

A Root Cause Analysis of REXIS Detection Efficiency Loss During Phase E Operations

by

Madeline Lambert

B.S. Physics, University of Wisconsin - La Crosse (2016)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2020

© Massachusetts Institute of Technology 2020. All rights reserved.

Author
Department of Aeronautics and Astronautics
May 19, 2020

Certified by.....
Rebecca Masterson
Principle Research Scientist
Thesis Supervisor

Accepted by
Sertac Karaman
Associate Professor of Aeronautics and Astronautics, Chair, Graduate
Program Committee

A Root Cause Analysis of REXIS Detection Efficiency Loss During Phase E Operations

by

Madeline Lambert

Submitted to the Department of Aeronautics and Astronautics
on May 19, 2020, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

The Regolith X-ray Imaging Spectrometer (REXIS) is a student-built instrument flown on NASA's Origins, Spectral Interpretation, Resource Identification, Safety, Regolith Explorer (OSIRIS-REx) mission. During main science operations, the instrument experienced detector efficiency loss in the form of loss of iron calibration source counts, which greatly affected the science output. In this thesis, a root cause investigation is performed on the loss of iron counts, and an optical light leak onto the edge of the instrument's detectors is identified as the most likely cause. A CAST analysis is then performed to identify possible organizational and cultural causes of the design that allowed for an optical light leak, and recommendations for future similar instruments (low-cost, high-risk) are made.

Thesis Supervisor: Rebecca Masterson

Title: Principle Research Scientist

Acknowledgments

I want to thank literally everyone who's ever helped me with this project, but I fear my acknowledgment section will soon become thesis-sized in length so I'll try to keep it brief. First of all, thank you to the whole REXIS team. Mark and Carolyn, thank you for making MIT such a fun and welcoming place. And Andrew, thanks for being such a great addition to the team (even if you violated the strict 5' 9" height limit for REXIS students). Becky, thank you for your constant advice and guidance. David and Dan, thank you for all the help with coding and the assistance with the many, many steps it took to approve each REXIS operation. Rick, thank you for your constant support and enthusiasm in your leadership. Thank you to the whole Harvard science team: Branden, Jae Sub, and Josh, you made operating this complicated little instrument a fun adventure. I would also like to thank everyone who assisted me in my CAST analysis, especially Kevin Ryu, whose knowledge proved invaluable. Also, thank you to my wonderful office mates, who made coming in to work (when I could) a wonderful experience. I'll always remember our competitive office soccer games, and the official Friendship Museum.

As always, thanks to my family for believing in me so much that I started to believe I could do anything (like finish a Master's degree at MIT). You're the most loving, supportive family anyone could ever wish for.

Contents

1	Introduction	23
1.1	OSIRIS-REx Mission Timeline Overview	24
1.2	REXIS Background	27
1.3	Motivation for Thesis	30
1.4	Thesis Roadmap	31
2	Background	33
2.1	REXIS Detectors	33
2.1.1	REXIS Image Processing	36
2.1.2	Optical-Blocking Filter	41
2.2	Root Cause Analysis Methods	41
2.2.1	5 Whys Method	42
2.2.2	Fault Tree Analysis	42
2.2.3	Ishikawa Diagrams	46
2.2.4	Failure Modes and Effects Analysis	46
2.3	CAST Analysis	48
2.3.1	STAMP Theory	48
2.3.2	Accident Analysis using CAST	48
3	REXIS Operations	55
3.1	Early Checkouts and Calibrations	58
3.1.1	Launch + 14 Days Payload Checkout	58
3.1.2	L+6 Checkout and Calibration	58

3.1.3	L+10 Checkout and SXM Calibration	59
3.2	Internal Calibrations	61
3.2.1	L+18 Checkout and Calibration	61
3.2.2	REXIS Internal Calibration	62
3.2.3	L+22 Checkout and Calibration	63
3.3	Cover Opening	64
3.4	External Calibrations	67
3.4.1	CXB Calibration	67
3.4.2	Crab Calibration: Part 1	70
3.4.3	L+30 Calibration	73
3.4.4	Crab Calibration: Part 2	77
3.4.5	OBF Verification	81
3.4.6	Mask Calibration	83
3.5	Asteroid Science	86
3.5.1	Orbital B	86
3.5.2	Orbital R	98
4	Root Cause Analysis	103
4.1	Iron Count Rates versus Total Event Rates	108
4.2	Potential Causes	109
4.2.1	Physical Causes	110
4.2.2	Software Causes	112
4.3	Less Likely Root Causes	112
4.3.1	Drift in CCD Voltages	113
4.3.2	Rise in CCD or MEB Temperatures	115
4.4	Most Likely Root Cause	118
4.4.1	Comparison of Optical Light to Iron Counts	121
4.4.2	Comparison of Stray Light Simulation to Flight Data	123
4.4.3	Artificial Grade Increases	127
4.4.4	Comparison of Orbital B to Orbital R	137

5	CAST Analysis	141
5.1	Part 1: Assemble Basic Information	142
5.2	Part 2: Model the Safety Control Structure	145
5.3	Part 3: Analyze the Loss at the Component Level	147
5.3.1	REXIS Detectors and DAM and FSW	148
5.3.2	Student Detector Team	149
5.3.3	Student Systems Engineering and I&T Team	151
5.3.4	Science/Instrument Architecture Team	154
5.3.5	Lincoln Laboratory	155
5.3.6	MIT Kavli Institute for Astrophysics and Space Research . . .	157
5.3.7	Project Management	159
5.4	Part 4: Analyze the Loss at the Control System Level	161
5.4.1	Summary of questions and answers	163
5.5	Part 5: General Recommendations	164
6	Recommendations and Conclusions	167
6.1	Summary of Root Cause Analysis	167
6.1.1	Limitations	168
6.1.2	Future Work	168
6.2	Summary of CAST Recommendations	169
6.2.1	Limitations	170
6.2.2	Future Work	170
A	REXIS Image Processing Information	173
A.1	Code used for grading simulation	177
A.2	Event grading code	182
B	CAST Materials	189
B.1	REXIS Level 4 Requirements pertaining to optical light	189
B.2	Email chain discussing decision to not paint the DAM	190
B.3	Plans to paint the DAM	194

B.3.1 Email chains discussing painting of the DAM 198

List of Figures

1-1	Graphic of OSIRIS-REx payload deck. Image from OSIRIS-REx official mission website.	24
1-2	OSIRIS-REx mission timeline. Image from OSIRIS-REx official mission website.	25
1-3	Model of the REXIS instrument [4]	28
1-4	Photograph of assembled REXIS instrument [2]	29
1-5	Timeline of students who worked on REXIS. Figure created by Professor Rick Binzel.	29
2-1	Model of DAM structure [4]	35
2-2	Physical layout of the REXIS CCDs in a drawing (left) and in flight data (right). In the drawing, the wider lines indicate the bounds of the full CCD, and the thinner lines mark the boundaries of the nodes within each CCD.	35
2-3	Diagram showing CCD wiring, with boxes around the lines cut during the accident. [7]	36
2-4	Image processing block diagram. Image from the Master's thesis of REXIS student Pronoy Biswas [12]	38
2-5	Grading scheme used for REXIS, adopted from Chandra's ACIS instrument [4]	40
2-6	Example of a 5-whys analysis, performed on a transistor chip with degraded memory [19]	43
2-7	Logic gate symbols used for FTA [22]	44

2-8	Event symbols used for FTA [22]	45
2-9	Example Fault Tree from official NASA FTA handbook [24]	45
2-10	Example Fishbone diagram used in root cause analysis of an optical throughput decrease for a NASA instrument [24]	47
2-11	Overview of CAST method [30]	50
2-12	Example control structure from a CAST analysis of a U.S. Coast Guard Aviation mishap [34]	52
3-1	REXIS operational timeline	56
3-2	Planned REXIS Operations throughout flight	57
3-3	Number of nodes used for each operation, in chronological order	57
3-4	Figure showing total downlinked events per second on August 15, 2017. The red line indicates the binary state of the filter, with the higher value indicating filtering is occurring, and lower value indicating no filtering is occurring.	61
3-5	Figure showing total downlinked events per second on October 10, 2017. The red line indicates the binary state of the filter, with the higher value indicating filtering is occurring, and lower value indicating no filtering is occurring. The limit set by the event grade filter should be 200 cps, but the black line (downlinked events) reaches far above that limit.	62
3-6	Effects of hot pixel mask application on flight data	64
3-7	Data products to determine the radiation cover opened	65
3-8	Plot of events over detector surface, and predictions of event gradients if cover only opened partially	66
3-9	(Top) plot showing event rates throughout cover opening activity. (Bot- tom) Plot showing comparison between CCD event rates and the zenith angle with relation to the sun.	68
3-10	Plot of energy (in ADU) versus time for CCD02 during all of flight. CXB Calibration data are outlined by the dotted line.	70

3-11	Comparison of energy histogram for CCD02 from before and during CXB calibration.	70
3-12	Plot of event rates from the science window on November 24, 2018, part of the Crab calibration. The filter is activated for the entire period, and portions of the events exceed the software limit of 200 counts per second, indicating a filter breach similar to L+10.	71
3-13	Plot of total downlinked and housekeeping event rates from September 2016 through Crab calibration in November 2018	72
3-14	Evidence of triangular cross-talk pattern on the CCDs from both the CXB and Crab calibrations	73
3-15	CCD physical layout. The nodes with the best resolution that were used for L+30 operations are colored in green.	74
3-16	Total event rates (downlinked and housekeeping) from internal calibration through L+30	77
3-17	Events for CCD 0 2, with grade differentiation. The dotted line indicates grades 6 and 7 events. No grades 6 and 7 events are recorded after EGF is changed.	77
3-18	L+30 housekeeping event rates per CCD compared to the predicted light curve (gray line near bottom of plot). Plot created by HCO researcher David Guevel.	78
3-19	Recorded spectrum from Crab Cal 2 (green line) compared to the absolute response function from before and after Crab Cal 2. The box indicates the excess low energy, and the brackets display the discrepancy in the energy cutoff around 75 ADU. Plot created by HCO researcher Daniel Hoak.	80
3-20	Bias maps downlinked from OBF verification. The second map has vertical and horizontal artifacts, and does not contain any hot pixels.	82
3-21	Comparison of bias map with subsequent event list data. The dotted outline shows an area of vertical distortion, and the arrows point out some of the horizontal artifacts.	83

3-22	Mask image created from one pointing target observation of Sco X-1 .	84
3-23	Event rates broken down by grade for CCD02 and CCD22. Arrows indicate when ST was doubled. Grade 0 events increase every time ST is doubled.	85
3-24	Ratio of CCD22 detector efficiency with doubled ST to detector efficiency with nominal ST during the mask calibration operation. The ratio of the two efficiencies is plotted from 0 to 6 keV. Plot created by HCO researcher David Guevel.	85
3-25	REXIS operations timeline, with notable anomalies marked by a star.	87
3-26	Block diagram of Orbital B block	90
3-27	Predicted iron count rates from calibration sources for CCD20 throughout flight. Actual measured values from each operation are represented by the dots. Plot created by HCO scientist Daniel Hoak.	93
3-28	Predicted iron count rates from calibration sources for CCD22 throughout flight. Actual measured values from each operation are represented by the dots. Plot created by HCO scientist Daniel Hoak.	93
3-29	Detected energy versus time for CCD22 from August 5, 2019	94
3-30	Plots showing change in total event rates and resulting spectra from before and after presumed SEU. The arrow on the top plot indicates where noise began. The red line on the bottom plot is the energy spectrum from before the SEU, and the black line is the spectrum from after the SEU.	95
3-31	Orbital B events divided by grades and time. The first red section indicates nominal performance between July 1 and July 5. The second green portion indicates the data affected by the SEU. The third purple portion is still during the SEU-affected data period, with the addition of varying noise. The final blue portion is after the first power cycle, when the flat noise from the SEU disappeared.	96
3-32	Top plot shows event rates throughout Orbital B. Bottom plot shows spectra collected at different times throughout Orbital B	97

3-33	Event rates and spectra for Orbital R. The spectrum from the day with higher event rates (November 14) is shown in black, and the two spectra of the periods preceding and following November 14 are shown in red and blue.	99
3-34	Predicted iron count rates and rates from each observation. Plot created by HCO scientist Jae Sub Hong.	100
3-35	Grades 0 through 5 iron counts on CCD22 from launch through Orbital R. Plot created by HCO scientist Jae Sub Hong.	101
3-36	Detected energy versus time for CCD22 from November 23, 2019. (The large gap in the data between 16,000 and 29,000 seconds is removed due to a bad bias map)	101
4-1	REXIS flight operations over time. Red stars indicate an anomalous behavior was observed, and black dotted line indicates high event rates.	104
4-2	REXIS housekeeping event rates throughout all of flight. Downlinked events are shown in red, and housekeeping events are shown in blue. Figure ends with Orbital B. X axis is elapsed time in seconds, and y axis is number of counts.	105
4-3	Continuation of REXIS housekeeping event rates plot. Orbital R data are shown, with downlinked events shown in red and housekeeping events shown in blue. X axis is elapsed time in seconds, and y axis is number of counts.	106
4-4	Event rates for one day of Orbital B. Spikes in event rates are indicated by arrows.	107
4-5	Plot of event energy vs time	107
4-6	Plots of iron event rates versus total event rates for all of Orbital B in log scale	109
4-7	Fishbone diagram for potential causes of CCD loss of iron counts . . .	110

4-8	CCD2 voltages for Orbital B and Orbital R. The yellow (and red, where applicable) high and low voltage alarm levels are indicated by the horizontal lines.	114
4-9	Voltages versus event rates for one day of Orbital B (July 29)	115
4-10	Voltages versus event rates for one day of Orbital B (July 31)	115
4-11	MEB (blue) and detector electronics temperatures throughout all of flight. The bottom plot contains all temperatures for all of Orbital B. Plot created by HCO scientist Jae Sub Hong.	117
4-12	MEB (blue) other detector electronics throughout all of flight, continued, depicting temperatures for Orbital R. Plot created by HCO scientist Jae Sub Hong.	118
4-13	Temperatures versus event rates for one day of Orbital B	118
4-14	Diagram showing possible pathways of light through DAM to readout edges of the nodes. Diagram created by HCO scientist Jae Sub Hong.	119
4-15	Picture of the bottom of the DAM, with arrows indicating where optical light could enter.	120
4-16	CCD2 event rates versus calculated optical brightness (calculated using relative albedo) from Bennu. Plot created by HCO researcher David Guevel.	122
4-17	Plots of iron event rates versus total event rates for all of Orbital B, with optical brightness ranges indicated by different marker shapes .	122
4-18	Comparison of just grade 0 iron counts and grades 0-5 iron counts to CXB exposure on nodes. The x-axes are CXB exposure as a percentage of node area, and the y-axes are iron counts in counts per kilosecond. Plots created by HCO scientist Jae Sub Hong	123
4-19	Example of event gradient across CCD22 seen throughout Orbital B .	124
4-20	Spatial OD maps of CCD absorption on one CCD with varying levels of light leak protection	125
4-21	LL stray light simulation spatial distribution compared to spatial distribution of events from Orbital B with loss of iron counts	126

4-22	Spatial distribution of events for Orbital B data with no loss of iron counts	127
4-23	Diagram showing steps involved in grading simulation.	128
4-24	Raw frame energy histograms for all nodes from L+30, shown in event list ADU. Each color depicts the energy distribution for the identified node.	129
4-25	Raw frame energy histogram for CCD22 from L+30, shown in event list ADU.	130
4-26	Energy histograms of event list data from August 02 (light leak) and July 26 (no light leak)	130
4-27	Scaled raw frame energy histogram for comparison to event list data fit	131
4-28	Three different bias maps for CCD22 recorded during Orbital B. The three energy histograms are shown on the left, and the images are shown on the right. X-axis is in units of raw frame ADU.	132
4-29	Bias map energy histogram for CCD22 from Mask Calibration. The energy histogram is on the left, and the image of the map is on the right. X-axis is in units of raw frame ADU.	133
4-30	Fitting process for 8-02 event list data to L+30 raw frame data energy distributions.	134
4-31	Visualization of event grid produced from grading sim	134
4-32	Result of grading simulation for approximate distribution with no added light	136
4-33	Result of grading simulation for approximate distribution from July 26 data	136
4-34	Result of August 2 approximate distribution with nominal and doubled ST.	137

5-1	Three identified paths for light to reach detector surface. (1) Light leak through the detector surface via pinholes created by particles. (2) Light leak through the edge of the detector. (3) Light leak up through the support wafer. [13]	145
5-2	Higher-level REXIS safety control structure. The dotted line indicates the boundary of the system under study in this analysis.	146
5-3	Zoomed-in REXIS safety control structure	147
B-1	Emails concurring on decision to not paint the DAM	193
B-2	Email chain indicating plans to paint the DAM, part 1 of 3	199
B-3	Email chain indicating plans to paint the DAM, part 2 of 3	200
B-4	Email chain indicating plans to paint the DAM, part 3 of 3	201

List of Tables

3.1	The six node enable maps cycled through during the L+30 operation	76
3.2	In-flight anomalies observed between Launch and Mask Calibration	86
3.3	Summary of anomalous behavior observed during Orbital B.	91
4.1	CCD video output and output gate voltage average comparisons for the six nodes active for Orbital B.	114
4.2	Summary of possible root causes for loss of iron calibration source counts	138
A.1	Energy ranges for ET, ST, and ULD	173
A.2	Different ADU values for a range of input energies. Actual values for different nodes were discovered to vary, but this provides a baseline.	173
A.3	CCD ET parameters used throughout flight, with changes in bold text, and explanation provided for each change. The values for the 10 nominal nodes are included, even though from L+30 onward, only five or six nodes were used	174
A.4	CCD ST parameters used throughout flight, with changes in bold text and explanation provided for each change. The values for the 10 nominal nodes are included, even though from L+30 onward, only five or six nodes were used	175
A.5	CCD ULD parameters used throughout flight, with changes in bold text and explanation provided for each change. The values for the 10 nominal nodes are included, even though from L+30 onward, only five or six nodes were used	176

List of Acronyms

CAST	Causal Analysis based on System Theory
CCD	Charge-coupled device
CfA	Center for Astrophysics
CXEL	candidate X-ray event list
DAM	Detector Assembly Mount
DD OBF	Directly deposited optical-blocking filter
DSN	Deep Space Network
EAPS	Earth, Atmospheric and Planetary Sciences
EGF	Event grade filter
ET	event threshold
FMEA	Failure modes and effects analysis
FPGA	Field-programmable gate array
FSW	Flight software
HCO	Harvard College Observatory
I&T	Integration and Testing
LL	Lincoln Laboratory
MKI	MIT Kavli Institute for Astrophysics and Space Research
NEM	Node enable map
OBF	Optical-blocking filter

OSIRIS-REx	Origins, Spectral Interpretation, Resource Identification, Security Regolith Explorer
OCAMS	OSIRIS-REx Camera Suite
OLA	OSIRIS-REx Laser Altimeter
OTES	OSIRIS-REx Thermal Emission Spectrometer
OVIRS	OSIRIS-REx Visible and Infrared Spectrometer
PI	Principle Investigator
PM	Project Manager
REXIS	Regolith X-ray Imaging Spectrometer
SSL	Space Systems Laboratory
ST	Split threshold
STAMP	System-Theoretic Accident Model and Process
ULD	Upper limit discriminator

Chapter 1

Introduction

The Regolith X-ray Imaging Spectrometer (REXIS) is a student-built instrument on board the Origins, Spectral Interpretation, Resource Identification, Security, and Regolith Explorer (OSIRIS-REx) mission, which launched in September of 2016. OSIRIS-REx is currently orbiting the asteroid 101955 Bennu (also known as 1999 RQ₃₆, but renamed Bennu for the OSIRIS-REx mission), and will return a sample of asteroid regolith to Earth in 2023 for further study. The five main science objectives of the mission are as follows [1]:

1. Return and analyze a sample of pristine carbonaceous asteroid regolith in an amount sufficient to study the nature, history, and distribution of its constituent minerals and organic material.
2. Map the global properties, chemistry, and mineralogy of a primitive carbonaceous asteroid to characterize its geologic and dynamic history and provide context for the returned samples.
3. Provide sample context by documenting the regolith at the sampling site in situ at scales down to the sub-centimeter.
4. Understand the interaction between asteroid thermal properties and orbital dynamics by measuring the Yarkovsky effect on a potentially hazardous asteroid and constraining the asteroid properties that contribute to this effect.
5. Improve asteroid astronomy by characterizing the astronomical properties of a primitive carbonaceous asteroid to allow for direct comparison with ground-




Figure 1-1: Graphic of OSIRIS-REx payload deck. Image from OSIRIS-REx official mission website.

based telescopic data of the entire asteroid population.

REXIS is one of five instruments on board OSIRIS-REx. The OSIRIS-REx Camera Suite (OCAMS) includes three types of cameras with differing resolutions for the various parts of the mission. The OSIRIS-REx Laser Altimeter (OLA) is a scanning light detection and ranging (LIDAR) to measure the distance from the spacecraft to the asteroid, and provide a high resolution map of the surface of Bennu. The OSIRIS-REx Thermal Emission Spectrometer (OTES) collects infrared spectral information from the surface of Bennu to assist in the elemental analysis of the regolith. The OSIRIS-REx Visible and Infrared Spectrometer (OVIRS) identifies elements on the surface using light in the visible and infrared range [2]. A graphic of the spacecraft deck with all the payloads is shown in Figure 1-1.

1.1 OSIRIS-REx Mission Timeline Overview

The OSIRIS-REx asteroid operations timeline is broken down into ten parts, as shown in Figure 1-2. The spacecraft launched in September of 2016, and was followed by a two year Outbound Cruise phase (not shown in the figure), which included regular

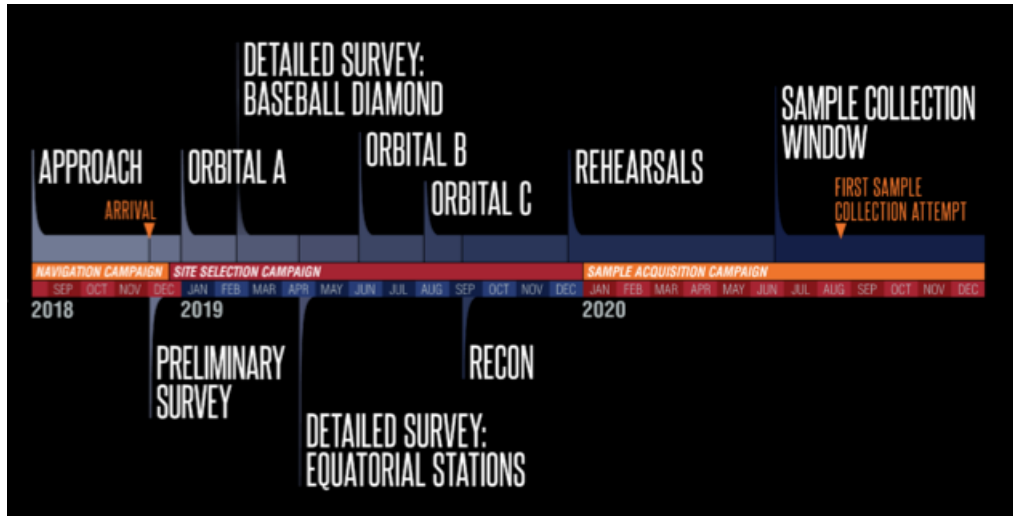


Figure 1-2: OSIRIS-REx mission timeline. Image from OSIRIS-REx official mission website.

instrument checkouts to monitor the health of all science payloads. In September of 2017, the spacecraft performed an Earth Gravity Assist to increase the spacecraft’s velocity on its trajectory to the asteroid. OSIRIS-REx first entered the Approach Phase in August 2018, once the spacecraft was about 2 million kilometers from Bennu. The goals of the Approach Phase were to locate the asteroid using the camera suite, and produce a preliminary shape model for the asteroid. Official asteroid arrival occurred on December 3, 2018. The asteroid arrival marked the beginning of the Preliminary Survey. Preliminary Survey involved hyperbolic orbits around the north and south poles of Bennu, at an altitude of 7 kilometers [2]. The goals of the Preliminary Survey were to provide data for a shape model of the asteroid with better resolution (75 centimeter), gather altimetry data, and determine the mass of the asteroid [2]. With this data, the mission team made a global coordinate system for use throughout the rest of the mission to identify features of interest and possible sample sites.

Next, the Orbital A phase marked the first time the spacecraft entered into an orbit around the asteroid. The goal of Orbital A was to practice the maneuvers necessary to enter a gravitationally-bound orbit, as needed for the later Orbital B global mapping phase. In Orbital A, OSIRIS-REx entered an orbit with a radius of 1.5 kilometers [2]. This phase also marked the transition from star-based to landmark-

based navigation [2]. Minimal science operations occurred during this period, with the emphasis placed more heavily on navigation and spacecraft dynamics in orbit around Bennu.

Orbital A was followed by the Detailed Survey Phase, which was further broken down into two phases: Baseball Diamond and Equatorial Stations. The Baseball Diamond Phase involved four passes over four areas of the surface of Bennu, the relative position of which resembled a baseball diamond (thus the name). The positions of the four areas also encompassed a wide range of viewing angles to capture the global properties of the asteroid. Each observation occurred at a range of 3.5 kilometers [2]. The goals of the Baseball Diamond Phase were to improve the resolution of the shape model to 35 centimeters, create digital terrain maps, and produce a global image mosaic of the surface of Bennu by combining images taken throughout the operation [2]. Equatorial Stations involved five kilometer passes over stations above the asteroid's equator [1]. This phase was designed to cover a wide range of illumination angles of the asteroid to image the asteroid in various lighting conditions [2].

Next, in the Orbital B phase, which occurred in July 2019, the spacecraft entered an even closer gravitationally-bound orbit around the asteroid, with an orbital radius of one kilometer [1]. Orbital B was the primary global mapping phase, with all instrument payloads operating and collecting data for a full month. The goal of this phase was to choose a primary and backup site from which to collect the regolith sample. Each potential sample site was ranked using Orbital B data to determine the safety of sampling from the site, the sampleability of the area, and the scientific value of the site [2]. Orbital B was followed by Orbital C, in which no science operations occurred. Instead, the camera suite took mosaic images of the sky surrounding the asteroid, to image any potential particles or natural satellites around Bennu.

Orbital C was then followed by the Reconnaissance phase, which was broken down into four parts (Recon A, Recon B, Orbital R, and Recon C). In Recon A, four initial candidates for sample site collection were flown over at an altitude of roughly 1200 kilometers for further assessment as to each site's viability as the primary sample site. Next, Orbital R contained operations that consisted of an extension of the

global mapping activities from Orbital B. Orbital R occurred at a higher altitude than Orbital B, with an orbital radius of about 1.2 kilometers compared to Orbital B's one kilometer. In Recon B, two additional flyovers were performed at a lower altitude than in Recon A (at roughly 525 meters), for further study of the two finalists for primary sample site [2]. Finally, Recon C consists of two flyovers at the lowest altitude so far (250 meters) for the primary site and backup site [2]. The first flyover for Recon C was performed in March 2020, and the second will occur in May 2020.

The final two phases of OSIRIS-REx asteroid operations will occur in 2020. First, the sample acquisition maneuver will be practiced in the Rehearsal phase. Finally, the Touch-And-Go (TAG) maneuver will be attempted in August 2020, with two backup attempts scheduled. The spacecraft will then begin the return journey to Earth beginning in 2021, and will drop off the sample in 2023 [1].

1.2 REXIS Background

REXIS is an X-ray fluorescence spectrometer, and is designed to detect X-rays incident from the surface of Bennu. X-ray fluorescence involves incoming high-energy photons exciting an atom, which then releases an X-ray photon. The energy of the released X-ray photon is characteristic of the element it originated from, so by measuring the energy of the released photon, the element of origin can be determined. By recording the energies of the fluoresced photons from a surface, the elemental ratios present in the surface can be determined. REXIS records X-ray fluorescence from the surface of Bennu in the soft X-ray band (0.5 to 7 keV) [3]. The REXIS instrument's main science objectives are to characterize Bennu among the known chondritic meteorite types, as well as create a global elemental abundance ratio map of the asteroid [4]. REXIS's main science objectives support goals two, three, and five of OSIRIS-REx (listed in the introductory paragraph).

The REXIS instrument consists of two main parts: the main spectrometer assembly (which uses charge-coupled devices (CCDs) for X-ray detection), and the solar X-ray monitor (which uses a Silicon Drift Diode (SDD) as an X-ray detector) [4]. The

final CAD model of the spectrometer assembly is shown in Figure 1-3 [4], and the fully assembled instrument is shown in Figure 1-4[2]. The CCDs are housed within the aluminum tower, and protected from radiation via an aluminum radiation cover. The radiation cover was held shut with a Frangibolt from launch through the Outbound Cruise, and then the cover was opened during the Approach Phase to begin collecting data and performing external calibrations. The main data collection period for REXIS was the Orbital B mission phase, so between cover opening and Orbital B, various calibrations were performed to optimize instrument parameters. Further discussion of REXIS flight operations are included in Chapter 3.




Figure 1-3: Model of the REXIS instrument [4]

From its proposal in 2010, to the end of its main science mission in 2019, over 80 undergraduate and graduate students at MIT and Harvard have worked on REXIS. Figure 1-5 shows a plot of students working on the project over time, from the proposal through January 2019. Besides being the second student-built instrument to fly on a NASA planetary mission (the first was the Student Dust Counter on the New Horizons mission to Pluto [5]), REXIS is the first planetary surface mapping instrument designed to utilize coded aperture imaging, and is the first instrument to




Figure 1-4: Photograph of assembled REXIS instrument [2]

fly CCID-41s with a directly deposited optical-blocking filter [4, 6].

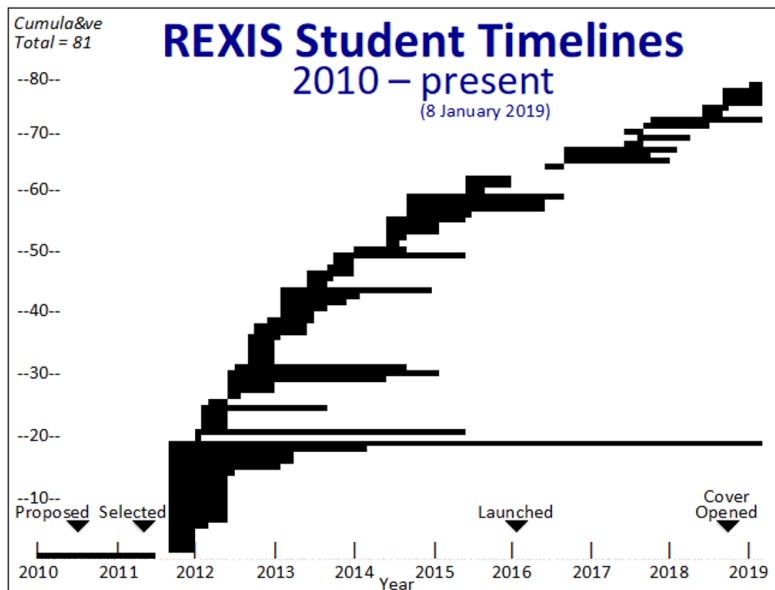


Figure 1-5: Timeline of students who worked on REXIS. Figure created by Professor Rick Binzel.

1.3 Motivation for Thesis

During the REXIS instrument’s main science period (Orbital B), the instrument experienced anomalous behaviors that resulted in the loss of roughly half of the usable science data collected. The remaining data were not sufficient to reach the REXIS science goals. There were periods of data collection where downlinked event rates from the on-board iron calibration sources decreased or disappeared from the data entirely. The loss of the calibration source counts resulted in higher uncertainty in the accuracy of the data taken during periods of time with fewer iron counts. This loss of iron counts served as a measurement of overall loss of detector efficiency during Orbital B. This thesis examines the decrease of detector efficiency by investigating the loss of iron calibration counts, and presents a root cause analysis of the issue.

NASA assigns a class specification (Class A, B, C, or D) to each mission, based on factors such as priority, significance, lifetime, and cost, with Class A as the highest priority and high cost, and Class D as lower priority and cost [7, 8, 9]. REXIS is considered a Class D instrument on Class B OSIRIS-REx mission, but is primarily a Do-No-Harm instrument. This designation means that the payload assumes a high technical risk, and should present no safety concern to the hosting spacecraft or other payloads [10]. The project has to balance higher technical risk in favor of staying within budget and schedule margins. If the REXIS instrument was not completed in time for the integration and launch of OSIRIS-REx, the mission had a mass model of REXIS to be attached to the spacecraft in its stead. Designing a Class D or Do-No-Harm instrument can present unique challenges in producing the instrument with limited resources. This thesis presents a CAST analysis of the identified root cause, with the goal of producing recommendations for future similar low-cost, high-risk instruments.

1.4 Thesis Roadmap

Chapter 2 contains background information about the REXIS detectors, Root Cause Analysis techniques, and an introduction to the CAST accident analysis method. Chapter 3 contains a timeline of REXIS in-flight operations since the launch of OSIRIS-REx in 2016, and discusses anomalous behaviors observed throughout flight. Chapter 4 contains the root cause analysis of the detector efficiency loss due to iron count dropout and the identification of the most likely candidate. Chapter 5 investigates the most likely cause identified through RCA, and contains the results of a CAST analysis performed on the REXIS project to determine other contributing factors to the loss, such as external or internal pressures, ineffective communication, or changes within the project over time. Chapter 6 summarizes the findings of the root cause analysis and the CAST analysis, and offers recommendations for future instruments.

Chapter 2

Background

To provide context for the root cause analysis and the CAST analysis, background information about the REXIS CCD design and image processing is outlined. Various root cause analysis methods and tools are described. The CAST method and underlying theory are also described.

2.1 REXIS Detectors

The REXIS detector plane is comprised of four back illuminated charge-coupled devices (CCDs) [4]. CCDs are detectors that utilize the semiconductor properties of silicon and the photoelectric effect to detect particles [11]. The devices are comprised of a grid of pixels. These pixels act as potential wells and are separated by charge-trapping gate structures on the silicon surface. Particles hit the silicon surface of the detector (each particle hit is called an “event”), and essentially knock electrons loose. The resulting electrons accumulate in each potential well. After a set amount of time, the CCD is clocked – a voltage is applied to each gate, which allows the charge to flow through to the readout area of the detector, emptying the pixels of charge. After each pixel is emptied out, the next clock cycle begins. The REXIS detectors constantly collect particles, and have a four second integration period after which the charges on the detectors are read out [12].

When the CCD is clocked, the charges in each pixel are converted to a digital pulse

[11]. The energy of that pulse is directly proportional to the incident energy of the particle that hit the detector. By looking at the pulses during each clock cycle, the initial energies of the incident X-ray photons can be determined. As stated in Chapter 1, the energy of a fluoresced X-ray is intrinsic to the material it originated from. Therefore, the CCDs capture the fluoresced X-rays, and on-board software compiles the X-rays into a list containing events with the desired energy range, specified by input parameters. The list is downlinked from the spacecraft, and run through a ground pipeline for final analysis and to create the final data products.

The CCDs are arranged in a two by two array, and are mounted in a Detector Assembly Mount (DAM), shown in Figure 2-1. Each CCD is 1024 by 1024 pixels, and is separated into four 256 by 1024 pixel nodes [6]. The physical layout of the CCDs and nodes are shown in Figure 2-2a, with a frame taken during flight shown for reference in Figure 2-2b. (For future references to nodes in this thesis, a shortened format is used, with the CCD indicated by the first number, and the node of the CCD indicated by the second number. For instance, CCD 2 node two is written as CCD02, or 02). Ten radioactive ^{55}Fe sources are mounted around the perimeter of the detector array. These sources are used for calibration of the CCDs both during ground testing and during flight operations [4]. There is another ^{55}Fe source mounted to the bottom of the radiation cover, which was used for in-flight calibrations while the cover was closed. Two months before official spacecraft arrival at the asteroid, the REXIS radiation cover was opened, taking the cover source out of view of the detectors. After cover opening, known astrophysical sources such as the Crab Nebula were used for broader spectra calibrations. The CCDs also have an optical blocking filter deposited directly on the detectors to prevent optical light from reaching the imaging surface.

REXIS CCD Damage

The CCDs are connected to the Video Board electronics via flexprints (flexible circuit board cable) [4]. During a vibration test performed on the fully-assembled spectrometer, two flexprints were accidentally torn. Figure 2-3 shows the CCD wiring. The





Figure 2-1: Model of DAM structure [4]

CCD23	CCD 2	CCD 3	CCD33
CCD22			CCD32
CCD21			CCD31
CCD20			CCD30
CCD13	CCD 1	CCD 0	CCD00
CCD12			CCD01
CCD11			CCD02
CCD10			CCD03

(a) Physical layout of the REXIS CCDs, with nodes labelled.



(b) Example of flight data showing physical distribution of the REXIS CCDs. Image created by HCO scientist Daniel Hoak.

Figure 2-2: Physical layout of the REXIS CCDs in a drawing (left) and in flight data (right). In the drawing, the wider lines indicate the bounds of the full CCD, and the thinner lines mark the boundaries of the nodes within each CCD.

specific wires affected by the tear in the flexprint are indicated by the boxes. CCD 3 was the most impacted by the tear, which caused two of the four nodes on the CCD to become non-functioning [7]. One node on CCD 2 was also rendered non-functional due to the flexprint tear. This accident brought the number of operational nodes to be used for data collection down from 16 to 13. To salvage the other CCDs and nodes, arathane was placed over the tears on the flexprints to prevent further tearing [7].




Figure 2-3: Diagram showing CCD wiring, with boxes around the lines cut during the accident. [7]

Additionally, during subsequent ground testing, some nodes were discovered to be too noisy to obtain satisfactory signal resolution. These nodes were not used for flight operations. In total, only 10 of the original 16 nodes were used for in-flight operations: nodes one through three on CCD 0, all four nodes on CCD 1, and nodes zero through two on CCD 2.

2.1.1 REXIS Image Processing

Image Processing Algorithm

REXIS has two operating modes, Science Mode and Safe Mode. All data collection and processing is done during Science Mode. Safe Mode is the mode entered upon powering the instrument on, and before powering the instrument off. REXIS data are downlinked when telemetry is on, but some housekeeping information, such as overall

event rate, is recorded when telemetry is off as well. Image processing parameters can be set either before or after entering Science Mode. A block diagram of the image processing algorithm is shown in Figure 2-4.

When REXIS is first put in Science Mode, a bias map is generated. This involves taking 10 frames, and calculating the median charge readout value for each pixel. This bias map is then subtracted from all subsequent images to remove instrument noise and dark current, isolating actual signal.

REXIS downlinks science data in the form of a candidate X-ray event list (CXEL). As full frames (all 1024 by 1024 pixels for each of the four CCDs) are read out, the image processing algorithm removes the energies that do not meet the criteria as set by the image processing parameters. When in Science Mode, REXIS data are constantly collected, but are only downlinked to the spacecraft if the telemetry flag is set. When the telemetry flag is set, all pixels of all four CCDs are read out at the designated clocking rate (four seconds is the default). For each frame, each pixel is first corrected for readout electronics drifts by adding an offset of 1.3 keV (to ensure the value remains positive as they are stored as unsigned integers) [12]. Next, the energies from any pixels specified by the HPM are discarded. Then each pixel has both its associated bias map value and overclock value subtracted from it. Finally, the remaining values are compared to the specified Event Threshold (ET) value. The values that are greater than ET are added to the CXEL, which is then passed to the flight software for further processing [12].

The flight software feeds the resulting CXEL through an event grading algorithm, using the grade definitions shown above in Figure 2-5. Before the data is processed by the algorithm, the CXEL is first reduced to contain only the maximum event value in a three by three grid, to reduce double counting of pixel energies as multiple events [4]. The software also removes any events that have energies higher than the designated Upper Level Discriminator (ULD) [12]. Then, that reduced CXEL is processed by the grading algorithm, using the input Split Threshold (ST) values to determine if an adjacent pixel has sufficient energy to be considered part of the central event, and each event is assigned a grade from 0 through 7.

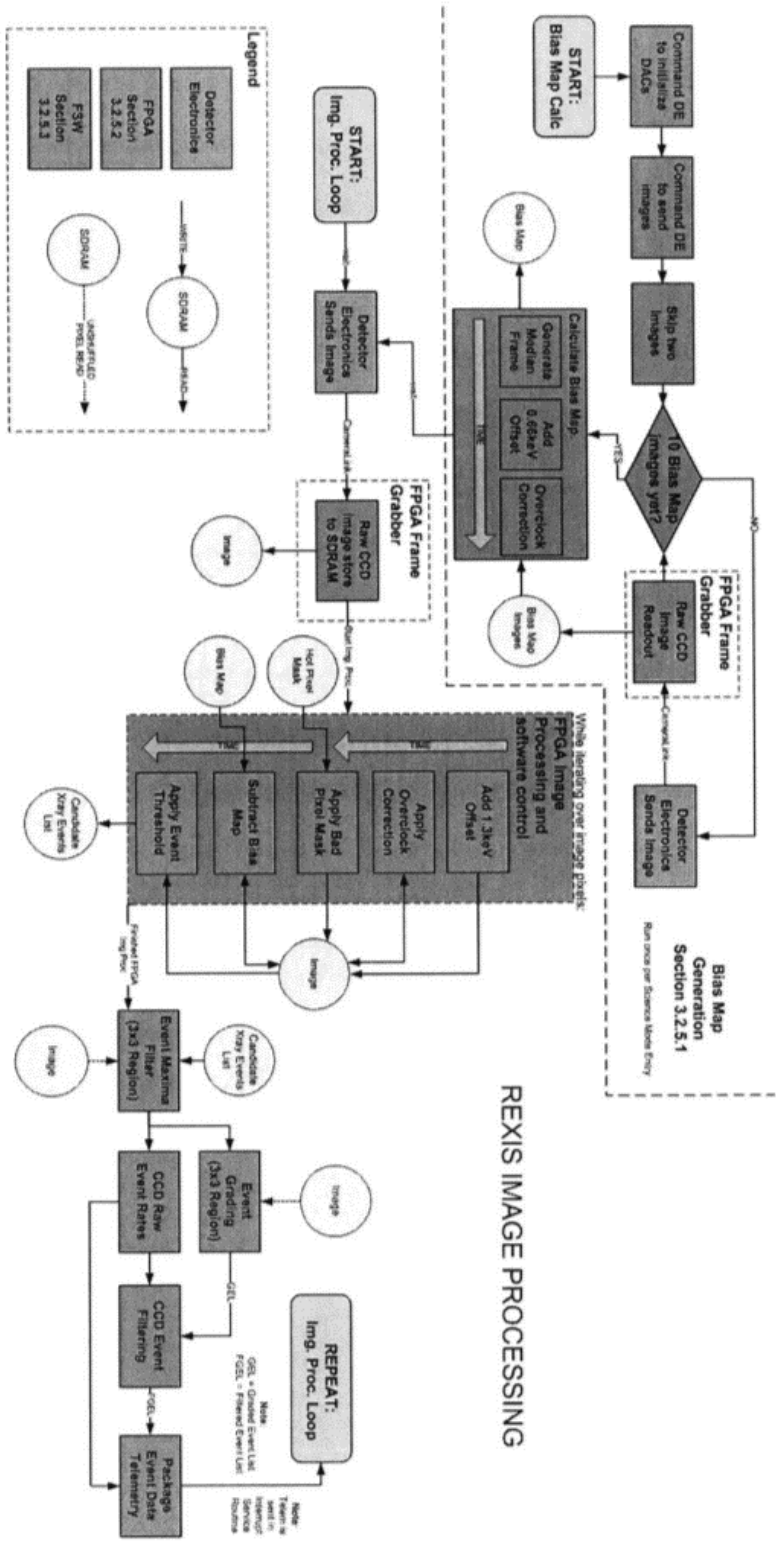


Figure 2-4: Image processing block diagram. Image from the Master's thesis of REXIS student Pronoy Biswas [12]

If the total amount of events per frame is higher than the specified event rate limit, which has a default value of 800 counts per frame (200 counts per second), then the software implements random event filtering [12]. When random event filtering occurs, events are randomly selected and discarded until the length of the CXEL reaches approximately the event rate limit. The final product downlinked from the spacecraft is an event list, which includes the position, energy, and grade of each event.

Image Parameter Settings

One full frame from the REXIS CCDs is 8 megabytes, so in order to achieve a smaller data volume, most of the REXIS image processing is done on board, using both the field-programmable gate array (FPGA) and the flight software [12]. The resulting downlinked data per four-second frame are stored as an x-ray event list that is only a few hundred bytes. The image processing algorithm, described in Section 2.1.1, involves setting multiple parameters for each CCD, as well as setting different parameters for the four nodes on the same CCD.

There are three parameter values that can be set differently for each node: the Event Threshold (ET), the Split Threshold (ST), and the Upper Level Discriminator (ULD) [12]. These parameters require 16 input values for each issued command. The ET sets the minimum energy of interest, the ULD sets the maximum energy of interest, and the ST sets the minimum energy a surrounding pixel needs to contain in order to be considered a part of the event on the central pixel.

The ST follows the grading scheme from ACIS, shown in Figure 2-5. Event grades indicate the physical spreading of the energy from an X-ray event between adjacent pixels. The software looks at a three-by-three grid surrounding the central pixel. If any pixels are above the energy set by ST for the node in which the pixels reside, they are considered part of the event. If there are no surrounding pixels with energies above ST, the event is marked as grade zero. If there is one or more pixels with sufficient energy, but none share a side with the central pixel, the event is marked as grade one. If there is one adjacent pixel above ST, the event is marked as grade two,




Figure 2-5: Grading scheme used for REXIS, adopted from Chandra’s ACIS instrument [4]

three, four, or five, depending on the position of the pixel relative to the central one. If more than one adjacent pixel in one corner of the grid has sufficient energy, the event is marked as grade six. All other combinations of triggered pixels are counted as grade seven, and are considered non-physical events.

There are two parameters that apply to all 16 CCD nodes: the Event Grade Filter (EGF), and the Node Enable Map (NEM). The EGF tells the flight software which grade events to exclude from the final downlinked CXEL. The NEM tells the software which nodes to include data from; if a node is not identified in the 16-bit NEM value, all events from that node are ignored by the image processing algorithm and do not continue through image processing.

One final image processing tool is the Hot Pixel Mask (HPM). The HPM consists of pixels that are “hot” - seeing more X-ray events per frame than seems physical - and is determined by looking through past observations and calibrations to identify

“hot” areas of the detector. Hot pixels can be due to radiation damage, charge traps, or pinholes in the directly deposited Optical Blocking Filter (OBF) [4].

2.1.2 Optical-Blocking Filter

The purpose of the optical blocking filter on the REXIS CCDs is to prevent optical light from reaching the detector surface, where the light would register as noise and make it difficult to isolate X-ray signal. The REXIS CCDs have 220 nanometers of aluminum deposited directly on the imaging surface to mitigate the effects of optical light [6]. Pinholes (microscopic areas with less or no OBF coverage) were observed during testing of the deposition methods, and were noted as a possible risk on the flight model. To mitigate the effects of optical light shining through pinholes to the detector surface, the REXIS flight software has the capability to ignore data from specific pixels through use of the Hot Pixel Mask software tool [12].

The directly deposited optical-blocking filter (OBF) was developed in parallel with the design process of REXIS through a partnership between the Lincoln Laboratory and the MIT Kavli Institute for Astrophysics and Space Research (MKI) [13]. Rather than have filter components mounted separately, as past missions had done [13], an aluminum coat was directly applied to the imaging surface of the back-illuminated CCDs [13]. Directly deposited filters had been used on front-illuminated CCDs as well as hybrid CMOS X-ray detectors [14, 15, 16]. However, REXIS is the first instrument to fly the back-illuminated CCD-41s with a directly deposited OBF [6].

2.2 Root Cause Analysis Methods

Root cause analysis is a blanket term that can describe many tools, methodologies, and systems by which to identify the cause (or causes) of a problem [17]. There are many tools and methods used for root cause analysis; sometimes multiple tools are even used for the same analysis. This section will briefly describe four commonly-used root cause analysis tools and methods: the 5 Whys method, failure modes and effects analysis (FMEA), fault tree analysis (FTA), and an Ishikawa (fishbone) diagram.

2.2.1 5 Whys Method

One commonly used method is called the “5 Whys” method. This method is explained by its name - once the problem is defined, ask “why?” as many times as necessary to get to the root of the issue (it is not always necessary to ask five times, and for more complex systems, five times may not be sufficient)[17]. After identifying the answer to the first “why”, using evidence or analysis, it is determined whether the issue could still persist if the identified cause was fixed. If the answer is still yes, another round of “why” begins, and if the answer is no, the root cause has been reached. The lowest level may be reached before five why’s; the question is asked until the answer is something that cannot be controlled [18]. Sometimes an issue will have multiple causes, which can branch off from one round of asking “why”. In that case, each separate branch is followed down to its conclusion. An example of a one-branched 5 whys analysis for a transistor chip with degraded performance is shown in Figure 2-6. Through asking “why” five times, the root cause of the degradation - that a wafer table seal had reached its end of life - was identified. Once the root cause or causes are identified, plans can be made to address and correct each cause [17]. The 5 whys method is not commonly used in aerospace, however. The method was refined and used by automotive companies [19], and has been used in debugging software [18] and industrial engineering [20].

2.2.2 Fault Tree Analysis

Fault tree analysis (FTA) is a top-down approach to root cause analysis [21]. FTA is a Boolean logic-based approach, and was initially developed by Bell Laboratories [21]. FTA diagrams are comprised of two main objects - logic gates and events. This analysis can be performed before an error occurs as a preventative measure, or as a measure to more quickly fix any issues that arise. First, an undesirable event is identified, which is known as the “fault condition” or the “top level failure” [21]. Next, possible causes of that top level failure are determined through analysis of the design of the system and the various failure probabilities of components, and are placed

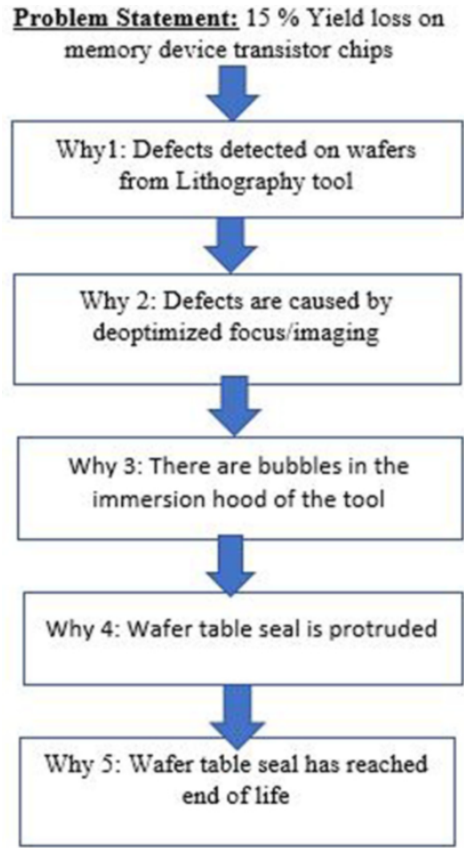


Figure 2-6: Example of a 5-whys analysis, performed on a transistor chip with degraded memory [19]

in the second level of the diagram and connected to the top level failure via the appropriate logic gate. Subsequent causes of each event at each level are determined until the analysis reaches a single component failure or error. From the completed diagram, subsets can be examined to determine the minimal set of adverse events or failures that could cause the top level failure.

The basic gate symbols used in FTA diagrams are depicted in Figure 2-7: the gumbdrop-type symbol represents an AND gate, the pointy symbol is an OR gate, the two gates merged with the triangular NOT gate are priority AND and exclusive OR gates, and the hexagon is an inhibit gate. The line at the top of all the figures indicate the output value, and the lines on the bottom indicate the inputs. Gates may have more than two inputs. AND gates are true when all inputs are true, while OR gates are true when one or more inputs are true. Priority AND gates are only

true if the inputs are true in a specified sequence. Exclusive OR gates are true only if only one of the inputs are true. The output of the inhibit gate is only true if a specific enabling event occurs [22].

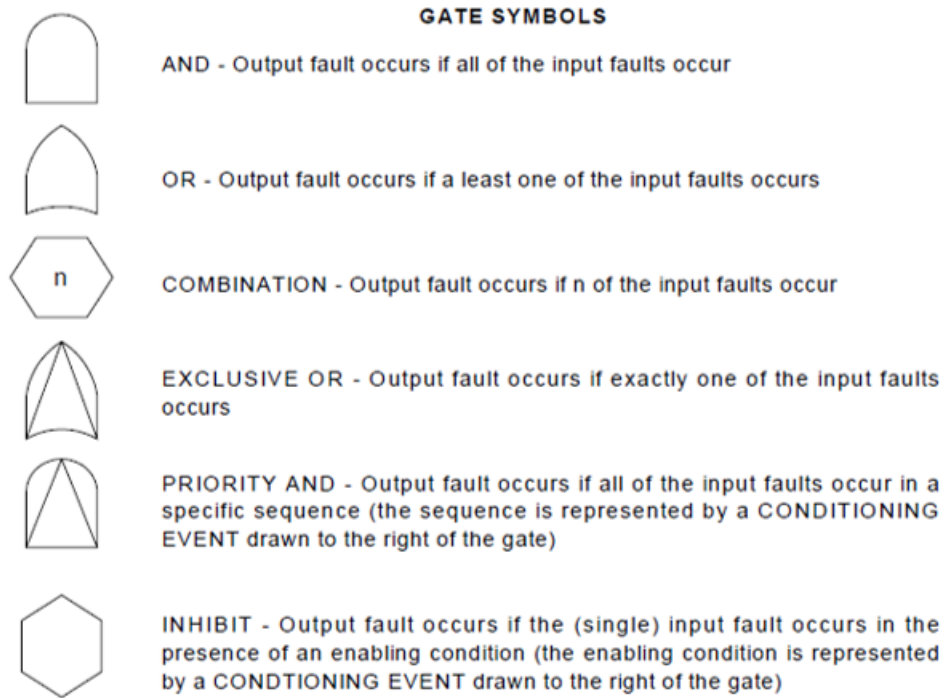


Figure 2-7: Logic gate symbols used for FTA [22]

The inputs and outputs of all gates are the different types of events shown in Figure 2-8. The two most commonly used events are the intermediate events, represented by rectangles, and the basic events, represented by circles [22]. Basic events require no further breaking down, and are considered the root causes. Intermediate events are failure events that occur between the top level failure and the lowest level basic events. The rectangular shape with a pointed top represents an external event, that is nominally expected to occur and not in itself considered a fault or failure. The diamond-shaped events are undeveloped events, about which insufficient information is available. An example of this would be identifying that the cause of a top level failure could be a computer system failing. There a multiple ways a computer system could fail, but for the purpose of the FTA and system safety development, it is only necessary at the moment to know that the computer system failing could be

a root cause. If the top level failure does occur during testing or operation, then the undeveloped events would be investigated further. Conditional events, shown by ovals, are the inputs to inhibit gates, and specify states that are necessary for an event to cause a failure event.

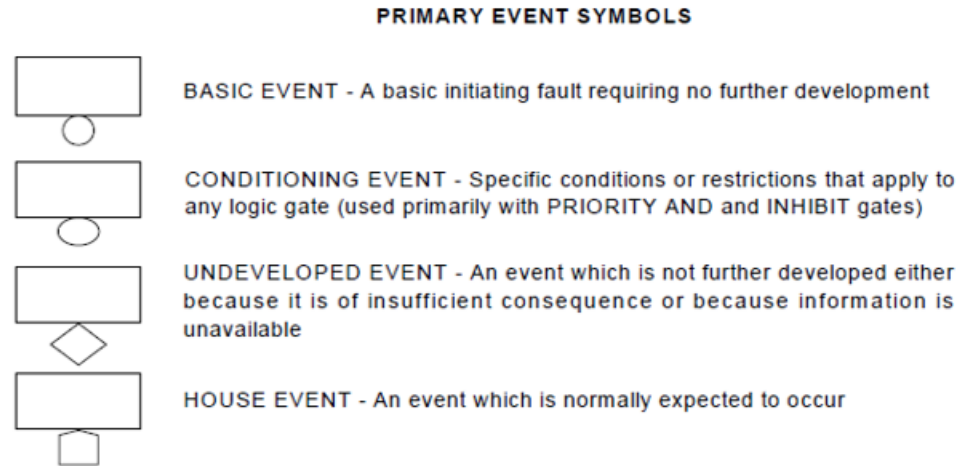


Figure 2-8: Event symbols used for FTA [22]

An example FTA diagram is shown in Figure 2-9. In this example, D fails when both A fails and either B or C fail. NASA missions utilize FTA in mission or component failures [23], and have a detailed handbook with instructions for how to perform the analysis [22].




Figure 2-9: Example Fault Tree from official NASA FTA handbook [24]

2.2.3 Ishikawa Diagrams

Ishikawa or Fishbone diagrams are another common tool used for root cause analysis. An example diagram is shown in Figure 2-10. Each initial branch off the main trunk (or 'bone' that branches off the 'body' of the fish) represents a type of issue. The most common categories used for these divisions are called the six M's: Man, Material, Machine, Method, Measurement, and Milieu (a.k.a. the environment) [21]. However, depending on the system or the issue, some of the six M's are inconsequential; each root cause analysis can have its own unique fishbone diagram layout. Once each branch is defined, different possible causes of the issue are placed along the bone with the category that most closely encapsulates the cause. Then, the team identifies which causes are most and least likely, through various analyses and tests [21]. Once the most likely cause or causes are identified, efforts to fix or mitigate the issue can begin.

The diagram shown in Figure 2-10 comes from a root cause analysis performed on an in-flight anomaly of the Visible Infrared Imager Radiometer Suite (VIIRS) [24]. The fishbone diagram was compiled to help determine the root cause of an observed optical degradation of the instrument. Through investigation of the possible causes identified on the fishbone diagram, the team was able to narrow down two separate causes for the anomaly, identified by red text on the diagram [24]. Fishbone diagrams are used in aerospace anomaly investigations as a beginning point for a root cause analysis, and the identified bones guide the investigation [24, 25].

2.2.4 Failure Modes and Effects Analysis

Another method commonly used for root cause analysis is called failure modes and effects analysis (FMEA). FMEA is often performed during the design process of a system, as a method to both drive safer system design and expedite any root cause analyses performed later in the life cycle of the system (i.e. during testing or operation). If a failure or error occurs, the root cause analysis process could be shortened by first looking through the already identified effects, and tracing back to the poten-

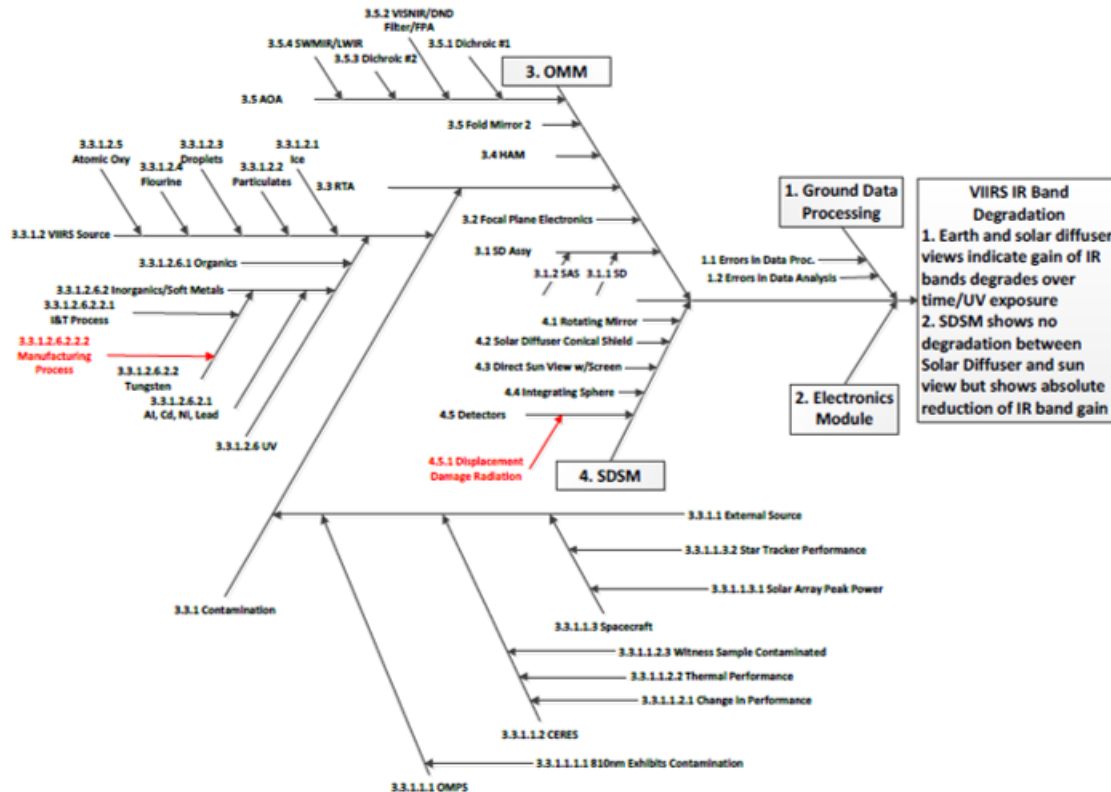


Figure 2-10: Example Fishbone diagram used in root cause analysis of an optical throughput decrease for a NASA instrument [24]

tial identified failure mode or modes. The primary focus of this method is identifying single failure modes and the associated effects on the system if the failure modes occur [21].

FMEA is a bottom-up approach to analysis; to perform a thorough FMEA, each part of the system is investigated, and failure modes are considered for every component [21]. The failure probability of each component is recorded as well [21]. Once failure modes are identified, the next step is to identify the associated failure mechanisms (causes of the failure mode), and finally identify the failure effects (behaviors of the system while experiencing a failure mode) [26]. FMEA can also be performed on a physical system (Design FMEA) or the manufacturing of the system (Process FMEA) [26]. FMEA was first developed by the United States Defence Department in 1949, and was used by NASA as a method of ensuring system safety for the Apollo missions [27, 28, 29], and is often used in conjunction with FTA [22].

2.3 CAST Analysis

A different approach to investigating accidents is called Causal Analysis based on System Theory, which comes from the STAMP theory (System-Theoretic Accident Model and Process), developed by Dr. Nancy Leveson of MIT [26]. STAMP theory, and by extension CAST, aim to learn more from incidents and accidents by studying a wider system boundary. Rather than focusing on design- or process-related root causes (such as the RCA methods discussed in Section 2.2), STAMP theory considers systemic causes, and looks within the organization, management, culture, and communication of the teams involved with the project. This section provides a brief overview of STAMP theory, and a discussion of the basic steps in a CAST Analysis.

2.3.1 STAMP Theory

STAMP is built upon systems theory, which began in the 1930s and 1940s, and was developed in response to systems becoming more complex and therefore difficult to analyze completely with the typically employed “divide and conquer” methodology [26]. Instead of breaking down systems on a component level, systems theory focuses on the emergent behavior of a system, and also considers how social aspects relate to the technical aspects [26]. STAMP utilizes systems theory to redefine safety as a control problem instead of a reliability problem, taking into account not only component failure, but component interaction failure, and considers the controls that exist between and within the technical and social aspects of a system [26].

2.3.2 Accident Analysis using CAST

CAST accident analyses use the concepts of STAMP to investigate not only what the root causes of an accident are, but why those events occurred [26]. The goal of a CAST analysis is to identify what went wrong, and to make recommendations for future similar systems in order to prevent similar accidents in the future [26]. In order to prevent hindsight bias in the analysis, emphasis is not placed on finding a specific person or people to blame; rather, the emphasis is placed on fully understanding the

various facets of why the accident occurred [30]. So rather than ask "who", CAST analyses more often ask "why" and "how". As the CAST Handbook says, "Blame is the enemy of safety" [30]. By placing blame on a person rather than finding intrinsic errors within an organizational structure, similar accidents may continue to occur.

In order to perform a successful CAST accident analysis, an important term to define precisely is "failure". For CAST analyses, failure is when a component does not act as designed, or breaks [30]. Therefore, people do not "fail" if they do not perform a specific task or follow an instruction. Companies do not "fail" unless they go bankrupt and out of business. Another misuse of the word "failure" is when it is applied to software. Software performs the logic it is provided. If software performance does not match what is expected, then the logic is flawed and must be edited; the software did not fail. Removing the word "failure" from accident analysis helps to move away from placing blame, and towards helping design better systems in the future.

Some other important terminology, with definitions taken directly from the CAST Handbook [30], are as follows:

- **Accident:** An undesired, unacceptable and unplanned event that results in a loss.
- **Loss:** The result of an accident. Loss can refer to many things; loss of life, loss of mission, loss of data, etc.
- **System Goals:** The reason the system was created in the first place.
- **System Constraints:** The ways that the system goals can acceptable be achieved.
- **Hazard:** A system state or set of conditions that, together with specific environmental conditions, can lead to an accident or loss.

Steps of a CAST Analysis

The five main parts of the CAST method are shown in Figure 2-11. The majority of the analysis is performed on the system safety control structure. A safety control

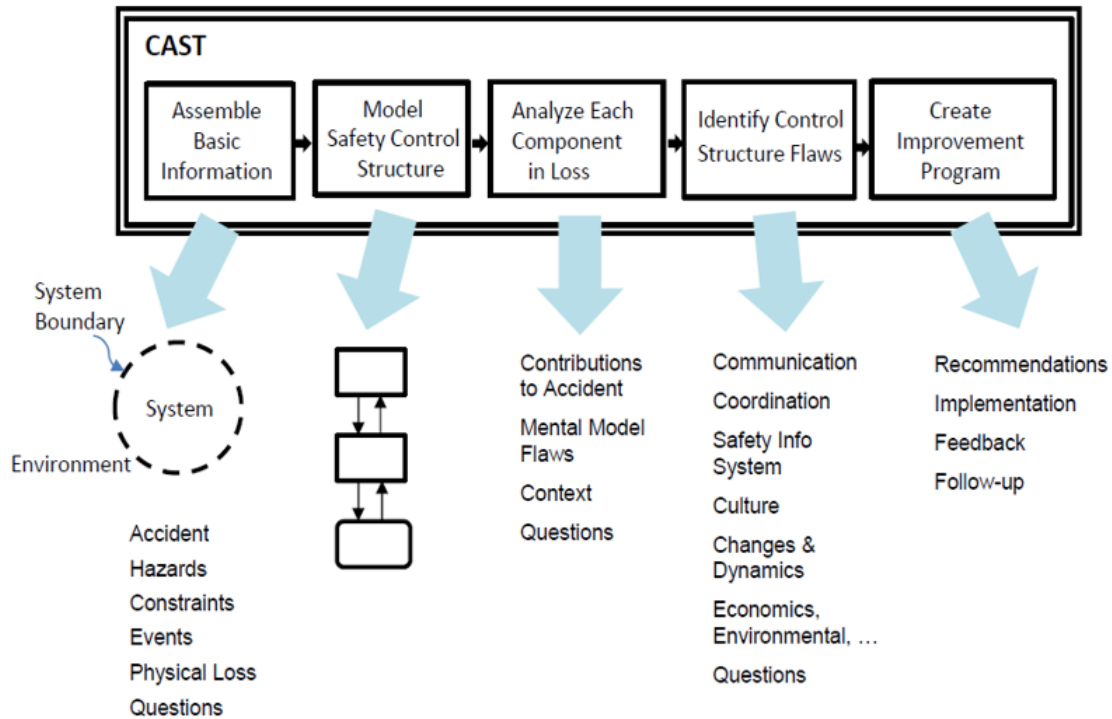


Figure 2-11: Overview of CAST method [30]

structure contains all technical and social components of the defined system, and includes the “controls” between each component. The controls are the actions or information that flow between two components within the safety control structure, and are represented by a directional arrow. A control may be an instrument providing its state, an operator sending a command, or information passed between two groups within the project hierarchy. The project hierarchy and the directness of control over the physical system are represented in the control structure. Items near the bottom of the structure are lower in the project hierarchy, and vice versa. Groups or parts of the system only directly interact if there is a control drawn between them. Accident analysis utilizing the CAST method begins with a high-level control structure, and then zooms in to the more involved and complex parts of the system. For instance, a block can simply be defined “Project Management” at first, but further iterations can include multiple personnel or subgroups within the “Project Management” block.

In the first step of the CAST method, basic information about the system and the accident under study are compiled [30]. The system and the boundary of the

analysis are defined, and the loss, and hazards that led to the loss, are described. Once the hazard(s) are defined, the associated system-level safety constraints required to prevent that hazard or hazards are identified. Next, the event is described as a sequence of events, and the person or team performing the analysis generates questions to answer about why the events occurred. Then, the physical loss is analyzed in terms of the physical components of the safety control structure, as well as the requirements on the physical design. Finally, any failures or unsafe interactions that led to the hazard, and missing or inadequate controls that could have prevented the accident are all identified.

In the second step, the safety control structure is modeled[30]. If a control structure does not already exist, a more abstract high level model can first be constructed, and then add more components and controls as the investigation proceeds. The control structure can change over the course of the accident investigation as the result of new information or new understanding of the system. An example control structure is shown in Figure 2-12. This control structure was part of an analysis of a U.S. Coast Guard aviation mishap. This control structure shows the divisions involved in the system development. The arrows between each block represent the necessary controls or interactions between the specified blocks to prevent the hazardous state identified in step one from occurring. The responsibilities of each component in the structure are also documented. The control structure can evolve throughout the process of the investigation, so as components or connections are added, the information in the subsequent steps of the analysis also evolve.

In step three of the CAST method, the components of the control structure are examined individually to determine why each component (also referred to as “controllers”) was not effective in preventing the loss or hazard [30]. Beginning from the bottom of the structure, the role of each component in the accident is documented, and then questions are composed as to why specific behaviors occurred, and why controllers acted the way they did. The goal is to answer each question raised in this step by the conclusion of the CAST analysis. Some questions may remain unanswered based on the level of information available at the time of the accident analysis; how-

CONTROL STRUCTURE – System Development (Detailed View)

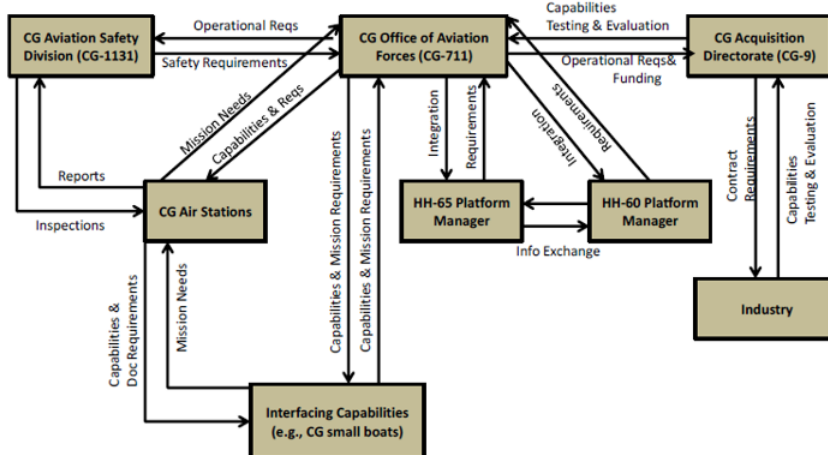


Figure 2-12: Example control structure from a CAST analysis of a U.S. Coast Guard Aviation mishap [34]

ever, even unanswered questions can result in recommendations for future projects.

Step four of the CAST method is to analyze the interactions between controllers, and look at the behavior of the control structure as a whole. To perform this part of the CAST method, systemic factors - factors such as communication and coordination between members of the project team, the overall safety culture of the organization, or how the organization changed over time - are considered [30]. This portion of the CAST method is performed in the same way as the previous two, through looking through documentation and interviewing various people within the project. For this step, it is important that the boundary of the system under study has been clearly defined, so the investigation can focus on the systemic factors closer to the portions of the system that led to the hazardous state or loss.

The final step of the CAST method is to create an improvement program. Once most of the questions have been answered, and the wide range of influencing factors towards the loss have been investigated, recommendations for future similar projects are made. This step may also include designing ways to implement the recommendations, giving feedback to the project that experienced the loss, and following up with the project in the future to determine the level at which the recommendations have been implemented.

The CAST method has been used in various industries. It has been used to investigate healthcare system anomalies [31], railroad or ferry accidents [32, 33] and airplane accidents [34]. However, the CAST method has been slow to be adopted by the aerospace industry. The majority of the analyses performed focus on large accidents. In this thesis, the analysis is performed on a smaller system (a spaceflight instrument), where the loss was a loss of science data.

Chapter 3

REXIS Operations

The general timeline for REXIS operations is shown in Figure 3-1. In addition to the REXIS-specific calibration activities, payloads on the spacecraft had checkout activities between launch and asteroid arrival to track the health of each instrument. For reference, these checkout and calibration activities are labelled as "L plus" some number. Unless otherwise indicated, the number roughly corresponds to the number of months since the launch of OSIRIS-REx. For instance, "L+6" indicates a period six months past the September 2016 launch date. The one exception to this nomenclature is L+14, which instead indicates that the operation occurred 14 days after launch, not months.

Figure 3-2 shows the plan for the instrument calibrations anticipated as necessary when first designing the operations plan for REXIS. All calibration activities needed to be completed before Orbital B operations. The internal calibration needed to be performed while the radiation cover was closed. Some of the external calibrations needed to be performed before a certain date due to sun angle and based on when certain objects in the sky would be visible. However, some operations, such as the CXB Calibration, were order agnostic as to when they occurred between cover opening and Orbital B.

As mentioned in Section 2.1, only 10 of the 16 nodes in the detector array were functioning well enough to be used for operations. As more data were obtained throughout flight, other nodes presented with some resolution issues, which resulted


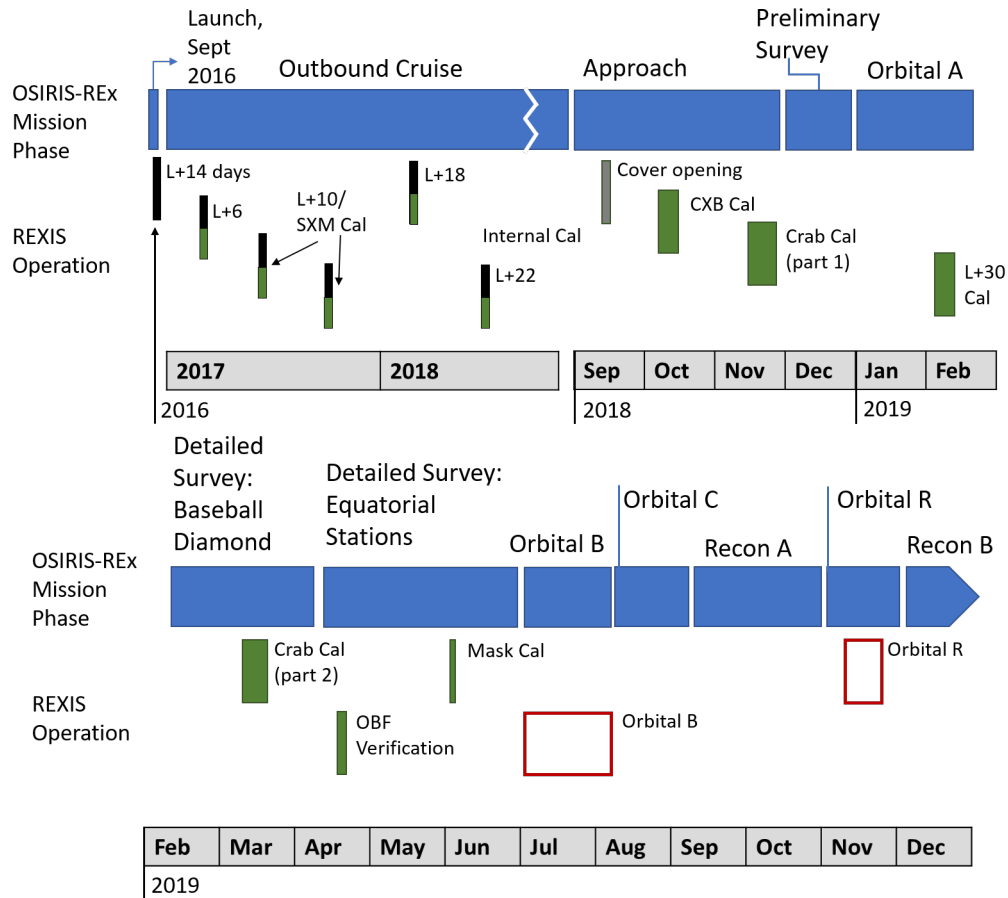


Figure 3-1: REXIS operational timeline

in a further down-selection of nodes. Figure 3-3 is a chart showing the total number of nodes used for each operation over time. If an operation used fewer than the nominal 10 nodes, the reasoning behind the decrease is described in the associated section.

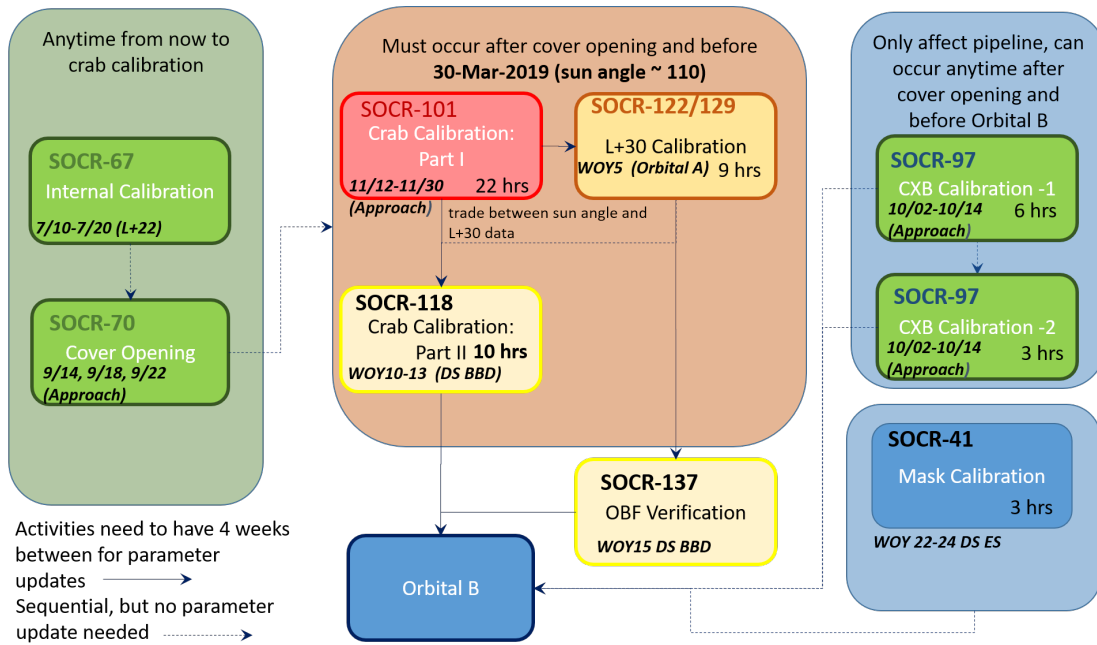


Figure 3-2: Planned REXIS Operations throughout flight

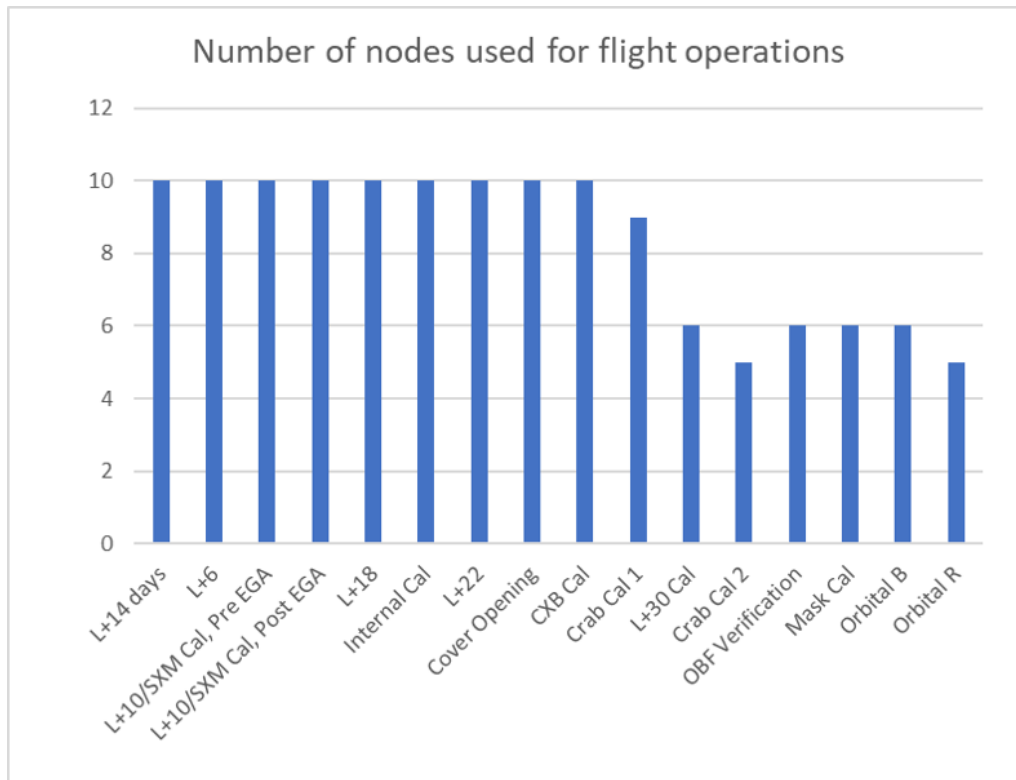


Figure 3-3: Number of nodes used for each operation, in chronological order

The nominal timeline for planning and implementing a sequence for an operation for OSIRIS-REx was eight weeks. At six weeks prior to the operation, no more command changes were allowed. At four weeks prior, no more parameter changes were allowed. Throughout the planning process, various products were passed between the REXIS team and the OSIRIS-REx planning team to ensure the operation proceeded precisely as planned. Some REXIS external calibration operations occurred closer together than eight weeks, meaning lessons learned about parameters from one operation could not be implemented in the operation immediately following.

3.1 Early Checkouts and Calibrations

3.1.1 Launch + 14 Days Payload Checkout

OSIRIS-REx launched on September 8, 2016. The first time that REXIS was turned on in flight was 14 days after launch, on September 21, 2016. The primary purpose of this payload checkout was to confirm functionality of the instrument, and get a preliminary measure of the performance of the CCDs and the SXM in flight. The payload checkout sequence consisted of two hours of functionality tests of changing different parameters on board, as well as taking a raw frame and downlinking some event list data to measure the energy resolution and response of the CCDs. The same checkout sequence was used in all subsequent checkout activities. The ten nominal nodes were used for this operation. From the results of this operation, a few hot pixels on the detector plane were identified, and functionality of REXIS was confirmed. No new anomalies were observed.

3.1.2 L+6 Checkout and Calibration

On March 14, 2017 another REXIS checkout was performed to track the overall health of the instrument. Between March 16 and 17, 2017 a separate instrument calibration sequence was run. The primary goal of the calibration sequence was to quantify the REXIS detector performance, and to search for the optimal CCD parameter settings

(ULD, ET, and ST), and optimal SXM threshold (TT). The calibration sequence consisted of ten total hours of event list data collection, and downlinking one full raw frame. Throughout the event list data collection, the instrument used various sets of ULD, ET, ST, and SXM TT. The set of ten nominal nodes was used for this operation. The L+6 checkout provided further information for long-term trending of the behavior of the CCDs by establishing baseline health and performance metrics to compare against future flight data. From the L+6 calibration activity, the instrument team's understanding of how individual node parameters - ET, ST, and ULD - affect the energy resolution and response of the detectors was improved, and these results informed future calibration operations.

The L+6 calibration also resulted in a better understanding of the SXM threshold. It was discovered that lowering the SXM threshold too much caused an abrupt cutoff of the detected solar signal, due to the threshold ending up well below the input solar signal. From these results, the REXIS team decided to continue using the original higher SXM TT value.

3.1.3 L+10 Checkout and SXM Calibration

The Earth Gravity Assist (EGA) of the OSIRIS-REx spacecraft occurred in September of 2017. In order to reach the correct velocity to get to Bennu, OSIRIS-REx used the gravity of Earth as a slingshot to increase the spacecraft's speed. The REXIS L+10 checkout and calibration was split into two parts, which occurred before and after the OSIRIS-REx EGA. For both of these operations, the ten nominal nodes were used.

The L+10 checkout sequence was performed on August 3, 2017, and the portion of the REXIS L+10 calibration that occurred before the OSIRIS-REx EGA happened between August 13 and August 17. The checkout involved the same activities as L+6 and L+14, and was performed to further the long-term tracking of the health and performance of the CCDs. The calibration activities consisted of further exploration of the CCD parameter settings in order to optimize them for asteroid science data collection. The calibration also included pointing requests for the spacecraft such

that the SXM was not entirely pointed at the sun. The purpose of this portion of the calibration sequence was to observe the effect of off-nadir pointing at the sun on SXM data, as it was possible that the SXM would not always be pointing at the sun during all science activities.

Due to the closer proximity of the spacecraft to Earth, as well as the relative position of OSIRIS-REx's solar panels, REXIS was operating at slightly higher temperatures than nominal; some yellow level temperature alarms were triggered (alarm levels are detailed more thoroughly in Section 4.3.2). These increased temperatures resulted in higher energy rates due to increased dark current, with the event rates reaching the amount specified by the event rate filter (200cps), shown in Figure 3-4. As discussed in Section 2.1.1, when the event rates surpass the event rate filter, the software begins randomly discarding events until the total events in the CXEL are around the set filter limit. If events remain above the filter limit for an extended period of time, potentially good data are effectively lost. However, even with the higher event rates, it was noted that there were no significant changes in CCD gain or energy resolution from prior calibration activities. The higher temperatures during this activity were expected, and present in the REXIS thermal prediction documents.

Between October 9 and 12, 2017, the same calibration sequence was repeated after the OSIRIS-REx EGA. On the second day of this operation, due to temperature increase, event rates on the CCDs reached significantly above the event rate filter, shown in Figure 3-5. The resulting downlinked data had a much higher volume than anticipated and contained multiples of some events. The noise from the additional events was too high to obtain good spectral resolution on the CCDs, which rendered the data essentially useless. Moving forward from this operation, the team continued to investigate why more events than commanded by the event rate filter in the flight software were being downlinked. If the problem persisted, and appeared during Orbital B data collection, the resulting data would be insufficient to meet the science goals of REXIS. The increased data volume also had the potential to surpass the REXIS data volume allotment on the OSIRIS-REx spacecraft. The high event rate and event rate filter breach appeared again during the first Crab Calibration (Section




Figure 3-4: Figure showing total downlinked events per second on August 15, 2017. The red line indicates the binary state of the filter, with the higher value indicating filtering is occurring, and lower value indicating no filtering is occurring.

3.4.2), and a mitigation method was developed, tested on the ground, and proven in flight during L+30 (Section 3.4.3).

3.2 Internal Calibrations

3.2.1 L+18 Checkout and Calibration

L+18 occurred on March 13, 2018. REXIS again performed the same checkout activities as in prior operations, and performed further CCD calibration activities. The focus was still on attempting to optimize the ET, ST, and ULD for all the nodes. The same ten nodes that had been used throughout flight were again used for this operation. Again, the goal of the checkout was to track the long-term trends in the instrument, and the goal of the calibration was to further investigate optimal parameter settings for each node for maximum possible science output upon reaching Bennu. Functionality of REXIS was confirmed via the checkout, and further data necessary for node image processing parameter optimization was obtained via the calibration

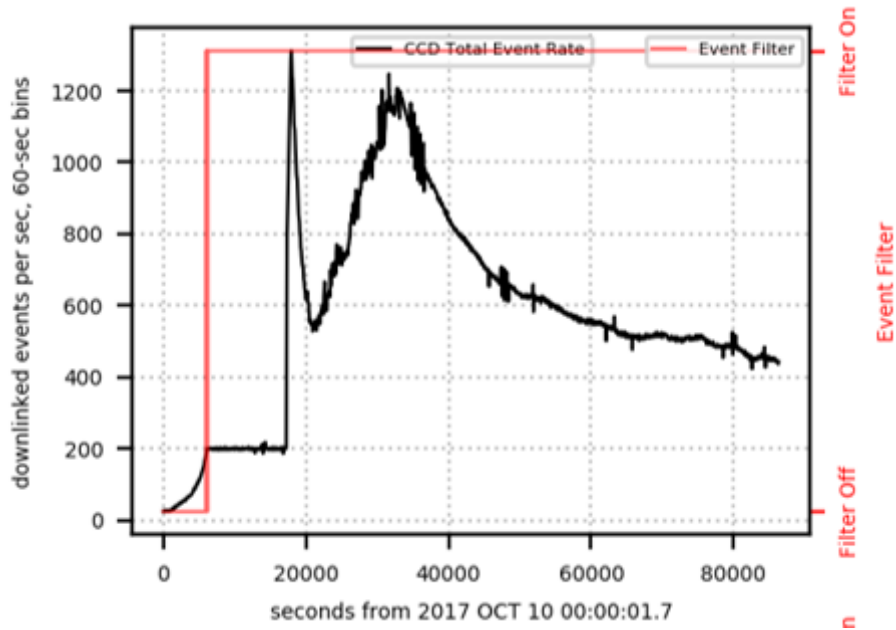


Figure 3-5: Figure showing total downlinked events per second on October 10, 2017. The red line indicates the binary state of the filter, with the higher value indicating filtering is occurring, and lower value indicating no filtering is occurring. The limit set by the event grade filter should be 200 cps, but the black line (downlinked events) reaches far above that limit.

activities. No new anomalies were observed.

3.2.2 REXIS Internal Calibration

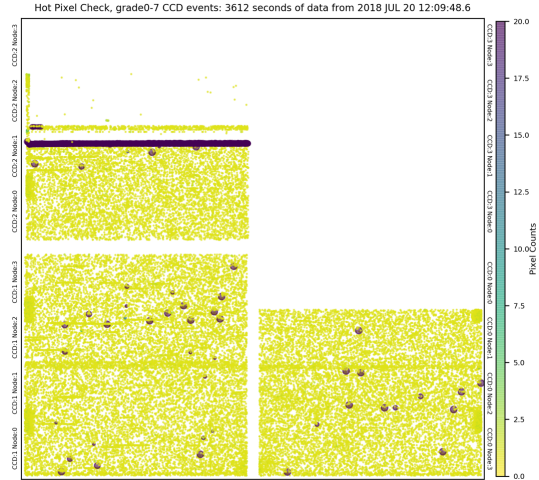
The REXIS Internal Calibration operation occurred between July 12 and 15, 2018. The main goals of the internal calibration were to characterize the internal background signal due to instrument noise, and to calibrate the CCDs using the on-board perimeter and cover Fe^{55} sources. The calibration activities consisted of setting the CCD parameters to the nominal values, as determined from the prior calibration activities, and setting the CCDs to downlink data for 80 consecutive hours. The ten nominal nodes were used for this operation. The calibration provided the team with a measurement of the internal background. This background was subtracted from the resulting data from REXIS operations after the radiation cover was opened, in order to isolate actual observed signal from instrument noise. No new anomalies were observed during the internal calibration activities.

3.2.3 L+22 Checkout and Calibration

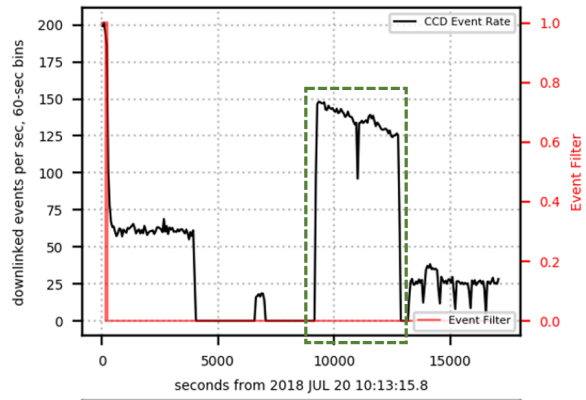
The L+22 Checkout and Calibration activities occurred on July 20 and 21, 2018, immediately following the REXIS Internal Calibration activities. The checkout was again performed to confirm REXIS functionality, and the L+22 calibration activities were performed to complete the optimization of the CCD parameter settings (ET, ST, and ULD) and mask the identified hot pixels using the hot pixel mask functionality for the first time in flight.

First, the checkout was performed at the beginning of the science operation. Upon completion of the checkout sequence, the L+22 calibration was performed. This calibration involved the optimization of CCD parameters by downlinking event list data for a designated period of time at different ULD, ST, and ET settings. Near the end of the L+22 calibration, a hot pixel mask was uploaded to mask known hot pixels as well as the areas of the detectors closest to the Fe^{55} sources, in order to verify correct physical correlation of the placement of pixels by the software. After an hour's worth of data were downlinked with the mask applied, the mask was then removed, and more event list data were taken with differing CCD parameters. All ten nominal CCD nodes were used for both of these activities.

The checkout returned nominal results, indicating good functionality of REXIS. During the L+22 calibration, some odd behavior was observed upon application of the hot pixel mask. As shown in Figure 3-6, once the hot pixels were masked by the software, random noise appeared in the event list data. There was not an identifiable pattern in this noise, nor did the sudden increase in noise seem physically feasible. Once the hot pixel mask was removed, the data returned to normal. Upon discovery of the anomalous behavior, the REXIS team began an investigation. From the investigation, it was determined that the noise appeared when a mask was applied across the boundary between some adjacent nodes. When a buffer was inserted near adjacent node boundaries, the noise no longer appeared. For future operations, the hot pixel masks applied were first checked to determine that a sufficient border between adjacent node boundaries was present.



(a) Detector plane showing increased grades 0 through 7 events while hot pixel mask was applied.



(b) Total event rates for 07-20-2018. Area outlined by dotted line indicates when hot pixel mask was applied.

Figure 3-6: Effects of hot pixel mask application on flight data

3.3 Cover Opening

In September 2018, REXIS performed its radiation cover opening operation, which was crucial for the remainder of the REXIS science mission. If the radiation cover did not open, the detectors would not be able to collect any X-ray data from the asteroid. Therefore, the REXIS team planned for three separate attempts to open the cover - September 14, 18, and 22. The radiation cover was held closed by a frangibolt, and was opened by applying voltage to a heater to warm the Frangibolt. The Frangibolt utilized a shape memory alloy (SMA) component, which elongated when heated, fracturing the bolt and releasing the cover. The cover opened on the first attempt. REXIS downlinked event list data for a half hour prior to and half hour after the cover opened, and a full raw frame was collected at the end of the operation. All ten nominal nodes were used for the cover opening activity.

Various data products were used to determine the state of the cover after firing the Frangibolt. The first data products that suggested the cover opened were the event rates and the event spectra, shown in Figure 3-7. The event rates, shown on the left in 3-7a, are consistently higher after the Frangibolt is fired than from before, which indicates that the detectors were seeing more signal. This increase in signal could

only have come from external sources, as the only signal from before came from the cover and side Fe^{55} sources and the instrument background noise, both predictably steady sources. The figure on the right, Figure 3-7b, shows the energy spectra from before and after the firing of the Frangibolt. The data from after the Frangibolt was fired not only show more total recorded signal, they also follow the dotted line on the chart, which shows the expected signal from the cosmic X-ray background (CXB). Signal from the CXB would only be visible to the REXIS CCDs if the cover was opened.

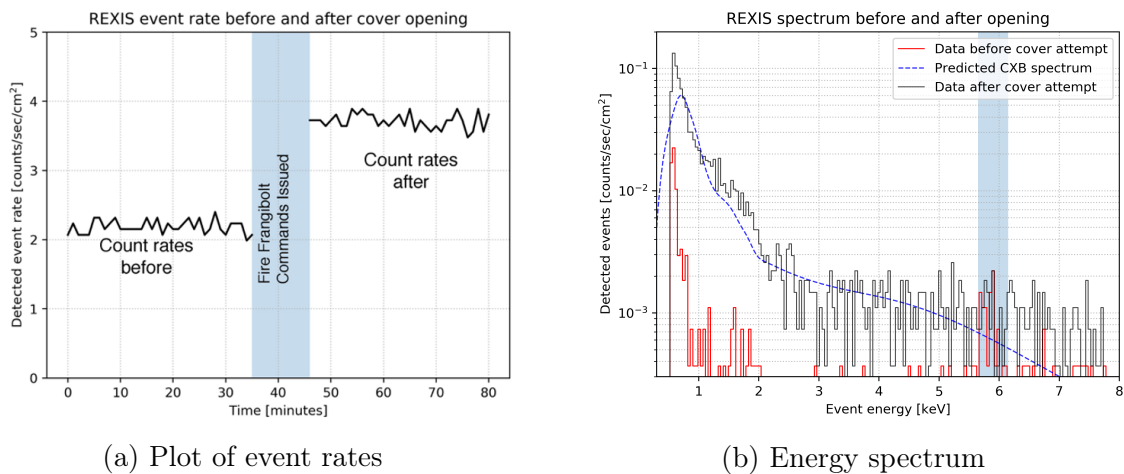


Figure 3-7: Data products to determine the radiation cover opened

Next, the team determined whether or not a gradient of events was present on the detector surface, to determine the angle at which the cover had opened. Analysis was performed to predict the gradient of events across the detector for varying angles of the radiation cover ranging from 13.75 degrees to the full 96.25 degrees. Figure 3-8 shows results of this analysis, along with the flight data from cover opening. Due to the detector layout, CCD 0 node 2 is one of the first nodes that observes the sky when the radiation cover is open, and so would be the first to register events, and would register more events than more inward nodes in the case that the cover did not open all the way. CCD 2 node 0 is one of the further nodes from the opening of the cover, and so would see fewer events if the cover was not opened all the way. The plot in Figure 3-8 shows the ratio of events on a portion of CCD2 node 0 to events on

a set area of equal size of CCD 0 node 2. The ratio of events on differing portions of CCD 0 node 2 and CCD 2 node 2 are roughly 1 or higher, meaning that there is not a noticeable decrease in events as the detector surface gets further from the opening of the radiation cover. Therefore, it was determined that the radiation cover opened fully after the first frangibolt firing attempt.

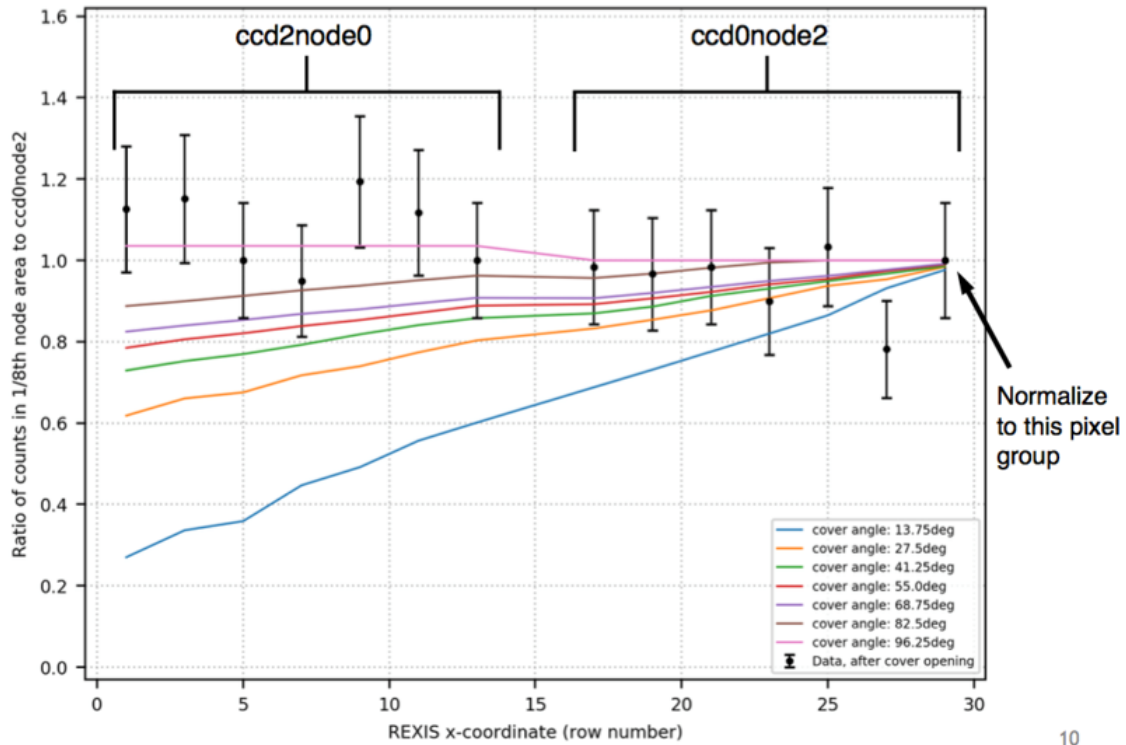


Figure 3-8: Plot of events over detector surface, and predictions of event gradients if cover only opened partially

Further analysis of the data showed an unexpected increase in events during a period when REXIS was not downlinking events, so the events only appeared in the housekeeping data. The top plot in Figure 3-9 shows the increase in events as observed in the REXIS housekeeping data. The events remain low following the cover opening (the small gap in the data around 400 seconds) up until 4000 seconds, where there is a large increase in events on all three functioning CCDs. A correlation was found between the event rates and the angle of the sun in relation to the REXIS boresite. The zenith angle of the sun (from the REXIS z-axis) is shown in the bottom plot in Figure 3-9, along with the corresponding total CCD event rates. The cover opening

operation occurred when the sun was at a zenith angle of 150 degrees with relation to the REXIS boresite. Upon the completion of the cover opening operation, the spacecraft slewed towards the sun. The total event rates begin increasing when the zenith angle between the REXIS z-axis and the sun reaches 95 degrees, spike between 93 and 92 degrees, and remain higher afterwards as the angle remains at 90 degrees. Due to this correlation between zenith angle relative to the sun and total event rates, the most likely hypothesis for the cause of the high event rates was detector sensitivity to spacecraft pointing, and the light environment at the spacecraft attitude.

3.4 External Calibrations

Once the radiation cover was opened, REXIS performed a series of calibrations using external sources rather than the iron source in the cover, as the cover source was now out of the field of view of the detectors. To be prepared for Orbital B science, the following calibrations (shown in the flowchart in Figure 3-2) were necessary:

- Measurement of the cosmic X-ray background (CXB Calibration)
- Calibration of the detectors using known external X-ray source, the Crab Nebula (Crab Calibration parts 1 and 2)
- First test of CCD parameters determined from internal calibrations in asteroid-like environment (L+30 calibration)
- Test of OBF, and identification of hot pixels (OBF Verification)
- Test of mask configuration to determine if thermal warping had occurred (Mask Calibration)

3.4.1 CXB Calibration

Between October 9 and 13 of 2018, REXIS performed the cosmic X-ray background (CXB) calibration. The goal of this calibration was to obtain a measurement of the CXB, and determine if the internal background of the CCDs had changed from internal calibrations. The CXB is the diffuse X-ray background present throughout the universe. The precise source of this background is yet undetermined, but its

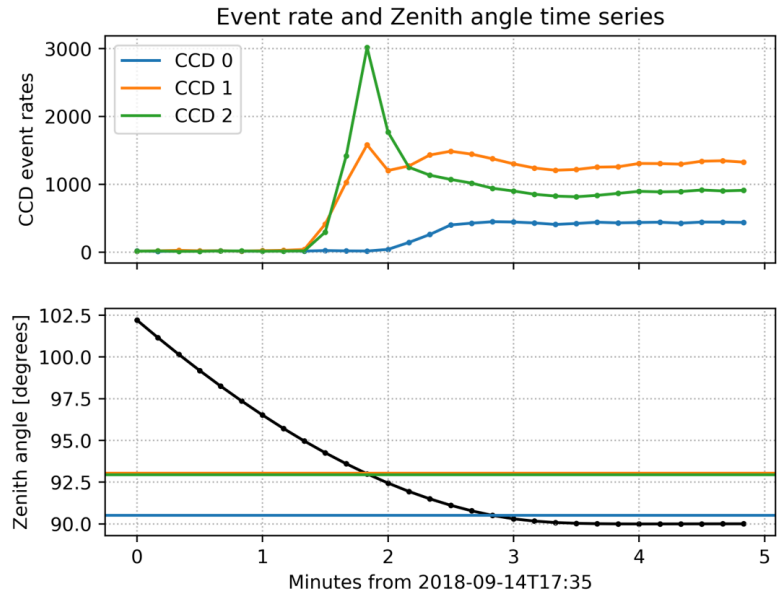
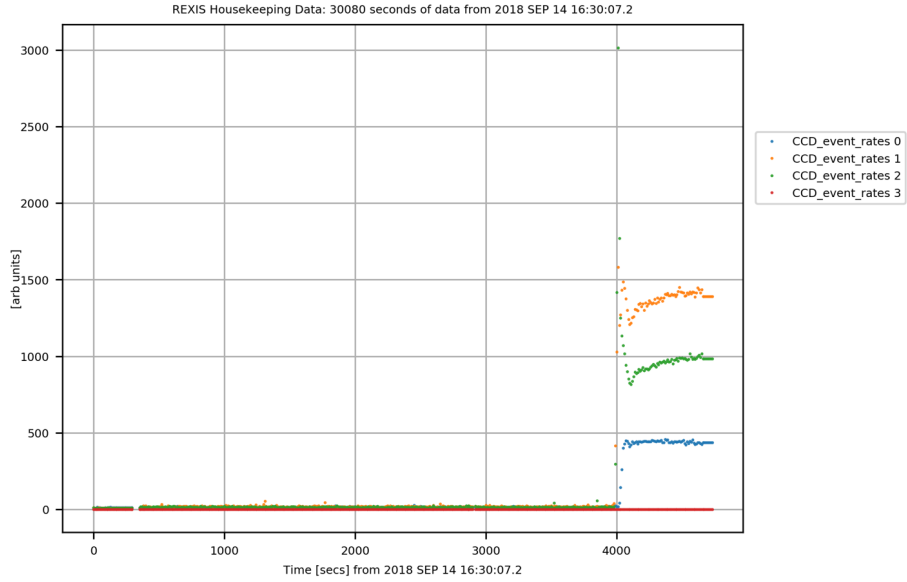


Figure 3-9: (Top) plot showing event rates throughout cover opening activity. (Bottom) Plot showing comparison between CCD event rates and the zenith angle with relation to the sun.

presence has been recorded since its discovery in a 1962 experiment [35]. The CXB is non-uniform at low energies, so two different areas of the sky with different CXB flux levels were chosen as the observation targets. The first target, named target zero for sequence purposes, was the North Polar Spur (Right Ascension 254.6 degrees, Declination 10.9 degrees). Target one was an area of the sky expected to have diffuse

cosmic background (Right Ascension 20 degrees, Declination 0 degrees). Event list data were recorded for six hours at target zero and three hours at target one, based on the different expected flux levels. The same nominal ten nodes were used for the CXB calibration. This observation resulted in a characterization of the REXIS internal background using external measurements.

During this observation, anomalous behavior in the measured energy of the Fe⁵⁵ calibration sources was observed, as shown in Figure 3-10. The densest part of the plot should appear between around 270 and 280 ADU (around 5.9 keV), corresponding with the known iron-L energy band for Fe⁵⁵. In all operations prior to cover opening, the detector response to the iron sources was as expected. However, during the CXB calibration, the plot shows a drift downward in that measurement, from the nominal 270 down to around 250-260 ADU. This drift downward indicates a change in either the energy gain or offset of the node. Figure 3-11 shows the energy histograms from CCD02 broken down into grades 0 and grades 2 through 5 from before and during the CXB calibration. There is a horizontal offset in the ADU where the peaks of the distributions occur, but no vertical change in the amplitudes of the peaks. The absence of a vertical change indicates that there was no change in the gain of the nodes, and instead there was a change in the offset. A potential cause for the energy offset was theorized to be stray light shining on the serial register of the CCDs or on the detector surface through pinholes.

Given the hypothesis that the high event rates observed during cover opening could be due to the spacecraft pointing, and the hypothesis that that the energy offset observed in the CXB calibration data could be due to light environment as well, measures were taken to attempt to mitigate the effects of the light environment on the data. Specifically, more frequent regeneration of the bias map while downlinking event list data increases the likelihood that the light environment present in the bias map is representative of the light environment in which the event list data are taken. The background caused by stray light is then subtracted from the data before further image processing. For future operations, a more frequent regeneration of bias maps was implemented. However, due to the eight week planning cycle of OSIRIS-REx,

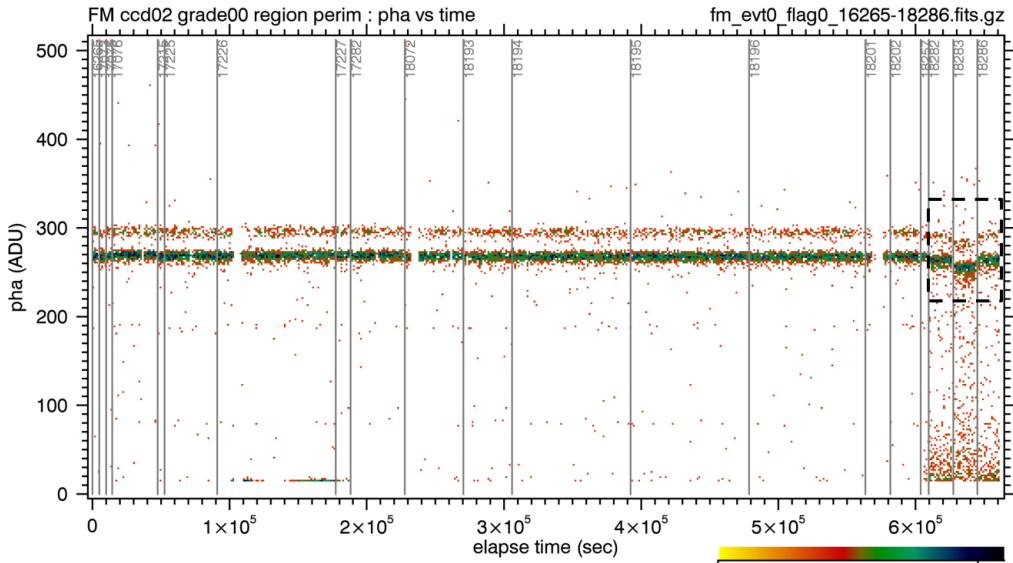


Figure 3-10: Plot of energy (in ADU) versus time for CCD02 during all of flight. CXB Calibration data are outlined by the dotted line.

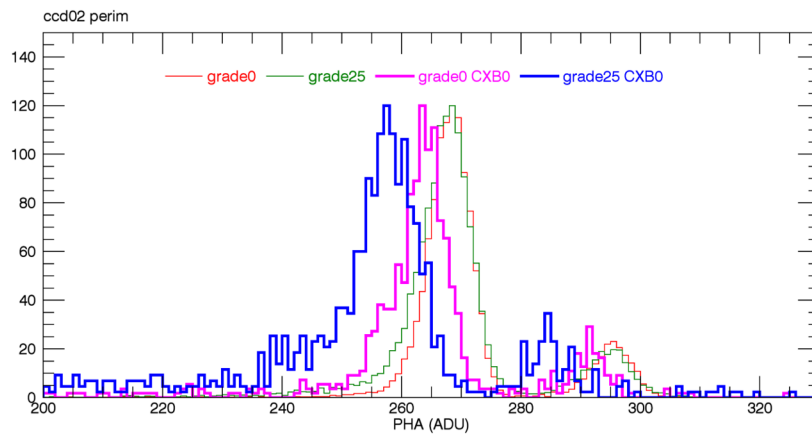


Figure 3-11: Comparison of energy histogram for CCD02 from before and during CXB calibration.

the first time to test this mitigation method was not until L+30.

3.4.2 Crab Calibration: Part 1

The goal of the Crab calibration, performed between November 14 and 26 of 2018, was to make several observations of the Crab Nebula, a known X-ray source, in order to measure the quantum efficiency of the CCDs. Through measurements of a source with known X-ray flux, the detector response was evaluated to determine the gain and offset of each node. REXIS was pointed at eight separate targets in the vicinity

of the Crab Nebula: target 0 was centered on the nebula, and the other seven targets created a 3.5 degree ring around the nebula. The ten nominal nodes were used for this observation.

However, the data from this observation was rendered unusable due to an event filter breach similar to the breach experienced during the L+10 calibration (Section 3.1.3). Figure 3-12 shows the event rates observed during one of the days of the Crab calibration, and the state of the event rate filter. The filter was activated during the entire data collection period of the Crab calibration operation, meaning that the total amount of events registered by the CCDs was over 200 counts per second for the entire observation period. Between zero and 5000 seconds, and between 13000 and 15000 seconds, the events downlinked are well above the software filter limit of 200 counts per second. The events exceeding the software limit indicates the filter is breached, and the resulting data are unusable.

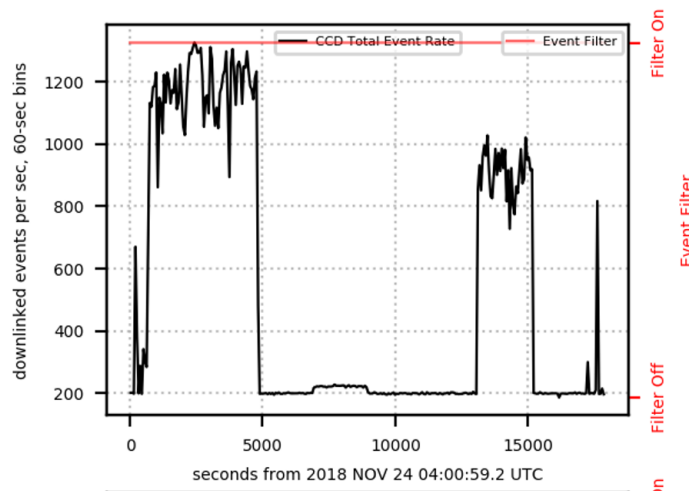


Figure 3-12: Plot of event rates from the science window on November 24, 2018, part of the Crab calibration. The filter is activated for the entire period, and portions of the events exceed the software limit of 200 counts per second, indicating a filter breach similar to L+10.

Figure 3-13 shows the long-term event rates from flight, including downlinked events (red) and housekeeping events that were recorded while REXIS was not downlinking data (blue). The Crab calibration event rates are noticeably higher than previous operations, and are similar to the brief spike in events seen at the end of

the cover opening operation. Again, sensitivity to scattered light was suspected as the most likely culprit. Due to the loss of data, the REXIS team requested and was granted a re-fly of the Crab calibration to be performed in March of 2019. The results of that operation are detailed in Section 3.4.4.

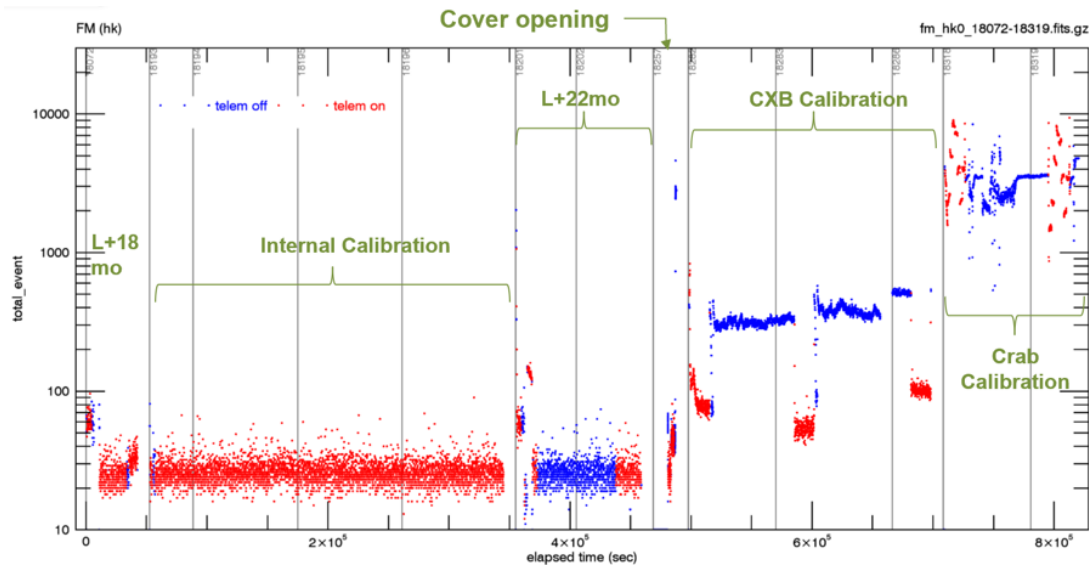




Figure 3-13: Plot of total downlinked and housekeeping event rates from September 2016 through Crab calibration in November 2018

Another anomalous CCD behavior was first observed during the CXB calibration and was again observed during the first Crab calibration operation. For portions of the observations, and in the bias map downlinked during the Crab calibration (the image on the right in Figure 3-14b), some nodes showed an odd triangular noise pattern (depicted on the left in Figure 3-14a). The noise pattern was not physical (i.e. could not have been due to actual X-ray photons due to its constant physical shape) and was present in the data from every day of the Crab calibration operation. Further investigation into this issue occurred from December 2018 through March 2019. A collaboration with an MIT Lincoln Laboratory researcher, Dr. Kevin Ryu, resulted in the hypothesis that the noise was due to cross-talk between the operational CCD1 and the non-functioning CCD3.



(a) Noise on CCD11 during CXB calibration



(b) Bias map from first Crab calibration showing evidence of cross-talk

Figure 3-14: Evidence of triangular cross-talk pattern on the CCDs from both the CXB and Crab calibrations

3.4.3 L+30 Calibration

L+30 calibrations occurred between February 16 and 20, 2018. This operation was the first time Bennu was in the field of view of REXIS. The goals of this observation were to test a hot pixel mask to ignore all previously identified hot pixels from the post-cover opening operations, while ensuring the increased random noise would not occur, and to test the feasibility of using the event grade filter (EGF) to reduce the total event rates to below the event rate filter. The success of the L+30 calibration operation was imperative to show that REXIS could downlink usable data with the use of some mitigation methods to avoid an event rate filter breach. If the event rates continued to be well above the filter rate limit to the point where a breach occurred, the REXIS data volume could possibly exceed its allocation, or delay the downlink of other instruments' data that was more directly necessary for the success of the OSIRIS-REx science mission. Additionally, due to the fact that REXIS was not necessary for OSIRIS-REx mission success, if it was shown that REXIS was not capable of downlinking usable data, it may have been decided to no longer operate the instrument in favor of freeing up operational resources for other instruments.

Due to the hypothesis that REXIS was sensitive to spacecraft pointing, as ob-

served during cover opening (Section 3.3) and the Crab calibration (Section 3.4.2), the REXIS team decided to regenerate the bias map more frequently to more accurately represent the light environment from which the event list data were to be collected. Another theorized effect of spacecraft pointing was the drastic increase in events such that the event rate filter was breached. To reduce the chances of the detectors recording enough events to trigger an event rate filter breach, the event grade filter (EGF) was used for the first time in flight to ignore grades 6 and 7 events (grades 6 and 7 events were highly unlikely to be actual X-ray events of interest, and so could be ignored without losing crucial science data). The node enable map (NEM) was adjusted to only use the six best-performing nodes: CCD02, CCD10 and 12, and CCD20, 21, and 22 (marked in green on Figure 3-15). The number of nodes used for downlinking data was reduced in order to decrease the total number of events downlinked by the CCDs in an attempt to avoid reaching the event rate filter. These specific six nodes were chosen because they had the best resolution. The entire L+30 operation sequence was tested on ground hardware at MIT using the thermal vacuum chamber and detector assembly to ensure all commands resulted in the expected behavior.

CCD23	CCD 2	CCD 3	CCD33
CCD22			CCD32
CCD21			CCD31
CCD20			CCD30
CCD13	CCD 1	CCD 0	CCD00
CCD12			CCD01
CCD11		CCD02	
CCD10			CCD03

Figure 3-15: CCD physical layout. The nodes with the best resolution that were used for L+30 operations are colored in green.

The L+30 observation was split into two separate observations - one pointed nadir

at the asteroid, and one during the high gain antenna (HGA) pass, while OSIRIS-REx downlinked data to the ground station through the use of the Deep Space Network (DSN). The nadir portion tested the nominal CCD parameter set in Orbital B-like conditions for the first time, and the portion during the HGA pass was to get another measurement of the CXB near the asteroid. To test the hot pixel mask and the EGF, as well as collect data with the nominal CCD parameters in Orbital B-like conditions, the nadir-pointed portion of REXIS L+30 observation was designed as follows:

- Two different hot pixel masks were applied and removed consecutively, and event list data were downlinked for ten minutes per mask. (The first mask was a repeat of the L+22 mask that caused high event rates, and the second mask consisted of identified hot pixels from prior external calibration operations, that followed the design rules discovered from ground testing).
- Once the hot pixel mask test was completed, the EGF was changed to only downlink grades 0 through 5 events.
- The initial bias map generated when REXIS first entered science mode was downlinked.
- Six different node enable maps were cycled through three times (the nodes enabled are detailed in Table 3.1). Data were collected for a total of 30 minutes with each node enable map for sets one through four, and a total of 15 minutes each for sets five and six.
- Partial raw frames were downlinked after every three node enable map change.
- A partial bias map was downlinked after every six node enable map change.
- At the end of the nadir-pointed observation, REXIS remained on and in science mode with telemetry off until the HGA pass observation.

The HGA pass operation was similarly designed, but times were shortened to fit within the smaller science window. It occurred four days after the nadir-pointed observation. At the beginning of the operation, the bias map was regenerated to better reflect the current background, and the EGF was reset to accept all grades, 0 through 7. The hot pixel mask test was repeated, but data were only collected for five minutes with each mask. Then the full bias map was downlinked, and the

Table 3.1: The six node enable maps cycled through during the L+30 operation

Set	Node(s) enabled
1	CCD02
2	CCD22
3	CCD20 and CCD22
4	CCD21 and CCD22
5	CCD02, CCD20, CCD21, and CCD22
6	CCD02, CCD10, CCD12, CCD20, CCD21, and CCD22

EGF was changed back to only allowing grades 0 through 5. The six different NEMs were cycled through, but only twice. Partial one-node raw frames were downlinked, and the bias map was only regenerated and nine nodes of it downlinked once within the sequence. Once the sequence was complete, REXIS went into safe mode and was powered off at the end of the HGA pass.

The hot pixel mask test was successful; each mask performed as expected, with the "bad" hot pixel mask inducing high event rates, and the "good" mask causing no increase in events. The EGF test was also successful. Event rates came nowhere near the 200 counts per second limit, and were lower overall, as shown in Figure 3-16. Figure 3-17 shows that grades 6 and 7 were no longer downlinked in the second half of the operation, as commanded. Additionally, regenerating the bias map more frequently had the desired effect; no energy offset shift was observed in the L+30 data. The successful collection of more event list data and some full bias maps helped identify more hot pixels to mask for future operations.

While the telemetry was turned off and REXIS was not downlinking event rates, the housekeeping data were still keeping track of the housekeeping event rates. The housekeeping event rates after the operation was complete were much higher than the downlinked event rates during the operation, showing that use of the EGF and decrease in total number of nodes used during the operation helped keep downlinked event rates below the filter limit. The housekeeping rates also showed fluctuations over time. In an attempt to explain the fluctuations, the REXIS team compared the predicted light curve to the overall event rates, show in Figure 3-18. There is good




Figure 3-16: Total event rates (downlinked and housekeeping) from internal calibration through L+30

REXIS CCD events: 11176 seconds of data from 2019 FEB 16 01:07:13.6 UTC




Figure 3-17: Events for CCD 0 2, with grade differentiation. The dotted line indicates grades 6 and 7 events. No grades 6 and 7 events are recorded after EGF is changed.

agreement between CCD 0 event rates and the predicted curve (indicated by the gray line near the bottom of the plot), and some correlation between CCD 2 event rates and light curve.

3.4.4 Crab Calibration: Part 2

Due to the event rate filter breach during the first Crab calibration attempt (3.4.2), the data was insufficient to achieve the calibration goals. A second attempt of the Crab calibration occurred between March 16 and 24 in 2019. The second operation had




Figure 3-18: L+30 housekeeping event rates per CCD compared to the predicted light curve (gray line near bottom of plot). Plot created by HCO researcher David Guevel.

the same goals as the first attempt, and followed a similar pattern to the observation sequence; eight different pointing targets in the vicinity of the Crab Nebula were observed for a half hour each per day. In total, four separate days of observations occurred for a total of 16 hours of downlinked event list data. The main differences between the first and second Crab calibration attempt were that the second calibration included a bias map regeneration once the spacecraft arrived at each new pointing target, the node enable map was changed to only include five of the best nodes (CCD02, CCD12, CCD20, CCD21, and CCD22), and the EGF was set to ignore grades six and seven. The purpose of these changes was to mitigate the effects of the instrument's sensitivity to the background observed at a particular spacecraft pointing, and to avoid downlinking high event rates. At the end of each day of operation, a full raw frame and the last bias map were downlinked. Only five nodes were used, instead of the six used for L+30, out of an abundance of caution to avoid reaching the event rate filter limit.

Attempt two of the Crab Nebula calibration was successful. No event rate filter breach occurred, and the boresite of REXIS was successfully refined in relation to

spacecraft pointing. REXIS collected spectra from each target pointing, which were used to measure the quantum efficiency (QE) of each node. In the calculation of the efficiency, however, there was a noted deficit of energies above 1 keV. In Figure 3-19, the absolute response function (ARF) for CCD22 from before the second Crab calibration is marked by the blue line, on the top. The adjusted ARF from the results of the second Crab calibration is shown in orange, and the data from the second Crab calibration are shown in green on top of the updated ARF. The brackets indicate one of the larger discrepancies in the ARFs; there is a larger than expected cutoff in the detector response around 75 ADU. The ARF from after the second Crab Cal is lower than the ARF from before, especially between 75 and 200 ADU. It was hypothesized that the difference between the two functions could be due to the lower than expected ST settings - the ST values commanded to the spacecraft were interpreted as lower energies by the software or detectors due to incomplete knowledge of each node's energy resolution. This unexpected discrepancy necessitated further in-flight testing of higher ST values; however, due to the eight week planning cycle, the next available operation in which parameters could be changed was not until the Mask calibration right before Orbital B.

Another anomaly was noticed in the spectra collected from the Crab Nebula calibration. There was an unexpected and unexplained excess in energies below 1 keV. The box drawn around the data in Figure 3-19 indicates the excess low energy down-linked during the operation. Rather than following the shape of the ARF and decreasing to around 100 counts as the ADU decreased from 50 to 0, the counts below ADU instead increase up to around 1000. Upon the discovery of this low energy excess, the team looked at past data and found a low energy excess in internal calibration data as well, and in ground testing data. Some theories as to the cause of the excess include an interaction occurring between incoming particles and the OBF layer, a flight software error, or cross-talk from the non-functioning CCD3. This low energy excess was tracked throughout the rest of flight, and is still under investigation. The low energy excess is not within the scope of this thesis.




Figure 3-19: Recorded spectrum from Crab Cal 2 (green line) compared to the absolute response function from before and after Crab Cal 2. The box indicates the excess low energy, and the brackets display the discrepancy in the energy cutoff around 75 ADU. Plot created by HCO researcher Daniel Hoak.

3.4.5 OBF Verification

The OBF verification took place between April 18 and 20 of 2019. The goal of this operation was to take data with the asteroid in the field of view to confirm that the hot pixel map worked as expected and captured most hot pixels in the asteroid light environment, and to identify any further pixels that required masking; this operation was the last opportunity to identify more hot pixels and add them to the final mask before Orbital B operations. The OBF verification operation used the six best resolution nodes and consisted of the following sequence:

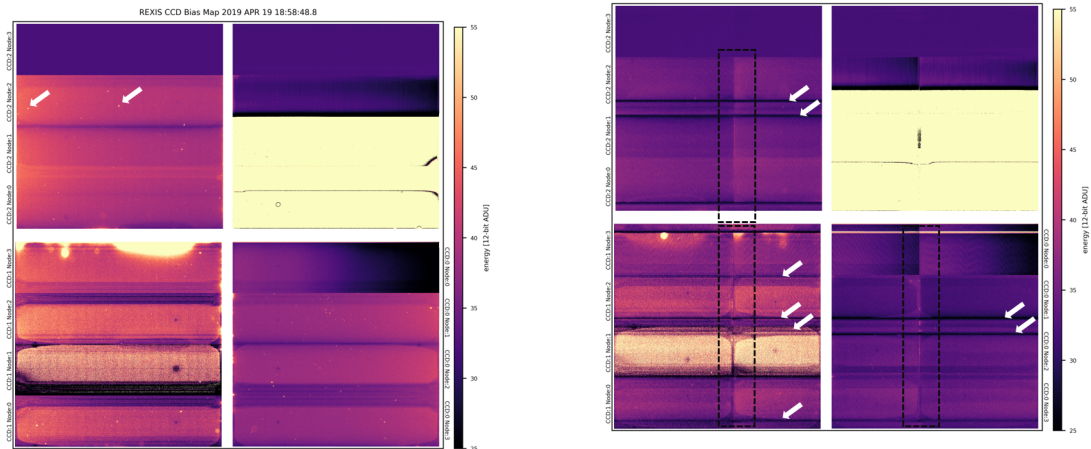
- A half hour of data was downlinked using nominal settings, with no mask applied.
- One full raw frame and one full bias map were downlinked. The bias map was then regenerated.
- The most up-to-date hot pixel mask was applied, and an hour of data was downlinked.
- One full raw frame and one full bias map were downlinked.

In order to test the detector response at lower energies, the ET for some of the better performing nodes (CCD02, CCD20, and CCD22) were lowered from the nominal setting for a portion of the data collection. The six best nodes were used for this operation and the EGF was set to ignore grades 6 and 7 events for the entirety of the operation.

The OBF verification was successful in testing the hot pixel mask - no anomalous behavior was observed when the mask was applied. The testing of lower ET settings was also successful. Periods with lower ET settings showed only minor increases in total event rates, which indicated that lower ET settings could be used during Orbital B to increase the range of observable energies without concern of breaking the event rate filter.

One anomaly was observed during the OBF verification. The first full downlinked bias map appeared as expected (shown in Figure 3-20a), but the second bias map contained some artifacts that did not appear physical (shown in Figure 3-20b). In

the figure on the right, the map shows non-physical horizontal and vertical distortions. The bias map on the left also has evidence of hot pixels, as is expected; however the corrupted map on the right shows no signs of hot pixels, a further clue that the map is not a physical representation of the detector plane.

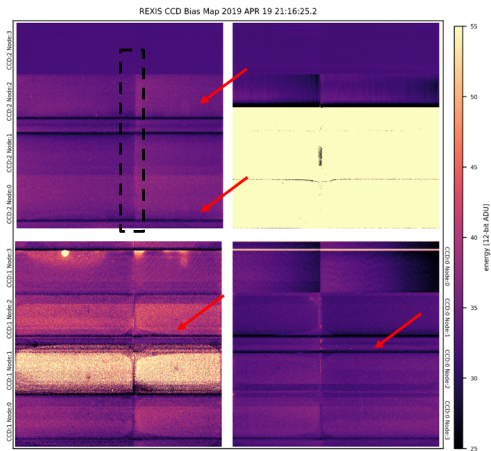


(a) Nominal bias map from OBF verification. Some hot pixels are identified with an arrow.


(b) Corrupted bias map from OBF verification. The vertical artifacts are outline by the dotted lines, and the horizontal artifacts are identified by arrows.

Figure 3-20: Bias maps downlinked from OBF verification. The second map has vertical and horizontal artifacts, and does not contain any hot pixels.

Initially, it was hypothesized that the corrupted image was due to an error during downlink. However, upon comparing the corrupted bias map to the event list data that followed, similarities in the physical layout of the downlinked event distribution and the bias map indicated that the event list data had been processed using the corrupted bias map. Figure 3-21 shows a side by side comparison of the bad bias map and the event list data, and indicates areas where the bias map artifacts appear in the event list data. The most likely cause of a bad bias map was hypothesized to be an on board software issue with the framegrabber. The exact mechanism by which the error occurs is yet undetermined, but a working theory is that something within the software times out, and the framegrabber loses its place on the detector plane. This bias map issue was tracked throughout Orbital B by downlinking a single node bias map for CCD22 for every time a bias map was regenerated.



(a) Bad bias map from OBF verification



(b) Event list data from period with bad bias map

Figure 3-21: Comparison of bias map with subsequent event list data. The dotted outline shows an area of vertical distortion, and the arrows point out some of the horizontal artifacts.

3.4.6 Mask Calibration

The mask calibration was designed to occur before Orbital B, but after the spacecraft performed observations of Bennu that resulted in the sun shining directly on the REXIS coded aperture mask. A concern was that the direct and prolonged exposure to the sun could cause physical distortion of the mask. The purpose of the mask calibration was to observe a bright, known X-ray source to create a projected image of the mask on the detector plane, and from that determine if the mask became thermally warped between cover opening and Orbital B. If the mask was warped, the physical distortion would have to be removed during ground data processing. The known X-ray source that was observed was Scorpius X-1 (Sco X-1).

The mask calibration operation utilized the six best nodes, the same ones as used in L+30. Due to the reduced number of nodes (down from the original ten to only six), multiple pointing targets were used to optimize how much of the mask was observed. Five total pointing targets in the vicinity of Sco X-1 were chosen. The sequence below was repeated at each pointing target:

- The spacecraft slewed to the RA and Dec coordinates.
- The bias map was regenerated.

- CCD parameters were set to nominal, and 27 minutes of event list data were downlinked.
- The ST values for CCD02 and CCD22 were doubled, and 27 additional minutes of event list data were downlinked.

The bias map was regenerated once the spacecraft finished its slew to each pointing target, to ensure the bias map had a similar background to the subsequent data. The last bias map was fully downlinked at the end of the observation. The higher ST settings were tested for nodes CCD02 and CCD22 during this observation to recover some of the signal above 1 keV that showed a deficit during Crab Cal (Section 3.4.4).

The mask calibration was successful. The observation produced clear images of different parts of the mask pattern on the detector plane. One of the five images produced is shown in Figure 3-22. In all five images, no obvious physical abnormalities are present. The bias map downlinked from this operation was nominal and no physical artifacts appeared in any of the other targets' event list data, which indicates that the odd bias map behavior did not occur during the mask calibration.




Figure 3-22: Mask image created from one pointing target observation of Sco X-1

The increase of the ST also produced the desired results of increased detector efficiency, and an expected increase in grade 0 events. Figure 3-23 shows the event rates broken down by grade for the two nodes that had their ST doubled for a portion of the operation. Every time ST was commanded higher, the grade 0 events increased,

as expected. The detector efficiency was also improved by the doubled ST, as shown in the plot in Figure 3-24. The plot shows the ratio between the high ST and low ST measured detector efficiency across the observed energy range for CCD22. The higher ST produced increased efficiency across the entire energy range, with the biggest improvement occurring around just under 2 keV, which corresponds with the silicon edge.




Figure 3-23: Event rates broken down by grade for CCD02 and CCD22. Arrows indicate when ST was doubled. Grade 0 events increase every time ST is doubled.




Figure 3-24: Ratio of CCD22 detector efficiency with doubled ST to detector efficiency with nominal ST during the mask calibration operation. The ratio of the two efficiencies is plotted from 0 to 6 keV. Plot created by HCO researcher David Guevel.

3.5 Asteroid Science

The noteworthy anomalies observed throughout flight before Orbital B are indicated by stars in the timeline figure (Figure 3-25, and outlined in Table 3.2). As anomalies occurred, and new information about the detectors was learned, the image processing parameters used evolved. Tables A.3, A.4, and A.5 (in Appendix A) show the evolution of the CCD image processing parameters throughout flight, beginning from the Internal Calibration operation, and ending with the parameters used for Orbital B and Orbital R.

Table 3.2: In-flight anomalies observed between Launch and Mask Calibration

Operation	Anomaly observed
L+10 (after OSIRIS-REx EGA)	Event rate filter breach
L+22 Calibration	Bad hot pixel mask-induced noise
Cover Opening	High event rates correlated with spacecraft pointing
CXB Calibration	Energy offset in recorded spectra
Crab Calibration part 1	Event rate filter breach, loss of usable data
OBV Verification	Bad bias map downlinked

In the original mission plan, Orbital B was intended to be the main science data collection period for REXIS. However, during Orbital B, the instrument experienced multiple anomalous behaviors that resulted in the loss of usable science data. The REXIS team petitioned for and was granted additional observation time, during Orbital R.

3.5.1 Orbital B

Orbital B occurred from July 1 to August 6, 2019. This phase was the main science observation for REXIS; all prior calibrations and operations were to optimize parameters for maximum Orbital B science. OSIRIS-REx was in a frozen orbit around Bennu with a radius of one kilometer, measured from the center of the asteroid. The goals of Orbital B for REXIS were the main science requirements of the instrument: to measure global elemental abundances ratios of Bennu to determine the type of

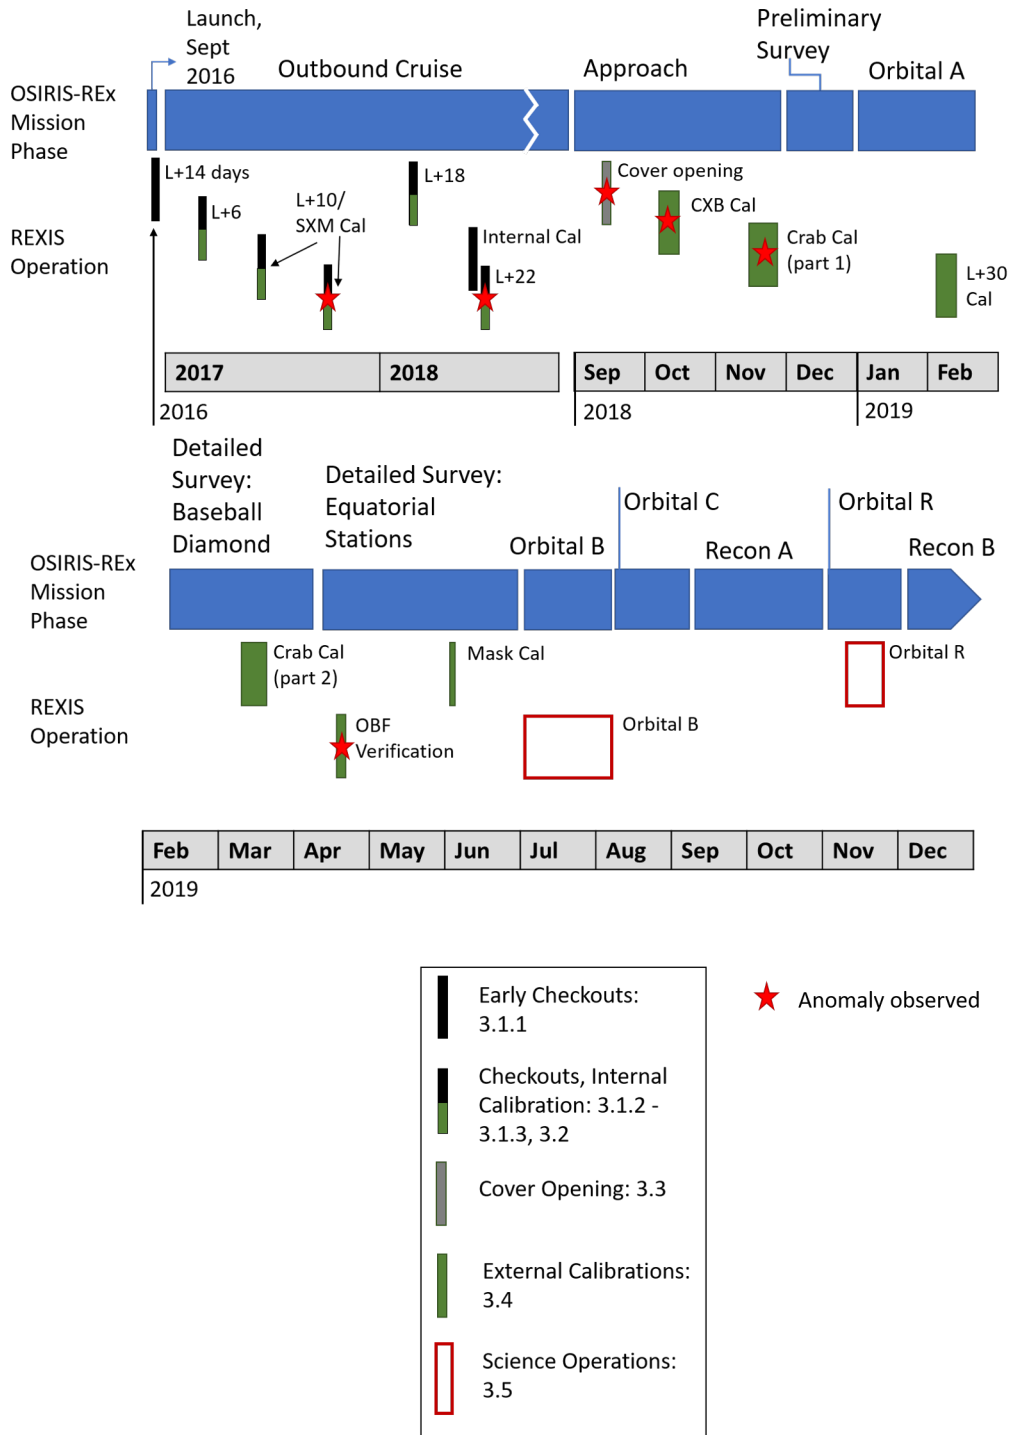


Figure 3-25: REXIS operations timeline, with notable anomalies marked by a star.

asteroid from the known meteorite groups, and to map the spatial distribution of those elemental abundances across the surface of Bennu, to help inform and provide context for the eventual sample site selection [4]. REXIS was designed for 420 total

hours of observation during Orbital B.

The Orbital B observations involved most of the instruments on board OSIRIS-REx. Due to the complexity of designing operational sequences for every instrument, the project instead designed reusable blocks for each instrument. Then, five separate sequences using each instrument block were written by the OSIRIS-REx mission operations team, each optimized for a different part of the orbit, or for a different combination of instruments. The instruments collected data during approximately 15-hour-long science windows every day. Due to the uncertainties of the position of the spacecraft at a given time, the actual schedule of sequences that was to be used for a given week was only chosen a few weeks prior to operation. The block development and testing occurred in early 2019 to ensure sufficient time for testing the full range of operation of each block.

Orbital B block

The REXIS team worked to develop a single block with enough flexibility to fit various operational needs. A diagram of the REXIS Orbital B block is shown in Figure 3-26. The instrument begins the loop by regenerating the bias map, and then either performing an initial sequence of six groups of varying ET and NEM settings, or skipping that first portion. Then, the image processing parameters are set, as indicated by the inputs to the command for the set of ET values, set of ULD values, NEM value, and EGF value. The ST parameters are hard-coded into the block, and so does not have an associated input. The hard-coded ST values were written before the outcome of the Mask Calibration was known, so they did not include the higher ST values tested in that operation. After the parameters are set, the block enters the “main” loop. In the main loop, the bias map is regenerated, telemetry is turned on for a specified amount of time, a raw frame of specified size is downlinked, and a bias map of specified size is downlinked. This main loop is repeated for a number of iterations as specified by one of the inputs to the block call. Within the main loop, the bias map regeneration, raw frame downlink, and bias map downlink frequency can be adjusted independent of each other; these three functions can occur every iteration

of the main loop, or every “x” number of loops, where “x” is specified by inputs to the block. There are separate frequency inputs for the bias map regeneration, raw frame downlink, and bias map downlink.

There are a total of twelve inputs to the block. In order of appearance in the input command, they are:

- a: The number of iterations of the “main” loop
- b: How often to regenerate the bias map (“every b loops”)
- c: How often to downlink the bias map (“every c loops”)
- d: How often to downlink a raw frame (“every d loops”)
- e: Length of time for which to leave telemetry on (per main loop), in seconds
- f: Run first cycle of 6 sets? (true/false)
- g: ET set to load (total of 10 different sets)
- h: ULD set to load (total of 4 different sets)
- i: NEM value to use
- j: EGF value to use
- k: Raw frame dimensions to downlink (6 different sets, also includes differing time to wait for downlink)
- m: Bias map dimensions to downlink (6 different sets, also includes differing time to wait for downlink)

The optional initial cycle of six sets in the block was for further exploration of ET and NEM configurations to optimize those settings for later observation. However, with the way Orbital B mission planning occurred, parameters could not be changed during Orbital B, so the parameters chosen two months prior to July 2019 were used for the entirety of the operation. The original parameters also included the nominal ST parameters used in all operations prior to mask calibration, which were hard coded into the block without any other sets to choose from. The parameter sets were chosen in February and March of 2019, before the Mask calibration operation was able to confirm that doubling ST for two nodes improved detector efficiency. For the majority of Orbital B observations, the six best performing nodes were enabled (CCD02, CCD10 and 12, and CCD20, 21, and 22).

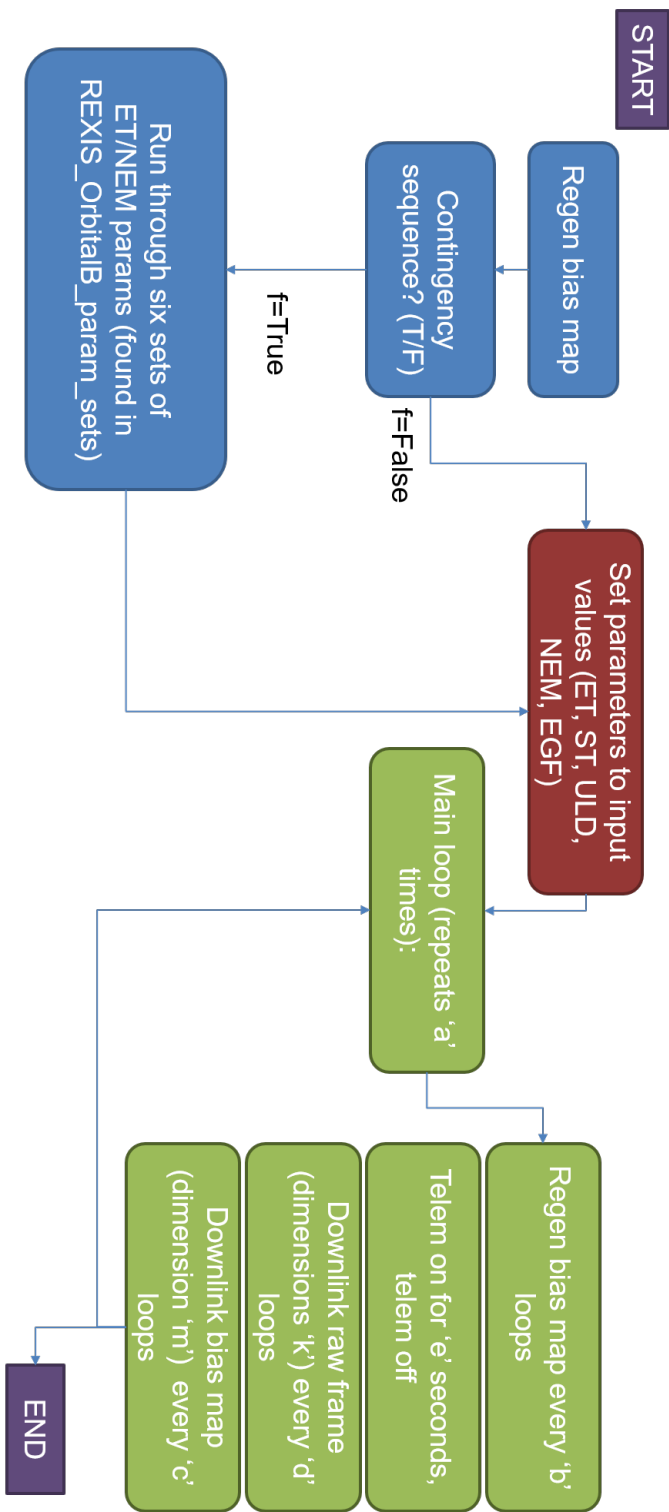


Figure 3-26: Block diagram of Orbital B block

Summary of Impact to REXIS Science

Four new anomalous behaviors were observed during Orbital B data collection. The anomalies and impact to data are summarized in Table 3.3. Two of the anomalous behaviors were observed in the data, and the other two affected the event list data. Additionally, the bad bias map behavior first observed in the OBF Verification (Section 3.4.5) continued; about 13% of the bias maps downlinked during Orbital B exhibited the odd behavior. Due to the anomalies that occurred during Orbital B data collection, the amount of usable science decreased drastically below the required amount. From the sequences scheduled for Orbital B, and accounting for only time when telemetry was on, the REXIS team began Orbital B expecting to downlink only 405 hours out of the required 420 hours. About 56 hours of science data were lost to bad bias maps, bringing the total to 349 hours. Due to a single event upset, an additional 146 hours were lost, leaving around 203 hours of usable data. Yet another hit to the data came from the detector efficiency losses, with some CCDs only operating at around 50-60 percent efficiency. In total, REXIS received less than half of the integration time it required to meet the science goals.

Table 3.3: Summary of anomalous behavior observed during Orbital B.

Anomaly	Date	Event list data lost (in hours)	Description
SXM histogram saturation	July 5-July 19	n/a	Most likely cause was multiple bits flipping. Fixed by instrument power cycle.
Excess flat noise	July 5-July 19	146	Most likely a single event upset. Fixed by instrument power cycle.
SXM decreased signal	All of Orbital B	n/a	Possibly a thermal effect on the MEB, still under investigation. Not fixed by power cycle.
Bad bias maps	All of Orbital B	56	Behavior observed in prior operations. Issue not fixed by power cycle.
Decreased detector efficiency	All of Orbital B	50-60% Efficiency	Not fixed by power cycle. Two of the nodes with the best resolution operated at 50-60% efficiency throughout Orbital B.

Detector Efficiency Decrease

Throughout Orbital B, the detectors were recording fewer iron counts than expected. Figures 3-27 and 3-28 show the predicted count rates from the Fe⁵⁵ calibration sources, and the actual measured rates for CCD20 and CCD22, respectively. During Orbital B, CCD20 detected only around 55.9 percent of the expected iron counts, and CCD22 only detected about 62.5 percent. Figure 3-29 shows a plot of count energy over time for CCD22 for one day of Orbital B. The solid line around 270-280 ADU is from the iron calibration sources. The boxes show times where the line disappears, which means at those times, REXIS is recording fewer or no counts from the iron calibration sources. Without the recorded energy from the calibration sources, the uncertainty with how the CCDs are reacting to the input signal increases, and the confidence in the resulting data decreases; since the iron source counts disappear, it is likely other X-ray counts were disappearing by the same mechanism. The iron counts disappearing serves as a measure for the loss of detector efficiency of the REXIS CCDs during Orbital B. A possible mitigation for the detector efficiency loss was to double the ST value for the affected nodes, as was done during the Mask calibration (Section 3.4.6). It was not possible to implement this parameter change during Orbital B, however.

SXM Data Issues

Throughout Orbital B, the SXM experienced two anomalies. The first occurred on July 5; there was a sudden saturation of high energy in the SXM data. This issue was likely due to bits flipping in the reset value for the energy histogram of the SXM, and was fixed when REXIS was power cycled. The second issue was observed throughout the entire phase. The overall signal level registered by the SXM was far lower than expected, given the solar state at the time. The most likely theory for the root cause of the decrease in signal is a thermal effect on the MEB. This signal decrease is still under study¹.

¹The root cause of the SXM signal decrease will be investigated in graduate student Andrew Cummings' Masters thesis in September of 2020.




Figure 3-27: Predicted iron count rates from calibration sources for CCD20 throughout flight. Actual measured values from each operation are represented by the dots. Plot created by HCO scientist Daniel Hoak.




Figure 3-28: Predicted iron count rates from calibration sources for CCD22 throughout flight. Actual measured values from each operation are represented by the dots. Plot created by HCO scientist Daniel Hoak.




Figure 3-29: Detected energy versus time for CCD22 from August 5, 2019

Single Event Upset

At the end of the day on July 5, another anomaly was noticed within the event list data, separate from the SXM histogram issue. Each node's event rate suddenly increased, and the noise remained flat for the following days. The top plot in Figure 3-30 shows the change in event rates before and after the noise appeared. Before the additional noise, count rates were consistently around two counts per second, and after the noise increase, the rates were consistently around 6 counts per second. The bottom plot shows the spectra observed before and after the appearance of the noise; the spectrum with the added noise is much flatter, making it difficult to discern any elemental lines. If the data continued to be collected with this increased noise, it would be much more difficult to determine the elemental ratios on the surface of Bennu. Due to the sudden appearance and persistence of the noise, it could not be due to variations in cosmic background - a varying background would have come on more gradually, and changed with time. Therefore it was hypothesized that the noise was due to a single event upset (SEU) to some value within the software that artificially added the noise. The noise was only observed on five of the six nodes




Figure 3-30: Plots showing change in total event rates and resulting spectra from before and after presumed SEU. The arrow on the top plot indicates where noise began. The red line on the bottom plot is the energy spectrum from before the SEU, and the black line is the spectrum from after the SEU.

- CCD21 did not record a noise increase - which further supported the hypothesis that the noise was artificial. Another clue that the root cause was internal to the instrument was the grade distribution, shown in Figure 3-31. During the period with the flat excess noise, grade 1 events show a similar sudden flat increase. Grade 1 events are often not physical X-ray events due to the physical layout, so seeing many grade 1 events over a short period of time is highly unlikely, which strengthened the hypothesis that the noise must be artificial and not physical. The exact place in memory where the bit or bits flipped was not determined. However, due to the severity of the consequences to the science goals of REXIS if the noise persisted, the REXIS team requested the instrument be power cycled to remove the effects of the SEU. The OSIRIS-REx mission planning team obliged, and four power cycles of REXIS were added in to the mission plan for the remainder of Orbital B to deter any further possible SEUs.

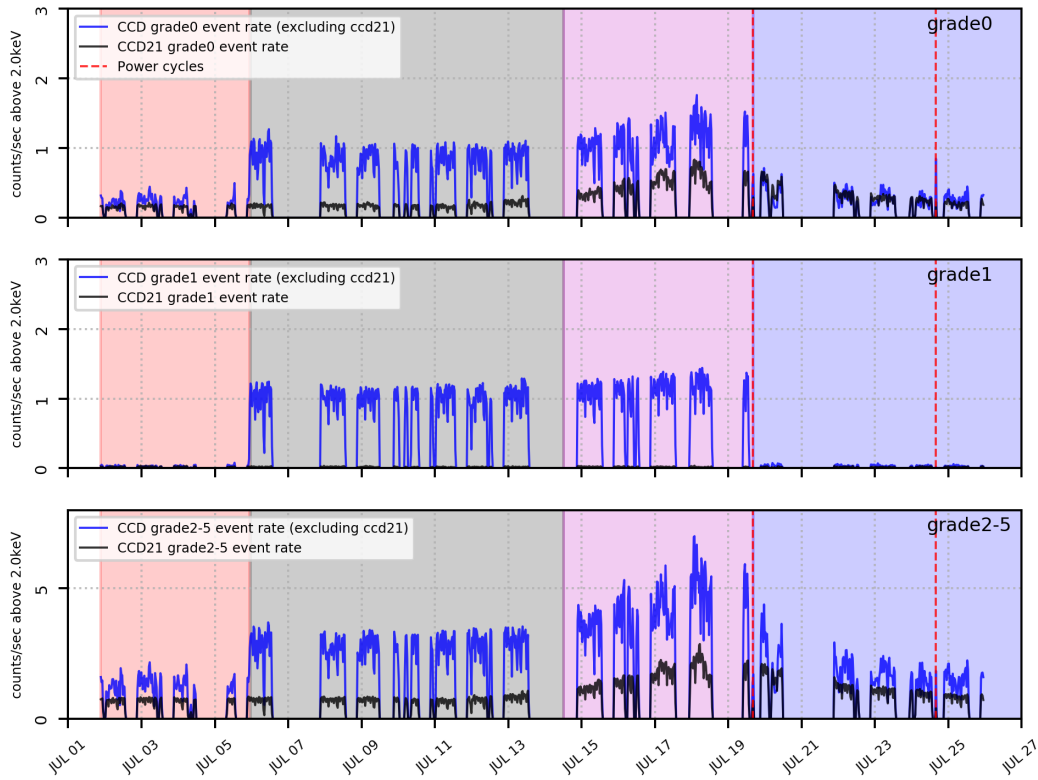


Figure 3-31: Orbital B events divided by grades and time. The first red section indicates nominal performance between July 1 and July 5. The second green portion indicates the data affected by the SEU. The third purple portion is still during the SEU-affected data period, with the addition of varying noise. The final blue portion is after the first power cycle, when the flat noise from the SEU disappeared.




Figure 3-32: Top plot shows event rates throughout Orbital B. Bottom plot shows spectra collected at different times throughout Orbital B

REXIS was power cycled on July 19, July 24, July 31, and August 4. The cadence of power cycles was chosen so that REXIS would reset roughly once per week. The total events and raw spectra collected throughout Orbital B are shown in Figure 3-32. The top plot shows the total grade 0 event rate throughout Orbital B, and the vertical dotted lines denote when REXIS was power cycled. The bottom plot shows the spectra from before, during, and after the excess noise. Immediately following the first power cycle, the flat excess noise no longer appears in the data, and the spectra recovers. The gradual increase and decrease in event rates seen between July 15 and July 23 is most likely due to variations in cosmic background, given its gradual appearance and disappearance. The excess noise no longer appears for the rest of the Orbital B operations, so power cycling REXIS more often was noted as a successful bit flip mitigation method for any future operations. However, due to the period with excess noise, REXIS lost about 146 hours of usable data.

3.5.2 Orbital R

Due to the loss of a significant portion of usable data in Orbital B, the REXIS team petitioned for and was granted more observation time. A second orbital period observation was planned for November 2019, named "Orbital R" due to its spot in the Reconnaissance phase of the OSIRIS-REx mission. This orbit was slightly larger, with a radius of 1.2 kilometers (measured from the center of the asteroid) compared to the 1 kilometer orbit of Orbital B. REXIS, along with other instruments, was turned on for two weeks, downlinking data during daily 16-hour-long science windows. The same structure of planning that was used for Orbital B was utilized for Orbital R; each instrument relied upon the use of blocks to create a few different sequences, and then the schedule of sequences was chosen based on orbital information a few weeks prior to the operation.

The two largest concerns from Orbital B were that another SEU could cause the loss of weeks of data, and that the detection efficiency loss indicated by the decrease of detected iron counts would continue, and REXIS would not collect sufficient data to achieve the science goals. To combat the former, the daily schedule was written so that REXIS would power cycle every day. To combat the latter, special commanding was written so that ST could be increased for three of the better performing nodes, CCD02, CCD20, and CCD22. To account for the additional event rates expected from increasing ST, only five nodes were used for Orbital R: CCD02, CCD10, CCD12, CCD20, and CCD22.

The total event rates for Orbital R are shown in Figure 3-33. There is no indication of noise similar to what was observed during Orbital B. One day, November 14, has a large increase in events, but the noise is not persistent past the power cycle, and did not cause the same grade distribution as the flat noise from Orbital B. The bad bias map issue persisted throughout Orbital R; the same percentage of bias maps downlinked were bad, roughly 13%.

The iron count rates recovered during Orbital B. Figure 3-34 shows the predicted iron count line, with boxes around the rates from Orbital B and Orbital R. The




Figure 3-33: Event rates and spectra for Orbital R. The spectrum from the day with higher event rates (November 14) is shown in black, and the two spectra of the periods preceding and following November 14 are shown in red and blue.

rates for Orbital R fall much closer to the predicted event rates than Orbital B rates did. Figure 3-35 shows the grades 0 through 5 iron count rates recorded on CCD22 throughout all of flight, without gaps. The green points on the left of the graph represent the data collected through internal cal, which were used to derive the theoretical decay curve of the iron sources. The diamonds on the bottom of the graph indicate the background counts in the source region of the detectors. During Orbital B (28 through 50 on the x axis), many of the event rates fall well below the predicted iron rate. Orbital R (50 and up on the x axis) iron event rates fall much closer to the predicted iron rate. Figure 3-36 shows a plot of energy versus time for CCD22 for one day of Orbital R. The iron line (outlined by the black line) remains visible throughout the whole day. From these improvements in iron count retention, it seems that changing the ST had the desired effect of recovering the CCD detection efficiency.




Figure 3-34: Predicted iron count rates and rates from each observation. Plot created by HCO scientist Jae Sub Hong.

The expected total integration time for Orbital R was about 220 hours. The bad bias map issue resulted in the loss of about 30 hours of data, resulting in roughly 180 hours of total usable science data. REXIS was powered off for the last time between 08:18 and 08:20 eastern time on November 25, 2019.




Figure 3-35: Grades 0 through 5 iron counts on CCD22 from launch through Orbital R. Plot created by HCO scientist Jae Sub Hong.

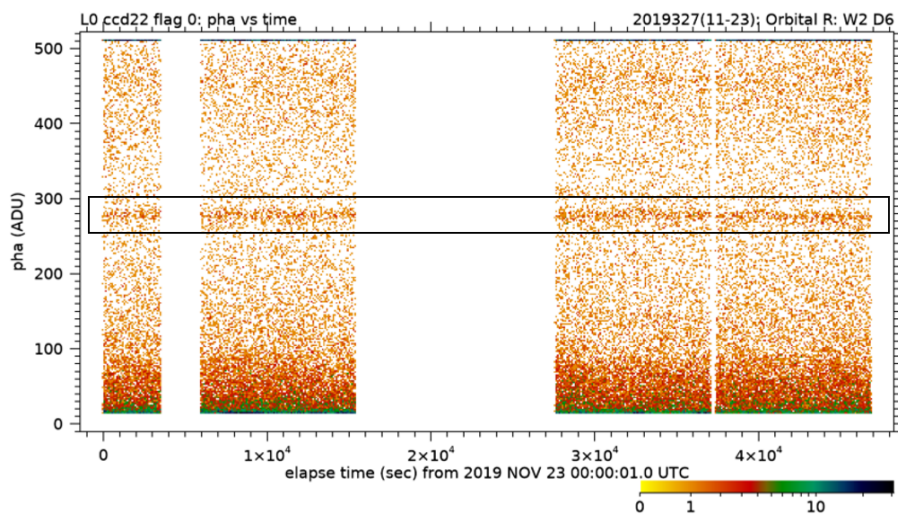


Figure 3-36: Detected energy versus time for CCD22 from November 23, 2019. (The large gap in the data between 16,000 and 29,000 seconds is removed due to a bad bias map)

Chapter 4

Root Cause Analysis

Upon cover opening, REXIS began to observe higher event rates, sometimes high enough to breach the event rate filter. The REXIS timeline is shown in Figure 4-1, and has the operations with higher event rates outlined; both the cover opening and the first Crab Nebula calibration recorded far higher event rates than expected, and the Crab calibration data had little scientific value because of an event rate filter breach due to event rates continuously surpassing the 200 counts per second filter limit. In Figures 4-2 and 4-3, the long-term event rates throughout all of flight are shown. There is a clear increase in events upon cover opening (marked by the vertical dotted line) with the noticeable outlier of the L+10 EGA calibration operations, where the events were increased due to high CCD temperatures. These periods where event rates were elevated were found to be related to spacecraft pointing, and suspected to be most likely due to scattered light on the detectors. Through various mitigation methods (changing EGF to only downlink grades 0-5, regenerating the bias map more frequently), the team was able to avoid filter breaches. During Orbital B, however, the spacecraft was at its closest distance to the asteroid yet, placing REXIS in a brighter optical environment.

During Orbital B, the event rate filter was not breached, due in part to changing the EGF to ignore grades 6 and 7 events. However, there were periods of time during Orbital B where events increased by over five times the baseline rate. An example of this behavior is shown in Figure 4-4. Additionally, during Orbital B the REXIS

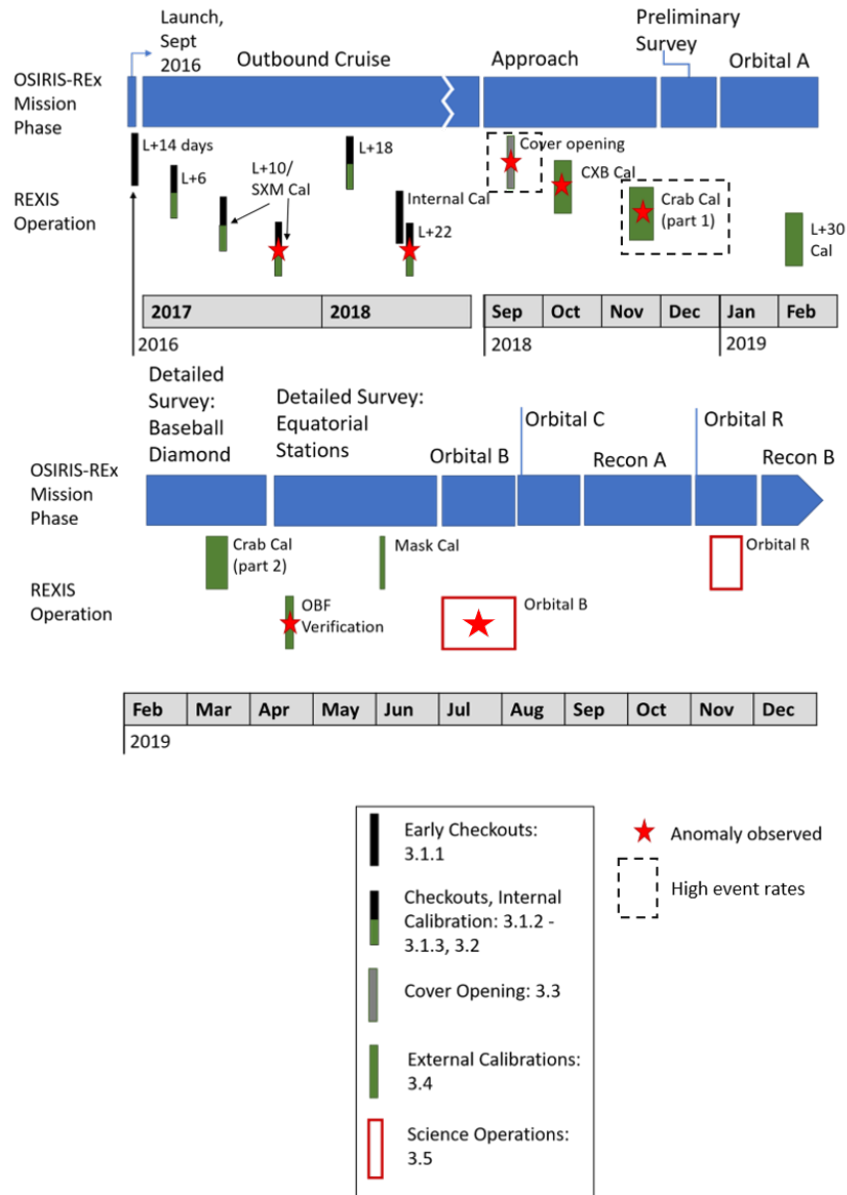


Figure 4-1: REXIS flight operations over time. Red stars indicate an anomalous behavior was observed, and black dotted line indicates high event rates.

instrument downlinked fewer iron calibration source counts than anticipated. This loss of iron counts indicates that the detector efficiency was decreased for Orbital B. Figure 4-5 shows the pulse height amplitude of the downlinked events over time, with time shown in seconds from midnight UTC, and covers one full day of Orbital B, August 5. The pulse height amplitude is measured in analog-to-digital units (ADU). ADU units represent how the incident photons are converted to digital signal, and are

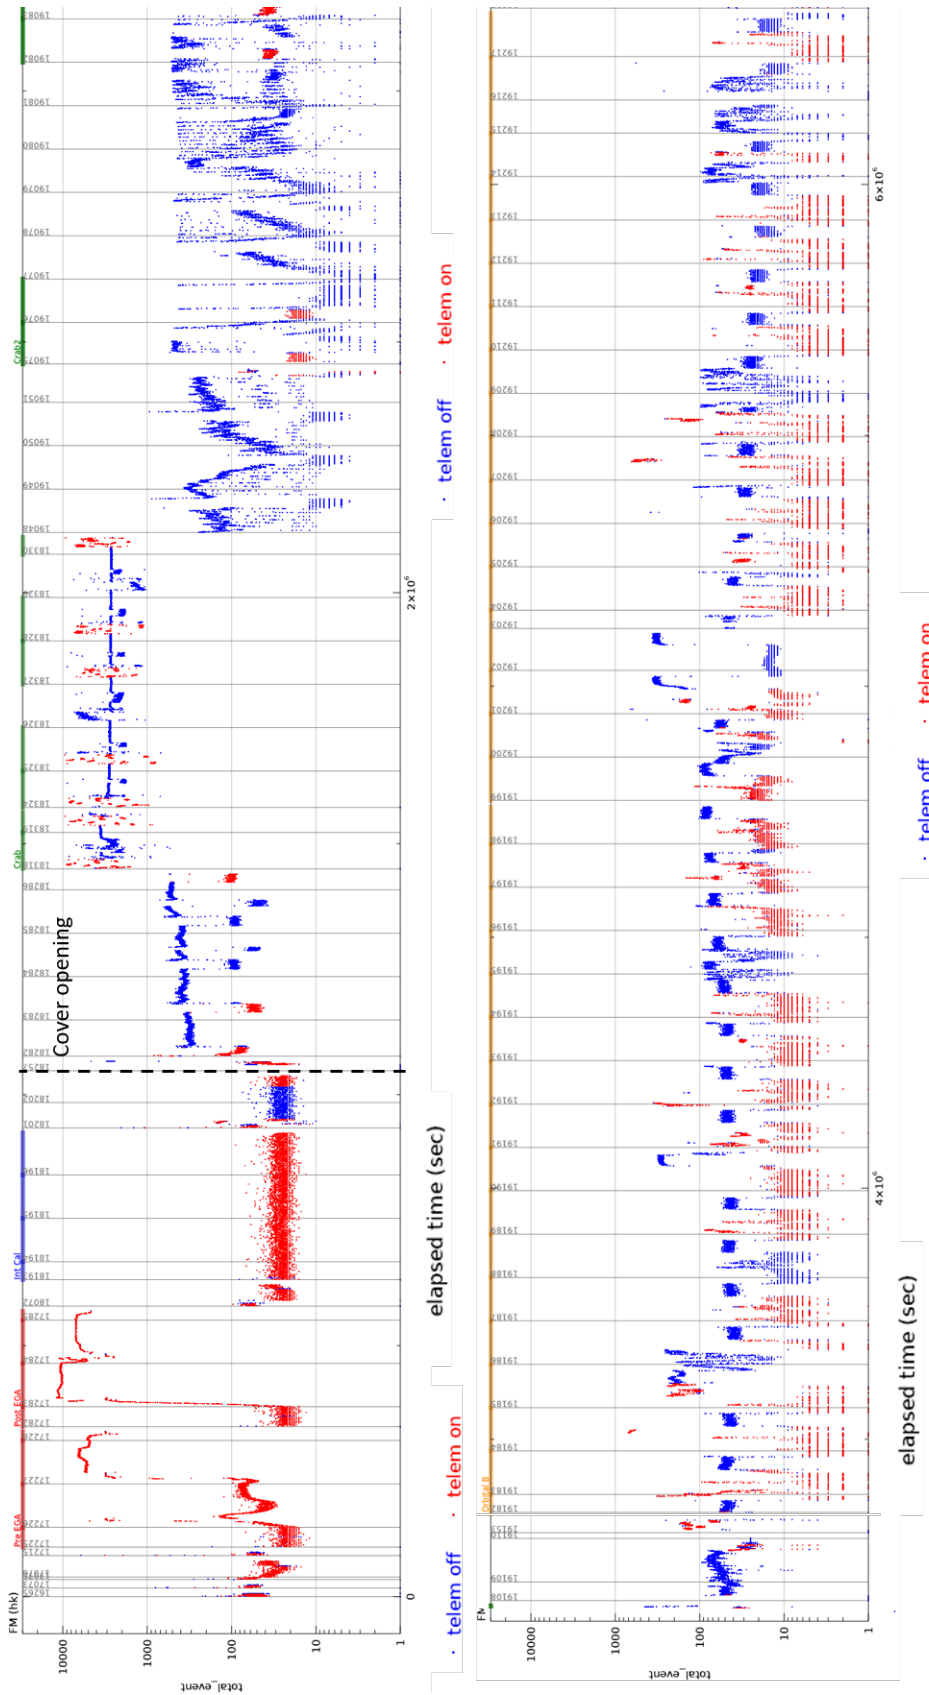


Figure 4-2: REXIS housekeeping event rates throughout all of flight. Downlinked events are shown in red, and housekeeping events are shown in blue. X axis is elapsed time in seconds, and y axis is number of counts.




Figure 4-3: Continuation of REXIS housekeeping event rates plot. Orbital R data are shown, with downlinked events shown in red and housekeeping events shown in blue. X axis is elapsed time in seconds, and y axis is number of counts.

converted to energy based on the gain and offset of the CCD node. The horizontal dotted line indicates the expected ADU value for the iron line around 5.9 keV, or around 280 ADU. In the periods of time indicated by the black boxes, the iron line disappears, meaning the instrument did not downlink iron counts during that period. Upon initial inspection of plots such as Figure 4-5 and 4-4, the periods of high event rates seem to correlate with the periods where iron counts decreased. If the periods of iron count loss do correlate with the higher event rates, the root cause of the loss of iron could be related to the other periods of high events from flight (cover opening, first Crab calibration).

REXIS CCD events: 52412 seconds of data from 2019 JUL 25 00:00:03.2 UTC




Figure 4-4: Event rates for one day of Orbital B. Spikes in event rates are indicated by arrows.




Figure 4-5: Plot of event energy vs time

In this root cause analysis of the iron count dropout, the correlation between higher count rates and loss of iron counts is investigated, and the potential causes of the high event rates are identified in a fishbone diagram. Each fishbone on the diagram is analyzed to determine its viability as the root cause, and the most likely root cause is presented with evidence.

4.1 Iron Count Rates versus Total Event Rates


To investigate the correlation between total event rates and iron count loss, total event rates are compared to total iron event rates throughout Orbital B. All Orbital B event list data, except the data affected by bad bias maps or the single event upset that contaminated the data between July 5 and July 19, are divided into bins of 600 seconds. Bins that are close to the end of an observation decrease in length to fit only the times of interest. Then, the iron counts are isolated from the event data by only looking at events in the physical area on the detector that is illuminated by the Fe^{55} sources, and choosing the events there within the correct energy range (around 270-280 ADU). For each bin, the event rates and iron event rates are calculated by dividing the respective total counts in that bin by the total length of time.

Figure 4-6a shows a plot of the iron event rates versus the total event rates for just grade 0 iron events, and Figure 4-6b shows a plot of grades 0 through 5 iron events versus total events on a log scale. The horizontal line marks the expected event rates throughout Orbital B, around 5,000 counts per kilosecond. This number is calculated by finding the average event rate per day of CCD22 during Orbital B. The vertical line marks the expected iron event rate for the respective grades. These values were found by calculating the average iron event rate over 600 second bins for all observations from before the cover was opened, and extrapolating using the known decay rate of Fe^{55} , found with equation 1.1 where A represents the number of decays per unit time, and λ is the decay constant. For Fe^{55} , the decay constant is 0.253 years^{-1} , corresponding to a half life of 2.737 years, using the relationship 1.2 where $t_{(1/2)}$ is the half life. For grade 0, on CCD22, the expected value during Orbital B was around 21 counts per kilosecond, and for grades 0 through 5, the expected rate was about 55 counts per kilosecond. The plot is divided into four quadrants. Quadrant I is when higher total event rates correlate with lower iron rates, quadrant II is where higher total event rates correlate with higher iron rates. Quadrant III represents where lower event rates correlate with higher iron rates. Quadrant IV represents where lower event rates correlate with lower iron rates. Both plots show an excess of events in quadrant

I. As total event rates increase, the iron rates decrease, which is not to the expected relationship - the iron count rates should remain relatively stable independent of total count rates, as the iron sources are not reliant on input X-ray flux from the sun to release any photons. There is a clear trend between high event rates and low iron rates.

$$A = Ae^\lambda \tag{4.1}$$

$$\lambda(\text{years}^{-1}) = \frac{\ln 2}{t_{(1/2)}} \tag{4.2}$$



(a) Total event rates vs grade 0 iron events

(b) Total event rates vs grades 0-5 iron events

Figure 4-6: Plots of iron event rates versus total event rates for all of Orbital B in log scale

4.2 Potential Causes

For this root cause analysis into detector efficiency loss as evidenced by loss of iron counts and high count rates, a fishbone diagram (described in Section 2.2) is used to identify all possible root causes. Once the causes are identified, they are systematically ruled out until the most likely root cause is identified. Figure 4-7 is the fishbone diagram that was compiled by the REXIS team and presented to OSIRIS-REx project management upon the discovery of the loss of iron counts. It lists the possible root

causes for the loss of iron counts, as first identified by the team upon discovery of the iron count loss. At the time of the diagram's creation, the REXIS team had observed the many other in-flight behaviors that indicated instrument sensitivity to the background at the spacecraft pointing, and stray light was already strongly suspected. Therefore, the fishbone diagram is relatively sparse, due to the biased mindset with which it was created. The possible root causes are divided into two categories, software and physical (here, physical indicates either a hardware issue or an astrophysical cause (like stray light)). The software-based possible root causes are either the ET or ST parameters being set incorrectly, or the bias map regeneration error resulting in image processing with a bad map. The possible physical causes identified were temperature variation of the CCDs or detector electronics, voltage variations, or an optical light leak.

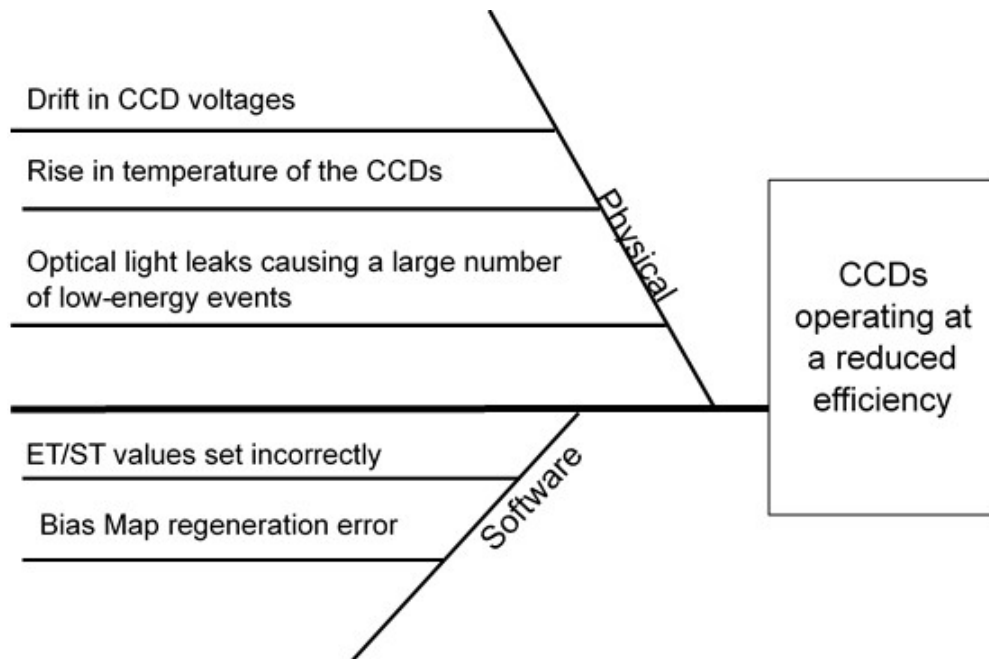


Figure 4-7: Fishbone diagram for potential causes of CCD loss of iron counts

4.2.1 Physical Causes

As explained in section 2.1, CCDs are controlled by changing different voltage levels to transfer charges from the pixels to the readout area. If a voltage drifts from its

nominal value, the charge transferred to the readout area and subsequently converted to X-ray events by the FSW could potentially not be representative of the actual charge present on the detector surface. Specifically, if the output gate voltages drift low enough, they could become too low to trigger the semiconductor properties of the silicon, leaving all the charge trapped in the affected pixels. Additionally, if the video output voltages drift outside their nominal operating range, the resulting output images could be corrupted. In either case, the resulting data would then contain a different pattern of X-ray counts than actually incident on the detector, which could cause the iron counts to be lost.

If the main electronics board (MEB) had a component or subsection with unexpected sensitivity to temperature, a variation in temperature could cause detector thresholds to shift, resulting in additional noise on the detectors. If the noise events coincided with an iron event and were above ST, the flight software would count the noise events as part of the iron count, and artificially increase the event from a grade 0 to higher. If the grade was increased to a 6 or 7, the event would not be included in downlinked data, resulting in the observed loss of iron calibration counts. CCD temperatures can also affect performance; if the temperatures are too high, the electronics begin to produce dark current that appears as global additional noise. However the CCD temperatures remained well below the temperature at which dark current began appearing (-40 degrees Celsius), so CCD temperatures can be ruled out immediately, but MEB temperatures still remain a possibility.

If optical light shines on the detector plane, the resulting energy events could produce noise localized to where the light hits the detectors. As optical light is lower energy than the soft X-rays REXIS was designed to detect, the optical light would show up as lower energy events. The majority of the detector plane had the DD-OBF to prevent optical light from reaching the detector surface; however, the edges of the CCDs that contain the wire bond pads that connect to the flexprints do not have OBF over the top of the wires. These sections are protected by portions of the DAM, but light could bounce off the reflective surfaces of the DAM and onto the uncovered readout edges of the nodes. If the light events appear next to or adjacent to actual X-

ray events from the calibration sources and increase the energy in the adjacent pixels to above ST, the FSW will count that ^{55}Fe event as a higher grade than nominal, and potentially (if the grade was artificially increased to 6 or 7) throw out the data from that event entirely.

4.2.2 Software Causes

If either ET or ST parameters for individual nodes are set to the incorrect values - i.e. if the value commanded to the instrument and the value used by the instrument differ - the image processing algorithm will process data from that node with the incorrect energy thresholds. If ET is set higher than the energy of the iron counts, those counts will not appear in the resulting X-ray candidate list. Similarly, if ST is set too low, the iron counts may be counted with surrounding pixels of lower energy, artificially increasing the grade of the event to an ‘ignored’ grade. However, this potential cause was quickly ruled out by comparing commanded values to the values in the downlinked parameter settings packets. All parameters were set to the expected values.

As mentioned in Section 3.4.5, during the OBF Verification, an issue with bias maps was discovered. This bad bias map is subtracted from all subsequent data until a new bias map is generated. If random noise is being subtracted from data instead of actual background, it is possible that the ^{55}Fe counts, as well as other real X-ray events, would be removed from the data. However, this possible cause was also quickly ruled out by removing the periods with bad bias maps during data processing performed on the ground after the first few days of Orbital B. Once the sections of data from bad bias map portions were removed, there was still a deficit in the iron counts.

4.3 Less Likely Root Causes


After identifying the possible causes and compiling them into a fishbone diagram, each bone on the chart is investigated more thoroughly to determine its viability as a root

cause for the behavior under study, specifically, the decrease in ^{55}Fe counts during periods of Orbital B science observations, as tracked by the behavior's correlation with higher event rates. This section contains the investigations for the causes that are less likely.

4.3.1 Drift in CCD Voltages

When in Science mode, REXIS records the voltage levels for defined channels every 64 seconds, and places the data in housekeeping packets. The housekeeping packets are downlinked from the spacecraft along with science data. A voltage average is taken from Orbital B and Orbital R to determine if there was a significant different in the video output voltages or the CCD output gate voltages. The long-term voltages are also investigated to determine if voltage drift is correlated with the increase in event rates and decrease in iron counts. Figure 4-8a shows the CCD2 video output voltages, and Figure 4-8b shows the CCD2 output gate voltage for all of Orbital B and Orbital R. The taller spikes that extend beyond the upper or lower limits of the vertical axes are artifacts that occur when REXIS is power cycled. The other noise is due to a periodic noise that occurs roughly every 5 hours (more clearly shown in Figures 4-9 and 4-10). This noise has been present since launch, and the root cause of it is not determined; however the noise does not have a significant effect on REXIS data. The average voltages for each of the operational CCD video outputs and output gate are shown in Table 4.1. Between Orbital B and Orbital R, there is very minimal change in the voltage for each channel, less than 0.03 percent for the VDC channels, and less than 0.1 percent for the output gate channel. The voltages in Table 4.1 fall well within the safe operational voltage ranges, indicated by the horizontal lines on the plots in Figure 4-8. The periodic noise occasionally results in a single voltage reading outside the yellow limit of the safe voltage range, but the reading is transient and has never persisted for more than one housekeeping packet, nor has it had any measurable effect on REXIS data.

To determine if the periodic noise is correlated with the increase of event rates, the voltages and event rates for two separate days of Orbital B are compared temporally.



(a) CCD2 video output voltages for Orbital B and Orbital R.

(b) CCD2 output gate voltage for Orbital B and Orbital R.

Figure 4-8: CCD2 voltages for Orbital B and Orbital R. The yellow (and red, where applicable) high and low voltage alarm levels are indicated by the horizontal lines.

Table 4.1: CCD video output and output gate voltage average comparisons for the six nodes active for Orbital B.

Channel	Orbital B avg (V)	Orbital R avg (V)	Percent Difference
CCD02 VDC	14.6587	14.6572	0.01023%
CCD10 VDC	15.5911	15.5947	0.02301%
CCD12 VDC	15.6044	15.6013	0.01987%
CCD20 VDC	14.7013	14.6970	0.02925%
CCD21 VDC	14.7287	14.7273	0.00951%
CCD22 VDC	14.8549	14.8575	0.01750%
CCD0 output gate	0.2890	0.2891	0.03460%
CCD1 output gate	0.3162	0.3163	0.03162%
CCD2 output gate	0.2456	0.2454	0.08145%

Figures 4-9 and 4-10 show the CCD2 video output voltages versus total event rates for July 29 and July 31 (the dates are chosen at random). The increases in event rates in Figure 4-9 occur after the noise in the voltage, while the increases in event rates in Figure 4-10 occur before or during the noise. There is no clear trend between the increase in event rates or the voltage noise. Therefore, the loss of iron calibration events is unlikely to be due to voltage changes. The voltage drift bone on the fishbone can be ruled out.




Figure 4-9: Voltages versus event rates for one day of Orbital B (July 29)




Figure 4-10: Voltages versus event rates for one day of Orbital B (July 31)

4.3.2 Rise in CCD or MEB Temperatures

To determine if a temperature change could be the cause of the iron count dropout, temperatures for the MEB, video board, and detector electronics for all of flight are plotted in Figures 4-11 and 4-12. The average temperatures from Orbital B are no more than three percent different from the prior external calibration operations, none of which recorded a significant decrease in iron count rates. To determine if the shorter-term variations in temperature correlate with event rates, the CCD temperatures are plotted against the event rates from a period of Orbital B, shown

in Figure 4-13. The temperatures do not vary with the event rates. Due to the fact that temperatures do not vary with event rate, and do not vary significantly from prior flight temperatures, temperature changes can be ruled out as a potential root cause for the loss of iron calibration counts.




Figure 4-11: MEB (blue) and detector electronics temperatures throughout all of flight. The bottom plot contains all temperatures for all of Orbital B. Plot created by HCO scientist Jae Sub Hong.




Figure 4-12: MEB (blue) other detector electronics throughout all of flight, continued, depicting temperatures for Orbital R. Plot created by HCO scientist Jae Sub Hong.




Figure 4-13: Temperatures versus event rates for one day of Orbital B

4.4 Most Likely Root Cause

All other bones on the fishbone diagram have been ruled out except for the potential light leak on the detectors. The hypothesis tested in this section is that the loss of

iron counts is caused by an optical light leak onto the detectors. As mentioned in Section 4.2, there were sections of the edges of the CCDs where the wires connecting to detector electronics were not coated with OBF, so as to not short out the wires. A diagram showing potential pathways for optical light through the DAM structure and onto those uncoated edges of the detectors is shown in Figure 4-14. A photograph of the underside of the DAM with arrows indicating the places optical light could shine through is shown in Figure 4-15.




Figure 4-14: Diagram showing possible pathways of light through DAM to readout edges of the nodes. Diagram created by HCO scientist Jae Sub Hong.

If light is hitting the readout edges of the detectors, there should be a large number of low-energy events with a non-uniform distribution across the detectors. There should be a higher concentration of low energy events nearest to the edge of the detector, and the distribution should attenuate as the distance from the readout edge of the node increases. The iron counts should decrease or disappear when there are spikes in the event rate of low energy, and the iron counts should be inversely proportional to the amount of optical light present around the detector. An operational mitigation method to recover iron counts could be to raise ST; by raising ST fewer




Figure 4-15: Picture of the bottom of the DAM, with arrows indicating where optical light could enter.

low energy events would be counted as part of true X-ray events, and so the detector would downlink more iron counts.

Three separate analyses are performed to investigate the light leak as a potential root cause. First, total event rates are compared to a predicted light curve from the asteroid to determine whether there is a correlation between higher count rates and optical light exposure. Then, flight data are compared to data from a stray light simulation performed at Lincoln Laboratory during the design phase of the REXIS instrument, to determine if the non-uniform distribution of low energy light events is similar to the distribution of events in flight data.

Finally, a simulation is created using Orbital B flight data to investigate the effect of low energy events on the grading of grade 0 events, to determine if the addition of low energy from a light leak could cause artificial grade increases. Flight data are used to create an estimate of the nominal low energy distribution and the light energy distribution, and the resulting energy is randomly selected from and placed into a three-by-three model of a CCD that contains a grade 0 iron event. The output of the model is a simulated frame. The frame is then run through the REXIS flight software, and the resulting simulated flight data file is parsed and examined using REXIS pipeline tools to compare the event grade distribution to the expected

distribution (which should be only grade 0 events).

4.4.1 Comparison of Optical Light to Iron Counts

High total count rates correlate with low iron rates, as shown in Figure 4-6. To investigate a light leak as a potential root cause, total event rates are compared to the calculated optical brightness observed by REXIS throughout Orbital B. The optical brightness prediction was calculated by HCO researcher Davied Guevel. Optical brightness is determined using the angle between the sun and Bennu, and the angle between Bennu and the spacecraft, and the relative albedo from the asteroid. The optical light prediction is limited, however; it does not account for reflections off the inside of the tower, and does not account for other instruments aboard OSIRIS-REx, or the surface roughness of Bennu. Figure 4-16 is a plot of housekeeping event rates versus the calculated optical brightness. Optical brightness is the fraction of optical light reflected onto the CCDs, and is calculated by taking the cosine of the illumination angle multiplied by the cosine of the emission angle, multiplied by the relative albedo from the surface of the asteroid. The plot shows some higher events at an optical brightness of zero, but those are most likely due to the increased housekeeping event rates that occur during a DSN pass. During DSN passes, the instrument may be pointed at other bright sources, but away from Bennu, so the optical brightness from Bennu is zero. The plot shows a trend between higher optical brightness and higher event rates.

Direct comparison of iron counts versus optical brightness does not show a clear correlation between amount of optical brightness and iron counts. Figure 4-17 is the comparison of total event rates to iron event rates, with different shapes depicting the amount of optical brightness during the period of time associated with each data point. Once the brightness reaches above 0.02, there does not seem to be a clear pattern. As stated before, the optical light curve prediction is simplified, and does not account for reflections of light off the inside of the REXIS tower, off any spacecraft structures, or from the irregular surface of Bennu.

A different way to quantify the light exposure of the detectors is to determine




Figure 4-16: CCD2 event rates versus calculated optical brightness (calculated using relative albedo) from Bennu. Plot created by HCO researcher David Guevel.


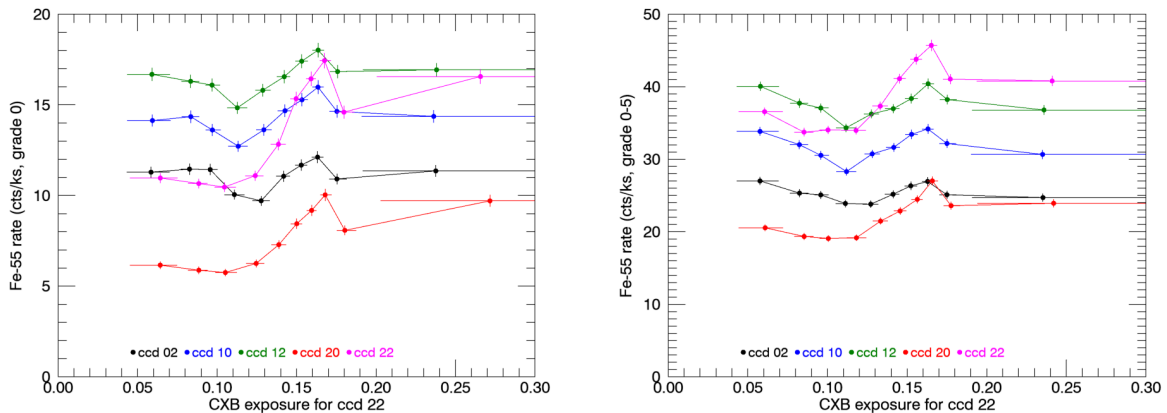


Figure 4-17: Plots of iron event rates versus total event rates for all of Orbital B, with optical brightness ranges indicated by different marker shapes

the amount of the detector surface exposed to the asteroid surface in comparison to the off-asteroid pointing. The off-asteroid exposed portions of the CCDs detect the CXB. An analysis performed by HCO research scientist Jae Sub Hong shows a positive correlation between the amount of detector surface exposed to CXB only and iron count rate. Figures 4-18a and 4-18b show the iron event rate for nodes CCD02,

CCD10, CCD12, CCD20 and CCD22 (the five best-performing nodes from Orbital B) as a function of the CXB exposure of CCD22. The iron event rates increase as more of the node is exposed to CXB. As the asteroid, the brightest source of optical light, leaves the field of view of the node, the iron events increase. The decrease in iron rates is correlated with increased optical brightness from the surface of Bennu.



(a) Comparison of grade 0 iron count rates on CCD02, CCD10, CCD12, CCD20 and CCD22 to the amount of CXB exposure of CCD22.

(b) Comparison of grades 0-5 iron count rates on CCD02, CCD10, CCD12, CCD20 and CCD22 to the amount of CXB exposure of CCD22.

Figure 4-18: Comparison of just grade 0 iron counts and grades 0-5 iron counts to CXB exposure on nodes. The x-axes are CXB exposure as a percentage of node area, and the y-axes are iron counts in counts per kilosecond. Plots created by HCO scientist Jae Sub Hong

4.4.2 Comparison of Stray Light Simulation to Flight Data

During Orbital B, the event distribution across CCD22 appeared non-uniform, by eye, with a greater concentration near the readout edge of the node and a lower concentration as the distance from the readout edge increased. An example of the event gradient is shown in Figure 4-19. This gradient appeared similar to the results of a stray light simulation performed at MIT Lincoln Laboratory in late 2014 and early 2015.

In the fall of 2014, after the first application of the OBF onto the CCID-41 devices, a test performed at MIT Lincoln Laboratory showed that there was potential for a light leak onto the imaging surface of the detectors. After the test, different




Figure 4-19: Example of event gradient across CCD22 seen throughout Orbital B

mitigation strategies were discussed. Further explanation of this process and the resulting decisions can be found in Chapter 5. As part of the investigation, a stray light model was constructed, and simulations were run to model the results of different detector protection strategies. The stray light simulation results also showed a gradient of events across a detector node, caused by optical light incident on the uncoated readout edge of the node, similar to what is observed in Orbital B flight data.

For the Lincoln Lab stray light simulation, a simplified model of the REXIS detector layout including the DAM was created in TracePro, a ray-tracing software program. Then, simulations were run using differing layers of OBF on the detector, and also the application of a black, light-absorbing paint (called Z306) underneath the DAM. The simulation results took the form of energy density maps of the surface of the CCDs. For simplicity of reporting, only one CCD is shown in the resulting image. The images shown in Figure 4-20 are the calculated spatial optical density (OD) map. The map shows how much energy is absorbed on the detector plane due to stray optical light. Optical density is a relation of the irradiance (energy density) of each pixel to the irradiance of the initial source of the light, weighted by the quantum efficiency of the detector. The OD is a measure of how much energy gets transferred from the incident optical light to the detector. Figure 4-20 shows the results of only

using an additional aluminum layer (i.e. using a thicker layer of the OBF) compared to what the detector optical light absorption would look like if the DAM was painted using the light-absorptive black paint. Figure 4-20b shows much less absorbed light on the surface, but Figure 4-20a is more representative of the actual flight-like configuration of the detectors. The determining factors for the decision to not paint the DAM with the black paint are discussed in Section 5.3.



(a) Detector plane with additional layer of Al OBF

(b) Detector plane with black paint on DAM

Figure 4-20: Spatial OD maps of CCD absorption on one CCD with varying levels of light leak protection

The data from the Lincoln Lab stray light simulation are compared to Orbital B flight. Two comparisons are made of the spatial distributions of the stray light sim and Orbital B data - one from a period of Orbital B data with iron count loss, and one from a period without iron count loss. The results of the comparison are shown in Figures 4-21 and 4-22. For the simulation data, the total absorbed power per column of the CCD are plotted from the readout edge inward. For the Orbital B data, the event distribution is plotted from the readout edge inward. The flight data are also broken down into two energy levels, one below 25 ADU, and one above, to see if the assumption that the gradient could be due to low energy optical light reaching the edge of the detector is plausible. If the gradient appears in both high and low energy

flight data, or in the data without iron count loss, it would be an indicator that light leak is not the issue, as optical light would only present as low energy in comparison to the rest of the electromagnetic spectrum under study.

Figure 4-21 shows the total absorbed power per column from the readout edge inward for the LL simulation data on the left. The plot on the right shows flight event list data in total events per column from the readout edge inward for CCD22 during a period of July 5 where there is a loss of iron counts. The top of the right plot shows events under 25 ADU, and the bottom shows events over 25 ADU. There is a gradient in the lower energy flight data (fit to a function) and no gradient in the higher energy data. The gradient in the distribution of the events over 25 ADU is overlain on the stray light simulation data (on the left in Figure 4-21, and scaled to account for amplitude. Figure 4-22 is the spatial distribution of events from Orbital B data with no loss of iron counts. There is no gradient in either low or high energy in the events from the period with no iron count loss. There is a gradient present in flight data suspected to be from a period of light leak similar to the gradient seen in data from the stray light simulation, and no gradient in flight data from a period with no suspected light leak. Therefore, a light leak onto the detectors via light reaching the readout edge of the nodes is plausible.

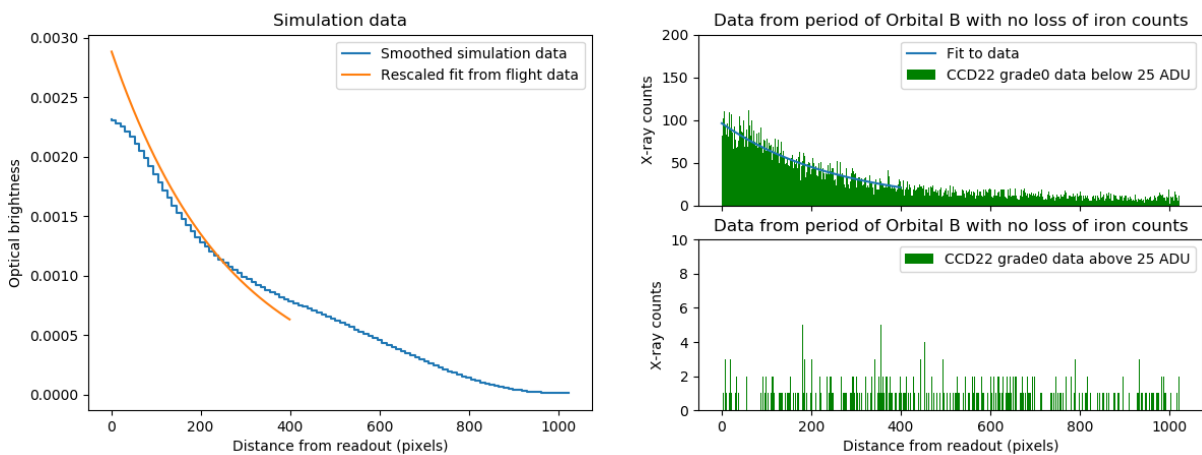


Figure 4-21: LL stray light simulation spatial distribution compared to spatial distribution of events from Orbital B with loss of iron counts




Figure 4-22: Spatial distribution of events for Orbital B data with no loss of iron counts

4.4.3 Artificial Grade Increases

A model was created to observe how additional low energy from a light leak affected the event grading algorithm. The hypothesis is that the additional low energy from a light leak fills the detector plane near the readout edge, and registers as energy on the pixels surrounding a grade 0 X-ray event. If the energy is higher than ST and lower than ET, the energy from the pixel (or pixels) is counted as part of the X-ray event, and the REXIS FSW grading algorithm registers the event as a higher grade than grade 0. For the periods where iron rates decrease drastically or disappear completely, the X-ray events are artificially increased to grades 6 or 7, which the EGF was set to ignore during Orbital B. The model consists of a three by three matrix to represent the pixel grid used for the REXIS FSW grading algorithm, with a single grade 0 iron event placed at the middle. The remaining eight pixels are then randomly assigned a value from a low energy distribution derived from various flight products. The distribution contains the nominal low energy recorded by the a single node of the detector (CCD22), with an additional distribution to represent the energy from a light leak. To test the hypothesis, three separate cuts of flight data are used to create low energy distribution estimates: one from Orbital B where iron counts disappear

completely, one from Orbital B where iron counts remain steady, and one nominal low energy distribution without additional energy from a light leak, as seen in the L+30 calibration. An overview of the process followed for this simulation is shown in Figure 4-23.

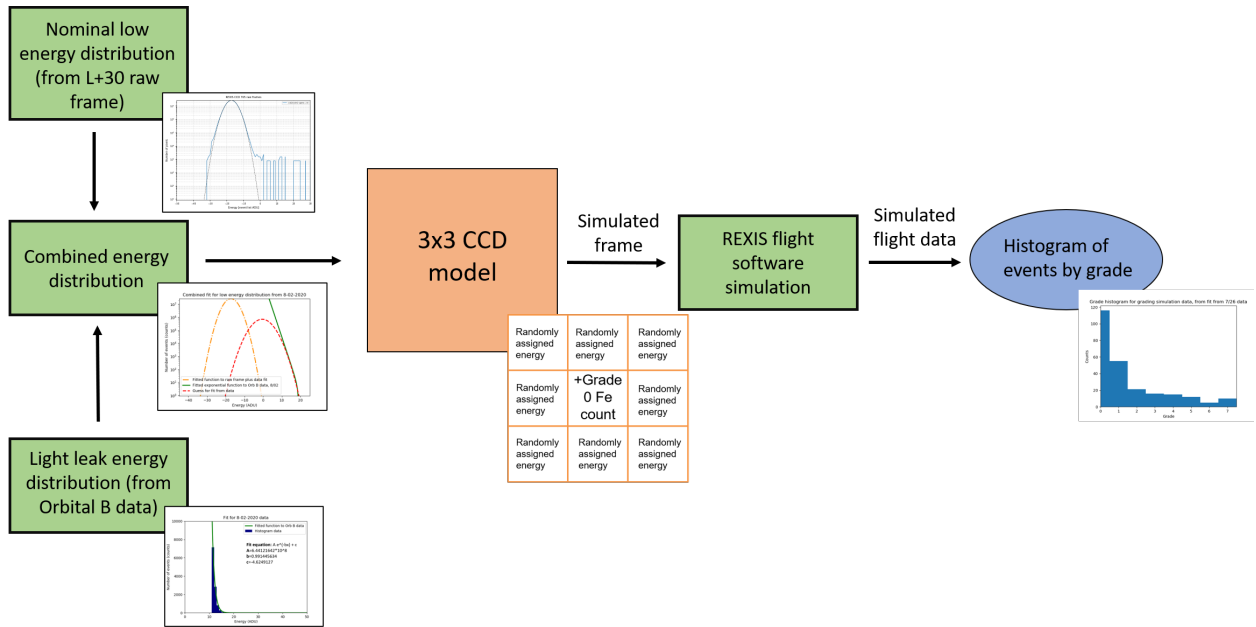


Figure 4-23: Diagram showing steps involved in grading simulation.

In the first step of the simulation, the nominal low energy distribution for CCD22 is estimated using a raw frame from L+30. Raw frames do not go through any image processing, so they contain all energies observed by the detector in a four second integration period, including values below ET. A raw frame taken during Orbital B would be preferable due to similarity of light conditions, but no raw frames were downlinked for the entirety of the operation. L+30 operations were nadir pointed and had Bennu in the field of view, but the spacecraft was further away from Bennu than it was in Orbital B. However, given that no raw frames were collected during any other asteroid-facing observations, a raw frame from L+30 is the closest match. The bias map from the same period as the L+30 raw frame was also downlinked, so the bias value for each pixel is subtracted from the frame, resulting in a low energy distribution from L+30. Raw frames and bias maps are stored using a different ADU scale than event list data. To convert between frame ADU units and the event list

ADU units, the values in the frame are divided by two - raw frame and bias map energies are stored as 12 bits (with a maximum value of 4096), and event list energies are stored as 9 bits (with a maximum value of 2048). The full conversions between ADU values are explained in Appendix A. Figure 4-24 shows the raw frame energy histograms from L+30 for all functioning nodes, and Figure 4-25 shows the energy histogram for just CCD22 in event list ADU units.




Figure 4-24: Raw frame energy histograms for all nodes from L+30, shown in event list ADU. Each color depicts the energy distribution for the identified node.

The second step of the simulation is to estimate the low energy added by the light leak by fitting distributions to event list data from the two identified periods of Orbital B and the data from L+30. Energy histograms of the event list data for the specified periods of time are created using REXIS ground data pipeline tools. Figure 4-26 shows the energy distributions for the two Orbital B data periods. The August 2 data (from a period with significant iron count dropout) has a steeper slope near the low energy cutoff from ET and has more low energy counts than the July 26 data (from a period with more prominent iron count detection). This excess of low energy is consistent with a light leak. A best-fit exponential function is fit to each distribution.

In the third step of the simulation, the nominal low energy distribution and the




Figure 4-25: Raw frame energy histogram for CCD22 from L+30, shown in event list ADU.




Figure 4-26: Energy histograms of event list data from August 02 (light leak) and July 26 (no light leak)

fits from the event list data are combined to create an approximation of the total low energy distribution. The energy distribution found from L+30 is scaled according to the amount of time spanned by the data in the event energy histograms. The August

2 data are from a period of 3,018 seconds in length, which is about 765 four-second frames, so the raw frame is multiplied by that number of frames. The raw frame energy histogram scaled for time is shown in Figure 4-27.




Figure 4-27: Scaled raw frame energy histogram for comparison to event list data fit

Next, the scaled raw frame energy distribution and the event list energy distribution are combined by fitting the right side of a Gaussian distribution to the exponential fit from event list energy data. The relative shape of the combined fit is determined by an investigation of bias maps throughout Orbital B, and a comparison of the Orbital B bias maps to each other and to a bias map from an operation with less optical light in the field of view. Bias maps from three separate times within Orbital B are shown in Figure 4-28. As the gradient across the bias map image increases in energy at the readout edge, the associated bias map histogram distribution shifts upward, and the lower shoulder to the right of the peak of the distribution increases in width. For comparison, a bias map taken during the mask calibration is shown in Figure 4-29, when there were no bright sources of optical light within the field of view of the detectors. The overall energy range of the energy histogram for the mask calibration bias map is far lower than and of the Orbital B maps. The bias map image also shows no sign of a gradient. A gradient of events from the readout node that attenuates inward was shown in Section 4.4.2 to be consistent with energy from an optical light

leak onto the readout edge of the nodes. Since Orbital B maps show an increase in the width of the shoulder of the distribution and an overall increase in energy as the gradient of events becomes larger, the two effects (increase in width of shoulder, and increase in overall energy) are considered a possible effect of the light leak.




Figure 4-28: Three different bias maps for CCD22 recorded during Orbital B. The three energy histograms are shown on the left, and the images are shown on the right. X-axis is in units of raw frame ADU.

With the estimate for how light leak presents in the low energy distribution, the raw frame histogram and low energy event list histograms are fit using a Gaussian distribution to connect the scaled raw frame distribution to the exponential function fit to the event list data. The height of the Gaussian is roughly matched to follow the relative peak-to-shoulder amplitude ratio exhibited by the bias maps in figure 4-28. For instance, the amplitude difference between the peak of bias map C and its shoulder is about one order of magnitude. The amplitude of the Gaussian distribution between the raw frame and event list energy distributions is estimated by multiplying the difference between the peak and the shoulder by the number of frames present in the specified data. For the August 2 data, the difference between the peaks is kept to roughly 1.5 orders of magnitude less than the raw frame distribution peak. The





Figure 4-29: Bias map energy histogram for CCD22 from Mask Calibration. The energy histogram is on the left, and the image of the map is on the right. X-axis is in units of raw frame ADU.

right hand side of the Gaussian is then fit to the exponential function fit from the event list data energy distribution, such that the values are within roughly 0.05% of each other for two thirds of the height of the distribution. The fit from the August 2 event list data is shown in Figure 4-30a, and the combined distribution is shown in Figure 4-30b.

Step four of the simulation involves creating simulated frames. Now that the estimate for the combined low energy distribution is defined, the probability of each energy range hitting a single pixel on the detector is determined by calculating the relative area under the curve of the energy range as a percentage of total area under the curve. An array is created with 100 indices. A number that is calculated to occur three percent of the time is placed into three indices, a number that is calculated to occur one percent of the time is only placed in one index, and so on. A random number generator is used to select eight different integers between 0 and 99, and



(a) Nominal (in orange) and light leak function (green) from August 2, connected by Gaussian estimate (red) (b) Combination of nominal and light leak distributions for August 2.

Figure 4-30: Fitting process for 8-02 event list data to L+30 raw frame data energy distributions.

the energies from the indices within the array indicated by the randomly generated numbers are used to fill the eight “pixels” around the center “pixel”, while a grade 0 iron count is placed in the center, as depicted in Figure 4-31. The code for creating the miniature frames, “create_frame.py” (shown in Appendix A) creates 50 frames at a time. For one energy distribution, the code is run five times, resulting in 250 frames total. The output data are saved in the format of a simulated frame from the REXIS CCDs.

Randomly assigned energy	Randomly assigned energy	Randomly assigned energy
Randomly assigned energy	+Grade 0 Fe count	Randomly assigned energy
Randomly assigned energy	Randomly assigned energy	Randomly assigned energy

Figure 4-31: Visualization of event grid produced from grading sim

In step five of the simulation, the simulated frames are sent through a C application written by HCO researcher David Guevel. This application treats the input data as a

frame from the REXIS CCDs, and sends it through the REXIS flight software grading algorithm. As an input to the command for the application (“img_proc”), the ET value, ST value, and ULD value to be used in the image processing are entered. For all runs of the simulation, ET is set to about 0.3 keV (the value used in Orbital B), ST is nominally set to 50 eV, and ULD is set to the maximum value. The application outputs a simulated flight data file, which is then parsed using the REXIS ground data pipeline. Once parsed, a grade histogram is created from the simulated flight data. The predicted behavior is that for the larger light leak (August 2 data), more of the input grade 0 events will be artificially increased to grades 6 and 7 than for the lower light leak data (July 26 data). The grade histogram should consist of more grades 6 and 7 events than grade 0. The L+30 data should show few to no artificial increases in grade, as the L+30 operation downlinked the expected iron count rates. So the L+30 grade histogram should consist of mostly grade 0 events. When the same data are run through the C application using a doubled ST value, the histograms should show an increase in the retention of grade 0 events.

The results of the three tests are shown in Figures 4-32, 4-33, and 4-34. Figure 4-32 shows the results from running the L+30 energy distribution for CCD22 through the grading simulation. As expected, without any added light, the resulting grade events are mostly grade 0, with only a few being artificially upgraded to other events. This is possible because in Figure 4-27, the right hand side of the scaled distribution stretches over the area of event list ADU that correspond with energy values including the nominal ST setting of 50 eV (roughly 2) and the nominal ET setting of 0.3 keV (roughly 11) (Conversions found in Table A.2 in Appendix A).

Figure 4-33 shows the results from the July 26 data. More grade 0 events get artificially upgraded, especially to grades 6 and 7, but the majority of the events are still within the grade range downlinked by the instrument during Orbital B, and so would not have produced a dramatic loss of iron counts, as observed. During the period of time of the data cut used for the July 26 simulation, CCD22 recorded a grade 0 iron event rate of roughly 24 counts per kilosecond, which is close to the expected 21 counts per kilosecond.




Figure 4-32: Result of grading simulation for approximate distribution with no added light





Figure 4-33: Result of grading simulation for approximate distribution from July 26 data

Figure 4-34a shows the results of running the August 2 approximate distribution through the grading sim. This histogram shows a clear increase in the frequency at which a grade 0 iron event is artificially increased to a grade 6 or 7 event. During the period of time of the data cut used for the August 2 simulation, CCD22 recorded a grade 0 iron event rate of roughly 3.4 counts per kilosecond, which is well below the expected 21 counts per kilosecond. This result is consistent with the majority

of grade 0 iron events being artificially upgraded to grades 6 and 7, and not being downlinked. Figure 4-34b is the result of running the same frames through the image processing application using a doubled ST value. More grade 0 events are recovered with the doubled ST, which matches expectations, and tracks with the understanding that doubling ST helped recover iron events during Orbital R. All of the results of the grading simulation match expected behavior for additional low energy from a light leak causing loss of iron events.



(a) Result of grading simulation for approximate distribution from August 2 data

(b) Result of grading simulation for approximate distribution from August 2 data, with double ST

Figure 4-34: Result of August 2 approximate distribution with nominal and doubled ST.

4.4.4 Comparison of Orbital B to Orbital R

The results of the grading simulation show that increasing ST helps recover iron counts during periods with optical light on the edge of the detectors. This result matches what was observed in Orbital R; ST was doubled for CCD02, CCD20 and CCD22, and the iron counts downlinked from those nodes increased to closer to the expected values. With only a 5 percent difference in temperature between the two operations and a less than 0.03 percent difference in voltage, this change in iron events would not be expected (as explained above in Sections 4.3.1 and 4.3.2). The major difference between the Orbital R and Orbital B operations is that the ST was altered

for Orbital R. The ST was raised for the most affected nodes, which helped ignore more of the low energy noise caused by light on the detectors, and isolated more iron counts. This recovery of iron counts observed in Orbital R supports the hypothesis that the initial loss of iron counts was caused by a light leak.

Table 4.2: Summary of possible root causes for loss of iron calibration source counts

Potential Cause	Status	Evidence
ET/ST set incorrectly	Ruled out	All parameters checked; expected energy threshold observed in data.
Bad bias maps	Ruled out	Data from periods with bad bias maps were removed during ground processing.
CCD voltage drift	Ruled out	Voltages remained stable throughout flight.
CCD/MEB temperature change	Ruled out	Temperatures in Orbital B remained stable, and similar to other operations without loss of iron counts.
Light leak	Most likely	Iron count loss correlated with high event rates and amount of CXB versus asteroid exposure. Gradient in events similar to stray light ground simulation. Grading simulation produced similar behavior to flight.

From the analysis presented in this chapter, the most likely cause for the loss of iron counts during Orbital B is a light leak (summarized in Table 4.2). Temperatures and Voltages did not vary enough, nor at the time scale of the variations in iron events. Bad bias maps were removed from the data and not included in the final analysis of data, and all ST and ET were set to the expected values throughout flight. The events on CCD22 are more densely distributed closest to the readout edge of the node, and dissipate as the distance from the readout edge increases, similar to the Lincoln Laboratory test that first showed an edge light leak could be an issue. The iron events decrease with high total event rates, and increase as the nodes are exposed to more CXB versus asteroid surface. The grading simulation shows that

the addition of lower energy from a light leak causes grade 0 events to be upgraded to higher events, resulting in an overall decrease in the number of iron events that are recorded, as was observed during Orbital B. Doubling the ST values result in a recovery of the iron counts. The high event rates that were correlated with loss of iron are most likely due to a light leak on the REXIS CCDs. A light leak can also explain why the instrument event rates were sensitive to spacecraft pointing (like during cover opening), and why the energy offset was observed during the CXB calibration.

Chapter 5

CAST Analysis

This chapter presents the findings of a CAST analysis performed on the REXIS instrument focusing on the presence of a light leak on the detectors. In Section 4.4, a light leak was shown to be a probable cause of the loss of iron events during Orbital B. This CAST analysis was performed to search for the physical and organizational factors that caused the design of REXIS to allow light to interfere with X-ray event detection. As described in Section 2.3.2, the five main steps of a CAST analysis are to assemble the basic system information, create a safety control structure model, analyze each component in the loss, identify control structure systemic-level flaws, and to create an improvement program from the findings of the analysis.

This analysis was conducted using past REXIS documentation, as well as interviews with current and past REXIS team members. To conduct these interviews, one or two members from each group within the safety control structure (shown in Figure 5-3) were contacted, with the exception of the MIT Kavli Institute (MKI) (the most relevant person to speak with was not able to be contacted). Two members of the Science and Instrument Architecture team, two members of Project Management, one member of the Student Systems Engineering and I&T team, one member of the Student Detector team, and one member of the Lincoln Laboratory team were all interviewed. The interviews were held in person or over the phone over the course of four months. Through the interview process, past documents, tests, presentations, and communications were uncovered that gave insight as to why and when decisions

were made, and how a light leak was able to interfere with X-ray events in flight.

5.1 Part 1: Assemble Basic Information

Definition of system and boundary of the analysis: For the purposes of this CAST analysis, the system under study is defined as the REXIS detectors, DAM, and flight software, the REXIS Instrument Design team (which includes both engineers and scientists, and professors, research scientists, and students from MIT and Harvard), Lincoln Laboratory, and the MIT Kavli Institute for Astrophysics and Space Research.

Loss and hazardous state(s) that led to the loss: For the purposes of this analysis, the **loss** in the analyzed accident is defined as the loss of detector signal, specifically the loss of iron calibration data during Orbital B. The **hazard** that led to the loss is stray optical light reaching unexpected parts of the instrument.

System-level safety constraints required to prevent the hazard: The system-level safety constraint required to prevent the hazard is straight-forward: the REXIS detectors should be shielded from optical light within some defined limit on flux.

Events leading to the loss, and related questions: Review of past materials and presentations as well as interviews with past and current REXIS team members revealed the following rough timeline for how the light leak was possible given the instrument design. Questions generated by each event are listed below the associated event.

1. June 2014: Edge light leak discovered on CCDs from test at LL. Also discovered light leak through underside of support wafer of detector, and concern about pinholes remained high¹.
 - How did the team react? Was it viewed as a large concern? Was this registered as an official risk?

¹Information from Interview with LL employee and REXIS Instrument Monthly Report presentations

2. July-September 2014: Investigated and tested light leak mitigation methods, and created stray light simulation to quantify the edge light leak effect². Three suggestions were made to handle the light leak: 1. Apply black paint (Z307) to smooth edges, and then apply additional OBF to the sides of the detectors to mitigate the edge light leak. 2. Coat the underside of the detectors with additional aluminum OBF to mitigate light through underside of support wafer. And 3. Paint the portions of the DAM that could reflect light to the wire bond pads of the CCDs with black paint.
 - Which methods ended up being implemented? What were the rationales behind the decisions to implement or not implement?
3. November-December 2014: Applied, verified, and characterized both side coating (option 1 from the above item) and underside coating (option 2) to detectors. DAM painting (option 3) at this point was still being considered as an option.
 - Were there downsides to applying the paint on the DAM? Did verification of the side coating and underside coatings show a mitigation of light leak?
4. March 3, 2015: MKI scientist performed a light leak test on detectors with additional side coating and underside coating, to determine if DAM painting was needed. The results of the test showed that light leak was no longer an issue, so the employee concluded that DAM painting was not necessary. This information was relayed to the Science Team and Project Management, and the executive decision was made to not paint the DAM³.
 - What was the setup of the test? What light source was used? Did other team members work to corroborate this conclusion?
5. March 18, 2015: DAM assembled
 - Was there any testing performed after assembly that could have shown the light leak?
6. August 2015: Thermal Vacuum testing performed on REXIS at Applied Physics Laboratory. There were no light sources in the chamber, but there were windows

²Information from Interview with LL employee and REXIS Instrument Monthly Report presentations

³Information from interviews with both project management and both science team members

in the chamber where light could have entered⁴.

- Was the light leak observed during the test? If not, why?

7. September 2016: Launch of OSIRIS-REx

8. September 2018: Cover opening - first indication that REXIS CCDs were sensitive to stray light

- Where was the optical light coming from? How was the light reaching the detectors?

9. July 2019: Orbital B: Observation of loss of iron counts due to suspected light leak.

Analysis of the physical loss in terms of the physical equipment and controls included in the design to prevent this type of accident: The REXIS detectors have a directly-deposited OBF (described in Section 2.1.2) to prevent optical light from reaching the detector. Black paint and additional OBF was applied to the edges and additional OBF was applied to the underside of the detectors upon a ground test in 2014 that indicated optical light was making it to the detector surface⁵. Three separate paths that light could take to reach the detector surface were identified, and are shown in Figure 5-1: path 1 indicates light coming in through pinholes on the OBF, path 2 indicates light shining onto the readout edge of the detector, and path 3 is light shining through the support wafer and up through the bottom of the CCD. Based on the shape of the additional low energy distribution observed during Orbital B (outlined in Chapters 3 and 4), the light leak was most likely on the edge of the CCD. The wire bond areas on the sides of the detectors were not painted over due to concerns the paint could strip the wires, so light could reach the unprotected areas of the detectors through that path (as described in Section 4.4).

Analysis of the physical loss in terms of failures and unsafe interactions leading to the hazard, missing or inadequate physical controls that may have prevented the accident, and any contextual factors that influenced the event: None of the physical components of REXIS failed; the CCDs correctly

⁴From interview with student systems engineer

⁵Interview with Lincoln Lab scientist




Figure 5-1: Three identified paths for light to reach detector surface. (1) Light leak through the detector surface via pinholes created by particles. (2) Light leak through the edge of the detector. (3) Light leak up through the support wafer. [13]

registered the energy that was deposited on the detector surface, and by looking at event list data, there was not an abundance of hot pixels caused by pinholes which means the OBF blocked most optical light from reaching the detector surface in places where the OBF was applied. There were no unsafe interactions that led to the hazard. Upon identification of the most likely cause as an edge light leak, an additional physical control that may have prevented the event was identified: using a light-absorbing black paint on the reflective parts of the bottom of the DAM. Painting the DAM was suggested during the design phase, but never implemented.

5.2 Part 2: Model the Safety Control Structure

The high-level safety control structure for the system under review is shown in Figure 5-2. The Instrument Design Team is broken down into subteams in Figure 5-3.

Figure 5-3 is the structure used for the remainder of this CAST analysis. The Instrument Design team is broken up into four subgroups: the Project Management, the Science and Instrument Architecture Team, the Student Systems Engineering and I&T Team, and the Student Detector Team. All four subgroups have safety-related responsibilities with respect to each other, but the Project Management team is the

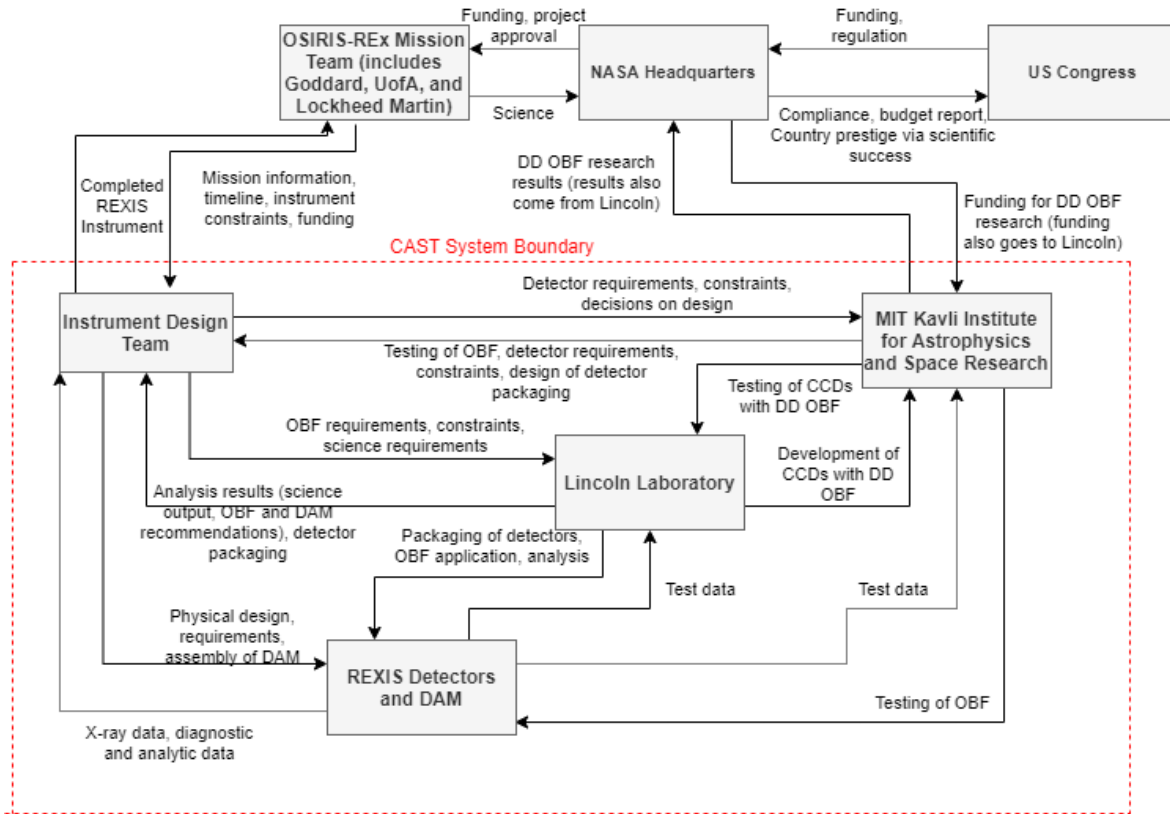


Figure 5-2: Higher-level REXIS safety control structure. The dotted line indicates the boundary of the system under study in this analysis.

group with the highest amount of control and communication with outside collaborators. The two groups outside of the Instrument Design team are the MIT Kavli Institute for Astrophysics and Space Research (MKI) and the MIT Lincoln Laboratory. These two groups worked on development of the DD-OBF, and communicated with the Project Management. Finally, the REXIS Detectors, DAM, and flight software are at the bottom of the structure. The types of interactions are divided into exchange of knowledge or ideas (depicted by the dotted green arrows), and a physical interaction (dashed blue arrow). Physical interactions only occur between groups and the physical system, the REXIS Detectors.

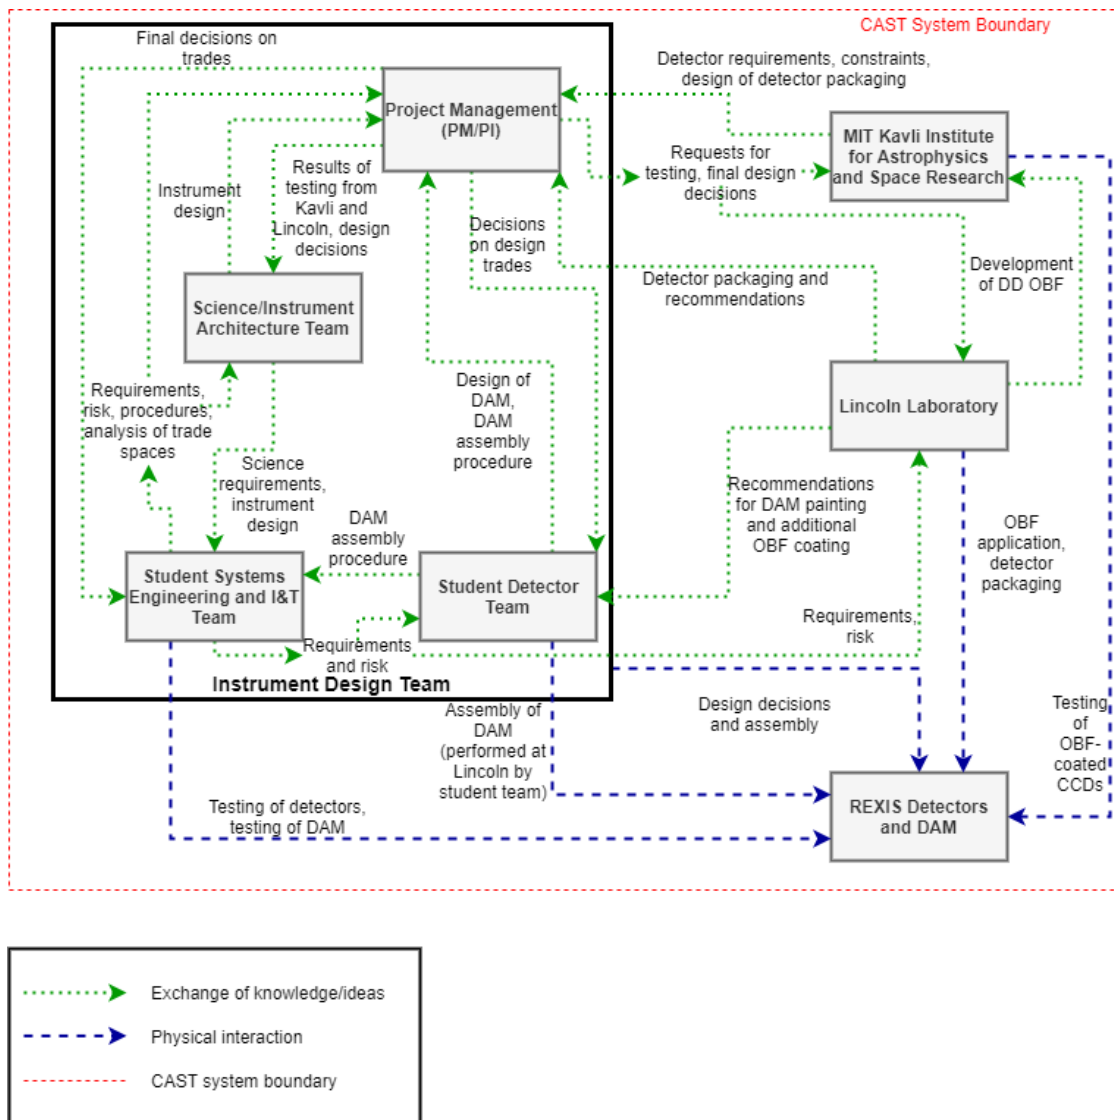


Figure 5-3: Zoomed-in REXIS safety control structure

5.3 Part 3: Analyze the Loss at the Component Level

In this section, starting from the bottom of the control structure, each component's behavior is examined in the context of how it helped contribute to the loss. The rationale behind each action or behavior is also documented, to provide further context and more information to create better recommendations for future projects. For each component, the following information is presented: the safety-related responsibilities

from Part 2, the contribution to the hazardous state (if any), flaws in the process model or mental model (thought process) that contributed to the behavior of the component, and contextual factors that explain the behavior of the component. A summary of each component's behavior, and recommendations on the safety control structure component level are also presented at the end of each subsection. These recommendations are also summarized at the end of this chapter, and in Chapter 6.

5.3.1 REXIS Detectors and DAM and FSW

Relevant safety-related responsibilities

- Collect X-ray data and output candidate X-ray event list
- (OBF) Prevent optical light from reaching the detector surface
- (DAM) Support the detector assembly
- (Flexprints) Connect CCDs to detector electronics.

Contribution to hazardous state

The hazardous state is reached when optical light reflects off the parts of the DAM that have a line of sight to the wire bond pads (where the CCDs are connected to the detector electronics) and serial register area of the CCDs. The detectors measure the optical light as a gradient of low energy events, and the FSW combines these low energy events with higher energy iron events, artificially increasing the events to grades 6 and 7, which the FSW then ignores, as commanded. The detectors and DAM, along with the FSW, caused the loss of iron events by throwing out the data from most of the iron counts during periods where light is shining on the detectors.

Flaws in the process model contributing to the actions

The process model of the REXIS Detectors and DAM contains no flaws; while the output was not desired, all parts of the system acted as expected given the inputs the system received. However, there were insufficient controls within the structure to prevent the hazard - the DAM was not painted with the light absorbing paint.

Contextual factors explaining the behavior

The flight software is written to account for various sources of constant noise (using bias maps, subtracting over- and underclocks, and utilizing a HPM to ignore hot pixels for an entire data taking period). The optical light leak resulted in varying noise on the detectors at random intervals, which the FSW was not designed to remove.

Summary of behavior with context

The detector assembly behaved as expected, given the design and software algorithms and the specific input. None of the parts of the assembly behaved in a manner counter to their original design; rather there was an unanticipated interaction of light within the design that resulted in the unwanted loss of iron counts. The DAM had reflective surfaces that directed optical light to the serial registers of the CCDs. The light reflected onto the readout edges of the nodes caused the exposed readout portion of the detectors to register more low energy events, and grade them as part of real X-ray events, increasing the grade of those X-ray events to grades 6 and 7, which were then not saved on the spacecraft.

Recommendations

An improvement on the design of the detector assembly would be to paint the parts of the DAM that had the potential to reflect light onto the wire bond pad area of the CCDs. A more broad recommendation is for instruments that utilize detectors with sensitivity to certain wavelengths, light mitigation methods should be thoroughly investigated and tested.

5.3.2 Student Detector Team

Relevant safety-related responsibilities

- Design DAM, packaging, and flexprints.
- Communicate designs to Lincoln Lab as needed.
- Create and document procedure for DAM assembly.

- Assemble DAM at Lincoln Lab facilities.

Contribution to hazardous state

The student detector team designed the DAM, and assembled it at Lincoln Lab facilities without applying any black paint to the DAM.

Flaws in the mental model contributing to the actions

The DAM was not designed to prevent optical light from reflecting onto the detector surfaces.

Contextual factors explaining the behavior

The student team was one of the lower components in the project hierarchy, and management had the final say as to whether or not to paint the DAM. Optical light was not considered in the design of the DAM. The student detector team was following instructions correctly by not applying the paint to the DAM. The student team did go so far as to draw up schematics of where on the DAM to apply the paint (shown in Appendix B). However, the project management team decided to forego painting the DAM. Additionally, at the time the light leak was discovered, it was too far into the design process to make edits to the overall DAM design.

At the time of the black paint suggestion, light leaking through the readout edges of the CCDs was not viewed as severe of a possible threat to data as pinholes through the main layer of OBF were⁶. Preliminary tests on spare CCDs at Lincoln Laboratory had shown far more pinholes than anticipated, which could have severely impacted the amount of usable science data collected, and would have contaminated the X-ray data even more than an edge light leak. Therefore, once the decision to not paint the DAM was made, other concerns and risks were given more attention, and the edge light leak was considered mitigated.

⁶Interview with student detector team member

Summary of behavior with context

The student detector team was in charge of designing and constructing the DAM. Members of the team also had many other responsibilities not pertaining to the DAM. LL and Kavli employees worked more closely with the design and testing of the CCDs with DD OBF than any of the students did⁷.

Recommendations

It would be useful to have a student whose job it was to assist more closely with the OBF and CCDs, or at least become the student subject expert, rather than having multiple students working on the subsystem along with other parts of the instrument design.

5.3.3 Student Systems Engineering and I&T Team

Relevant safety-related responsibilities

- Document system, subsystem, and science requirements.
- Keep track of risk.
- Perform trade analyses.
- Document and track procedures for assembly and testing.
- Perform testing on detector assembly during integration.

Contribution to hazardous state

There was no initial requirement specifically mentioning light leak. There were two requirements based on OBF performance - the first requirement focused on the flux of photons on the imaging surface only, and the second requirement was on the percentage of pinholes on the imaging surface (both requirements are listed in Appendix B).

A light leak through the side of the detectors was recorded in the overall risk matrix for the instrument, with likelihood three and consequence four (on a scale of

⁷Interview with student detector team member

one to five). Upon application of additional side and underside coating, the risk for a light leak was reduced to likelihood one, and marked as low criticality⁸. Additionally, throughout environmental testing of the detector assembly, an edge light leak was not registered as an issue.

Flaws in the mental model contributing to the actions

The team assumed that edge light leak was not a likely occurrence, and recorded the risk criticality as low. There was no requirement written specifically pertaining to light leak.

Contextual factors explaining the behavior

Since this specific type of effect due to light leak was not encapsulated in a requirement, the urgency in responding to it was lower than other issues that had a more direct link to instrument requirements.

During the time the various paths for light leak were brought into the team's discussion, many other risks came up as well, some with more severe or more immediate potential harm. For example, the aluminum tee's used in the detector structure were repeatedly running into thermal bonding issues, and coming out of place⁹. If the tee bonding issue was not resolved, the instrument would potentially have not made it onto the spacecraft. That was a far more immediate issue than being concerned about an edge light leak in the detectors during flight. The team was also running into schedule concerns with the main electronics board, so that subsystem took higher precedence at the time as well.

There was a test performed by an MKI scientist on a single CCD in ambient light that indicated an edge light leak was no longer an issue after the first two mitigation methods were implemented, so the team as a whole focused attention on other pressing issues and risks. The edge light leak risk was not marked as closed, but it was reduced in likelihood, bringing the criticality down from medium to low.

⁸From risk matrix documentation

⁹REXIS Instrument Monthly Report presentation slides

Environmental testing of the instrument did not reveal the edge light leak, because most tests were either done with the cover closed, or when there were no sources of optical light in the field of view of the instrument¹⁰. Additionally, most testing was performed before the flight software was complete. There was not time nor opportunity to perform an instrument-level light leak test. Due to the instrument status as a “Do-no-harm” instrument aboard a Class B mission, most of the testing was performed on interfaces between REXIS and the spacecraft, rather than on REXIS science performance.

There were two tests where the light leak could potentially have been discovered¹¹. In late August 2015, an instrument level performance test in a flight-like thermal environment was performed at the Johns Hopkins Applied Physics Laboratory. The vacuum chamber used for the test had windows that could have let in optical light; however during this test, the REXIS detectors were registering an overwhelming amount of noise due to issues with the electronics. REXIS was then taken apart, and the electronics were altered to reduce the noise. Due to this, the instrument almost did not make it onto the spacecraft. The other test with possible light was the spacecraft-level thermal vacuum test in February and March of 2016, where there were sun bulbs in the chamber. However, the cover of REXIS was closed for the entire test. With the way the REXIS test procedure was scheduled, the instrument would not have been collecting data while the sun bulbs were shining on REXIS.

Summary of behavior with context

At the time the potential for an edge light leak was brought to the attention of the team, the project was suffering from schedule crunch and other risks with more immediate impact than the potential edge light leak. A test indicated that edge light leak was not a significant concern, and there was no other test performed on the instrument before OSIRIS-REx launched that could have revealed the edge light leak.

¹⁰Interview with student systems engineer

¹¹Interview with student systems engineer

Recommendations

When designing a light-sensitive instrument, a requirement(s) should be written with direct language about how much light flux is acceptable, so other subsystems can consider optical light in the subsystem design.

5.3.4 Science/Instrument Architecture Team

Relevant safety-related responsibilities

- Communicate instrument design requirements and science requirements to project management and student team.
- Provide input for trade space decisions.
- Provide input to plan tests to verify requirements.

Contribution to hazardous state

The science team did not provide the systems engineering team with a light-leak specific requirement. They also accepted the results of the test performed by the Kavli team that indicated no DAM painting was necessary.

Flaws in the mental model contributing to the actions

There was very little information about the test communicated to the team, such as the light source used, or the exact physical setup, but the science team looked at the data from the test, and agreed with the conclusion that painting the DAM was not necessary.

Contextual factors explaining the behavior

Like the rest of the project, the science team was under schedule pressure. There were multiple issues that arose around this time period (discussed in the Systems Engineering subteam section), so the mentality of trying to fix or cross out any issue

possible could have been a strong influencing factor. The stray light simulation showed some improvement by painting the DAM, but it didn't look like a huge improvement¹².

Summary of behavior with context

The science team agreed that the DAM did not need to be painted, due to the test performed by Kavli. The agreement with the results, even though the test was not fully documented, was due to the amount of other risks and schedule pressure at the time¹³.

Recommendations

The criticality of the risk for edge light leak was recorded as medium (and then got reduced to low upon implementation of the application of side and underside coating), but the consequence remained at a four on a scale from one to five. For items with consequences of four or five, there should be thorough documentation of tests pertaining to the risk, and any decision to downgrade a risk should be carefully thoroughly vetted and reviewed by all relevant team members. Also, when first designing an instrument, the team should thoroughly research previous similar instruments and focus on looking for issues experienced by those instruments. The team might have determined an optical light leak was a large enough issue to merit writing a light-leak specific requirement, and determining the amount of testing necessary to ensure a more optical-light resistant system.

5.3.5 Lincoln Laboratory

Relevant safety-related responsibilities

- Provide selection of detectors for REXIS instrument design team to choose from.
- Develop DD OBF method and apply to REXIS CCDs.
- Analyze optical light effect on DD OBF-coated CCDs.

¹²Interview with science and instrument architecture team members

¹³Interview with science and instrument architecture team members

- Suggest and analyze various light mitigation techniques.
- Package REXIS detectors, oversee student assembly of DAM.

Contribution to hazardous state

Lincoln Lab applied the DD OBF to the REXIS detectors, performed stray light testing and simulations, and reviewed the design of the DAM.

Contextual factors explaining the behavior

The DD OBF was initially planned to be developed before REXIS began integration, but a lapse in funding caused the schedule to be delayed, resulting in the DD OBF being developed in parallel with REXIS¹⁴. Other OBF methods were suggested to REXIS project management, but due to system complexity and difficulty of testing a standalone OBF, the management decided to go forward with the DD OBF method¹⁵. Therefore, the issues such as pinholes, underside light leak, and edge light leak were discovered as REXIS was entering the integration and test period, rather than far before when mitigation methods could have been more thoroughly researched.

Other mitigation methods were also researched, such as inserting a mechanical structure to block light from certain areas, but found to be infeasible or too expensive. Another suggestion was to directly paint the wire bond pads with the black paint. The reason this option was not implemented was concern that the paint may have eventually worn through the protective layering around the wires and caused a short in the CCDs, which would have resulted in loss of all data¹⁶.

Summary of behavior with context

Lincoln Laboratory developed the directly deposited optical-block filter for REXIS, ending up in tandem with the REXIS instrument being designed and built¹⁷. Funding

¹⁴Interview with LL scientist

¹⁵Interview with LL scientist and interview with project management

¹⁶Interview with LL scientist

¹⁷Interview with LL scientist

for the project lapsed, causing the schedule slip back to closer to the REXIS timeline. Therefore, the mitigation methods for light leak were developed close to when REXIS was being assembled, rather than farther in advance when the methods could have been accounted for in the schedule. Lincoln Laboratory created the stray light simulation that indicated edge light leak was an issue, and that black paint on the DAM could help mitigate the issue, and presented this information to the project management.

Recommendations

Schedules slipping is a common occurrence, especially in space missions, so making recommendations such as “plan for more time” are unhelpful. Given that other optical light blocking methods were more expensive or time consuming, the REXIS Project management made the best choice possible with the given information for which optical blocking filter method to use, and Lincoln Laboratory made the best recommendations they could, as well. Perhaps a long-term environmental test could be performed by painting the wire bond pads as suggested, and observing the effect of the paint over time, to see if that could be a viable option in the future. Another recommendation for project development would be to have different options as backup for technologies that are being developed in parallel with the main project, in case the new technology proves infeasible or is significantly delayed.

5.3.6 MIT Kavli Institute for Astrophysics and Space Research

Relevant safety-related responsibilities

- In collaboration with Lincoln Lab, test detectors with DD OBF applied.
- Test optical light mitigation techniques suggested by Lincoln.
- Provide recommendation to REXIS Project Management about necessity of the various optical light mitigation methods.

Contribution to hazardous state

MKI performed a test on a CCD with additional side coating and underside coating (mitigation methods to reduce edge and underside leak). The results of the test indicated that light leak through the edge of the CCDs was not an issue. The test was not well-documented, however; the only information available from it was provided in an email to project management (shown in Appendix B).

Flaws in the mental model contributing to the actions

The test was not fully documented, so important information about the geometry of the setup, and the light source used were not made available to the team¹⁸. Those factors could have influenced the final decision on whether or not to paint the DAM.

Contextual factors explaining the behavior

The MKI team members were also working on other projects at the same time as the REXIS project. The REXIS use of DD OBF was one application of the overall project for MKI.

Summary of behavior with context

The MKI team performed a test on a CCD, which indicated that an edge light leak was no longer an issue after some additional OBF and paint was added to the sides, and some OBF was added to the underside of the support wafer. This test was not thoroughly documented, but after it was sent out to the REXIS project management, that subteam accepted the test results and made their decision to not paint the DAM based on the test.

Recommendations

There should be strong enforcement of careful and thorough documentation for tests, especially when the results are used to make decisions for a flight instrument, and the

¹⁸Interviews with science and architecture team members and project management

issue could affect flight data. Additionally, there should be team input and agreement with a test setup before the test occurs.

5.3.7 Project Management

Relevant safety-related responsibilities

- Make final decisions for design and trade spaces, specifically the mitigation methods suggested for optical light on the CCDs.
- Communicate test information from Kavli and Lincoln to Science team and Student team.
- Track budget for instrument and schedule for instrument delivery to OSIRIS-REx team.

Contribution to hazardous state

Project management had the final say in design and trade space decisions. After receiving the information about the CCD light leak test from MKI, management sent out the information to the science team as well. Upon receiving affirmation from the science team, all members of project management agreed to not paint the DAM¹⁹.

Flaws in the mental model contributing to the actions

The decision to not paint the DAM was made assuming the test was an accurate representation of the type of light seen during flight, and that it would have revealed an edge light leak if there were one. The decision was also made without consulting all involved parties; Lincoln Laboratory team members were not involved in the final decision to not paint the DAM²⁰.

¹⁹Email chain (in Appendix B)

²⁰Email chain (in Appendix B)

Contextual factors explaining the behavior

As stated in other sections, at this point of the project there were many other risks and concerns, as well as some worry about schedule overrun. The tee bonding issue was the most pressing immediate concern, because if it didn't get resolved, the team would be delayed in delivering REXIS to the spacecraft²¹. If the instrument was delayed enough, since it was not necessary for mission success, there was a mass model that would have been installed on the spacecraft instead. The main electronics board progress was also a major cause of stress, and was one of the larger threats to the schedule as well. The paint chosen for the DAM had a curing time of about a week, which would have delayed the assembly of the instrument by at least a week, if no other unforeseen issues arose²².

Finally, one member of the management team was also majorly involved in another space-based mission that reached its main science phase around the time the light leak was first introduced as an issue, and was offsite for most of fall 2014 and spring 2015²³.

Summary of behavior with context

Project management was given a test that showed favorable results, and was able to cross off one item in a list of exponentially growing concerns in a time of schedule overrun, so it is understandable that the positive news was readily welcomed. Members of the science team did their own quick analysis of the test results and came to the same conclusion, so there did not seem to be a reason to question the validity of the test.

Because the REXIS instrument is a “Do-no-harm” instrument aboard a Class B mission, there was much pressure on the project team to balance schedule with risk. If the instrument fell behind on schedule, the spacecraft would not delay any milestones to wait for REXIS; there was a mass model ready to be placed on the spacecraft

²¹Instrument Monthly Report presentations and Risk Matrix documentation

²²Interview with student detector team member

²³Interview with project management

instead. Therefore, project management often had to make difficult decisions about what was “good enough” in favor of getting the instrument on board the spacecraft²⁴.

Recommendations

Again, it is difficult to plan around unknowns, like which parts of a project will produce the most risk or cause for concern, or when schedule overrun will occur. However, when the consequence of a risk is a four or five, perhaps there should be additional checks performed on tests pertaining to that risk, and additional documentation required on tests pertaining to the risk. Another recommendation is at the very beginning of the project, the team should carefully construct an exhaustive list of all tests necessary. These tests should all map directly to requirements.

5.4 Part 4: Analyze the Loss at the Control System Level

In this section, system-level interactions and factors are considered. The categories covered in this section can include communication and coordination, the overall safety culture of the project, design of the safety management system, changes and dynamics over time, and internal and external economic and related factors. For this analysis, potential issues were identified in only the communication and coordination, changes and dynamics, and internal and external factors categories. Notes about each category are detailed below.

Communication and coordination There were many loops of communication within the REXIS team, as shown in the safety control structures in Figures 5-2 and 5-3. This ensured quicker communication within the subgroups of the team, but sometimes rather than expediting communication, there were bottlenecks. Specifically, the communication between Kavli and the entire Instrument Design team about the light leak test was funneled through one member of the project management team, who

²⁴Interview with project management

then passed on the information to the rest of the team²⁵. The test results were also not sent to LL for comment or review.

Changes and dynamics over time Since REXIS was primarily a student-based instrument, the greatest change over time was which students were working on a subsystem. For instance, one of the main students who worked on the DAM graduated and left after the DAM was assembled and before environmental testing was performed. Often undergraduate student interns would work on the project for only one or two semesters. With the constant turnover and change in workforce, the transfer of ideas and information was not entirely smooth, with nuanced knowledge of a subsystem leaving as the students who worked on it left. To ensure a smoother transition, careful documentation should be kept throughout the design process, and individual students should provide their notes as project heritage for future project members to reference. The REXIS team did a fair job in documenting most presentations and notes in their online repository, but individual contributions were sometimes harder to find.

Internal and external economic and related factors The grant funding the DD OBF research at LL and Kavli lapsed, which caused development of the DD OBF to occur in parallel with the development of REXIS, rather than the initial plan of having the technology more mature at the time for implementation in REXIS.

The REXIS team also was under schedule pressure to deliver the instrument to the spacecraft by the fall of 2015 so it could undergo environmental testing with the assembled spacecraft. If the instrument was not delivered, there was a mass model that would be flown on the spacecraft instead, so there was not a chance of delaying the schedule of the OSIRIS-REx mission for the REXIS instrument alone.

REXIS is defined as a “Do-no-harm” instrument, which means it was allocated a smaller budget and fewer resources, and is considered a low-cost, higher-risk instrument[7]. For “Do-no-harm” instruments or missions, cost and schedule risks are considered as having equal if not greater importance than technical risks. So by design, these instrument teams will usually accept greater risks to instrument performance in favor

²⁵Interview with project management

of keeping within the cost and schedule budgets. With that knowledge, it is understandable that the instrument team was more concerned about the instrument being assembled in time, rather than spending resources on an issue they believed was resolved²⁶.

Due to the REXIS instrument’s “Do-no-harm” definition, most of the resources were spent on testing the interfaces between REXIS and the spacecraft. There was far more emphasis placed on ensuring that REXIS would not harm the OSIRIS-REx spacecraft, and very few resources left to test the science output of REXIS²⁷.

5.4.1 Summary of questions and answers

The initial questions from Section 5.1 to be answered throughout this analysis were: *How did the team react (to the discovery of the light leak)? Was (the light leak) viewed as a large concern? Was this registered as an official risk?*

The team responded by investigating different light leak mitigation methods. It was registered as a risk with a relatively high consequence (four), so the team understood the consequences a light leak could have on data.

Which (light leak mitigation) methods ended up being implemented? What were the rationales behind the decisions to implement or not implement? What were the downsides of applying the paint on the DAM? Did verification of the side coating and underside coatings show a mitigation of light leak?

The additional coating of the sides and undersides of the CCDs was implemented, and were shown to mitigate the optical light leak, but the DAM was not painted due to schedule concern and a test that indicated the painting was not necessary.

What was the setup of the test (that indicated DAM painting was not necessary)? What light source was used? Did other team members work to corroborate this conclusion?

The complete setup of the test is unknown, as it was not fully documented. (Or if it was, the notes were not shared with the rest of the team). It is unknown which light

²⁶Interview with project management

²⁷Interview with project management

source was used. Members of the science team analyzed the data from the test and came to the same conclusion.

Was there any testing performed after assembly that could have shown the light leak? Was the light leak observed during the (environmental) tests (performed at APL and Lockheed)? If not, why?

None of the tests performed on the instrument after full assembly could have shown the light leak, due to either test setup (cover was closed, no light in chamber), or situational causes (sun bulbs kept going out, REXIS detectors had overwhelming electrical noise present).

Where was the optical light coming from? How was the light reaching the detectors?

There was not a complete stray light model of the entire spacecraft with the REXIS radiation cover opened, so it is difficult to say precisely where the optical light came from. Comparisons between predicted light curve from the asteroid and total event rates from L+30 and Orbital B show some correlation, but there is a more complicated relationship at play. The optical light most likely came through the reflective portions of the DAM and shone onto the readout edge of the nodes.

5.5 Part 5: General Recommendations

Many of the issues experienced through the design, integration, and testing phases of REXIS were due to schedule or budget restrictions, an issue that is nigh unavoidable in space missions, and especially in instruments designed to accept more risk. However, some recommendations for future similar projects were made throughout this analysis, and are summarized below:

- For instruments that utilize detectors with sensitivity to certain wavelengths, light mitigation methods should be thoroughly investigated and tested.
- Parts of the instrument that have the potential to reflect light on sensitive portions of the detectors should be painted with light-absorbing paint, or other methods should be utilized.

- When designing a light-sensitive instrument, a requirement(s) should be written with direct language about how much light flux is acceptable, so other subsystems can consider optical light in the subsystem design.
- There should be a person on instrument design team with the job of assisting closely with or in close communication with the team working on new or developing technology.
- The design team should have different options as backup for technologies that are being developed in parallel with the main project, in case the new technology proves infeasible or is significantly delayed.
- For projects using new technologies, additional resources should be allocated in the budget to sufficiently vet those new technologies.
- For items with consequences of four or five, there should be thorough documentation of tests pertaining to the risk, and any decision to downgrade a risk should be carefully thoroughly vetted and reviewed by all relevant team members.
- When first designing an instrument, the team should thoroughly research previous similar instruments and focus on looking for issues experienced by those instruments.
- There should be enforcement of careful and thorough documentation for tests, especially when the results are used to make decisions for a flight instrument.
- For tests pertaining to requirements, there should be team input and agreement with a test setup before the test occurs.
- At the beginning of a project, the team should carefully construct an exhaustive list of all tests necessary. These tests should all map directly to requirements.
- For projects with frequent transitions of team members, ensure individual knowledge is captured in documentation to ease knowledge transfer to new members.

Chapter 6

Recommendations and Conclusions

6.1 Summary of Root Cause Analysis

After going through each bone on the fishbone diagram compiled by the REXIS operations team, the only potential root cause that did not get ruled out was a light leak onto the detectors. High event rates on the detector were shown to have a positive correlation with the decrease of recorded iron counts. An attempt was made to connect optical light directly to the loss of iron counts, but the overall picture is more complicated than a simple one-to-one comparison. There seemed to be a connection between high event rates and optical brightness on the CCDs, but the data did not show a clear trend in increased optical brightness and decreased iron event rates. However, a comparison between the portion of the node exposed to CXB and iron rates showed a positive correlation. The more a node is exposed to the CXB, the higher the iron count rates get. Therefore, as the bright asteroid surface leaves the field of view of the node, the iron counts recover.

After comparing the event gradient on CCD22 to a stray light simulation performed by Lincoln Laboratory in 2014-2015, it was determined that the physical distribution of events on the node appeared similar to the physical distribution of energy absorbed onto the node from an edge light leak.

Finally, a simple simulation of the detector was performed using flight data from L+30 calibrations and Orbital B to determine if the suspected mechanism by which

the iron counts were disappearing was plausible. The results from using varying amounts of additional low energy show that if excess low energy is present (from a suspected edge light leak), and falls above ST or even ET, grade 0 iron events are artificially promoted to higher grades. The more low energy is added, the more counts are promoted to grades 6 and 7, which were then not downlinked by the instrument during Orbital B due to concerns that high enough event rates would break the internal count rate filter, and downlink significantly larger data volume than predicted, with less usable data. The results of the grading simulator analysis show that it is plausible that an optical light leak could lead to loss of iron counts.

6.1.1 Limitations

Without actual raw frame data from Orbital B, estimates were made as to what the low energy distribution seen by the detectors looked like. Had there been a more accurate low energy estimate, the effect of added energy by presumed optical light could have been more thoroughly explored. Additionally, a precise stray light model for REXIS with the radiation cover open in relation to the rest of the instruments on board OSIRIS-REx was not available, nor was it feasible to create one. The optical brightness calculation also did not account for the rough surface of the asteroid. Therefore, there could be subtle interactions of optical light with the surface of the asteroid, the other instruments onboard the spacecraft, or the inside of the REXIS tower that change the optical brightness or angle at which light enters the REXIS instrument that could help better explain the exact relationship between optical light and high event rates.

6.1.2 Future Work

At the time of this thesis, lab space was not available to perform ground testing on the REXIS spare flight hardware. An interesting and potentially useful experiment that could be performed would be to put optical light sources within a thermal vacuum chamber, and observe the effect on the data. The simple grading simulation could also

be expanded to create full 1024 by 4096 pixel frames with varying additional energy to further investigate the effects of optical light on overall event recording. The fuller model could imitate a flight node, with only a small area of the detector surface registering the iron calibration source counts, and the rest of the node capturing X-ray data from other incoming photons. By recreating the gradient of added low energy, the effect on overall data collection, not just iron counts, could be studied more in depth. Estimating the additional noise as a Gaussian was also a limiting factor; if there was a measurement of the energy distribution added by optical light, that distribution could instead be run through the grading simulation and the effects on event list data observed.

6.2 Summary of CAST Recommendations

After analyzing the REXIS team infrastructure, the following recommendations were made for future instruments:

- For instruments that utilize detectors with sensitivity to certain wavelengths, light mitigation methods should be thoroughly investigated and tested.
- Parts of the instrument that have the potential to reflect light on sensitive portions of the detectors should be painted with light-absorbing paint, or other methods should be utilized.
- When designing a light-sensitive instrument, a requirement(s) should be written with direct language about how much light flux is acceptable, so other subsystems can consider optical light in the subsystem design.
- There should be a person on instrument design team with the job of assisting closely with or in close communication with the team working on new or developing technology.
- The design team should have different options as backup for technologies that are being developed in parallel with the main project, in case the new technology proves infeasible or is significantly delayed.
- For projects using new technologies, additional resources should be allocated in

the budget to sufficiently vet those new technologies.

- For items with consequences of four or five, there should be thorough documentation of tests pertaining to the risk, and any decision to downgrade a risk should be carefully thoroughly vetted and reviewed by all relevant team members.
- When first designing an instrument, the team should thoroughly research previous similar instruments and focus on looking for issues experienced by those instruments.
- There should be enforcement of careful and thorough documentation for tests, especially when the results are used to make decisions for a flight instrument.
- For tests pertaining to requirements, there should be team input and agreement with a test setup before the test occurs.
- At the beginning of a project, the team should carefully construct an exhaustive list of all tests necessary. These tests should all map directly to requirements.
- For projects with frequent transitions of team members, ensure individual knowledge is captured in documentation to ease knowledge transfer to new members.

6.2.1 Limitations

Because this investigation occurred five years after the main time period under study, some information was unavailable, and some people were not able to be contacted. The investigation was also limited in scope, focusing solely on the behavior suspected to have caused the iron loss in Orbital B, an edge light leak. A more thorough analysis of the project may uncover further recommendations.

6.2.2 Future Work

As with any CAST analysis, the next step is to present the recommendations to the team, and check back in with the team or other projects in the future to see what recommendations were implemented. From that, the efficacy of the recommendations could be studied, possibly leading to further or amended recommendations. The ultimate goal of using the CAST method for an accident analysis is to hopefully

help improve the system through which space instruments are created, leading to progressively safer and better instruments.

Appendix A

REXIS Image Processing Information

Table A.1: Energy ranges for ET, ST, and ULD

Parameter	Maximum value	Default value	Multiplication factor
ET	65536	500	1x
ST	65536	50	1x
ULD	256	100	256x

Table A.2: Different ADU values for a range of input energies. Actual values for different nodes were discovered to vary, but this provides a baseline.

keV	Raw ADU	Event list ADU	Raw frame/bias map ADU	ULD ADU
42	65536	2048	4096	256
10	15254	476	953	60
9	13714	429	857	54
8	12203	381	762	48
5.9	9000	281	562	35
4.0	6102	191	381	24
3.5	5339	167	333	21
3.0	4576	143	286	18
2.0	3051	95	191	12
0.33	500	15	31	n/a
0.03	50	1	3	n/a

Table A.3: CCD ET parameters used throughout flight, with changes in bold text, and explanation provided for each change. The values for the 10 nominal nodes are included, even though from L+30 onward, only five or six nodes were used

CCD Image Processing Parameters		
Operation	ET (raw ADU)	Rationale
Internal Calibration	966, 966, 966, 1237, 2542, 1247, 1983, 1965, 1319, 966	ET were set to around 0.6 keV for each node.
Cover Opening	No change	n/a
CXB Calibration	966, 966, 966, 1104 , 2542, 1144 , 1983, 1197 , 1319, 966	Wanted to have lower ET to capture more of the low energy flux from the CXB.
Crab Calibration, part 1	No change	n/a
L+30 Calibration	No change	The image processing parameters did not change, but the NEM did - only 6 nodes were used.
Crab Calibration, part 2	966, 966, 966, 1104, 2542, 1913 , 1983, 1197, 1319, 966	Did not need to capture the lower energy flux from the Crab Nebula, and were concerned about reaching the event rate limit.
OBF Verification set 1	966, 966, 966, 1916 , 2542, 1913, 1983, 1197, 1319, 966	Testing different lower ET settings for better performing nodes.
OBF Verification set 2	966, 925 , 966, 1916, 2542, 1913, 1983, 1197, 1319, 966	Lower ET for CCD02
OBF Verification set 3	966, 887 , 966, 1916, 2542, 1913, 1983, 1197, 1319, 966	Lower ET for CCD02
OBF Verification set 4	966, 966, 966, 1916, 2542, 1913, 1983, 1155 , 1319, 912	Lower ET for CCD20 and CCD22

Continuation of Table A.3		
OBF Verification set 5	966, 966, 966, 1916, 2542, 1913, 1983, 1133 , 1319, 868	Lower ET for CCD20 and CCD22
Mask Calibration	966, 966, 966, 1916, 2542, 1913, 1983, 1197, 1319, 966	Same as nominal set from OBF Verification
Orbital B set 1	966, 966, 966, 1916, 2542, 1913, 1983, 1197, 1319, 966	Nominal set, roughly 0.6-0.66 keV.
Orbital B set 2	966, 887 , 966, 1916, 2542, 1913, 1983, 987 , 1319, 868	CCD02 and CCD22 set to roughly 0.27 keV, CCD20 set to roughly 0.3keV to capture lower energy, to increase chances of lower energy element detection.
Orbital R set 1 and 2	Same as used in Orbital B	n/a
End of Table		

Table A.4: CCD ST parameters used throughout flight, with changes in bold text and explanation provided for each change. The values for the 10 nominal nodes are included, even though from L+30 onward, only five or six nodes were used

CCD Image Processing Parameters		
Operation	ST (raw ADU)	Rationale
Internal Calibration	250, 150, 150, 350, 850, 350, 962, 150, 812, 150	ST were set to around 50 eV for each node.
Cover Opening	No change	n/a
CXB Calibration	No change	n/a
Crab Calibration, part 1	No change	n/a
L+30 Calibration	No change	The image processing parameters did not change, but the NEM did - only 6 nodes were used.

Continuation of Table A.4		
Crab Calibration, part 2	No change	n/a
OBF Verification	No change	n/a
Mask Calibration set 1	250, 150, 150, 350, 850, 350, 962, 150, 812, 150	Same as nominal set.
Mask Calibration set 2	250, 300 , 150, 350, 850, 350, 962, 150, 812, 300	Doubled ST for CCD02 and CCD22 to try to recover detector efficiency.
Orbital B	250, 150, 150, 350, 850, 350, 962, 150, 812, 150	Nominal values because it was too late to change parameters, and ST was not a changeable variable within the Orbital B block.
Orbital R	250, 300 , 150, 350, 850, 350, 962, 150, 812, 300	Doubled ST for CCD02 and CCD22 to try to recover detector efficiency, and recover loss of iron counts.
End of Table		

Table A.5: CCD ULD parameters used throughout flight, with changes in bold text and explanation provided for each change. The values for the 10 nominal nodes are included, even though from L+30 onward, only five or six nodes were used

CCD Image Processing Parameters		
Operation	ULD (raw ADU)	Rationale
Internal Calibration	37, 47, 42, 47, 49, 47, 48, 47, 46, 47	ULD were set to 8 keV
Cover Opening	No change	n/a
CXB Calibration	No change	n/a

Continuation of Table A.5		
Crab Calibration, part 1	No change	n/a
L+30 Calibration	No change	n/a
Crab Calibration, part 2	No change	n/a
OBF Verification	37 57, 53, 58, 63, 59, 60, 59, 57, 59	Increased ULD to test the higher values.
Mask Calibration	37, 47, 42, 47, 49, 47, 48, 47, 46, 47	Back to nominal set; there was not enough time between OBF Verification and Mask Calibration to change the parameters once the use of the higher ULD was confirmed to not cause any issues.
Orbital B and Orbital R	37 57, 53, 58, 63, 59, 60, 59, 57, 59	Higher ULD values to increase the detected energy range, and potentially detect gamma ray bursts, or detect higher energy lines in the case of a large solar flare.
End of Table		

A.1 Code used for grading simulation

The frames were made using a python code which formatted the data into binary files. The binary files were then fed through an application written in C, that ran the files through the REXIS flight software grading algorithm, and output a data file with the same format as flight data.

Frame creating code

```
import numpy as np
from random import seed
from random import randint
import os

#make the dummy array of data w/rows = 1024, and columns = 4096,
#in node order
fakeFrame = np.zeros((1024, 4096), dtype='uint16')

#define an array of possible energy values, based on the fits to
#the energy histograms
energies = np.zeros(100, dtype='uint16')

,,,,,

#Distribution with no light leak, obtained from L30 raw frame
#approximation
for i in range(0, energies.size):
    if (i < 97):
        energies[i]=0
    else:
        energies[i]=(8 << 5)
,,,,,

#Small/no light leak energy distribution, from 7/26 fit to raw frame
#approximation
for i in range(0, energies.size):
    #Include bit-shifting in this array due to half-sized bins
    if (i < 73):
        energies[i]=0
    elif (i < 76):
```

```

        energies[i]= (2 << 5)
elif (i < 79):
        energies[i]=(2 << 5) + 16
elif (i < 81):
        energies[i]=(3 << 5)
elif (i < 83):
        energies[i]=((3 << 5) + 16)
elif (i < 85):
        energies[i]=(4 << 5)
elif (i < 87):
        energies[i]=((4 << 5) + 16)
elif (i < 88):
        energies[i]=(5 << 5)
elif (i < 90):
        energies[i]=((5 << 5) + 16)
elif (i < 91):
        energies[i] = (6 << 5)
elif (i < 93):
        energies[i] = ((6 << 5) + 16)
elif (i < 94):
        energies[i] = (7 << 5)
elif (i < 96):
        energies[i] = ((7 << 5) + 16)
elif (i < 97):
        energies[i] = (8 << 5)
elif (i < 98):
        energies[i] = ((8 << 5) + 16)
else:
        energies[i] = (9 << 5)
, , , ,

```

*#Light leak energy distribution, from 8/02 fit to raw frame approximation
for i in range(0, energies.size):*

#Include bit-shifting in this array due to half-sized bins

if (i < 73):

energies[i]=0

elif (i < 76):

energies[i]= (2 << 5)

elif (i < 79):

energies[i]=(2 << 5) + 16

elif (i < 81):

energies[i]=(3 << 5)

elif (i < 83):

energies[i]=((3 << 5) + 16)

elif (i < 85):

energies[i]=(4 << 5)

elif (i < 87):

energies[i]=((4 << 5) + 16)

elif (i < 88):

energies[i]=(5 << 5)

elif (i < 90):

energies[i]=((5 << 5) + 16)

elif (i < 91):

energies[i] = (6 << 5)

elif (i < 93):

energies[i] = ((6 << 5) + 16)

elif (i < 95):

energies[i] = (7 << 5)

elif (i < 97):

energies[i] = ((7 << 5) + 16)

elif (i < 98):

```

        energies[i] = (8 << 5)
    elif (i < 99):
        energies[i] = ((8 << 5) + 16)
    else:
        energies[i] = (9 << 5)
    , , , ,
#array of indices to choose from energies
indices = np.zeros(8, dtype='uint16')
#start the random number generator, choose a different seed for
#every 50 files
seed(30)

#run this loop 50 times

for j in range(0, 50):
    #randomly select 8 values from random distribution
    for i in range(0, 8):
        indices[i] = randint(0, 99)

#Put an iron count in the center of a defined grid within CCD22,
#and vary energy around it
#CCD22 start column = 2560 to 2815
#select energy from each defined index
#center row
fakeFrame[200][2800] = 8928
fakeFrame[200][2799] = (energies[indices[0]])
fakeFrame[200][2801] = (energies[indices[1]])
# 'top' row
fakeFrame[199][2800] = (energies[indices[2]])

```

```

fakeFrame[199][2799] = (energies[indices[3]])
fakeFrame[199][2801] = (energies[indices[4]])
#'bottom' row
fakeFrame[201][2800] = (energies[indices[5]])
fakeFrame[201][2799] = (energies[indices[6]])
fakeFrame[201][2801] = (energies[indices[7]])

#code to increment the last number of the saved file:
#convert to binary file and save as new file
with open((os.path.join('test' + str(j + 1) + '.dat')), 'wb') as f:
    fakeFrame.tofile(f)

f.close()

#end of file

```

A.2 Event grading code

This section only includes the C code written to use the REXIS flight software grading algorithm, not all the associated REXIS flight software. The code was written by Harvard Center for Astrophysics employee David Guevel.

```

#include <stdio.h>
#include <unistd.h>

#include "img_cxel_processing.h"
#include "telem.h"
#include "static_assert.h"
#include "img_defs.h"

```

```

static Image_t frame; //make frame global
static uint16_t data[4096*1024];
static XRayCandidateRecord_t cxel[65536];
static struct ImgGradeArgs_t imgGradeArgs;

```

```

void generate_frame()
{
    int row, col;
    int counter = 0;
    for (row=0; row<1024; row++){
        for (col=0; col<4320; col++){
            if (counter == 100) {
                frame[row][col] = 100;
                counter = 0;

            } else
                frame[row][col] = 0;
            counter += 1;
        }
    }
}

```

```

uint32_t generate_cxel(uint16_t et) {
    //frame is in ccd coordinates
    uint32_t offset, index, row, col;
    index = 0;
    for (row=0; row<1024; row++){
        for (col=0; col<4320; col++){

```

```

        if (frame[row][col] >= et){
            offset = row * 4320 + col;
            if (index < 65536)
            {
                cxel[index] = sizeof(Pixie_t) * offset;
                index += 1;
            }
            //printf("%d\n", offset);
        }
    }
}
return index;
}

```

```

void read_frame(char* fname){
    //open the frame which is in image coordinates (1024, 4096)
    FILE *fstream;
    fstream = fopen(fname, "rb");

    //read 1024*4320 numbers from the file
    size_t size;
    size = fread(&data, sizeof(uint16_t), 1024*4096, fstream);

    // initialize array to zeros
    int i;
    int j;
    for (i=0;i<1024;i++){
        for (j=0;j<4320;j++){
            frame[i][j] = 0;
        }
    }
}

```



```

    }
}

//shape data into 2d array
for (i=0; i<1024; i++)
{
    for (j=0; j<4096; j++)
    {
        frame[i][img2ccdCols(j)] = (uint16_t) data[i*4096+j];
    }
}
}

```

```

int main(int argc, char **argv)
{
    // read in cxel and frame from file and put into appropriate
    // data structures
    telemInit();

    Time_t time;
    time.seconds = 0;
    time.subseconds = 0;

    char* end;
    uint16_t et = strtol(argv[1], &end, 10);
    uint16_t st = strtol(argv[2], &end, 10);
    uint16_t uld = strtol(argv[3], &end, 10);
    printf("ET:_%i\n", et);
    printf("ST:_%i\n", st);
    printf("ULD:_%i\n", uld);
}

```

```

//et = et << 8;

// set parameters
uint8_t ccd , node;
for (ccd=0;ccd<4;ccd++){
    for (node=0;node<4;node++){
        imgSetCcdUld(ccd , node , uld << 8);
        imgSetCcdSplitThresh(ccd , node , st );
    }
}
imgSetGradeFilter(255);
imgSetMaxEvent((uint8_t) 800);

//generate_frame ();
int nframe;
for (nframe=4;nframe<argc;nframe++){
    printf ("%s\n" , argv [nframe ]);

    read_frame(argv [nframe ]);
    uint32_t numrecs = generate_cxel(et);
    imgGradeArgs.frame = &frame;
    imgGradeArgs.cxel = &(cxel [0]);
    imgGradeArgs.numRecs = numrecs;
    imgGradeArgs.diagMode = false;

// put data into gradeImage function
uint32_t output;
void *outargs;
output = gradeImage(&imgGradeArgs , outargs , false);

```

```
time.seconds += 4;
telemSendCcdScienceTelemetry(false, &time);

}

// write output to file; using actual telemetry functions to be
// read by parser
writeQueue("output.dat");

return 0;
}
```


Appendix B

CAST Materials

B.1 REXIS Level 4 Requirements pertaining to optical light

"ID: REX-52 Optical Light Reduction

The OBF shall reduce the flux of photons between 100 and 10000 nm that encounter the CCD top surface such that less than 1e-7% of photons in that band are detected by the CCD.

Parent: REX-8

Rationale: Optical photons contribute to the dark current and therefore need to be limited."

"ID: REX-227 Pin Holes in OBF

The percentage of detector area made unavailable for data collection due to pinholes in the OBF shall be less than 0.75%.

Parent: REX-11

Rationale: Excess pin-holes will reduce the effective area of the CCDs, potentially be-low required levels."

Both requirements are located in the REXIS Level 4 Detector and SXM Requirements Document.

B.2 Email chain discussing decision to not paint the DAM

From: Branden Allen <ballen@cfa.harvard.edu>

Subject: Re: Fwd: Dam(n?) Painting

Date: March 3, 2015 at 7:13:16 PM EST

To: Rebecca A Masterson <becki@mit.edu>

Cc: "rexis-science@mit.edu" <rexis-science@mit.edu>, "Branden T. Allen" <ballen@cfa.harvard.edu>

Here is a bit more,

Here I looked for regional variation over groups of 128 x 128 pixels on the CCD: of course this can be adjusted and will be the basis for our uniformity examination.

Two features popped out at me:

1. We seem to have many less events on one CCD node by about a factor of 3. I don't see anything wrong in my event extraction but it would be good if Steve could confirm for me.

2. The gain on one edge of the detector appears systematically slightly lower but at a level this is actually very tiny.

None of this shows a dependence on the light source: see included plots.

Best,
Branden

Hi,

running both runs with the illuminated and non-illuminated runs through my analysis pipeline does not reveal any significant difference in detector performance.

Included are the spectra with the light on (green) and light off (blue) with the Fe-55 lines for grade 0 events in the rexis grading scheme. The energy resolution at 5.9 keV is about 180 eV and there is no gain shift.

I'll have a closer look at this data set but things appear to be in great shape for this CCD.

Best,
Branden

Hi,

Sounds good if we just leave it as is to me: if it isn't broken don't fix it. I'll have a look at the data that Steve just sent to be sure but I don't anticipate coming to a different conclusion.

Branden

On Tue, Mar 03, 2015 at 08:45:21PM +0000, Rebecca A Masterson wrote:

report from Steve Kissel

Josh/Rick - pls respond with concurrence not to paint the DAM (or an objection to that path...)

thanks,
Becky

Begin forwarded message:

From: "Steve E. Kissel" <sek@space.mit.edu<<mailto:sek@space.mit.edu>>>
Date: March 3, 2015 1:36:09 PM EST
To: Rebecca A Masterson <becki@mit.edu<<mailto:becki@mit.edu>>>
Subject: Dam(n?) Painting

To test the sensitivity of the CCD output register to stray visible light the aluminum frame store cover was removed from the standard test holder. When exposed to visible light (standard test intensity TBD) no evidence of light leakage in the serial register or frame store was seen. The data from this test are found under device w53-1-14-1 in the archive which is currently being created.

Light leak via pinholes or edge defects in the frame store would present as uniform vertical "streaks" while light leak into output register would create a uniform bias shift (including overclock) and enhanced noise due to photon statistics. This is because every pixel spends equal times in all frame store rows and all output register pixels. Light leak presenting as a bright pixel can only occur in the image array. No vertical smears or increased noise are seen. Furthermore, resolution and count rate from Mn_k line is unaffected.

It appears that no special effort to exclude light from the lower portion of the CCD is warranted.

From: "Prof. Richard Binzel" <rpb@mit.edu>
Subject: **Re: Dam(n?) Painting**
Date: March 3, 2015 at 5:13:26 PM EST
To: Rebecca A Masterson <Becki@mit.edu>
Cc: rexis-science@mit.edu>

We are Go! for no painting of the DAM.

From: "Grindlay, Josh" <jgrindlay@cfa.harvard.edu>
Subject: **Re: Dam(n?) Painting**
Date: March 3, 2015 at 4:28:09 PM EST
To: Rebecca A Masterson <becki@mit.edu>
Cc: "rexis-science@mit.edu" <rexis-science@mit.edu>

Terrific, Steve.


Yes, Becky, I concur that painting the dam is not needed and thus should not be done.

Josh

Figure B-1: Emails concurring on decision to not paint the DAM

B.3 Plans to paint the DAM

The following pages show sketches done on parts of the DAM design to indicate where the black paint was to be applied (the black diagonal lines indicated surfaces to be painted). After the drawings, there are email conversations indicating the plans to paint the DAM had been in the works since the discovery of the potential for an edge light leak.



ISOMETRIC VIEWS
(FOR REFERENCE ONLY)

COMPENSATION STATUS: UNCONTROLLED
DATE: 13-10-30
SCALE: 1:1
SHEET 2 OF 2

MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY
SPACE SYSTEMS
LABORATORY
77 MASSACHUSETTS AVENUE
CAMBRIDGE, MA USA 02139

SHIELD, RADIATION, END,
BOTTOM



SEE CHARGE DRAWING NO. REX-DA-142
DATE: 13-10-30

B.3.1 Email chains discussing painting of the DAM

Chris,

I've attached a marked up drawing and some photos to help show where we would like paint; it is the the inside faces of the opening gap for the flexprint (see attached) (I've also added an extra "bottom" view on the -141 drawing which I shaded to indicate paint).

There is one remaining question which I list in the slides which is whether the overlap of -141 and -142 should be painted. I think the entire underside can be painted if the following conditions are met:

- paint thickness is negligible (it sounds this way if the thickness is 50um, although how thick is the primer?)
- there are no contamination issues with flaking of the paint during assembly
- cheaper to paint the whole underside instead of having to deal with masking

If the cost of masking is not an issue, then we might as well go ahead and mask off the overlap areas because we guarantee avoiding potential issues described in the first two bullets.

I'm not sure who is able to answer the questions above, but again, we can avoid having to answer them by doing the masking and eating whatever costs may be associated.

David

From: Krebs, Chris [ckrebs@aurora.aero]
Sent: Wednesday, November 26, 2014 10:16 AM
To: Rebecca A Masterson
Cc: Mark A Chodas; Michael Paul Jones; David Brad Carte
Subject: RE: Questions on REXIS DAM in preparation for procurement

The black paint drawing note says to apply "IAW Manufacturer's instructions", so that would cover any primer requirements. It also includes thickness and cure requirements, based on Kevin's remarks and the black paint datasheet that Mike sent me.

If we need to revise the drawing after fabrication that's fine, but I don't think we'll need to.

All I need now is for David and/or Mike to identify the surfaces on -141 and -142 that need the paint applied.

Chris

From: Rebecca A Masterson [mailto:beck@mit.edu]
Sent: Tuesday, November 25, 2014 4:36 PM
To: Krebs, Chris
Cc: Mark A Chodas; Michael Paul Jones; David Brad Carte
Subject: Fwd: Questions on REXIS DAM in preparation for procurement

Chris et al:

I'd like to close the loop on some of the actions from the DAM CCB:

- Keith Warner replied separately and agreed that helicoils are not needed for the DAM assembly, so no additional changes there
- Kevin Ryu replied below and we just spoke as well. He agreed that we should "not" change the size of the flexprint gap - the black paint will give us what we need, so no additional change to DA-141 and DA-142
- for the black paint spec: use 50 um thickness. Mike - did you send the spec along? there is also a primer that goes under it - did Mike send you that? Note that I can envision changes to the black paint primer/thickness etc, but I think we can release now for Saliga and then rev the paint spec between build and coating if we need to. Do you agree with that Chris?

Figure B-2: Email chain indicating plans to paint the DAM, part 1 of 3

• black paint should be added to DA-141 and DA-142. Mike/David - can you guys mark up the drawing to show Chris where? I know Kevin's response below indicates otherwise, but we just had a phone conversation and when I better explained what parts the are he changes his response.

I think that covers all of it for now. Any chance these parts went out for quoting yet?

thanks,
Becky

Begin forwarded message:

From: "Ryu, Kevin - 0887 - MITLL" <kevin.ryu@ll.mit.edu>
Date: November 24, 2014 4:42:42 PM EST
To: "Warner, Keith - 0887 - MITLL" <warner@ll.mit.edu>, Rebecca A Masterson <becki@mit.edu>
Cc: "ckrebs@aurora.aero" <ckrebs@aurora.aero>, Michael Paul Jones <mpj@mit.edu>, Mark A Chodas <mchodas@mit.edu>, David Brad Carte <dbcarte@mit.edu>
Subject: RE: Questions on REXIS DAM in preparation for procurement

Hi Becky,

Please see answers below:

- DA-141 and DA-142 are the parts that come together to form the gap that the flexprint goes through. When doing the review we discussed the possibility of narrowing this gap to cut down on the stray light that enters the DAM. Kevin - when Mike did the stray light analysis did he look at that variable at all? I'm hesitant to make a change unless there's a real benefit in doing so. Do you know if we win anything on the stray light problem over and above the paint application with this change? Mike's simulations showed that we can achieve the desired OD with the application of Z307 to the DAM underside, so I would agree with you that we should not change the DAM design.

All the others are related to the addition of black paint to some of the DAM parts:

- parts that we are flagging for black paint include: DA-126 (Side shield), DA-403 (55Fe top end collimator), DA-404 (55Fe bottom end collimator). These three parts all need to be gold coated as well. Is there any issue with applying the black paint over the gold coating?
 - There should be no issue as long as proper primer is used.
- We need a spec and thickness of the black paint for our drawings. We can send the relevant drawings to you to confirm location of black paint.
- Thickness is usually about 1-2 mils (25-50 um) per coat of paint.
- My assumption is that LL will do the painting at the same time as DAM assembly, but I need to confirm that with you Kevin and discuss any needed contact updates. When should we discuss? Should I be looking for alternatives?
 - The Z307 paint requires 7 days to cure at controlled humidity and temperature so I would recommend having it scheduled accordingly.
- Do you guys have a preference for painting at part level or some higher assembly level? The issue of paint flaking during assembly came up if done at the part level.
 - This is a question for Peter Anderson, who is the subject matter expert. If you provide me with mock-up DAM parts and instructions of surfaces to be painted, I can pass it along to Peter Anderson to get his thought on how to best achieve what we want to do.
- on parts DA-141, DA-142: how important it is to paint these? There is no direct line of sight to the CCDs, so black paint would be over indle here. This question is related to the second one above re: trade off between paint and gap size
 - I don't know exactly which parts are referenced here, but in the simulation only the DAM underside facing the CCD was painted, so I don't think we have to worry about painting these parts.

Above answers are my best guesses to the questions and should be verified by Peter Anderson.

Best regards,
Kevin Ryu

Figure B-3: Email chain indicating plans to paint the DAM, part 2 of 3

From: Rebecca A Masterson [mailto:becki@mtl.edu]
Sent: Monday, November 24, 2014 12:07 PM
To: Ryan, Kevin - 0887 - MITLL; Warner, Keith - 0887 - MITLL
Cc: ckrebs@aurora.aero; Michael Paul Jones; Mark A Chodas; David Brad Carter
Subject: Questions on REXIS DAM in preparation for procurement

Kevin/Keith

On Friday the REXIS team meant to review the DAM parts. We need to get the procurement package out next week to stay on schedule. A few questions arose from that review that we wanted to run by you guys:

- Most important questions time-wise: i.e. we need to confirm this one before sending out the package for quote
- Mark Chodas noted that the majority of the holes on the DAM are threaded and not helicoil inserts. We have used helicoils for the parts that we expect to take in and out a bit - namely the 55Fe source holders. However all the other DAM parts are simply threaded holes. We wondered if there is any concern about damage to the CCDs to do particulate generation during assembly. This question is directed mostly to Keith since he has assembled two of these with us already.
 - DA-141 and DA-142 are the parts that come together to form the gap that the flexprint goes through. When doing the review we discussed the possibility of narrowing this gap to cut down on the stray light that enters the DAM. Kevin - when Mike did the stray light analysis did he look at that variable at all? I'm hesitant to make a change unless there's a real benefit in doing so. Do you know if we win anything on the stray light problem over and above the paint application with this change?

- All the others are related to the addition of black paint to some of the DAM parts:
- parts that we are flagging for black paint include: DA-126 (Side shield), DA-403 (55Fe top end collimator), DA-404 (55Fe bottom end collimator). These three parts all need to be gold coated as well. Is there any issue with applying the black paint over the gold coating?
 - We need a spec and thickness of the black paint for our drawings. We can send the relevant drawings to you to confirm location of black paint.
 - My assumption is that LL will do the painting at the same time as DAM assembly, but I need to confirm that with you Kevin and discuss any needed contact updates. When should we discuss? Should I be looking for alternatives?
 - Do you guys have a preference for painting at part level or some higher assembly level? The issue of paint flaking during assembly came up if done at the part level.
 - on parts DA-141, DA-142: how important it is to paint these? There is no direct line of sight to the CCDs, so black paint would be over indle here. This question is related to the second one above re: trade off between paint and gap size

Please let me know if a phone call would be better to talk through these. I'd like to discuss or get a response back today or tomorrow in preparation for getting the quote package out - at least for the top two questions above. the black paint can wait if needed.

thanks
Becky

Rebecca Masterson, Ph.D.
SSL Research Engineer
REXIS Instrument Manager

Bibliography

- [1] E. Beshore, D. Lauretta, W. Boynton, C. Shinohara, B. Sutter, D. Everett, J. Gal-Edd, R. Mink, M. Moreau, and J. Dworkin. The OSIRIS-REx Asteroid Sample Return Mission. In *2015 IEEE Aerospace Conference*, pages 1–14, 2015.
- [2] D.S. Lauretta et al. OSIRIS-REx: Sample Return from Asteroid Bennu. *Space Science Reviews*, 212:925–984, 2017.
- [3] R. A. Masterson, M. Chodas, L. Bayley, B. Allen, J. Hong, P. Biswas, C. McMenamin, K. Stout, E. Bokhour, H. Bralower, D. Carte, S. Chen, M. Jones, S. Kissel, F. Schmidt, M. Smith, G. Sondecker, L. F. Lim, D. S. Lauretta, J. E. Grindlay, and R. P. Binzel. Regolith X-Ray Imaging Spectrometer (REXIS) Aboard the OSIRIS-REx Asteroid Sample Return Mission. *Space Science Review*, 214, 2018.
- [4] M. Jones, M. Chodas, M. Smith, and R. Masterson. Engineering Design of the REgolith X-ray Imaging Spectrometer (REXIS) Instrument: An OSIRIS-REx Student Collaboration. In *SPIE Astronomical Telescopes and Instrumentation*, 2014.
- [5] M. Horányi et al. The Student Dust Counter on the New Horizons Mission. *Space Science Reviews*, 140, 2007.
- [6] K. Ryu, B. Burke, H. Clark, R. Lambert, P. O’Brien, V. Suntharalingam, C. Ward, K. Warner, M. Bautz, R. Binzel, S. Kissel, and R. Masterson. Development of CCDs for REXIS on OSIRIS-REx. In *SPIE Astronomical Telescopes and Instrumentation*, 2014.
- [7] Laura Bayley. Integration and Test of the REgolith X-ray Imaging Spectrometer and Recommendations for Low-Cost, High-Risk Spaceflight Programs. Master’s thesis, Massachusetts Institute of Technology, 2016.
- [8] Office of the Chief Engineer. NASA Space Flight Program and Project Management Requirements w/Changes 1-14. *NASA Procedural Requirements (NPR) 7120.5 E*, 2012.
- [9] Office of Safety and Mission Assurance. Risk Classification of NASA Payloads. *NASA Procedural Requirements (NPR) 8705.4*, 2014.

- [10] Michael P. Jones. The Engineering Design of the REXIS Solar X-ray Monitor and Risk Management Considerations for Resource Constrained Payload Development. Master's thesis, Massachusetts Institute of Technology, 2015.
- [11] Daniel Durini. *High Performance Silicon Imaging: Fundamentals and Applications of CMOS and CCD Sensors*, chapter 3. Woodhead Publishing, 2014.
- [12] Pronoy Biswas. Radiation Management, Avionics Development, and Integrated Testing of a Class-D Space-Based Asteroid X-ray Spectrometer. Master's thesis, Massachusetts Institute of Technology, 2016.
- [13] K. Ryu, M. Bautz, S. Kissel, P. O'Brien, and V. Suntharalingam. Directly deposited optical-blocking filters for single-photon X-ray imaging spectroscopy. *Journal of Astronomical Telescopes, Instruments, and Systems*, 3, 2017.
- [14] J. W. den Herder et al. The Reflection Grating Spectrometer on board XMM-Newton. *Astronomy and Astrophysics*, 365:L7–L17, 2001.
- [15] H. Tsunemi et al. In-orbit performance of the MAXI/SSC on board the ISS. *Publications of the Astronomical Society of Japan*, 62:1371–1379, 2010.
- [16] A. D. Falcone et al. Recent progress on developments and characterization of hybrid CMOS x-ray detectors. In *SPIE Astronomical Telescopes and Instrumentation*, 2012.
- [17] B. Anderson and T. Fagerhaug. *Root Cause Analysis Simplified Tools and Techniques*. ASQ Quality Press, second edition, 2006.
- [18] T. Katatoka, K. Furuto, and T. Matsumoto. The Analyzing Method of Root Causes for Software Problems. *SEI Technical Review*, 73, 2011.
- [19] Prashant Gangidi. A systematic approach to root cause analysis using 3 x 5 why's technique. *International Journal of Lean Six Sigma*, 10:295–310, 2019.
- [20] W. Perry and N. Mehlretter. Applying Root Cause Analysis to Compressed Air: How to Solve Common Compressed Air System Problems with the 5-Whys. *Energy Engineering*, 115:56–62, 2018.
- [21] Duane Kritzinger. *Aircraft System Safety: Assessments for Initial Airworthiness Certification*, chapter 4-5, pages 59–132. Woodhead Publishing, 2016.
- [22] William Vesely, Joanne Dugan, Joseph Fragola, Joseph Minarick III, and Jan Railsback. *Fault Tree Handbook with Aerospace Applications*. 2002.
- [23] J. Soeder, R. Schiedegger, L. Pinero, A. Birchenough, and J. Dunning. NASA's Evolutionary Xenon Thruster (NEXT) Power Processing Unit (PPU) Capacitor Failure Root Cause Analysis. In *American Institute of Aeronautics and Astronautics 10th International Energy Conversion Engineering Conference*, 2012.

- [24] G. Iona, J. Butler, B. Guenther, L. Graziani, and E. Johnson. VIIRS on-orbit optical anomaly: investigation, analysis, root cause determination and lessons learned. In *SPIE Optical Engineering and Applications*, volume 8510, 2012.
- [25] J. Leitner, B. Sood, E. Isaac, J. Shue, N. Lindsey, and J. Plante. Risk-based safety and mission assurance: Approach and experiences in practice. *Quality Engineering*, 30:648–662, 2018.
- [26] Nancy Leveson. *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press, 2011.
- [27] J. Huang, J. You, H. Liu, and M. Song. Failure mode and effect analysis improvement: A systematic literature review and future research agenda. *Reliability Engineering and System Safety*, 199, 2020.
- [28] A. Pllay and J. Wang. Modified failure mode and effects analysis using approximate reasoning. *Reliable Engineering System Safety*, 79:69–85, 2003.
- [29] JB Bowles and CE. Pelaez. Fuzzy logic prioritization of failures in a system failure mode, effects, and criticality analysis. *Reliable Engineering System Safety*, 50:203–213, 1995.
- [30] Nancy Leveson. *CAST Handbook: How to Learn More from Incidents and Accidents*. 2019.
- [31] Meaghan O’Neil. Application of CAST to Hospital Adverse Events. Master’s thesis, Massachusetts Institute of Technology, 2014.
- [32] Airong Dong. Application of CAST and STPA to Railroad Safety in China. Master’s thesis, Massachusetts Institute of Technology, 2012.
- [33] Tae-eun Kim, Salman Nazier, and Kjell Ivar Øvergård. A STAMP-based causal analysis of the Korean Sewol ferry accident. *Safety Science*, 83:930–101, 2016.
- [34] Jon Hickey. A System Theoretic Safety Analysis of U.S. Coast Guard Aviation Mishap involving CG-6505. Master’s thesis, Massachusetts Institute of Technology, 2012.
- [35] Riccardo Giacconi, Herbert Gursky, Frank R. Paolini, and Bruno B. Rossi. Evidence for X-rays From Sources Outside the Solar System. *Physical Review Letters*, 9:439–443, 1962.