Large-Scale Simulation of a Process for Cataloguing Small Orbital Debris

Don Phillion, Alex Pertica, Ben Fasenfest, Matt Horsley, Wim De Vries, H. Keo Springer, David Jefferson, and Scot Olivier
Lawrence Livermore National Laboratory

Keric Hill
Pacific Defense Solutions

Chris Sabol
Air Force Maui Optical and Supercomputing Center

ABSTRACT

We demonstrate a methodology for establishing orbits for the abundant, un-catalogued, yet dangerous, small orbital debris that will become observable with planned improvements to the Space Fence. Although roughly 15,000 orbital objects are present in the SSN catalog, it is believed that at least 200,000 objects that are massive enough to cause significant damage are in Earth orbit. With improvements to the Space Fence, LEO debris down to 5 cm in size may become observable. The additional hundreds of thousands of observations a day of mostly un-catalogued objects will present a significant data processing challenge. Of particular concern are the large numbers of observations that are uncorrelated either to a known object or to a single object. To deal with the large-scale uncorrelated track (UCT) problem, we have ported the Covariance-Based Track Algorithm (CBTA) into the supercomputer-based Testbed Environment for Space Situational Awareness (TESSA) in order to perform simulations at scale.

CBTA bins UCTs for which initial orbits and initial covariance matrices could be determined back to a common epoch and then uses a statistical measure to see if they correlate given the state vectors and covariance matrices at that common time. If they do, the observations from the two tracks are combined and orbit determination (OD) is used to attempt to fit an orbit to the combined tracks. If OD converges, a new UCT hypothesis is created and the state and covariance of that hypothesis is saved with the other pre-existing UCTs. If a certain number of tracks are successfully combined then they are used to create a new catalog object. Old UCTs are weeded out of the pool of hypotheses when they become obsolete, or when at least some of the observations are used to create a new catalog object.

For the simulation, we developed a Radar detection model simulating the performance of a notional new Space Fence. We propagated thousands of objects over a several day period creating a large number of observations. The methodology that we employed first attempts to match tracks to known orbits using an orbit determination process. Most of the observations cannot be correlated to known orbits and these are routed to the CBTA. We will report on the efficiency with which this hybrid process is able to catalog new objects and on the computational requirements necessary to deal with the problem at scale.

1. INTRODUCTION

The recent Cosmos 2252 – Iridium 33 collision demonstrated the catastrophic effect of large-body collisions. The Space Surveillance Network (SSN) is still tracking thousands of pieces of debris from this collision. The closing velocity from that impact was 11.7 Km/s which is typical for engagements in Low Earth Orbit (LEO). At closing velocities exceeding 10km/s, even much smaller debris is capable of incapacitating an operational satellite. Furthermore, the ensuing debris from such a collision would pose a threat to other satellites. To understand the effects of small debris collisions onto satellites, Lawrence Livermore National Laboratory performed a series of simulations employing the PARADYN hydrocode on a high-performance Linux cluster to generate debris from a matrix of varying size debris onto a satellite target at different closing velocities. The methodology behind this work is the subject of a separate paper presented at this conference. The results of these simulations show that debris in


Figure 1. Simulation of a 10cm radius aluminum ball impact onto a Iridium satellite at 10Km/s closing velocity

the few-centimeter to few-tens-of centimeter range can cause catastrophic collisions. The simulated result of a 10cm radius aluminum sphere onto an Iridium satellite is show in Fig. 1.

The above results motivate the need to be able to track and produce actionable collision warnings for small debris. The SSN maintains an orbital catalog of satellites and large debris using mainly ground-based radars for LEO objects and Ground-based telescopes for GEO objects. Many of the new object detections are made with the Space Fence, a collection of bistatic radars spread roughly in a line across the United States. The development of a next-

generation Space Fence is underway with contracts in place for the initial design, systems engineering, and trades analysis. Although the final configuration of the New Space Fence is not known at this time, the expectation is that this sensor system will have greater sensitivity than the current sensor and thus be able to detect crossings by objects only a few centimeters in size. Additionally, it is expected that the radar positional accuracy and the radar fan width will be such as to enable accurate enough initial orbits to be determined so as to be able to use the CBTA to join space fence UCT's. According to the NASA Orbital Debris Program Office, the estimated population of particles between 1 and 10 cm in diameter is approximately 500,000 and the number of particles smaller than 1 cm probably exceeds tens of millions[1]. Based on estimates for the density of small debris in LEO and MEO, the next generation Space Fence is likely to produce hundreds of thousands of tracks per day. Initially, at least, a large portion of these tracks will be from un-catalogued objects and this will present a significant challenge to sort out. In 2008, Hill et al. published the details of the Covariance-based Track Association algorithm for efficiently processing UCTs [1]. In this paper, we have implemented an enhanced version of the CBTA in a process that is scalable to the magnitude of the problem presented by future availability of New Space Fence UCTs.

2. ALGORITHM DESCRIPTION

Covariance Based Track Association (CBTA) joins together uncorrelated tracks for which state vectors and covariance matrices have been determined [2,3]. The initial two-body determined orbits for the uncorrelated tracks are refined by batch least squares fitting to SGP4 orbits. The CBTA is given a pool of uncorrelated tracks with these refined SGP4 orbits and the covariance matrices obtained from the batch least squares fitting. The tracks in this pool are time-ordered. SGP4 is adequate for joining tracks. Either SGP4 or force model orbits can be obtained for the combined orbits for the joined tracks. Only tracks for which a reasonably good initial orbit could be determined are put into the pool. The CBTA first finds all tracks that can be successfully joined with the first track in the pool. These tracks are then removed from the pool and the process is repeated with the next track left in the pool and so on. Combined orbits are then fit to the joined tracks. A joined orbit is valid if the observations are consistent with the combined orbit given the observational errors. If a joined orbit is not valid, then consistent subsets of the tracks in that joined orbit are sought. An invalid joined orbit should be a rare occurrence. Occasionally, the tracks for an object are pieced together into more than one joined track. The handling of invalid joined orbits, merging multiple joined tracks for the same object, and parallelization are areas of further work on the CBTA that are now being done. Parallelization will be done by dividing orbital space into overlapping orbital regions. Parallelization requires merging joined tracks for the same object so that the same object isn't added to the catalog more than once. We have already done some preliminary work on parallelization with pthreads to test these ideas. The algorithm that we have implemented is designed for speed. At the least, it needs to work for 100000 objects and 1000000 tracks.

The CTBA algorithm has been implemented in C++, so far on a single processor. Nevertheless, the CBTA joined 6000 tracks for 783 discovered objects in just 1290 seconds. For this run, the catalog had 1037 objects, of which 789 were observed. Not all of the objects were observed because some were in GEO orbits and some had inclinations less than the 33 degree latitude of the space fence. The simulation time-period was 1 day. Only one joined fence was invalid and it was correctly recognized as invalid. In all of the runs that have been done, it has never happened that an invalid track has wrongly been marked as valid. There were three broken joined tracks. This happens when the second track for an object is many orbits after the first track for the object. There is a maximum amount by which the mean longitudes can differ even if that difference is consistent with the covariance matrices. If a many orbit gap in the observations occurs after the second track, the orbital period can be accurately determined from the earlier tracks and the joined track is not broken. Multiple joined tracks for the same object can easily be joined afterwards, although this is not implemented yet.

The process employs a very fast filter which determines whether a track might be joined to another track. Provided this filter makes few mistakes and if the cpu time for this filter is negligible, then the scaling of cpu time with the number of tracks should be approximately linear. In fact, if the time to fit the combined orbits dominates then the

scaling should be linear as the number of objects. Eventually, for an exceedingly great number of objects, the scaling becomes quadratic with the number of tracks. Our study of the scaling so far indicates that this will not become a problem until the number of objects is in the millions. Studies of the scaling with very large numbers of objects (>100,000) will not be practical until MPI parallelization has been done.

Each of these steps will now be explained in detail. The first step is a very fast filter that screens which tracks might be joined to a selected track. It does so by comparing the equinoctial orbital elements. No propagation is done. The first five equinoctial coordinates are compared. The sixth equinoctial coordinate is not looked at because this would involve propagation and the idea is for the pre-screening to be extremely fast.

If a track passes this filter test for possibly being joined to a selected track, then the state vectors and covariance matrices are propagated to a middle time. Both are in equinoctial coordinates. The sum of the square roots of the two diagonals of the equinoctial covariance matrices are computed. If the absolute values of the differences of the two state vectors in equinoctial coordinates is less than several times the corresponding sums of square roots of the diagonal covariance matrices, then the tracks are marked as being joinable. There are maximum values for the allowable differences in the equinoctial coordinates regardless of what the covariance matrices are.

This process results in a number of joined tracks. However, a combined orbit must yet be fit to each joined track and it must be checked that the combined orbit fits the observations within several times their one-sigma errors. The first step to fitting a combined orbit is to find an initial combined orbit that will be refined with batch least squares. This initial combined orbit cannot be just one of the track initial orbits because batch least squares using all of the observations will frequently diverge.

We use Lambert's method [4] and the first and last tracks in the joined track to obtain this better initial combined orbit. Each track in a joined track has an initial orbit determination. A position at a time is obtained from the first track's orbit and another position at another time is obtained from the last track's orbit. The continuous angle theta between the two vectors is made to be nearly as possible either 90 degrees plus a multiple of 360 degrees or 270 degrees plus a multiple of 360 degrees. This is because Lambert's method requires that the whole number of completed orbits m be known and whether theta is in quadrants I and II or not be known. When theta has one of these two sets of values, these two things are determined the most certainly. Lambert's method may have more than one two-body solution even given m and whether theta is in quadrants I and II or not. Which solution is correct is determined by which best fits all the observations.

Once a two-body orbit is found which goes through the two positions at the two times, we want to find an SGP orbit which goes through the two positions at the two times. This is done by gradually walking the two-body solution to the SGP4 solution. Finally, we have an initial combined orbit that is good enough for iterative differential batch least squares orbit refinement. A joined track is a valid track if its refined combined orbit fits the observations within several times their one-sigma errors. If so, it can be added to the catalog as a newly discovered object. The batch least squares fitting can done using either the force model or SGP4. If SGP4, the fitting can determine if Bstar, the drag term, needs to be fit in addition to the six orbital parameters.

The CBTA can be run either in the SatOrb, which is a standalone Windows program with an elaborate GUI, as part of TESSA, which is an integrated SSA modeling environment, or stand-alone on the Livermore Computing high performance computers. The SatOrb CBTA GUI is ideal for testing purposes. TESSA is the Testbed Environment for Space Situational Awareness and runs on massively parallel Linux clusters using PDES (Parallel Discrete Event Simulation). The CBTA is one component in TESSA. Each component communicates with other components using events. In TESSA, the CBTA is given processed track reports from the ODAA (Orbit Determination And Association) module. These processed track reports include the initial orbit determination and the covariance matrix. For each processed track report, there is the track file in the same folder with a file name that can be determined from that of the file name of the processed track report. The CBTA joins tracks when it receives an

AssociateUCTsRequestSimEvent message and it generates both a report and a discovered object catalog that it sends in a DiscoveredObjectCatalogSimEvent event to the MTA (Multiple Track Aggregation) component, which maintains the updated catalog. There is a <CBTA> block in the TESSA setup XML file that has all the options and parameters for the CBTA. One option is for the CBTA to treat all tracks as uncorrelated tracks. This is what is used for testing purposes.

When the CBTA operates stand-alone, it can be given either a list of track file names or a list of processed track file names in a text file. If given a list of these track file names, all of these track files are processed to obtain the processed track report files. From then on, the CBTA operates as it does in TESSA. If given the name of a file containing a list of processed track file names, the track files must also be present and named according to a certain convention. In TESSA, the CBTA options are specified in the <CBTA> block in the setup XML file. When operating stand-alone, these options can also be specified in an XML file. If this file is not provided, the default options and parameters are used. This XML file can be a TESSA setup file.

The data products of the CBTA are a report and a discovered object catalog which includes the covariance matrices. There is an entry for each joined track followed by performance diagnostics. Even though all the tracks were treated as uncorrelated, in fact it is known what the ID's were for all the tracks. This information was not used, of course. Indeed, the CBTA does not even use a catalog. The ID's, if available, are only listed in the reports. No other use is made of them. The ID's are known because we are doing testing. In the field, the object for a joined track made up of uncorrelated tracks has no id. In this example, the id = 60000 was assigned, but it is actually the catalog object with ID=5. This joined track is marked as valid. That determination was made based on how well the combined orbit listed at the top fit all the observations. There is a line listing the average fitting errors to the observations followed by another line indicated the average observational errors. Sample output from the CBTA is shown in Fig. 2. Notice that the various tracks in this joined track all have different epoch mjd values. This is the reason that the mean anomalies have very different values.

```
Possible fuse track 60000
Lamberts method combined orbit:   a=8624519.455  e=0.184725  i=34.257  Omega=77.804
            omega=-222.173  M=663.689  Mjd=55207.2639814815
TWO LINE TLE
1 60000U 00000000 10011.26398148 .00000000 00000-0 00000-0 0  03
2 60000  34.2574  77.8039 1847253 -222.172 663.6893 10.83927919  04
RAt_Diff_Deg_Rms_batchlsq=0.004297   DECt_Diff_Deg_Rms_batchlsq=0.003095
Observational errors: rms_angle_error=0.004456   rms_range_error=3.17
JordanLake/Track_0494.xml  a=8624297.901  e=0.184590  i=34.263  Omega=77.828
            omega=-222.214  M=303.693  Mjd=55207.2639814815  Id=5
JordanLake/Track_0674.xml  a=8624431.438  e=0.184724  i=34.259  Omega=77.533
            omega=-221.767  M=301.970  Mjd=55207.3557986111  Id=5
GilaRiver/Track_0673.xml  a=8624585.758  e=0.184746  i=34.259  Omega=77.525
            omega=-221.760  M=305.587  Mjd=55207.3567245370  Id=5
LakeKickapoo/Track_0696.xml a=8624428.564  e=0.184717  i=34.257  Omega=77.516
            omega=-221.751  M=308.067  Mjd=55207.3573611111  Id=5
JordanLake/Track_0680.xml  a=8624025.342  e=0.184690  i=34.258  Omega=77.517
            omega=-221.737  M=317.904  Mjd=55207.3598842593  Id=5
GilaRiver/Track_0830.xml  a=8624513.674  e=0.184712  i=34.257  Omega=77.236
            omega=-221.348  M=304.775  Mjd=55207.4487731481  Id=5
LakeKickapoo/Track_0851.xml a=8624703.149  e=0.184751  i=34.257  Omega=77.223
            omega=-221.320  M=322.024  Mjd=55207.4531944444  Id=5
GilaRiver/Track_0841.xml  a=8624718.366  e=0.184760  i=34.257  Omega=77.233
            omega=-221.326  M=323.108  Mjd=55207.4534722222  Id=5
LakeKickapoo/Track_0858.xml a=8624507.721  e=0.184742  i=34.257  Omega=77.217
            omega=-221.303  M=331.819  Mjd=55207.4557060185  Id=5
JordanLake/Track_0840.xml  a=8625067.606  e=0.184767  i=34.258  Omega=77.211
            omega=-221.319  M=333.507  Mjd=55207.4561342593  Id=5
JordanLake/Track_0842.xml  a=8626238.753  e=0.184883  i=34.258  Omega=77.205
            omega=-221.315  M=339.746  Mjd=55207.4577314815  Id=5
GilaRiver/Track_1004.xml  a=8624172.847  e=0.184698  i=34.257  Omega=76.933
            omega=-220.890  M=334.979  Mjd=55207.5487731481  Id=5
LakeKickapoo/Track_1026.xml a=8615316.357  e=0.184170  i=34.252  Omega=76.954
            omega=-220.597  M=338.621  Mjd=55207.5497685185  Id=5
VALID_TRACK_FLAG=1
```

Figure 2. Sample output from CBTA algorithm

Orbital space is six-dimensional, so the number of objects can be enormous without having any two objects in close to the same orbit. Except for formation flying objects, it is highly unlikely that any long-term space objects will be in the same orbit. However, the debris objects from a break-up or collision event are initially in a three-dimensional sub-manifold. Early on, chances are that there will be objects in nearly identical orbits. However, objects that are in the nearly the same orbit but with very slightly different semi-major axis lengths will gradually move apart due to their different orbital periods. Even if there are debris objects in totally the same orbit, chances are that they will experience different atmospheric drag and their orbits will evolve differently. There may be some debris objects whose orbits cannot be determined soon after the event, but for which we must wait awhile. The test that joined tracks must match the observations within observational errors should prevent us from making any mistakes. We will be doing simulations with debris objects to learn more about this.

3. SPACE FENCE RADAR MODEL

The tracks used for the CBTA were generated using the high fidelity physics-based radar simulator which is part of TESSA. This simulator is designed in a modular fashion, where each module describes a particular physical process or radar function. It includes modules for atmospheric conditions, radio wave propagation, monopulse error estimation, antenna pattern effects, waveform generation, noise sources, and radar cross section computation. For each of these modules, multiple versions with varying levels of fidelity and speed are available. To accelerate the radar model when many computationally expensive modules are used, it is parallelized using MPI.

This radar model was used to implement a notional S-band space fence radar located at four installations. This radar is not meant to represent a specific future space fence radar, but to be similar to one possible design. The radar sites included three locations in the United States roughly along the 33rd parallel (Gila River, Lake Kickapoo, and Jordan Lake) as well as one Australian site, Canbera, located at latitude 35.3 South. Only the three US space fence sites were used in the TESSA runs. The radar operating frequency was 3.5 GHz, with an opening fan in the East-West direction from 10 degrees above both horizon. It was endowed with the ability to track objects through 22.5 degrees of arc when passing through the fan. The radar was modeled to use a monopulse array for the actual tracking of objects once detected. The positional accuracy was 0.03 degrees in angle and 10 meters in range. The radar power was tailored to detect objects as small as 5cm at a range of 1000 km, and yield a maximum effective range of 20,000 nautical miles. These parameters give it better sensitivity and slightly better range than the existing VHF space fence, while maintaining the same level of accuracy. Because this radar model operates in monostatic mode with some ability to track an object that penetrates its beam, preliminary object orbits can be determined from the tracks. A graphical portrayal of the notional new space fence is shown in Fig. 3.

Figure 3. Graphical portrayal of a notional new space fence

To increase turn around for the CBTA studies, most of the modules were set to the faster options. Because of this, the only parallelization needed was to run each radar site on its own processor. Atmospheric effects and noise sources were turned off. However, because the notional S-band space fence operated at 3.5 GHz, the degradation of signal-to-noise ratio due to the galactic noise is only around 0.02 dB, while the atmospheric effects are also small except at angles near the horizon. Therefore turning these options off has relatively little impact on the simulation. The new space fence antenna was modeled as an 800x800 dipole array. When the array is used to form a pencil beam, the result is extremely small and leads to low errors in positional accuracy. The radar cross section for objects in the simulation was taken from the Satellite Situation Report; where no RCS was available in the Satellite Situation Report, a default value of 0 dbsm was used.

Detection and accuracy of a given observation depends on the signal-to-noise ratio of the observation. The radar was tuned to have a skin track range of 5225 km, which allows it to detect a 20cm object at 2200km or a 5cm object at around 1000km. The minimum signal-to-noise ratio used for detection is 12 dB. During the simulation, the signal-to-noise ratio for a given target is used to determine the error in the observation, which is assumed to be Gaussian. A Cramer-Rao lower bound model for error from monopulse radar is used. This model results in angular accuracy of around 1/6th of the beamwidth for 12 dB targets at the limit of detection, and around 1/30th of a beamwidth for extremely strong targets.

During the simulation, the radar is scheduled to continuously observe the sky. A simple pulse-scheduling algorithm was assumed, which allowed the radar to update the position of every object within its field of view once a second. For every observation, the signal-to-noise ratio was computed taking into account radar parameters, such as power and antenna pattern, distance to target, and object radar cross section. This signal-to-noise ratio then determined the accuracy of the observation and if the object was detected. For all objects detected, the observation plus appropriate error was passed to the tracking algorithm.

A simple algorithm was used to assign observations to tracks. This algorithm supports only a single hypothesis and relies on proximity in space and time to assign observations to tracks. This has been found to be adequate as long as objects are well separated compared to the antenna beamwidth. Finally, the assembled tracks were passed to the CBTA in .xml format.

The 22.5° north-south arc width of the radar fans means that orbits are followed only for very small portions of their orbital arcs. For example, an object in a polar orbit at 700 km altitude will be tracked for only 2.21° of its orbital arc

and for only 36.4 seconds. Non-polar orbits will be tracked longer than polar orbits. Of course, orbits with inclinations less than the latitude of the space fence will never be observed.

4. RESULTS AND ALGORITHM SCALING

A series of CBTA runs were performed with varying amounts of orbital objects. The objects were randomly selected from the NORAD catalog and not all objects were observed. Objects for which a high quality joined track was produced are denoted as "discovered objects". The ratio of discovered objects to observed objects exceeds 97% in all runs. The results of these runs are summarized in Table 1. Of particular note, there were no joined tracks that were marked as valid that were not valid. A histogram of the joined track size for a typical run is shown in Fig 4. The CBTA algorithm is expected to scale quadratically as a worst case for large numbers of tracks. For the runs with moderate amounts of tracks that were performed to date, the scaling is effectively linear (Fig. 5).

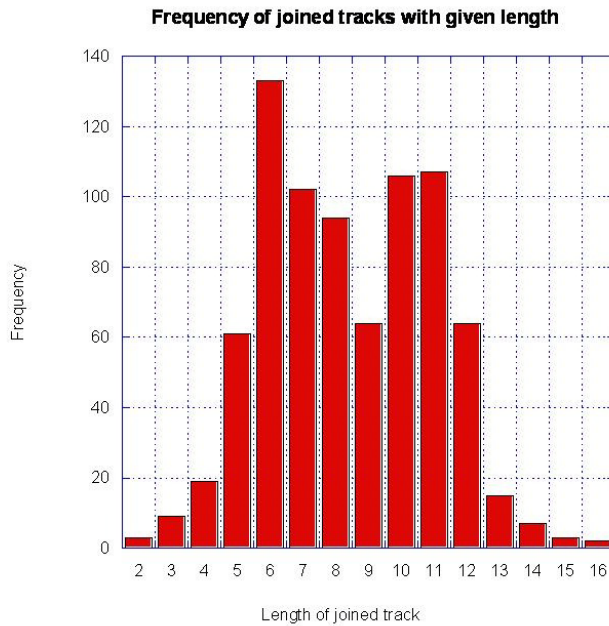| Total objects	Observed objects	Discovered objects	Total tracks	Joined tracks	Broken joined tracks	Invalid joined tracks	Joined tracks erroneously marked as valid	Standalone CPU seconds
1037	789	783	6659	6614	3	2	**0**	1851
1642	1081	1071	8723	8637	4	5	**0**	2036
1472	751	734	6332	6149	5	17	**0**	1258
3112	1832	1805	15055	14786	9	22	**0**	3386

Table 1. Results for 4 CTBA runs.

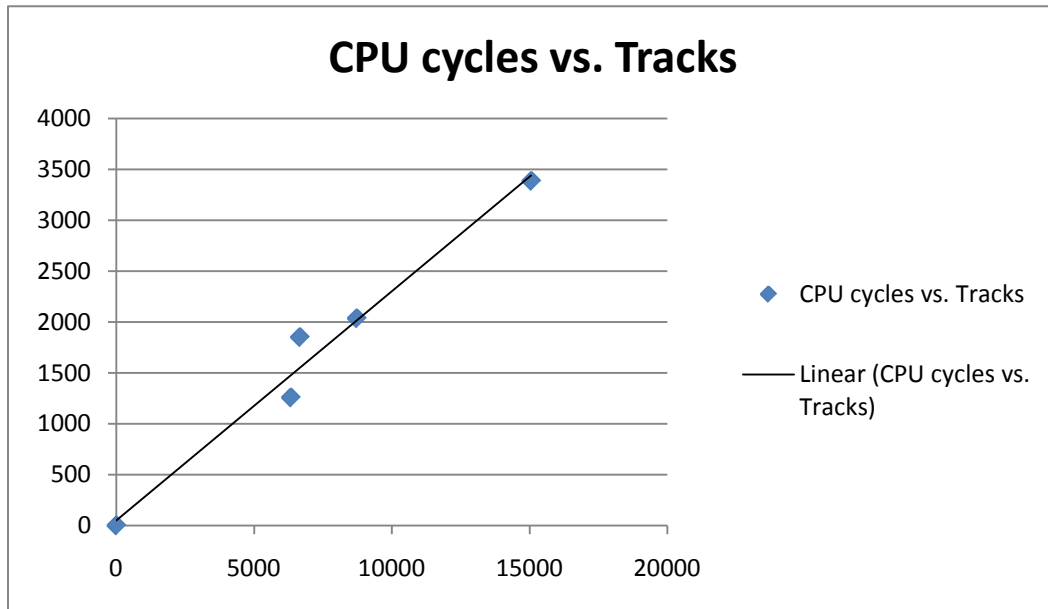Figure 4. Histogram of joined track length

Figure 5. Scaling trend for CBTA algorithm for moderate track quantities

5. STATUS AND FUTURE WORK

The next step in the development of the CBTA UCT process will be to implement an efficient parallelization scheme. The first dimension of parallelization will be to divide the pool of UCTs into orbital regimes. The granularity of orbital regime decomposition must be carefully tailored to maximize parallelization while minimizing the occurrence of tracks that belong to a single object being assigned to different UCT pools. Further parallelization may be obtained through the process of sequentially assigning pool tracks to different processors, with all processors operating on a common list. With an efficient parallelization scheme in place, we plan to run simulations with hundreds of thousands of UCTs, consistent with the observations that would result from a new Space Fence. We will then explore the efficiency of the process in terms of processing UCTs and cataloguing new objects.

6. ACKNOWLEDGEMENTS

7. REFERENCES

1. *NASA Orbital Debris Program Office web site:http://orbitaldebris.jsc.nasa.gov/faqs.html*

2.   *Covariance-based Uncorrelated Track Association* by Keric Hill, Kyle T. Alfriend, and Chris Sabol Paper AIAA 2008-7211 presented at the August 2008 AIAA/AAS Astrodynamics Specialist Conference and Exhibit in Honolulu, Hawaii

3. *Covariance-based Track Association Archival Report – DRAFT* by Keric Hill, April 16, 2010, Pacific Defense Solutions LLC, 1300 N. Holopono St., Suite 116, Kihei, HI 96753.

4. NASA Technical Note D-5368, *A Unified Form Of Lambert's Theorem* by E. R. Lancaster and R. C. Blanchard, Goddard Space Flight Center, Greenbelt, Md, September, 1969.