

**Methods to Detect Minor Debris Strikes in
Spacecraft Telemetry**

by

Anne Aryadne Bennett

B.S. and M.S. Mechanical Engineering

California Polytechnic State University, 2012

M.S. Astronautical Engineering,

University of Southern California, 2018

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Aerospace Engineering Sciences

2022

Committee Members:

Prof. Hanspeter Schaub

Dr. Russell Carpenter

Prof. Jay McMahon

Dr. Dan Kubitschek

Dr. Timothy Maclay

(Ph.D., Aerospace Engineering Sciences)

Methods to Detect Minor Debris Strikes in
Spacecraft Telemetry

Thesis directed by Prof. Hanspeter Schaub

Hazardous non-trackable debris poses a significant threat to active satellites. This is debris that is too small to track and avoid, but large enough to damage a satellite if it hits a sensitive component. However, multiple events have occurred where satellites are struck by debris but damage is minor and operations continue nominally. A 2017 NASA report recommends monitoring satellites for abrupt, unexpected momentum changes to collect data on debris populations. Following the recommendation of the 2017 NASA report, this research develops methods to detect minor debris strikes in typical spacecraft telemetry.

The intention is not to develop new hardware for spacecraft to fly, but rather to develop new processing techniques to determine if *in situ* data on minor debris strikes can be obtained using telemetry that satellites already collect. While several efforts have investigated telemetry when an anomaly occurs and found indications of a debris strike, this method proactively looks for perturbations indicative of minor debris strikes, even if the strike does not cause an anomaly in spacecraft performance.

Four key developments are made in this research. First, algorithms are developed to detect subtle, unexpected perturbations in attitude control system telemetry. A spacecraft dynamics model is developed to simulate the effects of a debris strike, then digital signal processing and change detection techniques are leveraged to develop algorithms for identifying subtle indications of strikes. Second, an extended Kalman filter with dynamic model compensation is developed and augmented with various techniques to identify subtle, abrupt changes in orbit produced by a debris strike. Third, these techniques are applied to the telemetry of a variety of NASA spacecraft to investigate the challenges of real-world application and obtain some preliminary results. Fourth, a tool is

developed to predict rates of detectable perturbations on various space systems and investigate the changes to underlying models and assumptions that could be informed by data collected using these techniques.

The emerging space era requires new methods for ensuring space sustainability with regard to debris; heritage methods are inadequate with dramatically increasing populations of satellites. To motivate appropriate design, operation, policy, regulation, mitigation, and remediation a thorough, validated understanding of the debris environment is required, particularly for non-trackable debris which cannot be avoided by active satellites. These methods develop a capability to use these proliferated satellite architectures to collect data on debris, producing a new type of debris measurement to aid in understanding the small debris environment thereby incentivizing appropriate space sustainability practices.

Dedication

To my dad, for inspiring my love of engineering.

*And to my mom, for dedicating her life to teaching me all the things,
and for giving me time and space to live childhood to its fullest.*

This dissertation could only be dedicated to the pair of you

דָּרָךְ-אֲמוּנָה בְּחֵרְתִי מִשְׁפָּטֶיךָ שְׁוִיתִי:
דְּבַקְתִּי בְּעִדּוֹתֶיךָ! הִנֵּה אֶל-תְּבִישָׁנִי:
דָּרָךְ-מְצוֹתֶיךָ אֲרוּץ כִּי תִרְחִיב לִבִּי:

Acknowledgements

I'm fortunate to have worked closely with two top-notch professionals in the field: Dr. Hanspeter Schaub, and Dr. Russell Carpenter. Thank you to Dr. Schaub for his extraordinary mentorship and support, I don't know how he maintains such insanely fast response times with his packed schedule. Thank you for teaching me to PhD, for your endless support of my chaotic military family life, for your guidance throughout my research, and for two excellent classes - I can already tell I'm going to spend the rest of my career wishing that everyone else in the industry had the opportunity to take 5010. Thank you to Dr. Carpenter for hosting me 17 weeks at SSMO across 4 trips, suggesting several of the key methods in this research, providing a gateway to the telemetry, and supporting and guiding me through all the challenges and inconvenience of doing research using real-world in addition to simulated data. And to both: I see how many ways you are pulled and yet I always feel like a priority, and I don't take that for granted. Thank you for all of the time you have invested in developing me as a professional.

To the other committee members: Thanks to Dr. Jay McMahan for two excellent classes, and most especially for understanding and accommodating during two personally challenging semesters - I've told stories about my experiences in your classes a couple times, to the tune of "This is what supporting pregnant women/moms looks like." Thanks to Dr. Dan Kubitschek for working with me as an RA the first year, and for his thoughtful follow-up comments after comps. And to Dr. Tim Maclay, for taking the time to visit and chat about my research at length, and his work as an industry leader in space sustainability.

Thanks also to Haisam Ido and the rest of the TaaS team, who manage the telemetry

database, for not only providing exceptional technical support to challenging questions ridiculously quickly, but also being really really nice about it when I conducted an accidental denial-of-service attack on the telemetry server.

Thanks to all my VTE hosts, for their hospitality and the enriching career experiences working onsite at so many great institutions: Dr. Russell Carpenter at NASA Goddard Spaceflight Center, Dr. Mark Matney and Dr. Alyssa Manis at NASA Johnson Space Center, Dr. Joel Williamsen at the Institute for Defense Analyses, Dr. Danielle Wood at MIT's Space Enabled Research group, and Phillip Rittmuller at Lawrence Livermore National Lab.

To everyone who took care of my children during this long task: exceptional childcare professionals Melanie, D'Maree, and Jada, for giving my kids all the TLC they needed to thrive while I worked. And also to all the relatives who traveled to watch the kids during various VTEs and conferences: Cynthia Walker (taking the gold star with 12+ weeks of care, mostly in hotels), Sandy Bennett, John Bennett, Stuart Allen, and Alison Barker.

To Brooke and Walker: I never could focus unless I knew you were thriving. Thanks for thriving so well in so many extraordinary circumstances. Hotels, planes, 12-hour driving days, tents, two out-of-state moves, empty houses, construction zones, and so many caretakers - you rock every bit of our chaotic military family life. Thanks for putting awesome into every single day, you are always my favorite thing.

To Steven: A few sentences can't come close to expressing how much I cherish you. Thanks for so often doing way more than 50% and taking ownership of so many critical elements of orchestrating our family life. Thanks for managing to rock your military career and still be there for our family. Thanks for listening to endless rants about Kalman filters, time constants, and inertias and making up silly songs like "♪Anne broke her, ♪batch filter! ♪." Thanks for responding to all my crazy ideas (like quitting my six-figure job to go back to school) with an enthusiastic "Do it, baby!" Thanks for always taking such great care of me and our kids. Perhaps I could have done this PhD without you, but it was definitely a helluva lot more fun with you.

Contents

Chapter		
1	Introduction	1
1.1	Prelude	1
1.2	Hazardous Non-Trackable Debris: an Overview	3
1.2.1	Origins of Debris	3
1.2.2	Fragmentation Event Case Study: Cosmos 1408 ASAT test	5
1.2.3	Fragmentation Event Mechanics	5
1.2.4	Sources of Debris Measurements	8
1.2.5	Classes of Debris Events	11
1.2.6	Effects of Untracked, Damaging Debris	13
1.2.7	Uncertainties in Assessing Damage	16
1.2.8	Lessons from Comparing Modeled Expectations to On-Orbit Experiences	19
1.2.9	Putting it All Together: Understanding HNT Debris Risks and Mitigation Limitations	24
1.2.10	Zooming Out: Framing research within Space Environment Management	27
1.2.11	State-of-the-Enterprise: Emerging Trends and Implications for Debris	30
1.3	Research Concept	32
1.3.1	Core Concept	33
1.3.2	Challenges Inherent in Research Concept	35
1.4	Related Efforts	38

1.4.1	Detecting Particle Impacts Using Specialized Hardware	39
1.4.2	Analyzing Observed Perturbations	42
1.4.3	Work Involving Detection of Momentum Perturbations	44
1.4.4	Peripherally Related Work	46
1.5	Research Overview	54
1.5.1	Rotational Change Detection	54
1.5.2	Orbit Perturbation Detection	55
1.5.3	Applying Techniques to Telemetry	56
1.5.4	Using the Data: Perturbation Rate Assessment Tool	57
1.5.5	Utility to Field	58
2	Developing Methods to Detect Rotational Momentum Transfer	61
2.1	Spacecraft Dynamics Simulation	61
2.1.1	Spacecraft Simulation	61
2.1.2	Applying Debris Strike to Simulation	63
2.2	Exploring Strike Detection and Estimation via Digital Signal Processing Techniques	66
2.2.1	Matched Filter Development	67
2.2.2	Characterizing Probability Distributions	68
2.2.3	Developing Likelihood Ratio Test	69
2.2.4	Using MMSE to Estimate Strike Parameters	71
2.3	Methods for Strike Detection	73
2.3.1	Refining Digital Signal Processing Techniques for Strike Detection	73
2.3.2	Change Detection Techniques	75
2.4	Algorithm Development	77
2.4.1	Matched Filter Development	77
2.4.2	Sequential Probability Ratio Tests (SPRTs)	80
2.4.3	Results from Algorithm Development	87

2.4.4	Discussion of Methods	90
2.5	Studying Effect of Trading Simulation Parameters	92
2.5.1	Simplifying Strike Detection Algorithms	92
2.5.2	Conducting Trades	92
2.5.3	Discussion	97
3	Methods to Detect Strikes via Translational Momentum Transfer	101
3.1	Concept	101
3.2	Filter Development	103
3.2.1	Generating Simulated Navigation Telemetry	103
3.2.2	Default Values for Parameters	104
3.2.3	Implementing EKF with DMC	106
3.2.4	Testing Filter on Telemetry with Simulated Debris Strike	109
3.2.5	Response of EKF with DMC to debris strike	111
3.2.6	Filtering Backward in Time	114
3.2.7	Derivation of Process Noise Covariance for Mirrored FOGM in Backward Filter	118
3.3	Developing Test Statistics	121
3.3.1	Test Statistic 1: Difference between forward and backward state estimates . .	121
3.3.2	Test Statistic 2: Mahalanobis distance	122
3.3.3	Test Statistic 3: Smoothed accelerations from fused backward and forward state estimates	123
3.3.4	Test Statistic 4: McReynold's filter-smoother consistency test	125
3.3.5	Test Statistic 5: Measurement residuals from propagated state estimate . . .	126
3.3.6	Test Statistic 6: Change in orbital constants	128
3.4	Tuning, Comparison, and Performance Assessment	130
3.4.1	Normalizing Test Statistics as SNR	130
3.4.2	Tuning τ and q	132

3.4.3	Performance When Detecting and Estimating Unknown Strikes	134
3.4.4	Performance in the Presence of Orbit Perturbations	135
3.4.5	A final comparison between measurement residuals and test statistic output .	139
3.5	Discussion	139
3.6	Conclusions	142
4	Applying Techniques to Telemetry from Active Satellites	143
4.1	Applying Techniques to Various Spacecraft	143
4.1.1	Telemetry Processing Architecture	145
4.1.2	Magnetospheric Multiscale Mission (MMS)	151
4.1.3	Mars Atmosphere and Volatile EvolutioN (MAVEN)	157
4.1.4	Lunar Reconnaissance Orbiter (LRO)	162
4.1.5	Fermi	176
4.1.6	Solar Dynamics Observatory (SDO)	178
4.2	Discussion of Lessons Learned	194
5	Development of a Perturbation Rate Assessment Tool	197
5.1	Context for Development Perturbation Rate Assessment Tool	197
5.2	Overview of Tool	200
5.3	Implementation Details	200
5.3.1	Definition of Frames Used for Perturbation Assessment	201
5.3.2	Modeling Strike Effects	203
5.3.3	Drawing Strike Parameters from Relevant Distributions	204
5.3.4	Implementing Trades, Assessing Results	215
5.4	Sample Results	218
5.4.1	Trading spacecraft size and orbit	219
5.4.2	Trading ORDEM uncertainty bounds and orbit	222
5.4.3	Trading MEF and L_c -to-mass conversion	223

5.5 Discussion	225
6 Conclusion	227

Tables

Table

1.1 Publicized debris strikes which perturbed but did not cause mission loss	43
2.1 Parameters used in spacecraft dynamics simulation	64
2.2 Particle impact events which impart 80 mN.m.s of angular momentum (assuming normal incidence)	65
2.3 Comparing error in estimated strike magnitude for each estimation method	80
3.1 Table of default parameters used in simulations throughout unless otherwise specified	106
4.1 Changes from initial to second generation telemetry processing architecture	146

Figures

Figure

1.1	Graphics published by ESA [2] showing effects from debris strike. Solar array was damaged and various telemetry points showed strike effects, but satellite operations continued without significant impacts	2
1.2	Graphic published by NASA's ODPO [3] showing number of tracked objects in orbit over time. Annotation 1 denotes the 2007 Chinese ASAT test, 2 denotes the Iridium-Cosmos collision, and 3 denotes the recent Russian ASAT test	4
1.3	USSF tweet regarding COSMOS 1408 breakup	5
1.4	Base graphic published by NASA ODPO [3] showing measurements of debris from Cosmos-1408 debris cloud, annotated with tracking capabilities and full estimated debris population (per ESA [5]) plotted on 2nd axis.	6
1.5	Graphic from NASA ODPO showing sources of data on which ORDEM is based. Some annotations added to highlight interpolation region and GEO debris [7]	9
1.6	Hypervelocity impact test against two cubesat mockups, [10]	10
1.7	Debris impacts are not always catastrophic, effects vary widely and can be benign .	11
1.8	Graphic published by Iridium[11] showing analysis results from Iridium-Cosmos collision with new system	12
1.9	DebrisSat: Hypervelocity impact test of large non-trackable object vs. smallsat target	14
1.10	Two hypervelocity impact tests against energetic spacecraft components: pressure vessel (left) and battery (right)	15

1.11 Schematic of Bumper’s inputs and analysis process for assessing spacecraft risk . . .	18
1.12 Illustration from [1] showing the process for determining rates of spacecraft anomalies	20
1.13 Illustration comparing model output with various shape assumptions to on-orbit data from [17]	22
1.14 Plotting the DebrisSat and SOCIT test fragments’ characteristic length vs. mass. Base graphic from [12], high-density sphere annotations added	23
1.15 Summarizing influences on risk to spacecraft from HNT debris. Boundaries are approximate.	25
1.16 Conceptual diagram organizing debris-related activities	28
1.17 Debris impacts spacecraft, causes rotational and translational momentum transfer .	33
1.18 Prisoner’s dilemma: currently a suboptimal data sharing norm is in effect	38
1.19 Various existing solutions for impact detection. Left: experimental result from Fraunhofer EMI investigating the ability of antennas to detect plasma clouds from hypervelocity impacts [24]. Center bottom: diagram of Invocon’s Distributed Impact Detection Sensor [23]. Center top and left: DLR’s solar-array based impact detector [29]	41
1.20 Graphic published by ExoAnalytic Solutions[33] showing filter residuals with and without an estimated MITE	44
1.21 Examples of BLEs and test results. Figures from Schonberg [45]	47
1.22 Examples of micrometeoroid flux and interplanetary dust impact data. left: [26], center and right: [47]	49
1.23 Examples of hypervelocity impacts, showing backscattered ejecta plume forming which enhances momentum, proportionate to amount of excavated material	51
1.24 Investigating derivation of ‘known’ MEF from Giotto encounter and comparing un- derlying assumptions with newer test data	53
1.25 Outline of Research	54
1.26 Comparing first generation and updated Iridium satellites	59

2.1	Response of spacecraft to simulated debris strike (60 mg at 8 km/s, 1 m from CoG, applied at $t=50$ s. $H_{\text{strike}} = 0.48$ N.m.s).	66
2.2	Example of matched filter accentuating a 50 mg debris strike applied at $t=50$ s. Note that when the strike is applied the telemetry increases slightly for a few points then decreases slightly, but it is difficult to see in the raw telemetry as it is beneath the noise floor. However, the filter produces a distinct spike in response.	67
2.3	Matched filter developed for each spacecraft axis. Differences in inertia result in slightly different wavelets for each axis.	68
2.4	Probability density functions for filter output	69
2.5	Highest LLRT curve corresponds to most likely strike size	70
2.6	Receiver Operating Characteristic curve for axis 1	71
2.7	LLRT curves with specified detection threshold and bins	71
2.8	Estimated strike size via MMSE vs. true strike size	72
2.9	Percentage of strikes detected vs strike size	73
2.10	Probability density functions of filter output with no strikes present and during a small strike	74
2.11	Development of filter thresholds via ROC curve and LRT	75
2.12	Example of the spacecraft's correction after a small maneuver imparts a slight rotation. Note similarities to simulated debris strike response.	77
2.13	Wavelet bank and results from running filters against a 50 mg strike at $t=30$ s. Minimum error is selected as most probable strike size (48 mg). Results from each estimator are compared.	79
2.14	Detection results per attitude matched filter and rate matched filter, and plot of estimated vs. true size. Non-detected strikes and false alarms shown as zeros on each axis.	81
2.15	Raw momentum telemetry with 10 mg strike at $t=500$ s	82
2.16	Output of SPRT filters with 5, 10, 15, and 20 mg strikes applied to the x-axis	83

2.17	Q-Q plots illustrating results when fitting various distributions to filter output . . .	85
2.18	Detection performance of each filter on 1,000 randomized strikes	86
2.19	True size of each strike in 1,000 run Monte-Carlo compared to estimated size using two estimation methods	88
2.20	Process for determining ΔH threshold. Note that this maps the filter threshold, which changes based on filter parameters, to a ΔH threshold which can be compared between filters and spacecraft.	93
2.21	Trends in detection threshold for various S/C sizes	95
2.22	Trends in detection threshold for various telemetry rates	96
2.23	Trends in detection threshold when trading telemetry noise	97
2.24	Trade noise in individual telemetry points (all data from SDO)	98
2.25	Trading miscellaneous spacecraft parameters	98
3.1	Table illustrating myriad situations which could produce a detectable ΔV for a hazardous non-trackable piece of debris. Scale refers to debris strike magnitude as specified throughout chapter. Debris listed as ‘potentially trackable’ is NOT tracked currently, but may be tracked as current state-of-the-art systems mature.	105
3.2	Performance of forward filter on telemetry with small debris strike	110
3.3	Measurement residuals from filter during debris strike show a minute feature which is not apparent via casual observation	110
3.4	Illustration and example of FOGM dynamic model compensation responding to im- pulsive ΔV	113
3.5	Example of changing q to tune process noise covariance matrix. Note scales on Y-axis indicating scale of unmodeled accelerations relative to Figure 3.4(b)	114
3.6	Comparison showing unmodeled acceleration estimation and propagation at the be- ginning of identical forward and backward filters (Note: This shows the filter initial- ization, no debris strike is shown)	116

3.7	Traditional FOGM used in forward filtering and mirrored FOGM for backward filter	117
3.8	Performance of backward filter on telemetry with small debris strike	118
3.9	Difference between forward and backward EKF state estimates shows significant spike in velocity without requiring knowledge of truth state	122
3.10	Mahalanobis distances	123
3.11	Smoothed state exhibits signature in unmodeled acceleration	124
3.12	McReynold's consistency statistic shows feature in all nine states. Red indicates scalar consistency statistic from Equation 3.60 plotted on separate axis to compare to each state from Equation 3.59	126
3.13	Residuals generated by comparing forward and backward orbit propagations to measurements	127
3.14	Orbital constants calculated from forward-filter state estimates show change at time of strike	129
3.15	Depiction of process for setting detection thresholds. Left: goal is to set threshold above baseline noise, which varies. Center: Probability density functions of filter output with and without debris strike indicate that threshold location is a trade between probability of false alarm and probability of missed detection. Right: Kernel distribution fitted to histogram of filter output with no debris strike used in conjunction with desired false alarm rate to establish threshold.	131
3.16	Trading τ and q to determine effect on SNR of various test statistics.	133
3.17	Superimposing SNR from various test statistics to compare results for given time constant τ	134
3.18	Monte-Carlo analysis indicating estimated size of random strike relative to true strike. Color indicates number of test statistics that detected each strike.	135

3.19	Simple perturbation models (<i>J3</i> and drag) added to truth state but not filter dynamics show that baseline filter noise is higher, but strikes are still detectable. This plot shows a 400 km orbit and a debris strike that is 2X the magnitude of the default strike	137
3.20	Superimposing SNR from various test statistics to compare results for given time constant τ . Note that preferred q is now larger than the trade without perturbations, and the SNR is smaller even with a larger strike	138
3.21	Trading τ and q again, but with additional unmodeled perturbations (<i>J3</i> and drag)	138
3.22	Comparing test statistic output to traditional measurement residuals in the presence of large and small debris strikes. Test statistic shows strong signal in response to strike far too small to see in measurement residuals.	140
4.1	Original telemetry processing architecture	147
4.2	Second generation telemetry processing architecture	148
4.3	Diagram of major code functions, indicating which are general between all spacecraft and which are developed specifically for each spacecraft.	150
4.4	MMS spacecraft formation studies electromagnetic reconnection	151
4.5	Effects of mission events on MMS inertial angular momentum	152
4.6	Results from MMS-4 telemetry in 2016	154
4.7	Flight regions of MMS	155
4.8	Results from MMS-4 telemetry in 2016	156
4.9	Example filter output from MMS-2 showing debris strike and other events	156
4.10	Results from MMS-2 telemetry in 2018	157
4.11	MAVEN spacecraft	158
4.12	MAVEN's RW telemetry	159
4.13	Running provided signatures against raw LPW data at time of known impact	160
4.14	Zooming out to entire day of data shows major features besides impacts	161

4.15 LRO spacecraft	162
4.16 LRO HGA in integration and test - annotations show gimbal rotation	164
4.17 Reaction wheel rotating about its center of mass, for calculating angular momentum contribution about point O.	166
4.18 Angular momentum of appendage about point O due to rotation about point A . . .	167
4.19 Schematic for calculating angular momentum of dual gimbal assembly with two massive elements in motion	168
4.20 Momentum contributions and raw telemetry for LRO - before pre-processing	170
4.21 Patching LRO wheel speed telemetry during zero crossings	171
4.22 Sample telemetry pre-processed with RW correction algorithms	172
4.23 Momentum contributions for LRO - after pre-processing RW speed measurements .	173
4.24 Momentum contributions for LRO during slew maneuver	174
4.25 Fermi spacecraft	176
4.26 Preliminary results from Fermi telemetry	178
4.27 SDO spacecraft	179
4.28 Filtering SDO angular momentum telemetry to remove periodic signatures	180
4.29 Daily number of perturbations tallied throughout year in 2016 and 2018	182
4.30 Results from SDO telemetry in 2016	182
4.31 Calculating square wave to fit telemetry and removing to produce quiescent data . .	184
4.32 Example of abrupt change in angular momentum	186
4.33 Graph of SDO's minor perturbations per day of year shows some correlation to micrometeoroid populations	187
4.34 Changes in flux of micrometeoroids, top is contribution from showers, bottom shows flux in sporadic background	188
4.35 Fitting various distributions to the filter output, and effect on thresholds	190
4.36 Daily trends in SDO inertial angular momentum telemetry	192
4.37 Zoomed in look at filter returns and data	193

4.38	Zoomed further to see small change	194
5.1	Graphic from [17] showing that the model matches the data fairly well	198
5.2	Schematic of architecture to assess expected perturbation rates	200
5.3	Coordinate frames used in perturbation prediction architecture	202
5.4	Drawing random number for velocity magnitude from ORDEM distribution	206
5.5	Velocity distribution from ORDEM	206
5.6	Characteristic length to mass mapping with various models	209
5.7	Step 1: Bus point cloud specifying surfaces	212
5.8	Step 2 and 3 of finding strike location on spacecraft	213
5.9	Step 4 and 5 of finding strike location on spacecraft	214
5.10	Step 6: Identifying final strike location on spacecraft	215
5.11	Plotting results from trade study, indicating relationship between change in angular and linear momentum	220
5.12	Comparing annual perturbation rate for each spacecraft and each detector	221
5.13	Comparing annual perturbation rate with ORDEM's error bounds	223
5.14	Comparing annual perturbation rate with various MEF and mass assumptions	225

Chapter 1

Introduction

1.1 Prelude

In April of 2017 Donald Kessler, the retired senior scientist from NASA's Orbital Debris Program Office, gave the keynote address at the 7th European Conference on Space Debris. He praised the worldwide efforts in space debris research, but said there were still gaps in research and technical capabilities that he was concerned about. In his presentation, he highlighted his concerns for the assembled global community of debris researchers.

In a nutshell, the issue is hazardous non-trackable (HNT) debris. This is debris that is too small to track and therefore spacecraft cannot maneuver to avoid it, but it is still large enough to damage a spacecraft. A substantial percentage of debris falls into this category, as less than 10% of the hazardous population is trackable. Kessler discussed a study by the NASA Engineering and Safety Center (NESC) which compared predictions of risks to on-orbit assets to anomalies experienced on orbit. The debris environment model is used to predict expected rates of perturbations and anomalies, and these predictions are compared to data from on-orbit operations. The results are hardly reassuring, with predictions sometimes differing from measured events by an order of magnitude [1]. In one case study an adjustment to two model parameters, the momentum enhancement factor and the assumed shape of the debris, is used to adjust the model output to obtain a much better match to the dataset.

In his keynote, Kessler discussed the methods that were used in the NESC report. In one study abrupt, unexpected momentum perturbations were cataloged by a spacecraft operator and

compared to model predictions. The NESC report consisted of a handful of limited datasets, but Kessler emphasized the implication that “Over 100 debris-producing collisions per year” are expected and “The environment could be measured by observing momentum changes . . . on existing intact population or possibly by ground sensors.” He concluded that “Monitoring the changing debris environment over all size ranges . . . continues to be a pressing need!”

A few months before the conference the European Space Agency had announced a debris strike on their Sentinel-1A spacecraft. Like many of the events cataloged in the NESC report, the strike did not cause substantial damage to the spacecraft and operations continued nominally. The operators shared telemetry, shown in Figure 1.1, which indicated an abrupt change in rotation rate that was corrected by the attitude control system. An orbit change was also observed. The Sentinel-1A spacecraft happened to have a camera for observing the solar array deployment, and this camera was used to verify the strike by observing damage to the array.

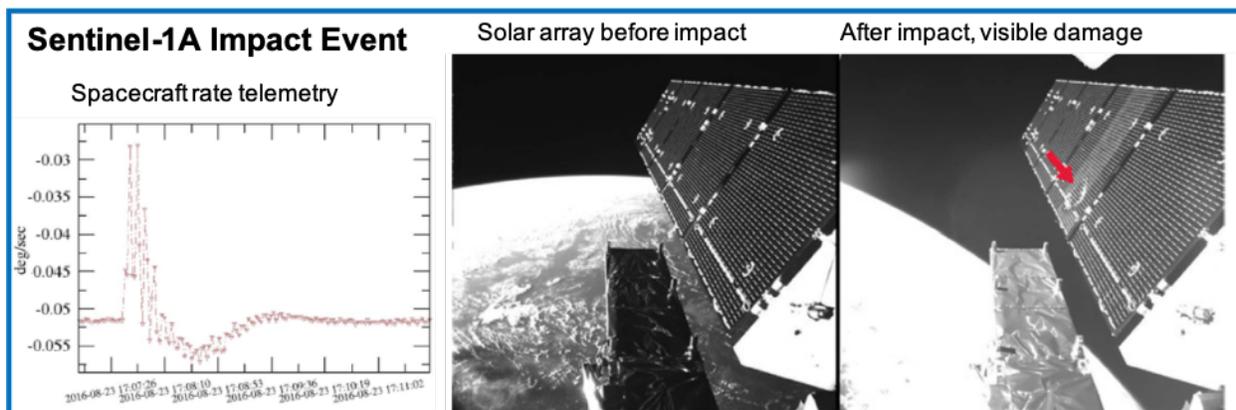


Figure 1.1: Graphics published by ESA [2] showing effects from debris strike. Solar array was damaged and various telemetry points showed strike effects, but satellite operations continued without significant impacts

The intention of this research effort, accepting Kessler’s challenge, is to develop methods to detect minor debris strikes on spacecraft using typical spacecraft telemetry. In studying typical spacecraft telemetry this research utilizes telemetry collected by most spacecraft, like attitude and orbit information; this approach is fundamentally different from developing specialized hardware to detect debris strikes. The Sentinel-1A telemetry provides an example of the type of features that are

expected in the event of a strike, and the NESC study illustrates an example of model adjustments that could be supported by data collected using these methods. At the 2017 conference several presentations expressed consternation over the proposed megaconstellations—a reality that we are now experiencing in 2022—and concerns about the implications for the debris environment. While these concerns still must be addressed, this research presents a silver lining to the emerging space era of proliferated satellite constellations: the satellites themselves may be used as sensors to monitor the small debris environment, and a validated understanding of the small debris environment can incentivize all actors to take appropriate actions for space sustainability.

1.2 Hazardous Non-Trackable Debris: an Overview

1.2.1 Origins of Debris

A growing public awareness of space debris provides some entertaining (and often frustrating) reading to technical professionals in the field. The articles typically begin by addressing where space debris comes from – often a cursory mention of paint flecks and collisions is made along with some fun whizz-bang facts: one time an astronaut lost their glove! Maybe some day we won't be able to launch through the 'debris cloud'! In the movie *gravity* (2013), a Russian anti-satellite (ASAT) test immediately sends large pieces of debris visibly whizzing past and impacting the International Space Station.

As usual, the reality is perhaps less glamorous but also far more complicated. The existing population of debris, shown in Figure 1.2, comes mostly from fragmentation events – human spaceflight activities are not a meaningful contributor. The mechanics of fragmentation events are discussed in Section 1.2.3, but the main contributors (roughly in order) are ASAT tests, the explosion of derelict satellites and rocket bodies, one major on-orbit collision, and several minor events that might be collision or explosion. Non-fragmentation debris is typically larger and easier to track, consisting of old rocket bodies, defunct spacecraft, and mission-related debris, i.e, debris that is intentionally released. Mission-related debris has been more-or-less static for over 20 years,

space players have mostly learned not to release intentional debris. Contributions from unintentional explosions leveled off in the late nineties as the global space enterprise learned to passivate spacecraft at end of life, dumping fuel and discharging batteries to avoid future explosions. In 2009 an active Iridium satellite collided with a defunct Russian satellite, creating over 2,000 pieces of tracked debris but also sparking a new era of collaboration between operators and government to ensure that the most relevant information is shared and tools are used to avoid collisions. With proliferated constellations space traffic management remains an ongoing concern, but collisions between well-tracked objects and maneuverable spacecraft can generally be considered avoidable.

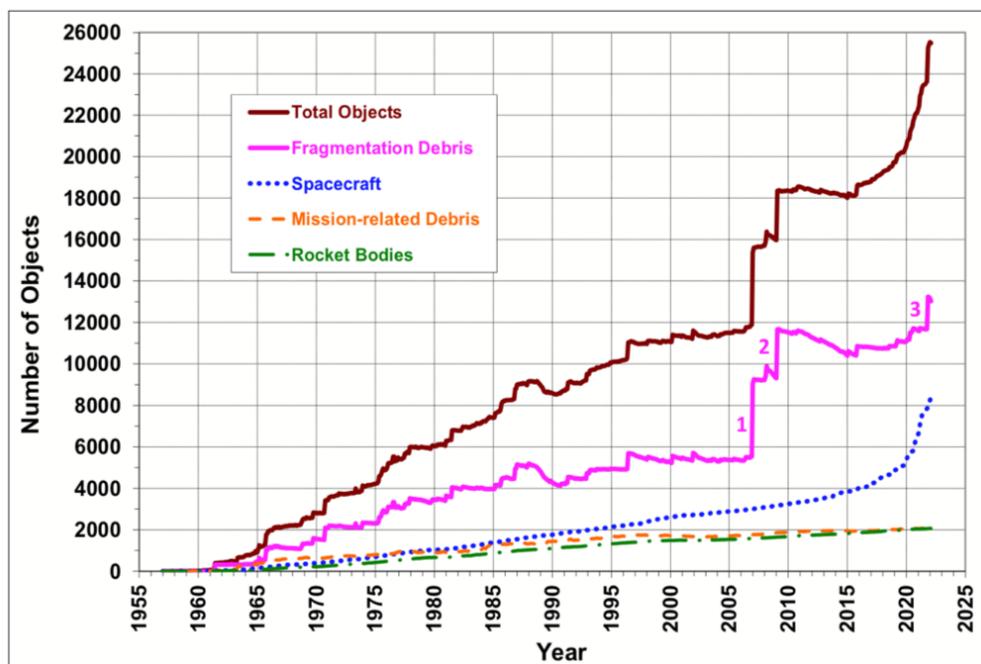


Figure 1.2: Graphic published by NASA’s ODPO [3] showing number of tracked objects in orbit over time. Annotation 1 denotes the 2007 Chinese ASAT test, 2 denotes the Iridium-Cosmos collision, and 3 denotes the recent Russian ASAT test

This leaves three problems: collisions with untracked (or poorly tracked) debris, collisions between derelict objects, and destructive ASAT activities. This research focuses on the collisions with untracked debris, and untracked debris mostly comes from fragmentation events. To start, consider a recent high-profile fragmentation event: The Cosmos-1408 Russian ASAT test.

1.2.2 Fragmentation Event Case Study: Cosmos 1408 ASAT test

On November 15th 2021 the USSF 18th Space Defense Squadron issued the tweet shown in Figure 1.3. After the event, NASA released a statement indicating that the ISS crew closed airlocks and took shelter in their spacecraft for the first couple orbits. NASA states that “As the Cosmos 1408 fragment cloud dispersed, risks to the ISS decreased and reached a quasi-stable state, which was about two times the pre-ASAT level” [3]. Com-



Figure 1.3: USSF tweet regarding COSMOS 1408 breakup

mercial and amateur observers began sharing their observations and analyses of the event, while the resultant trackable debris gradually showed up in the public catalog. The Commercial Space Operations Center (COMSPOC) predicted based on analysis that the event would produce ‘conjunction squalls’ for active satellites, and within a few months these conjunction squalls were observed, as the planes of active satellites began periodically lining up with the trackable debris’ orbits.

While the event was extremely concerning, one silver lining in the reporting was that several sources specifically mentioned the small debris hazard. Notably, in a State Department press briefing the spokesperson stated “The test has so far generated over fifteen hundred pieces of trackable orbital debris and hundreds of thousands of pieces of smaller orbital debris that now threaten the interests of all nations” [4]. This high-profile acknowledgement of HNT debris is refreshing, because historically discourse has focused on the trackable debris, while the untrackable debris presents a much more significant risk to spacecraft.

1.2.3 Fragmentation Event Mechanics

To illustrate, note that Figure 1.2, which is a classic figure used in many presentations at debris conferences, only represents tracked debris. Unfortunately, fragmentation events, either

collisions or explosions, create debris such that the relationship between the number of pieces and the size of the debris tends to follow a logarithmic function, i.e, **far** more small pieces are created in an event than large pieces. NASA's Standard Satellite Breakup Model (SSBM) is typically used to model these fragmentation events, and is shown applied to the Cosmos-1408 event as the black line in Figure 1.4.

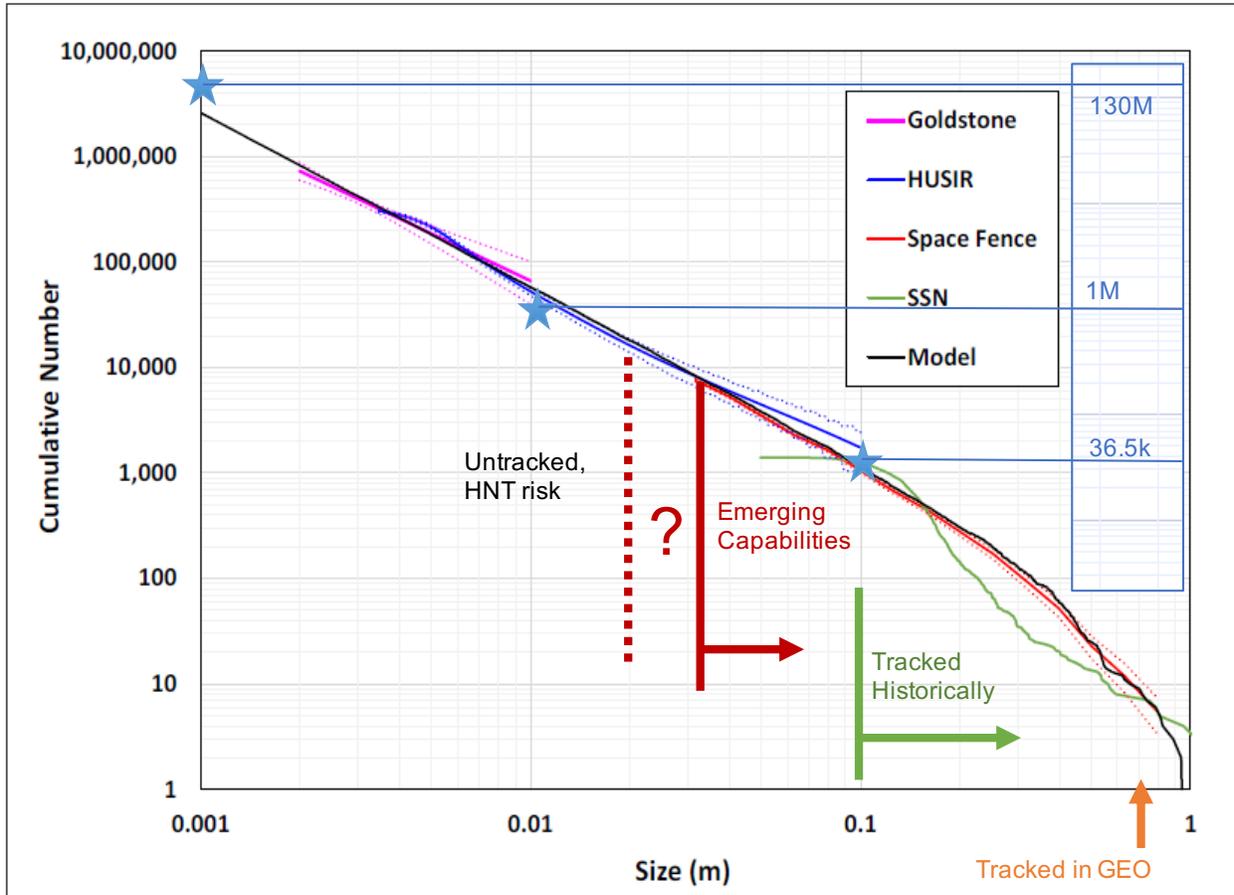


Figure 1.4: Base graphic published by NASA ODPO [3] showing measurements of debris from Cosmos-1408 debris cloud, annotated with tracking capabilities and full estimated debris population (per ESA [5]) plotted on 2nd axis.

Note the numbers on the horizontal and vertical axis of the graphic - with about 1,500 pieces of trackable debris, the event likely created 200,000 pieces 5 mm or larger. While the black line shows the SSBM model results, the green line represents the legacy tracking system the Space Surveillance Network, with a clear sensor roll-off causing the count to plateau at around ten

centimeters. Somewhat better, the red line shows the results from the newest ground radar of the US Space Force, Space Fence. This represents the current state of the art in small debris tracking, although one commercial operator (LeoLabs) describes a tracking capability down to 2 cm. Below these thresholds debris is untracked, so active satellites cannot be maneuvered to avoid it. This population, therefore, represents the hazardous non-trackable (HNT) debris risk. The blue and pink lines represent two radars which are used for characterizing but cannot track debris – these are discussed further in the next section.

Figure 1.4 also shows some aggregate debris population numbers plotted on a second log axis. The decades of the axis (in blue) are offset from the primary axis, but the size of each decade is matched (i.e, the whole line would be moved up by about an inch to plot on the same chart, but it is lowered to see the correspondence to the slope of the Cosmos-1408 debris). These debris population numbers are published routinely by the European Space Agency[5], these estimates are current as of 10 May 2022. These proportions provide a useful example of the quantity of small and large debris in a typical fragmentation event, although there may be additional sources for the small particles, like material degradation or solid rocket motor slag.

This chart highlights an important nuance in terminology: the difference between tracking debris and characterizing debris. The green and red lines indicate debris sizes that are trackable by legacy and emerging systems. For a well-tracked piece of debris, the orbit is known well enough that active satellites can maneuver to avoid collisions. The debris too small to track must be characterized. Since the small debris cannot be avoided, the goal is to design and operate spacecraft such that the probability of a significant collision is below a certain threshold. This risk assessment is conducted using a combination of debris environment models, which predict the expected flux, and risk assessment models, which assess the risk of spacecraft damage from impacts. Section 1.2.7 discusses this process in more detail while Section 1.2.10 discusses how tracking and characterization fit into the bigger picture of debris efforts. The next section discusses the measurements that are used to produce these debris environment models.

1.2.4 Sources of Debris Measurements

The sources of debris measurements are an important consideration for understanding HNT debris. Figure 1.5 shows the heritage sources of data for NASA’s orbital debris environment model, ORDEM. The teal area shows the debris that is tracked. The tracked debris is included in ORDEM, although obviously it represents a very low flux and can be avoided by maneuverable satellites. Note the altitude on the (log) y-axis. The radar measurements and *in situ* measurements provide information for the models, but there is a data gap known as the ‘interpolation region’ between the *in situ* and radar measurements, in the 1-3 mm range [6]. At this size low-density materials do not present a serious hazard to well-designed spacecraft, but ORDEM models a significant population of high-density particles in this range which has a substantial impact on overall risk to spacecraft. Also note the dates on the sample return missions. Additional *in situ* measurements from the ISS have been conducted, but in general *in situ* data is often stale and is always constrained to the altitude of the object which was returned to the ground for analysis.

An even more concerning measurement gap occurs in the GEO region, at the top of the chart. The space surveillance network tracks debris down to about 70 cm [8], but the SSN data is augmented by the MODEST telescope measurements which can observe debris down to about 30 cm. The MODEST telescope measurements of GEO debris occurred in multiple campaigns, the first set from 2004 to 2009 and the second from 2013 to 2014. The first campaign was used to build the populations, which were then validated using the second campaign a few years later. Concerningly, in the years between the campaigns the debris populations increased noticeably, so much so that ODPO added two additional ‘simulated’ breakups to the model to account for the new debris observed in the second campaign[9]. Fortunately, a new GEO observation asset, ES-MCAT, is coming online which will provide additional debris data collection in the GEO regime on a more persistent basis. Unfortunately, this asset’s detection thresholds are only a modest improvement over the previously used MODEST telescope, so debris less than about 10 cm will still be almost entirely unobserved in the GEO region.

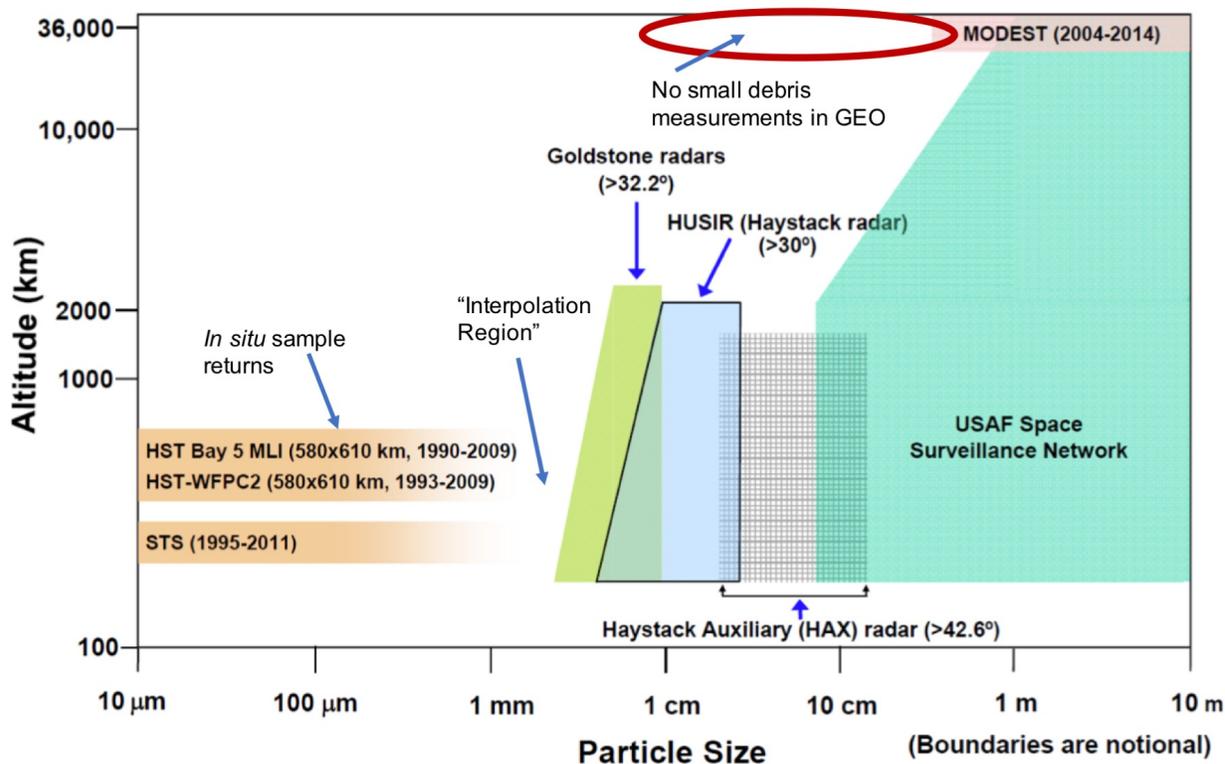


Figure 1.5: Graphic from NASA ODPO showing sources of data on which ORDEM is based. Some annotations added to highlight interpolation region and GEO debris [7]

"We can never measure what we want to measure"

-Mark Matney, NASA ODPO, 2019 presentation at Spacecraft Anomalies and Failures Workshop

As a final note about debris measurements, it is important to consider the specific measurements that are returned by each system. Radars send out pulses of radio waves and then receive a reflection, from which an object's radar cross section can be deduced, and from that an estimate of its size. Optical sensors observe light reflected from an object, likewise deducing an optical cross section that corresponds roughly to size. Returned samples are assessed for craters, and an estimate of the particle's size can be made but is likewise uncertain. Further analysis can ascertain the material composition of the impacting particle.

However, none of these measurements correspond directly to an object's ability to damage spacecraft, which is a complex, probabilistic interaction dependent on the particle's mass, shape,

material, size, and density. Therefore, assessing a spacecraft's risk of damage based on debris flux models which are built from radar, optical, and *in situ* measurements is an exercise which relies heavily on a long chain of assumptions. This is discussed further in Section 1.2.7.

To illustrate, Figure 1.6 shows debris resulting from two hypervelocity impact tests against cubesats, indicating an abundance of flat and/or thin fragments. These fragments are mostly CFRP and MLI, which are both low-density materials. An aluminum sphere, or perhaps even a low-density sphere, with an equivalent radar cross section would produce substantially more damage if impacting a satellite. This clearly indicates the complexity in shape, material, and aspect ratio of debris created in a fragmentation event with modern satellite construction.

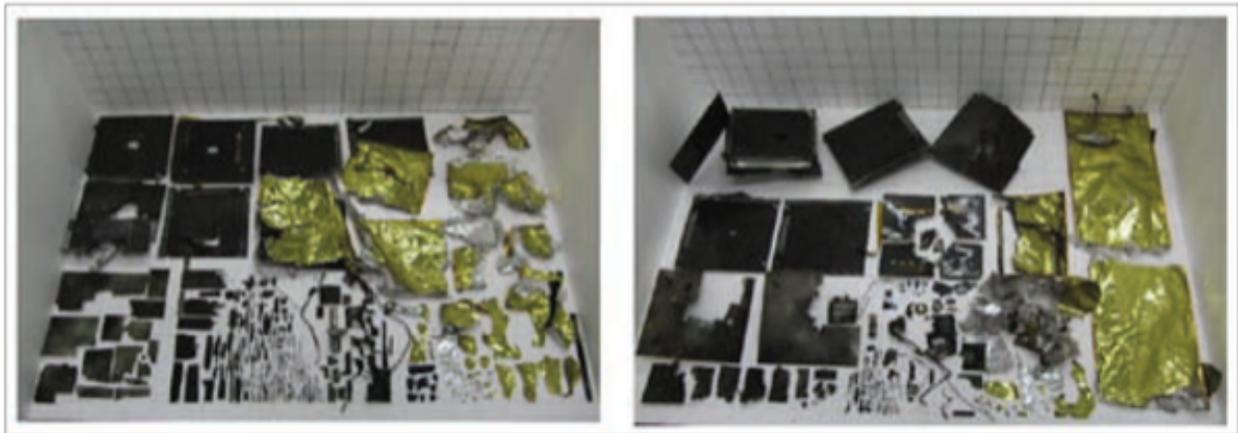


Figure 1.6: Hypervelocity impact test against two cubesat mockups, [10]

In ORDEM, the debris flux is returned in terms of characteristic length, since that is the variable which corresponds most closely to the available measurements. ORDEM then splits debris into various density classes, including high- medium- and low-density objects, but the decision of how to estimate damage from each characteristic length for assessing spacecraft risk is left to the analyst. Sections 1.2.7 and 1.2.8 discuss this topic in more detail. This brings up an important concept: determining what effect a piece of HNT debris might have on a satellite.

1.2.5 Classes of Debris Events

Determining what effects a debris impact will produce on a spacecraft is a complicated question, and the answer, as with many questions, is ‘it depends’. To begin, *in situ* sample return missions indicate vast numbers of minor strikes which did not affect mission performance. The Hubble solar arrays came down with more than 3,600 craters after flying for 4 years, while the Long Duration Exposure Facility (LDEF) had over 600 after 6 years. However, Hubble continued (and still continues) to function, 32 years later. Therefore, while objects on orbit are routinely impacted by small particles, the results are generally benign. Figure 1.7 shows LDEF and other examples of each class of debris events.

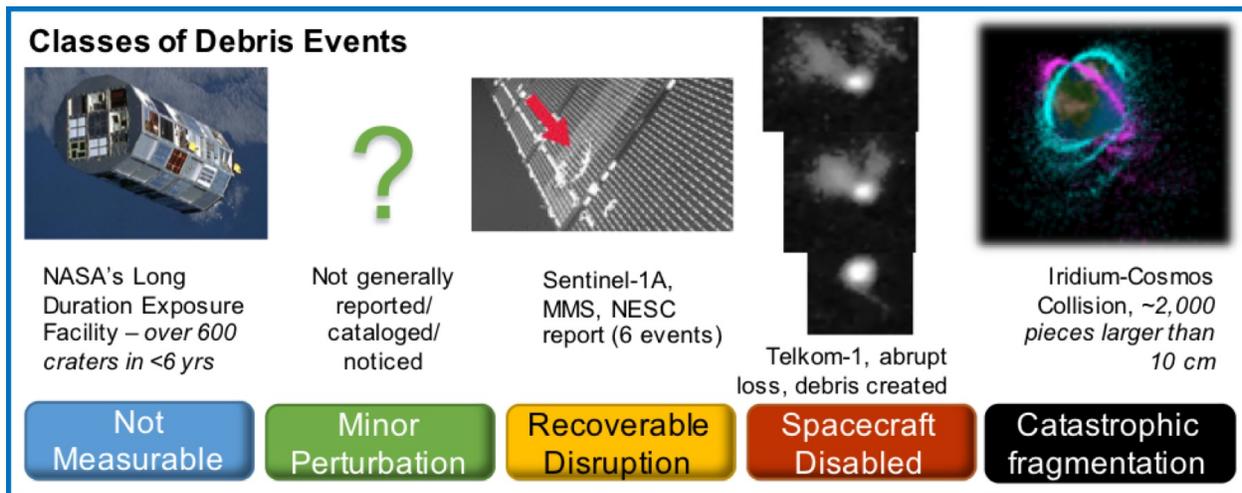


Figure 1.7: Debris impacts are not always catastrophic, effects vary widely and can be benign

Moving up the spectrum, the Sentinel-1A event provides one well-studied example of a debris impact that caused enough damage to produce off-nominal behavior such that the event was noticed, but nevertheless operations continued nominally. Other events in this class include strikes on NASA's Magnetospheric Multiscale Mission (MMS), 6 of the strikes cataloged in the NESC report, a strike on Maxar's Worldview-2 satellite, and a strike on NASA's Terra satellite. However, it's important to remember that not all anomalies get reported. In conversations with six major operators, some government and some commercial, I have not yet talked to any who have never

seen an event like this.

If a debris strike permanently disables a spacecraft it is often difficult to attribute that anomaly definitively. There are several abrupt major anomalies that might plausibly be caused by a debris strike, including on Telkom-1, AMOS-5, and Intelsat-29e. Sometimes debris is created during these events, and it was also created during the Sentinel-1A strike and the Worldview-2 strike. However, the quantity of debris is relatively low compared to events which completely fragment a satellite.

The primary example of a major unintentional fragmentation event is the Iridium-Cosmos collision in 2009. The operational Iridium satellite collided with a defunct Russian satellite, creating over 2,000 pieces of trackable debris and of course many more pieces of non-trackable debris. This event emphatically disproved the ‘big sky’ theory popular in the early days of spaceflight, that space was too big for satellites to collide and there are no concerns about creating debris. However, it’s important to note that the failure which led to this event was a matter of policy, not technology. At the time the US Air Force space tracking service only shared general perturbation ephemerides, not the more accurate special perturbation ephemerides. Also, operators did not typically share their maneuver plans, and in the Iridium-Cosmos event the spacecraft inadvertently maneuvered onto a collision trajectory during a routine station keeping maneuver. After this event policy was changed to share special perturbation ephemerides and maneuver plans, the resultant system returned a chance of collision (Pc) of 56% for this particular maneuver. The maneuver would never have been conducted. This distinction - that the largest accidental debris-causing event in spaceflight history was a policy failure, not a technical failure - is worth noting.

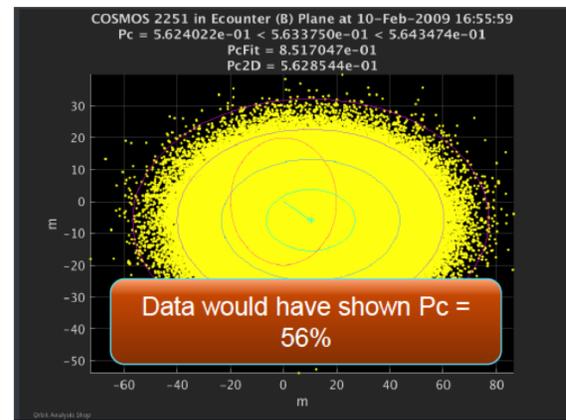


Figure 1.8: Graphic published by Iridium[11] showing analysis results from Iridium-Cosmos collision with new system

“U.S. Government policy did not then allow sharing SP data with commercial operators who had no choice but to accept the ‘Big Sky Theory’”

-Ryan Sheppard, Iridium SSA lead

International Conjunction Assessment Workshop, Paris, France, 2019

1.2.6 Effects of Untracked, Damaging Debris

With effective space traffic management preventing collisions with trackable debris, risk to active spacecraft comes from debris too small to track but still large enough to cause damage. The exact damaging boundary is difficult to specify precisely, especially for smaller debris sizes.

On the lowest end of the spectrum, impacts on *in situ* sample return missions have been carefully cataloged to facilitate understanding of the environment of very small debris. These sample return missions are fairly rare and have only been recovered from low orbits, mostly the 400 km ISS orbit, though a couple from the Hubble orbit were returned via the space shuttle. *In situ* sample returns inherently have a limited collection area, so data provided by these missions has a maximum size of around 1 mm or a little below. This extremely small population is unlikely to cause significant damage to a spacecraft.

On the other end of the HNT spectrum, a full scale hypervelocity impact test, DebrisSat, was conducted to characterize the fragments produced in an impact from a large non-trackable piece of debris (Figure 1.9). The projectile was a hollow aluminum cylinder, 8.6 cm x 9 cm, weighing 570 g. Note that this is below the trackable limit of the legacy space surveillance network but may be trackable at some altitudes with emerging capabilities. The target was a small 56 kg spacecraft model, with a high degree of fidelity in its construction including solar panels, instruments, multi-layer insulation, CFRP structure panels, etc. The impact occurred at 6.8 km/s, which was the maximum practical velocity for launching such a large projectile but still well below a typical LEO impact of 10-12 km/s [12].

The test resulted in a massive fragmentation event, producing far more debris pieces than predicted via the standard satellite breakup model (SSBM), which was tuned using fragmentation

data from legacy satellites. Newer construction involving CFRP and other lightweight materials tends to produce more fragments, as has been observed in on-orbit fragmentation events like the 2007 Chinese ASAT test. The SSBM predicted 85,000 fragments larger than 2 mm, as of FY19 193,000 fragments had been collected from the test. Concerningly, rigorous characterization of these fragments indicates a significant population of high-mass metal fragments as well as small lower density fragments. As of 16 Jan 2020, over 1,000 pieces heavier than 1 g had been cataloged, mostly metal [13].

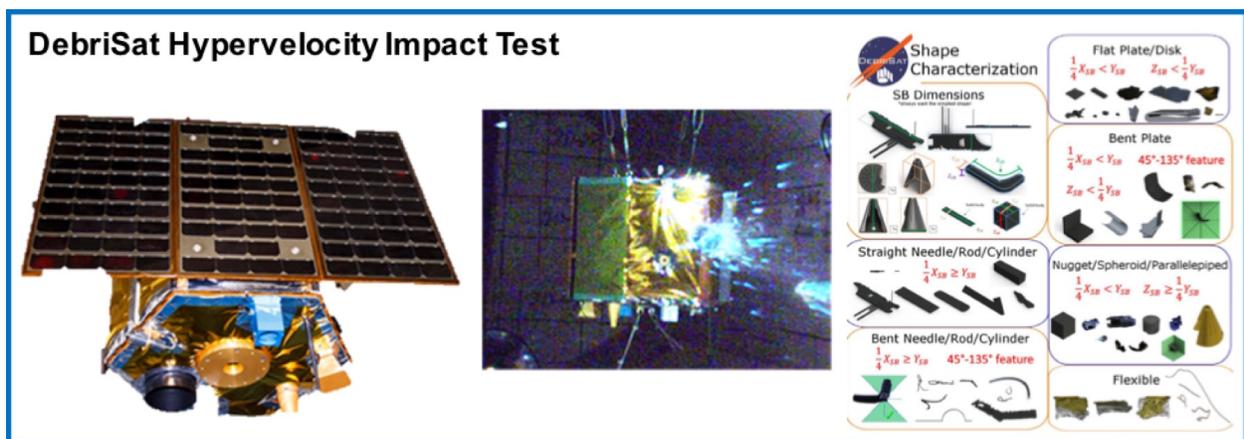


Figure 1.9: DebrisSat: Hypervelocity impact test of large non-trackable object vs. smallsat target

These two examples bound the problem of hazardous non-trackable debris. While spacecraft are impacted routinely by small objects the resultant damage is inconsequential. However, large objects just barely too small to track can cause catastrophic fragmentation events. As two further examples, consider Figure 1.10. On the left a composite overwrapped pressure vessel (COPV) tank experiences a rupture event when a 2 mm aluminum sphere impacts at 7 km/s. Note that this COPV tank, which might typically be used as a pressurant tank for a satellite's propulsion system, has no structure between the impacting particle and the exterior surface of the tank, as if the tank were mounted on the external surface of the satellite. This is thus a conservative worst case scenario, as placing a pressurant tank on the external surface of a satellite is a highly questionable design practice, although SMEs insist that some spacecraft fly like this. This serves as a useful

example of the importance of designing spacecraft with passive shielding, to protect the sensitive components of the satellite by placing them behind less-sensitive components or in places less likely to be struck.

The right side of the figure shows a Li-ion battery impacted by a 1 cm aluminum sphere at 6.86 km/s. In contrast to the pressurant tank, the battery is shielded by a typical structure panel and some layers of fabric to provide additional shielding. A typical spacecraft might have a structure panel like this with some exterior MLI, but no fabric. The test was designed to overwhelm the shield and assess the reaction of the batteries to an impact, and observed that “The impacted cell contents were energetically ejected ... Hundreds of metal fragments, many with dimensions greater than one centimeter, littered the target chamber floor after the test.” However, the batteries did not rupture or explode. Note that with the structure panel and some light shielding it took a much larger impacting projectile to produce a significant event from the energetic spacecraft component. Referring to Figure 1.4, the models predict roughly an order of magnitude more 2 mm debris pieces than 1 cm, so this improvement in shielding represents a significant reduction in risk.

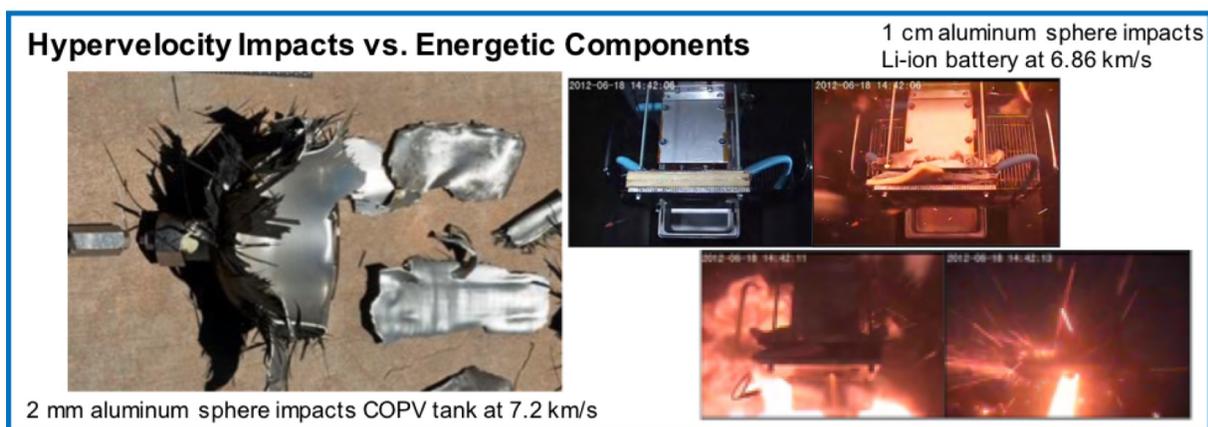


Figure 1.10: Two hypervelocity impact tests against energetic spacecraft components: pressure vessel (left) and battery (right)

Note that the impacting particle in both of these tests is an aluminum sphere, which brings up an important aspect of identifying the hazard that any given piece of debris presents to a spacecraft: A sphere maximizes the mass for a given characteristic length, and thus represents the

most conservative assumption. A scrap of MLI or CFRP which is 2 mm or 1 cm long will return a radar signature similar to an aluminum sphere of the same size, but would produce substantially less damage if impacting a spacecraft. Determining the exact extent of that damage is difficult, because launching a piece of MLI in a light gas gun would be a non-trivial challenge. However, debris researchers are moving out strongly to address these concerns. Recent efforts by NASA's hypervelocity Impact Technology (HVIT) team have involved launching cylindrical CFRP and high-density metal objects to study the effects of shape, and "On-going efforts within the NASA ODPO to update the next version of ORDEM include a new parameter that highly affects the damage risk – shape." [13].

The Sentinel-1A strike provides a useful example of this uncertainty. Based on the transferred momentum the calculated mass is 0.2 g, which would imply an aluminum sphere of diameter 5.2 mm. However, this projectile would tend to perforate the array [14], but camera observations indicate "A notable absence of visible penetration damage." [2]. The damaged area appears to be about 35-40 mm, but this may be due to spallation and cracking. NASA's SSBM predicts a range of about 4 mm to 40 mm for a debris object of 0.2 g, with a most probable object being around 1 cm. Thus, the Sentinel-1A strike was likely a larger, flatter projectile than an aluminum sphere, with less ability to perforate panels and damage the spacecraft than might be implied by the mass of the object. Hypervelocity impact test results as shown in Figure 1.6 remind us that larger, lighter debris pieces are to be expected, while the DebrisSat results remind us that nugget-shaped metal may also be present in abundance, especially for objects with larger masses. These examples illustrate the complexity of the problem, since the ultimate goal is to assess risks to spacecraft accurately to support and motivate safe, sustainable spacecraft operations, and all of these uncertainties complicate that task.

1.2.7 Uncertainties in Assessing Damage

A 2017 JPL paper [15] provides a useful study on typical conservative debris risk assessments and the limitations thereof. In this paper the authors conduct a risk assessment on an electronics

box that is known to be operational after 18 years on orbit. They assess various conservative and less-conservative assumptions and calculate the expected chance of failure using the then-current NASA debris environment model ORDEM3 and penetration risk assessment tool Bumper3-lite. The authors test various assumptions for assessing the risk, including treatment of the MLI covering the box as a dual-wall bumper vs. single wall, assuming that a detached spall causes failure vs. complete perforation, and assuming that an estimated 16% of the circuit card area is susceptible to failure vs. 100% of the area resulting in failure. The results show a widely varying chance of failure, with the most optimistic indicating a 20% chance of failure after 15 years and the most conservative assumptions producing an incredibly high 99.999% chance of failure for an electronics box that is known to be still functional. The authors conclude:

“Often when performing impact risk assessments on spacecraft, configurations are encountered with geometries and materials that fall outside the set of those included in the tests. This makes the applicability of the tools more uncertain, which often results in several plausible treatments of the problem. The predicted chance of failure can vary significantly from treatment to treatment, and the most conservative analyses often predict incredibly high probabilities of failure. Optimistic assumptions become necessary to achieve reasonable risks, potentially leaving the spacecraft vulnerable to failure.”

By a similar token, overly optimistic assumptions by emerging commercial companies may be used to check the box on licensing requirements to allow flight for a system that does not really meet the required specification. This may not even be intentional on the part of the operators, the point is that the correct answer is simply not known. At the Kinetic Space Safety Workshop (Lausanne, Switzerland, 2022) Prof. Shannon Ryan’s presentation provided examples even more compelling than the ones listed here. Transitioning to the next section Dr. Tim Maclay offered a tongue-in-cheek conclusion: “So the biggest risk is LNT, and the way to fix it is to change your model.” The question that remains is how to change the model - if the current assumptions are overly conservative, new methods are required to produce more accurate results without over-correcting the problem and underestimating risk to spacecraft.

*“It’s not based on measurements, it’s based on assumptions. How correct are those assumptions? *shrugs*”*

-Prof. Shannon Ryan, Deakin University

Kinetic Space Safety Workshop, Lausanne, Switzerland, 2022

The risk assessment tool used by the JPL authors (and many others in the space industry) is Bumper, produced by NASA’s hypervelocity impact technology (HVIT) team. A diagram of Bumper’s inputs, outputs, and analysis process is shown in Figure 1.11 (from reference [16]). Like any model, Bumper is only as good as its inputs. The right side of the figure shows the inputs traded by JPL to produce these drastically different results. The failure criteria were tuned - perforation vs. rear-side spall, 16% vulnerable area vs. 100%.

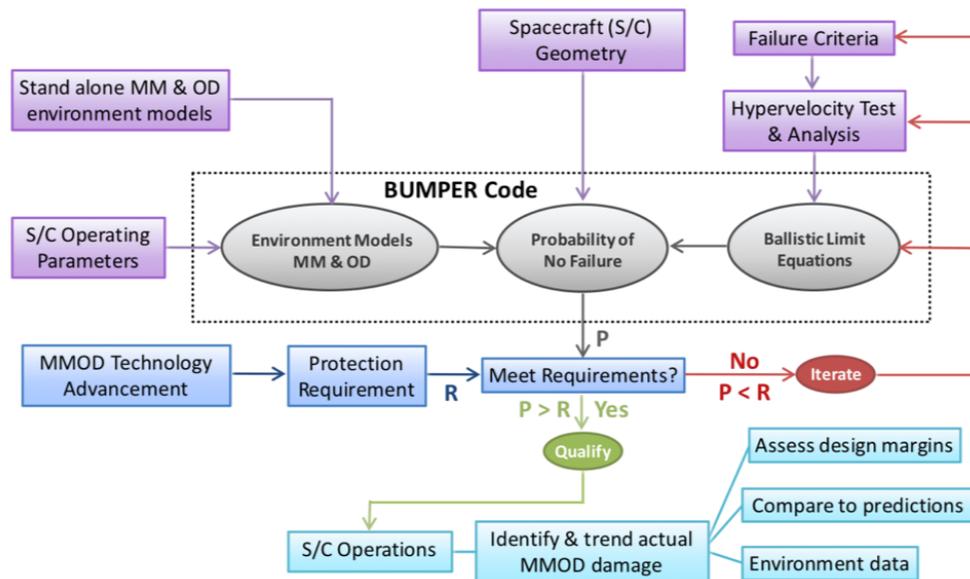


Figure 1.11: Schematic of Bumper’s inputs and analysis process for assessing spacecraft risk

Bumper relies fundamentally on Ballistic Limit Equations (BLEs) for assessing whether a given debris impact will damage the spacecraft. BLEs are described in more detail in Section 1.4.4.1, for the purposes of this discussion it is sufficient to understand that a BLE describes whether or not a certain projectile is expected to overcome a given target.

Bumper has an extensive library of BLEs, but the JPL authors discuss the limitations of

the available BLEs at length. In particular, they lament the lack of relevant test results for their spacecraft configuration. They use a fairly typical industry configuration, but test data on this configuration is limited, as many BLE developments are focused on shielding for human spaceflight. The authors also mention uncertainty in environmental models as a potential contributor in addition to the uncertainty in failure criteria and BLEs. One uncertainty contributor not addressed by the authors is that most BLEs are developed with spherical projectiles, which maximize mass for a given diameter, whereas in reality the shape of the object has a significant effect on its damage potential. Fortunately, HVIT is hard at work conducting hypervelocity tests with non-spherical, non-aluminum projectiles to support ORDEM's upcoming shape updates. However, as if things weren't already complicated enough, hypervelocity impacts are inherently probabilistic: the same impact might generate a different result if repeated, but since tests are expensive repetitions are rare, and BLEs approximate a deterministic solution to a probabilistic phenomenon.

As a final note, Bumper assesses the risk based on the provided environmental models, and can be used to assess micrometeoroid models as well as debris risks. Depending on the orbit, micrometeoroid risks are often more substantial than debris risks, but of course all spacecraft have been flying successfully in the micrometeoroid environment for all of spaceflight, and that environment will likely be stable in the future. Contrast this to debris, which is produced by human activity, and thus opportunities exist to shape the future debris environment in various ways. Micrometeoroid impacts can be difficult to distinguish from debris impacts, but some orbits and directions are more susceptible to one than the other. This topic of micrometeoroids vs. debris is referenced throughout, but some key notes are discussed in Section 1.3.2.1, after the research concept is introduced.

1.2.8 Lessons from Comparing Modeled Expectations to On-Orbit Experiences

With this background, it is time to circle back to the NESC report mentioned by Kessler at the 2017 debris conference, as described in the Prelude. This report introduced a few key concepts which motivate this research effort: First is the chain of uncertainties inherent in predicting spacecraft

risks, second is the dependence on shape for predicting anomaly rates, third is some findings regarding risk drivers, and fourth is a recommendation which this research is directed toward.

The NESC report provides the graphic shown in Figure 1.12 to illustrate the multistep process for assessing spacecraft risks, and one observation after all the above background information is that every single one of these steps provides the opportunity for uncertainty to creep in. With heritage assessment strategies the response to this uncertainty is typically to make conservative assumptions. This was an adequate approach when there were small numbers of satellites mostly financed by governments, but in the emerging space era it is insufficient for motivating adequate space sustainability practices. As the previous section illustrates, there are multiple plausible assumptions that can be made in spacecraft risk assessments, and without better data assumptions can be chosen to obtain the desired risk profile. Whether or not the calculated risk is accurate is an open question.

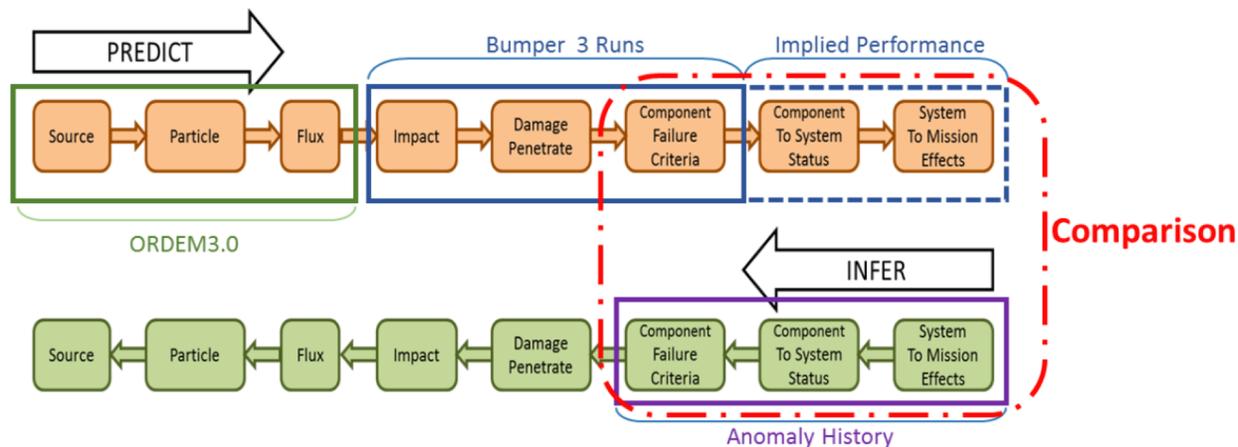


Figure 1.12: Illustration from [1] showing the process for determining rates of spacecraft anomalies

The bottom row of Figure 1.12 shows the process of attributing mission anomalies to debris. Comparing anomalies was the primary focus of the NESC report, the perturbation dataset was only one of several studies. As mentioned in Section 1.2.5, when an event produces an abrupt failure of the satellite it is often difficult to attribute the cause definitively. This may result in debris-related anomalies being undercounted, although the opposite error, misattributing other issues like

manufacturing defects to debris impacts, is certainly possible as well. An important note from this graphic is that perturbations occur ‘upstream’ from anomalies, so there will be higher rates of perturbations than anomalies, as not every impact will necessarily cause an anomaly. Some implications of this are discussed further in Section 1.3.2.2.

The second key finding, as pertains to this research, is the impact that shape assumptions have on the correlation between expectations and experiences. As previously mentioned, the report cataloged seven perturbations on a constellation, while the models predicted 24-164. This range in prediction exists because momentum enhancement factor (MEF) is highly uncertain. The MEF captures the effect that, in a hypervelocity impact, target material is expelled ‘backward’ from the impact location which enhances the total momentum delivered to the satellite - this concept is discussed further in Section 1.4.4.3. If an MEF of one is assumed (i.e, no momentum enhancement), the models predict 24 impacts. If MEF=2 is assumed, the prediction is seventy - a clear order of magnitude above the number of actual events. This is likely the closest to the true MEF value based on test and simulation results. Assuming MEF=3 to bound the problem results in a prediction of 164 perturbations.

However, even with an order-of-magnitude mismatch between expectations and experiences, it takes a surprisingly small change in the shape assumption to tweak the model output such that it matches the data. The NESC published a paper to accompany their report, and this paper shows (Figure 1.13) that by modifying the shape of the debris a much closer match between the model output and the data can be obtained. The black lines show various voided sphere assumptions, where the sphere is assumed to be hollow to decrease the mass per characteristic length. In the report, a more precise voided sphere relationship derived from the SOCIT debris tests also shows a reasonable relationship between ORDEM output and on-orbit data.

This is encouraging, and means that there is definitely progress to be made in tweaking assumptions that go into these risk assessment efforts to bring predictions closer to reality. But of course, not every piece of debris is a voided sphere, as shown in Figure 1.6, any more than it is a sphere. Debris comes in all sorts of shapes, and a given characteristic length can map to a

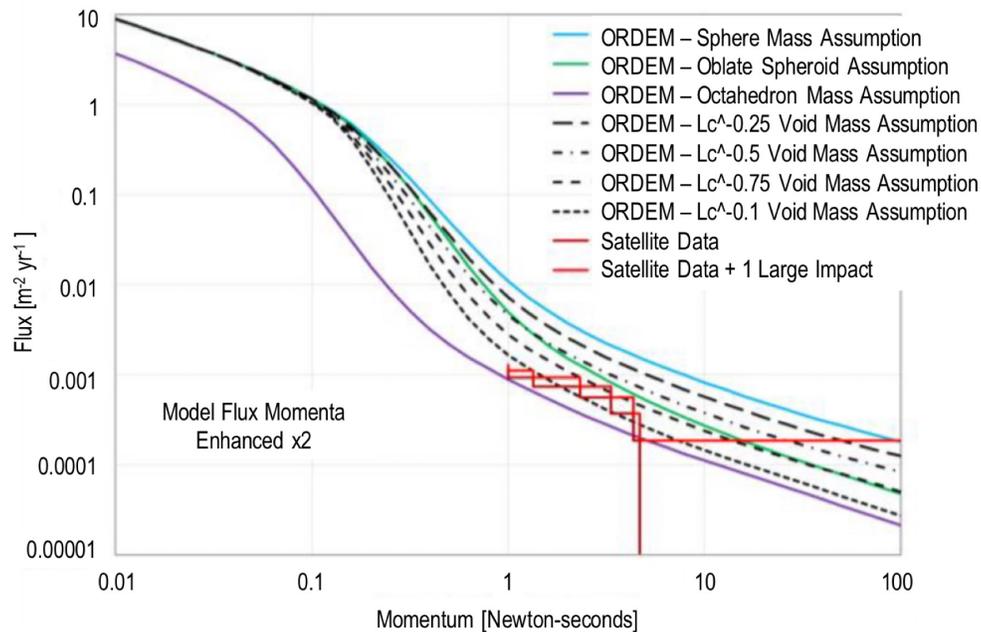


Figure 1.13: Illustration comparing model output with various shape assumptions to on-orbit data from [17]

wide range of shapes, masses, and therefore damage potential. Reality is more well-modeled by a Monte-Carlo analysis, but the parameter distributions to use for this Monte-Carlo analysis are not easy to come by. Figure 1.14 shows results from the DebrisSat experiment, which are an extremely useful guide in determining the fragment distribution for the breakup of a modern smallsat during a hypervelocity impact. It also plots (in red) the results from the older SOCIT test series, which conducted hypervelocity impacts against older styles of satellite construction.

These test results are a good starting point, but it would be useful to compare them to on-orbit data to validate that the debris effects experienced by spacecraft align with the expectations from these ground test results. To illustrate, when NASA upgraded the debris model to ORDEM 3.0 analysts found that it predicted substantially higher risks to assets - this motivated the JPL electronics box study [15] and the NESC report [1]. A major driver to these increased risks is a high-density population of very small pieces, 1-5 mm. Referring to the sources on which ORDEM is based, it's worth noting the lack of measurements in the 1-3 mm regime (see Figure 1.5). This

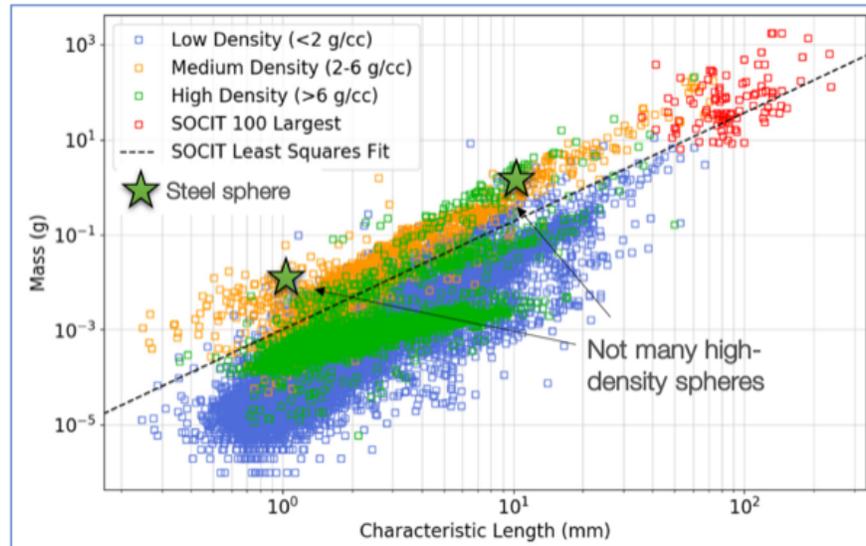


Figure 1.14: Plotting the DebrisSat and SOCIT test fragments' characteristic length vs. mass. Base graphic from [12], high-density sphere annotations added

steel population is based on *in situ* measurements of smaller particles, measured from sample return missions, and the model assumes that these smaller particles imply larger particles at higher altitudes. This is a pretty major assumption for something that affects the risk to spacecraft so dramatically, so one of the recommendations of the NESC report is to identify a source for this high-density population. This example illustrates that assumptions based on measurements, since “We can never measure what we want to measure,” may not be aligned with on-orbit experiences, similar to the JPL electronics box example which returned “incredulously high” probabilities of failure for a box that is known to be operational.

Based on the damage prediction uncertainties and associated potential corrections, the utility of the NESC report's recommendation is established: satellites can be monitored for momentum perturbations to collect more data about the HNT debris environment, and this data could be used to refine uncertain assumptions for damage analysis. This recommendation is the inspiration for this research.

1.2.9 Putting it All Together: Understanding HNT Debris Risks and Mitigation Limitations

Concatenating these concepts a few things become apparent. First, tracking capabilities are not sufficient to alleviate risks from debris. Due to much larger numbers of small objects than large, untracked debris represents more than 90% of the collision risk to spacecraft. Therefore, space traffic management efforts – maneuvering satellites to avoid tracked debris – are not sufficient to mitigate risks from debris.

Second, spacecraft cannot be fully shielded from this hazardous non-trackable debris. The International Space Station, which is the largest and likely most well-shielded object on orbit, can withstand impacts from debris up to about 1 cm for some modules. For a robotic spacecraft the risk threshold is much lower, and high density 1-3 mm particles are a major driver to overall risk [6].

Third, the local debris environment becomes dramatically worse every time there is a major on-orbit fragmentation event. This can be mitigated significantly with debris mitigation practices like passivating and deorbiting spacecraft at end of life, and continued efforts in space traffic management can (in theory) prevent collisions between active spacecraft and tracked objects. However, the large population of derelicts already on orbit results in a situation where, even with no additional launches, the debris population will continue to grow in some orbits. Active debris removal can mitigate this situation by removing high-risk derelicts, but this is a high-cost endeavour and thus needs a very compelling case to motivate this behavior. It is also a difficult, risky mission for the biggest and most critical derelicts.

Maclay and McKnight [18] developed a useful graphic for visualizing this, Figure 1.15 shows a slightly modified version of their concept. The horizontal and vertical axes illustrate size and number of objects, similar to the typical cumulative number graphics like Figure 1.4. The solid black line illustrates the baseline population of debris objects of each size. This population is currently expected to increase over time (upward arrow showing expected growth), but this may be

mitigated by active debris removal or just-in-time collision avoidance (JCA). However, if additional major breakups occur, the population could increase dramatically. Examples include kinetic ASAT activities (i.e, not effectively deterred), ineffective space traffic management, or derelict-on-derelict collisions. ADR/JCA are designed to mitigate this risk and keep debris population growth in check, but it's important to remember that two out of the three biggest debris-causing events were deliberate, and accept the geopolitical/military forces that must be applied to prevent those types of activities.

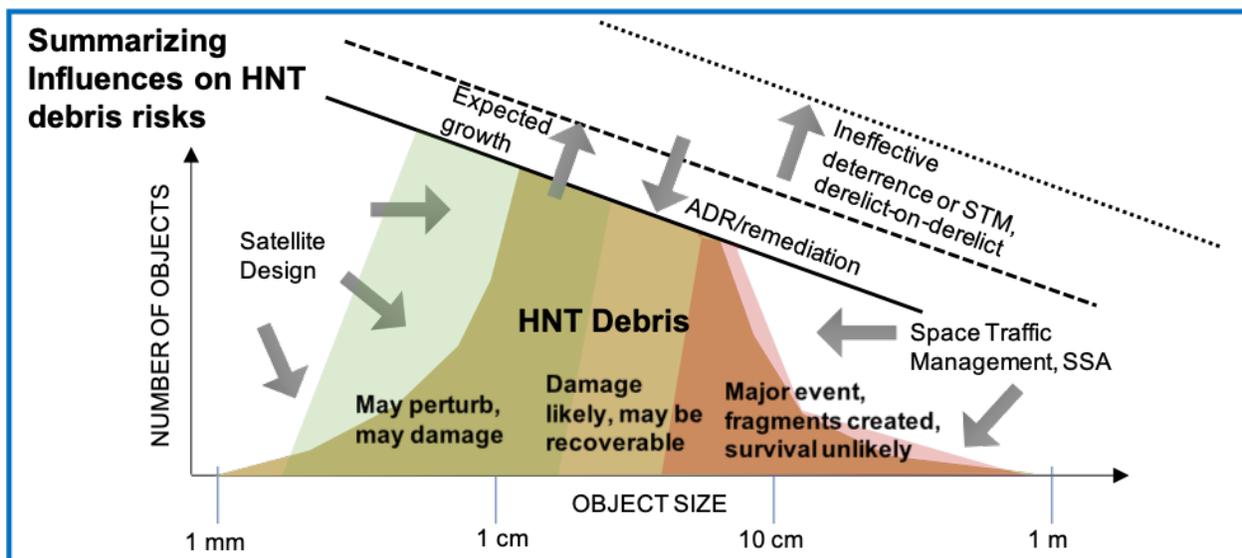


Figure 1.15: Summarizing influences on risk to spacecraft from HNT debris. Boundaries are approximate.

Beneath this line indicating the population are a few shaded areas representing various classes of debris. The large gold shaded area in the center is the hazardous non-trackable debris population. The curves on either side are not straight lines because, as previously discussed, a 5 mm piece of debris may or may not be damaging depending on the shape and material of the object. However, a very small steel or copper sphere may be damaging. Also, the impact location is important - if a spacecraft is flying with pressurant tanks, batteries, or instruments exposed in the ram direction it will be much more susceptible to damage than a spacecraft with these sensitive components within the structure or on a less impact-prone surface. Therefore, the location of this curve may

be influenced by satellite design, as shown in the graphic.

The right side of the HNT debris region represents the threshold of trackability. This is orbit and system dependent, heritage systems track down to 10 cm in LEO, 70 cm in GEO, while emerging systems track down to (optimistically) 2 cm in LEO, maybe 30 cm in GEO. If these emerging systems can maintain custody of these smaller objects with a sufficiently small covariance that operators can take action to avoid them, the curve on the right side of the chart will move toward the center significantly. However, it is important to remember that even a ‘trackable’ object can be so poorly tracked that it’s covariance is larger than an operator’s control box. This means first that the probability of collision is never large enough to warrant a response from the operator, depending on the threshold, and even if they did respond they would have to take the satellite off station to prevent the collision, potentially leaving their users with compromised capability during a critical situation. Therefore, an object needs to not just be tracked, but be tracked well enough that space traffic management can successfully mitigate collision risks. A recent example illustrates the risks, as China’s Yunhai-1 (02) satellite broke up last year (March, 2021) due to a collision with an small object associated with a Russian launch. The impacting object was not sufficiently well tracked to avoid, but the fact that the breakup was attributed to a collision with a certain object indicates that it was at least somewhat trackable.

Note that most researchers refer to the hazardous small debris population as ‘Lethal Non-trackable’ (LNT) debris. This term is explicitly avoided in this research, replaced with ‘Hazardous non-trackable’ (HNT) debris, because the entire research concept focuses on debris which is large enough to affect a spacecraft measurably, but not so large that it causes a major anomaly. An important aspect of this is that the same piece of debris which might cause devastating damage to an instrument could impact a solar array or antenna with only minor effects to the mission, as demonstrated in the Sentinel-1A strike. This concept is depicted in the green and red shaded regions on the chart. Smaller hazardous non-trackable debris may cause a minor perturbation, which may go unnoticed, or a recoverable disruption, if it hits a location without sensitive components (which is a large percentage of spacecraft area). This portion of the population corresponds to

the green shaded area. The red shaded area, on the larger size of the HNT region, represents debris that is just barely too small to track and thus large enough to cause substantial damage to a spacecraft, creating a significant population of additional fragments, and likely causing the demise of the satellite. One note is that this type of event is difficult to attribute definitively as a debris strike, many spacecraft have failed spontaneously, sometimes creating fragments, but with limited telemetry available from the event it is difficult to tell if an onboard issue or external forces were to blame.

It is important to note, putting emerging sensor system capabilities together with the NASA battery test in Figure 1.10, that there remains a gap between shielding capabilities and debris trackability. The gap is smaller on emerging systems than heritage systems, but a gap remains, and hazardous non-trackable debris will continue to be a concern for the foreseeable future. Since non-trackable debris is handled with statistical models to assess the probability of mission failure, characterizing these risks accurately will remain important. To characterize this risk effectively, accurate models and methods are necessary, and therefore additional measurements are necessary.

1.2.10 Zooming Out: Framing research within Space Environment Management

Pivoting from the technical aspects of understanding HNT debris risks to the management aspects of the debris environment, Figure 1.16 shows a schematic organizing various debris-related efforts under an umbrella of space environment management. This is similar to content in [18], but the organization is shifted around somewhat. With the broader optic of space environment management, this research falls as a tiny microcosm under ‘characterization’.

Two bold statements are made in this graphic, deliberately. First, small debris removal is listed as “Not technically feasible.” Nobody will be more thrilled than this researcher if the laser researchers (or someone else) figures out how to remove small debris. But there are substantial challenges standing between the state of the art and operational technology capable of having a significant impact on the small debris population, i.e. leading to a meaningful reduction in spacecraft risk. This bold claim is made because some sources [19] insist that a ‘dragnet’ solution

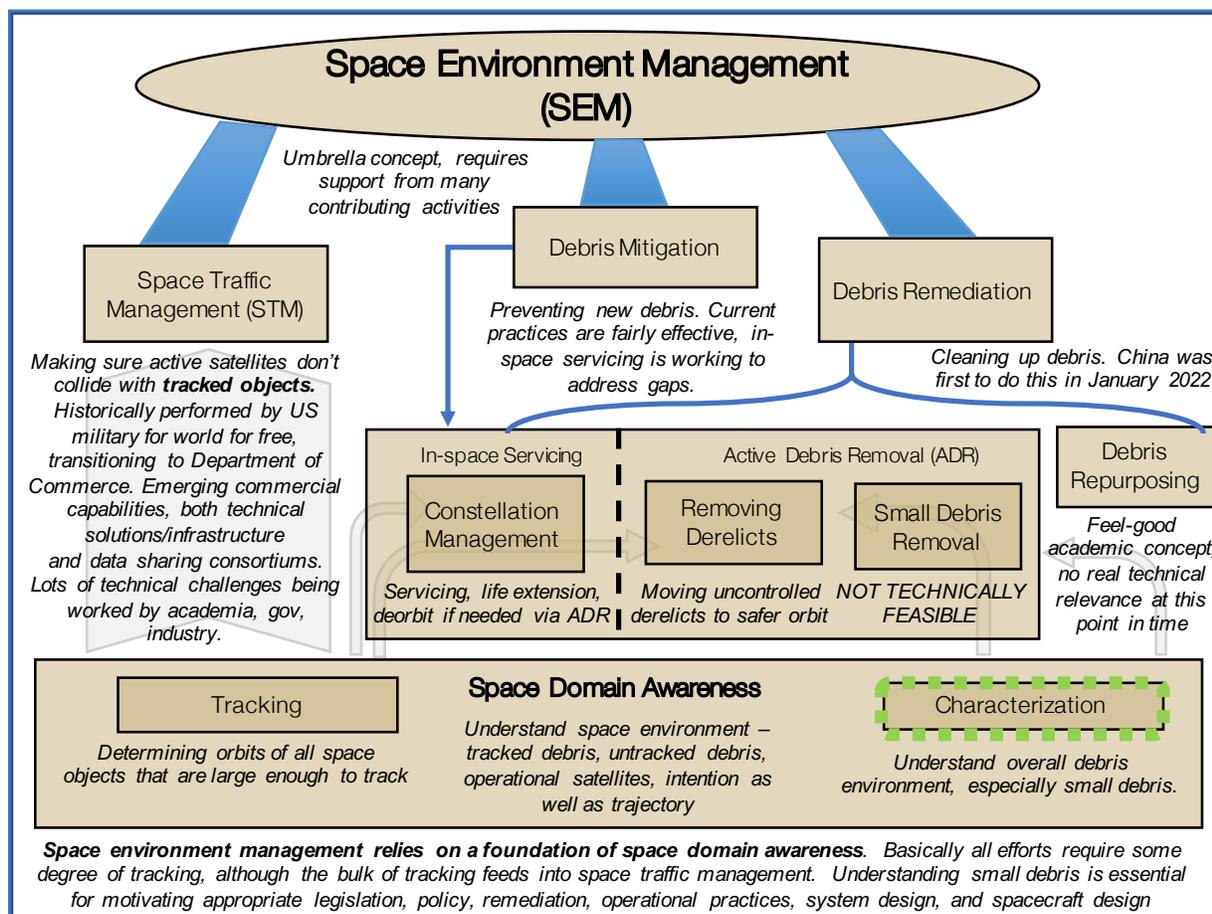


Figure 1.16: Conceptual diagram organizing debris-related activities

to debris, i.e, removing small objects after fragmentation, will undoubtedly be developed if the need arises. In fact, they even state that if debris were really that bad commercial industry would already have solved it, so it must not be that bad. It is dangerous to make predictions about technology development, but it is the opinion of this researcher that these ‘dragnet’ solutions are not technically feasible, and are so far over the horizon that policy should unequivocally assume that small debris will not be removable other than by natural perturbations. Similarly, debris repurposing may be nice in theory but the practicalities of the technical developments render it well over the horizon. Lumping recycling in with ADR tends to make ADR seem futuristic, not urgent and within reach, which is not helpful for motivating necessary remediation investments and activities.

An additional aspect of this chart needs to be highlighted: the distinction between derelict ADR and constellation management. Government officials from both the White House (in 2019) and DoD [20] have made statements indicating that they expect commercial industry to finance ADR efforts—they expect the business case to close. More recent efforts, like the Orbital Prime initiative, have begun providing tiny amounts of funding to small business to try to kickstart the industry, applied broadly to in-space servicing, manufacturing, debris removal etc. The structure outlined in Figure 1.16 very deliberately separates constellation management—including servicing, life extension, and using ADR to deorbit failed satellites **within the constellation**—from the more difficult problem of removing existing derelict satellites and rocket bodies. ADR technology is maturing rapidly and on-orbit capabilities will quickly follow the current suite of on-orbit demonstrations, but it is important for policymakers to understand the differences between constellation management ADR and derelict ADR.

Removing a failed satellite from a constellation is good business, and will likely be affordable. But this is a very different task from removing derelict satellites and rocket bodies that have been in space for decades and were never designed with removal in mind. But these derelicts present a substantial risk to the future debris environment as they have a significant probability of colliding and could potentially double the population of trackable debris in a single event. Unfortunately these are massive objects weighing sometimes thousands of kilograms, a system that can easily remove a failed modern commercial satellite is not the same system that is required to remove (or maneuver) these massive derelicts. Policymakers need to take this into account - commercial industry can be expected to do constellation management ADR, but governments need to finance derelict ADR before collisions render orbits unsafe to use.

Unfortunately, a lack of consensus within the space community impedes efforts to make progress on debris-related policy, regulation, and technology development. Brian Weeden, an industry expert who lives at the intersection between space operations and policy, described this as a situation where the policy makers, prepared to write legislation, ask him, “What do we do?” Weeden tells them “Well, we have twelve ideas, the scientists don’t entirely agree on which ones

are highest priority.” At which point the policymakers lay their pens down: “So you’re telling me it’s not that important.” Weeden’s response, to a room full of industry experts at the Kinetic Space Safety Workshop in Switzerland, was that we need to have a clear consensus on what to tell policymakers. To establish clear consensus, unequivocal, consistent data on debris is required.

“If I told you I could remove 3,000 pieces of small debris, would you pay me to do it? Well I can deorbit one object before it becomes 3,000 objects”

-ADR panelist at Kinetic Space Safety Workshop, Lausanne, Switzerland, 2022

1.2.11 State-of-the-Enterprise: Emerging Trends and Implications for Debris

It’s unlikely that anybody reading this dissertation needs to be told that the space industry is changing. The exponential growth in number of spacecraft shown in the past 5 years of Figure 1.2 is unlikely to slow down, and the global space enterprise needs to figure out how to fly that many spacecraft safely and sustainably. Fortunately, operators are mostly doing their part and sharing spaceflight safety data like precision ephemerides. The Department of Commerce has rolled out their initial Open Architecture Data Repository to begin the process of transitioning space traffic management capabilities from the DoD to Commerce. Commercial operators are delivering exquisite Earth observation and communications capabilities previously unique to governments, enhancing global transparency and communication and providing significant value from space to various industries which support humanitarian endeavors, environmental efforts, safety, etc. Commercial SSA providers are serving a similar role on orbit, using their observation systems to call out misinformation and expose on-orbit shenanigans, providing an unprecedented level of public attribution for questionable space activities. The era of commercial on-orbit servicing has begun, and ADR demonstrations are proceeding apace.

Concerningly, the Cosmos 1408 shutdown demonstrated that geopolitical instability and conflict can have drastic effects for the debris environment, potentially impacting all space operators. An interesting tidbit from the State Department press briefing which decried the event is that

one of the reporters argued with the State Department spokesperson at length about their statements on the event - this started with an incredulous "Are you doing space diplomacy now?" and involved no less than four interruptions in a lengthy back and forth, to address the question: Why is the State Department talking about this? Why not the Pentagon or NASA? Another reporter presses the spokesperson on his claim that "We will not tolerate this kind of behavior", asking him to be explicit: "What does not tolerating this actually mean?" So far as I can tell we're still waiting on an answer to that question, although Vice President Kamala Harris's announcement that the United States unilaterally "commits not to conduct destructive, direct-ascent anti-satellite (ASAT) missile testing" is a step in the right direction.

For the purposes of this research, the point of this is that space is *bigger* than it used to be. Space decisionmakers are no longer limited to niche military and civil agency roles. The global economy is coupled to space capabilities, and maintaining safe spaceflight is inherently a transnational effort involving both public and private entities. The traditional instruments of national power - diplomatic, information, military, and economic - must all be applied, effectively and in balance, to maintain and enhance global space capabilities underpinned by spaceflight safety. This research represents a tiny but potentially useful cog in this massive machine: obtaining more information about hazardous orbital debris.

"Engineering is the art of modelling materials we do not wholly understand, into shapes we cannot precisely analyze, so as to withstand forces we cannot properly assess, in such a way that the public has no reason to suspect the extent of our ignorance."

-Dr. A. R. Dykes, British Institution of Structural Engineers

1.3 Research Concept

From the recommendation of the NESC report, collecting data on momentum changes for comparison to modeled expectations, a straightforward research goal is articulated:

**Develop methods to detect minor debris impacts
using typical spacecraft telemetry.**

There are four core aspects to this concept that differentiate this research from other efforts.

- This research requires no specialized hardware, no specific instrumentation. The goal is to develop techniques using normal telemetry, of types that lots of spacecraft collect, to develop a broadly applicable capability.
- This research looks for perturbations instead of analyzing a known anomaly. There are several examples of cases where anomalous behavior sparked a deep investigation into spacecraft telemetry and found a debris strike. There are fewer examples of actively looking for debris strikes even when they do not cause anomalies. The details of these and the difference between them and this research are discussed in Section 1.4.3.
- This research applies techniques that have not been applied to this topic. Several techniques from various fields are applied to this problem: digital signal processing, change detection algorithms, and modifications to a navigation filter.
- This research is a general solution. These techniques are applied via simulation to telemetry that is fairly typical for many spacecraft. The base techniques can be applied on various spacecraft configurations, and applications to a few unique spacecraft are shown. Simulations and tools are designed such that they can be rapidly modified for alternate spacecraft. Key considerations for broader applications and lessons learned from applying these techniques to real telemetry are discussed throughout.

1.3.1 Core Concept

With this research goal, the first task is determining how a minor debris strike affects the spacecraft. In several cases (details provided in Section 1.4.2), a change in rotation rate and often a change in orbit are described. These features correspond to a change in angular momentum and change in linear momentum. When a piece of debris impacts a spacecraft it imparts momentum to a spacecraft, as shown in Figure 1.17, since momentum is, of course, always conserved. In a hypervelocity impact the momentum can be enhanced by a factor of two or more due to backscattered ejecta from the impact. This concept is discussed in detail in Section 1.4.4.3. The linear and angular momentum transferred to the spacecraft are calculated as a function of the mass of the debris, the relative velocity of the debris, the momentum enhancement factor, and the strike location on the spacecraft.

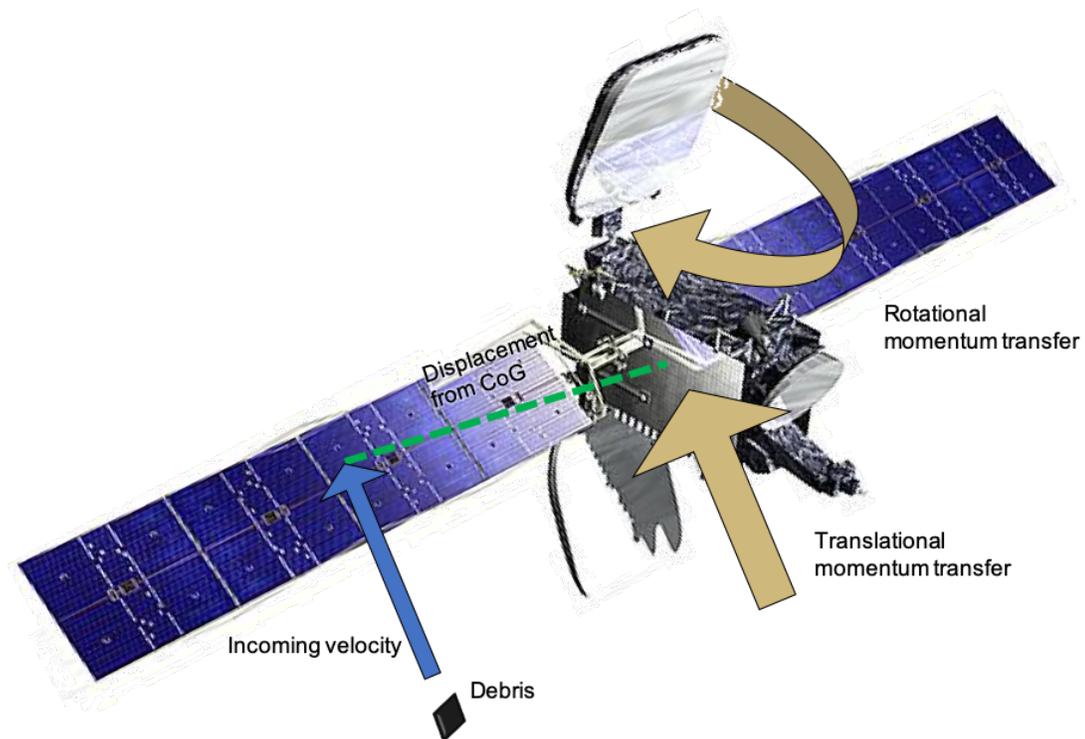


Figure 1.17: Debris impacts spacecraft, causes rotational and translational momentum transfer

Based on the strike's effects, the measurable features in spacecraft telemetry can be assessed. Strike effects vary, and while all impacts produce some level of momentum transfer other features can also occur. Various interplanetary spacecraft have had instrumentation such that the plasma produced in hypervelocity impacts from micrometeoroids can be detected. However, the intention of this research is not to detect impacts using specialized hardware, but using typical spacecraft telemetry. As such, the two types of information that nearly all spacecraft collect are information on pointing, rate and momentum and information on orbit state. The first comes from attitude control system (ACS) telemetry, the second from navigation telemetry, which can be either on-board or ground based solutions.

This research applies various techniques to this telemetry to develop methods to detect subtle perturbations from debris impacts. After a survey of digital signal processing techniques, a matched filter is selected for applying to spacecraft rate telemetry. Various change detection techniques are applied to look for subtle changes in the inertial angular momentum. For navigation telemetry, an extended Kalman filter with dynamic model compensation is augmented with various test statistics to process the telemetry and identify subtle, abrupt orbit changes. These techniques are developed and assessed using simulated data.

But the problem isn't solved with a nice method. Applying these techniques to real telemetry is complicated, murky, and full of surprises. Several case studies on applying these techniques to various spacecraft are discussed with the resultant data and lessons learned. Furthermore, while these techniques can return perturbations the next question becomes how to use that data. A perturbation assessment tool is developed to explore the tradespace implied by all the uncertainties discussed throughout Section 1.2, in the context of using spacecraft perturbation data to learn about the hazardous non-trackable debris environment.

1.3.2 Challenges Inherent in Research Concept

1.3.2.1 Local-Scale Challenges

“Awesome in theory, kind of a mess in practice”

-Unacknowledged source

Applying these methods to simulated telemetry is beautiful. Each of the algorithms shows promising abilities to accentuate subtle effects from strikes. However, applying these to real telemetry is substantially more challenging; each spacecraft has its unique idiosyncrasies which need to be filtered or accommodated in the algorithms in order to obtain clean data for strike detection. Effects that may cause false alarms include on-board events like thruster firing, planned and unplanned appendage motions which change the apparent angular momentum of the spacecraft, component data fluctuations or biases caused by thermal or other effects, electrostatic discharges which can actually impart momentum to the spacecraft, external perturbations like drag or solar radiation pressure, or unanticipated perturbations like, in one case, the spacecraft apparently interacting with Earth’s electromagnetic field. This finding, that this concept is challenging to apply in the real world, is consistent with efforts by other operators to apply techniques, which are discussed in more detail in Section 1.4.2.

Another major factor in using these techniques to gather data about the debris population is micrometeoroids. In this research the impact effects are uniformly referred to as debris strikes, but in reality a micrometeoroid strike would look very similar in spacecraft telemetry. For a single strike differentiating definitively whether it’s a debris strike or a micrometeoroid may not be possible. However, with a sufficiently large body of data from diverse orbits it would be possible to understand more about the various environments spacecraft fly through. Interplanetary missions could be used to understand micrometeoroid perturbations, and then compared with data from LEOs and GEOs to understand the differences between those orbits. Also, micrometeoroids and debris both have their own expected directions, so some strikes may at least be able to determine which population is most probable depending on the impact direction.

Applying these techniques to a single day or week of data is not too complicated, but in order to develop a meaningful body of data they must be applied over long time spans - years of telemetry at data rates of 1 Hz or so. This is a solvable problem, but it does introduce an extra level of complexity to the process, as it involves far more data than can fit in memory. This data must be stitched together, not processed in discrete chunks, since most techniques require a span of data so results would be lost at the junctions between datasets. The algorithms must be robust to idiosyncracies in the data, gaps in collection, or unanticipated spacecraft behaviors. All this infrastructure makes debugging a lot harder, and renders any sort of data cleaning issues (i.e, data gaps, etc) complicated to deal with. Again, this is not a technical roadblock, it just adds to the challenge.

A final challenge is determining the utility of the data. With so many uncertainties in the chain of analysis—debris size, shape, mass, velocity, strike location, estimation of imparted momentum, etc—with all of these uncertainties in every strike, fully determining the parameters of a single strike is not possible based on perturbation data. However, as the NESC report demonstrated, the data can be combined with models to learn about the consistency of the models compared to expectations. By changing the underlying assumptions, the the model output can be changed, which provides insight into what might be causing the discrepancies between expectations and experiences. From a handful of datapoints on one system the NESC report learned valuable lessons about the effect of momentum enhancement factor and the relationship between mass and characteristic length. From larger datasets on more diverse systems more could be learned. The NESC report offered a single case study, a more general capability would expedite lessons as additional data is collected.

This outlines the challenges expected for a single operator applying these techniques on an individual system. There is utility for an operator collecting this type of data, as it can be used for state of health monitoring, anomaly resolution, and to enhance understanding the local debris environmnet, especially for constellations. However, to fully leverage the potential of this data broader data collection and sharing would provide more impactful insights, but there is an

additional layer of challenges pertaining to sharing data.

1.3.2.2 Global-Scale Challenges

Traditionally, space operators (and manufacturers) are extremely averse about sharing data. This absolutely applies to anomalies, as it is assumed that acknowledging issues will make the operator or manufacturer look bad and impede their future business opportunities. People also frequently cite space insurers as a reason to avoid discussing anomaly data, though in my experience insurers tend to have a practical attitude toward anomalies and behave in ways that make them part of the solution, not part of the problem. Note that this hesitation about anomaly sharing is in contrast to the aero industry, which has established and required mechanisms for reporting and resolving anomalies to learn and share industry best practices supporting flight safety.

By studying perturbations instead of anomalies, this may reduce some of the barriers to reporting. Instead of operators having to disclose a major issue which impacts mission performance, they can deliver information indicating that their spacecraft encountered external hazards, but was robust enough to keep flying. This may be more palatable to operators, but is probably not a sufficient solution.

“Who’s gonna be brave enough to step out first?”

-Comment during discussion at Kinetic Space Safety Workshop, 2022, regarding using satellites as sensors to return data relevant to environmental modeling and safety efforts

Commercial operators are reluctant to share data because they are concerned about the impact it may have on their business. In some ways this is typical of a ‘Prisoner’s dilemma.’ In a prisoner’s dilemma, shown in Figure 1.18 two prisoners have the option of either choosing to testify or not. If neither testifies, there is only a weak case, and both serve minimal sentences. This is the best outcome for both. However, if only one testifies, that one gets off free, while the one who chooses not to testify serves a long sentence. Therefore, both typically testify, which results in a sub-optimal solution of medium sentences for both. Similarly, operators are currently

taking the less optimal solution of not sharing spacecraft data, which protects themselves at a cost of more uncertainty in spaceflight safety. This is a difficult situation to mitigate, as whichever operator chooses to share data takes on a risk to their business, so there is limited incentive to do so. A mutually beneficial solution would be to share data, as better information about the debris environment would allow operators to manage their risk appropriately.

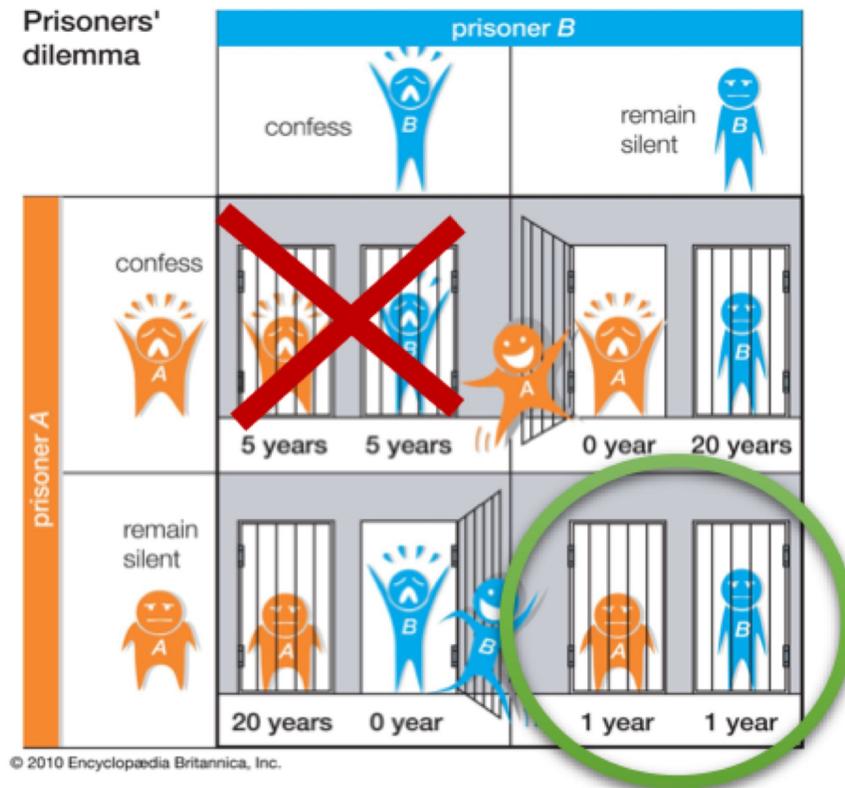


Figure 1.18: Prisoner's dilemma: currently a suboptimal data sharing norm is in effect

1.4 Related Efforts

The first section provided a narrative to motivate this research topic and highlight some important elements of the problem. This section takes a deeper dive into several sub-topics that are relevant to this research. Note that several of these are entire research topics in themselves, the points that are summarized here are the concepts that are relevant to future developments in this research.

1.4.1 Detecting Particle Impacts Using Specialized Hardware

As previously mentioned, the NESC report made a recommendation to monitor existing assets for momentum perturbations, and that is the recommendation toward which this research is directed. Another recommendation was “Consider adding technologies to future spacecraft to detect and characterize MMOD impacts, especially for missions above 600 km. These technologies include external high-fidelity video cameras, micro-electro-mechanical systems (MEMS) accelerometers, acoustic sensors, flash detectors, dedicated *in situ* impact sensors, etc.” However, this is specifically not the targeted topic for this research, the goal of this research is to determine what can be learned from data that is already typically collected in the course of missions, not designing additional hardware to fly. Per SME request, this section provides a list of references for various hardware that has been used to detect strikes on various missions.

The NASA Orbital Debris Program Office flew a dedicated impact detector on the International Space Station. This space debris sensor was a 1 m² technology demonstration consisting of a resistive-acoustic grid to provide information on the sub-mm scale orbital debris environment. It was designed to capture data on the projectile’s direction, speed, size, and density, but was unfortunately lost due to an anomaly after approximately 26 days of on-orbit operations [21]. With lessons learned from this experiment the underlying sensor technology may be flown again. The benefit of a detector like this is the exquisite information obtained about multiple debris parameters, the drawback is that the limited collection area makes it difficult to obtain significant information about debris large enough to harm a spacecraft due to the relatively low flux of larger debris (i.e., > 1 mm).

One concept from the German Aerospace Agency (DLR) is to add hardware to a solar array to allow detection of impacts, as shown in Figure 1.19. The idea is that basically all satellites have solar arrays, frequently with large areas, so adding a minimal amount of hardware to a solar array has little impact on size and weight but can provide data on the debris environment. DLR has conducted ground-based hypervelocity impact tests and a 4-year on-orbit experiment with a small

satellite, with a total detection area of 0.755 m^2 . The pre-flight analysis showed a fairly high probability of detectable impacts, but this was based on future model assumptions that introduced an expected breakup which did not actually occur. With the environment that the spacecraft experienced about 2 detectable impacts were expected, zero impacts are confirmed. Some changes in panel characteristics were observed, but are “quite likely” due to a hardware related failure, not an impact [22]. However, despite these limited results, with lessons learned from the first flight this represents a strong concept for operators/manufacturers who want to incorporate an *in situ* detection capability for future flights.

There are also concepts involving hardware that can be placed on a spacecraft to detect impacts at other locations via various phenomenologies. Invocon offers a few technologies for impact detection for manned spaceflight and unmanned spacecraft. Examples of technologies include acoustic emission transducers, accelerometers, grid-based systems, and grid-less acoustic and RF sensing systems[23]. Applications seem to be mostly on systems with a smooth outer structure, i.e, rockets, planes, and shuttle, but this provides one option to look into if there is interest in placing dedicated hardware for impact detection. A similar concept, developed by Fraunhofer EMI and funded by ESA, suggests an array of small antennas for *in situ* detection of hypervelocity impact effects [24].

A dedicated dust detector, SPADUS, was flown in the early 2000s and detected 368 dust impacts over an interval of 739 days in an 850 km nearly polar orbit. The instrument had a collection area of 576 cm^2 , and collected data on particles in the μm size range. Some velocities were recorded, some of which were clearly interplanetary particles (20-36 km/s) [25].

When a hypervelocity impact occurs it produces a various effects which can be detected by non-purpose-built instruments. Frequently, the plasma produced by the impact can be measured by sensors with a sufficiently high data rate - an experimental measurement of this expanding plasma cloud is shown in Figure 1.19. The Wind spacecraft and the STEREO mission both recorded impacts by observing jumps in the relative potential on their antennas [26]. Wind recorded impact rates at a couple dozen per day, and variations in impact rate showed a consistent annual fluctuation

on all three spacecraft. The MAVEN spacecraft in Mars orbit also recorded dust impacts using the Langmuir Probe and Wave instrument, at rates of nearly 100 per day, with the hit counts at lower altitudes even higher due to atmospheric dust in addition to the interplanetary dust [27]. Fortunately, the majority of interplanetary dust impacts have masses around 10^{-9} g or less [28] so the ability of this population to harm spacecraft is limited.

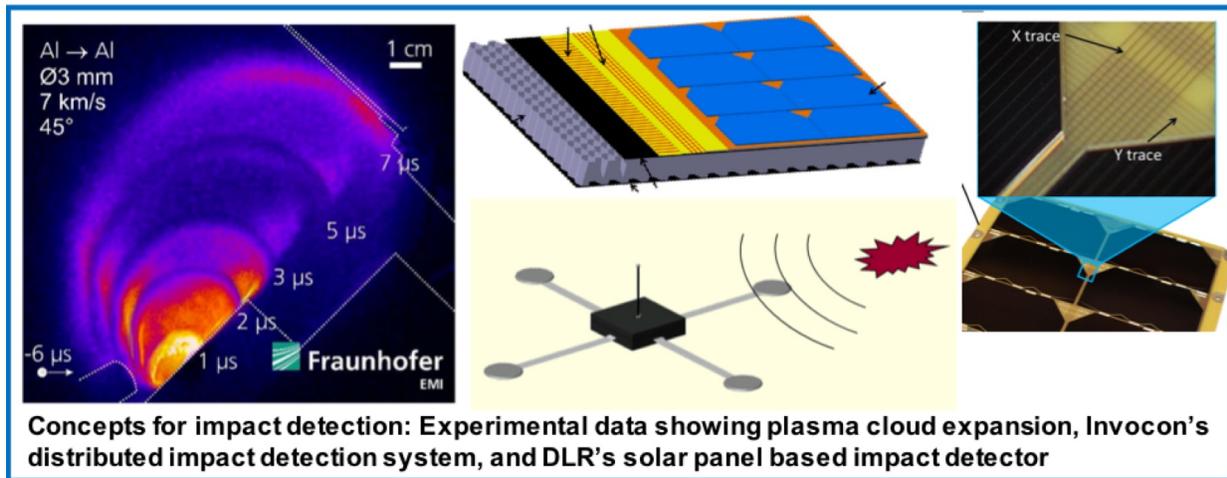


Figure 1.19: Various existing solutions for impact detection. Left: experimental result from Fraunhofer EMI investigating the ability of antennas to detect plasma clouds from hypervelocity impacts [24]. Center bottom: diagram of Invocon's Distributed Impact Detection Sensor [23]. Center top and left: DLR's solar-array based impact detector [29]

All of these types of hardware provide useful data for understanding various aspects of the environment, but the mission first has to decide to fly the hardware, which is a non-trivial effort. These capabilities definitely have the potential to provide complementary data for validation, by establishing when impacts happen and seeing if those match up to perturbations that are near the noise floor and therefore difficult to definitively establish as an impact. However, extra hardware is not the focus of this dissertation.

One interplanetary spacecraft demonstrated an ability to detect dust impacts without using specialized hardware. The Juno mission found routine observations of non-stellar objects appearing in the field-of-view of its startrackers. Attributing this to ejecta from micrometeoroid impacts, the

tally per day was logged and found that variations in the number of events per day correlated somewhat with known dust distribution features. The paper describes collection of data to characterize the distribution and proposes some effects to explain the observed dust distribution [30]. Similar events have been observed by similar sensors on at least two non-Earth-orbiting spacecraft. This concept provides an interesting opportunity to potentially allow confirmation of debris impact events, if lighting angles, camera positions, etc are favorable.

1.4.2 Analyzing Observed Perturbations

Several missions have observed anomalous events and found a debris strike as the root cause. Similar to the Sentinel-1A anomaly, NASA's Magnetospheric Multiscale (MMS) constellation experienced anomalous behavior that was attributed to a particle impact but the mission continued nominally [31]. The strike, which may have been debris or a micrometeoroid, was observed as a loss of one shunt resistor and a spike in the ambient plasma coupled with a dynamic event. On the spin-stabilized MMS spacecraft the dynamic event caused ringing in the accelerometer telemetry, excitation of boom vibration resulting in atypical nutation of the transverse rotation rates, and temporary loss of valid startracker attitude fix. With the lessons learned from the anomalous particle impact, a few other impact events have been observed in the accelerometer telemetry. After the strike, one MMS operator began leveraging digital signal processing techniques to try to find additional impacts, but the work was never published.

On another NASA spacecraft, the Terra Earth observation spacecraft observed anomalous cooling of their batteries, indicating a fault in battery temperature control. One of the battery cells also failed shortly thereafter. Close-out photos showed that the affected heater harness bundle was close to the affected cell. After multiple heater control resets the flight operations team "asked to examine observatory 3-axis rate information for timeframe of anomaly." IRU rate data showed a significant rotation increase and correction, similar to the Sentinel-1A telemetry, and the anomaly was attributed to a micrometeoroid or orbital debris impact. With a 700 km high-inclination orbit, an HNT debris impact is fairly likely.

The Sentinel-1A strike has already been discussed in detail, but a similar lesson is learned, that anomalous behavior prompted a thorough investigation of the telemetry and was ultimately attributed to a debris strike. Similar to the MMS operator, Sentinel-1A operators leveraged their datamining tool for anomaly investigation[32] to try to identify additional strikes that had not caused enough damage to be noticed. The results were inconclusive, and the number of events found varied with the threshold applied, without a concrete way to distinguish between events which might be a small impact or might just be noise. Like MMS, the results of these efforts were never published.

“We have tried this, and I can confirm that it is indeed difficult ... The findings were not very satisfying or conclusive.”

-Dr. Vitali Braun, Space Debris Engineer, ESA/ESOC Space Debris Office

Besides these, several other operators have mentioned similar events on other spacecraft, both government and commercial, which have not always appeared in the public literature. Several strikes which do appear in the public literature are listed in Table 1.1, which highlights the common feature of angular rotation and any anomalous behavior, if disclosed. However, smaller strikes which do not damage components and thereby cause anomalous behavior may go unnoticed.

Mission (Number of Strikes)	Anomalous Behavior	Additional Telemetry Features
MMS (2 known strikes + 3 likely strikes)	Shunt resistor failure, wireboom breakage, loss of startracker attitude fix	Plasma variation, change in rotation rates , accelerometer variation, apparent momentum oscillation (boom vibration)
Sentinel-1A (1 known strike, other indeterminate features)	Solar array power degradation, debris pieces	Angular rate increase and correction , orbit change, damage confirmed with on-board camera
NESC Report Spacecraft (7 known strikes)	Not disclosed?	Angular rate increase , orbit perturbation (change in satellite mean altitude)
Terra EOS AM-1 (1 known strike)	Failure of battery thermal control harness and one battery cell	Angular rate increase and correction

Table 1.1: Publicized debris strikes which perturbed but did not cause mission loss

1.4.3 Work Involving Detection of Momentum Perturbations

There are a handful of efforts from various entities to detect unexpected momentum perturbations. These are described in sufficient detail to establish the differences between these efforts and this research. The NESC report covered both anomalous events and one dataset on momentum perturbations, but it has already been discussed in detail as it is the primary motivation for this research.

“Anne’s work did not slip our attention ...

Her paper is sitting on my desk”

-Operator who contributed data for NESC report

Some key related work has been done by Exo-Analytic Solutions on detecting momentum impulse transfer events (MITEs) on GEO objects using their global telescope network. They have demonstrated a capability to detect orbit perturbations with in-track velocity changes of 0.2 - 10 mm/s [33]. Some of these are explained with rigorous high-fidelity modeling of solar radiation pressure, but others remain unexplained and could be caused by on-board systems or, potentially, debris impacts. Figure 1.20 shows an example of one MITE, as indicated by the measurement residuals from ExoAn-

alytic’s ground-based network, where modeling a small $\Delta\mathbf{V}$ resulted in an improved fit for the residuals. This allows estimation of the size of the event, in this case a $\Delta\mathbf{V}$ of 0.7 mm/s is estimated. While ExoAnalytic obtained these impressive results using external observations of GEO objects from their ground based optical sensors this research focuses on using a system’s knowledge of its own orbit, and develops methods to enhance detectability beyond what can be attained by

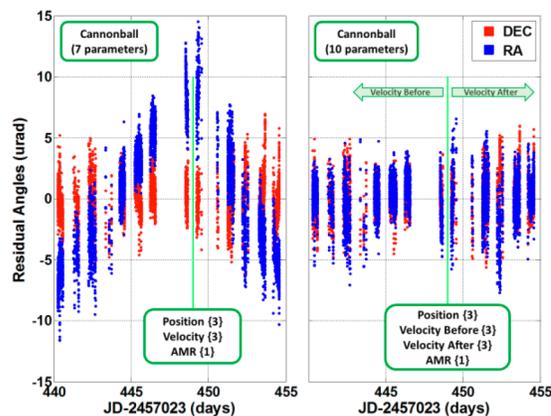


Figure 1.20: Graphic published by ExoAnalytic Solutions[33] showing filter residuals with and without an estimated MITE

observing measurement residuals.

A recent study by ESA conducted at the Fraunhofer Ernst-Mach-Institut assesses the momentum transferred during small hypervelocity impacts on spacecraft materials via modeling and test and then develops software to simulate the dynamic response of the a spacecraft to these disturbances. Measurements are planned for an upcoming mission which will derive the momentum transmitted into the spacecraft by microparticle impacts based on these models and analysis [34]. On the LISA Pathfinder, located around L1, mission data was used to identify 54 micrometeoroid impact events imparting momentum between 0.2 and 230 μNs [35]. This is a similar concept to this research but uses LISA's exquisite ACS precision to identify strikes instead of post-processing typical telemetry with dedicated algorithms to accentuate subtle features.

Another ESA spacecraft, GAIA, located at L2, has released occasional information indicating that it observes higher than expected rates of micrometeoroid impact events by observing small changes in the spacecraft rate and sunshield thermal profile, and giving an example of one rate perturbation event likely from the Persids meteor shower [36, 37, 38]. Specific comparisons of impact rates, detection thresholds, and comparison to models have not yet been found, although it appears that work may still be ongoing [39].

Similarly, a paper by the Institute for Defense Analyses (IDA) discussed methods to monitor satellites for small changes in orbit[40] with the intent of cataloging minor debris impacts and feeding back the results to debris models. Methods proposed for detecting orbit changes include changes to satellite mean altitude (dSMA), crosslink disruption, and coarse GPS single-point measurements assessing a discrepancy from the satellite's nominal orbit. The concept in this paper is the most similar to the topic of this research, but the analytical developments are dramatically different. This research develops specific algorithms to process simulated telemetry, whereas the IDA paper speculated generically on feasibility based on coarse estimates of capabilities. Also, this research applies the concept to a swath of on-orbit telemetry, in particular considering detection of rotational momentum transfer in ACS telemetry which is not addressed by the IDA paper.

Included in this category of similar research are several papers published in the course of this

research effort. An October 2019 conference paper applied change detection algorithms to simulated attitude control system telemetry, showing that the algorithms have a strong ability to accentuate minor fluctuations in the inertial angular momentum, which would indicate the presence of an external torque[41]. This development is described in Section 2.3.2 of this dissertation. A follow-on conference paper in December 2019 applied these techniques to NASA spacecraft and found several interesting features [42], described in Section 4.1. A February 2020 conference paper presented a trade study to characterize the sensitivity of these methods to various system parameters [43], results are shown in Section 2.5.

Two journal papers have discussed the primary developments of this work. A 2020 *Acta Astronautica* paper discussed the development of methods for detecting rotational momentum [44], and a paper describing development of methods to detect translational momentum changes is currently in peer review with the *Journal of Spacecraft and Rockets* (as of October 2022).

1.4.4 Peripherally Related Work

This section includes brief summaries of topics that are entire research areas in their own right, but insight from these topics is relevant to the development of these methods.

1.4.4.1 Ballistic Limit Equations

The ballistic limit represents the ability of a shield to withstand an impacting projectile. The term has been used in literature dating back to at least the 1930s to describe work characterizing an armor's effectiveness in stopping bullets. In space, ballistic limit equations (BLEs) saw use in the 1960s to estimate the micrometeoroid hazard to the Apollo command service module, with continuing development since that time. As Section 1.2.7 described, ballistic limit equations still underpin debris risk analyses, albeit now wrapped in software, not typewritten.

Ballistic limit equations (BLEs) are used to express whether or not a given projectile will penetrate a given target. For a given projectile/target combination, i.e, aluminum sphere vs. CFRP panel, BLEs are typically plotted with velocity on the x-axis and the projectile diameter on the

y-axis. The ballistic limit is plotted such that a projectile with characteristics above the curve (i.e, velocity and diameter) is expected to overcome the target, while projectiles below the curve are not expected to overcome the target. The failure criteria can be defined in different ways, either complete penetration (hole through the target) or rear-side spall (material ejected from back of target due to shock wave). Reference Figure 1.23, which shows an attached spall caused by stress waves propagating through the target from the impact. When this spall detaches, it can be considered a target failure, even when there is no penetration of the target.

Various BLEs exist with various input parameters, which can include a variety of specifications like material densities, strength, shield configuration, and impact angle. These equations are typically empirically derived based on test data, and are thus specific to certain shield/projectile combinations. They can be applied under different conditions, but ‘results may vary.’ Figure 1.21 shows an assortment of ballistic limit equations for various shapes, indicating the variability in the results due to projectile shape. Also note that this graph plots the ballistic limit for a variety of shapes against one target with normal impact, and yet it is only for one set of BLEs. Multiple sets of BLEs exist, and which to use is a choice by the analyst.

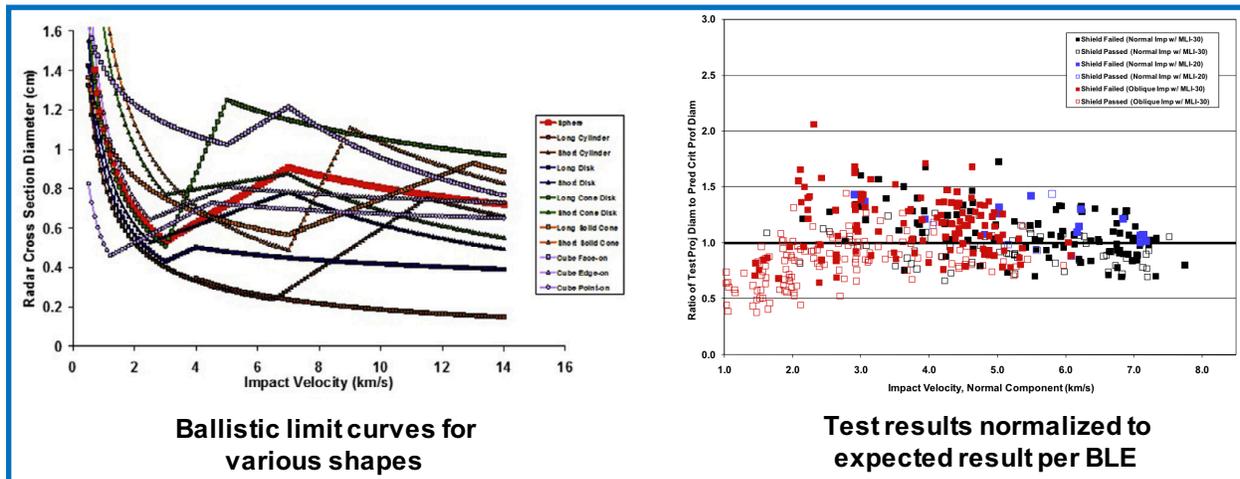


Figure 1.21: Examples of BLEs and test results. Figures from Schonberg [45]

Also, while the BLEs generate a deterministic solution, where one set of inputs provides one result, the actual results are inherently probabilistic. Repeated tests to characterize this variance

are rare, as tests are very expensive. Figure 1.21 shows test data normalized to a ballistic limit curve. If the equation perfectly predicted target response (fail/no fail), then all the points below the line would be unfilled, all points above the lines would be filled. Unfortunately, results vary widely, showing projectiles which differ from anticipated behavior by as much as a 30%.

A further complication is that the maximum velocity for ground tests is around 7 km/s for most test equipment, some can push to around 10 km/s with a very small projectile. Orbital collisions can be in excess of 14 km/s, so test results aren't able to fully replicate orbital conditions. In the testable regime there are multiple non-linearities in ballistic limit as the physics of the impact changes with increasing energy, and there is an inherent assumption that target response in non-testable regimes is similar to target response in testable regimes.

1.4.4.2 Micrometeoroids

Micrometeoroids offer an additional complication. A micrometeoroid impact can deliver momentum to a spacecraft similar to a debris impact. Therefore, to draw useful conclusions about the debris environment from perturbation data the micrometeoroid environment must also be understood. A few relevant aspects of the micrometeoroid environment are discussed here.

Like debris, the micrometeoroid flux model is based on various measurements and has a degree of uncertainty. MEM3 [46] is NASA's latest micrometeoroid environment model, and is used for modeling micrometeoroid impact risks to spacecraft. MEM3 incorporates flux down to a size of 10^{-6} g because it is concerned with damage, but as mentioned above there are surprisingly high rates of dust impacts detected on several missions, which are thought to be much smaller particles. The micrometeoroid environment has preferred directions, as does the debris environment, and this can be used to identify the likelihood of a particular strike being debris or micrometeoroid if the strike direction can be established. Figure 1.22 shows some measurements of variations in sporadic micrometeoroid flux. These annual variations can also give some clues regarding whether detected impact populations are debris or micrometeoroids.

The momentum imparted by the strike is of course a function of the micrometeoroid's mass

and velocity. The density of micrometeoroids is bimodal, with some micrometeoroids having approximately the density of water and some being 3-5X heavier. With the density and the size the mass can be computed, and the shape of micrometeoroids is likely closer to spherical than the shapes of debris, which are often flat or elongated due to manufacturing features. The velocities of micrometeoroids, however, are thought to vary between 12 km/s and 72 km/s. Therefore, micrometeoroid velocities are much higher than the regime where impact mechanics can be tested with large particles, some implications of this are discussed further in the next section.

Another peculiarity is the variation in micrometeoroid flux. While easily-visible meteors come from comet tails and have known showers, like the Perseids, there is also a ‘sporadic annual variation in micrometeoroid flux’ which is separate from known comet tails. This has been measured via radar, and it is found that the variation is somewhat repetitive year over year[47]. A similar annual variation is shown in dust impacts on the two Stereo spacecraft and the Swift spacecraft [26]. While a major meteor shower may increase the flux temporarily by 10X or so, the timespan of these showers is not long so the total risk is driven by the sporadic background. Attitude disturbance rates and likelihood of damage may be lower, only increasing by about 2X for the duration of a meteor shower, based on examination of the momentum-limited and mass-limited flux [48].

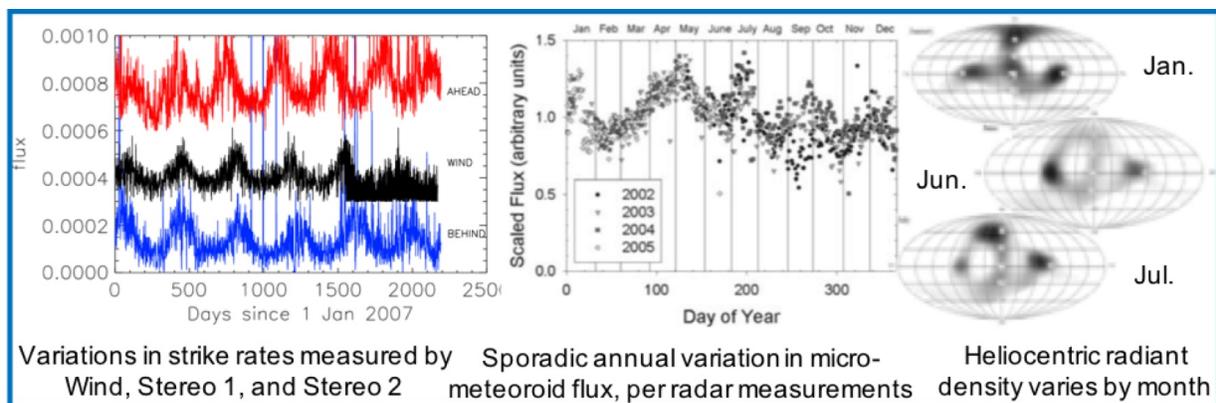


Figure 1.22: Examples of micrometeoroid flux and interplanetary dust impact data. left: [26], center and right: [47]

The takeaway is that yes, separating debris impacts from micrometeoroid impacts is a chal-

lenge, but discriminating features between the two populations exist such that with enough data from diverse orbits, especially interplanetary missions, conclusions could be drawn about each environment. Often, the two populations are grouped as micrometeoroids and orbit debris (MMOD) for the purposes of impact risk assessments.

1.4.4.3 Momentum Enhancement

Momentum enhancement in hypervelocity impacts has already been alluded to, but it is here discussed in more detail to understand the limitations of what is known about this important mechanic. A hypervelocity impact is defined as an impact in which the speed of the projectile is faster than the speed of sound in the material. The impact thus has a large amount of kinetic energy which is delivered to the target in a very short amount of time. This leads to two interesting effects, where a shockwave propagates through the target, reflecting at boundaries and coming back toward the point of impact, as shown in Figure 1.23, and a potentially large amount of material is expelled from the target in the direction of the impact. Figure 1.23 also shows the material expelled from a large monolithic aluminum target. This backscattered ejecta, thanks to conservation of momentum, results in more momentum being delivered to the target than was originally present in the impactor.

It is important to understand the limitations of momentum enhancement factor testing and characterization to appreciate the uncertainties involved. As previously discussed, hypervelocity impact tests to derive ballistic limit equations are very expensive and can't quite test to orbital debris velocities. This test limitation is even more significant for micrometeoroids, since their velocities are 12 km/s-72 km/s [48]. Also, hypervelocity impact testing to date has mostly focused on ballistic limit, damage, etc, there is a much smaller set of hypervelocity impact tests which measure momentum enhancement. Out of these, the majority measure momentum enhancement of hypervelocity tests against rock-like targets, for modeling impacts on meteors. Published results from hypervelocity tests which measure the momentum enhancement of strikes against spacecraft materials are still lacking in the literature, although pending results have been mentioned [34].

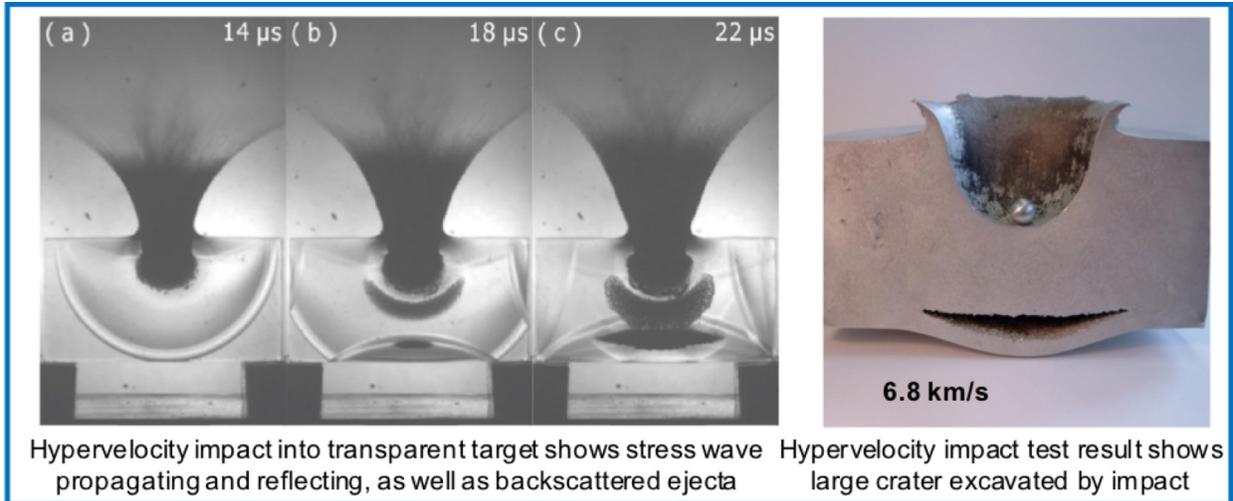


Figure 1.23: Examples of hypervelocity impacts, showing backscattered ejecta plume forming which enhances momentum, proportionate to amount of excavated material

Some relief from this test capability deficit comes from computational techniques, one popular approach for hypervelocity impacts is spherical particle hydrodynamics. Simulations can be created to model the effects of hypervelocity impacts on various targets, obtaining a solution from which momentum enhancement can be derived [6]. These results anticipate an MEF between 1 and 3 for a debris impact on a dual-wall spacecraft structure, and this is what is used in the NESC report for counting perturbations. However, a solar array might see an MEF of less than one if the strike bursts through the array. If the array is extremely thin, like the original Hubble arrays, the MEF will be very low, but for a more typical CFRP array it might be higher, around 0.45 for one simulation [14].

It is important to validate SPH and related computational techniques against measurement data to ensure the validity of the simulation. This presents a difficulty at non-testable velocities, and for impacts of widely varying debris shapes and materials versus diverse spacecraft surfaces. Also, the amount of computation time involved results in limitations on the configurations that can be simulated. For example a simplified simulation obtaining the aforementioned debris strike against a CFRP panel took 140 hours to complete using ANSYS Autodyn on a student's laptop,

so meaningful trade studies on more sophisticated models would require nontrivial computational resources [14].

Another relevant test capability is microparticle accelerators, which can accelerate very small particles to micrometeoroid velocities. However, the masses of particles involved are on the order of 10^{-12} g or smaller, so the results need to be extrapolated several orders of magnitude to reach a particle which could impart detectable momentum to a target. Some attempts have been made to characterize the momentum enhancement using data from these tests, but measurements are difficult given the small masses involved. Measurements have been obtained indicating momentum enhancement factors of around 10 at a velocity around 10 km/s, but the uncertainty bounds on these measurements are extremely large. Various attempts have been made to fit equations to these measurements resulting in a variety of momentum enhancement predictions at high velocities. McDonnell [49] describes two such fits, one predicts an MEF of 5-20 at micrometeoroid velocities, while one predicts an MEF of 3-10 at these velocities. This second one is typically used, since the 68.4 km/s Giotto encounter is ‘known’ to have a momentum enhancement factor of around 10.

However, it’s worth digging back through the literature to determine how this MEF is established for the Giotto encounter. Reference [50] describes an MEF of 11 used to interpret the science data, but digging back to the cited reference [51] shows that this is hardly unequivocal. The data from two types of tests is conflicting, and McDonnell picks the lower of the estimates. The lower estimate is derived from ballistic pendulum data, with data ranging from 1-8 km/s. An additional dataset is obtained from the Canterbury microparticle accelerator, but the measurement data is hard to obtain at higher velocities and subject to bias, with large error bars, so McDonnell discards this data and uses the lower estimate, justifying this with a derivation of an upper bound using conservation of momentum and energy. To derive this upper bound McDonnell assumes a relationship between projectile diameter and crater diameter derived in 1970. More recent data, shown in Figure 1.24, disproves this relationship and reveals that the actual correlation is higher, which results higher momentum. Figure 1.24 shows a quote from the McDonnell abstract, indicating that he selected an MEF of 12 based on ballistic pendulum (BP) data but the electrostatic (ES)

accelerator data “indicates a figure considerably higher.” The graph on the left, from McDonnell, shows his selected best fit from the BP data leading to a modest MEF, but the curve is well below the ES test data. The best fit from the ES data curves up to incredibly high MEFs, but no data is available in the range of fast micrometeoroids, the shape of the curve in the 20-70 km/s regime is unknown. The figure on the right from [52] shows some more recent test data with an annotation added showing McDonnell’s assumption for the relationship between crater diameter and projectile diameter. This shows that McDonnell underestimated the mass ejected in a hypervelocity impact, which would move his ‘theoretical upper bounds’ upward, and of course still assumes that the data at higher velocities can be extrapolated from data at lower velocities, which is always questionable. The relevance to this research is that we really have no idea what the MEF is at micrometeoroid velocities, but it is fairly likely to be higher than the prevalent models.

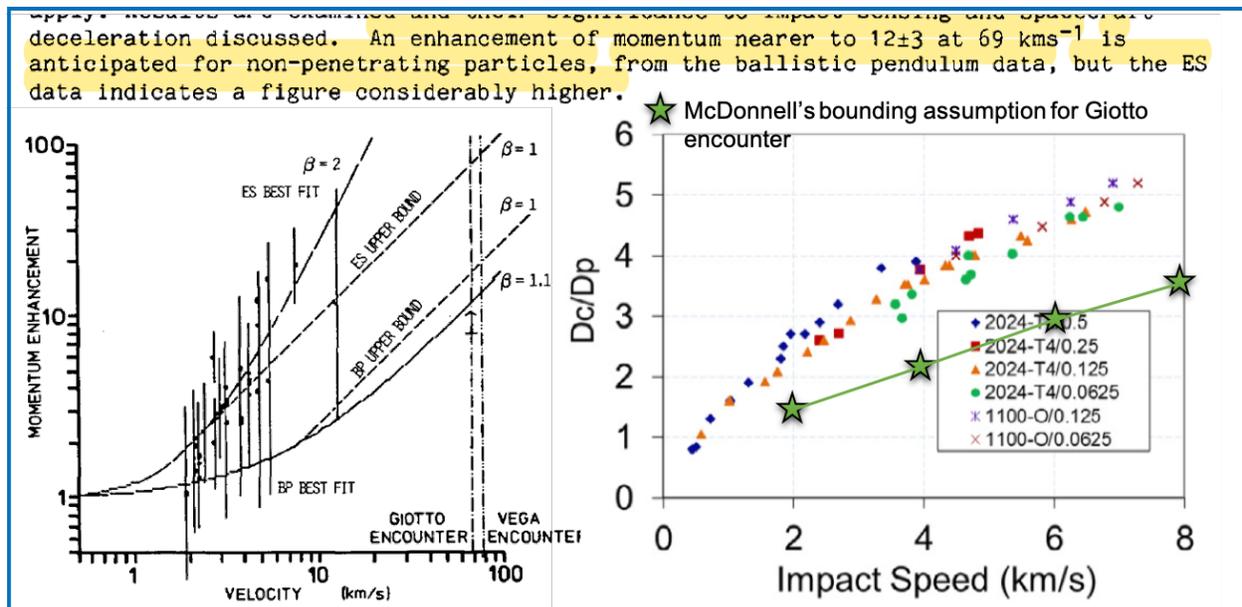


Figure 1.24: Investigating derivation of ‘known’ MEF from Giotto encounter and comparing underlying assumptions with newer test data

1.5 Research Overview

With this extensive background on relevant aspects, it is finally time to discuss the content of this research effort. The goal of this research is to develop methods to detect minor debris impacts using typical spacecraft telemetry. Efforts toward that goal are split into three sections: The first develops methods to identify the rotational momentum transfer using telemetry from a spacecraft's attitude control system (ACS), the second detects translational momentum transfer using navigation telemetry, and the third explores two developments related to applying these techniques. The first part (Chapter Four) applies the techniques to on-orbit telemetry, while the second part (Chapter Five) investigates the effect of various parameters on the expected rate of perturbations.

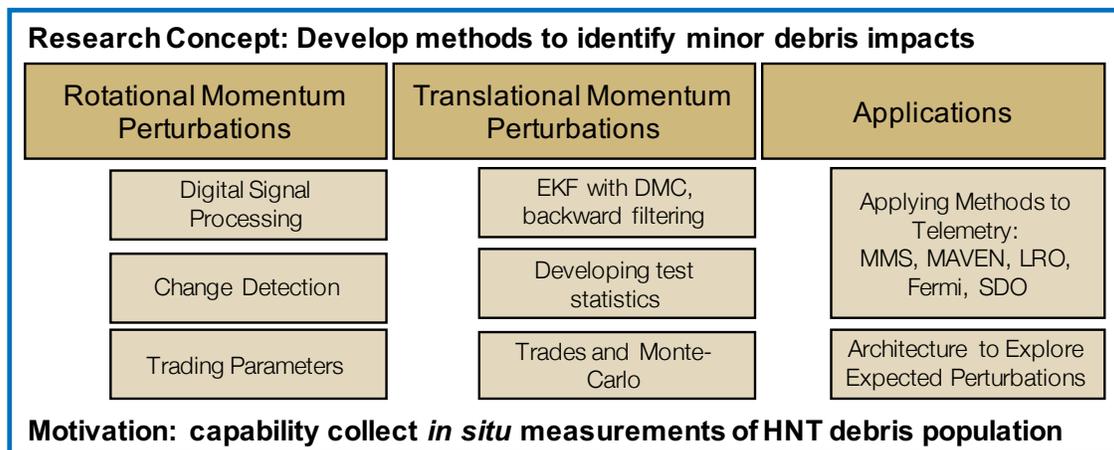


Figure 1.25: Outline of Research

1.5.1 Rotational Change Detection

A spacecraft dynamics simulation is used to model the effect of debris strikes on ACS telemetry. A numerical simulation is developed which models a typical spacecraft configuration, a 3-axis controlled rigid-body spacecraft with four reaction wheels which maintains pointing along a reference trajectory. Debris strikes are applied to the simulation while noise obfuscates the spacecraft's response to the strikes. Matched filters are applied to identify the 'signals' produced in spacecraft

state telemetry during a debris strike. This includes developing thresholds for detection, developing methods to estimate strike parameters, and using a Monte-Carlo analysis to assess algorithm performance. After this, change detection techniques are applied to the momentum telemetry, which exhibits lasting changes when a debris strike produces a change in the spacecraft's inertial angular momentum, and these two methods for strike detection are compared. The specific algorithms selected for these developments are chosen based on a survey of canonical digital signal processing (DSP) and change detection (CD) techniques. These canonical techniques are chosen because this research effort is intended to develop a robust, well-understood, and well-characterized method for an initial study in applying this concept to improve understanding of the overall problem of debris impact detection.

Once these algorithms are developed, an analysis is performed to compare their effectiveness on various simulated systems. Spacecraft and telemetry parameters are traded to determine the sensitivity of the techniques to various parameters, like spacecraft size and data rate.

This chapter's core development is producing techniques to process telemetry to identify subtle effects on a spacecraft's rotational momentum state. The algorithms are developed and applied to simulated telemetry to compare their performance against known perturbations allowing a definitive assessment of each algorithm's accuracy.

1.5.2 Orbit Perturbation Detection

This chapter develops some techniques for processing spacecraft navigation telemetry to identify subtle, abrupt deviations in a spacecraft's orbit state. An extended Kalman filter (EKF) with dynamic model compensation (DMC) is employed with the unmodeled accelerations estimated using a first-order Gauss-Markov (FOGM) process. This filter is then augmented with various test statistics to identify an abrupt deviation in state. The filter is applied forward and backward across the data and then smoothed, with the backward filter employing a FOGM that is mirrored in time. Various test statistics are developed and compared to identify which ones can produce the highest signal-to-noise ratios (SNR) in the presence of a small change in velocity. Because the

backward filter employs a mirrored FOGM model for the unmodeled accelerations and produces a state estimate using only ‘future’ data, and the forward filter only uses ‘past’ data, this produces a discrepancy in between the two filters at the time of the strike. This approach enhances debris strike detectability by accentuating the difference between forward and backward states at the point of the debris strike, and differs from many smoothers which produce a recursive estimate using all the data and thus tend to smooth out a debris strike. This fusing and smoothing approach also causes a much larger feature to appear in the unmodeled accelerations, which produces a strong feature when a debris strike occurs. Various test statistics are developed from filter output to show a distinct feature when a debris strike occurs.

Simulated onboard navigation telemetry is employed to develop and test the filter; very basic orbit propagation and measurement noise models are used since the focus is on methods development and characterization. A trade study is conducted to determine the effect of filter parameters on performance, and characterize the various methods relative to each other. A Monte-Carlo illustrates the filter’s ability to detect and estimate unknown strikes. An additional study adds perturbations to the orbit propagation but not the filter to investigate the performance in the presence of additional unmodeled accelerations, and the filter is run against on-orbit telemetry from various spacecraft to investigate real-world considerations for applications.

1.5.3 Applying Techniques to Telemetry

After developing the techniques on simulated data, the most promising techniques are applied to flight data from several spacecraft to explore their effectiveness on real-world data and the associated challenges. The first task in applying these techniques to telemetry is to develop an architecture that can ingest telemetry and then apply the techniques. This architecture is built to be agile, with the ability to quickly change between various spacecraft and telemetry types and experiment with various algorithms. The architecture is versatile to ingest telemetry from various systems, and can apply processing algorithms to telemetry to identify subtle strikes or use other telemetry (i.e, thruster firings) to ignore ‘detections’ when planned activities result in spacecraft

momentum perturbations. The results are then aggregated to identify potential strike times and perturbation rates. This processing architecture, of course, must accommodate large quantities of data in order to process the swaths of telemetry required to identify unknown debris strikes.

The architecture is used to process the data from several spacecraft in various orbits. The specific techniques required to process data from each spacecraft is discussed, as well as the results from applying these techniques.

1.5.4 Using the Data: Perturbation Rate Assessment Tool

A perturbation rate assessment tool is developed which uses environment models to predict rates of perturbations on spacecraft. This tool is intended to gain insights about how the data resulting from these techniques could be used to improve knowledge of the debris environment by studying the effects of various assumptions and treatments of the problem on the expected perturbation rates. The NESC report fit the curves from debris environment model predictions to a dataset that included seven events on one system. From this limited dataset, it can be seen that heritage risk assessment methods are likely overly conservative but adjusting the shape assumption and refining other uncertain parameters provides a path to improved risk assessments for more effective space environment management.

This research topic generalizes the NESC report techniques to investigate the tradespace of perturbation rates based on models. A tool is developed where a spacecraft can be specified by its dimensions and the appropriate debris environment can be read in for a specified orbit. Various parameter assumptions, like MEF or shape, can be traded, along with various spacecraft sizes and orbits, to determine the effects on predicted perturbation rates. Exploring these relationships gives more insight into the best areas to apply these techniques to find a useful population of debris strikes, and the results could be compared to on-orbit data to assess the realism of the various assumptions underlying the model. Various distribution parameters can be implemented in a Monte-Carlo analysis to gauge the effect they have on perturbation rates, or potentially investigate the match between model predictions and on-orbit experiences.

1.5.5 Utility to Field

Getting useful perturbation data from satellites is certainly challenging, but the NESC report and Exoanalytic's MITEs study prove that it is possible to obtain useful data and gain more insight. The techniques developed here offer ways to improve the detection thresholds for momentum perturbation events, which can be leveraged to produce a larger dataset of perturbations to assess. With more perturbation data, the models and risk assessment methods can be improved such that expectations are consistent with operator experience. This will aid in providing the needed consensus for space environment management efforts and safe satellite design and operation, as well as motivating necessary policy and regulation.

While the full utility of these techniques would come from operators making data available to modelers and researchers, there is also utility in applying these techniques as a single operator. In the Terra anomaly the ground crew tried various responses before looking at the attitude telemetry and realizing that the anomaly was caused by a debris strike. Having an immediate notification to operators that a telemetry feature indicates a debris strike could be very useful in correctly responding to an anomaly and choosing the best course of action to assure spacecraft safety. Also, an improved knowledge of the local debris environment can motivate safe design practices from operators. The Iridium constellation, launched in 1997, unfortunately experienced two of the three worst debris causing events near their orbit. However, when they performed a technology refresh with launches beginning in 2017 they remained in the same orbit, but their re-designed satellites have a much lower cross sectional area in the ram direction (where most debris strikes occur), with the primary mission antenna oriented parallel to the velocity to minimize vulnerable area. The block 1 configuration is shown compared to the Iridium NEXT configuration in Figure 1.26.

Based on this, it would seem that Iridium's on-orbit experiences flying through one of the more severe LEO regimes were tolerable, but risk is sufficiently high to motivate a safer flight profile. This provides an example of an action that manufacturers/operators can take to operate safely and responsibly in more congested orbits, based on an improved understanding of the risks

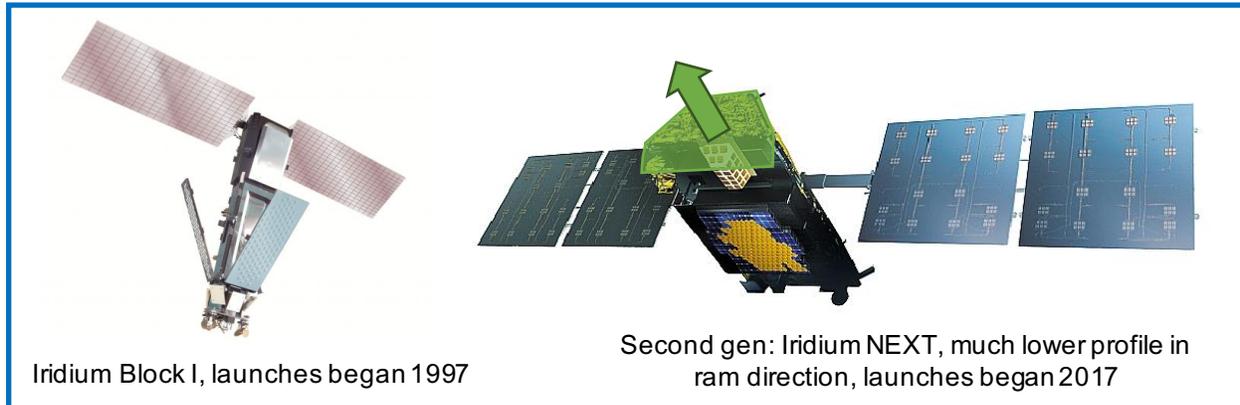


Figure 1.26: Comparing first generation and updated Iridium satellites

which is gained by flying a large number of satellites in those orbits for a long period of time. If more protection was desired there are various options, like the fabric in Figure 1.10, to add even more shielding against HNT debris fairly cheaply.

In approaching the problem from the reverse direction, identifying subtle strikes instead of responding to anomalous behaviors as they occur, these methods allow spacecraft operators and the debris community to proactively improve models and methods for assessing debris risks and attributing anomalies. Monitoring perturbation events offers a unique optic for characterizing the small debris environment since the measurement depends explicitly on the mass of the particle, which is more closely correlated to the damage potential than the size measurements obtained via traditional radar and optical systems.

Various efforts have suggested using satellite-as-a-sensor [53] techniques to improve understanding of various orbital phenomena to improve spaceflight safety. Commercial operators seem interested in this approach, as a Project Kuiper representative recently stated that they would soon have “3,236 sensors on orbit” to aid in improving models. A NASA technical fellow, Dr. Joe Minnow, mentioned that this interest extends to the federal level, where space weather action plans have elements that include telemetry monitoring and using telemetry data towards mitigating space environment risks. This research was presented at the Spacecraft Anomalies and Failures workshop, and during his outbrief Dr. Minnow gave several complimentary thoughts on this research and its

utility, concluding as follows:

*“But then somebody’s got to figure out what to do with [the data], and I think
this is a good start”*

-Dr. Joe Minnow, NASA Technical Fellow, describing these techniques during his outbrief remarks
at the Spacecraft Anomalies and Failures Workshop, 2021

Chapter 2

Developing Methods to Detect Rotational Momentum Transfer

Following the recommendations in the NESC report, the first major task in this research is to develop methods to detect rotational momentum transfer due to minor debris impacts. The Sentinel-1A debris strike provides an exemplar of the type of event these algorithms should detect. However, assessing the performance of the algorithms—the detection and estimation accuracy—requires *a priori* knowledge of the debris strike characteristics, and therefore the ability to apply strikes to telemetry. As such, the first subtask is to build a spacecraft dynamics simulation to model the effects of a debris strike on the rotational dynamics of the spacecraft (Section 2.1). Next, digital signal processing algorithms are explored and developed (Section 2.2) for application to this problem. Change detection techniques are also explored and applied, and then the performance of both types of algorithms in detecting and assessing debris strikes on simulated telemetry is characterized (Section 2.4.1 and 2.4.2). Finally, a trade study is conducted that trades various specifications to understand the performance of the algorithms when applied to various systems (Section 2.5.2).

2.1 Spacecraft Dynamics Simulation

2.1.1 Spacecraft Simulation

A numerical spacecraft dynamics simulation is used to investigate a spacecraft's response to debris strikes. The simulated spacecraft is configured as a rigid-body 3-axis controlled inertially-pointing GEO satellite with four reaction wheels (RWs). The state vector consists of the spacecraft's attitude and rate along with the rotation speeds of the four reaction wheels. A Runge-Kutta fourth-

order integrator is used to integrate the equations of motion given in Equations (2.1), (2.2), and (2.3) [54, ch. 8]. Note that Equations (2.2) and (2.3) must be solved simultaneously.

$$\dot{\mathbf{q}} = \frac{1}{2} \begin{bmatrix} -q_1 & -q_2 & -q_3 \\ q_0 & -q_3 & q_2 \\ q_3 & q_0 & -q_1 \\ -q_2 & q_1 & q_0 \end{bmatrix} \boldsymbol{\omega} \quad (2.1)$$

$$[I_{sc}]\dot{\boldsymbol{\omega}} + I_{ws}[G_s]\dot{\boldsymbol{\Omega}} = -[\tilde{\boldsymbol{\omega}}][I_{sc}]\boldsymbol{\omega} + \mathbf{L}_{ext} - \sum_{i=1}^N [\hat{\mathbf{g}}_{ti}I_{ws}\Omega_i\hat{\mathbf{g}}_{gi}^T\boldsymbol{\omega} - \hat{\mathbf{g}}_{gi}I_{ws}\Omega_i\hat{\mathbf{g}}_{ti}^T\boldsymbol{\omega}] \quad (2.2)$$

$$I_{ws}[G_s]^T\dot{\boldsymbol{\omega}} + I_{ws}[I_{N \times N}]\dot{\boldsymbol{\Omega}} = \mathbf{u}_s \quad (2.3)$$

The spacecraft's rotation rate, $\boldsymbol{\omega}$, is the rate of the body frame relative to an inertial frame, and is expressed in body frame coordinates. The attitude quaternion of the body relative to inertial is \mathbf{q} , and $\boldsymbol{\Omega}$ is an $N \times 1$ matrix of wheel speeds for N reaction wheels. The external torque applied by the debris strike is \mathbf{L}_{ext} while $\hat{\mathbf{g}}_{si}$ is the unit vector of the spin axis of the i th reaction wheel. The transverse directions are $\hat{\mathbf{g}}_{ti}$ and $\hat{\mathbf{g}}_{gi}$, and the $[G_s]$ matrix is defined as $[G_s] = [\hat{\mathbf{g}}_{s1} \dots \hat{\mathbf{g}}_{sN}]$. The spacecraft inertia (with reaction wheels included) is given as $[I_{sc}]$, and the spin-axis inertia of each reaction wheel is I_{ws} . An identity matrix is denoted as $[I_{N \times N}]$, and the applied control torques for each reaction wheel are specified in \mathbf{u}_s . These control torques are calculated as shown in Equation (2.4) to apply a desired torque, \mathbf{L}_{des} , which is calculated per Equation (2.5).

$$\mathbf{u}_s = [G_s]^T([G_s][G_s]^T)^{-1}\mathbf{L}_{des} \quad (2.4)$$

Note that the Modified Rodrigues Parameters (MRPs), denoted as $\boldsymbol{\sigma}$, are used as the attitude coordinate for the control law and represent the attitude error between the current body frame and the desired reference frame. Similarly, $\delta\boldsymbol{\omega}$ represents the difference between $\boldsymbol{\omega}$ and the desired reference angular velocity, $\boldsymbol{\omega}_r$ [54].

$$\mathbf{L}_{des} = K\boldsymbol{\sigma} + [P]\delta\boldsymbol{\omega} - [\tilde{\boldsymbol{\omega}}][I_{sc}]\boldsymbol{\omega} - [I_{sc}](\dot{\boldsymbol{\omega}}_r - \boldsymbol{\omega} \times \boldsymbol{\omega}_r) \quad (2.5)$$

The proportional control gain on attitude error is denoted as K , and $[P]$ must be a positive definite matrix to guarantee stability via Lyapunov functions.

State noise is added to the system at each timestep in a Gaussian random walk. Gaussian measurement noise is added to $\boldsymbol{\omega}$ and $\boldsymbol{\sigma}$ when calculating the control torques and is also added to the torque command. The state is saved throughout the simulation and then truncated to a lower rate to represent downlinked telemetry. Gaussian measurement noise is added to this telemetry prior to processing. Table 2.1 contains the parameters used in this simulation. All noise is zero mean.

The spacecraft's inertial angular momentum is calculated from the noisy state telemetry via

$${}^{\mathcal{B}}\mathbf{H} = {}^{\mathcal{B}}[I_{\text{sc}}]{}^{\mathcal{B}}\boldsymbol{\omega} + \sum_{i=1}^N [BW]_i \begin{matrix} \mathcal{W} \\ \left[\begin{array}{c} I_{ws}\Omega_i \\ 0 \\ 0 \end{array} \right] \end{matrix} \quad (2.6)$$

$${}^{\mathcal{N}}\mathbf{H} = [NB] {}^{\mathcal{B}}\mathbf{H} \quad (2.7)$$

where ${}^{\mathcal{B}}\mathbf{H}$ is the spacecraft angular momentum in the body-fixed frame, $[BW]_i$ is the direction cosine matrix to go from the i th wheel's coordinate frame to the body frame, $[NB]$ is the direction cosine matrix to go from the body frame to the inertial frame, and ${}^{\mathcal{N}}\mathbf{H}$ is the spacecraft's inertial angular momentum. Note that the momentum telemetry does not have any noise added to it, but it incorporates the noise from the attitude telemetry, rate telemetry, and wheel speed telemetry.

2.1.2 Applying Debris Strike to Simulation

The debris strike is applied to the spacecraft attitude dynamics simulation as a brief torque \mathbf{L} , computed as follows. The net change in momentum imparted by the debris strike is denoted as $\mathbf{H}_{\text{strike}}$ where $\mathbf{R}_{d/s}$ is the location of the strike relative to the spacecraft center of mass, m_d is the mass of the debris, and $\mathbf{V}_{d/s}$ is the velocity of the debris relative to the spacecraft.

$$\mathbf{H}_{\text{strike}} = \mathbf{R}_{d/s} \times m_d \mathbf{V}_{d/s} \quad (2.8)$$

$$\mathbf{H}_{\text{strike}} = \int \mathbf{L} dt \quad (2.9)$$

The torque, \mathbf{L} , is applied in the simulation for one time step, and the torque magnitude in each axis is calculated such that the total change in momentum, $\mathbf{H}_{\text{strike}}$, is as specified in the problem

Table 2.1: Parameters used in spacecraft dynamics simulation

Parameter	Value	Unit
S/C Inertia _{xx}	3000	kg.m ²
S/C Inertia _{yy}	2500	kg.m ²
S/C Inertia _{zz}	3500	kg.m ²
RW Inertia	0.16	kg.m ²
RW max torque	0.2	N.m
RW 1 pointing vector	[1, 0, 0]	-
RW 2 pointing vector	[0, 1, 0]	-
RW 3 pointing vector	[0, 0, 1]	-
RW 4 pointing vector	[1, 1, 1]	-
State noise in attitude	$\sigma=1\text{E-}8$	-
State noise in rate	$\sigma=1\text{E-}8$	rad/s
Meas. noise in attitude sim	$\sigma=1\text{E-}6$	-
Meas. noise in rate sim	$\sigma=1\text{E-}6$	rad/s
Noise in applied torque	$\sigma=0.001$	N.m
Meas. noise in attitude tlm.	$\sigma=7\text{E-}5$	-
Meas. noise in rate tlm.	$\sigma=2\text{E-}4$	rad/s
Meas. noise in RW speed tlm.	$\sigma=0.25$	rad/s

setup. For convenience, the traded debris strike variables are the direction of $\mathbf{H}_{\text{strike}}$ and the mass of the debris. The velocity of the debris is fixed at 8 km/s normal to $\mathbf{R}_{d/s}$, which is 1 m. The mass of the debris is traded to investigate different magnitudes of debris strike, and is generally between 1 and 100 mg. In early exploration the diameter of an aluminum sphere is varied as the independent variable for changing the debris strike magnitude. However, this introduces an unnecessary nonlinearity to the magnitude so in continued developments the mass of the debris is traded instead. Note that the 8 km/s and 1 m are placeholders, not assumptions. The rotational momentum of a strike at 8 km/s and 1 m is equivalent to a strike at 4 km/s and 2 m, or 1 km/s and 8 m, etc. These placeholders are simply convenient to provide some intuition about the sizes of debris strikes that are being discussed, not an assumption on which the analysis depends.

In reality the velocity of debris would be dependent on the debris (or micrometeoroid) population it originates from, and the strike could be anywhere on the spacecraft. The variables are fixed for simplicity and clarity since this chapter focuses on modeling the effect on the spacecraft and developing detection methodologies, not on modeling the debris population. Table 2.2 lists

a variety of impacting particles in various orbit regimes which would impart angular momentum equivalent to a 10 mg particle striking 1 m off the center of mass (CoM) at 8 km/s with normal incidence, imparting 80 mN.m.s of angular momentum to the satellite.

Table 2.2: Particle impact events which impart 80 mN.m.s of angular momentum (assuming normal incidence)

Object	Density	Size
Case #1: 100 mg particle impacting 2 m from CoM at 400 m/s, MEF = 1 (i.e, GEO)		
MLI square ¹	0.04269 g/cm ²	1.53 cm
CFRP square ²	1.8 g/cm ³	0.75 cm
Mylar square ³	0.0032 g/cm ²	5.59 cm
Case #2: 3.64 mg particle impacting 1 m from CoM at 11 km/s, MEF = 2 ⁴ (i.e, LEO)		
Al. sphere	2.7 g/cm ³	1.37 mm dia.
Steel sphere	8.05 g/cm ³	0.95 mm dia.
Copper wire ⁵	8.96 g/cm ³	1.29 cm
CFRP square ²	1.8 g/cm ³	1.42 mm
Mylar square ³	0.0032 g/cm ²	1.07 cm

It is important to note that a strike which breaks through a solar array will impart less momentum than the total relative momentum of the debris, while a hypervelocity strike which breaks up on the surface will cause a plume of ejecta in the reverse direction which can increase the imparted momentum by a factor of two or more [57] due to momentum enhancement. The momentum enhancement factor is included in Table 2.2, but is assumed to be one throughout the developments in this chapter. While this effect is critical for accurately assessing impact events, the emphasis in this chapter is on developing detection techniques, not debris population modeling. When these techniques are applied to on-orbit telemetry and used to estimate parameters in real debris strikes these effects must be accounted for.

When a debris strike is applied to a truth simulation with no noise, the results are as shown

¹ Square of multi-layer insulation, areal density of standard JPL blanket per reference [55]

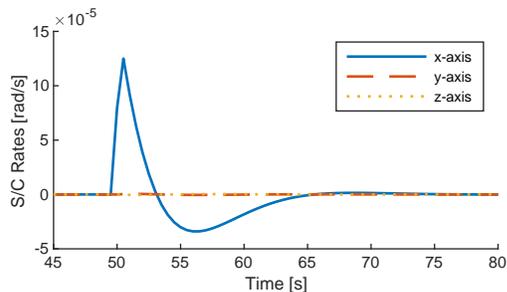
² Square of carbon fiber reinforced polymer, 1 mm thick, density per [56]

³ Double aluminized perforated Mylar, single outer blanket layer per [55]

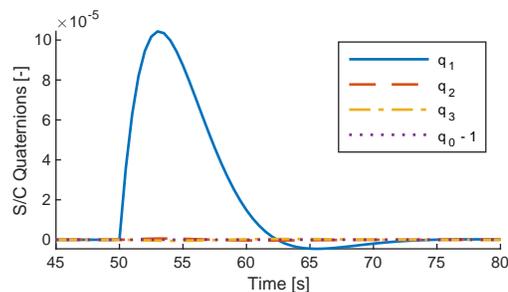
⁴ MEF based on findings in reference [57]

⁵ Length of copper wire 0.2 mm diameter

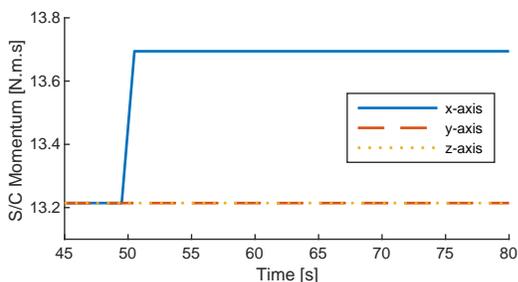
in Figure 2.1. The strike induces a rotation in the spacecraft which is corrected by the attitude control system. It manifests as a spike in the spacecraft rate, a drift and correction in the spacecraft attitude, and a net increase spacecraft momentum. With this simulation of spacecraft dynamics during a debris strike the next task is exploring ways to identify these features in telemetry.



(a) Rate Telemetry



(b) Attitude Telemetry



(c) Momentum Telemetry

Figure 2.1: Response of spacecraft to simulated debris strike (60 mg at 8 km/s, 1 m from CoG, applied at $t=50$ s. $H_{\text{strike}} = 0.48$ N.m.s).

2.2 Exploring Strike Detection and Estimation via Digital Signal Processing Techniques

Digital signal processing is an entire field pertaining to getting information out of noisy data. Techniques from this field are explored here for application to the debris strike detection problem.

2.2.1 Matched Filter Development

Since the response of the spacecraft's state to a certain external impulsive torque can be simulated fairly accurately and the resultant 'signal' parameters are then well known, a matched filter can be applied to telemetry to maximize the signal-to-noise ratio of the spacecraft's response to simulated debris strikes. A matched filter functions by taking a known wavelet (the anticipated signal) and cross-correlating it with a sequence of measurements [58, ch. 5]. The filter output will increase when the signal is present in the noise, as shown in Figure 2.2.

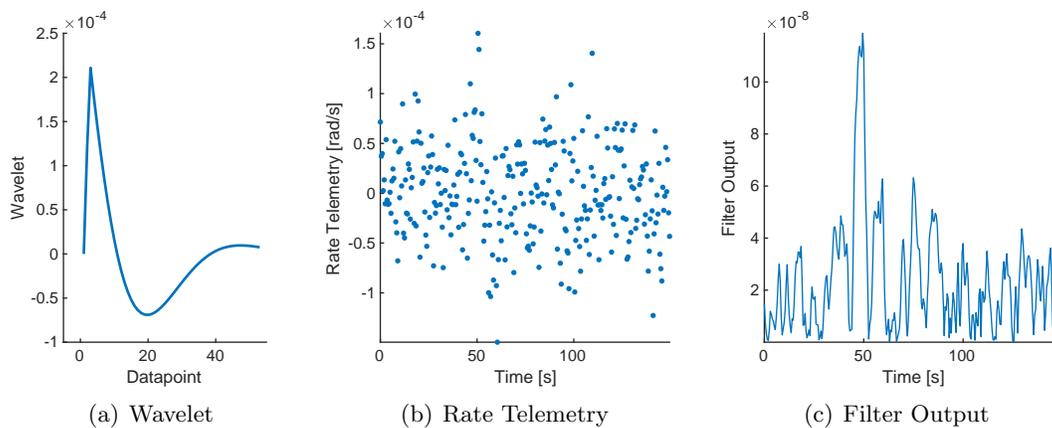


Figure 2.2: Example of matched filter accentuating a 50 mg debris strike applied at $t=50$ s. Note that when the strike is applied the telemetry increases slightly for a few points then decreases slightly, but it is difficult to see in the raw telemetry as it is beneath the noise floor. However, the filter produces a distinct spike in response.

The matched filter is developed by applying a debris strike to the simulation with no noise, and the resulting wavelet is saved for the rotation rate in each spacecraft body frame axis. During initial experimentation it was determined that, since the frequency of the response remains the same until the reaction wheel maximum torque is reached, a large wavelet which envelopes smaller effects produces essentially the same SNR in the filter than a wavelet with the same magnitude of dynamic effect. Thus, for computational efficiency, three wavelets were developed to represent the largest debris strike that could be applied to each individual axis without saturating the reaction wheels. Each wavelet is applied to the corresponding axis' telemetry and captures the effects of

debris strikes that are smaller than the one used to generate the wavelet. This produces the results shown in Figure 2.3. The first plot shows the three wavelets for each axis, the second shows the spacecraft telemetry and true state for a debris strike with the momentum applied to axis 1, and the third plot shows the results from the wavelets cross-correlated with the telemetry, so the output of the matched filter. Note the coupling in the control response between the three axes: even though the strike is applied strictly to axis one, the gyroscopic effects of the wheels induce a slight rotation in the other axes as the spacecraft corrects the rotation caused by the strike. This effect will be seen in future results as well.

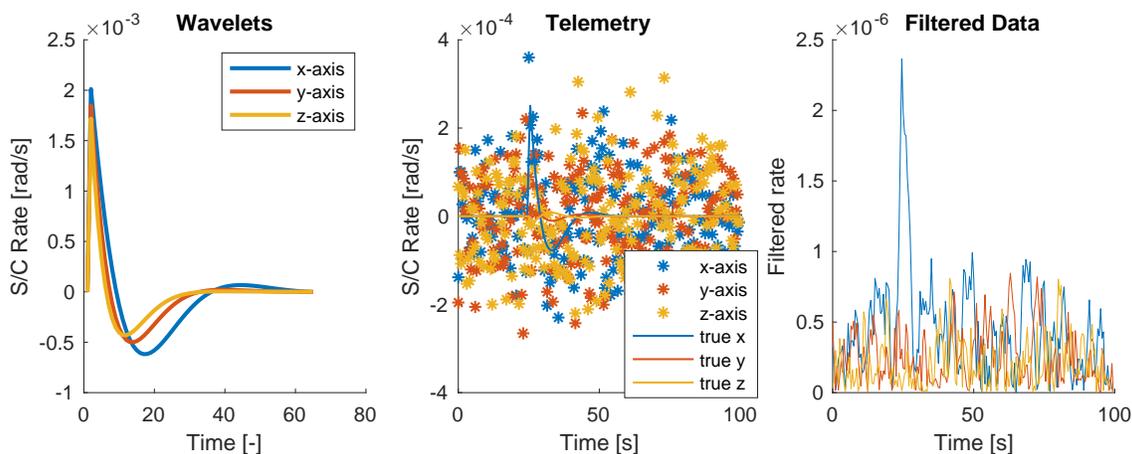


Figure 2.3: Matched filter developed for each spacecraft axis. Differences in inertia result in slightly different wavelets for each axis.

2.2.2 Characterizing Probability Distributions

Since the simulation can be run as many times as necessary, it is possible to characterize the statistical distributions of the system. The probability density functions (pdfs) are developed for both a filter run against the telemetry with no debris strikes present, and to find the characteristics of the filter's spike when debris strikes of certain magnitudes are present. The resulting pdfs are shown in the Figure 2.4 and are used to develop a likelihood ratio test in the next section. Each debris strike pdf indicates the distribution of the maximum spike value when the given size of debris

strike is applied on each axis across 1000 trials. In this graphic the independent variable used to trade strike magnitude is debris diameter assuming an aluminum sphere.

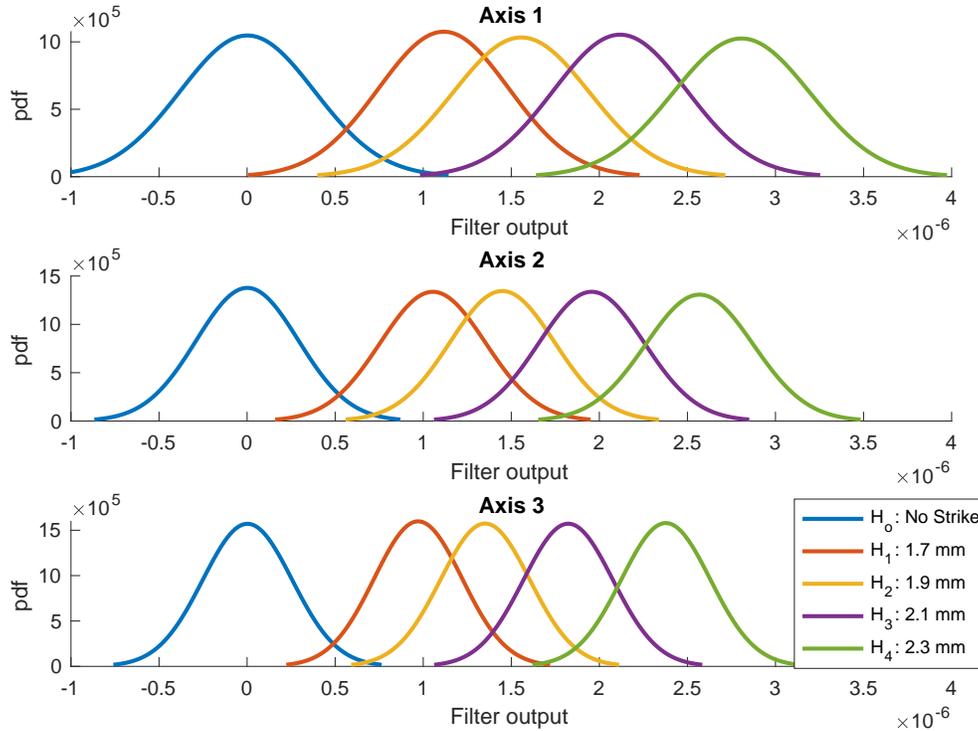


Figure 2.4: Probability density functions for filter output

2.2.3 Developing Likelihood Ratio Test

From the probability density functions a likelihood ratio test (LRT) is developed. The ratio of the probability of no debris strike to the probability of a debris strike of the sizes used to develop the pdfs is assessed for each potential filter output, and plotted in Figure 2.5. The log-LRT is used since the results become quite large, eventually overwhelming the computational capabilities of Matlab. As expected, the likelihood ratio of each hypothesis is the highest in the range where the pdf from that size strike is centered.

To select a threshold for the LRT the receiver operating characteristic (ROC) curves are developed for each hypothesis. Using the Neyman-Pearson Lemma, a maximum Probability of False

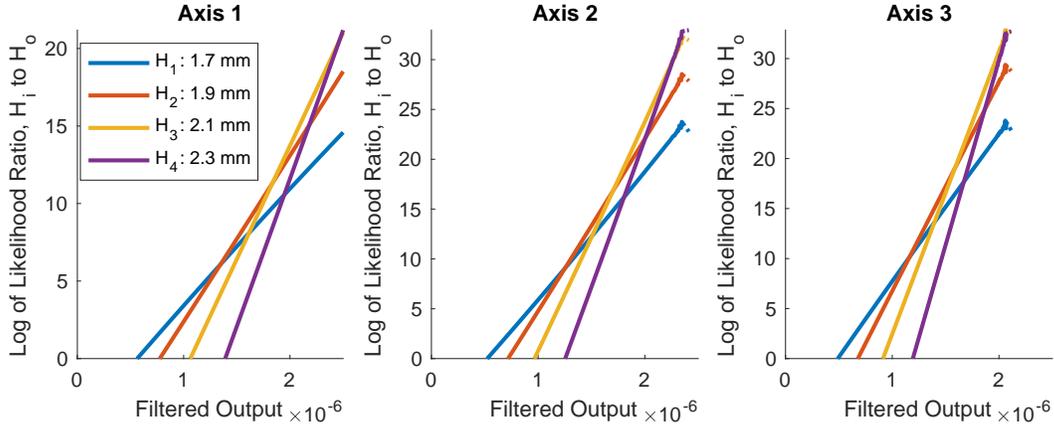


Figure 2.5: Highest LLRT curve corresponds to most likely strike size

Alarm (P_F) is selected as 0.1%. 5% was originally tried, since that is often used as a threshold for rejecting the null hypothesis, but that threshold produced far too many false alarms due to the high data rates under consideration. Later work explores various ways to treat this problem, eventually establishing a rate of false alarm which takes into account the data rate. Figure 2.6 shows the ROC curve for the first axis, both the overall curve and the curves zoomed in to show the probabilities of detection for each hypothesis at the selected false alarm threshold. Note that at the established probability of false alarm there is a 40% chance of detecting a small 1.7 mm strike, but an 80% chance of detecting a 1.9 mm strike.

The slope of the lowest ROC curve at the desired P_F is used as a threshold for the likelihood ratio test. This produces the following thresholds derived from the LRT curves, as shown in Figure 2.7.

This results in thresholds to perform M-ary hypothesis testing on the filter output from the matched filters previously generated. Based on the filter output, the debris strikes can be detected when they cross the lower threshold and assigned to one of the hypotheses based on spike size, assigning each detected strike to one of the size classes and therefore producing a coarse estimate of strike size in each axis.

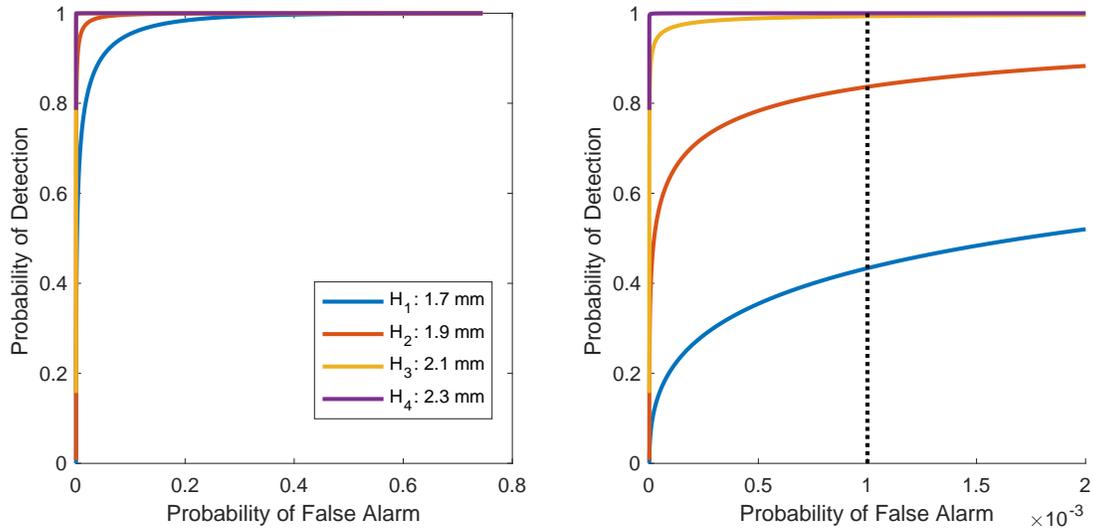


Figure 2.6: Receiver Operating Characteristic curve for axis 1

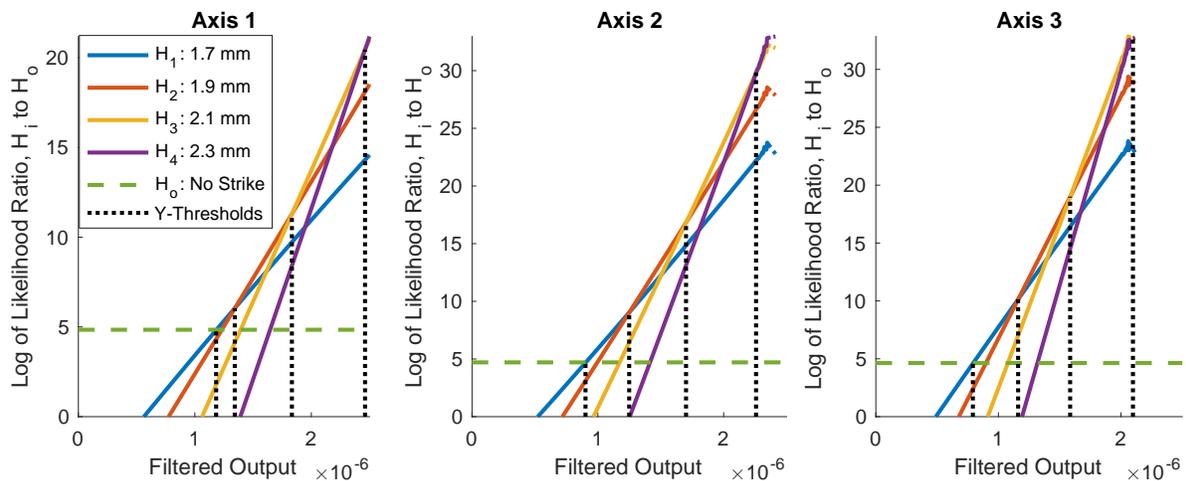


Figure 2.7: LLRT curves with specified detection threshold and bins

2.2.4 Using MMSE to Estimate Strike Parameters

In order to produce a more refined estimate of the strike size, the minimum mean-squared-error is used. This is accomplished by developing a bank of wavelets corresponding to the telemetry response to various sizes of debris strikes, stepping up from small strikes below the detection threshold to large strikes which nearly cause the reaction wheels to saturate. When a strike is

detected using the M-ary detection techniques, the bank of wavelets is run on the telemetry at the time at which the strike is detected, and the mean-squared-error (MSE) between each wavelet and the telemetry is calculated. This MSE is at a minimum when the size of the wavelet is closest to the size of the spacecraft's response to the debris strike, so the strike size which was used to create that wavelet is selected as the most probable size of strike. Since it's unlikely, in reality, that debris strikes would be confined to strictly one axis, when a debris strike is detected in one axis the wavelet bank is applied to all three axes, to determine if momentum has been transferred to the spacecraft on those axes too.

This technique is looped over 1000 runs using a randomized strike magnitude and direction, to characterize the algorithm's ability to correctly detect strikes and assess strike size. Figure 2.8 indicates the relationship between the true strike size and the estimated strike size on each axis for each run. Undetected strikes show as zeros. The algorithm consistently overpredicts the size of the strike by an amount corresponding to a 0.65 mm increase in the diameter of the aluminum sphere. This is due to the coupling between the dynamics, since a strike applied to one axis also produces a response in the other axes, the response in each axis to off-normal strikes is consistently larger than the response to an equivalent strike on only that axis. To improve the accuracy models should be developed which take this coupling into account.

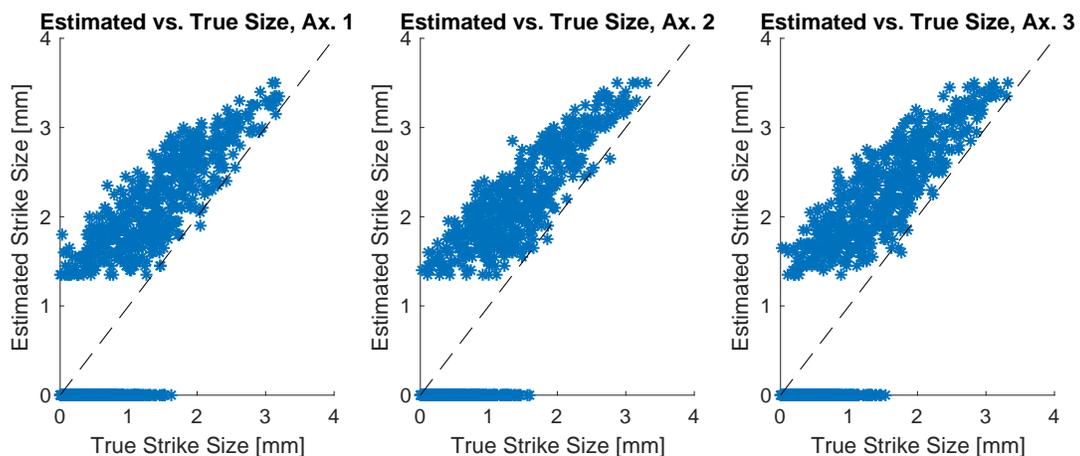


Figure 2.8: Estimated strike size via MMSE vs. true strike size

To further characterize the performance, the rate of detection vs strike size is shown in Figure 2.9. As expected, the rate of detection improves dramatically for larger strikes, with all strikes over 2 mm (diameter, aluminum sphere) detected. This corresponds well to the probabilities of detection expected from the receiver operating characteristic curve.

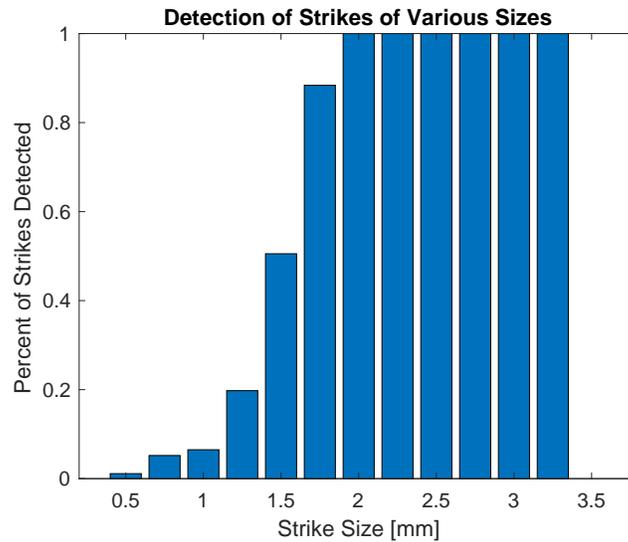


Figure 2.9: Percentage of strikes detected vs strike size

These results indicate that signal processing techniques can be used to detect debris strikes in spacecraft telemetry and assess the magnitude of those strikes.

2.3 Methods for Strike Detection

2.3.1 Refining Digital Signal Processing Techniques for Strike Detection

These matched filters are applied to the rate and attitude telemetry to identify debris strikes, since the spacecraft's dynamic response to the strike produces a known 'signal' in the telemetry. A matched filter maximizes the output signal-to-noise ratio (SNR) for a known signal in independent and identically distributed Gaussian noise.

From the last section, to determine the threshold for detecting a debris strike based on filter output, a receiver operating characteristic (ROC) curve is developed empirically and a desired

probability of false alarm, α , is selected. The ROC curve shows the probability of false alarm (P_{FA}) vs. the probability of detection (P_D), which is based on two probability density functions (pdfs). The first pdf characterizes the filter output when no debris strike is present (i.e., the null hypothesis, \mathcal{H}_0). The second pdf characterizes the filter's peak when a certain debris strike is present (\mathcal{H}_1). Figure 2.10 shows the pdfs underlying the ROC curve, which is a more intuitive way to understand the effect.

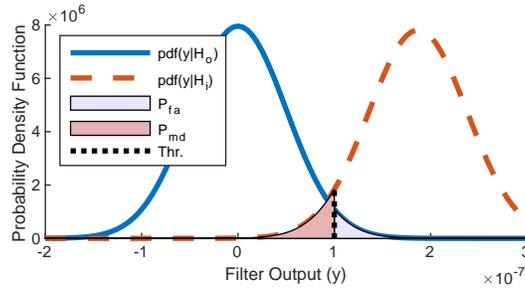


Figure 2.10: Probability density functions of filter output with no strikes present and during a small strike

Note that in order to decrease the P_{FA} the threshold is moved to the right, but that necessarily decreases the P_D for that strike and increases the probability of a missed detection, P_{MD} . Thus, selection of a desired P_{FA} is a design choice for a given application. A low P_{FA} will detect significant perturbations with little risk of false alarms, but will miss small but detectable perturbations. A higher P_{FA} will produce more false alarms but will also catch the small, subtle perturbations. Also note that while Gaussian pdfs are shown in the figure they will only generate accurate thresholds if the data is Gaussian. Therefore, the distribution of the data should always be checked and appropriate pdfs selected before applying this method.

From the P_{FA} and P_D the ROC curves in Figure 2.11(a) are developed. A desired P_{FA} is selected, and the slope of the ROC curve at the point where $P_{FA} = \alpha$ is the threshold for the likelihood ratio test, τ_{thresh} . The likelihood ratio test, \mathcal{L} , is also based on the two pdfs of the filter output where

$$\mathcal{L}(\mathbf{y}) = \frac{p(\mathbf{y}|\mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \tau_{\text{thresh}} \quad (2.10)$$

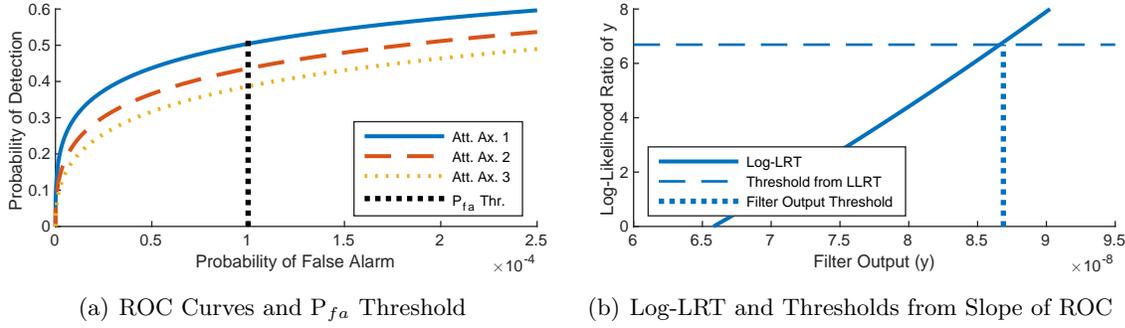


Figure 2.11: Development of filter thresholds via ROC curve and LRT

Where $p(\mathbf{y}|\mathcal{H}_i)$ is the probability of \mathbf{y} given hypothesis \mathcal{H}_i . The resulting $\mathcal{L}(\mathbf{y})$ is compared to the threshold τ_{thresh} from the ROC curve and the appropriate hypothesis is selected. This threshold is mapped back to the filter output by finding the point at which τ_{thresh} intersects the $\mathcal{L}(\mathbf{y})$ line and choosing the corresponding $\mathbf{y}_{\text{thresh}}$ as the filter output threshold (Figure 2.11(b)).

$$\mathcal{L}(\mathbf{y}_{\text{thresh}}) = \tau_{\text{thresh}} \quad (2.11)$$

Note that the log of the likelihood ratio test (LLRT) is plotted along with the log of the threshold since the numbers become large. Under the Neyman-Pearson Lemma, this is the most powerful test that maximizes P_D under the constraint $P_{FA} = \alpha$ [59, ch. 2]. For the purposes of debris strike detection the filter output is compared to the threshold $\mathbf{y}_{\text{thresh}}$. If the filter output is above this threshold then a strike is declared and the strike parameters are estimated.

2.3.2 Change Detection Techniques

The second set of techniques applied to attitude control system telemetry are change detection techniques, which are described here. The inertial angular momentum of a spacecraft is quiescent in the absence of external forces, but a debris strike imparts an abrupt change in the momentum. Three techniques for detecting changes in quiescent but noisy data are applied to the spacecraft momentum: a simple summation filter, a more refined cumulative sum (CUSUM) sequential probability ratio test (SPRT), and a Shiryaev SPRT.

The simple summation filter detects changes by adding the sum of all datapoints after zeroing the data to the expected average. For zero-mean noise, this sum trends to zero over time so that a value above a certain threshold indicates that a fault (or change) has occurred.

The CUSUM algorithm utilizes the log of the likelihood ratio test from equation (2.10), summed sequentially to give the test statistic S_n [58, ch. 9].

$$S_n = \sum_{k=1}^n \ln(\mathcal{L}(\mathbf{y})) \quad (2.12)$$

The parameter S_n trends negative when the samples, as a whole, are more likely to be from \mathcal{H}_o than \mathcal{H}_1 , and trends positive when they are more likely from \mathcal{H}_1 . This change in drift is detected through

$$W_n = S_n - \min_{0 \leq k < n} S_k \quad (2.13)$$

where W_n stays close to zero while S_n trends downward, then grows if it trends upward.

The multi-hypothesis Shiryaev SPRT is implemented for m alternative hypotheses per Malladi and Speyer [60]. They define

$$\phi_{ki} = F_{ki} + \tilde{p}_i \cdot (1 - F_{ki}) \quad (2.14)$$

where F_{ki} is the cumulative distribution function (CDF) expressing the probability that, at datapoint k , hypothesis i is true. \tilde{p}_i is the *a priori* probability that a transition to hypothesis i occurs. $F_{k+1,i}$ is computed recursively given each new datapoint x_{k+1} via equation (2.15), where $f_i(x_{k+1})$ is the pdf given hypothesis i evaluated at x_{k+1}

$$F_{k+1,i} = \frac{\phi_{ki} \cdot f_i(x_{k+1})}{\sum_{j=1}^m \phi_{kj} \cdot f_j(x_{k+1})} \quad (2.15)$$

The CDFs are initialized to F_{oi} , which is the probability that a transition to hypothesis i has already occurred.

2.4 Algorithm Development

2.4.1 Matched Filter Development

2.4.1.1 Develop Filters to Identify Strikes

As described in Section 2.2, a truth simulation is used to develop matched filters for the spacecraft rate and attitude when working with simulated telemetry. To develop wavelets for a real spacecraft, the spacecraft dynamics simulation used during development can be used to determine initial wavelets. Once the satellite is in orbit, the spacecraft's response to other momentum-changing events can be used to tune the wavelets. For example, a spacecraft's response to a slight angular momentum imparted during a small maneuver has similar characteristics to the response to a debris strike. Figure 2.12 shows some telemetry from NASA's Solar Dynamics Observatory (SDO) recovering its attitude after a maneuver.

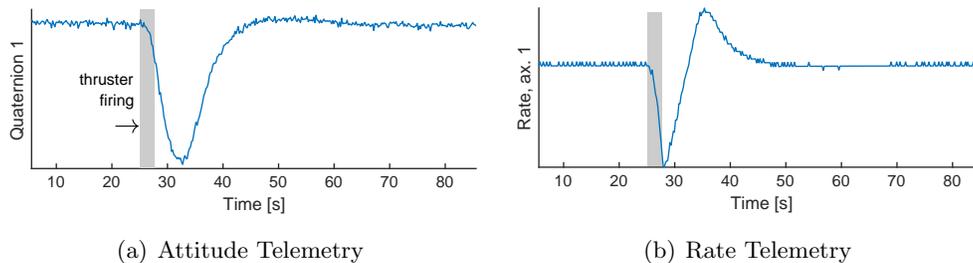


Figure 2.12: Example of the spacecraft's correction after a small maneuver imparts a slight rotation. Note similarities to simulated debris strike response.

Note that the attitude wavelets are only run on the vector components of the error quaternion, since the scalar component will always be close to one for small errors. Under the small angle approximation, the scalar component will remain near one but the vector components will vary linearly with $\frac{\phi}{2}$, where ϕ is the angle of the error in each axis.

2.4.1.2 Developing Filter Thresholds

The filter thresholds are developed using a Receiver Operating Characteristic (ROC) curve. To develop this curve, pdfs of the filter output are developed empirically. The first pdf is the filter output with no strikes present, to characterize the noise in the filter. Then a specified strike is applied to the simulation 1000 times and the filter response to each strike is logged and used to generate a pdf of filter response to that strike. The strike size is selected such that the center of the bell curve is in the vicinity of the desired threshold on the no-strike pdf. To achieve this, 30 mg strikes are applied to generate the filter response for attitude thresholds, and 40 mg strikes are applied to generate the filter response for rate thresholds. From the two pdfs the ROC curves are generated, and the slope of the ROC curve at $P_{fa} = 0.01\%$ corresponds to the threshold for the likelihood ratio test [59]. The likelihood ratio test is applied to possible filter outputs and the filter output at which the likelihood ratio test exceeds the threshold from the ROC curve is selected as the filter output threshold for strike detection. See Figure 2.11 for a graphical representation of this process.

2.4.1.3 Determining Strike Parameters

Once a debris strike is detected using the matched filter output and associated threshold, the task is to assess the magnitude of the debris strike. This is accomplished using a bank of matched filter wavelets and determining which filter wavelet shows the closest match to the telemetry. Then the size of strike used to generate that wavelet is selected as the estimated strike size. Refer to Section 2.1.2 and Table 2.2 for a detailed discussion on the relationship between strike size as traded in this chapter and debris strike parameters.

Four methods are compared to determine which wavelet is closest to the telemetry. The first method takes the root-sum-squared error (RSSE) between the telemetry and each wavelet then selects the wavelet where this is a minimum. The second takes the minimum mean-squared error (MMSE), the third the minimum absolute error (MAE), and the fourth the sum of the absolute

value of the error cubed ('Cubic'). The results show that the MMSE and RSSE have identical performance, which makes sense since the errors are penalized identically in both methods - as a function of the square of the error. The absolute error is penalized linearly while the fourth method penalizes the error cubed. Across 1000 Monte-Carlo runs with randomized debris strikes, the error in estimated strike magnitude for each method is shown in Table 2.3. Note that these errors are for accurately detected strikes only, Section 2.4.1.4 discusses the overall detection performance of the algorithm. Based on these results, the MMSE is selected as the estimation method for the remainder of this study. Figure 2.13 shows a wavelet bank, simulated telemetry with the true state overlaid, and the results of each estimation method along with the true strike size.

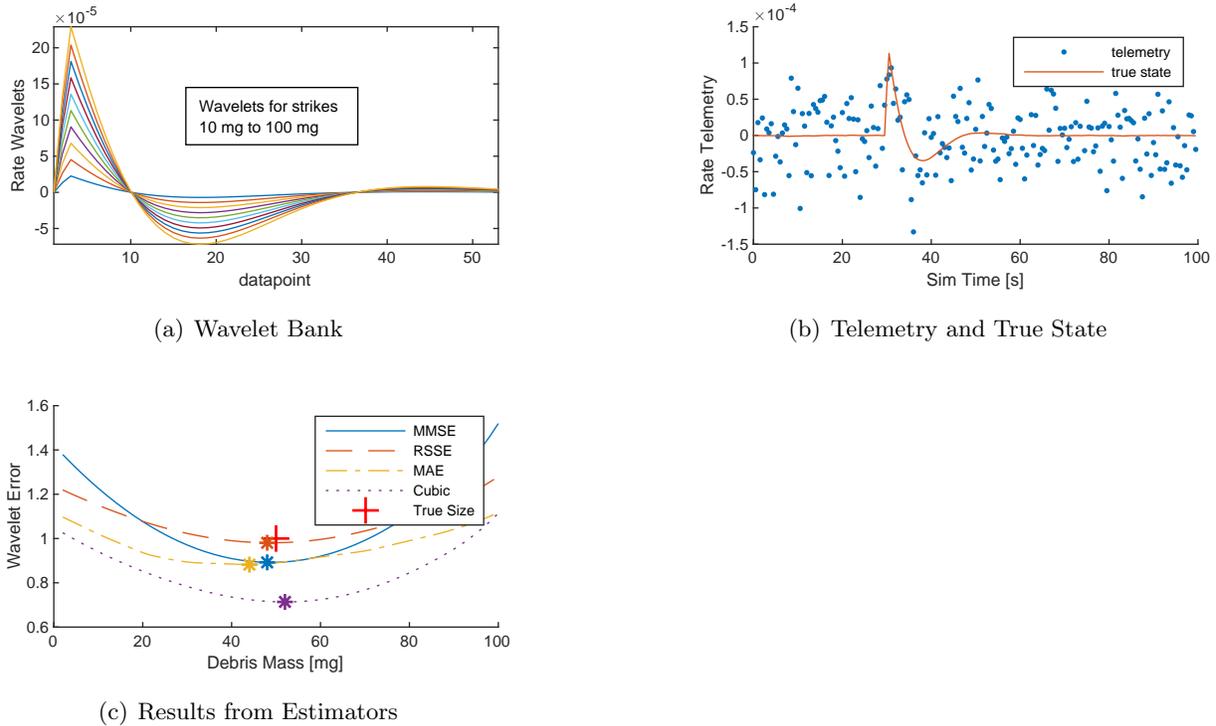


Figure 2.13: Wavelet bank and results from running filters against a 50 mg strike at $t=30$ s. Minimum error is selected as most probable strike size (48 mg). Results from each estimator are compared.

Table 2.3: Comparing error in estimated strike magnitude for each estimation method

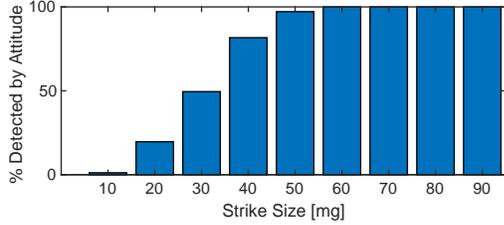
Estimation Method	Mean Error	Standard Dev.
Using Attitude		
MMSE	1.4 mg	8.7 mg
RSSE	1.4 mg	8.7 mg
MAE	2.0 mg	10.9 mg
Cubic	1.5 mg	9.1 mg
Using Rate		
MMSE	-4.46 mg	13.1 mg
RSSE	-4.46 mg	13.1 mg
MAE	-2.95 mg	14.3 mg
Cubic	-4.50 mg	13.2 mg

2.4.1.4 Assessing Algorithm Performance

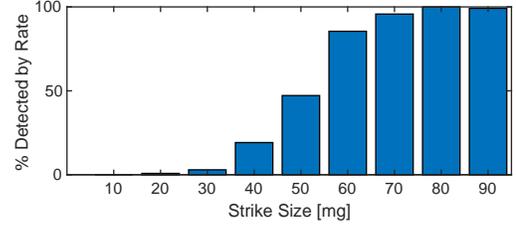
To assess the algorithm's performance, a 1000 run Monte-Carlo is conducted randomizing the strike time, size, and direction. The enveloping matched filters from Section 2.4.1.1 are run against the telemetry to detect strikes per the thresholds developed in Section 2.4.1.2. When they detect a strike the wavelet bank is run at that time and the MMSE is computed. The closest-matching wavelet is selected as the strike's estimated size, and that result is compared to the true size. The results, shown in Figure 2.14, indicate that the attitude filter consistently detects strikes larger than 60 mg, while the rate filter consistently detects strikes larger than 80 mg. Both filters also detect smaller strikes, but the percentage detected decreases with smaller strikes. The estimated strike sizes cluster in the vicinity of the true strike sizes, but vary by around ± 20 mg. A fair number of outliers are under-predicted by more than that, especially via the rate telemetry.

2.4.2 Sequential Probability Ratio Tests (SPRTs)

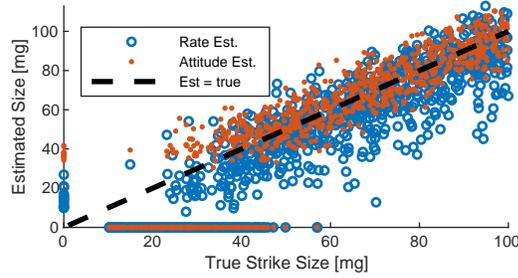
Three SPRTs are implemented on the angular momentum telemetry, as the inertial angular momentum is theoretically quiescent in the absence of external torques while a debris strike produces an abrupt change in momentum. SPRTs compare the noisy output data to an expected probability distribution and the filter output increases when a prolonged deviation from the expected output occurs. Since the momentum changes gradually due to effects like solar radiation pressure the



(a) Strikes Detected by Attitude Filters



(b) Strikes Detected by Rate Filters



(c) True Size vs. Estimated Size

Figure 2.14: Detection results per attitude matched filter and rate matched filter, and plot of estimated vs. true size. Non-detected strikes and false alarms shown as zeros on each axis.

parameters of the expected distribution change over time. To apply SPRTs to the debris strike problem a sliding window filter is used. This sliding window incorporates a ‘pre-window’ which is used to compute the expected distribution parameters, providing a fading estimate of the most recent distribution parameters. In the ‘post-window’ the SPRT algorithms are applied to detect changes from the distribution characterized in the pre-window. This causes the filters to peak at the debris strike time, when all data in the pre-window is before the strike and all data in the post-window is after the strike.

The simple summation filter subtracts the average of the pre-window from the post-window data and then sums the zeroed post-window data. When the mean of the post-window is similar to the mean of the pre-window the filter output, y_{sum} , is near zero. When a change occurred at the junction between the windows the filter output reaches a maximum. For m datapoints,

$$y_{\text{sum}} = \sum_{i=1}^m (x_i - \mu_{\text{pre-window}}) \quad (2.16)$$

The CUSUM algorithm sums the likelihood ratio of each datapoint given the μ and σ calcu-

lated from the pre-window compared to two alternative hypotheses: a displacement of $\Delta\mu$ N.m.s in either the positive or negative direction. Then the test statistic W_n from Equation (2.13) is used as the filter output. This also produces a maximum at approximately the time of the strike.

The Shiryaev algorithm applies the recursive relation in Equation (2.15) to each data point for j alternative hypotheses, where each \mathcal{H}_j is a displacement of $n\Delta\mu$ N.m.s from the μ of the pre-window and \mathbf{n} is a 1D array of integers from -10 to 10. The final CDFs ($F_{j,\text{end}}$) at the end of the post window are used as a weights to produce the following weighted sum as the estimated change in μ based on the entire post-window measurement sequence

$$\Delta\mu_{\text{est}} = \sum_{i=1}^j n_i \Delta\mu F_{i,\text{end}} \quad (2.17)$$

Figure 2.15 shows the inertial angular momentum telemetry with a small, subtle debris strike. The change in distribution is not readily apparent from the raw data, but the highlights in Figure 2.15(b) show that the distribution underlying the noisy data has changed abruptly. This telemetry shows a 10 mg strike, and Figures 2.16 and 2.18 show that the filters respond favorably even to subtle changes like this.

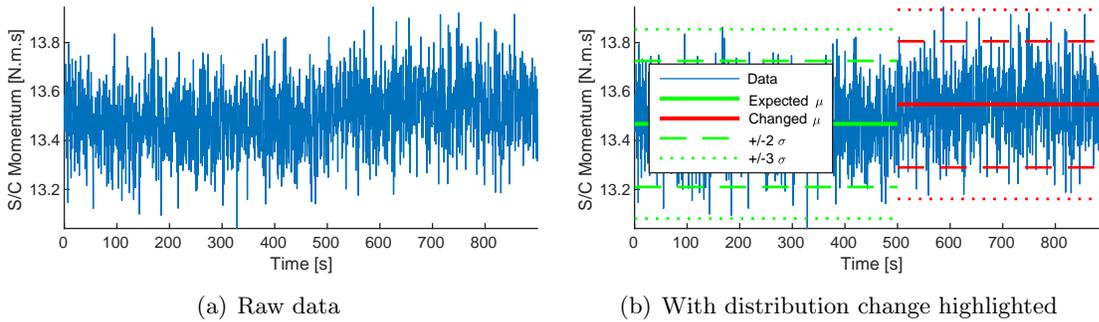


Figure 2.15: Raw momentum telemetry with 10 mg strike at $t=500$ s

2.4.2.1 Selecting Filter Parameters

The performance of each of the SPRTs is dependent on the filter parameters such as the length of the windows and the choice of $\Delta\mu$. The length of the pre-window is fixed at 200 datapoints and

the post-window at 180 datapoints to allow reasonable computation times during long Monte-Carlo runs. The *a priori* probability of change for the Shiryaev algorithm, \tilde{p} , is fixed at 1×10^{-6} .

A $\Delta\mu$ of 0.1 N.m.s is chosen for the CUSUM algorithm, which corresponds to a 12.5 mg strike. This threshold could be lowered to detect smaller strikes, but that would result in additional noise in the absence of strikes. Selection of the $\Delta\mu$ threshold is a design choice for a given application, depending on whether it is desirable to detect any potential small strikes with a higher probability of false alarm or preferable to detect only the larger strikes with more confidence. The Shiryaev $\Delta\mu$ array is chosen to achieve the desired granularity in strike size estimates. A granularity of 5 mg in strike size is selected as the granularity for the Shiryaev algorithm, which corresponds to a $\Delta\mu$ of 0.04 N.m.s.

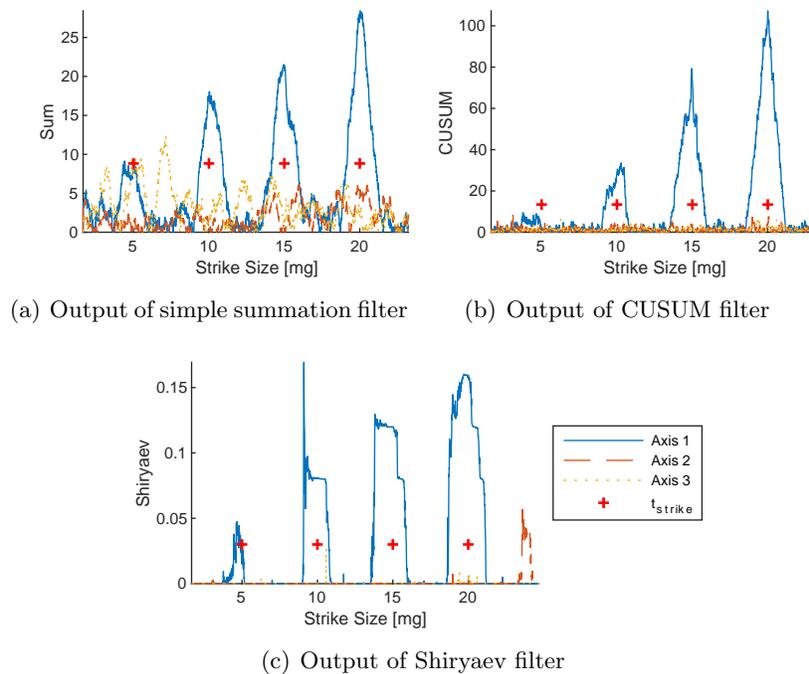


Figure 2.16: Output of SPRT filters with 5, 10, 15, and 20 mg strikes applied to the x-axis

With these parameters, the filters produce the results shown in Figure 2.16 when strikes with magnitudes 5, 10, 15, and 20 mg are applied in the x-axis. As shown, the filters struggle to detect the 5 mg strike but the 10 mg strike and higher are clear. The CUSUM algorithm and summation

filter both show peaks at approximately the time of the strike, but the Shiryaev algorithm tends to have a plateau in the vicinity of the strike. This is because this algorithm is developed to provide quick change detection, and estimates the parameters directly. As soon as the end of the post window starts incorporating the change the algorithm begins predicting that the change has occurred, then the estimates for the change are refined, becoming less noisy as more of the post-window includes the change. When the window passes the point of the change the estimates for the change drop off as the pre-window starts incorporating the change in its fading assessment of the \mathcal{H}_o parameters. The simple summation algorithm has a noticeable peak for the smallest strike, while the CUSUM does not because it is well below its detection threshold. However, the noise in the simple summation algorithm is nontrivial, while the CUSUM has an excellent SNR for the three larger strikes. The Shiryaev algorithm seems to have a high SNR, but it also has false alarms just within this short run - a large false alarm on axis 2 at the end of the run, and a small spike on axis 3 that doesn't quite cross the threshold just after the 10 mg strike.

2.4.2.2 Developing Filter Thresholds

To develop detection thresholds the pdf of filter output is used, similar to Section 2.4.1.2. However, the filter output for the SPRT filters is not always Gaussian like the matched filter output, so each algorithm's threshold development is tailored to accommodate its unique output distribution parameters. The noise distribution for the output of each algorithm when no debris strikes are present is developed using a simulation with 40,000 filtered datapoints, while the noise distribution for the output when a strike is present is developed by applying a 10 mg strike to the simulation 1,000 times and measuring the peak response. Figure 2.17 shows the Q-Q plots of the filtered data using distributions specific to each algorithm.

2.4.2.3 Summation Filter Output Distributions

The output of the summation filter with no debris strikes is modeled fairly well by a Gaussian distribution. The filter's output has slightly lighter tails than a Gaussian distribution, so the

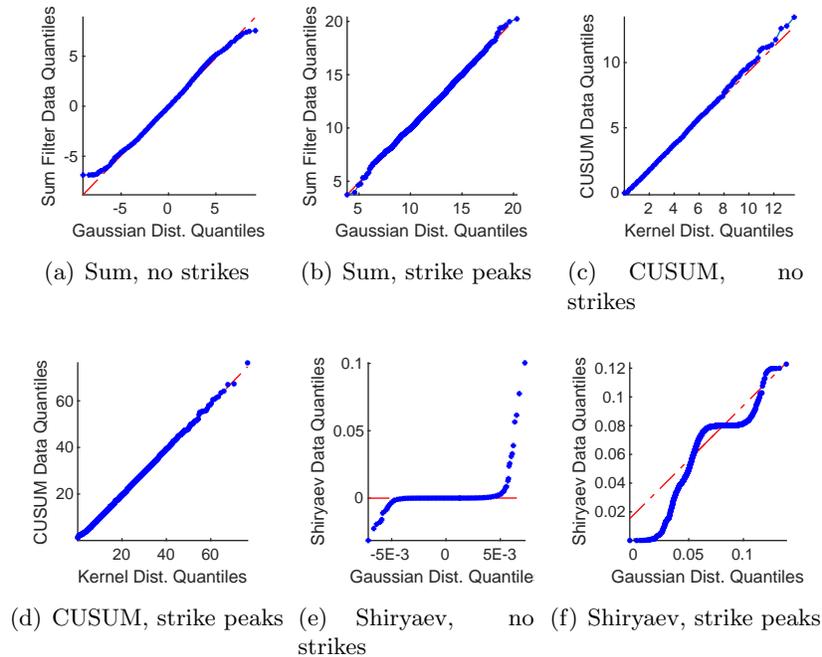


Figure 2.17: Q-Q plots illustrating results when fitting various distributions to filter output

thresholds developed using a Gaussian are somewhat conservative in terms of false alarm rate. The peak of the filter output when a debris strike is applied is also fairly Gaussian, so the method outlined in Section 2.3.1 is used to develop thresholds for strike detection using a desired P_{fa} of 0.01%. Based on the ROC curve, the selected threshold for the summation filter is 8.8.

2.4.2.4 CUSUM Output Distributions

The CUSUM algorithm has a one-sided output and thus requires a one-sided pdf. However, standard one-sided pdfs have lighter tails than the data from the CUSUM algorithm producing more false alarms than desired. To avoid this, a kernel distribution is used. When a debris strike is present the variation in the peak of the CUSUM algorithm is two-sided but skewed, and a kernel distribution is also used for that model. The pdfs generated by these distributions are used to develop the detection thresholds for the CUSUM algorithm per Section 2.3.1. This results in a threshold of 13.4 for the CUSUM algorithm. Note that 10 mg is slightly below the size of strike that the CUSUM algorithm is tuned to detect, so the output is variable resulting in a broad pdf.

2.4.2.5 Shiryaev Output Distributions

The Shiryaev algorithm's output is highly non-normal and defied all attempts to model it with standard pdfs. This applies to both the no-strike noise parameters and the output during debris strikes. Therefore, instead of using a poorly-fitted and non-representative ROC curve the thresholds are developed empirically from the no-strike noise data. Based on this data, a threshold of 0.03 N.m.s, or 75% of the smallest strike the algorithm is tuned to detect is selected. This threshold produced a P_{FA} of 0.4% in one long simulation, but the threshold's true P_{FA} is not assessed analytically due to the lack of accurate pdfs characterizing the data.

2.4.2.6 Assessing Detection Performance

A 1,000-run Monte-Carlo is conducted to assess each algorithm's ability to detect debris strikes using the thresholds developed in Section 2.4.2.2. The Monte-Carlo randomizes the strike magnitude and direction, so it is applied on any combination of the three axes. Figure 2.18 shows the detection accuracy of each algorithm for various strike sizes. These results indicate that detection performance is excellent for all three algorithms for strikes greater than 10 mg. Note that strike direction is randomized as well as strike magnitude, so each axis sees only a fraction of each strike.

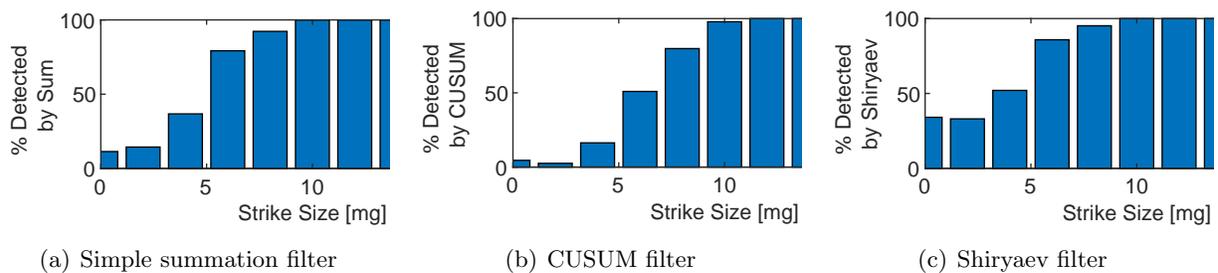


Figure 2.18: Detection performance of each filter on 1,000 randomized strikes

The detection results for the CUSUM algorithm taper off more quickly than the Shiryaev or summation filter results, but this is likely due to the filter being tuned to detect strikes slightly larger than 10 mg. If the filter were tuned to detect smaller strikes it would likely perform a

little better, as the SNR is exemplary as shown in Figure 2.16. However, the filter output noise would increase if it were tuned to detect smaller strikes so the current tuning is maintained as a conservative and reliable filter.

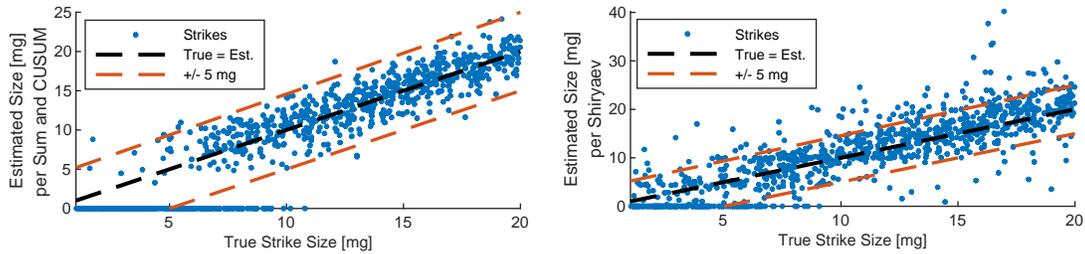
The Shiryaev and summation filters show an ability to detect some strikes even in the < 3 mg range. It is likely that some of these strikes are false alarms. In a run with no strikes and 20,000 datapoints the Shiryaev filter produced eight false alarms while the summation filter produced one and the CUSUM filter produced zero. In a run the length of this Monte-Carlo simulation there would likely be well over 100 false alarms for Shiryaev and a dozen or so for the summation filter, although many of those would be hidden within the response to real strikes.

2.4.2.7 Estimating Strike Parameters

When a strike is detected by the simple cumulative sum filter or the CUSUM algorithm an estimate of strike parameters is obtained by comparing the average of the pre-window data to the average of the post-window data for each axis. The difference in averages is used as an estimate of the momentum imparted by the strike. With the Shiryaev algorithm's weighted sum the output of the filter corresponds to the estimated strike size, so the filter output is used as an estimate of strike parameters. Figure 2.19 shows the accuracy of each filter's estimate of strike magnitude by plotting the true vs. estimated strike size. Non-detected strikes are shown on the x-axis. These results show that both methods do an exemplary job of estimating strike parameters. The estimates from the Sum and CUSUM algorithms tend to be more accurate, with nearly all the errors within ± 5 mg. The results from the Shiryaev algorithm are good as well but have significantly more outliers with poor estimates.

2.4.3 Results from Algorithm Development

The matched filters perform well for detecting strikes larger than about 40 mg, and can estimate the strike size to within around ± 20 mg. Matched filters offer the advantage of Gaussian filter output which allows high confidence in the filter's false alarm rate. This is key when these methods



(a) Estimation performance of CUSUM and sum filters (b) Estimation performance of Shiryaev filter

Figure 2.19: True size of each strike in 1,000 run Monte-Carlo compared to estimated size using two estimation methods

are used to evaluate debris populations using on-orbit telemetry, since the level of confidence in the results must be understood to evaluate whether the measured strikes are reasonable based on the modeled fluxes. Note that the exact filter thresholds and estimation accuracy are dependent on the spacecraft parameters and telemetry noise parameters used in this assessment, and will vary with different parameters for different spacecraft. The next section discusses some preliminary results when trading spacecraft parameters.

The performance of the SPRT methods run against the inertial angular momentum telemetry is outstanding. They are able to detect debris strikes reliably down to 10 mg and occasionally down to 3 mg or less. Their estimation performance is also excellent, generally estimating strike magnitudes to within ± 5 mg. This excellent performance compared to the matched filters is especially noteworthy because the momentum telemetry is calculated from the noisy attitude, rate, and noisy wheel speed telemetry, whereas the matched filters only see either rate or attitude noise. This shows that the SPRT algorithms' ability to identify subtle changes in noisy telemetry is exceptional.

Each SPRT algorithm has unique strengths and weaknesses. The summation algorithm offers extremely low computational requirements, and very little tuning or *a priori* knowledge is required to run it effectively. Its noise parameters are also fairly well modeled which allows the selection of thresholds with predictable false alarm rates. A drawback is that it has the lowest SNR of any

of the algorithms which may degrade its performance on orbit when the measurement and state noise is not necessarily Gaussian. However, its extreme computational simplicity offers a potential on-board application. State-of-health telemetry is typically downlinked at substantially lower rates than on-board data processing to conserve bandwidth. Higher rate data could assist with accurate detection of small strikes. This simple algorithm could be used to screen for strikes using the high-rate on-board telemetry, and buffered telemetry could be downlinked at a higher rate if a probable strike is detected. This would offer more information to definitively establish if a strike has occurred and to assess any effects of the strike.

The CUSUM algorithm offers consistent performance and a very high SNR, with noise parameters that, while not Gaussian, did allow modeling via kernel distributions. One feature of this algorithm is its tunability: it allows a design choice about the size of strike to detect which must be balanced against the risk of false alarms. A drawback is that it has a non-trivial computational burden as the likelihood ratios must be computed for every datapoint in every sliding window. However, for running on telemetry as it is downlinked or running on historical telemetry it performs very well.

The Shiryayev algorithm is challenging in its highly non-Gaussian distribution parameters which make it difficult to assess the false alarm rate for a given threshold. It has very good performance but is more finicky than the CUSUM algorithm, and it is difficult to determine the time of the strike precisely. This could be a detraction if trying to identify the effects of the strike in multiple telemetry streams. However, this difficulty is partly due to its rapid response to strikes, which could be an asset in some applications where only a small window of data is usable or when quick strike detection is desired. Additional tuning may help to improve the false alarm characteristics of this algorithm by adjusting \tilde{p} and experimenting with various distributions for strike size hypotheses.

2.4.4 Discussion of Methods

Applying these techniques to on-orbit telemetry is a non-trivial challenge, as the noise and events experienced by a real spacecraft can trip the detection algorithm thresholds even when no strike is present. While the theoretical inertial angular momentum is perfectly quiescent in the absence of external forces, the reality is more variable. A drift in the inertial angular momentum will be interpreted as a debris strike by the SPRTs if it produces a significant change between the pre-window and post-window. Therefore, the windows must be kept short enough that any unmodeled drift in momentum is small relative to the size of debris strike that the filters are tuned to detect. Accommodating these unmodeled drifts must take into account varying solar activity, gravity gradient torques and atmospheric drag for LEO orbits, and any changing conditions encountered by satellites in elliptical orbits.

Spacecraft events which trip the detectors need to be accommodated in any algorithms applied to on-orbit telemetry. Planned events can be accommodated either by blanking results during occasional events like thruster firings, or the SPRTs could be run on a residual generated by a spacecraft model that includes all on-board events which change the measured angular momentum, such as solar array tracking. This would highlight any subtle unexpected perturbations.

However, other unexpected perturbations may be interpreted as particle impacts. For example, Bogorad showed that an electrostatic discharge (ESD) event can impart momentum to the discharging object [61]. This is especially difficult to mitigate as a hypervelocity impact can also create an ESD event, so differentiating an ESD event that causes momentum transfer from a debris strike that causes ESD is a unique challenge. Spacecraft that are prone to ESD may struggle to separate strikes from other events, but well-built spacecraft with little risk of ESD could use the plasma generated by an impact as validation that an impact has occurred. In either case, the data should be examined for correlation between probable strike events and solar activities or orbit location/charging conditions to determine if ESD is a probable cause of any detected perturbations. Similarly, a satellite may experience a dynamic event due to thermal changes as it enters or exits

eclipse which could manifest as a change in the observed inertial angular momentum. Fortunately eclipse dynamics are easy to screen for in an algorithm based on distinct changes in other telemetry (i.e., solar array power).

Another complication is the fact that spacecraft are not actually rigid bodies. In the Sentinel-1A debris strike the spacecraft rate telemetry exhibited a significant amount of ‘ringing’ which was probably due to the fundamental frequency of the solar array being excited by the strike. A similar phenomenon was observed in the MMS strikes [31, 42]. This ringing obfuscates the effects of the strike as the calculated spacecraft momentum oscillates instead of showing an abrupt change, and this feature has an effect on the SPRT algorithms’ output. However, the oscillation itself might also provide an effect that could be detected with the right filter. A spring-mass-damper appendage could be simulated for detailed assessment of the effects of appendage vibration.

Additional telemetry types should be considered for their contributions to identifying and assessing debris strikes. If a debris strike is large enough to change the orbit measurably the linear momentum may be estimated, which is the topic of the next chapter. Alternately, hypervelocity impacts produce a change in the ambient plasma surrounding a spacecraft [27], and mm-size impacts on solar arrays cause degradation in the power generation capabilities [62]. These and related features may be detectable.

Alternate data processing strategies may also prove useful for detecting impact events. For example, various research efforts continue to investigate the practicality of machine learning approaches for various space applications [63]. Machine learning techniques may be one interesting approach to solving this problem, but at present they often suffer from a lack of clear methods to verify and validate their performance. For this initial development, canonical techniques are used to provide a well-understood and well-characterized methodology for solving the problem, rather than a ‘black box’ solution which may include unknown undocumented features.

An essential element of validating the data returned by these methods is to compare the strikes predicted via models to detected strikes, and to compare results between spacecraft. This model comparison should take into account micrometeoroid populations as well as debris. This

method returns the angular momentum of the strike, so a small and fast micrometeoroid would manifest similarly to a slower and larger piece of orbital debris. The effects of the time-variant flux density should also be considered when evaluating these methods. Polar regions are expected to exhibit higher debris fluxes for LEO spacecraft and ‘rush hours’ are anticipated for GEO spacecraft [64].

2.5 Studying Effect of Trading Simulation Parameters

The previous section explored and developed algorithms for detecting angular momentum transfer using one specific spacecraft. This section trades various parameters to examine how detection thresholds vary based on an individual spacecraft’s characteristics.

2.5.1 Simplifying Strike Detection Algorithms

The algorithms developed in the previous section, the matched filter and the SPRT, are referred to in this trade study as the ‘refined strike detection algorithms.’ These refined strike detection algorithms rely on exquisite knowledge of the spacecraft’s response characteristics or on computationally intensive algorithms. While a higher telemetry rate would improve performance, state of health telemetry is typically downlinked at fairly low rates but processed at much higher rates on board. Therefore, simple algorithms are also investigated which are computationally simple enough to easily be run onboard and require minimal tuning. For the attitude and rate telemetry, the simple filter just sums the attitude or rate error over a window corresponding to the duration of the initial peak in the response. For the momentum telemetry, the average momentum over 200 datapoints is differenced with the average momentum for the preceding 200 datapoints, so a higher difference in averages indicates that a change has occurred.

2.5.2 Conducting Trades

To characterize the performance of the algorithms under various combinations of parameters a process is developed which outputs the ΔH detection threshold and estimation accuracy for a

given scenario. First, a simulation is run in which a series of strikes with increasing magnitude is applied to the simulation (Figure 2.20(a)). The filters are run and each filter's peak output is recorded along with the applied strike size. A positive 2nd-order polynomial is fitted to this data to develop a relationship between the strike size and filter output (Figure 2.20(b)). Peaks below the noise floor are ignored when fitting. Then a long simulation is run with no strikes. The filters are run on this data producing a population of filter output data when no strikes are present. A Kernel distribution is fitted to this data since standard distributions often fit poorly, especially on the tails. A false alarm rate of one per day is selected as the strike detection threshold - for a given telemetry data rate, a probability of $1/(\text{number of telemetry points in a day})$ is used as the desired probability of false alarm (P_{fa}). The threshold for strike detection is calculated based on the fitted probability density functions given the desired P_{fa} (Figure 2.20(c)).

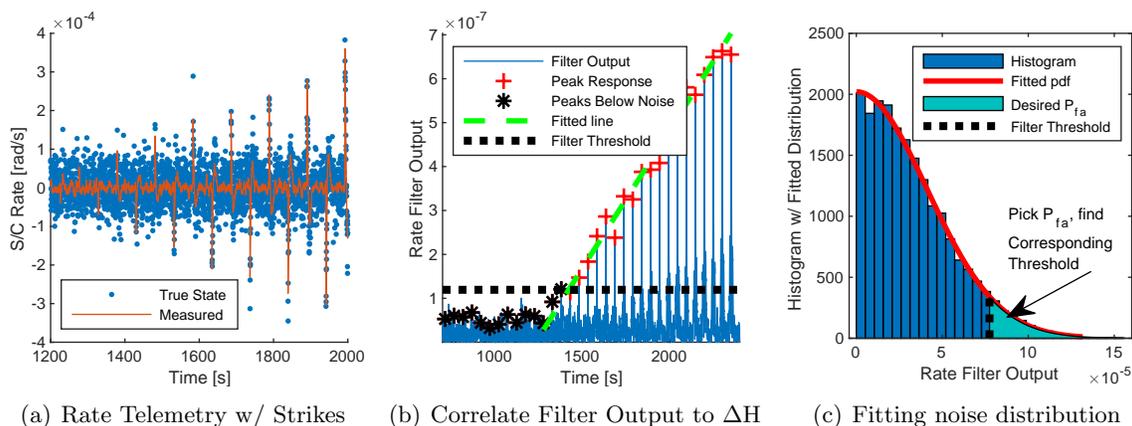


Figure 2.20: Process for determining ΔH threshold. Note that this maps the filter threshold, which changes based on filter parameters, to a ΔH threshold which can be compared between filters and spacecraft.

This filter output threshold is mapped to a threshold for detecting a change in angular momentum, ΔH . Depending on the spacecraft, this ΔH threshold can be either the three components of the vector $\Delta \mathbf{H}$ or the magnitude and a scalarized direction, ΔH_{mag} and ΔH_{dir} , referred to generically as just ΔH . The ΔH threshold mapping is accomplished using the relationship between strike size and filter output. Note that minor changes in the probability density functions or curve fitting

can lead to changes in the thresholds, leading to somewhat noisy measurements of the threshold. Therefore, when evaluating the results the major trends are the desired outcome, exact thresholds vary slightly across runs producing some noise in the results. To quantify the estimation accuracy a series of strikes is applied to the simulation that are 1.5x the ΔH threshold, then the estimation accuracy for these strikes is calculated as described in the section on Strike Size Estimation Algorithms.

Note that some parameters are fixed in order to reduce the scope of the problem to clearly show the influence of traded parameters. The strikes are parameterized in terms of mass only, assuming a fixed relative velocity and for each spacecraft size a fixed lever arm from the center of mass (which is based on the overall size of the spacecraft). Also, the strikes are always applied about the x-axis. This allows the spacecraft response to be characterized in one dimension as a function of the traded parameters only, allowing clear trends to emerge.

2.5.2.1 Spacecraft Size

Three design point spacecraft are developed to study the effects of spacecraft size on strike identification and assessment. The large spacecraft is NASA's Solar Dynamics Observatory, SDO, at approximately 2,000 kg. The medium spacecraft is Firesat, a notional spacecraft design of 215 kg. The third is a 12U cubesat at 24 kg. The detailed parameters used for each of the spacecraft are given in the appendix. The noise parameters are consistent across the three spacecraft and are loosely derived from noise observed in SDO's telemetry. Note especially that strikes are applied in terms of particle mass, and the distance between the strike and the center of mass (CoM) of the spacecraft is different for the different spacecraft sizes. Thus, a particle of a certain size striking SDO will impart substantially more angular momentum than the same particle striking the 12U cubesat since it is likely to hit significantly further from the center of gravity.

The results in Figure 2.21 show that smaller spacecraft can detect substantially smaller ΔH s. This is expected, since a smaller spacecraft has a lower moment of inertia so the same ΔH will impart more spin and be more easily detected. The detection rates in terms of particle size

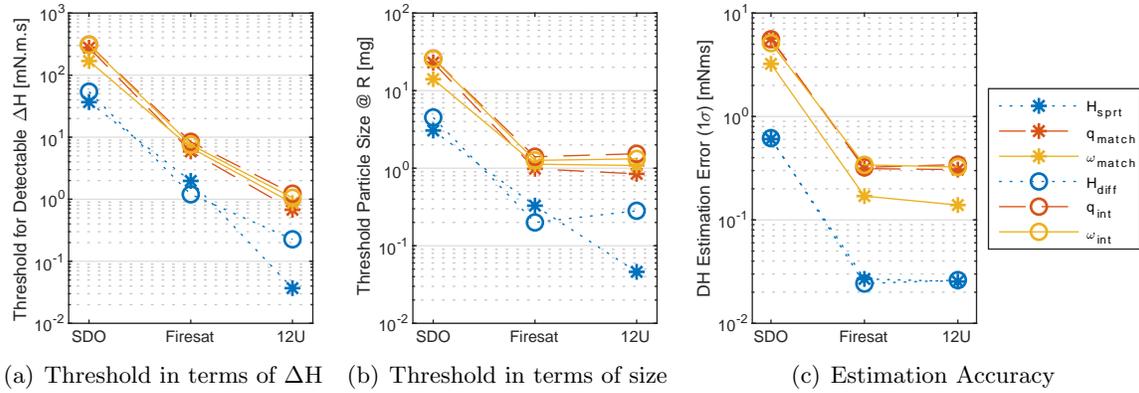


Figure 2.21: Trends in detection threshold for various S/C sizes

are also plotted, which take into account the shorter lever arm of the smaller spacecraft. The smaller two spacecraft still have lower detection thresholds than SDO, though the effects are less significant. The estimation accuracy is shown in Figure 2.21(c). Unsurprisingly, the spacecraft that can detect smaller strikes can also characterize the size of those strikes more accurately. Note that the discrepancies between simple and refined algorithms are minimal while the algorithms using momentum outperform the algorithms using attitude and rate by nearly an order of magnitude in all cases. This is especially noteworthy because calculating the inertial momentum incorporates the errors from both attitude and rate as well as wheel speed.

2.5.2.2 Telemetry Rate Trade

The simulation always runs at 100 Hz, but the telemetry is reduced to a lower rate representing downlinked telemetry. The default telemetry rate is 10 Hz. Since the dynamic spacecraft response to strikes lasts for about 25 seconds this provides a 250 point wavelet for the matched filters. The telemetry rate is traded but the length of the matched filters is left at 25 seconds, so the matched filters have more or fewer points at higher or lower telemetry rates. The momentum algorithms, however, work with a fixed length of 200 datapoints. The results are shown in Figure 2.22. As expected, the attitude and rate telemetry show notable degradation at lower telemetry rates since they have less data to work with. However, the momentum algorithms show little vari-

ation with telemetry rate. The attitude and rate filters approach the accuracy of the momentum filters at high data rates, but note that even with substantially more data (2500 points at 100 Hz vs 200 for the momentum algorithms) the momentum algorithms still perform better.

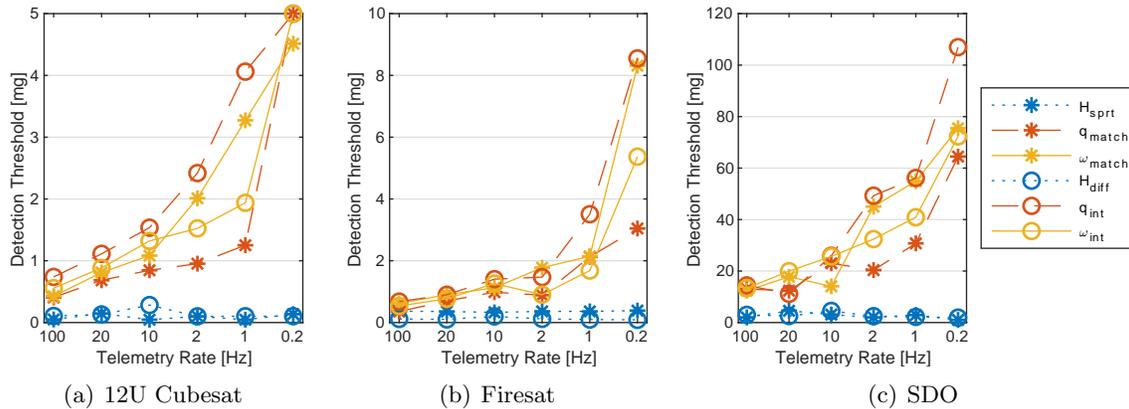


Figure 2.22: Trends in detection threshold for various telemetry rates

2.5.2.3 Telemetry Noise Trade

First, the noise levels in all telemetry points are scaled to gain an understanding of the global effect of telemetry noise on algorithms. The noise is scaled up and down by a percentage of default noise applied on each telemetry point. The resulting thresholds show the expected trends with less noise resulting in lower thresholds and more noise resulting in higher thresholds (Figure 2.23). The results are somewhat muddy due to noise in the detection threshold measurement methods.

Next, the noise in the individual telemetry points is traded. While the momentum processing algorithms outperform the state error algorithms with the default noise levels, there are substantial differences in the accuracy of the rate measurement of various inertial measurement units. Therefore, trades are conducted where the noise levels in individual telemetry types are scaled to determine if the state algorithms can perform better than the momentum algorithms given rate and attitude telemetry with lower noise, or wheel speed telemetry with higher noise. Figure 2.24 shows the results. The momentum processing algorithms still perform better with 10x the wheel speed noise ($1\sigma = 24$ RPM), after that performance degrades abruptly. The state algorithms begin

to perform better with about order of magnitude reduction in noise ($1\sigma = \sim 0.825$ arcsec/sec for rate, $1\sigma = \sim 2.9$ arcsec for attitude).

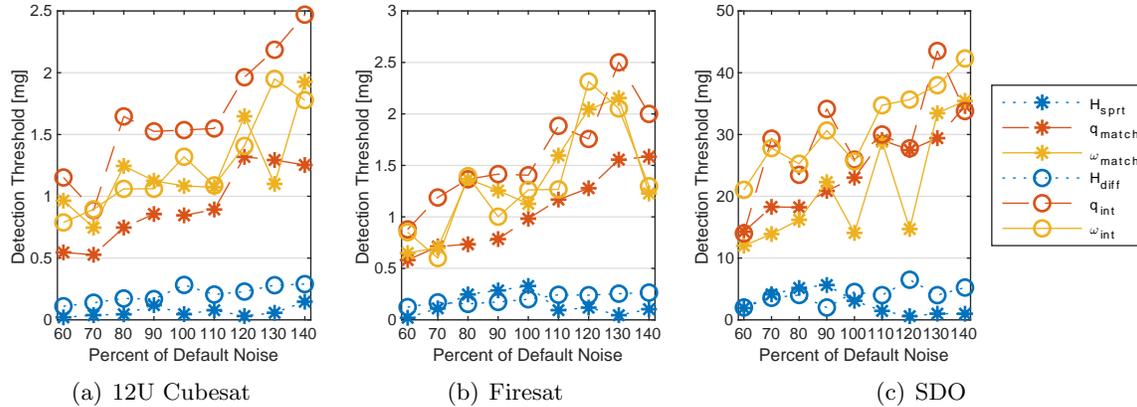


Figure 2.23: Trends in detection threshold when trading telemetry noise

2.5.2.4 Other Trades

Three additional trades are conducted to see if they affect the results: the initial reaction wheel speed, the reaction wheel inertia, and the spacecraft inertia (Figure 2.25). The reaction wheel starting speed and inertia have no significant effect on the threshold, variations seem to be due to inherent noise in the process of setting the thresholds. The similarities between Figure 2.25(a) and Figure 2.25(b) are due to the random number generator in both runs having the same default seed, so the algorithms are working on the same random data. The spacecraft inertia changes the shape of the response wavelets which produces some minor changes in the thresholds, but these are relatively small for reasonable changes in inertia and are mostly overwhelmed with noise in Figure 2.25(c). All trades are made with SDO's default parameters.

2.5.3 Discussion

2.5.3.1 Overall Trends

The results show that small spacecraft can detect substantially smaller strikes than large spacecraft, that a higher telemetry rate and lower noise can improve the strike detection threshold,

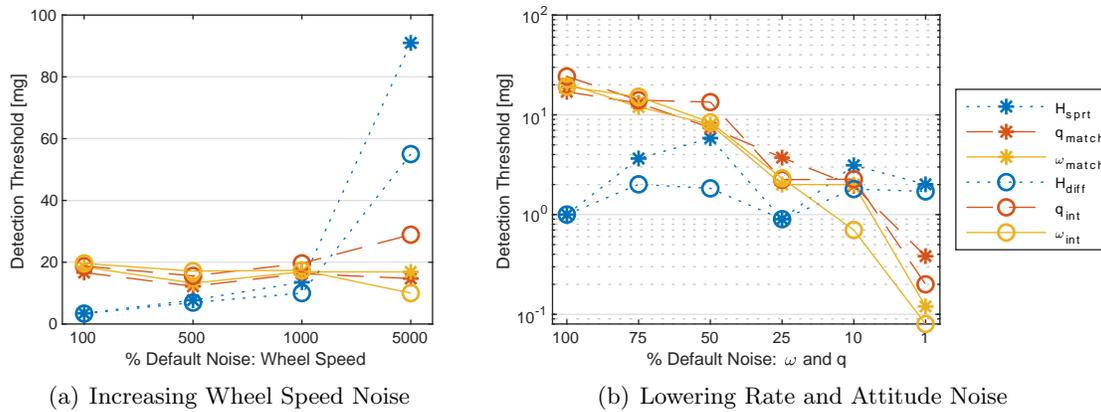


Figure 2.24: Trade noise in individual telemetry points (all data from SDO)

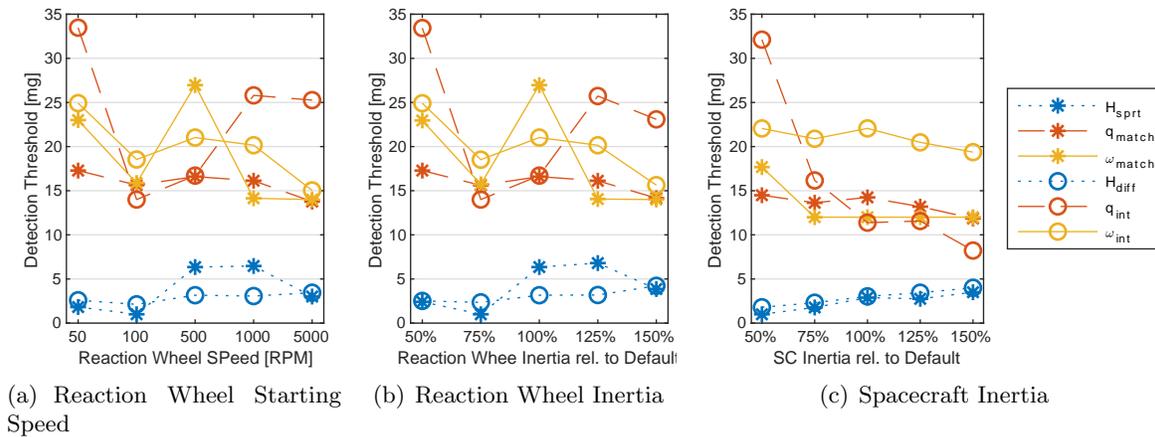


Figure 2.25: Trading miscellaneous spacecraft parameters

and that the starting reaction wheel speed and reaction wheel inertia have minimal effect on the thresholds.

The results also showed that, while the momentum processing algorithms outperform the state error processing algorithms at the selected default noise levels, with substantially noisier measurements of reaction wheel speed the momentum algorithm performance is degraded. Similarly, reducing the noise in the attitude and rate measurements improves their performance significantly.

Surprisingly, the simple algorithms performed as well or sometimes better than the refined algorithms. The attitude matched filter was consistently the best among the state error processing algorithms, but not overwhelmingly so, and the SPRT performed comparably to differencing the

averages in the momentum telemetry. This result should be caveated by mentioning that the non-Gaussian noise seen in real telemetry complicates setting an appropriate threshold, and the SPRT, with its ability to be tuned to ignore small effects, returns a substantially better signal-to-noise ratio on real data which is a non-trivial advantage. For this reason the SPRT is not discarded, even though performance is comparable on simulated telemetry while processing time is significantly slower.

2.5.3.2 Considerations for Real-World Application

It is important to remember that the noise in these simulations is highly idealized. Perfect zero-mean white Gaussian noise is applied during the simulation to perturb the true state, and then perfect zero-mean white Gaussian measurement noise is added to the true state to produce the telemetry. In reality noise is much more complicated and can be colored, biased, non-Gaussian, and vary based on component temperature, wheel speed, and miscellaneous other effects which may or may not be fully characterized. Because of this, **these thresholds should be considered a best case scenario**; a real spacecraft's noise would be much harder to characterize accurately and consistently across all mission scenarios and therefore the detection thresholds would probably need to be raised to accommodate these variations.

In addition to difficulties in characterizing the baseline noise, experience has shown substantial unexpected events when processing on-orbit telemetry which obfuscate the debris strike detection results. For example, the momentum calculation assumes rigid body, so appendage vibration causes an apparent oscillation in the inertial angular momentum which trips the momentum detection threshold. Section 4.1 discusses the various effects seen.

Similarly, the momentum processing algorithms assume that the inertial angular momentum is unchanging relative to the minimum detectable threshold. In the simulated telemetry this is true, but in reality a spacecraft's momentum is constantly changing due to solar radiation pressure, etc. These changes must be negligible relative to the detection threshold over the length of the window for the momentum processing algorithms to work effectively, which caps the length of window that

can be used in practice.

Translating from a ΔH measurement to specific debris parameters is non-trivial, since ΔH is the cross product of the strike location relative to the center of mass crossed with the mass times the relative velocity. This results in one equation with three unknowns. However, a three-axis measurement of the ΔH vector can define a plane of the impact, which helps to define possible strike location, and orbit debris is most likely to hit from the ram direction while micrometeoroids also show preferred directions, although they vary annually. So one ΔH measurement cannot be used to solve for strike parameters, but combining a population of ΔH measurements with known fluxes and spacecraft geometry it might be possible to fit a mass distribution to the debris flux that maximizes the likelihood of obtaining the observed data.

The trades performed in this study illustrate the effect that various spacecraft and noise parameters have on the effectiveness of these algorithms. While this study makes some idealized assumptions the general trends inform understanding the strike detection capabilities of various spacecraft with various ACS suites.

Chapter 3

Methods to Detect Strikes via Translational Momentum Transfer

3.1 Concept

The translational momentum imparted by a strike produces a small ΔV , which produces a change in orbit. Orbits are perturbed continuously, but most perturbations occur over longer timescales than the perturbation caused by a debris strike, which will typically occur in less than a second[65]. This distinguishing feature, that a debris strike is more abrupt than most perturbations, enables the development of techniques designed to look for this abrupt change in the truth state and identify strikes in the presence of typical perturbations and sensor noise. This chapter describes the development some methods which employ an extended Kalman filter with various adaptations to accentuate the abrupt velocity change expected in a debris strike using simulated spacecraft navigation telemetry.

Many authors have applied a broad smorgasbord of techniques to a similar problem: detecting unknown maneuvers using orbit state determination. Highlights include the Orbit Determination Toolkit by Analytical Graphics, Inc., which employs a backward-running fixed-interval smoother to estimate velocity changes due to an impulsive maneuver[66], as well as various flavors of fixed-epoch smoother. Ko and Scheeres investigate the use of thrust fourier coefficients to detect unknown maneuvers[67], while Holzinger et. al. detect and characterize maneuvers using control-distance metrics[68]. Kelecey and Jah develop techniques to detect and reconstruct maneuvers of low-thrust LEO satellites[69], and Jiang et. al. employ a Mahalanobis distance-based technique to detect unknown maneuvers[70].

Similar concepts that have been implemented in the literature are referenced in the introduction: The NESC report studied changes in satellite mean altitude during the seven momentum perturbation events, the Sentinel-1A strike sparked an effort by multiple entities to estimate the imparted ΔV with a high degree of consistency between the estimates, and ExoAnalytic Solutions studied momentum impulse transfer events on their GEO satellites detecting ΔV s as small as 0.7 mm/s in the in-track direction. The primary contribution of this research is to develop techniques that are specifically designed to look for debris strikes, instead of observing debris strike features in the course of other operations. The Exoanalytic approach is closest to this, as they observed poorly fit residuals from their sensor network and used their filters to estimate a ΔV as well as solar radiation pressure coefficients. This chapter extends that concept with dedicated algorithms.

This chapter investigates a technique that is similar to but unique from each of these, wherein an extended Kalman filter (EKF) with dynamic model compensation (DMC) is used to estimate unmodeled accelerations as a first-order Gauss-Markov (FOGM) process and then augmented with various test statistics. The filter is applied forward and backward across the data then smoothed, and various test statistics are developed and compared to identify the highest signal-to-noise ratios (SNR) in the presence of a small change in velocity. The backward filter employs a mirrored FOGM model for the unmodeled accelerations and produces a state estimate using only ‘future’ data, instead of a smoothed state estimate using both past and future data as with many smoothers. This modification enhances debris strike detectability by accentuating the difference between forward and backward states at the point of the debris strike, and presents a key innovation relative to the literature. The concept is a modified version of the smoother developed by Fraser and Potter[71], which also fused a state estimate from a past-data-only estimate and a future-data-only estimate. Wall et al [72] discuss the concept of using a mirrored model in a backward filter, but the implementation differs from what is developed here.

In this chapter a filter is developed to produce a state estimate appropriate for detecting orbit perturbations due to debris strikes. Simulated onboard navigation telemetry is employed to develop and test the filter; very basic orbit propagation and measurement noise models are used since the

focus is on methods development and characterization. A ‘truth’ orbit is generated with a simulated debris strike applied as an impulsive ΔV , then this truth state is used to produce simulated GPS telemetry (position and velocity). An EKF with DMC is developed and applied to this telemetry, where the DMC models the acceleration state as a first-order Gauss-Markov (FOGM) process. For debris strike detection the filter is modified such that it can be applied backward in time, and then a smoother is applied to fuse the forward and backward state into a smoothed state estimate. Several test statistics are developed to post-process the filter output and accentuate the effects of the debris strike, including a Mahalanobis distance and a McReynold’s filter-smoother consistency test. The filter performance is assessed using a trade study to characterize the filter’s ability to detect strikes while trading the FOGM time constant and the power spectral density of the state process noise. A Monte-Carlo analysis is performed to characterize the performance when detecting and estimating the size of a randomized debris strike. Finally, a couple orbit perturbations (drag and J3) are added to the truth state but not the filter dynamics, to illustrate the filter’s performance in the presence of additional unmodeled accelerations. Discussion includes expected challenges when the filter is applied to on-orbit telemetry, and the potential utility of these techniques for understanding the hazardous non-trackable debris environment.

3.2 Filter Development

3.2.1 Generating Simulated Navigation Telemetry

Simple two-body orbit propagation is used throughout this chapter, except as noted in Section 3.4.4. While modeling higher-order orbital perturbations will be critical in applying these techniques to on-orbit telemetry tools already exist for precision orbit determination and these tools can be leveraged in real-world applications. The focus here is on methods development, so a two-body orbit propagator is used.

A debris strike is applied as an instantaneous change in velocity ($\Delta \mathbf{V}$). Since the momentum transfer effects of a hypervelocity impact typically resolve in less than a second[65] a strike will

essentially occur between two telemetry points. The momentum imparted by the debris strike is calculated as

$$\Delta\mathbf{p} = \beta m_d \mathbf{v}_d \quad (3.1)$$

where $\Delta\mathbf{p}$ is the momentum transferred to the spacecraft, β is the momentum enhancement factor [6], m_d is the mass of the debris, and \mathbf{v}_d is the velocity of the debris relative to the spacecraft. This is applied in the inertial frame to produce the $\Delta\mathbf{V}$, but \mathbf{v}_d can be specified in either the inertial or orbit frame and then converted if needed. By conservation of momentum, the change in spacecraft velocity is then

$$\Delta\mathbf{V} = \frac{\Delta\mathbf{p}}{m_s} \quad (3.2)$$

where m_s is the mass of the spacecraft. At the specified strike time this change in velocity is applied to the spacecraft state and the new post-strike orbit state is propagated for the rest of the specified simulation time. The default values for the debris strike parameters are provided in Section 3.2.2, and any changes from these values are specified. Note that when a change is specified in the text as a scaled proportion to the ‘default’ strike, the mass is scaled by that amount. The default debris mass is 0.2 g, so a ‘5X strike’ specified in the text means that the mass of the debris is set to 1 g instead.

From the spacecraft orbit state full-state measurements are derived, intended as a basic simulation of GPS telemetry. These position and velocity measurements consist of the truth state with Gaussian noise superimposed in accordance with the noise parameters specified in Section 3.2.2. As with orbit propagation this model is far simpler than real-world noise characteristics, but the focus in this chapter is on methods development so a simplified model is appropriate. Comments on applying these methods to real-world telemetry are included in the discussion, Section 3.5.

3.2.2 Default Values for Parameters

Debris strike parameters are inherently probabilistic and vary by orbit, so specifying a value often sparks debate which distracts from more relevant debates about methods. Therefore, an entire

table of values is provided here in Figure 3.1 to specify debris strike and spacecraft parameters which could produce the baseline ΔV and the scaled ΔV s used in various simulations throughout this chapter. The table is shaded in accordance with the likely hazard presented by the piece of debris and its detectability. As shown, the hazardous non-trackable debris pieces (orange) represent a large percentage of the strikes which are applied in this chapter. Some systems could push the detectability threshold low enough that even non-damaging strikes could be detected (blue), and some systems (large spacecraft and/or spacecraft with substantial unmodeled perturbations) could have detection thresholds corresponding to debris which is potentially trackable. However, steel debris, typically modeled as spheres, currently drives the risk to LEO spacecraft, and this table indicates that detection thresholds are likely in the range where this type of debris could be detected to acquire more information about the hazardous non-trackable debris population.

		Densities, in g/cm ³ :					2.7	8.05	0.04269	1.8	8.96	
Regime	Scale	Mass S/C (kg)	Mass db (g)	velocity (km/s)	MEF	DV (cm/s)	Aluminum sphere diameter (cm)	Steel sphere diameter (cm)	Square of MLI (cm, on edge)	1 mm thick square CFRP (cm)	Length of 2 mm dia. copper wire (cm)	
LEO	Default	200	0.2	12	2	2.4	0.52	0.36	2.16	1.05	0.71	
LEO		100	0.1	12	2	2.4	0.41	0.29	1.53	0.75	0.36	
LEO		10	0.01	12	2	2.4	0.19	0.13	0.48	0.24	0.04	
LEO		1000	1	12	2	2.4	0.89	0.62	4.84	2.36	3.55	
LEO	2X	200	0.4	12	2	4.8	0.66	0.46	3.06	1.49	1.42	
LEO		100	0.2	12	2	4.8	0.52	0.36	2.16	1.05	0.71	
LEO	5X	200	1	12	2	12	0.89	0.62	4.84	2.36	3.55	
LEO		100	0.5	12	2	12	0.71	0.49	3.42	1.67	1.78	
LEO	1/10th	200	0.02	12	2	0.24	0.24	0.17	0.68	0.33	0.07	
LEO		1000	0.1	12	2	0.24	0.41	0.29	1.53	0.75	0.36	
GEO	Default	200	6.857143	0.7	1	2.4	1.69	1.18	12.67	6.17	24.36	
GEO		2000	68.57143	0.7	1	2.4	3.65	2.53	40.08	19.52	243.60	
Key:							Likely too small to cause damage	Hazardous nontrackable	Potentially trackable			

Figure 3.1: Table illustrating myriad situations which could produce a detectable ΔV for a hazardous non-trackable piece of debris. Scale refers to debris strike magnitude as specified throughout chapter. Debris listed as ‘potentially trackable’ is NOT tracked currently, but may be tracked as current state-of-the-art systems mature.

These parameters represent fairly small debris strikes, and many systems may not be able to detect strikes on this scale. However, spacecraft have detected strikes with no specialized filtering, so introducing a filter which can accentuate the subtle effects of strikes can only improve their performance as detectors and result in a larger body of data for assessing the debris environment. The parameter values used throughout this chapter are provided in Table 3.1, any deviations from

these values (i.e, for trade studies, etc) are specified in the text.

Table 3.1: Table of default parameters used in simulations throughout unless otherwise specified

Parameter	Value	Unit
Spacecraft Parameters		
S/C Mass	200	kg
Orbit Parameters		
Orbit altitude	800	km
Orbit eccentricity	0	[-]
Orbit inclination	40	deg
Orbit RAAN	30	deg
Orbit arg. of perigee	50	deg
Measurement Parameters		
Measurement Noise in Position $\sigma_{\text{meas, pos}}$	10	m (1σ)
Measurement Noise in Velocity $\sigma_{\text{meas, vel}}$	0.25	m/s (1σ)
Data rate	1	Hz
Debris Parameters		
Debris mass	0.2	g
Debris velocity	12	km/s
Momentum Enhancement Factor[57]	2	[-]
Strike direction	[.5774; .5774; .5774]	In inertial frame
Filtering Parameters		
Time constant τ (Initial, before tuning)	10	s (for FOGM)
Process noise spectral density q (Initial)	2e-18	km ² /s ⁵
Time constant τ (Final, after tuning)	1	s
Process noise spectral density q (Final)	2.15e-17	km ² /s ⁵
Measurement noise matrix \mathbf{R}	$[\sigma_{\text{meas, pos}}^2 \mathbf{I}_{3 \times 3}, \mathbf{0}_{3 \times 3}; \mathbf{0}_{3 \times 3}, \sigma_{\text{meas, vel}}^2 \mathbf{I}_{3 \times 3}]$	measurement units ²
Orbit Perturbation Parameters		
Coefficient of Drag C_D	2.2	[-]
Reference altitude r_o	per Vallado[73]	Table 8-4
Atmospheric scale height H	per Vallado[73]	Table 8-4
Reference density ρ_o	per Vallado[73]	Table 8-4
Area of satellite bus	1	m ²
Area of solar array (tracks sun, area rel to $\hat{\mathbf{v}}$ varies)	5	m ²

3.2.3 Implementing EKF with DMC

The extended Kalman filter is implemented per the algorithm in Tapley, Schutz and Born (TSB) [74], pg. 212. The filter uses dynamic model compensation, estimating unmodeled accelerations as part of the filter state. Thus the state vector consists of position \mathbf{r} , velocity \mathbf{v} , and unmodeled accelerations \mathbf{w} . With position and velocity measurements the measurement partials,

denoted $\tilde{\mathbf{H}}$, are straightforward:

$$\tilde{\mathbf{H}}_i = [\mathbf{I}_{6 \times 6} \quad \mathbf{0}_{6 \times 3}] \quad (3.3)$$

where $\tilde{\mathbf{H}}$ is defined and used per the TSB formulation, $\mathbf{I}_{6 \times 6}$ is a 6x6 identity matrix, and $\mathbf{0}_{6 \times 3}$ is a matrix of zeros. The covariance of the measurement noise (the \mathbf{R} matrix per TSB notation) is provided in Section 3.2.2.

The truth orbit is propagated with Matlab's ODE45 function, but within the filter the state is propagated with a fourth-order Runge-Kutta (RK4) and the state transition matrix (STM) is propagated with a 2nd order Runge-Kutta (RK2), in accordance with typical practices. Propagating the state and STM both with ODE45 also works as well but is significantly slower, while propagating the state with an RK2 was tried and also works well but is less typical in operational filters. The state propagation RK4 calculates the time derivative of the state as follows

$$\dot{\mathbf{X}} = \begin{bmatrix} \mathbf{v} \\ \mathbf{a} \\ \dot{\mathbf{w}} \end{bmatrix} \quad (3.4)$$

where the accelerations are the accelerations from orbital dynamics plus the unmodeled accelerations estimated with DMC.

$$\mathbf{a} = -\frac{\mu}{r^3} \mathbf{r} + \mathbf{w} \quad (3.5)$$

The dynamic model compensation estimates the unmodeled accelerations using a first order Gauss-Markov (FOGM) process[75]. The FOGM model is

$$\dot{w}(t) = -\frac{1}{\tau} w(t) + \eta(t) \quad (3.6)$$

with $\eta(t) \sim N(0, \sqrt{q\delta(t-s)})$. Note that the correlation time, τ , is the time at which the magnitude of the acceleration will fade to $1/e$ of its prior value, in the absence of additional excitation. Since the expected value of η is zero the state propagation uses

$$\dot{\mathbf{w}} = -\frac{1}{\tau} \mathbf{w} \quad (3.7)$$

The state transition matrix is also propagated using an RK2 integrator, specifically Lear's formulation[76] for an RK2 propagation of the covariance with two evaluations of \mathbf{A} .

$$\Phi(t_{i+1}, t_i) = \mathbf{I}_{9 \times 9} + \frac{\mathbf{A}_i + \mathbf{A}_{i+1}}{2} \Delta t + \mathbf{A}_i \mathbf{A}_{i+1} \frac{\Delta t^2}{2} \quad (3.8)$$

where

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ -\frac{\mu}{r^3} + \frac{3\mu X^2}{r^5} & \frac{3\mu XY}{r^5} & \frac{3\mu XZ}{r^5} & 0 & 0 & 0 & 1 & 0 & 0 \\ \frac{3\mu XY}{r^5} & -\frac{\mu}{r^3} + \frac{3\mu Y^2}{r^5} & \frac{3\mu YZ}{r^5} & 0 & 0 & 0 & 0 & 1 & 0 \\ \frac{3\mu XZ}{r^5} & \frac{3\mu XY}{r^5} & -\frac{\mu}{r^3} + \frac{3\mu Z^2}{r^5} & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{\tau} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{\tau} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{\tau} \end{bmatrix} \quad (3.9)$$

For dynamic model compensation the process noise covariance matrix \mathbf{S} is added during the time update of the covariance matrix, with $\bar{\mathbf{P}}$ as the *a priori* covariance and \mathbf{P} as the *a posteriori*.

$$\bar{\mathbf{P}}_i = \Phi(t_i, t_{i-1}) \mathbf{P}_{i-1} \Phi^T(t_i, t_{i-1}) + \mathbf{S}_i \quad (3.10)$$

Where \mathbf{S}_i is calculated from the process noise covariance integral. Note that this process noise covariance integral represents the process noise spectral density \mathbf{Q} transformed into the state space via \mathbf{B} and then integrated from the previous measurement epoch t_{i-1} to the current epoch t_i using the state transition matrix Φ .

$$\mathbf{S}_i = \int_{t_{i-1}}^{t_i} \Phi(t_i, \epsilon) \mathbf{B}(\epsilon) \mathbf{Q}(\epsilon) \mathbf{B}^T(\epsilon) \Phi^T(t_i, \epsilon) d\epsilon \quad (3.11)$$

In this chapter, \mathbf{S}_i is calculated per the formulation in NASA's Navigation Filter Best Practices[75]. The \mathbf{Q} is the power spectral density of the acceleration noise, and is applied in the orbit frame (radial, tangential, normal). The \mathbf{B} matrix maps the orbit frame to inertial frame (so that

it can be added to the state vector). Thus,

$$\mathbf{Q}_{\text{rtn}} = \begin{bmatrix} q_r & 0 & 0 \\ 0 & q_t & 0 \\ 0 & 0 & q_n \end{bmatrix} \quad (3.12)$$

$$\mathbf{B} = \begin{bmatrix} \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} \\ \mathbf{M}_{\text{rtn}} \end{bmatrix} \quad (3.13)$$

where \mathbf{M}_{rtn} is the direction cosine matrix to map the orbit frame to the inertial frame. For convenience, the following definition of $\tilde{\mathbf{Q}}$ is established.

$$\tilde{\mathbf{Q}} = \mathbf{M}_{\text{rtn}} \mathbf{Q}_{\text{rtn}} \mathbf{M}_{\text{rtn}}^T \quad (3.14)$$

With this $\tilde{\mathbf{Q}}$ the process noise covariance \mathbf{S}_i is calculated as

$$\mathbf{S}_i = \begin{bmatrix} \gamma_{rr} \tilde{\mathbf{Q}} & \gamma_{rv} \tilde{\mathbf{Q}} & \gamma_{rw} \tilde{\mathbf{Q}} \\ \gamma_{rv} \tilde{\mathbf{Q}} & \gamma_{vv} \tilde{\mathbf{Q}} & \gamma_{vw} \tilde{\mathbf{Q}} \\ \gamma_{rw} \tilde{\mathbf{Q}} & \gamma_{vw} \tilde{\mathbf{Q}} & \gamma_{ww} \tilde{\mathbf{Q}} \end{bmatrix} \quad (3.15)$$

with the γ coefficients defined as follows.

$$\gamma_{rr,\text{fwd}} = \frac{\tau^5}{2} \left\{ \left(1 - e^{-2\Delta t/\tau}\right) + \frac{2\Delta t}{\tau} \left(1 - 2e^{-\Delta t/\tau}\right) - 2 \left(\frac{\Delta t}{\tau}\right)^2 + \frac{2}{3} \left(\frac{\Delta t}{\tau}\right)^3 \right\} \quad (3.16)$$

$$\gamma_{vv,\text{fwd}} = \frac{\tau^3}{2} \left\{ \left(1 - e^{-2\Delta t/\tau}\right) - 4 \left(1 - e^{-\Delta t/\tau}\right) + 2\Delta t/\tau \right\} \quad (3.17)$$

$$\gamma_{ww,\text{fwd}} = \frac{\tau}{2} \left(1 - e^{-2\Delta t/\tau}\right) \quad (3.18)$$

$$\gamma_{rv,\text{fwd}} = \frac{\tau^4}{2} \left\{ \left(e^{-2\Delta t/\tau} - 1\right) - 2 \left(e^{-\Delta t/\tau} - 1\right) + \frac{2\Delta t}{\tau} \left(e^{-\Delta t/\tau} - 1\right) + \left(\frac{\Delta t}{\tau}\right)^2 \right\} \quad (3.19)$$

$$\gamma_{rw,\text{fwd}} = \frac{\tau^3}{2} \left\{ \left(1 - e^{-2\Delta t/\tau}\right) - \frac{2\Delta t}{\tau} e^{-\Delta t/\tau} \right\} \quad (3.20)$$

$$\gamma_{vw,\text{fwd}} = \frac{\tau^2}{2} \left(1 - e^{-\Delta t/\tau}\right)^2 \quad (3.21)$$

3.2.4 Testing Filter on Telemetry with Simulated Debris Strike

A debris strike is applied to the spacecraft and simulated measurements are generated per the parameters specified in Section 3.2.2. Figure 3.2 shows this simulated debris strike applied at

$t = 7500$ seconds, and the results show that an abrupt change in velocity causes the filter state to diverge from the truth state for about 50 minutes. Of course, producing this graph requires knowledge of the truth state to calculate the error in the filter's state estimate, which is not available in a real-world application. The unmodeled accelerations, which would be available in a real-world application, show a tiny feature barely distinguishable from the baseline noise.

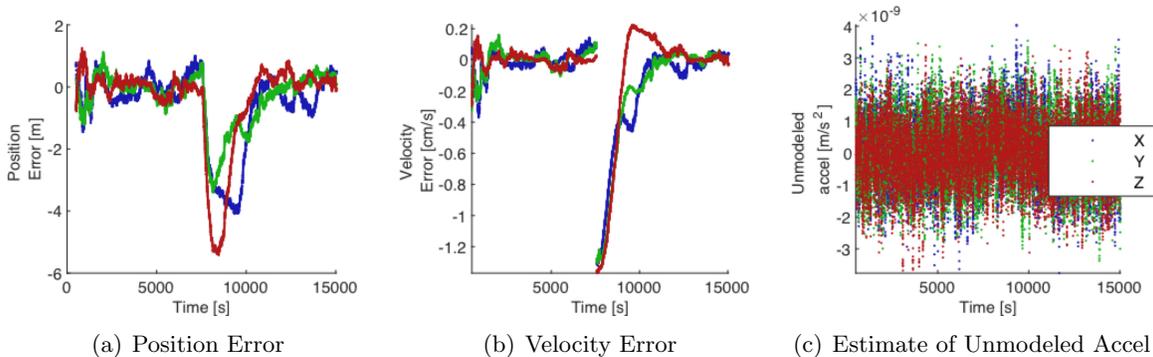


Figure 3.2: Performance of forward filter on telemetry with small debris strike

Figure 3.3 shows the measurement residuals, which are typically used for assessing filter performance, indicating that this debris strike is too small to spot using the typical filter assessment method of looking for patterns in residuals.

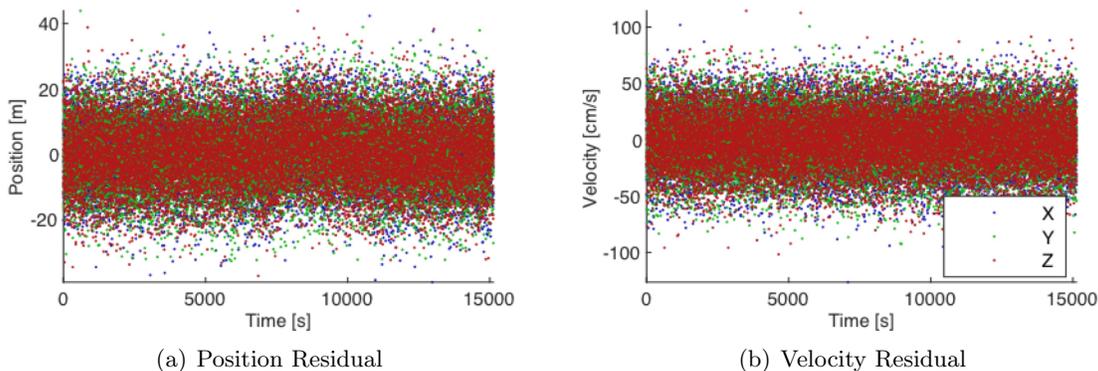


Figure 3.3: Measurement residuals from filter during debris strike show a minute feature which is not apparent via casual observation

3.2.5 Response of EKF with DMC to debris strike

3.2.5.1 A closer look at DMC and FOGM.

To understand the implications of this dynamic model when applied to a debris strike some further discussion of DMC and the FOGM model is warranted. The unmodeled accelerations are initially set to zero. At each filter measurement update step, the Kalman gain maps the measurement innovation (position and velocity) into the state space (position, velocity, and unmodeled acceleration) to update the estimated state. The magnitude of this update is proportionate to the size of the measurement innovation, and is scaled by the Kalman gain. As shown in Equation 3.22, the strength of the Kalman gain is effectively proportionate to the ratio of the state uncertainty (covariance) over the state uncertainty plus the measurement uncertainty (measurement noise covariance, \mathbf{R}). Therefore, a large state uncertainty and small measurement uncertainty results in measurement innovations being mapped into a large state update. *Vice versa*, a small state uncertainty and large measurement uncertainty results in a smaller Kalman gain and a smaller state update. Essentially, via the Kalman gain the filter is estimating that some unmodeled acceleration has acted on the state based on the measurement innovation it sees, and updating its estimate of the unmodeled acceleration accordingly.

$$\mathbf{K}_i = \mathbf{P}_i^- \mathbf{H}_i^T / (\mathbf{H}_i \mathbf{P}_i^- \mathbf{H}_i^T + \mathbf{R}_i) \quad (3.22)$$

With the unmodeled acceleration modeled as a FOGM process, Equation 3.6 shows that the time derivative of the unmodeled acceleration, $\dot{w}(t)$, consists of two components: the first term produces a function which decays exponentially in accordance with the time constant τ . The second term, $\eta(t)$ is the random noise term. The first term indicates that the unmodeled acceleration at the last timestep will also appear in the current timestep, but will have decayed slightly in accordance with the time constant τ . The measurement update calculated using the Kalman gain is added to this current estimate of the unmodeled acceleration. This allows the unmodeled acceleration to grow if measurement innovations persistently indicate that a larger unmodeled acceleration has acted on the process, and then decay away toward zero when the state has been corrected sufficiently

that measurement innovations are small. Simultaneously, if the measurement innovations indicate that the state is ‘over-corrected’ the measurement update will proactively reduce the size of the unmodeled accelerations by subtracting a state correction.

The random noise term, $\eta(t)$, comes into play via the covariance matrix. In the propagation step the process noise covariance \mathbf{S}_i is added to propagated covariance from the previous step (Equation 3.10). This accounts for the integrated effect of process noise acting on the state from one timestep to the next, with the magnitude specified by q , the power spectral density of the acceleration noise. Thus, q provides a ‘knob’ to tune the strength of the Kalman gain such that the filter behaves as desired by inflating the covariance. In this case the desire is for the filter to detect when a subtle debris strike has influenced the filter state. Note that this deviates from typical desired filter performance, which is producing the best possible state estimate based on the most recent measurements available.

3.2.5.2 Effect of debris strike on filter.

In the presence of an impulsive $\Delta\mathbf{V}$ the truth velocity changes abruptly from the estimated state, and as the state propagates forward the position also begins to deviate. The EKF observes these measurement innovations and applies them via the Kalman gain to the *a priori* state estimate, nudging the state estimate in the direction of the truth state. At the next time step, thanks to the FOGM model, the acceleration state still incorporates that $w(t_i)$ but now at t_{i+1} it has decayed slightly in accordance with the time constant τ . This is important, because now if the state estimate has not fully corrected, i.e. $w(t_i)$ was not big enough to correctly capture the unmodeled acceleration, then a larger $w(t_{i+1})$ is assumed to have occurred, and is added to the acceleration state during the measurement update. This continues to happen while large measurement innovations indicate that the estimated state is not reflective of the truth state. Once innovations show the the state has converged the unmodeled accelerations decay away, or, if innovations show it is becoming over-corrected, the measurement update will subtract an unmodeled acceleration to accelerate the decay.

Figure 3.4 depicts this concept, showing a conceptual diagram as well as a sample of the filter’s response to an instantaneous change in $\Delta\mathbf{V}$. Note that this strike is 20x the magnitude of the default strike specified in Section 3.2.2 to show the pattern of the unmodeled acceleration response in the forward filter. The filter’s response to the default strike is smaller and is hidden within the noise, so the pattern of the response is not clear.

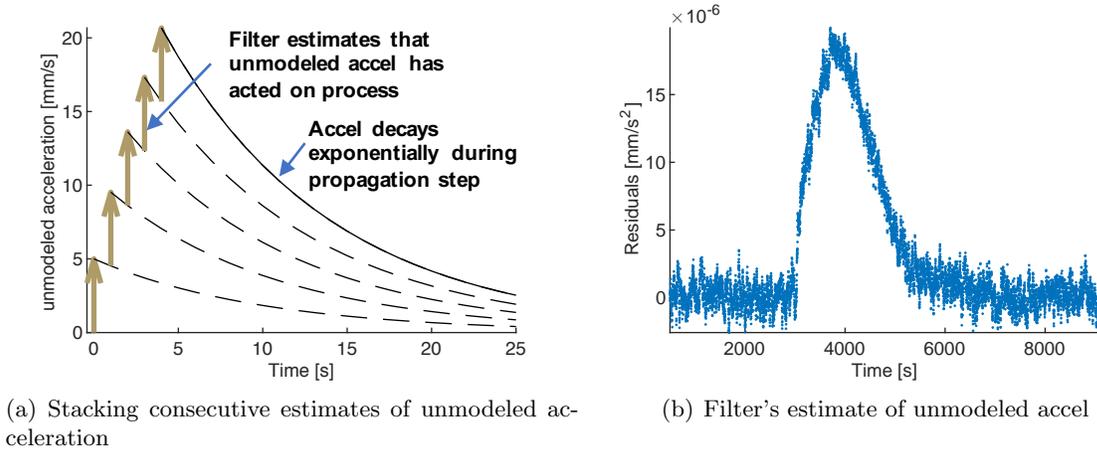


Figure 3.4: Illustration and example of FOGM dynamic model compensation responding to impulsive $\Delta\mathbf{V}$

Recalling that $\eta(t) \sim N(0, \sqrt{q\delta(t-s)})$ a designer might be inclined to increase q , the noise intensity, to increase the state covariance and allow the unmodeled acceleration to correct the trajectory more rapidly via a larger Kalman gain. In practice, however, a large q produces higher noise in the estimate of the unmodeled accelerations as the filter responds aggressively to noisy measurements. A smaller q improves the signal-to-noise ratio of the feature, but a q that is too small causes the filter to lose the state estimate when the debris strike happens, as the covariance is no longer large enough to produce a sufficiently large Kalman gain able to correct the state in the presence of the perturbation. Figure 3.5 shows the same debris strike as Figure 3.4, but with q scaled to $100\times q$ (left) and $q/100$ (center) to illustrate the resultant filter issues. Figure 3.5(c) shows the position errors associated with the small q , indicating that the filter loses the state estimate for a prolonged period of time when q is too small (relative to Figure 3.4(b), where the state estimate reconverges in a reasonable amount of time). Changing τ also produces some non-

intuitive behaviors. Section 3.4.2 shows some trades to tune filter parameters to maximize strike detection for this application. Section 3.2.2 gives the baseline ‘initial’ τ and q used to produce the results shown here, as well as the ‘final’ τ and q after tuning.

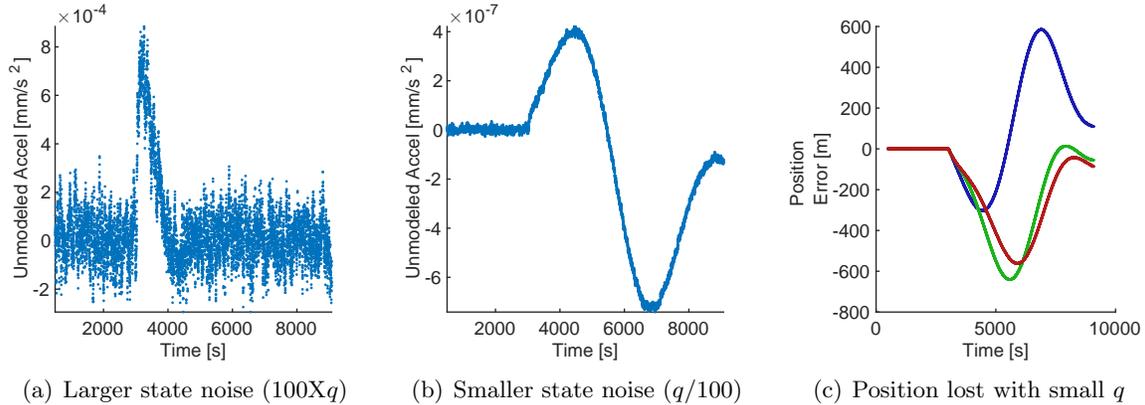


Figure 3.5: Example of changing q to tune process noise covariance matrix. Note scales on Y-axis indicating scale of unmodeled accelerations relative to Figure 3.4(b)

3.2.6 Filtering Backward in Time

Most of the test statistics developed in Section 3.3 for identifying abrupt changes in velocity employ a state estimate obtained by running the EKF backwards in time. Note that this is not the same as a backwards-smoothing extended Kalman filter or RTS smoother, where the fusing process incorporates all the data and thus would tend to smooth out the debris strike impulse instead of detecting it. In this application a backwards-running filter that incorporates only ‘future’ observations is required. Therefore, with the previous section’s background on using a FOGM model to estimate unmodeled accelerations, the task is now to replicate this behavior when the filter moves backward in time.

The objective is to run the EKF with DMC backwards through time, effectively filtering the data in the mirror-image of the forward-time filter to produce a new state estimate at each data point. This backwards state estimate is derived from the data that comes after the datapoint (t_{eval}) in time and does not incorporate data that precedes t_{eval} in time¹. Thus, two state estimates can

¹ Note that the backward filter is seeded with a state estimate from the forward filter, so the backward propagation

be compared at each epoch t_{eval} , one seeing only the data after t_{eval} and one seeing only the data before t_{eval} . Thus, if a debris strike occurs at t_{eval} , there is a discrepancy in the velocity between the two state estimates, representative of the $\Delta \mathbf{V}$ that is imparted by the strike.

To accomplish this two key aspects of the forward-filter performance need to be reproduced in backward time: the dynamic model and the unmodeled acceleration behavior. Implementing a backward dynamics model is straightforward. A negative Δt inherently causes the spacecraft to ‘fly backwards’ along its trajectory, and no changes to the basic model dynamics are required. Replicating the behavior of the unmodeled acceleration dynamics backward in time is more complicated. While a FOGM process is used to estimate the unmodeled acceleration, it is important to note that the underlying dynamics of the acceleration are not, in fact, a FOGM process. It is illustrative to examine the performance of the filter when the measurement data is reversed in time and fed into the unchanged filter. The filter’s estimate of the unmodeled acceleration during the initial filtering steps is shown in Figure 3.6.

Figure 3.6(a) shows that in the forward filter the estimated unmodeled acceleration decays during state propagation between timesteps. The *a priori* value at the next timestep is smaller than at the previous timestep, then the filter adjusts that estimate up or down by applying the Kalman gain to the measurement deviation. In the reversed filter, however, the behavior is different. The state propagation causes the estimate of the unmodeled acceleration to increase exponentially during the time update, a somewhat unstable behavior resulting in large corrections which flip between negative and positive values. Note especially the scale at the top of the y-axis: Just 35 datapoints in, the backward filter is already applying corrections to the unmodeled acceleration which are several orders of magnitude larger than the forward filter. Note that the rate of decay is set according to the time constant, τ , so for the short τ s used here the effect is pronounced. A longer τ which is more appropriate for a measurement bias or some other unmodeled dynamic

window needs to be long enough that any state corrections have been accomplished and the backwards EKF’s state estimate is converged and tracking the measurement data, the influence of the initial state estimate has faded. This is analytically unsatisfying, but appears to work pretty well. The backward filter is initialized with a large covariance, which tends to de-emphasize any correlations.

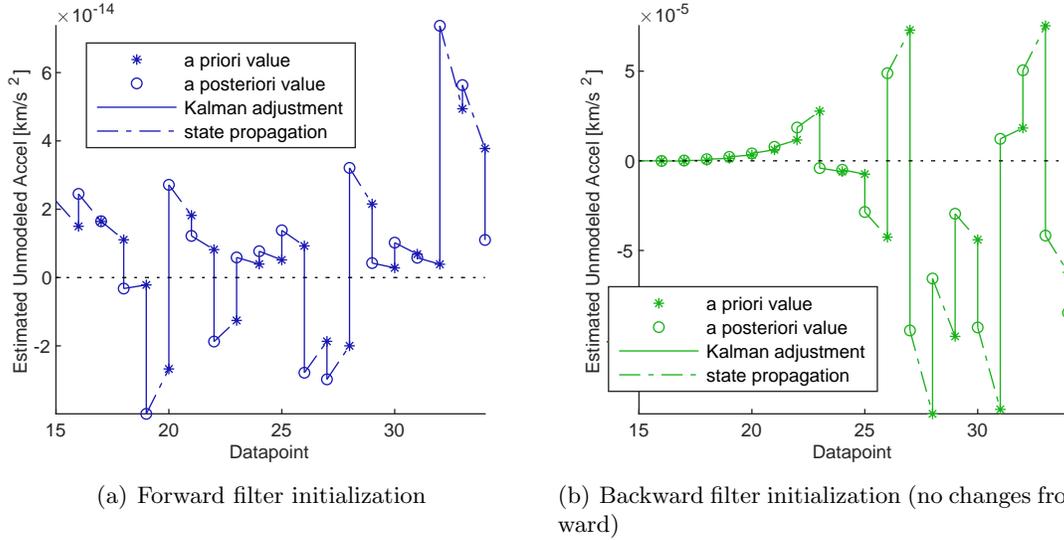


Figure 3.6: Comparison showing unmodeled acceleration estimation and propagation at the beginning of identical forward and backward filters (Note: This shows the filter initialization, no debris strike is shown)

effect would have a less significant change in performance between forward and backward filters.

To explain this asymmetric behavior recall Equation 3.7 and Figure 3.4. For a given timestep Δt , Equation 3.7 produces Equation 3.23, showing the behavior seen in Figure 3.4 where the FOGM process decays exponentially as time moves forward.

$$\mathbf{w}(t + \Delta t) = \mathbf{w}(t)e^{-\frac{\Delta t}{\tau}} \quad (3.23)$$

Figure 3.6 shows this behavior in the forward filter, with the unmodeled accelerations decaying exponentially toward zero during the state propagation. In the backward filter, however, they are increasing exponentially, leading to excessively large accelerations applied to the filter state and requiring additional large accelerations to correct them.

To correct this, the backward filter is adjusted to employ a slightly different FOGM process: one that decays as the filter moves backward in time, in a mirror-image to the forward filter behavior. This change obtains the desired FOGM behavior, with the accelerations decaying as the state progresses from one timestep to the other in both the forward and backward filters.

For this mirrored model, the relevant equations are as follows. Note that Δt could be either

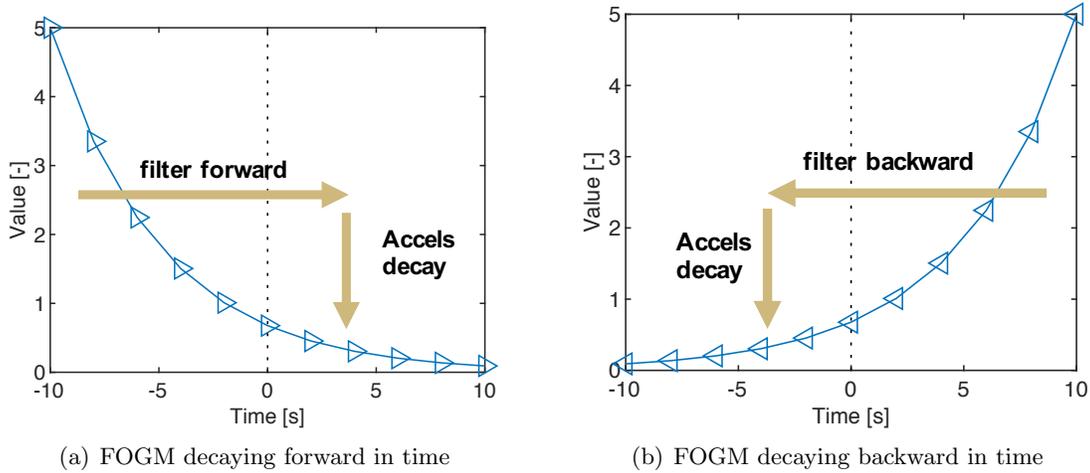


Figure 3.7: Traditional FOGM used in forward filtering and mirrored FOGM for backward filter

positive or negative, but this is used when Δt is negative and the filter is moving backward through the measurements. The \mathbf{A} matrix in Equation 3.9 also changes, with the sign on the $\frac{1}{\tau}$ now positive instead of negative in the 3x3 block on the lower right.

$$\mathbf{w}(t + \Delta t) = \mathbf{w}(t)e^{\frac{\Delta t}{\tau}} \quad (3.24)$$

$$\dot{\mathbf{w}}(t) = \frac{1}{\tau}\mathbf{w}(t) \quad (3.25)$$

To implement this mirrored FOGM the time constant τ is set to be negative in the backward filter. Comparing the relevant equations shows that this change essentially replaces Equation 3.7 with Equation 3.25 in the state propagation, and makes the necessary change to the \mathbf{A} matrix used to propagate the state transition matrix. The remaining function where τ appears is calculating the state process noise covariance \mathbf{S}_i , and Section 3.2.7 shows that with a $-\tau$ each γ_{fwd} term is equal to the analogous γ_{bkwd} term derived using the mirrored FOGM indicated in Equation 3.25. For a more general solution, the key principle is that the dynamic behavior of the processes needs to be time-symmetric for the backward filter to function correctly. For the FOGM a $-\tau$ accomplishes this, for other models alternate modifications may be needed.

One further change is required for the backward filter: similar to Fraser's[77] developments, the state process noise covariance is subtracted when filtering backward in time. This can be seen

by referring to the process noise covariance integral in Equation 3.11, note that \mathbf{S}_i is derived from the integral from t_{i-1} to t_i . For the backward filter, the state process noise covariance moves from $i + 1$ to i instead of from $i - 1$ to i . Integrating backward in time results in a negative integral, but the state process noise variance needs to be added to capture the increasing uncertainty in the state estimate due to process noise from one filter step to the next. Therefore, in the backward filter \mathbf{S}_i is subtracted.

$$\mathbf{S}_{i,\text{bkwd}} = - \int_{t_{i+1}}^{t_i} \Phi(t_i, \epsilon) \mathbf{B}(\epsilon) \mathbf{Q}(\epsilon) \mathbf{B}^T(\epsilon) \Phi^T(t_i, \epsilon) d\epsilon \quad (3.26)$$

Note that this concept can alternately be applied by subtracting the state process noise covariance \mathbf{Q} , as Fraser and Potter[78] do when filtering backward.

With these adjustments, the backward filter performance, shown in Figure 3.8, is comparable to the forward filter performance (Figure 3.2), although it is mirrored as the filter proceeds through the data in the opposite direction.

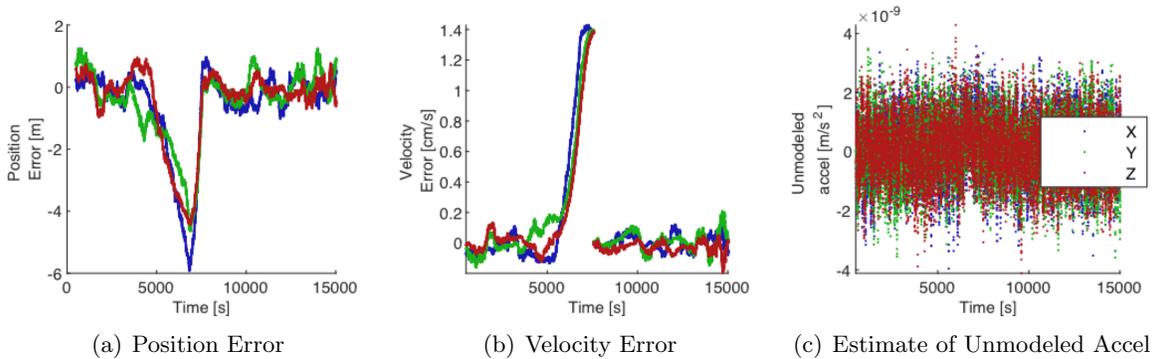


Figure 3.8: Performance of backward filter on telemetry with small debris strike

3.2.7 Derivation of Process Noise Covariance for Mirrored FOGM in Backward Filter

To ensure that the mirrored FOGM in the backward filter is implemented correctly, this section rederives the γ terms with the mirrored FOGM for comparison to the forward γ terms. First, the FOGM acceleration is integrated twice to establish the acceleration, velocity, and position

at $t + \Delta t$. Note that this is the kinematic effect of the process noise only, it does not incorporate system dynamics. For a keplerian orbit and a short timestep this approximation is sufficient for characterizing the effect of noise on the state, but if longer timesteps or stronger system dynamics were involved it should be revisited.

$$\mathbf{w}(t + \Delta t) = \mathbf{w}(t)e^{\Delta t/\tau} \quad (3.27)$$

$$\mathbf{v}(t + \Delta t) = \mathbf{v}(t) + \mathbf{w}(t)\tau(e^{\Delta t/\tau} - 1) \quad (3.28)$$

$$\mathbf{r}(t + \Delta t) = \mathbf{r}(t) + \mathbf{v}(t)\Delta t - \mathbf{w}(t)\tau\Delta t + \mathbf{w}(t)\tau^2(e^{\Delta t/\tau} - 1) \quad (3.29)$$

Where the constants of integration are solved for by noting that at $\Delta t = 0$, $\mathbf{w}(t + \Delta t) = \mathbf{w}(t)$, $\mathbf{v}(t + \Delta t) = \mathbf{v}(t)$, and $\mathbf{r}(t + \Delta t) = \mathbf{r}(t)$.

Next, form the state transition matrix at $t + \Delta t$ by taking the partial derivatives of these equations with respect to the state. Note that this assumes that Δt is small enough for a first-order Taylor series truncation to be used. The STM is found by

$$\Phi(t + \Delta t, t) = \frac{\partial \mathbf{X}(t + \Delta t)}{\partial \mathbf{X}(t)} \quad (3.30)$$

$$\Phi(t + \Delta t, t) = \begin{bmatrix} \mathbf{I}_3 & \Delta t \mathbf{I}_3 & [\tau^2(e^{\frac{\Delta t}{\tau}} - 1) - \tau \Delta t] \mathbf{I}_3 \\ \mathbf{0}_{3 \times 3} & \mathbf{I}_3 & \tau(e^{\frac{\Delta t}{\tau}} - 1) \mathbf{I}_3 \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & e^{\frac{\Delta t}{\tau}} \mathbf{I}_3 \end{bmatrix} \quad (3.31)$$

For convenience, the following definitions are formed from this state transition matrix

$$\Phi_{12} = \Delta t \quad (3.32)$$

$$\Phi_{13} = \tau^2(e^{\frac{\Delta t}{\tau}} - 1) - \tau \Delta t \quad (3.33)$$

$$\Phi_{21} = \tau(e^{\frac{\Delta t}{\tau}} - 1) \quad (3.34)$$

$$\Phi_{33} = e^{\frac{\Delta t}{\tau}} \quad (3.35)$$

This approximation of the state transition matrix is then used to form the integrand for the state process noise covariance matrix. With Equations 3.12 and 3.13, the integrand of the process noise covariance matrix can be assembled, denoted as \mathbf{N}_i . Recalling Equation 3.11, the integrand

is

$$\mathbf{N}_i = \Phi(t + \Delta t, \epsilon) \mathbf{B}(\epsilon) \mathbf{Q}(\epsilon) \mathbf{B}^T(\epsilon) \Phi^T(t + \Delta t, \epsilon) \quad (3.36)$$

such that

$$\mathbf{S}_i = \int_t^{t+\Delta t} \mathbf{N}_i d\epsilon \quad (3.37)$$

Multiplying this out and approximating \mathbf{N}_i for one timestep from t to $t + \Delta t$ reduces to

$$\mathbf{N}_i \approx \begin{bmatrix} \Phi_{13}^2 \tilde{\mathbf{Q}} & \Phi_{13} \Phi_{21} \tilde{\mathbf{Q}} & \Phi_{13} \Phi_{33} \tilde{\mathbf{Q}} \\ \Phi_{13} \Phi_{21} \tilde{\mathbf{Q}} & \Phi_{21}^2 \tilde{\mathbf{Q}} & \Phi_{33} \Phi_{21} \tilde{\mathbf{Q}} \\ \Phi_{13} \Phi_{33} \tilde{\mathbf{Q}} & \Phi_{33} \Phi_{21} \tilde{\mathbf{Q}} & \Phi_{33}^2 \tilde{\mathbf{Q}} \end{bmatrix} \quad (3.38)$$

For the filter, the Φ coefficients of \mathbf{N}_i are multiplied out and then integrated to form the γ coefficients for calculating \mathbf{S}_i . The constants of integration are found by setting $\gamma = 0$ when $\Delta t = 0$.

$$\Phi_{13}^2 = \left(-\Delta t \tau - \tau^2 + e^{\Delta t/\tau} \tau^2 \right)^2 \quad (3.39)$$

$$\gamma_{rr, \text{bkwd}} = \frac{\Delta t^3 \tau^2}{3} + \Delta t^2 \tau^3 + \Delta t \tau^4 - 2e^{\Delta t/\tau} \Delta t \tau^4 - \frac{\tau^5}{2} + \frac{1}{2} e^{\frac{2\Delta t}{\tau}} \tau^5 \quad (3.40)$$

$$\Phi_{21}^2 = \tau^2 \left(\Delta t - 2e^{\Delta t/\tau} \tau + \frac{1}{2} e^{\frac{2\Delta t}{\tau}} \tau \right) \quad (3.41)$$

$$\gamma_{vv, \text{bkwd}} = \frac{1}{2} \tau^2 \left(2\Delta t + \left(3 - 4e^{\Delta t/\tau} + e^{\frac{2\Delta t}{\tau}} \right) \tau \right) \quad (3.42)$$

$$\Phi_{33}^2 = e^{\frac{2\Delta t}{\tau}} \quad (3.43)$$

$$\gamma_{ww, \text{bkwd}} = -\frac{\tau}{2} + \frac{1}{2} e^{\frac{2\Delta t}{\tau}} \tau \quad (3.44)$$

$$\Phi_{13} \Phi_{21} = \left(-1 + e^{\Delta t/\tau} \right) \tau \left(-\Delta t \tau - \tau^2 + e^{\Delta t/\tau} \tau^2 \right) \quad (3.45)$$

$$\gamma_{rv, \text{bkwd}} = \frac{\Delta t^2 \tau^2}{2} + \Delta t \tau^3 - e^{\Delta t/\tau} \Delta t \tau^3 + \frac{\tau^4}{2} - e^{\Delta t/\tau} \tau^4 + \frac{1}{2} e^{\frac{2\Delta t}{\tau}} \tau^4 \quad (3.46)$$

$$\Phi_{13} \Phi_{33} = e^{\Delta t/\tau} \left(-\Delta t \tau - \tau^2 + e^{\Delta t/\tau} \tau^2 \right) \quad (3.47)$$

$$\gamma_{rw, \text{bkwd}} = -e^{\Delta t/\tau} \Delta t \tau^2 - \frac{\tau^3}{2} + \frac{1}{2} e^{\frac{2\Delta t}{\tau}} \tau^3 \quad (3.48)$$

$$\Phi_{33} \Phi_{21} = e^{\Delta t/\tau} \left(-1 + e^{\Delta t/\tau} \right) \tau \quad (3.49)$$

$$\gamma_{vw, \text{bkwd}} = \frac{1}{2} \left(-1 + e^{\Delta t/\tau} \right)^2 \tau^2 \quad (3.50)$$

For the backward filter, the text established that setting the time constant τ to be negative is an effective means to adjust the code to use the mirrored FOGM in the backward filter. For the \mathbf{S}_i matrix, $\tilde{\mathbf{Q}}$ is not a function of τ , so the question is whether using a negative τ in the forward-filter γ_{fwd} coefficients renders an equivalent result to the γ_{bkwd} coefficients derived above. Comparing Equations 3.40, 3.42, 3.44, 3.46, 3.48, 3.50 to Equations 3.16-3.21 with a negative τ , it does render an equivalent result, so the state process noise covariance calculation can be left intact in the backward filter, no additional sign changes are required within the γ terms beyond setting τ to be negative.

3.3 Developing Test Statistics

Various test statistics are developed to accentuate the effects of a debris strike relative to the baseline noise. The intent of these test statistics is to explore various ways of post-processing the filter state estimate to enhance the detectability of debris strike features. It is important to note that these results all show simulated telemetry, and as such the measurement noise is perfectly Gaussian and there are no additional orbit perturbations acting on the state, except in Section 3.4.4. In applying these techniques to real-world measurements all the various sources of difficult-to-characterize noise and perturbations will have a substantial effect on filter performance. It is expected that some test statistics which produce a better signal-to-noise ratio (SNR) on simulated data may perform worse on real-world data, if they require a longer period of data to function, for example. Therefore, several test statistics are investigated to produce an *a la carte* selection of techniques for on-orbit application.

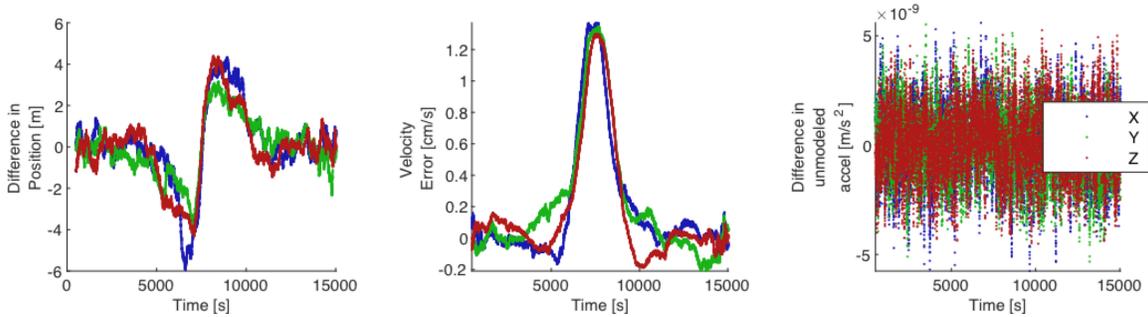
3.3.1 Test Statistic 1: Difference between forward and backward state estimates

The results in Figures 3.2 and 3.8 depend on knowledge of the truth state to produce the deviation in truth state, but a proxy can be obtained by differencing the states estimated by the forward and backward filters. Figure 3.9 shows the difference between the forward and backward state estimates and indicates a significant feature in the velocity at the time of the debris strike.

This is calculated using

$$\mathbf{X}_{\text{diff}} = \mathbf{X}_{\text{bkwd}} - \mathbf{X}_{\text{fwd}} \quad (3.51)$$

The velocity difference serves as the first test statistic evaluated for strike detection effectiveness. It also provides the estimate of ΔV , which is used even when other test statistics are used to identify the presence of a strike.



(a) Difference in Position Estimate (b) Difference in Velocity Estimate (c) Difference in Unmodeled Acceleration

Figure 3.9: Difference between forward and backward EKF state estimates shows significant spike in velocity without requiring knowledge of truth state

3.3.2 Test Statistic 2: Mahalanobis distance

The second test statistic is the Mahalanobis distance. The Mahalanobis distance expresses the generalized distance between two vectors, providing a multidimensional generalization of the distance between a point and an associated distribution. This is calculated per Equation 3.52, which calculates the vector difference between forward and backward states, followed by Equation 3.53, which calculates the Mahalanobis distance of that state difference relative to the expected uncertainties. Note that Equation 3.52 uses the forward *a posteriori* state and the backward *a priori* state. Thus, both states are at the same epoch, but only the forward state incorporates the measurement data at that epoch.

$$\mathbf{X}_{M,i} = \mathbf{X}_{F,i} - \bar{\mathbf{X}}_{B,i} \quad (3.52)$$

$$D_{\text{MH}} = \sqrt{\mathbf{X}_{M,i}^T \mathbf{P}_{S,i}^{-1} \mathbf{X}_{M,i}} \quad (3.53)$$

This chapter investigates both full-state and velocity-only Mahalanobis distances to characterize the discrepancies between the forward and reverse filters. The full-state Mahalanobis distance is calculated using the distance between the forward and backward state estimates and the smoothed covariance (see Equation 3.53, next section), while the velocity-only distance replaces \mathbf{X} with just the velocity substate, and \mathbf{P}_S with the corresponding central 3x3 block of the smoothed covariance. The results are illustrated in Figure 3.10, which shows both full-state and velocity-only Mahalanobis distances. This overlay indicates that the baseline noise is higher in the full-state distance, so only the velocity distance is used as a test statistic as it has a higher signal to noise ratio.

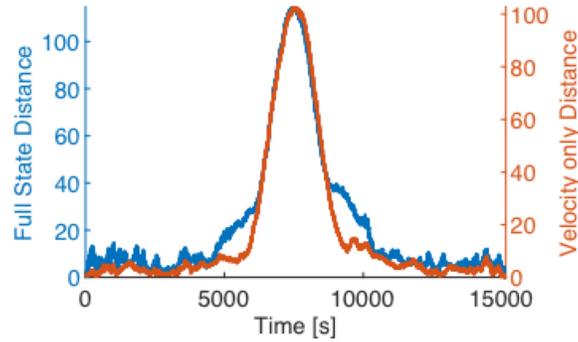


Figure 3.10: Mahalanobis distances

3.3.3 Test Statistic 3: Smoothed accelerations from fused backward and forward state estimates

The state estimates and covariances from the forward and backward filters are fused into a smoothed estimate using a modified Fraser-Potter fixed-interval smoother[75, 79]. Note that the *a posteriori* estimates from the forward filter ($\mathbf{X}_{F,i}$, $\mathbf{P}_{F,i}$) are fused with the *a priori* estimates from the reverse filter ($\bar{\mathbf{X}}_{B,i}$, $\bar{\mathbf{P}}_{B,i}$) at the same measurement time (i). Thus, like the Mahalanobis distance, all the estimates are at the same time but only the forward filter incorporates the measurement update at that time. The smoothed state estimate and covariance are calculated per

$$\mathbf{X}_{S,i} = \mathbf{W}_{F,i}\mathbf{X}_{F,i} + (\mathbf{I} - \mathbf{W}_{F,i})\bar{\mathbf{X}}_{B,i} \quad (3.54)$$

$$\mathbf{P}_{S,i} = \mathbf{W}_{F,i} \mathbf{P}_{F,i} \mathbf{W}_{F,i}^T + (\mathbf{I} - \mathbf{W}_{F,i}) \bar{\mathbf{P}}_{B,i} (\mathbf{I} - \mathbf{W}_{F,i})^T \quad (3.55)$$

Where $\mathbf{W}_{F,i}$ is

$$\mathbf{W}_{F,i} = \bar{\mathbf{P}}_{B,i} (\mathbf{P}_{F,i} + \bar{\mathbf{P}}_{B,i})^{-1} \quad (3.56)$$

The smoothed state estimates for the velocity and acceleration are shown in Figure 3.11. The smoothed velocity exhibits an intuitive behavior where the smoothed estimate lies between the forward and backward filter states with the high-frequency noise removed, as shown in Figure 3.11(a). Note that this figure employs the truth orbit to display the errors in state estimates.

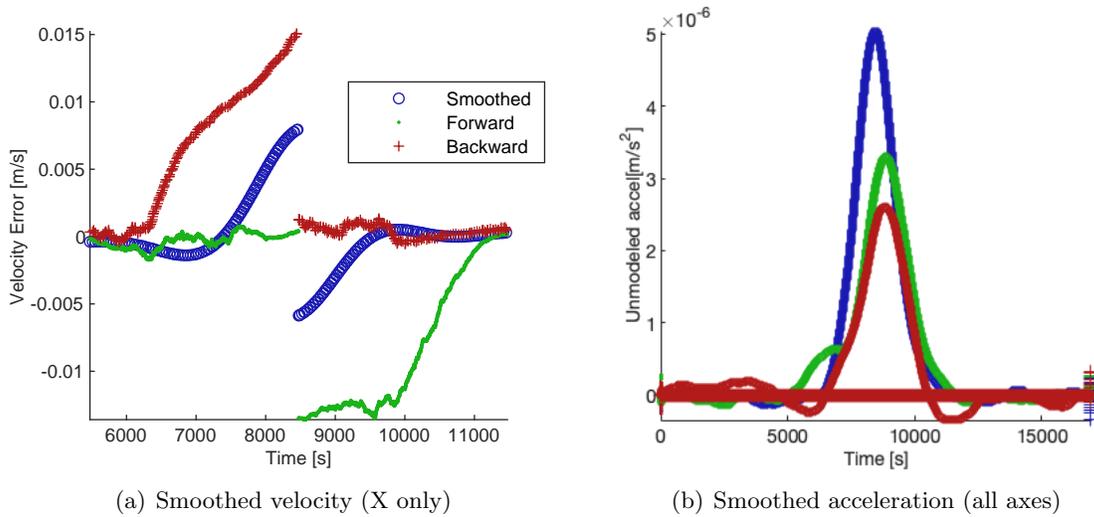


Figure 3.11: Smoothed state exhibits signature in unmodeled acceleration

An interesting feature appears in the smoothed accelerations (Figure 3.11(b)), which are derived from the filtering process and do not rely on knowledge of the true state. Recall that the unmodeled acceleration estimates in both the forward and backward filter are close to zero ($\sim 10^{-9}$ m/s²) with no distinct features (reference Figure 3.9). In the smoothed acceleration, however, a distinct feature appears on the order of $\sim 10^{-6}$ m/s². To explain this recall Equation 3.56, where a weighting matrix, \mathbf{W}_F is calculated based on the ratio between the forward and backward covariances. The off-diagonal terms in this weighting matrix allow each smoothed state to be affected when a discrepancy occurs between the forward and backward states (this can be seen in Equation 3.54). In this case, a discrepancy in the velocities is differentiated into the estimate

of the unmodeled acceleration via the terms $\mathbf{W}_{F(7,4)}$, $\mathbf{W}_{F(8,5)}$, and $\mathbf{W}_{F(9,6)}$ (specifying the ij th term of the \mathbf{W}_F matrix as $\mathbf{W}_{F(i,j)}$). This unmodeled acceleration feature is the third test statistic evaluated for debris strike detection performance.

3.3.4 Test Statistic 4: McReynold's filter-smoother consistency test

The fourth test statistic is derived from the McReynold's Filter-Smoother Consistency Test. In this test, an abrupt deviation from expected state can be identified by characterizing the consistency between the forward filter state estimate and the smoothed state estimate. First the difference between the forward state estimate and the smoothed state is found (per Wright[66])

$$\mathbf{X}_D = \mathbf{X}_F - \mathbf{X}_S \quad (3.57)$$

$$\mathbf{P}_D = \mathbf{P}_F - \mathbf{P}_S \quad (3.58)$$

Then the McReynold's test statistic (R_m) is calculated as the ratio of each individual state ($X_{D,j}$ for state j) over the square root of the corresponding diagonal element in the covariance ($\sigma_{D,j} = \sqrt{P_{D,j,j}}$)

$$R_{m,j} = |X_{D,j}|/\sigma_{D,j} \quad (3.59)$$

An alternative formulation calculates a single scalar test statistic to express the consistency between the filter and smoother, rather than a separate statistic for each state. This full-state McReynold's consistency statistic is calculated via

$$R_{\text{scalar}} = \sqrt{\mathbf{X}_D^T \mathbf{P}_D \mathbf{X}_D} \quad (3.60)$$

The resultant McReynold's filter-smoother consistency statistic shows signatures in all nine states, as shown in Figure 3.12. This indicates that all nine states show a feature at the time of the debris strike, but the acceleration states show the strongest feature. The scalar consistency statistic is plotted on each state on the second axis, and shows that it is slightly better than the position states and nearly comparable the acceleration states (assessing the quality of the test statistic by the signal to noise ratio, it can be seen that with the axis scaled such that the peaks match

the baseline noise is higher, relatively speaking, in the position states than the scalar consistency statistic, for example).

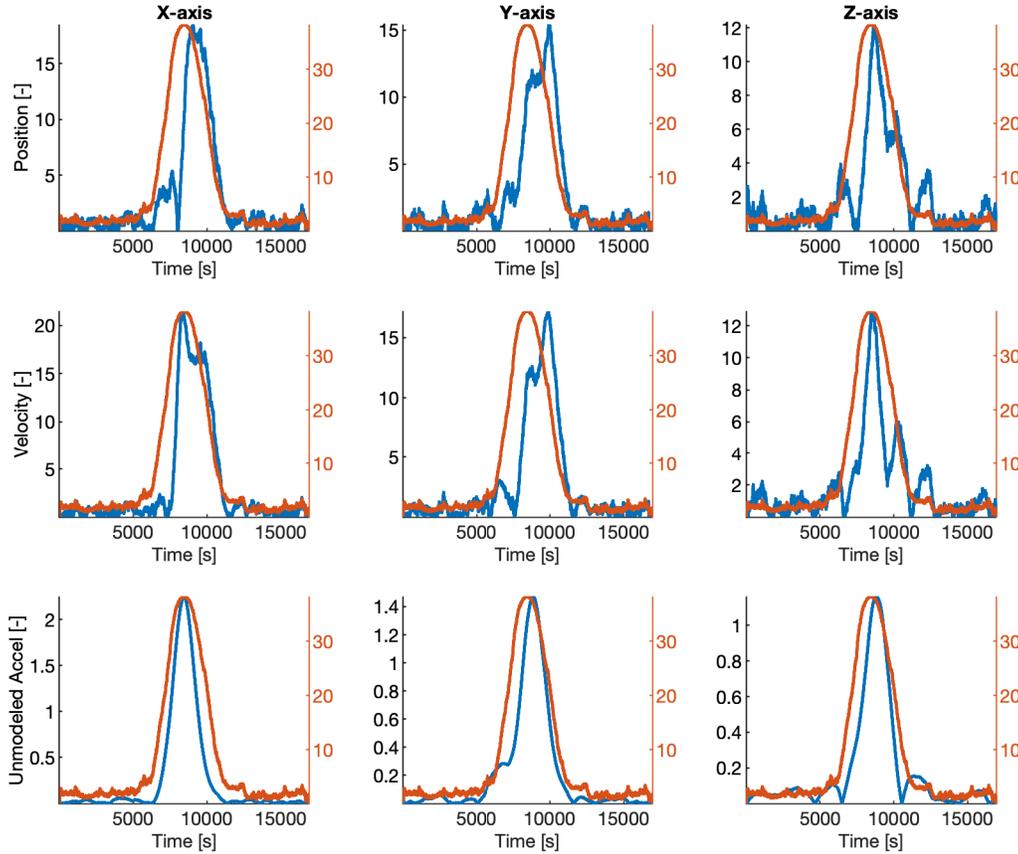


Figure 3.12: McReynold's consistency statistic shows feature in all nine states. Red indicates scalar consistency statistic from Equation 3.60 plotted on separate axis to compare to each state from Equation 3.59

3.3.5 Test Statistic 5: Measurement residuals from propagated state estimate

The fifth test statistic applies change detection methods to assess the measurement residuals relative to a nominal orbit propagated forward and backward from the current state estimate. The state just before the strike is propagated forward in time and the measurement residuals relative to the state estimate are determined (shown in the plots on the right side of Figure 3.13). Simultaneously the backward state estimate at $t_i = t_{\text{strike}} + 1$ is propagated backward toward $t = 0$

and compared to the measurements before the strike. Thus, both state estimates are propagated across the strike such that the state estimate does not incorporate the change in velocity, and the deviation in orbit due to the strike is observable in the measurement residuals. For reference, a simulation with no strike is also plotted in Figure 3.13 to show the magnitude of the residual patterns created by the effects of the debris strike. These results indicate that the position deviations tend to be greater than the velocity deviations.

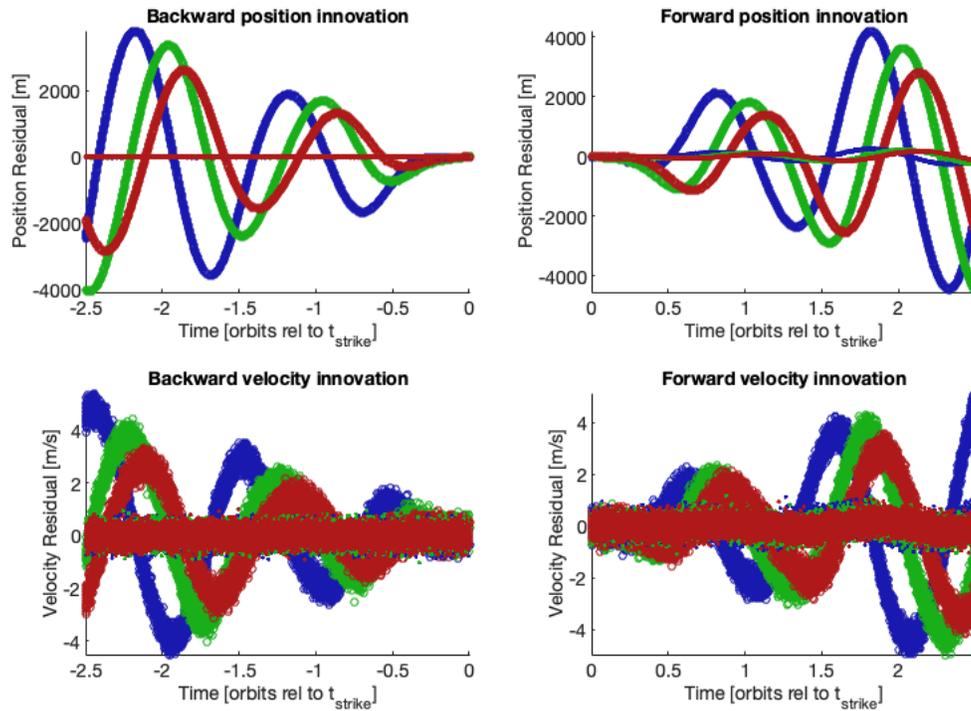


Figure 3.13: Residuals generated by comparing forward and backward orbit propagations to measurements

Note that the example in Figure 3.13 is for a debris strike in the $-\hat{v}$ direction which produces a change in semi-major axis and causes the residuals to build up quickly due to asynchronous periods. A debris strike in a different direction would have a less dramatic effect, but the majority of debris strikes in LEO tend to be roughly in the $-\hat{v}$ direction[80] so this provides a useful example.

Note that this test statistic is especially sensitive to orbit perturbations which may perturb the state more slowly than a debris strike, but still significantly, and its applicability to real-world

data will likely be orbit-dependent. This is because a longer period of time is used to generate the test statistic. Also, this test statistic takes by far the longest to process, since the orbit must be propagated forward and backward from each point at which the test statistic is evaluated, and then the resultant noise in the distribution of those residuals needs to be characterized to develop a meaningful test statistic using a log likelihood ratio test. Since this test statistic adds substantially to processing time and is unlikely to generate results comparable to the real world, it is not included in the trade studies and comparisons of the other test statistics, but is described here. It might be well applied on systems with precision orbit determination filters/propagators or limited perturbing effects.

3.3.6 Test Statistic 6: Change in orbital constants

In prior work[6, 17] the change in satellite mean altitude (dSMA) is used to identify orbit perturbations. For comparison to these efforts, the change in semimajor axis is used for the sixth test statistic. However, a dSMA will only occur if a strike changes the magnitude of the spacecraft velocity. Out-of-plane ΔV components will change the orbit plane, but not the semimajor axis. ΔV components which change the direction of the velocity but not the magnitude will likewise not change the semi-major axis. Therefore, for this test statistic three orbital constants are used: the traditional dSMA, the eccentricity (e), and the direction of the orbital angular momentum (h_{dir}). It is worth noting that in many LEO orbits the prevalent debris flux is near-head-on,[80] so the dSMA alone would likely be a strong indicator for many collisions, but for other orbits, like GEO, prevalent fluxes may occur in other directions, so the more general solution using three orbital constants is developed here. See discussion for additional cautions on using an orbit change detector with highly directional sensitivity.

The orbital constants are calculated using only the forward-filter state, to represent a straightforward application requiring minimal post-processing of existing data. Using the filter's estimate of the position \mathbf{r} and velocity \mathbf{v} , the semi-major axis a , eccentricity e and orbital angular momentum \mathbf{h} are calculated as follows. The magnitudes of the position and velocity are denoted r and v ,

respectively, while μ represents the geocentric gravitational constant.

$$a = \frac{1}{\frac{2}{r} - \frac{v^2}{\mu}} \quad (3.61)$$

$$\mathbf{e} = \frac{(v^2 - \frac{\mu}{r})\mathbf{r} - (\mathbf{r} \cdot \mathbf{v})\mathbf{v}}{\mu} \quad (3.62)$$

$$\mathbf{h} = \mathbf{r} \times \mathbf{v} \quad (3.63)$$

To reduce the direction of the angular momentum to a scalar test statistic the azimuth and elevation of the angular momentum unit vector is calculated (i.e, in spherical coordinates as expressed in the inertial frame), and then the root-sum-square of these values is taken to obtain a scalar representing the current direction of the orbital plane. Note that this is an approximation that assumes small angles, but since the intention is just to reduce the change in orbit plane to a scalar this is adequate for this application.

With these scalar test statistics, changes in the orbital state due to impacts in any direction in the orbit frame are rendered observable. A strike in the velocity direction produces changes in the semimajor axis and eccentricity, a strike in the radial direction produces a change in eccentricity, and an out-of-plane strike produces a change in the direction of the orbit plane. In previous test statistic examples the strike is applied evenly across each of the three axes in inertial space, but Figure 3.14 shows the strike is applied evenly across the three directions of the orbit frame to produce a change in each of these test statistics.

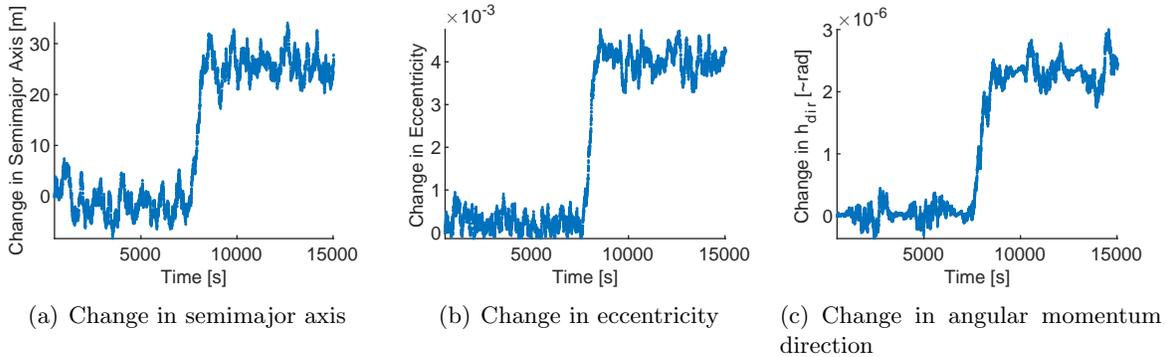


Figure 3.14: Orbital constants calculated from forward-filter state estimates show change at time of strike

Since these test statistics can be applied without reversed filters or smoothing the filter parameters are adjusted to be more typical of standard operations. In Figure 3.14 the process noise q is inflated to be more representative of a typical navigation filter. The value used is 100X the initial default value for q shown in Section 3.2.2. In the prior test statistics the filter is designed to allow the state to deviate to make the debris strike more observable, but the intention of this test statistic is to compare to “typical” methods so the filter is re-tuned toward a more traditional performance, where the covariance is tuned to correct the state quickly in the presence of perturbations to maintain an accurate state estimate based on the available data.

3.4 Tuning, Comparison, and Performance Assessment

To evaluate the filter’s performance a streaming implementation is introduced to analyze a data series and identify strikes without *a priori* knowledge of the strike time. The forward filter is applied once across the entire dataset (in on-line analysis this could be updated persistently as telemetry is received). Then the test statistics are calculated using a sliding window where all test statistics are evaluated at t_{eval} , which is situated a fixed interval behind the leading datapoint in the window. Thus, the backward filter can be applied backward through the data and across t_{eval} , then the smoother, allowing the calculation of each test statistic at t_{eval} based on data windows surrounding t_{eval} . Since the test statistics show a spike in the presence of a debris strike an ‘integrated’ test statistic is stored as well as the peak value, to determine if that improves the detectability of the debris strike feature.

3.4.1 Normalizing Test Statistics as SNR

Each test statistic is calculated at each t_{eval} . The peak value of each test statistic is stored and also a window of data is empirically integrated to form an integrated test statistic, to see if that produces a stronger signal. To compare test statistics against each other each value must be normalized. This is done by calculating the signal-to-noise ratio (SNR) of each test statistic, where the baseline noise is calculated in a long simulation with no debris strike applied.

Calculating the detectability threshold T_{thres} from the baseline noise is a nontrivial exercise on simulated data and is even more complicated on real data. The core of the challenge lies in assessing the desired probability of false alarm vs. probability of missed detection. A higher detection threshold will result in fewer false alarms but may miss some detectable strikes, while a lower threshold will detect more strikes but also return a high rate of false alarms. Figure 3.15 shows a diagram of this decision, which is inherently application dependent: for some applications a low false alarm rate may be paramount, while for others a high chance of detecting small strikes may be preferred. The blue curve shows the probability density function (pdf) of the test statistic output when no debris strike is present. The red shows the pdf when a given debris strike has occurred. Depending on where the threshold is set (dashed black line), the probability of false alarm can be smaller or larger, and the probability of missed detection will become larger or smaller according to the same threshold.

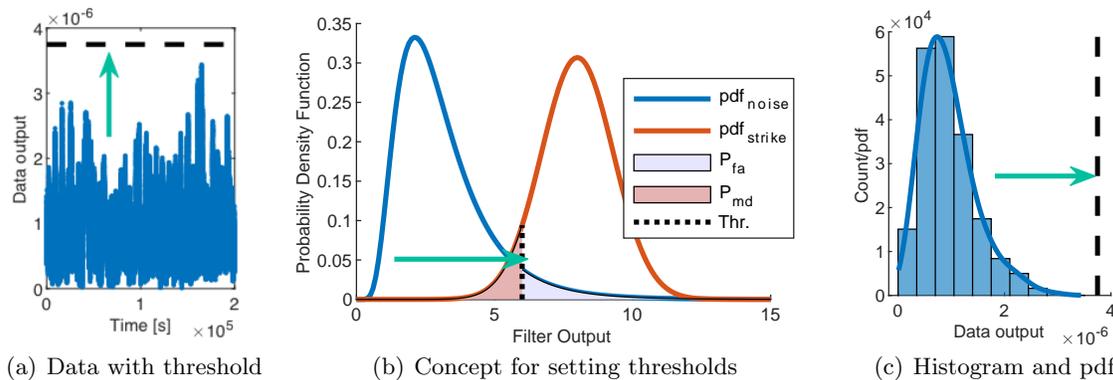


Figure 3.15: Depiction of process for setting detection thresholds. Left: goal is to set threshold above baseline noise, which varies. Center: Probability density functions of filter output with and without debris strike indicate that threshold location is a trade between probability of false alarm and probability of missed detection. Right: Kernel distribution fitted to histogram of filter output with no debris strike used in conjunction with desired false alarm rate to establish threshold.

For this chapter, the thresholds are determined by setting a desired rate of false alarm of one per week. With this rate of false alarm and a given telemetry rate a target probability of false alarm is established. A long simulation is conducted to establish the baseline noise inherent in each test statistic when no debris strike is present. An example is shown in Figure 3.15(a), with

the test statistic output evaluated at 200,000 datapoints and the resultant baseline noise, with no debris strike, plotted. The characteristics of the distribution of this noise varies depending on test statistic, so a kernel distribution is fitted to the noise, as shown in Figure 3.15(c).

From this distribution and the desired probability of false alarm the detection threshold T_{thres} is set - i.e, it is set at a point where the remainder of the tail of the distribution, above the threshold, corresponds to the desired false alarm rate. With this detection threshold the SNR of a given test statistic in the presence of a given debris strike is calculated. A debris strike is applied repeatedly to the simulation in a Monte Carlo, with randomized strike time and orbit location, to establish the magnitude of each test statistic's response to the debris strike by averaging the response across all Monte Carlo runs. The magnitude of this response divided by the detection threshold is saved as the SNR for the test statistic. Thus, the various test statistics can be compared to each other to assess their relative performance in detecting strikes.

3.4.2 Tuning τ and q

Typically q is tuned to obtain a covariance that bears the desired scale relative to the state errors (i.e, an appropriate number of datapoints fall within the 3σ boundary) or to minimize state errors. In this application, however, the goal is not to obtain the best possible state estimate at each time, instead the goal is to detect debris strikes. A traditionally tuned filter might tend to correct the state too quickly and smooth out the debris strike, whereas in this case it is desired to let the effects of the debris strike affect the state for long enough to cause a detectable signal. τ is also a tuning parameter, and while intuitively a short τ might mimic a short debris pulse, in practice it can fail to accumulate a substantial enough acceleration quickly due to the mechanics depicted in Figure 3.4(a). These two influences interact and a trade study is required to determine which combination of τ and q produces the best SNR.

An initial trade study is conducted to understand the relationships between τ , q , and the SNR for the various test statistics. A subset of results are shown in Figure 3.16. These indicate that there is a clear zone where τ and q produce a higher SNR, but that this zone varies somewhat

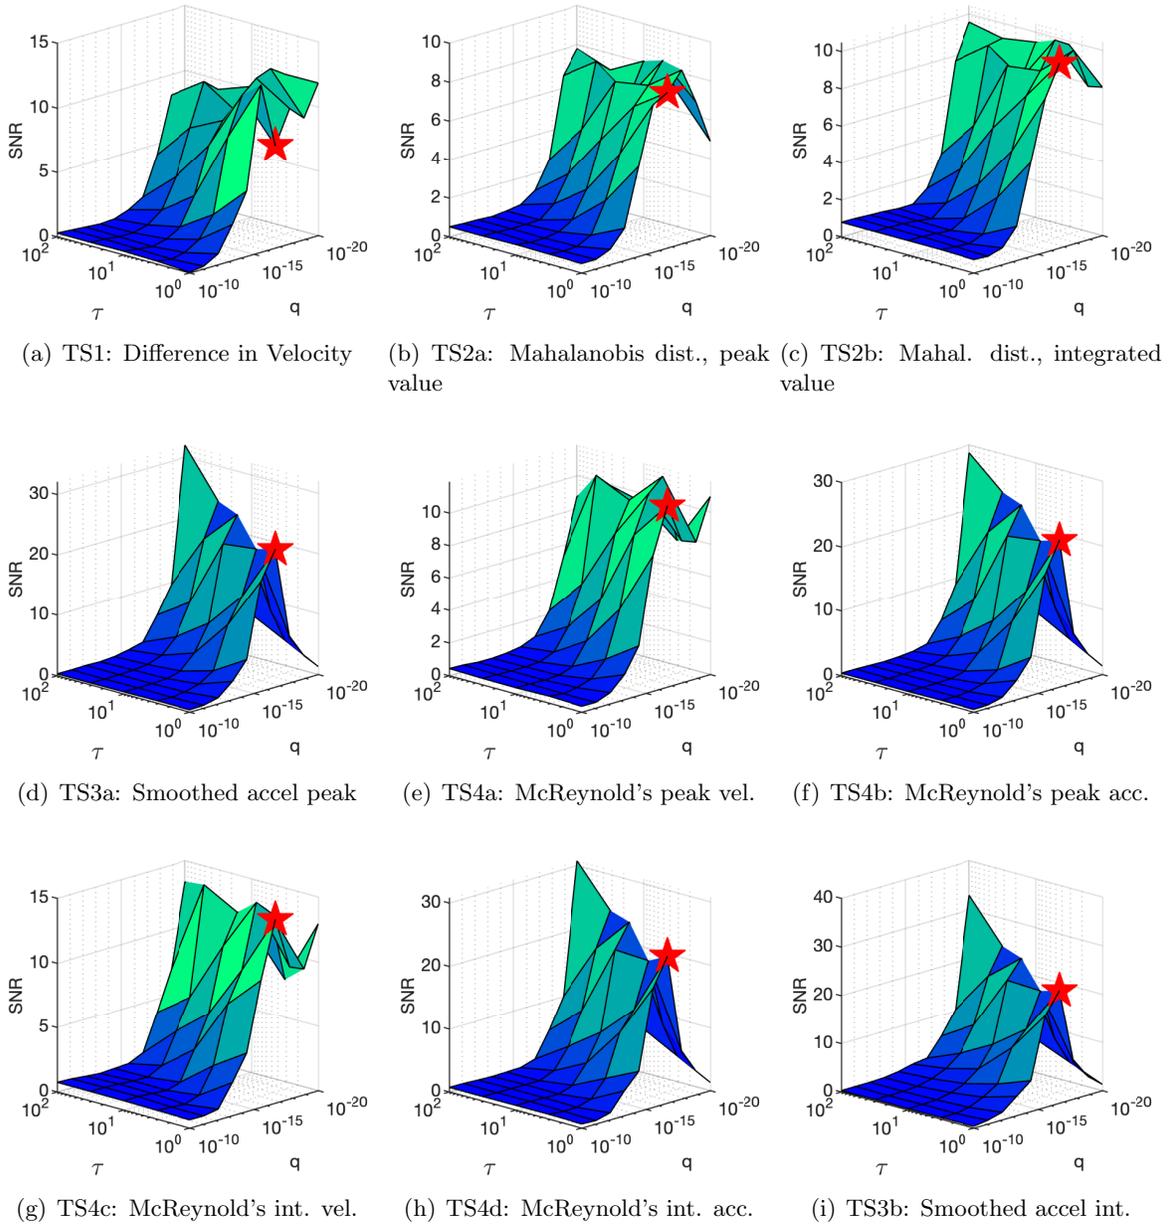


Figure 3.16: Trading τ and q to determine effect on SNR of various test statistics.

between the test statistics. In general, a smaller q and larger τ produce an equivalent or sometimes improved test statistic, but with a smaller q the filter tends to diverge from the truth state for longer when an unexpected acceleration occurs, so the larger q and smaller τ is selected. The red star in the figure is at $\tau=1$ s and $q=2.15e-17$ m²/s⁵, which produces a strong SNR in most test statistics. Note that the acceleration test statistics produce a noticeably higher SNR than the

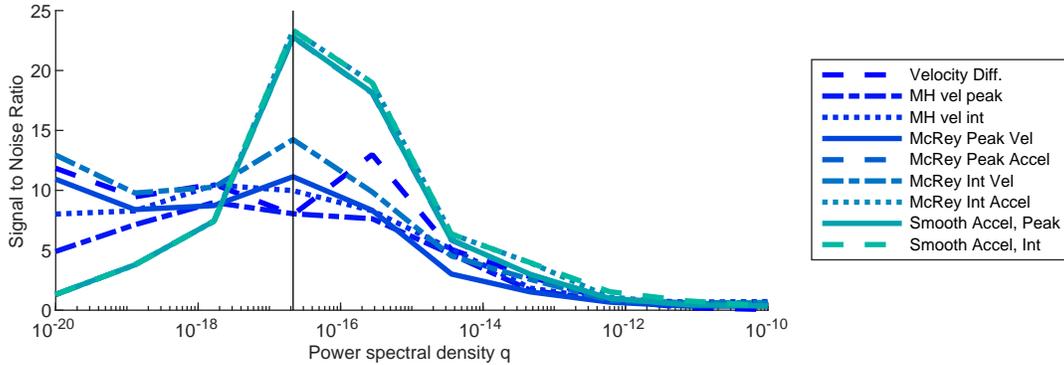


Figure 3.17: Superimposing SNR from various test statistics to compare results for given time constant τ .

velocity-related test statistics.

Figure 3.17 shows these trade studies superimposed for a single τ , to illustrate the relationship between q and SNR for each of the test statistics. This shows that the best test statistics are the smoothed accelerations and the McReynold's consistency statistic for accelerations, but it is important to remember that some other test statistics may perform better in the real-world where additional orbit perturbations may affect the state. Orbit perturbations typically affect the state more slowly than a debris strike, so test statistics which require a long span of data may perform well on simulated telemetry but struggle in real-world applications.

3.4.3 Performance When Detecting and Estimating Unknown Strikes

A Monte-Carlo is conducted which randomizes the strike magnitude and direction and orbit parameters (inclination and true anomaly) such that the strike is applied randomly relative to the orbit frame of the satellite. For each strike each test statistic is calculated, and the test statistics which detect the strike are recorded along with the estimated size of the strike, if it is detected, as shown in Figure 3.18. Note that false alarms are not characterized in this analysis, but missed detections are shown along the x-axis. The colors of the plot indicate the number of test statistics which successfully identified the strike on each axis, noting that the RSS axis also includes the scalar test statistics: Mahalanobis distance and change in orbital constants. This Monte-Carlo uses

the τ and q identified in the trade study in the prior section.

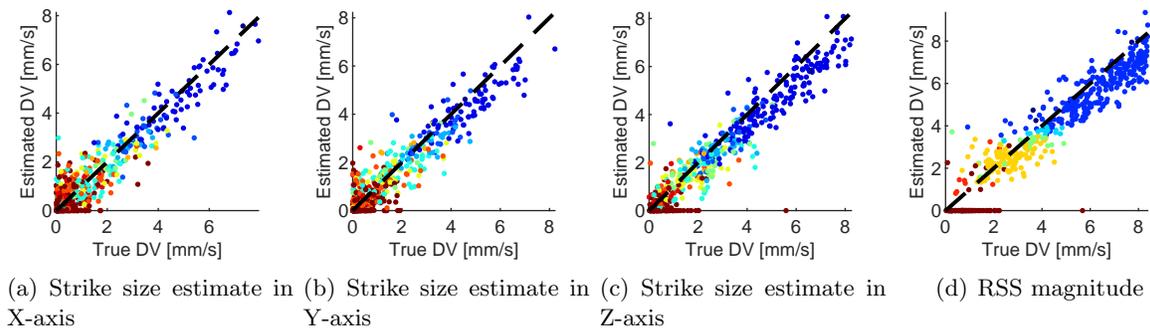


Figure 3.18: Monte-Carlo analysis indicating estimated size of random strike relative to true strike. Color indicates number of test statistics that detected each strike.

While the detectability of some strikes seems improbably small, similar real-world capabilities have been demonstrated by Exoanalytic Solutions, who identified an unexpected in-track ΔV of approximately 0.7 mm/s.[33]. Also, characterizing the noise baseline using a kernel distribution provides a decent proxy for the detectability threshold, but with a fairly short data set to develop the no-strike noise distribution false alarms may still occur. Conducting trade studies with long datasets requires excessively long computation times without adding much to this content, so shorter runs are used since the intention of this chapter is to explore the problem using simulated data. Specific noise levels will change, likely dramatically, with each individual on-orbit system.

3.4.4 Performance in the Presence of Orbit Perturbations

The primary challenge of applying these techniques to real-world systems will likely be separating debris strike perturbations from other orbital perturbations and telemetry idiosyncrasies. In prior work, algorithms have been applied to detect translational momentum changes in the attitude telemetry of NASA spacecraft and found abundant unexpected features that complicate detection of debris strikes.[42]

While modeling the unknown unknown is always a challenge a quick proxy assessment is performed by adding J3 perturbations and drag perturbations to the simulated truth state, but not

to the filter dynamics. To accentuate the effects of these perturbations the orbit is lowered to 400 km (from the default 800 km) and the eccentricity raised from 0 to .001 (the lower orbit is necessary to show the effects of drag on the short time scales seen here). To produce a signature higher than the noise in the presence of these perturbations, the debris strike magnitude is increased by 2X.

The acceleration due to drag is added to the orbit state per Vallado Ch. 8,[73]. The parameters used to calculate drag are specified in the Section 3.2.2. Note that the solar array is rotated relative to the velocity throughout the orbit, as if it were tracking the sun, which causes the drag applied to the spacecraft to vary as it would on orbit, with ϕ and θ determined such that the array always points in the inertial X direction regardless of the satellite location. However, the density of the atmosphere only varies with altitude, without seeing the variability throughout the orbit that would be experienced in reality. All vectors here are expressed in the inertial frame.

$$\mathbf{a}_{\text{drag}} = -\frac{1}{2} \frac{C_D A_{\text{tot}}}{m} \rho v_{\text{rel}}^2 \frac{\mathbf{v}_{\text{rel}}}{v_{\text{rel}}} \quad (3.64)$$

$$A_{\text{tot}} = A_{\text{bus}} + A_{\text{array}} \cos(\phi) \cos(\theta) \quad (3.65)$$

$$\mathbf{v}_{\text{rel}} = \mathbf{v} - \begin{bmatrix} 0 \\ 0 \\ \omega_e \end{bmatrix} \times \mathbf{r} \quad (3.66)$$

$$\rho = \rho_o e^{-\left(\frac{r-r_o}{H}\right)} \quad (3.67)$$

The following acceleration due to J_3 perturbations is added to the orbit state, per Equation 11.65 in Schaub and Junkins.[54].

$$\mathbf{a}_{J_3} = \frac{1}{2} J_3 \frac{\mu}{r^2} \left(\frac{r_{\text{eq}}}{r} \right)^3 \begin{pmatrix} 5 \left(7 \left(\frac{z}{r} \right)^3 - 3 \left(\frac{z}{r} \right) \right) \frac{x}{r} \\ 5 \left(7 \left(\frac{z}{r} \right)^3 - 3 \left(\frac{z}{r} \right) \right) \frac{y}{r} \\ 3 \left(1 - 10 \left(\frac{z}{r} \right)^2 + \frac{35}{3} \left(\frac{z}{r} \right)^4 \right) \end{pmatrix} \quad (3.68)$$

While J_2 is neglected it is also straightforward to capture J_2 in Kalman filter dynamics. Capturing an estimate of J_3 and drag is only a little less straightforward, so this represents the efforts of a rather mediocre filter, and shows that a small strike is still detectable even in the presence of unrealistically severe unmodeled orbit perturbations.

The results in Figure 3.19 show that adding these unmodeled dynamics causes additional noise in the test statistics, as expected. The filter's estimate of the unmodeled acceleration successfully estimates these unmodeled perturbations so that the filter stays converged, but this produces additional noise in the filter output and increases the noise floor of the filter, obfuscating the effect of the debris strike. However, a somewhat larger strike is still detectable, which is still well into the size regime of hazardous nontrackable debris. It is worth noting that these $J3$ and drag perturbations are currently completely unmodeled in filter dynamics, but these and more are accounted for in precision orbit determination filters, so the performance of real-world systems could be better than what is shown here.

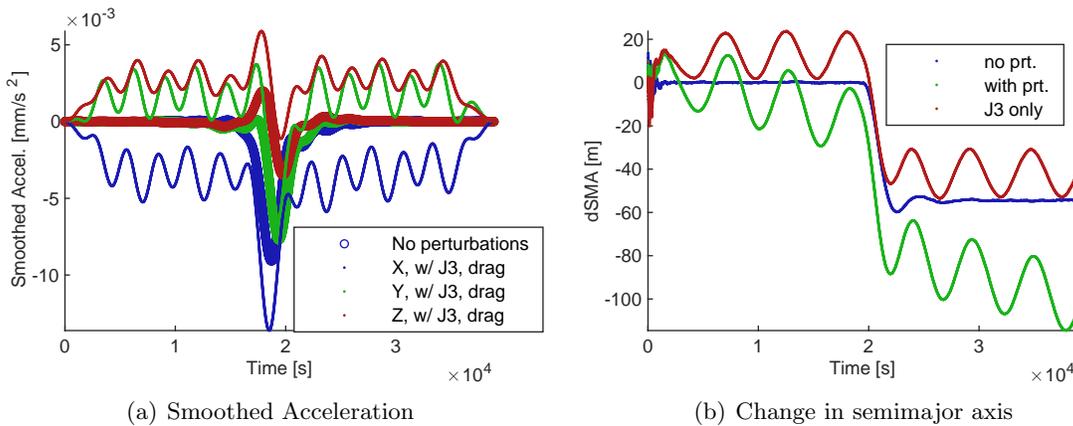


Figure 3.19: Simple perturbation models ($J3$ and drag) added to truth state but not filter dynamics show that baseline filter noise is higher, but strikes are still detectable. This plot shows a 400 km orbit and a debris strike that is 2X the magnitude of the default strike

Figure 3.20 repeats the trade from Figure 3.17, but with the perturbations added to the state and the orbit lowered to 550 km and debris strike size increased by 5X. Comparing to the previous trade, it can be seen that the best performance occurs with a somewhat larger q , as a larger q allows the filter to track the perturbed state more accurately. With a well converged filter state the larger debris strike still causes a detectable signal in the presence of noise due to $J3$ and drag. Figure 3.21 shows the trade study results with the perturbations.

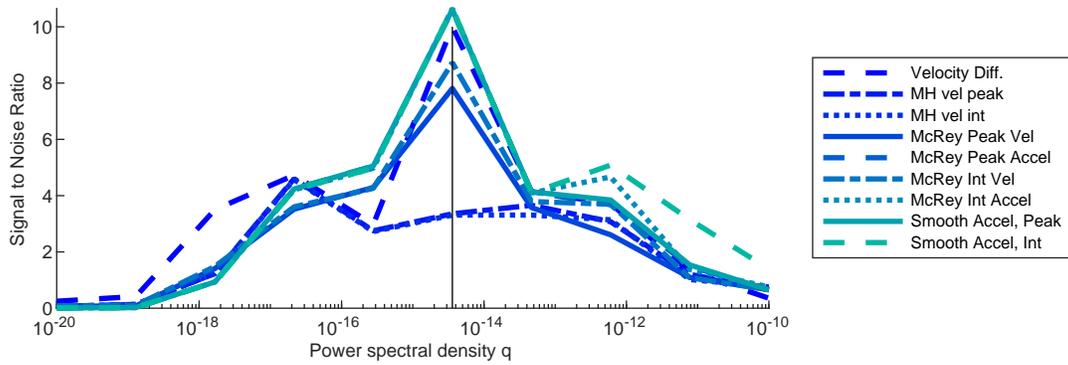


Figure 3.20: Superimposing SNR from various test statistics to compare results for given time constant τ . Note that preferred q is now larger than the trade without perturbations, and the SNR is smaller even with a larger strike

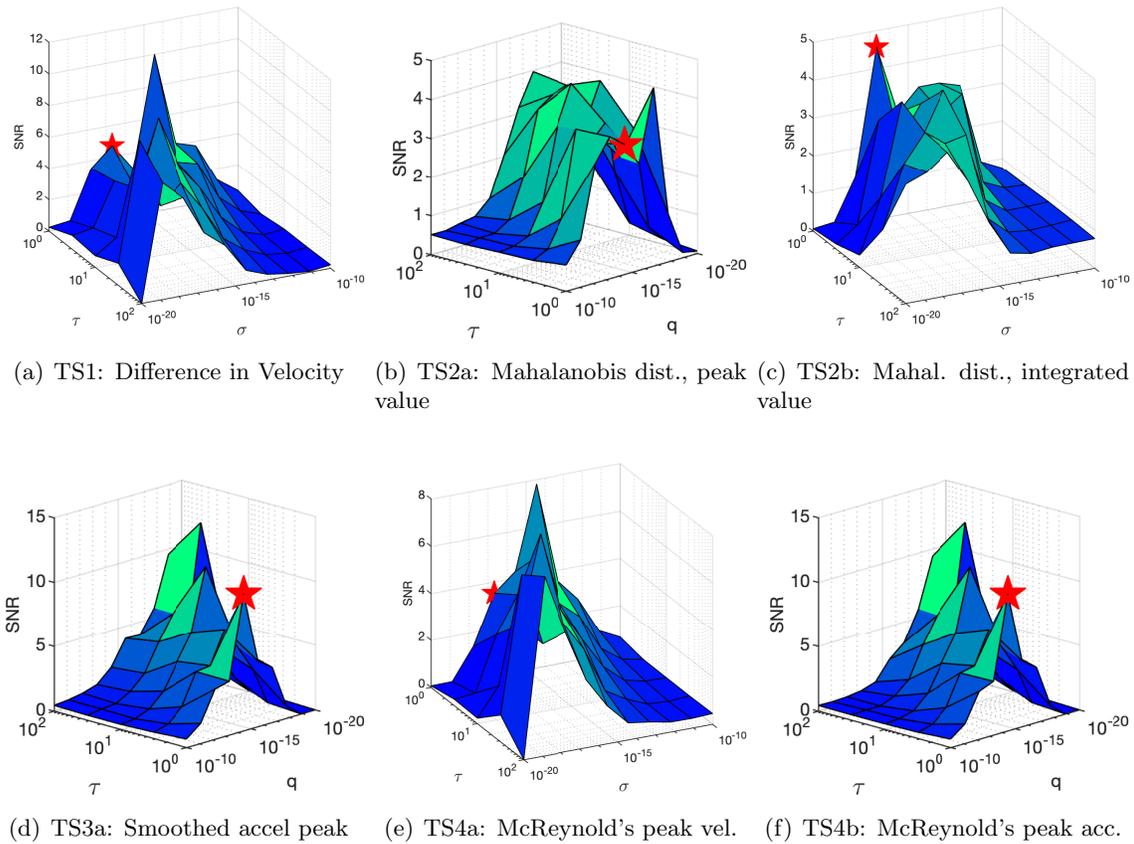


Figure 3.21: Trading τ and q again, but with additional unmodeled perturbations ($J3$ and drag)

3.4.5 A final comparison between measurement residuals and test statistic output

As a final observation, Figure 3.22 shows the pattern of measurement residuals from the forward EKF with no test statistics or post-processing. Patterns of filter residuals are typically used to assess filter performance and screen for unexpected features. These residuals indicate that standard practices for observing filter performance are much less sensitive to debris strike features than the test statistics developed herein. The output from the smoothed acceleration test statistic is compared to the measurement residuals for two debris strikes: the small one is 1/10th the default strike specified in Section 3.2.2, the large one is 2X the default strike. This shows that a large strike begins to show up in the measurement residuals of a typical filter, but the test statistic produces a much stronger signal for detecting subtle strikes.

3.5 Discussion

The magnitude of detectable strikes is highly dependent on system parameters. These numbers are only placeholders and do not prove a capability; real system data must be used to determine detectable ΔV thresholds. Therefore, while these metrics are promising in terms of being able to detect minor impacts from hazardous non-trackable debris the quality and usefulness of data obtained will likely be system-dependent, and not all systems may have the precision orbit determination required to separate abrupt unexpected changes produced by debris strikes from routine changes due to gravity effects, drag fluctuations, etc.

While implementing these methods in the presence of real-world orbit perturbations will be a challenge, tools exist for providing precision orbit determination for some missions. These tools can be leveraged when implementing these methods on real-world data, so this development focuses on methods rather than on simulating and reinventing precision orbit determination. It is more efficient to apply these methods to real data and characterize the results rather than exhaustively modeling orbit perturbations, even though they are critical to detection thresholds. Sensor noise characteristics present a similar complication, where real-world sensor noise and irregularities can be

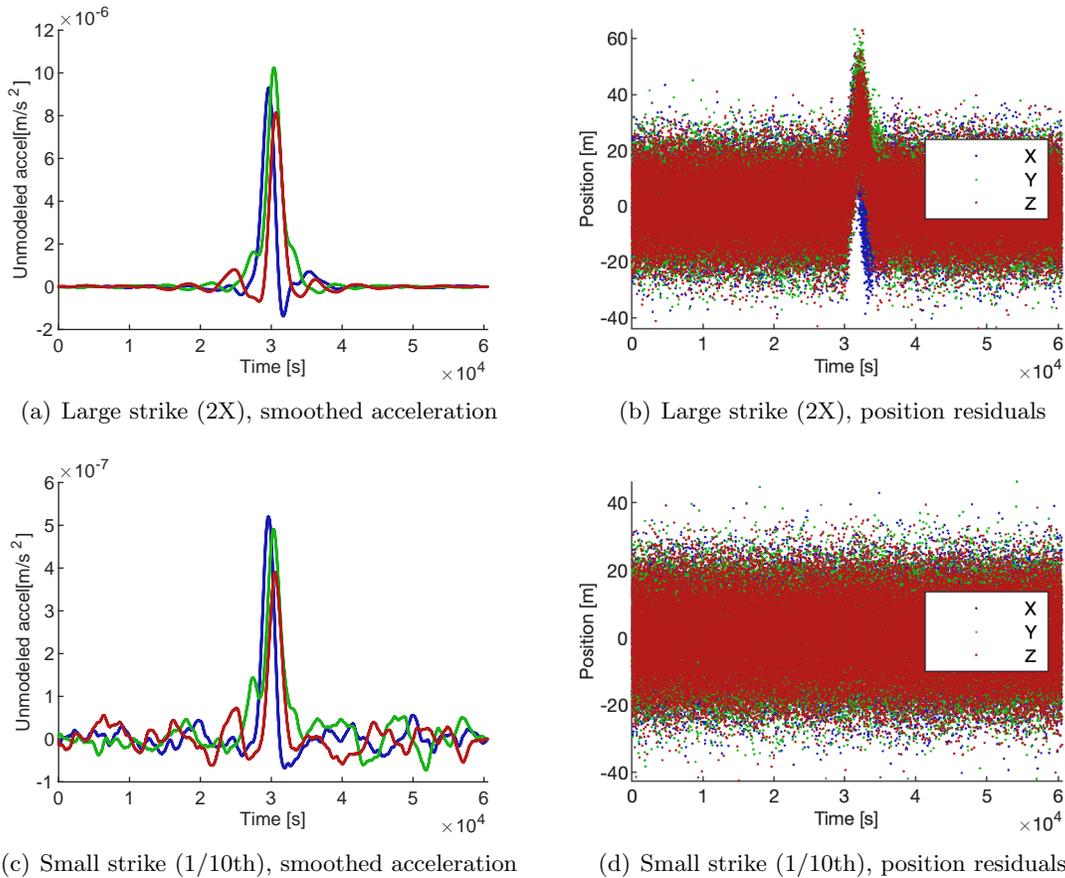


Figure 3.22: Comparing test statistic output to traditional measurement residuals in the presence of large and small debris strikes. Test statistic shows strong signal in response to strike far too small to see in measurement residuals.

very challenging to model accurately but are a critical element of accurate prediction of detectable thresholds for these and similar techniques.

While the exact magnitude of detectable orbit changes is questionable on-orbit experience has proven that the capability is possible. Several spacecraft [17, 2] have detected orbit changes in conjunction with minor anomalous events and thereby detected debris strikes, so it seems likely that improving these orbit change detection capabilities, as illustrated in this chapter, will result in an improved ability to detect minor debris impacts in the typical course of operations.

While this chapter leverages simulated GPS telemetry these methods could conceivably be applied to non-cooperative and uncontrolled objects to obtain a larger population of *in situ* mea-

surements. The detection thresholds would likely suffer, but with LEOLabs and other companies providing increasingly precise orbit determination services it is possible that the increasing precision required for space traffic management in the emerging space era will allow methods like this to be applied broadly to populations of objects on orbit. For example, ExoAnalytic Solutions has already demonstrated an impressive capability to detect similar events on GEO satellites, and maintains exquisite and persistent coverage of objects in GEO. Combining capabilities and datasets from various providers could likely generate a rich dataset for validating and tuning debris environment models.

The heritage method of detecting debris strikes via a change in satellite mean altitude produces a strong test statistic, but it is primarily sensitive to debris strikes in the \hat{v} direction. With current LEO populations this may be justified, as models predict that most debris strikes will occur in this direction. However, it has limitations as radial or out-of-plane components of strikes are not observable. Breakups occurring in GTO or Molniya orbits are already difficult to characterize, so using a debris strike detection system that can only detect the expected debris strikes introduces a strong confirmation bias within the measurement system, which should make any scientist cringe. To maintain a clearer picture of the debris environment it is important to have direction-agnostic and orbit-agnostic debris sensing systems, and this chapter outlines one method that could be used by satellites to detect impacts and characterize the small debris environment.

Note that this filter is developed using GPS telemetry, since this research is focused on typical telemetry, but the same techniques could be applied in filters that use alternate forms of navigation data. Ground based sensors may have to deal with measurement gaps, but ExoAnalytic's work proves that a sufficiently sophisticated algorithm can detect events of this nature. LEO is a harder task since the perturbations are more significant, especially drag which varies, but it might not be out of the question particularly for larger strikes. Given the large number of objects tracked, useful strike rates might be detectable even if the detection thresholds were significantly higher. The perturbation assessment tool developed in the third research segment is designed to answer questions of this nature, determining how many detectable strikes might be expected with various

systems and sensor capabilities.

One intriguing possibility is the concept of using crosslinked constellations that can incorporate crosslink data into the navigation filter. The GRACE mission measures gravity perturbations with exceptional accuracy using a sophisticated inter-satellite ranging system, demonstrating the concept that exquisite measurements of inter-satellite ranging can attain extremely precise measurements of orbit perturbations. With increased use of intersatellite crosslinks in emerging constellations, incorporating this data may give the orbital accuracy needed to see small debris strikes as well as other perturbation sources, like drag, which could also use improvements to models for future space traffic management needs.

3.6 Conclusions

An extended Kalman filter with dynamic model compensation can be adapted to accentuate abrupt orbit changes that are too small to show a feature in the pattern of filter residuals typically used for filter performance assessment. Various test statistics derived from the filter output provide strong signals to allow detection of subtle changes indicative of minor debris strikes. The strongest of these test statistics are derived from the smoothed unmodeled acceleration, obtained when the smoother fuses estimates from separate forward and backward filters. This is similar to the Fraser-Potter smoother, and a necessary adaptation is ensuring the backward filter employs a mirrored model for the unmodeled acceleration behavior such that filter performance is symmetric when filtered forward and backward in time. These methods offer a capability for spacecraft operators to monitor the state of health of their spacecraft and, with a sufficient population of data, could provide *in situ* data for tuning debris environmental models and risk assessment processes.

Chapter 4

Applying Techniques to Telemetry from Active Satellites

“All happy spacecraft are alike; every unhappy spacecraft is unhappy in its own unique way.”

-Dr. Russell Carpenter, paraphrasing Leo Tolstoy's *Anna Karenina*

4.1 Applying Techniques to Various Spacecraft

The change detection algorithm developed in Chapter 2.4.2 is applied to several NASA spacecraft. Applying the change detection algorithm to active spacecraft proves to be a non-trivial challenge. While the inertial angular momentum of a system is theoretically quiescent in the absence of an external torque, due to conservation of momentum, measuring the reality is much more complicated. The main challenge turns out to be obtaining a sufficiently quiescent estimate of the inertial angular momentum from the telemetry and parameters which are available for a given spacecraft. As opposed to simulated telemetry, on-orbit telemetry is full of idiosyncrasies which are often unique to an individual spacecraft. These idiosyncrasies frequently affect the calculation of the inertial angular momentum, producing perturbations which trip the detection threshold and result in false alarms.

Some of these idiosyncrasies can be compensated for or removed from the data, but many are difficult to model effectively or identify definitively. The challenge lies in cleaning the data sufficiently to apply the algorithms with a high degree of confidence that meaningful results are being returned. This finding is consistent with comments by other operators who have tried techniques

like this: implementing methods like this is more difficult than it may seem, doubt remains about the accuracy of the results, and there are features where it is difficult to definitively assess whether a strike has occurred or whether it is just noise.

The following sections describe the processes developed to apply these techniques to several NASA spacecraft, incorporating the techniques that seemed promising and identifying the limitations that were discovered as well as the lessons learned. It is important to note that in reality this process of development was nowhere near as linear as what is described here. Many things were tried, many did not succeed, and many solutions which seem clear in retrospect were much less clear during the development. The NASA spacecraft investigated and the justification for why each spacecraft is of interest are as follows. A variety of spacecraft are considered to explore how the algorithms function on a range of spacecraft types. Types include 3-axis stabilized, spinners, flexible appendages, articulated appendages, LEO, GEO, HEO, non-Earth orbiting, and various mission profiles, with some spacecraft slewing frequently and others with quieter attitude requirements.

- **Magnetospheric Multiscale Mission (MMS):** MMS consists of four spacecraft in a highly elliptical orbit. Multiple MMS spacecraft have detected impacts via their highly sensitive accelerometer telemetry, so they are examined to determine if these techniques can identify known impacts and/or see other impact events.
- **Mars Atmosphere and Volatile Evolution (MAVEN):** The MAVEN spacecraft, in orbit around Mars, studies the Martian atmosphere. It is selected for analysis because one of its instruments is sensitive to the plasma effects of dust impacts. Identifying supplementary indications of an impact at the same time as a perturbation would provide compelling evidence that subtle perturbations are actually particle impacts.
- **Lunar Reconnaissance Orbiter (LRO):** LRO is a spacecraft in lunar orbit which observes the lunar surface. It is selected for analysis because on-board instruments have detected anomalous events possibly indicative of micrometeoroid impacts.
- **Fermi Gamma-ray Space Telescope:** Fermi is a space-observing spacecraft in low-Earth

orbit. It is selected for analysis because it flies in one of the more debris-rich environments among NASA SSMO spacecraft which have telemetry available for analysis.

- **Solar Dynamics Observatory (SDO):** SDO resides in GEO and observes the sun. It is selected for analysis because it is in GEO, which represents an undermeasured debris environment, and because it is a fairly quiescent spacecraft.

4.1.1 Telemetry Processing Architecture

The bulk of telemetry processing was conducted using NASA Space Science Mission Operation's (SSMO) telemetry database, referred to as TaaS (Telemetry-as-a-Service). This cloud-based database stores telemetry from several SSMO missions, including those listed above. Telemetry can be requested via cURL commands, which for this effort were implemented via a Matlab architecture which can request telemetry from TaaS and then operate on that telemetry and store the results. The key goals for the telemetry processing architecture are as follows:

- Process years of telemetry efficiently
- Investigate multiple spacecraft each with unique telemetry set
- Run different detection algorithms on various telemetry points. i.e, detect that a change in angular momentum has occurred, but ignore its effects because the thrusters fired at that time (for example, a when a wheel desaturation is conducted)
- Be able to modify algorithms, processing techniques, thresholds, window sizes, spacecraft, telemetry sets, etc. rapidly as techniques are developed and real-world considerations are accounted for.

The architecture to accomplish these goals was developed in two generations. The first generation was developed in summer 2019, and used to conduct an initial study on several spacecraft. This architecture was updated in summer 2022 to conform with SSMO best practices, which had changed during the pandemic, while addressing some speed issues and revisiting some tradeoffs

that were made with the first generation. Table 4.1 describes the differences between the first and second generation and the reason for the changes.

Table 4.1: Changes from initial to second generation telemetry processing architecture

Item Changed	First Generation	Second Generation	Reason for Change
Use of datastores, tall arrays	Create tall array from datastore, operate on tall array	Read data directly from datastore with in-memory chunks	Tall array operations excessively slow relative to in-memory operations
Data storage method	cURL raw data to csv on cloud or local, then process csv files to apply algorithms, store processed data on cloud or local	Use custom datastore which reads cloud data directly to Matlab for processing, can optionally store raw and/or processed data on local, data no longer stored in two places on cloud	Bring into accordance with new SSMO best practices, remediate inefficient use of cloud storage (\$\$\$) with first generation
Degree of generality	Prioritize general architecture to rapidly retool between spacecraft	Prioritize efficiency to investigate and respond to peculiarities of current spacecraft	Bulk of work involves dealing with spacecraft idiosyncracies, so customization is more efficient than forcing a one-size-fits-all solution since specific requirements are redefined with every new spacecraft

Figure 4.1 shows the initial architecture. The TaaS telemetry database is sent to files on the S3 cloud, and those files are used to generate a MATLAB datastore. A datastore is a collection of data which may be too big to fit in memory, but creating the datastore object allows the information in the datastore to be processed. In the Gen 1 architecture this processing was accomplished via ‘tall’ arrays. These tall arrays (or tall tables, for various datatypes) are arrays too big to fit in memory, but Matlab has a capability to operate on them and return results via the ‘transform’ function (for operating on a single row) or the ‘movingWindow’ function (for operating on a sliding window of data). In theory, MATLAB will process the tall arrays in the fewest passes practical to apply the functions called via transform or movingWindow, but in reality the tall arrays get

processed many times, and the tall array functions appear to have a large amount of processing overhead time. This makes them much less efficient than originally envisioned. The other issue with this setup is that the data is stored multiple places: in the TaaS telemetry database, in separate cloud-based folders (the reading and writing of which take time), and sometimes also in local folders. In the interim between the 2019 visiting technologist experience and the 2022 visiting technologist experience the SSMO team reconfigured their telemetry processing practices, introducing a custom Matlab datastore type that reads directly from TaaS to prevent having the same data stored in multiple places on the cloud.

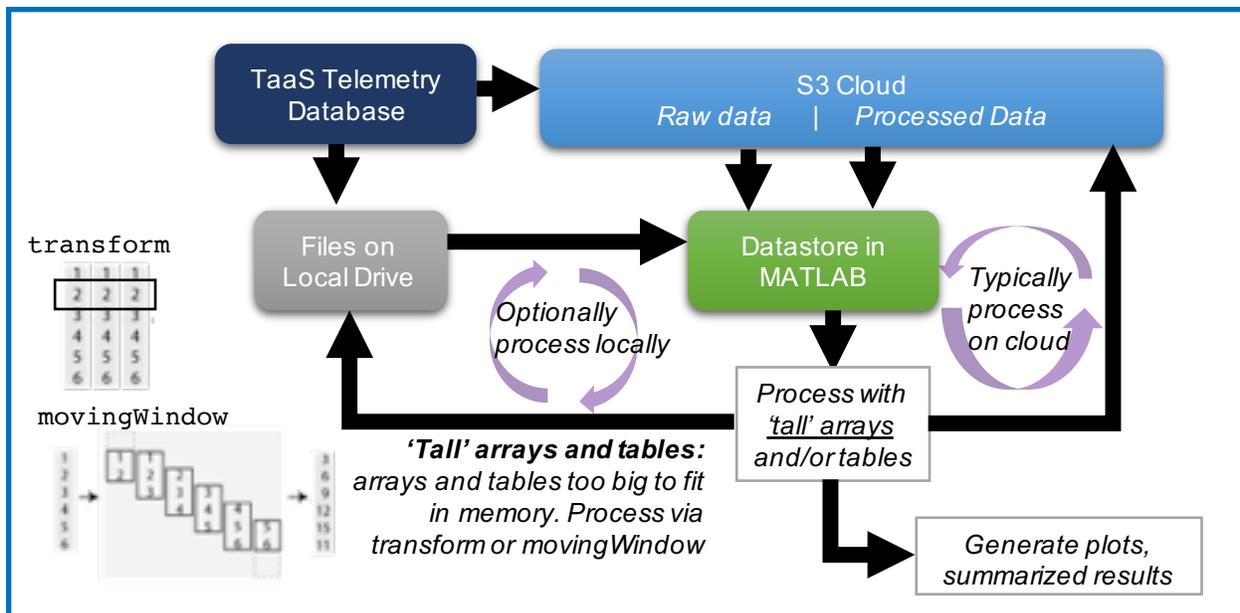


Figure 4.1: Original telemetry processing architecture

This architecture proved fairly effective, and between 2019 and 2022 data from three different systems besides TaaS were ingested, one being the MAVEN spacecraft data from the web-based Planetary Data System. Retooling the architecture between spacecraft was fairly straightforward, as intended, although reformatting the new data to mesh with the architecture sometimes took time. However, a tradeoff was determined: each new spacecraft presents unanticipated challenges in terms of the overall processing methods that need to be applied, various timespans of data that need to be dealt with differently, and especially various plots that need to be made as the processing

methods are developed. A variety of plots are necessary to observe the interim results and move toward the desired performance. While a general architecture sounds nice in theory, in practice it needed constant updates to accommodate each new spacecraft's idiosyncrasies, and of course these updates need to still be backward compatible with previously processed spacecraft. With the 2022 VTE needing some updates to conform with SSMO best practices, a slight retooling of the architecture was performed to allow more customization for individual spacecraft with fewer adjustments to the architecture, which slowed the process of adding a new spacecraft slightly but dramatically increased the efficiency of exploring and developing methods to deal with unique spacecraft idiosyncrasies. Also, since tall arrays had proven slower than expected, some speed tests were conducted to determine the difference in processing time between continuing with tall arrays or reading data into memory before processing, and reading into memory turns out to be the faster option. The resultant second generation architecture is shown in Figure 4.2

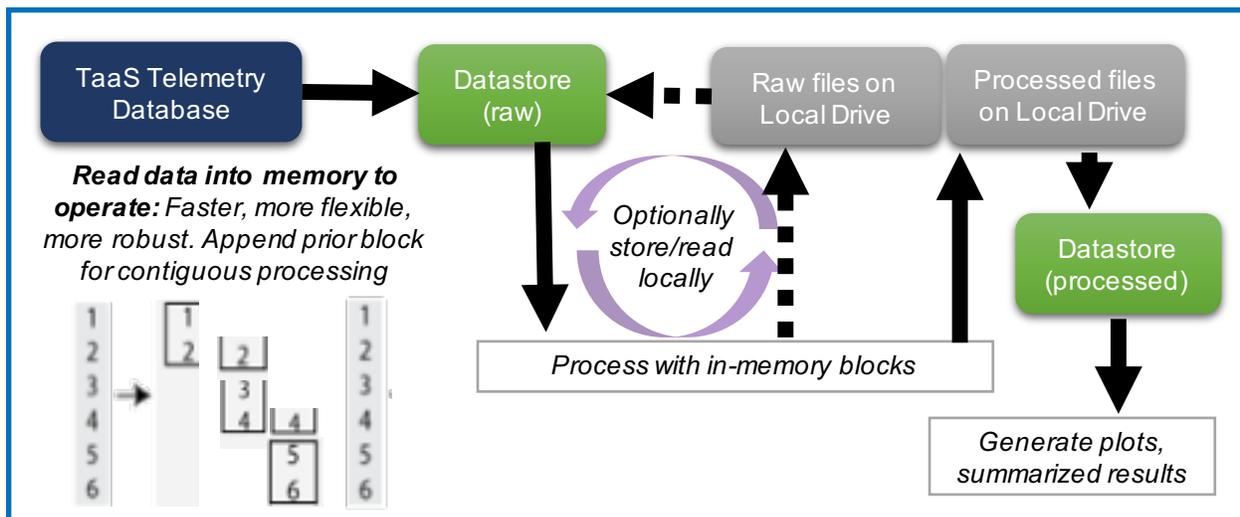


Figure 4.2: Second generation telemetry processing architecture

As shown, the cloud storage is eliminated entirely, TaaS reads directly into a datastore instead of routing through intermediary files, but the ability to process telemetry locally is maintained. This is important as iterating on processing methods is much faster with local files than reading the same data from the cloud again and again as new methods are tested. Tall arrays are also eliminated

entirely since performance overall was lackluster, which has the primary advantage of being faster with the added nontrivial advantage of making debugging easier, as values can be spot checked without processing the entire tall array. At first reading into memory was avoided because many processing techniques require a sliding window of data, and so losing data at the junction between blocks is undesirable. In the new architecture this is avoided by saving a buffer from the previous block and appending it to the current block so the junctions between the blocks are processed smoothly, with no datapoints lost.

While not shown in this high-level view, some of the generalized functions (plotting, results concatenation) were also eliminated. This makes it take a little longer to start up on a new spacecraft, but overall efficiency is improved as the general architecture doesn't need to be retooled and have backward compatibility checked every time a new spacecraft needs a new type of data processing, different plotting, etc. Many functions are still generic and used for all spacecraft, but the plotting functions are now spacecraft-specific. The efficiencies gained in exploring the features of and developing solutions for a new spacecraft easily overwhelm the additional time cost of adding a new spacecraft, and the code is simpler and more clear even though it's less general.

To put the speed considerations in context, consider the amount of data that needs to be processed. For SDO, the data rate is 0.25 Hz, so the 10 years of data processed to obtain Figure 4.33 mean that over 500 million datapoints are downloaded, processed, and filtered. The results from that filter are summarized into the figure, and the big events are checked individually to see which look like debris strikes and which may have another cause (data gap, component idiosyncrasy, etc). For a spacecraft with a faster data rate, like LRO or Fermi, which also require several more telemetry types to obtain a quiescent approximation of the inertial angular momentum, the amount of datapoints would be in the tens of billions. This is a solvable problem, but speed considerations for large-scale data are an important element of applying these techniques.

Figure 4.3 shows a high-level description of the main functions. A top-level function, generic between spacecraft, is used to specify the run parameters. This runs a second general function which reads in the data from either TaaS or local storage (per run parameters set in top-level

function), saves the raw data to local if specified, runs the functions that operate on the telemetry, and saves the processed results. The functions that operate on the telemetry had a high degree of commonality in Gen 1, in Gen 2 more flexibility in this function is allowed and the sub-functions to do the individual operations are generally spacecraft specific, although some (like applying the CUSUM filter) are the same for all spacecraft. Once the processed telemetry has been stored, a separate function (which also used to be general but is now spacecraft-specific) plots the results (or a subsection of the results) for examination, this is used extensively in development on small datasets as the processing methods are developed. After the processing methods are developed larger datasets are run, and these results are processed in a second function, also spacecraft specific, which can pull out the desired summarized results. The following sections discuss the specific steps taken for each spacecraft, the results seen, and the ensuing adventures. As mentioned earlier, many approaches besides these were tried, these sections summarize the final and most relevant status of each effort.

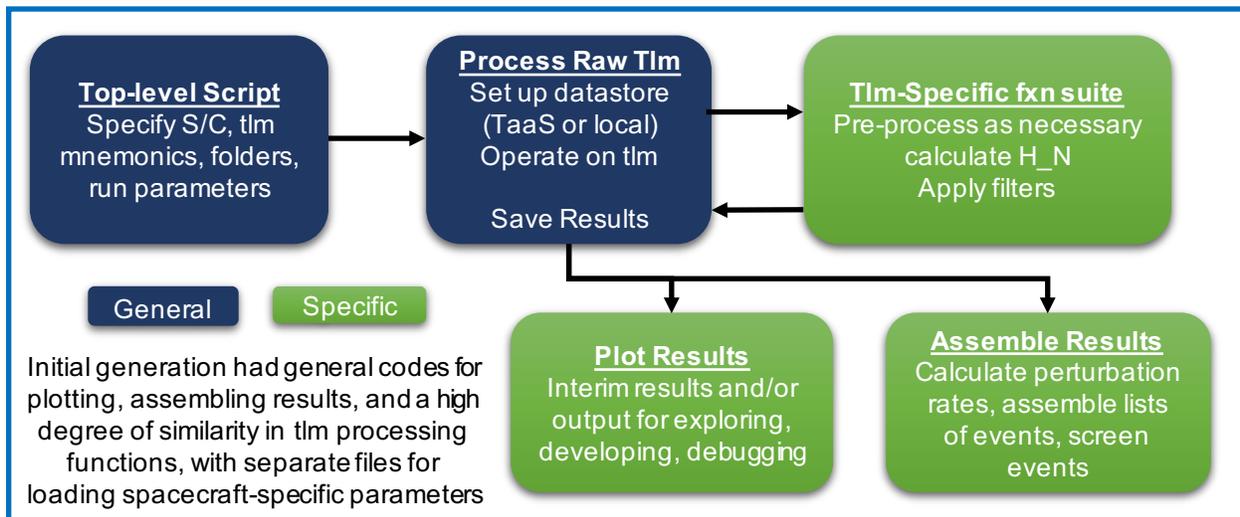


Figure 4.3: Diagram of major code functions, indicating which are general between all spacecraft and which are developed specifically for each spacecraft.

4.1.2 Magnetospheric Multiscale Mission (MMS)

4.1.2.1 Processing MMS Telemetry

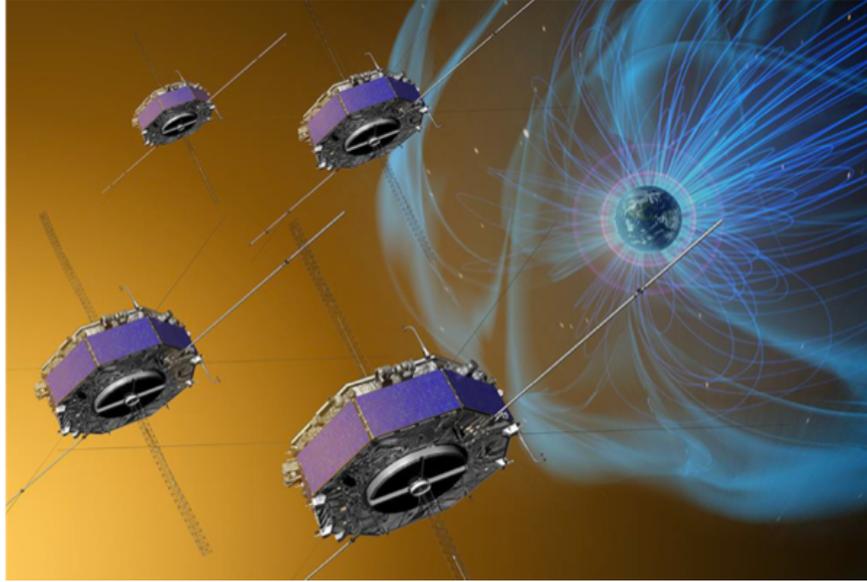


Figure 4.4: MMS spacecraft formation studies electromagnetic reconnection

The MMS mission consists of four spacecraft flying in formation in a highly elliptical orbit (Figure 4.4). These spacecraft are spin-stabilized and have 60 m long wire booms and other appendages. They have thrusters for constellation resizing and other maneuvers, but they do not have reaction wheels. Startrackers provide quaternions and rates are derived from attitude data. With this telemetry the inertial angular momentum (${}^{\mathcal{N}}\mathbf{H}$) of each MMS spacecraft is calculated via ${}^{\mathcal{N}}\mathbf{H} = [NB][I_{sc}]{}^{\mathcal{B}}\boldsymbol{\omega}$, where $[NB]$ is the direction cosine matrix to convert from the body frame to the inertial frame, $[I_{sc}]$ is the inertia of the spacecraft, and ${}^{\mathcal{B}}\boldsymbol{\omega}$ is the angular rate of the spacecraft expressed in body-frame coordinates.

The MMS spacecraft's appendages cause unique behaviors that manifest as MMOD strike detections. Whenever the spacecraft goes through perigee the inertial angular momentum experiences a change, as shown in Figure 4.5(a). Whenever the thrusters fire these appendages begin to oscillate, and it takes several days for those oscillations to damp out such that the observed rigid-body momentum calculation no longer trips the debris strike detection thresholds. Figure 4.5

illustrates both of these effects. To mitigate them, the telemetry is blanked during perigee passages, and for 10 days after each thruster firing.

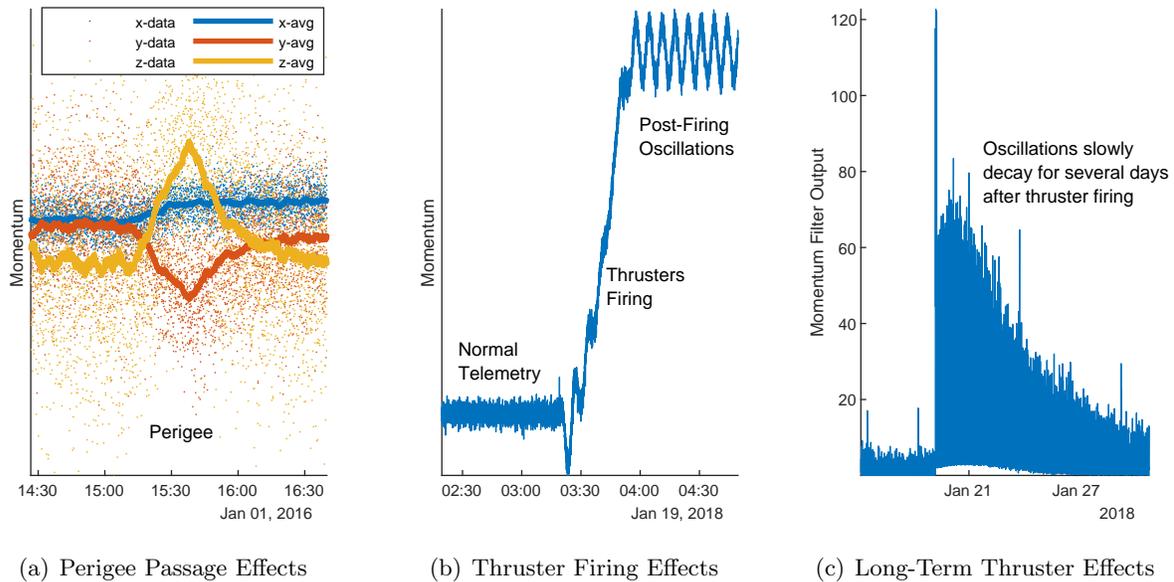


Figure 4.5: Effects of mission events on MMS inertial angular momentum

One unique feature of MMS is its highly sensitive accelerometers. In two events, a minor loss of spacecraft capability (i.e, loss of a shunt resistor) has been coupled with a dynamic event evidenced by the accelerometer telemetry and a change in spacecraft nutation, and the damage is attributed to an MMOD strike. In other instances, an oscillation in accelerometer telemetry is observed and attributed as a minor strike, but no damage is evident.

To set MMS thresholds the perigee effects are removed to establish a quiescent dataset. The thresholds are set as discussed in Section 2.4.2.2, although each axis has a different threshold as this spin-stabilized spacecraft exhibits a different noise level in each inertial axis. A CUSUM sequential probability ratio test is applied to the inertial angular momentum in a sliding window of 500 datapoints, with a change in mean of $\Delta\mu = 0.05$ Nms between the null hypothesis and the alternative hypothesis (indicating that a debris strike has occurred as indicated by a change in the distribution underlying the inertial angular momentum). A second filter was applied to look for

an increase in the variance of the accelerometer readings, and a third filter looks for plasma probe readings for variations seen in known strikes, but the results of these filters were noisy and did not show a strong correlation to the results from the momentum filters, so they were investigated but not used extensively. Orbit telemetry (to obtain altitude) and thruster telemetry are used to identify periods of telemetry where the filter results need to be blanked as the thresholds are tripped by known non-debris-strike events.

4.1.2.2 MMS Results

MMS has some limitations because of its long appendages and lack of 3-axis stabilization. The known debris strikes manifest as increased noise due to wire boom oscillation causing changes in the apparent inertial momentum. Specifically, the momentum calculation assumes a rigid body, so an oscillating appendage causes the calculation of $\mathcal{N}\mathbf{H}$ to oscillate. However, this is not a direct measurement of the angular momentum imparted by the strike, and thus doesn't really include information on the magnitude or direction of the strike. The noise levels are fairly high compared to spacecraft that will be discussed in follow-on sections.

The MMS-2 spacecraft and MMS-4 spacecraft are each evaluated. For MMS-4 300 days from 2016 are examined and for MMS-2 all of 2018 is evaluated. These years are selected because there is a known debris strike in each year which can be used to validate the debris detection algorithms. Due to the effects discussed in Section 3.2, significant portions of the MMS telemetry must be blanked to avoid false alarms. In these results 10 days are blanked after each thruster firing and all perigee passages are blanked.

The MMS-4 spacecraft's results in 2016 are as follows. Figure 4.6(a) shows a plot of the results from all 300 days analyzed. The output of the CUSUM filter is plotted whenever any of the three axes are above the detection threshold. The grey bars are times when the output was blanked due to mission events, either perigee passages or thruster firings. The known debris strike in February causes a clear blip, shown to scale in Figure 4.6(a) and zoomed in for Figure 4.6(b). A distinct signature appears in the data and persists for some time because the strike excites

appendage oscillation causing the observed inertial momentum to oscillate. It takes a couple days for those oscillations to damp out. The raw momentum from the strike is shown in Figure 4.6(c).

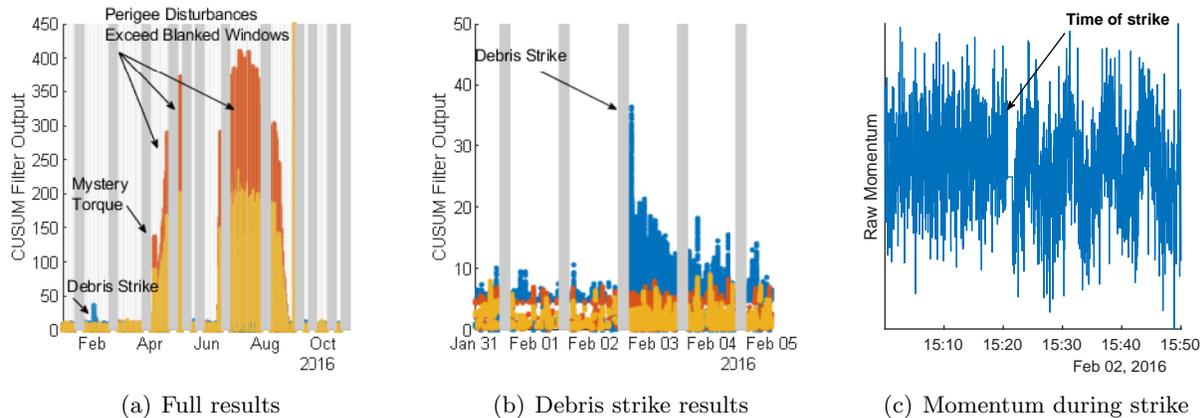


Figure 4.6: Results from MMS-4 telemetry in 2016

However, this debris strike is overshadowed by large signatures resulting from perigee disturbances where the disturbance extends beyond the blanked window. In April through September the typical perigee disturbances shift so that they are outside the blanked window. This is surprising, because at first the perigee disturbances were assumed to be gravity gradient torques. However, the blanked windows are set based on spacecraft altitude, and the effect of gravity gradient torques would not be expected to show significant variation with respect to altitude, especially with variations occurring from one season of the year to the next. To blank all of these events, given their inconsistent locations, would require substantially longer windows. This would result in the loss of a large percentage of the data. An example of these perigee disturbances exceeding the blanked window is shown in Figure 4.8(a).

There are also occasional significant features of unknown origin, one from April 6th is shown in Figure 4.8(b). In this feature the inertial momentum of the spacecraft increases for about an hour then decreases for about an hour, implying that a mild torque is applied to the spacecraft during this time period. This occurs while the spacecraft is near apogee, at an altitude of 76,000 km. The thruster firing telemetry does not indicate any firing events, so the origin of the torque

is puzzling. One hypothesis is that MMS's long 60 m wire booms are charged, and this charged, spinning spacecraft is interacting with Earth's changing electromagnetic field; as MMS flies through these electromagnetically active regions a torque is exerted on the spacecraft. MMS's apogee is planned to study electromagnetic reconnection, and flies through electromagnetically active regions, shown in Figure 4.7. Its perigee would experience a stronger and more consistent magnetic field closer to Earth, but still potentially subject to changes due to the influence of solar storms.

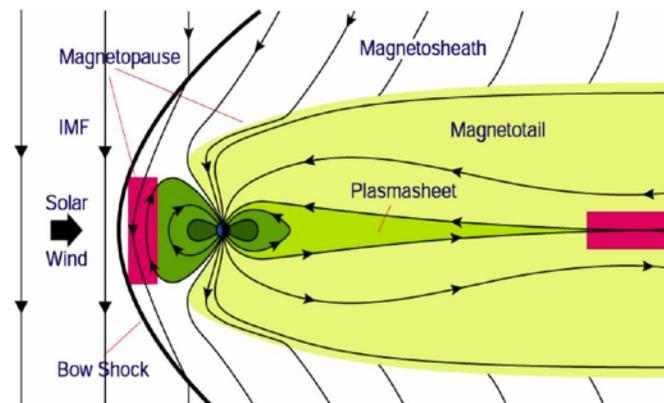


Figure 4.7: Flight regions of MMS

Figure 4.8(c) shows the momentum telemetry during this event and also during a normal perigee passage. As shown, this event causes a slight dip and then return to typical values, while the perigee passage produces a more distinct feature and a net change in inertial momentum. While these features may be scientifically interesting their inconsistency makes strike detection more difficult, as the changing inertial angular momentum trips debris strike detection thresholds, and large portions of the data must be blanked to accommodate this effect.

The MMS-2 results in 2018 are similar to the MMS-4 results but somewhat cleaner, with only a few days of perigee effects that extend past the blanked window. Figure 4.9 shows some example filter output, with a known debris strike observable but overwhelmed by the wire boom oscillations caused by thruster firing. The known debris strike occurs at a time when the telemetry was blanked due to a previous thruster firing, so it would have been missed without *a priori* knowledge of the strike. Figure 4.9 shows the filter output decaying over time as the oscillations from the previous

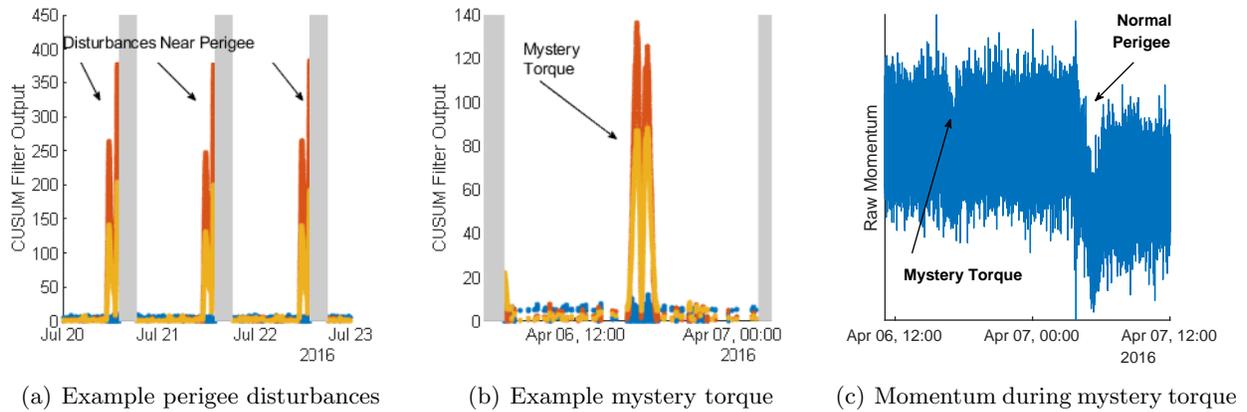


Figure 4.8: Results from MMS-4 telemetry in 2016

thruster firing damp out, then an increase in oscillations from the debris strike, then a huge spike followed by decaying oscillations from another thruster firing.

The summarized results are shown in Figure 4.10. These results are not particularly compelling. The magnitude of the ‘detections’ is very consistent, and also high. This may indicate that the filter is just detecting noise, as strikes should exhibit a variety of magnitudes. This is also indicated by the relatively consistent number of events throughout the year, in contrast to the results described in future sections, which show fluctuations that are somewhat consistent with expected variations in micrometeoroid flux.

The summarized MMS-4 results exhibit significantly higher strike counts, around 40 per day. However, they do not exhibit clear spikes in filter output or patterns in strike count per day, so it is expected that this 40/day is just noise. This implies that the noise levels are not consistent from spacecraft to spacecraft

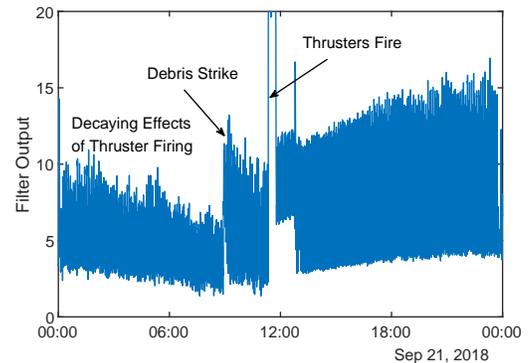


Figure 4.9: Example filter output from MMS-2 showing debris strike and other events

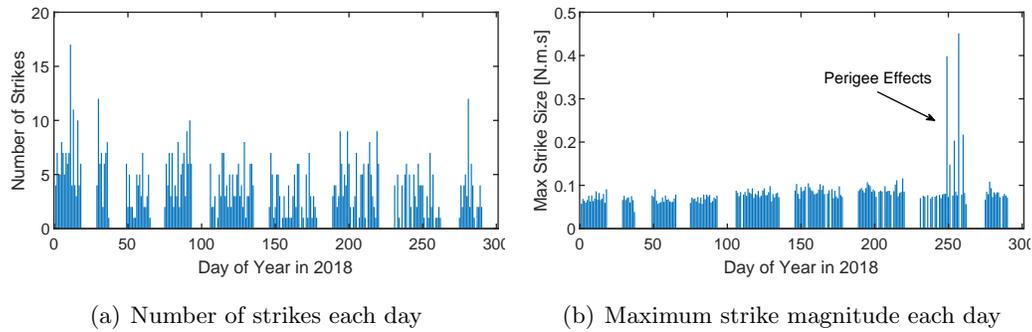


Figure 4.10: Results from MMS-2 telemetry in 2018

within the constellation, or perhaps from year to year.

While there are additional science studies that could be attempted or processing methods that could be applied to MMS, it is decided based on these results to focus on other spacecraft. For one, MMS is an extremely unique spacecraft, so methods developed to accommodate the idiosyncracies in MMS telemetry would not transfer more generically to other spacecraft. More importantly, since MMS appears to sense strikes based on wireboom oscillation, this may indicate that a strike has occurred but does not provide information on magnitude or direction. The detection threshold appears to be fairly high, so the utility of further analysis is questionable. For these reasons, MMS analysis is concluded and other spacecraft are investigated.

4.1.3 Mars Atmosphere and Volatile EvolutionN (MAVEN)

The MAVEN spacecraft is a Mars-orbiting spacecraft which investigates the Martian atmosphere (Figure 4.11). Its instrument suite includes a Langmuir Probe and Wave instrument (LPW), which measures the local electric field at a variety of data rates, some very high (65+kHz). Signatures in the LPW data have been observed and attributed to dust impacts, similar to other interplanetary spacecraft [26]. A lengthy investigation of the MAVEN spacecraft was conducted, since if changes in momentum are detected they could be checked against LPW data, if available, and this could be used to provide confidence that an impact has in fact occurred.

First, the housekeeping telemetry was used to try to generate a quiescent estimate of the



Figure 4.11: MAVEN spacecraft

spacecraft's inertial angular momentum for impact detection. MAVEN's telemetry and science data can be obtained from the Planetary Data System (PDS), making it a rare case where a spacecraft's data is publicly available. The PDS data is reformatted to interface with the telemetry processing architecture. Then the spacecraft dynamics are modeled.

Modeling the spacecraft dynamics requires knowledge of the spacecraft parameters (inertia, reaction wheel configuration, etc). These parameters were obtained from a design document provided by the manufacturer, but the properties provided appeared to be estimates from about a year prior to launch, listing both CBE and margin masses/inertias with significant variations between them. The reaction wheel inertias were the nominal estimate, not an as-built measurement reflected in the end item data packages, and the reaction wheel misalignment transformation matrices were placeholders, set to the identity matrix, not updated with as-built measurements. There were additional inconsistencies, but overall the best spacecraft parameters available for this analysis had fairly high error bars.

Unsurprisingly, based on this, obtaining a quiescent estimate of the spacecraft's inertial angular momentum was challenging. The MAVEN spacecraft has a dual-gimballed articulated payload

platform (APP), and any time the APP moves (as it does routinely) or the spacecraft slewed there were significant deviations in the calculated inertial angular momentum, due to the parameters for the calculation being inexact. The analysis and results are very similar to the Lunar Reconnaissance Orbiter analysis and results, as both analyses entail modeling the dynamics of dual-gimballed appendages, see Section 4.1.4 for the details of how this is implemented. To further complicate matters, MAVEN's reaction wheel telemetry is sporadically discontinuous, with the speed sometimes moving in discrete steps and other times behaving more smoothly. When the data is stepped the inertial angular momentum shows discontinuous lines, indicating that the true motion of the RWs may not be consistent with the stepped telemetry, but may instead interpolate between the steps. However, even if this issue is solved the inertial angular momentum sees massive changes when the spacecraft slews, indicating that the inertia estimate is far from precise, which is not too surprising given that the parameters received have pretty major variations between the current best estimate (CBE) and margined mass, indicating they were developed fairly early in the design process. With these issues, an inertial angular momentum estimate quiescent enough to apply change detection algorithms and generate meaningful results was not obtained.

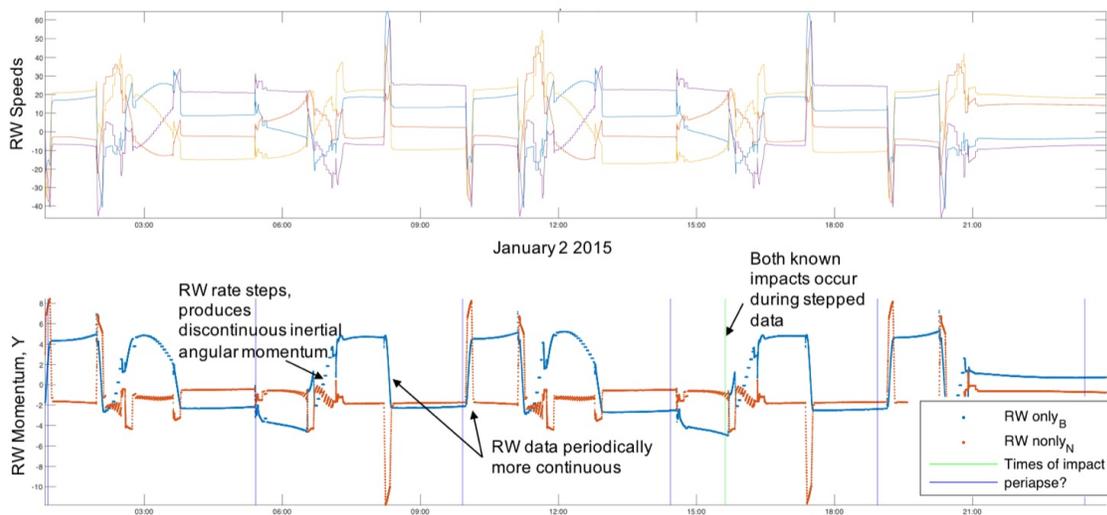


Figure 4.12: MAVEN's RW telemetry

Taking a different tack, the next effort was to look at the dust impact times and see if they occurred during a quiescent period for the spacecraft, allowing the identification of perturbations at the time of known strikes. Unfortunately, it was discovered that the method for identifying the dust impacts was tuned to the small signatures observed in the telemetry. Example dust impact signatures were obtained, and these sample signatures were run against the LPW data to identify times at which a comparable feature occurred by examining the correlation coefficients. Unfortunately, any list of strike events examined for reference [27] seems to have been lost in the years since the paper has been published. The research team provided the signatures which were used to filter the data, and they were applied to the raw LPW data using a cross-correlation filter to determine if known strikes could be identified. The results are shown in Figure 4.13.

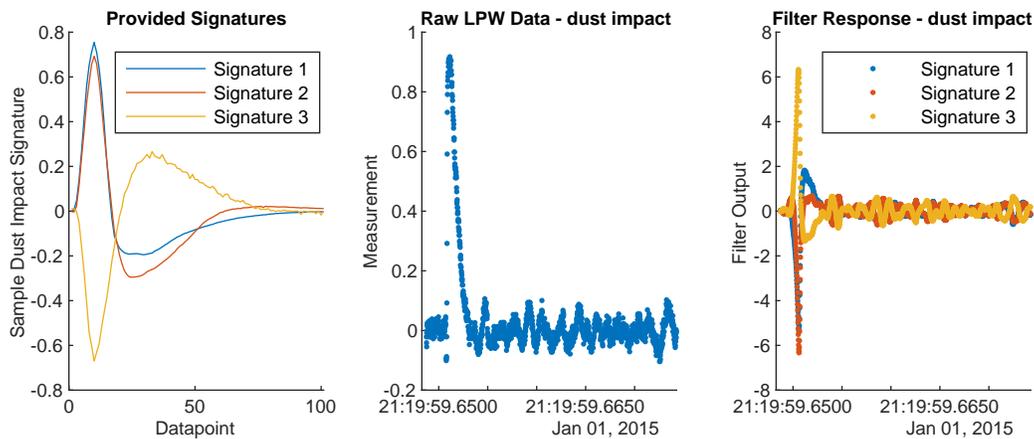


Figure 4.13: Running provided signatures against raw LPW data at time of known impact

Unfortunately, the noise in the data was substantially higher than the signature of the dust impacts, as shown in Figure 4.14. While the provided filters respond to the debris strikes, they also respond to the other noise. It may be that the research team had a second method of cleaning the data to mitigate the effects of these larger features. Also, in the context of this effort, it is not expected that every single dust impact can be identified as a detectable perturbation on the spacecraft. The perturbation would be higher with an impact far out on the solar array, which may not be observed by spacecraft instrumentation, or smaller impacts may not cause a detectable pertur-

bation. If a dust impact/micrometeoroid strike is large enough to cause a detectable perturbation, it may not even trip the detection filters since they are tuned to identify the small dust impacts, a larger impact might create a larger, differently-shaped signature which may be filtered out of the data along with all the other larger features in the LPW data. Figure 4.14 shows an example of a day of data from the LPW, and the filter output when the provided signals are cross-correlated with this data. The lower charts show the horizontal axis zoomed in to see the data in more detail, and the patterns of noise do not look much like sporadic dust impacts. Note especially the scale on the vertical axis in Figure 4.14 vs Figure 4.13: the cross-correlated example signatures return a much stronger response in the presence of noise than in the presence of a dust impact.

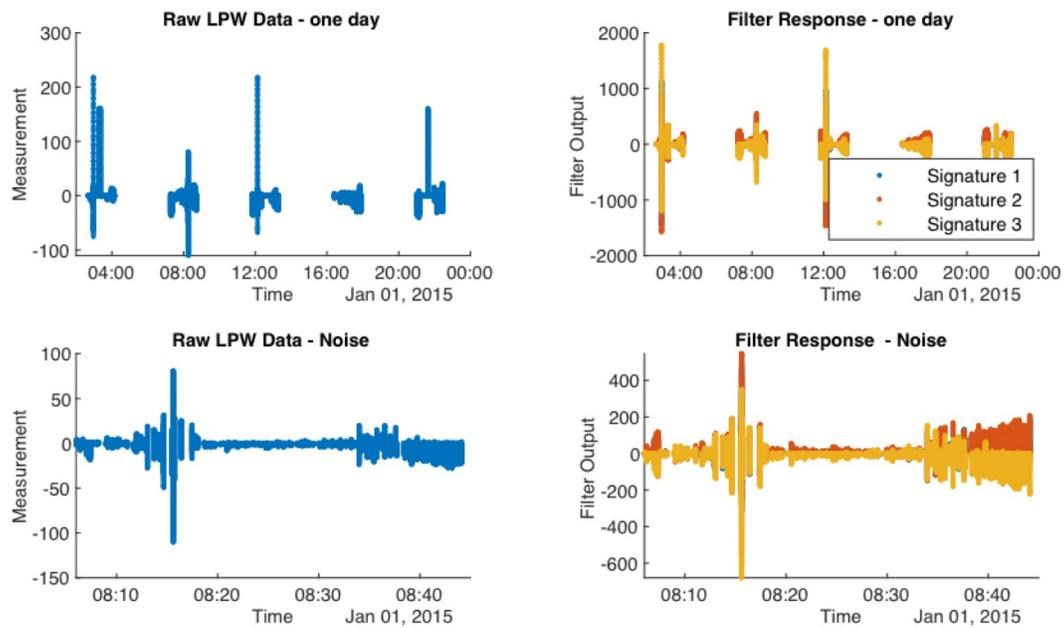


Figure 4.14: Zooming out to entire day of data shows major features besides impacts

Based on these limitations—difficulties in obtaining a sufficiently quiescent estimate of the spacecraft’s angular momentum, and lack of a high-confidence secondary detector that is sensitive to the strikes that are expected to cause detectable perturbations—the MAVEN analysis was discontinued. The next steps in this analysis would involve determining how to filter the LPW data to remove the non-impact noise and identify only impacts, without removing any larger impacts which

may cause a larger signature. An assessment of whether an impact out on the solar array would be measured by the LPW would probably also be advisable. Beyond this, a filter designed to estimate uncertain spacecraft dynamics parameters based on flight data would be useful in updating the as-designed parameters such that a quiescent estimate of the inertial angular momentum could be obtained. This recommendation is relevant to most of the spacecraft assessed, not just MAVEN, and is discussed further in Section 4.1.4.2.

4.1.4 Lunar Reconnaissance Orbiter (LRO)

LRO is a spacecraft in Lunar orbit which examines the Lunar surface (Figure 4.15). It is three-axis controlled using reaction wheels, and slews frequently to observe various targets. LRO's Narrow Angle Camera (NAC) has occasionally downloaded images where a dynamic effect happens while the image is being taken. One theory for the cause of these is a micrometeoroid impact, another is some sort of thermal flexing anomaly. Since the times of these image events are available, this presents an opportunity to see if a dynamic event is observable in spacecraft ACS telemetry at the time when the NAC observed a dynamic anomaly during image collection.

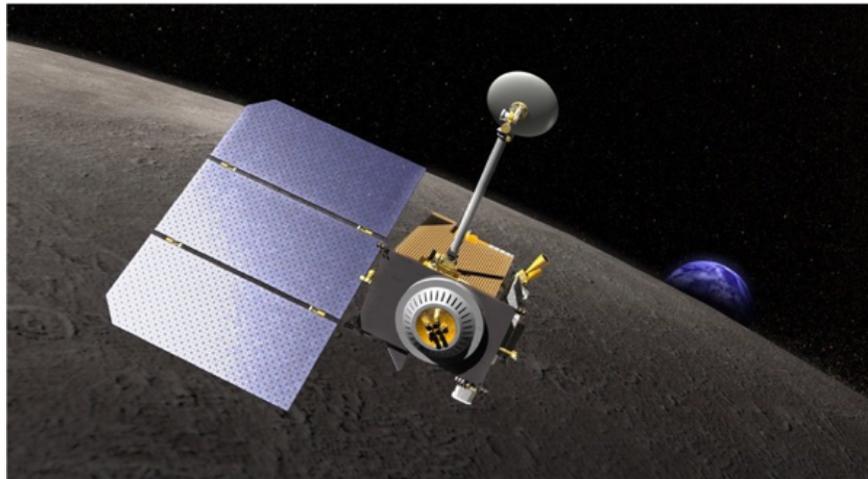


Figure 4.15: LRO spacecraft

4.1.4.1 LRO Dynamics Calculation

Like other three-axis controlled spacecraft with reaction wheels, the first task is to calculate the system's inertial angular momentum based on spacecraft rates, inertia, and reaction wheel speeds. LRO has two moving appendages, the solar array (SA) and high-gain antenna (HGA), each of which has two gimbals. During some seasons the solar array is fixed, during others it tracks the sun. The high gain antenna telemetry typically indicates movement. Since this is one of the more complicated spacecraft, the method for calculating the dynamics is outlined here, simplified versions are applied on the other spacecraft.

The total angular momentum of the system \mathbf{H}_{sys} is a combination of the angular momentum of the entire body, the angular momentum of the reaction wheels, and the angular momentum of moving appendages, in this case the HGA and SA. Thus

$$\mathbf{H}_{\text{sys}} = \mathbf{H}_{\text{sc}} + \mathbf{H}_{\text{rw}} + \mathbf{H}_{\text{hga}} + \mathbf{H}_{\text{sa}} \quad (4.1)$$

Each of these \mathbf{H} components is calculated in the body frame, since that is how most telemetry is specified, and then converted to the inertial frame. The inertial frame is more quiescent and generally better for comparison than the body frame which is often rotating. The direction cosine matrix (DCM) for the frame conversion is calculated from the quaternions provided in telemetry. The momentum contributions from the reaction wheels and spacecraft motion are calculated as follows, where $[\mathbf{I}_{\text{sys}/\text{CoM}}]$ denotes the inertia of the system about the center of mass.

$$\mathbf{H}_{\text{sc}} = [\mathbf{I}_{\text{sys}/\text{CoM}}]\boldsymbol{\omega}_{\text{sc}} \quad (4.2)$$

$$\mathbf{H}_{\text{rw}} = \sum_{i=1}^N [BW]_i \begin{bmatrix} I_{ws}\Omega_i \\ 0 \\ 0 \end{bmatrix} \quad (4.3)$$

$$(4.4)$$

Although instead of summing the reaction wheel contributions individually spacecraft typically provide a $3 \times n$ tensor with n reaction wheels, typically 4, to convert the individual wheels speeds

directly into the spacecraft body frame. The total inertia of the spacecraft is the inertia of the spacecraft ‘hub’, or main, non-articulated body, plus the inertia of any moving appendages at their current position. The contributions of these appendages are calculated as shown in Equations 4.6 and 4.7. The HGA is shown as an example because the inertia of both gimbal assemblies are specified for the HGA, making it the more complicated solution. I.e, the inner gimbal, closest to the spacecraft, has a mass and inertia that is specified in the spacecraft parameters separate from the mass and inertia of the outer gimbal plus HGA subassembly. For the solar array the mass and inertia of the inner gimbal subassembly is set to zero in the spacecraft parameters, so the analysis could be simplified, although the same process is followed. A picture of LRO’s HGA is shown in Figure 4.16 during range-of-motion tests, annotations are added to show the gimbal rotations.

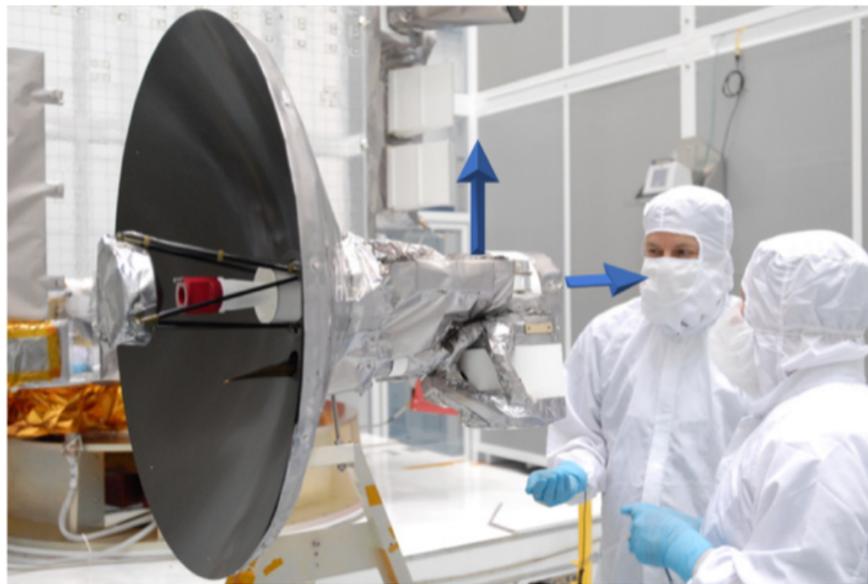


Figure 4.16: LRO HGA in integration and test - annotations show gimbal rotation

Equation 4.5 sums the various inertia contributions for the moving components. Here, ‘CoM’ denotes spacecraft center of mass, the center of mass specified in the spacecraft parameters is used. For LRO, the spacecraft parameters came from SSMO’s flight dynamics simulator known as the ‘HiFi’, which is used for analysis of flight data and tends to correspond well to on-orbit behavior. Note that the reaction wheels are included in the spacecraft hub inertia, because although they

spin their inertia contribution to the spacecraft does not change as they move.

$$\mathbf{I}_{\text{sc/CoM}} = \mathbf{I}_{\text{hub/CoM}} + \mathbf{I}_{\text{hga/CoM}} + \mathbf{I}_{\text{sa/CoM}} \quad (4.5)$$

The parallel axis theorem is used to calculate the contribution of each appendage to the spacecraft inertia based on its current position. The constant inertia of the appendage is specified in the appendage frame, so a frame conversion is used based on the current gimbal positions, using the gimbal positions to generate a current appendage-to-body DCM via Euler angle rotations. For the HGA, where the first gimbal assembly is massive enough to be treated separately, this is done twice to calculate the inertia contribution of the first gimbal assembly based on its current rotation and also the contribution of the HGA and second gimbal assembly based on both rotations. Here, the gimbal frame is denoted as \mathcal{A} with the frame origin at A. $[BA]$ denotes the DCM to convert from the \mathcal{A} frame to the body frame \mathcal{B} . The pre-super-script denotes the frame that the inertia matrix is expressed in, and a $\tilde{\mathbf{R}}$ denotes the skew-symmetric matrix of vector \mathbf{R} , in this case a position vector denoting the location of A relative to the spacecraft's center of mass. 'app' refers generically to a moving appendage, this process is used for MAVEN's APP as well as LRO's HGA and SA.

$${}^{\mathcal{B}}\mathbf{I}_{\text{app/A}} = [BA]{}^{\mathcal{A}}\mathbf{I}_{\text{app/A}}[BA]^T \quad (4.6)$$

$${}^{\mathcal{B}}\mathbf{I}_{\text{app/CoM}} = {}^{\mathcal{B}}\mathbf{I}_{\text{app/A}} + M_{\text{app}} {}^{\mathcal{B}}[\tilde{\mathbf{R}}_{\text{app/CoM}}]{}^{\mathcal{B}}[\tilde{\mathbf{R}}_{\text{app/CoM}}]^T \quad (4.7)$$

When considering the angular momentum contributions of the articulated appendages, first recall the angular momentum contribution of a reaction wheel. A reaction wheel's momentum contribution is the same regardless of where it is situated relative to the center of mass of the spacecraft, but this is because it is symmetric and rotating about its own center of mass. Figure 4.17 shows a diagram illustrating the notation used in Equations 4.8-4.11, which determine the reaction wheel's contribution to momentum of the total system.

For the reaction wheel, the momentum contribution of the two differential mass elements

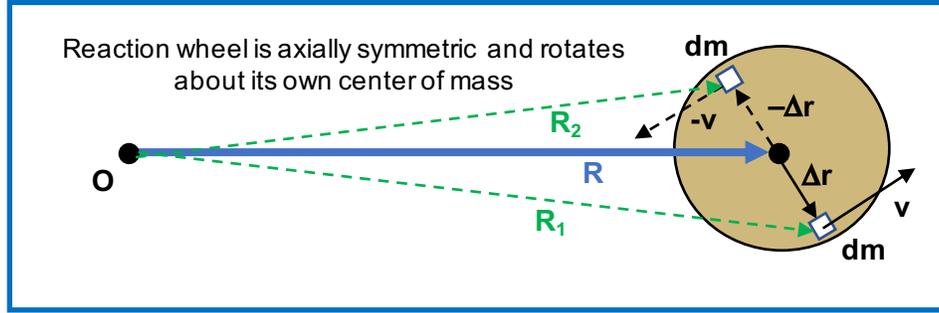


Figure 4.17: Reaction wheel rotating about its center of mass, for calculating angular momentum contribution about point O.

highlighted in Figure 4.17 to the total angular momentum about point O is:

$$\mathbf{H}_{2m/O} = \mathbf{R}_1 \times m\mathbf{v}_1 + \mathbf{R}_2 \times m\mathbf{v}_2 \quad (4.8)$$

$$\mathbf{H}_{2m/O} = (\mathbf{R} + \Delta \mathbf{r}) \times m\mathbf{v} + (\mathbf{R} - \Delta \mathbf{r}) \times m(-\mathbf{v}) \quad (4.9)$$

$$\mathbf{H}_{2m/O} = \mathbf{R} \times m\mathbf{v} + \Delta \mathbf{r} \times m\mathbf{v} + \mathbf{R} \times m(-\mathbf{v}) - \Delta \mathbf{r} \times m(-\mathbf{v}) \quad (4.10)$$

$$\mathbf{H}_{2m/O} = \Delta \mathbf{r} \times m\mathbf{v} + \Delta \mathbf{r} \times m\mathbf{v} \quad (4.11)$$

This shows that, since the terms with the large \mathbf{R} cancel, an axially symmetric object rotating about its center of mass will contribute the same amount of angular momentum to the system, regardless of where it is situated relative to the system's center of mass. If a solar array rotates about only about its center axis it might have an equal-and-opposite mass element, depending on construction, but LRO's dual gimballed array might not for some rotations, and the high gain antenna definitely does not, so the location of the object relative to the spacecraft's center of mass is relevant to the calculation of the total inertial angular momentum contribution of these appendages.

Figure 4.18 shows a simplified sketch of an appendage rotating around one axis that is not at its center of mass, and derives the associated momentum contribution from this rotation. In addition to the typical $[\mathbf{I}]\boldsymbol{\omega}$ inertia term, a second term is derived for the momentum contribution of an appendage rotating about a point displaced from its center. Note that this term is sometimes denoted $m\mathbf{R} \times \dot{\mathbf{R}}$ [81]. Deriving from first principles, per the notation specified in Figure 4.18, the angular momentum integrated over the rotating mass is calculated for each mass element i . In the

final step, the integral over each mass element m is approximated with the total point mass for the moving appendage M .

$$\mathbf{H}_{\text{app}/O} = \int_m \mathbf{R}_i \times m \mathbf{v}_i \quad (4.12)$$

$$\mathbf{H}_{\text{app}/O} = \int_m \mathbf{R}_i \times m(\boldsymbol{\omega} \times \mathbf{r}_{i/A}) \quad (4.13)$$

$$\mathbf{H}_{\text{app}/O} = \int_m (\mathbf{R}_{A/O} + \mathbf{r}_{i/A}) \times m(\boldsymbol{\omega} \times \mathbf{r}_{i/A}) \quad (4.14)$$

$$\mathbf{H}_{\text{app}/O} = \int_m \mathbf{R}_{A/O} \times m(\boldsymbol{\omega} \times \mathbf{r}_{i/A}) + \int_m \mathbf{r}_{i/A} \times m(\boldsymbol{\omega} \times \mathbf{r}_{i/A}) \quad (4.15)$$

$$\mathbf{H}_{\text{app}/O} = \mathbf{R}_{A/O} \times M(\boldsymbol{\omega} \times \mathbf{r}_{\text{appCoM}/A}) + \mathbf{I}_{\text{app}/A} \boldsymbol{\omega}_{\text{app}/A} \quad (4.16)$$

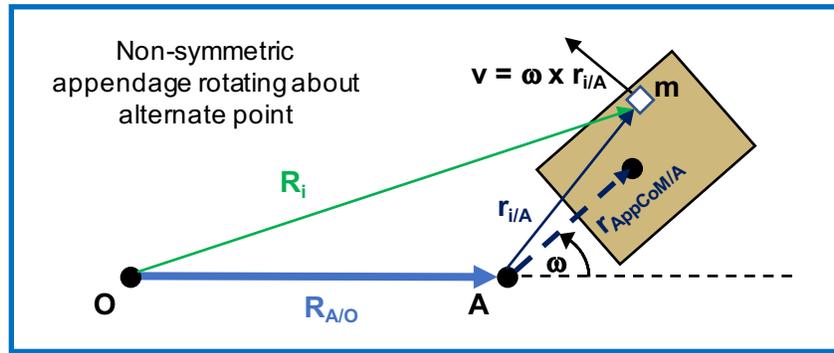


Figure 4.18: Angular momentum of appendage about point O due to rotation about point A

With this formulation, the total contribution of the motions of the HGA to the overall momentum of the spacecraft is calculated as shown in Equations 4.17 - 4.18. Note that the angular momentum contribution of each component is broken down into two pieces, reflecting the two terms in Equation 4.16: the total contribution is the inertia of the appendage times the rotation of the appendage, plus the point mass of the appendage and its angular motion relative to the point about which the momentum is calculated. The same calculation is used for LRO's SA, and for MAVEN's APP. In these equations, G1 indicates the gimbal 1 subassembly, which is the inner gimbal, with the rotation axis denoted A. G2 indicates the subassembly moved by the outer gimbal, which is the outer gimbal plus the HGA, with its rotation point denoted B. Note that the center of mass of each

subassembly also factors into the calculation of the momentum contribution. The frames in which the vectors are expressed are left unspecified in these equations for clarity, but of course converting each appendage's vectors and inertias into the correct frames for addition is a critical element of this calculation. Figure 4.19 identifies the notation used in this calculation with a simplified 2D case.

$$\mathbf{H}_{G1/O} = \mathbf{R}_{A/O} \times M_1(\boldsymbol{\omega}_1 \times \mathbf{r}_{G1CoM/A}) + \mathbf{I}_{G1/A}\boldsymbol{\omega}_1 \quad (4.17)$$

$$\mathbf{H}_{G2/O} = \mathbf{R}_{A/O} \times M_2(\boldsymbol{\omega}_1 \times \mathbf{r}_{G2CoM/A} + \boldsymbol{\omega}_2 \times \mathbf{r}_{G2CoM/B}) + \mathbf{I}_{G2/B}\boldsymbol{\omega}_2 + \mathbf{I}_{G2/A}\boldsymbol{\omega}_1 \quad (4.18)$$

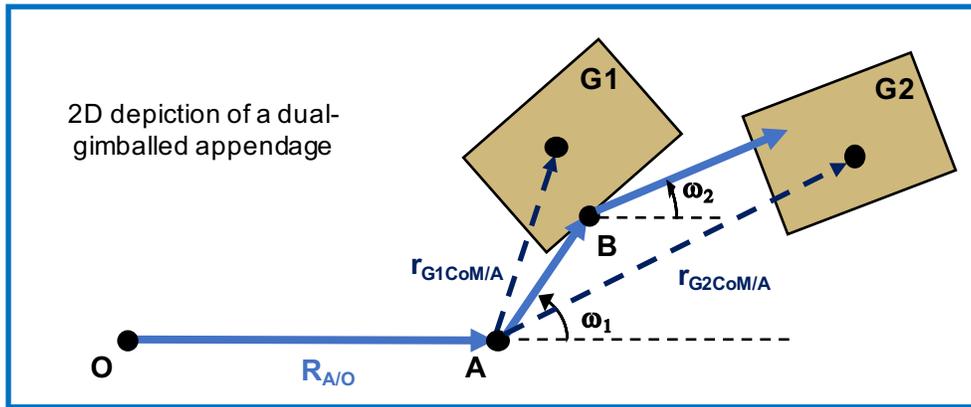


Figure 4.19: Schematic for calculating angular momentum of dual gimbal assembly with two massive elements in motion

This method is used to calculate \mathbf{H}_{hga} and \mathbf{H}_{sa} for Equation 4.1. Note that a typical spacecraft with a single-gimbal solar array would have a similar but simpler analysis to account for the motions of that appendage. A spacecraft with flexible appendages like MMS has significantly more complex dynamics which are harder to predict accurately, the effects of these are seen in the MMS analysis.

4.1.4.2 LRO Results

To begin, a dataset is selected that has limited solar array motions, and the momentum contribution of the high gain antenna is calculated as well as the spacecraft and reaction wheel

momentum, and these are compared to the on-board calculation of the system momentum. Everything is converted to the inertial frame, where results can be compared with less of an effect due to the spacecraft's changing attitude throughout the orbit. Figure 4.20 shows LRO's version of a set of plots which is typically used to explore the telemetry and develop processing methods. The top rows are the momentum contributions of each element of the momentum. The system inertial angular momentum is at the top, the goal is to make this quiescent enough to apply the change detection filter. Below it are the contributions from each subsystem: 'Hscrate' is the momentum of the rotating spacecraft body, note that this takes into account the current position of the HGA and SA when calculating the inertia. Then comes the momentum contribution of the reaction wheels, then the momentum contribution of the HGA. The SA contribution is omitted from this plot as it is not moving in this dataset. The colors represent each chunk of data plotted from the datastore, in this case the data was processed in 30 minute increments although larger datasets are typically processed in 6 hour increments - increments larger than 6 hr overwhelm memory, since LRO has a high (0.2 Hz) data rate.

Below these the relevant raw telemetry is plotted: attitude, rates, RW speeds, gimbal position and rate, etc. The axes are linked between all the plots so that if a feature is examined in one plot (i.e, zoomed in) then all the plots adjust to show the same time span. This makes it easier to see if a particular feature is due to a particular component's activities. With the data displayed this way, it is easy to see that oddities in the reaction wheel momentum are driving quite a lot of noise in the inertial angular momentum.

As shown in Figure 4.20, there are significant and frequent features in the inertial momentum, indicative that the measurement of the inertial momentum is not accurately capturing the features of the spacecraft. With the telemetry plotted this way, the culprit is obvious: the momentum contribution of the reaction wheels contains significant noise. However, since this is not reflected in a mirrored momentum perturbation measured by the spacecraft rate, this implies that it is a measurement idiosyncrasy, not a true oscillation.

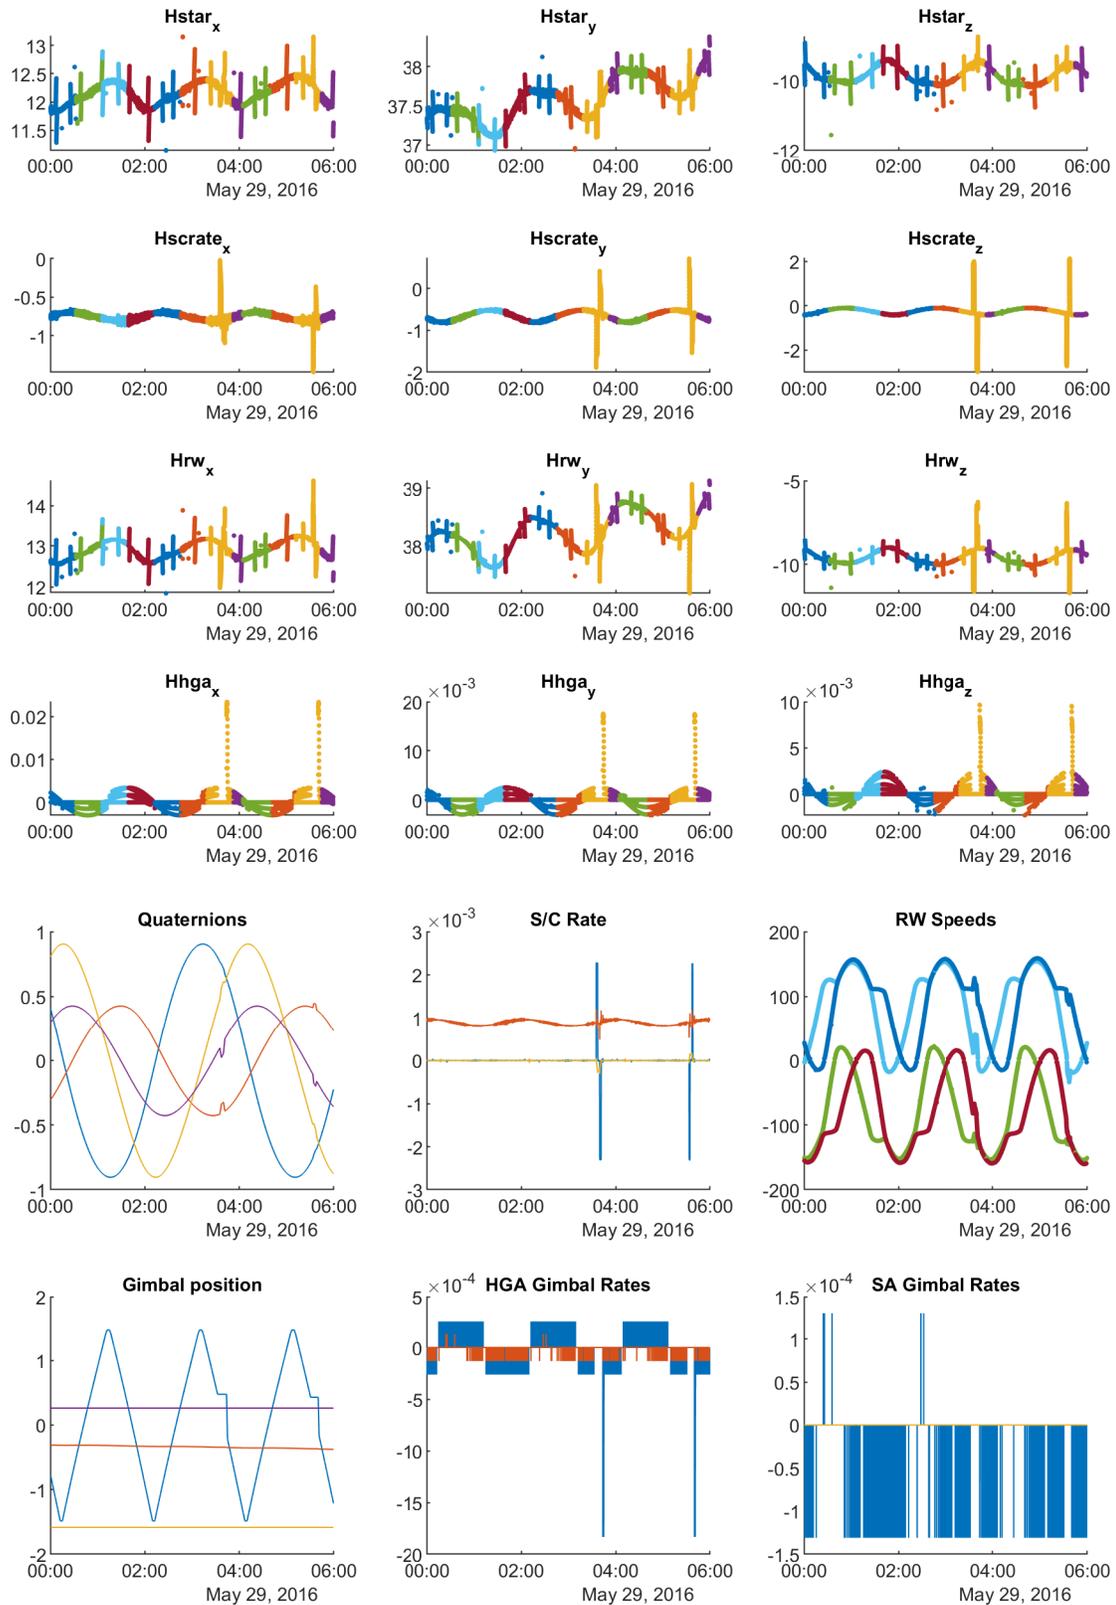


Figure 4.20: Momentum contributions and raw telemetry for LRO - before pre-processing

Examining the spacecraft telemetry and chatting with SMEs, it turns out that the spacecraft holds the reaction wheel speed telemetry during zero crossings (instead of letting it go through zero), which causes a discrepancy in the calculation of the inertial angular momentum because the spacecraft's actual motion differs from the calculated motion with the held value. A fairly straightforward fix is introduced to accommodate this, where the last few measurements before and after each value is held are used to fit a line between the data, and the reaction wheel speed is assumed to follow that linear fit. Figure 4.1.4.2 shows the telemetry values held as the wheels pass through zero, and one spurious measurement which effects the reaction wheel and therefore system inertia calculations. Figure 4.1.4.2 shows the wheel momentum telemetry before and after the RW speeds are corrected. As shown, the spurious measurement and zero crossing are fixed, but there is still a substantial fluctuating amount of noise in the data, which stems from noise in the reaction wheel speed measurements. This feature shows up routinely, its cause has not been identified.

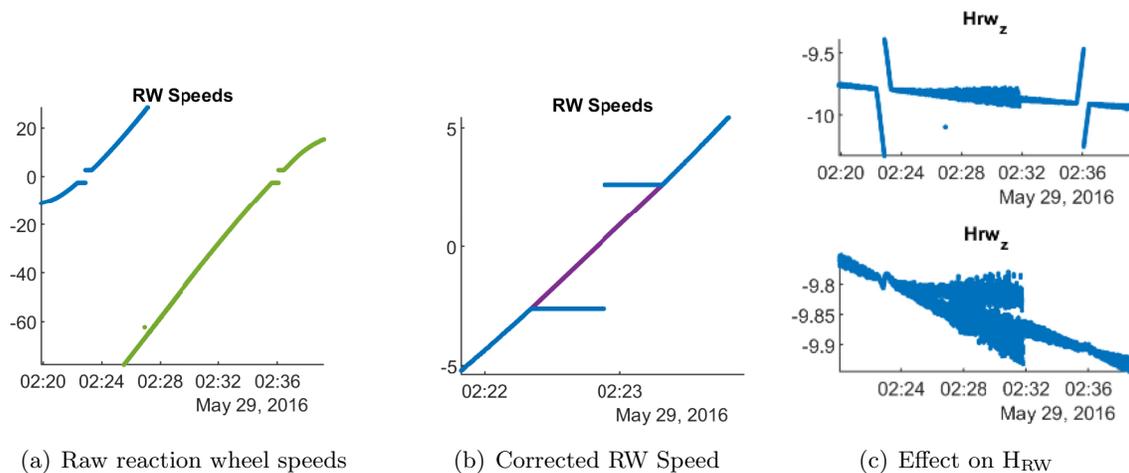


Figure 4.21: Patching LRO wheel speed telemetry during zero crossings

This algorithm performs well at estimating the true wheel speed during zero crossings and correcting spurious measurements. In Figure 4.22, the algorithm shows that it can approximate the telemetry even when slews occur during the zero crossings. This slew is shown around datapoint 2E5, where the reaction wheels spin up to absorb more angular momentum, rotating the spacecraft presumably to image a specified target, then pause to hold at the new attitude, then dump the

momentum into the spacecraft to spin it back to its original attitude. This type of event happens frequently on LRO. Earlier in the chart two zero crossings and a couple of spurious measurements can be seen which have been cleaned - the ‘initial’ telemetry (small dots) shows irregularities, but the ‘clean’ telemetry (circles) is steady.

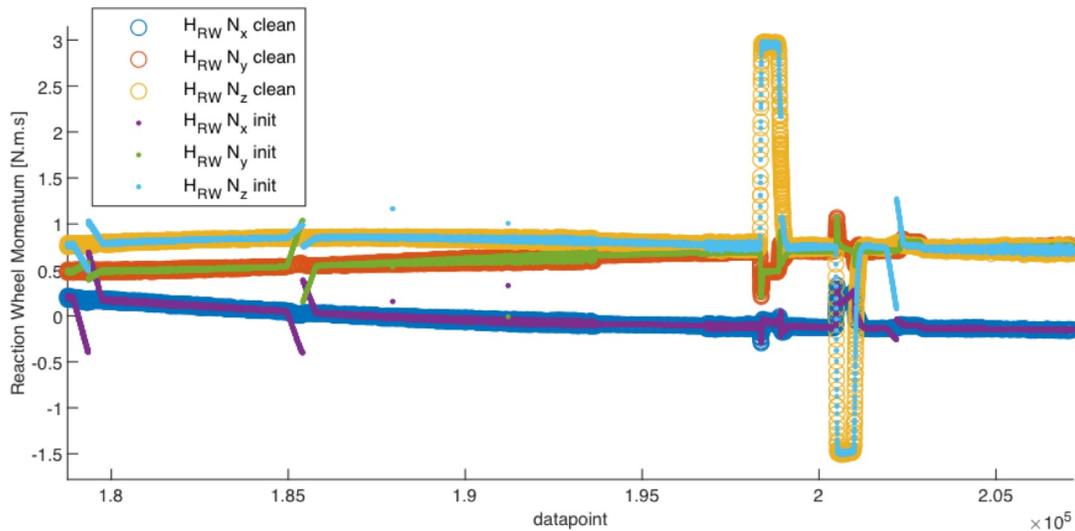


Figure 4.22: Sample telemetry pre-processed with RW correction algorithms

With these adjustments the telemetry gets much cleaner. The telemetry shown in Figure 4.20 is now repeated with the cleaned wheel speeds, and shown in Figure 4.23. With these improvements, the driving features are the sine wave in the orbit that appears to track with the orbit period, and the the large features that occur when the spacecraft slews.

Figure 4.24 takes a closer look at these slew maneuvers, to see what happens in the angular momentum at these times. The most notable feature is that the spacecraft rate contribution (i.e, the total inertia of the spacecraft, including appendage configuration, times the angular rate telemetry) is typically the mirror inverse of the reaction wheel contribution (the inertial angular momentum contained in the reaction wheels). This is exactly as expected, as the momentum transfers between reaction wheels and spacecraft body to slew the spacecraft as desired. However, if everything was perfectly measured conservation of momentum would ensure that the exact momentum absorbed/released by the reaction wheels would be the momentum released/absorbed by

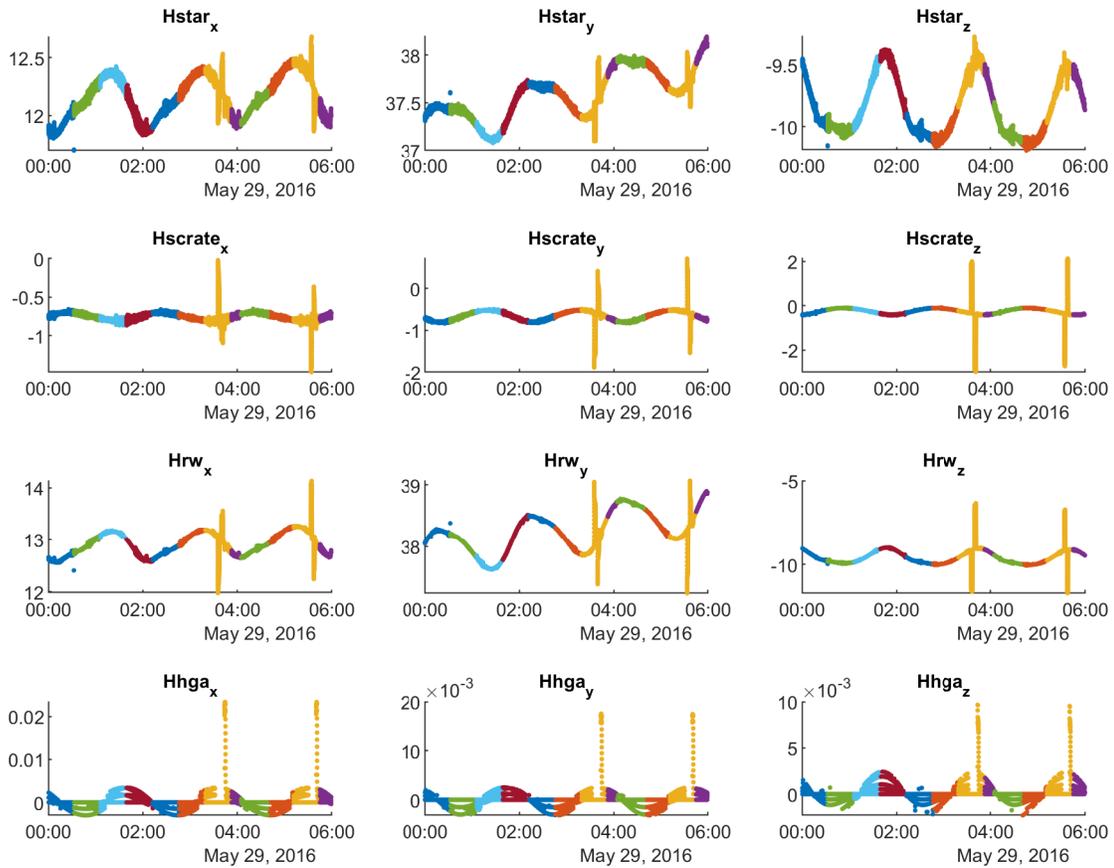


Figure 4.23: Momentum contributions for LRO - after pre-processing RW speed measurements

the spacecraft body, subject to the additional motions of the appendages. This is what is shown in \mathbf{H}^* , which should in theory be perfectly quiescent in the absence of an external torque.

Instead, features are abundant, which prevents the application of change detection algorithms at these times. While the RW telemetry and the S/C rate telemetry are more-or-less mirrored, they are not exactly mirrored, even accounting for the contributions of the HGA, and therefore features remain in \mathbf{H}^* . These discrepancies are due to how the momentum is measured, and contributors may be reaction wheel inertia discrepancies, pointing discrepancies, and especially spacecraft inertia estimate discrepancies. Typically reaction wheel EIDPs specify an as-built value for each reaction wheel, and alignment cubes are used after manufacture to measure precisely the alignment of the reaction wheel relative to all the other relevant spacecraft features (star tracker boresight,

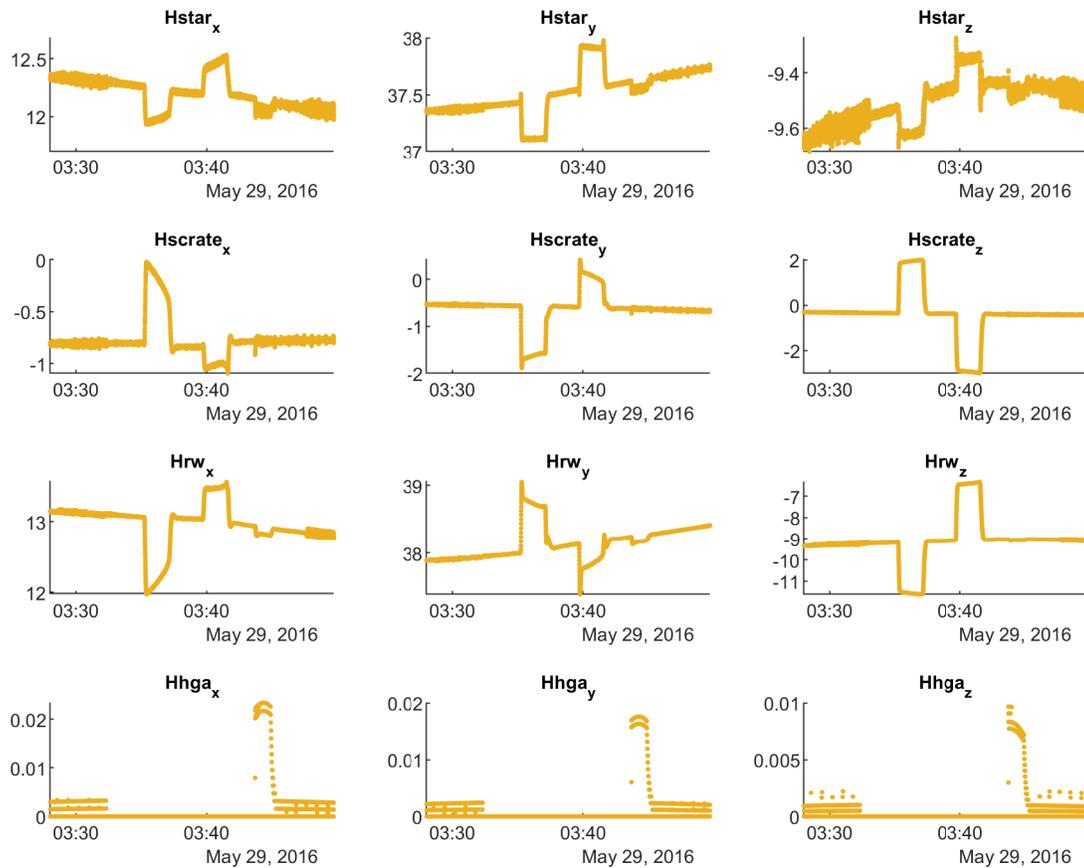


Figure 4.24: Momentum contributions for LRO during slew maneuver

instrument boresight, IRUs, thrusters, etc). These are used to correct the nominal DCMs, and two spacecraft (LRO and MAVEN) had placeholders for these matrices without populated values, although on LRO the specification states that the mounted configuration was within 1 degree of the nominal configuration. However, the elephant in the room is likely the estimate of inertia. There's not a great way to measure the spacecraft inertia precisely, leaving CAD estimates as the best source for that parameter, and based on the work in this research it is suspected that CAD estimates might not be as accurate as may be desirable for this type of effort. Fuel slosh may also contribute to this uncertainty, and while the effects of fuel slosh on performance can be modeled the precise distribution of fuel at a particular point in time is difficult to assess. In typical spacecraft operations, where sensor pointing accuracy and low jitter are sufficient for mission performance,

the robustness of the control algorithms will take care of any inertia uncertainties and converge to the desired solution based on spacecraft attitude and rate, such that mission performance is not significantly affected by the uncertainties in these estimates.

Because of these findings, two algorithms were developed to try to improve the estimate of the spacecraft's inertia. Neither worked satisfactorily. First a batch filter was tried at a time when the spacecraft was slewing, but the filter converged to a solution which did not significantly improve the quiescence of \mathbf{H}^* . A second filter, an EKF from [82] was also developed. This filter also did not improve the estimates of the inertias, it uses a FOGM-like process to assign a large time constant to its estimate of the inertias, similar to estimating an instrument bias, but the state noise covariance matrix used by the authors doesn't produce a strong coupling between the rate discrepancies and the inertia parameters, so the filter produces an estimate of the rate but doesn't correct the inertia parameters in any meaningful way. Several other filters are available in the literature, many of them require certain maneuvers by the spacecraft, which of course is not available for processing historical telemetry. Developing a new filter which is capable of accurately estimating inertias and other relevant parameters is identified as a primary research goal if these efforts are continued. Psiaki has a method which looks promising [83]. Unfortunately, by the time this limiting factor was identified as a substantial source of issues across multiple spacecraft the time remaining was too short to develop a filter of this complexity.

Finally, the database showing anomalous events during narrow angle camera telemetry was used to plot the telemetry at the time of these events and see if any unexplained features showed up. A change detection filter was run on the telemetry to plot filter output at the time of the events, to see if a small change might occur at the time of the event indicating a particle impact. While there are sometimes features in the vicinity of these NAC anomalies, there was nothing that stood out as being potentially related, given how frequently features appear in this telemetry.



Figure 4.25: Fermi spacecraft

4.1.5 Fermi

4.1.5.1 Processing Fermi Telemetry

The Fermi spacecraft is located in LEO, it is three-axis controlled using reaction wheels and slews constantly to observe various parts of the sky (Figure 4.25). It uses magnetic torquer rods to control its angular momentum which presents a unique challenge as the angular momentum is changing constantly. Therefore a proxy momentum, \mathbf{H}^* , is calculated to negate the effects of the torquer bars and produce a relatively quiescent momentum for debris strike analysis. First, the inertial angular momentum, ${}^{\mathcal{N}}\mathbf{H}$ and change in momentum due to the torquer bars, ${}^{\mathcal{N}}\dot{\mathbf{H}}$, are

calculated as follows

$$\mathcal{N}\mathbf{H} = [NB][I_{sc}]\mathcal{B}\boldsymbol{\omega} + [NB] \sum_{i=1}^N [BW]_i \begin{bmatrix} I_{ws}\Omega_i \\ 0 \\ 0 \end{bmatrix} \quad (4.19)$$

$$\mathcal{B}\mathbf{L} = d_{\text{mt}} \begin{bmatrix} I_x \\ I_y \\ I_z \end{bmatrix} \times \mathcal{B}\mathbf{B} \quad (4.20)$$

$$\mathcal{N}\dot{\mathbf{H}} = [NB]\mathcal{B}\mathbf{L} \quad (4.21)$$

where $\mathcal{B}\mathbf{L}$ is the torque applied to the spacecraft in body frame coordinates, d_{mt} is the magnetic dipole of the torque rods, I_x , I_y , and I_z are the current applied to each torque rod (which are aligned with the body frame axes), and $\mathcal{B}\mathbf{B}$ is the magnetic field as measured by on-board sensors in body frame coordinates. Then, \mathbf{H}^* is calculated by applying a sliding window to the data and propagating each $\mathcal{N}\mathbf{H}$ forward in time using $\mathcal{N}\dot{\mathbf{H}}$. It is propagated for the length of the CUSUM algorithm's sliding window, and then the difference between the propagated $\mathcal{N}\mathbf{H}$ and the true $\mathcal{N}\mathbf{H}$ is used as \mathbf{H}^* . In theory, given accurate propagation the difference is near-zero, but if a debris strike imparts an unmodeled torque a difference will persist until the window moves past the debris strike. This would produce a sustained change in \mathbf{H}^* for detection via the CUSUM SPRT.

4.1.5.2 Fermi Results

The Fermi processing algorithm described in Section 4.1.5.1 represents a first-order solution, and the results indicate that more refinement is needed before the data can be used to detect debris strikes. Figure 4.26 shows the preliminary results. The inertial angular momentum is changing due to use of the torquer bars, but the proxy momentum, \mathbf{H}^* , is relatively quiescent in comparison as desired. However, there are still significant features in \mathbf{H}^* that manifest as debris strikes in the detection algorithms, so more refinement is needed before \mathbf{H}^* can be used to detect debris strikes.

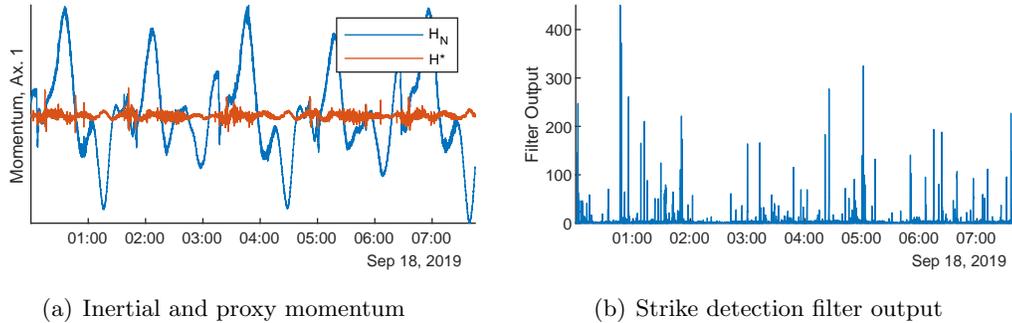


Figure 4.26: Preliminary results from Fermi telemetry

At first it was assumed that the magnetorquers were the source of these features, but after processing the MAVEN and LRO telemetry and applying the lessons learned, it appears that uncertainties in the spacecraft inertia are the primary culprit, as the main features occur when the spacecraft is slewing rather than when the magnetorquers are applied. An effect from the magnetorquers can be seen, but the inertia uncertainties appear to dominate. It's worth noting that Fermi is the only spacecraft so far that sends down a current estimate of its inertia in telemetry, presumably accommodating the effects of solar array motions, fuel remaining, etc, so it is relevant that these effects persist even with the spacecraft's own best estimate of its inertia, which lays to rest a concern that algorithms to calculate the inertia are incorrect.

Since Fermi flies in LEO with a larger concentration of debris than most spacecraft it would be a strong candidate for continued analysis, if the inertia uncertainties could be resolved.

4.1.6 Solar Dynamics Observatory (SDO)

4.1.6.1 SDO Telemetry Overview

The Solar Dynamics Observatory (SDO) is a three-axis stabilized spacecraft in GEO orbit which stares at the sun (Figure 4.27). Since its body frame is aligned with the sun vector it rotates very slowly with respect to inertial, which makes it a very quiescent spacecraft for analysis. The spacecraft's estimate of its body-frame angular momentum and attitude quaternions are used to

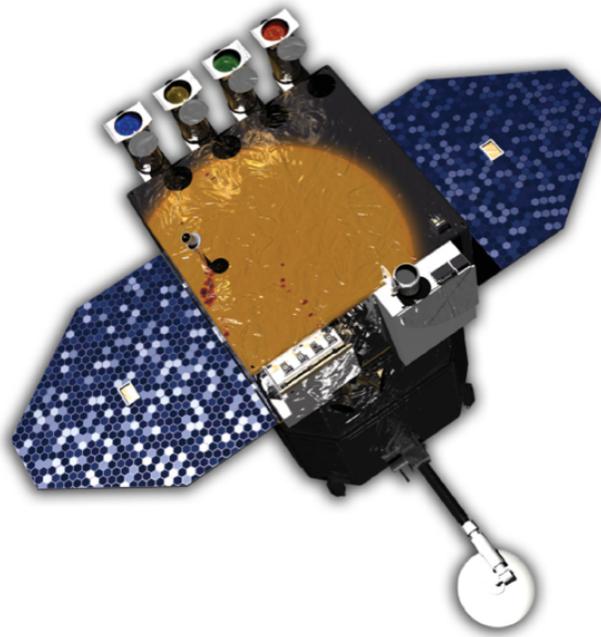


Figure 4.27: SDO spacecraft

calculate the inertial angular momentum. Among all the spacecraft examined, this is the minimal amount of telemetry required to compute the inertial angular momentum, and provides the cleanest estimate.

However, SDO's angular momentum telemetry exhibits a square wave with a period of approximately 13 minutes. This square wave is fairly consistent across years of data, but each occurrence manifests as an abrupt change in angular momentum of exactly the type that these filters are designed to detect. The theory is that this square wave is due to the motions of the high-gain antenna (HGA) as it tracks the ground station from SDO's inclined GEO orbit. However, the HGA gimbal telemetry which was examined did not show features corresponding to the square wave.

Since the abrupt momentum change trips the debris strike detection algorithms, this feature needs to be removed before the detection algorithms can be applied. Two methods of removing this feature were developed. The first, during the 2019 effort, involved applying a digital signal processing technique to notch out the frequencies corresponding to the square wave and thus smooth the data. The second, during the 2022 effort, took a simpler approach of generating a square wave

to match the telemetry and removing just the square wave.

4.1.6.2 Initial Method and Results

In the first method, this feature is removed from the data by designing a notch filter that eliminates the frequencies for both the square wave and a low-frequency (2 per day) oscillation that the data displays. A Kaiser window finite impulse response filter is used, with the frequencies associated with the undesired square wave specified as the frequencies to notch based on the fast Fourier transform (FFT) of the data as shown in Figure 4.1.6.2, with the resultant filtered data shown in Figures 4.1.6.2 and 4.1.6.2. The edges of this filter show some effects that manifest as debris strikes, so the filter is run on 10 days of data at a time and the junction between the data sets is blanked.

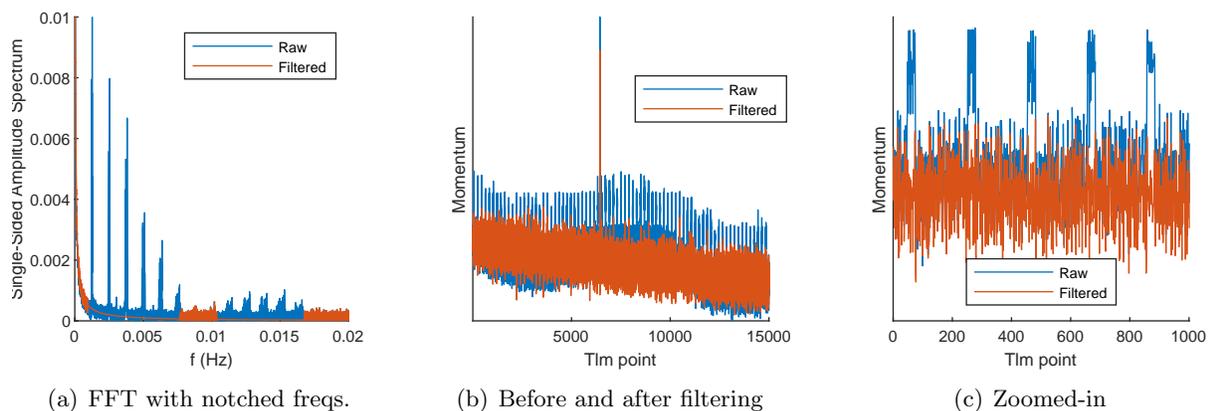


Figure 4.28: Filtering SDO angular momentum telemetry to remove periodic signatures

SDO's strike detection filters trip whenever the thrusters fire for a station-keeping or momentum-dumping maneuver, so all thruster firings are blanked from the data. Slewing maneuvers also trip the threshold due to uncertainties in the spacecraft's inertia, as the spacecraft slews the calculated inertial angular momentum varies even though the true inertial angular momentum is constant. Fortunately, slewing maneuvers and thruster firings occur infrequently. The spacecraft performs reaction wheel reconfigurations to avoid zero crossings, but the filters are robust to these and they do not need to be blanked.

For the 2019 analysis, the detection threshold is set by plotting a day of data and removing obvious spikes to show the no-strike noise floor. A kernel distribution is fitted to each axis and the resulting probability density function is used to develop the threshold. The thresholds are set such that the probability of exceeding the threshold equates to one in the number of telemetry points downlinked each day, for an approximate false alarm rate of one per day. This produces a similar threshold in each axis since the noise levels are similar, so an average value is selected as the threshold for all three axes.

This method produced some interesting results. The CUSUM algorithm applied to SDO's telemetry indicates frequent small perturbations and occasional larger perturbations. The algorithm is run with a pre-window of 200 datapoints to establish a baseline noise distribution, then a post-window of 100 datapoints to evaluate whether the underlying distribution has changed. A $\Delta\mu$ of 20 mNms is used to define the changed distribution for the alternative hypothesis in the change detection algorithm.

Figure 4.29 shows the summarized results for 2016 and 2018, with the number of small perturbations per day tallied throughout the year. These results exclude anything detected in the vicinity of mission events that cause an apparent change in body-frame angular momentum. If these detections were false alarms, a roughly constant rate would be expected throughout the year. Instead, there are periods of higher activity and lower throughout the year. These might just be due to mission activities, component idiosyncracies, or noise levels, except that these periods of higher activity and lower activity are found to correlate somewhat between 2016 and 2018.

Figure 4.30 shows the raw telemetry from the two largest events, on day 247 and 270 of 2016. The event on day 247 exhibits an abrupt momentum change (Figure 4.30(a)) but then an unexpected oscillation. The reaction wheel telemetry (Figure 4.30(b)) shows odd data preceding the time of the event. Based on this, this event is not perceived to be a strike but is some other phenomenon. Its cause has not been identified.

The second event, on day 270, exhibits telemetry that is more aligned with expectations. The angular momentum changes abruptly (Figure 4.30(c)), and the spacecraft exhibits an increase

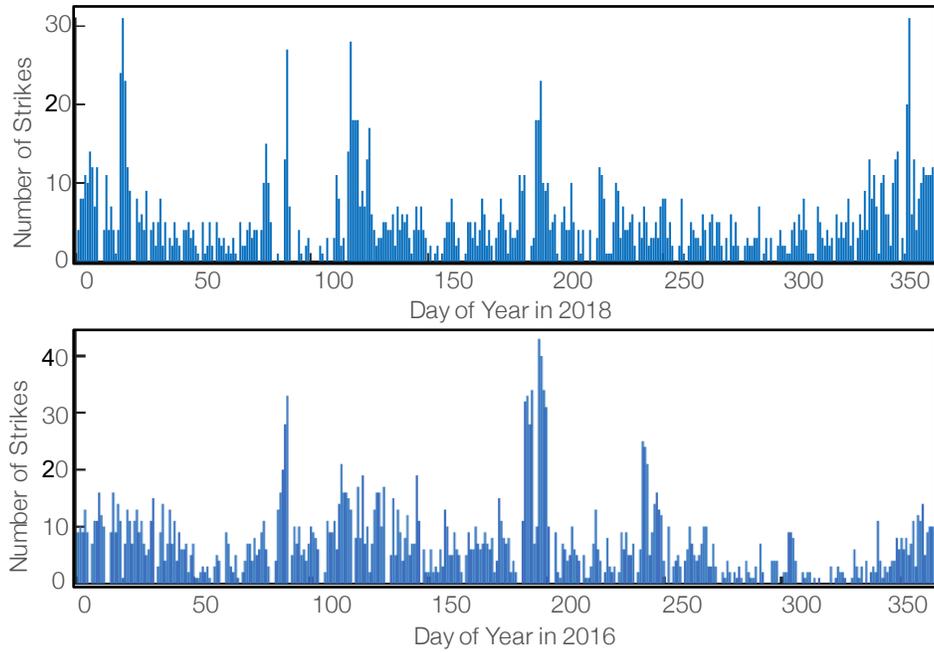


Figure 4.29: Daily number of perturbations tallied throughout year in 2016 and 2018

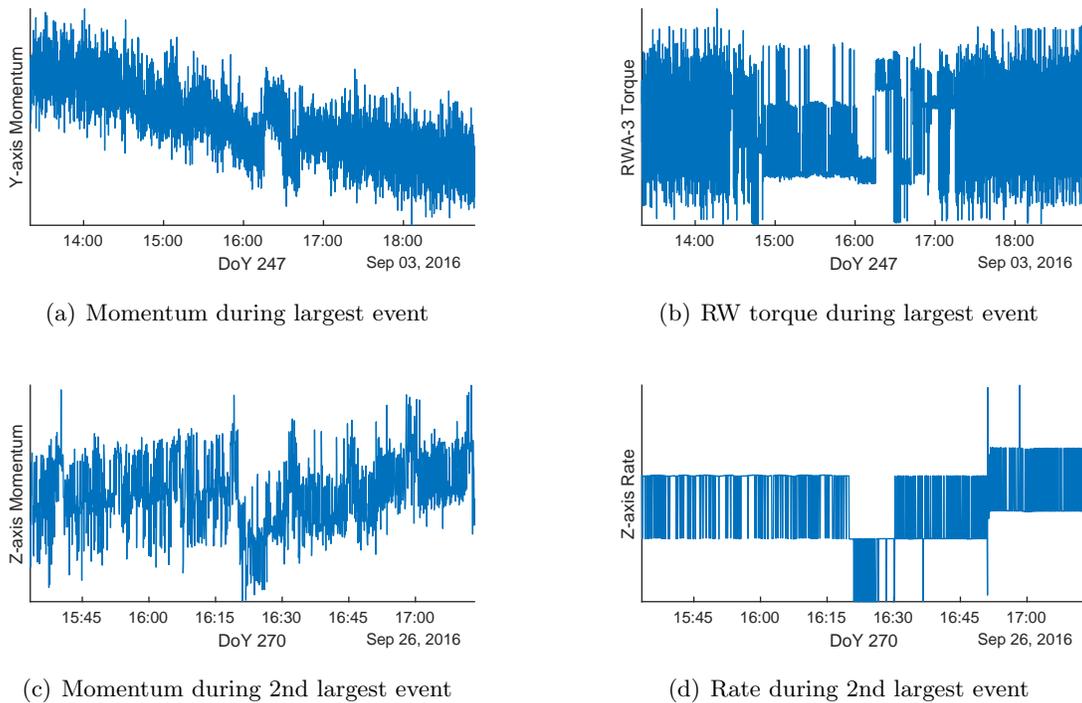


Figure 4.30: Results from SDO telemetry in 2016

in angular rate (Figure 4.30(d)) but it is corrected by the attitude control system. However, when shown this even the spacecraft operators noted that the IRU heaters had been turned on a short time prior, so it is more likely due to a change in the IRU measurement due to thermal effects than a true MMOD strike.

It is important to note that these detections are not validated against other methods, they are just the data returned by this specific algorithm. The filtering required to clean the telemetry has been known to produce false alarms at the junction between two sets of filtered data. This is rectified in this telemetry by filtering 10 days at a time and blanking the data at the intersection between filter sets, so filtering residuals are not known to cause false alarms in this telemetry. However, the potential for filtering to induce transients that manifest as MMOD strikes should not be ignored. Checkpointing provides some confidence that these detections are real momentum transfers by showing that the raw, unfiltered momentum data does exhibit step changes at the times indicated by the filter. However, the estimated magnitudes are highly suspect. The magnitudes are estimated by differencing the average between the pre-window and the post-window. The noise floor of this estimate, even when no strikes are present, is significant, and the estimated sizes are only marginally higher than the noise. Therefore these estimates of strike size may or may not be accurate, especially for smaller strikes.

Also, per analysis by Marshall Space Flight Center’s Meteoroid Environment Office using MEM3 the frequency of impacts is far too high for impacts of this magnitude - a handful of impacts per year of this magnitude may be expected, but not several per day. This is based on MEM3 and momentum enhancement factor relationship developed by McDonnell [49], although as [51] noted some measurements indicated an MEF figure “Considerably higher.”

4.1.6.3 Alternate Method for SDO Data Processing

Therefore, the 2022 data processing effort introduced an alternate method for filtering out SDO’s square wave, intended to be more predictable. Instead of employing a digital signal processing filter to notch out the frequencies corresponding to the square wave, a square wave is generated

and then its parameters are adjusted to fit the data for each window. For each data window, the correct phase offset is found by cross-correlating the square wave with the signal, then the period and duty cycle are traded to find the combination which minimizes residuals (i.e, provide best fit to square wave present in data).

After these parameters are set, the amplitude of the square wave is found by taking the average of the telemetry data during the peaks of the square waves and subtracting the average of the telemetry data during non-peak periods, to find the difference in magnitude between the square wave peaks and valleys for each individual dataset. This process is necessary because the magnitude of the square wave varies significantly, while the period and duty cycle vary somewhat over the course of time. This process is applied to data windows of less than an hour, if longer windows are processed the period and duty cycle are inconsistent across the window and result in bad 'corrections' which introduce discrepancies to the data instead of eliminating them.

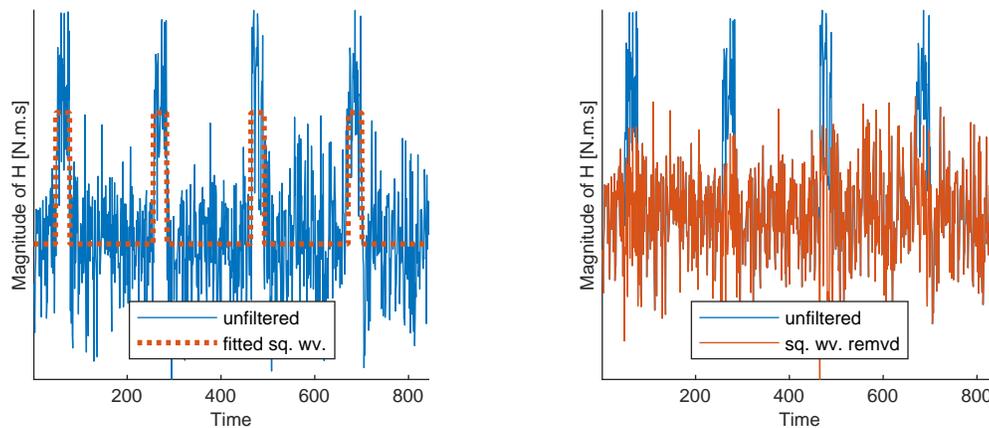


Figure 4.31: Calculating square wave to fit telemetry and removing to produce quiescent data

This algorithm is applied to pre-filter the data, then the CUSUM change detection filter is applied, and a high and low threshold is applied to identify both small and large events. See Section 4.1.6.5 for a detailed description of setting these thresholds. The large events are cataloged in a 'big event list'. When scrolling through and plotting the events in this list the performance of this algorithm can be observed across many datasets throughout SDO's lifetime. The pattern of the

square wave definitely varies over time, but the algorithm generally works satisfactorily to correct it, except when data gaps occur in the telemetry. Data gaps throw off the algorithm's ability to estimate the square wave accurately and produce false alarms, but these seem to occur relatively infrequently.

This filtering method avoids the uncertainty of the prior method, which appeared to work well on the datasets which it was tested against but represented a somewhat more 'black box' approach which may respond differently to datasets with other features. This produced the potential for the filtered data to either introduce apparent strikes when there were none, or potentially filter out strikes that did exist. The checkpointing of the large strikes in Figure 4.32 provides some confidence that changes in momentum are being observed correctly by the filter, but there is still uncertainty associated with the process, especially in estimating the magnitude of the event, which makes it hard to compare the data collected to expected perturbation rates.

Two other changes are introduced. One is to apply a dual threshold to detect both small, subtle and large, obvious perturbations; this is described in detail in Section 4.1.6.5. The other is to calculate and apply the filter to the magnitude of the inertial angular momentum and the direction of the inertial angular momentum, instead of the three axes of the angular momentum where it was applied previously. The direction is calculated using the small angle approximation, as described in Section 3.3.6. In the CUSUM change detection filter, $\Delta\mu$ applied to define the distribution for the alternative hypothesis was 10 mNm for the change in magnitude of angular momentum, and 0.004 radians for the change in direction of angular momentum.

With this new filtering method, the final results are produced. Per the goals identified in the comprehensive exam, 10 years of available data are plotted, and the list of big events is identified. The largest of these were examined and spacecraft activities or events that don't exhibit the characteristics of a debris strike were removed from the big event list. One example of a large, clear event showing an abrupt change in angular momentum is shown in Figure 4.32. There are several events of this nature, and there are some where the momentum moves around like in Figure 4.30(a) but more subtly. There are also many events where it is hard to tell by eye that

anything is going on.

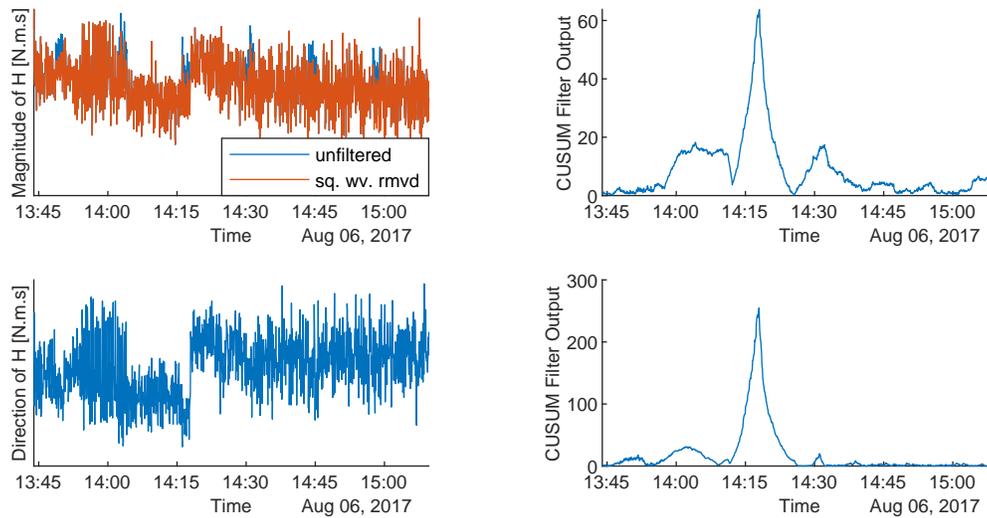


Figure 4.32: Example of abrupt change in angular momentum

4.1.6.4 Discussion of SDO results

Based on the 2016 and 2018 results, the first priority of the 2022 effort was to see if other years showed similar patterns. Fortunately, with a 2010 launch SDO offers over a decade of on-orbit data for analysis. Similar to previously, the number of times per day that the filter output exceeded the threshold is plotted throughout the year and shown in Figure 4.33. With a shorter timeline available for conducting this analysis, the thruster firings and slew maneuvers are not blanked from this dataset, since they only occur occasionally and would add one event per day, which is not particularly consequential to this dataset. The ‘Trend’ variable is the median value averaged over 5 days in a sliding window.

Similar to the 2016 and 2018 data, the 10 year dataset does seem to show a correlation in variation of perturbation counts from one year to the next. Most years show an increase in perturbation counts in the vicinity of day 130 and day 220. Three years early in operations show tremendous spikes in hit counts, but since these patterns do not repeat year-over-year they are

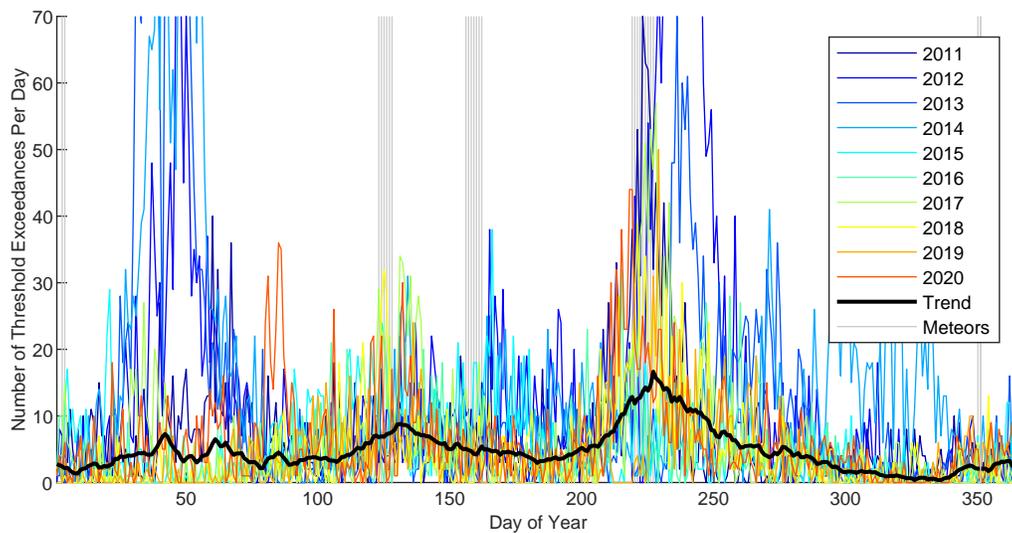
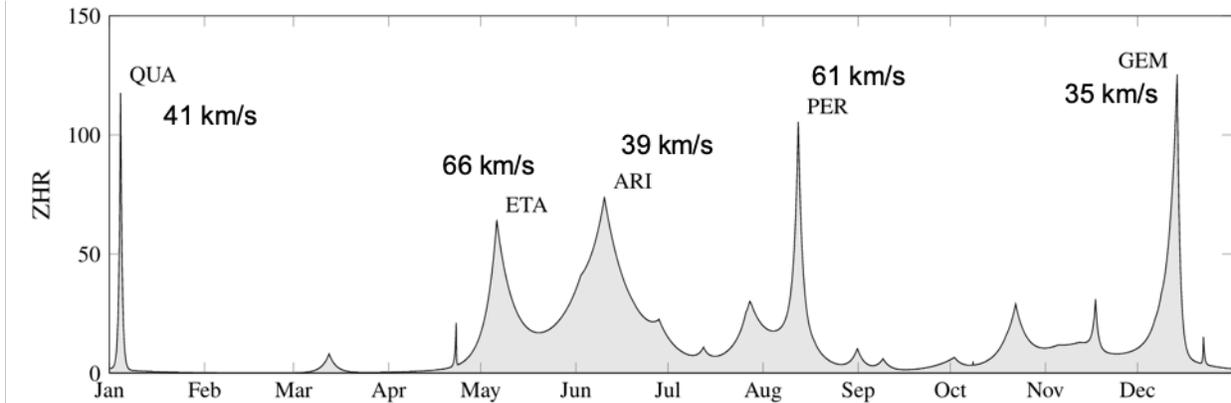


Figure 4.33: Graph of SDO’s minor perturbations per day of year shows some correlation to micrometeoroid populations

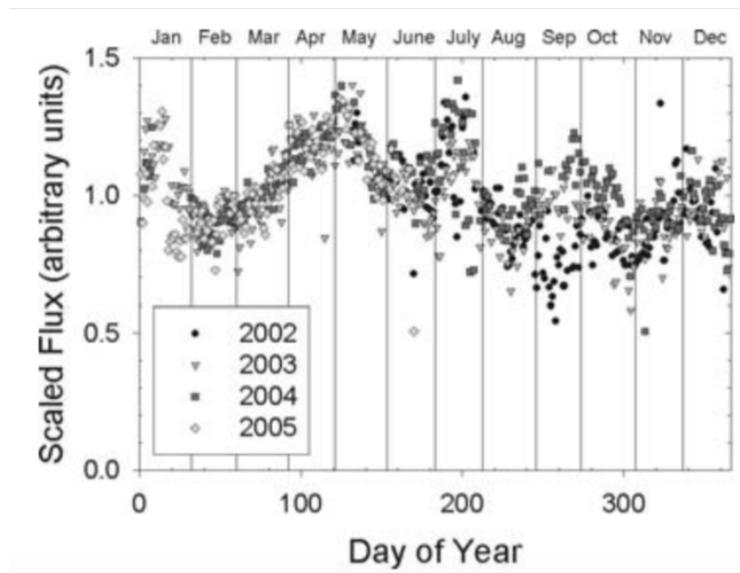
likely due to spacecraft activities, component noise, or other features rather than micrometeoroids.

Compare the patterns in Figure 4.33 to charts showing the known micrometeoroid fluxes. Figure 4.34(a) shows the flux of micrometeoroid streams, annotations are added showing estimated shower velocity. Grey lines are added to Figure 4.33 at the peaks of the major showers, which shows that for some showers there is a distinct uptick in SDO’s cataloged rate of small perturbations, while other showers see upticks only in some years, and others don’t show up at all. However, the two most well-defined peaks across most years correspond fairly well to the Perseids and Eta Aquariids. These are the showers with faster relative velocities which can therefore impart more momentum to a spacecraft, especially if the momentum enhancement factor is higher than expected at faster velocities (reference Section 1.4.4.3).

Besides the showers, another factor is the sporadic micrometeoroid flux [47], shown in Figure 4.34(b). The sporadic background is by definition those micrometeoroids that do not belong to known showers, but these make up the bulk of meteoroids encountered, and therefore are the focus of most models. The data shown in this figure is obtained using radar measurements, as opposed



(a) Meteor showers throughout year, measurement of zenith hourly rate



(b) Daily flux of sporadic meteors, measured via radar

Figure 4.34: Changes in flux of micrometeoroids, top is contribution from showers, bottom shows flux in sporadic background

to the typical rate of observed meteors (zenith hourly rate) as plotted in Figure 4.34(a), and is cataloged over several years. This shows that trends sometimes repeat year over year, especially the peak around DoY 125, while other trends have a looser correlation from year to year. For example, most years show a small peak around DoY 280, but in 2002 the counts dipped around that time.

While this data is encouraging, it is not yet validated and must be treated with caution. The

most compelling argument against this data is that these perturbation rates are much too high for the magnitude of perturbations that are being measured. Spacecraft activities, data gaps, and component idiosyncracies are all known to produce false alarms in the data, so while there's not an immediately obvious mechanism that could produce false alarms with these annually-repeating patterns it also can't be ruled out at this point. Obtaining data from additional spacecraft that also show these patterns would be compelling, but until that is accomplished these results must be treated with skepticism. Other validation methods could include a separate measurement that can detect impacts, some options for this type of validation are outlined in Section 1.4.1.

4.1.6.5 Threshold Setting

For the SDO analysis, the filter output is evaluated against two thresholds. Since a multitude of small events are observed and tracking the rate of these events is of interest, one threshold is set fairly low, with a false alarm rate of 3/day. However, since these are difficult to observe definitively in the raw data, a second, higher threshold with a false alarm rate of 1/30 days is introduced to identify larger events that can be seen clearly in the data, to evaluate the appearance of these larger events and whether they seem to be debris strikes or other idiosyncracies.

An in-depth look at setting these two thresholds provides a case study in some of the challenges of setting thresholds, and a caution about the pitfalls of casual use of statistical distributions. The top three figures in Figure 4.35 show a histogram of the filter output from SDO when no overt strikes or events are observed ('Data'). Three different probability density functions (PDFs) are fitted to this data: a Rayleigh distribution, a Gamma distribution, and a Lognormal distribution. Based on the two false alarm rates for the high and low threshold described above, the two thresholds are calculated using the fitted PDFs from each of the distributions. I.e, with a false alarm rate (RFA) of 3/day, the pdf is integrated into a cumulative distribution function (CDF), a probability of false alarm (PFA) is established corresponding to 3 data points divided by the number of data points in a day, and the CDF value of 1-PFA is mapped to the desired output threshold. Therefore, at the core of this process is how well the selected PDF fits not just the data in general, but

especially the tail of the data, as that is the point which must be accurate to obtain a reasonable threshold.

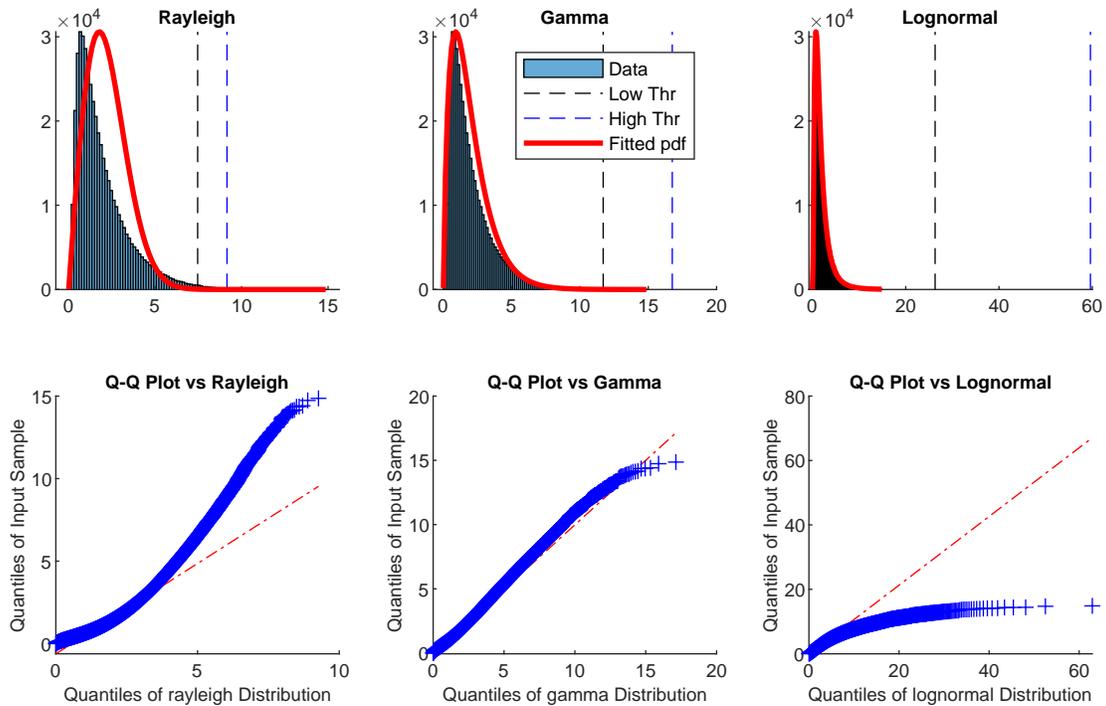


Figure 4.35: Fitting various distributions to the filter output, and effect on thresholds

Figure 4.35 shows how significantly the threshold can vary based on how well the selected distribution fits the data. The Rayleigh distribution has a very light tail which visibly over-predicts data near the mean and under-predicts the amount of data on the tail, so the thresholds are extremely low. If these thresholds were used false alarms would be seen at rates substantially higher than the desired false alarm rate, due to the poorly fit distribution tail. At the other end of the spectrum, the lognormal distribution has a long, heavy tail, pushing the thresholds unreasonably far to the left. If these thresholds were used the output data would never be sufficiently high to trip the thresholds, even though they were based on the same false alarm rate, because this distribution also models the tail poorly expecting far more measurements displaced from the mean than actually occur. The Gamma distribution appears to perform adequately for this dataset,

and while the main portion of the distribution isn't quite a perfect fit, the tail and the resulting thresholds seem reasonable. This distribution is used to generate the thresholds to produce the data shown in Figure 4.33.

In the bottom row of Figure 4.35, the Q-Q plots provide a useful tool for assessing how well each distribution fits the data. If the data is well represented by the statistical distribution, the expected quantiles based on the distribution will correspond to the sample's quantiles, and the data will fall along the red line. As shown, the data curves upward relative to the Rayleigh distribution indicating that the data is spread further or is heavier tailed than the distribution. It curves steeply downward relative to the lognormal distribution, indicating that it is lighter tailed or clustered more closely to the mean than expected for this distribution. The Gamma distribution isn't a perfect fit, it shows a little variation relative to the expected line, but it overall sticks pretty close and thus this distribution is a reasonable choice for setting the detection thresholds. This exercise demonstrates the importance of checking the data versus the statistical distribution used to model it, and the significant negative effects that can occur if the noise is poorly modeled by the selected distribution.

4.1.6.6 Remaining Limitations of SDO Processing Methods

While most spacecraft show much more dramatic variations in inertial angular momentum during their active operations, even SDO exhibits some unexpected variations. These variations were not immediately apparent looking at small data sets, but in scrolling through and checking the 'big events list', patterns were noted that appear to be a limiting factor in setting the thresholds and identifying strikes. These patterns had been noticed before, but the degree to which they are affecting the data processing may be higher than originally expected.

Figure 4.36 shows the magnitude and direction of the inertial angular momentum along with the associated CUSUM filter output. As well as a secular drift probably due to solar radiation pressure, there is a sine wave in the data with a period of approximately 2/day. The source of this is a little puzzling, as SDO stares at the sun and orbits at a rate of 1/day, in GEO. The

effect of SRP torques should be consistent, and any uncertainties in inertia shouldn't effect the measurement of the inertial angular momentum periodically like this. A nadir-pointing spacecraft like LRO sees a sinusoidal variation at the rate of the orbit period, probably due to uncertainties in spacecraft inertia effecting the measurement of the inertial angular momentum vector based on the spacecraft's attitude as it goes throughout its orbit, but since SDO rotates only very slowly relative to the inertial frame these inertia uncertainties shouldn't be displaying a sinusoidal 2-per-orbit effect. One possible explanation is that SDO's two HGA's, which track the ground station to downlink high-rate data on the sun, are moving in such a way as to produce this 2/orbit sine wave.

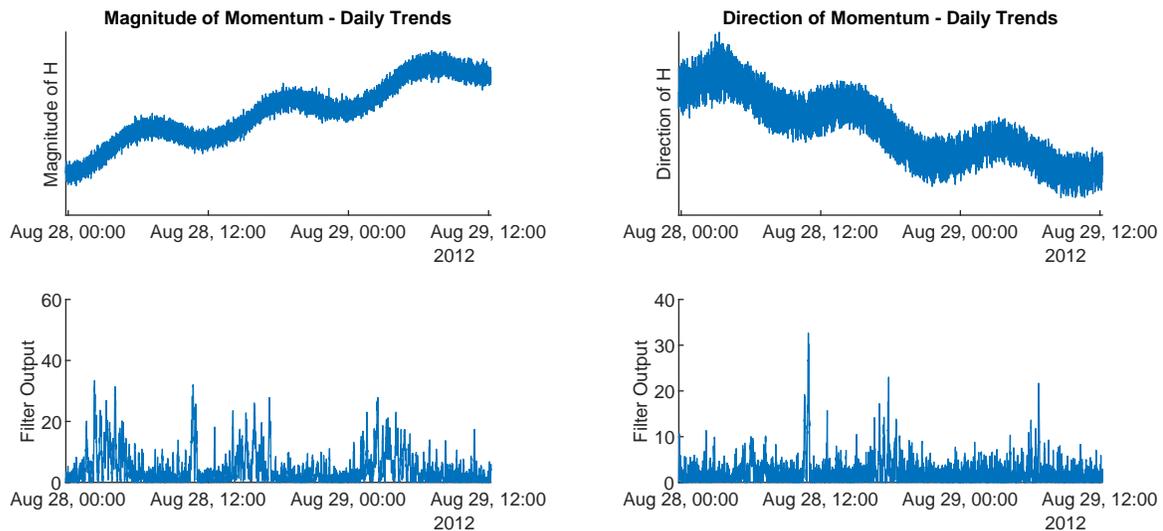


Figure 4.36: Daily trends in SDO inertial angular momentum telemetry

Figure 4.36 shows that the filter output responds to this sine wave, and the magnitude of the filter output increases when the telemetry has a steeper slope. The magnitude of this is accounted for when setting the thresholds, and with a 3/day rate of false alarm the threshold is set such that these peaks will exceed the threshold around 3x per day. Therefore, a small strike occurring during a steeper slope that increases the magnitude of the inertial momentum might trip the threshold, even though it would be too small to trip it if it had occurred in a 'valley'. Therefore, the magnitude of

measured strikes might be lower than anticipated. Also, there is always a risk that the filters are measuring noise due to some other effect, not strikes, and this graphic shows an example for how that may occur, although a mechanism that could produce a pattern of erroneous measurements like the pattern seen in Figure 4.33 has not been identified.

Figure 4.37 shows the same data, but zoomed in on the large spike that occurs during a quiescent period around 10:45. The feature is large enough that it trips the threshold for a change in the direction of momentum, even though it occurs during a quiescent period in the data, it is higher than the increased noise during the periods where the data has a steeper slope. However, to the right at 12:30 a smaller feature can be seen which is clear during a quiescent period but too small to trip the thresholds that are derived from the noisier data. At this scale this smaller peak doesn't show a distinct feature in the telemetry, but zooming in still further (Figure 4.38) a change can be seen. Note that this feature exhibits a single peak and a clear change, which is expected in a debris strike, whereas the larger feature exhibits a double peak which may be attributable to some other oddity in the telemetry, it is difficult to say definitively.

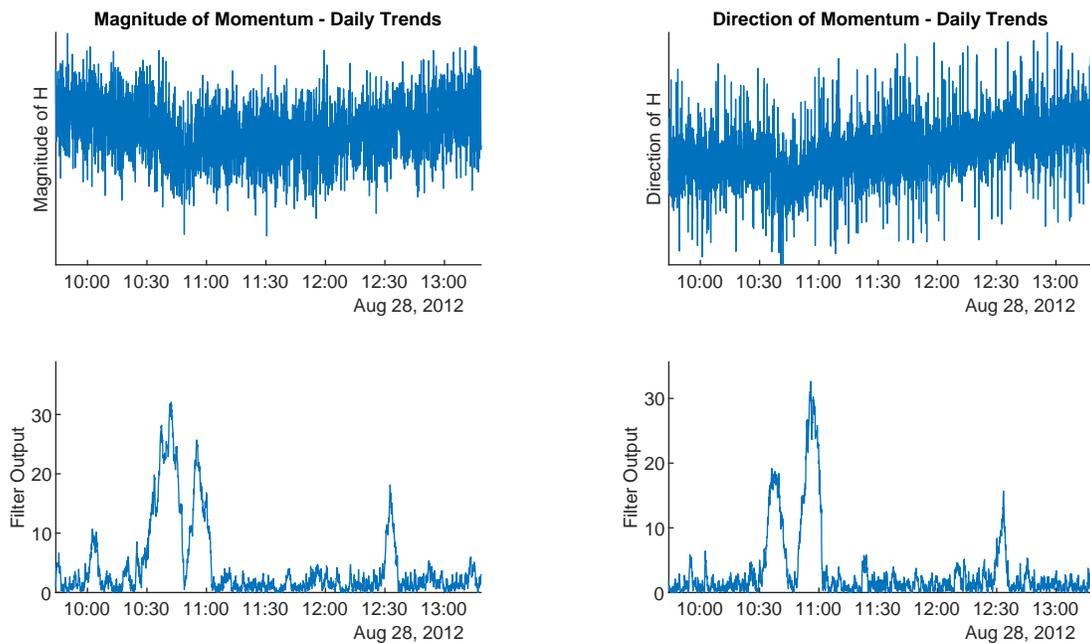


Figure 4.37: Zoomed in look at filter returns and data

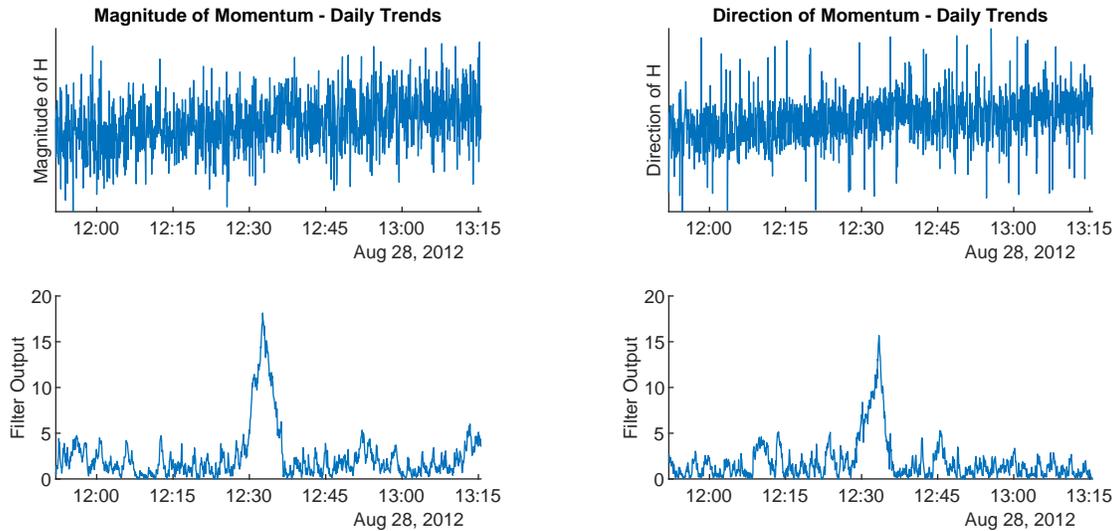


Figure 4.38: Zoomed further to see small change

These findings show that the non-quiescent trends in the data have an effect on the process of filtering and detection. Some strikes may be missed, some noise may be attributed as strikes. Originally, an attempt was made to filter out this trend and sine wave but the initial attempts were unsuccessful (due to the shorter window that ended up being required to fit the square wave effectively) so it was dropped from the code to perform the analysis shown in Figure 4.33. If continuing this analysis, the next step would be revisiting this aspect of the datacleaning and then re-running the analysis with the sine wave and daily trend removed. This may return some cleaner results with higher confidence in the perturbation rates and magnitude of perturbations, but this would only offer a marginal improvement to the overall findings as it would still be the same data on one spacecraft and therefore still subject to skepticism. Focusing on achieving similar results with a second spacecraft or external validation of events is really the compelling evidence needed to confirm these results.

4.2 Discussion of Lessons Learned

While there is always more that could be done, these results present a preliminary study on applying these methods to on-orbit data. The overall lesson is that on-orbit data is much more

challenging than simulated data, as reality typically is, and the detection thresholds identified using simulated data should generally be considered a best case scenario. Sorting out or accommodating spacecraft idiosyncracies to obtain a quiescent estimate of the inertial angular momentum is a substantial challenge to applying these methods, but the data returned by SDO illustrates that perhaps information can be gleaned from these techniques. Also, the unforeseen events detected on other spacecraft, like the MMS ‘mystery torques,’ indicate that these methods can identify events that might be of interest to operators or scientists even if they are not necessarily debris strikes.

Two findings that were consistent with other work is that accurate system documentation is key and that definitely attributing events is challenging. On two spacecraft the provided system documentation had placeholders for reaction wheel misalignment matrices, but the as-designed parameters had not been updated with the as-built parameters. Another repeated finding is that there are a lot of small features but they are difficult to confirm definitively as impacts. Consistent with the ESA findings, a threshold can be set low to see a lot of features, or high to see fewer features, but in the absence of other compelling evidence it is difficult to definitively state whether or not these features are strikes.

One finding that is new for this method is that estimating the inertia is key to obtaining a quiescent angular momentum. Every spacecraft showed variations in its inertial angular momentum, and these were often quite severe during slews, likely due to errors in the estimate of the spacecraft’s inertia. Similarly, many spacecraft showed sinusoidal variation at the same period as the orbit, likely also due to errors in the inertia estimate. The appendage motions and external torques from SRP or magnetic torquer bars sometimes also had an effect, but on multiple spacecraft the limiting factor seemed to be the errors in the inertia estimate. Therefore, future efforts should focus on developing methods to estimate these effects and account for them in the calculation of the inertial angular momentum. This seems feasible, and would help to push the actual detection threshold down toward the nominal detection threshold which is based on the inherent noise in the data. An issue that is harder to accommodate is component idiosyncracies, which are more spacecraft dependent, and these might turn into a limiting factor if the estimates of system and appendage

inertia and external torques were obtained.

The largest source of concern bringing skepticism to the SDO findings is that the magnitude of the events seems to be much higher than anticipated. This might very well mean that these are not true strike detections and are only measuring noise or other system idiosyncracies, but a review of the literature indicates that perhaps the MEF is larger than predicted by the prevalent models at higher velocities. Additional validation is needed to adjudicate this uncertainty.

It is important to remember that these missions are all exquisite one-offs. They all perform their primary mission extraordinarily well, and have dramatically improved humanity's understanding of the Sun's dynamics, the Lunar surface, and the Martian atmosphere. These spacecraft were never designed with this data in mind, so when references are made to errors in inertia measurements or other uncertainties, this is only an issue because the spacecraft data is being pushed for measurements that it was never designed for. The spacecraft are performing exceptionally well at their designed mission, so the fact that there might be a path for some of them to also contribute data on the MMOD environment is a neat side study, but in critiques of their capabilities in this area it must be remembered that this is not relevant to the quality of their contributions in their primary mission area.

Overall, the takeaway from these studies is that the methods are harder to apply than anticipated, the results are murky but may indicate an interesting measurement capability, and additional developments are available that could potentially improve the results. Also, while the primary objective of this research is to obtain measurements of the HNT debris population, for an individual spacecraft operator the ability to detect not just debris strikes but also other subtle component idiosyncracies in real time could be useful for state of health monitoring and anomaly attribution.

Chapter 5

Development of a Perturbation Rate Assessment Tool

5.1 Context for Development Perturbation Rate Assessment Tool

To summarize the research thus far, Chapter Two developed methods to detect angular momentum transfer in debris strikes, ΔH . Chapter Three developed methods to detect linear momentum transfer in debris strikes, ΔV . Chapter Four applied these methods to telemetry from active satellites, investigating the challenges of implementation and beginning to understand what sort of detection thresholds may be attainable. Together, they demonstrate a potential capability to detect minor debris strikes in the telemetry of active satellites, even if the strikes do not cause an anomaly in spacecraft performance. This final chapter gets after the ‘So what?’: What data could be collected by various space systems if these algorithms were applied successfully? And if that data were collected, how could it be used to improve debris risk assessments?

Recall Section 1.2.8, which describes a 2017 report by the NASA Engineering and Safety Center (NESC) which compares the predicted risks per debris environment models to actual anomalies experienced on orbit. The report found that models frequently overpredicted the risks to space systems. In one study, the models predicted between 24 and 164 perturbation events caused by debris strikes, but operators only experienced seven events. However, the report described an effort to adjust the assumptions underlying the predictions, changing the assumed MEF to two and reducing the assumed mass of the debris pieces by changing their shape, and *voilà*, the data matched on-orbit experiences fairly well. Figure 1.13, showing the improved fit, is repeated here for convenience.

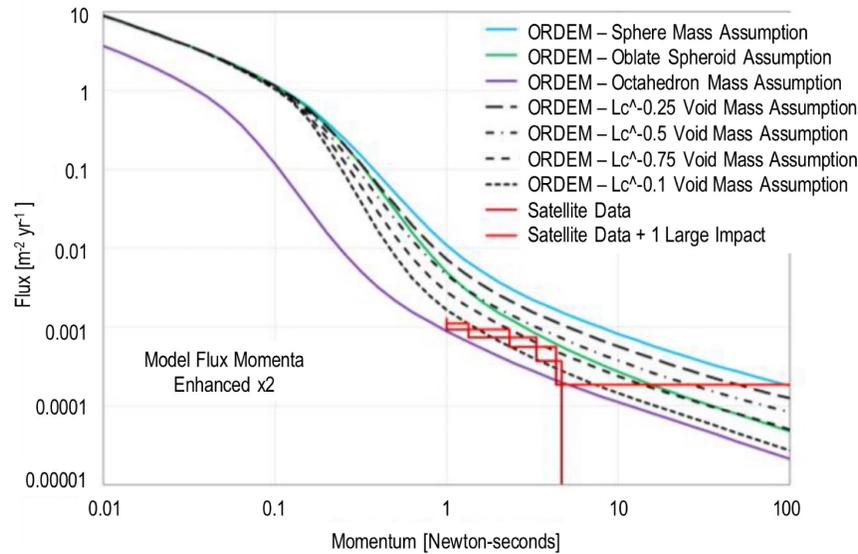


Figure 5.1: Graphic from [17] showing that the model matches the data fairly well

However, while these tweaks resulted in a nice fit for this small dataset, the underlying distributions are still unknown. The MEF is likely around two at debris velocities, based on test and analysis, but may be much higher at micrometeoroid velocities and may be different against various spacecraft materials. While not every debris piece is a sphere, and a sphere is the most conservative shape assumption since it maximizes the mass per characteristic length, not every piece of debris is a voided sphere which is the way it is modeled in Figure 5.1. DebrisSat and other hypervelocity impact tests prove that debris comes in widely varying shapes and materials, but identifying the statistical distributions which represent these shapes and materials is a major challenge, and yet necessary to produce accurate debris risk assessments.

Also, the system that is being used to collect the data matters. The trades in Section 2.5 illustrate that a smaller spacecraft may have a lower detection threshold, but also its lower cross-sectional area means that fewer debris strikes will occur. Several questions remain: which is a better detector, a large spacecraft with large solar arrays, but also large reaction wheels which are insensitive to minor changes? Or a small spacecraft with a smaller collection area but a more significant response to strikes? How much do solar arrays matter? Will strikes breaking through

arrays make the data irrelevant? The NESC report indicated that all seven impacts were likely bus strikes, but the Sentinel-1A spacecraft saw a major debris strike which did not break through the array, possibly because the projectile had a large area for its mass, like a scrap of MLI or CFRP.

The tool developed in this chapter puts all these considerations together to predict the rate of detectable perturbations on a specified spacecraft in a specified orbit. In several places the models available to predict things like MEF or predict the mass of a given characteristic length are unsatisfactory. Instead of using one model for these important assumptions, the tool incorporates several models that can be traded to assess the effect on expected perturbation rates. If better models are developed, it is straightforward to add them to this tool and determine their effect on expected perturbation rates. The idea is that if several datasets are available the various underlying assumptions could be tested against these datasets, to determine which seem to be valid and which are inconsistent with on-orbit experiences.

This tool uses NASA's Orbital Debris Environmental Model, ORDEM, to predict the flux of debris particles. ORDEM returns debris fluxes in expected number of particles per meter squared per year, according to the user-specified orbit. It bins these fluxes in a variety of ways, such as by velocity, the directions which the debris is likely to strike the spacecraft, and density class (i.e, flux of high-density particles (copper, steel) vs. medium density (aluminum)). This data is typically used to assess debris risks for a mission and demonstrate compliance with debris mitigation practices - i.e, determining the probability of mission failure due to debris strike that could prevent post-mission disposal. This tool switches that up, and instead uses the ORDEM data to conduct Monte-Carlo draws of debris strikes against satellites, to measure the perturbation caused by those strikes and determine if the strike can be detected by a specified threshold. Using the expected flux per area per year, an estimated rate of detectable strikes is then established for a specified spacecraft in a specified orbit based on the magnitude of strike that the spacecraft can detect.

5.2 Overview of Tool

The intention of this tool is to not just predict perturbation rates for a given spacecraft in a given orbit, but also establish a trade architecture to study how these predictions change depending on the model assumptions used. As such, the core of the analysis involves determining how a given piece of debris will affect a spacecraft, but this is wrapped in two layers to obtain the desired results. The first layer applies strikes to the spacecraft in a Monte Carlo, so that strike parameters can be pulled from specified distributions. This Monte Carlo is wrapped in a trade architecture which can specify which models the Monte Carlo should draw from, and summarizes the results of the Monte Carlo to understand the relationship between the models used and the expected rate of perturbations. Figure 5.2 shows a diagram of this architecture.

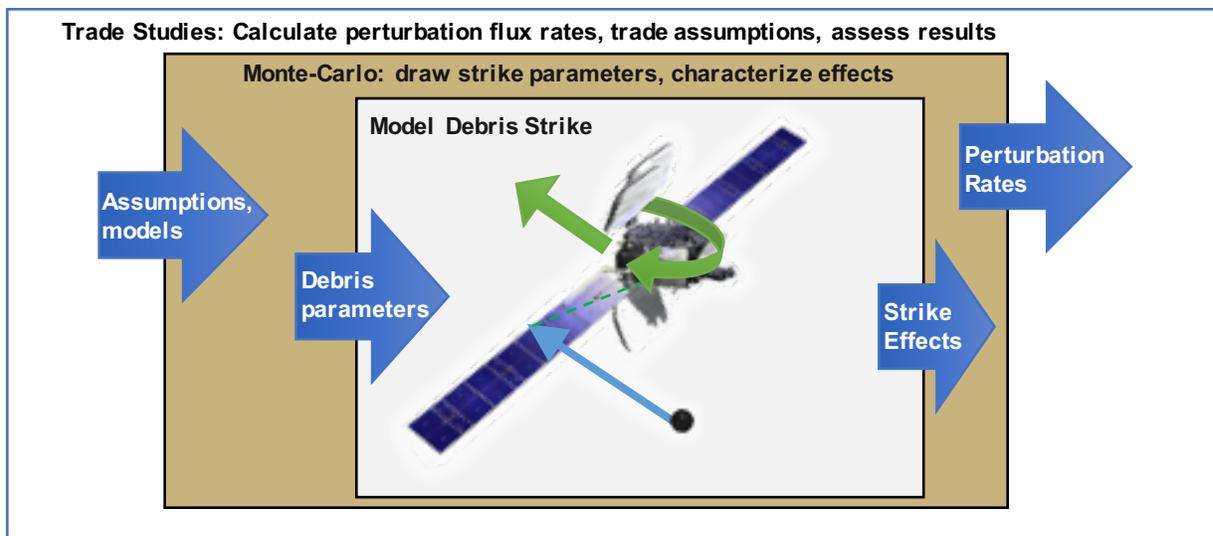


Figure 5.2: Schematic of architecture to assess expected perturbation rates

5.3 Implementation Details

Given the uncertainties outlined in Section 5.1 the architecture is developed to be agile, with a capability to be adapted with various methods and models. For example, the momentum enhancement factor can be assigned a static number such as one, two, or three, or be a velocity-

dependent relation with a maximum of around 10 at 68 km/s, or a higher velocity-dependent relation with a maximum around 20 at 68 km/s [49]. This provides a built-in capability to test the effects of various theories on the expected number of perturbations, to characterize the influence of common assumptions and the discrepancies that erroneous assumptions might produce between model predictions and the experiences of satellite operators.

The architecture is developed in three pieces, beginning with the basic mechanics of the debris strike and its effect on a spacecraft, then augmenting that function to draw representative values out of parameter distributions and apply them in a Monte Carlo, then wrapping the capability into a structure that allows efficient and agile trade studies between various model assumptions to determine expected perturbation rates.

5.3.1 Definition of Frames Used for Perturbation Assessment

Several reference frames are defined as shown in Figure 5.3 for use throughout the architecture development. The frames associated with the orbit are shown on the left: the Earth-centered inertial frame \mathcal{N} , the orbit frame \mathcal{O} , and the geocentric solar ecliptic frame \mathcal{G} . The zoomed-in view on the right shows the frames associated with the spacecraft orientation and configuration: the spacecraft body frame \mathcal{B} and the solar array frame \mathcal{S} . Each frame's basis vectors are defined as follows:

- Earth-centered inertial: $\mathcal{N} = \{O, \hat{n}_1, \hat{n}_2, \hat{n}_3\}$. This is the core frame, other frames are defined relative to this one.
- Orbit frame ('Hill frame'): $\mathcal{O} = \{C, \hat{o}_r, \hat{o}_\theta, \hat{o}_h\}$ defined by radial vector \hat{o}_r from O , the center of the Earth to C , the center of mass of the satellite. The third basis vector \hat{o}_h is defined as normal to the plane of the satellite's motion, so the direction of $\mathbf{h} = \mathbf{r} \times \mathbf{v}$, and then $\hat{o}_\theta = \hat{o}_h \times \hat{o}_r$. Debris flux is typically defined in this frame.
- Geocentric solar ecliptic frame: $\mathcal{G} = \{O, \hat{g}_1, \hat{g}_2, \hat{g}_3\}$. \hat{g}_1 points from Earth to the sun, \hat{g}_3 is perpendicular to the plane of the ecliptic, and $\hat{g}_2 = \hat{g}_3 \times \hat{g}_1$. This frame is relevant to

determining solar panel pointing and, would be relevant to micrometeoroid flux directions, if included in flux, especially if the variability in micrometeoroid flux was well modeled.

- Spacecraft body frame: $\mathcal{B} = \{C, \hat{b}_x, \hat{b}_y, \hat{b}_z\}$. This frame is fixed to the spacecraft body, its orientation relative to the orbit frame is defined by mission requirements - nadir pointing is common. This frame is relevant to determining the effects of the strike on the spacecraft state.
- Spacecraft solar array frame: $\mathcal{S} = \{C, \hat{s}_1, \hat{s}_2, \hat{s}_3\}$. This frame defines the rotation of the solar arrays relative to the spacecraft body frame. The \hat{s}_1 direction is normal to the solar array face and \hat{s}_2 is aligned with \hat{b}_y , so an Euler rotation around the 2nd axis defines the orientation of the array relative to the spacecraft. This frame is relevant to determining the cross-sectional area of the spacecraft relative to the debris and micrometeoroid fluxes.

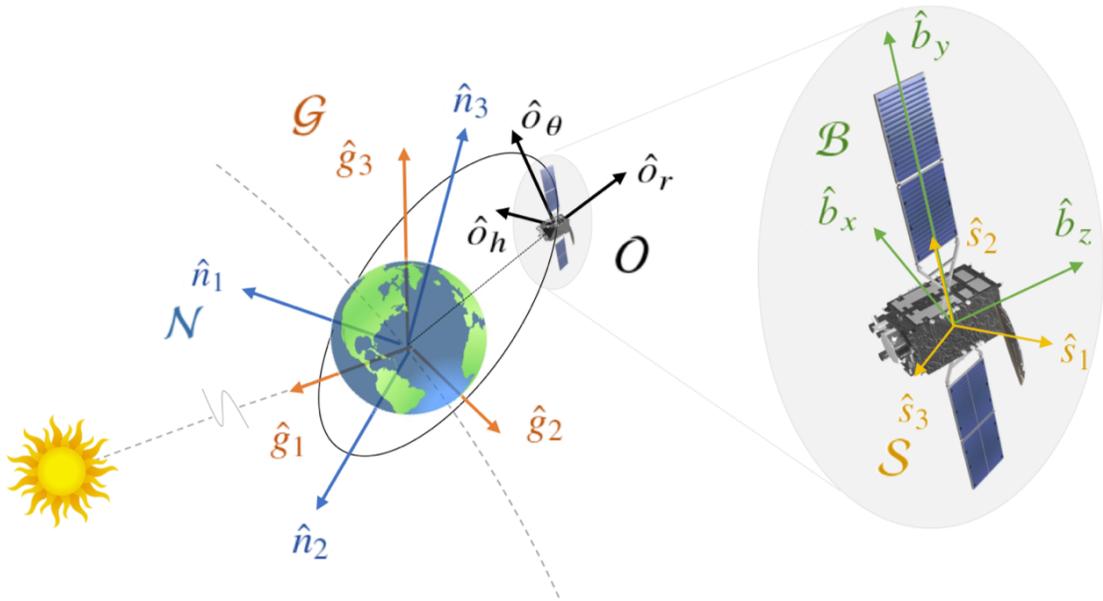


Figure 5.3: Coordinate frames used in perturbation prediction architecture

5.3.2 Modeling Strike Effects

The first task is to determine the response of the spacecraft to a given strike. When a piece of debris strikes a spacecraft it produces a change in both rotational and translational momentum. Based on the previous developments it is clear that the exact detectable threshold will vary based on the spacecraft considered, and that several parameters inherent in the analysis are not well understood.

5.3.2.1 Change in Translational and Rotational Momentum

The change in translational momentum is denoted as $\Delta\mathbf{P}$. This $\Delta\mathbf{P}$ depends not only on the incoming momentum of the debris relative to the spacecraft, $m_d\mathbf{V}_{d/s}$, but also on the momentum enhancement factor (MEF), which is not particularly well characterized for this impact scenario. For this effort, the change in translational momentum is calculated as

$$\Delta\mathbf{P} = \beta m_d \mathbf{V}_{d/s} \quad (5.1)$$

The imparted momentum also causes a change in the angular momentum of the spacecraft, assuming the strike is displaced from the spacecraft center of gravity by some lever arm, $\mathbf{R}_{d/s}$. The change in angular momentum of the satellite $\Delta\mathbf{H}$ is calculated as

$$\Delta\mathbf{H} = \mathbf{R}_{d/s} \times \Delta\mathbf{P} \quad (5.2)$$

Each of these variables, β , m_d , $\mathbf{V}_{d/s}$, and $\mathbf{R}_{d/s}$ is inherently probabilistic, and the method for determining each for the Monte Carlo analysis is discussed in the Section 5.3.3

5.3.2.2 Detectable Effects

These changes in the translational and rotational momentum of the satellite manifest in several ways, and detecting these changes is the topic of the preceding developments. For this effort, these changes are simplified to $\Delta\mathbf{V}$ and $\Delta\mathbf{H}$, as the detectability thresholds developed in previous sections can generally be reduced to these parameters.

The imparted translational momentum causes an abrupt change in the velocity of the satellite, $\Delta\mathbf{V}$. Using conservation of momentum, this is calculated via

$$\Delta\mathbf{V} = \frac{\Delta\mathbf{P}}{m_{\text{sat}}} \quad (5.3)$$

This applied $\Delta\mathbf{V}$ produces a change in the orbit of the satellite. The orbit change may manifest in different ways depending on the direction of the $\Delta\mathbf{V}$. For in-track changes, Williamsen et. al. investigated the magnitude of the change in satellite mean altitude [6], which proves to be a strong detection method, and fairly easy to implement. Therefore, the function outputs the vector change in angular momentum $\Delta\mathbf{H}$, the magnitude ΔH , the magnitude of the change in velocity ΔV , and the change in satellite mean altitude Δa for evaluating the detectability of the strike applied in each Monte Carlo run. Δa is calculated via the vis viva equation, where the semi-major axis before and after the strike is calculated as

$$a = \frac{\mu r}{2\mu - rv^2} \quad (5.4)$$

where r and v are the scalar position and velocity of the satellite. Thus, the r just before and just after the strike is the same, but the velocity changes by $\Delta\mathbf{V}$, leading to a change in a , Δa .

5.3.3 Drawing Strike Parameters from Relevant Distributions

While Section 5.3.2 can determine the strike effects based on prescribed debris parameters, in reality the parameters of a debris strike are highly variable, and a variety of different plausible strikes can map to the same effect on the satellite. However, if overly conservative assumptions are used to determine those parameters, the resulting predictions may vary widely from reality, as shown in the NESR report [1].

Therefore, this architecture is intended to capture and trade nuances like debris shape and MEF to use for developing estimates for frequency of detectable perturbations. This section describes the process used to randomize the strike parameters based on available datasets and models for determining the relevant characteristics of debris populations. This randomization process allows a Monte-Carlo analysis of the effects of debris strikes, and by changing the underlying processes

and assumptions the resultant changes to predicted perturbation rates can be observed.

5.3.3.1 Drawing Random Values from Distributions

The debris strike parameters are drawn using inverse transform sampling to generate random values from specified empirical distributions. To accomplish this, the ORDEM output is normalized to form a probability density function (pdf), then a cumulative sum is taken to produce a cumulative distribution function (CDF). Then a uniform random variable is generated on the interval from zero to one, and the CDF is used to map that number to a variable which is now representative of the original empirical distribution. This process is depicted graphically in Figure 5.4, where an example of drawing random numbers from the velocity distribution is shown.

5.3.3.2 Velocity of Debris

The distribution for the velocity of debris is pulled from the data in ORDEM, NASA's Orbital Debris Environment Model. When ORDEM is run in spacecraft mode (as opposed to sensor mode, which predicts observations) it returns a variety of tabulated information regarding the debris flux that the specified spacecraft is expected to experience. In spacecraft mode the orbit of the spacecraft is specified, so the results need to be re-generated for each individual orbit under consideration. ORDEM provides files which detail the cumulative distribution of the characteristic length of debris vs. expected flux, and the velocity distribution of the debris, and also the expected prevalent directions of the flux relative to the orbit frame.

Figure 5.4(a) shows the distribution of the velocity magnitude retrieved from ORDEM. This distribution is scaled into an empirical pdf, i.e, scaling such that the total area under the curve is one, and then summed into an empirical CDF, as shown in Figure 5.4(b). A uniform random number between zero and one is generated and mapped to a statistically appropriate velocity magnitude using this CDF. The characteristic length of the debris is selected in a similar manner based on ORDEM data.

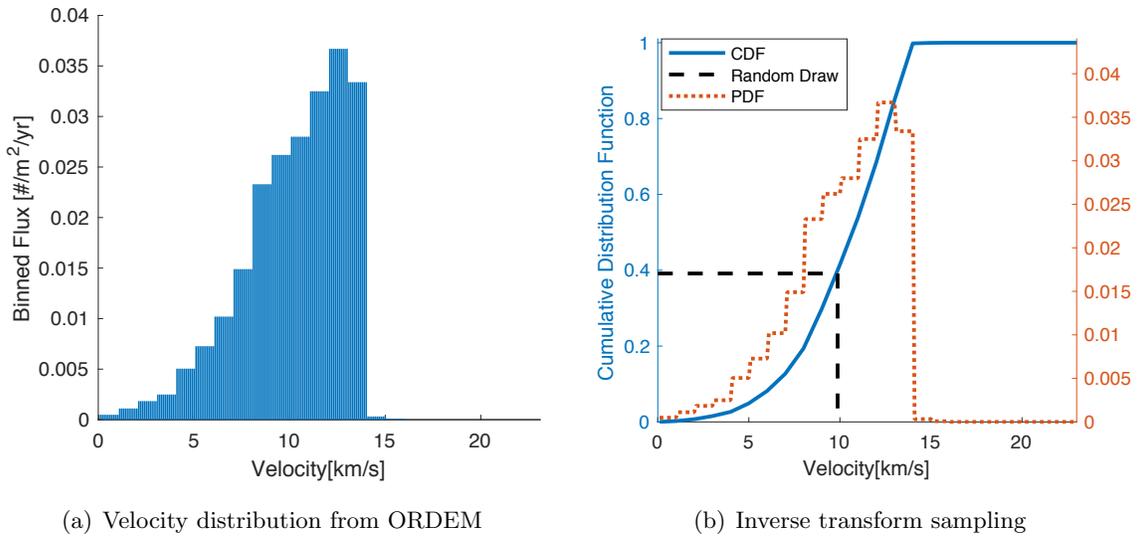


Figure 5.4: Drawing random number for velocity magnitude from ORDEM distribution

5.3.3.3 Direction of Debris

Similarly, the direction of the debris in the orbit frame is randomized based on ORDEM flux data. ORDEM provides an igloo file which breaks the orbit frame up into boxes and provides the flux expected in each of those boxes. For LEO orbits, the strikes mostly come from the ram direction. Figure 5.5 shows a graphical output from ORDEM indicating a heat map for the prevalent flux directions

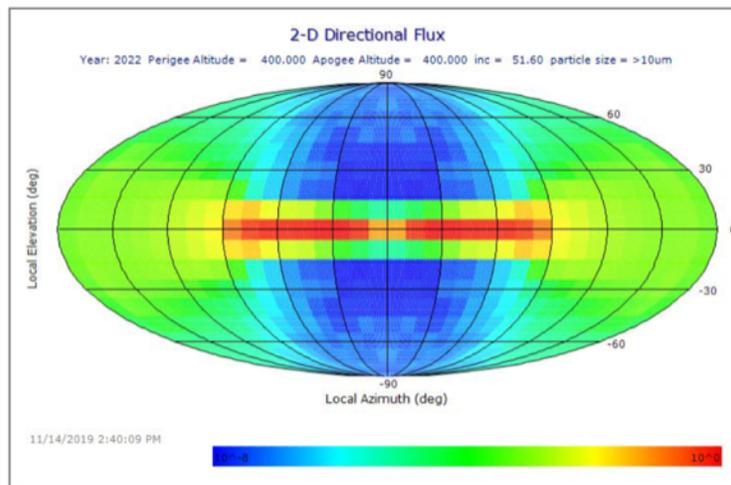


Figure 5.5: Velocity distribution from ORDEM

5.3.3.4 Mass of Debris

The NESC report suggests that uncertainties in the relationship between characteristic length (the variable provided in ORDEM) and mass may be a contributor to discrepancies between expected perturbations and experienced perturbations. This is fundamentally due to significant uncertainties in small debris measurement. Radar measures the radar cross section of debris which is then correlated to characteristic length, while sample return missions permit identification of the impacting materials for very small particles but no estimate of size or data about the larger debris fragments. A common assumption is to use a spherical shape and the density class specified in ORDEM, but this results in a conservative risk assessment as it maximizes the impacting mass for each characteristic length. An alternate assumption is the voided sphere assumption, which reduces the volume based on the area-to-mass ratios identified in the SOCIT test series. However, these tests used older manufacturing methods for satellites and the debris generated by these older satellites exhibits very different characteristics relative to the debris from newer satellites [12].

DebrisSat is in the process of characterizing fragmentation debris from newer satellites to an unprecedented level of detail, but the diversity in shapes and materials means that there is no one-size-fits-all relationship for converting characteristic length to mass, it is dependent on the properties of each object. However, in a Monte-Carlo analysis these nuances can be captured by inputting probabilistic distributions of debris properties, the trick is to develop accurate distributions for this complex and difficult-to-measure population.

Three options are investigated for this architecture, although it would be straightforward to add more options if additional models were developed:

- **Aluminum Sphere Assumption:** The conservative sphere-mass assumption is used to derive the mass assuming that the characteristic length is the diameter of an aluminum sphere. Thus,

$$m_{\text{deb}} = \frac{4}{3} \rho_{\text{al}} \pi \left(\frac{L_c}{2} \right)^3 \quad (5.5)$$

- **ORDEM-provided density classes with sphere assumption** ORDEM provides the

flux binned by density class, specifying the flux of NaK droplets (0.9 g/cc), low-density materials (1.4 g/cc), medium-density materials (2.8 g/cc), and high-density materials (7.9 g/cc). This option calculates the mass by determining, for this L_c , what percentage of the flux is expected to be in each category, and then randomly drawing the category accordingly and assigning a density based on that.

- **SOCIT voided sphere assumption:** The SOCIT fit for aluminum pieces is applied to debris objects if the characteristic length larger than 0.79 mm. The density of the voided sphere (in g/cm³) is specified for aluminum as

$$\rho_{\text{voided}} = 2.8 \frac{L_{C, \text{mm}}^{-0.62}}{.79} \quad (5.6)$$

and this adjusted density is used to determine the mass of the debris based on the characteristic length for low and medium density materials. For high density materials, this alternate relationship is used

$$\rho_{\text{voided}} = 7.9 \frac{L_{C, \text{mm}}^{-1}}{1.02} \quad (5.7)$$

- **Rough estimate of DebrisSat results:** Referring back to Figure 1.14, it is clear that while the SOCIT fit approximates the relationship between characteristic length and mass, in reality a given characteristic length can map to a wide range of masses, spanning an order of magnitude or more. To capture this a probabilistic distribution is added that generates a random number of the mass based on the characteristic length in a distribution intended to be a rough approximation of the DebrisSat data. The following algorithm is used to determine effective density based on characteristic length, similar to the SOCIT

approach but modified to be probabilistic:

$$\rho_{\text{voided}} = \frac{2.8 \frac{L_{C, \text{mm}}^{-0.62}}{.79}}{5} (1 + 5\sigma) \quad \text{for } L_c < 7\text{mm} \quad (5.8)$$

$$\rho_{\text{voided}} = \frac{2.8 \frac{L_{C, \text{mm}}^{-0.62}}{.79}}{5} (1 + 10\sigma) \quad \text{for } 7\text{mm} < L_c < 20\text{mm} \quad (5.9)$$

$$\rho_{\text{voided}} = 2.8 \frac{L_{C, \text{mm}}^{-0.62}}{.79} \quad \text{for } L_c > 20\text{mm} \quad (5.10)$$

$$(5.11)$$

where $\sigma \sim \mathcal{N}(0, 1)$. Note that this is a very approximate fit based on the graphic in [12]. The intention is just to illustrate that this architecture can accommodate the probabilistic nature of debris and apply it via a Monte Carlo, more granularity on the expected distribution of debris would be needed to bring any reality to these numbers.

Figure 5.6 shows the characteristic length-to-mass relationship for a Monte Carlo using each of these relationships. The probabilistic relationship is clear, with a large spread of data, while the other relationships fall on lines. Some have two lines representing two different density classes. The low density and NaK populations appear to represent an extremely low percentage of the flux in these size regimes, so while they are theoretically possible they don't appear to be represented in these Monte Carlo draws.

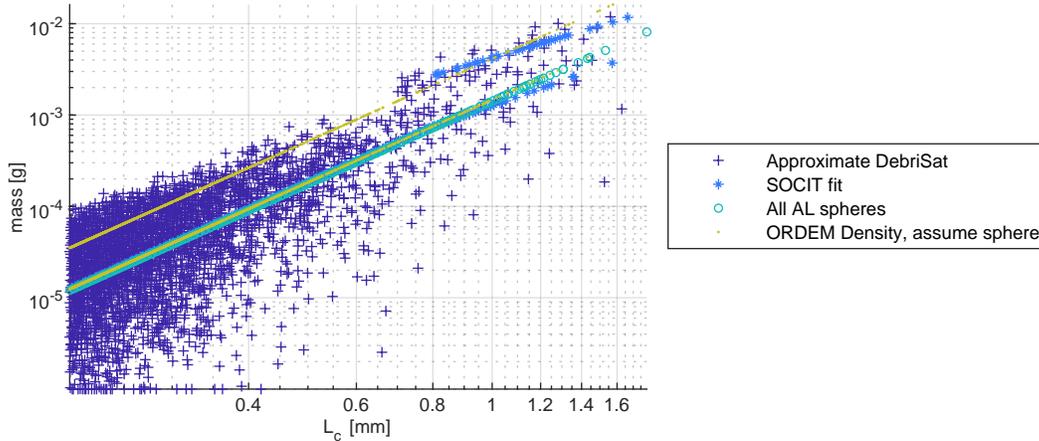


Figure 5.6: Characteristic length to mass mapping with various models

5.3.3.5 Momentum Enhancement Factor

As discussed in Section 1.4.4.3, the momentum enhancement factor for hypervelocity impacts against spacecraft materials is not well known. Therefore, several options are provided for how to calculate momentum enhancement.

- A. **Static MEF:** The NESC report bounds the results with an MEF of 1 and an MEF of 3, this code offers the option to set the MEF to 1, 2, or 3. Using 1 and 3 to bound a debris strike is reasonable, or an approximate median value of 2 can be used.
- B. **Nysmith-Denardo MEF:** Some efforts indicate a velocity dependence to MEF, although this is difficult to test for high velocities and macro-sized projectiles. As discussed in Section 1.4.4.3, there is reason for skepticism about the existing velocity-dependent MEF relationships because they disregarded the set of data indicating a higher MEF and based the model on only the data indicating a lower MEF. This is the lower of the two relationships from McDonnell's paper [49], with a maximum MEF around 10 at 68 km/s. Subsequent work has shown that some of the assumptions underlying these lower estimates were erroneous, and would indicate a higher MEF.

$$\beta = 1 + 0.144 \frac{(v_{\text{rel}} - .55)^2}{v_{\text{rel}}} \quad (5.12)$$

- C. **Rembor's MEF:** This is the higher of the two relationships from McDonnell's paper [49], with a maximum MEF around 20 at 68 km/s. This takes into account measurements from a microparticle accelerator to obtain data at around 25 km/s, but may underpredict MEF at lower velocities somewhat.

$$\beta = 1 + 0.3(v_{\text{rel}} - 2) \quad (5.13)$$

- D. **Rembor with solar array:** This relation applies the Rembor equations, but when the strike impacts a solar array, the MEF is 0.5. This is approximately in the ball park for a small sphere against a typical array [14], but the exact value will depend strongly on

a number of parameters, such as how close the piece of debris is to the ballistic limit of the solar array. This is not a well-substantiated value, it is a rough estimate based on one simulation for a parameter that hasn't be adequately measured in test and is difficult to simulate effectively.

The true relationship for MEF may well be discontinuous, with different velocity regimes fitted by different functions, similar to the ballistic limit equation. It will likely change depending on the ballistic limit of the particular projectile and impact location, where a particular velocity doesn't produce a particular MEF, but involves a complex relationship between the projectile and the target. Projectiles that are close to but don't exceed the ballistic limit will likely have the highest MEF against a penetrable structure like a solar array. The intention of this exercise is not to provide a correct answer, which is unknown, the intention is to illustrate how the model output is affected by uncertainty in these parameters. One concept on the horizon is applying machine learning approaches to replace the legacy ballistic limit equations for predicting the outcome of hypervelocity impact events [84]. With modern and emerging computational capabilities, substantial improvement on heritage methods is likely possible, but is still a nontrivial development due to limitations in test capabilities coupled with a problem that is inherently highly variable.

5.3.3.6 Strike Location on Spacecraft

The final debris strike parameter is determining the location of the strike on the spacecraft. Prior efforts assume that all spacecraft are 1 m^2 [6], or that strikes only occur on the bus body [1], or construct a detailed CAD-like model of the spacecraft. For this architecture, a middle ground between these methods is appropriate. Treating all spacecraft as 1 m^2 will substantially change the fluxes experienced by the spacecraft: a cubesat will see dramatically fewer impacts while a large comm satellite will see more. Since the solar arrays are generally the largest surface on the spacecraft ignoring those strikes will reduce the amount of data collected significantly, even accounting for HVI impact potentially breaking through a solar array such that the MEF is less

than one. On the other side of the spectrum, implementing CAD-like spacecraft model details would increase the time required to conduct Monte Carlo/trade studies and substantially increase the overhead work to analyze different mission profiles. Therefore, the strike location is determined by defining a point cloud based on simple spacecraft parameters: bus dimensions, solar array length and width, and number of solar arrays.

This 3D point cloud is derived to represent each spacecraft surface with an array of points at a specified grid spacing based on the spacecraft size specifications. The solar arrays are rotated and the body frame of the spacecraft is aligned relative to the orbit frame in accordance with the specified spacecraft configuration. The process for defining this point cloud and randomizing the strike location is as follows.

1. Define point cloud on spacecraft surfaces: Use specified bus and solar array dimensions and configuration to generate a 3D point cloud specifying the bus surfaces. The desired grid spacing is specified, then the generated bus grid is slightly smaller than the desired projected grid (see step 2) to ensure that each point in the projected grid is populated correctly and none are missed. Then the grid spacing is shrunk slightly if necessary to ensure that the edges of the bus point cloud correspond to the exact edges of the bus using an array of points spaced evenly across each bus surface. The same concept is applied to the solar arrays, taking into account the rotation of the solar arrays specified in the spacecraft configuration function, to form the point cloud of solar array points, shown in yellow.

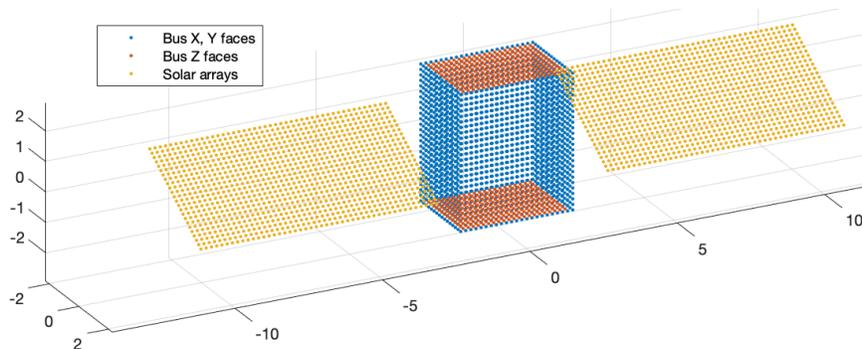
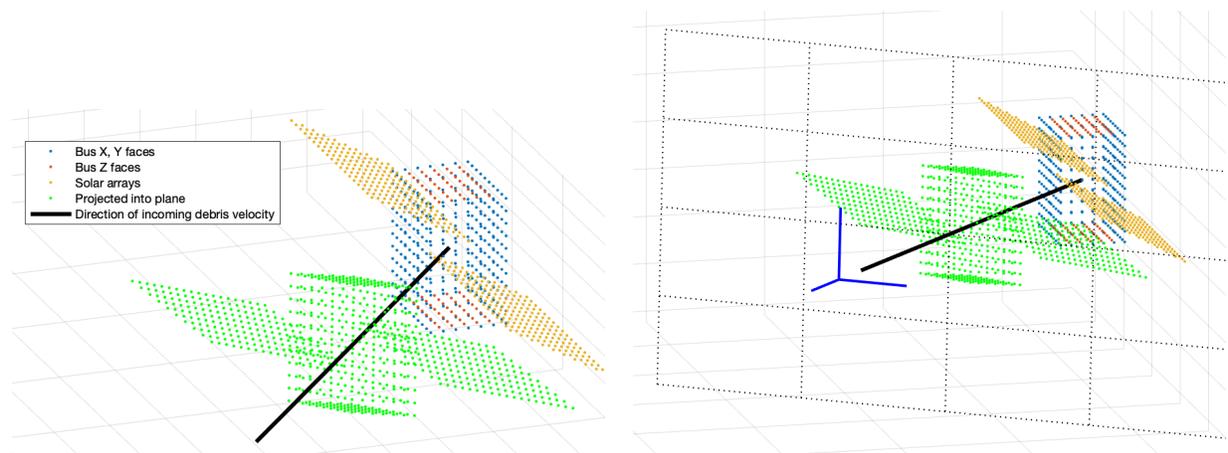


Figure 5.7: Step 1: Bus point cloud specifying surfaces

2. Project bus point cloud onto plane perpendicular to the incoming debris:

The debris is equally likely to hit anywhere on the bus body, *as viewed from the direction of incoming velocity*. Therefore, in order to randomize the strike point the silhouette of the spacecraft is determined in a plane normal to the velocity direction, \hat{v} , which is the velocity of the debris relative to the spacecraft expressed in the body frame and normalized to unit length. Each point in the 3D point cloud is projected into this plane, producing the points in the projected plane, which are shown in green.

3. Define frame and grid for projection plane: The velocity frame is denoted as $\mathcal{V} = \{O, \hat{v}_1, \hat{v}_2, \hat{v}_3\}$, where \hat{v}_1 is the negative velocity unit vector, \hat{v}_2 is aligned with the X-Y plane (in the body frame) and orthogonal to \hat{v}_1 , and $\hat{v}_3 = \hat{v}_1 \times \hat{v}_2$. \hat{v}_2 is referred to as the ‘Q-dimension’ for the projected plane, and \hat{v}_3 is referred to as the ‘R-dimension’.



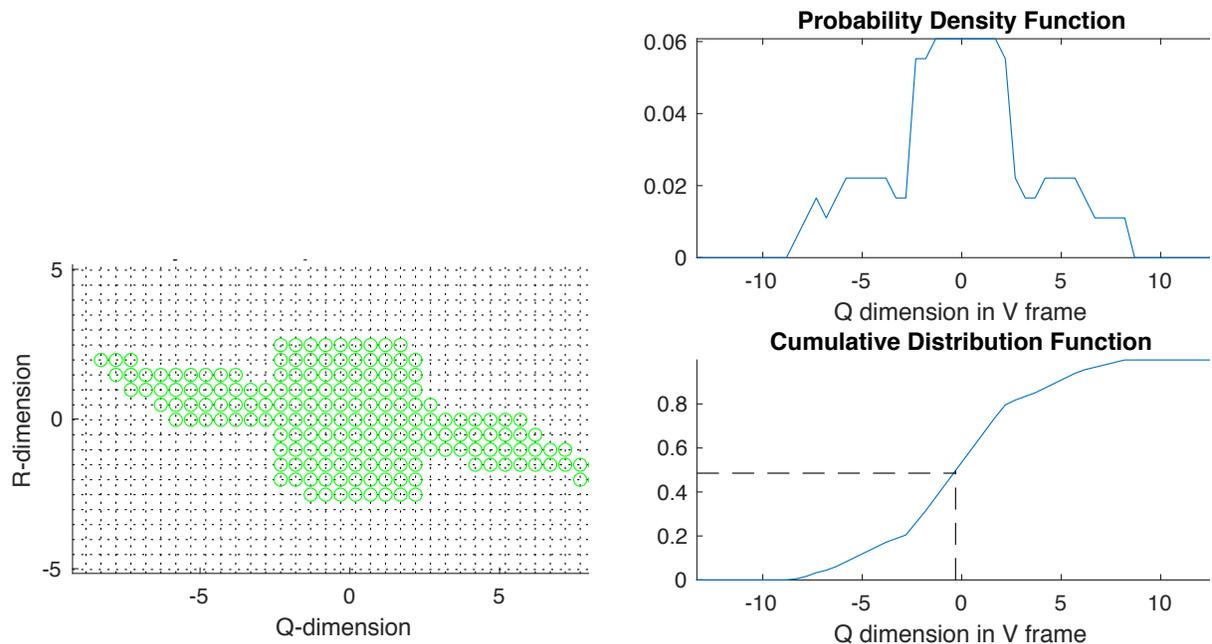
(a) Step 2: Bus points projected onto plane normal to incoming velocity vector. (b) Step 3: Basis vectors for \mathcal{V} frame in blue, projection plane in black

Figure 5.8: Step 2 and 3 of finding strike location on spacecraft

4. Determine which grid points have projected points: The projection plane is now subdivided into a grid that surrounds the spacecraft silhouette, with the grid spacing at exactly the specified grid spacing, no longer truncated. Each point from the 3D spacecraft point cloud is mapped to a specific index in the QR plane, which is flagged as a potential strike location. In other words, the segments of the QR plane which can ‘see’ the spacecraft are now identified, and

the projected bus area is now normalized into the frame of the incoming debris, so that the actual location that is struck can now be found without a bias toward the bus surfaces which are at an angle toward the incoming debris. To illustrate, note the top and bottom surfaces of the bus in step 3, which have a higher concentration of points, but are normalized to an equal hit probability in step 4.

5. Determine pdf and CDF of Q-dimension from silhouette: Each ‘slice’ of the QR plane is summed to indicate the number of segments in that slice which can ‘see’ the spacecraft. This indicates the areal density of spacecraft bus within that slice of the QR plane. This summation is normalized into a pdf and the cumulative sum is taken to produce a CDF. A random number $\sim \mathcal{U}(0, 1)$ is drawn to determine which Q slice is struck using inverse transform sampling. A second random number, this time an integer, is drawn to determine which of the allowable R indices in that Q-slice is struck, i.e, which out of the R indices in the selected Q slice that can ‘see’ the spacecraft. This defines the strike location in the QR plane.



(a) Step 4: Projected bus points mapped to bus silhouette in \mathcal{V} frame, now 2D (b) Step 5: pdf and CDF showing randomly selected Q-index

Figure 5.9: Step 4 and 5 of finding strike location on spacecraft

6. Finalize bus strike location: Multiple bus point cloud points may map to the selected QR location; the forwardmost point (towards the incoming velocity) is selected as the strike location and its position in the body frame is selected as $\mathbf{R}_{d/s}$, the vector expressing the location of the strike relative to the center of mass of the spacecraft. The type of strike location, bus or solar array, is later used to determine the MEF. In the figure, the green stars indicate 3D point cloud points which map to the selected QR location, the forwardmost point is selected as the final strike location and the vector $\mathbf{R}_{d/s}$ is shown as a dashed blue line.

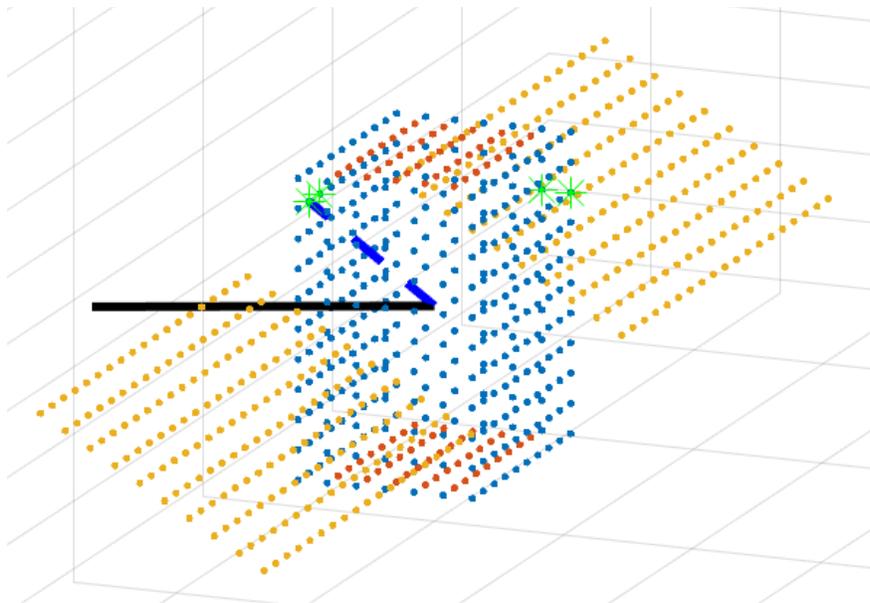


Figure 5.10: Step 6: Identifying final strike location on spacecraft

5.3.4 Implementing Trades, Assessing Results

With these methods for drawing the strike parameters from specified distributions, the trade architecture is implemented as follows:

1. Set up traded variables. Any two variables can be traded any number of times, out of the following: spacecraft size, orbit, MEF model to use, characteristic length to mass conversion model to use, flux data to use, and grid spacing (for calculating the strike location). For relatively clear results, one variable is typically traded at 3 values and a second variable is

typically traded with 2 values, as shown in Section 5.4.

2. Load spacecraft and orbit data as specified in the traded variables, set up arrays and files to store Monte Carlo and trade results
3. Run Monte Carlo for each trade (i.e, trading 3 values of one variable and 2 values of another variable results in 6 separate Monte Carlos to perform, store, and analyze).
 - 3.1. Establish spacecraft configuration. For these trades, the spacecraft body frame is aligned with the orbit frame with the solar array rotation angle randomized between zero and $\frac{\pi}{2}$, but different mission configurations could be defined in relation to the orbit and inertial frames. The mean anomaly of the orbit is randomized between 0 and 2π
 - 3.2. Calculate debris characteristic length, velocity magnitude, and velocity vector by drawing random numbers vs. the ORDEM flux data for this orbit
 - 3.3. Determine the strike location on the spacecraft
 - 3.4. Calculate mass and momentum enhancement factor based on the specified models and previously determined characteristic length, velocity, and strike location (i.e, bus or solar array)
 - 3.5. Calculate effect of strike on spacecraft, i.e, change in rotation and translational momentum and associated measurable parameters.
4. Calculate the total expected flux per year. The architecture truncates the ORDEM flux data at a specified characteristic length to avoid applying excessive numbers of small strikes that do not perturb the spacecraft sufficiently for detection. Therefore, the total ORDEM-predicted flux in particles-per-meter-square-per-year, f_{ORDEM} , can be calculated at the size range considered. Using an estimate of the spacecraft's effective area, A_{eff} , the number of particle hits expected per year, $R_{\text{strikes/yr}}$, is calculated.

$$R_{\text{strikes/yr}} = A_{\text{eff}} f_{\text{ORDEM}} \quad (5.14)$$

From this, the Monte Carlo data is used to predict the number of expected strikes at a certain detectability threshold or above. This provides the expected perturbation rate for the current trade case, as a function of detectability threshold. This is shown in the summarized results, to explore the changes in perturbations expected as a function of spacecraft and model used.

For item 3.1, note that since ORDEM is an orbit-averaged flux this doesn't fully replicate actual spacecraft behavior. In reality more strikes would be expected near the poles than near the equator. Note that this caveat (orbit-averaged flux) applies to the publicly available version of ORDEM, with assistance from ODPO more resolution on things like flux variation with orbit location may be obtainable. If fitting expected flux to on-orbit experiences things like orbit location may be relevant and a more sophisticated set of ORDEM data could be implemented to capture these variations.

The traded variables in the results below are as follows, but note that the architecture is agile to sub in additional models for each of these variables fairly quickly.

- Spacecraft size: NASA's Fermi spacecraft (4000 kg), generic LEO smallsat (200 kg), 6U cubesat (12 kg).
- Spacecraft orbit: Two orbits are used for the sample results, a circular orbit at 550 km and 28.5 degrees inclination, and a circular orbit at 800 km and 90 degrees inclination, both using ORDEM 3.1 from 2020. However, any orbit from LEO to GEO can be used, each desired orbit just needs to be run in ORDEM to generate the raw flux data, then the analysis can be run using that data.
- Momentum enhancement factor model: options are described in Section 5.3.3.5.
- Characteristic length-to-mass conversion model: options are described in Section 5.3.3.4.
- Flux data to use: This can be either ORDEM, ORDEM flux plus $2\times$ the specified σ , or ORDEM flux minus $2\times$ the specified σ , since ORDEM also provides σ values for each

bin expressing the uncertainty in modeled flux. In theory, this could potentially pull data from the micrometeoroid environment model MEM3 or the European flux model MASTER which would be an interesting comparison study, but these have different interfaces and the necessary conversion architecture has not yet been designed.

- Grid spacing: This defines the grid spacing for calculating the strike location. This can be set arbitrarily small but run times increase substantially, so the default is to use 1/6th of the smallest bus/solar array dimension which provides adequate resolution without excessively long run times.

5.4 Sample Results

These results are included to illustrate the capabilities of the tool and provide general guidance on the expected perturbation rates. Additional refinement to models would be required to produce more precise results. For example, the effective spacecraft area used to calculate the results is approximate, calculated as

$$A_{SA, \text{eff}} = \frac{1}{2}(L_{SA}W_{SA}) \quad (5.15)$$

$$L_{c, \text{bus}} = \frac{1}{3}(X_{\text{bus}} + Y_{\text{bus}} + Z_{\text{bus}}) \quad (5.16)$$

$$A_{\text{bus}, \text{eff}} = L_{c, \text{bus}}^2 \quad (5.17)$$

$$A_{SC, \text{eff}} = A_{\text{bus}, \text{eff}} + A_{SA, \text{eff}} \quad (5.18)$$

Monte Carlos indicate that the true value of effective area is higher, but establishing a precise value requires a very small grid spacing, which makes the computation time excessively long, so this approximation is used. The following trades are performed to explore the tradespace using this tool:

- Trading spacecraft size and orbit: This trade investigates which classes of spacecraft (small, medium, large) can detect higher perturbation rates via translational and rotational momentum changes in two different orbits: the 800 km orbit with a fairly high flux of small

debris from two major debris-causing incidents, and the less severe 550 km orbit.

- Trading orbit and uncertainty bounds: This trade examines the uncertainty specified in ORDEM and its effect on overall perturbation rate
- Trading MEF and L_c -to-mass conversion: This trade investigates the effect of difficult-to-define parameters on overall perturbation rates

5.4.1 Trading spacecraft size and orbit

Based on the results in the spacecraft parameter trade study in Section 2.5 and the analysis in [6], it is anticipated that smaller spacecraft might have lower thresholds for detecting perturbations, depending somewhat on their instrument suite and the filter applied to identify strikes. However, a smaller spacecraft will also experience fewer debris strike due to its smaller collection area, and for rotational change detection a smaller spacecraft will generally have a smaller lever arm at which the debris strike occurs, and thus experience less of a change in rotational momentum.

In order to investigate the relationships between these variables, this first trade examines three spacecraft in two different orbits. The small spacecraft is a 6U cubesat, 12 kg, with bus dimensions 20x30x10 cm and two deployed solar arrays which are 20x30 cm. The medium spacecraft is 200kg, has a 1 m³ bus, and two 1mx2m solar arrays. The large spacecraft, approximating Fermi, is 4000kg, with a 2.5x2.5x2.8 m bus, and two 2x6m solar arrays. The orbits are 800 km, which is one of the most hazardous orbits in terms of non-trackable debris thanks to two large debris-causing events, and the 550 km orbit, which was previously a more benign orbit prior to the Russian ASAT test. These debris flux populations are from ORDEM 3.1 in 2020, so they do not include the Russian ASAT test.

Figure 5.11 shows the scalar change in rotational momentum ΔH plotted relative to the scalar ΔV that occurs for each debris strike in each of 10,000 Monte Carlo runs for each of 6 trade cases. As shown, there are three distinct populations of strikes for each of the three spacecraft. The largest spacecraft, in blue, sees substantially less change in velocity as its mass is orders of

magnitude larger. However, it also has a longer lever arm, so it sees more of a change in angular momentum. Note, of course, that these are plotted on log-log axes, so even a fairly small change is actually substantial. The smallest spacecraft, in contrast, sees a significant fraction of strikes which can be detected by a change in velocity (as indicated by the horizontal black line), but few strikes which could be detected by a change in angular momentum (vertical black line). This assumes that the spacecraft have the same detectability thresholds, which they might not, depending on the specific sensor suite which each spacecraft flies (or the quality of its external navigation solution, for detecting orbit changes).

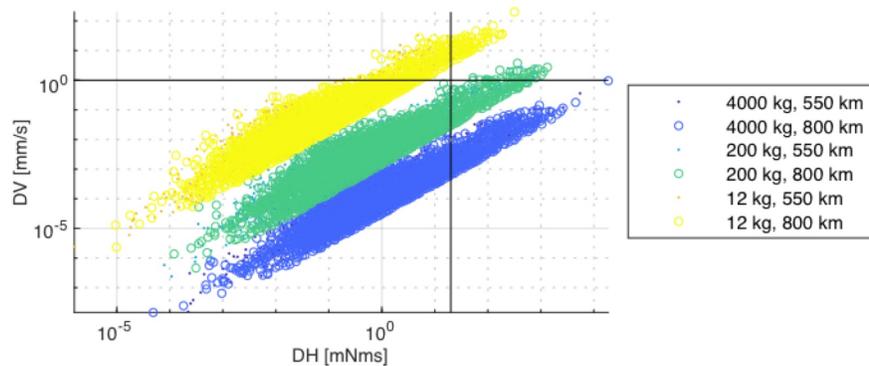


Figure 5.11: Plotting results from trade study, indicating relationship between change in angular and linear momentum

Note that the population from the 800 km orbit and the population from the 550 km orbit are comparable on this plot. This is because this shows all applied Monte Carlo runs, and while the flux at 800 km is larger than the flux at 550 km, but this is not captured as the same number (10,000) strikes are plotted at each altitude. However, the small spacecraft would take many times longer to accumulate this number of strikes, because of its smaller surface area.

Figure 5.12 addresses this by plotting the annual perturbation rate expected on each spacecraft. The perturbation rate is plotted on the vertical axis while the detection threshold is plotted on the horizontal axis, allowing a mapping where if a spacecraft had a detection threshold identified on the horizontal axis, it would be expected to see the number of annual perturbations plotted on the vertical axis. Now, with the Monte Carlo results normalized to account for the smaller size of

the spacecraft and differences in flux between the two orbits, there are clear differences between the cases investigated.

Figure 5.12(a) shows that when linear momentum transfer is detected via orbit changes, the 550 km orbit clearly has lower perturbation rates than the 800 km orbit. Interestingly, the interplay between the larger mass resulting in less total ΔV but the larger spacecraft area resulting in more strikes appears to almost cancel out, and the differently-sized spacecraft in each orbit see roughly similar numbers of strikes compared to the spacecraft in the other orbit. This graph is a little choppy due to the limitations of the Monte Carlo. The lower boundary for the characteristic length is set to $3E-5$, and since the Monte Carlo draws proportionate to the debris population many more small pieces are drawn than large pieces. With 10,000 Monte Carlo runs the largest spacecraft sees barely any strikes that are detectable at these thresholds, which is why the blue lines are choppy.

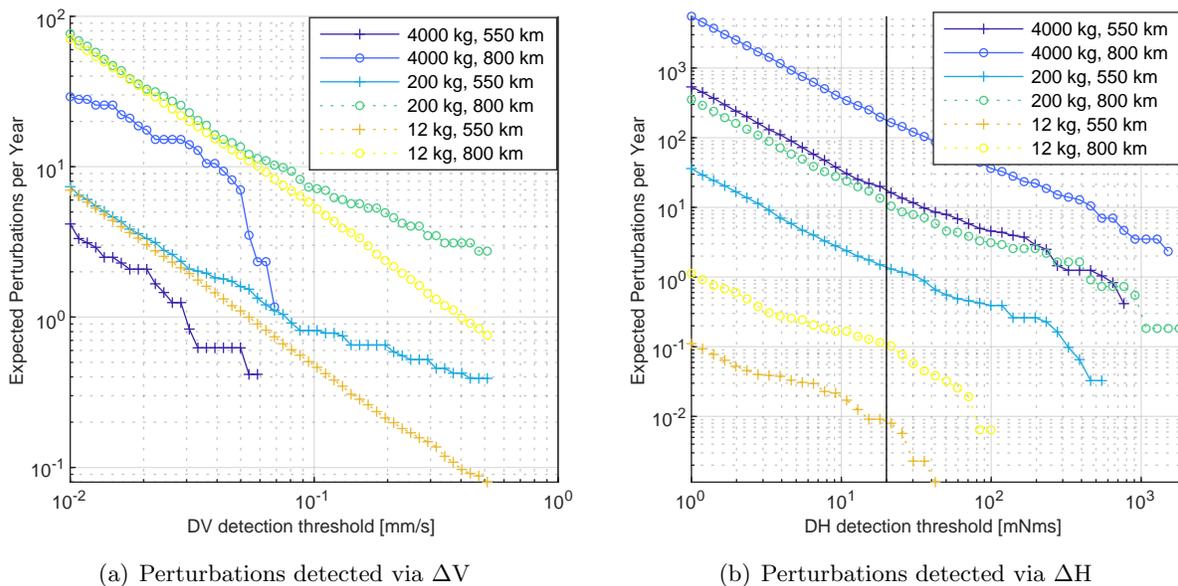


Figure 5.12: Comparing annual perturbation rate for each spacecraft and each detector

Figure 5.12(b) shows the comparable graph, but for strikes detected via angular momentum change, ΔH . With this detection method the largest spacecraft shows substantially more detectable strikes due to its longer lever arm, as it both sees a stronger effect from the strikes and sees more

strikes due to having a larger surface area. Again, this assumes that the detection thresholds are the same between spacecraft, although this graph can also show that a 200 kg smallsat at 550 km with a threshold of 20 mNms would see about the same number of perturbations as a large 4000 kg satellite in the same orbit with a 400 mNms threshold. In an odd happenstance, a 4000 kg spacecraft in the less hazardous 550 km orbit sees a nearly identical perturbation rate to a smaller 200 kg spacecraft in the more hazardous 800 km orbit. With these more easily detectable rotational perturbations, the lines are smoother than the ΔV plot.

Overall, the takeaway from these figures is that the rotational change detection algorithms are a good fit for a large spacecraft and would likely generate a useful population of data, likewise for a constellation with large numbers of satellites in the 200 kg range. Applying them to a cubesat is likely not worthwhile, unless very low detection thresholds are attainable and large numbers of satellites are flown. However, for algorithms detecting ΔV the detection thresholds modeled here would not be expected to show large number of strikes, unless lots of spacecraft were available with high-precision orbit determination solutions. As space traffic management and space situational awareness capabilities improve this might become possible, some large constellations might be nearing a useful capability here depending on the accuracy of their orbit determination.

5.4.2 Trading ORDEM uncertainty bounds and orbit

This trade is intended to investigate the uncertainty inherent within ORDEM. ORDEM publishes a σ value for flux along with the calculated flux values, so the uncertainty in the model can be seen. For this trade, the same two orbits are investigated with the 200 kg spacecraft, but instead of just taking ORDEM's baseline flux the flux $+2\sigma$ and the flux -2σ are used as well. Thus, the variation in the perturbation rate can be seen as a function of the uncertainty in the environmental flux model. The results are about as expected, with the baseline flux bounded by the upper and lower uncertainty bounds, although the lower uncertainty bound grows significantly for larger perturbations, which is interesting given the mismatch in the NESC report between the perturbations expected and the perturbations experienced. Toward the right hand side of the plot,

as per usual, results get choppy as they become highly dependent on a small number of Monte Carlo draws. It would appear that there is a high degree of confidence that the flux is not substantially higher than what is modeled in ORDEM, but it might be quite a lot lower. This is likely due to the difficulties in obtaining data at this size regime and altitude. For comparison, it would be interesting to see if ESA’s MASTER provided a perturbation rate that is within ORDEM’s error bounds when run against this problem.

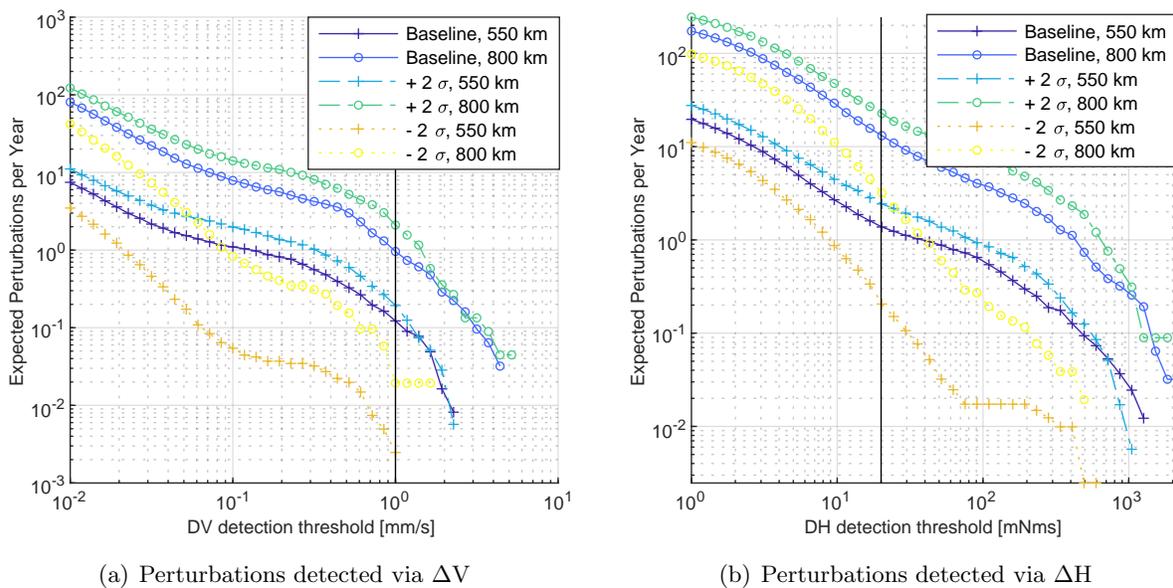


Figure 5.13: Comparing annual perturbation rate with ORDEM’s error bounds

5.4.3 Trading MEF and L_c -to-mass conversion

After investigating the relationships between spacecraft sizes and orbits and ORDEM uncertainty bounds, the inherent uncertainty in perturbation rates due to uncertainty about model parameters is investigated. In this trade, the models for calculating the MEF and the characteristic length-to-mass conversion are investigated. Figure 5.6 showed the distribution of L_c to mass relationships using the various models. Figure 5.14 shows the MEF relations and the results from the trades, based on a Monte Carlo run for each trade.

Figure 5.14(c) shows each of the relations between MEF and velocity. Recall that these

are the ones available in the literature, but reality is probably more complicated, discontinuous, and probabilistic, especially for debris objects against spacecraft materials. The yellow shows the Rembor relation adjusted to reduce the MEF to 0.5 for all solar array strikes, but in reality some strikes would not penetrate the array and would maintain an MEF closer to the Rembor fit, while others punch through easily and deliver less MEF.

Figure 5.14(a) shows the rate of detectable perturbations at various thresholds for linear momentum transfer, the main takeaway is that these model uncertainties cause results to vary by a factor of 5 or so, and a high MEF results in the highest perturbation rates.

Figure 5.14(b) shows a similar plot but with angular momentum transfer. This plot shows a spread of about 10X due to these uncertainties. Unsurprisingly, this plot shows the Rembor relation as the highest perturbation rate, since it has the highest MEF (light blue family of curves). However, when the Rembor relation is modified to set all solar array impacts to an MEF of 0.5, it becomes the lowest rate of perturbations (yellow family). The Nysmith-Denardo relation (green family) shows marginally higher perturbation rates relative to using a constant MEF of 2, since it is effectively averaged by the constant MEF of 2. It is slightly higher because faster velocities are more common in this orbit, which biases the MEF upward.

Comparing the results from the various mass assumptions, the marker styles indicate that using the ORDEM density classes and assuming spheres results in the highest estimates ('+' sign), while assuming that all spheres are all modeled as aluminum ('x' sign) results in the lowest estimates for a given MEF assumption. This is a little surprising, and indicates that the high density population is a significant contributor to the perturbation rate, as it increases it by a factor of 2 or so. This is relevant because the high density population is a major contributor to spacecraft risk, so accurately modeling its contribution is imperative for accurate risk assessments. Overall, the takeaway is that these uncertain model parameters are significant in correctly predicting perturbation rates.

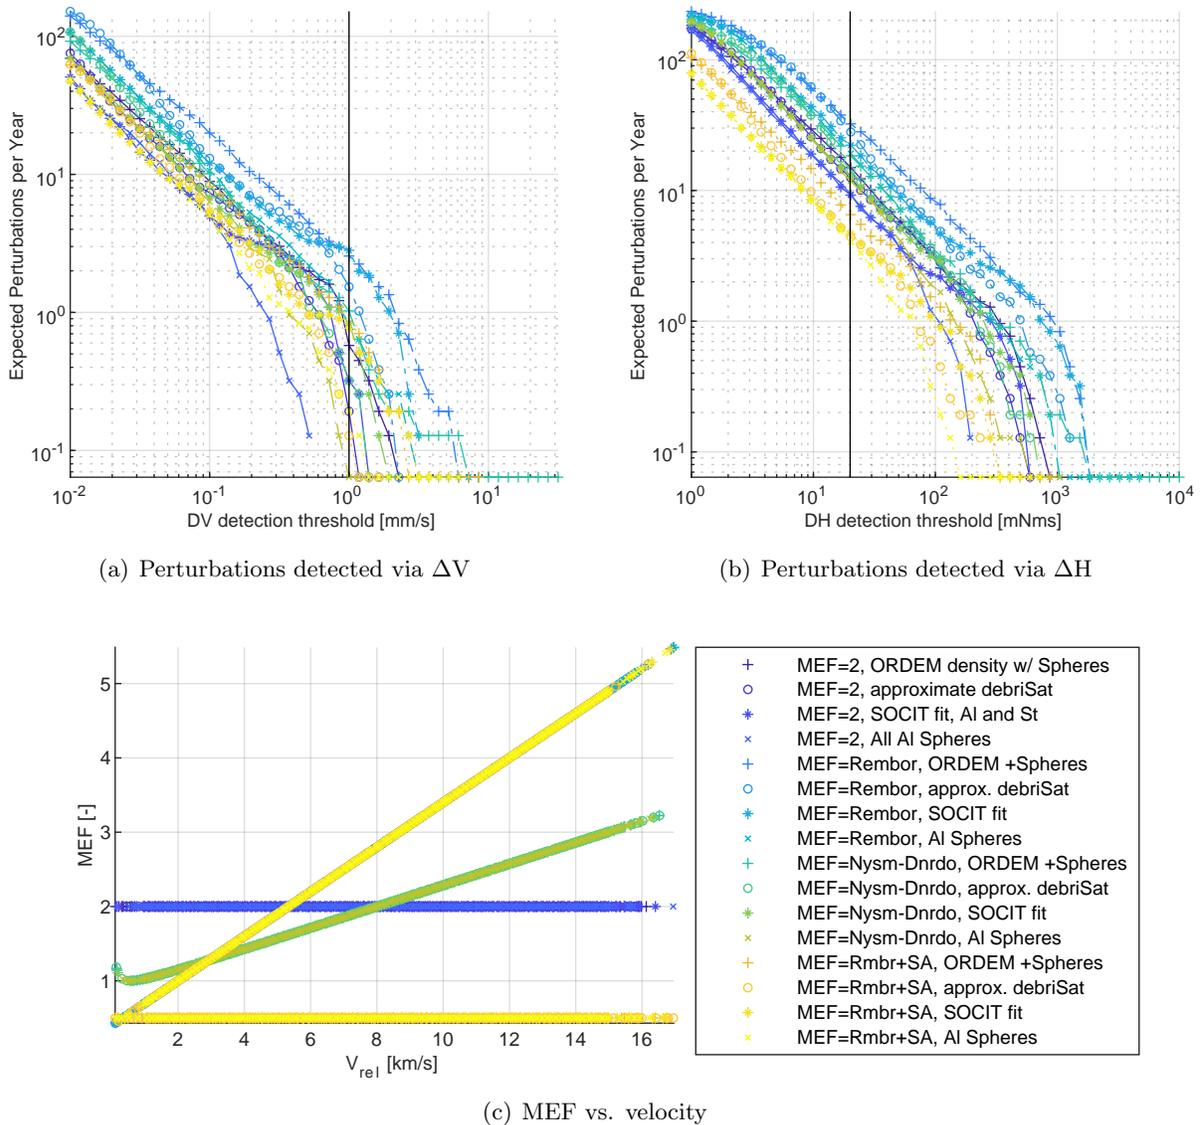


Figure 5.14: Comparing annual perturbation rate with various MEF and mass assumptions

5.5 Discussion

This tool demonstrates the variability in perturbation rate due to a variety of factors. The orbit and size of the spacecraft can cause widely varying differences in the rate of expected perturbations, but the uncertainty in ORDEM is also significant, especially when considering larger perturbations. This tool represents a preliminary solution to investigate the problem, it would be

interesting to add the micrometeoroid model MEM and compare to the European model MASTER, but since the three tools have completely different interfaces meshing them is a nontrivial undertaking. A 2015 paper compared the results of ORDEM and MASTER across a couple of design cases [85], but this was written by the teams which develop the tools as comparing the output is nontrivial due to the tools' significant structural differences. It is also worth noting that MEM provides the average sporadic micrometeoroid flux, so capturing variations like the ones seen in the SDO data would not be attainable with MEM.

Of course, ideally this tool would be used to compare to a measured population of perturbations on an active system. If multiple systems were contributing significant data sets lessons could be learned about the relationship between on-orbit experiences and model predictions, potentially leading to a better understanding of things like MEF and helping to resolve some of the uncertainties driving overly conservative spacecraft risk assessments, without underpredicting the risks as increasing numbers of large constellations are licensed and flown.

In making these comparisons, it is important to remember that ORDEM is the most sophisticated and data-rich model in the world. However, it is fundamentally limited by data collection systems, so improving data collection especially in undermeasured orbits can provide additional insight into debris populations. Validation between measurement types, as opposed to the typical validation between different measurement series of the same type, also offers a unique opportunity to gain insight into relevant aspects of the hazardous non-trackable debris population.

Chapter 6

Conclusion

Based on the recommendation in the 2017 NASA report [1], this research investigates methods to detect abrupt, unexpected momentum transfer to identify minor debris strikes on active spacecraft. Several methods to identify strikes are developed and compared, then the methods are applied to NASA spacecraft to identify performance capabilities and investigate challenges, and a tool is developed to investigate the relationship between debris environment models and data that could be collected using these methods.

Some of the key findings and contributions of this research are as follows:

- Digital signal processing techniques, specifically a matched filter, can be applied to the spacecraft rate and attitude telemetry to identify perturbations like the one seen in the Sentinel-1A debris strike.
- However, applying a change detection algorithm to the inertial angular momentum is generally even more effective for identifying subtle perturbations due to debris strikes.
- Changes in translational momentum can be detected using navigation telemetry, and applying post-processing techniques to a navigation filter improves the ability to detect these subtle, abrupt changes.
- Applying these techniques to simulated data demonstrates a strong capability, but applying them on real telemetry is much more challenging

- While most spacecraft collect telemetry from which their inertial angular momentum can theoretically be derived, in practice most have several idiosyncrasies which complicate the application of these techniques. Sometimes additional processing can be added to accommodate these idiosyncrasies and still obtain useful data.
- One limiting factor common across several spacecraft seems to be the estimate of the spacecraft inertia. A spacecraft's typical operations are fairly robust to inertia uncertainties, but the inertia uncertainties affect the calculation of inertial angular momentum making it difficult to apply change detection algorithms effectively. The path forward is to improve estimation of the spacecraft's inertia, and designing a filter to do so is the recommended next step for developing this capability. Reference [83] looks like a promising approach for this.
- Setting detection thresholds is an important element of applying these techniques, and must be performed carefully. If the threshold is set too high detectable strikes may be missed, while if the threshold is set too low abundant false alarms will occur. The desired behavior will be dependent on the goals of the application, a higher threshold gives more confidence that detections are correct while a lower threshold may provide additional data for studies.
- Two known debris strikes on the MMS mission were successfully identified, and data returned by the Solar Dynamics Observatory (SDO) indicates small perturbations with annual variations that are somewhat consistent with micrometeoroid flux variations, although the magnitude of the events seems larger than expected. The filters also identified instances of abrupt momentum transfer on SDO consistent with debris strikes, demonstrating that the detection algorithms are working effectively.
- The findings from these applications are broadly consistent with other efforts, indicating that applying concepts like this is more challenging than it may seem, and it is difficult to definitively assess whether a feature in telemetry is a debris strike or just noise.

- A tool is developed to explore the potential use cases for these techniques. The tool assesses the expected perturbation rate based on information from debris environment models. This can be used to identify potential contributions from systems by using these algorithms, i.e determining which systems could potentially return useful amounts of data. It also offers the ability to investigate the effect of underlying assumptions on model predictions, showing a path toward using perturbation data to improve assumptions and reduce uncertainty in applying models, which would improve spacecraft risk assessments.

With these developments significant progress has been made toward realizing the recommendation of the 2017 NASA report and developing a capability to use active satellites as *in situ* debris detectors, without having to fly any extra hardware. Accurate debris models and risk assessment processes are required in the emerging space era so that operators and governments worldwide can make good decisions for space sustainability. A sophisticated, validated understanding of the debris environment is necessary to build consensus and motivate appropriate legislation, operations practices, and technology development to enable a sustainable future for the global space enterprise.

Bibliography

- [1] **Evaluation of Micrometeoroid and Orbital Debris (MMOD) Risk Predictions with Available On-orbit Assets**. NASA Engineering and Safety Center Technical Assessment Report NESC-RP- 14-01000. NASA Engineering and Safety Center, Oct. 2017.
- [2] H. Krag et al. “A 1 cm space debris impact onto the Sentinel-1A solar array”. en. In: **Acta Astronautica** 137 (Aug. 2017), pp. 434–443. ISSN: 00945765. DOI: 10.1016/j.actaastro.2017.05.010. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0094576517304125> (visited on 09/11/2019).
- [3] NASA ODPO. “The Intentional Destruction of Cosmos 1408”. In: **Orbital Debris Quarterly News** 26.1 (Mar. 2022). URL: <https://orbitaldebris.jsc.nasa.gov/quarterly-news/pdfs/odqnv26i1.pdf>.
- [4] Ned Price. **State Department Press Briefing - November 15, 2021**. Tech. rep. US Department of State, Nov. 2021. URL: <https://www.state.gov/briefings/department-press-briefing-november-15-2021/>.
- [5] **Space Debris by the Numbers**. URL: https://www.esa.int/Our_Activities/Space_Safety/Space_Debris/Space_debris_by_the_numbers.
- [6] Joel Williamsen et al. **Characterizing the Orbital Debris Environment Using Satellite Perturbation Anomaly Data**. Tech. rep. IDA Document NS D-10643. Institute for Defense Analyses, June 2019.
- [7] **NASA Orbital Debris Program Office, Debris Measurements**. US Government. URL: <https://orbitaldebris.jsc.nasa.gov/measurements/>.

- [8] James Shell. “Optimizing orbital debris monitoring with optical telescopes”. In: **Advanced Maui Optical and Space Surveillance Technologies Conference**. Maui, HI, Sept. 2010.
- [9] A Manis et al. “The Updated GEO Population for ORDEM 3.” en. In: **First International Orbital Debris Conference**. Sugarland, TX, Dec. 2019, p. 10.
- [10] **Orbital Debris Quarterly News**. Jan. 2009. URL: <https://orbitaldebris.jsc.nasa.gov/quarterly-news/pdfs/odqnv13i1.pdf>.
- [11] Ryan Shepperd and Kristy DiOrio. “The History of Satellite Collision Assessment and Risk from the Iridium Perspective”. en. In: **SSA Operators’ Workshop**. Boulder, CO, 2022, p. 13.
- [12] James Murray et al. “Analysis of the DebrisSat Fragments and Comparison to the NASA Standard Satellite Breakup Model”. en. In: **First International Orbital Debris Conference**. Sugarland, TX, Dec. 2019, p. 10.
- [13] Heather M. Cowardin et al. “Optical Characterization of DebrisSat Fragments in Support of Orbital Debris Environmental Models”. en. In: **The Journal of the Astronautical Sciences** 68.4 (Dec. 2021), pp. 1186–1205. ISSN: 0021-9142, 2195-0571. DOI: 10.1007/s40295-021-00278-9. URL: <https://link.springer.com/10.1007/s40295-021-00278-9> (visited on 06/02/2022).
- [14] Conor Ryan. “Momentum Transfer due to Hypervelocity Impacts into Spacecraft Solar Arrays”. en. PhD thesis. TU Del, Nov. 2021.
- [15] James Chinn and Martin Ratliff. “Challenges of Debris-Impact Risk Assessment for Robotic Spacecraft”. en. In: **Procedia Engineering** 204 (2017), pp. 437–444. ISSN: 18777058. DOI: 10.1016/j.proeng.2017.09.738. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1877705817342923> (visited on 05/18/2020).
- [16] Dana M Lear, Eric L Christiansen, and James L Hyde. “BUMPER: A TOOL FOR ANALYZING SPACECRAFT MICROMETEOROID AND ORBITAL DEBRIS RISK”. en. In:

- NASA JSC ARES HVIT Reference Document** (Dec. 2020), p. 10. URL: [https://ntrs.nasa.gov/api/citations/20205011690/downloads/Bumper%20-%20A%20Tool%20For%20Analyzing%20Spacecraft%20MOD%20Risk%20\(paper\).pdf](https://ntrs.nasa.gov/api/citations/20205011690/downloads/Bumper%20-%20A%20Tool%20For%20Analyzing%20Spacecraft%20MOD%20Risk%20(paper).pdf).
- [17] Michael Squire. **EVALUATING MICROMETEOROID AND ORBITAL DEBRIS RISK ASSESSMENTS USING ANOMALY DATA**. Tech. rep. NASA Langley Research Center.
- [18] T. Maclay and D. McKnight. “Space environment management: Framing the objective and setting priorities for controlling orbital debris risk”. en. In: **Journal of Space Safety Engineering** 8.1 (Mar. 2021), pp. 93–97. ISSN: 24688967. DOI: 10.1016/j.jsse.2020.11.002. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2468896720301415> (visited on 06/08/2022).
- [19] Dave Baiocchi and William Welser. **Confronting space debris: strategies and warnings from comparable examples including Deepwater Horizon**. en. RAND Corporation monograph series. Santa Monica: Rand Corporation, 2010. ISBN: 978-0-8330-5056-4.
- [20] Jeff Foust. “Space Force backs development of commercial orbital debris removal systems”. In: **SpaceNews** (Sept. 2021). URL: <https://spacenews.com/space-force-backs-development-of-commercial-orbital-debris-removal-systems/>.
- [21] P Anz-Meador et al. “The Space Debris Sensor Experiment”. en. In: **1st International Orbital Debris (IOC) Conference**. Sugarland, TX: NASA Technical Report JSC-E-DAA-TN74830, Dec. 2019, p. 10.
- [22] Waldemar Bauer et al. “In orbit testing of SOLID debris detector”. en. In: **Acta Astronautica** 197 (May 2022), pp. 235–245. ISSN: 00945765. DOI: 10.1016/j.actaastro.2022.05.024. URL: <https://linkinghub.elsevier.com/retrieve/pii/S009457652200220X> (visited on 06/15/2022).
- [23] Aaron Trott. **Flight Systems and Technologies for Impact Detection and Location**. In Space Inspection Workshop, 2014. URL: <https://www.nasa.gov/sites/default/files/>

- files/A_Trott-Flight_Systems_and_Technologies_for_Impact_Detection_and_Location.pdf.
- [24] Martin Schimmerohn, Max Gulde, and Alain Hilgers. “USING ANTENNAE FOR IN-SITU MEASUREMENTS OF MICROMETEOROID AND SPACE DEBRIS IMPACTS”. en. In: (), p. 6.
- [25] A.J. Tuzzolino et al. “Final results from the space dust (SPADUS) instrument flown aboard the earth-orbiting ARGOS spacecraft”. en. In: **Planetary and Space Science** 53.9 (Aug. 2005), pp. 903–923. ISSN: 00320633. DOI: 10.1016/j.pss.2005.03.008. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0032063305000772> (visited on 06/15/2022).
- [26] P. J. Kellogg, K. Goetz, and S. J. Monson. “Dust impact signals on the wind spacecraft”. en. In: **Journal of Geophysical Research: Space Physics** 121.2 (Feb. 2016), pp. 966–991. ISSN: 2169-9380, 2169-9402. DOI: 10.1002/2015JA021124. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/2015JA021124> (visited on 03/19/2020).
- [27] L. Andersson et al. “Dust observations at orbital altitudes surrounding Mars”. en. In: **Science** 350.6261 (Nov. 2015), aad0398–aad0398. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aad0398. URL: <https://www.sciencemag.org/lookup/doi/10.1126/science.aad0398> (visited on 05/07/2020).
- [28] M. Landgraf et al. “Aspects of the mass distribution of interstellar dust grains in the solar system from in situ measurements”. en. In: **Journal of Geophysical Research: Space Physics** 105.A5 (May 2000), pp. 10343–10352. ISSN: 01480227. DOI: 10.1029/1999JA900359. URL: <http://doi.wiley.com/10.1029/1999JA900359> (visited on 05/07/2020).
- [29] Waldemar Bauer. **SOLID-A solar panel based impact detector**. URL: <https://indico.esa.int/event/370/contributions/5922/attachments/4239/6338/SOLID-A%20solar%20panel%20based%20impact%20detector.pdf>.
- [30] J. L. Jorgensen et al. “Distribution of Interplanetary Dust Detected by the Juno Spacecraft and Its Contribution to the Zodiacal Light”. en. In: **Journal of Geophysical Research:**

- Planets** 126.3 (Mar. 2021). ISSN: 2169-9097, 2169-9100. DOI: 10.1029/2020JE006509. URL: <https://onlinelibrary.wiley.com/doi/10.1029/2020JE006509> (visited on 06/15/2022).
- [31] Trevor Williams, Joseph Sedlak, and Seth Shulman. “Magnetospheric Multiscale Mission Micrometeoroid/Orbital Debris Impacts”. en. In: **Spacecraft Anomalies and Failures Workshop 2017**. Chantilly, VA, Dec. 2017.
- [32] Jose Martinez. “DrMUST - a Data Mining Approach for Anomaly Investigation”. en. In: **SpaceOps 2012 Conference**. Stockholm, Sweden: American Institute of Aeronautics and Astronautics, June 2012. DOI: 10.2514/6.2012-1275109. URL: <http://arc.aiaa.org/doi/10.2514/6.2012-1275109> (visited on 03/19/2020).
- [33] Brien Flewelling et al. “Explained and Unexplained Momentum Impulse Transfer Events”. en. In: **7th European Conference on Space Debris**. Darmstadt, Germany: ESA Space Debris Office, Apr. 2017.
- [34] Robin Putzar et al. “How hypervelocity impacts can affect the LISA mission – The MIRAD study”. In: **International Astronautical Congress**. Washington, D. C., Oct. 2019.
- [35] J. I. Thorpe et al. “Micrometeoroid Events in *LISA Pathfinder*”. In: **The Astrophysical Journal** 883.1 (Sept. 2019), p. 53. ISSN: 1538-4357. DOI: 10.3847/1538-4357/ab3649. URL: <https://iopscience.iop.org/article/10.3847/1538-4357/ab3649> (visited on 01/17/2022).
- [36] Will Gater. **Galaxy mapper’s first discovery: surprise space debris**. July 2014. URL: <https://www.newscientist.com/article/dn25925-galaxy-mappers-first-discovery-surprise-space-debris/>.
- [37] **Image of the Week: GAIA Intersects a Perseid Meteor**. Jan. 2016. URL: https://www.cosmos.esa.int/web/gaia/iow_20160106.

- [38] GAIA Flight Control Team. **Meteoroid Impact Detections by the GAIA Spacecraft at L2**. URL: https://www.cosmos.esa.int/documents/653713/1049906/02_serpell_meteoroids_2016.pdf/aa361b27-367e-4b0b-9865-3c63e508a470.
- [39] Marie-Liis Aru. **Using a star-mapping telescope to measure tiny space rocks**. Nov. 2021. URL: <https://space-travel.blog/gaia-rocks-96ec6d70ae90>.
- [40] Joel Williamsen et al. “Characterizing the Orbital Debris Environment Using Satellite Perturbation Anomaly Data”. In: **First International Orbit Debris Conference**. Houston, TX, Dec. 2019.
- [41] Anne Aryadne Bennett, Hanspeter Schaub, and Russell Carpenter. “Assessing Debris Strikes in Spacecraft Telemetry: Development and Comparison of Various Techniques”. en. In: **70th International Astronautical Congress**. Washington, D. C., Oct. 2019.
- [42] Anne Aryadne Bennett and Hanspeter Schaub. “Identifying and Assessing Debris Strikes in NASA Spacecraft Telemetry”. en. In: **First International Orbital Debris Conference**. Sugarland, TX, Dec. 2019, p. 10.
- [43] Anne Aryadne Bennett and Hanspeter Schaub. “Effect of Spacecraft Parameters on Identification of Debris Strikes in GN&C Telemetry”. en. In: **Guidance Navigation and Control Conference**. Breckenridge, CO: American Astronautical Society, Feb. 2020, p. 12.
- [44] Anne Aryadne Bennett, Hanspeter Schaub, and Russell Carpenter. “Assessing debris strikes in spacecraft telemetry: Development and comparison of various techniques”. en. In: **Acta Astronautica** 181 (Apr. 2021), pp. 516–529. ISSN: 00945765. DOI: 10.1016/j.actaastro.2020.09.009. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0094576520305452> (visited on 06/15/2022).
- [45] William P. Schonberg. “Concise history of ballistic limit equations for multi-wall spacecraft shielding”. en. In: **REACH** 1 (Mar. 2016), pp. 46–54. ISSN: 23523093. DOI: 10.1016/j.reach.2016.06.001. URL: <https://linkinghub.elsevier.com/retrieve/pii/S2352309316300050> (visited on 10/06/2021).

- [46] A V Moorhead. “NASA Meteoroid Engineering Model (MEM) Version 3”. en. In: **Marshall Space Flight Center**. Huntsville, AL (Jan. 2020), p. 62.
- [47] M. D. Campbell-Brown and J. Jones. “Annual variation of sporadic radar meteor rates”. en. In: **Monthly Notices of the Royal Astronomical Society** 367.2 (Apr. 2006), pp. 709–716. ISSN: 0035-8711, 1365-2966. DOI: 10.1111/j.1365-2966.2005.09974.x. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2005.09974.x> (visited on 03/25/2020).
- [48] Althea V. Moorhead et al. “Meteor Shower Forecasting in Near-Earth Space”. en. In: **Journal of Spacecraft and Rockets** 56.5 (Sept. 2019), pp. 1531–1545. ISSN: 0022-4650, 1533-6794. DOI: 10.2514/1.A34416. URL: <https://arc.aiaa.org/doi/10.2514/1.A34416> (visited on 06/21/2022).
- [49] J A M McDonnell. “Momentum Enhancement in Hypervelocity Impacts: Parameters for Space Debris and Meteoroid Perturbations”. en. In: (), p. 10.
- [50] J. A. M. McDonnell et al. “Dust density and mass distribution near comet Halley from Giotto observations”. en. In: **Nature** 321.S6067 (May 1986), pp. 338–341. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/321338a0. URL: <http://www.nature.com/articles/321338a0> (visited on 06/26/2022).
- [51] J.A.M. McDonnell et al. “The impact of dust grains on fast fly-by spacecraft: Momentum multiplication, measurements and theory”. en. In: **Advances in Space Research** 4.9 (Jan. 1984), pp. 297–301. ISSN: 02731177. DOI: 10.1016/0273-1177(84)90046-2. URL: <https://linkinghub.elsevier.com/retrieve/pii/0273117784900462> (visited on 06/26/2022).
- [52] James D. Walker et al. “Scale Size Effect in Momentum Enhancement”. en. In: **Procedia Engineering** 58 (2013), pp. 240–250. ISSN: 18777058. DOI: 10.1016/j.proeng.2013.05.028. URL: <https://linkinghub.elsevier.com/retrieve/pii/S187770581300934X> (visited on 06/21/2022).

- [53] Chris Tschan. “What Do We Need for Satellite Event Attribution?” en. In: **Spacecraft Anomalies and Failures Workshop**. Chantilly, VA, July 2014.
- [54] Hanspeter Schaub and John L. Junkins. **Analytical mechanics of space systems**. Third edition. AIAA education series. Reston, Virginia: American Institute of Aeronautics and Astronautics, Inc, 2014. ISBN: 978-1-62410-240-0.
- [55] Colin H Smith et al. “Performance of Multi-Layer Insulation for Spacecraft Instruments at Cryogenic Temperatures”. en. In: **46th International Conference on Environmental Systems**. Vienna, Austria, July 2016, p. 20.
- [56] Frank Schäfer et al. “Hypervelocity Impact Testing of CFRP/AL Honeycomb Satellite Structures”. en. In: **Fourth European Conference on Space Debris**. Darmstadt, Germany, Apr. 2005, p. 6.
- [57] Joel Williamsen and Steven Evans. “Orbital debris momentum transfer in satellite shields following hypervelocity impact, and its application to environment validation”. en. In: **Procedia Engineering** 204 (2017), pp. 500–507. ISSN: 18777058. DOI: 10.1016/j.proeng.2017.09.747. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1877705817343011> (visited on 09/11/2019).
- [58] Pierre Moulin and Venugopal V. Veeravalli. **Statistical inference for engineers and data scientists**. Cambridge, UK ; New York, NY: Cambridge University Press, 2019. ISBN: 978-1-107-18592-0.
- [59] Bernard C. Levy. **Principles of signal detection and parameter estimation**. eng. OCLC: 254592856. New York, NY: Springer, 2008. ISBN: 978-0-387-76542-6 978-0-387-76544-0.
- [60] Durga P Malladi and Jason L Speyer. “A Generalized Shirayev Sequential Probability Ratio Test for Change Detection and Isolation”. en. In: **IEEE TRANSACTIONS ON AUTOMATIC CONTROL** 44.8 (1999), p. 13.

- [61] Alexander L. Bogorad et al. “Electrostatic Discharge Induced Momentum Impulse From Charged Spacecraft Surfaces”. en. In: **IEEE Transactions on Nuclear Science** 53.6 (Dec. 2006), pp. 3607–3609. ISSN: 0018-9499. DOI: 10.1109/TNS.2006.885108. URL: <http://ieeexplore.ieee.org/document/4033649/> (visited on 10/03/2019).
- [62] Xu Kunbo et al. “Investigation on solar array damage characteristic under millimetre size orbital debris hypervelocity impact”. en. In: (), p. 8.
- [63] Andrew Harris and Hanspeter Schaub. “Deep On-Board Scheduling for Autonomous Attitude Guidance Operations”. en. In: **AAS Guidance, Navigation and Control Conference**. Breckenridge, CO, Feb. 2020, p. 19.
- [64] Darren S. McKnight and Frank R. Di Pentino. “New insights on the orbital debris collision hazard at GEO”. en. In: **Acta Astronautica** 85 (Apr. 2013), pp. 73–82. ISSN: 00945765. DOI: 10.1016/j.actaastro.2012.12.006. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0094576512004869> (visited on 09/10/2019).
- [65] Er kai Watson and Martin Steinhauser. “Discrete Particle Method for Simulating Hypervelocity Impact Phenomena”. en. In: **Materials** 10.4 (Apr. 2017), p. 379. ISSN: 1996-1944. DOI: 10.3390/ma10040379. URL: <http://www.mdpi.com/1996-1944/10/4/379> (visited on 01/06/2022).
- [66] James Wright. **Orbit Determination Toolkit Theory & Algorithms**. Tech. rep. Analytical Graphics Inc, Sept. 2020.
- [67] Hyun Chul Ko and Daniel J. Scheeres. “Maneuver Detection with Event Representation Using Thrust Fourier Coefficients”. en. In: **Journal of Guidance, Control, and Dynamics** 39.5 (May 2016), pp. 1080–1091. ISSN: 0731-5090, 1533-3884. DOI: 10.2514/1.G001463. URL: <https://arc.aiaa.org/doi/10.2514/1.G001463> (visited on 01/18/2022).
- [68] Marcus J. Holzinger, Daniel J. Scheeres, and Kyle T. Alfriend. “Object Correlation, Maneuver Detection, and Characterization Using Control Distance Metrics”. en. In: **Journal of Guidance, Control, and Dynamics** 35.4 (July 2012), pp. 1312–1325. ISSN: 0731-5090,

- 1533-3884. DOI: 10.2514/1.53245. URL: <https://arc.aiaa.org/doi/10.2514/1.53245> (visited on 01/18/2022).
- [69] Tom Kelecy and Moriba Jah. “Detection and orbit determination of a satellite executing low thrust maneuvers”. en. In: **Acta Astronautica** 66.5-6 (Mar. 2010), pp. 798–809. ISSN: 00945765. DOI: 10.1016/j.actaastro.2009.08.029. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0094576509004287> (visited on 01/18/2022).
- [70] Yuzi Jiang et al. “Extended Kalman Filter with Input Detection and Estimation for Tracking Manoeuvring Satellites”. en. In: **Journal of Navigation** 72.3 (May 2019), pp. 628–648. ISSN: 0373-4633, 1469-7785. DOI: 10.1017/S037346331800098X. URL: https://www.cambridge.org/core/product/identifier/S037346331800098X/type/journal_article (visited on 01/18/2022).
- [71] D. Fraser and J. Potter. “The optimum linear smoother as a combination of two optimum linear filters”. en. In: **IEEE Transactions on Automatic Control** 14.4 (Aug. 1969), pp. 387–390. ISSN: 0018-9286. DOI: 10.1109/TAC.1969.1099196. URL: <http://ieeexplore.ieee.org/document/1099196/> (visited on 05/24/2022).
- [72] Joseph E. Wall, Alan S. Willsky, and Nils R. Sandell. “On the fixed-interval smoothing problem”. en. In: **Stochastics** 5.1-2 (June 1981), pp. 1–41. ISSN: 0090-9491. DOI: 10.1080/17442508108833172. URL: <https://www.tandfonline.com/%20doi/full/10.1080/17442508108833172> (visited on 06/25/2022).
- [73] David A. Vallado and Wayne D. McClain. **Fundamentals of astrodynamics and applications**. eng. 4. ed. Space technology library 21. Hawthorne, Calif: Microcosm Press, 2013. ISBN: 978-1-881883-18-0.
- [74] Byron D. Tapley, Bob E. Schutz, and George H. Born. **Statistical orbit determination**. Amsterdam ; Boston: Elsevier Academic Press, 2004. ISBN: 978-0-12-683630-1.
- [75] Russell Carpenter and Christopher D’Souza. **Navigation Filter Best Practices**. Tech. rep. NASA/TP-2018-219822. NASA, Apr. 2018.

- [76] William Lear. **Kalman Filtering Techniques**. Tech. rep. JSC-20688. Houston, TX: NASA Johnson Space Center, 1985.
- [77] Donald Fraser. “A New Technique for the Optimal Smoothing of Data”. PhD thesis. Cambridge, Massachusetts: Massachusetts Institute of Technology, Jan. 1967.
- [78] D. Fraser and J. Potter. “The optimum linear smoother as a combination of two optimum linear filters”. en. In: **IEEE Transactions on Automatic Control** 14.4 (Aug. 1969), pp. 387–390. ISSN: 0018-9286. DOI: 10.1109/TAC.1969.1099196. URL: <http://ieeexplore.ieee.org/document/1099196/> (visited on 05/24/2022).
- [79] Robert Grover Brown and Patrick Y. C. Hwang. **Introduction to random signals and applied Kalman filtering: with MATLAB exercises**. 4th ed. Hoboken, NJ: John Wiley, 2012. ISBN: 978-0-470-60969-9.
- [80] **NASA Orbital Debris Engineering Model ORDEM 3.1 - Software User Guide**. Dec. 2019.
- [81] Cody Allard, Hanspeter Schaub, and Scott Piggott. “General Hinged Rigid-Body Dynamics Approximating First-Order Spacecraft Solar Panel Flexing”. en. In: **Journal of Spacecraft and Rockets** 55.5 (Sept. 2018), pp. 1291–1299. ISSN: 0022-4650, 1533-6794. DOI: 10.2514/1.A34125. URL: <https://arc.aiaa.org/doi/10.2514/1.A34125> (visited on 07/22/2022).
- [82] Abdellatif Bellar and Mohammed Arezki Si Mohammed. “Satellite Inertia Parameters Estimation Based on Extended Kalman Filter”. en. In: **Journal of Aerospace Technology and Management** (2019). ISSN: 2175-9146. DOI: 10.5028/jatm.v11.1016. URL: <http://ref.scielo.org/8s336s> (visited on 07/22/2022).
- [83] Mark L. Psiaki. “Estimation of a Spacecraft’s Attitude Dynamics Parameters by Using Flight Data”. en. In: **Journal of Guidance, Control, and Dynamics** 28.4 (July 2005), pp. 594–603. ISSN: 0731-5090, 1533-3884. DOI: 10.2514/1.7362. URL: <https://arc.aiaa.org/doi/10.2514/1.7362> (visited on 07/20/2022).

- [84] Shannon Ryan, Stephen Thaler, and Sevvandi Kandanaarachchi. “Machine learning methods for predicting the outcome of hypervelocity impact events”. en. In: **Expert Systems with Applications** 45 (Mar. 2016), pp. 23–39. ISSN: 09574174. DOI: 10.1016/j.eswa.2015.09.038. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0957417415006673> (visited on 07/29/2022).
- [85] P.H. Krisko et al. “ORDEM 3.0 and MASTER-2009 modeled debris population comparison”. en. In: **Acta Astronautica** 113 (Aug. 2015), pp. 204–211. ISSN: 00945765. DOI: 10.1016/j.actaastro.2015.03.024. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0094576515001241> (visited on 07/29/2022).