

National Aeronautics and Space Administration



NASA Spacecraft Conjunction Assessment and Collision Avoidance Best Practices Handbook

Volume 2 - Technical Appendices

Office of the Chief Engineer

NASA-SP-20205011318/REV2-VOL2

August 2026

NASA-SP-20205011318/REV2-VOL2

This volume continues the NASA Spacecraft Conjunction Assessment and Collision Avoidance Best Practices Handbook by providing the technical background material and references supporting the best practices (contained in Volume 1, Best Practices.) The two volumes may be updated separately. For convenience, the appendices for acronyms and the glossary are duplicated between the volumes.

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

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

Table of Contents	iv
Change Log	viii
Appendix A. Master List of Acronyms.....	1
Appendix B. Glossary.....	4
Appendix E. Use of Analytic Theory Orbital Data in Conjunction Assessment....	7
Appendix F. Expected Conjunction Event Rates	11
Appendix G. Orbital Debris Density	17
Appendix H. Long-Term Collision Risk and Satellite Colocation Analysis	22
Appendix I. Satellite Covariance Realism Assessment Procedures	32
Appendix J. CARA Conjunction Risk Assessment Tools Validation	44
Appendix K. R-15 Message	45
Appendix L. Commercial Data in NASA Conjunction Assessment	48
Appendix M. Use of the Probability of Collision (Pc) as the Risk Assessment Metric for Conjunction Assessment	53
Appendix N. Pc Calculation Approaches.....	61
Appendix O. Collision Consequence	105
Appendix P. Event Actionability	116
Appendix Q. Notional Display Flow for Conjunction Event Processing.....	130
Appendix R. List of Works Cited	144

The lists of best practices (Appendices C and D) are available in Volume 1, Best Practices.

List of Figures

Figure E-1 Miss-Distance Distributions for Critical ($P_c > 1E-04$) and Non-critical ($P_c < 1E-04$) Events	8
Figure E-2 eGP and Regular SGP4 Errors: LEO > 500 km (log scale)	9
Figure E-3 Expansion of Figure E-1 to Include CDFs for P_c Thresholds of $1E-03$ and $1E-02$	10
Figure F-1: Event Rates, 0.5 Day MCP, Starlink Included	13
Figure F-2: Event Rates, 1 Day MCP, Starlink Included.....	14
Figure F-3: Event Rates, Highest P_c in Event, Starlink Included	14
Figure F-4: Event Rates, 0.5 Day MCP, Starlink Excluded	15
Figure F-5: Event Rates, 1 Day MCP, Starlink Excluded	15
Figure F-6: Event Rates, Highest P_c in Event, Starlink Excluded	16
Figure G-1 Orbiting Object Flux in 2025 as a Function of Orbit Inclination, Orbit Altitude, and Object Size.....	20
Figure H-1 Kessler Collision Rates between NewSat and HST	25
Figure H-2 Total NewSat Collision Rates Among All Cataloged Satellites	26
Figure H-3 NewSat Collision Rates Excluding the Starlink Constellation	26
Figure H-4 NewSat Conjunction Rates with and without Starlink.....	28
Figure H-5 Predicted Rates for NewSat and HST Satellite Conjunctions	29
Figure H-6 Predicted Rates for NewSat and FGST Satellite Conjunctions	30
Figure H-7 Predicted Rates for NewSat and Swift Satellite Conjunctions	31
Figure N-1 Relative Positions and Uncertainties of Two Objects	62
Figure N-2 Combining the Sizes of the Primary and Secondary Objects	63
Figure N-3 Geometry at TCA.....	63
Figure N-4 Two Equivalent Representations of the Conjunction Plane	64
Figure N-5 P_c Rate and Cumulative P_c for a Non-Rectilinear Conjunction	69
Figure N-6 P_c Rate and Cumulative P_c for a Low Velocity Conjunction.....	70
Figure N.7 Schematic Flowchart of the Multistep P_c Algorithm	72
Figure N-8 Comparison of Monte Carlo Collision Probabilities with the Two-Dimensional P_c Method (left), the Three-Dimensional N_c Method (center), and the Two-Dimensional N_c Method (right) for a Large Set of CARA Conjunctions	74

Figure N-9 Cumulative Three-Dimensional N_c Risk Over Four Repeating Conjunctions.....	76
Figure N-10 Behavior of Relative Density Error by Perigee Height and Solar Activity	78
Figure N-11 Two-phase Application of Consider Parameters	80
Figure N-12 Profiles of 250 Conjunctions with Primary and Secondary Satellites of Non-negligible Drag	83
Figure N-13 Three-Dimensional N_c Temporal Risk Plot with Monte Carlo from Epoch Result Overlay (in Pink).....	95
Figure N-14 Mismatch between Elongated In-track Covariances and Forced Cartesian Ellipsoidal Representation	97
Figure N-15 Comparative Results of 373 High- P_c Conjunctions	99
Figure N-16 Comparative Results of 63,603 Conjunction Events	100
Figure O-1 Range of Debris Production Possibilities.....	107
Figure O-2 Fragment Production Levels for 1 km/s Relative Velocity.....	108
Figure O-3 CDF of the Ratio of Calibrated HBR Values to Known HBR Values for 586 Satellites	110
Figure O-4 CDF of the Ratio of Calibrated RCS+BC Mass Estimates to Known Masses for 554 Satellites.....	111
Figure O-5 CDF of the Ratio of Calibrated RCS+SRPC Mass Estimates to Known Masses for 303 Satellites	112
Figure O-6 Collision and Fragmentation Probabilities for Known-on-Unknown Conjunctions.....	114
Figure P-1 SWTS Plot for Invariant Situation.....	126
Figure P-2 SWTS Plot for “Ridge” Situation	126
Figure P-3 SWTS Plot for Variation Situation	127
Figure P-4 LUPI Upper Bound for Certain EDR Values	129
Figure Q-1 Ground Trace Plot	133
Figure Q-2 Conjunction Plane Plot.....	134
Figure Q-3 Probability of Collision Time-History Plot.....	135
Figure Q-4 Total Miss Distance Plot.....	136
Figure Q-5 Componentized Miss Distance Plot	137
Figure Q-6 State Comparison Consistency Plot	139

Figure Q-7 Solar Storm Prediction	140
Figure Q-8 Space Weather Trade-Space Plot	141
Figure Q-9 Maneuver Trade-Space Plot	142

List of Tables

Table P-1 Geopotential, Atmospheric Drag, and Solar Radiation Pressure	128
Table P-2 Ballistic Coefficient and Solar Radiation Pressure Coefficient Reasonability	128
Table P-3 LUPI Minimum and Maximum Values	129
Table Q-1 Primary Satellite Vital Statistics	131
Table Q-2 Secondary Satellite Vital Statistics	132

Change Log

Date	Description	Revision
Aug 2026	Initial release as the second volume comprising the set of technical appendices (including Appendix A and B for convenience.) Updated Appendix G <i>Orbital Debris Density</i> with data from 2025. Updated Appendix M <i>Use of the Probability of Collision (Pc) as the Risk Assessment Metric for Conjunction Assessment</i> . Updated Appendix N <i>Pc Calculation Approaches</i> . Added to Appendix O <i>Collision Consequence</i> a section (new O.5) on availability of satellite size and mass estimates	—

Appendix A. Master List of Acronyms

2D	Two-Dimensional
3D	Three-Dimensional
ASW	Astroynamics Support Workstation (tool)
BC	Ballistic Coefficient
BRDF	Bidirectional Reflectance Distribution Function
CARA	Conjunction Assessment Risk Analysis (Program)
CCSDS	Consultative Committee for Space Data Systems
CDF	Cumulative Distribution Function
CDM	Conjunction Data Message
COLA	Collision on Launch Assessment
COMBO	Computation of Miss Between Orbits (algorithm)
CONOPS	Concept of Operations
DISCOS	Database and Information System Characterizing Objects in Space
DOD	Department of Defense
EDR	Energy Dissipation Rate
FOD	Flight Operations Directorate at NASA JSC
GEO	Geosynchronous Orbit
GPS	Global Positioning System
HAC	High Accuracy Catalog
HBR	Hard-Body Radius
HEO	High Earth Orbit
HST	Hubble Space Telescope
ISS	International Space Station
JSC	Johnson Space Center
LCOLA	Launch Collision Avoidance
LEO	Low Earth Orbit
LUPI	Length of Update Interval

MADCAP	Multi-mission Automated Deep-space conjunction Assessment Process
MC	Monte Carlo
MVN	Multi-Variate Normal (function)
Nc	Number of Collisions (statistically expected)
NASA	National Aeronautics and Space Administration
Nmc	Number of Monte Carlo (trials)
NPR	NASA Procedural Requirements
ODM	Orbit Data Messages (CCSDS standard)
ODMSP	Orbital Debris Mitigation Standard Practices
ODPO	Orbital Debris Program Office
O/O	Owner/Operator
Pc	Probability of Collision
PDF	Probability Density Function
R-15	Ready Minus 15 (days, message format)
RCS	Radar Cross Section
SDK	Software Development Kit
SDS	Space Defense Squadron (U.S. Space Force)
SGP4	Simplified General Perturbation Theory #4
SPD-3	Space Policy Directive-3, the National Space Traffic Management Policy
SRP	Solar Radiation Pressure
SRPC	Solar Radiation Pressure Coefficient
SSA	Space Situational Awareness
SSN	Space Surveillance Network
SWTS	Space Weather Trade Space
TBMC	Two-Body Monte Carlo
TCA	Time of Closest Approach
TLE	Two-Line Element
URL	Uniform Resource Locator

U.S.	United States
USSPACECOM	U.S. Space Command
VIM	Vehicle Information Message
VSFB	Vandenberg Space Force Base
WRMS	Weighted Root Mean Square

Appendix B. Glossary

Cislunar Space. Cislunar space is the three-dimensional volume of space beyond Earth's geosynchronous orbit that is mainly under the gravitational influence of the Earth and/or the Moon. Cislunar space includes the Earth-Moon Lagrange point regions (defined below), trajectories utilizing those regions, and the Lunar surface. (Taken from the 2002 "National Cislunar Science & Technology Strategy" from the National Science & Technology Council's Cislunar technology strategy Interagency Working Group.)

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

Collision Avoidance. The planning and execution of risk mitigation strategies to avoid a collision between two space objects. (See also conjunction mitigation.).

Colocation. Being in the vicinity of another operational spacecraft close enough that systematic conjunctions occur.

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 protected asset (also called "primary object").

Conjunction analysis. The process of predicting a close-approach event by screening the ephemeris of the protected asset against the space object catalog (i.e., conjunction assessment) and then analyzing the event to determine the associated threat to the asset (risk assessment).

Conjunction Assessment. The identification of close approaches using ephemeris screening against a catalog of 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 secondary owner/operators so that they can perform an avoidance maneuver.

Conjunction Risk Assessment. The process of assessing a conjunction (predicted close approach) to determine the likelihood of two space objects colliding and to determine the expected consequence if they collide in terms of spacecraft inoperability and expected debris production.

Covariance. Characterization of uncertainty components and their interactions surrounding a space object's estimated state at a given time.

Demise. When a spacecraft is no longer in space such as a complete burn up on reentry. Spacecraft passivation and migration to disposal orbits occur prior to demise.

Disposal. An end-of-mission process involving a spacecraft's passivation and reentry into the atmosphere for its ultimate demise or its movement (if necessary) to an orbit or trajectory considered acceptable for orbital debris limitation. For purposes of this NPR, "disposal" includes

the reorbiting or deorbiting of a spacecraft at the end of mission. Disposal occurs prior to spacecraft demise.

Ephemeris (plural: ephemerides). A file containing a time-ordered set of position and velocity measurements describing an object's predicted trajectory.

Geosynchronous 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 USSPACECOM, which provides conjunction assessment screenings.

Hard-body radius. The radius of a circle equal to the sum of the circumscribing radii of both the protected plus the secondary spacecraft.

High Earth Orbit (HEO). An orbit with period greater than 225 minutes and eccentricity greater than 0.25. Note that this definition is chosen to align with the definition used by USSPACECOM, which provides conjunction assessment screenings.

Low Earth Orbit (LEO). An orbit with an orbital period less than 225 minutes and eccentricity less than 0.25. Note that this definition is chosen to align with the definition used by USSPACECOM, which provides conjunction assessment screenings.

Maneuver Plan. The specific parameters that represent a planned spacecraft maneuver, including execution time, burn duration, and delta- v . The position and velocity of a single object at a specified epoch is specified in an Orbit Parameter Message (OPM). The industry message standard for communicating maneuver plans is the Consultative Committee for Space Data Systems (CCSDS) standard CCSDS 502.0-B-3, "Orbit Data Messages," Section 3 located at <https://ccsds.org/publications/bluebooks/entry/3073/>.

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

Maneuverable spacecraft. A spacecraft that has capability permitting the manipulation of the spacecraft's trajectory in a non-Keplerian fashion.

Non-cooperative tracking data. Data obtained describing the location of space objects that do not themselves actively provide location data. For example, a piece of orbital debris that has no transponder could not provide ephemeris data. Other examples include natural objects such as asteroids, and spacecraft belonging to operators who do not choose to share an ephemeris for their spacecraft.

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

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

Orbital debris. In this NPR, orbital debris is defined as any object placed in space by humans that no longer serves any useful function. Objects range from spacecraft to spent launch vehicle

stages to components and include materials, fragments, trash, or other objects that are intentionally or inadvertently cast off or generated. (Derived from NPR 8715.6.)

Payload. A specific complement of instruments, sensors, equipment, and support hardware carried into space to accomplish a mission or a discrete activity in outer space. Personnel are not considered a payload or a part of a payload.

Protected asset. The asset of focus (also called “primary object”) for which risk assessments of potential collisions are being performed and collision avoidance mitigation activities are being considered and possibly performed.

Resident space object. An artificial object that orbits Earth.

Safety volume. A volume defined around a protected asset as a zone for predicting close approaches with a “secondary” space object entering that volume.

Secondary object. Any cataloged object residing in space that passes through the safety volume of a protected (i.e., primary) asset when the protected asset with its safety volume is flown along its trajectory. The resident space object is identified as a “secondary object” and is the conjuncting object in a predicted future close-approach event. The secondary object may be an operated spacecraft.

Secondary payload. A payload that is manifested subordinate to a primary or co-manifested payload and is, therefore, subordinate in launch date and orbit selection. (See also “payload,” “primary payload,” and “rideshare.”)

Space Flight Mission. NASA space flight programs, projects, and activities (including spacecraft, launch vehicles, payloads, and 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.

Space situational awareness. The knowledge and characterization of space objects and their operational environment to support safe, stable, and sustainable space activities. (From Space Policy Directive-3.)

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 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 for conjunction assessment, 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 Object Catalog, although higher-order theory data can also be furnished from other sources) should be employed in the conjunction 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 N 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 through a covariance matrix, at the time of closest approach (TCA).

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 (or other statistically sophisticated) conjunction assessment calculation and thus for conjunction assessment applications will yield only a miss distance at the time of the two objects' TCA. Because no uncertainty information on this miss distance is provided, no probabilistic conclusion can be drawn. A miss-distance threshold can be (arbitrarily) selected to enable a binary "safe / not safe" conjunction evaluation, but there is no ability to determine, for example, how often the estimated 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. The lack of any strong correlation between P_c and miss distance is clear: the span of miss distances for the critical events (less than 100 m to more than 10 km) is extremely broad and overlaps significantly with the results for the non-critical events. For example, to eliminate 95% of the high- P_c events using a miss-distance criterion alone, a 10 km miss-distance threshold would need to be chosen—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 conjunction assessment. 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, the latter of 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. Most, but not all, TLEs are generated with this latter method. Performance data for the orbital regime corresponding to Figure E-1 above (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 timeframe during which most conjunction risk mitigation 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)¹

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, $1E-02$. Figure E-3 suggests 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 conjunction assessment: remediate any conjunctions with a miss distance less than 1-2 km. This move is ill-advised 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 $1E-03$). In addition, 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 conjunction risk assessment is not tenable. Thus, 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.

¹ 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 data, usually furnished as TLEs, does have its uses in conjunction assessment; some of these have included the following:

- 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 maneuverability of a payload that appears as a conjuncting secondary, as such 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 sufficient to enable their repair, can at least make clear that such cases constitute non-actionable conjunction assessment data.

Although the above uses of analytic theory data can be tangentially helpful to the conjunction assessment enterprise, 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 Conjunction Event Rates

The conjunction assessment 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 to model spacecraft fuel consumption, identify staffing needs, and determine the appropriate degree of conjunction assessment-related tool automation. The OCAP (Orbital Conjunction Assessment Plan) effort defined and required in NPR 8079.1, which results in a document produced before design time in NASA mission development, will perform specialty analyses that determine for a particular mission the expected conjunction assessment and risk analysis burden; the graphs contained in this section are intended to provide at least some first-order information for orientation. These data summarize three years' recent conjunction assessment 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). In November 2021, Russia conducted a kinetic anti-satellite (ASAT) test that destroyed a derelict Russian satellite at approximately 480km, which generated an additional 1500 cataloged debris fragments in LEO. This increase is reflected in the three-year averages given below, although by 2024 most of these fragments had decayed.

The graphs presented in the following figures 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. Although different 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). Hard-body radius (HBR) values for the primary satellite of 0.1m, 0.31m, 1m, 3.1m, and 10m are used, with the size of the secondary object assigned by the CARA secondary size assignment mechanism described more fully in Appendix N. All the CARA events were reprocessed with these different HBR 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 a situation in which the 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 in which the 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 do not manifest a $P_c > 1E-07$. Based on this definition, all red and yellow events are also green events.

The above written definitions of the different color risk levels were kept ambiguous in one respect, namely which of the many CDMs that a single event generates is used to determine whether a particular color threshold is met. For most objects, events are discovered at their first screening, which is seven days to TCA for LEO and ten days to TCA for other regimes; but the event color often changes as the event progress (and regresses). The different schools of thought in locating the temporal point of evaluation are given below:

- Maneuver commitment point (MCP). The maneuver commitment point is the time before TCA at which the O/O must make a decision whether to maneuver the spacecraft in order to have enough time to finalize any residual maneuver planning and be able to uplink and execute the maneuver before TCA. This time can also be somewhat inflated to consider contingencies, such as a failed contact during a scheduled contact session. Maneuver commitment points for modern satellites typically fall in the 12–24-hour range, or from 0.5 to 1 days, before TCA. Because this represents the state of the CA event at which a decision must be made, it is perhaps the most logical point at which to evaluate the event severity with the previously described color scheme.
- Any point in event history. For NASA CARA processing, a yellow status identifies an event for further investigation, meaning OD propriety, tracking adequacy, etc., with the opportunity to remedy any shortcomings that may be identified through OD modification and/or requests for additional tracking; and of course, any event that first presented itself as red would receive this level of scrutiny, and then some. So, it is of some interest which events ever manifest P_c values of the yellow and red level throughout their CDM histories.

Each of these two event evaluation criteria is used in the subsequent graphs.

Since the height of each color bar indicates the number of events of that color, the overall height of each bar thus indicates the CDM generation rate. Note that this number, while in some ways a useful leading figure, is a function of the size of the screening volume that is employed and thus is not that meaningful in itself. For a screening volume similar in size to those used by NASA CARA, 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 the volume used is reasonably sized.

There is additionally the question of whether to include events that result from conjunctions with satellites in the Starlink constellation. With over 6000 active satellites presently in orbit, Starlink by far constitutes the largest group of single-purpose satellites; and because officially they would

prefer to take CA-related mitigation actions themselves, it is straightforward for an O/O to enter into an agreement with them to allocate this responsibility to them (NASA has in fact done so through a Space Act Agreement). Because this allocation is a live possibility, the forthcoming graphs show results both including and ignoring Starlink-related conjunctions. It should be pointed out that actual differences on this score for any particular O/O will depend on the degree of overlap between its orbit and a Starlink shell.

Five different orbital regions are represented:

- LEO1: eccentricity < 0.2 and perigee height $< 500\text{km}$,
- LEO2: eccentricity < 0.2 and perigee height between 500km and 750km ,
- LEO3/4: eccentricity < 0.2 and perigee height between 750km and 1400km ,
- HEO: eccentricity > 0.2 ; and
- GEO: eccentricity < 0.2 , orbital period between 1300 and 1800 minutes, and inclination < 35 degrees.

There are notable differences in event frequency among the different regimes, although some of that variation is due to the different levels of sampling among them, given that NASA operates substantially different numbers of payloads in each of these regions. These results are more notional than precise; nonetheless, they can give some level of orientation to the expected loading for each event type (green, yellow, and red).




Figure F-1: Event Rates, 0.5 Day MCP, Starlink Included




Figure F-2: Event Rates, 1 Day MCP, Starlink Included




Figure F-3: Event Rates, Highest Pc in Event, Starlink Included




Figure F-4: Event Rates, 0.5 Day MCP, Starlink Excluded




Figure F-5: Event Rates, 1 Day MCP, Starlink Excluded




Figure F-6: Event Rates, Highest Pc in Event, Starlink Excluded

Appendix G. Orbital Debris Density

Orbital debris can be generated by a number of processes that have conspired to produce the present debris environment:

- Detritus cast off as part of launch and injection activities,
- 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.

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 can thus be helpful to examine the current debris density situation when choosing an operational orbit because the orbital debris density does influence the rate of close approach events as well as the likelihood of satellite damage due to collision with an object too small to be tracked—the threat posed by so-called “lethal non-trackable debris.”

The graphs in Figure G-1 are produced from the NASA Orbital Debris Engineering Model (ORDEM) 3.2 release developed by the NASA Orbital Debris Program Office (ODPO). Generated for the 2025 debris environment, these graphs give orbiting object flux (per satellite unit area per time on orbit) as a function of orbital altitude (a circular orbit is presumed) and inclination. The ORDEM model includes all intact operational and non-operational objects larger than 10 cm (including spacecraft, rocket bodies, payload adapters), not just debris.


Given that satellites often have many years of orbital lifetime, it could be argued that the graphs should be projected further 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 to try to guess at the number, severity, and location of future debris-producing events. Since the purpose here is to compare the relative debris densities of different regimes, the simplicity of a single rather than probabilistic presentation has governed the selection.

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. However, variation with orbital altitude 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 orange (dashed) line represents objects of sizes greater than 10 cm, which is the general size to which the publicly released SSN catalogue is considered complete.
- The grey (solid) line is for objects of sizes greater than 5 cm, which represents a size to which tracking can in some cases be successful if conditions are favorable.
- The blue (dotted) 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 orange (dashed) and blue (dotted) lines thus represents the flux level of objects that are too small to be tracked yet large enough to potentially be lethally damaging. The gap represents a collision risk that cannot be currently mitigated and so is simply accepted as the cost of operating a satellite in space.



**Figure G-1 Orbiting Object Flux in 2025 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 question why conjunction assessment is pursued at all, given that so much of the collision risk cannot be remediated and simply must be accepted. To this question there 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." An unjustified over-investment in safety should be avoided, and mission objectives should not be regularly and substantially impaired by conjunction assessment 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 conjunction assessment operations according to the CARA mission statement. Space actors thus employ the conjunction assessment 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 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 conjunction assessment 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. (Appendix O gives a more detailed treatment of the collision debris production topic.)

Choosing a mission trajectory that occupies a region of space with a lower debris density is thus a favorable selection for two reasons:

- It lowers the risk of a lethal collision with an untracked object; and
- In general, it reduces the rate at which serious conjunctions with cataloged objects will occur.

In addition to achieving the above two goals (as the curves in Figure G-1 demonstrate), choosing orbits with lower orbital altitudes also facilitates compliance with the 25-year disposal guidelines.

While there may be a general correlation, it is important to recognize that 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, yet mitigation actions would be warranted only a few times per year. 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 known) is smaller than the hard-body radius (HBR) 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 serious conjunction events exceeds that projected by the debris flux value: in only 1 out of 10,000 cases would the penetration actually occur 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 conjunction events.)

Appendix H. Long-Term Collision Risk and Satellite Colocation Analysis

This section describes an approach developed by NASA CARA that uses long-term collision risk assessment methodologies applied to proposed satellites to predict:

- Long-term average conjunction and collision rates that a new satellite can expect to experience; and
- The potential frequency of repeating conjunctions (i.e., multiple close approaches on successive orbits) experienced by a new satellite due to being collocated with another satellite.

These methods provide the means to avoid deploying new satellites into crowded regions of orbital space, and to prevent orbital collocation with one (or more) pre-existing satellites. Although the methods also apply to satellite mid-mission repositioning or end-of-life disposal efforts, they are demonstrated here using an example of a newly proposed satellite.

With the ongoing deployment of large constellations, such as the Starlink or OneWeb systems, avoiding overly crowded regions of space is becoming an increasingly important consideration. In addition, avoiding orbital collocation prevents problematic repeating conjunctions with other active satellites. To this end, when a new operational orbit is being selected, a long-term collision risk and orbit collocation analysis should be performed. Fortunately, it is often possible to make sufficient modifications to an intended nominal orbit to avoid these unfavorable situations, yet still meet mission objectives.

H.1 Long-Term Risk Assessment Theory and Algorithmic Description

Both semi-analytical and Monte Carlo methods provide the means to estimate statistically expected long-term collision rates between satellites. Monte Carlo methods also provide estimates of the overall expected frequency of conjunctions, as well as the frequency of repeating conjunctions which can be used to identify potentially collocated satellites.

Analytic Long-term Collision Rate Analysis

The analytic theory developed by Kessler (1981) focuses on estimating collision rates between either spherical objects, or, for this application, circumscribing spheres. It provides estimates of such collision rates by neglecting all orbital perturbations except the orbital precession induced by Earth's J_2 gravitational term (see Vallado 1981 for a description of orbital precession). For each pair of analyzed satellites, the Kessler rate represents an average over six angles (Hall and Matney 2000). More specifically, the method averages collision rates over the three Keplerian orbital element angles for each of the two interacting objects. These three angles are the orbital right ascension of ascending node (RAAN, or Ω), the argument of perigee (ω), and the mean anomaly (M). In other words, Kessler rates represent six-dimensional phase space averages, which are equal to long-term temporal averages in idealized situations for which the ergodic hypothesis is valid. For this reason, Kessler rates are often referred to as “long-term” estimates,

even when calculated in more realistic, non-idealized situations to which the ergodic hypothesis does not necessarily strictly apply.

The Kessler (1981) formulation expresses the average collision rate between two satellites as an integral over the volume of their overlapping orbital regions, as follows:

$$\dot{N}_c^K = \sigma_{1,2} \int_V [S_1 S_2 v_{rel}] dV \quad (\text{H-1})$$

In this equation, V indicates the overlapping orbital volume shared by the primary and secondary satellites; $\sigma_{1,2} = \pi(R_1 + R_2)^2$ their combined collisional cross section; S_1 and S_2 their average spatial densities, respectively; and v_{rel} their average relative velocity during conjunction events. (The superscript K indicates that the Kessler method has been used to estimate the collision rate, and the dot above the symbol N_c represents the first time derivative.) The Kessler collision rate is zero for orbits that do not cross one another at all over long-term J_2 -only propagations, for which the overlap volume V vanishes. The integrand factors within the square brackets of equation (H-1) depend on first three Keplerian orbital elements of the primary and secondary satellite orbits, i.e., the semi-major axis (SMA), the eccentricity, and the inclination of each. Calculating the integral over volume in equation (H-1) entails two-dimensional numerical integration over geocentric latitude and radial distance.

Monte Carlo Colocation Conjunction Rate Analysis

Monte Carlo methods provide another means to estimate collision and conjunction rates, and allow for more generalities, including:

- Trajectories that account for orbital perturbations other than just those induced by the J_2 gravitational term,
- Both spherical and non-spherical shapes, used to estimate conjunction rates that arise from ellipsoidal or box-shaped screening volumes; and
- The ability to estimate the frequency of repeating conjunctions, used to assess satellite colocation.

The CARA Monte Carlo long-term rate estimation software uses the U.S. Space Force's Computation of Miss Between Orbits (COMBO) astrodynamics standards algorithm to determine close approaches between trajectories propagated using the SGP4 orbital theory. Input consists of a two-line element (TLE) set of mean orbital elements for the primary object plus a catalog of TLEs for the secondary objects being considered. Each Monte Carlo trial entails propagating trajectories over a duration of several days (typically seven or ten), beginning each propagation at or near the epoch of the TLE catalog. The initial mean orbital elements for the trial propagations are the same as specified by the input TLEs, except that newly sampled (ω, Ω, M) angles are generated for both objects, each randomly drawn from a uniform distribution spanning 0 to 2π . The Monte Carlo algorithm registers a "hit" for each instance that the trial trajectories approach one another close enough to penetrate a screening volume surrounding the primary object. For sufficiently large screening volumes, multiple hits can

potentially occur per trial, so the algorithm also registers the earliest hit that occurs for each primary-secondary pair as a “first-contact” event.

Dividing the total number of hits by the product of the number of trials and the propagation duration yields the Monte Carlo estimate for the collision rate:

$$\dot{N}_c^{MC} = \frac{N_{hit}}{N_{trial} T_{prop}} \quad \text{and} \quad \Delta \dot{N}_c^{MC} \approx \frac{\dot{N}_c^{MC}}{\sqrt{\min(N_{hit}, N_{trial} - N_{hit})}} \quad (\text{H-2})$$

The second equation above provides an approximation of the 1-sigma uncertainty of the estimate due to Monte Carlo counting uncertainties. Rates of first-contact events are similarly calculated and are always less than or equal to the overall rate, i.e., $\dot{N}_{fc}^{MC} \leq \dot{N}_c^{MC}$. As demonstrated later, for colocated orbits that produce mostly repeating conjunctions, the rate of first-contact conjunctions is significantly less than the total rate.

H.2 Long-Term Risk Assessment Methodology Demonstration

To demonstrate long-term risk and satellite colocation assessments, this section introduces a hypothetical proposed satellite called “NewSat” with a nominal initial orbit intentionally designed to be colocated with that of the Hubble Space Telescope (HST).

In this hypothetical scenario, NewSat has a hard-body radius of 3.5 m, and is to be launched and reach its mission orbit on 2022-11-18 00:00:00 UT. The mean orbit of HST at this epoch, propagated from a TLE cataloged shortly before that time, has a mean motion equal to 15.11279168 revs/day, an eccentricity of 0.0002381, and an inclination of 28.4696° — corresponding to an equatorial perigee altitude of about 530.6 km and an apogee altitude of 532.2 km. To create a colocated orbit for this demonstration, the TLE for NewSat’s nominal orbit has been made identical to that of HST in all respects except that the inclination has been rounded to the value of 28.5° and the eccentricity increased to 0.001. (These changes prevent the two orbits from being identical but are not sufficiently large to prevent them from being colocated, as will be demonstrated later.) This nominal NewSat deployment orbit corresponds to an equatorial perigee altitude of about 523.7 km and an apogee of 537.5 km, which is 13.8 km higher than perigee. However, for this hypothetical scenario, NewSat is given the flexibility to deploy into an alternate orbit with the same inclination, but with perigee altitudes anywhere from 500 to 600 km, and with an apogee up to 20 km higher than perigee. As discussed below, the long-term risk assessment methodology estimates rates between NewSat and HST over this entire region of orbital parameter space, as well as for all the other satellites contained within the 2022-11-18 TLE catalog, which represent a total of 36,368 potential secondary objects, a population that includes active and retired satellites, spent rocket-bodies, and orbital debris.

Figure H-1 shows a plot of Kessler collision rates between NewSat and HST, calculated with equation (H-1), and using hard-body radii of 3.5 m and 10 m for the two satellites, respectively. Specifically, the plot shows the collision rate (on the color axis), as a function of perigee altitude (horizontal axis) and apogee-perigee difference (vertical axis). The white star marks NewSat’s nominal deployment orbit. However, as mentioned previously, NewSat could deploy into an alternate orbit at any location within the bounds of the plot. Unsurprisingly, the nominal orbit

corresponds to a non-zero long-term collision rate between the NewSat and HST. However, much of the area on the plot corresponds to zero collision rate (shown as the darkest blue color on the plot); these regions represent candidate orbits for NewSat that do not overlap at all with HST’s orbit because they are sufficiently higher or lower in altitude. NewSat could change deployment plans by shifting its orbit into one of these regions, but that would only mitigate the long-term collisional threat due to HST and not necessarily that due to other cataloged satellites.




Figure H-1 Kessler Collision Rates between NewSat and HST

Figure H-2 shows collision rates summed over all the satellites contained within the full 2022-11-18 TLE catalog. For the nominal NewSat orbit marked by the white star, a total of 2,183 secondary objects were found to have non-zero Kessler rates among the 36,368 cataloged objects. The calculation uses hard-body radii for these secondary objects taken from a table of published dimensions for known satellites, when available, and based on radar cross section values for unknown objects such as orbital debris (Hall and Baars 2023). Figure H-2 shows that the nominal NewSat orbit is adjacent to (but not within) the region affected by cataloged Starlink constellation satellites, which occupy several orbital shells at higher altitudes (the regions affected by three of these shells are marked in Figure H-2). So, if the NewSat mission were to change its deployment plans, doing so to avoid these already crowded regions would largely eliminate the expected number of Starlink conjunctions. For clarity, Figure H-3 shows collision rates for all cataloged objects, but with Starlink constellation satellites excluded. A comparison of the color axis scales on Figures H-2 and H-3 illustrates how relatively large the long-term rates due to the Starlink constellation were in this specific region of orbital space at the time of this analysis (i.e., 2022-11-18). For constellations in the process of deployment, such rates tend to grow as more satellites are inserted into existing and additional orbital shells.




Figure H-2 Total NewSat Collision Rates Among All Cataloged Satellites




Figure H-3 NewSat Collision Rates Excluding the Starlink Constellation

Figures H-1, H-2, and H-3 graphically illustrate how long-term collision rate analyses can be used to avoid deploying satellites into crowded regions, and thereby help minimize the long-term collision risk among tracked objects. For the nominal deployment orbit, Figure H-2 indicates that NewSat’s predicted collision rate is $\sim 6 \times 10^{-4}$ per year, corresponding to a statistically expected collisional “lifetime” (expected time until a collision) of $\sim 1,700$ years among cataloged satellites. (As mentioned previously, other physical effects such as active maneuvering, orbital decay, etc., are likely to occur on shorter timescales than this, meaning that the ergodic hypothesis does not strictly hold for this system; nevertheless, long-term collision rates and collisional lifetimes are still used in these situations as indicators of statistical risk.) Notably, shifting NewSat’s deployment orbit into one of the regions most affected by the Starlink constellation would lead to a significantly shorter collisional lifetime of 30 to 100 years—in the absence of any risk mitigation maneuvering, which fortunately is being performed by Starlink constellation satellites. However, a careful inspection of Figures H-1, H-2, and H-3 indicates that changing the planned NewSat deployment to an orbit with a mean equatorial perigee altitude of about 558 km and an apogee ≤ 7 km higher would avoid both the HST and Starlink regions simultaneously and extend NewSat’s collisional lifetime to $\sim 3,000$ years. Furthermore, an extended analysis that includes all future Starlink constellation deployments (as tabulated by Bassa et al. 2022) indicates that the collision rate predicted for this specific orbital region should remain relatively low, at least among currently cataloged objects and currently planned Starlink satellites. Of course, the conclusions drawn above are predicated on major constellations’ remaining in their current operating orbits for a reasonable period of time. Current FCC licensing practices can permit rather large tolerances in constellation orbital location parameters (such as average orbital altitude); because operators can choose to exercise this discretion and relocate their constellations, conclusions such as these cannot be more than provisional. Nevertheless, they remain an important activity in working to avoid orbital placements expected to be problematic in terms of collision event rates.

H.3 Orbital Collision Avoidance Plan Analysis

As specified in NPR 8079.1, missions that are owned, developed or operated by NASA are required to complete a pre-launch Orbital Collision Avoidance Plan (OCAP). For this, in addition to an acceptable altitude range, a prospective mission might specify nominal orbital inclination and eccentricity values, along with upper and lower acceptable eccentricity bounds. For instance, instead of the nominal deployment orbit described above, the NewSat mission’s pre-launch plan might specify a nominal eccentricity of $e = 0.0006$, with an operationally acceptable range of $0.0002 \leq e \leq 0.001$. In this case, Figure H-4 shows the results of a pre-launch long-term collision risk analysis over a candidate perigee altitude range of 500 to 600 km. The left panel of Figure H-4 plots statistically expected collision rates for the low, nominal and high eccentricities, calculated for all cataloged objects including Starlink satellites, showing curves dominated by relatively high peaks near the Starlink constellation’s orbital shells. The right panel plots rates excluding Starlink, allowing a more accurate evaluation of rates caused by the remainder of the cataloged satellites.




Figure H-4 NewSat Conjunction Rates with and without Starlink


H.4 Colocation Analysis Methodology

During their on-orbit residence, most spacecraft come into conjunction with other, secondary, objects with some frequency. However, if two active satellites are maintained in similar orbits or in orbits that have consistent intersections, synodic alignment can occur that produces multiple close approaches on successive orbits. If the alignment is very close, such repeating conjunctions can persist over extended periods (e.g., longer than typical seven- or ten-day conjunction screening durations) and could also recur periodically over a satellite's orbital lifetime. Because conjunction management between active payloads is more complex, requiring coordination of maneuver plans and joint decisions about near-term courses of action, it is desirable to minimize such situations.

In the example introduced in the previous section, the nominal NewSat orbit was constructed specifically to be highly collocated with HST's orbit. This section demonstrates how Monte Carlo analysis can be used to identify such collocated orbits, which are characterized by a high frequency of repeating conjunctions, or, equivalently, a relatively low frequency of first-contact conjunctions. For this, the Monte Carlo method is used to predict the average rate of conjunctions and the fraction expected to be first-contact events. This specific analysis estimates conjunction rates averaged over a screening period of seven days (which is the same screening duration used for most LEO CARA satellites). Conjunctions occur in the Monte Carlo simulation when an ellipsoidal screening volume centered on the primary satellite is penetrated by the trajectory of a secondary satellite. The ellipsoid used for the present analysis has principal axes that measure $2 \times 25 \times 25$ km, aligned with the primary object's radial, in-track and cross-track directions.

The goal of colocation analysis is to determine the orbital conditions that create repeating conjunctions, which can be accomplished by analyzing the frequency of first-contact events. As mentioned previously, first-contact conjunctions correspond to the initial instance that a conjunction screening volume is penetrated. This means that pairs of satellites that experience

high frequencies of repeating conjunctions necessarily have low first-contact conjunction rates relative to their total conjunction rates. The colocation analysis graph in Figure H-5 demonstrates this effect by comparing total conjunction rates and first-contact rates, plotted as a function of Δa , which represents the offset of NewSat's orbital semi-major axis (SMA, or a) from its nominal value. (Note, the analysis assumes that NewSat's eccentricity and inclination remain at their nominal values of 0.001 and 24.5°, respectively). The black dashed curve marked with circles in Figure H-5 plots the total conjunction rate, $\dot{N}_c^{MC}$, calculated using equation (H-2). This total includes both first-contact conjunctions and any subsequent conjunctions (i.e., repetitions). The red dotted curve marked with diamonds plots the first-contact rate, $\dot{N}_{fc}^{MC}$, which is always less than or equal to the total rate. The black curve shows that conjunctions of either type are predicted to occur for SMA offsets within about ± 10.5 km of the nominal deployment value, i.e., in the range $-10.5 \text{ km} \leq \Delta a \leq +10.5 \text{ km}$. However, the red curve indicates that the rate of first-contact conjunctions is noticeably lower than the total over the range $-6.1 \text{ km} \leq \Delta a \leq +6.1 \text{ km}$. Furthermore, the first-contact rate is less than 50% of the total rate over the more restricted range of $-2.2 \text{ km} \leq \Delta a \leq +2.2 \text{ km}$.



Nmc = The number of Monte Carlo trials


RIC = Radial – In-track – Cross-track

Figure H-5 Predicted Rates for NewSat and HST Satellite Conjunctions

This analysis adopts this 50% reduction (i.e., $\dot{N}_{fc}^{MC} / \dot{N}_c^{MC} < 0.5$) as a means to identify highly colocated orbits, meaning that, for this specific screening volume and under these conditions, NewSat would need to shift its deployment SMA by at least 2.2 km in either direction from the nominal value to avoid a high level of colocation with HST. Notably, shifting by about 6.1 km would largely eliminate the occurrence of repeating conjunctions, but a shift of 10.5 km or more would be required to largely eliminate all conjunctions between the two satellites. The one-dimensional analysis shown in Figure H-5 assumes only SMA orbital adjustments, holding the deployment eccentricity and inclination constant at their nominal values. A similar two-dimensional analysis holding only the inclination constant can also be performed to determine

how much the SMA and/or eccentricity would need to change to eliminate high levels of colocation (or, equivalently, how much the perigee and/or apogee altitudes would need to change). The generalized method can be extended to even higher dimensions to include all relevant, adjustable orbital elements.


If the NewSat mission were to change the planned deployment orbit to avoid colocation with HST, as described above, then it would be worthwhile to do so in a manner that does not result in it being collocated with any other high value satellites. For example, the Fermi Gamma-ray Space Telescope (FGST) and the Swift satellite, both NASA missions, occupy orbits similar to NewSat's nominal proposed orbit, so it is conceivable that shifting the NewSat deployment orbit could inadvertently create a high level of colocation with one of these satellites. Figures H-6 and H-7 demonstrate this by showing one-dimensional colocation analysis graphs for FGST and Swift, respectively (each made using the same assumptions as described above for Figure H-5). Figure H-6 shows that FGST's red curve with diamonds is less than half of the black curve with circles for $-7.0 \text{ km} \leq \Delta a \leq -2.5 \text{ km}$, so if NewSat were to shift into this range, then the analysis predicts it would become highly collocated with FGST. Similarly, Figure H-7 indicates that shifting NewSat's orbit into the range $+7.8 \text{ km} \leq \Delta a \leq +12.4 \text{ km}$ would result in high colocation with the Swift satellite.



Nmc = The number of Monte Carlo trials

RIC = Radial – In-track – Cross-track

Figure H-6 Predicted Rates for NewSat and FGST Satellite Conjunctions



N_{mc} = The number of Monte Carlo trials

RIC = Radial – In-track – Cross-track

Figure H-7 Predicted Rates for NewSat and Swift Satellite Conjunctions

H.5 Long-Term Conjunction Rate and Colocation Analysis Software

The NASA CARA software instantiation of the long-term risk and colocation analysis capability is an ongoing tool development effort, currently being refactored to execute faster (which is especially required for the Monte Carlo analysis modes), and eventually to be more portable (because it currently uses the U.S. Space Force COMBO tool, which has restricted distribution). The main CARA software tool entitled “*LongTermRisk*” uses TLE catalogs as input and allows users to estimate long-term Kessler collision rates as well as Monte Carlo conjunction rates and associated first-contact rates. (The *LongTermRisk* tool was used to create all the figures shown in this appendix.) The tool also has the capability to ingest auxiliary TLE catalogs containing predicted future satellite constellation(s), for instance, or model populations of untracked orbital debris. Because the tool uses the limited-distribution AstroStandards version of the U.S. Space Force COMBO close approach analysis function, and because the associated execution environment is relatively complicated to set up, it is envisioned that the ability to run this tool will remain within the NASA CARA organization.

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.

The U.S. Space Force 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 satellite O/O ephemerides, the situation is much more uneven. While it was previously true that in the majority of conjunctions the covariance for the secondary object predominated, given that now the majority of conjunctions for a primary are typically not with debris but rather other intact spacecraft, O/O-generated covariances often define the entire uncertainty volume for a conjunction. This development makes 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 the Ansys Orbit Determination Tool Kit (ODTK), 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., onboard 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 thus to determine how well actual state errors conform to the Gaussian distribution that the covariance specifies. Such a statement requires additional 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 conjunction assessment applications, and for many conjunction 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 conjunction 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;² if it is propagated, it is intended to be a reflection of the state estimate 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 or narrow ranges of states. Many aspects of the U.S. Space Force's 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 conjunction assessment applications, but current trends are pushing conjunction assessment maneuver commitment times to 1 or even 0.5 days, indicating that realism optimization at those particular propagation states is perhaps a better choice.

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

² It is questionable whether the formal 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.

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 time between ephemerides 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

Satellite 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: onboard 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 interpolation 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, for 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.

Although the RIC frame has been traditionally used, there is a growing movement to modify the relative frame to one that is defined by the in-track and not the radial component. The RIC frame fixes the radial error component to align with the vector from the earth’s center to the

satellite, the in-track component is normal to that component in the orbital plane, and the cross-track component is normal to the orbital plane. This approach is a helpful mapping of errors to their intuitive sources for circular or near-circular orbits, but this intuition is lost for eccentric orbits. An alternative formulation that is more widely applicable is to fix the in-track component to align with the satellite's velocity vector, define the "radial" component as normal to that in the orbital plane, and construct the cross-track as normal to the orbital plane. This aligns one of the components with the velocity, which is very close to the direction of action of the largest error source (atmospheric drag acceleration), and therefore remains meaningful for orbits of all eccentricities. This alternative formulation, often called "VNB" (aligns with Velocity, then Normal to velocity, then normal to the plane [B]), is identical to RIC for circular orbits and therefore provides the same intuitive framework for the most common orbit type.

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 covariance realism 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 as a vector 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):

$$\frac{\varepsilon_u^2}{\sigma_u^2} + \frac{\varepsilon_v^2}{\sigma_v^2} + \frac{\varepsilon_w^2}{\sigma_w^2} = \chi_{3\text{dof}}^2 \quad (\text{I-3})$$

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:

$$C = \begin{bmatrix} \sigma_u^2 & 0 & 0 \\ 0 & \sigma_v^2 & 0 \\ 0 & 0 & \sigma_w^2 \end{bmatrix} \quad (\text{I-4})$$

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:

$$\varepsilon = \begin{bmatrix} \varepsilon_u & \varepsilon_v & \varepsilon_w \end{bmatrix} \quad (\text{I-6})$$

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:

$$\varepsilon C^{-1} \varepsilon^T = \begin{bmatrix} \varepsilon_u & \varepsilon_v & \varepsilon_w \end{bmatrix} \begin{bmatrix} 1/\sigma_u^2 & 0 & 0 \\ 0 & 1/\sigma_v^2 & 0 \\ 0 & 0 & 1/\sigma_w^2 \end{bmatrix} \begin{bmatrix} \varepsilon_u \\ \varepsilon_v \\ \varepsilon_w \end{bmatrix} = \frac{\varepsilon_u^2}{\sigma_u^2} + \frac{\varepsilon_v^2}{\sigma_v^2} + \frac{\varepsilon_w^2}{\sigma_w^2} = \chi_{3\text{dof}}^2 \quad (\text{I-7})$$

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 $\epsilon C^{-1} \epsilon^T$ 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 ϵ_x and ϵ_y and the correlation coefficient is ρ , the entire equation, with correlation terms included, assumes the form:

$$\epsilon C^{-1} \epsilon^T = \frac{1}{(1-\rho)^2} \left(\frac{\epsilon_x^2}{\sigma_x^2} + \frac{\epsilon_y^2}{\sigma_y^2} - \frac{2\rho\epsilon_x\epsilon_y}{\sigma_x\sigma_y} \right) \quad (\text{I-8})$$

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 $\epsilon C^{-1} \epsilon^T$ 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 personal 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”. Although counterintuitive, this approach is not unreasonable: 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 observed 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—if they do, then the covariance properly represents the actual observed error 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 first to calculate and tabulate the differences between the Cumulative Distribution Function (CDF) of the sample distribution and that of the hypothesized distribution, then to calculate a goodness-of-fit statistic from these differences, and finally 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 the above calculations and reside on the NASA CARA public-facing 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 do, to some degree, 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 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 previously, 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) 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.

An extended example of the application of the technique described here is available in Zaidi and Hejduk (2016).

Appendix J. CARA Conjunction Risk Assessment Tools Validation

CARA and JSC FOD are chartered to perform conjunction assessment 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 conjunction risk assessment evaluation tools. However, cases can occasionally arise in which it is not practical or desirable for CARA/JSC FOD to perform these calculations, such as for potential autonomously controlled missions that will perform onboard conjunction assessment. In such cases, the needed calculations or tools will be validated by CARA/JSC 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. conjunction 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

[illegible]

MET - Mission Elapsed Time

SMA - Semi-Major Axis

RAAN - Right Ascension of the Ascending Node

Appendix L. Commercial Data in NASA Conjunction Assessment

The breadth and quality of commercial SSA data has developed in the last decade from essentially no commercial presence at all to the presence of a number of different vendors, each specializing individually in radar, optical, or passive radio frequency satellite tracking. In addition to tracking measurement data, vendor offerings now include additional SSA products such as vectors, ephemerides, and even conjunction risk assessment analysis outputs such as Conjunction Data Messages (CDMs).

In examining and facilitating the use of commercial data in NASA conjunction risk assessment calculations, CARA and JSC FOD adhere to the following principles:

1. Use raw observation (measurement) data only and combine with Space Surveillance Network (SSN) observations for a single solution.
2. Validate all non-SSN data.
3. Undertake a cost/benefit analysis before purchasing commercial SSA data.

More details on these principles are described below.

L.1 Use raw observation (measurement) data only and combine with SSN observations for a single solution

Of all the data types offered by commercial providers, only the raw satellite tracking data are at present fully and unambiguously of use to the NASA conjunction assessment enterprise.

Because NASA has access to both the USSPACECOM satellite tracking data from the SSN and the USSPACECOM space catalog maintenance operational system, NASA has the ability to include raw commercial observation data along with SSN tracking data into a combined close approach prediction. This process involves combining all of the measurement data on an object of interest from both the SSN and commercial sources to perform a single orbit determination. The predicted close approaches can then be calculated from this combined source. The CARA and JSC FOD personnel who are embedded in the 18 SDS work center at VSBF presently have this capability in a “sandbox” area on the operational system that is not used for routine operations. The sandbox area enables CARA/JSC FOD 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.

Work is also progressing on a regularization of the reception of commercial data into the Mission Delta 2 enterprise; when such a capability is in place, conjunction assessment products deriving from combined commercial and DOD tracking data will be possible routinely. But receiving higher-level products from commercial vendors, such as vectors or even conjunction assessment outputs like CDMs, introduces multiple solutions into the enterprise and creates significant operational difficulties in trying to adjudicate among them.

Event evaluation in the presence of multiple solutions

There is presently a debate within the industry regarding the best way to proceed with event evaluation in the presence of multiple solutions. CARA/JSC FOD encounters a version of this situation in miniature each day in that most conjunction events possess two solutions: one that uses the DOD predicted position for the primary object state and a second that uses the O/O predicted ephemeris for that state. Counterintuitively, it is not always clear which of these two sources for the primary object predicted state is superior.

O/Os typically have more and better past position data on their satellite, either from telemetry tracking or onboard GPS fixes, and they also understand the configuration of their satellite better and thus should be able to calculate the components of the ballistic coefficient with more precision. The O/Os should then be able to come up with a better epoch solution on which to base prediction, and they are positioned to do a better job in prediction as well because they potentially possess a superior ballistic coefficient. However, O/O orbit determination software atmospheric density models often introduce more error into predicted states than the model used by the 18 SDS and 19 SDS.

In the situation that CARA faces, the choice of a preferred data source is thus not straightforward because a number of factors influence the selection: the capabilities of the O/O's orbit determination software, the particular orbit that the satellite occupies, the degree to which the space weather situation is perturbed, and the span from the present time to TCA, among others.

Establishing a rubric of precedence

To reliably choose between these two information sources, CARA has performed a Verification and Validation (V&V) activity of the orbit determination products for selected NASA O/Os and has developed rubrics for selecting a particular solution (DOD- or O/O-based) source in particular contexts. A V&V activity of this type requires several months' worth of data and is complicated to perform, necessitating specialized tools and subject matter expert interpretation of the results. With multiple providers of information for both the primary and secondary objects involved in a conjunction, the situation would be that much more complicated.

Assimilative solution selection

An alternative to establishing a rubric of precedence for the multi-solution situation is using an assimilative modeling approach similar to that used for hurricane modeling.

The results from multiple providers are collected and comparatively analyzed, both visually and with formal clustering algorithms. Due to the large representation among the submissions (i.e., most of the individual results clustering around a single result), the hope is to be able to establish a supermajority result that suggests itself as the favored solution. Although such a procedure appears attractive at first, it encounters difficulties in implementation.

- First, it requires a sufficient number of providers supplying solutions to obtain meaningful data clusters. Clustering among, say, ten submitted solutions has the propensity to identify meaningful clusters and thus enable the drawing of certain

conclusions. With only two or three providers, this is not really possible. At present, the number of different providers is too small to use a clustering approach.

- Second, because of explainable and therefore expected discontinuities in space surveillance data sets, it is not unusual for minority submissions to an assimilative framework to represent the superior solution. The common situation of satellite maneuvers is a good example: only one of a set of providers may have detected a maneuver for one of the spacecraft in a conjunction and updated that spacecraft's state appropriately. This solution is clearly superior to other submissions that did not obtain tracking in time to identify the maneuver and correct the trajectory. But as a single solution representing this modified state, this solution will appear assimilative as an outlier and is likely to be rejected. This scenario has played out repeatedly at the Sprint Advanced Concept Training (SAC-T) exercises conducted by the U.S. Space Force's Joint Commercial Operations (JCO) cell: a single vendor's submission differs appreciably from the other submissions yet is shown later to be the preferred solution.

In short, while assimilative solution selection remains a potentially helpful construct for employing multiple solutions contemporaneously to guide a single operational decision, more research and refinement are required before such a paradigm can be employed in an operationally satisfactory manner to support conjunction risk assessment.

Direct fusion of orbital solutions

Finally, there are proposals in the academic literature for the direct fusion of different orbital solutions, including CDMs from different providers, into a single solution that in principle combines the information of all the individual submissions.

Carpenter (2020) represents the most promising of such proposals, which is also nicely tailored for the conjunction assessment problem. While the theoretical development is indisputable, a couple of practical problems remain:

- First, the combination methodology relies directly on realistic covariances accompanying each state estimate. If these covariances are not representative, then the fused product is errant not just in its own covariance but in its synthesized state estimate. Because covariance realism is an elusive achievement for many providers, and because a rigorous V&V effort is required to establish the abiding presence of realistic covariances in a provider's product set, direct fusion is not a mechanism that can be easily or immediately employed.
- Additionally, the theory requires the cross-correlations between the covariance matrices arising from different providers to be characterized and deployed in the fusion paradigm, yet there is no obvious method for establishing these cross-correlations between the providers' orbital determination uncertainties. There is an argument that, due to the use of different input data with different errors, the correlated error should be negligible; at the same time, it could reasonably be expected that some shared model error, such as geopotential error or atmospheric density forecast error, would appear in a cross-correlation analysis.

So, for the direct fusion methodology described here, more focused study is needed before an approach of this type can be employed operationally.

Conclusion about the use of commercial SSA data

Given the present difficulties with the presence of multiple SSA and conjunction assessment products, NASA has concluded that it is necessary to limit the use of commercial SSA data to the combination of commercial and DOD tracking information.

L.2 Validate all non-SSN data

All non-SSN sources supplying data to be used in orbit determination solutions to maintain the catalog must be regularly calibrated to ensure the removal of any data biases that might skew the conjunction assessment screening results.

Using raw observation data from commercial vendors (as opposed to post-processed solutions) enables straightforward data characterization and numerical validation.

Commercial vendors supporting NASA are regularly asked 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 U.S. Government 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.

L.3 Undertake a cost/benefit analysis before purchasing commercial SSA data

While following the recommendations of L.1 and L.2 above would allow the introduction and proper use of commercial SSA data in NASA conjunction assessment, a third consideration focuses on whether the use of these data would in fact result in a significant operational benefit.

NASA, as a U.S. Government entity, follows the guidance of the National Space Policy to receive conjunction assessment services from the 18 SDS and 19 SDS, who maintain 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 would rarely, if ever, result in a different operational decision regarding whether to require and execute a conjunction mitigation action.

Most conjunction 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 to understand the financial obligation appropriate to purchase commercial data to augment the existing NASA conjunction assessment process 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 taken 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 a 1 in 10,000 chance of a collision versus the 1 in 20,000 chance calculated with the SSN data alone.

The benefit would depend on how many times per year these decisions would be altered and thus how much risk is mitigated. By some estimates, over 90% of the objects larger than 1 cm and thus able to render a satellite inoperative in a collision are too small to be tracked and thus are not considered in the conjunction risk assessment process. Therefore, part of the trade space is to consider how much adding commercial data reduces collision risk compared to the background risk. This decision would likely vary by orbit altitude based on trackability of objects of various sizes and the 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 other spacecraft. Given this, it is best left to the individual mission project office to choose appropriate conjunction assessment 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; and it is thus in response to this second goal that actual conjunction risk assessment and mitigation thresholds are articulated. Mission operators can always take a more conservative posture than what is appropriate for debris mitigation, but at the least they must 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 to 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. With tracking improvements and the launch of additional satellites, one can expect a regularly maintained space catalog of perhaps 30,000 to 50,000 objects. This means that, roughly speaking, 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 is thus 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 conjunction 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 excessively 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, and mitigation. First, because collisions between protected payloads and large secondary objects will produce by several orders of magnitude the largest number of debris fragments, 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 understood, these also represent straightforward situations for which due diligence counsels 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 determinacy 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 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 encumbers science mission objectives.

M.2 Point of Reference: Conjunction Risk Assessment Based on Miss Distance

In the early days of conjunction assessment, satellite position uncertainty data were not regularly available, so the only orbital data product on which to base conjunction risk mitigation decisions was the predicted miss distance at TCA between the two conjunctioning 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 that 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 presented an undue burden to mission operations. Methods were thus developed that considered the characterized error in the state estimates to try 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 attempt to give a 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 only operationally useful export from such an analysis is an understanding of how the O/O’s risk mitigation burden might have been reduced were the maneuver commitment point able to be placed much closer to TCA. 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 Pc for Conjunction Risk Assessment Requirements

Many conjunction 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 collision likelihood metric for requirements specification purposes, different considerations, technical and otherwise, are relevant:

- Suitability to NASA’s risk posture regarding conjunction 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 conjunction assessment industry.

Based on these considerations, CARA presently employs 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 decision makers. It can be turned into a statement about miss distances if this framework is desired. Finally, it is the standard metric typically used by the military, civil, and commercial conjunction assessment industries.

The P_c metric, as calculated by analytic methods, was first explained formally in 1992 (Foster and Estes). A comprehensive treatment of the calculation and associated issues can be found in Chan (2008), although this monograph is now somewhat dated and does not consider many of the newer collision likelihood calculation proposals. (See Appendix K for an explanation of the presently embraced 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 conjunction assessment 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. There are a number of objections to the parameter, a precis of which will be given below; but in the main the issues center on 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 why it has driven consideration of alternative risk metrics.

As stated previously, probabilistic conjunction assessment 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 predicted 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' positions 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 calculated probability 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. The cause of these dilution situations is presumed to be inadequate tracking of the objects, thus producing large covariances that reflect a poor understanding of their present and therefore predicted states.

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.

M.5 Responses to Probability Dilution: Alternative Metrics

The discipline has advanced a number of different responses to probability dilution, and one such response is to define alternative risk assessment metrics that attempt to avoid this problem. A full survey and treatment of the mainstream proposed alternatives is not possible in this short appendix, but a precis and description of many of the proposals is offered below.

Maximum P_c . Advanced by S. Alfano in the same article that introduced the concept of the dilution region to the industry (Alfano 2005b), the “Max P_c ” is, for dilution region situations, the probability of collision generated by systematically shrinking the covariance and calculating the P_c with each discrete reduction; the Max P_c is the largest such P_c generated from this covariance-shrinking exercise. The thinking is that, in the presence of increased tracking, the covariance would be smaller; so the Max P_c is the highest P_c that could in principle be attained if more data were available to help combat the dilution region situation. One could implement this construct operationally by acting on the Max P_c rather than the standard P_c calculated for a particular

event. While there is an argument within the industry whether the P_c is actually a probability, the $MaxP_c$ is one degree further removed from such a definition: it represents what the P_c would in principle be if the tracking levels actually were such that the real covariance and the manipulated covariance that produced the $MaxP_c$ were to agree in size, and therefore at best it is a conditional probability. However, despite its further distancing from an actual representation of collision likelihood, to be conservative one could use the $MaxP_c$ in lieu of the standard P_c as a method of managing potentially understated risk from the standard P_c in dilution region situations.

Ellipsoid Overlap. Balch et al. (2019) published an extended critique of the standard P_c , considering the dilution effect and other issues, and further saw the situation as a repudiation of certain Bayesian approaches. Their recommendation in other publications was to take the P_c out of service entirely and replace it with more conventional statistical methods, such as ensuring that the three-sigma ellipsoids of the two objects not overlap. This recommendation is easy enough to implement, to wit: if the two three-sigma ellipsoids overlap, then a maneuver should be pursued to ensure that they not overlap.

Plausibility/Belief. Sanchez et al. (2024) proposed a framework in which both collision plausibility and belief would be separately calculated and tracked during the evolution of a conjunction event. Collision belief essentially shows the durability of data indicating that a collision is likely, and plausibility showed the antecedent conditions that could make a collision possible. When both plausibility and belief are low, the situation is safe; when they both are high, the situation is dangerous and requires a mitigation action; and when they diverge (and the only way this can happen is with high plausibility and low belief), then it is not entirely clear how to proceed, but at least one knows that, despite little durable evidence, there is a real possibility of a collision. This presentation in principle gives the analyst additional information, for the standard P_c on its own does not indicate whether a probability dilution situation inheres.

Confidence Interval on Miss Distance. Elkantassi and Davison (2022) assembled a construct that, instead of calculating a collision probability, stays within a frequentist framework and instead estimates a significance probability that the miss distance is less than the combined hard-body radius, as derived from a statistical evidence function. This construct is immune to the probability dilution problem; furthermore, an approach is sketched by which its insights could potentially be deployed in a Bayesian framework to “rescue” the P_c , presuming that one could come up with an appropriate conjunction assessment Bayesian prior.

M.5 Issues with Alternative Risk Assessment Metrics

Given the four alternatives to the P_c enumerated above, all of which at least either point out or avoid the dilution phenomenon entirely, why is there any desire to remain with the P_c at all? Unfortunately, the alternatives listed above all introduce similar practical and conceptual difficulties for conjunction assessment mitigation operations because they essentially remediate the dilution region situation by requiring mitigation actions in all such cases.

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. A recent CARA profiling exercise revealed that, in the main, use of these

alternatives multiplicatively increases the number of mitigation actions from a factor of somewhat larger than unity for the Alfano MaxPc (and therefore perhaps operationally tractable), to a factor of a little under 8 for the Elkantassi and Davison approach, to a factor of 9 for the Balch-recommended ellipse overlap approach. 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 methods such as minimizing covariance ellipsoid overlap rather than reducing the Pc value to an acceptable level will require very large maneuvers—much larger than the typical satellite maneuvers that are presently performed. Large maneuvers expend additional fuel and, if pursued more frequently, may 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/uncatalogued 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 worsened, the satellite covariance(s) increase in size, and the Pc decreases until ultimately it results in a value of 0 (to machine precision). This asymptotic value of nullity aligns nicely with the zero-value Pc that is, at least indirectly, imputed to conjunctions with untracked/uncatalogued 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 Pc 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 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 disjunctive move from taking larger and larger actions to suddenly doing nothing. Given that, as stated previously, 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 Pc, 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 Pc 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 Pc 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—a situation that could potentially be revealed by the summoning of additional tracking data. It is appropriate to recognize that this situation (state estimate with large covariance) is not dissimilar to that of a conjunction with an untracked/uncatalogued 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/uncatalogued 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 apply the Max P_c method, as the one with the lowest false-alarm rate, and use the Max P_c as the test statistic to compare to the mitigation threshold. At a minimum, if it were possible to marshal additional tracking data to improve the state estimate, then an approach such as this would be quite appropriate to identify conjunctions that would benefit from such additional data.

Finally, it should be noted that CARA continues to evaluate both existing and new proposals for collision risk assessment, and it is quite possible that an approach that favors one such method over the standard P_c will emerge and become operationally deployed. Such an approach, however, will need to avoid producing excessive additional false alarms that, to this point, have been justified through an appeal to excessive safety conservatism. Any approach embraced by CARA will need to align with the mission statement's mandate not to constitute an undue burden to satellite mission operations.

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 may always 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. 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 may always 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 conjunction assessment requirements compliance. While this parameter is widely used in the conjunction assessment industry, issues related to its calculation exist that merit extended discussion. The most frequently used analytical techniques to calculate the Pc are well established and computationally efficient but include assumptions that restrict their use and make the calculations vulnerable to error for a small fraction of conjunctions. Numerical techniques also exist, such as Monte Carlo Pc estimation. 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 Pc estimates, 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 goal, this appendix will address the following technical areas:

- A step-by-step description of the conjunction plane 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 orbital state covariance matrices for the two objects in conjunction,
- Discussion of the use of Monte Carlo techniques, which is the numerical method used for high-fidelity 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,
- Discussion of alternative analytic Pc calculation methods, specifically the two- and three-dimensional “Nc” (as opposed to “Pc”) estimation methods that address conjunctions affected by curvilinear trajectories and non-Gaussian distribution effects; and
- Discussion of correcting Pc estimates for covariance cross correlation effects, which affects some conjunctions involving high-drag satellites.

N.2 Conjunction Plane Analytic Pc Calculation

The conjunction plane method of Pc calculation, which is by far the most widely used approach in the conjunction assessment industry, was developed for the Space Shuttle Program and first described 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), Garcia-Pelayo (2016), and Elrod (2019) — but all rely on the same basic methodology: applying reasonable assumptions to enable analytical approximations. While the particulars vary, these approaches all share the same concept of calculating the Pc estimate by integrating over a two-dimensional region on a conjunction encounter plane.

Conjunction plane Pc analysis begins with the states and covariances for the primary and secondary objects' orbital states at 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 a subsequent 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 Pc calculations depend 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 in this manner.




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 often via an estimation technique as the secondary is usually a debris object for which there is no definitive size information. If these two spheres approach and begin to overlap one another at any point during the encounter, then 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 super-sphere (also 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 collision sphere (also 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 (e.g., the z 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 z-component probability of the dot’s falling within the straw. As such, the z-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 x- and y-components. The entire situation can thus be reduced from three to two dimensions and analyzed as a phenomenon that occurs in a plane that is orthogonal to the soda-straw direction, which is the direction of the relative velocity vector. This procedure defines the “conjunction plane,” which can be viewed in two equivalent representations as discussed by Chan (2008, see Figure 5.1), and as shown in Figure N-4.




Figure N-4 Two Equivalent Representations of the Conjunction Plane

In the first representation, shown in the left panel of Figure N-4, the center of the secondary object is placed at the origin, and the local (x, y) coordinate system configured so that the miss vector (which extends a distance of r_{miss} from the secondary to the primary object) lies along the horizontal x 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 (illustrated here using a 1σ 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 two-dimensional integral (Foster and Estes 1992)

$$P_c = \frac{1}{\sqrt{\det(2\pi\mathbf{C})}} \iint_A \exp\left(-\frac{\mathbf{r}^T \mathbf{C}^{-1} \mathbf{r}}{2}\right) dx dy \quad (\text{N-1})$$

in which $\mathbf{r}$ is the 2×1 conjunction-plane miss vector, $\mathbf{C}$ is the 2×2 covariance matrix, and A is the hard-body-radius circular area. There are several different approaches to evaluating this integral. Foster and Estes (1992) applied a two-dimensional quadrature technique; this works well, and MATLAB®’s³ adaptive *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. Garcia-Pelayo et al. (2016) also derives a multi-term series approximation. However, even though these series approximations are relatively computationally efficient, experience indicates that they occasionally produce inaccuracies. Elrod (2019) formulates the integral in terms of complementary error functions (to improve accuracy for conjunctions with small P_c values) and uses Gauss-Chebyshev quadrature with nodal symmetry (which improves the efficiency of the numerical integration considerably). The Foster/Estes and Elrod approaches (the former being the established standard, and the latter being extremely fast, especially with a vectorized MATLAB implementation) are included as the *Pc2D_Foster* and *PcElrod* functions 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.)

In the second conjunction plane representation, shown in the right panel of Figure N-4, the secondary object location again is placed at the origin, but in this case the local (x', y') coordinate system is configured so that the principal axis of the covariance ellipse lies along the horizontal axis. The center of the primary object lies at (x_{miss}, y_{miss}) , and the axes are oriented so that this miss position lies within the 1st quadrant, so that $x_{miss} \geq 0$ and $y_{miss} \geq 0$. Using this representation, the P_c is given by a one-dimensional integral involving error functions (Alfano, 2005a)

$$P_c = \frac{1}{\sqrt{8\pi\sigma_x}} \int_{-R}^R \{\text{erf}(y_+) - \text{erf}(y_-)\} \exp\left[-\frac{(u + x_{miss})^2}{2\sigma_x^2}\right] du \quad (\text{N-2})$$

with R indicating the combined hard-body radius, $u = x' - x_{miss}$ the integration variable, and $y_{\pm} = (y_{miss} \pm \sqrt{R^2 - u^2})/(\sqrt{2}\sigma_y)$. The integrand factor in the curly brackets represents the difference of two error functions, $\Delta = \text{erf}(y_+) - \text{erf}(y_-)$. In many software environments (e.g., MATLAB), this factor often can be computed significantly more accurately as a difference of complementary error functions, i.e., $\Delta = \text{erfc}(y_-) - \text{erfc}(y_+)$, especially when calculating very small P_c values (Elrod, 2019). See Abramowitz and Stegun (1970) and Press et al. (1992) for details on computing $\text{erf}(-)$ and $\text{erfc}(-)$ functions.

For most conjunctions, Gauss-Chebyshev quadrature provides an efficient and accurate means to calculate the one-dimensional integral in equation (N-2), using an approach similar to that

³ MATLAB is a registered trademark of The MathWorks, Inc.

described by Elrod (2019). However, for conjunctions involving relatively large hard-body radii (or, equivalently, small covariance ellipse σ_x dimensions), this method can potentially become inaccurate. Specifically, Gauss-Chebyshev quadrature becomes increasingly inaccurate as the hard-body radius grows beyond the limit $R_{lim} = 4\sigma_x - \min(R)$. In such cases, which rarely occur in practice, an adaptive numerical integrator can be used to calculate an accurate estimate, e.g., MATLAB's *integral* function. (The function *PcCircle* in the NASA CARA SDK repository implements a vectorized algorithm that automatically determines which of these two integration methods should be used to compute Pc values accurately for all input hard-body radii and covariance values, with a computation speed comparable to that of the *PcElrod* function for most conjunctions. Also, for cases that have hard-body radii well below the limit given above, testing indicates that the SDK functions *Pc_Foster*, *PcElrod*, and *PcCircle* all output nearly identical numerical Pc values, i.e., that typically agree to five digits of precision or more.)

The second conjunction plane representation shown in the right panel of Figure N-4 also provides an extremely efficient means to calculate an upper limit estimate for the Pc value. This upper bound corresponds to the two-dimensional (2-D) integral over the square that circumscribes the hard-body radius circle (as shown in Figure N-4), which has the following analytical solution

$$P_c < P_{sq} = \frac{[\text{erf}(X_+) - \text{erf}(X_-)] [\text{erf}(Y_+) - \text{erf}(Y_-)]}{4} \quad (\text{N-3})$$

with

$$X_{\pm} = \frac{x_{miss} \pm R}{\sqrt{2}\sigma_x} \quad \text{and} \quad Y_{\pm} = \frac{y_{miss} \pm R}{\sqrt{2}\sigma_y} \quad (\text{N-4})$$

Again, in many cases the erf(-) differences in equation (N-3) can be computed more accurately using erfc(-) differences. Notably, the circumscribing square probability estimate does not require any numerical integration at all, but instead only requires the computation of four erf(-) or erfc(-) functions, usually making it relatively easy to program into software and computationally efficient. These considerations could be important in some circumstances (e.g., for computations performed on orbiting satellites), prompting the use of P_{sq} as an approximation for P_c itself. For instance, CDMs generated by the USSPACECOM ASW processing system often report P_{sq} as the estimate of conjunction's collision probability. (An optional, non-default mode of the function *PcCircle* in the NASA CARA SDK repository implements an efficient vectorized algorithm that calculates P_{sq} values).

It is perhaps helpful at this point to review the four assumptions employed by the conjunction plane P_c calculation methods given by equations (N-1) through (N-4), because alternative P_c estimation techniques will be needed when these assumptions do not inhere:

1. Statistical Independence: 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 as an isolated topic in a subsequent section of this appendix.
2. Rectilinear Motion: 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 between Earth-orbiting satellites. 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 conjunction plane P_c approximation will miscarry.
3. Gaussian Position Distributions and Negligible Velocity Uncertainties: The position vector state errors for each satellite at TCA follow Gaussian distributions, and the velocity vector state errors are negligibly small. When combined, these assumptions lead to the conjunction plane representations shown in Figure N-4. Conjunctions that do not satisfy these assumptions are addressed by alternate analytical and Monte Carlo methods, discussed in the sections below.
4. Temporally Isolated Event: The conjunction presents a single, well-defined event so that the collision likelihood can be ascertained by examining that single instance. Objects that stay in close proximity for extended periods accumulate risk throughout long interactions, rather than quickly accumulating risk at or near a well-defined TCA. A different approach is also required for their evaluation, which is discussed in the sections below.

As mentioned previously, the issue of statistical independence (i.e., assumption 1 above) will be discussed in a subsequent section of this appendix. To address conjunctions that do not satisfy any of the other assumptions, two alternative (but more computationally intensive) analytical methods are available: the “three-dimensional N_c ” method, which research indicates can be applied to conjunctions that violate assumptions 2, 3 and/or 4 above, and the “two-dimensional N_c ” method, applicable to temporally isolated conjunctions that violate assumptions 2 and/or 3. Notably, the three-dimensional N_c method could, in principal, be applied to all conjunctions to estimate P_c values. However, this is not justified because most LEO satellite conjunction P_c values can be estimated accurately using the much more efficient conjunction plane methods described above. Additionally, most of the remaining conjunctions can be estimated accurately using the two-dimensional N_c method, which, although much slower than the conjunction plane methods, is still significantly faster than the full three-dimensional N_c method. (The function *UsageViolationPc2D* soon to be posted on the NASA CARA Software Development Kit (SDK) repository provides an algorithm that determines if a given conjunction violates one or more of conjunction plane P_c estimation assumptions, and if so, which of the other available P_c estimation methods are appropriate.)

N.3 Three-Dimensional Nc Method Analytic Pc Calculations

The relatively infrequent conjunctions that do not satisfy the conjunction plane method 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. 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 integration over the surface of a unit sphere. When combined with a one-dimensional time integration, this yields a “three-dimensional Pc” 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 Pc 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 Pc calculations, even though all the required three-dimensional Pc assumptions are satisfied.

Shelton and Junkins (2019) provided a key insight into why the original three-dimensional Pc approximation fails in certain situations. Their analysis indicates that accurate Pc 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 Coppola (2012) three-dimensional Pc formulation did not incorporate this concept, but Hall (2021) reformulated the method to do so explicitly. 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 for the two satellites. The analysis demonstrates that such “peak overlap” states can be determined using an iterative calculation that converges quickly. The formulation derives an expression for “Nc” — the statistically expected number of collisions — which equals Pc for single, temporally isolated conjunctions, but that may exceed Pc for multi-encounter interactions. The resulting “three-dimensional Nc” method (also known as the “3D-Nc” method) entails a total of three numerical integrations, one over time and two over the surface of a sphere. The outermost, time integration expression has the form:

$$N_c(\tau_a, \tau_b) = \int_{\tau_a}^{\tau_b} \dot{N}_c(t) dt \quad (\text{N-5})$$

This expression estimates the number of collisions statistically expected to occur at some time during the risk assessment interval $\tau_a \leq t < \tau_b$, which can represent either a short duration closely bracketing a single close-approach encounter, or an extended duration spanning multiple encounters. The expected collision number, N_c , is closely related to the collision probability, P_c . In fact, they are equal for single, isolated encounters between well-tracked satellites. The collision rate for such a temporally isolated conjunction is expressed as a two-dimensional integral over the unit sphere:

$$\dot{N}_c(t) = R^2 \int_0^{2\pi} \int_0^\pi [v_t(\hat{r}) MVN(R\hat{r}, \tilde{\mathbf{r}}_t, \tilde{\mathbf{A}}_t)] \sin(\theta) d\theta d\varphi \quad (\text{N-6})$$

In this equation, the radial unit vector is given by $\hat{r} = [\cos(\varphi) \sin(\theta), \sin(\varphi) \sin(\theta), \cos(\theta)]^T$, so the surface of the unit sphere is spanned by the azimuthal angle $0 \leq \varphi < 2\pi$ and the axial angle $0 \leq \theta \leq \pi$. The leading factor of R^2 represents the square of the combined hard-body radius, meaning that the expression actually represents an integration over the surface area of the collision sphere. As explained in Hall (2021), the integrand function in the square brackets is the product of two factors. The first is an average velocity term, $v_t(\hat{r})$, which is a function of time (as indicated by the t subscript) and calculated using the TCA states and covariances of the primary and secondary objects. The second factor, $MVN(R\hat{r}, \tilde{\mathbf{r}}_t, \tilde{\mathbf{A}}_t)$, represents a multi-variate normal (MVN) function which depends on $\hat{r}$ and R , as well as a 3×1 relative position/velocity state vector, $\tilde{\mathbf{r}}_t$, and an associated 3×3 $\tilde{\mathbf{A}}_t$ covariance matrix, both of which are also calculated from the TCA states and covariances. Lebedev quadrature (Lebedev and Laikov, 1999) provides an efficient method for numerical integration over the unit sphere. (The function `Pc3D_Hall` of the NASA CARA SDK repository computes conjunction Pc estimates, along with associated Nc and Nc rate estimates, using the Hall (2021) three-dimensional Nc method as summarized in equations (N-5) and (N-6) above.)




Figure N-5 Pc Rate and Cumulative Pc for a Non-Rectilinear Conjunction

Figure N-5 shows the time-dependent collision rate (left) and the cumulative Pc (right) calculated using the three-dimensional Nc method, as applied to an archived CARA conjunction that has a high relative velocity (~ 13 km/s), but that fails to satisfy the rectilinear motion and Gaussian PDF assumptions. This specific event does not satisfy these assumptions because it involves an object in a highly eccentric orbit with a conjunction occurring in the relatively tightly curved part of the trajectory near perigee, a phenomenology discussed by Hall (2018). In this case, the three-dimensional Nc method calculates a Pc value which accurately matches the Monte Carlo from TCA method Pc estimate, but that is a factor of fifteen larger than the erroneous conjunction plane Pc estimate, calculated in this case using the two-dimensional Foster and Estes (1992) method. It is of interest that the peak point of risk accumulation occurs

nine seconds before the TCA, with essentially all the risk accumulated by 8 seconds before TCA. Such a result can seem counterintuitive at first, for one would initially expect that the point of highest risk would always be at the TCA. However, what is operative is the alignment between the geometric miss distance and the covariance. If at TCA very little of the position uncertainty lies along the relative miss vector, then that miss vector is a strong statement of the actual miss; and if the vector is somewhat larger than the hard-body radius (HBR), then the collision risk at that point is quite low. If, however, somewhat earlier or later than TCA the combined covariance (which, it must be remembered, is constantly changing position) does more substantially align along the miss vector, then more possible instantiations of the true miss are likely to be smaller than the HBR (even with the nominal miss vector larger than the expected miss at TCA), so the risk at that point is actually higher. In the more extreme cases, such as that shown in the above example (and even more strongly in the one below), most of the risk accumulation can be relatively far from TCA, meaning that methods that examine the situation only at TCA can misrepresent this risk.




Figure N-6 Pc Rate and Cumulative Pc for a Low Velocity Conjunction

Figure N-6 shows the collision rate (left) and the cumulative Pc (right) for low relative velocity conjunction (~ 3 m/s), that fails to satisfy the rectilinear motion, temporally isolated, and Gaussian PDF assumptions. In this case, the three-dimensional Nc algorithm estimates a Pc value which matches the Monte Carlo from TCA method Pc estimate, but that is a factor of 330 larger than the erroneous conjunction plane Pc estimate.

N.4 Two-dimensional Nc Method Analytic Pc Calculations

As formulated in equations (N-5) and (N-6), the three-dimensional Nc (or 3D-Nc) method requires the computation of a time-series of many unit-sphere integrations, each calculated using Lebedev quadrature. However, as described by Hall et al. (2023), for temporally isolated conjunctions (i.e., those that are not too “extended” or “blended” as described by Hall 2021), the integration over time can be approximated analytically, ultimately yielding a single unit-sphere integral. This relatively efficient “two-dimensional Nc” (or 2D-Nc) method yields accurate Pc estimates for high velocity, temporally isolated conjunctions that are affected significantly by

curvilinear trajectory effects (such as the example shown in Figure N-5). However, 2D-Nc is not applicable to some low velocity events that are too temporally extended or blended (such as that as shown in Figure N-6); these kinds of conjunctions still require the use of the 3D-Nc method or Monte Carlo estimation. (The function *Pc2D_Hall* of the NASA CARA SDK repository computes Pc estimates for temporally isolated conjunctions using the two-dimensional Nc method.)

N.5 Multistep Algorithm for Pc Calculations

Hall et al. (2023) describe a multistep algorithm that uses the 2D-Pc, 2D-Nc, and 3D-Nc semi-analytical methods along with the from-TCA and from-epoch Monte Carlo methods, in order to calculate accurate Pc values regardless of whether they are affected by curvilinear trajectory effects or by low-velocity encounters. Figure N.7 shows a flowchart of the multistep algorithm which sequentially applies these five estimation methods, but only as required for computational efficiency. The gray rectangles represent main computation modules, and the blue diamonds indicate decision points. For each conjunction, the process starts with the 2D-Pc calculation, and identification of any associated “usage violations” that indicate potential inaccuracies due to the 2D-Pc method assumptions not being satisfied. The algorithm then sequentially applies more advanced Pc estimation methods, but only as required. It next tries the 2D-Nc method, which is significantly more efficient than the 3D-Nc method. If necessary, it then tries the 3D-Nc method itself, a step required relatively infrequently and primarily for long duration conjunctions caused by low relative velocities. Finally, in the rare cases that all of these semi-analytical methods register usage violations, or if the user desires further validation of the results, the algorithm performs from-TCA MC estimation followed as required by from-epoch MC estimation. (Note: Figure N.7 refers to from-TCA estimation as “SDMC” meaning *Simplified Dynamics Monte Carlo*, and from-epoch estimation as “BFMC” meaning *Brute Force Monte Carlo*.)




Figure N.7 Schematic Flowchart of the Multistep Pc Algorithm⁴

The rectangles on the left side of Figure N.7 show the fractions of conjunctions processed at each of the refinement steps, based on testing by Hall et al. (2023) of 1.93 million conjunctions archived in the CARA database during the years 2021 and 2022. Among this original set, 7.3% of the conjunctions possess orbit determination or filtering conditions that prevent further processing. For the remaining 1.79 million conjunctions, the multistep algorithm reports estimates based on the 2D-Pc method for 94.7%, the 2D-Nc method for 5.0%, and the 3D-Nc method for 0.3%. Relatively few conjunctions ultimately require MC methods: 0.006% for SDMC and 0.002% for BFMC. By default, the algorithm invokes successive estimation methods only as required, but allows processing with any or all of the methods, useful for more complete analysis of high interest events (HIEs).

Detecting potential 2D-Pc method usage violations represents one of the most important steps of the multistep algorithm, and entails comparing Mahalanobis distances between the two satellites calculated using the rectilinear vs curvilinear (two-body) motion approximations. Large Mahalanobis distance differences indicate likely 2D-Pc method failures, and usage violations correspond to cases for which the differences exceed a conservatively small threshold. The comparative Mahalanobis distance analysis also provides an approximate correction factor for the 2D-Pc estimate, as well as estimates for the duration of the conjunction and the time of peak collision risk, measured as an offset from the nominal conjunction TCA. By default, the multistep algorithm only proceeds to other estimation methods for conjunctions that register Mahalanobis distance-based usage violations and corrected 2D-Pc estimates that are too large, or that have durations that are too extended, or offset times that are too long. Hall et al. (2023) provide further details on evaluating usage violations for the 2D-Pc method and the other Pc estimation methods shown in Figure N.7. The CARA team is currently positioning to use a multistep approach for operational processing. (The function *PcMultiStep* currently pending

⁴ From Hall et al. (2018)

addition the NASA CARA SDK repository computes P_c estimates for temporally-isolated conjunctions using the first three steps of the multistep algorithm, which combines software modules for the 2D- P_c , 2D- N_c , and 3D- N_c methods, along with their associated usage violation modules.)

N.6 Comparison of P_c Estimates for Temporally Isolated Conjunctions

Figure N-8 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 and for events with $2D\ P_c > 10^{-7}$ (Hall 2021). The vertical axes on all three panels plot Monte Carlo (MC) estimates for the collision probability—specifically, Monte Carlo from TCA method P_c estimates, which are also referred to in this figure as two-body Monte Carlo method P_c estimates (i.e., TBMC- P_c estimates, as described in more detail in section N.13). The error bars in Figure N-8 show 95% confidence intervals estimated using the Clopper-Pearson (1934) method; several of the conjunctions had zero hits registered in the Monte Carlo simulations, which are represented in Figure N-8 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, three-dimensional N_c in the center, and two-dimensional N_c on the right. The colored points on each plot indicate the results of a binomial proportion statistical test that evaluates the agreement between the estimates. Specifically, black points in Figure N-8 indicate analytical P_c estimates that agree reasonably well with the Monte Carlo estimates 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 conjunction plane P_c comparison plotted in the left graph 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 (~1%) of the original conjunctions. These disagreements represent conjunctions that violate one or more of the assumptions required for conjunction plane P_c estimation. The center graph clearly shows that the three-dimensional N_c method matches the Monte Carlo P_c estimates better, producing only 66 yellow and zero red points. Finally, the two-dimensional N_c method plotted on the right produces very nearly the same results as the three-dimensional N_c method but requires much less computation time. The three comparisons shown in Figure N-8 indicate that, for temporally isolated conjunctions, the two- and three-dimensional N_c methods are consistent with one another, and match Monte Carlo P_c estimates significantly better than the conjunction plane P_c estimation method.




Figure N-8 Comparison of Monte Carlo Collision Probabilities with the Two-Dimensional P_c Method (left), the Three-Dimensional N_c Method (center), and the Two-Dimensional N_c Method (right) for a Large Set of CARA Conjunctions

N.7 P_c Calculations for Multi-encounter Interactions

An additional feature of the two- and three-dimensional N_c methods is their ability to address explicitly the amalgamated risk of repeating conjunctions. While most conjunctions are temporally isolated, there are two conjunction types that exhibit different behaviors. The first type is produced by objects in orbits in synodic alignment that generate a temporal sequence of conjunctions once every revolution or multiple of a revolution (e.g., two nearly circular orbits at different inclinations that produce conjunctions at one or both nodal crossing points). The second type involves objects that orbit close to each other for extended periods, generating extended interactions with multiple close approaches (e.g., two satellites in nearly the same orbit, but with only a slight difference in inclination and/or eccentricity). In both cases, if each of the multiple encounters is considered separately, situations can arise in which each encounter in the series falls below a P_c mitigation threshold, but the combined risk of all of the encounters exceeds that threshold. The two- and three-dimensional N_c methods account for such multi-encounter interactions by providing estimates for the total expected number of collisions, and upper and lower bounds for the probability of collision (Hall, 2021; Hall et al., 2023). The total statistically expected number of collisions for a multi-encounter interaction is given by a summation of N_c values for each conjunction in the sequence:

$$N_c(\tau_a, \tau_b) = \sum_{k=1}^K N_{c,k} \quad (\text{N-7})$$

with the index $k = 1 \dots K$ indicating the close approaches applicable to the risk assessment interval $\tau_a \leq t < \tau_b$, and $N_{c,k}$ representing the associated expected number of collisions for each. The upper and lower P_c bounds for the combined interaction are given by:

$$\max(N_{c,k}) \leq P_c(\tau_a, \tau_b) \leq 1 - \prod_{k=1}^K (1 - N_{c,k}) \quad (\text{N-8})$$

which also implies that $P_c(\tau_a, \tau_b) \leq N_c(\tau_a, \tau_b)$. Figure N-9 shows an example of a multi-encounter interaction, in which a pair of satellites experience four conjunction events over about a five-hour period. Each of the individual conjunctions produces a P_c value in the upper yellow region (between 10^{-5} and 10^{-4}), as plotted in the bottom panel. The solid line in the top panel shows the upper limit of the cumulative P_c for the interaction; notably, after the third conjunction, the cumulative P_c exceeds 10^{-4} — a value frequently selected as a risk 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 P_c value generated by the method 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 found to be below the mitigation threshold, then there is no need to marshal computationally intensive methods (such as Monte Carlo), for it has already been demonstrated no mitigation action is warranted.




Figure N-9 Cumulative Three-Dimensional N_c Risk Over Four Repeating Conjunctions

N.8 Input Covariance Data Considerations

Most calculations are only as good as the input data that drive them, and P_c calculation is no exception. Appendix P 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 conjunction risk mitigation actions. The purpose of this section is to address the routine improvement and basic error checking extended to covariances as part of the P_c 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.8.1 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 conjunction assessment 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 orbit prediction 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 conjunction assessment practitioners receive from the DOD, 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 give 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 that affects 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 Ap (major magnetic storms list) and Dst (disturbance storm-time) space weather indices. Figure N-10 shows the behavior of these polynomial fits divided into four different classes of Ap/Dst activity; y-axis values are omitted here to allow full releasability of the figure:




Figure N-10 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:

$$B = C_D \frac{A}{M} \quad (\text{N-9})$$

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), the effect can be irrelevant during the fit (because the rotation is so slow that it does not affect near-term look-back and look-forward), or the effect can be such that the frontal area assumes a particular value for much of the fit-span and then a very different value during a subsequent prediction interval. 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-11 illustrates the two-phase application of these consider parameters:



DCP = Dynamic Consider Parameter

Figure N-11 Two-phase Application of Consider Parameters

If the CDMs generated by the DOD are used for conjunction assessment applications, the good news is that all 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.8.2 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 conjunction assessment 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 conjunction risk 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 orbital state covariance matrix is not positive semidefinite, then Monte Carlo sampling on the entire matrix may not be 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.8.3 Covariance Correlation

For nearly all the broader conjunction assessment 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 conjunction assessment 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

$$P_m = P_s + P_p - \sigma_{s/g} \sigma_{p/g} G_s G_p^T - \sigma_{s/g} \sigma_{p/g} G_p G_s^T \quad (N-10)$$

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 conjunction assessment practitioner would come upon all of the needed data to affect the proposed compensation. As described in detail in the next section, 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, also 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.) Hall (2021) and Hall et al. (2023) describe how the density forecast errors and sensitivity vectors can be used to estimate decorrelated joint covariances for the two- and three-dimensional Nc methods.

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, as expected, also 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-12 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

P_c calculated with the decorrelated joint covariance to that of the P_c 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 P_c is notable, in many instances more than an order of magnitude. For the remaining cases, there is an increase in P_c from a factor of 1.5 to 5.




Figure N-12 Profiles of 250 Conjunctions with Primary and Secondary Satellites of Non-negligible Drag⁵

While the P_c value changes by less than a factor of 1.5 for most of the conjunctions, a sufficient number are affected more substantially therefore justifying the integration of this additional consideration into the P_c computation, especially because it is a straightforward calculation from data provided directly in the CDM.

N.8.4 Conjunction Data Message DCP Uncertainty and Sensitivity Vector Explanation

As mentioned above, the content of Conjunction Data Message files produced by the United States Space Force's ASW software has been enhanced to enable conjunction risk assessments that account for covariance correlation. Specifically, ASW release versions 19.2 and after insert additional information within CDMs to enable users to include the atmospheric drag covariance cross-correlation effect when estimating a conjunction's probability of collision, as calculated from DCP uncertainty values and sensitivity vectors (as described in the previous section). This section provides an overview of these CDM modifications, and a brief description of their usage.

CDM Mean Position/Velocity State Vectors

CDMs can contain the mean position/velocity state vectors of the primary and secondary satellites at TCA in a couple of reference frames, as described in more detail in the CDM

⁵ From Casali et al. 2018.

specification documents by Consultative Committee for Space Data Systems (specifically, CCSDS 502.0-B-3, listed in the references section). For Pc computation, an inertial reference frame must be used, sometimes requiring these original CDM states to be converted into the desired inertial reference frame. This section denotes the resulting mean inertial position and velocity vectors at TCA for the primary and secondary objects as $(\bar{\mathbf{r}}_p, \bar{\mathbf{v}}_p)$ and $(\bar{\mathbf{r}}_s, \bar{\mathbf{v}}_s)$, respectively.

CDM Position/Velocity State Covariance Matrices

CDMs specify position/velocity state covariance matrices for the primary and secondary satellites using the radial-transverse-normal (RTN) coordinate frame. A CDM file specifies this symmetric 6×6 covariance for each object using keyword values for the 21 non-redundant matrix elements as follows:

$$\mathbf{C} = \begin{bmatrix} \mathcal{C}_{R,R} & \mathcal{C}_{T,R} & \mathcal{C}_{N,R} & \mathcal{C}_{\dot{R},R} & \mathcal{C}_{\dot{T},R} & \mathcal{C}_{\dot{N},R} \\ \mathcal{C}_{T,R} & \mathcal{C}_{T,T} & \mathcal{C}_{N,T} & \mathcal{C}_{\dot{R},T} & \mathcal{C}_{\dot{T},T} & \mathcal{C}_{\dot{N},T} \\ \mathcal{C}_{N,R} & \mathcal{C}_{N,T} & \mathcal{C}_{N,N} & \mathcal{C}_{\dot{R},N} & \mathcal{C}_{\dot{T},N} & \mathcal{C}_{\dot{N},N} \\ \mathcal{C}_{\dot{R},R} & \mathcal{C}_{\dot{R},T} & \mathcal{C}_{\dot{R},N} & \mathcal{C}_{\dot{R},\dot{R}} & \mathcal{C}_{\dot{T},\dot{R}} & \mathcal{C}_{\dot{N},\dot{R}} \\ \mathcal{C}_{\dot{T},R} & \mathcal{C}_{\dot{T},T} & \mathcal{C}_{\dot{T},N} & \mathcal{C}_{\dot{T},\dot{R}} & \mathcal{C}_{\dot{T},\dot{T}} & \mathcal{C}_{\dot{N},\dot{T}} \\ \mathcal{C}_{\dot{N},R} & \mathcal{C}_{\dot{N},T} & \mathcal{C}_{\dot{N},N} & \mathcal{C}_{\dot{N},\dot{R}} & \mathcal{C}_{\dot{N},\dot{T}} & \mathcal{C}_{\dot{N},\dot{N}} \end{bmatrix} \quad (\text{N-11})$$

The RTN-frame covariance $\mathbf{C}$ can be transformed into a 6×6 inertial frame position/velocity covariance $\mathbf{P}$ by applying the following equation:

$$\mathbf{P} = \mathbf{M} \mathbf{C} \mathbf{M}^T \quad (\text{N-12})$$

with the 6×6 transformation matrix $\mathbf{M}$ having the form:

$$\mathbf{M} = \begin{bmatrix} \mathbf{M} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{M} \end{bmatrix} \quad (\text{N-13})$$

with $\mathbf{0}_{3 \times 3}$ representing a 3×3 matrix of zeros. The 3×3 orthonormal matrix $\mathbf{M}$ rotates vectors from the pseudo-inertial (non-rotating, instantaneously frozen) RTN frame into the inertial frame:

$$\mathbf{M} = [\hat{\mathbf{R}} \ \hat{\mathbf{T}} \ \hat{\mathbf{N}}] \quad (\text{N-14})$$

with column vectors given by the three RTN unit vectors, calculable from the object's inertial mean position and velocity vectors as follows:

$$\hat{\mathbf{R}} = \bar{\mathbf{r}}/|\bar{\mathbf{r}}| \quad \text{and} \quad \hat{\mathbf{N}} = (\bar{\mathbf{r}} \times \bar{\mathbf{v}})/|\bar{\mathbf{r}} \times \bar{\mathbf{v}}| \quad \text{and} \quad \hat{\mathbf{T}} = \hat{\mathbf{N}} \times \hat{\mathbf{R}} \quad (\text{N-15})$$

The 6×6 covariance in eq. (N-12) can be decomposed into three 3×3 sub-matrices:

$$\mathbf{P} = \begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{C} \end{bmatrix} \quad (\text{N-16})$$

with $\mathbf{A}$ representing the marginalized covariance of the position vector, $\mathbf{C}$ the marginalized covariance of the velocity vector, and $\mathbf{B}$ position-velocity cross correlations.

Eqs. (N-11) to (N-16) can be used to calculate inertial frame position/velocity state covariance matrices at TCA for the primary and secondary objects involved in a conjunction, $\mathbf{P}_p$ and $\mathbf{P}_s$, respectively, as well as the marginalized position covariance matrices, $\mathbf{A}_p$ and $\mathbf{A}_s$. (Note: because RTN is an object-specific frame of reference, these calculations must employ different rotation matrices, $\mathbf{M}_p$ and $\mathbf{M}_s$, respectively.)

Uncorrelated and Correlated Relative Position Covariance Matrices

Collision probability estimation using the 2D-Pc method requires the conjunction's inertial relative position miss-vector, $\bar{\mathbf{r}}_m = \bar{\mathbf{r}}_s - \bar{\mathbf{r}}_p$, along with the associated miss-vector covariance matrix, $\mathbf{A}_m$. If the primary and secondary position vectors are statistically independent (i.e., uncorrelated), then the relative position miss-vector covariance is given by their sum:

$$\mathbf{A}_m = \mathbf{A}_p + \mathbf{A}_s \quad (\text{N-17})$$

This approach provides a viable approximation for 2D-Pc estimation for conjunctions in which $\mathbf{A}_p$ and $\mathbf{A}_s$ have sufficiently weak correlation. However, as described in the previous section, Casali et al. (2018) have demonstrated that some conjunctions have stronger covariance correlations, due to shared atmospheric density forecast components arising from the ASW global density portion of the Dynamic Consider Parameter. In these cases, the miss-vector covariance can be corrected by removing the cross-correlated components as follows:

$$\mathbf{A}_m = \mathbf{A}_p + \mathbf{A}_s - \sigma_{p/g}\sigma_{s/g}[\mathbf{G}_p\mathbf{G}_s^T + \mathbf{G}_s\mathbf{G}_p^T] \quad (\text{N-18})$$

with $\sigma_{p/g}$ and $\sigma_{s/g}$ denoting the atmospheric density 1-sigma relative uncertainties for the primary and secondary, respectively. The vectors $\mathbf{G}_s$ and $\mathbf{G}_p$ represent the sensitivity of the miss-vector covariance on global density relative uncertainties. (Note: see eq. (11) of Casali et al. 2018 and the related discussion for a derivation of eq. (N-18), and a more detailed explanation of its components. Also, this section uses the symbol “ $\mathbf{A}$ ” for 3×3 position covariances, instead of the symbol “ $\mathbf{P}$ ” which this section uses to represent full 6×6 position/velocity covariances. This change of notation means that eq. (N-18) is simply a restatement of eq. (N-10).)

CDM files produced by ASW versions 19.2 and later provide the sigma values and sensitivity vectors required to calculate corrected relative position covariance matrices using eq. (N-18). These data have been added within “comment” lines in the primary and secondary object portions of the CDM, with the format:

COMMENT DCP Density Forecast Uncertainty = 2.45030416E-01

COMMENT DCP Sensitivity Vector RTN Pos = -1.67441647E+01 3.68889831E+02 1.63797508E-01 [m]

COMMENT DCP Sensitivity Vector RTN Vel = -3.98670591E-01 1.09452965E-02 -4.83454839E-04 [m/sec]

The first occurrence of parameters with this specific format appears in the primary object section of the CDM and provides $\sigma_{p/g}$ (the DCP density forecast 1-sigma uncertainty), $\mathbf{G}_p^{RTN}$ (the 3×1 DCP position sensitivity vector, expressed in the primary's RTN frame), and $\mathbf{H}_p^{RTN}$ (the 3×1 DCP velocity RTN-frame sensitivity vector). The second occurrence provides the corresponding DCPs for the secondary, i.e., $\sigma_{s/g}$, $\mathbf{G}_s^{RTN}$ and $\mathbf{H}_s^{RTN}$. The RTN frame sensitivity vectors provided

in the CDM can be converted to inertial frame vectors using the transformation matrix defined in eq. (N-14) separately for each object:

$$\mathbf{G}_p = \mathbf{M}_p \mathbf{G}_p^{RTN} \quad \text{and} \quad \mathbf{G}_s = \mathbf{M}_s \mathbf{G}_s^{RTN} \quad \text{and} \quad \mathbf{H}_p = \mathbf{M}_p \mathbf{H}_p^{RTN} \quad \text{and} \quad \mathbf{H}_s = \mathbf{M}_s \mathbf{H}_s^{RTN} \quad (\text{N-19})$$

Uncorrelated and Correlated Relative Position/Velocity Covariance Matrices

For statistically independent primary and secondary states, the relative position/velocity miss-state covariance is given by the sum of the 6×6 covariances for the two objects:

$$\mathbf{P}_m = \mathbf{P}_p + \mathbf{P}_s \quad (\text{N-20})$$

This approach provides a viable approximation for conjunctions in which the primary and secondary position/velocity states have sufficiently weak correlation. In other cases, the miss-vector covariance can be corrected by removing the cross-correlated components as follows:

$$\mathbf{P}_m = \mathbf{P}_p + \mathbf{P}_s - \sigma_{p/g} \sigma_{s/g} [\mathbf{\Gamma}_p \mathbf{\Gamma}_s^T + \mathbf{\Gamma}_s \mathbf{\Gamma}_p^T] \quad (\text{N-21})$$

with $\sigma_{p/g}$ and $\sigma_{s/g}$ again denoting the atmospheric density 1-sigma relative uncertainties for the primary and secondary objects, respectively. The 6×1 vector $\mathbf{\Gamma}_p = [\mathbf{G}_p^T \mathbf{H}_p^T]^T$ represents the sensitivity of the primary object's miss-state covariance on global density relative uncertainties. The secondary object's 6×1 sensitivity vector, $\mathbf{\Gamma}_s$, is defined similarly. (Again, see Casali et al. (2018) for more detail on these DCP uncertainties and sensitivity vectors.) The RTN frame sensitivity vectors provided in the CDM can be converted to inertial frame vectors using the transformation matrix defined in eq. (N-13) for each object:

$$\mathbf{\Gamma}_p = \mathcal{M}_p \mathbf{\Gamma}_p^{RTN} \quad \text{and} \quad \mathbf{\Gamma}_s = \mathcal{M}_s \mathbf{\Gamma}_s^{RTN} \quad (\text{N-22})$$

Example Calculation of Uncorrelated and Correlated Covariance Matrices

This section examines a mid-2022 CARA conjunction selected to demonstrate how to calculate and compare uncorrelated and correlated covariance matrices—intended as a reference guide for interested users. The next four pages provide the CDM in text file format for this example conjunction,⁶ which involves LEO primary and secondary satellites that have a combined HBR of 4.5 m. Using the 2D-Pc estimation method along with the original covariance given by eq. (N-17) yields 2D-Pc = 1.60×10⁻⁴. Substituting the covariance corrected by removing the cross-correlated components, as given by eq. (N-18), yields 2D-Pc = 8.04×10⁻⁴. (Note: for this specific conjunction, the multistep algorithm described in section N.5 indicates no 2D-Pc method usage violations; in addition, these 2D-Pc estimates accurately match those calculated with all of the more advanced semi-analytical and Monte Carlo Pc estimation methods discussed in section N.5.) So in this example case, accounting for the covariance correlation effect reduces the collision probability estimate by about a factor of two. For other conjunctions, accounting for the effect can decrease or increase Pc estimates to a lesser or greater degree than in this specific example (see Figure N-12 and the next section of this appendix). The remainder of this section summarizes the calculation and comparison of the covariances given by eqs. (N-17) and (N-18), as derived from the CDM file provided for this example conjunction.

⁶ This CDM cleared for public release by NASA on April 03, 2024 via Orbital Data Request #24-001-1.

NASA Spacecraft Conjunction Assessment and Collision Avoidance Best Practices Handbook – Volume 2, Technical Appendices

```

CCSDS_CDM_VERS                = 1.0
CREATION_DATE                  = 2022-03-11T22:52:43.000
ORIGINATOR                     = CARA
MESSAGE_FOR                    = ICON
MESSAGE_ID                     =
000044628_conj_000027127_20220313_181420_20220311_225243
COMMENT SCREENING_OPTION = Covariance
TCA                            = 2022-03-13T18:14:20.971
MISS_DISTANCE                  = 1939 [m]
RELATIVE_SPEED                 = 13335 [m/s]
RELATIVE_POSITION_R            = -57.3 [m]
RELATIVE_POSITION_T            = 917.5 [m]
RELATIVE_POSITION_N            = -1707.3 [m]
RELATIVE_VELOCITY_R            = -10.1 [m/s]
RELATIVE_VELOCITY_T            = -11736.0 [m/s]
RELATIVE_VELOCITY_N            = -6332.5 [m/s]
COLLISION_PROBABILITY          = 1.601e-04
COLLISION_PROBABILITY_METHOD   = FOSTER-1992
COMMENT HBR = 4.5 [m]
OBJECT                         = OBJECT1
OBJECT_DESIGNATOR              = 000044628
CATALOG_NAME                   = SATCAT
OBJECT_NAME                    = ICON
INTERNATIONAL_DESIGNATOR       = 2019-068A
EPHEMERIS_NAME                 = NONE
COVARIANCE_METHOD              = CALCULATED
MANEUVERABLE                   = N/A
REF_FRAME                      = ITRF
GRAVITY_MODEL                   = EGM-96: 36D 360
ATMOSPHERIC_MODEL              = JBH09
N_BODY_PERTURBATIONS           = MOON,SUN
SOLAR_RAD_PRESSURE              = YES
EARTH_TIDES                     = YES
INTRACK_THRUST                 = NO
COMMENT COVARIANCE_SCALE_FACTOR = 1.000
COMMENT EXCLUSION_VOLUME_RADIUS = 5 [m]
COMMENT OD_DATA_SOURCE = ASW
RECOMMENDED_OD_SPAN            = 6.34 [d]
ACTUAL_OD_SPAN                 = 6.34 [d]
OBS_AVAILABLE                  = 140
OBS_USED                       = 140
TRACKS_AVAILABLE               = 38
TRACKS_USED                    = 38
RESIDUALS_ACCEPTED              = 99.2 [%]
WEIGHTED_RMS                   = 1.115
COMMENT Apogee Altitude = 621 [km]
COMMENT Perigee Altitude = 575 [km]
COMMENT Inclination = 27.0 [deg]
CD_AREA_OVER_MASS              = 0.028466 [m**2/kg]
CR_AREA_OVER_MASS              = 0.006874 [m**2/kg]
SEDR                           = 0.000568 [W/kg]
COMMENT DCP Density Forecast Uncertainty = 2.45519905000000105e-01
COMMENT DCP Sensitivity Vector RTN Pos = -3.718622471813500852e+01
3.499336708808674757e+03 2.307466130887458888e+00 [m]
COMMENT DCP Sensitivity Vector RTN Vel = -3.805998793451176976e+00
2.254928072259999861e-02 -2.698557836399999906e-03 [m/sec]
X                               = -5.771809181165332120e+03 [km]
Y                               = 3.716587497753967000e+03 [km]
Z                               = 1.098406950424936895e+03 [km]

```

NASA Spacecraft Conjunction Assessment and Collision Avoidance

Best Practices Handbook – Volume 2, Technical Appendices

X_DOT	= -2.995090631353613642e+00 [km/s]
Y_DOT	= -5.615424864506517011e+00 [km/s]
Z_DOT	= 3.227036277207581794e+00 [km/s]
CR_R	= 2.536723712995373319e+02 [m**2]
CT_R	= -7.820021346272073060e+03 [m**2]
CT_T	= 8.060409152666834416e+05 [m**2]
CN_R	= -1.052519669690997972e+01 [m**2]
CN_T	= 3.820856351588464008e+02 [m**2]
CN_N	= 6.035679921552716110e+01 [m**2]
CRDOT_R	= 8.675863248707885944e+00 [m**2/s]
CRDOT_T	= -8.762210717913994813e+02 [m**2/s]
CRDOT_N	= -4.248532765564856928e-01 [m**2/s]
CRDOT_RDOT	= 9.527151535479747579e-01 [m**2/s**2]
CTDOT_R	= -2.343743211674224058e-01 [m**2/s]
CTDOT_T	= 4.491561324993874393e+00 [m**2/s]
CTDOT_N	= 9.204011711869856255e-03 [m**2/s]
CTDOT_RDOT	= -5.070741091757811556e-03 [m**2/s**2]
CTDOT_TDOT	= 2.299529516857000096e-04 [m**2/s**2]
CNDOT_R	= 9.066947510392180457e-03 [m**2/s]
CNDOT_T	= -5.152612790665972442e-01 [m**2/s]
CNDOT_N	= -2.690884470369837855e-02 [m**2/s]
CNDOT_RDOT	= 5.816996267946000299e-04 [m**2/s**2]
CNDOT_TDOT	= -6.987286540099999741e-06 [m**2/s**2]
CNDOT_NDOT	= 6.518335978740000280e-05 [m**2/s**2]
OBJECT	= OBJECT2
OBJECT_DESIGNATOR	= 000027127
CATALOG_NAME	= SATCAT
OBJECT_NAME	= PSLV DEB
INTERNATIONAL_DESIGNATOR	= 2001-049BX
EPHEMERIS_NAME	= NONE
COVARIANCE_METHOD	= CALCULATED
MANEUVERABLE	= N/A
REF_FRAME	= ITRF
GRAVITY_MODEL	= EGM-96: 36D 360
ATMOSPHERIC_MODEL	= JBH09
N_BODY_PERTURBATIONS	= MOON,SUN
SOLAR_RAD_PRESSURE	= YES
EARTH_TIDES	= YES
INTRACK_THRUST	= NO
COMMENT COVARIANCE_SCALE_FACTOR = 1.000	
COMMENT EXCLUSION_VOLUME_RADIUS = 1 [m]	
COMMENT OD_DATA_SOURCE = ASW	
RECOMMENDED_OD_SPAN	= 2.68 [d]
ACTUAL_OD_SPAN	= 2.68 [d]
OBS_AVAILABLE	= 92
OBS_USED	= 92
TRACKS_AVAILABLE	= 31
TRACKS_USED	= 31
RESIDUALS_ACCEPTED	= 99.4 [%]
WEIGHTED_RMS	= 1.335
COMMENT Apogee Altitude = 608 [km]	
COMMENT Perigee Altitude = 581 [km]	
COMMENT Inclination = 97.6 [deg]	
CD_AREA_OVER_MASS	= 0.064060 [m**2/kg]
CR_AREA_OVER_MASS	= 0.004320 [m**2/kg]
SEDR	= 0.001988 [W/kg]
COMMENT DCP Density Forecast Uncertainty = 2.439348709999999976e-01	
COMMENT DCP Sensitivity Vector RTN Pos = -1.460809382574394988e+02	
1.162086123869119001e+04 2.328580003735519099e+00 [m]	

NASA Spacecraft Conjunction Assessment and Collision Avoidance Best Practices Handbook – Volume 2, Technical Appendices

```
COMMENT DCP Sensitivity Vector RTN Vel = -1.266771874844411983e+01
9.210463809680000113e-02 2.896832152199999792e-03 [m/sec]
X = -5.772760285552115420e+03 [km]
Y = 3.715340023984103027e+03 [km]
Z = 1.097267173051542386e+03 [km]
X_DOT = -1.623105746025805263e-01 [km/s]
Y_DOT = 1.909435378341858058e+00 [km/s]
Z_DOT = -7.411893064038498125e+00 [km/s]
CR_R = 1.301605442775661004e+03 [m**2]
CT_R = -1.016274964117690106e+05 [m**2]
CT_T = 8.089918981945027597e+06 [m**2]
CN_R = -2.429169841325214918e+01 [m**2]
CN_T = 1.622404315843846007e+03 [m**2]
CN_N = 5.295809561338548122e+01 [m**2]
CRDOT_R = 1.107819221070270004e+02 [m**2/s]
CRDOT_T = -8.818682020919026399e+03 [m**2/s]
CRDOT_N = -1.774955774407904974e+00 [m**2/s]
CRDOT_RDOT = 9.613108641163835344e+00 [m**2/s**2]
CTDOT_R = -8.323103272534031483e-01 [m**2/s]
CTDOT_T = 6.408717576511170932e+01 [m**2/s]
CTDOT_N = 1.710008118112443104e-02 [m**2/s]
CTDOT_RDOT = -6.985970698234267307e-02 [m**2/s**2]
CTDOT_TDOT = 5.374822317214000313e-04 [m**2/s**2]
CNDOT_R = -2.223716517953015087e-02 [m**2/s]
CNDOT_T = 2.006038160752205091e+00 [m**2/s]
CNDOT_N = -8.268733544575995742e-03 [m**2/s]
CNDOT_RDOT = -2.190799544182556007e-03 [m**2/s**2]
CNDOT_TDOT = 1.263112609890000012e-05 [m**2/s**2]
CNDOT_NDOT = 4.640061346839999712e-05 [m**2/s**2]
```

Processing the CDM given above indicates the following J2000 Earth-centered inertial frame TCA position covariance matrices for the primary and secondary objects (in m²)

$$\mathbf{A}_p = \begin{bmatrix} +4.6400\text{e}5 & -3.0303\text{e}5 & +2.5858\text{e}5 \\ -3.0303\text{e}5 & +1.9815\text{e}5 & -1.6896\text{e}5 \\ +2.5858\text{e}5 & -1.6896\text{e}5 & +1.4421\text{e}5 \end{bmatrix} \quad (\text{N-23})$$

and

$$\mathbf{A}_s = \begin{bmatrix} +3.1240\text{e}5 & +4.8579\text{e}4 & +1.5580\text{e}6 \\ +4.8579\text{e}4 & +7.5911\text{e}3 & +2.4237\text{e}5 \\ +1.5580\text{e}6 & +2.4237\text{e}5 & +7.7713\text{e}6 \end{bmatrix} \quad (\text{N-24})$$

Summing these according to eq. (N-17) yields the original, uncorrected relative position miss-vector covariance matrix

$$\mathbf{A}_m = \begin{bmatrix} +7.7639\text{e}5 & -2.5445\text{e}5 & +1.8166\text{e}6 \\ -2.5445\text{e}5 & +2.0574\text{e}5 & +7.3414\text{e}4 \\ +1.8166\text{e}6 & +7.3414\text{e}4 & +7.9155\text{e}6 \end{bmatrix} \quad (\text{N-25})$$

Notably, the numerical values of the elements of these covariance matrices depend on the specific choice of the J2000 inertial reference frame; using another inertial frame will yield different numerical values. However, the determinants and eigenvalues of these matrices are frame-independent quantities and provide a means of comparing the consistency of covariances calculated in different inertial reference frames. The Pc values calculated from these covariances are also frame-independent quantities.

Eq. (N-18) provides the means to calculate the corrected relative position miss-vector covariance matrix. The atmospheric density 1-sigma relative uncertainties for the primary and secondary can be extracted directly from the CDM

$$\sigma_{p/g} = 0.24552 \quad \text{and} \quad \sigma_{s/g} = 0.24393 \quad (\text{N-26})$$

After conversion to the J2000 inertial frame, the corresponding sensitivity vectors are (in m)

$$\mathbf{G}_p = \begin{bmatrix} +2.6566\text{e}3 \\ -1.7317\text{e}3 \\ +1.4799\text{e}3 \end{bmatrix} \quad \text{and} \quad \mathbf{G}_s = \begin{bmatrix} -2.2834\text{e}3 \\ -3.5515\text{e}2 \\ -1.1390\text{e}4 \end{bmatrix} \quad (\text{N-27})$$

Inserting these values into eq. (N-18) yields the relative position miss-vector covariance that has been corrected by removing the cross-correlated components (in m²)

$$\mathbf{A}_m = \begin{bmatrix} +1.5030\text{e}6 & -4.3476\text{e}5 & +3.8312\text{e}6 \\ -4.3476\text{e}5 & +1.3207\text{e}5 & -1.0764\text{e}6 \\ +3.8312\text{e}6 & -1.0764\text{e}6 & +9.9346\text{e}6 \end{bmatrix} \quad (\text{N-28})$$

This corrected covariance matrix differs significantly from the uncorrected covariance given by eq. (N-25), which leads the derived 2D-Pc estimates to differ by a factor of two, as previously noted. The uncorrected covariance in eq. (N-25) has the following frame-independent eigenvalues (in m²): $\lambda_{m,1} = 8.3512\text{e}6$, $\lambda_{m,2} = 5.4626\text{e}5$, and $\lambda_{m,3} = 2.0282\text{e}2$. The corresponding

determinant is $9.2523\text{e}14 \text{ m}^6$. The square roots of the eigenvalues correspond to the following sigma values (in m): $\sigma_{m,1} = 2.8898\text{e}3$, $\sigma_{m,2} = 7.3909\text{e}2$, and $\sigma_{m,3} = 1.4241\text{e}1$. The determinant for the corrected covariance in eq. (N-28) is $8.1020\text{e}13 \text{ m}^6$, with sigma values of $\sigma_{m,1} = 3.3961\text{e}3$, $\sigma_{m,2} = 1.8959\text{e}2$, and $\sigma_{m,3} = 1.3980\text{e}1$. As mentioned previously, these quantities can be used to check for consistency of covariances calculated using a different inertial reference frame.

Calculating the 6×6 miss-state covariance matrix, $\mathbf{P}_m$, requires the use of eqs. (N-19) through (N-22), but otherwise is similar to the calculation described above. The 2D-Nc and 3D-Nc semi-analytical Pc estimation methods require such full 6×6 corrections (Hall 2021), as does the from-TCA Monte Carlo method.

Pc Estimates with and without Covariance Correlation Correction

Conjunction 2D-Pc values calculated using the miss-vector covariance in eq. (N-18) represent collision probabilities corrected for global atmospheric cross-correlation effects. These can differ from the uncorrected 2D-Pc values calculated using the covariance in eq. (N-17). As shown in Figure N-12, analysis of archived conjunctions indicates that this correction usually does not change Pc values appreciably, except in a minority of conjunctions that have both elevated drag-related energy dissipation rates and an appropriate combination of orbital geometries. Among this minority, however, the corrections can potentially elevate Pc values by a factor of 1.5 or more (as shown on the right side of the plot in Figure N-12), meaning that accurate and conservative risk assessments for these cases rely on applying the covariance cross-correlation corrections made possible by the ASW system's recent CDM modifications.

For low-velocity or multi-conjunction interactions, the collision probability using the 2D-Nc or 3D-Nc methods (described previously in section value) can be calculated using a miss-state covariance matrix of the same form as given in eq. (N-21), which then can be used to estimate the Pc value for the interaction, as explained in more detail by Hall (2021).

N.9 Monte Carlo Pc Calculation Techniques

Analytic approaches to Pc calculation are more computationally efficient than Monte Carlo methods, especially the conjunction plane two-dimensional Pc method. However, as discussed previously, analytic methods require certain enabling assumptions that are not necessarily valid for all conjunctions. Monte Carlo approaches require fewer enabling assumptions, but they are not typically employed as the first method of computation. Instead, Monte Carlo methods are usually reserved for cases that are suspected of violating the enabling assumptions of the analytic methods. As described by Hall et al. (2018) and Hall (2021), there are two strains of the Monte Carlo method that are regularly employed:

- “Monte Carlo from epoch,” in which the orbit-determination-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 relatively short propagations to each trial’s new TCA are required.

Each of these two approaches will be discussed in turn.

N.9.1 Monte Carlo from Epoch

The principal appeal of calculating the Pc using the Monte Carlo from epoch approach is that it requires almost no simplifying or restrictive assumptions, making it as close to a “gold standard” for Pc 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 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 results in a simulated collision and a “hit” is registered; if not, it is considered a “miss.” This sampling, propagation, and 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 Pc estimate. There are algorithms that can be applied to estimate the confidence interval on the Monte Carlo Pc after a given number of trials due to event counting uncertainties (e.g., MATLAB’s *binofit* function).

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” property of the result is lost.

- Correlation between the primary and secondary covariances, as described in the previous section, should be considered. This correlation can be modeled during the random draw process by forcing correlation in the sample set of 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 occur and therefore 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 (e.g., 10^{-5} or above). Computation times for smaller P_c events can be prohibitively long. For example, validating a P_c estimate of $\sim 10^{-7}$ at the 95% confidence level for a conjunction with TCA five days from the orbit determination epoch time would require an estimated two years of execution time on a 20-CPU, reasonably modern server (Hall et al. 2018). For this reason, it is typically necessary to reserve the Monte Carlo from epoch 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 such reuse of samples violates the conditions for the proper application of the formulae to calculate P_c uncertainty 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 such 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 Pc estimation.

A study effort discussed in greater detail in the next section determined that the Monte Carlo from epoch method 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 Pc values using methods formulated for temporally isolated close approaches. 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 Pc calculation for that encounter, there is no guarantee that another encounter in the sequence may not actually possess a higher Pc due to different covariance orientations even though it has a larger nominal miss distance. Second, calculating a single encounter Pc at each of the close approaches and then amalgamating these using the following formula (derived from DeMorgan’s Law of Complements)

$$P_{Tot} = 1 - \prod_{i=1}^n (1 - P_{ci}) \quad (N-29)$$

potentially overestimates the overall collision likelihood because it presumes that the individual events in the sequence are statistically independent, but in fact they may not be, especially if blended in time. This total probability estimate matches that given in equation (N-8) and provides an upper bound for the amalgamated risk over the time interval of interest. For maximum efficiency, Monte Carlo from epoch would then be run operationally only in cases in which this upper bound is above a mitigation threshold and there is interest in determining whether the higher-fidelity Monte Carlo 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 overall rectitude of the Pc calculation. As stated earlier, lower-Pc conjunctions may present intractable Monte Carlo execution times, but if one wishes only to ensure that the Monte Carlo Pc fall below the mitigation threshold (rather than establish a high-fidelity Pc value), this can usually be accomplished with far fewer Monte Carlo trials.

Optimal application of the Monte Carlo from epoch Pc estimation method does not entail Pc evaluation over a relatively short time interval bracketing a single conjunction’s TCA, but rather over a more extended interval that spans several close approach encounters. For example, the collision likelihood between two objects would not be evaluated at the nominal TCA for a single conjunction, but, perhaps, over a risk assessment interval projecting forward seven days. This multiday interval would not only 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 (Hall et al. 2018 and Hall 2021). 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-13 (which shows the same conjunction sequence as Figure N-9 from the

earlier discussion of two- and three-dimensional N_c calculation methods). The pink shaded area in Figure N-13 shows the Monte Carlo P_c estimation confidence region, and the pink line shows the best estimate Monte Carlo result. As can be seen, the black upper line, which is the upper bound 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-13 Three-Dimensional N_c Temporal Risk Plot with Monte Carlo from Epoch Result Overlay (in Pink)

Because it is complicated to set up the execution environment for the Monte Carlo from epoch calculation, and because “gold standard” results require assembling extensive 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 estimation will remain with NASA CARA. However, a more computationally efficient mode of Monte Carlo estimation, which is serviceable for several different applications and is easier to obtain and employ operationally, is described in the next section.

N.9.2 Monte Carlo from TCA

A much more computationally efficient variant on Monte Carlo from epoch, which has been used by conjunction assessment practitioners for some time, is Monte Carlo conducted from TCA, or two-body Monte Carlo (TBMC) estimation. 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. This means that an efficient two-body motion propagation scheme usually provides an accurate trajectory approximation, 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). This specific method of Monte Carlo from TCA is also referred to as “two-body Monte Carlo” Pc (TBMC-Pc) estimation. To use the TBMC-Pc method, 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-14. 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-14 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 rendered 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. Hernando-Ayuso et al. (2023) discuss using equinoctial elements for near-linear orbit uncertainty propagation in more detail.

The detailed from-TCA Monte Carlo method equinoctial element-based sampling and propagation procedure is summarized as follows:

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 conjunction assessment applications, and they are described below.

The first of these study efforts was performed as part of the previously cited analysis 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-15. The top panel shows a scatter plot of the P_c calculated by Monte Carlo from epoch versus that from Monte Carlo from TCA. The intersection of each “plus” sign gives the scatter-plot point, and the length of the plus-symbol bars indicates the uncertainty of the calculation. One can see that the agreement is strong because all the points are close to the dashed line that indicates perfect equality. The bottom panel plots the base-ten logarithm of the ratio of P_c values estimated using the Monte Carlo from epoch method to those estimated with the Monte Carlo from TCA method. 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 10^{-3} , 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 method and Monte Carlo from epoch method results.




Figure N-15 Comparative Results of 373 High-Pc Conjunctions⁷

The second study effort involved a comparison between the Monte Carlo from TCA Pc estimation method and the two-dimensional Pc method (Hall 2019b), which was later extended to compare the from TCA method to two- and three-dimensional Nc methods (Hall 2021). In these studies, 63,603 temporally isolated conjunction events with two-dimensional Pc values greater than 10^{-7} were subjected to Pc calculation by the two-dimensional Pc method, as well as the two- and three-dimensional Nc methods, and all of these were compared to the Monte Carlo from TCA method. The comparative results are shown in Figure N-16 (which refers to Pc values estimated using the Monte Carlo from TCA method as TBMC-Pc values). These plots are similar to those already shown in Figure N-8, except in this case the colored symbols represent the worst kind of Pc estimation failure, i.e., conjunctions in which the analytical calculations significantly underestimate the Monte Carlo from TCA method Pc values. The left panel shows that, although the two-dimensional Pc method performs reasonably well for the vast majority of temporally isolated conjunctions, it underestimates from TCA Pc values by a factor of 1.5 or more in 0.258% of the investigated cases. (See the top legend of the left graph.) For those cases that showed 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.

⁷ From Hall et al. 2018.

As an aside, some of these differences between the conjunction plane two-dimensional Pc estimates and the Monte Carlo method estimates are considerable; there are several cases in which the two-dimensional Pc estimate understates the Monte Carlo from TCA estimate by more than an order of magnitude, as shown in the red in the left panel of Figure N-16. The center and left panels of Figure N-16 show that the three- and two-dimensional Nc methods do not similarly underestimate the from TCA Pc method values, except in a few statistically insignificant cases.




Figure N-16 Comparative Results of 63,603 Conjunction Events⁸

Originally, Monte Carlo from TCA was advanced as a robust and computationally tractable way to ensure reliable Pc calculations in the face of occasional miscarriage of the two-dimensional Pc algorithm. However, Figure N-16 demonstrates that the development and testing of the two- and three-dimensional Nc analytic methods have led NASA CARA to recommend that these methods supplant routine use of Monte Carlo from TCA. One may still resort to this more computationally efficient Monte Carlo method if desired, but testing indicates that the Nc calculation methods outperform Monte Carlo from TCA (in that it accurately matches Pc estimates for isolated conjunctions and also provides a cumulative risk upper bound for repeating conjunctions) and is at least an order of magnitude more computationally efficient, especially for events with small Pc values. Monte Carlo from TCA is a capability available for download from the NASA CARA software repository,⁹ and O/Os who are currently using it or a similar implementation are not urged to take it out of service, but as a Pc calculation approach, it does not seem to offer any enduring advantage over the two- and three-dimensional Nc methods.

Finally, a dedicated study on covariance Gaussianity was also conducted (Lechtenberg 2019b). The same set of 44,000 conjunctions used by Hall (2019b) 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

⁸ From Hall 2019b and Hall 2021.

⁹ Specifically, the *MCWorkbench* SDK function estimates TBMC-Pc conjunction values using CDM test files as input. See Section 7, Contact Information in this document for the CARA software repository URL.

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 such a large fraction of the investigated cases show good agreement between the Monte Carlo at TCA and the two-dimensional P_c , clearly the strict 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 since 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 conjunction assessment. Of course, fresh tracking data are always appreciated as they shorten the propagation interval and lend additional confidence to the solution.

N.10 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 because 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 P_c calculation, however, it is also one of the governing parameters of the calculation: the P_c 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 increases the calculated P_c by a factor of nine, or nearly one order of magnitude. Because of this sensitivity, it is important not to overstate or artificially enlarge the hard-body radius simply for

“conservatism” because 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) study several 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 conjunction assessment, 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 recent growth through the deployment of large constellations, this particular strategy merely creates additional false alarms that needlessly 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 size could be obtained either from known or published dimensions (for intact spacecraft or rocket bodies) or estimated from remote sensing data, such as radar cross-section or photometric brightness measurements. In the latter case, the estimated hard-body radius size of the secondary may also have an associated uncertainty estimate, which can also be incorporated into the Pc estimation process (as discussed in more detail below). There are also proposed techniques that use encompassing shapes other than spheres; for example, Alfano et al. (2022) propose using a system of bounding rectangular prism dimensions for many active satellites, as this type of shape more properly aligns with current spacecraft design.
3. **Maximum projected area into any plane.** Since the circumscribing sphere described above most often ends 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 Pc 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 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 CAD¹⁰ 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 Pc 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 Pc 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) describes a method to decompose complex shapes into individual triangles and use an equivalent-area method to evaluate the Pc for each triangle; the composite Pc is simply the sum of the Pc 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.

To facilitate Pc estimation, CARA has undertaken an effort to estimate hard-body radii of unknown orbiting objects based on radar cross section (RCS) measurements obtained by the Space Fence radar system (Baars and Hall 2022; Hall and Baars 2023). Even after accounting for radar calibration irregularities and occasional outlier data points, such RCS measurements have considerable point-to-point scatter, and the process of converting the RCS data into hard-body radius estimates introduces considerable additional uncertainty. The resulting hard-body radius uncertainty PDFs are non-Gaussian, although they can be roughly approximated using a log-normal distribution. Notably, RCS-based hard-body radius estimates are typically uncertain by factors of two to three, so it is essential that Pc estimates account for these uncertainties. For an unknown secondary involved in a conjunction, the RCS-based uncertainty PDF can be used to

¹⁰ Computer-aided Design

calculate a mean hard-body radius estimate, $\bar{R}_2$, and an associated variance, $\sigma_{R_2}^2$. As mentioned previously, for many conjunctions, the estimated Pc value varies roughly in proportion to the square of the hard-body radius. In these cases, the appropriate hard-body radius to use for Pc estimation is given by the effective combined hard-body radius

$$R_{\text{eff}} = \sqrt{(R_1 + \bar{R}_2)^2 + \sigma_{R_2}^2} \quad (\text{N-30})$$

with R_1 representing the primary object's hard-body radius, usually established using methods 2 or 3 listed above. (The current recommendation for the less often used method 4 is to increase the projected shape of the primary object outwardly on the conjunction plane by the length $\bar{R}_2 + 3\sigma_{R_2}$ to produce a conservative Pc estimate, although CARA continues to research a more accurate approach in these cases.) So for most cases, CARA recommends calculating Pc values using the combined hard-body radius given in equation (N-30) for all conjunctions involving secondary objects with uncertain sizes, partly because of the formula's relative simplicity, but primarily because using the effective radius approach reproduces reasonably well the more accurate RCS-based Pc estimates calculated using Monte Carlo analyses. (See Hall and Baars 2023 for details.)

Currently, Space Fence RCS measurements sufficient for size estimation are available for over 90% of CARA conjunctions involving unknown secondary objects (Baars and Hall 2022; Hall and Baars 2023). Analysis indicates that about 98% of such conjunctions involve secondary objects with mean hard-body radius less than 35 cm, and only about 0.3% have $\bar{R}_2 > 1.5$ m. CARA uses a multi-faceted approach to estimate secondary sizes. For known secondary objects for which published dimensions are available (e.g., active or retired payloads, rocket-bodies), the approach uses the known hard-body radii with zero uncertainty, tabulated along with a reference for the source of the size information. For unknown objects with sufficient RCS data (which includes a large fraction of tracked LEO orbital debris objects), the CARA approach tabulates the mean estimated hard-body radius and the 1-sigma uncertainty, which can be used in equation (N-30). For reference, this approach also tabulates the median hard body radius values, as well as the associated 95% and 99% confidence intervals, as estimated from the actual RCS-based uncertainty PDF. For unknown objects with insufficient RCS data, one can assign a reasonably conservative default size estimate of 1.5 m. (Published size estimates are included in the *DataFiles/HBRMassEstimates* directory in the NASA CARA software repository. See Section 7, Contact Information in this document for the specific URL.)

Appendix O. Collision Consequence

O.1 Introduction

It is common in conjunction assessment 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) and 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. Conjunction risk assessment metrics such as the P_c are simply measures of collision likelihood, which is only part of the overall risk assessment. Because conjunction assessment 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 mission operators have 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 conjunction assessment activities during mission lifetimes. One of the purposes of the present document is to lay out best practices for conjunction assessment 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, should it result 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; Krisko 2011). 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 not the other, resulting in less debris. In the ODPO model, the catastrophic/non-catastrophic distinction is governed by the following relationship, based on the relative kinetic energy of the encounter:

$$\varepsilon = \frac{\min(M_1, M_2)}{\max(M_1, M_2)} \left(\frac{V_{rel}^2}{2} \right) \quad (\text{O-1})$$

in which V_{rel} is the relative speed of the two satellites, M_1 is the mass of the primary satellite, and M_2 is the mass of the secondary satellite. When the relative kinetic energy exceeds 40,000 Joules per kilogram, ε_* , the collision is expected to be catastrophic and produce much larger amounts of debris. The ODPO has further developed a relationship to estimate the number of debris pieces greater than a specified size that a collision will generate, which is used by the EVOLVE orbital debris population model (Johnson et al. 2001; Krisko 2011):

$$F_{EV}(M_1, M_2) = \begin{cases} 0.1(V_{rel}^2 \min(M_1, M_2))^{0.75} L_c^{-1.71} & \text{if } \varepsilon \leq \varepsilon_* \text{ (noncatastrophic)} \\ 0.1(M_1 + M_2)^{0.75} L_c^{-1.71} & \text{if } \varepsilon > \varepsilon_* \text{ (catastrophic)} \end{cases} \quad (\text{O-2})$$

in which L_c is the characteristic length (size) above which one wishes to determine the number of debris pieces. Lechtenberg and Hejduk (2019) and Lechtenberg (2019a) classify “trackable” fragments as those having characteristic lengths exceeding a cutoff of $L_c = 0.05$ m. The present discussion adopts this value; other analyses may use different cutoffs, such as $L_c = 0.1$ m, which changes the fragment numbers estimated using equation (O-2). For catastrophic collisions, the number of fragments is determined by the sum of the two satellites’ masses; otherwise, the ODPO EVOLVE break-up model uses the product of the mass of the lighter satellite and the conjunction’s squared relative velocity. (Note: proper usage of equation (O-2) requires V_{rel} to be specified in km/s, masses in kg, and L_c in m, as described by Krisko 2011.)

The relationships in equations (O-1) and (O-2) conform to first-order intuition regarding expected relative levels of debris production. Any child who has played with dirt clods can testify that a greater relative velocity between two colliding 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 cratering the larger one and, while itself fully fragmenting, leaving the larger object largely intact. This full versus partial fragmentation division introduces a discontinuity in the debris production relationship given in equation (O-2). For instance, Figure O-1 shows the extremely broad range of debris production from a collision involving a 3000 kg primary satellite. In this case, the expected number of trackable ($L_c > 5$ cm) fragments indicated by the ODPO model varies from near zero up to many thousands—all of which represent real possibilities because the range of secondary object masses and relative velocities shown in Figure O-1 are encountered frequently in practice.

Since the debris production potential can vary so greatly among different conjunctions, it makes sense that conjunction remediation requirements should take great care when estimating the parameters required for its calculation. The conjunction relative velocity is easy to calculate accurately and is in fact provided directly as part of the CDM. Similarly, the HBR and mass of

the primary object are usually known accurately by the satellite O/O. The HBR and mass of the secondary object, however, are more elusive; in the majority of conjunctions, this object is a debris fragment for which there is no *a priori* information. In general, one must look to remote characterization methodologies to estimate unknown object HBR and mass values. For low-Earth orbits, fragmentation debris comprise the majority of unknown objects and are characterized primarily by radar observations, often resulting in highly uncertain sizes and masses (as discussed in more detail in the following sections).




Figure O-1 Range of Debris Production Possibilities¹¹

In formulating a collision consequence estimation methodology, an approach is desired that is informed by and accounts for the estimated sizes and masses of the secondary objects, and their associated uncertainties. For example, as a conservative approach to catastrophic collision determination and fragment production, Lechtenberg (2019) recommends using the upper 99.9th percent uncertainty quantile of secondary mass estimates. Alternatively, collision consequence metrics can be estimated as statistical expectation values to account for the full uncertainty distributions. Along these lines, Hall and Baars (2023) formulate a “fragmentation probability” metric that measures a conjunction’s probability of producing more than a threshold number of trackable fragments, F , given by:

$$\bar{P}_F = \mathbb{E}[P_F(R_1, R_2, M_1, M_2)] \approx \mathbb{E}[P_c(R_1 + R_2) \times U(F_{EV}(M_1, M_2) - F)] \quad (\text{O-3})$$

where R_1 and R_2 are the primary and secondary object’s hard-body radii (HBR), respectively, $U(-)$ is the unit step function, and $\mathbb{E}[-]$ is the expectation value operator, which indicates statistically expected results by averaging over all relevant random variables, thereby accounting for uncertainties in object HBR values and masses estimated using radar cross section (RCS) measurements.

¹¹ This figure represents a revised version of Figure 1 from Hejduk *et al* 2017, updated to incorporate the proper implementation of the NASA ODPO debris production model as described by Krisko (2011).

Figure O-2 illustrates the effect of secondary object mass uncertainties for a conjunction with a relative velocity magnitude of $V_{rel} = 1$ km/s. The left panel shows the number of collision fragments with $L_c > 5$ cm predicted by the equation (O-2), plotted as a function of the primary and secondary object masses. The center diagonal region shows primary and secondary mass combinations that produce catastrophic collisions. The off-diagonal regions correspond to non-catastrophic collisions. The right panel shows potential results for collisions with a hypothetical secondary object with a best-estimate mass of ~ 4 kg, but with significant mass estimation uncertainties, shown as vertical error bars corresponding to 70%, 95% and 99.9% confidence intervals. (As explained later, these error bars represent the typical large-amplitude and asymmetric uncertainties that result from radar-based mass estimation.) If this secondary object were to collide with a primary object known to have a mass of 10 kg, then the ODPO model would indicate a catastrophic collision, even accounting for the uncertainty of the secondary’s mass. The left-most point plotted with error bars illustrates this “statistically catastrophic” situation. However, if the primary instead were known to have a much larger mass of 800 kg, then the resultant collision would be “statistically non-catastrophic” as illustrated by the right-most point plotted with error bars. Primary objects with intermediate masses would produce a statistical mixture of catastrophic and non-catastrophic model collisions, as illustrated by the middle two points plotted with error bars in Figure O-2.




Figure O-2 Fragment Production Levels for 1 km/s Relative Velocity

The conservative approach of based on the 99.9% uncertainty mass quantiles to assess collision consequence risk uses the fragment production levels located at the top of the error bars in Figure O-2. A statistical expectation analysis averages risk metrics over the uncertainty distribution, accounting for the variations along the error bars in the plot. Both methods rely on knowledge of the underlying uncertainty distribution of mass estimates. The following sections summarize radar-based HBR and mass estimation for unknown objects in low-Earth orbit, focusing in particular on the resultant uncertainty distributions.

O.3 Estimating the Sizes of Unknown Objects

Stokely *et al* (2006) describe the NASA size estimation model (SEM), which provides a semi-empirical method to convert radar cross-section (RCS) measurements for an unknown satellite into a statistical ensemble of estimates for the “characteristic length” of the object. Notably, this characteristic length is not equivalent to the hard-body radius, but instead represents the average of the largest dimensions of an object measured along three orthogonal axes, which corresponds to the diameter for spherical objects. Applying the SEM to the n^{th} RCS measurement available for the j^{th} satellite yields:

$$D_{j,n} = \begin{cases} \lambda_{j,n} \sqrt{4z/\pi} & z > 5 \\ \lambda_{j,n} \sqrt[6]{4z/(9\pi^5)} & z < 0.03 \\ \lambda_{j,n} g_{SEM}(z) & \text{otherwise} \end{cases} \quad (\text{O-4})$$

with $D_{j,n}$ denoting the characteristic length (m) corresponding to the RCS measurement, $\lambda_{j,n}$ the radar wavelength (m), and $z = Y_{j,n}/\lambda_{j,n}^2$, with $Y_{j,n}$ representing the measured RCS value (m^2). Stokley *et al* (2006) provide an interpolation table for the function $g_{SEM}(z)$. Notably, the NASA SEM provides no direct means to estimate the error or uncertainty of RCS-based characteristic length estimates. This can be especially problematic if only one representative (e.g., median) RCS value is available per satellite. As described in Hall and Baars (2023), ideally numerous RCS measurements should be used per satellite over extended time periods (e.g., ~1,000 RCS measurements acquired over ~2 years) to analyze the resulting characteristic length distributions and estimate both the unknown hard-body radii and the associated uncertainties.

For CARA’s HBR estimation process (see Hall and Baars 2023), Space Fence RCS data are collected for up to two years prior to the analysis epoch date. The analysis is restricted to objects with twenty or more well calibrated and quality-curated Space Fence system RCS measurements, as described by Baars and Hall (2022). Next, the NASA SEM is used to estimate RCS-based characteristic length estimates for each satellite, which typically have considerable point-to-point variability or scatter. Sources of this scatter include the object’s aspect-dependent radar reflection properties and projected area (which vary in time due to rotation and changing observational geometry), as well as RCS measurement noise (which is often relatively minor).

The process uses a set of relatively small satellites that all have convex cuboid shapes in order to calibrate and de-bias the estimated HBR values. Specifically, the calibration analysis employs sizes obtained from the *Database and Information System Characterising Objects in Space* (DISCOS), which tabulates the bounding dimensions and general shape descriptions of a large number of known satellites.¹² The calibration analysis to convert characteristic sizes into estimated hard-body radii is restricted to satellites labeled as “box” shaped in the DISCOS database.¹³ Figure O-3 compares the empirical CDF of the ratio of calibrated RCS-based HBR values to the known values for 586 box-shaped satellites on the calibration epoch date of 2022-01-15. The CDF indicates that, among this population, calibration uncertainties limit the 90%

¹² From Flohrer *et al* (2013) and Mclean *et al* (2017).

¹³ <https://discosweb.esoc.esa.int/>

confidence accuracy range of the estimated HBR values to within a factor of 2.4 for potential underestimations, and to within a factor of 3.1 for potential overestimations.




Figure O-3 CDF of the Ratio of Calibrated HBR Values to Known HBR Values for 586 Satellites¹⁴

O.4 Estimating the Masses of Unknown Objects

Hall and Baars (2023) describe two approaches to estimate the mass of an unknown object; the first involves the use of the orbit determination-based ballistic coefficient (BC) while the second uses the orbit determination-based solar radiation pressure coefficient (SRPC). For satellites that experience measurable atmospheric drag orbital perturbations, the BC estimation method is preferred. However, for satellites with a perigee height above 450 km in altitude, the SRPC method mass estimates are somewhat more accurate.

O.4.1. RCS + Ballistic Coefficient Method

Combining RCS-based sizes and orbit determination based ballistic coefficients provides a means to estimate the masses of LEO satellites that experience measurable atmospheric drag. This “RCS+BC” mass estimation method approximates the characteristic mass of a LEO satellite using an expression formulated previously¹⁵:

$$\mathcal{M}_{D,j} = A_j B_j C_{D,j} \approx (\pi D_j^2 / 4) B_j C_{D,j} \quad (\text{O-5})$$

In this expression, $\mathcal{M}_{D,j}$ indicates the mass that characterizes the atmospheric drag experienced by the j^{th} satellite, which is proportional to the product of the projected area, A_j , the inverse ballistic coefficient, B_j , and the drag coefficient, $C_{D,j}$. The analysis uses RCS-based characteristic projected areas to estimate satellite cross sectional areas projected normal to the local atmospheric flow, $A_j \approx (\pi D_j^2 / 4)$, which is a rough approximation, especially for highly elongated or flattened objects (see Hall and Baars 2023). The mass calibration process uses box-

¹⁴ From Hall and Baars (2023)

¹⁵ From Lechtenberg 2019a.

shaped DISCOS satellites, as described previously for the HBR calibration process. For the calibration epoch date of 2022-01-15, sufficient RCS data and VCM BC estimates are available for $J = 554$ of the DISCOS box-shaped satellites.




Figure O-4 CDF of the Ratio of Calibrated RCS+BC Mass Estimates to Known Masses for 554 Satellites¹⁶

Figure O-4 compares the empirical CDF of the ratio of calibrated mass estimates to the known values for the 554 satellites. The CDF indicates that, among this population, calibration uncertainties limit the 90% confidence accuracy range of the RCS+BC mass estimation process to within a factor of 4.1 for potential mass underestimations, and to within a factor of 10.4 for potential mass overestimations. (This skewed uncertainty distribution provides the basis for the large-amplitude and asymmetric representative error bars shown in the left panel of Figure O-2.)

O.4.2. RCS + Solar Radiation Pressure Coefficient Method

Combining RCS-based sizes and orbit determination-based solar radiation pressure coefficients provides a means to estimate the masses of satellites that experience measurable solar radiation pressure (SRP) perturbations. This “RCS+SRPC” mass estimation analysis approximates the characteristic mass of a LEO satellite using an expression similar in form to that used for RCS+BC based mass estimation given by equation O-5 (see Hall and Baars 2023):

$$\mathcal{M}_{R,j} \approx (\pi D_j^2 / 4) G_j C_{R,j} \quad (\text{O-22})$$

In this expression, $\mathcal{M}_{R,j}$ indicates the mass that characterizes SRP perturbations experienced by the j^{th} satellite, which is proportional to the product of three quantities: the area projected towards the incident sunlight (again approximated using $\pi D_j^2 / 4$), the inverse solar radiation pressure coefficient, G_j , and the reflectivity coefficient, $C_{R,j}$. See Vallado (2001) for more details on orbital SRP perturbations. The Hall and Baars (2023) analysis restricts RCS+SRPC mass estimation to satellites with perigee altitudes above 450 km altitude to ensure that imperfectly modeled atmospheric drag perturbations do not bias the solar radiation pressure coefficients estimated in the orbit determination analysis, which has been observed to occur for LEO

¹⁶ From Hall and Baars (2023)

satellites at lower altitudes. For the calibration epoch date of 2022-01-15, sufficient RCS data and VCM SRPC estimates are available for $J = 303$ of the DISCOS cuboid satellites with perigees above 450 km altitude.




Figure O-5 CDF of the Ratio of Calibrated RCS+SRPC Mass Estimates to Known Masses for 303 Satellites¹⁷

Figure O-5 compares the empirical CDF of the ratio of calibrated mass estimates to the known values for the 303 satellites. The CDF indicates that, among this population, calibration uncertainties limit the 90% confidence accuracy range of the RCS+SRPC mass estimation process to within a factor of 3.8 for potential mass underestimations, and to within a factor of 8.4 for potential mass overestimations. (Again, this skewed distribution roughly corresponds to the large-amplitude, asymmetric error bars shown in Figure O-2.) For satellites with perigee altitudes above 450 km, the RCS+SRPC method provides somewhat more accurate mass estimates than the RCS+BC estimation method.

O.5. Availability of Satellite Size and Mass Estimates from NASA

As a service to the community, NASA CARA makes their satellite size and mass estimates available on their public Github site; this information can be accessed with the following link: https://github.com/nasa/CARA_Analysis_Tools/tree/master/DataFiles/HBRMassEstimates. The Github location includes the database itself, conference papers that explain the calculation methodologies in even greater data than the treatment in this present appendix, a users' guide, and software that automatically downloads size information from the ESA DISCOS database. Due to licensing restrictions, CARA is not permitted to distribute DISCOS information as part of the NASA size/mass database; but after establishing a free DISCOS account, users can employ the provided software to download all the needed DISCOS information to complete the size-mass database. This publicly-available version of the NASA size-mass database imposes a lower-limit size floor of 10cm (which will appear as 5cm in the database, since it gives hard-body radii, which are half of the estimated characteristic dimension of the object). US government users can request from NASA a special version of this database that does not impose this size floor.

¹⁷ From Hall and Baars (2023)

O.6. The Probability of Exceeding a Threshold Number of Trackable Fragments

Recall equation O-3 that defines a fragmentation probability metric, which is calculated using an expectation value to determine the probability that a conjunction will generate more than a threshold number of trackable fragments. As described by Hall and Baars (2023), for known-on-unknown object conjunctions, calculating this expectation value represents integrating over a set of five random variables that account for the RCS-based size and mass estimation uncertainties of the secondary object. Using the Monte Carlo method, these random variables are used to generate a set of R_2 and M_2 values for a number of Monte Carlo trials, N_{MC} . The fragmentation probability is then estimated as:

$$\bar{P}_F \approx \frac{1}{N_{MC}} \sum_{n=1}^{N_{MC}} [P_c(R_1 + R_{2,n}) \times U(F_{EV}(M_1, M_{2,n}) - F)] \quad (\text{O-24})$$

For $F = 0$, it is straightforward to show that $\bar{P}_F = \bar{P}_c$, i.e., the expected collision probability. More generally, a conjunction's fragmentation probability cannot exceed the collision probability, i.e., $\bar{P}_F \leq \bar{P}_c$. For groups like CARA that use the collision probability to assess close approach risks, this inequality property makes $\bar{P}_F$ a convenient choice to assess environmental risks, for two reasons. First, if the screening process indicates that $\bar{P}_c$ is less than a negligibly small level (i.e., 10^{-10}), then $\bar{P}_F$ also must be less than that level. Second, the inequality implies that all conjunctions that represent unacceptably high environmental risks (such as those with $\bar{P}_F \geq 10^{-4}$) are naturally contained within the set of conjunctions identified as high-level collision risks (with $\bar{P}_c \geq 10^{-4}$).

To demonstrate how collision and fragmentation probabilities can be used together, a threshold of $F = 1000$ trackable fragments was used to assess an unacceptably high level of risk to the orbital environment. Figure O-6 compares fragmentation and collision probabilities for known-on-unknown CARA conjunctions, calculated using this fragment production threshold. Among the conjunctions analyzed, 1,736 represent red-level close approach risks with $\bar{P}_c \geq 10^{-4}$. Among these, 121 (or 7%) also have $\bar{P}_F \geq 10^{-4}$, so they represent red-level environmental risks as well. Figure O-6 plots these dual high-risk conjunctions as red stars, illustrating how collision and fragmentation probabilities can be used together to identify the subset of events with the highest spacecraft and environmental risks. Pink diamonds and triangles represent conjunctions with red-level close approach risks, but lower environmental risks.




Figure O-6 Collision and Fragmentation Probabilities for Known-on-Unknown Conjunctions¹⁸

As discussed above, a fragment production threshold of $F = 1000$ indicates that 7% of the analyzed known-on-unknown conjunctions with $\bar{P}_c \geq 10^{-4}$ also have $\bar{P}_F \geq 10^{-4}$. This analysis uses $F = 1000$ as a threshold number of trackable fragments because such a collision would potentially represent a major source of LEO debris generation, comparable to the Iridium-Cosmos collision in 2009 which produced about 1,200 trackable fragments that remained in Earth orbit for longer than one year. Other production thresholds provide different assessments of environmental risk in a non-linear fashion. For instance, about 30% of the studied conjunctions have $\bar{P}_F \geq 10^{-4}$ for a fragment production threshold of $F = 100$. So, considering the range of $10^2 \leq F \leq 10^3$ indicates that between 30% and 7% of the high mission risk conjunctions also could be assessed to pose an elevated risk to the orbital environment.

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 posture generally, missions certainly may embrace the 10^{-4} P_c 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, especially in the presence of other relevant considerations, such as the inability to identify a truly adequate mitigation action that did not impose non-trivial mission compromise. Similarly, it could be concluded that the P_c threshold should be made more demanding for conjunctions that

¹⁸ This figure represents a revised version of Figure 10 from Hall and Baars (2023), updated to incorporate the proper implementation of the NASA ODPO debris production model as described by Krisko (2011).

would create very large amounts of debris were they to result in a collision. Finally, the use of fragmentation propensity may 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.

Appendix P. Event Actionability

P.1 Introduction

Although widely used, the term “conjunction event actionability” is somewhat unfortunate because it implies a broader scope than it merits. Its precise meaning is the assessment of the propagated 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 conjunction 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 thus are enabling credible collision risk calculations.

The propriety of considering the quality of the event’s input orbital data is obvious because 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 because, in the absence of durable information, 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, out of conservatism, be afforded mitigation actions if the calculated risk parameters exceed the threshold for mitigation actions. Rather, it is necessary 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 tracking-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 conjunction 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 conjunction assessment 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 expedited document update should these values be modified. 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 timespan of the tracking data for the fit, data density, tracking data sensor 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 can 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 and 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 conjunction 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 positions of a particular satellite. 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 determined by considering characteristics of the satellite's orbit, 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 conjunction assessment 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 conjunction assessment 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 satellite's 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 associated 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 (SRP), which is a function of a solar radiation pressure coefficient (SRPC) (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 the details of this retrieval may at first not seem important, in fact they 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 or solar radiation pressure 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 and ensuring a reasonable solution for the non-conservative force parameters (drag and solar radiation pressure). The algorithm operates by beginning with the maximum acceptable LUPI for a particular orbit and attempting to reduce it as much as possible while preserving sufficient observational data to endure a quality update, and if the maximum LUPI itself 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 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 conjunction risk assessment. 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 discarded 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 conjunction risk assessment, 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 conjunction risk assessment.

A table of these maximum and minimum LUPI lengths, as a function of orbital parameters, 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 by the data being assigned 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 (WRMS), 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 average of these ratios should equal unity (as will its square root). 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. Thresholds for 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 conjunction mitigation actions, so if one of these default covariances is encountered, the state data that accompany it are not suitable for conjunction 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 conjunction risk assessment and, if necessary, mitigation action planning.

P.9 Evaluation of Propagation from Epoch to TCA

Most conjunction 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 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 conjunction risk assessment, except in the few narrow instances called out below.

P.9.1 Propagated Covariance is Defective

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

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 conjunction 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 inactionable. 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 squared 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 Pc to zero (and sometimes introduce convergence problems in Pc 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 Pc of zero (because this is what very large covariances do). Embracing 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 inactionability; non-positive-semidefinite covariances do not.

P.9.2 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 fully to represent curvilinearly produced results. As discussed at greater length in Appendix N, while this debate continues in the critical literature, the conjunction assessment 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 conjunction assessment, based on NASA CARA’s analysis, this problem is not crippling for the categories of events that actually matter to conjunction risk assessment, and in any case, it can be addressed with a particular Pc calculation technique that is both straightforward and tractable.

P.9.3 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 in order to reach the current time. 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 that tracking data be received frequently because 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 were 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/LUPI. It is generally considered dangerous to try to propagate for a longer period of time 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.9.4 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. Having said that, it is also the case that the performance of the Anemomilos model during the most recent solar cycle (solar cycle 25) has been disappointing: normed on data from earlier solar cycles, the model is not producing representative results presently. A number of different mitigation activities are presently being considered, including simply turning off this particular modeling feature.

Unfortunately, it is not possible during regular operations to determine the degree to which Anemomilos is affecting prediction outcomes, although this model will engage only when the predicted Disturbance Storm Time (Dst) parameter value is smaller than -75 nanoTeslas, so one can consult space weather predictions to determine whether there is any possibility of the model's being employed. This general issue is addressed at greater length in the walk-through of a sample event in Appendix Q.

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 collision 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, the y-axis that for the secondary object, and the color indicating the resultant P_c value. However, 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 P_c changes little. This event is not sensitive to space weather mismodeling, and thus actionability is not strongly affected by solar storm developments and potential mismodeling.




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 should still be safe even in the presence of model error.




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.

¹⁹ From Hejduk et al. 2017.

²⁰ From Hejduk et al. 2017.




Figure P-3 SWTS Plot for Variation Situation²¹

P.10 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 conjunction 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 reliable conjunction risk assessment.

²¹ From Hejduk et al. 2017.

P.11 Annex

This annex 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	N	Y
		Rocket Body / Platform	24	N	Y
		Debris/Unknown	24	N	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 > 10000	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

Attribute	Object Type	Energy Dissipation Rate (W/kg)			
		< 0.0006		>= 0.0006	
		Min	Max	Min	Max
Drag	Payload	0.001	0.1	0.001	0.1
	Rocket Body	0.001	0.2	0.001	0.2
	Debris	0.001	10	0.001	10
Solar Radiation Pressure	Payload	0.001	0.1	0.001	0.1
	Rocket Body	0.001	0.2	0.001	0.2
	Debris	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. Notional Display Flow for Conjunction Event Processing

Q.1 Introduction

While appendices E through P in this Handbook address many of the technical aspects of the conjunction event processing discipline, this appendix illustrates how the individual pieces come together to perform on-orbit conjunction risk assessment. The operational orbital safety process is illustrated through a series of graphical displays that conjunction risk assessors use to determine whether a mitigation action is warranted for a particular conjunction and, if it is, what sorts of such actions might be effective in adequately reducing the collision risk.

The graphical displays in this appendix are a subset of those used by the NASA CARA group to perform risk assessment in protection of NASA's fleet of non-human space flight satellites, and the examples presented are for a recent actual event that resulted in the planning and execution of a mitigation action. To facilitate public release of the information, neither spacecraft is specifically identified, and the orbital data have been slightly altered.

NASA executes conjunction assessment screenings of its assets against the DOD space catalog once every eight hours, and for the orbit regime that contains the particular protected asset in use in this example, a screening volume of 0.5 km x 17 km x 20 km was used. This particular conjunction appeared at the 6.5 days-to-TCA point, which means that it was essentially identified as early as possible (using 7-day predicted ephemerides for the primary object). Because the initial P_c was greater than $1E-07$, this event was in "yellow" status from the very beginning, meaning that the collision likelihood was high enough to merit close attention as the event developed. As such, the conjunction risk assessment process was engaged from the beginning.

The displays in this appendix follow this event from initial detection to the mitigation action commitment point, which for this event was one day prior to the time of the closest approach. Commentary is provided for each display, and since there is always more that can be said about each display, the commentary is reasonably abbreviated and focused on the most salient points of each presentation.

Q.2 Primary and Secondary Object Background Information

In making a risk assessment determination, it is important to have information about the orbits and maintenance history of the primary and secondary objects as background. Different orbits have different properties and expected stability, and maintenance histories and parameters testify to the quality of the orbit determination available to analyze this event. It is helpful to display these data in consolidated form so that the parameters of interest can be viewed and apprehended quickly.

Q.2.1 Primary Object Information

Table Q-1 provides the orbital and orbit determination estimation information for this conjunction's primary object.

Table Q-1 Primary Satellite Vital Statistics

Orbital Information at TCA		Estimation Specifics	
Orbital Parameter	Value	Parameter	Value
Period (min)	97.6	Avg. Tracks Per Day	5.7
Perigee Height (km)	681.3	Num. Obs. In Span	237
Apogee Height (km)	721.7	Time of Last Observation	<24 hours
Inclination degrees)	98.2	WRMS	1.55
EDR (W/kg)	2.77e-04	Ballistic Coefficient (m ² /kg)	0.027
RCS	Large (>1m ²)	SRP Coefficient (m ² /kg)	0.005

In most cases, the primary object is much better maintained than the secondary object, meaning that the primary object's orbit determination information is generally good and therefore not determinative in assessing the event's overall data quality. Furthermore, if the object is screened using the owner/operator (O/O) ephemeris, there will typically not be any available orbit determination estimation information at all. But it is helpful to examine whatever information is available for the primary object to make sure the situation is well understood.

One observes that the primary is in a nearly circular, sun-synchronous LEO orbit with an average altitude of about 700km. The energy dissipation rate (EDR) value indicates that, while drag has an effect, it is relatively mild. The altitude is such that both drag and solar radiation pressure (SRP) should be modeled, and the two resultant solved-for values are reasonable. (Consult Appendix P for acceptable ranges of values for parameters such as this.) The weighted root mean square (WRMS) value here is right on the edge of being too high but probably still within tolerance. The tracking density is reasonably large, and tracking data have been included in the orbital solution that was collected within the last day, so propagation to TCA will not have to extend very far into the future.

Overall, this primary satellite is well maintained and not propagated excessively, so there is nothing to indicate any problems with the fit or propagation.

Q.2.2 Secondary Object Information

Basic orbital and orbit determination estimation information for the secondary object is shown in Table Q-2.

Table Q-2 Secondary Satellite Vital Statistics

OBJECT - Orbital Information @ TCA

Orbital Parameter	Value
Period (min)	98.7
Perigee Height (km)	674.9
Apogee Height (km)	718.5
Inclination (degrees)	81.2
EDR (W/kg)	5.19e-03
RCS	Small ($< 0.1\text{m}^2$)

OBJECT - Estimation Specifics

Parameter	Value
Avg. Tracks Per Day	0.5
Num. Obs. in Span	15
Time of Last Observation	< 48 hours
WRMS	3.41
Ballistic Coefficient (m^2/kg)	0.682
SRP Coefficient (m^2/kg)	0.285

OBJECT - Event Flags

Event Flag	Status
Single Station Tracking	Y
Increased Tasking Requested	Y
Increased Tasking Received	Y
State Update Consistency Check	Y

The orbit of the secondary object is not substantially different from that of the primary, although the inclination is somewhat lower. However, due presumably to a different area-to-mass ratio, the drag influence is much higher for the secondary object: the high EDR value indicates a significant drag effect, and this is confirmed by the large solved-for ballistic coefficient. The orbit altitude indicates that SRP should be solved for as well, and it is: the large value here also indicates a high area-to-mass ratio. The tracking levels for the secondary object are poor: 0.5 tracks/day is an order of magnitude smaller than the tracking density for the primary.

Examining the Conjunction Data Message (CDM) (not shown here) reveals that the 15 observations (obs) used in the latest orbit determination update for this object span a 9.5-day period. Thus, while the overall tracking situation is not the best, the length of the update interval (LUPI) is only half the maximum value and therefore tolerable. (Consult Appendix P for information about the LUPI parameter.) The fact that all the tracking comes from a single sensor is undesirable, but it is a circumstance that is observed occasionally with small debris objects, which this is given its radar cross section (RCS) value. The WRMS is high but below the quality threshold of 5 for a debris object. Data in the CDM also indicates that the propagation time to TCA is not all that long (a few days) and certainly much shorter than the length of the update interval.

Overall, while the maintenance history and propagation situation for this object is not the best, it is good enough to serve as a basis for risk assessment, presuming that other issues (e.g., space weather) do not ultimately jeopardize the credibility of the propagation.

Q.3 Primary and Secondary Object Trajectories

The ground trace plot shown in Figure Q-1 depicts the two satellite trajectories in relation to each other and identifies the geographical position of their point of closest approach.




Figure Q-1 Ground Trace Plot

In the above encounter, the closest-approach point—over Antarctica—is far away from the field of regard of any of the SSN ground sensors, meaning that this part of both satellites’ trajectories will not be modeled as well as if the ground sensor were directly under the point of closest approach. The relative velocity value indicated in Figure Q-1, almost 14 km/sec, is both very fast and typical for LEO conjunctions; it is certainly large enough to allow hyperkinetic assumptions about the conjunction to inhere and thus to allow the use of simplified approaches to calculating the probability of collision, such as the 2D Pc.

The approach angle may also be of some interest. A zero-degree approach angle describes a situation in which one object “chases” another, usually presenting lower-velocity encounters; a 180-degree approach angle constitutes a “head-on” collision. In both such instances, in-track error plays a smaller role in the conjunction dynamics because it controls the timing of any collision more than whether a collision would in fact take place. (This is of course only partially true because in-track error is usually correlated with radial error and thus will engender a radial component.) Values in between these extremes describe the typically encountered oblique situation, which is the case in this particular event being evaluated.

Q.4 Conjunction Plane Information

The Figure Q-2 visual, called a Conjunction Plane Plot, communicates certain information about the event. In interpreting this representation, it may be helpful to review the development of the 2D P_c calculation methodology outlined in Appendix N.




Figure Q-2 Conjunction Plane Plot

The relative miss vector lies along the x-axis, with one end at the origin and the other terminating at the blue dot. This blue “dot” is in fact a circle with a radius equaling the hard-body radius (HBR) that represents the combined sizes of the primary and secondary objects. The ellipses represent the combined uncertainty of the primary and secondary objects’ positions, here resolved from an ellipsoid to an ellipse in the plane in which a collision, should it occur, would take place.

For low-likelihood conjunctions, the blue dot is far outside of the three-sigma ellipse, and frequently the ellipse itself is collapsed onto the x-axis, indicating that the in-track error is dominating the conjunction. In the case being discussed, the situation is different: the blue dot is within the one-sigma error ellipse, and the ellipse itself has visible definition, indicating that the uncertainties in this situation are more varied.

Q.5 Collision Probability Evolution

The Figure Q-3 graphic gives a time history of the probability of collision (P_c) during this event’s development.




Figure Q-3 Probability of Collision Time-History Plot

While it is true that in most cases operational practice is grounded correctly on the P_c at the “mitigation action commitment point” (namely, that temporal point before TCA at which a decision must be rendered regarding whether a mitigation action will take place), the behavior of the P_c over the event’s history does help to indicate whether the event’s risk assessment data are stable enough to allow a durable risk assessment decision to be made.

Figure Q-3 shows the P_c time history:

- For the primary and secondary objects’ states and covariances based on the skin-tracking results (indicated as “ASW” in the legend), and
- For the primary object’s state and covariance derived from a submitted owner/operator ephemeris (indicated as “O/O” in the legend).

An additional feature of Figure Q-3 is that the graphic markers (square or circle) are filled in for the particular updates that received new tracking.

Figure Q-3 makes clear that the secondary object is not particularly well tracked; most of the updates, noted by the open markers, are based only on the update of the space weather coefficients and do not constitute a new state estimate resulting from fresh tracking. However, the significant state update that does contain new tracking (just a little less than two days to TCA) is barely under the mitigation threshold (red line), and after additional space weather

updates, the last such update before the mitigation action commitment point (one day) exceeds the red threshold. While there is a large variation observed in the Pc history, it is not so extreme as to indicate to an analyst that the situation is so unstable that a mitigation action cannot be planned or counselled wisely. As such, based (provisionally) on Figure Q-3, it would be recommended for this O/O to take a mitigation action.

Q.6 Miss Distance Evolution

While the Pc is the collision likelihood metric of choice, it is helpful to examine the miss distance history as well to see if it is manifesting a dramatically different pattern from the Pc history.




Figure Q-4 Total Miss Distance Plot

Miss distance is one of the principal inputs to the Pc, but the combined covariance and HBR are also important. There are cases in which the miss distance behaves in a manner that would seem to reduce collision risk but the covariance evolution is such that the collision likelihood remains high; in such situations, further investigation is prudent.

When comparing Figure Q-4 to the Pc history, one observes a situation in which the miss distance evolution is not in obvious lock-step with the Pc evolution, which is not an unusual outcome given what appears to be happening with the covariance evolution at the same time. But one also observes a situation for which the miss distance has moved to a very low value in the

last update. Even though this is not brought about via a tracking update but just space weather model updates, it is small enough to have driven the P_c above the mitigation threshold. While in principle it could change at the next update and push the P_c below the threshold, the situation is not so contradictory that one would assert that it has no predictive force.

As such, the miss distance behavior does not work against a conclusion that collision mitigation is appropriate.

Q.6.1 Miss Distance Component Behavior

In Figure Q-5, the individual miss distance components are examined, which is a more meaningful display of the time history of this particular datum.




Figure Q-5 Componentized Miss Distance Plot

Figure Q-5 also gives a time-history of miss distance behavior, but it separates the miss distance, helpfully, into three components centered on the primary object:

- The radial miss (in the direction toward the center of the Earth),
- The in-track miss (perpendicular to the radial miss in the plane of the orbit), and
- The cross-track miss (the direction perpendicular to this plane).

The blue circles represent the signed miss distance in the particular component of interest (and if filled in, tracking was received as part of the state estimate update that produced that miss distance), and the red squares represent the one-sigma uncertainty for that miss component.

Ideal behavior is a set of blue circles that essentially stay at the same value update to update, while the red squares shrink closer to that value as the propagation length is decreased from 7 days to ~1 day.

The in-track results above (middle graph) essentially exhibit this behavior: the in-track miss does not fluctuate wildly (although the y-axis scale should be kept in mind), and the uncertainties shrink nicely as the event develops.

The situation is somewhat less well-behaved for the cross-track graph (far right), but in the main, the actual miss does converge to a value while the uncertainties shrink.

The radial component, however, does not behave so nicely; and since this miss component is the most determinative for the Pc, it needs to be watched carefully.

- First, one should notice how much smaller the radial miss values are compared to the other two components. It is because of this small radial miss that this event is a high-risk event in the first place.
- Second, its evolution would have to be characterized as at least somewhat unstable: it changes sign and does not uniformly focus on a single value over time. However, such behavior is not truly unusual for the radial component.
- Third, when fresh tracking is finally obtained, the miss distance does jump and changes sign, but it jumps to a worrisome value (nearly zero) and does not stray far from that value as the event pushes forward to the mitigation action commitment point.

The fact that this key component of the high-risk designation stays reasonably stable coming into the mitigation action commitment point gives additional confidence that the reasonably high Pc value should be taken seriously.

Q.6.2 Normalized Component Miss Distance Value Comparison

Figure Q-6 is a different way to render the development of the componentized miss distance as the event builds to TCA. Rather than actual component miss distance values, Figure Q-6 represents normalized values: the change in componentized miss distance divided by the expected error in that component from the previous update. In essence, Figure Q-6 reflects how well or poorly each update behaves with regard to its expected error.

Ideal performance would be for updates to sit very close to the zero line: there would be strong agreement update to update, and any actual difference between predictions would be small compared to the prediction error or, to be more precise, ~68% of such ratios would fall within the +1 to -1 regions.




Figure Q-6 State Comparison Consistency Plot

The behavior shown in Figure Q-6 is quite desirable. Most values are very near the zero-line, and the few that are not do not deviate much beyond the one-sigma region. One can encounter situations in which gyrations beyond the three-sigma region occur throughout; in comparison to that, the above is very well behaved. It is also not unusual to see notable changes introduced as the result of additional tracking; the value for the radial component right after the two-days-to-TCA vertical line (red filled dot) is an example of this.


Overall, this is a reasonably well-behaved event whose final result can be trusted to yield an actionable result.

Q.7 Space Weather Considerations

The space weather situation is important to consider when ascertaining whether the orbital dynamics are expected to be sufficiently stable that the data examined at the mitigation action commitment point are expected to inhere until TCA.

Q.7.1 Solar Storms

The main area of interest in Figure Q-7 is the expectation of effects from solar storms such as coronal mass ejections (CMEs), and these are best predicted by examining the disturbance storm time (Dst) parameter prediction.



(Image from http://sol.spacenvironment.net/~sam_ops/current_data/Dst_streamB_forecast.jpg)

Figure Q-7 Solar Storm Prediction

The upper part of Figure Q-7 gives both a history and prediction of this parameter: the historical data are in black, the one-hour Anemomilos (the model used by the DOD to predict this parameter) given in blue, and the longer-term Anemomilos forecasts given in purple; the current time is indicated by the vertical dotted line.

Dips in Dst are the behavior of significance with a threshold of -75 nanoTeslas being the threshold used by CARA to indicate a disturbed situation. In Figure Q-7, one can observe a dip below this value, then a “recovery” increase, then some residual dipping before steadying off and heading to the zero value. The prediction shows a linear increase to the zero value, which indicates a stable situation.

The takeaway from Figure Q-7 is that, while the space weather situation was somewhat perturbed in the event's recent history, at the current time (indicated by the vertical dotted line), it appears to have largely stabilized, and the future situation looks to be only more stable.

Therefore, one can be reasonably comfortable acting on the present data at the maneuver commitment point since space weather is not likely to cause unexpected changes in drag predictions through the time of closest approach.

Q.7.2 Space Weather Modeling

Figure Q-8 communicates the expected effect on the probability of collision if the space weather situation has been mismodeled in the predicted object trajectories.




Figure Q-8 Space Weather Trade-Space Plot

Note that Figure Q-7 and the outcome of this particular event outline a situation in which Figure Q-8 probably would not be used operationally, but in a discussion of this type, it seems appropriate to include this information nonetheless.

To determine whether space weather has been mismodeled, one first recognizes that, in the drag equation, the atmospheric density estimate and the ballistic coefficient are multiplicatively coupled; this means that the effect of changing one of these parameters can be emulated by changing the other. To produce Figure Q-8, the ballistic coefficients for the primary and secondary objects were increased and decreased incrementally by up to half an order of magnitude in each direction (e.g., the “Primary Drag Scale Factor” on the x-axis spans from -0.5 to 0.5 in log space). This produces the same effect on predicted position if the atmospheric density were mismodeled by that amount. The colors on the intensity plot are the base-ten logarithm of the P_c .

What is important in the presentation, however, is not the actual P_c values but the morphology of the plot. If the midpoint (which represents the current actual solution) is on a “ridge,” as in Figure Q-8, any mismodeling of the atmospheric density will result in a lower P_c value. Thus, if the current P_c is below one’s threshold for action, one can embrace a decision to dismiss the

conjunction; i.e., if the current P_c value is below the threshold, and any mismodeling is expected only to lower the P_c from that value, then the event can be considered safe.

A presentation that is essentially one-color results in a similar conclusion because in such a case, the presentation is indicating that the event is simply not sensitive to atmospheric density mismodeling. If, however, the midpoint dot is on a “slope” (e.g., if it were on one of the light blue lines in Figure Q-8), then mismodeling could potentially either increase or decrease the P_c .

Thus, in such a case, the event is sensitive to atmospheric density mismodeling, and changes may be observed in future orbit determination updates in either direction that could affect the decision to maneuver. If the event is close to the red threshold, CARA would typically recommend maneuvering despite the space weather uncertainty for conservatism.

Q.8 Mitigation Maneuver Planning

To assist O/Os in planning mitigation actions, a “Maneuver Trade-Space” plot similar to Figure Q-9 is produced.




Figure Q-9 Maneuver Trade-Space Plot

A potential mitigation burn, taken from a variety of different burn sizes at a variety of different execution times, is placed into the primary object’s ephemeris, the altered ephemeris is propagated forward, new P_c s for both the main conjunction under consideration and any other conjunctions that arise from this altered trajectory are calculated, and the P_c that is an amalgamation of these individual P_c values is calculated. This calculation is run over and over, working through the entire range of burn intensities and burn times, and from the results, an

intensity plot like Figure Q-9 is constructed. The colors indicated on the second y-axis at right indicate the P_c ($1E-X$, where X is the value on the axis) that will be achieved if a burn of the magnitude chosen on the x axis and the phasing time chosen on the right axis is executed. Choosing a maneuver time and magnitude that result in a P_c of $1E-07$ (dark blue) would be considered to fully mitigate the conjunction risk.

A number of interesting conclusions can be drawn from a Maneuver Trade-Space plot. If it is the case that the O/O wishes to mitigate the risk down to the CARA-recommended level of $3E-06$, then a maneuver should be chosen that lies along the dashed white line above. The longer one waits to perform the maneuver, the larger the maneuver will need to be; the maneuver possibility shown in the bottom right of the graph is twice as large as the one shown at the top right, so maneuvering early allows a smaller maneuver to fully mitigate the conjunction. However, it is also often the case that waiting an additional half-day will allow more tracking data to be collected that will shrink the covariance and potentially push the P_c below the red threshold, meaning no maneuver would be necessary at all. This tension ultimately must be resolved by the O/O: does the O/O want to commit early and lock in a smaller maneuver, or does the O/O want to wait to act until the maneuver commitment point, hoping that the need for the maneuver will evaporate completely? The risk of waiting is that a large maneuver would be required to mitigate the conjunction, or that it may not be possible to fully mitigate the risk given the size maneuver that the spacecraft has the capability to perform.

Note that Figure Q-9 does not dictate any specific action to the O/O; instead, it provides general guidance for the sizes and timing of maneuvers that meet chosen mitigation requirements. Based on these data, the O/O selects a promising maneuver (or set of maneuvers), builds them into predicted ephemerides, and submits them to the conjunction assessment screening authority for screening. This latter action is necessary because it is the only way to give assurance of both the safety of the proposed maneuver (or set of maneuvers) relative to other on-orbit objects and the actual amount of risk reduction that the maneuver (or set of maneuvers) can be expected to yield.

In this example, the collision likelihood is above the CARA mitigation threshold ($1E-04$) at the maneuver commitment point used by the O/O for this spacecraft, and mitigation actions that appropriately reduce the risk are available and not large enough to negatively affect the satellite's main mission appreciably. Therefore, the O/O made the straightforward decision to perform a mitigation action.

Appendix R. List of Works Cited

U.S. Government References

NPR 8079.1, NASA Spacecraft Conjunction Analysis and Collision Avoidance for Space Environment Protection.

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 U.S. Government Orbital Debris Mitigation Standard Practices.

NASA Technical Publication NASA/TP–20220014814, “Astrodynamics Convention and Modeling Reference for Lunar, Cislunar, and Libration Point Orbits.” Version 1.1, 1 December 2022. David Folta, Natasha Bosanac, Ian Elliott, Laurie Mann, Rebecca Mesarch, and Jose Rosales. <https://ntrs.nasa.gov/citations/20220014814>.

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>.

National Geospatial-Intelligence Agency (NGA), Cis-Lunar SSA Technical Steering Group. “Proposed Message Set For All Altitude Regimes.” (DRAFT) November 2021. Butt, Y. et al. https://www.nasa.gov/sites/default/files/atoms/files/cislunar_ssa_proposed_message_set-nov2021_rev2_without_emailspdf.pdf

U.S. Government Orbital Debris Mitigation Standard Practices, November 2019 Update.
https://www.nasa.gov/wp-content/uploads/2018/01/usg_orbital_debris_mitigation_standard_practices_november_2019.pdf.

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

White House National Science & Technology Policy Council, “National Cislunar Science & Technology Strategy,” November 2022.

Combined Force Space Component Command. *Spaceflight Safety Handbook for Satellite Operators: 18 & 19 SDS Processes for On-Orbit Conjunction Assessment & Collision Avoidance*. Version 1.7, April 2023. Vandenberg Space Force Base, California, USA.
https://www.space-track.org/documents/SFS_Handbook_For_Operators_V1.7.pdf.

United States Space Command. *Launch Conjunction Assessment Handbook*. Version 3.1. November 2025. Vandenberg Air Force Base, California, USA. https://www.space-track.org/documents/Launch_Handbook_For_Operators_V3.1.pdf.

Authored and Other References

1. Abramowitz, M. and Stegun, I. A. 1970, *Handbook of Mathematical Functions*, New York, NY: Dover Publications, Inc.
2. 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.
3. Alfano, S. 2004. “Orbital Covariance Interpolation.” 14th AAS/AIAA Space Flight Mechanics Conference, Maui HI, February 2004.
4. 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.
5. 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.
6. Alfano, S. et al. 2022. “Operators’ Requirements for SSA Services.” *Journal of the Astronautical Sciences*. <https://doi.org/10.1007/s40295-022-00346-8>.
7. Baars, L., Hall D., and Casali, S. 2019. “Assessing GEO and LEO Repeating Conjunctions using High-Fidelity Brute Force Monte Carlo Simulations.” 2019 AAS Astrodynamics Specialist Conference (Paper #19-612), Portland ME, August 2019.
8. Baars, L. and Hall, D. 2022 “Processing Space Fence Radar Cross-Section Data to Produce Size and Mass Estimates” 2022 AAS Astrodynamics Specialist Conference (Paper #22-586), Charlotte, N.C., August 2022.
9. 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>.
10. Butt, Y. et al. 2021. See above: National Geospatial-Intelligence Agency (NGA), Cis-Lunar SSA Technical Steering Group. “Proposed Message Set For All Altitude Regimes.” (DRAFT). November 2021.
https://www.nasa.gov/sites/default/files/atoms/files/cislunar_ssa_proposed_message_set-nov2021_rev2_without_emailspdf.pdf.
11. Bassa, C. G., Hainaut, O. R., and Galadi-Enriquezz, D. 2022. “Analytical Simulations of the Effect of Satellite Constellations on Optical and Near-infrared Observations.” *Astronomy & Astrophysics* Vol. 657, P. A75, Jan 2022.
12. 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.
13. Carpenter 2019.
14. Carpenter 2020.

15. 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.
16. Consultative Committee for Space Data Systems (CCSDS) 508.0-B-1. Conjunction Data Message (CDM) Blue Book. Issue 1. Recommended Standard for Conjunction Data Messages. Washington, D.C.: CCSDS, June 2013.
<https://ccsds.org/publications/bluebooks/entry/3064/>.
17. CCSDS 502.0-B-3. Orbit Data Messages Blue Book. Issue 3. Recommendation for Space Data System Standards²². Washington, D.C.: CCSDS, April 2023.
<https://ccsds.org/publications/bluebooks/entry/3073/>.
18. Chan, F. C. 2008. *Spacecraft Collision Probability*. El Segundo, CA: The Aerospace Press.
19. 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.
20. Chan, F. C. 2019. “Collision Probability for Polygonal Cross Sections.” 2019 AAS Astrodynamics Specialist Conference (Paper #19-911), Portland ME, August 2019.
21. 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.
22. 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.
23. D’Agostino, R. B. and M. A. Stephens. 1986. *Goodness-of-Fit Techniques*. New York: Marcel Dekker, Inc.
24. 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.
25. Elkantassi, S. and Davison, A. C., 2022. “Space Oddity? A Statistical Formulation of Conjunction Assessment.” *Journal of Guidance Control and Dynamics*, Vol. 45 No. 12 (2022), pp. 2258-2274. <https://doi.org/10.2514/1.G006282>.
26. Elrod, C. 2019. “Computational Bayesian Methods Applied to Complex Problems in Bio and Astro Statistics.” Doctoral Dissertation, Baylor University, July 2019.
27. Flohrer, T., Lemmens, S., Bastida Virgili, B., Krag, H., Klinkrad, H., Parrilla, E., Sanchez, N., Oliveira, J., Pina, F. 2013. “DISCOS – Current Status and Future Developments.” Proc. 6th European Conference on Space Debris, (ESA SP-723), Darmstadt, Germany, April 2013.

²² There is a version 3 Orbit Data Messages (ODM) that will be published relatively soon (early 2023), but it doesn't change the Orbit Ephemeris Message (OEM) in a material way.

28. Folta, David, Natasha Bosanac, Ian Elliott, Laurie Mann, Rebecca Mesarch, and Jose Rosales. 2002. See above: NASA Technical Publication NASA/TP–20220014814, “Astrodynamics Convention and Modeling Reference for Lunar, Cislunar, and Libration Point Orbits.” 1 December 2022.
29. 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).
30. 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.
31. Garcia-Pelayo, R. and Hernando-Ayuso, J. 2016. “Series for the Collision Probability in Short-Encounter Model,” *Journal of Guidance Control and Dynamics*, Vol. 39, No. 8, pp. 1904-1912, Aug. 2016.
32. 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.
33. Grubbs, F.E. 1969. “Procedures for detecting Outlying Observations in Samples.” *Technometrics* 11 (1969), 1-21.
34. 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.
35. 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.
36. 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.
37. 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.
38. 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.
39. Hall, D.T. and Matney, M.M. 2000. “Collision Probabilities for Keplerian Orbits.” *Space Debris*, Vol. 2. (2000), pp.161-198.
40. Hall, D.T. 2021. “Expected Collision Rates for Tracked Satellites,” *Journal of Spacecraft and Rockets*, Vol. 58 No. 3 (2021), pp. 715-728.

41. Hall, D.T. 2022. “Semi-Empirical Astronomical Light Pollution Evaluation of Satellite Constellations.” *Journal of the Astronautical Sciences*, .69, pp.1893-1928, 2022.
<https://doi.org/10.1007/s40295-022-00358-4>.
42. Hall, D. and Baars, L. 2023, “Collision and Fragmentation Probabilities for Satellites with Size and Mass Estimates Based on Radar Observations.” *Journal of Spacecraft and Rockets*, Vol. 60 No. 4 (2023), <https://doi.org/10.2514/1.A35697>.
43. Hall, D. T, Baars, L. B., and Casali, S. J. 2023 “A Multistep Probability of Collision Computational Algorithm.” 2023 AAS/AIAA Astrodynamics Specialist Conference (Paper #23-398), Big Sky MT, August 2023.
44. 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.
45. Hejduk, M.D. 2008. “Space Catalogue Accuracy Modeling Simplifications.” 2008 AAS Astrodynamics Specialists Conference, Honolulu, HI, August 2008.
46. 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.
47. 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.
48. 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.
49. 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.
50. 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.
51. 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.
52. J. Hernando-Ayuso, G. Bau, and A. Martinez-Cacho. 2023 “Near-Linear Orbit Uncertainty Propagation Using the Generalized Equinoctial Orbital Elements,” *J. of Guidance, Control and Dynamics*, Vol. 46, No. 4, 2023, pp. 654-665. <https://doi.org/10.2514/1.G006864>

53. Hoots, F. R. and R. L. Roehrich. 1980. "Spacetrack Report No.3 – Models for Propagation of NORAD Element Sets," December 1980.
54. 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.
55. 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.
56. Kaplan, S. and B. Garrick. 1981. "On the Quantitative Definition of Risk." *Risk Analysis*, Vol. 1 No. 1 (1981), pp. 11-27.
57. Kessler, D. J. 1981. "Derivation of the Collision Probability between Orbiting Objects: The Lifetimes of Jupiter's Outer Moons." *Icarus*, Vol. 48 (1981), pp. 39-48.
58. Krisko, P. 2011. "Proper Implementation of the 1998 NASA Breakup Model." *NASA Orbital Debris Quarterly News*, Vol. 15, Issue 4, October 2011, pp. 4-5.
59. Lebedev, V., and Laikov, D. 1999. "A Quadrature Formula for the Sphere of the 131st Algebraic Order of Accuracy." *Doklady Mathematics*, Vol. 59, No. 3 (1999), pp. 477-481.
60. 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.
61. Lechtenberg, T. 2019a. "An Operational Algorithm for evaluating Satellite Collision Consequence." 2019 AAS Astrodynamics Specialist Conference (Paper #19-669), Portland ME, August 2019.
62. 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.
63. Li, Z., Ziebart, M., Bhattarai, S., Harrison, D., Grey, S. 2018. "Fast Solar Radiation Pressure Modelling with Ray Tracing and Multiple Reflections." *Advances in Space Research*, Vol. 61, No. 9, pp. 2352-2365.
64. 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.
65. Mclean, F., Lemmens, S., Funke, Q., Braun, V. 2017. "DISCOS 3: An Improved Model for ESA's Database and Information System Characterising Objects in Space." *Proc. 7th European Conference on Space Debris*, Darmstadt, Germany, June 2017.
66. Patera, R.P. 2001. "General Method for Calculation Satellite Collision Probability. 2001." *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 4 (July-August 2001), pp. 716-722.

67. Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. B. 1992. *Numerical Recipes in FORTRAN: The Art of Scientific Computing*, 2nd ed., Cambridge University Press, New York, NY, 1992.
68. 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.
69. Rosner, B. 1975. “On the Detection of Many Outliers.” *Technometrics* 17 (1975), 221-227.
70. Rosner, B. 1977. “Percentage Points for the RST Many Outlier Procedure.” *Technometrics* 19 (1977), 307-312.
71. Sabol, C., T. Sukut, K. Hill, K. Alfrend, 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.
72. 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.
73. 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.
74. 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.
75. Sivia, D.S. and Skilling, J. 2006. *Data Analysis: A Bayesian Tutorial*. 2nd ed., Oxford, UK, Oxford University Press.
76. Stokely, C.L., Foster, J.L., Stanbery, E.G., Benbrook, J.R., and Juarez, Q. 2006. “Haystack and HAX Radar Measurements of the Orbital Debris Environment.” National Aeronautics and Space Administration Report, Orbital Debris Program Office, (JSC-62815), Nov. 2006.
77. Tarzi, Zahi B., Brian T. Young, David S. Berry. 2022. "Deriving Event Thresholds and Collision Probability for Automated Conjunction Assessment at Mars and the Moon", AAS 22-042, 2022 American Astronomical Association (AAS)/American Institute of Aeronautics and Astronautics (AIAA) Astrodynamics Specialist Conference, Charlotte, NC, August, 2022. <https://www.nasa.gov/cara/madcap>.
78. Tobiska, W.K., D. Knipp, W. J. Burke, D. Bouwer, J. Bailey, D. Odstreil, 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.
79. Walker, C. et al. 2020, “Impact of Satellite Constellations on Optical Astronomy and Recommendations Toward Mitigations.” *SatCon-1 Workshop Report*, NSF NOIRLab, Aug 2020.

80. Walker, C. et al. 2021. “Executive Summary.” *Report of the SatCon-2 Workshop*, NSF NOIRLab, Jul 2021.
81. Vallado 1981.
82. Vallado, D.A. 2001. *Fundamentals of Astrodynamics and Applications*. 2nd ed. El Segundo, CA, Microcosm Press.
83. 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.
84. Zaidi, W.H. and Hejduk, M.D. 2016. “EOS System Covariance Realism.” 2016 AAS/AIAA Astrodynamics Specialist Conference, Long Beach CA, September 2016.

National Aeronautics and Space Administration

Mary W. Jackson Headquarters
300 E Street SW
Washington, DC 20546

[nasa.gov](https://www.nasa.gov)