The Sensor Test for Orion RelNav Risk Mitigation (STORRM) Development Test Objective

John A. Christian, [*] Heather Hinkel, [†] and Christopher N. D'Souza [‡]

NASA Johnson Space Center, Houston, TX, 77058

and

Sean Maguire[§] and Mogi Patangan [¶]

Jacobs Technology, Houston, TX, 77058

The Sensor Test for Orion Relative-Navigation Risk Mitigation (STORRM) Development Test Objective (DTO) flew aboard the Space Shuttle Endeavour on STS-134 in May-June 2011, and was designed to characterize the performance of the flash LIDAR and docking camera being developed for the Orion Multi-Purpose Crew Vehicle. The flash LIDAR, called the Vision Navigation Sensor (VNS), will be the primary navigation instrument used by the Orion vehicle during rendezvous, proximity operations, and docking. The DC will be used by the Orion crew for piloting cues during docking. This paper provides an overview of the STORRM test objectives and the concept of operations. It continues with a description of STORRM's major hardware components, which include the VNS, docking camera, and supporting avionics. Next, an overview of crew and analyst training activities will describe how the STORRM team prepared for flight. Then an overview of in-flight data collection and analysis is presented. Key findings and results from this project are summarized. Finally, the paper concludes with lessons learned from the STORRM DTO.

I. Introduction

The Orion/Multi-Purpose Crew Vehicle, hereafter referred to as Orion, is a new crewed vehicle being designed by NASA and Lockheed Martin to ferry humans to/from the International Space Station (ISS) and other possible exploration destinations. One particularly challenging flight phase of the Orion mission is rendezvous, proximity operations, and docking (RPOD). The unique requirements and constraints of the Orion vehicle make off-the-shelf relative navigation sensors inadequate for achieving the required relative navigation performance. Therefore, a new relative navigation sensor is being developed. The Vision Navigation Sensor (VNS) is a flash LIDAR and will be the primary navigation instrument used by the Orion during the RPOD flight phase. Past experience suggests that on-orbit testing is critical to gaining confidence that the VNS will perform as expected. In fact, as is shown in Fig. 1, many previous relative navigation sensors have experienced difficulties on-orbit that were not predicted by ground testing. This need led the Orion GNC team to propose the Sensor Test for Orion Relative-Navigation Risk Mitigation (STORRM) Development Test Objective (DTO).The STORRM DTO flew aboard the Space Shuttle Endeavour on STS-134 (16 May - 1 June 2011), and tested Orion's VNS and high-resolution Docking Camera (DC) in an on-orbit environment. The objective was to collect data from these sensors to mitigate the loss-of-mission risk in upcoming Orion test flights and to increase the technology readiness level (TRL) of these sensors.

---

[*]Engineer, GN&C Autonomous Flight Systems Branch, and AIAA Member.
[†]STORRM Principal Investigator, GN&C Autonomous Flight Systems Branch.
[‡]Deputy Branch Chief, GN&C Autonomous Flight Systems Branch, and AIAA Member.
[§]GN&C Analyst, On-Orbit GN&C, Engineering & Science Contract Group.
[¶]GN&C Analyst, On-Orbit GN&C, Engineering & Science Contract Group.

American Institute of Aeronautics and Astronautics

Condition	Sensor												
	SVS	TCS	Rendezvous Radar	Star Tracker	ESA's RVS	HHL	OPTEC LIDAR	VGS/AVGS	Sandia LDRI	Neptec LCS	Kurs	RGPS-ESA	RGPS RME
Orbital lighting/reflections/shadows	X			X		X							
Background interference				X									
Longer range performance		X						X					
Acquisition range			X	X			X						
Higher velocities at longer range					X								
Modeling of real target vehicle in lab								X					
Potential hot spots			X	X	X	X							
Obscuration		X											
Multi-path												X	
Long-term exposure of targets to space environment		X											
Vehicle dynamics		X				X							
Dynamics of transitions during prox ops trajectory													
Performance in space environment		X			X			X				X	X
Coupling between dynamic and environmental effects								X					X
Sensor moding/transitioning								X				X	X

Figure 1. Historical survey of difficulties experienced by previous relative navigation sensors.

As of the writing of this paper, STORRM has completed its mission onboard Endeavour on STS-134. In general, STORRM system performance during rendezvous, docking, undocking, flyaround and re-rendezvous was excellent! Some minor issues were encountered during rendezvous, but timely response from the STORRM team put onboard contingency procedures in motion resulting in minimal disruption to STORRM activities. Unfortunately prior to undock, it was discovered that the Data Recorder Unit (DRU) which records DC images lost the ability to initialize, resulting in the inability to operate the DC or record its images. The cause of this issue will be investigated following the retrieval of all the data off the memory boards within both DRUs. STORRM collected a total of 368 GBytes of VNS data and 232 GB of DC data on-orbit. The STORRM team was able to observe through periodic data snapshots that the VNS performed as expected in terms of moding and laser firing. The VNS performed phenomenally in imaging the ISS, and exceeded expectations for its long-range acquisition of the ISS. The team was also able to see magnificent images from the docking camera in real-time. Details on data analysis are included in this paper.

II.   Mission Objectives & Concept of Operations

The overall STORRM Mission Objectives were to test the on-orbit performance of the VNS and the DC on three different trajectories: (1) a nominal Orbiter RPOD trajectory to the ISS,[1] (2) a nominal Orbiter undocking trajectory from the ISS, and (3) an Orion-like flight rendezvous/proximity operations trajectory to the ISS. To meet the third objective the Space Shuttle separated from the ISS and then performed an unprecedented re-rendezvous with the ISS to approximately 956 feet. The Undock and re-rendezvous profile is shown in Figure 2.

The mission objectives are achieved by meeting the flight test objectives identified for each of the sensors. A summary of these flight test objectives are shown in Table 1 for the VNS, and in Table 2 for the DC.

To achieve the flight test objectives, plans were put into place with the Shuttle Program to allow for command and control of the sensors and recording rates from a laptop computer in the crew cabin. Various parameters that will eventually be controlled internally to the sensors were externally controlled for STORRM. To facilitate the commands to the sensors, a STORRM Software Application (SSA) was developed. This laptop application was monitored by Mission Specialist Drew Feustel during the mission. One of the goals for the SSA was to minimize crew time and interaction, which led to the implementation of a range-based control of the sensor settings. The range was obtained from Orbiter onboard filtered range at far ranges then from the Orbiter laser sensor, the Trajectory Control Sensor (TCS) at closer ranges. The TCS is a scanning LIDAR system used by the Shuttle during proximity-operations and docking.[1] The data
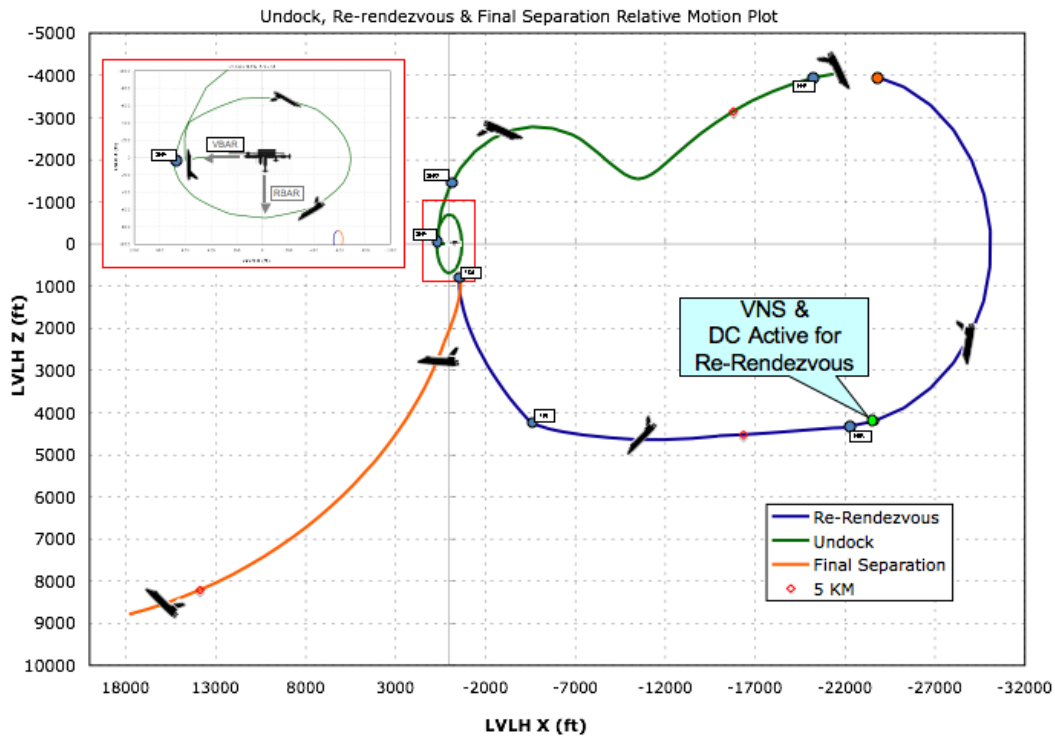
American Institute of Aeronautics and Astronautics

Figure 2. Relative Motion Profile for the STORRM Undock and Re-Rendezvous

collected from the TCS will be used as "truth" measurements for comparison with the VNS data.

Primary STORRM data collection opportunities took place on flight day 3 (day of rendezvous with ISS) and flight day 15 (day of undock) of the STS-134 mission. Preceding the first data collection, a special procedure and software phase was executed. This procedure was called "Tools Checkout." Tools Checkout consists of a timed set of operations which mode the sensors and data recorders in a predetermined sequence. Completion of Tools Checkout assures the proper function of the STORRM equipment and that all the associated connections are properly made from onboard through the ground. A second Tools Checkout was planned prior to Undock, however, the activity was deleted from the timeline due to the DRU issue.

Data collection for the STORRM experiment was divided in to phases: Rendezvous, Undock, Re-

Table 1. Summary of VNS objectives.

Objective	Primary	Secondary
Characterize the ISS in the VNS wavelength (hot spots, reflectance map, etc.)	X	
Determine initial acquisition range	X	
Prove operation with large relative velocities	X	
Collect data through known break track and reacquire conditions	X	
Prove tracking through accelerations while maneuvering	X	
Demonstrate transition between 3-DoF and pose modes	X	
Allow for overlap with ground testing facilities	X	
Collect data to support pose calculation	X	
Characterize the ISS PMA2 augmented docking target	X	
Determine range of pose acquisition		X
Prove tracking through TPI-type maneuver		X
Long range "loss of signal"		X

American Institute of Aeronautics and Astronautics

Table 2. Summary of DC objectives.

Objective	Primary	Secondary
Prove operation in harsh orbital lighting conditions	X	
Collect images to support assessment of validity for piloting cues	X	
Assess DC as a potential backup to star tracker for mid-range bearings		X
Collect images to support Natural Feature Image Recognition analysis		X

Rendezvous and Separation. The phases, which are bounded by significant environmental or operational events, represent the highest level of decomposition of the mission timeline. The Orbiter trajectory, look-angle to ISS, ISS reflector assets, orbital lighting, sensor configuration, and various other parameters are unique to each phase. These differences provide a variety of collection opportunities from which to better evaluate the sensors. Astronaut action is required to place the SSA into the proper phase of data collection. Within each phase of data collection, the sensor settings were automatically controlled through range bins. Upon entering each range bin, a set of commands were sent to the sensors and data recorders to properly configure the system for data collection over that range. Sensor settings such as VNS gain, detector settings, field-of-view (FOV) size, and DC data recording rate are examples of range-based commands. Within each range bin, timed commanding also took place. The timed commands control the region of interest upon which the DC calculated automatic gain and exposure settings. This allowed the STORRM team to collect valuable performance data on different region of interest (ROI) sizes that will feed into Orion decisions for DC operations. Two of the three ROI sizes were tested during Shuttle rendezvous and docking. The third ROI was intended to be tested during the flyaround, however, was not tested due to the DRU failure.

An early decision made on the STORRM project was that during STORRM operations, a ground support team, located in the Mission Control Center (MCC) would be able to monitor certain, critical health and status items. In order to accommodate this request, the SSA rendered a secondary screen, transparent to the crew operator, containing these particular items. This secondary screen was sent to the ground over the Shuttle Sequential Still Video (SSV) system. The STORRM ground team was able to monitor data at a 40-80 second update rate.

III.   Overview of STORRM Hardware

The major components on the STORRM experiment were the Sensor Enclosure Assembly (SEA), Avionics Enclosure Assembly (AEA), Payload General Support Computer (PGSC), Ground Operations Station (GOS), Orbiter Crew, and Docking Target Reflective Elements. A high level diagram of system data flow is shown in Fig. 3. Each of these components and the major subsystems are described in the following paragraphs.

The SEA was located on the Orbiter Docking System (ODS) Truss adjacent to the TCS, as shown in Fig. 4. The SEA houses the STORRM sensors, the VNS and the DC, and keep-alive heaters. These sensors were designed and built by Ball Aerospace Technologies Corporation in Boulder, Colorado. More detail on the specifications of these sensors may be found in an earlier paper by Miller, et al.[2]

The overarching goal of the experiment was to test the VNS and DC in the same environment that the sensors would experience on the first Orion rendezvous to the ISS. On Orion, these sensors are expected to be located in the hatch window of the crew module. However, at the outset of the STORRM design phase the Orion docking hatch design was still in flux, and providing an enclosure with a "shirtsleeve environment" in the Orbiter Payload Bay was not realistic. Because of this, the following factors were considered early to set the path of the SEA design:

1. Location of the SEA in the Orbiter Payload Bay: The SEA was co-located with the TCS for better truth comparison and to avoid line-of-sight limitations. In the early design phases of STORRM, trade studies were conducted to determine which mounting location would allow sufficient data to be taken on the ISS docking axis for comparison with ground testing. These studies showed that the STORRM Reflective Elements would be in the VNS field of view until 13.7 meters.
2. Pressurization of the SEA: The STORRM project resided within a class of Orbiter payloads that were not permitted to use active cooling systems. Because only passive cooling systems could be used,

it was found that the desired "shirtsleeve environment" was not attainable. Therefore there was no justification to design the SEA as a pressurized volume.

3. Number of panes of glass on the SEA: Much consideration was given to the number of panes of glass and types of coatings used on the SEA. At the time of this decision, the Orion Docking Hatch had three panes of glass with anti-reflective coating that were approximately 25 cm in diameter. Mass and volume constraints as well as uncertainty in the final Orion window design led the team to make the decision that a single pane of glass was to be used on the SEA.

Unlike many other payloads in the Orbiter Payload Bay, the SEA was not covered in Multi-layer Insulation (MLI) blanketing. This aspect of the design was forced by a 'thermal isolation' requirement levied by the Shuttle Program. This requirement stated that the SEA could not use the ODS Truss or TCS as a path of thermal conductance. Indeed, the STORRM project was very sensitive to not perturbing TCS operation, so heat rejecting materials and keep-alive heaters were chosen as a solution.

The VNS housing was highly integrated into the design of the overall the SEA. The VNS was cantilevered on the +Z axis side of the SEA in the Orbiter Body Frame. Additionally, the VNS connectors and -Z side (Orbiter Body Frame) were designed such that the unit could be easily accessed without a complete disassembly of the SEA.[1] Internally to the VNS housing resides a 256×256 pixel focal plane array. The

---

[1]It should be noted that the packaging strategies used for the VNS within the SEA were unique to the STORRM DTO project.
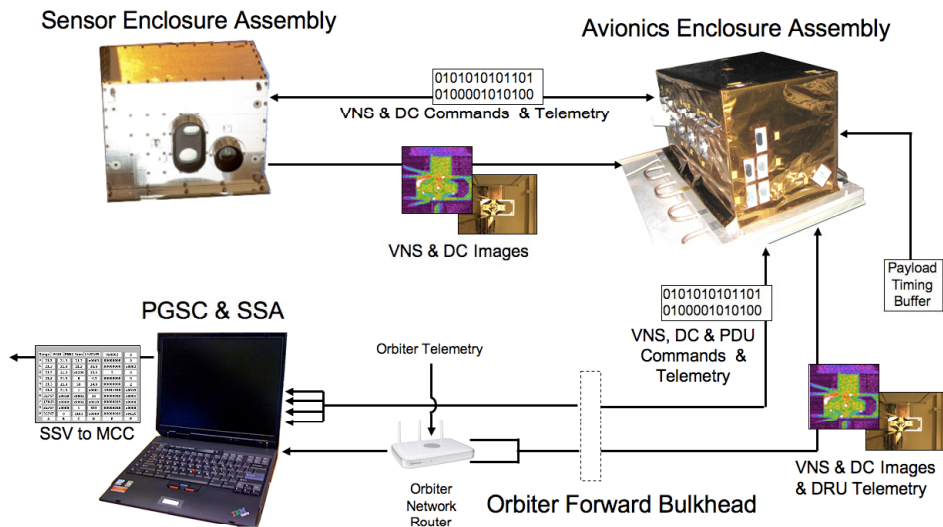

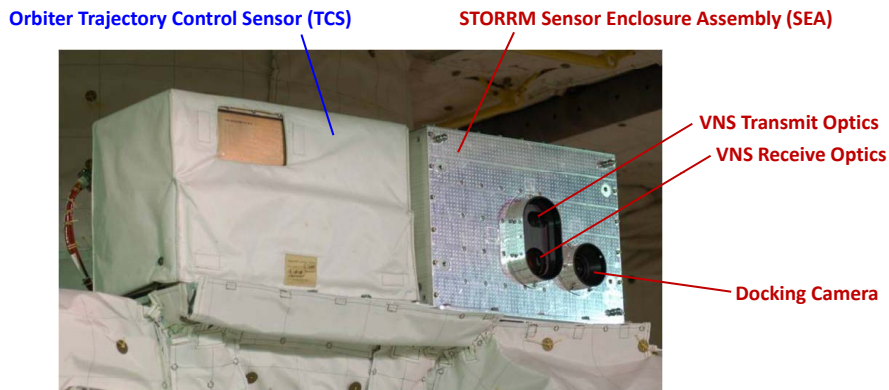Figure 3. STORRM System Components and Data Flow Diagram

Figure 4. Image of Orbiter TCS and STORRM SEA as mounted on the ODS.

American Institute of Aeronautics and Astronautics

VNS has a variable field of illumination that can be selected to be 12 degrees or 20 degrees. Additionally the VNS FOV is 20 degrees, full cone.[2] An exploded view of the VNS is provided in Fig. 5.

The DC was also cantilevered on the +Z axis, Orbiter Body Frame, of the SEA. Although the DC focal plane array has many more pixels, at 2592×1944 pixels, the DC is much smaller in mass and volume than the VNS and was not a driving factor in the thermal budgets or power consumption of the system.[2] The STORRM DC does not use all of these pixels, and a typical image has dimensions of 2400×1800 pixels. The placement of the docking camera within the SEA was also benefited by the SEAs location on the ODS truss. This location allowed the the docking camera to be positioned such that the camera boresight was on the zero station of the Orbiter Body Frame Y-axis. This placement on the Orbiter centerline allowed for simple, in-plane, comparisons of Orbiter Centerline Camera Images and STORRM Docking Camera Images. The DC has a 38.7×29.3 degree rectangular FOV.[2]

Cables from the back of the SEA were bundled and routed along a cable tray fastened to the ODS Truss. These cables carried power, sensor commands, telemetry, and high-speed serial image data. From the cable tray on the truss the cables followed standard paths until they reached the next hardware element along the STORRM data path, the Avionics Enclosure Assembly.

The Avionics Enclosure Assembly (AEA) was located on the port sidewall of Payload Bay 3. The AEA consists of two Data Recorder Units (DRUs) that receive and store data from the sensors in the SEA, and a Power Distribution Unit (PDU) that distributes power to the SEA and AEA. The AEA also passes sensor configuration commands and telemetry to/from the PGSC and SEA.

The Power distribution Unit (PDU) served two important roles in STORRM. The first was to open and close power relays on command, allowing systems to be powered automatically via command to a set of computer controlled relays. Physical inhibit switches for PDU power were also available to the crew. The second role was that the PDU was the central hub of all the low-level system health and status telemetry. The PDU provided 72 telemetry items that described currents, voltages, temperatures, and op-codes that represented the state of the system.

The next sub-system within the AEA was the DRUs. The primary responsibility of each DRU was to receive and store raw images from the sensor at extremely high data rates. The bandwidth requirements and
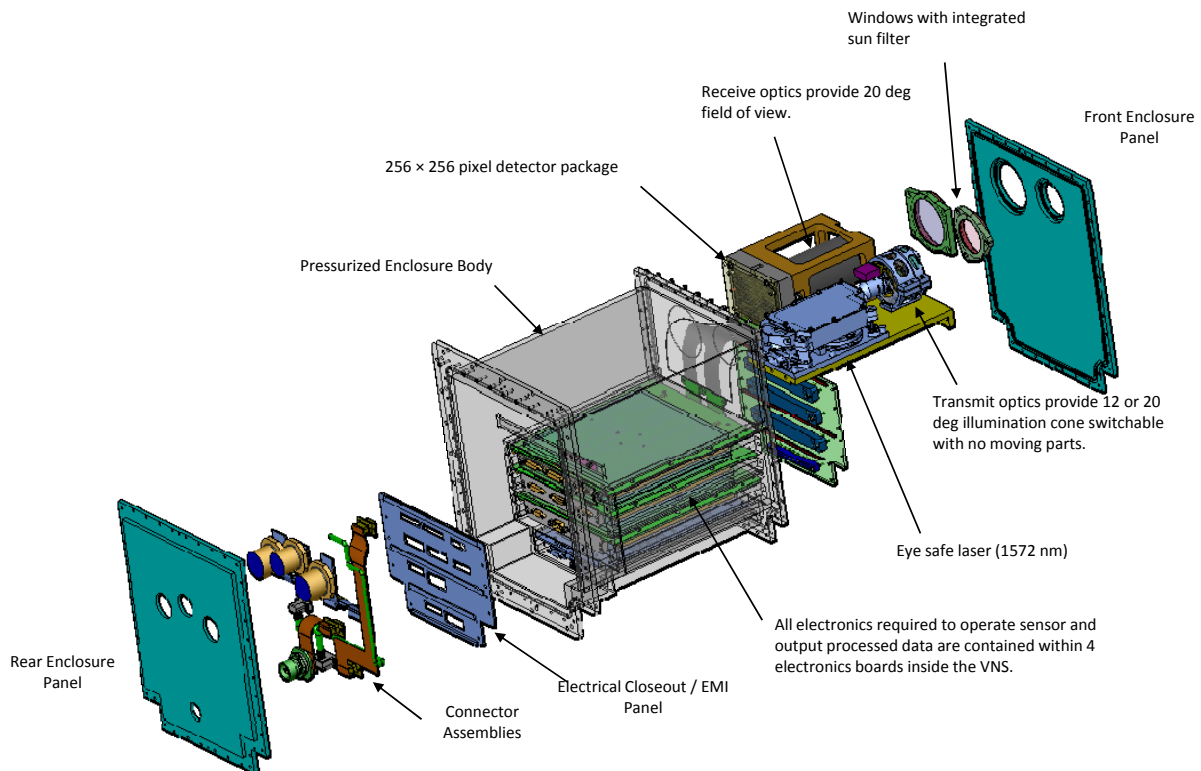

Figure 5. Exploded view of STORRM VNS.

American Institute of Aeronautics and Astronautics

Table 3. STORRM bandwidth requirements and expected data volume.

Sensor	Bandwidth (MB/sec)	Expected Data Volume (GB)
VNS	15.9	432
Docking Camera	29.1	284

expected data volumes are summarized in Table 3. The large volume and high data rate needs for each sensor was the primary driving decision to place the AEA on the Bay 3 sidewall location of the Orbiter. The Orbiter forward bulkhead is not equipped with infrastructure that could support those data rates. Modifications to the Orbiter forward bulkhead were briefly entertained as an option in early design discussions, but ultimately rejected because of costly leak checks that would need to be re-performed on the vehicle. The DRUs were developed with solid state memory technology and they were connected to the sensors in a single-string fashion; each DRU being responsible for recording from only one of the sensors. The two different DRUs were designated "DRU1" and "DRU3." DRU1 was used to record VNS data, and DRU3 was used to record DC data.

The secondary responsibility of the DRUs was to forward a single image frame, or 'snapshot', of data to the STORRM PGSC upon request. Alternatively the DRUs could be placed in a mode called Data Retrieval Mode in which raw blocks of memory from the DRUs could be copied via a network connection to the STORRM Payload General Support Computer. Details of the types of data received from the DRUs will be discussed more detail in section V. The DRUs provided their own set of telemetry and also acted as a pass-through for commands and telemetry to/from the STORRM PGSC to the sensors. Lastly, Orbiter time was received by the DRUs via the Payload Bay Timing Buffer. Time was logged and regular time sync signals were provided to the sensors.

Cables were routed from the AEA through the Orbiter forward bulkhead to the next node in the data path, the STORRM Payload General Support Computer (PGSC). The PGSC was located in the flight deck of the Orbiter crew cabin. This laptop computer was loaded with the STORRM Software Application (SSA) that initiated data collection from the sensors, moded the DRUs, and allowed the crew and the ground to monitor the health and status of the systems. The experiment was controlled via crew inputs to the SSA. The SSA connected to the orbiter system via seven separate paths:

1. RS-422 serial line for PDU command and telemetry
2. RS-422 serial line for VNS command and telemetry
3. RS-422 serial line for DC command and telemetry
4. TCP/IP network line for DRU1 command and telemetry
5. TCP/IP network line for DRU3 command and telemetry
6. TCP/IP network line connecting to Orbiter Onboard network (via shared hub)
7. S-Video line for near real-time telemetry observation.

Although the configuration of the STORRM PGSC was extensive, it ultimately resulted in serving its function as a nucleus of information and data transfer from the components in the payload by to the Crew and to the Ground Operations Station during the mission.

The Ground Operations Station (GOS) consisted of computer terminals located in the Mission Control Center (MCC) at the NASA Johnson Space Center in Houston, Texas, which also has access to real-time Orbiter telemetry and voice loops. The STORRM personnel monitored the sensor experiment data on the ground at the MCC. The primary vehicle used for data observation was as follows. The STORRM PGSC was configured to render a virtual screen and send the contents of that virtual screen, or extended desktop into the Orbiter Photo-TV stream via the S-Video cable. The Orbiter Photo-TV stream is downlinked, decoded and broadcast on the internal TV monitors within MCC. Using this method a table of parameters could be rendered by the STORRM PGSC and ultimately it would be visible to the GOS personnel for recording and trending.

The last piece of the system on board the Orbiter was the Mission Specialist responsible for carrying out STORRM activities. An Orbiter crew member was assigned to configure the STORRM PGSC and initiate a software application to allow data to be collected automatically. The mission specialist also configured the system as necessary when phases of the experiment concluded. Crew and ground personnel training is discussed in greater detail in section IV.

A description of the suite of STORRM hardware is incomplete without including the set of five reflective elements attached to the docking target at the Pressurized Mating Adapter 2 (PMA2) of the ISS. These reflective elements, built at the NASA Langley Research Center, were designed to be highly reflective in the VNS laser wavelength, but opaque in the TCS and visible wavelengths. Each of the STORRM reflective elements are housed in a titanium bracket that fastens to the docking target. The reflective element housings also have fiducial markings to aid in the accurate measurement of the reflector locations. The reflective material was chosen after extensive ground bake-out tests found the other, runner-up, materials were susceptible to deterioration and decay in the space environment. The placement of the reflectors was also chosen after considerable trades showed that alternative patterns did not perform as well with respect to the criteria of separation, visibility and non-ambiguous relative geometry. The reflective elements were attached to the existing ISS docking target during STS-131, as shown in Fig. 6.

Location knowledge of these reflective elements was needed to quantify the bias error contribution from reflective element location uncertainty. This information will be critical during post-flight data analysis in order to correctly evaluate VNS measurement accuracy. The reflector position knowledge requirement is 0.3 mm (1-sigma). The location of the STORRM reflective elements was obtained using an imagery analysis technique called photogrammetry. In photogrammetry, a number of images of an object are taken from various vantage points and the information in these images is used to reconstruct the geometric properties of the observed objects. In the case of STORRM, this information is used to determine the exact, as-installed location of the reflectors. Precision measurements were enabled by using the fiducial markings on the reflective element housings, along with fiducial markings on two scale bars that were taped to the docking target prior to capturing the photogrammetry images. Two photogrammetry sessions were performed during the STS-134 mission: one right after docking before removal of the stand-off cross and the other after installation of the stand-off cross before undocking. In each of these photogrammetry sessions, the crew took a photo at each of nine locations: eight photos at 45 degree intervals around docking target and one photo centered on the docking target.

IV. Crew and Analyst Training Activities

The STORRM experiment interfaced with numerous Orbiter systems and levied unique trajectory requirements on the Orbiter. Because STORRM was highly embedded in the STS-134 mission plan, the training activities for the experiment required a three pronged strategy and a flexible testbed on which to insert malfunction scenarios. The first strategy in training was one-on-one sessions with the mission specialist. The second was a set of Stand-alone Simulations with all crew members. The last last method was through Joint Integrated Simulations with all the STS-134 ground personnel. Each method used for training depended heavily on a software simulation called the STORRM Avionics Module. The Avionics Module was a functional equivalent of all of the physical and electrical interfaces to/from the STORRM Software Application (SSA). These training methods and the role of the Avionics Module will now be described in greater detail.

a) Approximate location of STORRM reflectors on docking target.

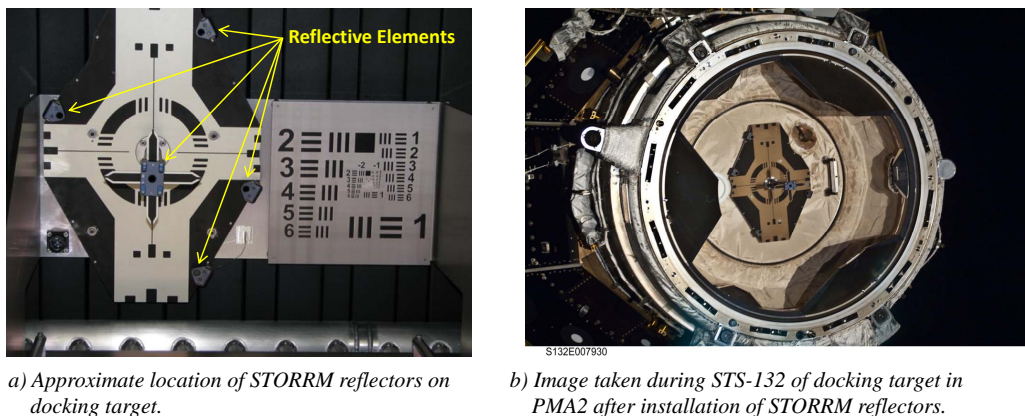b) Image taken during STS-132 of docking target in PMA2 after installation of STORRM reflectors.

Figure 6. Location of the five STORRM reflectors on the PMA2 docking target.

The one-on-one sessions with the Mission Specialist were the primary opportunity to explain the experiment background, LIDAR operation, and the hardware architecture. Additionally, sensitivities relating to the mission objectives and operational timeline were conveyed in these sessions. There were 6 one-on-one sessions conducted in total. The first sessions dealt with long-lead items such as gathering feedback on the graphical user interface preferences on the SSA. The final one-on-one sessions covered "worst-case" scenarios in which the Crew would be required to deploy a low-level set of commands to the VNS and associated hardware. The Avionics Module was extremely useful in one-on-one sessions because of it's portability; trainers could look over the shoulder of the crew and watch the Mission Specialist configure the laptop and launch the application. Furthermore, "what-if" scenarios were carried out in real-time with the controlled and known simulation environment afforded by the Avionics Module. This portability proved critical to mission success during STS-134. The Avionics Module was set up in the MCC and utilized numerous times to check out specific screen layouts, assist in writing and verifying quick-turnaround crew procedures, and to allow the STORRM main mission communication point-of-contact team to run the procedures for familiarity.

The second training method used was Stand-alone Simulations. Stand-alone Simulations were crew training activities that involved all the Shuttle Crew, executing the Flight Plan as it would happen on the Orbiter during the mission. Stand-alone Simulations took place in the fixed base Shuttle Mission Simulator (SMS) at the NASA Johnson Space Center. The curriculum of these training sessions focused on execution of the STORRM procedures and proper interface with the STORRM PGSC. For stand-alone simulations, the STORRM PGSC was located on the flight deck and the Avionics Module was located out of sight on the simulator mid-deck. The Avionics Module in these training exercises allowed the trainers indirect insight into the crew's interaction with the SSA. Stand-alone training sessions were the primary opportunity for inserting failures that impacted the STORRM experiment. In addition to direct malfunctions of STORRM hardware, other inserted malfunctions included failures of Orbiter relative navigation instruments, thrusters and electrical systems. The stand-alone training environment was an excellent platform, providing insight into the intersystem dependencies and sensitizing the STORRM team of actual crew workload issues.

The third training approach, the Joint Integrated Simulations (JISs), were large scale training opportunities for the entire ensemble of Flight Controllers and analysts involved in the mission. Because of the integrated nature of the JISs, training objectives were weighted equally among the Crew, Flight Controllers and other ground support teams. STORRM ground personnel used these opportunities to become accustomed to the voice loop protocols, the Sequential Still Video (SSV) telemetry path and various Mission Control resources provided for situational awareness. The Avionics Module was used in the same fashion that it was used for in the Stand-alone Simulations. Joint Integrated Simulations provided end-to-end training for the STS-134 flight and exposed STORRM GOS sensitivities that would have otherwise gone undetected.

The biggest challenge in Crew training precipitated from the fact that the Crew procedures were developed in parallel with training activities, hardware testing, and software development. This parallel development environment led to significant iteration on, and accumulation of, contingency procedures. Furthermore, the nuances of STORRMs use of resources within the greater STS-134 mission were discovered via the various simulations and training exercises. The result of this parallel development led the team to choose a procedure logic that was highly dependent on particular on-orbit scenarios. The lessons learned in procedure development will be discussed in further detail in Section IX.

In addition to the suite of training activities described above, STORRM was given the opportunity to demonstrate the operation of the STORRM flight hardware and software to the crew members of STS-134. This was unique because the original training plans were based on the assumption that the actual flight hardware would not be available. This training experience was extremely valuable to the crew in that it provided insight into the system timing and sensor image quality that was unachievable in other ground simulations. Ultimately the extensive training left the Orbiter Crew, Mission Control Flight Controllers STORRM team well prepared for the unique mission plan. In a post-flight crew de-brief, Drew Feustel stated that "training was great... there was nothing I felt unprepared for with respect to STORRM." The crew felt the amount of time spent up-front to make the procedures straightforward and explicit was justified, as there were a number of times that Feustel was not available and STORRM procedures were carried out by astronaut Mark Kelly who had not received STORRM-specific training. Additionally, the time spent up front allowed the crew to successfully perform the a new procedure that was uplinked before undocking.

V.   Data Storage, Retrieval & Downlink

The entire design of the STORRM experiment was based around two fundamental needs. The first was that the ground personnel should have insight into the system operation in near real-time.[1] The second fundamental need was that the system should allow the flexibility to a change configuration of the sensors during the mission. Prior experience has shown that simple insight into system status and the ability to reconfigure a setting can be the determining factor in avoiding failure. This section will describe the varieties of data that the system produced, how it was stored, and how the team went about obtaining and decomposing that data.

As described in Section III, during Orbiter Rendezvous and Proximity Operations, raw data from the VNS and DC were stored as binary blocks on DRUs in the payload bay. This process, called Data Recording Mode, also yielded three other types of data:

1. Single sensor image frames, or 'snapshots', forwarded to the STORRM PGSC at 30 second, or 60 second intervals for VNS and DC, respectively.
2. Command and telemetry logs stored from each subsystem on the STORRM PGSC in a Log folder.
3. A set of critical telemetry items rendered to the STORRM PGSC virtual screen, interleaved into the Orbiter Sequential Still Video (SSV) and broadcast to TV monitors in the MCC.

It is important to note a few of the implications of these different types of data. First, because the STORRM PGSC was connected to the Orbiter Network the first two types of data could be available for downlink when the communication link was present and not in use by another Orbiter system. Next the broadcast SSV screen allowed multiple groups in different areas of the MCC building to view the data simultaneously. The STORRM team took advantage of this and enlisted a group of interested young engineers to oversee individual data elements on the SSV screen and track time trends of the data. Finally, by trending the data on a set of shared on-line spreadsheets, STORRM personnel could monitor the status of the experiment from anywhere with an internet connection. After thorough testing, training, and adjustment, this system of SSV and on-line trending of critical telemetry met the needs of having near real-time insight to all critical system parameters without any modifications to the Orbiter standard telemetry stream.

The SSA was designed with the capability to change the sensor settings. The desire to be flexible to change sensor settings gave rise to the next type of data needed for analysis during the mission: DRU block files. DRU block files contain sequential images from the sensor at the full data rates of the sensors. While snapshot images allowed some qualitative assessments on the data, the block files contain sequential images that allow detailed analysis to be conducted (data analysis itself is discussed in section VI). The DRU block files are attained by a retrieve and downlink process, called Data Retrieval Mode. Data Retrieval Mode was employed in the ISS-Attached phase of STS-134. In Data Retrieval Mode, the SSA first copied unprocessed blocks of memory from the data storage drives in the Payload Bay on to the PGSC local disk drive. The blocks of data were then downlinked from the PGSC disk drive to data storage in MCC via the Orbiter on-board network. Once the block files were copied to data storage drives in MCC they were parsed into individual frames and distributed to analysts.

The difficulty in the block file retrieval and downlink process was that these tasks took a substantial amount of time to execute. The need then arose to correlate the distance to the ISS with specific write addresses from each DRU. Ultimately data from the Rendezvous & Proximity Operations Program (RPOP) onboard the Orbiter was used to provide the rendezvous Range vs. Time profile. Time stamps from RPOP were converted into the epoch used in the STORRM log files and then the DRU write addresses of interest were interpolated for a set of 35 different events of particular interest. Shortly after the block file data was downlinked, parsed and analyzed, it was determined that the estimates of the data retrieved was within 15 seconds of each desired event; well within the acceptable error.

With a combination of relatively fast SSV data, medium data rate snapshots, system logs, and the sequential frames contained in the block files, the information received at all phases during the mission gave the STORRM subsystem leads and data analysts confidence that the system was well understood and that the mission objectives were protected.

---

[1]Although near 'real-time' was never explicitly defined, it was largely accepted to be on the order of 2-3 minutes. Actual system performance showed data latencies sometimes up to 5 minutes through the Orbiter Sequential Still Video path.

American Institute of Aeronautics and Astronautics

VI. Prediction of Reflector Visibility

The prediction of reflector visibility allowed the STORRM team to refine the data retrieval-downlink requests during the docked timeframe in verifying that the data being retrieved had available reflectors in the VNS FOV. More importantly, it allowed data analysts to verify that the outputted VNS centroids were that of reflectors rather than spurious returns in the VNS intensity image. The predicted visible reflectors were obtained using the Reflector Visibility Tool.

The Reflector Visibility Tool made use of the ISS and Orbiter relative position and attitude obtained from RPOP to check for line of sight visibility between reflectors and the VNS. The tool accounted for reflector obscuration and the narrow and wide VNS laser field of illumination modes (recall from Section III that the field of illumination may be commanded to 12 degrees or 20 degrees). However, the tool did not account for reflector return signal strength and signal interference. Reflector locations with respect to the Space Station Analysis Coordinate System (SSACS) were obtained from the ISS 3D CAD Team. Locations of the STORRM reflective elements were calculated using the known PMA2 docking target location and the proposed installation locations for the STORRM reflective elements.

There were three flags in considering a reflector to be visible: (1) the VNS must be in the reflector's unobscured field of regard, (2) the reflector must be in the VNS 256×256 pixel image plane, and (3) the reflector must be illuminated by the VNS laser. Once these checks were satisfied, the reflector was projected on the VNS image plane and output VNS intensity image. Figure 7 shows a screenshot of the output panel in the Reflector Visibility Tool's GUI. The output panel includes the predicted reflectors in a VNS intensity image, the relative motion plot of the Orbiter with respect to the ISS CG in a LVLH reference frame, time tags, calculated and raw TCS or SV range, vehicle position and attitude inputs, and a table of visible reflectors and their predicted pixel locations.

The tool's accuracy in predicting exact reflector pixel location was dependent on several factors.

1. Vehicle state inputs from RPOP contained uncorrected data
2. Detector alignment to VNS receive optics
3. ODS pressurization effect to ODS truss alignment

The vehicle state inputs from preliminary corrected ground RPOP data could account for misalignments between the predicted and actual reflector centroids. The ISS attitude was an assumed fixed attitude, and some TCS reflector tracking biases were not yet corrected. A post-mission best estimated trajectory (BET) might correct for some of the misalignments from these inputs. The second factor was the alignment of the detector behind the VNS optics. Although the VNS boresight translational alignment was mapped to
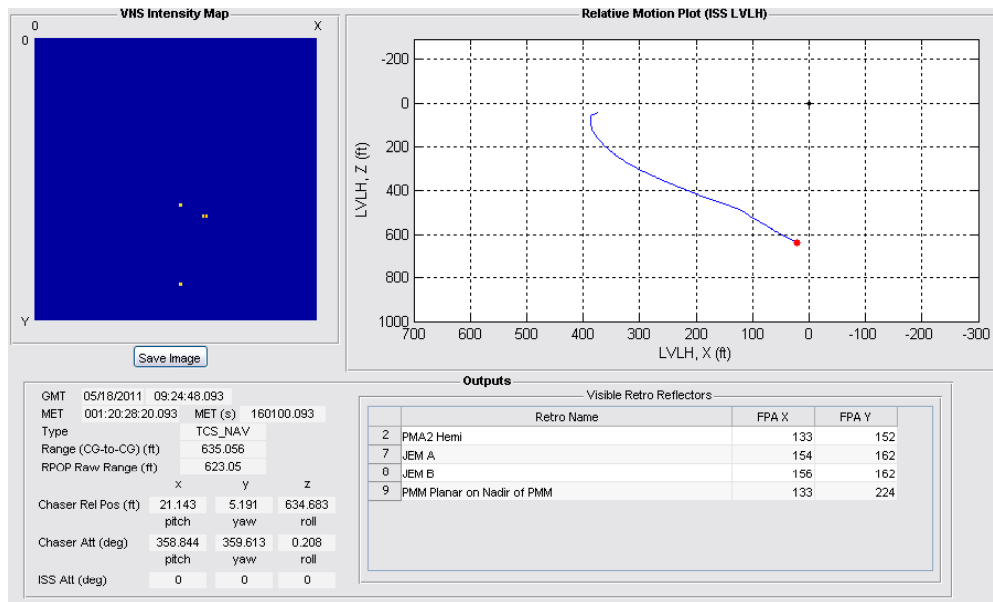

Figure 7. Reflector Visibility Tool Output Screenshot.

American Institute of Aeronautics and Astronautics

the detector, the rotational alignment was unknown for the flight. A VNS to TCS alignment analysis will provide this missing piece of information. The last factor was due to the fact that the ODS truss alignment might have changed as a function of the ODS vestibule pressurization. The ODS pressurization data can be obtained post-flight to account for this misalignment in the model.

VII.   Data Analysis

A number of tools were developed to quickly process the STORRM data collected and downlinked during the mission to provide near real-time feedback on the performance of the VNS and DC. Upon receiving the data on the ground, data integrity checks are performed and the data is parsed into a "snapshot" format. This snapshot format means that each individual VNS or DC data file contains data collected at a single time, along with the appropriate header file. This snapshot data is then fed into either the VNS Quick Look Tool or the DC Quick Look Tool.

VII.A.   Docking Camera Quick Look Tool

The DC Quick Look (DCQL) Tool is capable of interpreting raw STORRM DC snapshot data and displaying raw and processed images. A screenshot of the DCQL Tool GUI is shown in Fig. 8.

The raw DC data contained in a snapshot file is still mosaiced, so the DCQL Tool contains a demosaicing algorithm to reconstruct the color image (the user still has the option to either view a grayscale image or a color image). The STORRM DC's detector uses a color filter array with a standard Bayer pattern as shown in Fig. 9. The demosaicing algorithm in DCQL Tool is based on the work of Alleysson, Süsstrunk, and Hérault.[3] This demosaicing algorithm was chosen because of (1) simplicity in implementation, (2) algorithm speed, (3) a better preservation of high spatial-frequency signals and reduction in aliasing when compared to simple bilinear interpolation. For the convenience of the reader, some of the main points from this approach are repeated here.

Recall that the human eye is about twice as sensitive to green light as red or blue light. A RGB image may be converted to grayscale according to[4]

$$L = 0.299\,R + 0.587\,G + 0.114\,B \tag{1}$$

where L is the luminance (or grayscale intensity), R is the intensity in the red channel, G is the intensity in the green channel, and B is the intensity in the blue channel. For the sake of the demosaicing algorithm, suppose this is approximated as

$$L \approx (R + 2\,G + B)/4 \tag{2}$$

From here, we seek a 3×3 image filter that, when convolved with the mosaiced image, will produce the luminance regardless of color of the pixel on which the filter is centered. Therefore, given a 3×3 filter, F, and a $n \times m$ mosaiced image, I', we seek the luminance image (i.e. grayscale image), L,

$$L = F * I' \tag{3}$$

where $*$ is the 2D convolution operator. In other words if

$$F = \begin{bmatrix} a & b & a \\ b & c & b \\ a & b & a \end{bmatrix} \tag{4}$$

then the following relations must be true

$$L = \begin{bmatrix} a & b & a \\ b & c & b \\ a & b & a \end{bmatrix} * \begin{bmatrix} G & R & G \\ B & G & B \\ G & R & G \end{bmatrix} = \begin{bmatrix} a & b & a \\ b & c & b \\ a & b & a \end{bmatrix} * \begin{bmatrix} B & G & B \\ G & R & G \\ B & G & B \end{bmatrix} = \begin{bmatrix} a & b & a \\ b & c & b \\ a & b & a \end{bmatrix} * \begin{bmatrix} R & G & R \\ G & B & G \\ R & G & R \end{bmatrix} \tag{5}$$

We may now obtain three equations, one for each of the color channels:

$$\begin{aligned} R: & \quad 2b = c = 4a \\ G: & \quad 4a + c = 4b = 4b \\ B: & \quad 2b = 4a = c \end{aligned} \tag{6}$$

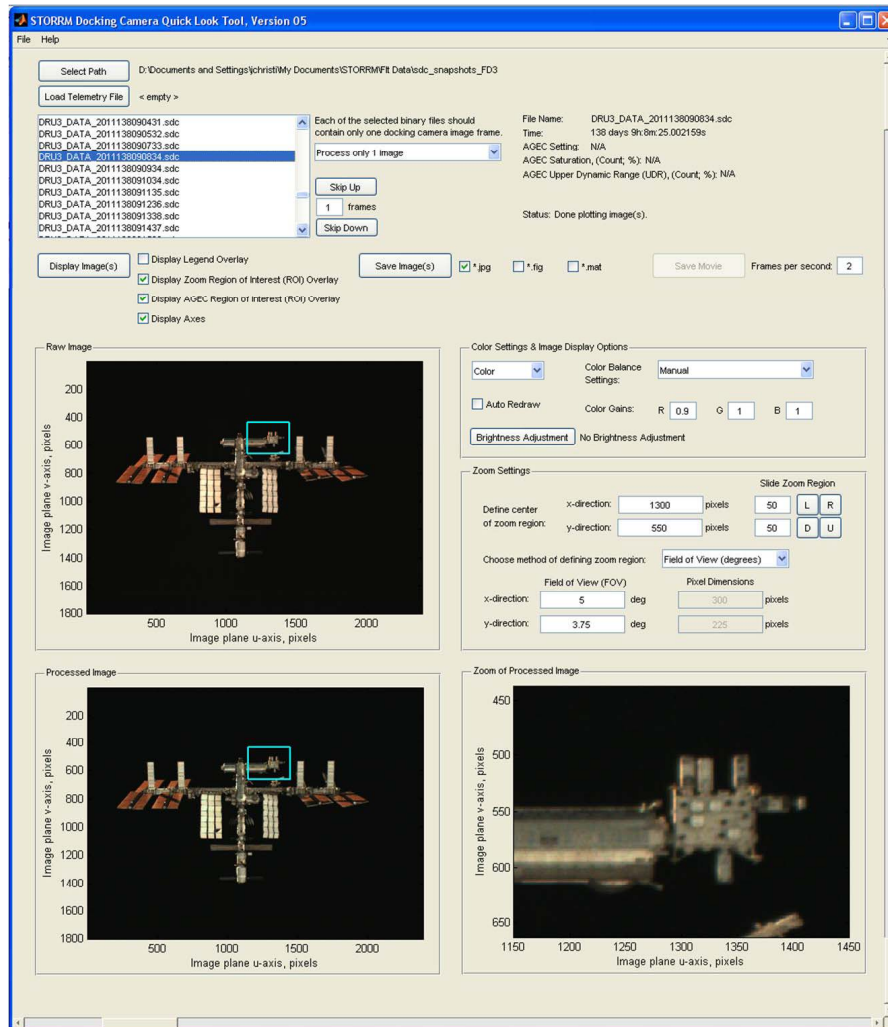American Institute of Aeronautics and Astronautics
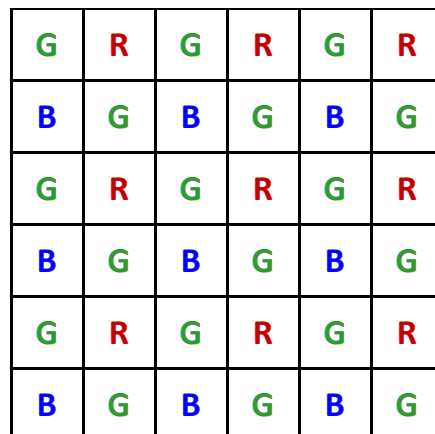
Figure 8. Screenshot of the DCQL Tool GUI.

Figure 9. Bayer pattern color filter array.

It is not surprising that the red and blue color channels produce a duplicate of the same relation. From Eq. 2, we also know that the coefficients must sum to unity,

$$4a + 4b + c = 1 \tag{7}$$

Solving the system of equations defined by Eq. 6 and Eq. 7 yields

$$
\begin{aligned}
a &= 1/16 \\
b &= 1/8 \\
c &= 1/4
\end{aligned}
\qquad \longrightarrow \qquad
F = \frac{1}{16}
\begin{bmatrix}
1 & 2 & 1 \\
2 & 4 & 2 \\
1 & 2 & 1
\end{bmatrix}
\tag{8}
$$

It is interesting to note that F is almost identical to a 3×3 Gaussian filter with a standard deviation of 0.85 pixels. The DCQL Tool uses this method for computing the grayscale image.

Alleysson, Süsstrunk, and Hérault demonstrate that the mosaiced image consists of the sum of the luminance image and a chrominance image,[3]

$$I' = L + C \tag{9}$$

where C is the chrominance image. Therefore, the chrominance image may be thought of as the result of applying a high-pass filter to the mosaiced image. One may now use bilinear interpolation to compute the two missing colors at each pixel location in the chrominance image,

$$\{C_R, C_G, C_B\} = \phi(C) \tag{10}$$

where $\phi(\ )$ is the bilinear interpolation function, C_R is the $m \times n$ red channel of the demosaiced chrominance image, C_G is the $m \times n$ green channel of the demosaiced chrominance image, and C_B is the $m \times n$ blue channel of the demosaiced chrominance image. Therefore, the RGB demosaiced image is given by

$$
\begin{aligned}
I_R &= L + C_R \\
I_G &= L + C_G \\
I_B &= L + C_B
\end{aligned}
\tag{11}
$$

This is how the DCQL Tool computes a full color image.

Once the image has been demosaiced, the DCQL Tool contains a number of image processing algorithms to adjust the color gains, brightness, and contrast. Individual adjustment of the color gains is extremely important because the STORRM DC was not white balanced under natural lighting conditions, rather it was tuned to the Orion docking lights which have a bluish tint. Therefore, adjustment of the color gains is necessary to recover a true color image.

In addition to individually varying the color gains, the DCQL Tool also has the capability to globally adjust brightness. This is done adjusting each color channel according to the common parameter b and the equation,

$$
J_c =
\begin{cases}
(1 + b)\, I_c & b \geq 0 \\
(1 - b)\, I_c + b & b < 0
\end{cases}
\tag{12}
$$

where I_c is the normalized intensity in color channel c (either R, G, or B) in the original image, J_c is the normalized intensity in color channel c in the brightness-adjusted image, and the parameter $b \in [-1\ 1]$ defines the brightness adjustment. A positive value of b creates a brighter image and a negative value of b creates a darker image. Note that there is no change in the image if $b = 0$. The term "normalized intensity" means that the intensity values have been normalized to be between 0 and 1.

Similarly, contrast is adjusted independently along each color channel according to the common parameter ξ and the equation,

$$J_c = 0.5 + (I_c - 0.5) \tan\left[(\xi + 1) \frac{\pi}{4} \right] \tag{13}$$

where J_c is the normalized intensity in color channel c in the contrast-adjusted image and the parameter $\xi \in [-1\ 1)$ defines the brightness adjustment. A positive value of ξ increases the contrast and a negative value of ξ decreases the contrast.

The DC Quick Look Tool is capable of processing and saving a user defined set of data (e.g. a single image, an arbitrary set of selected images, all the images in a target folder, etc.). Additionally, the user can

specify region of interest and the DC Quick Look Tool will provide a zoomed-in image of this region. This is particularly useful in assessing the quality of the DC for piloting cues.

Some example DC images that have been processed by the DCQL Tool are shown in Fig. 10 and Fig. 11. These images were taken during rendezvous with the ISS on flight day 3 of STS-134.
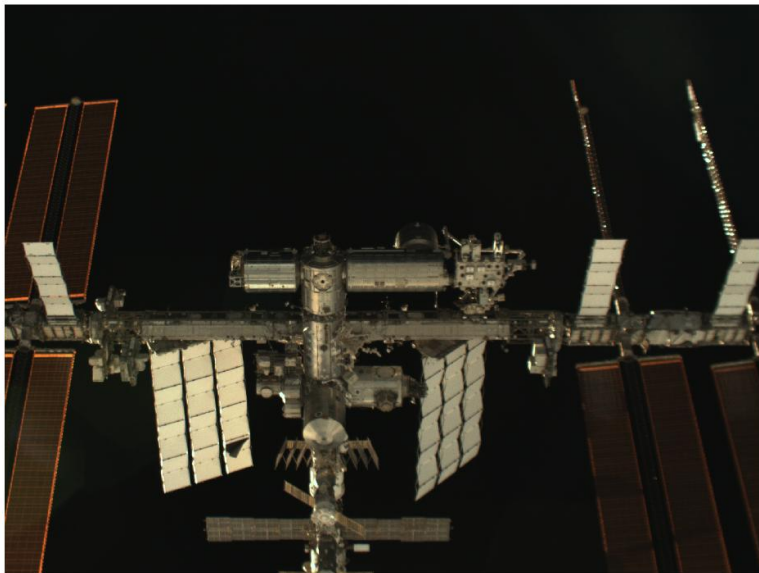

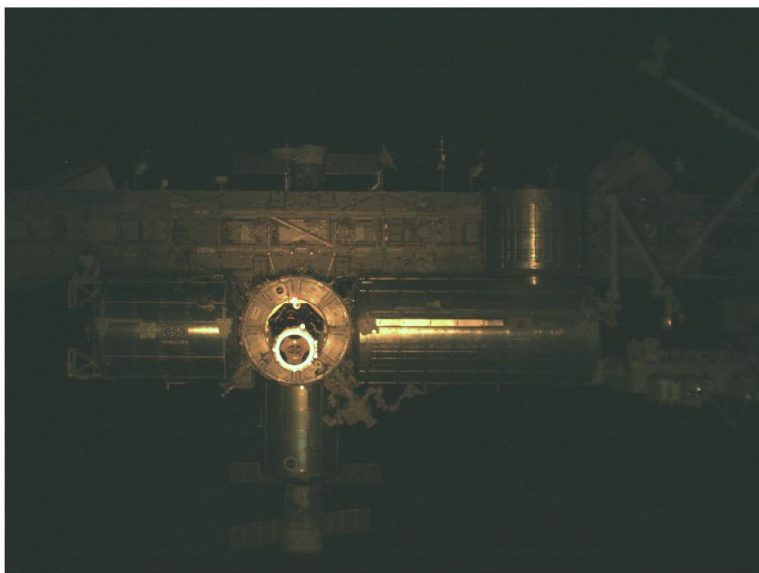**Figure 10. Example docking camera image #1. Image taken during STS-134 flight day 3 rendezvous.**

Figure 11. Example docking camera image #2. Image taken during STS-134 flight day 3 rendezvous.

VII.B. VNS Quick Look Tool

The VNS Quick Look (VNSQL) Tool is capable of interpreting raw STORRM VNS snapshot data in order to generate a 256×256 range image and intensity image. It is also capable of performing image processing on raw STORRM data to provide bearings and ranges to centroids of candidate reflectors at a user specified rate. A screenshot of the VNSQL Tool GUI is shown in Fig. 12.

Data may be processed to produce only a single set of reflector range and bearing measurements, or a large volume of data may be processed in a batch format. The ability to generate a single set of reflector
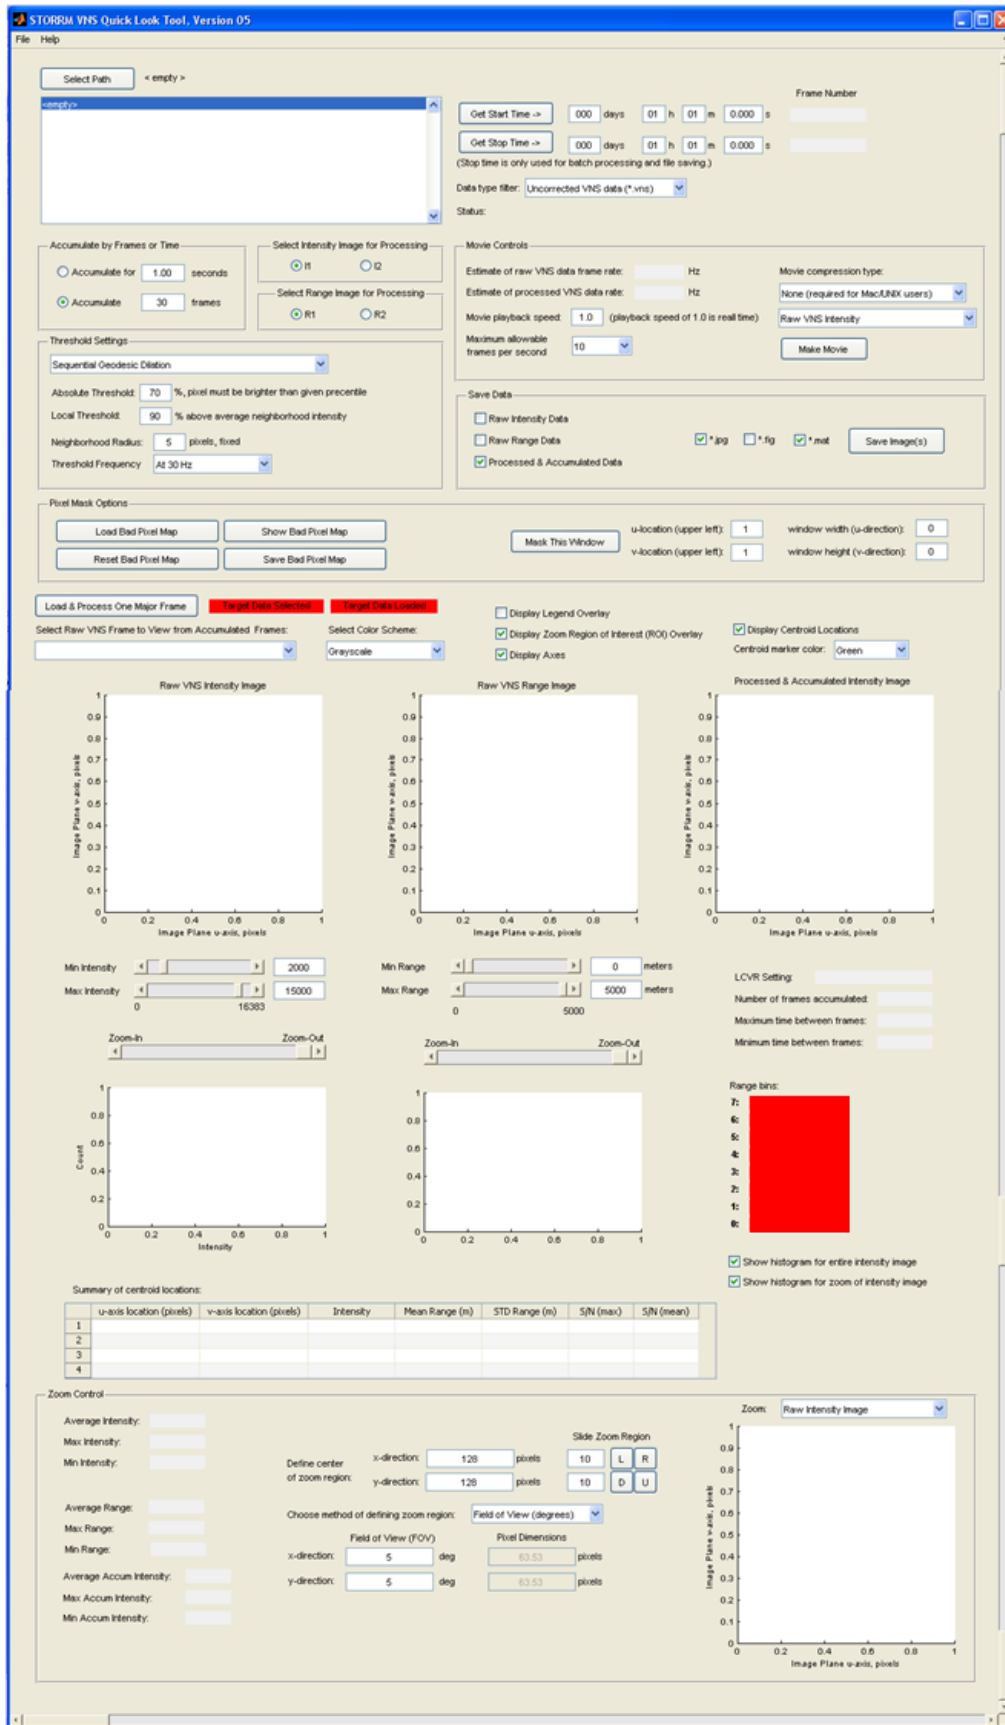
American Institute of Aeronautics and Astronautics

Figure 12. Screenshot of the VNSQL Tool GUI.

measurement (which is actually created by processing a number of STORRM VNS images) is useful in helping the analyst understand the performance of the sensor. The VNSQL contains five different methods for automatically finding centroids in a VNS image. The identification of pixel blobs belonging to a reflector are based on the concepts of thresholding, morphological image processing, and image accumulation. The largest difference between the five centroiding methods is related to how the raw image is converted to a binary image for morphological image processing and/or accumulation. The five methods available in the VNSQL Tool are:

1. Global threshold
2. Local and global threshold
3. Thresholding with Otsu's method[5]
4. Thresholding with Otsu's method on edges
5. Sequential geodesic dilation

Preliminary analysis indicates that the "local and global thresholding" and "sequential geodesic dilation" methods provide the best results. Once the centroids have been computed, a summary of the results is displayed in a table in the VNSQL Tool GUI. For each centroid found, this table contains the centroid $[u, v]$ coordinates in the image, the mean reflector intensity, the mean range, the range standard deviation, and two estimates of the intensity signal-to-noise ratio. An example VNSQL Tool centroiding result using data from the flight day 3 rendezvous is shown in Fig. 13.

A user-defined region of interest may be defined in the VNSQL Tool and the analyst can zoom in on a particular portion of the image. The tool provides summary statistics of the raw intensity, raw range, and processed image in this region of interest. For either the single-measurement or batch-processing modes, it is possible to save raw data, JPEG pictures, and movies. On its own, the output of the VNSQL Tool may be used to qualitatively assess the performance of the STORRM VNS one snapshot image at a time. A more quantitative assessment may be performed by (1) comparing the VNS centroid results with the TCS measurements or by (2) comparing many VNS centroid results from data collected throughout the mission. Further, the outputs may be fed into a reflector identification algorithm, coupled with a relative navigation filter, to obtain an assessment of the resulting relative navigation performance. One post-processing task is to generate a VNS-derived BET to compare with the TCS-derived BET.

VNS data analysis during the mission focused on determining the range at which the VNS acquired the ISS and an assessment of reflector observability at various ranges. After early results demonstrated that the acquisition range exceeded expectations and that reflectors were clearly identifiable and matched predictions
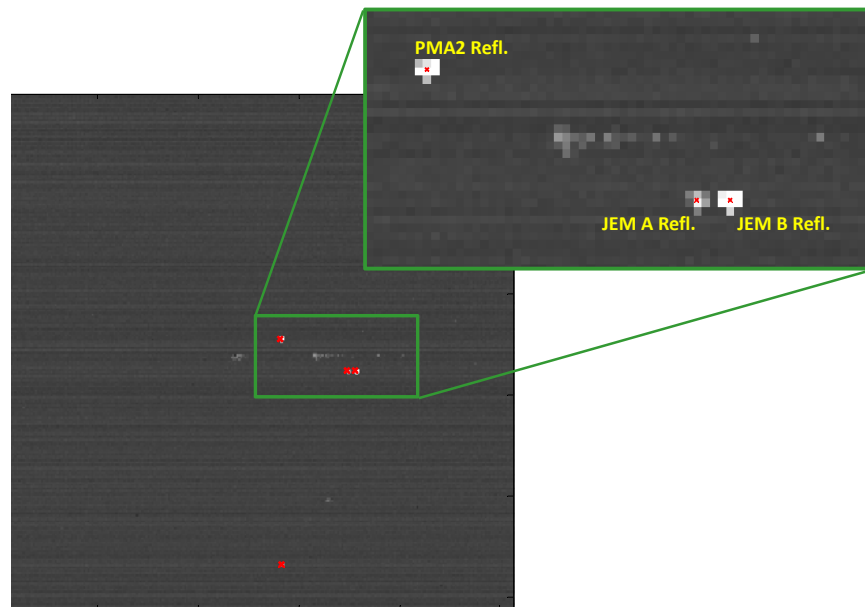

Figure 13. Reflector centroids (denoted by a red ×) are automatically detected in a VNS image using one of five different methods in the VNSQL Tool. This image was taken during rendezvous with the ISS on flight day 3 of STS-134.

American Institute of Aeronautics and Astronautics

at long-range and mid-range, focus was shifted to close-range performance. Of the 35 events where DRU block files were retrieved and downlinked during the mission, nine of them corresponded to points along the v-bar approach from 100 ft to 20 ft (at 10 ft intervals). This provided the data analysis team with a substantial amount of close-range data to assess the reflectivity of the ISS components at close-range and to tune the VNS centroiding algorithms. Some example VNS images that have been processed by the VNSQL Tool are shown in Fig. 14 and Fig. 15. These images were taken during rendezvous with the ISS on flight day 3 of STS-134.

The VNS image shown in Fig. 14 is characteristic of what was seen during much of flight day 3 and flight day 15. Of particular note is that the ISS reflectors were the easily detectable and were the brightest objects in the image at mid-range and long-range. Automatically detecting reflectors in these types of VNS images is relatively robust and straightforward. As the VNS gets closer to the ISS, spurious returns from other non-reflector objects on the ISS become significant. As may be seen in the intensity image in Fig. 14, there are significant returns from the docking ring and other features on the PMA. Additionally, the range point cloud in Fig. 14 clearly demonstrates that flash LIDAR data will be a powerful tool for relative navigation.

*a) Actual VNS image of ISS taken during Twice
Orbital Rate v-bar Approach (TORVA) maneuver.*

*b) Predicted reflector locations as produced by the
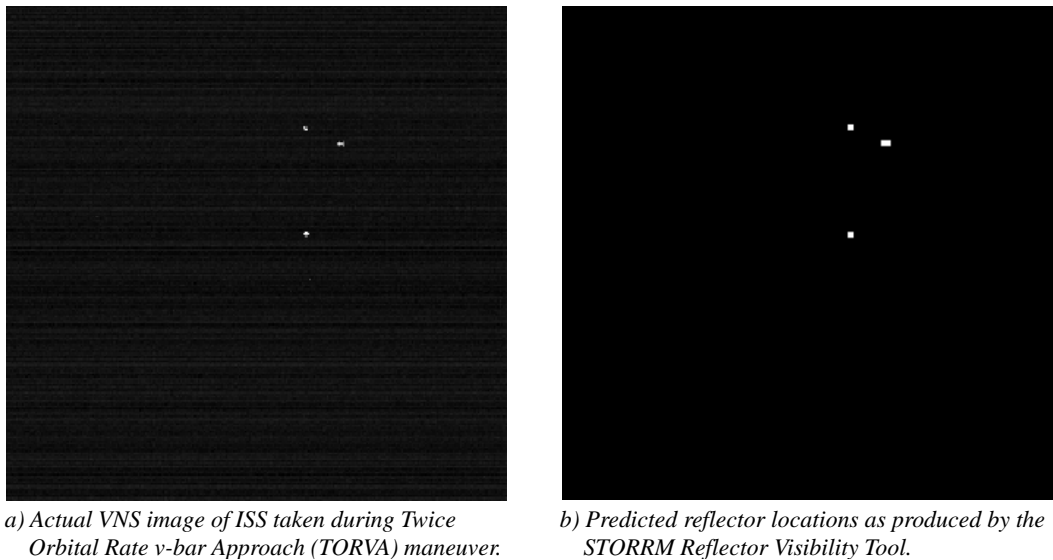STORRM Reflector Visibility Tool.*

Figure 14. The ISS reflectors appear in the expected configuration and are clearly visible in the VNS intensity images at intermediate range. There are few spurious reflections at these ranges. Clockwise from the top left in these images are the PMA2 reflector, the JEM A/B reflector pair, and the PMM reflector.
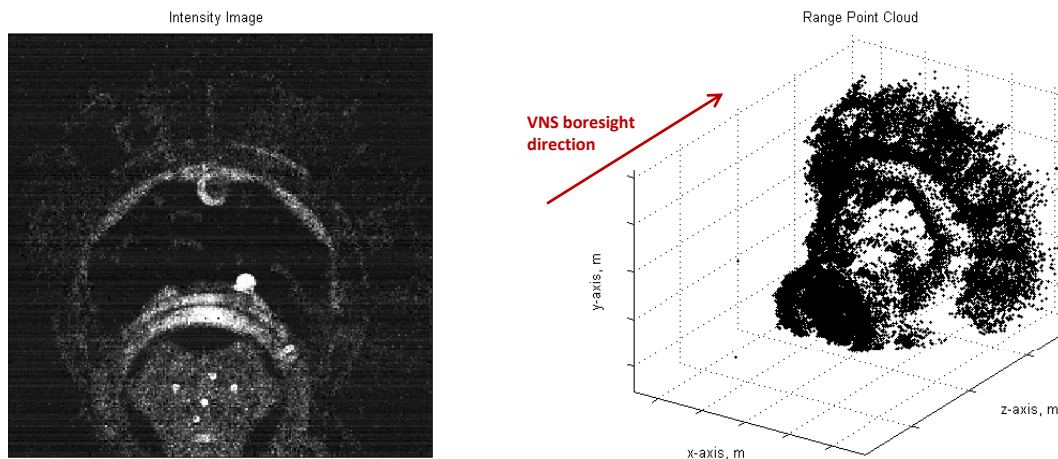

Figure 15. VNS intensity image and range point cloud taken during docking on STS-134 flight day 3. The 3D structure of the ISS PMA) is clearly visible.

VIII. Role of the VNS, Docking Camera, & STORRM Data in the Orion Relative Navigation Architecture

The heart of the Orion Relative Navigation hardware is the VNS and the DC. At long ranges, these sensors are supplemented with bearing measurements from a star tracker. Current plans do not have the DC being used directly as a relative navigation sensor; rather it will be used for piloting cues. In addition to these sensors, the Orion will have a Vision Processing Unit (VPU), comprising a set of FPGAs, to process the image data and output a relative state (range and bearing) and pose (relative position and attitude) as measurements. The Relative Navigation System uses measurements from these sensors to estimate the relative position of the two vehicles.

The primary components of the relative navigation software is a set of two Extended Kalman Filters. The first filter is a 'Translation' Filter, called RTFILT. It is a dual inertial state Extended Kalman Filter. The state space of this filter comprises the inertial positions and velocities of both the Orion vehicle and the target vehicle. It also includes the attitude of the Orion vehicle (as a modified Rodrigues parameter, a la the Multiplicative Extended Kalman Filter). Included in the state space are the measurement biases (range, azimuth and elevation) and gyro and accelerometer biases (15 states) for a total of 32 filter states. This filter also uses IMU data and GPS measurements. The IMU data are processed as dead-reckoning, i.e. the filter navigates the IMU case; hence IMU data are not processed as 'measurements'. Rather they are incorporated directly, and compensated by directly using the IMU states in the filter. GPS measurements are processed as PV (position and velocity) every few minutes; this is done so that the measurements are 'uncorrelated.' When the vehicles are in LEO, this serves to 'anchor' the inertial state. Outside of LEO, it is anticipated that the inertial state will drift unless a ground-update is performed. Ground Updates are processed as measurements.

When the two vehicles are within 20 m of each other, pose measurements from the VPU become available and the relative attitude filter (RAFILT) is initialized. The relative attitude filter is a Multiplicative Extended Kalman filter and only consists of the relative attitude of the target. It is assumed that the Orion attitude state is known.

RTFILT uses range and bearing information from the available sensors. The VNS produces range and bearing for up to 10 centroids on the target. Hence, before RTFILT can use the range and bearing, the reflectors must be identified. Reflector ID is a non-trivial task. When one reflector is visible, it is impossible to identify which reflector is being used (usually the c.g. is assumed). However, with two or more reflectors in the field of view, this task is easier. Algorithms to identify two or more reflectors have been developed and will be used on Orion. With the range and bearing to reflector centroids computed and the identification of these reflectors specified, RTFILT can go about the task of estimating the relative position of the two vehicles.

When the two vehicles are less than 20 m and four or more reflectors (on the docking target) are visible, the VPU uses VNS measurements and a RANSAC-based algorithm to compute a relative position and attitude (pose) measurement. RAFILT uses the relative attitude measurements produced by the VNS to compute a filtered relative attitude of the two vehicles. Included in the state-space of this filter is the target attitude rate.

From the flight data, it is apparent that the VNS has different performance depending on the range bin and selected mode. Based upon limited analysis, both these modes performed as expected. This will go into further refinement of the VNS design. The tight schedule of STORRM precluded extensive calibration and performance testing of the VNS. The Orion VNS units will be fully characterized, calibrated, and tested.

The relative navigation team will be using the VNS data obtained from STORRM to estimate what the measurement noise (range and bearing) was from the actual sensor in the space environment. Whereas truth was not known, a best estimated trajectory (BET) is being developed to characterize 'truth'. This will then be used to evaluate the noise characteristics of the VNS under the different modes, settings, and range gates. With a better understanding of the noise and bias characteristics of the VNS, the relative navigation system can be 'tuned' to the expected VNS performance. This, in concert with future ground tests, will allow the relative navigation system to be robust as well as meet the stringent requirements of executing rendezvous and docking of two vehicles in space.

IX. Lessons Learned

Many lessons were learned throughout the STORRM project. A summary of some of these is as follows:

1. Requirements should be written with emphasis on system capability, not for a specific concept of operations. STORRM's mission concept was envisioned in the early developmental stages of the project. For the project to gain momentum, this concept was socialized to the different stakeholders and involved parties. When design began, many decisions were framed in terms of the specific mission concept. This ultimately became a constraint that could have been avoided. The Mission itself wasn't affected by this, but it complicated the design and planning stages of the project.
2. All procedures should be bookended (begin and end) at some stable state. They should be compartmentalized into indivisible sequences of events. The STORRM procedures were ultimately sculpted into this modular form. The benefit of arriving at some stable and known condition at the conclusion of each procedure is that it allows unambiguous off-ramps for troubleshooting and contingency measures. The alternative is a set of instructions that are highly scenario dependent.
3. The philosophy of fault detection and troubleshooting measures should be chosen in the design phase. When a fault is found in a system there is a fundamental choice to be made: how deep in the layering of systems should one go to isolate the issue? Because this was never discussed and agreed upon, the software parameters that guide fault detection and isolation were inconsistent. In addition, the use of a global system reset, or soft cycle, was the first line of defense in problem resolution, but no automatic processes were put in place to initiate a soft cycle from any and all system states/configurations. Early philosophical discussions in the design phase could have vetted these types of artifacts of the final implementation.
4. Parallel test ports should be implemented on all flight hardware elements. In the case of STORRM, the throughput rate from the avionics to the STORRM software was limited by the vehicle bulkhead and became a design point for the avionics. This was very short-sighted as it became a significant driver during the environmental test campaign. Since data retrieval off the data recorder units was so restricted, extremely limited data was pulled to make assessments and time spent retrieving data was a hit to schedule.
5. A simulation to drive the software with realistic data and to manipulate command and telemetry parameters of the system to reproduce failures is critically important. Even on a project with limited schedule and budget, this is a worthwhile investment. For STORRM, mission success would have been at risk without this benefit to software testing, crew training, procedure development and checkout, and gaining a deep understanding of the interfaces often uncovering errors in documents that were not previously noticed. The portability of the simulation was an important added benefit for crew training.

X. Acknowledgements

References

[1]Goodman, J., "History of Space Shuttle Rendezvous and Proximity Operations," *Jounral of Spacecraft and Rockets*, Vol. 43, No. 5, Sept-Oct 2006, pp. 944–959.

[2]Miller, K. L., Masciarelli, J., Rohrschneider, R., and Gravseth, I., "Critical Advancement in Telerobotic Servicing Vision Technology," *AIAA SPACE 2010 Conference & Exposition*, Anaheim, CA, 30 Aug - 2 Sept 2010.

[3]Alleyson, D., Süsstrunk, S., and Hérault, J., "Color Demosaicing by Estimating Luminance and Opponent Chromatic Signals in the Fourier Domain," *Proc. Color Imaging Conf.: Color Science, Systems, Applicaitons*, 2002, pp. 331–336.

[4]Bradski, G. and Kaehler, A., *Learning OpenCV: Computer Vision with the OpenCV Library*, O'Reilly Media, Inc., Sebastopol, CA, 2008.

[5]Gonzalez, R. and Woods, R., *Digital Image Processing, 3rd Ed.*, Pearson Prentice Hall, Upper Saddle River, NJ, 2008.