

Near Earth Asteroid Scout CubeSat Science Data Retrieval Optimization Using Onboard Data Analysis

Jack Lightholder
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
626-710-3246
Jack.A.Lightholder@jpl.caltech.edu

Julie Castillo-Rogez
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-640-9231
Julie.C.Castillo@jpl.caltech.edu

David R. Thompson
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-2200
David.R.Thompson@jpl.caltech.edu

Christophe Basset
Jet Propulsion Laboratory,
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-2405
Christophe.Basset@jpl.caltech.edu

Abstract— Small spacecraft are continually evolving in capability and mission complexity. As spacecraft size decreases, physical limitations present new challenges for mission designers. These include limited instrument aperture, low communications bandwidth, and reduced attitude control. Software techniques can address these limitations to retain the capabilities of larger spacecraft, in a small form factor. These techniques move the first order science analysis, which is traditionally completed on the ground, onboard the spacecraft. This can minimize the amount of data volume required for first order decision making. We present a collection of techniques designed for mitigating limited pointing stability for target acquisition, onboard image calibration and decision making, in a low bandwidth environment. These technologies will fly on the Near Earth Asteroid Scout (NEA Scout) interplanetary CubeSat mission in 2019 [1]. To support verification and validation, the flight software implementations of these algorithms have been run on New Horizons, Rosetta and terrestrially acquired data sets. These experiments validate the capability to detect the target, through noise and target location uncertainty.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. CHALLENGES FACED BY NEA SCOUT	2
3. COMMAND DEVELOPMENT.....	3
4. CALIBRATION RESULTS	5
5. CONCLUSIONS.....	6
ACKNOWLEDGEMENTS	7
REFERENCES.....	7
BIOGRAPHY	7

1. INTRODUCTION

As part of the Space Launch System (SLS) Exploration Mission 1 (EM-1) mission, the Near-Earth Asteroid Scout (NEA Scout) CubeSat mission will fly to about 1 AU from Earth to conduct a flyby of a near Earth asteroid (NEA) less than 100 m across. A 6U CubeSat, NEA Scout will be guided by a solar sail, towards its target asteroid 1991 VG [1]. Due to its small size and low albedo, complete characterization of 1991 VG from Earth is difficult. A combination of target orbit uncertainty and long lead times for solar sail trajectory correction maneuvers drive a requirement to identify the target in optical navigation imagery at a distance of about 60,000 km. At closest approach, the same imager will be used for near field imaging of the target. Figure 1 summarizes the NEA Scout concept of operations.

Traditional large spacecraft accomplish these imaging objectives using long exposures to increase SNR and identify the low albedo target. Due to the pointing drift and jitter inherent in a small platform, long exposure imaging is less feasible for NEA Scout. Onboard image processing overcomes this challenge. The spacecraft aligns and combines a stack of rapidly acquired images, resulting in a single image with a higher SNR than its constituent images. We filter the aligned images using a temporal median. This solution fits within the memory constrained onboard context. Prior to alignment, each image undergoes a first order image calibration, onboard, to improve the results of the alignment. This calibration consists of a dark current subtraction, flat field adjustment and bad pixel mask application. The temporal median has the added benefit of removing transient imaging artifacts, such as cosmic rays. Interplanetary CubeSats, such as NEA Scout, are additionally physically constrained by the size of their antenna and available transmission power, which is a major challenge for science-driven CubeSat missions and the basis for the work presented in this paper. At closest target approach, NEA Scout will be constrained to approximately <1 kbps downlink bandwidth. We address this limitation with automatic image cropping algorithms

and software routines which downlink image statistics, giving operators a better understanding of the image content before committing it to the downlink queue. Alternatively, operators can command specific cropping operations, or as a window around the brightest point in the image. The combination of these techniques enables early target detection in an onboard context, without stringent pointing requirements, in a low bandwidth mission scenario. These capabilities leverage onboard data processing to

distill decision making data to tenable size for low bandwidth deep space communication paradigms. The demonstration of this novel science software on an actual science-driven mission will pave the way for future smallsat missions to distant destinations [2].

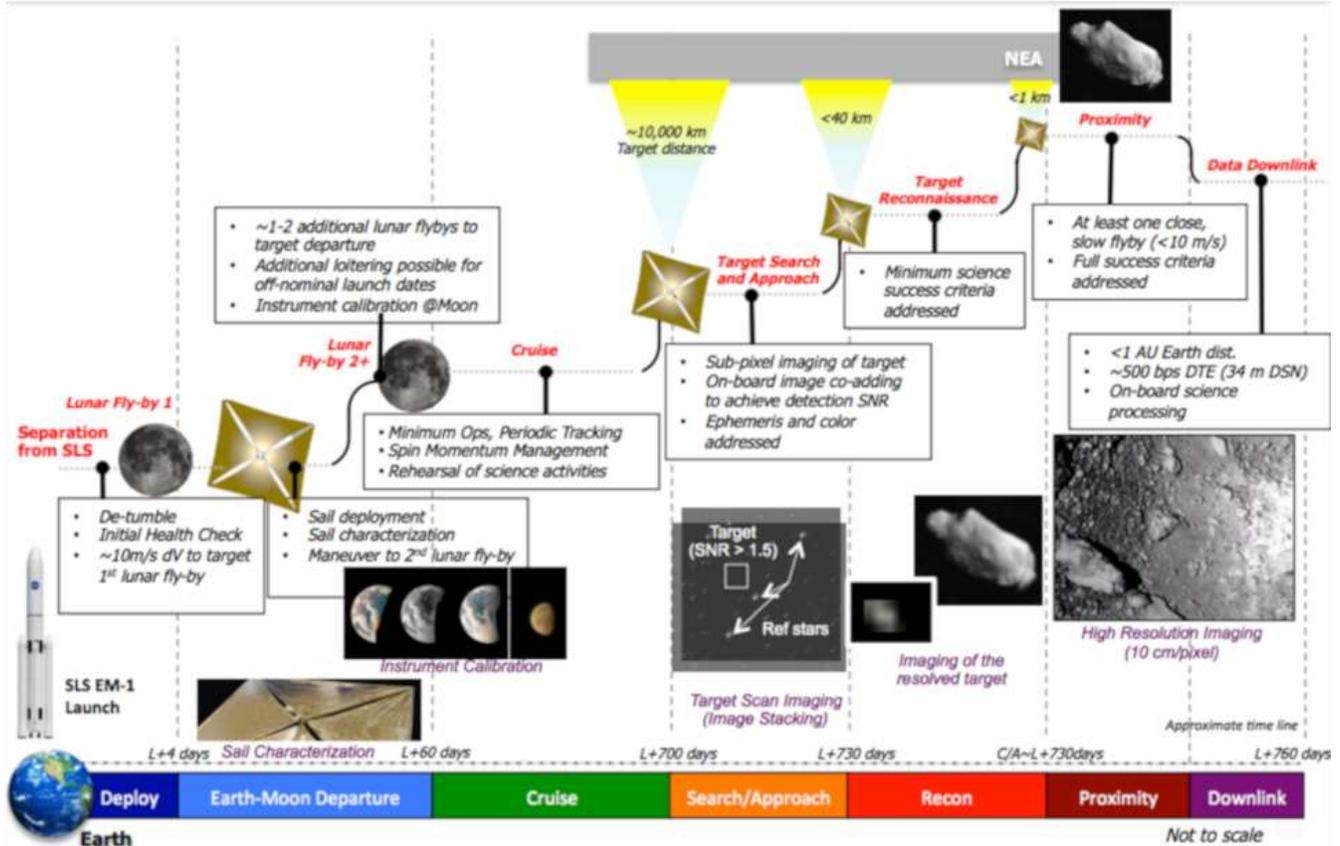


Figure 1. Summary of NEA Scout's activities throughout the ~2 year mission.

2. CHALLENGES FACED BY NEA SCOUT

This section provides context for understanding the challenges faced by the NEA Scout mission with relevance to many other science driven CubeSats. The majority of the CubeSats currently under development for flight in the near term are based on a 6U form factor. In the case of NEA Scout, about half of that volume is allocated to the propulsion system (solar sail and thrusters), 1/3rd to the avionics and instrument, and the rest is used by the power system, the antenna, structure, and harness. The instrument is a monochromatic camera certified for deep space and that acts both as a science instrument and an optical navigation camera. NEA Scout's science objectives are to retire strategic knowledge gaps for Human exploration and increase our understanding of near earth asteroids by focusing on a class of targets (<100 m) that has not been covered by previous and ongoing missions. Specific measurement objectives include global shape determination and regional morphology mapping, determination of rotational parameters, including whether the object is a single axis rotator or a tumbler, albedo mapping on a global scale, and high-resolution imaging of a fraction of the surface. At closest approach, the resolution is projected to be <10 cm/pix.

The size of the NEAScout reference target, 1991 VG, is between 5 and 17 meters. Although ground-based observations acquired 25 years apart have provided relatively accurate ephemeris for that body, its small size and potentially low albedo, make it a challenging target for approach observations. The encounter is planned at about 1 AU from Earth.

The NEA Scout camera detector is similar to that used for the navigation cameras on the Mars 2020 rover [3]. This camera takes advantage of the new generation of arrays with a frame size of ~14 MPx that enables good spatial resolution of the target images while preserving a large field of view necessary for target search and optical navigation. A major drawback is the large volume of the raw data, 225 Mb per image for an imaging depth of 16 bits. In absence of a priori knowledge on the target, it is not possible to predict the parameters for lossless compression. However, it is understood that the target fills in only a small fraction of the field of view. During the Approach phase (see Fig. 1), this fraction, including margins based on the target position uncertainty and the spacecraft attitude uncertainties, is about 0.28%. During the science phase of the mission (Reconnaissance and Proximity) the NEA occupies about 7% of the field of view. The total downlinked

data volume is about 200 Mb, the bulk of which is acquired during the science phase. This corresponds to a downlink time of 60 hours at a rate of 1 kbps projected at 1 AU.

Besides data volume, the NEA Scout mission is facing another key challenge in the form of pointing conflicts among various subsystems: camera, solar panels, medium-gain antenna, and solar sail. Ground contacts are limited to about 50 minutes, driven by the secondary batteries, followed by recharge periods of about 8 hours. When all constraints are accounted for, the 60-hour downlink has to be broken down over a period of 30 days.

Pointing performance meets the requirement to stay within 0.2 pixel over an integration time of 0.7 sec and to stay within a box of 100x100 pixels during the acquisition of 20 images for the target search activity described in this paper. Also, use of JPL's small computer, the Sphinx [4], provides the data storage and computing performance necessary to implement the data management strategies presented in this paper.

3. COMMAND DEVELOPMENT

Onboard Image Calibration

Onboard data analysis begins with the application of a first order image calibration. This includes the application of a flat field, for calibrating away gain difference across the detector elements, and a dark current subtraction, to reduce detector electronics noise. Additionally, a bad pixel mask is applied to eliminate known defective pixels from being used in subsequent analysis. These three components are applied using the algorithm outlined in equation 1.

$$\begin{aligned} \text{calibrated pixel}_i = & \quad (1) \\ & (\text{raw pixel}_i - \text{dark field pixel}_i) \\ & * \left(\frac{\text{flat field pixel}_i}{\text{flat field mean value}} \right) \\ & * \text{radiometric calibration coefficient} \end{aligned}$$

If bad pixel mask_i == 1; calibrated pixel == 0

A calibration of this type would traditionally be the first processing done to data records once they are downlinked. Due to the extremely constrained downlink of the mission, this calibration must be moved onboard the spacecraft. In flight, a minimum signal to noise between the target and background of 7 is required. This low SNR, coupled with the subpixel size of the target at first acquisition, requires this level of image calibration to successfully identify the target.

The ability to do image calibration in flight allows for tighter coupling between the onboard science software and imaging hardware. This calibration equation could be extended to handle additional detector characteristics, based on mission needs. Additional information about the NEA Scout camera and calibration considerations for our mission can be found in section 4.

Onboard Frame Coaddition

Traditional target acquisition paradigms leverage the use of long exposure imaging to increase the SNR between their target and background. Small spacecraft have limited ability to reach this pointing precision for long periods of time. Without a stabilized imaging platform, long exposure imaging will introduce sufficient blurring to obscure a subpixel target. Our approach leverages a collection of rapidly collected, low exposure time images, which are subsequently combined to emulate a single long exposure image. This coaddition process requires the quantification of shift between images, and an approach for combining images once they are spatially aligned. This activity levies requirement on C&DH performance (computing and memory), which in part drove the design of the Sphinx C&DH used on NEA Scout and other CubeSat missions.

Shift Calculation

Detection of shift occurs between each image, and the base image of reference. Using the first image as a reference base, the (x, y) position of the N number of brightest stars are calculated. This is accomplished by determining the brightest pixel, and isolating an area around it. This process is repeated for N stars, until the matrix of star positions is known. The isolation boxes around each star ensure subsequent new pixels are not part of an already located star. Due to the natural Gaussian behavior of star brightness, it is assumed the brightest pixel will be near the center of each star detected.

For all subsequent images in the reference stack, we then determine the position shift using the reference stars from the base image for reference. For each image, the known reference star locations are visited. A bounding box height and width are defined by the user, which is used to sweep the around the expected star location. The sliding window sweeps the region around the expected star position, performing 2-D cross correlation at each window location.

$$r = \frac{\sum_m \sum_n (A_{mn} - \bar{A})(B_{mn} - \bar{B})}{\sqrt{(\sum_m \sum_n (A_{mn} - \bar{A})^2)(\sum_m \sum_n (B_{mn} - \bar{B})^2)}} \quad (2)$$

Where \bar{A} = mean(A), \bar{B} = mean(B), m = row index and n = column index

This cross correlation compares the window segment in the new image against the box surrounding the known star location in the base image. Once the window has explored the defined search area around the expected star location, the 2-D cross correlations are compared. We take the window location with the highest cross correlation to be the new location of the star. The (x, y) shift between that bounding box location and the original reference star location dictates the shift of that star in the frame. We calculate this value for each star and take the mean as the shift for the image [5].

Frame Coaddition

After calculating the total image position offsets, we can then compare space-aligned pixels from subsequent frames. Memory constraints preclude using all images at once. Instead, for a stack of n images to which we are appending image n+1, we calculate a three-value temporal median combining images n-1, n, and the new frame n+1. As this process is performed for each new image,

the resulting median calculations are added to the base of output image. The result is a spatially aligned image, median filtered, with a higher SNR than the constituent images. Figure 2 portrays the median filter calculation algorithm.

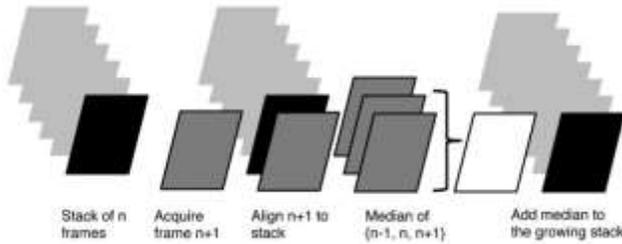


Figure 2. Visual representation of image coaddition process [5].

In addition to combining the image stack into a single high SNR image, the median filtering in this process removes artifacts which do not exist in the majority of image frames. This is of particular relevance for cosmic ray hits, which can have high DN values and confuse subsequent target detection algorithms.

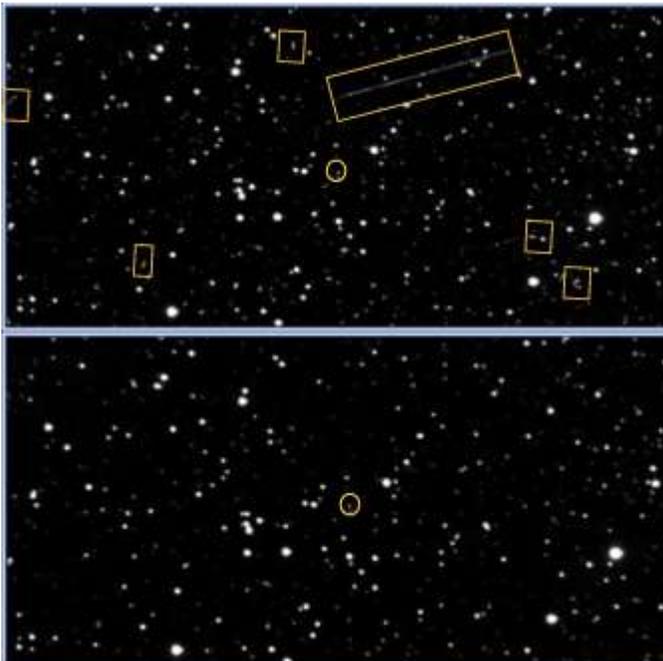


Figure 3. Rosetta OSIRIS Narrow Angle Camera detection of 2867 Steins. Raw image with artifacts in boxes and Steins circled (top) and image processed using coaddition pipeline and artifacts removed (bottom) [6].

The resulting frame now encompasses the data from the calibrated images, with increased SNR and noise reduction.

Target Detection

Downlinking full frame coadded images still exceed the NEA Scout mission bandwidth, hence further onboard data refinement is necessary. Since the target is subpixel during approach, downsampling will obscure the small feature scale too significantly

to be of use. Instead, intelligent cropping will be used to isolate the area around the target for downlink.

The first mode of cropping is with a user-defined image region. This region can be fit to the target uncertainty ellipse, removing image regions where the target is known to not be. More aggressive cropping can be performed by taking two co-added image stacks and calculating the shift between them. Due to the relative proximity to the target, compared to background stars, the target will shift different than the background stars between these two images. Performing an aligned subtraction on these images predominantly removes background stars, while leaving the target unimpacted. Cropping can then be done directly around the target by specifying the size of a window around the brightest pixel in the resulting subtraction.

Occlusion of the target by reference stars remains a limitation of this approach. In cases where the target has been obscured by a star, target information will be lost in star removal. This can be mitigated through camera repeating the image coaddition process after allowing for star position to move. In the case of the NEA Scout mission trajectory, 20% of the imaging field is expected to be populated by stars. Repeating the activity of multiple occurrences mitigates the concern of occlusion of the target behind a star. The small data product size resulting from this activity, allows for the activity to be repeated multiple times, with minimal amounts of data to be downlinked to verify successful target acquisition.

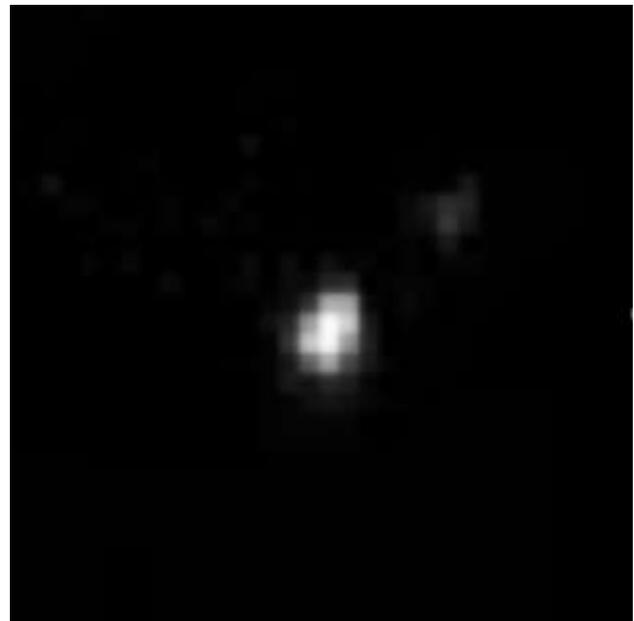


Figure 4. Pluto/Charon detection using the described pipeline on raw LORRI images, 265 million miles from Pluto [7].

During closest approach, the same camera will be used for near field imaging. At closest approach, the spacecraft will be within ~800 meters of the target. Full target imaging at this distance will cover approximately 7% of the detector field. Knowing this, combined with target refinement information collected in cruise, the system can calculate a bounding box to ensure the full target is in a resulting crop.

Image Metrics and Compression

While calibration, coaddition and cropping strategies significantly reduce data volume for downlink, the rate of image acquisition still outstrips the downlink capabilities. The mission plans to acquire N images during closest approach, but projects a downlink data volume sufficient to transmit M of them. The resulting challenge is deciding which images are of highest prioritize for downlink, without downlink subsampled images. To this end, the mission will use commands for generating image statistics onboard, which are believed to be suggestive of image quality, with minimal bandwidth. Statistics include a histogram of pixel values, a contrast calculation, mean, standard deviation, minimum and maximum pixel values. Additionally, header values relating the camera performance at the time of image acquisition are downlinked. These statistics provide insight into the order for which images should be prioritized for downlink, without significant bandwidth utilization.

Finally, ICER compression is utilized to further compress image segments identified for downlink. ICER supports lossless and lossy compression, in a scene specific manner [8]. Due to the scene-specific nature of the ICER compression algorithm, simulation was necessary to quantify the expected compression ratio for near field imaging and cruise images. Using Rosetta OSIRIS Narrow Angle Camera approach images for target 67P, 12-bit images were compressed with an average compression of 6 bits per pixel, for lossless compression. Studies of ICER performance with Mars imagery report an upper bound of 9 bits per pixel for lossless compression of 12-bit source images [8]. These metrics inform the bandwidth which should be booked for each image segment compressed during the NEA Scout mission.

4. CALIBRATION RESULTS

In flight calibration improves the input quality of data for target detection, through the removal of known detector characteristics. The most prominent characteristic for this detector is the impact of temperature on dark current and detector response. In an effort to fully characterize these responses, and integrate them into the onboard calibration, a laboratory calibration of the flight camera was conducted. In an effort to achieve the highest fidelity temperature calibration, while minimizing software complexity, the decision was made to store dark current and flat field calibration products onboard for each 5-degree temperature step in the allowable flight temperatures. This tradeoff increases calibration product storage size, while eliminating the need to fit a response model to the temperature term of our onboard calibration. Laboratory calibration products were collected before camera integration and will be loaded into the spacecraft before launch.

Flat and dark current fields were acquired in laboratory conditions, with the intention of use in flight. In an effort to encapsulate the greatest amount of camera characterization, these fields were taken at a variety of temperatures and exposures, to capture temperature dependencies in our calibration. These temperatures ranged from -25C through 50C, the operational range of the camera, in 5C increments. At each temperature, 5 images were captured at exposures of 200, 400, 600 and 800ms. Figures 5 and 6 outline the dark and flat field response, as a function of temperature and exposure. For each temperature and exposure combination, the 5 frames were averaged to reduce the impacts of photon shot noise on the final calibrated product.

Camera allowable flight temperatures support imaging between -25C and 50C, however mission simulations place all imaging

sequences in an expected thermal window of -5C to 5C. The average dark current response below 20C has linear dependence with temperature. Due to this characteristic, linear interpolation between dark current fields will be used when determining the dark current calibration value, which best matches the temperature of the raw image, for the calibration equation in figure 2.

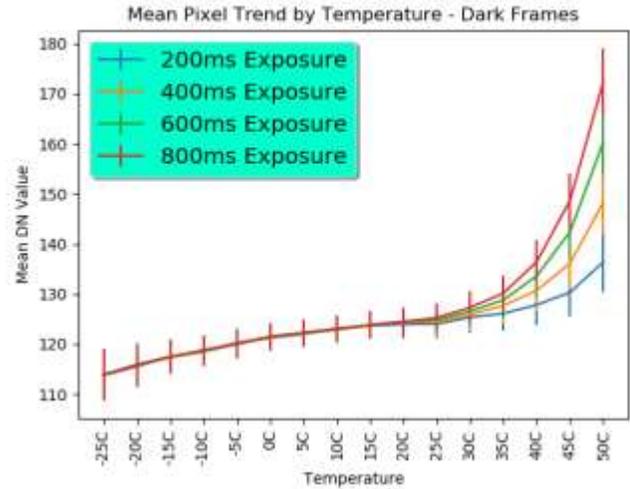


Figure 5. Dark field mean pixel trend with temperature.

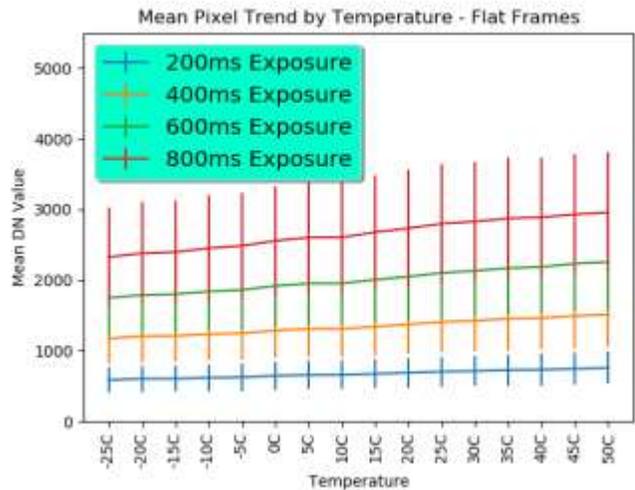


Figure 6. Flat field mean pixel trend with temperature.

Detector response in flat field imaging shows strong linearity both in exposure time and temperature relationships. Exposure time linearity is leveraged to create a single flat field for each temperature. Individual pixel response within the resulting flat field is representative of the response ratio between pixels at any exposure setting within the flight exposures. Pixel response as a function of temperature is less pronounced than the trend observed in dark current temperature trends. The existing software infrastructure and onboard storage availability supports the decision to include a collection of temperature differentiated flat fields in flight and perform linear interpolation between frames, identical to the process for dark current value interpolation.

In addition to flat field and dark current calibration files, bad pixel masks are included onboard. Bad pixels carry more significant risk, due to the subsequent onboard analysis of calibrated images. Leaving pixels with high artificial DN values in the calibrated

images has the potential to interfere with frame alignment or cause the cropping algorithms to lock onto pixels other than the target. Pixels with flat field response values more than 3σ from the field mean are identified as bad pixels in the base bad pixel mask included onboard. Throughout the mission, these masks can be updated via a command to signify a pixel should no longer be considered in analysis.

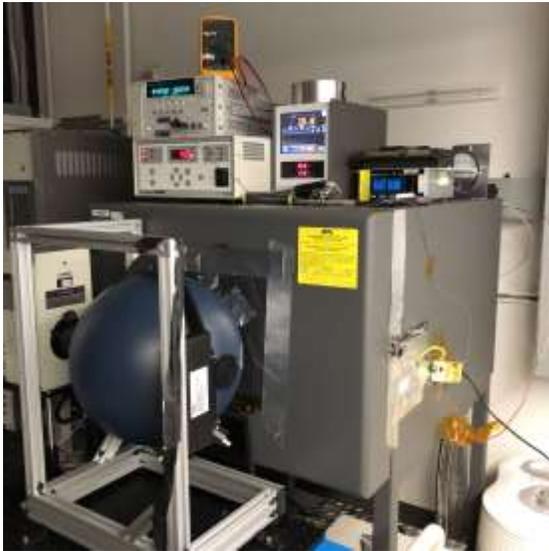


Figure 7. Experimental setup for camera calibration.

NEA Scout Camera

To reduce development time of the camera while reducing the overall cost and risk inherent to creating a new design, the NEA Scout camera takes advantage of an existing modular camera platform implemented for the Orbiting Carbon Observatory 3 (OCO-3) context cameras [9]. The OCO-3 context camera provided the electronics body while allowing mission-specific customizations. To meet the signal to noise ratio and field of view requirements, the NEA Scout camera integrates the monochrome version of the CMV20000 CMOS detector used in the OCO-3 implementation. For the optics, a ruggedized commercial lens was procured that meets the speed and field of view necessary for the object detection and close-up imaging: $f/2.8$, 50.2mm focal length and 27° field of view. The image circle projected onto the detector from the lens is 24mm, reducing the useful detector window to 3840×3840 pixels. In practice, the target detection only needs of a reduced size so the detector windowing capability is used to capture a smaller size image for each frame (3840×2184 pixels).



Figure 8. NEA Scout camera before spacecraft integration.

Table 1. NEAScout camera physical specifications.

Sensor Capabilities	
Type	20M pixel CMOS image sensor
Useful array size	3840×3840 pixels
Pixel size	$6.4\mu\text{m}^2$
Full well	$15,000e^-$
Dark noise	$8e^-$ RMS
Windowing	Y-direction only
Shutter	Global
Color	Monochrome (with microlenses)
Quantization	12-bit per pixel
Electrical interface	
Physical	LVDS
Protocol	Spacewire RMAP
Power	< 3 Watts
Memory	64Mbits
FPGA	Microsemi Rad-tolerant ProASIC3
Camera Specifications	
Mass	390g
Volume	$63\text{mm} \times 63\text{mm} \times 71\text{mm}$
Operating temperature	-25C to $+50\text{C}$
Survival temperature	-35C to $+70\text{C}$
Optics	27° FOV, $f/2.8$, 50.2mm $i\text{FOV}=0.09\text{mrad/pix}$

5. CONCLUSIONS

Software techniques can supplement spacecraft hardware limitations to achieve comparable science objectives in small spacecraft, as in large. These limitations include reduced pointing precision and limited bandwidth volume.

Technologies such as onboard image calibration and analysis metric reporting eliminate the need to downlink raw images for immediate calibration. Moving this well characterized analysis onboard the spacecraft makes these calibrated data products available for subsequent onboard analysis. The ability to detect shift between images, along with temporal median filtering allows stacks of rapidly acquired images to be combined into a single high SNR image, emulating long exposure optimal navigation images, on a lower precision pointing platform.

Onboard alignment of images for spatially aligned pixel-wise subtraction enables background star removal, increasing the chances of onboard target detection. Dynamic cropping around the target, after background star removal, allows for target location quantification and imaging using small amounts of bandwidth. Decreases in bandwidth support interplanetary smallsats, where communications passes are small and infrequent.

Increasing the paradigm of onboard data analysis enables new mission profiles which are not possible with traditional methods for analyzing science return. Distilling science data return enables increased focused of attention by human operations, reducing turnaround time for critical decision making. Missions like NEA Scout, where target characterization at launch might be limited, require in flight characterization of the target and subsequent operations agility to respond to the changing operational requirements. Onboard data analysis allows for increased target observation data, earlier in cruise, giving the operations team

valuable extra time to perform course corrections for the navigation system.

Increased insights, focused data, optimized bandwidth and early target detection enable missions like NEA Scout on bandwidth limited and pointing constrained platforms. Utilizing software solutions to enable new mission platforms increases exploration potential, allowing smallsats to increase science return and expand beyond low Earth orbit.

ACKNOWLEDGEMENTS

© 2018 California Institute of Technology. Government sponsorship acknowledged.

This work is being carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract to NASA. The Near Earth Asteroid Scout mission is a project managed by the Marshall Space Flight Center. Spacecraft avionics and software are being developed at the Jet Propulsion Laboratory. Government sponsorship acknowledged.

REFERENCES

- [1] L. McNutt, L. Johnson, P. Kahn, J. Castillo-Rogez, and A. Frick, "Near-Earth Asteroid (NEA) Scout," AIAA SPACE 2014 Conference and Exposition, 2014.
- [2] S. Chien, B. Bue, J. Castillo-Rogez, D. Gharibian, D. R. Thompson, and K. L. Wagstaff, "Agile Science for Primitive Bodies and Deep Space Exploration," SpaceOps 2014 Conference, Feb. 2014.
- [3] J. N. Maki, C. M. McKinney, R. G. Sellar, R. G. Willson, D. S. Copley-Woods, D. C. Gruel, D. L. Nuding, M. Valvo, T. Goodsall, J. McGuire, J. Kempenaar, and T. E. Litwin, "Enhanced Engineering Cameras (EECAMs) For The Mars 2020 Rover," *3rd International Workshop on Instrumentation for Planetary Missions*, 2016.
- [4] B. Sherwood, S. Spangelo, A. Frick, J. Castillo-Rogez, A. Klesh, E. J. Wyatt, K. Reh, and J. Baker, "PLANETARY CUBESATS COME OF AGE," *66th International Astronautical Congress*, 2015.
- [5] W. Huffman, D. R. Thompson, B. Bue, J. Castillo-Rogez, and J. Boland, "Autonomous Onboard Point Source Detection by Small Exploration Spacecraft," *Publications of the Astronomical Society of the Pacific*, vol. 127, no. 958, pp. 1279–1291, 2015.
- [6] "Hviid, S. F., Keller H. U. and the OSIRIS Team, OSIRIS ROSETTA EXPERIMENT DATA RECORD V1.4, RO-A-OSINAC-2-AST1-STEINSFLYBY-V1.4, ESA Planetary Science Archive and NASA Planetary Data System, 2011."
- [7] Cheng, A., NEW HORIZONS Raw LORRI PLUTO CRUISE V1.0, NH-X-LORRI-2-PLUTO CRUISE-V1.0, NASA Planetary Data System, 2015.
- [8] A. Kiely and M. Klimesh, "The ICER Progressive Wavelet Image Compressor," IPN Progress Report, vol. 42, no. 155, p. 1, Nov. 2003.
- [9] C. McKinney et al., "Context cameras for the Orbiting Carbon

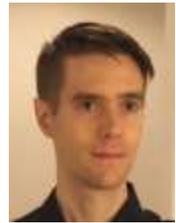
Observatory 3 (OCO-3) instrument", IEEE Aerospace conference, 2018:

<https://ieeexplore.ieee.org/document/8396759>

BIOGRAPHY



Jack Lightholder received a B.S. in Computer Science from Arizona State University in 2016. He is a member of the Machine Learning and Instrument Autonomy Group at NASA's Jet Propulsion Laboratory. Jack supports a variety of missions in concept, development and operational phases, including instrument flight software development for the Near Earth Asteroid CubeSat mission and engineering camera operations for the Mars Science Laboratory, Curiosity.



David R. Thompson is a research technologist with the Jet Propulsion Laboratory, California Institute of Technology, in Pasadena, CA. He is Investigation Scientist for NASA's AVIRIS imaging spectrometers, Instrument Scientist for the EMIT mission, and a Group Lead in JPL's Imaging Spectroscopy Group. He received a Ph.D. in Robotics from the Robotics Institute at Carnegie Mellon University in 2008.



Julie Castillo-Rogez is a planetary scientist with the Jet Propulsion Laboratory, California Institute of Technology, in Pasadena, CA. She is currently the Project Scientist and Deputy Principal Investigator for NASA's Dawn mission. She is also the Science PI for the Near Earth Asteroid Scout mission developed jointly by Marshall Space Flight Center and the Jet Propulsion Laboratory. She received a Ph.D. in Geophysics from U. of Nantes, France, in 2001.



Christophe Basset is an Instrument System Engineer in the Flight Instrument Detectors and Camera Systems group at JPL. He joined the CMOS imaging group at JPL in 2000 to research computational imagers for his doctoral thesis at the California Institute of Technology. He received his Ph.D. in 2007. He is now involved in the support of the development of imaging systems and other instruments for the Mars 2020 program and the NEA Scout CubeSat.