AAS 13-435

OPERATING CHARACTERISTIC APPROACH TO EFFECTIVE SATELLITE CONJUNCTION FILTERING

Salvatore Alfano[1], David Finkleman[2]

This paper extends concepts of signal detection theory to examine the performance of conjunction screening techniques. The most effective way to identify satellites likely to collide is to employ filters to isolate orbiting pairs that cannot come close enough over a prescribed time period to be considered hazardous. Such pairings can then be eliminated from further computation to quicken the overall processing time. The three most common filters are the apogee/perigee, orbit path, and time. The apogee/perigee filter eliminates pairings that lack overlap in the respective ranges of radius values regardless of planar orientation. The orbit path filter (also known as the geometric pre-filter) takes planar orientation into account to eliminate pairings where the distance (geometry) between their orbits remains above some user-defined threshold, irrespective of the actual locations of the satellites along their paths. The time filter identifies pairs that have survived other screening processes but are unlikely to be in close proximity during the analysis interval. The workings of each filter are summarized and then tested with various threshold and pad settings. Every filtering process is vulnerable to Type I and Type II errors, admitting infeasible conjunctions and missing feasible conjunctions. Documenting Type I and Type II errors using an Operating Characteristic approach guides selection of the best operating point for the filters. This work provides a formalism for selecting filter parameters.

INTRODUCTION

Dealing with an ever more crowded space environment requires identifying potentially dangerous orbital conjunctions and executing a suitable course of action. The problem of on-orbit collisions or near misses is receiving increased attention in light of the collision between an Iridium satellite and COSMOS 2251. More recently, the International Space Station conducted a debris avoidance maneuver on 31 October 2012 to avoid a piece of Iridium 33 debris. This text applies principles of signal detection to identify potential orbital conjunctions with quantifiable confidence, guiding rapid and efficient computation. [1]

The identification of potentially dangerous conjunctions requires finding pairs of satellites that are likely to be very close to each other and the time at which the close approach occurs. Once high-risk conjunctions are identified, the probability of collision can be determined if the uncertainties in each orbit and the relative motion of the satellites can be determined. The linear separation between the mean orbits of two satellites is the most-used initial discriminant. Mean orbits are not really the satellite trajectories on the time and distance scales of conjunctions and

[1] Senior Research Astrodynamicist, Analytical Graphics Inc., Center for Space Standards and Innovation, 7150 Campus Dr., Suite 260, Colorado Springs, Colorado, 80920-6522.

[2] Senior Scientist, Analytical Graphics Inc., Center for Space Standards and Innovation, 7150 Campus Dr., Suite 260, Colorado Springs, Colorado, 80920-6522 and Academician of the International Academy of Astronautics

collisions. The satellite trajectories differ from the mean orbits because there are deviations due to physical perturbations.

COMMON FILTERING TECHNIQUES

For a problem containing only two objects, orbital conjunctions are identified by computing the distance between the two objects at all points in time during the analysis period and determining if the distance ever falls below a selected threshold. This paper deals with selecting those thresholds. Since applying this methodology to the problem of a single object versus the entire space catalog of nearly 20,000 objects (or worse yet to the problem of all catalog objects *vs.* all other catalog objects) is a computational challenge, this also accelerates executing time-critical and time-fragile operations.

As explained by Woodburn: "Hoots et al.[2] designed a series of three filters through which candidate objects have to pass before a final determination of the close approach distance is made. Two of the filters are purely geometrical and one uses the known properties of simplified two body orbital motion of the two objects. These filters eliminate the majority of the objects in the catalogue from intense scrutiny thereby reducing the number of computations needed. The trajectories of the remaining candidate objects are sampled to determine the actual close approach periods. The three filters are the apogee/perigee filter, the orbit path filter, and the time filter. These filters have the advantage of being easy to understand, but in practice they can be inadequate when implemented as originally described.[3] Additionally, the filters were initially designed for use with a space catalog comprised solely of two line element (TLE) sets. Now that a special perturbations version of the space catalog is being generated, any dependence on TLE specific information must be eliminated.

"Recognizing the practical limitations of Hoots' filtering technique, Healy[4] took advantage of parallel processing to expedite computation. The computations in Healy's method are based on sampling the relative distance between each pair of orbiting objects and using a simplified model of the relative motion to identify potential conjunctions efficiently during a time step. The candidate conjunctions are then subjected to a more detailed analysis using the full fidelity orbit model. Rodriguez[5] et al. refined Healy's approach in addition to employing a form of the apogee/perigee filter described by Hoots et al. Rodriguez points out the complexity of the orbit path and time filters and implies a lack of robustness." All of these approaches must suffer Type I and Type II errors, which can be mitigated somewhat by "padding" orbits with times and distances that might encompass orbit uncertainties.

We will examine the rationale of several filters and demonstrate error cases for each. Requirements for the use of computational *pads* will then be examined. The reliability of the filtering process will be evaluated using "all-on-all" examples where results will be generated with and without the use of conjunction filters. Computational time savings will not be addressed due the tremendous variability associated with platforms, operating systems, and hardware used for data storage.

All results in this paper were generated using the publicly available catalog of TLEs. Errors were initially assessed by finding all unique conjunction pairings for 14,546 objects in the public TLE catalog with minimum ranges of less than 5, 10, 30, and 50 km for the three day period beginning at 22 Feb 2012 19:00:00.000. Although TLEs were used exclusively in this study, the analysis method that follows will work with any trajectory source. The approach that follows can be used with any type of trajectory source data involving natural satellite motion.

CONJUNCTION FILTERS

As described by Woodburn et al, "Orbital conjunctions are identified through the examination of pairs of trajectories of orbiting objects. The goal of the process is to find *all*[1] conjunctions between a set of objects of interest, referred to as primary objects, and the set of all cataloged orbiting objects, referred to as secondary objects. Note that each entry in the set of primary objects is also a member of the set of secondary objects and that the set of primary objects can contain all of the secondary objects to create the so called "all-on-all" conjunction problem. Conjunction filters provide an efficient mechanism for finding conjunctions by providing quick identification of primary/secondary pairings which cannot come close enough together to yield a conjunction.

"In general, conjunction filters utilize approximations based on known characteristics of orbital motion to maximize the efficiency of the computation. Each filter defines a proxy for the distance between two orbiting objects, a candidate conjunction pair, and then either eliminates the pair from any further consideration or limits the time periods which require further analysis. We must be careful, however, that the accuracy of the results is not compromised by these approximations. One simple method of accounting for approximations is to use distance pads which increase the size of the conjunction threshold distance during the filtering process to cover the effects of the filter approximations. The effect of padding is to increase the number of candidate pairings which pass through each filter as possibly having conjunctions in order to reduce the likelihood that a candidate pair will be improperly eliminated which could lead to a missed conjunction. *(Trading more false detections for the sake of omitting fewer true detections.)* The difficulty associated with using pads lies in selecting pad values that preserve the accuracy of the computation while still providing computational benefit. The length of the analysis interval is an important consideration when using filters in the conjunction detection process. Since the filters assume a simplified motion model to expedite computation, the errors imparted into the filtering process can increase with the length of the analysis interval unless mitigating measures are taken. This is of most concern when padding is being used as the length of analysis interval tends to affect the size of pads required to obtain accurate results."

These filters are subject to false positive identifications (Type I errors) as well as false negative identifications (Type II errors). A Type I error occurs when the filter determines the satellite pair should be assessed further but no conjunction is found. This is to be expected because the screening is based on aggregated orbital characteristics and not precise satellite positions over the time interval. The ratio of non-conjuncting pairs forwarded by the filter to the total number of non-conjuncting pairs assessed is a figure of merit. Specificity, or false positive rate, is defined as one minus this ratio. A Type II error occurs when the filter determines the satellite pair should not be assessed further but a conjunction exists. This reveals that the predicted satellite positions will come closer than the minimum distance between orbits determined by the filter. The sensitivity, or true positive rate, is the ratio of actual conjuncting pairs forwarded by the filter to the total number of conjuncting pairs. A Type II error is one minus this ratio.

Type II errors result from the simplifying assumptions that went into the development of the various filters. By adding a pad or buffer to the desired distance threshold one can reduce such errors and avoid missing possible conjunctions. Reducing or eliminating Type II errors may require large pads which will, in turn, increase the computational burden. The pads should

[1] It is statistically impossible to identify "all." That fact is the substance of this paper.

account for the precession of the nodes and apsidal rotation as well as the secular and short-periodic variations for all the orbital elements (semi-major axis included). Such padding should strike a balance between timeliness and accuracy as determined by the user.

It is important for the user to configure the pad settings such that operational requirements are met in a timely fashion with an acceptable limit of Type II errors. This requires both step size control and distance bounds based on natural motion. Pad settings should be chosen in a manner that allows the end user to know whether the tool was exercised with a focus on accuracy or speed. If accuracy alone is important, one must be certain that the filters will not prematurely eliminate a pair of satellites from consideration that might be found to have a conjunction; Type II errors are shunned[1]. Such assurance comes with increased processing time as the computations would take much longer due to very conservative (large) pad settings. To maximize positive detections, no Type II errors, the simplest approach is to turn the filters off and assess every pair, accepting the increase in downstream computations. The padding characteristics presented in the following test cases assume natural relative motion and should not be used if the ephemeris of either satellite contains maneuvers unless those maneuvers are included in process models.

The Apogee/Perigee Filter

"The goal of the apogee/perigee filter is to eliminate pairings which cannot produce conjunctions due to a lack of overlap in the range of radius values experienced by the two trajectories, Figure 1. A simple example of a pairing which would be eliminated by this filter is a GEO satellite vs a LEO satellite.[3]

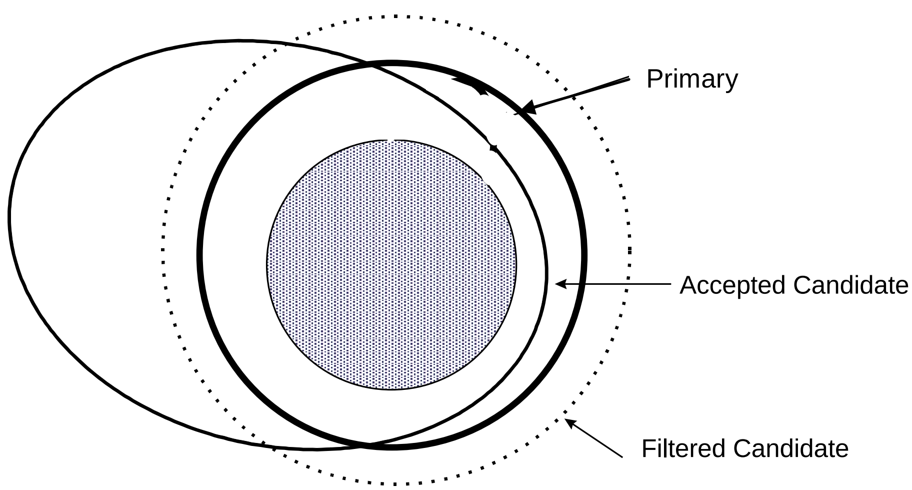

Figure 1. Apogee/Perigee Filter Principle[9]

"The original description of the apogee/perigee filter given by Hoots et al. recommends computing the apogee and perigee radii for the primary and secondary objects at the start of the analysis interval. The reasoning was that the perigee radius would remain constant and the apogee radius would decrease over the interval. The method for determining the apogee and

[1] It is impossible to eliminate all false detections while still enjoying any detections at all. Zero false alarms is absolutely unachievable.

perigee values at the start time was based on determining orbital elements at the midpoint of the analysis interval and using approximate rates of those elements to project values at the start time. The approximate rates of the mean motion and eccentricity were computed using the time derivative of the mean motion, which is part of a TLE but is not readily available for other forms of orbit data. This makes an exact implementation of the filter as previously described difficult for the case of varying ephemeris sources. While not explicitly stated, we believe that apogee and perigee radii were computed based on mean orbital elements.

"There is one extremely important feature of the apogee/perigee filter that is critical to understanding its potential benefit: the computations associated with the apogee/perigee filter need only be performed once for each object, not once for each pair of objects. The application of the computed apogee and perigee values to pairs of objects only requires simple comparisons. This is not important for the case of one primary, but is a huge distinction for the case of "all-on-all" conjunction analysis since the computational load increases by order N instead of N^2 as is the case with the other filters."

To characterize the errors associated with pad sizes for this filter, we conducted a large scale conjunction screening analysis using STK 9.2.3. This involved finding all unique conjunction pairings for 14,546 objects in the public TLE catalog with minimum ranges of less than 5, 10, 30, and 50 km for the three day period beginning at 22 Feb 2012 19:00:00.000. 50km is considered a reasonable upper limit for distance screening based on probability-based actions[6]. For this analysis, three days prior to a predicted conjunction is deemed sufficient to assess a conjunction threat, consider the advantages and detriments of various avoidance actions, plan the control sequence for a specific evasive maneuver, uplink and confirm said sequence to the satellite, execute the maneuver, and confirm its successful completion.[1] Over this three day interval 30,496 actual pairings were predicted to come within 5km, 121,913within 10km, 1,078,404 within 30km, and 2,948,777 within 50km.

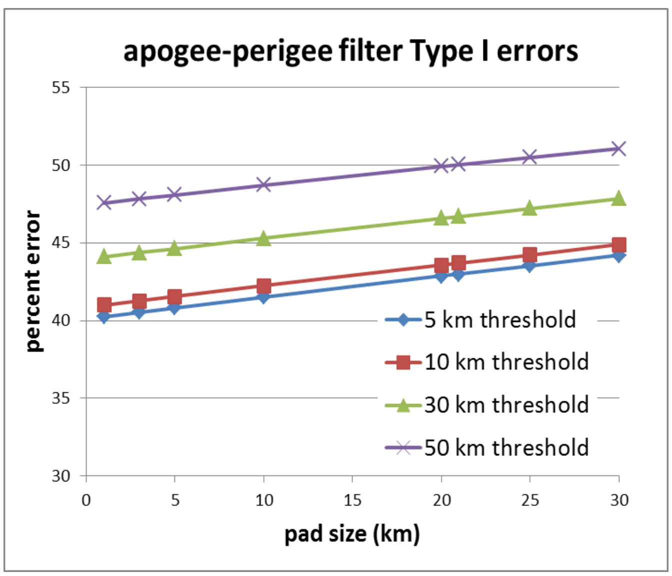

Figure 2a. Apogee/perigee filter Type I errors

[1] This conclusion is based on conjunction warning experience of the Space Data Center. Many phenomena that can change orbits by amounts comparable to conjunction tolerance thresholds occur on time scales of a few days or less.

As one would expect, Figure 2a shows that the Type I (false positive) errors increase as the pad size increases regardless of the screening threshold. The largest error (over 51%) occurs for a 50km threshold with a 150km pad, meaning that about 49% of all non-conjuncting pairs were properly eliminated by this filter while 51% were erroneously determined to require further processing.

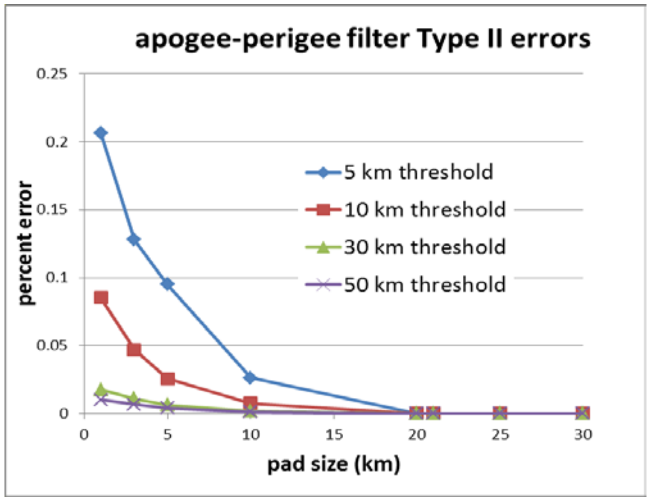

Figure 2b. Apogee/perigee filter Type II errors

Figure 2b shows that a pad size of only 20 km was sufficient to eliminate all Type II (missed detections) errors for all the distance thresholds considered. The largest error (0.2%) occurs for a 5km threshold with a 1km pad, meaning that for every 1000 conjuncting pairs two were erroneously eliminated by this filter. Although it appears that using a pad of 20 km would be sufficient to capture all cases successfully, it is necessary to point out that the results of Figure 2 are provisional for the 22 Feb 2012 data. A more comprehensive investigation with many different epochs is planned to guide fairly universal selection of detection thresholds.

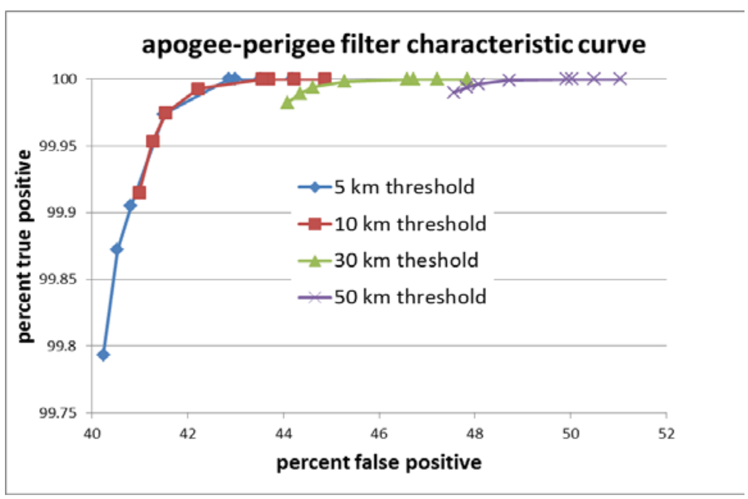

Figure 2c. Apogee/perigee filter true versus false positive rates

Figure 2c reveals that an increase in false positive rates will be accompanied by an increase in true positive rates until the latter reaches 100%. Figure 2c is also known as an Operating Characteristic curve[7], comparing two operating characteristics as the criterion changes. Such curves provide a simple yet direct way to examine costs and benefits for decision making. Ideally we desire 100% true positive rates and 0% false positive rates. Therefore, the closer a given curve is to the upper left-hand corner, the more effective the filter. Based on the above, 5km and 10km thresholds provide a better tradeoff than 30km or 50km when dealing with apogee/perigee filters. This, of course, assumes one is willing to accept some Type II errors. An alternative operating point can be found where the tangent to the curve has unit slope. At that point false positives and missed detections are balanced. Moving thresholds in either direction allows more of one at the expense of much more of the other. It is the law of diminishing returns. Insisting on absolutely capturing of as many positive detections as possible makes the process vulnerable to an exceptional numbers of false alarms.

The Orbit Path Filter

As described by Woodburn, "The goal of the orbit path filter, called the geometric pre-filter in Hoots et al., is to eliminate pairings which cannot produce conjunctions because the distance between their orbits remains above the conjunction threshold, irrespective of the actual locations of the objects along the paths. For the case of two circular orbits, the solution of the minimum distance between the paths is simple and is computed at the relative line of nodes, Figure 3. This solution was used by Hoots et al, as the starting point for a Newton iteration scheme to solve the more general problem where the orbit paths are elliptical." Woodburn and Dichmann[8] presented an improved algorithm for computing the distance between the elliptical paths that are not coplanar. Alfano developed a focus-centered, elliptical-ring torus, path filter that is valid for coplanar and non-coplanar orbits.

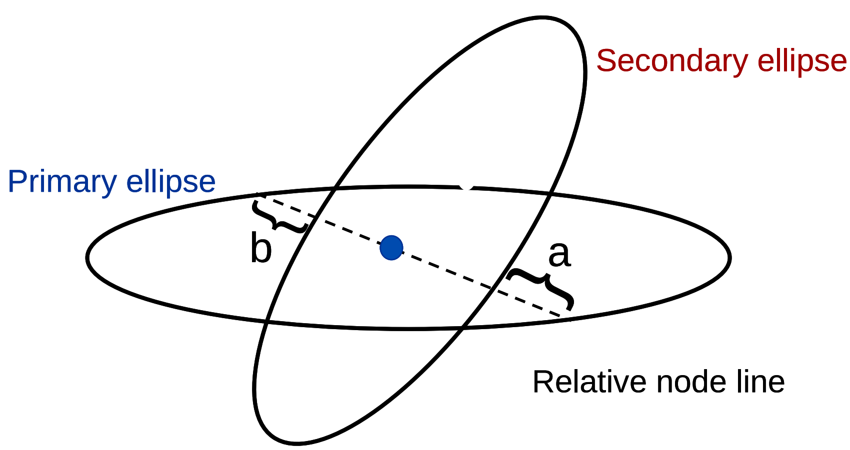

Figure 3. Orbit Path Filter Principles[9]

Using the exact same data and approach from the previous section, we computed the filter values over the three day interval to produce Figures 4a, 4b, and 4c.

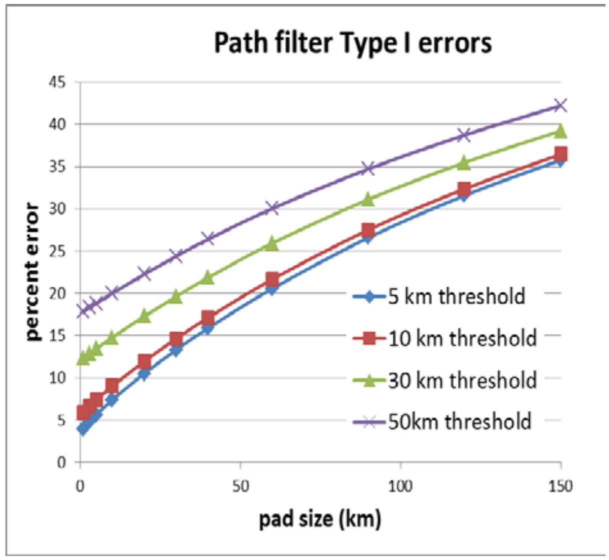

Figure 4a. Orbit-path filter Type I errors

All errors were below 43%. This resulted in fewer pairs being forwarded for further evaluation and is an improvement over the Apogee/perigee filter.

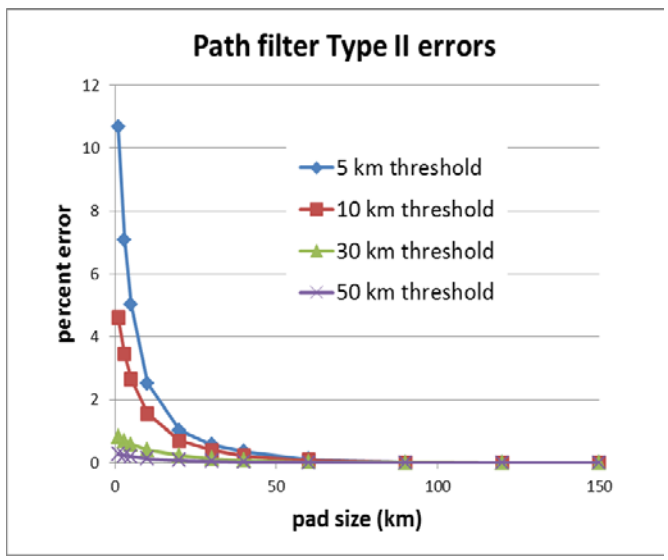

Figure 4b. Orbit-path filter Type II errors

Figure 4b shows that the errors are quite large for small pad sizes. A pad size of 90 km was sufficient to eliminate almost all Type II errors for the thresholds considered. Although impossible to see due to the figure's resolution, Type II errors did not go to zero for the 50 km threshold. A 90 km pad produced 30 errors from almost 3 million conjunctions (.001% error). A 120 km pad produced 4 errors and a 150 km pad produced 2 errors. In comparison to the previous filter, this filter performed poorly when addressing Type II errors. Simply applying a pad for this filter, even when accounting for the differences between mean and osculating orbits, would not sufficiently and properly eliminate candidate pairs.

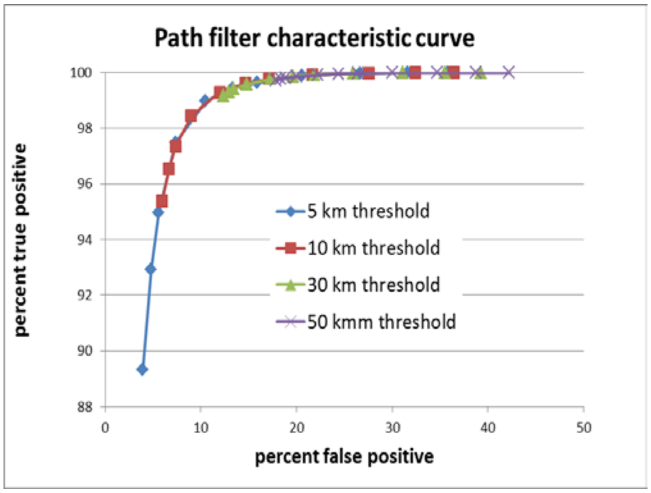

Figure 4c. Orbit-path filter true versus false positive rates

Figure 4c shows the trade space between false positive rates and true positive rates. Again, although impossible to see due to the figure's resolution, the true positive rate did not reach 100% for the 50 km threshold.

The Time Filter

As explained by Woodburn, "The goal of the time filter is to identify time intervals when each object in a pairing is close enough to the elliptical representation of the other objects trajectory to have a conjunction. The actual metric computed is the distance of primary object from the orbit plane of the secondary object, Figure 5. Since this distance is always guaranteed to be less than the actual distance between objects, the metric is a conservative proxy for the separation distance. Pairings are eliminated from further consideration when there is no overlap between the intervals associated with the primary object and those associated with the secondary object. If a pairing does pass through this filter, the amount of time over which more detailed event detection must be is typically limited to a fairly small number of relatively short intervals. This filter works on the premise that both objects have to be in the right place at the right time for a conjunction to occur.

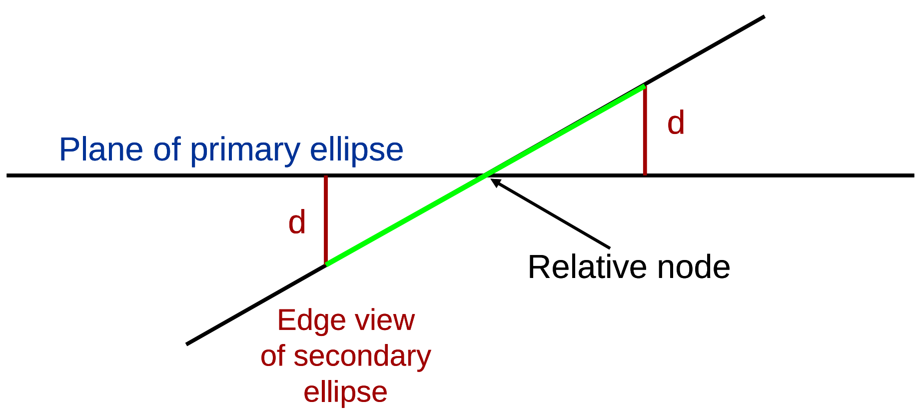

Figure 5. Distance from Primary to orbit plane of Secondary[9]

"Intervals are generated for each object independently and are anchored to the relative node between the two orbits. The condition for a possible conjunction is that an interval around the ascending node for the primary object overlaps with an interval around the descending node for the secondary object or vice versa, Figures 6 and 7. Like the orbit-path filter, the time filter is limited to cases where the trajectories are not coplanar and neither object is maneuvering.

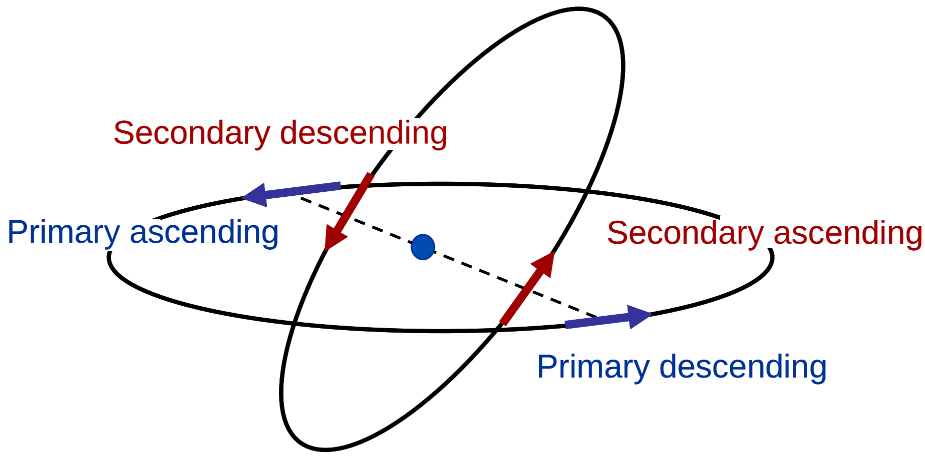

Figure 6. Relative node crossings[9]

Figure 7 notionally displays when each satellite will be within a specified distance of the path (not the position on the path). The cases where there is no overlap between the reference (primary) object and the candidate (secondary) object intervals are the cases that can be eliminated from further processing

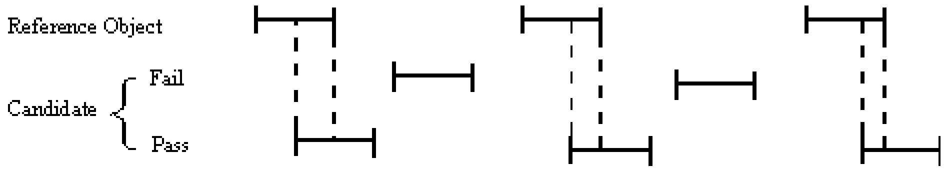

Figure 7. Time pre-filter (intervals near relative node) [9]

As before, the exact same data and approach from the previous section is used to produce Figures 8a and 8b.

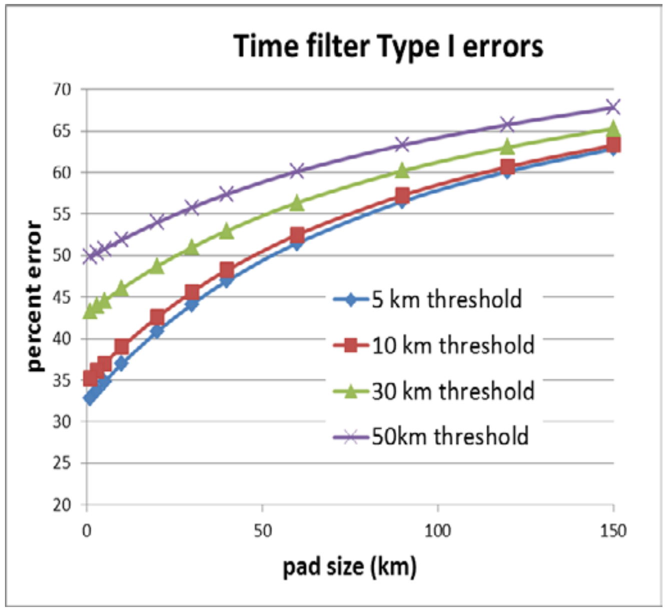

Figure 8a. Time filter Type I errors

Figure 8a reveals that this filter identified many prospective conjunctions that never came to fruition. The upside to these errors is seen in the following chart.

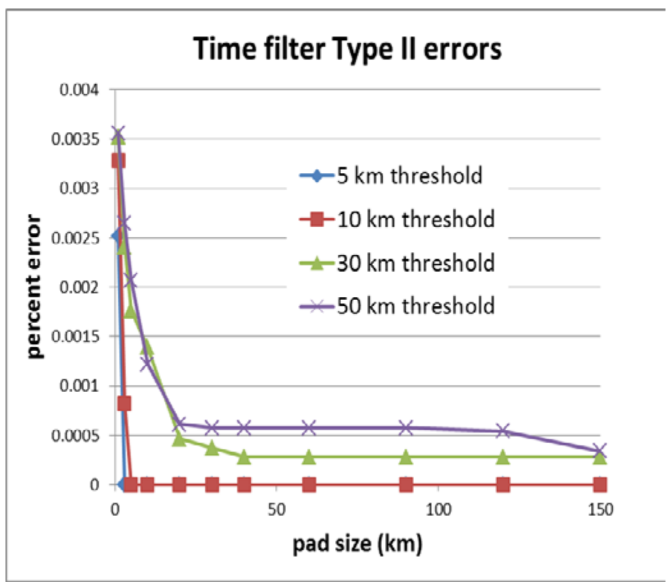

Figure 8b. Time filter Type II errors

Figure 8b shows that the time filter did quite well keeping Type II errors low even with small pads, but did not completely eliminate them as the apogee/perigee filter did.

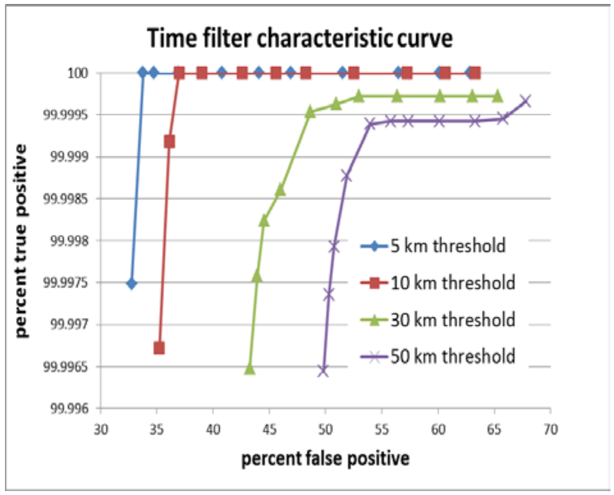

Figure 8c. Time filter true versus false positive rates

Figure 8c displays very favorable Operating Characteristic curves, even besting the apogee/perigee filter for false positive rates below 45%. Unlike the apogee/perigee filter, however, the time filter never reaches a 100% true positive rate for thresholds of 30 km and 50 km even with a pad setting of 150 km.

Conjunction Range Threshold Detection

When using a minimum range threshold for identifying conjunctions, one must determine the actual instances where the range between satellites is less than the operator specified value. It is necessary to determine the time periods where the distance is below the conjunction threshold and/or to determine the time(s) of closest approach within those periods. Even when employing probability-based conjunction assessment methods, a user can prescreen on distance. Alfano developed an analytical approximation that relates maximum probability to a conjunction threshold, thus ensuring that the resulting pairs screened by distance include all pairs that meet a probability threshold. It is therefore beneficial to identify these minimum-range time periods regardless of the risk metric.

EXAMPLE OF USE

This section takes the reader through an exercise in pad selection. Once the errors for a single catalog have been characterized, those results can be used to determine initial pad filter settings. The chosen settings should then be tested against other catalogs to determine adequacy for operational use. The pads can be tailored for a specific conjunction distance or chosen to accommodate a range of distances. The user may wish to test using an entire catalog or a subset addressing some specific orbit regime. To minimize computation time, it is prudent to execute the filters in their order of ascending complexity. This means the apogee/perigee filter should be examined first, followed by the path filter, then by the time filter.

It is instructional to examine the relative effectiveness of all the filters. The following figure shows the operating characteristics of all three filters (Figures 2c, 4c, and 8c) on a single chart.

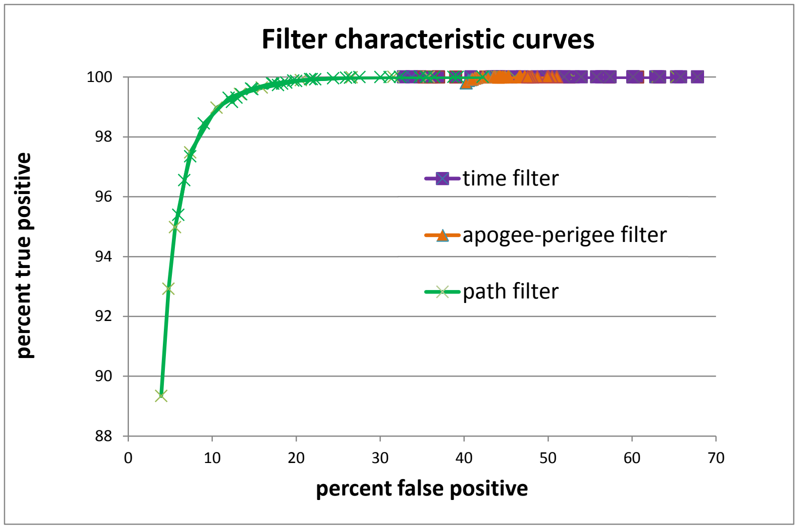

Figure 9. True versus false positive rates for all filters

As is readily seen, the path filter has a significantly lower true positive rate than the other filters. Because of its poorer relative performance, we have chosen to exclude this filter in the example that follows. Examining Figure 2b, Type II errors go to zero for all curves when the pad is 20 km or greater. We select an apogee/perigee filter pad setting of 20km accordingly. Examining Figure 5b, there appears to be a sharp bend in the curves at 20 km. We therefore choose this value for our time filter pad setting. These settings will be used for conjunction screening distances from 5 to 50 km against the full catalogs; they are by no means definitive. Each user should select filter pad settings based on desired operational performance and acceptable error. Having chosen the pad settings, their errors can be assessed against historical catalogs under the assumption that past performance is indicative of future performance. If the errors are unacceptable, the pads should be adjusted and the performance reassessed. Although not part of this analysis, one may also wish to address overall computation time as part of this assessment.

Figure 10 shows the evolution of the public satellite catalog up to October 2011.

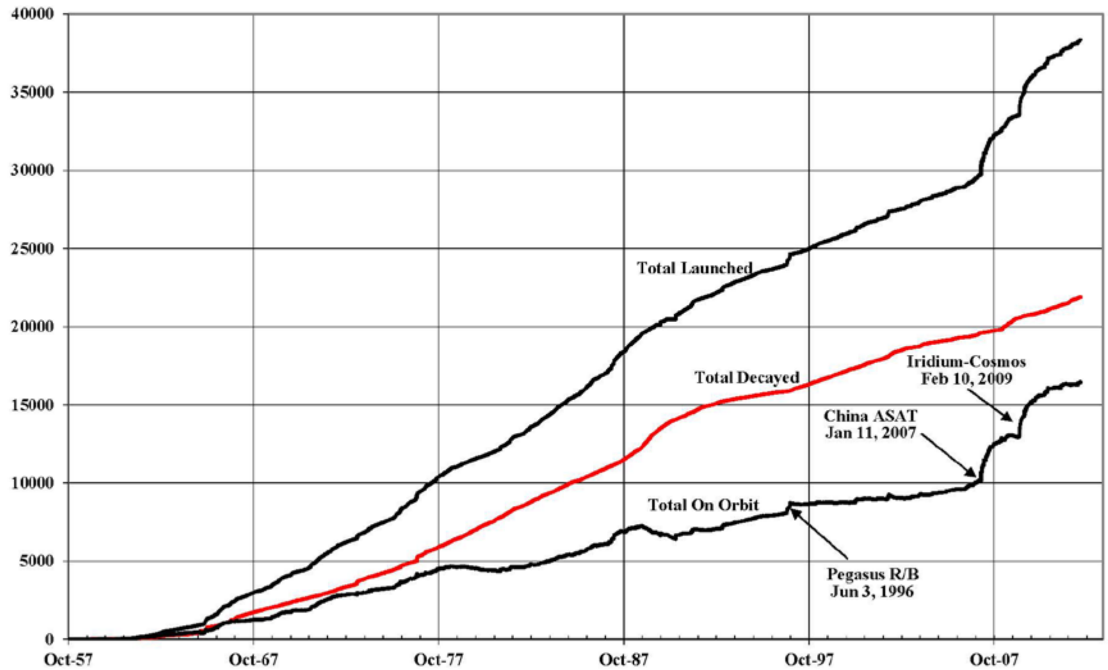

Figure 10. Satellite Population Growth [10]

Two significant debris-generating events stand out: the Chinese ASAT test and the Iridium-Cosmos collision. We chose to test our pad settings using some random catalogs spanning the years 2000 to 2012. To this, we added the catalog for the day prior to the Chinese ASAT test and also the day prior to the Iridium-Cosmos collision. Because this example is intended only for illustration, our tests use the same pad settings for all conjunction screening distances. Again, what we have done is meant only as an example demonstrating how to create and evaluate an Operating Characteristic curve and is not to be considered definitive.

Figure 11a shows the percentage of false positive identifications by date in groupings of the conjunction distance thresholds of 5, 10, 30, and 50km.

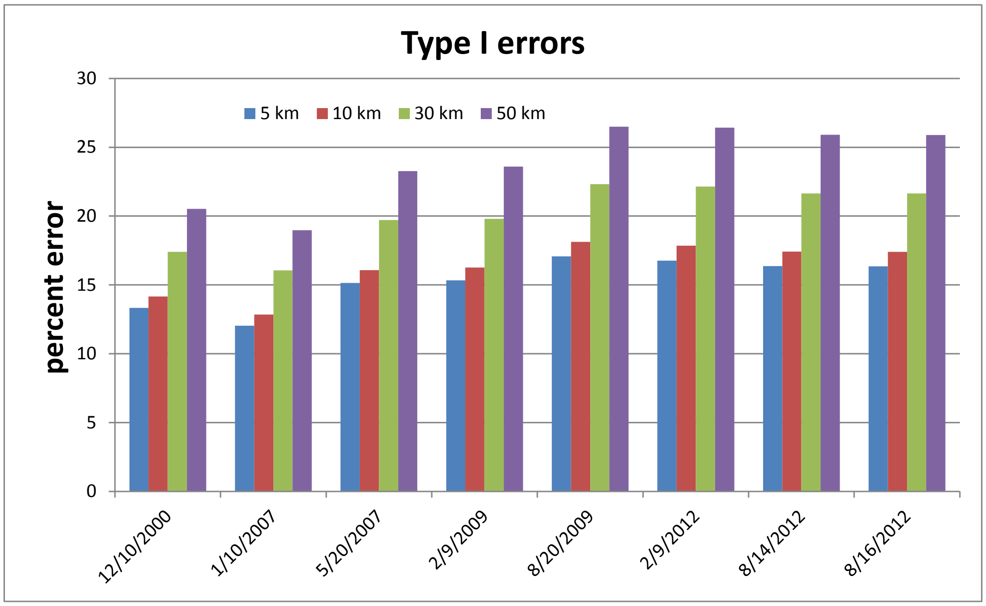

Figure 11a. Type I errors

It appears the percent of Type I errors grows after each major debris-generating event. This suggests that, at the very least, the pads should be re-adjusted after such events if one desires to maintain a consistent level of Type I errors. Also, the larger the conjunction screening distance the larger the error. These trends are not of concern provided the errors are acceptable to the user for the intended operation. If deemed inadequate, different pad setting should be chosen and this assessment re-accomplished.

Figure 11b shows the same groupings for false negative identifications.

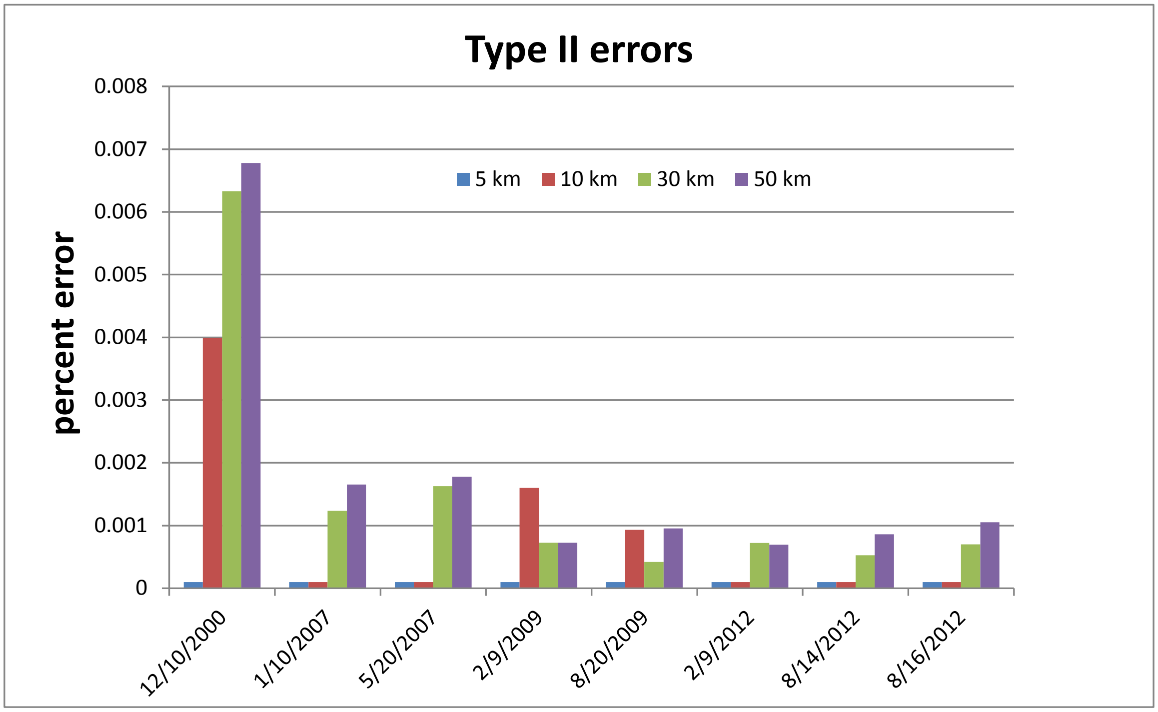

Figure 11b. Type II errors

As expected, the growth of Type I errors over time corresponded with a decrease in Type II errors. For a 5 km conjunction screening distance, the error was zero for all 8 catalogs tested. The largest raw number of missed conjunctions for all test cases was 41 out of 604,809 on December 10, 2000, with a 50 km conjunction screening distance. If these errors are acceptable for the intended operation, then the pad settings are suitable.

This is important knowledge and aids in understanding. Sequential filters can improve detection performance and, although not demonstrated here, the order in which filters are applied can matter from a computational loading perspective. Different orbit classes respond differently. Therefore, an effective conjunction assessment process demands exceptional operational judgment and experience based on concrete mathematical principles and a documented logical thread that all stakeholders can understand.

These analyses required extensive computations. In principle, similar analyses should be conducted regularly, and screening parameters should be adjusted to maintain a constant rate of Type I errors. The radar analogy of this process is called “constant false alarm rate.”[11] The a priori computational burden is offset by computational savings in real time operations and greater confidence in the outcomes.

CONCLUSION

The most computationally and operationally effective way to identify satellites likely to collide is to employ filters to identify orbiting pairs that cannot come close enough over a prescribed time period to be considered hazardous. Such pairings can then be eliminated from further computation to quicken the overall processing time. In this work we described and tested the three most common filters: apogee/perigee, orbit path, and time.

Every filtering process is vulnerable to Type I and Type II errors, admitting infeasible conjunctions and missing feasible conjunctions. We have extended a concept of signal detection theory to predicting the performance of conjunction screening techniques. Documenting Type I and Type II errors using an Operating Characteristic approach guides selecting the best operating point for the filters. Such studies are common in radar detection, presenting the likelihood of detection to the likelihood of false alarm as one varies detection thresholds. This work established and demonstrated such an approach in determining conjunction filter parameters based on characteristic tradeoffs. It was shown that the orbit-path filter suffers significant Type I errors even with extremely large "pads" and has poorer performance when compared to the other filters.

The approach is valid for all types of filters individually or convolved in any order. This work can also be used to assess "sieve" filters that eliminate trajectories through sequential screening for threshold crossings along track, radially, and co-normal or equivalent orthogonal coordinate systems. This work is exceptionally computationally intensive a priori. Once conducted, it need not be repeated very often, and operational decision making should not be impeded.

ACKNOWLEDGMENTS

We wish to thank Dr. James Woodburn (AGI Chief Orbital Scientist), Dr. Vincent Coppola (AGI Senior Astrodynamics Specialist), and Frank Stoner (AGI Astrodynamics Engineer) for allowing us to use their filter descriptions verbatim from reference AAS 09-372.

REFERENCES

[1] Finkleman, D., Berry, D., "Requirements and Guidance for Conjunction Assessment," AAS 11-434, AAS/AIAA Astrodynamics Specialists Conference, Girdwood, AK, Aug 2011

[2] Hoots, F.R., Crawford, L.L. and Roehrich, R.L. "An Analytic Method to Determine Future Close Approaches Between Satellites", *Celestial Mechanics*, Vol. 33, 1984, pp. 143-158.

[3] Woodburn, J., Coppola, V, Stoner, F. "A Description Of Filters For Minimizing The Time Required For Orbital Conjunction Computations," AAS 09-372, AAS/AIAA Astrodynamics Specialist Conference, Pittsburgh, Pennsylvania August 9-13 2009.

[4] Healy, L.M., "Close Conjunction Detection on Parallel Computer," *Journal of Guidance, Control and Dynamics*, Vol 18, No. 4, July-August 1995.

[5] Rodriguez, J.R.A., Fadrique, F.M.M, Klinkrad, H., "Collision Risk Assessment with a Smart Sieve Method", Proceedings of the Joint ESA-NASA Space-Flight Safety Conference, ESTEC, Noordwijk, Netherlands, June 2002.

[6] Alfano, S. "Determining A Probability-Based Distance Threshold For Conjunction Screening" AAS 13-352, AAS/AIAA Space Flight Mechanics Meeting, Kauai, Hawaii, February 10-14, 2013.

[7] Operating Characteristic reference – could use Finkleman Girdwood paper, although that is not a fundamental reference.

[8] Woodburn, J. and Dichmann D., "Determination of Close Approaches for Constellations of Satellites", Paper No. C-5, IAF International Workshop on Mission Design and Implementation of Satellite Constellations, Toulouse, France, November 1997.

[9] http://www.agi.com/resources/help/online/stk/source/extfile/cat/cat02-04.htm#catprefilters January 10, 2013.

[10] Vallado, D., 2013, *Fundamentals of Astrodynamics and Applications, 4th ed.*, Space Technology Library, Microcosm Press, Hawthorne, California, p. 850.

[11] Eaves, J.L., Reedy, E.K, 1987, *Principles of Modern Radar*, Chapman & Hall, New York, p. 368.