MATURING METHODS FOR IN SITU DEBRIS STRIKE DETECTION USING ON-ORBIT TELEMETRY

Kristia Kyerra Harris[(1)], Dr. Anne Aryadne Bennett[(2)], Stanley Smeltzer[(3)], Gayle Zheng[(4)]

[(1)] Northrop Grumman Systems Corporation, 45101 Warp Drive, Dulles, VA, 20166, USA, Email: kristia.harris@ngc.com
[(2)] Northrop Grumman Systems Corporation, 45101 Warp Drive, Dulles, VA, 20166, USA, Email: anne.bennett@ngc.com
[(3)] Northrop Grumman Systems Corporation, 45101 Warp Drive, Dulles, VA, 20166, USA, Email: stanley.smeltzer@ngc.com
[(4)] Northrop Grumman Systems Corporation, 45101 Warp Drive, Dulles, VA, 20166, USA, Email: gayle.zheng@ngc.com

ABSTRACT

With the acceleration in space activity, the need for accurate and reliable models for orbital debris risks becomes increasingly critical to the safety of current and future operational satellites. This study investigates methods for using typical spacecraft telemetry to improve the understanding of risks from hazardous non-trackable debris.

Previously, methods developed for identifying indications of a minor debris strike have worked well on simulated telemetry; however, these methods are challenging to apply effectively to on-orbit satellite telemetry. This follow-on research aims to refine the prior methods and develop new scalable techniques to progress toward a capability to collect *in situ* debris data from a larger selection of satellites. Techniques used to detect debris strikes include assessing the spacecraft for abrupt, subtle attitude and momentum changes, including identifying these features in noisy data and filtering out expected motions from day-to-day spacecraft activities. This paper presents interim results of this ongoing study.

1 INTRODUCTION

As there continues to be an acceleration in the amount of space activity, including the population of trackable fragmentation debris more than doubling in the past 25 years [1], accurate and reliable models for orbital debris risks are becoming increasingly essential to ensure the safeguarding of forthcoming operational satellites. Due to the nature of debris causing events, debris of varying sizes are produced from even a single event. The larger, trackable debris population is of lower concern due to the ability to predict and manoeuvre to avoid this debris; however, small debris currently remains non-trackable and thus the associated risk is difficult to mitigate.

While many debris strikes cause minimal operational effects, this population of debris has the potential to be fatal to the struck spacecraft. A piece of debris smaller than 1 centimetre (cm) has the potential to cause serious damage which was demonstrated in [2] and can be seen in Fig. 1. Currently tracking capabilities of smaller debris are very limited. The Space Surveillance Network tracks debris down to around 10 cm in low earth orbit (LEO) and 70 cm in geostationary orbit (GEO) [3], and the National Aeronautics and Space Administration's (NASA) Orbital Debris Engineering Model (ORDEM) version 3 (ORDEM3.0) only models debris down to 10 cm in GEO, with the GEO debris models based on measurements of pieces 30 cm and larger [4]. Per [5], as of March 2025 there are approximately 40,500 debris objects greater than 10 cm in orbit, but there are approximately 1,100,000 objects between 1 cm and 10 cm. This equates to less than 10% of this hazardous non-trackable (HNT) debris category being tracked.

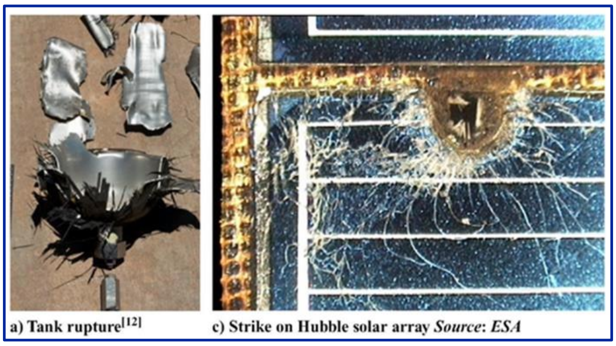

Figure 1. The impacts effects from small debris. The tank in a) was impacted by a 2 mm aluminium sphere. The hole in c) is 2.5 mm [2].

While it can be challenging to attribute an anomaly definitively to a debris strike, there have been several instances where the anomalous event experienced by a spacecraft may be caused by a debris strike. Some examples of anomalous events include Telkom-1, AMOS-5, NSS-806, and Intelsat-29e [6]. In 2017, the NASA Engineering and Safety Center (NESC) produced a report on a study conducted using similar events to compare the observed anomalies to the correlating ORDEM3.0 predictions and the typical risk assessment

Proc. 9th European Conference on Space Debris, Bonn, Germany, 1–4 April 2025, published by the ESA Space Debris Office
Editors: S. Lemmens, T. Flohrer & F. Schmitz, (http://conference.sdo.esoc.esa.int, April 2025)

methods. The NESC report found a low correlation between the on-orbit events experienced and the ORDEM3.0 predictions, with ORDEM3.0 predicting a significantly higher risk of failures and perturbations than LEO systems have experienced [7].

One of the recommendations put forth by the NESC report is to collect data on satellite orbital perturbations and momentum changes [4]. This recommendation led to an effort for using existing satellites as *in situ* sensing platforms to improve knowledge about the local debris environment. Detection of the momentum perturbations produced by minor debris strikes and the associated data could be used to tune debris environment and risk models, refining estimates of uncertain parameters [6].

To reduce the strain that carrying additional dedicated sensors for the purpose of detecting strikes has on satellites, an alternative solution is to develop ground-based methods to process the standard, existing telemetry collected from operational spacecraft. This telemetry can then be used to develop algorithms that can detect minor momentum perturbations that are generally too subtle to be observed in the normal course of operations [6].

This type of concept has been explored by multiple researchers including ExoAnalytic Solutions, Fraunhofer Ernst Mach-Institut, and the Institute for Defense Analyses. ExoAnalytic Solutions worked to detect momentum impulse transfer events on GEO objects using their global telescope network [8]. The Fraunhofer Ernst Mach-Institut research assesses the momentum transferred during small hypervelocity impacts on spacecraft materials via modelling and test, followed by simulating the dynamic response to these disturbances [9]. The Institute for Defense Analyses, intending to catalogue and incorporate results into debris models, discusses methods to monitor satellite perturbations [10].

Bennett et al. developed methods to identify subtle perturbations in spacecraft angular momentum [11] and applied these techniques to NASA spacecraft [12], successfully identifying instances of unexpected abrupt angular momentum perturbations [6]. The concepts developed in the work by Bennett et al. simulated the spacecraft dynamics to model the effect of debris strikes on attitude control system telemetry to identify subtle debris strikes applied to the spacecraft dynamics simulation. However, when Bennett et al. applied the methods to on-orbit telemetry it proved to be challenging to apply effectively, with almost all spacecraft experiencing an overwhelming amount of noise and other activities which obfuscate the strikes. Some of these challenges are thought to be caused by unidentified perturbations seen in the on-orbit telemetry, unmodelled momentum changes, and the estimation of inertia.

This follow-on research aims to refine the prior methods employed by Bennett et al. and develop new techniques to progress toward a capability to collect *in situ* debris data from active, on-orbit satellites. The approach of this study starts with on-orbit telemetry and develops methods to process that telemetry. This includes a focus on cleaning the telemetry, anonymizing the spacecraft, and developing the associated methods based on the nature of on-orbit telemetry.

The methods developed in this research are based on using typical data, such as attitude control system telemetry from on-orbit spacecraft, to create scalable techniques that expand the ability to use a wide variety of existing telemetry data as a diverse dataset. Techniques used to detect subtle debris strikes include assessing the spacecraft momentum for abrupt, subtle changes, including identifying these features in noisy data, and filtering out expected perturbations from day-to-day spacecraft activities. Methods employed include Matched Filters and Comb Filters, Savitsky-Golay smoothing, and sequential probability ratio tests.

This research underscores the importance of collaboration and data-sharing throughout the community. To promote data-sharing among various entities, a key focus of this research is anonymization of data. Multiple layers of anonymization are applied throughout this study to ensure that no specific spacecraft or proprietary information can be identified through the findings of this study. Developing methods to anonymize spacecraft data provides confidence that scientific data can be shared without adverse effects to the satellite owners and facilitates the inclusion of data from non-NASA spacecraft to increase available datasets.

The eventual goal of this line of effort is to assess the prevalence of debris strikes and compare the findings to predictions made using debris environment models, providing an additional dataset to validate and tune debris risk assessment models and methods. Through improving the debris risk assessment precision and confidence, we can work to manage the risks posed to current and future space operations, contributing to the long-term safety of orbital environments.

2 TECHNICAL APPROACH

The technical approach for this study leverages lessons learned in prior work. Prior work began developing methods for debris strike detection with simulation data and then was applied to on-orbit telemetry. The simulation data was found to be very "clean," allowing strike detection algorithms to show strong responses to subtle strike features, while the on-orbit telemetry contained many unexpected perturbations and idiosyncrasies which posed a major challenge to obtaining results from the debris strike detection algorithms [12]. Therefore, this work focuses on developing methods to work with the complications of on-orbit telemetry by advancing the data processing architecture and telemetry cleaning.

One major challenge in applying debris strike detection methods to on-orbit telemetry in previous studies was hypothesized to be due to the uncertainty in the spacecraft inertia. Normal operations of a spacecraft are generally unaffected by small uncertainties in inertia, but these uncertainties pose a significant challenge when applying debris detection algorithms [12]. In Section 3.1, a spacecraft dynamics simulation is used to study the effect of spacecraft inertia uncertainties on the telemetry.

After assessing that the spacecraft inertia uncertainty is demonstrably a significant factor in producing unexpected features in the spacecraft telemetry, the spacecraft's on-orbit data is used to develop estimated inertia values that may be more consistent with observed data. In Section 3.2**,** an optimization technique is applied to identify an adjustment to the inertia which produces simulated telemetry similar to the on-orbit telemetry, with the goal that the adjusted inertia can be used to "clean" some of the perturbations from the telemetry to enable the detection of smaller debris strikes.

Section 0 discusses the data processing architecture that is optimized to become more agile and have increased functionality. This allows for easy integration of new capabilities as they become available while dealing with different spacecraft idiosyncrasies. The spacecraft telemetry mneumonics that are necessary for the various algorithms are identified, and then those datapoints are brought into the architecture for processing. This augmentation to a more robust architecture is imperative to enable applications to large datasets without having to manually fix one-offs and rewrite code for individual spacecraft. Since this application relies on processing a long duration of data for several spacecraft, it is critical to have a robust and agile data-processing capability to acquire a meaningful population of debris data.

Section 3.4 discusses the telemetry cleaning techniques that are developed to reduce the unexpected features and perturbations in on-orbit data, which has a wide range of known and unknown features due to various effects. Some are predictable and attributable while some are difficult to predict, attribute, and clean. These developments mitigate the effects of spacecraft activities on the debris strike detection algorithms using an assortment of filter designs. Once the majority of these effects are mitigated, the post-processed telemetry is smoother and more stable to enable the detection of subtle features from debris strikes.

Section 3.5 develops a synthetic debris strike that is applied to telemetry data. These synthetic debris strikes are used to test the cleaning and debris detection algorithms to ensure that these algorithms don't "clean out" the debris strike feature. This test verifies that the strike detection algorithms do respond to debris strikes and show a distinct response at the time of the strike.

It is important to understand how much data can be gleaned using these techniques. In section 3.6, an analysis is performed to evaluate how many debris strikes are expected on a specific spacecraft by applying strikes to the spacecraft in Monte Carlo simulations and using a debris flux model (ORDEM) to assess the expected annual rate of strikes. From this analysis, the sensitivity of the expected debris strike rates under the different model assumptions are assessed.

Prior work encountered persistent challenges with operators being able to share data from spacecraft operations without fear of adverse consequences. Abundant data has been collected and continues to be collected from a wide variety of orbits, but it is challenging to gain access to this data to use for public science purposes. Therefore, the anonymization of commercial satellite telemetry is paramount in enabling the use of a broad selection of data. Section 3.7 describes the anonymization process that is established early in this study to ensure anonymization techniques are utilized on the forefront of the data processing architecture and throughout the study, to develop methods which protect the anonymity and performance capabilities of specific spacecraft so that scientific studies can be performed and findings shared with the community without risking adverse consequences to the operators contributing data.

3 DEVELOPMENTS

3.1 Testing Inertia Hypothesis

3.1.1 Observations from Prior Work

In the prior NASA-funded study at the University of Colorado, all five NASA spacecraft investigated showed variability in their inertial angular momentum which presented challenges to filters intended to identify strikes.

For example, on the Lunar Reconnaissance Orbiter (LRO), in orbit around the moon, the spacecraft system's angular momentum is a combination of the angular momentum due to the rotation of the spacecraft body, the rotation of the spacecraft's reaction wheels (RW), the rotation of the spacecraft's solar array, and the motions of the spacecraft's high gain antenna. These contributions are summed into a total angular momentum, denoted "Hstar". In theory, when this momentum is expressed in the inertial frame, due to conservation of momentum the contributions of the various moving components should sum to zero. If the reaction wheel spins up, it produces and equal-and-opposite torque on the spacecraft, which would be reflected in the dynamic response and shown in the spacecraft rate telemetry. When the momentum is expressed in the body frame the magnitudes of the various components (x, y, z) will change due to rotation of the spacecraft, but when the momentum is expressed in the inertial frame the components should, in theory, be constant unless an external torque is applied.

In practice, this is far more challenging. Fig. 2 shows the telemetry from LRO after some initial cleaning (to patch some idiosyncrasies in reaction wheel speed reporting in the telemetry). The full details of this analysis are presented in [13], key graphics are repeated here with annotation to illustrate the findings. The sinusoid observed at the orbit period was seen in every spacecraft that points in a consistent direction throughout its orbit (e.g., nadir). The fluctuations in momentum during slew manoeuvres were also seen in every spacecraft.

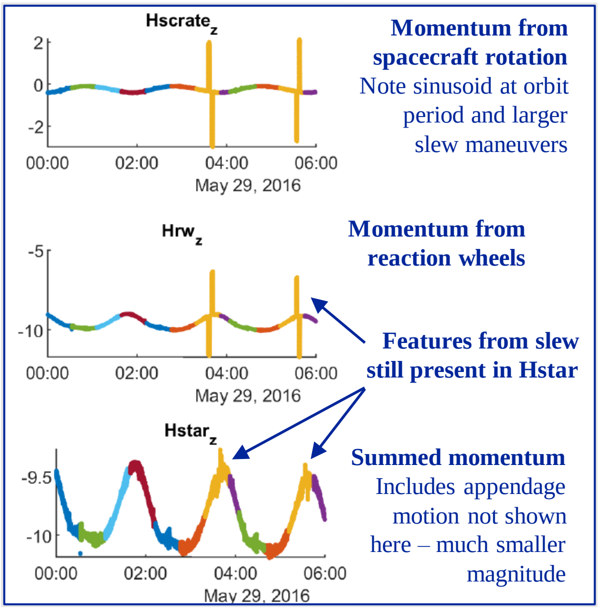

Figure 2. Typical LRO telemetry shows oscillation and other features in total system angular momentum, even when expressed in inertial frame

To take a closer look at the slew manoeuvre, Fig. 3 zooms in on the telemetry during one of the slews. The reaction wheels are spun to produce an equal and opposite torque on the spacecraft, so the momentum of the rotating spacecraft body should exactly offset the momentum of the spinning reaction wheels – but it does not quite match, as seen in the final graphic. A potential cause for this mismatch is the estimate of the spacecraft's inertia. In a simplified example, shown on the right side of Fig. 3, if the inertia of a rotating body changes, then the rotation rate would also change, but the total angular momentum would be conserved ($H_1 = H_2$). However, if the measurement of the inertia (I_2) is inaccurate, then the calculation of the momentum (H_2) will be inaccurate as well and no longer match H_1.

Since all the spacecraft assessed to date show idiosyncrasies in their inertial angular momentum, particularly during slews and in patterns commensurate with the orbit period, this study began with a hypothesis that uncertainties in the spacecraft inertia could be a major contributor to the idiosyncrasies observed. Therefore, a primary line of effort for this study is assessing this hypothesis via simulation and then developing methods to improve estimates of the spacecraft inertia, with the intention that an improved estimate could result in smoother inertial angular momentum allowing the detection of smaller strikes than could otherwise be observed.

It is important to note that these observations in no way impact the ability of these spacecraft to do their missions. All of these spacecraft are high-performance vehicles performing critical science operations for years with a stellar track record. The design of their attitude control and other systems is perfectly sufficient to perform their missions, even with the idiosyncrasies and uncertainties discussed in this work. This study is pushing well beyond what the housekeeping telemetry was ever intended to do to try to squeeze some additional unplanned science measurements out of these spacecraft, which presents these additional challenges.

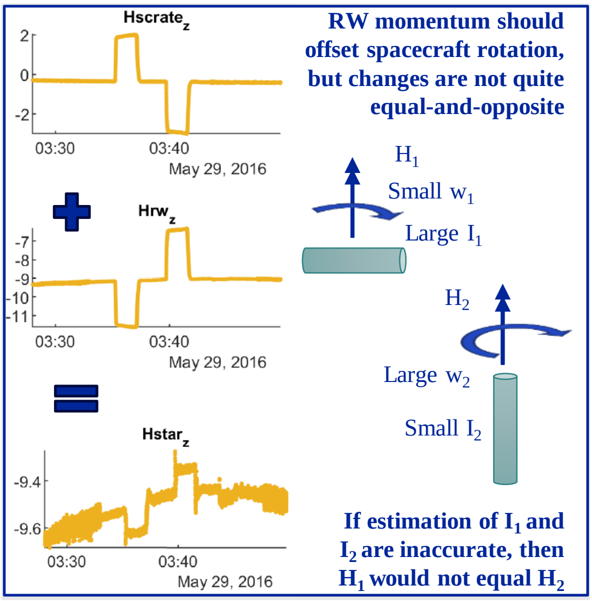

Figure 3. Taking a closer look at slew telemetry shows idiosyncrasies potentially due to uncertainties in the estimate of spacecraft inertia

3.1.2 Simulating Inertia Uncertainties

A spacecraft dynamics simulation is used to test the effects that uncertainties in the estimated inertia have on the measurement of the inertial angular momentum. The simulation models a three-axis controlled spacecraft with reaction wheels, which is controlled to a specified trajectory for rate and attitude. The spacecraft dynamics are propagated with a fourth-order Runge-Kutta integrator. This propagation represents the "true" dynamics of the simulated satellite.

From these truth parameters, "telemetry" is generated by adding noise and reducing the data rate. The noise models used are straightforward, adding Gaussian noise to rate and attitude, and using a probabilistic model to convert true reaction wheel speed into tachometer measurements (referred to as "discretized" speed to represent discrete measurements instead of continuous). The spacecraft momentum is calculated from these noisy measurements.

The spacecraft parameters are configurable, so the simulation can be updated to represent various spacecraft by changing the inertia, the reaction wheel configuration, etc. This produces an estimate of the spacecraft's response to dynamic events, which is used to assess the response to debris strikes (discussed in section 3.5) as well as the effects of inertia uncertainty.

For assessing the effects of inertia uncertainty, on-orbit data is used to define the slew trajectory for the spacecraft's manoeuvre. Telemetry for rate and attitude is upsampled to produce the desired attitude and rate of the spacecraft over time. Then the simulation is run to control the simulated spacecraft's rate and attitude to follow this trajectory. Fig. 4 shows a sample slew derived from on-orbit data that is used as the simulation slew profile for the examples in this section.

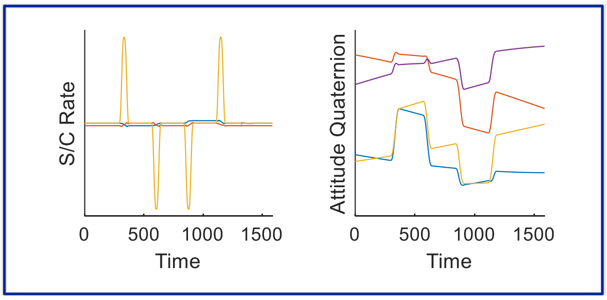

Figure 4. Trajectory used for the simulation during the examples in section 3.1.2, derived from on-orbit slew

With this simulation capability, a simulated spacecraft which manoeuvres similarly to the on-orbit spacecraft, the hypothesis that inertia uncertainties are causing many of the idiosyncrasies in the inertial angular momentum telemetry can be tested.

To assess this hypothesis, the simulation is modified to have a "truth" inertia tensor that is different from the "estimated" inertia tensor. The truth inertia is used during the simulation to simulate the dynamics of the spacecraft during slews. This produces the "truth" rate, RW speed, and attitude, which are downsampled into telemetry with noise applied. Then, when the momentum is calculated from this data, the momentum calculation uses the estimated inertia instead of the truth inertia. This estimated momentum is assessed to see if it displays similar behaviours to the on-orbit telemetry.

The "estimated" inertia is set to equal the specified inertia of the spacecraft under consideration, since this is the program's estimate of the spacecraft's inertia. The simulated "truth" inertia is then derived from that estimate by perturbing the magnitude of the diagonal terms of the inertia of the frame expressed in the principal frame. In other words, the eigenvectors of the estimated inertia are used to convert the estimated inertia into the principal frame, where the off-diagonal elements are zero. Then the diagonal elements are perturbed – in this example each diagonal element is decreased by 5%.

In order to perturb the principal frame of the inertia tensor and still generate a valid inertia, an additional transformation is added to the body-frame-to-principle-frame transformation. The frame is offset by 5 degrees in yaw, pitch, and roll, which perturbs the off-diagonal terms while still maintaining a valid inertia tensor. This perturbed inertia is used as the "truth" inertia for the simulation, representing a case where the spacecraft's estimated inertia is 5% higher in each diagonal term and offset 5 degrees in each rotation axes from the true inertia of the spacecraft. Results are shown in Fig. 5.

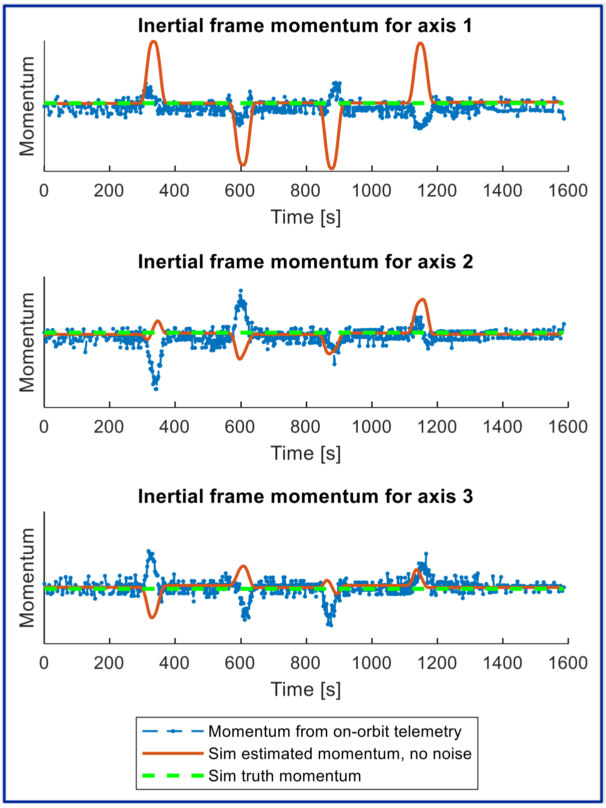

Figure 5. Simulating spacecraft manoeuvres and comparing "estimated" momentum to "true" momentum to on-orbit momentum calculation

Fig. 5 shows the simulation output compared to on-orbit data during a slew. The vertical axes scale is matched between the subplots in the figure. The green line shows the truth momentum of the sim, which, as expected, is constant when expressed in the inertial frame. However, the on-orbit calculation of momentum shows spikes whenever the spacecraft is slewing. The simulated estimate of the momentum, which uses an estimate of the spacecraft inertia which is offset from the 'true' spacecraft inertia shows similar spikes at similar scales, though distributed differently among the three axes.

This observation shows strong evidence lending credibility to the hypothesis that uncertainties in the estimate of the spacecraft's inertia are contributing significantly to the features in the spacecraft momentum which pose challenges for the detection of small debris strikes. The simulation successfully replicates the patterns of behaviour seen in on-orbit telemetry, so the next task is to figure out how to improve the estimation of the spacecraft inertia, which could clean several challenging features out of the on-orbit telemetry and render small debris strikes easier to detect.

3.2 Estimating Inertia

To begin, particle swarm optimization (PSO) is implemented to identify an inertia that improves the simulation's ability to replicate the behaviour seen in the on-orbit telemetry. There are several papers in the literature on various methods to estimate the inertia and other spacecraft parameters, these techniques can be leveraged as needed. The PSO represents a quick 'first cut' at the problem, to see if it performs well enough for this application.

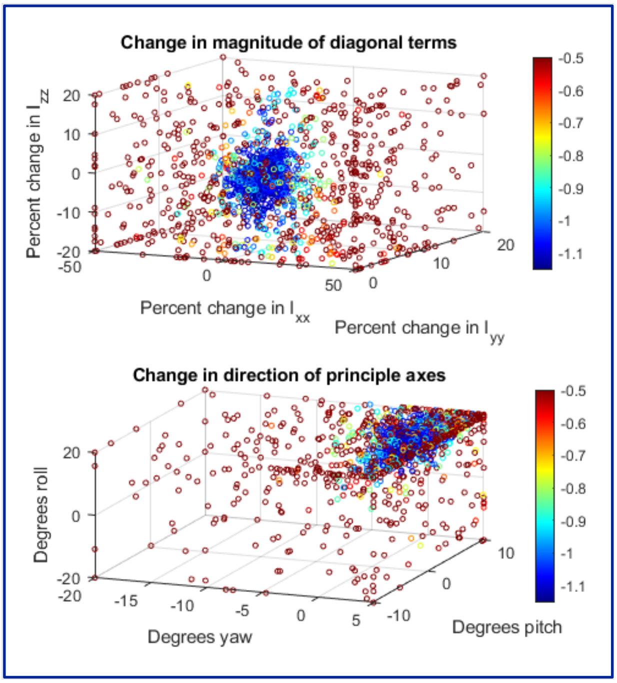

Figure 6. First pass of particle swarm optimizer, querying the tradespace to dial in on promising regions

In a PSO, a "swarm" of particles is initialized around the tradespace and then each particle is scored on its performance (ability to minimize the objective function). After this the swarm is iterated, updating each particle's position and velocity within the tradespace, and each successive iteration converges toward the best-known position within the swarm, which eventually tends to converge to a local minimum but doesn't guarantee convergence to a global minimum. One advantage is that the function does not need to be differentiable, as is required for many optimization problems.

For this application, the swarm is initialized in a 6-dimensional tradespace to explore various corrections to the spacecraft's "truth" inertia, using the specified inertia as the "estimated" inertia. These corrections are the same as the perturbed inertia described in section 3.1.2, with three parameters representing a scaling of each of the three diagonal terms of the principal inertia, and three angles representing an Euler angle rotation to correct the direction of the principal axis. The objective function is simply the mean squared error between the on-orbit telemetry and the simulated telemetry – in other words, the difference between the red and blue data in Fig. 5.

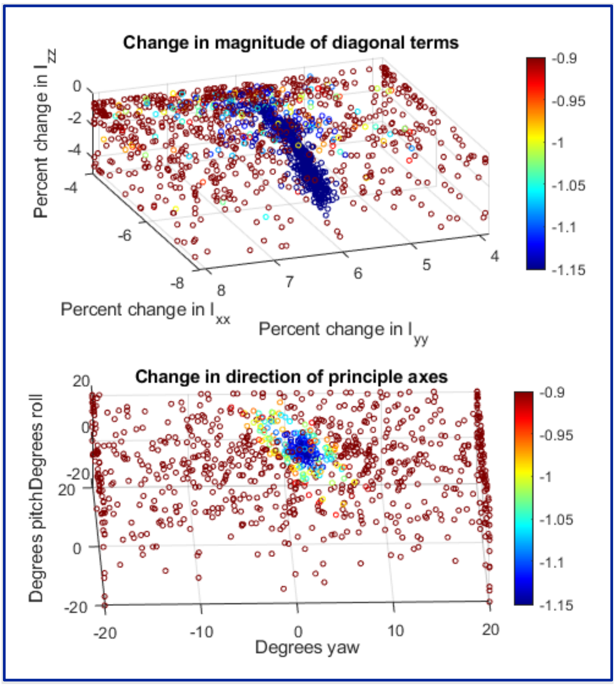

Figure 7. Second pass of particle swarm optimizer, assessing performance of objective function for various inertia adjustments

The PSO is conducted in two passes (Fig. 6 and Fig. 7). The objective function is more sensitive to changes in the magnitude of the diagonal terms than changes in the direction, reducing the likelihood that the converged solution is close to the global minimum since some adjustments to principal axis angle may score well even if they are not close to the correct value. In the second pass (Figure 7) the bounds are reduced, and the direction appears to converge nicely to a single region in space while the magnitude of the diagonal terms struggles a little more. In this case, there is an arc in the x-z plane that seems to score fairly well, without much differentiation along the arc, suggesting that the data used in this run of the PSO may not render these inertia terms as observable as would be ideal to converge on a distinct solution. Referring back to the slew trajectory in Fig. 3, the spacecraft rotates primarily about one axis, so finding a slew manoeuvre with some more diversity in rotation axis would likely improve convergence to a local minimum as the inertia would be more fully observable based on the data.

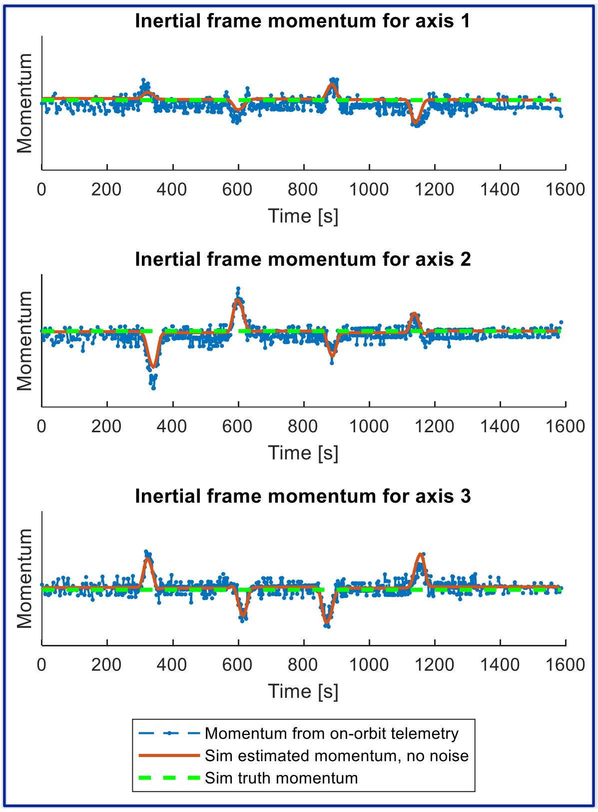

Figure 8. Simulated momentum compared to on-orbit momentum after PSO corrects "true" inertia in sim

Even with these limitations, the end result is satisfying, as shown in Figure 8. When the simulation is run using the adjusted inertia, the simulated values track much more closely with the observed values on orbit, suggesting that many of the idiosyncrasies in telemetry which cause challenges for running strike detection algorithms may be alleviated by improving estimates of the spacecraft's inertia and post-processing the telemetry to produce an improved measurement of the spacecraft momentum. In this case, the adjustments to the magnitude of the diagonal terms were single-digit percentages, while the adjustments to the angles were single-digit degrees in yaw and pitch, with a 15-degree adjustment in roll. These are reasonably small corrections, aligned with expectations about the uncertainty in inertia estimates depending on the spacecraft.

3.3 Data Processing Architecture

The data processing architecture provides the foundation on which the methods and techniques used to detect debris strikes are built. The architecture consists of the initial formatting of the retrieved on-orbit telemetry through to the data load-in and preparation of the data for cleaning and calculations. There are two different sets of code for data processing.

The first set of code is the Comma-Separated Values (CSV) Generation code repository. This repository of code is the standalone portion that uses retrieved raw telemetry from the archiver as an input and converts that data to specifically formatted CSV files for this study.

The CSV Generation code has two separate inputs, the pseudonym decoder file and the raw telemetry files. The decoder file stores each mneumonic that is to be processed and the associated pseudonym for that mneumonic and the spacecraft. This is used by the code to determine the replacement names for the spacecraft and mneumonic as a first step anonymizing the data. The code has a generalized input function to handle a wider variety of vehicle and database nuances allow it to accommodate a plethora of different spacecraft and archivers. The flow of the CSV Generation code is shown in Fig. 9.

There are five segments to the reformatting of the raw telemetry data. These segments are applying pseudonyms, updating the time stamps to a more Matrix Laboratory (MATLAB) friendly format, separating the data by date to create individual date files, initial cleaning of the data, and creating the new CSV files.

One of the first things the code does is switch out the spacecraft name and the mneumonic name of the data being processed to the pseudonyms defined in the decoder file. The next big tasks that the code completes are based in optimizing the CSV files for MATLAB and creating an effective and organized storage solution for the telemetry. One of these tasks is to update the time stamps for the telemetry. This takes the time in any format and updates it to the MATLAB serial date number format. This makes the timestamp completely numeric leading to increased processing speed within MATLAB. Part of the reformatting of the raw telemetry data includes cleaning of the telemetry. This does not alter the data in any way, it simply removes artifacts from the archiving and retrieval process. Cleaning of the telemetry includes removing empty data rows, removing placeholders such as undefined numeric results, and removing duplicated timestamps. Once the data is clean, each single and complete day is separated to be saved in a unique file. Each file of saved data consists of a single day of data, but that file can house multiple mneumonics that are at the same telemetry rate.

The second set of code for data processing is the Calculation Initialization repository. This portion of the code sets the telemetry in the proper format for the telemetry cleaning, filtering, and calculation process along with handling the saving and organization of data and results. This code is built upon the minimum working example (MWE) which was developed following the efforts in [6]. It takes the initial MWE and optimizes the code to allow for streamlined scalability of the code, data processing, and thus the study as a whole. There is a focus on making it simple and straightforward to process additional spacecraft and add new algorithms and filters.

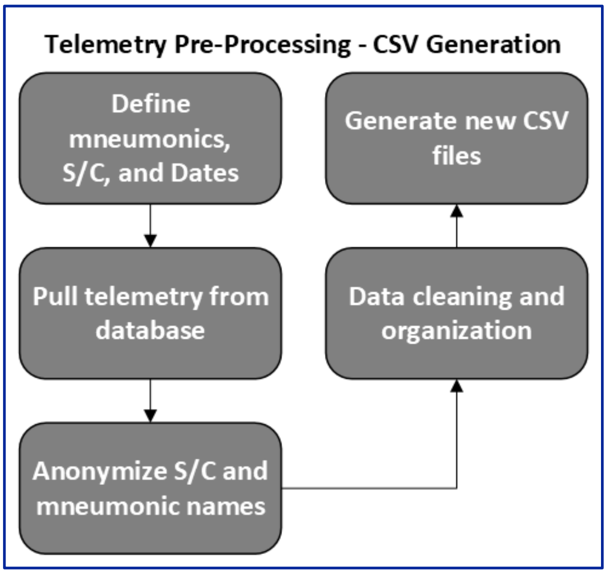

Figure 9. A graphical representation of the code flow of the CSV Generation code

The improvements to the code allow for the simple processing of multiple and new spacecraft and the ability to process larger datasets, whether that is more mneumonics, higher rate telemetry, or longer timespans of data. The process of adding additional spacecraft, mneumonics, algorithms, and filters now has a plug-and-play approach which increases the efficiency of exercising varying test cases and iterating through methods. This section of the code starts with the individual CSV files generated in the CSV Generation code repository as an input and ends with the synchronized timetable of mneumonics that is passed on to the filters and algorithms.

As shown in Fig. 10Figure 10, there are a few main sections of the Calculation Initialization repository code portion of the data processing architecture. Walking through Fig. 10, left to right, describes the data processing architecture and flow of the Calculation Initialization repository. The data processing flow begins with initializing the workspace and setting up the initial parameters needed to run, which includes spacecraft pseudonym, date range to run, raw on-orbit telemetry or synthetic telemetry, etc. this then extends to defining the spacecraft and calculation parameters, such as inertias and transformation matrices, at a lower level of the code. Following the initialization and setup, the mneumonics to load into the workspace are defined. These can be defined and loaded in large groups or as single mneumonics based on the requirements of the subsequent calculations. The next big process in the flow is the loading in a combination of the telemetry from the CSV files. Using the defined mneumonics, spacecraft, and date ranges the individual CSV files are loaded and interpolated to matched data rates allowing for all the data to exist in a single timetable.

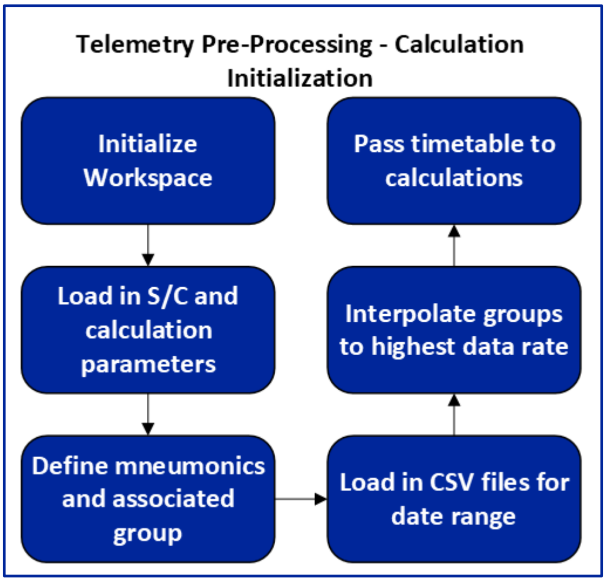

Figure 10. A graphical representation of the code flow of the Calculation Initialization code

An inherent process has been designed for determining the optimal interpolation method for each individual mneumonic to ensure the accuracy and authenticity of the data is maintained. Within the code there is an option to enable the chunking of data. This feature allows for data to be loaded, synchronized, cleaned, and run through the calculations in shorter timetables. With this feature, each chunk is processed and then an option is presented to combine the chunks back together or maintain them as chunks. This allows for less computer memory usage and faster processing when working with larger datasets. At this point, the single complete timetable of required mneumonics, or timetable chunk if requested, is completed. All of the data loaded and processed can be saved at this point, as well as the interpolation method choice for each mneumonic. This allows for rapid iteration and testing of calculations and filters while basically removing the time intensive data loading and synchronization process for subsequent runs. Once the final timetable is created, it is available to use in all of the telemetry cleaning techniques and various methods discussed in the subsequent sections.

3.4 Telemetry Cleaning and Strike Detection

The crux of the effort relies on detecting very subtle features in telemetry, which has proven challenging in on-orbit telemetry due to an abundance of features and idiosyncrasies. These larger features, motions, and adjustments in the telemetry often trigger a response in the debris strike detection filters which overwhelms potential responses to subtle debris strikes, making it difficult to detect strikes. For example, Fig. 11 shows the filter's response to a known strike on the Magnetospheric Multiscale mission (MMS). MMS experienced large and unexpected disturbances ("mystery torques") exciting responses in the filters, as well as thruster firings which are removed from the data (grey bars).

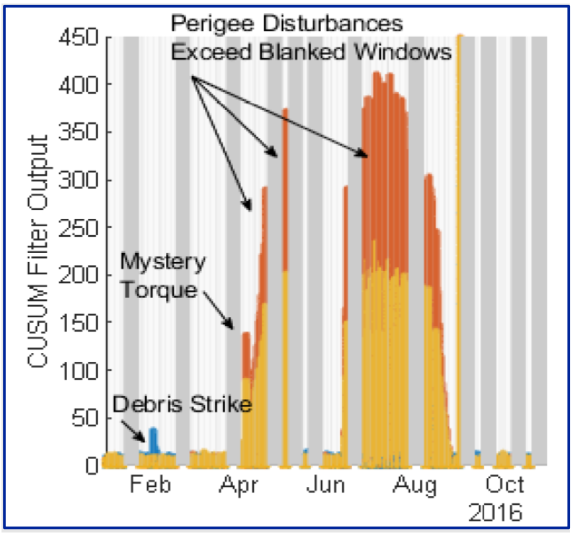

Figure 11. Data from prior study showing debris strike with other large, unexpected features [13]

Therefore, a primary effort in this study is developing methods to pre-process telemetry from spacecraft to remove (or "clean") features which are not debris strikes but which trigger a response in the strike detection algorithms. These techniques, along with the strike detection techniques, are described in the order in which they are generally applied to the telemetry. Note that various spacecraft have various idiosyncrasies which may render some of these techniques unnecessary or require additional cleaning techniques to be developed.

3.4.1 Cleaning Techniques

A majority of satellites use reaction wheels to control their attitude, and the reaction wheels often measure speed using a tachometer which produces a discrete number of pulses per wheel revolution. Since precise wheel speed is often unimportant to ground operators, the wheel speed telemetry is sometimes reported in coarser increments, with discrete rather than continuous readings. These "discretized" measurements of the wheel speed produce a discretized estimate of the spacecraft momentum, from which it is more difficult to see abrupt jumps in the momentum if they happen between two discretized RW readings.

To smooth these readings, a Savitsky-Golay filter is implemented. A Savitsky-Golay filter operates by fitting a polynomial of specified order to a series of adjacent datapoints and then taking the midpoint of that polynomial to represent the filtered value of the data. This technique is applied in a sliding window across the data to smooth noise from the data without distorting the underlying signal. A running average is a trivial example of a Savitsky-Golay filter.

Fig. 12 shows the output of a running average compared to a higher-order Savitsky-Golay filter for estimating the value of RW speeds between the discretized wheel speed outputs in telemetry. This chart shows a Savitsky-Golay filter operating on a window of 7 datapoints, with a 4th-order polynomial fit, compared to a running average operating on a window of 11 datapoints. Having a shorter window is an advantage, if the window is too long the smoothed RW speed introduces variability into the momentum which trips the strike detection filters by over-smoothing the data and changing the RW speed data too soon, before a change has actually occurred. Note that the Savitsky-Golay filter is capable of estimating speeds that are outside bounds of the raw discretized telemetry, which is probably more reflective of the true wheel speed. The filter parameters for running average and Savitsky-Golay can be tuned for specific spacecraft to optimize ability to detect minor debris strikes by choosing an appropriate window size to capture abrupt changes in RW speed on scales that allow the application of the change detection algorithms, without over-smoothing and losing the strike feature or under-smoothing and creating an excessively noisy/coarse filter output.

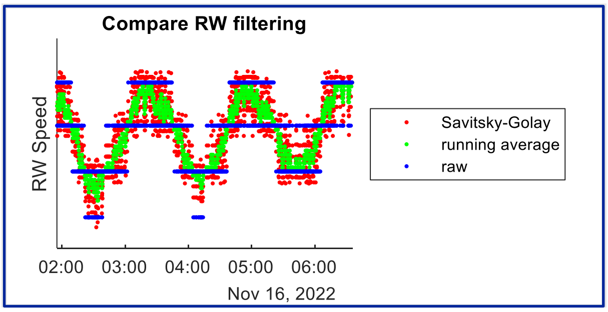

Figure 12. Comparing raw RW speed telemetry to running average and Savitsky-Golay filter

As a next step in pre-processing, a comb filter is used to remove periodic signals from the data. In this implementation, rather than adding a time-delayed version of the signal, a Fourier series is fit to the data to estimate the periodic signal in the data. Then this Fourier series is subtracted from the data, to remove the periodic signal and produce a smoother dataset. Fig. 13 shows the data before and after the comb filter is applied. Note that many of the periodic features are reduced significantly while the non-periodic features remain intact – this is desired in this application because the goal is to remove any expected changes in the data while leaving unexpected changes intact.

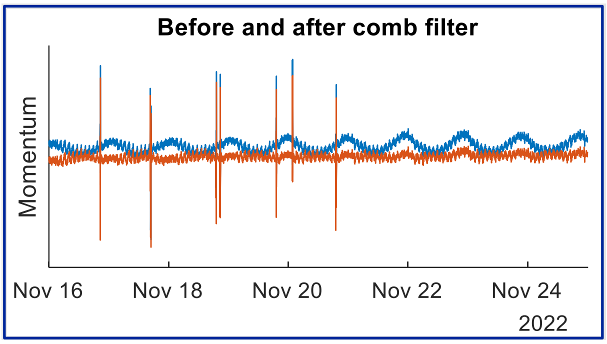

Figure 13. Applying comb filter to a long data series

In this example, a 7-term Fourier function is fit to the data using Eq. 1

$$f(x) = \sum_{n=1}^{4} A_n \cos\left(2\pi \frac{n}{4} x\right) + B_n \sin\left(2\pi \frac{n}{4} x\right) \quad (1)$$

Then, this function is subtracted from the data via Eq. 2, where α = 1 to subtract the fitted signal from the data.

$$y = x - \alpha f(x) \quad (2)$$

In this example the main periodic feature is pretty much removed, but some higher-frequency features and a remnant of the feature remain where the Fourier fit didn't completely capture the shape of the signal. The inertia estimation work is intended to improve the ability to remove some of these signals by correcting the underlying model, leaving fewer unknown periodic signals for the Fourier function/comb filter to remove.

There are an array of digital signal processing techniques that can be applied to remove signals like this, in prior work a notch filter [12] has been used or a square wave fitted to and then removed from the data [13] to capture and correct the unique periodic features on other data.

3.4.2 Strike Detection Filters

After preprocessing, it is time to run filters which are designed to detect debris strikes on the cleaned, filtered data. Two primary methods are used in this work, and both are developed in prior studies [11, 4, 13]. They are repeated briefly here for convenience.

First, a matched filter can be applied to spacecraft rate and attitude telemetry to detect the subtle response when the spacecraft experiences a minor debris strike. Debris strikes impart rotation which causes the attitude to drift, then the spacecraft control algorithms autonomously spin up the wheels to correct the transient rate and bring the pointing back to the desired attitude. A matched filter looks for an expected signal in noisy data, which is implemented by cross-correlating two signals: the noisy data and the expected signal, or "wavelet", which shows the expected response of the spacecraft to a debris strike. A graphic of this is shown in Fig. 14.

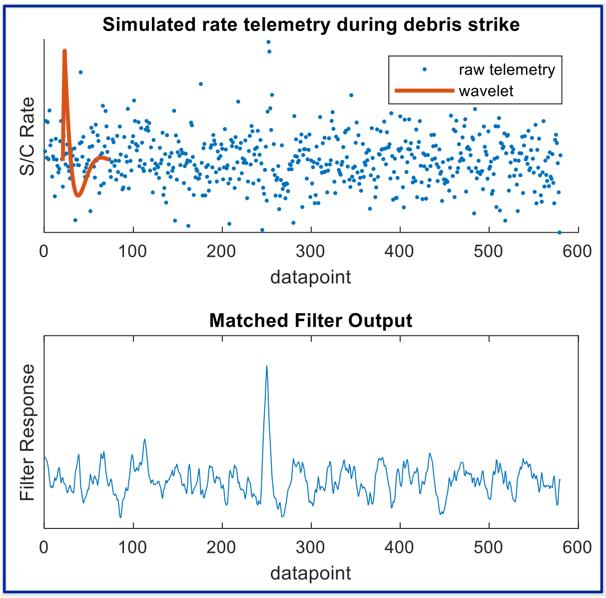

Figure 14. Raw telemetry with subtle debris strike at time = 250 seconds, showing rate wavelet and matched filter output from cross-correlation

The second primary detection technique is a change detection method referred to as a Sequential Probability Ratio Test (SPRT). This technique is used to find points at which the distribution of underlying noisy data has experienced a change. This is implemented in a sliding window which consists of a "pre-window" and a "post window," to look for a strike in the middle at the junction between the two windows.

The pre-window is used to establish the expected distribution (mean and variance) of the data. Then, each datapoint in the post window is evaluated sequentially to determine if, as a whole, the data set is more likely to be from the same distribution as the pre-window or from a specified changed distribution. This application looks for a change in mean, and evaluates each datapoint on the likelihood that it is drawn from the same distribution as the pre-window or from a distribution with a higher or lower mean value. This is evaluated via a likelihood ratio test, which is summed across the datapoints, so that the resulting test statistic stays near zero when the data is collectively more likely to be from the pre-window distribution, and spikes upward when the data is collectively more likely to be from a changed distribution. Fig. 15 shows an example of implementing this technique.

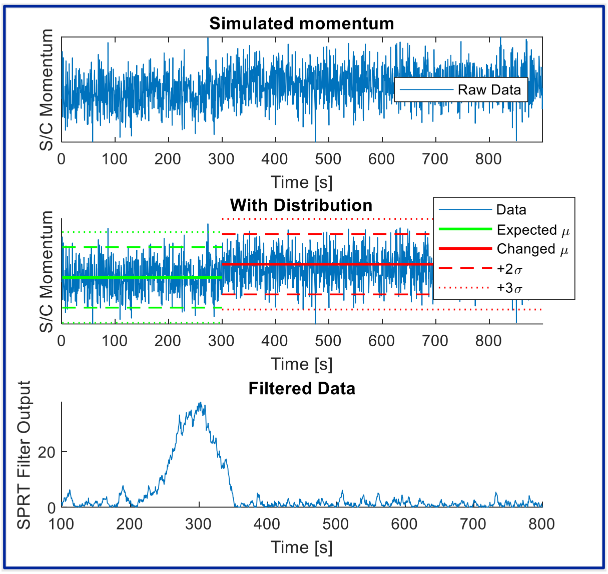

Figure 15. Momentum data with a small change indicative of a minor debris strike, SPRT response shows clear feature at strike

3.5 Adding Synthetic Strike to Telemetry

In order to assess the performance of the detection algorithms, this study develops a capability to insert a "synthetic" debris strike into the regular on-orbit telemetry data of the spacecraft. This manifests as a sudden external torque upon the spacecraft dynamics simulation. By mimicking the effects expected in a strike, the synthetic telemetry can be run through the filters, fine-tuning them so that they behave as expected: cleaning expected noise in the telemetry while accentuating the features expected in a strike.

The spacecraft dynamics simulation described in Section 3.1.2 is used to assess the spacecraft's response to the debris strike. A reference trajectory is created using the telemetry so that the simulation is following the attitude and rate profiles of the on-orbit spacecraft during this timeframe. The magnitude and direction of the strike are specified, various magnitudes can be tested to see the effects of various sizes of strikes. Parameters for the spacecraft are specified in the simulation to assess how it would react to the small angular momentum transfer produced in a strike. To show the impact, the simulation is run without the strike and again with the strike.

The effects of the strike are then incorporated into the telemetry to be returned as "synthetic" telemetry which is on-orbit telemetry plus the effects of a synthetic debris strike. This includes the rotation rate of the spacecraft with the response to the strike, the RW speed plus the change in speed due to the strike, and momentum. Depending on the input format of the raw RW telemetry, discrete steps or continuous estimate, it is adjusted to output data in the same format. The simulation outputs a continuous RW speed, so if the spacecraft downlinks discrete measurements of RW speed (tachometer count or similar) then the tach count readings need to be re-generated from the truth RW speed plus the change in speed from the added momentum of the debris strike.

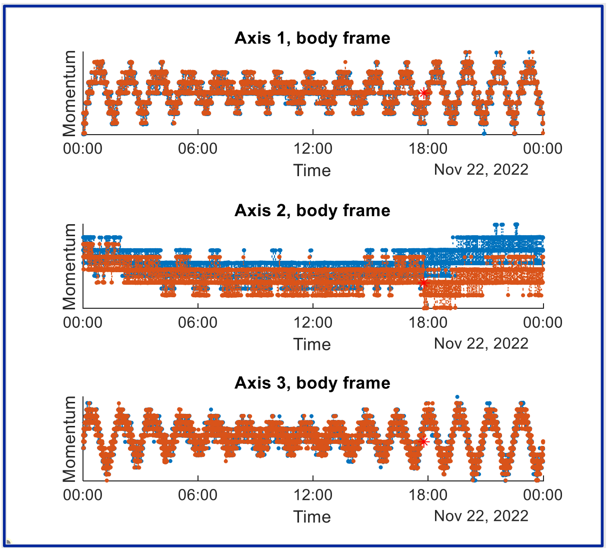

Figure 16. Comparing the telemetry in all three axes of body-frame momentum after the synthetic strike has been inserted on axis 2, with raw data in blue and synthetic in orange on top

Once the synthetic strike has been added, the changes in telemetry can be observed. Fig. 16 and Fig. 17 illustrate how the momentum telemetry looks with an inserted synthetic strike, superimposed on plots of the original data with no strike. Note that the body frame momentum shows oscillations in the first and third axis commensurate with the orbit period. The second axis, however, is fairly quiescent since the spacecraft is pointed such that the momentum in that axis is consistent throughout the orbit. The strike is placed on the second axis, showing a very obvious signature.

For the inertial frame momentum, the oscillations from the slew to maintain pointing in orbit frame are reduced, but not eliminated entirely. The work on inertia estimation in Section 3.2 is intended to reduce these remaining movements as well as other features, and techniques like the comb filter can also be used to remove periodic features like this without affecting the strike signature. In the inertial frame, the strike is partially on axis two and partially on axis three. This is a simplified case for this example, since the strike is placed on a straightforward axis, the effects are a little more nuanced when the strike is placed at other locations.

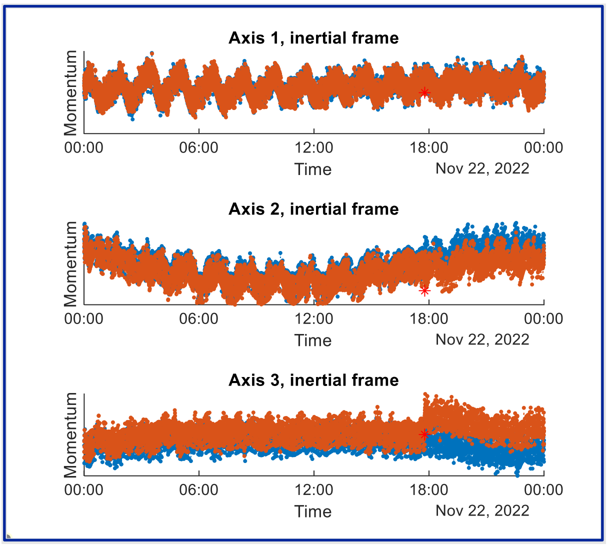

Figure 17. Comparing the telemetry in all three axes of inertial-frame momentum after the synthetic strike has been inserted. Raw data in blue and synthetic in orange

The next step is running this synthetic telemetry through the filters, to see how they respond: the filters that clean the telemetry should not reduce the signature of the debris strike, and the filters that detect the strike should show a clean response to the strike. Currently, a matched filter and SPRT have been applied to the synthetic data. The mechanics of these filters are discussed in more detail in Section 3.4.2.

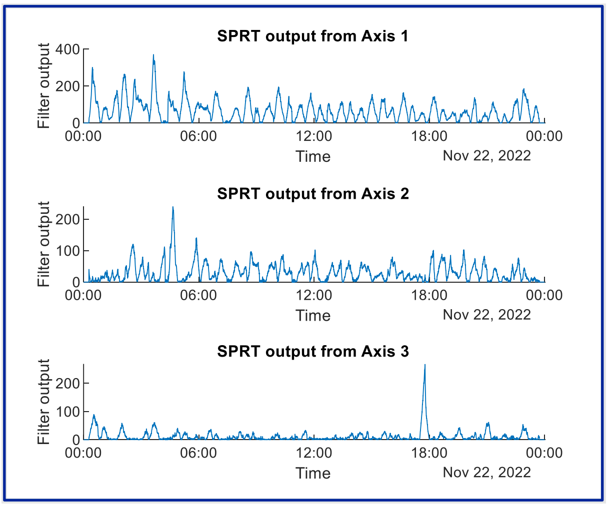

Figure 18. SPRT applied to momentum telemetry

Fig. 18 shows the output from the SPRT filter when applied to the momentum telemetry in Fig. 17. This illustrates the importance of using "clean" momentum telemetry for the SPRTs. Axis one shows minimal effect for the strike but has an oscillation from the spacecraft rotation around its orbit, and the SPRT results are noisy. Axis 2 shows a small effect from the strike, but it is overwhelmed by the noise from the oscillations. Axis 3 has clean, quiescent data and shows a very clear response at the time of the strike. Note the y-scales. If the threshold to detect the clear strike signature on axis 3 was applied to axis 2 without regard to the noise in the data, there would be abundant false alarms due to the oscillations, for example one spike just before 6 AM. This illustrates the importance of obtaining clean telemetry in order to use the strike detection algorithms effectively, which is a major focus of this effort.

Fig. 19 shows the synthetic debris strike in the rate telemetry and the matched filter applied. This illustrates a very strong response to the perturbation, since the rate telemetry for this spacecraft has high accuracy and low noise relative to the rotation produced in response to the strike. However, the relatively low data rate could reduce the sensitivity of the matched filter significantly if the strike occurred at a different time relative to the telemetry output, the majority of the rotation and correction could occur before the next telemetry point is logged which would reduce the size of the signal substantially.

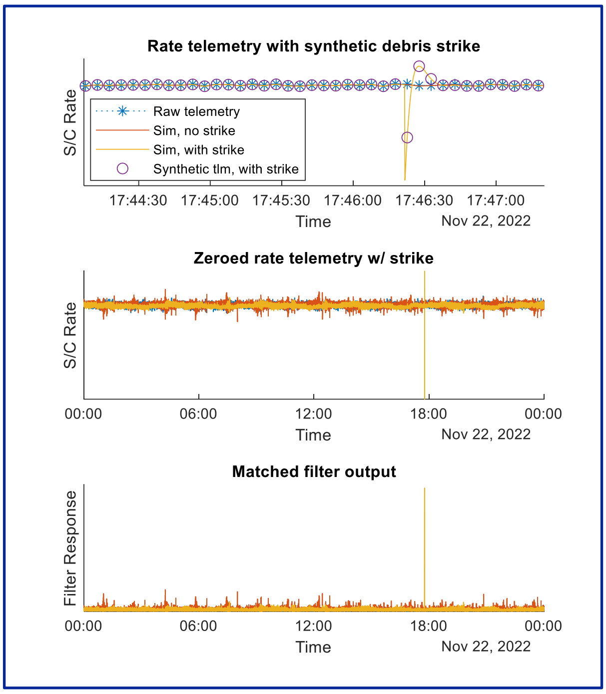

Figure 19. Rate telemetry with synthetic strike, then with matched filter applied.

3.6 Assessing Expected Perturbations

A key question central to understanding the potential utility of this concept is understanding the relationship between the number of debris strikes which can be detected by a given spacecraft, the threshold of perturbation which a spacecraft can detect, and the parameters underlying debris analysis which may influence the expected results. While many of these variables are uncertain, this study conducts trade studies to assess the relationship between and characterize the effects of these variables.

For example, a spacecraft in an orbit with significant populations of hazardous non-trackable debris, such as an 800 kilometre (km) high-inclination orbit, would expect to encounter more strikes than a spacecraft in a less-concerning orbit, such as a 500 km low-inclination orbit. A larger spacecraft would expect to experience more strikes than a smaller spacecraft, but this effect could be offset by a smaller spacecraft experiencing a more dramatic (and therefore more detectable) change due to the debris strike, such as a larger change in rate (due to the smaller inertia) or a larger change in reaction wheel speed (due to smaller RWs). However, a smaller spacecraft may have lower rate telemetry or less sensitive instrumentation, possibly making strikes harder to detect.

It is critical to understand the sensitivity of the expected result to the parameters underlying the analysis. In the 2017 NESC report [7] the shape of the debris piece was adjusted to show the variation in expected strikes and its relationship to on-orbit observations. It was also shown that changing the momentum enhancement factor from one to three changed the predicted number of perturbations by a factor of more than 6 times. These variables (shape, density, and momentum enhancement) are very difficult to characterize definitively for all the potential configurations of debris impacts, so the idea of this component of the study is to model a series of assumptions for these variables in a trade study to gain an understanding of the potential effect on predicted number of strikes for a given spacecraft in a given orbit.

To accomplish this, this study leverages the Mission Architecture Resilience and Survivability in Debris Environment Tool (MARSDET). MARSDET assesses the effects of the debris population on satellites. It consists of a model of a debris strike, which is wrapped in a Monte Carlo, which is wrapped in a trade study. It applies a specified debris strike to a specified satellite and calculates the resulting effects (change in angular and linear momentum, damage, and secondary debris created). It repeats this assessment across a Monte Carlo, to assess various sizes and velocities of debris striking a satellite at various locations. For this Monte Carlo, the primary parameters of each debris strike are drawn from NASA's ORDEM model (for direction, characteristic length, density, and velocity) and then the secondary parameters can be selected via a variety of models (momentum enhancement factor (MEF) and relationship between mass and characteristic length). Finally, these Monte Carlos are conducted as a series of parametric trades to assess the influence of various analysis parameters on the results. Traded variables for this analysis are the underlying models for MEF and characteristic length-to-mass conversion. A more detailed description of MARSDET and its analytical capabilities is provided in [14].

The expected rate of minor debris strikes on NASA's Orbiting Carbon Observatory-2 (OCO-2) satellite is assessed. OCO-2 is modelled in a 700 km orbit at an inclination of 95 degrees, which is one of the more congested orbits in terms of small debris populations. A family of scenarios are traded to assess the effects of various models for momentum enhancement factor and the relationship between characteristic length and mass. This approach explores the uncertainty in these parameters and its effects on the expected number of strikes that a spacecraft like OCO-2 might detect.

The results are plotted in Fig. 20 with the detection threshold on the x-axis and the expected perturbation rate on the y-axis. There is an inverse relationship between the size of strike which can be detected and the number of strikes which are expected. If a spacecraft can detect smaller strikes, then more strikes are expected.

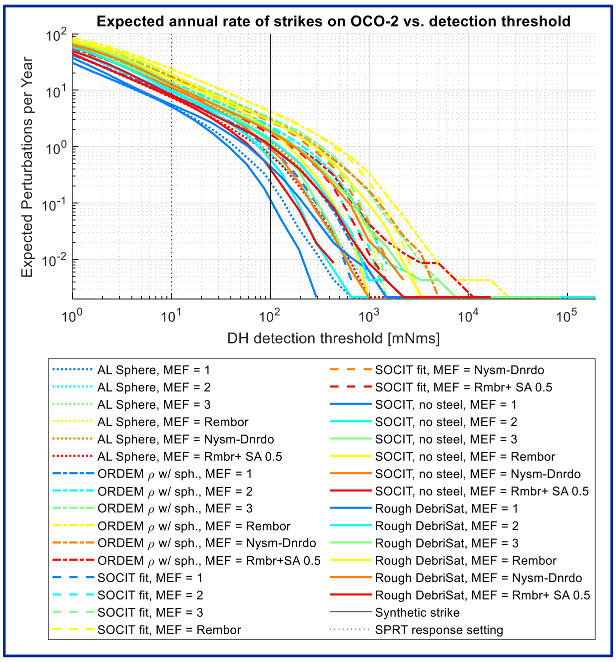

Figure 20. OCO-2 expected strikes vs threshold when various models are applied

The models for the characteristic length to mass conversion are as follows: "AL Sphere" assumes that every debris piece is an aluminium sphere, "ORDEM rho w/ sph" uses the ORDEM density classes and assumes a spherical shape, "SOCIT fit" uses the voided mass relationship derived in the Satellite Orbital Debris Characterization Impact Test (SOCIT) hypervelocity impact experiments [7, 13], "SOCIT no steel" indicates that the SOCIT relationship is used but the high-density population is assumed to SOCIT characteristic length (Lc)-to-mass relationship for aluminium, and "Rough DebriSat" is an example of a probabilistic draw which roughly approximates the graph of the DebriSat Lc-vs-mass data, as described in [13].

The models for MEF are: setting MEF to a constant 1, 2, or 3, using Rembor's fit, using the Nysmith-Denardo fit, or finally using the Rembor fit for bus strikes but assuming that all solar array impacts have an MEF of 0.5, similar to a simulation by Ryan [13].

There are a few interesting takeaways from Fig. 20. For one, modifying the underlying models to use various reasonable assumptions for debris risk analysis results in an order of magnitude difference in the number of strikes expected. This is concerning, because often missions are assessing their risks using assumptions like these, and with the results so strongly dependent on underlying assumptions it is difficult to accurately assess the risks with reasonable confidence, which complicates licensing, insurance, and system resiliency decisions.

For the purposes of this study, Fig. 20 shows that the expected data collection (number of detectable strikes) is, as expected, strongly correlated to the detection threshold. A detection threshold of 100 milli-Newton metre second (mNms) would allow OCO-2 to see approximately one strike per year, while if that threshold can be lowered to 10 mNms then OCO-2 would expect to see about 10 strikes per year, resulting in a much richer dataset for assessing models and trends over time, especially when extrapolated across larger populations of contributing spacecraft.

3.7 Anonymization

To set the stage for scaling up the capabilities of the methods and techniques previously described, this study places a focus on anonymization of data and results. This focus on anonymization promotes the sharing of findings between researchers, government, and commercial stakeholders. It also opens the door for increased data sharing potentially allowing researchers access to telemetry data to use in refining the debris assessment methods and techniques as well as giving additional test cases for proving the methods against a diverse dataset. It is understood that if this study cannot protect customer data, then customer data cannot be used for this study.

The mantra for anonymization in this study is that no spacecraft should be able to be identified through the results of this study. This section discusses the current plan for the development of anonymization techniques including the identification of specific data that needs to be protected and our approach to detailing methods and results without disclosing sensitive information. As this study progresses, the techniques will adapt to ensure that the mantra holds.

The first step in the anonymization process is internal to the research team. The naming of the spacecraft and the spacecraft mneumonics are the first items to be anonymized. During the initial intake of telemetry, the processed and stored telemetry files are renamed to our naming convention, allowing the spacecraft names or identifying markers to be removed from the filenames. Pseudonyms are used in place of the actual spacecraft names. Additionally, within the files, the mneumonic names are also changed to pseudonyms to negate the ability to tie a telemetry set to a spacecraft or company through the mneumonic naming. These pseudonyms are stored in a single file that can be used as an actual name to pseudonym decoder. This file is only accessible by the primary researchers and is not accessible to the telemetry cleaning, processing, and calculation part of the code.

A spacecraft identification analysis is performed to determine what information can trace back to a specific spacecraft and consequently what information requires protection. Some examples of information that would fall into this category includes altitude, orbit slot, times and spacecraft locations, hardware suite, mission, and years of operation. This is all information that needs to be protected to ensure spacecraft anonymity. In addition to singular items of information, specific combinations of information also need to be protected.

Another part of the anonymization process is the information that could reveal competition-sensitive capabilities. Items such as Guidance Navigation and Control algorithms, anomaly details, and specific detection thresholds are all information that demonstrates performance capability and thus needs to be protected in the interest of industrial and international competition concerns.

An additional level of anonymization is planned to be incorporated in the reporting of the methods and findings for the final report associated with this study. The report will discuss the methods developed to execute these techniques, as well as the process for presenting results in a way that shows relevant information without disclosing sensitive information.

There are a few techniques that will be applied in the final report to ensure that the data and methods share do not inadvertently allow for derivation of performance or correlation to a specific spacecraft. To protect the spacecraft performance and hardware, the report will refrain from using numbers on the y-axis of graphs and remove ties to spacecraft lifetime, spacecraft time, and orbit, and generally normalizing the data. Another alternative that will be used, when it fits the data and the graph, is to rescale the axis. Debris strike occurrence times will not be shown in the report. Plot patterns will be normalized, for example data can be shown relative to latitude as opposed to time. Data can be plotted using multiple spacecraft normalized to their unique customized detection thresholds. Additionally, the data can be plotted relative to the expected perturbations instead of the raw number of perturbations.

This will remove to ability to trace back the spacecraft via correlations such as number of strikes and its relationship to particular orbits and altitudes. Depending on the findings, additional methods will be developed to allow for the sharing of those findings without impacting the spacecraft anonymity.

The goal is to show key information and methods without revealing specific spacecraft, proprietary performance, or spacecraft identifying items using anonymization techniques throughout the execution of this study and in the products produced by this study.

4 CONCLUSIONS

This study makes significant strides in the process of identifying subtle debris strikes using active on-orbit satellites as *in situ* debris detectors. Prior work is leveraged and expanded upon to provide meaningful insights and a path forward for many challenges experienced in prior work. Overall, the results to the investigation of these challenges are optimistic with strong evidence to support the hypothesis that the challenges in the detection of small debris strikes previously experienced is due, in part, to uncertainties in the estimate of the spacecraft's inertia. Progress has been made in using an adjusted inertia estimation in simulated spacecraft to provide a closer match to the spacecraft's on-orbit experience. Progress has also been made in the telemetry cleaning and strike detection algorithms as applied to on-orbit telemetry. While there are many challenges to using on-orbit active satellite telemetry as *in situ* debris strike detectors, progressing the capability to refine debris risk assessment models and methods is intended to aid in motivating safer satellite operational practices and improved understanding of the debris environment.

5 ACKNOWLEDGEMENT

This work is supported by NASA grant 80NSSC24K0710. The authors are grateful for Drake Gates' contributions to developing some of the techniques described herein.

6 REFERENCES

1. (2019). Monthly Number of Objects in Earth Orbit by Object Type, *Orbital Debris Quarterly News*, **23**(1), 13.

2. Christiansen, E., Lyons, F., Davis, B. & Lear, D. (2017). Reaction of Spacecraft Batteries to Hypervelocity Impact, *Orbital Debris Quarterly News* **21**(1), 7–9.

3. Shell, J. (2010). Optimizing Orbital Debris Monitoring with Optical Telescopes, *Advanced Maui Optical and Space Surveillance Technologies Conference*, Maui Economic Development Board, Maui, HI.

4. Bennett, A.A. & Schaub, H. (2020). Effect of Spacecraft Parameters on Identification of Debris Strikes in GN&C Telemetry. *Advances in the Astronautical Sciences AAS Guidance, Navigation, and Control*, AAS 20-132, 172.

5. European Space Agency Space Debris Office (2025). Space Debris by the Numbers. Online at *https://www.esa.int/Our_Activities/Space_Safety/Space_Debris/Space_debris_by_the_numbers* (as of 9 March 2025)

6. Bennett, A.A., Carpenter, R. & Schaub, H. (2023). Methods to Detect Impact-Induced Orbit Perturbations Using Spacecraft Navigation Data. *Journal of Spacecraft and Rockets.* **60**(4).

7. NASA Engineering and Safety Center (2017). Evaluation of Micrometeoroid and Orbital Debris (MMOD) Risk Predictions with Available On-orbit Assets, NESC-RP-14-01000, NASA Engineering and Safety Center.

8. Flewelling, B., Hendrix, D., Bantel, M., Cunio, P., Therien,W., Jeffries, M. & Clark, C. (2017). Explained and Unexplained Momentum Impulse Transfer Events. In *Proc. 7th European Conference on Space Debris*, ESA Space Debris Office, Darmstadt, Germany.

9. Putzar, R., Watson, E., Schimmerohn, M., Karrang, P. & Millinger, M. (2019). How Hypervelocity Impacts Can Affect the LISA Mission—The MIRAD Study, *International Astronautical Congress*, *International Astronautical Federation*, Paris, France.

10. Williamsen, J., Pechkis, D., Balakrishnan, A. & Ouellette, S. (2019). Characterizing the Orbital Debris Environment Using Satellite Perturbation Anomaly Data. *First International Orbit Debris Conference, NASA Orbital Debris Program Office* 6065.

11. Bennett, A. A., Schaub, H. & Carpenter, R. (2019). Assessing Debris Strikes in Spacecraft Telemetry: Development and Comparison of Various Techniques. *70th International Astronautical Congress*, *International Astronautical Federation*, Paris, France.

12. Bennett, A. A. & Schaub, H. (2019). Identifying and Assessing Debris Strikes in NASA Spacecraft Telemetry. *First International Orbital Debris Conference, NASA Orbital Debris Program Office* 6014.

13. Bennett, A.A. (2022). Methods to Detect Minor Debris Strikes in Spacecraft Telemetry. *Graduate School of the University of Colorado*.

14. Bennett, A.A. (2024). Investigating the Prediction of Debris Risks Given Uncertainties in Models and Assumptions. *International Astronautical Congress,* IAC-24,A6,2,1,x85287, Milan, Italy.