Orbital Express AVGS Validation and Calibration for Automated Rendezvous

Andrew F. Heaton[1], Richard T. Howard[2], Robin M. Pinson,[3]
NASA Marshall Space Flight Center
Huntsville, AL 35812

From March to July of 2007, the DARPA Orbital Express mission achieved a number of firsts in autonomous spacecraft operations. The NASA Advanced Video Guidance Sensor (AVGS) was the primary docking sensor during the first two dockings and was used in a blended mode three other automated captures. The AVGS performance exceeded its specification by approximately an order of magnitude. One reason that the AVGS functioned so well during the mission was that the validation and calibration of the sensor prior to the mission advanced the state-of-the-art for proximity sensors. Some factors in this success were improvements in ground test equipment and truth data, the capability for ILOAD corrections for optical and other effects, and the development of a bias correction procedure. Several valuable lessons learned have applications to future proximity sensors.

INTRODUCTION

NASA supplied the Advanced Video Guidance Sensor (AVGS) to DARPA and prime contractor Boeing as part of the Orbital Express mission that flew in March of 2007. The Orbital Express mission was a demonstration of automated satellite servicing. Orbital Express accomplished several firsts in automated spaceflight, including the first successful automated American propellant transfer, the first automated transfer of batteries, the first automated transfer CPU transfer, and the first American Automated Rendezvous and Capture (AR&C). The AVGS was a critical component for the AR&C operations. Boeing built and supplied the chase vehicle, called ASTRO, while Ball Aerospace provided the NEXTSat target vehicle.

The AVGS Sensor is a combination of video and lasers that uses corner cube retro-reflectors on the target vehicle. The lasers are fired at two different wavelengths. One laser wavelength is filtered by a lens on the reflectors. The different-frequency lasers are fired asynchronously at 5 Hz each, with an offset of 0.1 second. The sensor stores the alternating images at 10 Hz internally. One 10 Hz image thus contains all background lighting plus the corner-cube reflections, while the next 10 Hz image has the background lighting and no corner-cube reflections. These images are subtracted, and spots from the corner-cube reflectors are readily visible in the resulting image with little or no noise spots. This method of locating the retro-reflectors is highly robust to a variety of on-orbit lighting conditions.

For the Orbital Express mission, the AVGS used four smaller corner-cubes for the Short Range Target (SRT), and four larger corner-cubes for a Long Range Target (LRT). Both of these target configurations are displayed conceptually in Figure 1. Once the target assemblies are captured in the imager by the "background-subtraction" process described above, the "Inverse Perspective" algorithm processes the 2-D

---

[1] Aerospace Engineer, Guidance, Navigation and Control Group, EV42, NASA Marshall Space Flight Center, Huntsville, AL 35812, andrew.f.heaton@nasa.gov, AIAA Member, AAS Member.
[2] Team Leader, Advanced Vehicle Sensors Team, NASA MSFC, Huntsville, AL 35812 / ES62, non-member.
[3] Aerospace Engineer, Guidance Navigation and Mission Analysis, NASA MSFC AL 35812 / EV42, non-member.

reflector image into a full 6-DOF attitude and range solution. The Inverse Perspective combines the spot data as input with the known locations of the corner-cubes on the target spacecraft to determine the solution.
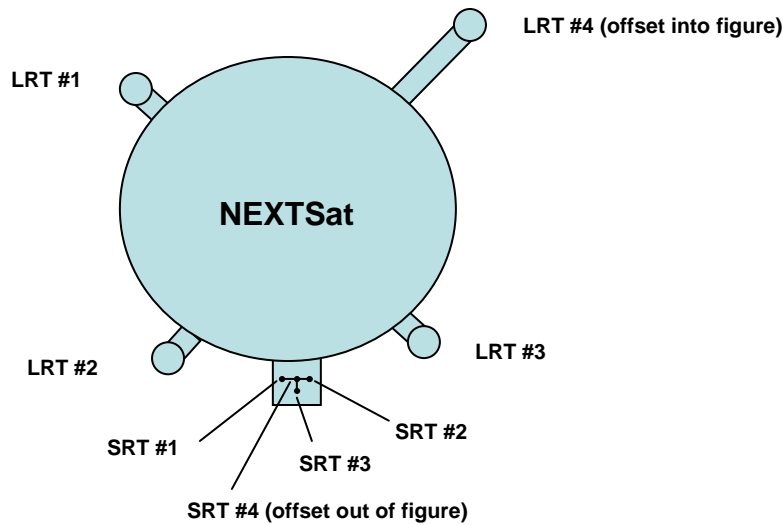

Figure 1) AVGS Corner-Cube Locations on NEXTSat

The AVGS has a long history of development and previous flight validation. AVGS flew on the Demonstration of Automated Rendezvous Technology (DART) in 2005. Although the mission experienced an anomaly and did not achieve the phase where AVGS was used in closed-loop fashion, valuable bearing data was collected in the AVGS "spot mode", and the AVGS performance was nominal. Prior to DART, an earlier model of the AVGS known as the Video Guidance Sensor (VGS) flew two open-loop experiments on the shuttle flights STS-87 and STS-95 (in 1997 and 1998 respectively). NASA developed the original concept and Engineering Development Unit (EDU) of the VGS in 1992, so the highly successful flight on Orbital Express is the result of 15 years of development. This is a timeframe commensurate with the development of the recent highly successful flight of the ESA Videometer on the ESA Automated Transfer Vehicle (ATV). The Videometer uses a technological approach similar to the VGS that flew on STS-87 and STS-95

Table 1) Orbital Express AVGS Specification

Target	Operating Range (m)	Range (mm)	Azimuth, Elevation (deg)	Roll (deg)	Pitch, Yaw (deg)
SRT	1-3	±12	±0.033	±0.13	±0.20
SRT	>3-5	±35	±0.033	±0.25	±0.33
SRT	>5-10	±150	±0.035	±0.45	±0.70
SRT	>10-30	±1500	±0.037	±1.30	±2.0
LRT	>10-30	±150	±0.027	±0.15	±0.70
LRT	>30-50	±400	±0.030	±0.25	±1.2
LRT	>50-100	±1666	±0.033	±0.50	±2.4
LRT	>100-300	$\pm15,000$	±0.035	±1.40	±7.0

Automated Rendezvous and Capture (AR&C) requires a relatively tight error budget for a proximity sensor. A manned spacecraft can react more quickly to unplanned contingencies and 3-sigma error dispersions. Therefore, a manned docking typically requires less stringent sensor accuracy. Thus, for

automated operations it becomes more important to calibrate the sensor and validate that it meets specification. Calibration and validation should be accomplished prior to flight on the ground, but the ability to adjust calibration on-orbit is also very valuable. For Orbital Express, the AVGS had a challenging specification (see Table 1).

PROXIMITY SENSOR ERROR SOURCES

Several factors affect proximity sensor performance. There is the "absolute accuracy" of the sensor itself, which can be affected by many internal characteristics of the sensor and/or any target it might use. These can be classified as "internal errors" and will vary from sensor to sensor. Error can also be introduced from external sources, of which there are two broad categories. The first category is misalignment (which includes the sensor and the target on the target spacecraft, if one is used). The second category is environmental factors that might affect performance such as lighting or thermal stresses. The effect of the environmental factors can vary from degraded performance to loss of signal.

The accuracy of the sensor can be tested extensively on the ground, in most cases. Some allowance must be made to account for vacuum conditions and orbital lighting. In our experience, the vacuum effect, while noticeable, does not appreciably affect performance for visual proximity sensors and so ground testing to validate truth can proceed with confidence (the presence of a vacuum is typically to improve, not degrade performance). Effective ground testing for validation of sensor accuracy is most strongly affected by the accuracy of the truth estimate used as a basis for comparison. While several factors affect truth validation, for AVGS, we found the biggest factors to be optics and test equipment

Misalignments can be corrected after the fact. That is, once the vehicles are docked, a bias can be estimated and removed via a software ILOAD procedure. (An ILOAD in this case is a text file that allows software parameters to be easily updated through a simple process that does not require recompiling code). The AVGS team on Orbital Express was fortunate in that the vehicles were launched in the mated position, and that the docking mechanism had a high degree of repeatability. Thus, we were able to execute a highly accurate bias calibration procedure in October of 2006, when the vehicles were mated in the clean room in Titusville. This will not be the case for every mission, but experience on Orbital Express suggests that a bias calibration procedure is possible for vehicles that dock for the first time on-orbit.

As for lighting, in the case of the AVGS it typically does not directly degrade the performance of the sensor itself, but can in some cases limit the operational envelope of the sensor. We say this not in the sense that the operational envelope is severely restricted (some proximity sensors require very tight lighting constraints in order to work), but in the sense that it might degrade the maximum range or lead to occasional dropouts at some of the longer ranges. The AVGS is designed to be (and is) very robust to on-orbit lighting conditions, and yet some problems were still encountered (in spite of a successful spaceflight of the sensor and two spaceflights of its predecessor the VGS). In fact, it's not a stretch to say that lighting is by far the biggest space environmental effect on vision-based proximity sensors. In the case of the AVGS some minor problems were encountered with lighting that can be attributed primarily to target design and secondarily to lens optics. Needless to say, very valuable lessons were learned in this area.

Thermal stress can also affect sensor performance by direct effects on the sensor and thermal bending that might affect sensor and target location. Thermal bending did not appear to be an issue for the AVGS on Orbital Express. It's possible that some future mission might have a target or sensor mounted in a location or in an environment where this would be a concern, but in the four flights to date of the AVGS and VGS, this has never been an issue. What thermal stress did affect on Orbital Express was the low-frequency noise signature of the AVGS, which was directly correlated with temperature variations in the sensor. While this characteristic may be specific to AVGS and so not applicable to future sensors, it did constitute by far the largest source of noise for the AVGS (though still well below specification) and so is important.

Finally, if a proximity sensor does use a target (the AVGS and the three Boeing camera-based sensors on Orbital Express all used a target), then the target design can greatly affect sensor performance. This was a very valuable lesson for the AVGS team on Orbital Express that re-enforced existing notions on target

design and also provided fresh insight, to the point that we feel justified in saying proximity sensor target design is now a science and not an art.

SENSOR VALIDATION TESTING

For AVGS ground testing, the emphasis of ground testing is typically hardware-related issues in the earlier phases of the program. For example, in the early stages there is a good deal of "Optical Characterization Testing" (OCT). OCT tests determine the proper settings (laser power, video chip integration time and video chip threshold) to allow the AVGS to image target spots at all operational ranges. Other tests might be performed to correct hardware issues, determine the boresight and/or focal length of the sensor, and so forth. During the later phases of a flight program, the test emphasis switches to validation of the performance specification. Accurate truth data for comparison is a key element of this test. This paper is primarily concerned with the tests that are directly related to validating the accuracy of the sensor.

In the past, the specification for the VGS and AVGS was more forgiving than the OE specification that appears in Table 1. The OE specification was derived directly from the DART specification as it existed in March of 2003, at the very beginning of NASA's Orbital Express effort. Unfortunately, the DART specification was developed prior to DART flight unit testing and proved to be optimistic. NASA relaxed the DART specification in fall of 2003, and recommended to Boeing and DARPA that the Orbital Express adopt the new standard. Boeing insisted that the original specification be honored, and so a more rigorous approach to truth validation was needed for OE than that required for DART.

More accurate truth data was possible for Orbital Express due to two advances in test equipment. First, the Flight Robotics Laboratory (FRL) at the Marshall Space Flight Center was able to purchase a Leica laser-range finder to produce accurate truth estimates for the ground testing. The Leica is a device that can measure several points on the AVGS sensor and the AVGS target and provide a range and pose estimate that is independent of the sensor itself. The user is even able to define different coordinate frames in which to take the measurement, which allows easily identifiable features on the sensor or target to be used for reference. We used a variety of reference points on the target and sensor for the tests, but for the Systems Performance Test in October of 2005, the reference points were the back plate of the AVGS sensor and the top of the glass for the offset corner cube on the SRT target.

Boeing also supplied NASA a test platform for short range testing called the Newmark. The Newmark consists of a pair of gimbaled test platforms that allow relative translation and rotation of the sensor and target with respect to each other in a tightly controlled environment. The Newmark functioned out to a range of about 55 meters, and was highly effective in the close-range calibration testing of the sensor. Specifically, the Newmark enabled tight control of the relative rotation target with respect to the sensor to a precision of 0.01 degrees. The target platform can rotate by command to any attitude within the operational constraints of the AVGS, and the commands can be automated with a script. The sensor platform was used to change the bearing of the target with respect to the sensor, and can also be rotated to within 0.01 degrees and commanded manually or with a script.

The Newmark was most useful at a relative range of less than 10 meters, where the SRT is required to function alone (the LRT is only required to Track to 10 meters). The most important testing position for the AVGS was the docked or mated position. On-orbit, the mated configuration of the vehicles was such that the AVGS was 1.219 meters from the SRT fixture. Pre-mission, this range was not always known, and so the "mated position" tests varied from 1.1 to 1.5 meters over the course of the project. However, in the long run, this was helpful, as the large number of tests performed between 1.1 and 1.5 meters helped to rigorously characterize the behavior of the AVGS in the capture envelope. Capture envelope performance was the most critical operation for the AVGS on-orbit.

The AVGS solution consists of Range, Azimuth, Elevation and a quaternion for the attitude. For testing, the attitude is typically translated into Euler angles for ease of analysis, since the specification is written in Euler angles (however, the quaternion is passed to the navigation filter on-orbit). Figure 2 illustrates the AVGS coordinate frame. Figure 2 shows a sample SRT image in the Field of View, which is eight deg by eight deg. Azimuth and Elevation are bearing angles defined as "positive-right" and "positive-up",

respectively, while Roll, Pitch, and Yaw are Euler rotations about the X, Y and Z axes of the AVGS reference frame, respectively.
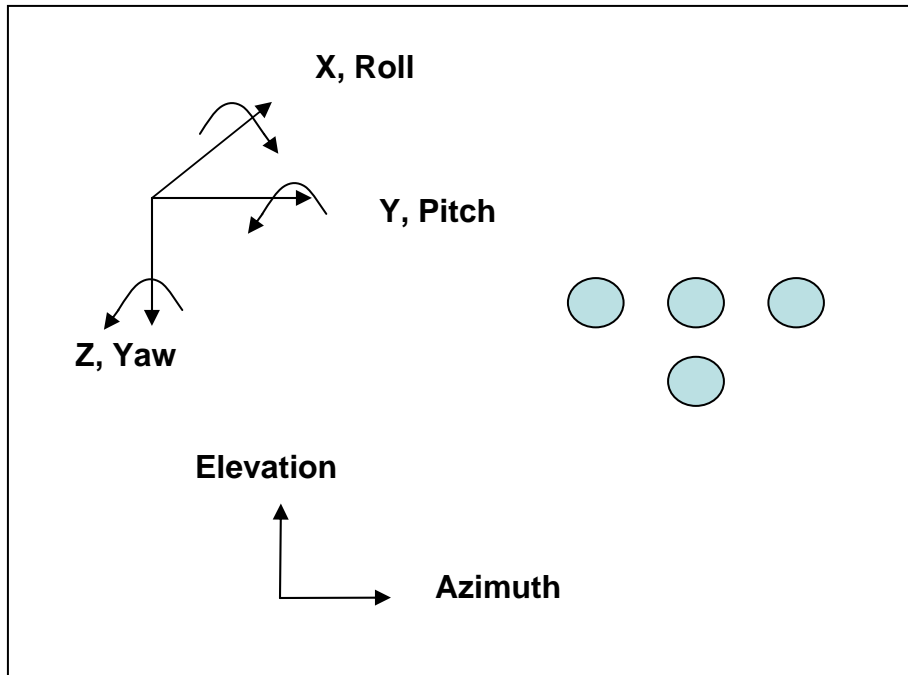

Figure 2) AVGS Coordinate Frame

SYSTEMS PEFORMANCE TEST

The Systems Performance Test (SPT) in October of 2005 was the primary source of accuracy validation for the AVGS sensor. The tests consisted of a series of 2-minute tests at ranges of 28, 11, 5, 1.5 and 1.3 meters. (Two minutes was thought to allow sufficient variation for the typical AVGS noise characteristics at the time the test was designed). Note that the "docked range" did not quite correspond to the actual docking range of 1.219 meters, because the location was not known precisely when the Systems Performance Test took place. In addition to testing at a number of ranges (of which the 1.5 and 1.3 meter tests were most critical), the Newmark test apparatus was also used to vary the target attitude (Pitch, Yaw and Roll) and sensor bearing angle (Azimuth and Elevation) through a series of 2-minute tests to thoroughly characterize the AVGS performance over the entire operational envelope.

Although the operational envelope of the AVGS extended to 120 meters, and the sensor was considered critical at ranges less than 60 meters, less rigorous ground testing was required outside of 30 meters. The less stringent approach was possible for two reasons. One, the specification at ranges greater than 30 meters is rather forgiving. Two, the Boeing sensors provided adequate overlap with the AVGS outside of 30 meters, for the given specification. Taking these factors into account, the testing for ranges greater than 30 meters did not need to be as rigorous. Here we will note that the AVGS was not tested at any range over 100 meters, even though technically the specification extended to 300 meters. Testing outside of 100 meters was not considered necessary by the customer (though greatly desired by NASA) because operationally AVGS was only required inside of 120 meters. (However, on the second unmated operation, Scenario 3-1, the AVGS was used as the primary sensor from 150 meters to dock and performed well, in spite of not having any OCT data for calibration at ranges greater than 100 meters). Since the most critical area that received attention pre-mission was the close proximity ranges, including and particularly the capture envelope range (anything less than 1.5 meters), this paper focuses primarily on the close-range testing (30 meters and less).

The attitude of the target and the bearing angle of the sensor were varied systematically for each of the ranges tested in the Systems Performance Test. Each parameter of the solution was tested over its operational range inside of 60 meters. The operational range for the Azimuth and Elevation was +/-4 degrees (based on the size of the ASTRO approach corridor). The operational range for the target Pitch, Yaw and Roll was also +/-4 degrees (based on the 3-sigma accuracy requirement on the NEXTSat attitude accuracy). Also at each test range, the range was also varied slightly using the Newmark precision translational capability (+/- 1 cm in the X direction, with a precision of 1 mm).

At each test range, each parameter was varied systemically. For instance, at a range of 1.3 meters, the sensor was rotated to Azimuths of 4, 3, 2, 1, 0, -1, -2, -3 and -4 with no other rotations or translations. Then this test pattern was repeated five times while varying *one* of the other parameters at each test Azimuth at a time. So at each test Azimuth, Pitch was varied, then reset to zero while Yaw was varied, and so forth. In all, at the 1.3-meter range 73 individual two-minute tests were performed while for the entire test, roughly 360 individual two-minute tests were performed. Some special manual testing also occurred at the end of the test phase to address particular problems.

Several lessons were learned in the course of the SPT that greatly enhanced the ability to determine and fine-tune the accuracy of the AVGS. One aspect explored in detail was the effect of the optical centers of the corner-cubes and sensor on the test, and how to compensate for inaccuracies or uncertainties in knowledge of their location. For the SPT, the test philosophy mandated by Boeing (Dean Hallmark and Jerry LeCroy of Boeing designed and executed the SPT) was to use easily identifiable physical marks on the target and sensor for truth measurements. These locations were then carefully adjusted to the optical centers of the corner cubes and sensor in post-test analysis. The test reference mark for the sensor was the backplate and the reference mark for the target was the top of the glass of SRT #4 (see Figure 1). Given this approach, knowledge of the location of the optical centers of both target and sensor was critical to accurate truth measurements.

The optical center of the sensor was known with a high degree of accuracy from previous testing. The optical center of the sensor was determined by Orbital Sciences Corporation, manufacturer of the sensor. This measurement was validated by earlier testing and considered to be known with an accuracy of +/- 1mm per axis. The Newmark test apparatus for the sensor was actually designed to rotate around the optical center of the sensor, and the uncertainty in the Newmark rotation was also on the order of +/ 1mm. Therefore the overall uncertainty of the optical center of the AVGS was known to about +/- 1.4 mm Root Mean Square (RMS), 1-sigma. As it turned out, even this small uncertainty had to be considered in the SPT data analysis.

The optical center of the corner cubes was not part of the specification from the manufacturer. That is, it was not known a priori and had to be determined via testing. Although some idea had been gleaned from earlier, less rigorous tests, the knowledge was much less precise than for the sensor. Furthermore, the SRT and LRT corner cubes are different sizes, meaning each has a different optical center. Knowledge of each was roughly on the order of +/- 5 mm prior to the SPT.

A lack of knowledge of the optical center of the corner cubes (and to a lesser extent, of the sensor) contributes directly to errors in Range, Azimuth, and Elevation, since the center directly affects each of those parameters. There is also an indirect effect on Pitch, Yaw and Roll, that is roughly proportionate to the Azimuth and Elevation errors. The uncertainty in optical center presented some special issues for testing using the Newmark test apparatus. Although the Newmark was very precise, an induced error can occur in test data when the target or sensor is not rotated about its respective optical center.

This effect is best described using the illustration in Figure 3. If the target is rotated about the exact optical center of corner cube SRT #4 (which is the overall reference point for the sensor solution), then the expected position of the target should match the center. If however, the target is rotated about a point other than the optical sensor, a translation error is introduced into the "truth position". In this idealized example, the error would be in Azimuth, which would couple into a Yaw error as well. In the real case, the error is actually cross-coupled in the full translational 3 DOF of the optical sensor, and those misalignments are then coupled with the rest of the solution. However, since in the SPT only one other parameter at a time

was varied with the main parameter being tested, it was possible to determine the optical center of the target corner cubes from experiment data, and adjust the data a posteriori. That is to say, in the portion of the test that featured the Azimuth at 1.3 meters, the other parameters could largely be ignored and a calibration based on just Azimuth could be achieved. As it turns out, only Azimuth and Elevation exceeded specification in the original testing, and it was eventually determined that the problem was largely due to rotation of the target assembly about a point other than its optical center, and to a lesser extent, some residual error in the location of the sensor optical center.
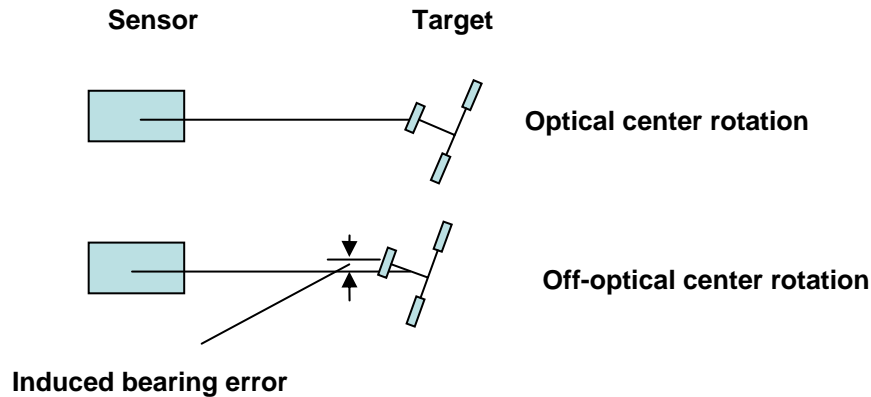

Figure 3) Test Error Induced From Test Apparatus

The problem of rotating about a point other than the optical center was known from earlier tests, and had previously been corrected, but the SPT needed a resolution of accuracy roughly an order of magnitude higher. In the previous tests, a rough estimate of the optical center sufficed, but for the SPT it had to be estimated to within +/- 1 mm for both sensor and target. The sensor optical center knowledge was already close to the needed accuracy, but an iterative effort was required to achieve sufficient knowledge of the optical center of the corner cube retro-reflectors.

Briefly, the procedure for correcting for the optical center of the corner cubes consists of two steps. The first step is to take the expressions that represent the known corner-cube locations in the AVGS flight software and vary them systematically in an independent simulation until an estimate of the corner-cube and sensor optical centers is found that adjusts the test data to be within specification. This adjustment is then applied to a full simulation of the actual flight software using the raw spot centroids locations fed back in as raw data, to verify the solution. What we found was that the independent simulation required several iterations but once it converged to a proper optical center for the target and sensor, the full simulation would match the result of the independent simulation.

An example of the correction process appears in Table 2. Some selected Azimuth data is shown from the original SPT test at 1.3 meters. The errors before and after the optical center correction are shown, and the improvement brings the sensor to within specification.

Table 2) Azimuth Errors with Optical Center Correction

Test Azimuth	Original Error	Corrected Error
2	-0.1097	0.0150
1	-0.0675	0.0006
-1	-0.0428	-0.0007
-2	-0.0284	-0.0138

In addition to the advantage of achieving flight certification, the process developed produced two ancillary benefits as well. Specifically, it enabled us to estimate the optical center of the SRT corner cubes to approximately 1 mm. This was beneficial to the overall performance of the sensor at all ranges, since this

reference point was used by the ASTRO flight software to transform the AVGS solution (which is sensor-to-target) into the ASTRO body coordinate reference frame with a high degree of accuracy. The second ancillary benefit is that we developed the second simulation described above, which essentially consisted of the actual flight software taking recorded spot data as input (as opposed to actual digital measurements) and this simulation proved very useful for many other forms of data analysis.

Although it was important and led to ancillary benefits, the optical center correction was only part of the SPT data calibration process. We also had to develop a process to correct for fixed sensor biases, since these could in theory exceed the specification (the optical sensor correction essentially was absorbed by this later process). Some discussion of the specification is necessary. The specification was written as a total error specification. During the DART mission, the specification was eventually adapted to one that included a bias and noise specification, which is typical of most navigation sensors. However, due to the Boeing insistence that we honor the original DART specification, we were unable to adjust the OE specification in a similar manner. Furthermore, Boeing desired that we design the sensor I-LOAD so that a fixed bias from all error sources could be removed. That is, in addition to the bias induced by the internal error of the sensor, it was also highly desired that any mounting misalignments of the sensor or target be correctable via an I-LOAD procedure (an I-LOAD is a writable portion of the flight software that contains constants and calibration data that can be replaced if necessary during the mission without modifying the actual flight software.) Therefore, we were tasked by the customer to develop a "Bias Calibration Procedure".

BIAS CALIBRATION PROCEDURE

In order to fully develop a bias calibration procedure, we performed some additional ad hoc SPT tests in November and December of 2005. These consisted of the following: A simulated docking position was selected at a range of approximately 1.3 meters. Using the Newmark test equipment, the target was placed near the docking position (range of about 1.3 meters, all other parameters near zero). Some error was deliberately allowed to remain, as a "simulated bias". The Leica was used to measure the truth location of the simulated docking position carefully. Then several minutes of data were taken and treated as a simulated docking position with bias and noise errors. This data was then analyzed to estimate the bias by eliminating the noise, and a procedure was developed to remove the bias.

The procedure consisted of several steps. First, the noise had to be separated from the bias. By taking a sufficiently long sample, the noise could be largely removed via simple averaging. It was only necessary to acquire a sufficiently long sample of data so that the sensor noise would fully manifest itself (that is, the data sample would contain 3-sigma max and min values as well as a representative sample of the overall noise pattern). In the process of analyzing the data, we found that the noise characteristics of the AVGS fall into two frequency categories, which we call low-frequency and high-frequency noise. The high frequency noise is fairly benign and follows an almost Gaussian distribution. The high-frequency noise was around 2 Hz and had small amplitudes (on the order of 1/10 the 1-sigma requirement). The sensor also had low-frequency noise (around 0.003 Hz) that was much more significant. The low-frequency noise had a maximum amplitude of nearly one-half the 1-sigma requirement. Thus low frequency noise is by far the dominant source of noise error, and in our ad hoc SPT tests we had to allow a sufficient amount of time for it to fully manifest. We found that 5 minutes was necessary for the low frequency noise to be fully estimated and removed from a given data sample.

Here we mention that one of the I-LOAD parameters is a matrix of Cartesian coordinates in of the AVGS corner-cube reflectors that is used in the Inverse Perspective algorithm in the AVGS flight software. For the purposes of this discussion we need only to say that they are effectively the known locations of the corner cubes, and that they can be adjusted via I-LOAD. We refer to this set of parameters as the "target matrix".

The original bias calibration procedure developed in January of 2006, based on the Systems Performance Test consisted of the following:

1. Data was taken at a simulated docking position as described above.

2. Simulated biases were estimated by removing the noise via simple averaging of the solution over a sufficient length of time.
3. The "target matrix" of the AVGS target was translated in the flight software I-LOAD to remove Azimuth and Elevation simulated biases.
4. The "target matrix" of the AVGS target was rotated with a rotation matrix consisting of the attitude biases as measured in step 2, and the flight software I-LOAD adjusted to the new corner cube positions.
5. The new I-LOAD from steps 3 and 4 was tested with the flight software simulation using raw spot data from the original test and the residual bias checked against the specification.
6. Steps 3 and 4 were repeated until the simulated bias was removed.

We found that this bias calibration procedure worked very well, but that it required numerous iterations and that convergence was difficult to achieve. The results of the accepted bias calibration are presented in Table 3 and Table 4. Table 3 is the original uncorrected error and Table 4 is the residual bias error following the calibration procedure.

Table 3) Original Errors at 1.3-meter SPT Test Position

Range (m)	Az (deg)	El (deg)	Pitch (deg)	Yaw (deg)	Roll (deg)
0.001	0.03209	-0.19652	-0.28617	0.636735	-0.10338

Table 4) Calibrated Errors at 1.3-meter SPT Test Position

Range (m)	Az (deg)	El (deg)	Pitch (deg)	Yaw (deg)	Roll (deg)
0.000	0.008724	-0.00126	0.002233	0.003075	-0.00028

The successful development of a bias calibration procedure, allowed the AVGS to "pass" the Systems Performance Test and the sensor Systems Readiness Review in January of 2006 at the Boeing facility in Anaheim, CA. However, the procedure left something to be desired, as it required a large number of iterations and even when successful, and allowed fairly large bias errors to remain in the sensor solution (Table 4 is the best result and not representative of most).

As the flight date approached, Boeing became concerned about the Calibration Procedure as the actual written procedure was ad hoc and not suitable for a flight project. Furthermore, there were concerns on the part of the NASA AVGS team about the amount of time the procedure took to execute and the residual biases that it allowed. Therefore, in March of 2006, at the request of Boeing, we attempted further improvements to the Bias Calibration Procedure in a week-long series of tests at the Flight Robotics Laboratory in Huntsville, AL. The goal was to reduce the number of iterations, improve residual biases, and streamline the written procedure into one suitable for a flight project.

After a full week of testing, we had great success with improving the accuracy and convergence time of the Bias Calibration Procedure. After the lessons learned from the March testing were applied, it was rare that more than one iteration of the Bias Calibration Procedure was required. Furthermore, the residual biases were reduced to the effective measurement precision of the sensor. Another way of stating the last sentence is to say that, with the refined procedure, we can arbitrarily reduce the bias to zero in attitude for any given data sample, and to an acceptably low value in Azimuth and Elevation (with further iterations Azimuth and Elevation can effectively be zeroed as well, but typically this is not done). Tables 5 and 6 show the improved performance from a March 2006 test sample.

Table 5) March Test Original Errors at 1.3-meter Range

Range (m)	Az (deg)	El (deg)	Pitch (deg)	Yaw (deg)	Roll (deg)
0.001	0.10485	-0.10313	-0.10492	0.080435	0.00000

Table 6) March Test Calibrated Errors at 1.3-meter Range

Range (m)	Az (deg)	El (deg)	Pitch (deg)	Yaw (deg)	Roll (deg)
0.000	0.00401	-0.00229	-0.00004	0.00000	0.00000

The improvement in performance indicated by Tables 5 and 6 was made possible by the following improvements in the Bias Calibration Procedure:

1. The AVGS solution quaternion had to be re-normalized due to lack of precision in the output solution. After normalization, the quaternion is then used to calculate the rotation matrix described in step 3 above. (The quaternion is actually used in the Bias Calibration Procedure, the Euler angles in the table above were calculated for "human eye" convenience and so a direct comparison to the specification could be made.)
2. The target matrix precision had to be expanded from 4 to 6 decimal places. Although the I-LOAD target matrix in the AVGS flight software only has four decimal places, the target matrix is part of the rotational correction, meaning that it is multiplied with a direction cosine matrix calculated from the attitude bias error. Thus the numerical error from lack of precision in the target matrix is amplified.
3. In determining the bias from data samples, the median proved a much more accurate gauge of the true bias than the average.

With these three improvements, the Bias Calibration Procedure was more than sufficient to support the OE flight. We note that on item 3, it was originally thought that a filter might be necessary, but the low-frequency noise characteristics of the AVGS were of such a nature that simply replacing the mean solution with the median worked well. This is due to the signature of the low-frequency noise, which has a pattern that has a very repeatable amplitude but varies in frequency. Using the median eliminates any "variable-frequency" effect from this error source.

The true test of the Bias Calibration Procedure occurred in October of 2006, when mated data from the vehicles in the clean room in Titusville, FL was obtained. This data was used for a Bias Calibration Procedure including a new I-LOAD file for the flight sensor software in December of 2006, which was immediately followed by another test to validate that the procedure worked. The test was performed by Boeing personnel remotely and the results sent to NASA for analysis. At this point in time, we also worked with Boeing to develop the process of delivering an I-LOAD to them during the flight, if necessary.

There was a concern that the bias would change from the bias measured (and then corrected) on the ground because the vehicles were launched mated together with a separation ring that was removed after launch. The removal ("blow-off") of the separation ring was known at a minimum to cause a 2.5 mm difference in the AVGS "X" direction (the direction along the line of sight), but effects on attitude and bearing were unknown. Therefore, since it was expected that the bias might actually change significantly between the ground and on-orbit operations, DARPA felt that if the bias calibration had to be repeated, it could wait till the vehicles were on-orbit and the separation ring removed.

NASA and the Boeing AVGS lead managed to persuade DARPA that it was better to perform the Bias Calibration Procedure in December of 2006 for two reasons. One, the initial biases on the sensor were very large, and it was possible that the on-orbit schedule might not allow an I-LOAD update. Two, the Bias Calibration Procedure had never been tested with the real flight computers and hardware, and it was risky to wait until the spacecraft were on-orbit to attempt it for the first time. DARPA eventually decided to allow the Bias Calibration Procedure in December 2006, in spite of an ambitious schedule to get the vehicles ready for flight. One additional factor that was persuasive was the thought that if the bias calibration on the ground went well enough, we might not need it on orbit (a strong possibility in spite of the separation ring blow-off).

As stated previously, the initial biases from the October 2006 mated vehicle test were quite large, violating the AVGS specification. Table 7 documents the bias from the original mated ground test sample. A comparison of Table 7 to the AVGS specification in Table 1 reveals that the initial biases greatly exceeded specification. For instance, the Pitch bias was - 0.54 deg compared to a specified accuracy of 0.2 deg.

Table 7) Preflight Mated Biases

Range (m)	Az (deg)	El (deg)	Pitch (deg)	Yaw (deg)	Roll (deg)
0.000	0.024637	-0.001719	-0.547150	0.298424	0.539984

The bias correction was tested in December of 2006 near the end of the time the mated stack was in the clean room in Titusville. The results appear in Table 8. As can be seen, the bias calibration correction was successful. It only remained to see how well it would hold up in flight, after the separation ring was removed. Note that the AVGS results in Table 7 display a 2.5 mm bias in Range. This was deliberately introduced as part of the Bias Calibration Procedure to compensate for the presence of the separation ring during launch and early operations.

Table 8) Preflight Corrected Mated Biases

Range (m)	Az (deg)	El (deg)	Pitch (deg)	Yaw (deg)	Roll (deg)
0.0025	0.006	-0.008	0.006	0.000	0.010

The correction held up very well, and the small differences were well within the internal tolerance established to decide if the Bias Calibration Procedure should be repeated, and were of course well within specification. Also, the small differences in Table 8 are commensurate with the typical variability of the sensor due to low frequency noise, which can be corrected to an arbitrarily small value for a given sample of data, but will vary some small amount from sample to sample. The next section goes into more detail on the low frequency noise and its implications for sensor accuracy.

The final and most important test for the Bias Calibration Procedure was on-orbit. Unfortunately, the sensor could not be tested prior to separation ring blow-off, because the separation ring included covers over all the proximity sensors. However, after the separation ring was removed on April 17, 2007, the AVGS was able to actively track the NEXTSat as it was berthed (the NEXTSat was removed and replaced using a robotic arm on the ASTRO in order to remove the separation ring). During this initial operation of the AVGS, no mated data was available. However, on April 18, 2007, a calibration test was performed for all sensors. During this test, the AVGS was able to obtain a sample of data. The biases in the data appear in Table 9.

Table 9) On-orbit Mated Biases

Range (m)	Az (deg)	El (deg)	Pitch (deg)	Yaw (deg)	Roll (deg)
0.000	-0.009	-0.024	0.028	-0.038	0.094

Although larger than the biases on the ground with the separation ring intact, the biases in Table 9 are well below specification. They did exceed the internal standard for the Bias Calibration Procedure. That is, the results in Table 9 would require another iteration of the Bias Calibration Procedure. However, the preflight agreement with DARPA and Boeing was that if the AVGS bias was within specification it would not be repeated on-orbit**.** Therefore an in-flight Bias Calibration Procedure was not performed. During nine additional mated data takes, the AVGS bias remained constant within +/- 0.05 deg in attitude and +/- 0.01 deg in bearing. No bias in Range was ever detected. The slight variations in bearing and attitude were due to the low-frequency noise, with perhaps some variation in the docking mechanism as well. The investment of time and energy before the flight in developing the Bias Calibration Procedure and effort expended in persuading DARPA to allow the procedure in the mated configuration in Titusville was well spent and paid great dividends on-orbit, as the AVGS was used for six unmated operations, including three

dockings, two berthings and the final departure of the NEXTSat from ASTRO during End of Life (EOL) operations.

LOW-FREQUENCY NOISE SIGNATURE

The largest source of error that is not corrected in the AVGS is low-frequency noise error. An example of the noise appears in Figure 4. In Figure 4, pitch is plotted for a 30-minute segment of ground test data conducted under very quiescent conditions. The data displays a strong periodic component with a duration of approximately five minutes.
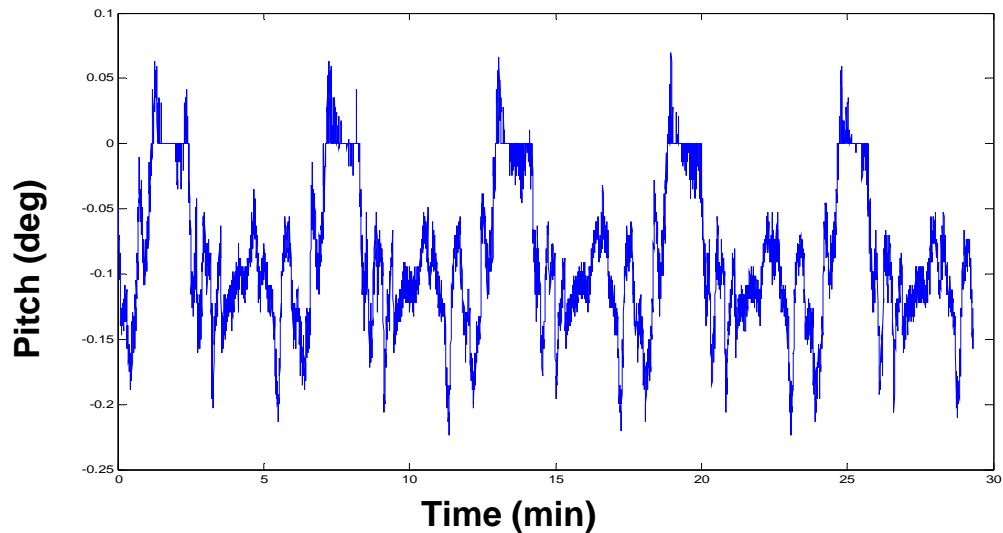

Figure 4) AVGS Pitch Behavior in Quiescent Ground Test
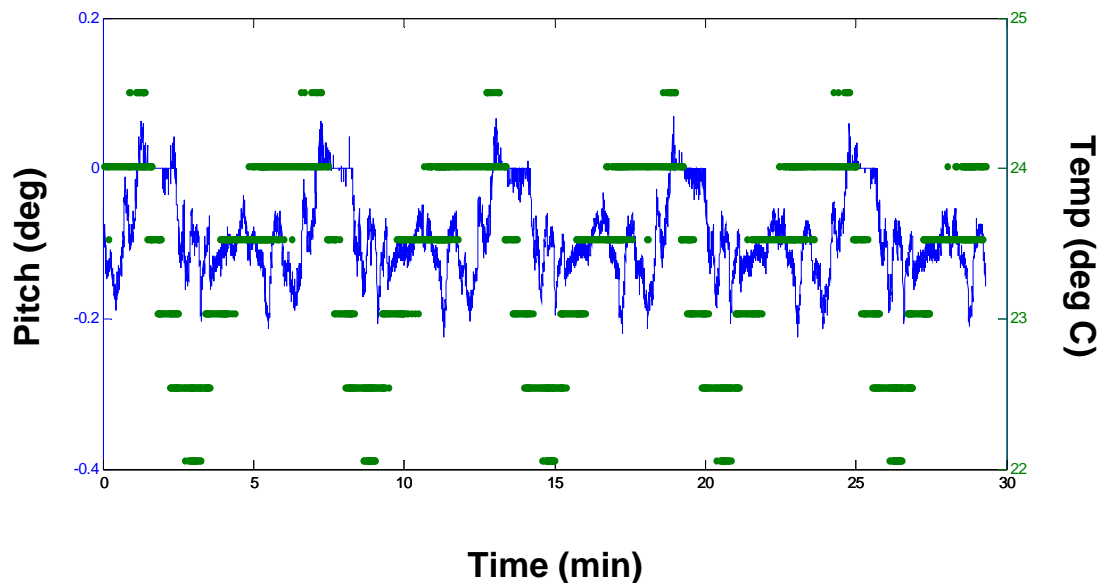
Figure 5) AVGS Pitch Noise Correlation to Temperature

Figure 5 is a plot of a temperature sensor vs. the same pitch data for the same time period. In Figure 4, the right axis shows the scale for temperature in deg C, and the temperature is represented in the plot with

points. The left vertical axis shows Pitch in deg and is represented by the line. While the temperature data lacks resolution and is limited to just six values, it is still clear that the noise signature displays a strong correlation to temperature data. The precise reason that the temperature and solution are correlated so tightly is unknown at this time. It could be the effect of the temperature on the lens, the focal length of the imager, the imager itself, or some other factor. The best theory to date is that the image chip is moving due to thermal stress in the pins that hold it in place.
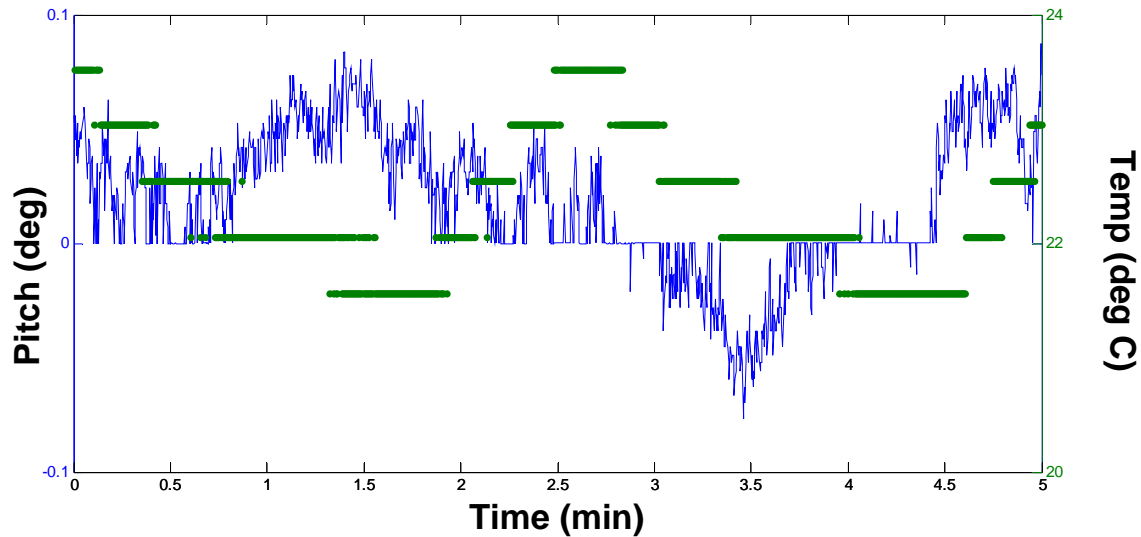

Figure 6) AVGS Solution-Temperature Correlation from Flight AC3 Test

The behavior in Figures 4 and 5 was repeated on-orbit. Figure 6 is a plot of the same two parameters for a long segment of mated data called the "AC3 Stress Test" (the AC3 is one of the ASTRO flight computers). Interestingly, the period in Figure 6 has changed to ~3 minutes from ~5 minutes in Figures 4 and 5. The reason is that the space environment induced a different thermal reaction from the sensor. The heaters were forced to turn on nearly twice as often on-orbit as they were at room temperature on the ground, thus nearly halving the period of the thermal cycle. Since the Pitch also displayed the same change of frequency, the on-orbit results strongly re-enforce the correlation between temperature and solution low-frequency noise.

In all three of Figures 4, 5, and 6, we can see that the variation in Pitch is roughly +/- 0.1 deg. Since this variation is well within the one-sigma specification of 0.2 degrees, no corrective action is required. However, if greater accuracy were needed, several methods of mitigation could be considered. If the problem is indeed caused by the mounting of the imager chip, then a structural change could fix the problem. If the problem cannot be traced specifically to structural issues, then the thermal cycle of the sensor could be adjusted with different heaters to control to a narrower temperature range, provided that the mass or power penalty of such a solution is not severe. If a hardware solution could not be found or was undesirable for some reason, it's very likely that a software change incorporating temperature data could be used to adjust the solution and at least partially remove the low-frequency noise.

Additionally in Figure 6 it can be seen that there is a phase shift when compared with Figure 5. The phase shift sometimes occurs as a function of time for a single long-duration sample. The most likely link between temperature and its effect on the solution is the pins holding the image chip bending under thermal stress. If so, such an explanation would be consistent with the effect showing the same frequency dependence but shifting in phase from one sample to another. While this is an attractive theory, at the current time the exact link between temperature and solution is not definitively known.

Another aspect of Figure 6 that contrasts with Figure 5 is that the behavior of the Pitch is much less periodic. This reflects the fact that the on-orbit environment was not as quiescent as the test data in Figures

4 and 5. In fact, the ground test represented by those figures is notable for how quiescent it was. Typically atmospheric turbulence and structural vibration play a role in adding environmental noise to a given AVGS data sample during any type of test. However, the test in Figures 4 and 5 was performed around 4 AM local time with minimal disturbances from the environment, and the results reflect this. Without this test, the phenomena of low-frequency noise would probably have escaped notice.

In summary, although the AVGS is highly accurate and well within specification, if required the accuracy could be increased by removing or greatly reducing the largest remaining source of error.

TARGET DESIGN LESSONS LEARNED

Target design is also very important to the accuracy and reliability of a proximity sensor. In the case of the Orbital Express AVGS, we learned many important lessons in this area. Some of these lessons could be highly beneficial to future missions that use a proximity sensor that uses a target.

A picture taken on-orbit of the SRT configuration is located in Figure 7, and the LRT configuration is in Figure 8.

Figure 7) AVGS SRT Assembly on the NEXTSat during Separation Ring Blow-off
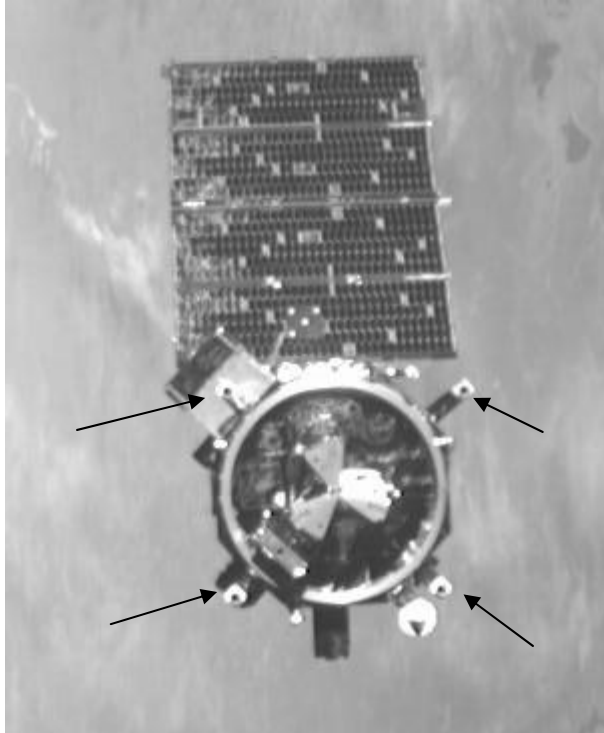
Figure 8) AVGS LRT Assembly during Proximity Operations (arrows indicate corner cube locations)

The traditional (classic) target configuration for the VGS and AVGS is to have three retro-reflectors in a row, equally spaced, with the middle retro-reflector offset on the side facing the chase vehicle (i.e., closer to the chase vehicle) in the docked position. The top three spots of the SRT in Figure 7 are an example of this pattern. This target design allows easy pattern identification for everything but Roll (the rotation axis perpendicular to the plane containing the base of the target). For Roll, another identifying feature must be used. SRT spot #3 (see Figure 1) was used for "Roll discrimination" (the process of identifying Roll, which can be ambiguous in the presence of target symmetry). In contrast, the OE LRT had no easy way to do Roll discrimination.

From Figure 8, it can be seen that the LRT corner-cubes on the NEXTSat are arranged in almost a square pattern. It is intuitively easy to see how this would cause problems for pattern recognition, as a square is symmetric, making pattern recognition and particularly Roll discrimination difficult.

The difficulties with the LRT pattern recognition forced the AVGS team to incorporate the SRT into the LRT pattern recognition algorithm. This solution worked, but also complicated the optical tuning of the sensor, as it was now necessary to view the SRT at all LRT operational ranges in order to acquire the pattern. In addition, at sufficiently high tilt angles the SRT can merge with one of the nearer LRT spots. Certain tilt angles for which the SRT is merely near one of the two nearest LRT spots can also cause ambiguities. These problems were eventually resolved, but added greatly to the complexity of the flight software.

Based on the above and other considerations, we present the following guidelines for future target design to aid pattern recognition, enhance solution convergence, and improve accuracy:

1. Have at least 3 spots and have a spot offset from the others in the direction of or away from the sensor along the sensor Line of Sight (LOS).
2. The offset spot should ideally be located between the other two spots, so that when viewed from the sensor with no tilt angle or rotation, the three spots should appear to form a straight line.

3. The amount the center spot is offset should be large enough that it will always have a range that is consistently either larger or smaller than that of the other two spots throughout the sensor operational envelope. (The offset spot should never have a range that is between the other two spots.)
4. For Roll discrimination, a 4th spot is required.
5. Symmetric arrangements of retro-reflectors such as squares or rectangles must be avoided.
6. The ratio of the offset spot to the two other spots has certain values that are desirable for pattern recognition.

In summary, target design is an important feature of sensor performance and should not be neglected. Not only can accuracy suffer with a poor design, pattern recognition can be adversely affected as well, with unnecessary and costly complications to flight software. It should also be noted that targets have been used with all visual proximity sensors that have achieved successful automated docking operations to date (the ATV ESA (European Space Agency) Videometer, the OE Boeing sensors and OE NASA AVGS). Based on a broad and deep experience, accomplishing an automated docking using a visual proximity sensor without a target would be incredibly difficult.

OTHER LESSONS LEARNED

The scope of this paper is limited to lessons learned that directly affected the accuracy and calibration of the paper. However, other lessons were learned in areas not directly related to calibration and validation that had an indirect effect on sensor accuracy. Here we present only a brief overview of these. More information is available in many cases from other sources.

One of the biggest factors in sensor accuracy is the optics of the lens. The lens designed for AVGS had optical defects that caused distortion of the image across the Field-of-View (FOV). These effects had to be corrected via empirical software corrections to the raw spot locations from the imager. This process is discussed in great detail in the paper by Howard, et al. Although this issue did affect sensor accuracy, it was resolved by the time the Systems Performance Test took place, and so was not an issue to be discussed in this paper.

Another issue that can affect sensor accuracy is also related to optics. The determination of the effective focal length of the sensor plays an important role in mapping the raw pixel data from the imager into spot locations that can be processed by the flight software. An incorrectly measured focal length can introduce errors that vary as a function of range into the solution. The AVGS had some early issues with a poorly determined focal length, but again these were resolved prior to formal calibration.

The AVGS uses the Texas Instruments (TI) TMS320VC33 card for the calculation of the solution. This card was selected based on throughput for imager processing. Unfortunately, due diligence was not completely performed and as it turned out the card was limited to single precision (32-bit) numbers. During the DART mission, this caused an error that grew with range to truncation issues with small angles between the spots. Fortunately, the card also offers an "extended precision" (40-bit) number, and this was sufficient to fix the problem. However, double precision numbers (64-bit) are still the preferred approach by a wide margin. In addition to precision issues, the card also had its own math package for processing that was not as precise as the typical industry standard and was replaced by a shareware math package.

Since the problem with the processor was solved as part of the DART project, it was not a primary focus of this paper, but is mentioned here to illustrate the law of unintended consequences. A decision might be good for one aspect of the sensor design, but negatively impact another.

CONCLUSIONS

Many valuable lessons have been learned from the validation and calibration of the AVGS sensor for Orbital Express. One major lesson is to be careful to assess the impact of one part of the design on another. The issues experienced with target design and the selection of the processing card strongly illustrate this lesson. For target design, some basic ground rules have been established to aid future design. We also

learned valuable lessons on the importance of optics to the accuracy of the sensor. Future projects could benefit greatly from the optics lessons in particular. We also gained insight into how the low-frequency AVGS noise behaves over time, as well as ideas on how to correct it. While this lesson is more specific to AVGS, it could greatly enhance future performance of the AVGS if necessary, and could give other sensors ideas on how to process noise as well. Overall, it should be stressed that a total systems approach for future vision-based proximity sensors should include optics, processor cards, target design, truth data determination, test design and rigorous data analysis.

Many years of development of the AVGS culminated in a highly successful flight on the Orbital Express mission that, among other accomplishments, achieved the first American automated rendezvous and docking in space. The insight gleaned from this experience provides essential guidance for future proximity sensors.

REFERENCES

[1]Howard, R., Bryan, T., Book, M, and Dabney, R., "The Video Guidance Sensor – A Flight Proven Technology", *Proceeding of the 22nd Annual American Astronautical Society Guidance and Control Conference*, Vol. 101, pp 281-298, 1999.

[2]Polites, M.M. "An Assessment of the Technology of Automated Rendezvous and Capture in Space", NASA/TP-1998-208528, July 1998

[3]Calhoun, P. C., and Dabney, R., "A Solution to the Problem of Determining the Relative 6 DOF State for Spacecraft Automated Rendezvous and Docking", *Proceedings of SPIE* – Volume 2466, 1995, pp. 175-184.

[4]Howard, R.T., Johnston, A. S., Bryan, T. C., and Book, M.L., "Advanced Video Guidance Sensor (AVGS) Development Testing", *Proceedings of SPIE Space Technologies and Operations Conference*, 2004.

[5]LeCroy, J. E., Hallmark, D. S., and Howard, R. T., "Effects of Optical Artifacts in a Laser-Based Spacecraft Navigation Sensor," in *Sensors and Systems for Space Applications*, edited by R. T. Howard and R. D. Richards, SPIE Conference Proceedings Vol. 6555, Bellingham, WA, 2007, pp. 1-11.