AAS 09-372

A DESCRIPTION OF FILTERS FOR MINIMIZING THE TIME REQUIRED FOR ORBITAL CONJUNCTION COMPUTATIONS

James Woodburn*, Vincent Coppola† and Frank Stoner‡

Classical filters used in the identification of orbital conjunctions are described and examined for potential failure cases. Alternative implementations of the filters are described which maintain the spirit of the original concepts but improve robustness. The computational advantage provided by each filter when applied to the one versus all and the all versus all orbital conjunction problems are presented. All of the classic filters are shown to be applicable to conjunction detection based on tabulated ephemerides in addition to two line element sets.

INTRODUCTION

The problem of on-orbit collisions or near collisions is receiving increased attention in light of the recent collision between an Iridium satellite and COSMOS 2251. More recently, the crew of the International Space Station was evacuated to the Soyuz module until a chunk of debris had safely passed. Dealing with the reality of an ever more crowded space environment requires the identification of potentially dangerous orbital conjunctions followed by the selection of an appropriate course of action. This text serves to describe the process of identifying all potential orbital conjunctions, or more specifically, the techniques used to allow the computations to be performed reliably and within a reasonable period of time.

The identification of potentially dangerous conjunctions is most commonly done by determining periods of time when two objects have an unacceptable risk of collision. For this analysis, we will use the distance between the objects as our proxy for risk of collision. We are interested in when two orbiting objects come closer than a minimum acceptable distance from each other. While other measures, such as probability of collision, may be the final desired metric of risk, the distance between two object provides an efficient proxy for identification of events of interest. Additional analyses may then be performed on identified events to provide supplemental information useful for making decisions regarding potential courses of action.

For a problem containing only two objects, the problem of identifying orbital conjunctions is simply solved 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. Application of this methodology to the problem of a single object vs the entire space catalog of nearly 20000 objects (or worse yet to the problem of all catalog objects *vs.* all other catalog objects) quickly illustrates the need for computational acceleration techniques.

* Chief Orbital Scientist, Analytical Graphics, Inc., 220 Valley Creek Blvd., Exton, PA, 19341-2380.
† Senior Astrodynamics Specialist, Analytical Graphics, Inc., 220 Valley Creek Blvd., Exton, PA, 19341-2380
‡ Astrodynamics Engineer, Analytical Graphics, Inc., 220 Valley Creek Blvd., Exton, PA, 19341-2380.

Hoots et al.[1] 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 the orbital motion of the two objects. These filters serve to “weed out” the majority of the objects in the catalogue from intense scrutiny thereby greatly reducing the number of computations needed. After the application of the filters, the trajectories of the remaining candidate objects are sampled to determine the actual close approach periods. The three filters will be referred to as 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 have been shown to be inadequate when implemented as originally described. Additionally, the filters were originally 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.

Healy[2] took a different approach to conjunction detection where filtering is forgone in an implementation designed to take advantage of parallel processing. There is also suggestion in this work that special care and additional analysis are required to use the filters described by Hoots et al. which makes their use cumbersome in a parallel processing application. 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 efficiently identify potential conjunctions during a time step. The candidate conjunctions are then subjected to a more detailed analysis using the full fidelity orbit model. Rodriguez[3] et al. provide a set of refinements of the method described by Healy but also incorporate a form of the apogee/perigee filter described by Hoots et al.. The discussion in this reference points out the complexity of the orbit path and time filters and implies a lack of robustness. While this approach taken in Rodriguez et al. certainly has its merits, we seek to explore the potential of the filters originally described by Hoots et al. to determine when and how the concepts behind those filters may be safely applied to improve the speed of conjunction detection.

For each of the filters put forth in Hoots et al., we will examine the premise of the filter, demonstrate failure cases and provide details of an improved implementation. The assumptions and potential failure cases for the improved implementations will be identified. Requirements for the use of computational ***pads*** will be examined along with the associated impact on computation time. The computational efficiency provided by reasonable filter combinations will be tabulated to demonstrate the relative value of each filter. Finally, the reliability of the filtering process will be evaluated using an “all on all” example where results will be generated with and without the use of conjunction filters.

TRAJECTORY SOURCE

The source of the trajectory information to be used for identifying orbital conjunctions is an important consideration in the configuration of the conjunction detection process. Historically the main source of ephemeris information used in orbital conjunction analyses has been the catalog of two line element (TLE) sets generated by the United States Space Command (USSPACECOM). Ephemeris information is generated from a TLE using the SGP4 General Perturbations (GP) propagation algorithm[4,5]. The SGP4 algorithm allows for the single point computation of satellite position and velocity at a time of interest directly from the input TLE. This feature of the analytic propagation method can be exploited to minimize the number of computations required during the conjunction filtering process.

In recent times, USSPACECOM has begun generating a second form of the space catalog where the trajectory of each satellite is represented by a state vector and force modeling settings to be used in conjunction with a specific Special Perturbations (SP) propagator. When using the SP version of the space catalog, it is necessary to numerically integrate the trajectories of each satellite either prior to or during the conjunction detection process. The resulting tabulated ephemeris may then be interpolated by standard techniques to compute the satellite position and velocity at a time of interest. The time required to compute the ephemeris is typically large when compared to the time required to perform the conjunction detection process. This is even true when the “all on all” conjunction problem is addressed[6]. The use of SP ephemeris need not change the way in which conjunction filters are used, but could result in shorter run times compared to the use of TLEs if the improved accuracy of the SP generated ephemeris is leveraged to reduce the size of the detection threshold.

Another viable source of trajectory information for use in conjunction detection is owner/operator provided data[7]. Tabulated ephemerides provide the most directly usable information. Some operators, especially those flying geosynchronous spacecraft, prefer to use simple analytical models for their trajectories. In this case, the propagation model can be integrated into the conjunction detection process or can be used external to the process to generate tabulated ephemerides. Owner/operator provided ephemeris information is typically more accurate than USSPACECOM provided information due to the inclusion of cooperative tracking data during the orbit estimation process.

It is appropriate to mention at this time there is an important nuance related to the source of ephemeris information that can affect the use of conjunction filters. The apogee/perigee and orbit path filters are geometric constructs which are based on assumptions of near two-body motion. The conversion between the Cartesian position and velocity of an object and an elliptical representation of its orbit requires the use of a gravitational parameter for the case of osculating orbital elements or a gravitational parameter, reference distance and a J2 coefficient of the gravity field for simple mean orbital elements. Obtaining the correct elliptical representation of the orbit requires that the appropriate gravity field values be used for all objects. In some cases this will require that multiple gravity field values be used in a single analysis. We also note that the mean elements used in this analysis are Kozai-Izsak mean elements[8,9]. Specifically, conversion to the selected mean elements removes the first order short periodic variations due to J2.

All results in this paper were generated using the publicly available catalog of TLEs for 11 February, 2009. This version of the catalog contained 11970 objects. The analysis runs covered time spans between one and five days beginning at 05:00 hours UTC. For cases where timing results involving the use of ephemeris files is presented, the original TLE information was used to produce ephemeris files for all objects in the catalog.

CONJUNCTION FILTERS

Orbital conjunctions are identified through the examination of pairs of trajectories of orbiting objects. The goal of the process is to find all 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 pairing 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. The difficulty associated with using pads lies in selecting pad values that preserve the accuracy of the computation while still providing computational benefit. Another technique to minimize the effects of filter approximations is to perform verification and rectification computations during the application of the filters. Such “trust but verify” strategies ensure that the filter processes do not walk away from reality. 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, 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.

One additional filter, not described in detail below, is typically applied during the prediction of conjunction events. The out of date filter is used to eliminate secondary objects from consider if their orbital information is not considered to be current enough to support accurate conjunction identification.

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.

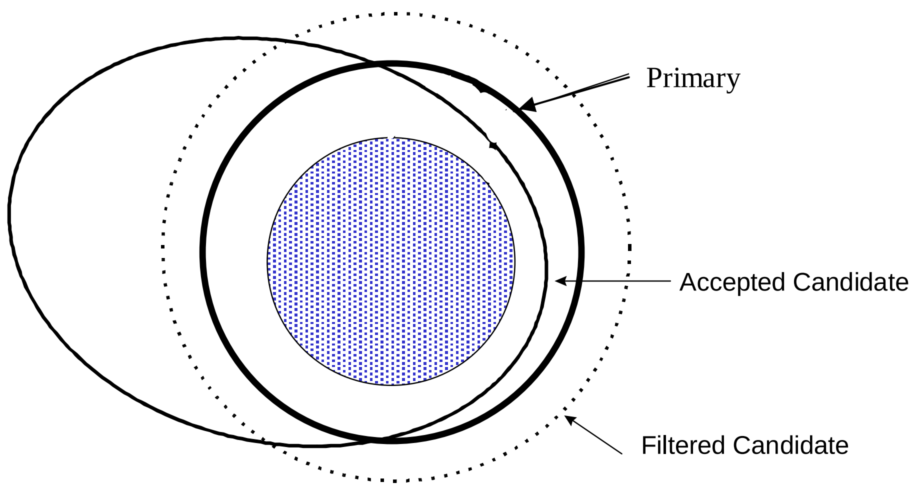

Figure 1. Apogee/Perigee Filter Geometry

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 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.

For the purpose of this study we computed the apogee and perigee values at the start time based on osculating and mean orbital elements. Testing the radial distance overlap condition based on information at a single epoch within the analysis span can be shown to lead to the erroneous elimination of candidate pairs. The approximation breaks down due to periodic variations in the elliptical representation of the orbits due to J2 and luni-solar perturbations. Table 1 contains an example test case which yields a conjunction based on a 5 Km detection threshold which is missed when applying the apogee/perigee filter based on the osculating orbit elements at the start time of the analysis period. In this case, the candidate pair was rejected by the apogee/perigee filter even though the actual minimum and maximum radius values for the primary object were used during the filtering process leaving the only source of error to be the computed apogee and perigee values for the secondary object.

Table 1. Apogee/Perigee Filter Failure Case

Object	Two Line Element Set
Primary	1 17191U 86097A 09042.77558376 .00000463 00000-0 33698-4 0 4241 2 17191 82.5015 115.7853 0012022 11.7757 348.3745 15.01450187205077
Secondary	1 26281U 99057FT 09042.35988229 .00004627 00000-0 30570-3 0 9896 2 26281 98.3425 293.5702 0025414 292.4435 67.4052 15.07724983479847

Figure 2 shows comparisons of computed apogee and perigee radius values for the Low Earth Orbiting (LEO) satellite, SSC #26281. The reason for the failure of the apogee perigee filter based on a single sample of the osculating orbit elements at the start time of the analysis is evident. The sampled value for the apogee at the start time of 6913 km is in error by approximately 20 km from the actual value of the apogee and causes the filter to determine that there is no overlap with the SSC #17191 which had a perigee value of 6929 km. For the case of SSC #26281, it appears that the apogee and perigee radius values as computed from mean orbital elements are better suited for conjunction filtering than those computed from osculating values.

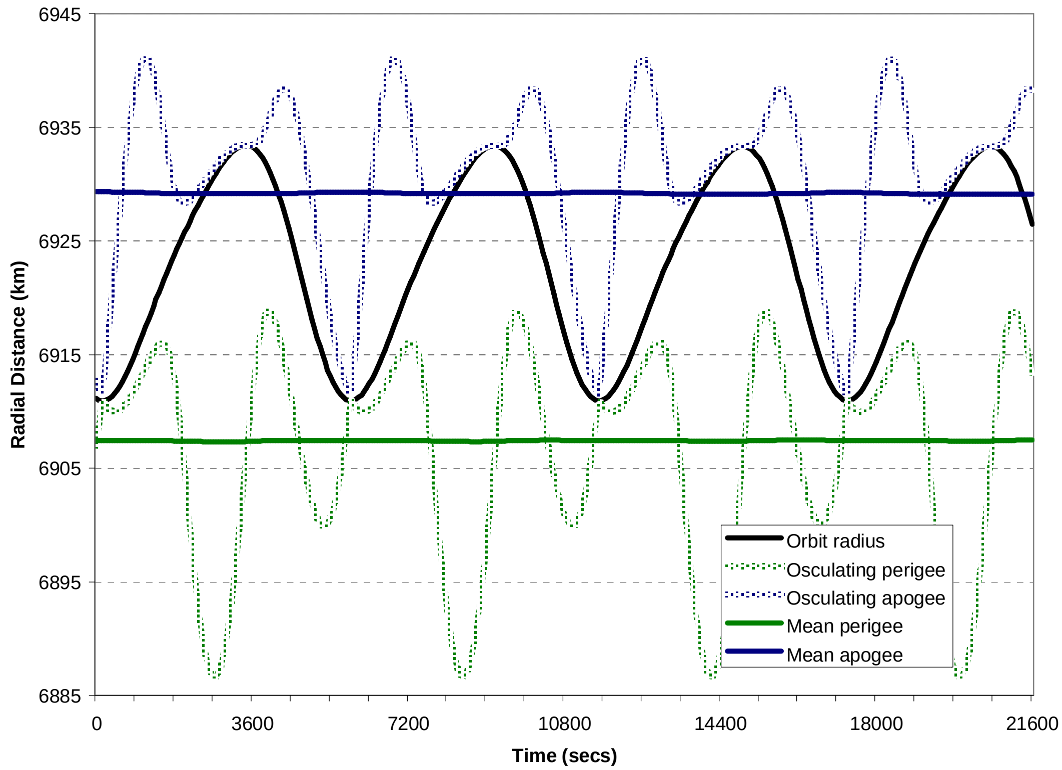

Figure 2. Comparison Of Computed Values Of Apogee And Perigee (LEO, SSC 26281)

The behavior of Highly Eccentric Orbits (HEO) is seen to be significantly different than that of the nearly circular LEO orbit. The HEO orbit used for this analysis is actually a Geostationary Transfer Orbit (GTO) with perigee in LEO and apogee near the geostationary belt. In this case, the computed osculating and mean apogee and perigee radius values both provide poor representations of the actual apogee and perigee values unless they are sampled near the time of the extremum, Figure 3. The difficulty in this case is caused by the effects of third body perturbations which are not removed in the selected mean element theory. The HEO result also demonstrates that the assumption of a constant perigee radius does not apply for this type of orbit.

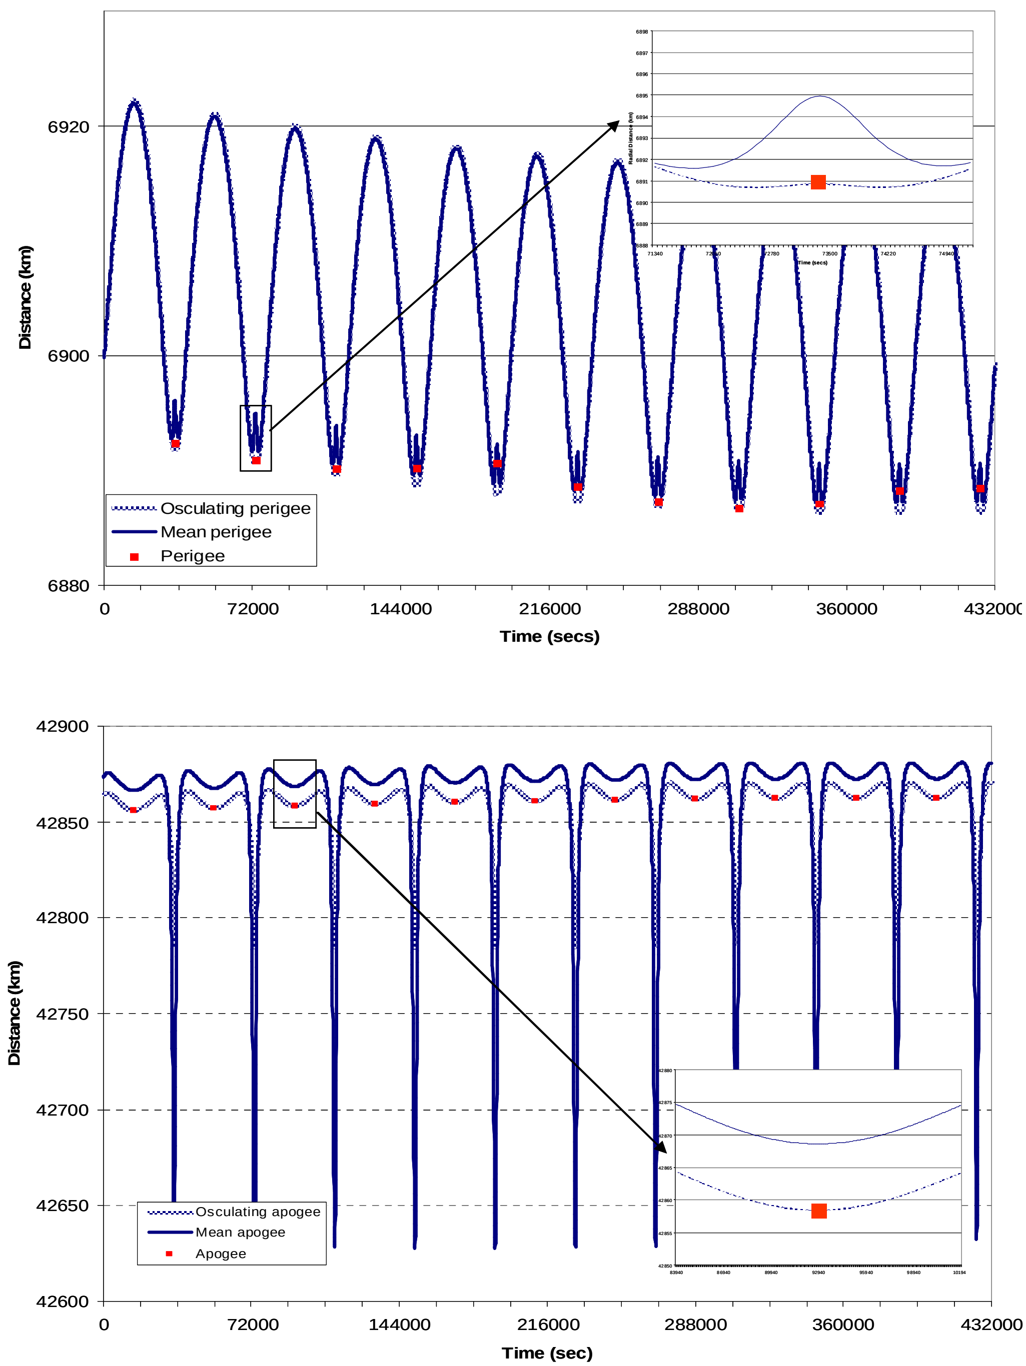

Figure 3. Comparison Of Computed Values Of Apogee And Perigee (HEO, SSC 23687)

Two methods for improving the apogee/perigee filter were investigated. The first method involves adding a pad to the detection threshold and sampling each trajectory at the beginning and end of the analysis interval to determine the range of radial distance for each object. The reason for sampling at both ends of the analysis interval is to capture the effects of trends in the apogee or perigee values as was observed in the case of the HEO. The use of multiple samples should also reduce the chance of using a particularly unfortunate sample such as at the bottom of a trough in Figure 3. The application of pads during the filtering process provides a simple means of accounting for periodic effects on the orbit trajectories. The goal in the selection of the pad value should be to select the minimum value that sufficiently covers the periodic effects. The minimum value is desirable to minimize the number of candidate pairing which will pass through the filter and thus reduce the required processing time. The selection of a minimum sufficient pad value is difficult, however, especially when Highly Eccentric Orbits (HEO) are involved. As a result, the pad needs to be set large enough to cover all test cases.

To provide an estimate of the required pad sizes, a simple analysis was performed on the set of catalog objects to determine the maximum error between the actual largest and smallest radii values for each object and what would be predicted by computing the osculating or mean orbit elements at the start and end times of the analysis interval. Based on the resulting graphs, shown in Figures 4 and 5, it appears that using a pad of approximately 30 Km per satellite would be sufficient to capture all cases successfully. We note that if a small number of primaries are being considered, then sampling of the primary trajectories can be done to determine their actual radial extents thus reducing the padding requirement to that of a single satellite. While this result is strictly valid only for the test case used, we have found it effective in practice. As shown in Figure 3, however, it is still possible for the apogee/perigee filter to fail under these conditions.

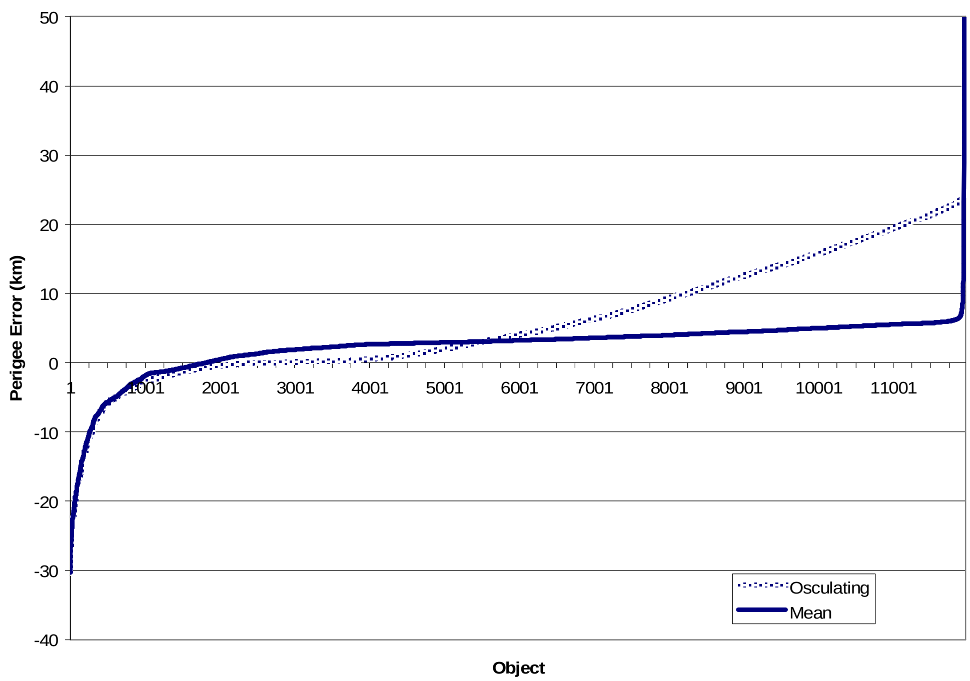

Figure 4. Perigee Approximation Errors

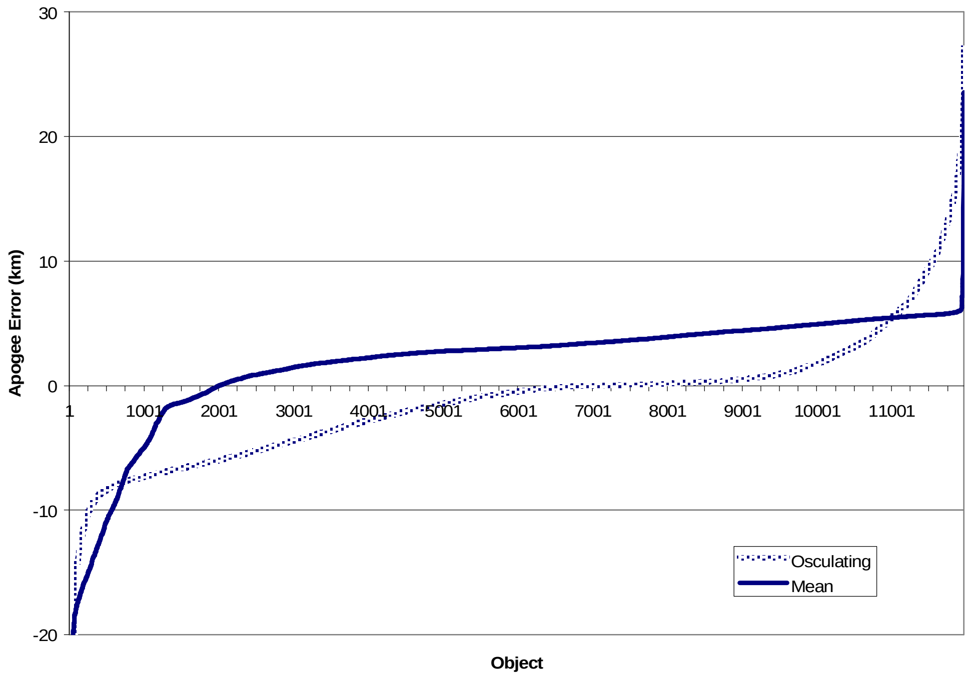

Figure 5. Apogee Approximation Errors

A more reliable method for improving the robustness of the apogee/perigee filter is to sparingly sample all trajectories to determine the minimum and maximum radial distance for each one. This method requires significantly more computation than the first method where the object positions and velocities were only required at the start and end of the analysis interval, but the resulting values may be used with much smaller pads, 1 km is sufficient, thus increasing the number of candidate pairs which are filtered. This trade off will become important as the number of primary objects is increased. The computational cost of this method is also highly dependent on the source of ephemeris information. It is much more costly when used with orbit information specified by TLEs than when tabulated ephemeris is used. Another advantage of the sampling method is that one need not make any assumptions about the type of motion which occurs over the analysis interval. This allows for the application of the filter to cases where the trajectories under consideration contain maneuvers. The efficiency of the technique by which we sample and detect the min/max radii will have a significant impact on the performance of the filter when using this method.

The Orbit Path Filter

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 6. 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. For cases where either orbit has moderate eccentricity, however, the Newton method usually requires an initial guess which is closer to the final solution than the points along the relative node in order to converge. Woodburn and Dichmann[10] presented an improved algorithm for computing the distance between the elliptical paths. The orbit path filter as described may not be applied to objects in coplanar orbits as the algorithm for finding the closest point between the paths becomes ill defined. For coplanar cases, the candidate pair is passed through for additional consideration.

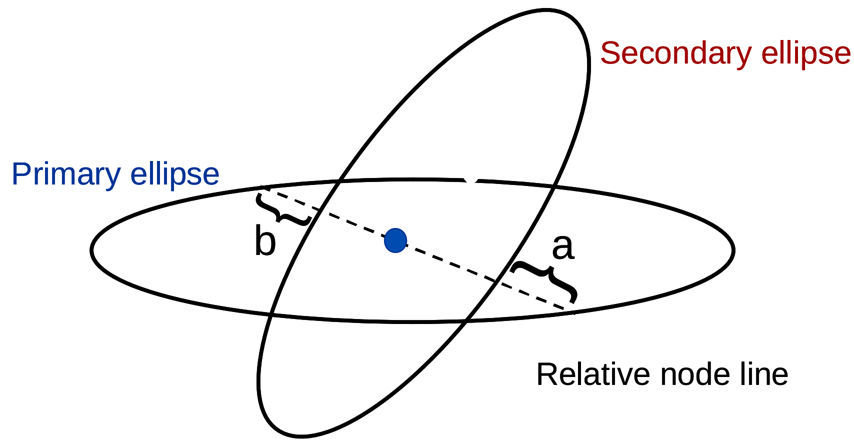

Figure 6. Orbit Path Filter Geometry

Success of the orbit path filter depends upon the time at which the filtering metric is computed. Hoots et al. recommend sampling once at the midpoint of the analysis interval and computing the rate of change of the path to path distance based on the secular rates of the orbit elements due to effects of oblateness. The rate at the midpoint is then projected to the start and end times of the interval. If the minimum resulting path to path distance based on the sampled value and the two projected values is less than the detection threshold then the candidate pair is passed through the filter. While not explicitly stated, it is assumed that path to path distances were computed based on mean orbital elements.

As was the case with the apogee/perigee filter, the computation of the rates of the orbital elements requires information which is not generally available. As an alternative to using the rates, we have chosen to compute sample the path to path distances at the start, middle and end of the analysis span. The pair of TLEs contained in Table 2 provide a test case for which the orbit path filter fails under this sampling strategy. In this case, all three samples are well above the detection threshold of 5 km, yet a conjunction occurs.

Table 2. Orbit Path Filter Failure Case

Object	Two Line Element Set
Primary	1 00130U 61015Q 09042.53163123 -.00000058 00000-0 13804-4 0 1891 2 00130 66.7709 101.1030 0080133 49.8006 311.0048 13.98086160426145
Secondary	1 10730U 75027E 09041.68856875 -.00000310 00000-0 -10589-3 0 6011 2 10730 114.9454 275.4040 0122342 287.9987 70.7850 13.92737619721619

Similar to our first approach with the apogee/perigee filter, we can apply a pad to the detection threshold to avoid missing conjunctions due to the simplifying assumptions of the path filter. To gain insight into the size of pad that may be required, a time history of the distance between the paths for the objects in Table 2 was generated. Figure 7 shows a comparison of the path to path distance computation for this case where the distance is computed using both mean and osculating orbital elements. The graph clearly demonstrates two things: the computation using osculating elements contains far too much frequency content to be useful under a condition of minimal sampling and the computation using mean elements can differ significantly from the actual instantaneous value of the metric.

It is clear based on Figure 7 that simply applying a pad, while it may account for the differences between the mean and osculating orbit representations, will not solve the problem of improper elimination of candidate pairs in the orbit path filter. A better sampling strategy is also required and needs to be combined with technique to find the minimum of the path to path distance over the analysis interval. We note that had we used the metric as proposed by Hoots et al., the filter would not have failed in this case, but it can clearly be seen that there are large areas of the curve where the filter would have failed had the sample been taken there.

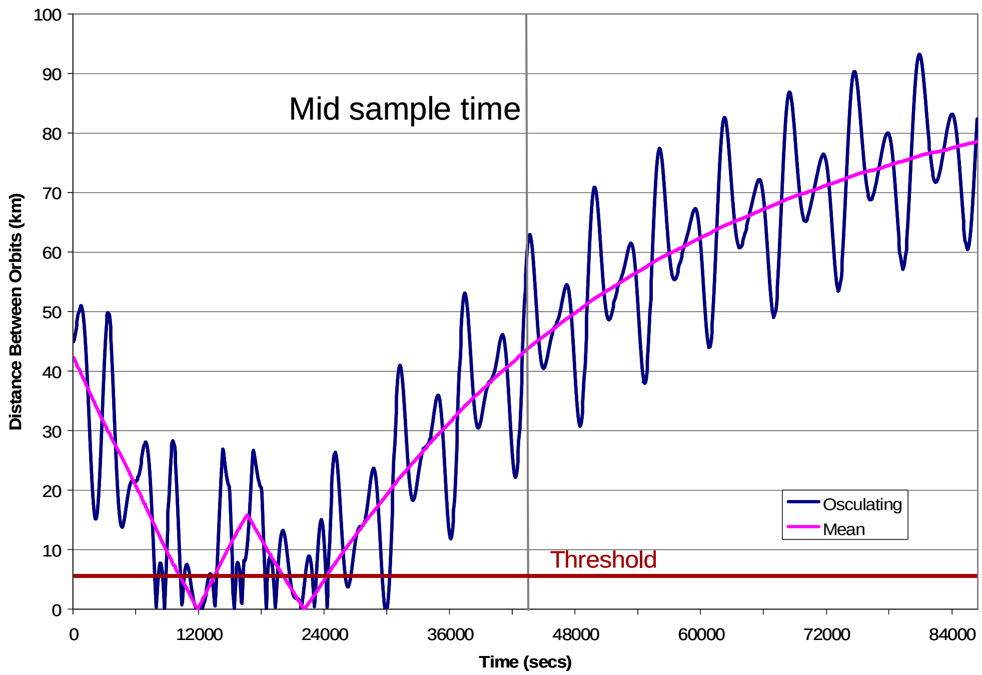

Figure 7. Computed Distance Between Orbits Using Osculating And Mean Elements

The orbit path filter requires a fairly complex computation, the path to path distance between two orbits, the result of which has been shown to have non-trivial behavior. This observed behavior of the path to path distance metric needs to be wrapped in additional event detection logic to identify cases of possible conjunction as shown in Figure 7. The event detection logic needs to include a sampling strategy which allows for the detection of the important trends in the sampled metric. While these additional refinements to the algorithm described here are possible, the coding complexity and computational load of the filter may out weigh its benefit. It should also be noted that use of the orbit path filter is restricted to cases where neither object is maneuvering.

The Time Filter

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 8. 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.

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 visa versa, Figure 9. Like the orbit path filter, the time filter is limited to cases where the trajectories are not coplanar and neither object is maneuvering.

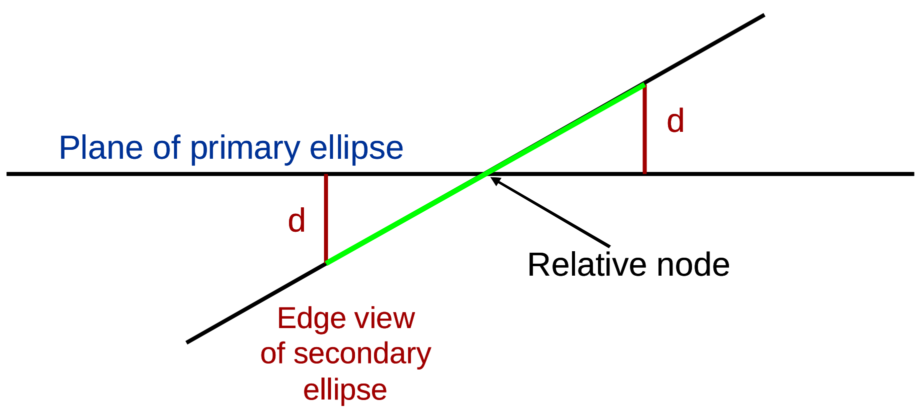

Figure 8. Distance From Primary To Orbit Plane Of Secondary

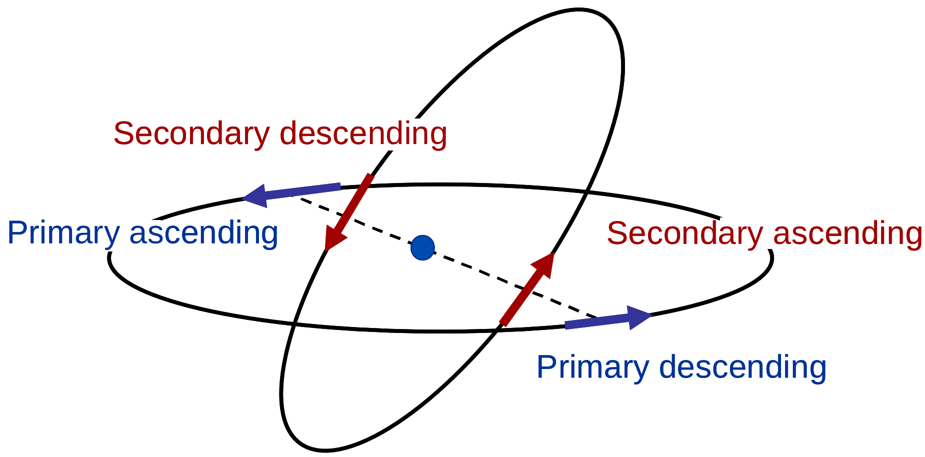

Figure 9. Relative Node Crossings

The original concept for generation of time intervals as presented by Hoots et al. involves generating a series of intervals for each object when the metric satisfies the threshold condition starting with an interval near the mid point of the analysis period. The relative nodal period, corrected for the effects of drag and oblateness, is then used to generate a series of intervals covering the analysis time period. The corrections for drag are computed in terms of the rate of the change of the mean mean motion, which is provided as part of a TLE. The idea of starting at the midpoint of the analysis interval is to reduce the effect of perturbations which are not accounted for in the formulation of the filter. While this methodology is very efficient at generating intervals, experience has shown that generated intervals do not necessarily represent the actual times near the relative node. We do not show a specific failure case for this filter since all failures are simply related to an inability to compute a relative nodal period which is valid over the entire analysis interval.

The robustness of the time filter can be improved significantly by modifying algorithm to use a “trust but verify” approach to the generation of node crossing intervals. The modification to the interval generation involves finding the precise period between the first two ascending relative node intervals and the first two descending relative node intervals for each satellite. The relative node times are computed using osculating elements since we are now interested in specific nodal crossing times. The next interval of each type (ascending or descending) is then generated by adding the previous period. The new relative node intervals are then tested and corrected which yields new estimates of the ascending and descending relative nodal periods. The testing of the relative node crossing times requires more computation time than the original method, but allows the time filter to be applied in a safe manner. One way to reduce the amount of time required by the new methodology is to relax the requirement of correcting every nodal crossing to one of correcting every Mth nodal crossing. The correction at the Mth nodal crossing can then be monitored and the algorithm can reduce the initial value of M if the error in the predicted nodal period is too large. A simple way to maintain the integrity of the interval solutions in this case is to lengthen both ends of the computed intervals by a time pad and then verify that the correction in the node time at the Mth nodal crossing is less than the time pad. The general algorithm for generating the intervals which bound the crossings of the relative nodes is depicted in Figure 10.

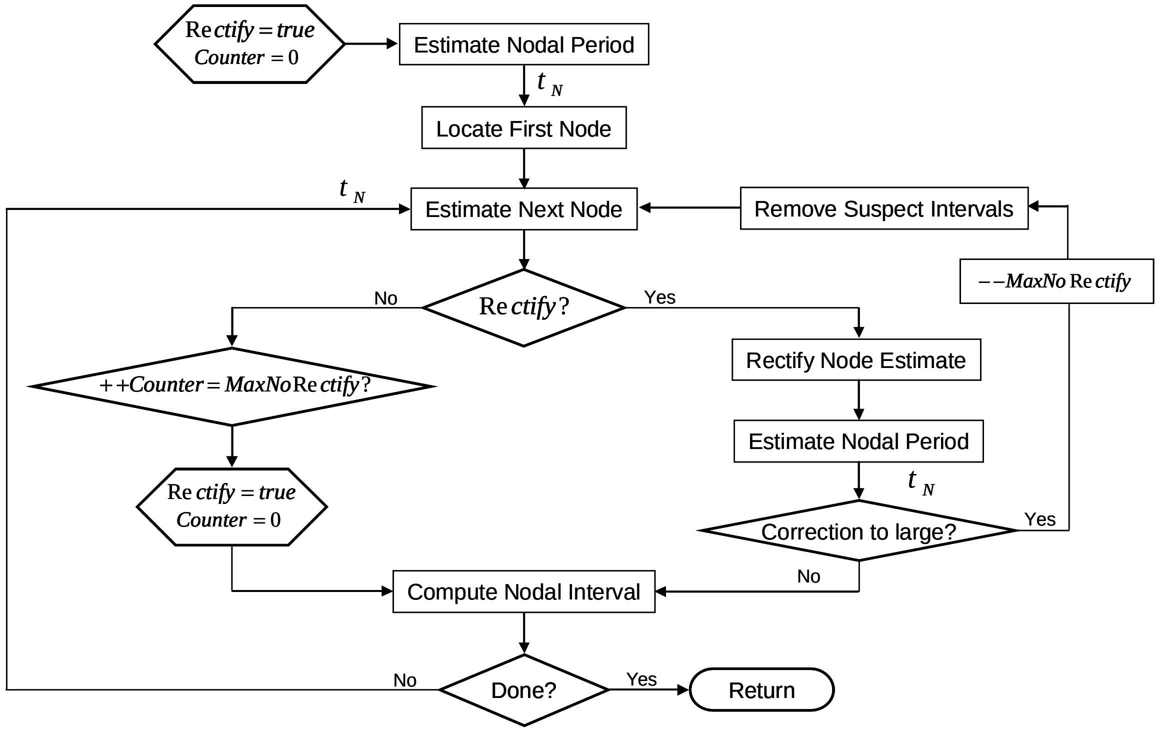

Figure 10. Relative Node Interval Generation

Range Threshold Detection

For the case where a minimum range threshold is being used as the measure of conjunction risk, it is ultimately necessary to determine the actual periods where the range is less than the operator specified value. In reality, these periods, or at least the time of closest approach, generally need to be identified regardless of the risk metric. It is therefore necessary to employ an event detection scheme to identify precise conjunction events. The event detection needs to operate over the intervals of time which cannot be eliminated by the filters. The efficiency of the of the event detection algorithm has a large impact on the overall performance of the conjunction identification process when the time filter is not used. When the time filter is used, however, the amount of time spent in event detection is reduced to a level where the efficiency of the event detection is much less important.

RESULTS

The series of filters described above were applied to the problems of conjunction detection for each of the following problems: one LEO primary, one GEO primary, one GTO primary, 10 LEO primaries, 10 GEO primaries and "all on all". The object ids used as primaries for each problem are listed in Table 3. A variety of filter configurations were tried for each problem. The combination of a problem and a filter configuration will be referred to as a test case. Descriptions of the filter configuration are given in Table 4. The filters were applied in the order listed in Table 5. Each test case was run with two ephemeris sources: TLEs and tabulated ephemeris. For each run, the computation times are normalized by the time required to perform the conjunction analysis without filters and are reported as speed up factors. A reported value of 2, for example, indicates that the computation required half the time required for the no filter case. The results are given in Table 5 with values listed inside parentheses indicating speed up factors when tabulated ephemeris was used. It is important to note that for cases where tabulated ephemeris was used, the speed up factors were computed after the common time to load the ephemeris into memory was subtracted from the overall process time. In these cases, the time required to load the ephemeris accounted for the vast majority of the overall process time so the practical effect of the filters is not nearly as great as indicated by the values in Table 5. The numbers in Table 5 should be thought of as being representative of the effectiveness of the various filter combinations, but will vary with specific primary selections and filter implementations. Due to the problems noted above, 4 out of 5338 total conjunctions were missed in the all on all test case when the orbit path filter was used. Associated timing results are shown in yellow in Table 5. The decision to include the results was to show the potential benefit of the orbit path filter if further improvements to its implementation were made.

Table 3. Descriptions Of Primary Object Sets Used In Timing Tests

Primaries	Description
1 LEO	SSC#10730
1 GEO	SSC#26352
1 GTO	SSC#23687
10 LEO	10 IRIDIUM satellites SSC# 24837, 24839, 24840, 24869, 24872, 24965, 25577, 25578, 25777, 25778.
10 GEO	10 INTELSAT satellites SSC# 20315, 20523, 22871, 23175, 24916, 25473, 26590, 26766, 28358, 32253.
All on All	Full catalog of 11970 objects

Table 4. Descriptions Of Filters Used In Timing Tests

Filter	Description
APP	Apogee/perigee filter using precise apogee/perigee computation for primary objects and sampling at the beginning and end times of the analysis interval for secondary objects. Pad of 30 km.
APS	Apogee/perigee filter using precise apogee/perigee computation for primary and secondary objects. Orbits sampled every 20 degrees of true anomaly. Pad of 1 km.
OP	Orbit path filter using samples at the beginning, middle and end of the analysis interval based on mean orbit elements. Pad of 30 km.
T	Time filter using a maximum number of 5 consecutive nodes without rectification. Pad of 30 km plus 10 seconds on each side of the computed interval.

Table 5. Conjunction Timing Results

Filters	1 LEO	1 GEO	1 GTO	10 LEO	10 GEO	All on All
APP	1.7 (2.2)	9.6 (11)	1 (1)	3 (2.6)	13 (12)	2.3 (2.5)
APP/T	12 (5.9)	26 (17)	9.3 (2.5)	40 (7.3)	57 (20)	17 (15)
APP/OP/T	16 (7.3)	31 (20)	34 (5)	54 (8.5)	57 (19)	40 (32)
APS	0.79 (0.98)	0.95 (0.74)	0.53 (0.57)	4.3 (2.8)	9.1 (5.6)	2.5 (2.5)
APS/T	1.2 (1.1)	1 (0.75)	0.96 (0.9)	15 (6.2)	13 (6.4)	18 (18)
APS/OP/T	1.2 (1.2)	1 (0.76)	1.1 (1.1)	16 (10)	14 (6.4)	42 (34)

When comparing the numerical results in Table 5, it is important to remember that the numbers represent computational acceleration factors. In this sense, a value of 15 is not five times better than a value of 10, it is 50% better. It should also be noted that there is significant uncertainty, probably on the order of 25%, in the reported values and that the specifics of a particular implementation will have a strong influence on the results as well.

Several things jump out immediately from the tabulated timing results. First, locating the precise apogee and perigee values for the secondaries (APS) in the apogee/perigee filter can result in longer run times than when no filters are used for cases where only a single primary satellite is considered. When larger numbers of primaries are considered, however, the up front computational hit of the APS filter is compensated for by the improved efficiency of the filter in terms of eliminated candidates. In the case of all on all conjunction processing, the APS filter provides the same performance as the less robust APP version. Second, the time filter is seen to provide the largest jumps in computational efficiency of all of the filters. This is due to the fact that the time filter not only eliminates candidate objects, it also greatly reduces the amount of time required for detailed analysis on the remaining conjunction candidates. Finally, the orbit path filter is seen to be most effective in reducing the computational time for highly eccentric orbits. This is mainly due to the fact that the apogee/perigee filter is not effective in reducing the candidate population for these types of orbits. In the case of near circular LEO and GEO primary objects, the orbit path filter is least effective of all of all of the filters but it does reduce the time required for all on all conjunction detection by a factor of two.

CONCLUSION

The use of a series of filters has been shown to be an effective means for reducing the computation time required for conjunction analyses. Two of the three classic filters, the apogee/perigee filter and the time filter, have been shown to be robust based on modifications to the implementation strategies. The orbit path filter is seen to be the most troublesome of the classic filters. Except for the case where the primary objects are in HEO or the all on all case, the associated computational benefit does not justify the additional effort required to improve the robustness of the filter. The recommendation of this paper is, therefore, to not use the orbit path filter. The time filter is seen to provide the greatest reduction in required computation time despite the more complicated logic involved in the modified implementation.

The filters as described here are equally applicable to conjunction analyses based on two line element sets or tabulated ephemerides. The efficiency gains differ depending on the ephemeris source due to the much higher cost of computing the position of an object using SGP4 compared to interpolating a provided ephemeris. The use of tabulated ephemerides results in an up front cost due to the time required to load the ephemerides into memory, which scales linearly with the number of secondary objects, N. The up front cost is quickly overcome as the number of primaries, M, is increased due to the improved efficiency of filter operations whose cost is proportional to N*M.

REFERENCES

[1] 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.

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

[3] 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.

[4] Hoots, F.R., and Roehrich, R.L., "Spacetrack Report #3: Models and Propagation of the NORAD Element Sets", U.S. Air Force Aerospace Defense Command, Colorado Springs, CO, 1980.

[5] Vallado, D.A., Crawford, P., Hujsak, R., Kelso, T.S., "Revisiting Spacetrack Report #3", AIAA 2006-6753, 2006 AIAA/AAS Astrodynamics Specialists Conference, Keystone, CO, August 2006.

[6] Coppola, V., Dupont, S., Ring, K., Stoner, F., "Assessing Satellite Conjunctions for the Entire Space Catalog using COTS Multi-core Processer Hardware", AAS 09-374, 2009 AAS/AIAA Astrodynamics Specialists Conference, Pittsburgh, PA, August 2009.

[7] Kelso, T.S., Vallado, D.A., Chan, J. Buckwalter, B., "Improved Conjunction Analysis via Collaborative Space Situational Awareness", 2008 AMOS Conference, Maui, Hawaii, September 2008.

[8] Kozai, Y., "The Motion of a Close Earth Satellite", The Astronomical Journal, No. 64, pgs 367-377, November 1959.

[9] Izsak, I.S., "A Note on Perturbation Theory" The Astronomical Journal, No. 68, pgs 559-561, October 1963.

[10] 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.