

High-Performance Computer Modeling of the Cosmos-Iridium Collision

Scot S. Olivier, Kem Cook, Ben Fasenfest, David Jefferson, Ming Jiang, Jim Leek, Joanne Levatin, Sergei Nikolaev, Alex Pertica, Don Phillion, Keo Springer, Wim De Vries

Lawrence Livermore National Laboratory

ABSTRACT

This paper describes the application of a new, integrated modeling and simulation framework, encompassing the space situational awareness (SSA) enterprise, to the recent Cosmos-Iridium collision. This framework is based on a flexible, scalable architecture to enable efficient simulation of the current SSA enterprise, and to accommodate future advancements in SSA systems. In particular, the code is designed to take advantage of massively parallel, high-performance computer systems available, for example, at Lawrence Livermore National Laboratory. We will describe the application of this framework to the recent collision of the Cosmos and Iridium satellites, including (1) detailed hydrodynamic modeling of the satellite collision and resulting debris generation, (2) orbital propagation of the simulated debris and analysis of the increased risk to other satellites (3) calculation of the radar and optical signatures of the simulated debris and modeling of debris detection with space surveillance radar and optical systems (4) determination of simulated debris orbits from modeled space surveillance observations and analysis of the resulting orbital accuracy, (5) comparison of these modeling and simulation results with Space Surveillance Network observations. We will also discuss the use of this integrated modeling and simulation framework to analyze the risks and consequences of future satellite collisions and to assess strategies for mitigating or avoiding future incidents, including the addition of new sensor systems, used in conjunction with the Space Surveillance Network, for improving space situational awareness.

1. INTRODUCTION

Freedom of operation in space is crucial to U.S. interests. In FY08, based on multiple requests from U.S. Defense, Intelligence and Congressional leadership, the DOE National Nuclear Security Administration (NNSA) Laboratories, Los Alamos, Sandia and Livermore, working with the Air Force Research Lab, began to investigate new ways to apply forefront science and technology to improve our country's capacity for detecting and monitoring threats to U.S. space operations. Our goal is to help catalyze the transformation of U.S. space situational awareness (SSA) by developing and demonstrating advanced capabilities to enable a new, real-time, net-centric paradigm for SSA.

As part of this program, in early 2008, a team of over a dozen LLNL scientists and engineers began working on the implementation of a comprehensive framework for modeling, simulation and visualization of space situational awareness (SSA). This work, supported by internal LLNL funding, has resulted in the successful establishment of a new, integrated modeling and simulation framework, encompassing the SSA enterprise, for quantitatively assessing the benefit of specific sensor systems, technologies and data analysis techniques: the Testbed Environment for Space Situational Awareness (TESSA). A paper [1] describing the initial implementation of TESSA was published in the proceedings of the September 2008 Advanced Maui Optical and Space Surveillance Technologies Conference.

TESSA is based on a flexible, scalable architecture, exploiting LLNL high-performance computing capabilities, to enable efficient, physics-based simulation of the current SSA enterprise, and to accommodate future advancements in SSA systems. TESSA includes hydrodynamic models of satellite intercept and debris generation, orbital propagation algorithms, radar cross section calculations, optical brightness calculations, radar system models, optical system models, object detection algorithms, orbit determination algorithms, simulation analysis and visualization tools. By effectively linking these detailed modules together, using an efficient "parallel discrete event simulation" architecture, TESSA provides a fundamentally new capability for the country that can be used to drive development of game-changing advances in SSA through dramatic increases in effective utilization of existing data sources and coherent planning for new sensor systems tuned to specific threat scenarios.

2. COSMOS-IRIDIUM COLLISION ANALYSIS

2.1 SUMMARY

An extraordinary space event occurred on February 10, 2009 when the defunct Russian satellite, Cosmos 2251, and the operational U.S. communications satellite, Iridium 33, collided over Siberia. These two satellites hit each other at a closing velocity of 11.6 km/s, resulting in large amounts of debris dispersed along the pre-collision orbital trajectories of both satellites, creating an increased hazard for other satellites.

Shortly after the collision, the SSA team at LLNL began to analyze the event, which provided a serendipitous opportunity to apply TESSA to a timely and interesting problem, to evaluate the current capabilities of TESSA and to identify areas for future development.

The analysis of the Cosmos-Iridium collision began with a hydrodynamic simulation of the impact in order to generate a detailed description of the physical and orbital characteristics of the debris that resulted from the collision. The orbital position of the debris was then calculated as a function of time based on the velocity vectors assigned by the hydrodynamic simulation to individual pieces of debris. Next, a simulation of radar and optical telescope detection of debris pieces over a 24 hour period was performed. Individual detections were associated to produce tracks for the debris pieces and an orbit determination algorithm was applied to these tracks. This procedure resulted in an estimate of the extent to which our postulated sensor suite could “recover” orbital debris tracks. Finally, a conjunction analysis was performed to predict the extent to which the debris from this collision represented a hazard for the remaining 65 active Iridium satellites.

The results of this exercise, obtained within several days of the collision, predicted that the debris from the impact produced an increase of 50% in the collision risk for the Iridium constellation, compared with the collision risk with objects already being tracked before the collision. This prediction was consistent with results that were computed several weeks after the collision using observations of the debris from the USAF Joint Space Operations Center. Details of the separate simulation components and the analysis processes, along with the visualization of the results, are described in Sections 2.2-2.8. An overview of the TESSA architecture used for this analysis is given in Section 2.9.

Future work will concentrate on addressing the advanced research topics necessary to enable effective utilization of TESSA simulation and modeling for providing actionable SSA, incorporating new capabilities into TESSA including advanced features for real-time analysis of data from existing sensors, and utilizing TESSA to inform the design new sensors optimized for specific threats to space operations. Further details of future work are discussed in Section 3.

2.2 DEBRIS GENERATION MODELS

Hypervelocity impact modeling for debris generation involves many complex phenomena that span several time- and length-scales. Advanced hydrodynamic codes capable of handling large deformations and fracture, statistical material failure models, massively parallel computers, and detailed satellite mesh models enable debris generation modeling. It is important to improve the understanding of the size and velocity of collision debris through detailed modeling, as it poses hazards to existing and planned space assets. LLNL has developed modeling tools that can predict the shape and velocity vector of the debris resulting from a hypervelocity collision. These tools have been extensively validated over the course of more than a decade with a variety of hypervelocity impact experiments using both sled tracks and light gas guns.

The current version of TESSA, used for the analysis of the Cosmos-Iridium collision, uses the explicit hydrodynamics code, ParaDyn (DYNA3D in parallel), enriched with smooth particle hydrodynamics (SPH) in order to simulate hypervelocity collision events. LLNL’s massively parallel, high-performance computer, Hera, is currently used for these simulations. An anisotropic, statistical fracture model, MOSSFRAC, is used to model the impact-induced fragmentation of the satellites. DFRAG (DYNA3D FRAGment post-processor) is used to extract detailed debris state vector information from collision simulations for orbital propagation, conjunction analysis, radar and optical signature modeling, as well as visualization.

For the initial analysis of the Cosmos-Iridium collision, detailed information on satellite geometry was not immediately available, so mesh models were based on estimates of overall geometry and mass. The Iridium satellite geometry consisted of a 2.0 meter long by 0.6 meter wide triangular bus with two 0.6 meter wide by 3.5 meter long solar panels (8.0 meter tip-to-tip) and three 1.2 meter long by 0.6 meter wide communication panels. The estimated and mesh model masses were 556 kg and 547 kg, respectively. The Cosmos satellite geometry consisted of a 2.0 meter diameter by 3.0 meter long cylindrical bus with a 5 meter long gravity boom. The estimated and mesh model masses were 900 kg and 920 kg, respectively. While the collision closing speed was known, the engagement geometry of the colliding satellite was not initially well understood. Therefore, several engagement geometries representing possible collision scenarios were considered. A full Iridium-Cosmos overlap scenario served as the upper bound for collision damage and momentum transfer, while a Cosmos boom on Iridium solar panel served as the lower bound. A number of intermediate overlap cases were also considered. A visualization of two of the collision simulations is shown in Fig. 1.




Fig. 1. Starting with high-fidelity solid models and calculated engagement conditions, LLNL’s PARADYN hydrodynamics computer code is used to simulate the hypervelocity collision of the two satellites. PARADYN post-processing routines then output a physical description and a state vector for thousands of debris pieces generated in the collision

2.3 DEBRIS PROPAGATION METHODS

Once debris is generated, the next step in our modeling process is to propagate the debris through space based on the location of the intercept, the orientation of the satellites and the debris velocity vector magnitude and direction. TESSA currently includes orbital propagation algorithms based on Kelso’s November 2007 SGP4 analytic propagator, as well as a force model. The simplified general perturbation (SGP) method is based upon the Gaussian or Lagrange Variation of Parameters theory [2, 3]. SGP4 includes the sun and moon third body perturbations. Atmospheric drag is modeled using the inverse ballistic coefficient. Our implementation of the SGP4 code has been validated by extensive comparison to the commercial Satellite Tool Kit (STK) from Analytical Graphics Inc. (AGI).

For higher accuracy, we have also implemented a force model. The force model numerically propagates the orbit taking into account all the important forces using a high-order geopotential model such as the 300x300 EGM96 model, sun and moon third body forces, solar radiation pressure, and atmospheric drag. The atmospheric drag models that have been implemented in the LLNL force model are NRL MSISE-00, NRL MSISE-90, GOST84, GOST2004, Harris-Priester, and Jacchia Roberts 1971. There are also the gravitational perturbations caused by the “solid” body earth tides and the ocean tides caused by the sun and moon. The LLNL force model uses time conversion and polar wander tables downloaded from the US Naval Observatory. It uses the 1980 IAU 106 term Earth nutation model and the 1976 IAU earth precession model [4, 5]. The IAU 2000A and IAU 2000B models [6] are required for better than milli-arcsecond accuracy, but these corrections are also available as dPSI and dEPS tables, which the LLNL force model uses. The LLNL force model has been validated by extensive comparison to AGI’s STK HPOP (High Precision Orbital Propagator).

A visualization of the debris propagation for the Cosmos-Iridium collision is shown in Fig. 2.




Fig. 2. The hydrodynamic collision simulation provides state vectors for each individual piece of debris. This allows the debris positions to be propagated forward in their orbits. Some of the debris quickly re-enters the atmosphere (orange dots).

2.4 RADAR MODELS

Radar simulations for TESSA are composed of three main segments. First, the radars to be modeled must be chosen and their performance parameters and locations entered into the TESSA configuration file. Next, the radar cross sections (RCS) of debris chunks must be computed from the results of the hydrodynamics code. Finally, the detection process is modeled, using the radar scheduler and radar pipelines to produce simulated radar track output.

A selection of radar sensors from the Space Surveillance Network and its affiliated sensors were chosen for the Cosmos-Iridium scenario. These radars include a mix of large phased-array antenna radars and mechanically steerable dish antenna radars. The phased-array radars used were the PAVE PAWS, Ballistic Missile Early Warning System (BMEWS), MPS-39 and FPS-85 radars. The Altair radar at Kwajalein was modeled as an example of a mechanically steerable radar. The three PAVE-PAWS and two BMEWS radars modeled are all multi-face UHF phased-array radars. Each face of the antenna covers an azimuth angle of 120 degrees. The PAVE-PAWS radars at Cape Cod Air Force Station, Beale Air Force Base, and Clear Air Force Station, as well as the BMEWS radar at Thule Air Base have two faces, giving them 240 degree coverage. The BMEWS radar at RAF Fylingdales has three faces, giving it full 360 degree coverage. These radars have an effective range of around 3,000 miles. The MPS-39 radar is a multiple object tracking radar, capable of tracking up to 40 objects within a 60 degree cone. It consists of a G-band phased array mounted on a mechanically steerable pedestal. It can track a one square meter object out to almost 500 nautical miles with range accuracy of a few meters and cross range accuracy around a hundredth of a degree. It is located at Vandenberg Air Force Base and typically used to support space launch activities. It was included in the simulation because of its high accuracy and its unique combination of mechanical steering/phased array technology. The final phased-array radar modeled was the FPS-85 radar, located at Eglin Air Force Base. This UHF radar has 45 degrees of azimuth coverage and can track 200 satellites simultaneously, with a maximum range of 22,000 nautical miles. ALTAIR, a UHF steerable dish radar located on the Kwajalein Atoll in the Pacific, was chosen for the simulation due to its location near the equator. The ALTAIR radar has a range that extends well into geosynchronous orbit and can track up to 32 targets simultaneously.

Once the radars to be modeled are chosen, the next step is to preprocess the debris chunks produced by the Dyna3D hydrodynamics code to determine their RCS. The RCS is a measure of how much of the incident power on an object is reflected back toward a transmitter. The RCS is the critical parameter for determining how visible an object is to radar; in general it depends on object size, shape, orientation, material, and frequency. A single object may have a different RCS for each incident angle. Because the orientation of the debris is not currently tracked within TESSA (although this is planned for the future), a sampling of around 1500 evenly spaced orientations are used to determine a median RCS.

Two types of debris objects are output from Dyna3D. The first type is large chunks, which were large enough to take up several mesh elements in the hydrodynamics simulation and are therefore output as a triangular surface mesh. The second type is very small chunks, which are output as only an equivalent volume. Different techniques

are used to compute the RCS of each type of object. The largest chunks, which consist of tens or hundreds of thousands of elements and are many wavelengths across, are solved using a physical optics approach. This approach assumes that the current excited on any facet visible to the incident wave is proportional to the tangential magnetic field incident on it. These excited currents are computed and then integrated against a Green's function to produce the scattered field and radar cross section. This technique is fairly accurate for large objects and is very fast. However, it is not applicable to small objects or objects near a wavelength in size. Medium-sized chunks, whose dimensions are on the order of a tenth of a wavelength to a few wavelengths, are solved by an integral equation method. EIGER (Electromagnetic Interactions GENEralized), developed at Lawrence Livermore National Laboratory in collaboration with Sandia National Laboratory and the University of Houston, solves the integral version of Maxwell's equations to find the induced currents on a debris chunk and the backscattered field and RCS from those currents. EIGER is extremely accurate for chunks in the resonant region. Integral equations solvers are fairly time consuming to run, a problem which EIGER solves through massive parallelization. An example of the results of an EIGER calculation for a debris object is shown in Fig. 3. The smallest chunks which have no mesh associated with them and are much less than a wavelength in size are assigned an RCS based on Rayleigh scattering theory. This approximation assumes the volume is distributed as a sphere and hence the RCS is the same for all incident angles. This approximation becomes more and more accurate as the size of the object decreases.

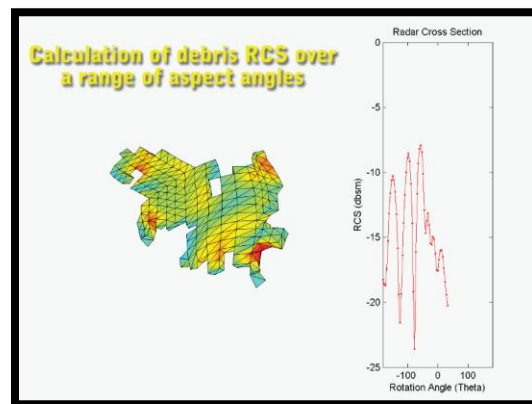


Fig. 3. The radar cross section (RCS) of each fragment is calculated over a dense sampling of aspect angles for use in modeling ground-based radar sensing of the debris.

Upon completing the RCS preprocessing, the radar detection simulation can be launched. The first task of the radar detection simulation is to set up the radar schedules. Each radar has its own scheduler, which takes as inputs the radar location, range, and coverage. For the Cosmos-Iridium collision simulation, the scheduler was given the orbit of the Cosmos satellite. This orbit was propagated forward for the duration of the simulation to determine when the satellite would have been in the field of view of each radar if there had not been a collision. A padding window of ten minutes on either side was added to each schedule to capture some of the faster or slower debris. The radars were then scheduled to look towards the satellite during these times. The schedule derived from this procedure was a toy model representing a simple approximation of how the radar systems were actually scheduled shortly after the Cosmos-Iridium collision. Schedule optimization is an important problem, which we plan to address in future work (see Section 3).

After schedules have been generated, they are handed to the radar pipelines. At minimum, a separate radar pipeline is used for each radar. If an additional speed increase is needed, each radar can be spread across multiple pipelines and schedules passed to them in a round-robin manner.

Each radar pipeline loops over the times given to it in its schedule. The time step used can be varied per radar to model differences in pulse repetition frequencies. At each time step, three main operations are performed: the orbit determination subroutine is called to find the satellites or debris chunks that are within the radar's beam, the objects within the beam are tested for detection, and the detected objects are passed to a tracking routine.

Radar detection is processed once the objects within the beam are determined. Objects are lumped together in range, azimuth, and elevation bins. Radar cross sections (RCS) of debris are determined by sampling the cumulative probability distribution of the RCS values pre-computed for each debris piece at a few thousand different angles. For catalog objects, a fixed RCS value is used. Objects within the same bin have their RCS summed to model incoherent reflection of power. The radar-range equation is used along with the distance and RCS of each object and the known sensitivity of the radar to determine if the radar detected the object.

Objects that are detected are passed on to the tracking algorithm. Currently, the algorithm used to determine if consecutive observations form a track relies only on the spatial proximity and velocity of observations. If a new observation is near the end of a current track in roughly the same direction as the track was moving, it is added to the track. If not, it starts a new track. The track determination could be improved by using RCS values and standard orbital motions to help identify to which track an observation belonged. Only a finite number of concurrent tracks are allowed, based on the number of simultaneous tracked objects possible for the radar being modeled. Once the simulation time steps are completed, all tracks containing at least two observations and spanning at least two seconds are written to disk and passed on to the orbit determination portion of TESSA.

2.5 OPTICAL MODELS

A part of TESSA, the Optical Detection Pipeline handles optical image simulations. The pipeline consists of two parts: image generation and image analysis. In the image generation stage, the pipeline models a variety of effects, including atmospheric light scattering, stellar component (including unresolved stars), tumbling motions of debris, light propagation through the telescope optics, tracking mode, charge transfer artifacts, weather conditions, various noise sources, etc. As a result, it produces a realistic-looking optical image of the star field and satellites (if they happen to traverse the field-of-view). The images, written as FITS (Flexible Image Transport System) are shipped to the next stage of the pipeline for processing. During the processing stage, the raw pixel values are analyzed to extract sources (stars and satellites), and fit their observed fluxes to recover astrometric positions. Combined with the knowledge of the WCS (World Coordinate System) of the image and the timing information, these positions can be converted into orbital data.

Due to slew time requirements on the optical instruments and the limitations on the accuracy of the optical imagery, the Optical Detection Pipeline has limited applicability to LEO objects (most of the astrometric data for LEO orbit determination comes from radar observations). Therefore, it was not directly relevant for modeling the Cosmos-Iridium collision. Nevertheless, to showcase the capabilities of the optical image simulations, the Optical Detection Pipeline was used to generate a sequence of image frames from an SSN system, which were combined into a video sequence. The video showed a train of 8 Cosmos fragments on close orbits, displaying varying optical brightness of the fragments as they rotated. A visualization of this optical image simulation is shown in Fig. 4.

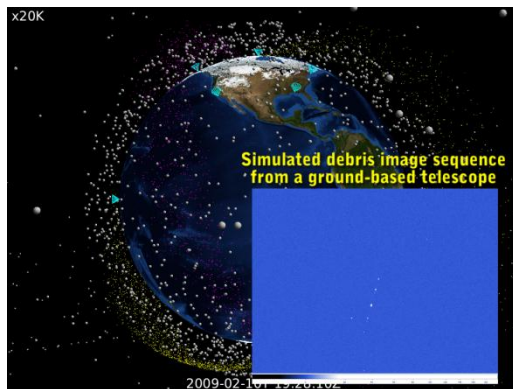


Fig. 4. Results of optical image simulation for model debris from the Cosmos-Iridium collision.

2.6 ORBIT DETERMINATION METHODS

Many first orbit determination algorithms have been implemented that either use line-of-sight data or use position data. All of these algorithms obtain the ideal two-body solution.

First orbit determination from position data is done using one of three algorithms. Lambert's method requires two positions, an elapsed time, and the number of completed revolutions. It is valid for hyperbolic as well as elliptic orbits. Gibb's method [7] uses three widely spaced positions. Three points on a branch of a conic and the focus uniquely specify the conic. The time information is not used at all. It also is valid for hyperbolic as well as elliptic orbits. Herrick-Gibbs method [7] uses three finely spaced positions. Taylor series expansions in the position vectors as a function of time are done. All three of these methods are very robust and require no initial guess.

First orbit determination from three line-of-sight measurements at known times has been implemented using Gauss's method [5], Laplace's method [7], Escobal's Double R method [7], Herget's method, and Gooding's method. All of these methods require an initial guess except for Laplace's method. Laplace's method uses three finely spaced lines of sight at known times. Laplace's method works well for orbital platforms as well as for earth-based sites and is robust. It can also use more than three closely spaced line of sight measurements at known times. In the TESSA implementation of first orbit determination, no single method is exclusively relied upon. Rather every method that might work is successively tried as needed. The method that is most likely to work is tried first. In the case of finely spaced line-of-sight data, Laplace's method is always tried first. In the case of widely spaced line of sight data, Gauss's method is tried first.

After a first orbit has been determined, this orbit is immediately refined using batch least squares [5]. Batch least squares can also be done on measurements obtained on multiple passes made by a tracking network. In order to accurately know where the orbiting object will be many orbits in the future, the orbital period must be determined with extremely high precision. This can never be determined with sufficient accuracy from measurements made on one arc segment on a single pass. The Extended Kalman filter [5] can be used to detect maneuvers and model unknown forces. Process noise is added to open the process window and to prevent the filter from becoming smug. This process noise gives the filter a "fading" memory.

For each track, association to a known catalog object is attempted using the raw observations. The initial orbit determination and the association are completely independent. First the catalog object that matches the raw observations in a least squares sense is determined. This is done taking into account the measurement 1 sigma errors in the weightings. Then this association is accepted as valid only if the fitting errors are less than some multiple of the observational 1 sigma errors. The next closest object is also determined. Uncorrelated tracks are associated using a covariance approach.

Comparison of orbits estimated from the simulated observations to the original simulated orbits can be used to assess the accuracy of the orbit determination methods. A visualization of this comparison is shown in Fig. 5.




Fig. 5. Sensor-derived orbits are compared to "truth" orbits, providing information on the accuracy and efficiency with which ground-based sensors can characterize the debris cloud.

2.7 CONJUNCTION ANALYSIS METHODS

In general, conjunction analysis can be divided into two distinct cases. The first is a scenario in which one wants to assess the odds of whether object A will hit object B, and the second, more general case, concerns the odds of a collision occurring between groups of objects. The requirements placed on precise knowledge of the orbital elements of the involved satellites are quite different. The first case can only be adequately addressed when one not only has very accurate orbit determinations of both objects involved (accurate to a few meters), but also a good handle on the uncertainties involved for each of the orbital elements (usually based on the covariance matrix). Unfortunately, meeting these quality conditions is rather hard, and this explains why satellite operators, in general, do not have actionable knowledge about upcoming close conjunctions.

The orbital fidelity requirements for the second case are more relaxed, and are within the accuracy of the low-accuracy Two Line Element (TLE) catalog. Even though the intrinsic positional accuracy of the low-accuracy Standard General Perturbation (SGP4) propagator is not better than about 1 km, its uncertainty is uncorrelated with the actual position. As such, it can be used to estimate the rate at which objects come close to each other, provided the separation is larger than the SGP4 accuracy. For instance, on average, the full Iridium constellation had about 600 encounters within 10 km each day prior to the Iridium 33 - Cosmos 2251 collision. This rate determination of about 600 close conjunctions per day does not depend on high accuracy orbital determinations.

If we define the encounter plane to be the plane perpendicular to the orbit of the object of interest, a second object that has a closest approach to the first one will intersect this plane at a certain location. Its orbital path at that time will also be perpendicular to this plane. The distribution of these intersection points of a large number of conjunctions (either against one satellite over a multi-orbit period, or a constellation of satellites) turns out to be uniform. In other words, the density of these points is constant, and represents a constant fluence. Therefore, the collision rate can be estimated based on relative cross-sections in this encounter plane. For instance, assuming an effective satellite surface area of 10 square meters for an Iridium satellite, the relative area ratio is 3×10^{-8} for the 10 km conjunction cut-off. The daily rate of 600 per day then factors this into a 2×10^{-5} chance of a collision against the constellation per day (or 7×10^{-3} per year).

Using the relative surface area numbers to compute collision rates can easily be checked against, for instance, measured impact rates of small debris on in-orbit satellites specifically flown for these measurements, or impacts against solar panels of the Hubble Space Telescope. The NASA Orbital Debris Program Office has published extensive studies assessing the risks of exposed astronauts, and their numbers are consistent with ours for similarly sized debris. This approach to collision analysis based on actual orbit propagation, in contrast to NASA's zonal debris density approach, has the potential to provide a better understanding of collision rates due to the significant spatial debris density variations caused by the limited and quantized distribution of the orbital elements of actual debris – most of the larger debris can be traced back to specific satellites with their specific orbital regimes.

The implementation of the TESSA conjunction analysis code that was used for the Cosmos-Iridium collision analysis utilized the low-accuracy SGP4 propagator to calculate where a particular object was located in space at a certain time, given the orbital elements derived from the public NORAD TLE catalog. It performed the NxM possible conjunctions on the fly, and could therefore account for conjunction rate variations due to orbital evolution. We found that we could run these conjunctions for at least a few weeks beyond the TLE time stamp without trouble. Beyond that, it was necessary to update the TLEs for the objects of interest. It is also rather trivial to continuously use the most appropriate TLE set for cases that have already taken place (like the Iridium 33 - Cosmos 2251 collision). TLE sets are typically updated three times a day, though not for all objects in the catalog.

Furthermore, we are able to randomly generate orbital elements, either replacing the whole catalog, or adding, for instance, debris generated by a collision event. It is then rather trivial to calculate the relative increase in collision rate for a particular asset given the new debris environment. It also allows us to calculate the relative risk due to debris that is too small to be tracked by the SSN, but is still large enough to potentially cause catastrophic damage to a satellite (a size range between 1 cm and 10 cm).

The USAF Joint Space Operations Center has been able to identify many (approximately 1000) pieces of debris due to the Iridium - Cosmos collision. Using our framework and computational infrastructure, we are able to run full

Monte-Carlo case studies for a range of debris cloud configurations. Matching the observed conjunction rate of the observed debris to the modeled debris rate, we are able to constrain the actual collision geometry. A full body-on-body impact would produce too much high-velocity debris to match the observed conjunction rate, so a more likely scenario involves a small overlap impact (e.g., a gravity boom against a solar panel) and subsequent, late-term satellite breakup.

Results for our analysis of the Cosmos-Iridium collision are shown in Fig. 6. These results, obtained within a few days of the collision, predict that the debris from the impact produced an increase of 50% in the collision risk for the Iridium constellation, compared with the collision risk with objects already being tracked before the collision. This prediction was consistent with results that were computed several weeks after the collision using observations of the debris from the USAF Joint Space Operations Center.

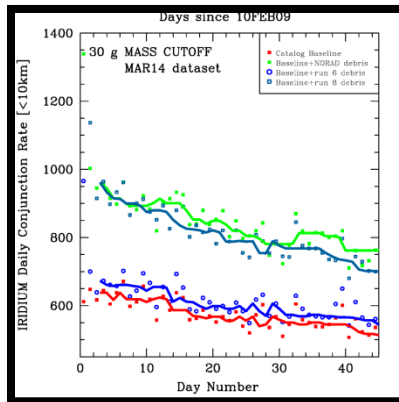


Fig. 6. The number of daily conjunctions to the Iridium constellation both before (red line) and after (green line) the collision, based on published data. This is compared to the estimation of the conjunction rate derived from our simulations for two postulated encounters; one that generates a large amount of debris (upper blue line) and one that generates a small amount of debris (lower blue line).

2.8 VISUALIZATION TOOLS

In order to develop an effective visualization system for the Cosmos-Iridium collision, we had to address several issues. First and foremost was the issue of managing large amounts of unstructured data that are generated from the debris simulation and orbital propagation. Next was the issue of how to visually convey the various physical interactions and phenomena that occur between the debris and their environment. Lastly was the issue of how to provide an interactive and collaborative environment for the physicists and engineers to evaluate and validate their results in real-time.

For the Cosmos-Iridium collision, there were several sources that contributed to the growth of data. The debris simulation produced a set of both large and small objects, called chunks and particles, respectively. Associated with the large chunks were actual surface mesh geometries that we had to process in order for them to be properly rendered. For each object in the space catalog and debris simulation, the orbital propagation produced a set of positions at 10 second intervals to represent the orbit of each object. Rather than using the traditional Two Line Elements (TLEs) to represent each object's position at each time step, we were able to reduce the data to only their positional values in Earth-Centered Earth-Fixed format. By using a specific binary file format, we drastically reduced the amount of data that need to be accessed in order to visualize the temporal evolution of the entire scene. Currently, we are developing data structures to efficiently represent the unstructured data in a multi-resolution format [8].

In our collision simulation, a debris piece is marked as being detected when it crosses the field-of-view of a radar station. A naïve strategy to visualize the detection result would be to change the representation of the debris at that time step in order to signify detection (e.g., change the debris to match that of the station color). However, this strategy does not work well in an animation setting where each time step is played back at a high frame rate, such as

QuickTime at 24 fps. In order to give each detection result more emphasis, we used the concept of persistence to highlight the detection for a certain amount of time after the actual detection. Using the size and color attributes, we provide a cross-fade (linear interpolation) between the debris' normal state and detected state. We have also successfully used the persistence concept to visually highlight the physical phenomena of debris objects that entered the Earth's atmosphere rather than orbiting around the Earth. We are in the process of developing strategies for visualizing the inherent uncertainties that are part of the collision simulation [9].

One of the goals for the visualization system is to provide a collaborative environment for the simulation scientists to evaluate the quality of their simulations. The ability to visualize and interact with the simulation results in real-time is the key to performing their evaluation. We have developed a graphical user interface that includes not only the Earth-centered visualization, but also various query and analysis capabilities that are statistical in nature. We are exploring techniques in information visualization that can present the query and analysis results in a more expressive and insightful fashion [10].

2.9 TESSA SIMULATION FRAMEWORK

Each of the advanced simulation and modeling techniques described above have been incorporated into the Testbed Environment for SSA (TESSA), a comprehensive SSA enterprise simulation framework, utilizing a parallel discrete event simulation architecture. TESSA is structured as a parallel discrete event simulation (PDES) primarily because the computation is dominated by simulation of the radar and optical sensors at the discrete moments when they make their observations of the sky. It is not, as one might initially surmise, dominated by the continuous computations required for orbital propagation. The simulator is not time-stepped because the simulation times at which the sensor observations occur are not statically predictable, and essentially never coincide. For such models time-stepping would be exceedingly wasteful. An event-driven architecture, which only executes at those simulation times and in those processes where the system state changes, is dramatically more efficient. The TESSA simulator, therefore, is a specially-designed (conservative) distributed discrete event simulation platform. This decision has important consequences, both for the performance of the TESSA model and for the process of developing it.

TESSA represents each sensor in the SSN as a separate logical process (LP), i.e., a unit of parallelism (in the language of discrete event simulation). Some of the sensors are so complex that they also have substantial internal parallelism as well, and so many of them are further decomposed into multiple processes (in the operating system sense), and may also be threaded to achieve further parallelism. This need for parallelism within LPs is an important reason we developed the TESSA simulator; other parallel discrete event packages available do not support events or LPs with internal parallelism. The output of the sensor LPs passes in pipeline fashion to other LPs that deal with reducing and interpreting the sensor data, identifying the orbiting objects observed, and perhaps cataloging them. Each pipeline stages executes in parallel with the others. Because each of the sensors is a separate LP, the sensor codes can easily be developed independently by different teams with different expertise. The code for each LP can run and be tested as a standalone model in its own right, because it does not have to link with the code for other LPs. LPs can even be written in different languages and still interoperate automatically in the TESSA simulator.

In event-driven simulations the way to achieve parallelism is (perhaps surprisingly) to allow some LPs to run ahead in simulation time while others lag behind, yet still maintaining the fundamental causality requirement that processes can only affect events in the simulation time future, and never in the past. There are two broad classes of parallel synchronization algorithms that can accomplish that, known as conservative and optimistic. The TESSA simulator is conservative, meaning that it does not resort to rollback in order to maintain the causality requirement.

The TESSA simulator is written in C++. It runs on top of Co-op, an LLNL componentization package that allows parallel components to be dynamically launched and to communicate via remote method invocation. Co-op is also what allows the various components of a TESSA simulation to be written in different languages (C, C++, Fortran 77 or 90, Python, or Java) through the use of Babel, LLNL's well-known language interoperability tool. TESSA event messages are encoded in XML and are transported from the sending LP to the receiving LP either through the file system or via Co-op remote method invocation. The simulator runs on any of LLNL's Linux or AIX clusters, and routinely runs TESSA simulations using between 200 and 300 cores for 6 to 12 hours, although there is no barrier to scaling up to at least thousands of cores when that becomes necessary.

3. STATUS AND FUTURE WORK

We have used the TESSA framework to analyze the February 10, 2009 collision of the defunct Russian satellite, Cosmos 2251, and the operational U.S. communications satellite, Iridium 33. This event provided a serendipitous opportunity to apply TESSA to a timely and interesting problem, to evaluate the current capabilities of TESSA and to identify areas for future development. The results of this exercise, obtained within several days of the collision, predicted that the debris from the impact produced an increase of 50% in the collision risk for the Iridium constellation, compared with the collision risk with objects already being tracked before the collision. This prediction was consistent with results that were computed several weeks after the collision using observations of the debris from the USAF Joint Space Operations Center.

In the future, we envision using TESSA in at least three specific ways to address space flight safety. First, in the case of future impact events resulting in substantial amounts of debris (for which the likelihood is presently significant on the time scale of a year), TESSA can be used to simulate the expected distribution functions of the debris and then to assess, based on the simulation results, whether any active satellites are threatened to the extent that they should be relocated, and furthermore, to determine how to more optimally schedule sensors to more quickly find and track the actual debris pieces. TESSA can also be used to interpret data from sensors to constrain the collision parameters, including properties of the colliding objects, when these are not well known initially. The key new capabilities needed for this application of TESSA to space flight safety are fast and accurate methods for simulating debris distribution functions [11], accurate and detailed models of the Space Surveillance Network sensors, a flexible parallel discrete event simulation architecture, which supports feedback for sensor scheduling, and new methods and algorithms for estimating collision parameters based on sensor data.

Second, in order to increase the fidelity of collision risk assessments, to the point where the probability of collision between individual objects at specific instances can be calculated to the level of probability required to justify repositioning an active satellite ($> \sim 1\%$ probability), TESSA can be used to support new methodologies involving the fusion of data from a combination of operational and existing, unexploited and/or new sensors along with innovative techniques for prioritized cueing of sensor observations and data analysis. For example, active satellites can be screened for close conjunctions with any other tracked object, and sensor observations and/or analysis of unexploited data can be prioritized based on the computed collision risk and the level of uncertainty, due to inaccurate knowledge of the relevant orbital parameters. The key new capabilities needed for this application of TESSA to space flight safety are advanced sensor models, new techniques for improving the accuracy of measured orbital parameters from relevant sensor data [12-14] and fast methods for computing collision risks along with rigorous uncertainty quantification [15-17].

Third, in order to develop effective strategies for detecting and tracking orbital objects that are too small to be routinely monitored with current operational sensors ($< \sim 10$ cm diameter), but are still large enough to cause significant damage if they were to collide with a satellite ($> \sim 1$ cm diameter), TESSA can be used to compare the performance of a range existing, planned and proposed sensors. Furthermore, in order to enable actionable collision prediction and avoidance for the smaller objects, TESSA can be used to guide the design of new sensors, scheduling algorithms and data fusion techniques optimized for this purpose. This application of TESSA requires further new capabilities along the lines of those described above, which are scaled to provide the required sensor sensitivity to observe smaller objects and to function effectively with the much larger population of smaller objects.

4. ACKNOWLEDGEMENTS

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

5. REFERENCES

1. Olivier SS, "A Simulation and Modeling Framework for Space Situational Awareness". Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, Wailea, Maui, Hawaii, 2008
2. Hoots F. R. and Roehrich, R. L., "Spacetrack Report No. 3: Models for Propagation of NORAD Element Sets", December 1980, Package Compiled by TS Kelso 31 December 1988.

3. Beutler, G., "Methods of Celestial Mechanics I: Physical, Mathematical, and Numerical Principles", Springer (2005), chapter 6.
4. Seidelmann, P. K., "1980 IAU Theory of Nutation; The Final Report of the IAU Working Group on Nutation", Kluwer Academic Publishers (1981), Provided by the NASA Astrophysics Data System.
5. Montenbruck, O. and Gill E., "Satellite Orbits: Models Methods Applications" Springer (2001)
6. McCarthy D. D. and Petit G., "IERS Conventions 2003", IERS Technical Note 2003
7. Vallado D. A., "Fundamentals of Astrodynamics and Applications Third Edition", Space Technology Library (2007)
8. Woodring, J., Chaoli Wang, and Han-Wei Shen, "High dimensional direct rendering of time-varying volumetric data", *Proceedings of IEEE Visualization 2003*, pp. 417-424, 2003.
9. Pang, A.T., Craig M. Wittenbrink, and Suresh K. Lodh, "Approaches to Uncertainty Visualization", *The Visual Computer*, 13:370-390, 1996.
10. Wong P.C. and Jim Thomas, "Visual Analytics", *IEEE Computer Graphics and Applications*, 24(5): 20-21, 2004
11. Grady, D., "Investigation of Explosively Driven Fragmentation of Metals - Two-Dimensional Fracture and Fragmentation of Metal Shells: Progress Report II", *UCRL-CR-152264 B522033*, February 2003
12. Ewe, W., L. Li, M. Leong, "Fast Solution of Mixed Dielectric/Conducting Scattering Problem Using Volume-Surface Adaptive Integral Method," *IEEE Transactions on Antennas and Propagation*, November 2004, 52, (11), pp. 3071-3077.
13. Seitzer, P., Abercromby, K., Barker, E., Rodriguez, H. "Optical Studies of Space Debris at GEO - Survey and Follow-up with Two Telescopes", Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, held in Wailea, Maui, Hawaii, September 12-15, 2007, Ed.: S. Ryan, p.E37
14. Bowman, B. R., W. Kent Tobiska, Frank A. Marcos, Cheryl Y. Huang, Chin S. Lin, and William J. Burke."A New Empirical Thermospheric Density Model JB2008 Using New Solar and Geomagnetic Indices" AIAA/AAS Astrodynamics Specialist Conference 18-21 August 2008, Honolulu, Hawaii AIAA 2008-6438
15. Chan, F. K. 1997, "Collision Probability Analyses for Earth-Orbiting Satellites"
16. Alfriend, K. T., Akella, M., Frisbee, J., Foster, J. L., Lee, D-J, and Wilkins, M. 1999, "Probability of Collision Error Analysis", in: *Space Debris*, 1, 21
17. Alfano, S. 2007, "Review of Conjunction Probability Methods for Short-Term Encounters", in: *AIAA Space Flight Mechanics Meeting*, Feb 2007