct mitigation focus to those conjunctions that hold the greatest debris production potential.

¹⁸ From Lechtenberg 2019a.



Appendix P. Event Actionability

P.1 Introduction

Although widely used, the term “CA event actionability” is a somewhat unfortunate term because it implies a broader definition than it merits. Its precise meaning is the assessment of the state and state uncertainty information for the two satellites in conjunction (but with the focus usually on the secondary object) to determine whether these data constitute a reasonable basis for CA risk assessment and mitigation actions. It does not include the further step of deciding, in the presence of a high-risk situation, whether the mitigation actions are excessively invasive or detrimental to the mission’s objectives to warrant their being set aside despite the collision risk; that is, whether an O/O would be justified in refraining from taking “action” to mitigate the conjunction and lessen the collision risk because of the mission impact of the mitigation action. This latter consideration is real and constitutes part of the risk assessment process, but one arrives at this decision point only after determining that the orbital data feeding the analysis are durable and are enabling credible collision risk calculations.

The propriety of considering the quality of the event’s input orbital data is obvious as an O/O will certainly wish to certify that the collision risk is truly high and that a proposed mitigation action will remediate the collision risk before considering its actual execution. The desire for a conservative risk approach does not make the problem simpler. If the input astrodynamics data are questionable, it is also questionable whether any proposed mitigation action will actually improve the situation as such an action could in fact make the situation worse. So it is not appropriate simply to say that situations of questionable data quality should receive such actions if the calculated risk parameters exceed the threshold for mitigation actions. It is important to establish a method for identifying those cases in which the input data may be sufficiently poor to make meaningful risk assessment impossible.

It is important to recognize that the question to be addressed here is whether the orbital data properly represent the two objects’ states and uncertainties and not whether the associated uncertainties are large or “too large.” If one embraces P_c as the principal (but not necessarily the sole) collision likelihood assessment metric, excessively large covariances have the effect of depressing the P_c value and thus giving a (much) lower statement of risk. This particular phenomenon, called the “dilution region” situation, is a known peculiarity of P_c and is discussed at some length in Appendix M, but there is general agreement that high P_c values do represent situations of high collision likelihood. In such situations, the fact that covariances may seem excessively large should not constitute an impediment to event actionability. If, despite a somewhat data-impooverished situation (hence the large covariance), a worrisome P_c is still produced, *a fortiori* the situation should be treated as serious and worthy of mitigation.

This area of CA risk assessment—the determination of event actionability—possesses the least precision in terms of clear, documentable standards. The certification of orbit-determination updates as valid and reasonable and propagation intervals and methods as sufficiently representative is, in the end, rendered not through conformity to strict, fully articulated rules but rather through the judgment of expert orbital analysts who have personally performed thousands of orbit-determination updates and propagations. Astrodynamics is often at least as much an art



as a science, and batch orbit-determination update particularly so. To be sure, most of the cases that are actually encountered in CA practice can be judged to be valid and acceptable through rule conformity alone, and there are certain cases that are so poor that there is little difficulty in recognizing their invalidity immediately. But there is also a relatively small input data group that cannot be directly assigned to either of these resolutions, and it is here that expert opinion is required.

There are two separate aspects of the astrodynamics input data to consider: the states and covariances for the primary and secondary objects at epoch as outputs of the orbit-determination process, and the states and covariances propagated forward to TCA. These two aspects will be treated separately below, with the concluding section outlining how considerations from each are combined into a judgment of actionability/non-actionability for a particular event. Some of the threshold values for the actual parameters used in the tests to be described below are placed in an annex to this appendix, both for ease of reading and to allow easier document update should these values be updated. Finally, while it is the case that in principle a judgment of the suitability of the primary object's state and covariance is required as part of the overall actionability assessment, in practice it is almost always the secondary object's data quality that is an issue. To simplify the treatment, the following discussion will frequently frame the discussion in terms of a single object that can be presumed to be the secondary.

P.2 Evaluation of the Propriety of the Orbit-Determination Fit

It should be obvious that a quality orbit-determination fit for the object is a foundational element in any use of its state and uncertainty data. If the fit itself is questionable, it is difficult to have confidence in any astrodynamics products propagated forward from it. Because the USSPACECOM catalog is at present usually the sole supplier for all data on the secondary object (especially in the case of debris objects with a small radar cross section value) and also provides solutions for the primary object, the focus has to be placed on the orbit-determination fit mechanism used for this catalog, which is not a sequential estimator but a minimum-variance batch process. This means that in addition to the force model settings for the dynamical models, the time-span of the data for the fit, density, composition, and retention of the sensor data used in each batch update have to be considered. If a different data provider for orbital information about the secondary object were to be used, a set of orbit-determination quality criteria tailored to that particular provider's orbit-determination methodology and engine would need to be developed and consulted. The remainder of this section could be seen as a proposed template for how that development and consultation might proceed.

If it were possible to generate hard-and-fast rules for determining the control settings and data handling for these batch updates, the assessment of fit acceptability would be much more straightforward. Despite decades of attempts, astrodynamics specialists for DOD have concluded that such a rule set is simply not possible, especially when densely packed debris clouds are considered. The only way truly to evaluate the propriety of an orbit-determination fit is a manual review by an expert. This individual has access to information of the full suite of input sensor raw and "management" data including:



- Tables of residuals,
- Residual plots in a variety of different constructions including the ability to edit out individual observables,
- Before-and-after values of key orbital elements in the context of long-term plots to evaluate trends, and
- The ability to attempt multiple corrections with different force-model settings and data contents to evaluate the results comparatively and choose the best outcome.

Such evaluations, most of them visual, simply cannot be replaced by rules themselves.

However, what is possible is to provide a set of guidelines that, if violated, would typically prompt a manual review by an expert. Within a NASA or USSPACECOM context, an expert could be summoned and, with access to the data products enumerated in the above paragraph, could render an expert opinion on the propriety of the orbit determination. Outside of these agencies, such data access is not available. Nonetheless, the rule set can be useful; for example, orbit-determination results that do conform to all of the general rules can, with a high degree of confidence, be accepted as a basis for CA risk assessment.

The following subparagraphs describe the particular orbit-determination input summaries, settings, and quality factors evaluated by the rule set, and the annex to this appendix gives threshold values that allow for the evaluation of an orbit determination's conformity to these rules. All the information needed to evaluate a particular orbit determination is present in the CDM.

P.3 Force Model Settings and Non-Conservative Force Model Parameter Values

Any space catalog has at its core a set of astrodynamics dynamical models that are used to predict the trajectories of space objects, and the orbit-determination process determines parameter values to use within these models to predict future object positions. Different physical phenomena and their effects on satellites are modeled, and for some of these models, varying levels of fidelity can be selected. The following paragraphs list each of the selectable force models that can be applied, identify the selection criteria, and give some indication of parameters and parameter values that might lead one to question the propriety of an orbit determination.

1. **Geopotential.** Two-body motion, the simplest of possible orbit solutions, is the foundation of the astrodynamics model set. The first major improvement in precision is the recognition that the Earth is not a uniform and uniformly dense sphere but varies from this idealization in many ways. A spherical harmonic representation of the actual geopotential is used to represent the Earth's bulge at the equator and other density non-uniformities; the spherical harmonic order is specified as part of the correction. The minimum required geopotential order (in terms of zonals and tesserals) for an acceptable fit is specified by orbit characteristics, and any orbit-determination update that uses a geopotential order lower (smaller) than this specified value is flagged for manual review.



The actual minimum geopotential order thresholds as a function of orbital parameters are given in the annex to this appendix.

2. **Lunar-Solar Perturbations.** While the Earth is obviously the most significant central attracting body for a satellite, the sun and moon are large enough to visit a non-negligible effect on a satellite's orbit. For precision solutions, this modeling (called "third-body effects") should always be engaged.
3. **Solid Earth Tides.** The gravitational force of the moon deforms the quasi-spherical Earth, and this deformity changes the geopotential and satellite orbits. Often, the effect is small and unlikely to alter CA results. Because it is not computationally difficult to characterize, it is standard practice to include this modeling input in precision solutions. In the main, however, it is unlikely that the failure to include this perturbation would make the resultant orbit determination questionable from a CA perspective although it could alter the results observably.
4. **Atmospheric Drag.** The major non-conservative force affecting many LEO orbits is atmospheric drag. The modeling of drag acceleration is a function of the satellite's velocity relative to the atmosphere, the atmospheric density, and the ballistic coefficient, which itself is a combination of the coefficient of drag, the satellite's frontal area, and its mass. Because for most satellites, especially debris objects, neither the drag coefficient nor the frontal area nor the mass is known, the entire ballistic coefficient is treated as a single solved-for parameter as part of the orbit determination. Certain LEO orbits should include atmospheric drag modeling and solutions for the ballistic coefficient, and for those that do, there are ranges of reasonable solutions for this solved-for parameter. The annex to this appendix delineates the LEO orbits that require a drag solution and the ranges of acceptable ballistic coefficient values.
5. **Solar Radiation Pressure.** While electromagnetic radiation has no mass, it does impart momentum, and the momentum carried by the sun's energy and incident upon the satellite has a discernable effect on the resultant orbit. The USSPACECOM dynamical model set employs a standard solar radiation pressure model, which is a function of a solar radiation pressure coefficient (analogous to the ballistic coefficient) and a scalar and vector distance to the sun. Certain orbits should have the solar radiation pressure model enabled, and for these there are acceptable ranges of values for the solar radiation pressure coefficient. The relevant orbits for solar radiation pressure compensation and the associated ranges of values for the coefficient are given in the annex to this appendix.

P.4 Orbit Determination Fit Span (or Length of Update Interval)

Sequential estimators process each measurement individually, usually as each is received; batch estimators collect a "batch" of measurements and process them as a group. When using batch estimation, it is necessary to identify the particular group of measurement data to be fit, which usually means determining the period back in time from the present for retrieval of data to use in an update. While this activity may at first not seem important, in fact it can greatly affect the quality of the resultant orbit determination. If the data period is too long, the prediction error from the orbit-determination fit state is increased; if the data period is too short, a robust drag



solution is not possible; and if there are too few measurements contained within the span, the solution is not reliably representative.

To set the data interval and density correctly, the USSPACECOM batch update process employs a complicated algorithm called the Dynamic Length of Update Interval (LUPI) Assignment (DLA) to examine the measurement data history for each fit and assign the proper LUPI to strike an acceptable balance between maximizing prediction accuracy while ensuring a reasonable drag solution. The algorithm operates by beginning with the maximum acceptable LUPI for a particular orbit and attempting to minimize it while preserving sufficient observational data to endure a quality update, and if the maximum LUPI still does not encompass enough tracking data, the algorithm will expand past this theoretical maximum value to attempt to bring more data into the correction. The goal of this approach is to provide quality updates most of the time and serviceable updates all of the time with “serviceable” defined here as sufficient to allow sensor reacquisition of the satellite. Updates that are merely serviceable are, however, not in themselves of sufficient quality for CA conduct. Also worrisome are LUPI lengths shorter than the default minimum, which occur when large amounts of measurement data are eliminated from a satellite’s history (due, for example, to a satellite maneuver after which the old observational data have to be thrown away because they no longer represent the satellite’s current trajectory). While a subject matter expert could, with broader data access, perhaps certify some of these results as acceptable for CA, as a general rule, orbit determinations with LUPI lengths that are either longer than the theoretical maximum or shorter than the theoretical minimum cannot produce reliable state updates and reliable covariances and thus cannot serve as a basis for CA. A table of these maximum and minimum lengths is given in the annex to this appendix.

P.5 Residual Acceptance

The batch solution generates an optimal solution by minimizing the sum-of-squares residual of the data chosen for the fit; therefore, large residuals against individual observables affect the fitting process substantially. Because it is not uncommon to receive “bad” measurement data from sensors, a stratagem has been implemented to limit the overall effect of potentially errant data: between iterations of the non-linear regressive solver, the measurement residuals against the provisional solution are examined, and data points that manifest excessively large residuals are eliminated. This allows the fit to proceed without being influenced disproportionately by only a few measurements.

The elimination of measurement data in a data-fitting context merely because they appear to encumber a more satisfactory fit is always somewhat disconcerting as most numerical analysts discourage the elimination of any given datum without some *a priori* reason to suspect that it may be errant. However, incidences of satellite cross-tagging (in which the assignment of data to the correct satellite miscarries with the data given assignment to some other, usually nearby satellite) are frequent enough that some provision has to be made for eliminating sensor observations that simply do not seem consistent with the rest of a satellite’s measurement data. There are also situations in which data mismatches are overtly recognized, such as after a satellite maneuver, because the trajectory has changed impulsively and observational data before the maneuver do not represent the current trajectory and have to be eliminated intentionally.



If too much of the measurement information is eliminated, the concern is that the fit may not be properly representative. Therefore, when the percent of the residuals eliminated exceeds a certain value, the fit merits manual review. A table of these values is given in the annex to this appendix.

P.6 Weighted Root-Mean-Square of Residuals

It is common when performing an iterative batch fit to use the root-mean-square of the residuals as an indication of fit quality, tracking this value to ensure that the fit is improving with each iteration. For an unweighted root-mean-square, the ideal value would be zero as it would indicate that the model fits the data exactly without any residual error. For the USSPACECOM minimum-variance batch solver, a weighted root-mean-square is used, in which the square of each residual is divided by the variance of the *a priori* expected error for that observable type (e.g., range to target, azimuth angle, elevation angle) from that particular sensor as established by a sensor calibration process separate from the orbit-determination activity. If the residual errors and actual sensor errors are Gaussian, then the sum of these ratios should equal unity (and so will its square root as well). The weighted root-mean-square is a better measurement of fit error because the best reasonable expected outcome is that the error of the fit will be of essentially the same magnitude as the inherent error in the observations.

If the fit's weighted root-mean-square is too large, then the fit is not particularly convincing, and manual review would generally be counselled. A weighted root-mean-square value less than unity, while not necessarily bad, nonetheless would also warrant a manual review. While a fit of a higher quality than the inherent error in the observations is possible, it is usually a result of too few data in the orbit determination and therefore should be investigated. Values of excessively high weighted root-mean-square are set by object type and given in the annex to this appendix.

P.7 Default Covariance

The DOD precision orbit-determination solution produces an estimation error covariance as a regular product. There are situations, however, in which the automatic precision orbit-determination solution cannot be made to converge, and in such cases, the system executes a general perturbations analytic orbit-determination solution, employing the same theory used to produce TLEs. This process as presently designed produces a state (although an inaccurate one with 1-2 km of expected error in LEO) but no covariance, and this is indicated by providing a diagonal covariance for which the diagonal elements for the position portion are each ten Earth radii squared.

General perturbation solutions are not sufficiently accurate to serve as a basis for CA mitigation actions, so if one of these default covariances is encountered, the state data that accompany it are not suitable for CA risk assessment.

P.8 Review Process

As stated above, if the full orbit-determination data are present and a subject matter expert is available to render an expert opinion, situations in which the general rules given above are violated can be investigated and an adjudication obtained. If this level of access is not available, the above rules can serve as a one-sided test; namely, if the current orbit determination conforms



to all of the rules, then it can be presumed that the orbit determination is durable and can serve as a basis for CA risk assessment and, if necessary, mitigation action planning.

P.9 Evaluation of Propagation from Epoch to TCA

Most CA risk assessment approaches make assessments using the states and covariances of the primary and secondary objects propagated to TCA, so in addition to considering the propriety of the orbit-determination fits themselves, one also has to consider whether any issues with the propagation itself might make the TCA products unsuitable for conducting risk assessment. Issues related to propagation are often called out as potentially problematic, and a treatment of each of these is given in this section. But the overall conclusion is that the USSPACECOM approach to producing CA astrodynamics products at TCA is robust and contains mechanisms to address the major sources of expected propagation error, and these can be augmented with additional processes within risk assessment. As such, if an orbit determination sustains the criteria for an acceptable fit, it can be presumed that this fit propagated to TCA renders a state and covariance acceptable for CA, except in a few narrow instances called out below.

P.10 Propagated Covariance is Defective

In some situations, the propagated covariance can be defective to the degree that it cannot be used as presented for CA.

1. **Null Covariances.** The first is a null covariance in which all the entries are zeroes. This is the equivalent of not having a covariance at all, and when described in the previous section that addressed orbit-determination fit results, it was pointed out that situations with such a covariance cannot be used for CA risk assessment. If there is no covariance at epoch, there cannot be a propagated covariance at TCA either, so such a situation can be considered to be nonactionable. It is possible to use maximalist techniques such as those developed by Alfano (2005b) and by Frisbee (2015) to proceed in situations in which only one of the two objects has a covariance, but these techniques are exceedingly conservative and therefore not recommended as a routinely used proxy for collision likelihood.
2. **Default Covariances.** The second is a default covariance in which the presence of diagonal elements of ten Earth radii indicates a general perturbation solution that produced no covariance. While in principle this covariance can be propagated to TCA, it will essentially be an “infinite” covariance at TCA—just as it was at epoch—and therefore force any calculated P_c to zero (and sometimes introduce convergence problems in P_c calculation routines). Situations such as this can be judged to be non-actionable (because the covariance for one of the objects cannot be established) or produce in every case a P_c of zero (because this is what very large covariances do). Following either possibility leads to the same conclusion of not requiring a mitigation action.
3. **Non-Positive-Semidefinite Covariances.** The third is for the TCA covariance to become non-positive-semidefinite. This issue is addressed at greater length elsewhere (Hall et al. 2017b), and the takeaway from that discussion is that the problem arises due to truncation



and interpolation errors (not fundamental problems with the orbit determination) and can be straightforwardly repaired. The article recommends certain repair techniques, and the source code for them is provided on the NASA CARA software repository. (See Section 7, Contact Information in this document for the specific URL.) A modification to the USSPACECOM operational system has been implemented that should essentially eliminate non-positive-semidefinite situations.

Of the mainstream direct defects in TCA covariances, null and default covariances do produce situations that result in non-actionability; non-positive-semidefinite covariances do not.

P.11 Excessive Propagation Interval to Retain Gaussian Covariance Behavior

The astrodynamics discipline has had to address the question of Gaussian uncertainty volumes and the difficulties introduced by the Cartesian framework's inability to represent curvilinearly produced results easily. As discussed at greater length in Appendix N, while this debate continues in the critical literature, the CA enterprise has been able largely to set it aside by recognizing that the issue is not really relevant to high-Pc events because of the amount of overlap of the "main" portions of the covariance necessary to produce a high Pc. Furthermore, the issue can be remediated easily by performing Monte Carlo from TCA with the random sampling conducted from the equinoctial representation of the covariance. While much has been made in the literature of the problem of non-Gaussian covariances for CA, based on NASA CARA's analysis, this problem is not crippling for the categories of events that actually matter to CA risk assessment, and in any case, it can be addressed with a particular Pc calculation technique that is both straightforward and tractable.

P.12 Excessive Propagation Interval Due to a Lack of Tracking Data

The USSPACECOM orbit-determination process places the epoch time of its orbit-determination fits at the time of the last observation. Since this is the time at which the fitted results are modeled, any propagation into the future takes place from that time. Because this orbit determination is used as the most current update until more tracking data are received and a new update executed, to propagate the state and covariance to TCA usually involves two stages: to propagate forward from the epoch time to the current time, and then to propagate from the current time forward to TCA. If no new tracking data are received for some time, then a quite long propagation of the original orbit-determination state may be required. For example, if the object has not been tracked for five days and TCA is three days from the present time, a total of eight days' propagation is required to produce the state and covariance at TCA.

It is best if tracking data are received frequently as this produces better data density for the fit and reduces the overall amount of propagation time. There is nothing about long propagation times *per se* that indicts the utility of the resultant products; as the propagation proceeds, the uncertainties in the propagated covariance grow appropriately to represent the expected uncertainties of the states at the propagation termination, which is here TCA. Additionally, adjustments to the covariance to account for atmospheric density forecast error and satellite frontal area uncertainty are included and contribute to the propagated covariance's growth (discussed in Appendix N). Finally, a large covariance matrix will in most cases depress the



calculated Pc value, naturally correcting for a situation in which the state uncertainties are so large that it may not be comfortable to use them as a basis of action. Given this dynamic, the lack of recent tracking data invalidates neither the original fit (presuming there was sufficient tracking in the fit-span to produce an acceptable fit) nor the propagation to TCA.

Nonetheless, a rule of thumb that has been used over the last few decades for orbit-determination batch updates is that the propagation interval should not be longer than the fit-span or LUPI. It is generally considered dangerous to try to propagate longer than the span of measurement data that directed the orbit-determination update. While there are undoubtedly situations in which propagations of this type could still be representative, it is believed that the reasonable span of the orbit determination has been exceeded in such a case, and it is thus reasonable to refrain from any mitigation action for conjunctions for which either object's orbit determination at TCA was generated with so extended a propagation interval.

P.13 Excessive Atmospheric Density Error Due to Solar Storms

As explained in Appendix N, for some years, the USSPACECOM orbit-determination process has adjusted the covariance in the CDM for anticipated atmospheric density forecast error through the application of a consider parameter. The Jacchia-Bowman-HASDM 2009 (JBH09) atmospheric density model used by USSPACECOM contains special functionality to reduce propagation error in the presence of Coronal Mass Ejections (CMEs), a commonly encountered type of solar storm. A special model given the name “Anemomilos” attempts to predict the size, time of arrival, and geoeffectiveness of CMEs, and the atmospheric density is modified appropriately during the prediction interval to reflect the storm's predicted arrival and effects. (Tobiska et al. 2013) Paired with the appropriate consider parameter value, which will compensate for the expected uncertainty in this type of modeling, the predicted product should in principle account for atmospheric density mismodeling. The fact that the solar environment is perturbed should not, therefore, result in an *a priori* conclusion that the predictions of an event's severity are errant and that a claim of event inactionability is appropriate.

Some additional insight into such situations can be gained by constructing what are called Space Weather Trade Space (SWTS) plots that explore how the collision likelihood of a particular event would change if the atmospheric density were mismodeled. Because atmospheric density and ballistic coefficient are multiplicatively coupled in the atmospheric drag acceleration equation, changing one of these parameters has the same effect as changing the other. It is thus possible to vary the primary and secondary objects' drag parameters to understand the effect of drag modeling on the event's conjunction likelihood.

The plots are constructed by varying the nominal ballistic coefficient by half an order of magnitude in each direction with the x-axis giving the primary object ballistic coefficient variation and the y-axis that for the secondary object and the color indicating the resultant Pc value, but these details are not so important because it is the overall morphology of the figure that is meaningful. In Figure P-1, despite large variations in the ballistic coefficient (as a proxy for atmospheric density), the Pc changes little. This event is not sensitive to space weather mismodeling, and thus actionability is not affected by solar storm developments.

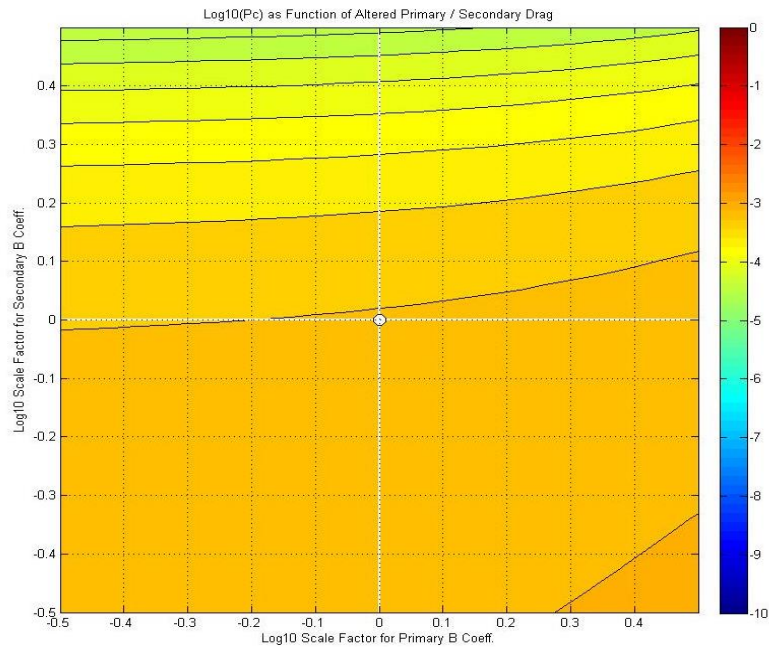


Figure P-1 SWTS Plot for Invariant Situation¹⁹

In Figure P-2, the current P_c , as represented by the center point of the figure, lies on a “ridge” of high P_c values; any atmospheric drag mismodeling that moves the situation off the center point will result in a lower rather than higher P_c . So, if the current P_c is below the threshold that would require a mitigation action, the event will still be safe even in the presence of model error.

¹⁹ From Hejduk et al. 2017.




Figure P-2 SWTS Plot for “Ridge” Situation²⁰

Finally, in situations in which the P_c is not on a ridge and where non-trivial variation is observed, then model error could either increase or decrease the P_c as shown in Figure P-3. In such cases, one is left without the reassurance of either of the two previous examples. However, given that the estimation process attempts to take account of drag uncertainty through modeling even of solar storms, the CDM result is presumed to represent the best estimate possible and, in the absence of other indicting data, is expected to be reliable and actionable.

²⁰ From Hejduk et al. 2017.

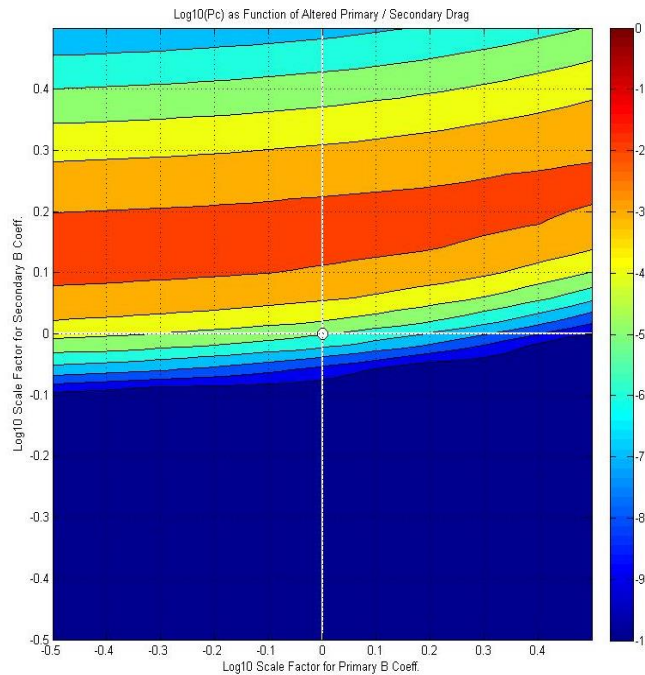


Figure P-3 SWTS Plot for Variation Situation²¹

P.14 Summary

As the previous discussion indicates, the orbit determination and propagation process nearly always produces data that can be considered actionable. There are certain indications that an orbit-determination fit may potentially fall short, and an even smaller number that would counsel the dismissal of an orbit determination as the basis for CA risk assessment. It is similarly infrequent that the propagation results to TCA can be questioned, and these situations can usually be addressed in other ways (e.g., the repair of a non-positive-definite covariance). In both situations, there are additional one-sided tests that, if passed, can reassure the user about the propriety of the data and its ability to subtend CA risk assessment.

²¹ From Hejduk et al. 2017.



P.15 Appendix P Annex

This annex to Appendix P provides specific thresholds for the orbit-determination quality checks.

Table P-1 Geopotential, Atmospheric Drag, and Solar Radiation Pressure

Perigee Height (km)	Eccentricity	Object Type	Geopotential Zonal and Tesseral	Solve for Drag?	Solve for Solar Rad. Pressure?
0 < p < 500	< 0.25	Payload	36	Y	N
		Rocket Body / Platform	36	Y	N
		Debris/Unknown	36	Y	N
500 < p < 900	< 0.25	Payload	36	Y	Y
		Rocket Body / Platform	36	Y	Y
		Debris/Unknown	36	Y	Y
900 < p < 2000	< 0.25	Payload	24	Y	Y
		Rocket Body / Platform	24	Y	Y
		Debris/Unknown	24	Y	Y
0 < p < 500	> 0.25	Payload	36	Y	N
		Rocket Body / Platform	36	Y	N
		Debris/Unknown	36	Y	N
500 < p < 1000	> 0.25	Payload	24	Y	Y
		Rocket Body / Platform	24	Y	Y
		Debris/Unknown	24	Y	Y
1000 < p < 2000	> 0.25	Payload	18	N	Y
		Rocket Body / Platform	18	N	Y
		Debris/Unknown	18	N	Y
2000 < p < 10000	any	Payload	12	N	Y
		Rocket Body / Platform	12	N	Y
		Debris/Unknown	12	N	Y
p > 2000	any	Payload	8	N	Y
		Rocket Body / Platform	8	N	Y
		Debris/Unknown	8	N	Y

Table P-2 Ballistic Coefficient and Solar Radiation Pressure Coefficient Reasonability

Energy Dissipation Rate (w/kg)	Object Type	Ballistic Coefficient		Solar Radiation Pressure Coefficient	
		Min	Max	Min	Max
< 0.0006	Payload	0.001	0.1	0.001	0.1
	Rocket Body / Platform	0.001	0.2	0.001	0.2
	Debris/Unknown	0.001	1	0.001	1
> 0.0006	Payload	0.001	0.1	0.001	0.1
	Rocket Body / Platform	0.001	0.2	0.001	0.2
	Debris/Unknown	0.001	1	0.001	1



Table P-3 LUPI Minimum and Maximum Values

EDR Min	EDR Max	e	Length of Update Interval (days)	
			Lower Bound	Upper Bound
0	0	any	14	See graph
0+	0.0006	< 0.25	3.5	18
		> 0.25	14	See graph
0.0006	0.001	any	1.5	17
0.001	0.0015	any	1.5	15
0.0015	0.002	any	1.5	14
0.002	0.003	any	1.5	12
0.003	0.006	any	1.25	11
0.006	0.009	any	1.25	10
0.009	0.015	any	1.25	8
0.015	0.05	any	1.25	8
0.05	0.05+	any	1.25	7

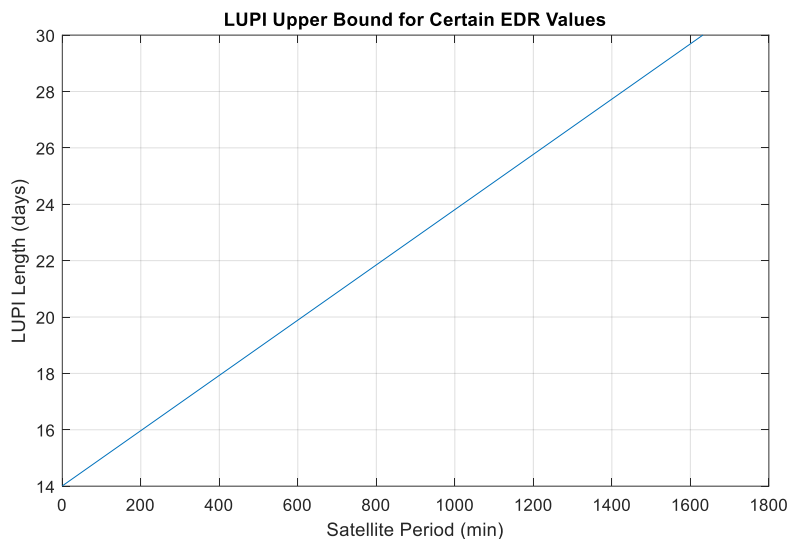


Figure P-4 LUPI Upper Bound for Certain EDR Values

- **Residual Acceptance.** Fewer than 80% residual acceptance would generally prompt manual review of an orbit determination.
- **Weighted Orbital-Determination Residual Root-Mean-Square.** Weighted root-mean-square values exceeding the following thresholds are considered excessive:
 - Payload: 1.5
 - Rocket Body / Platform: 2.0
 - Debris/Unknown: 5.0

No thresholds for sub-unity weighted root-mean-square levels have yet been established.



Appendix Q. List of Works Cited

USG References

NPR 7020.7, NASA Information Technology Program and Project Management Requirements

NPR 8715.6, NASA Procedural Requirements for Limiting Orbital Debris and Evaluating the Meteoroid and Orbital Debris Environments

NASA Standard 8719-14B, Process for Limiting Orbital Debris, which defines how NASA implements the ODMSP

NASA Examples of Information to Expedite Review of Commercial Operator Applications to Regulatory Agencies. May be downloaded from:

<https://www.nasa.gov/recommendations-commercial-space-operators>

U.S. Government Orbital Debris Mitigation Standard Practices, November 2019 Update.

https://orbitaldebris.jsc.nasa.gov/library/usg_orbital_debris_mitigation_standard_practices_november_2019.pdf

Presidential Memorandum. Space Policy Directive-3, National Space Traffic Management Policy. June 18, 2018.

18th Space Control Squadron Combined Force Space Component Command. *Spaceflight Safety Handbook for Satellite Operators: 18 SPCS Processes for On-Orbit Conjunction Assessment & Collision Avoidance*. Version 1.5, August 2020. Vandenberg Air Force Base, California.

Authored and Other References

1. Akella, M. R. and K. T. Alfriend. 2000. "The Probability of Collision Between Space Objects." *Journal of Guidance, Control and Dynamics*, Vol. 23 No. 5 (September-October 2000) pp. 769-772.
2. Alfano, S. 2004. "Orbital Covariance Interpolation." 14th AAS/AIAA Space Flight Mechanics Conference, Maui HI, February 2004.
3. Alfano, S. 2005a. "A Numerical Implementation of Spherical Object Collision Probability." *Journal of the Astronautical Sciences*, Vol. 53 No. 1 (January-March 2005), pp. 103-109.
4. Alfano, S. 2005b. "Relating Position Uncertainty to Maximum Conjunction Probability." *Journal of the Astronautical Sciences*, Vol. 53 No. 2 (April-June 2005), pp. 193-205.
5. Baars, L., D. Hall, and S. Casali. 2019. "Assessing GEO and LEO Repeating Conjunctions using High-Fidelity Brute Force Monte Carlo Simulations." 2019 AAS Astroynamics Specialist Conference (Paper #19-612), Portland ME, August 2019.
6. Balch, M. S., R. Martin, and S. Ferson. 2019. "Satellite Conjunction Assessment and the False Confidence Theorem." *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 475, Issue 2227, 2019. (<https://doi.org/10.1098/rspa.2018.0565>).



7. Carpenter, J. R. and F. L. Markey. 2014. "Wald Sequential Probability Ratio Test for Space Object Conjunction Assessment." *Journal of Guidance, Control, and Dynamics* Vol. 37 No. 5 (September-October 2014), pp. 1385-1396.
8. Casali, S., D. Hall, D. Snow, M. Hejduk, L. Johnson, B. Skrehart, and L. Baars. 2018. "Effect of Cross-Correlation of Orbital Error on Probability of Collision Determination." 2018 AAS/AIAA Astrodynamics Specialist Conference (Paper #18-272), Snowbird UT, August 2018.
9. CCSDS 508.0-B-1 Conjunction Data Message (CDM) Blue Book
10. CCSDS 502.0-B-2 Orbit Data Messages Blue Book
11. Chan, F. C. 2008. *Spacecraft Collision Probability*. El Segundo, CA: The Aerospace Press.
12. Chan, F. C. 2015. "Formulation of Collision Probability with Time-Dependent Probability Distribution Functions," AAS/AIAA Space Flight Mechanics Meeting (Paper #15-233), Williamsburg, VA, January 2015.
13. Chan, F. C. 2019. "Collision Probability for Polygonal Cross Sections." 2019 AAS Astrodynamics Specialist Conference (Paper #19-911), Portland ME, August 2019.
14. Clopper, C. and E. S. Pearson. 1934. "The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial." *Biometrika*, Vol. 26 (1934), pp 404-413.
15. Coppola, V.T. 2012. "Including Velocity Uncertainty in the Probability of Collision between Space Objects." AAS/AIAA Spaceflight Mechanics Meeting, Charleston SC, Paper 12-247, February 2012.
16. D'Agostino, R. B. and M. A. Stephens. 1986. *Goodness-of-Fit Techniques*. New York: Marcel Dekker, Inc.
17. DeMars, K., Y. Cheng, and M. Jah. 2014. "Collision Probability with Gaussian Mixture Orbit Uncertainty." *Journal of Guidance, Control, and Dynamics*, Vol. 37 No. 3 (2014), pp. 979-985.
18. Elrod, C. 2019. "Computational Bayesian Methods Applied to Complex Problems in Bio and Astro Statistics." Doctoral Dissertation, Baylor University, July 2019.
19. Foster, J. L. and H. S. Estes. 1992. "A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles." NASA/JSC-25898 (August 1992).
20. Frisbee, J. 2015. "An Upper Bound on High-Speed Satellite Collision Probability when only One Object has Position Uncertainty Information." 2015 AAS/AIAA Astrodynamics Specialist Conference (Paper #15-717), Vail CO, August 2015.
21. Ghrist, R. W., and D. Plakalovic. 2012. "Impact of Non-Gaussian Error Volumes on Conjunction Assessment Risk Analysis." 2012 AIAA/AAS Astrodynamics Specialist Conference, Minneapolis MN, August 2012.
22. Grubbs, F.E. 1969. "Procedures for detecting Outlying Observations in Samples." *Technometrics* 11 (1969), 1-21.



23. Hall, D.T., M. D. Hejduk, and L. C. Johnson. 2017a. “Time Dependence of Collision Probabilities During Satellite Conjunctions.” AAS Space Flight Mechanics Meeting (Paper # 17-271), San Antonio, TX, February 5-9, 2017.
24. Hall, D.T., M. D. Hejduk, and L. C. Johnson. 2017b. “Remediating Non-Positive-Definite State Covariances for Collision Probability Estimation.” 2017 AAS/AIAA Astrodynamics Specialist Conference, Stevenson WA, August 2017.
25. Hall, D.T., S. J. Casali, L. C. Johnson, B. B. Skrehart, and L. G. Baars. 2018. “High-Fidelity Collision Probabilities Estimated using Brute Force Monte Carlo Simulations.” 2018 AAS/AIAA Astrodynamics Specialist Conference (Paper #18-244), Snowbird UT, August 2018.
26. Hall, D.T. 2019a. “Determining Appropriate Risk Remediation Thresholds from Empirical Conjunction Data using Survival Probability Methods.” 2019 AAS Astrodynamics Specialist Conference (Paper #19-631), Portland ME, August 2019.
27. Hall, D.T. 2019b. “Implementation Recommendations and Usage Boundaries for the Two-Dimensional Probability of Collision Calculation.” 2019 AAS Astrodynamics Specialist Conference (Paper #19-632), Portland ME, August 2019.
28. Hametz, M. E. and B. A. Beaver. 2013. “A Geometric Analysis to Protect Manned Assets from Newly Launched Objects – COLA Gap Analysis.” AAS 13-360.
29. Hejduk, M.D. 2008. “Space Catalogue Accuracy Modeling Simplifications.” 2008 AAS Astrodynamics Specialists Conference, Honolulu, HI, August 2008.
30. Hejduk, M.D., S. J. Casali, D. A. Cappellucci, N. L. Ericson, and D. E. Snow. 2013. “A Catalogue-Wide Implementation of General Perturbations Orbit Determination Extrapolated from Higher-Order Theory Solutions.” 2013 AAS/AIAA Space Flight Mechanics Meeting, Kauai, HI. February 2013.
31. Hejduk, M.D., D. Plakalovic, M. E. Hametz, L. K. Newman, J. C. Ollivierre, B. A. Beaver, and R. C. Thompson. 2014. “Launch COLA Operations: An Examination of Data Products, Procedures, and Thresholds, Revision A.” NASA Special Technical Report. 12 January 2014.
32. Hejduk, M.D., L. K. Newman, R. L. Besser, and D. A. Pachura. 2015. “Predicting Space Weather Effects on Close Approach Events.” 2015 AMOS Technical Conference, Kihei HI, September 2015.
33. Hejduk, M.D., F. Laporte, M. Moury, T. S. Kelso, L. K. Newman, and R. Shepperd. 2017. “Consideration of Collision ‘Consequence’ in Satellite Conjunction Assessment and Risk Analysis.” 26th International Symposium on Space Flight Dynamics, Matsuyama, Japan, June 2017.
34. Hejduk, M.D., D. E. Snow, and L. K. Newman. 2019. “Satellite Conjunction Assessment Risk Analysis for ‘Dilution Region’ Events: Issues and Operational Approaches.” Fifth Annual Space Traffic Management Conference Progress through Collaboration, Austin TX, February 2019.



35. Hejduk, M.D. and D. E. Snow. 2019. "Satellite Conjunction 'Probability,' 'Possibility,' and 'Plausibility': A Categorization of Competing Conjunction Assessment Risk Analysis Paradigms." 2019 AAS/AIAA Astrodynamics Specialist Conference (Paper #19-652), Portland ME, August 2019.
36. Hoots, F. R. and R. L. Roehrich. 1980. "Spacetrack Report No.3 – Models for Propagation of NORAD Element Sets," December 1980.
37. Jenkin, A.B., J. P. McVey, and G. E. Peterson. 2020. "Solution of long-coast re-entry COLA problems with COLA gap methods" *Journal of Space Safety Engineering* Volume 7, Issue 2, June 2020, Pages 105-112.
38. Johnson, N.L., P. H. Krisko, J. C. Liou, and P. D. Anz-Meador. 2001. "NASA's New Breakup Model of EVOLVE 4.0." *Advances in Space Research*, Vol. 28 No. 9 (2001), pp. 1377-1384.
39. Kaplan, S. and B. Garrick. 1981. "On the Quantitative Definition of Risk." *Risk Analysis*, Vol. 1 No. 1 (1981), pp. 11-27.
40. Lechtenberg, T. F. and M. D. Hejduk. 2019. "Assessment and Validation of Collision 'Consequence' Method of Assessing Orbital Regime Risk Posed by Potential Satellite Conjunctions." 2019 Hypervelocity Impact Symposium, Destin FL, March 2019, HVIS2019-061.
41. Lechtenberg, T. 2019a. "An Operational Algorithm for evaluating Satellite Collision Consequence." 2019 AAS Astrodynamics Specialist Conference (Paper #19-669), Portland ME, August 2019.
42. Lechtenberg, T., 2019b. "Multivariate Normality of Cartesian-Framed Covariances: Evaluation and Operational Significance." 2019 AAS Astrodynamics Specialist Conference (Paper #19-671), Portland ME, August 2019.
43. Mashiku, A.K. and M. D. Hejduk. 2019. "Recommended Methods for Setting Mission Conjunction Analysis Hard Body Radii." 2019 AAS Astrodynamics Specialist Conference (Paper #19-702), Portland ME, August 2019.
44. Patera, R.P. 2001. "General Method for Calculation Satellite Collision Probability." *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 4 (July-August 2001), pp. 716-722.
45. Rosenberg, F. 2003. "Radar Characterization of RORSAT Coolant Spheres and their Potential Use for Calibration and Validation of Debris Tracking Sensors." MIT Lincoln Laboratory Project Report STK-264, November 2003.
46. Rosner, B. 1975. "On the Detection of Many Outliers." *Technometrics* 17 (1975), 221-227.
47. Rosner, B. 1977. "Percentage Points for the RST Many Outlier Procedure." *Technometrics* 19 (1977), 307-312.



48. Sabol, C., T. Sukut, K. Hill, K. Alfriend, B. Wright, Y. Li, and P. Schumacher. 2010. "Linearized orbit covariance generation and propagation analysis via simple Monte Carlo simulations." AAS/AIAA Space Flight Mechanics Meeting (paper AAS-10-134), San Diego, CA, February 2010.
49. Seitzer, P. 2020. (U MI) "Large Constellations of LEO Satellites and Astronomy." Advanced Maui Optical and Space (AMOS) Surveillance Technologies Conference, September 15-18, 2020.
50. Shelton, C.T. and J. L. Junkins. 2019. "Probability of Collision between Space Objects Including Model Uncertainty." *Acta Astronautica*, Vol.155 (2019), pp.462-471.
51. Schilling, B., Y. Taleb,, J. Carpenter, M. Balducci, and T. Williams. 2016. "Operational Experience with the Wald Sequential Probability Ratio Test for Conjunction Assessment from the Magnetospheric MultiScale Mission." 2016 AIAA/AAS Astrodynamics Specialist Conference, Long Beach CA, August 2016.
52. Tobiska, W.K., D. Knipp, W. J. Burke, D. Bouwer, J. Bailey, D. Odstrcil, M. P. Hagan, J. Gannon, and B. R. Bowman. 2013. "The Anemomilos Prediction Methodology for Dst." *Space Weather* (Vol. 11, No. 9 (September 2013) pp 490-508. DOI: 10.1002/swe.20094.
53. Woodburn, J. and S. Tanygin. 2014. "Coordinate Effects on the use of Orbit Error Uncertainty." International Symposium on Space Flight Dynamics 24 (Baltimore, MD), May 2014.

National Aeronautics and Space Administration

Headquarters
300 E Street SW
Washington, DC 20546

www.nasa.gov

NASA/SP-20205011276