Design and Performance of the BBRISON UV-VIS Fine Pointing System

Jed Diller
Southwest Research Institute
1050 Walnut St.
Boulder, CO 80302
303-546-9670
dillerj@boulder.swri.edu

Kevin Dinkel
Southwest Research Institute
1050 Walnut St.
Boulder, CO 80305
303-546-9670
dinkelk@boulder.swri.edu

Zach Dischner
Southwest Research Institute
1050 Walnut St.
Boulder, CO 80305
303-546-9670
dischnerz@boulder.swri.edu

Nick Truesdale
Southwest Research Institute
1050 Walnut St.
Boulder, CO 80305
303-546-9670
truesdalen@boulder.swri.edu

Eliot Young
Southwest Research Institute
1050 Walnut St.
Boulder, CO 80302
303-546-9670
eliot.young@swri.edu

***Abstract*—Balloon-borne telescopes have the potential to provide space-quality imaging, but long exposures require pointing stability on the order of the optical point spread function (PSF). Consequently, one of the main goals of the BRRISON (Balloon Rapid Response for ISON) payload was the characterization of a fine steering system. The BRRISON telescope will be stabilized at the 5 arcsecond (RMS) level; the job of the fine steering system will be to improve the pointing by another two orders of magnitude, from ±5" to better than ±0.05". Due to a flight anomaly, we will discuss the design and ground test performance of the two main components of the fine steering system: the fine guidance system (FGS) and the fine steering mirror (FSM). The FGS was designed with the goal of providing a 20 Hz optical reference signal with 0.01" positional acuity. The FSM is specified to provide 100 nRad control of the mirror, translating to 0.0024" on the sky. We report on the measured quality of the error signal and widths of observed FWHMs during ground testing.**

TABLE OF CONTENTS

1. INTRODUCTION AND BACKGROUND……1

2. HARDWARE……...2

3. SOFTWARE……3

4. TESTING……..……..7

5. RESULTS……...…….8

6. CONCLUSIONS……..……...….9

REFERENCES……9

BIOGRAPHY……..9

1. INTRODUCTION AND BACKGROUND

The BRRISON (Balloon Rapid Response for Comet ISON) project was conceived with the science goals of collecting infrared, ultra-violet, and visible image data on Comet ISON and other planetary targets. The performance characterization of a fine steering motion compensation system was a primary project goal and is the focus of this paper. The BRRSION high altitude balloon gondola is composed of a 0.8m telescope, three axis control system, and two instruments. The John Hopkins Applied Physics Laboratory (APL) built the gondola and BIRC infrared instrument. The UVVis instrument housing near ultra-violet and visible detectors and the fine steering system was built by Southwest Research Instrument in Boulder, CO.

The primary challenge while gathering planetary target data is the small, far away, and dull nature of the targets. To get high quality images from a balloon in the stratosphere, a large aperture telescope must be pointed to the arcsecond level while being blown around by high altitude winds and acting like a 600 ft pendulum. The BRRISON gondola attempted to achieve this level of pointing using the three part control system highlighted in Figure 1 below.

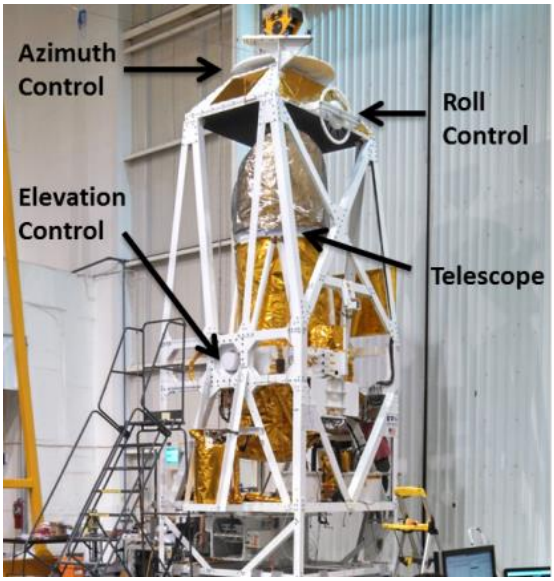

Figure 1: BRRISON gondola and course control components

The gondola is decoupled from rotation in the flight train and balloon by a large disk mounted on top of the gondola. This large disk controlles the yaw of the gondola. The pitch of the telescope is controlled by two axial motors at the telescope and instrument's center of mass. Many balloon borne gondolas use just these two systems. BRRISON's predecessor gondola (and the project from which the mirror was reused) the Stratospheric Terahertz Observatory (STO) used just these two systems. The BRRISON course control system included an additional roll compensation wheel near the top of the gondola. These systems together allow for <5 arcsecond pointing of the telescope.

978-1-4799-1622-1/14/$31.00 ©2014 IEEE

The UVVis fine steering system was designed to compensate for the motion left over by the course pointing system and improve the overall performance by two orders of magnitude. Using STO flight data as a guide, the residual motion from the course control was expected to be less than 5 arcseconds in elevation and azimuth and almost all of the power at frequencies less than 1Hz. Because BRRISON had heritage in STO, it was assumed that the BRRISON performance would be on the same order or better with the addition of the roll compensator. Using this information the UVVis fine steering system was conceived with the goals of operating at 20Hz and being able to compensate for motion at <5Hz.

The UVVis instrument is composed of two mechanisms, two detectors, static optics, an aluminum honey core bench, struts connecting it to the telescope, heaters, and a pressure vessel full of electronics containing the power system, computer, and electronics to control the detectors, mechanisms, and fast steering mirror.

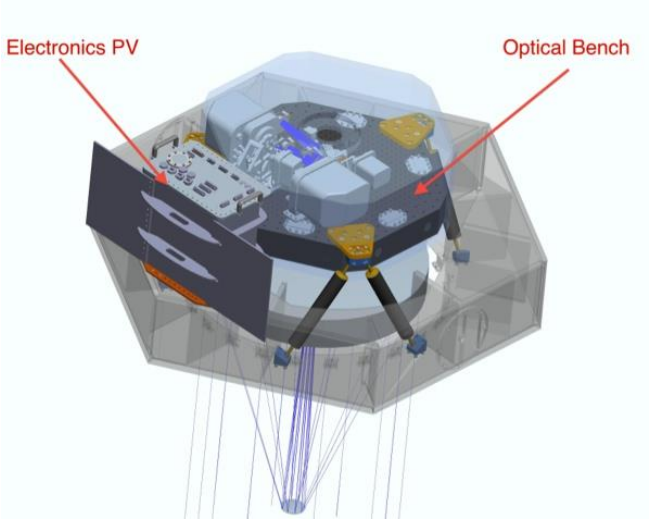

Figure 2: UVVis instrument model including telescope structure and ray tracing.

2. HARDWARE

Control Loop Components

The control loop is composed of six major pieces of hardware and large amount of software. Following the data path in from the telescope, light from a star or other target coming through the telescope is bounced of a retractable fold mirror and then bounced off of a high performance fast steering mirror (FSM). After bouncing off the FSM, light passes through the rest of the optics and split through a diacrylic. Ultra-violet light passes through to the science detector and redder wavelengths are directed to the guide camera. Images from the guide camera are read in by a frame grabber interfacing to the mother board via PCIe. In CPU memory the images are processed through a digital control law that determines what voltages the DAC should output to the FSM controller board to move the FSM and minimize error. Those voltage are output to the FSM's controller board before actually being sent to the FSM itself. The FSM tips and tilts, changing the direction of the light coming in and changing where the target is located on the guide camera's detector.

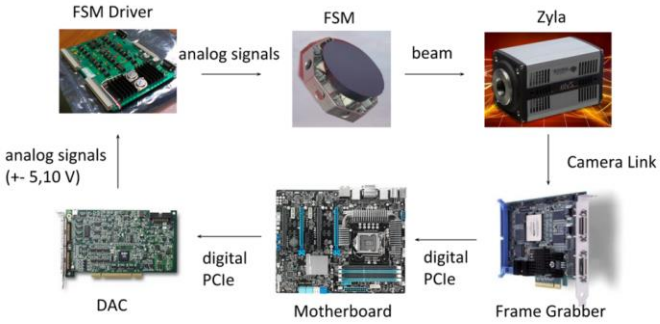

Figure 3: Fine steering control loop hardware.

FSM and Controller Board— The FSM used for fine pointing in the UVVis instrument is made by Left Hand Design Corporation in Longmont, CO. For fine steering purposes a mirror was chosen with the smallest range and jitter to maximize step resolution. The FSM chosen was a two-axis, 60mmx80mm mirror with a ±5 mrad range, a jitter of 100 nrad, and a bandwidth of 2 kHz. For the f/17 BRRISON telescope, the jitter translates to 0.00243" on-sky. The FSM is actuated by feeding ±10V differential voltage commands and logic signals to the FSM's controller board. The FSM was mounted to the optics bench while the controller board was housed in the electronics pressure vessels with the rest of the electronics.

Guide Camera— The Andor Zyla was chosen as the UVVis guide camera The Zyla was an attractive candidate for its high frame rate, pixel size of 6.5x6.5 micron (and plate scale of 0.09 pixels/arcsecond), as well as prior team experience with the sCMOS, Camera Link, and frame grabbers. The Zyla uses two-tap Camera Link to transfer 16bit, 5.5 Mpixel images as fast as 100Hz. Shrinking the region of interest down from full frame allowed the SwRI team to capture images as fast as 655Hz during testing.

Frame Grabber— Images taken by the guide camera are transferred into computer memory using a frame grabber and its software development kit (SDK). The frame grabber reassembles images after they are transferred from the Zyla via two-tap Camera Link protocol. The frame grabber interfaces to the motherboard via a PCIe slot. A BitFlow Karbon-CL was chosen for its two-tap configuration and compatibility with Linux.

Motherboard— The mother board was a consumer COTS gaming motherboard, the ASUS P8 Z77 WS, with four PCIe slots and an Intel i7 3770K processor.

Analog I/O Board— Analog and digital signals sent to the FSM's controller board needed to be generated quickly and with high resolution. With a goal of fine steering correction on the 0.05" level, the resolution of a FSM step was desired

to be a factor of 10X greater than that on-sky. The FSM's ±5mrad range and ±10V command inputs defined what level of resolution was needed from the analog I/O board. Over ±10V, 16bits resolution results in a 305uV step resolution which converts to 0.0315" in mirror tilt and 0.0037" on-sky step size. A Sensoray Model 826 PCI Express I/O Board was chosen for its 16 bit resolution and compatibility with Linux. The I/O board was used for mechanism control in additionally to FSM commanding.

3. SOFTWARE

The software comprising the digital control law determines how the FSM should be moved to compensate for motion detected between guide camera images and how to do that to achieve 0.05" error. The digital control law can be broken into two main parts, start identification and centroiding and FSM command determination.

Digital Control Law Overview

The logic of the control law is as flows. In a given image candidate stars are searched out and a rough center for each candidate star is found. These candidates are then filtered. If the candidate is too big, small, misshapen, etc. based on variable parameters, the star is rejected. If no candidates make it through this filtering the control law restarts, there is no useable candidate in the field of view. If a candidate star does make it through this filtering it is then compared to the other successful candidates. A refined centroid if then found for the brightest candidate. This centroid is compared to the first iteration in a given run of the control law and the difference in centroids determines the current error. Using the filtering it can be ensured that the same star is used between frames. Once the error between frames is found, that error is handed over to an algorithm that determines what the FSM motion should be in response. That motion is then converted into command voltages that the I/O board sends to the FSM controller. The control law was designed to operate at a minimum of 20Hz. To achieve the desired performance of <0.050" error with a relatively low bandwidth compared to the frequency being compensated for, a large part of the command determination was a predictive algorithm to determine FSM commands based on information from past and current iterations. It should further be introduced that the speed at which the control law could run was dependent on the time elapses taking and transferring an image, doing the command calculations, and the time between issuing commands to the I/O board and it changing of the position of the FSM. This iteration time was dominated by the time elapsed taking and acquiring images.

Star Finding and Centroiding

For each image a robust median of the background was calculated. This robust median consisted of finding the median of the background and then rejecting outliers outside of a variable sigma threshold and recalculating the median. Because median background finding is computationally costly, a histogram median method was used in addition to varying the density of pixels sampled in the image array.

With the background median of the image found, candidate stars were searched out using a depth first search for pixels sufficiently above the background, as defined by another variable parameter. In the search if a pixel was found above the background, it would be considered a "good pixel." This good pixel would then be flagged as visited so it was not revisited and its neighbors would be searched in the same manner. Using this recursive search, blobs of good pixels would be identified as candidate stars.

These candidate blogs were then filtered. If a blob was too small, say a single hot pixel, it would be rejected. If it was too big, say a planet in the FOV rather than the planet's satellite that is the desired target for centroiding, it too would be rejected. An oblong check ensured that candidate was near round and not a cosmic ray streak. Candidates could also rejected if they had saturated peaks because it would throw off the refined centroid. Candidates intersecting an edge of an image were also rejected. Only candidates that passed all of these filters were candidates for centroiding. The brightest candidate was then chosen because the higher signal to noise resulted in a better centroid.

The single brightest target chosen was subframed by a factor of its height and width. In that subwindow an IWC2 centroid of the target was calculated and related back to pixel space on the detector. This centroid compared to the original frame in a run of the control law gave the current iteration error.

During instrument operation the parameters that determined the target for centroiding could be changed on the fly. This gave the UVVis team a robust and adaptive way chose an object for centroiding in any viable field of view. Depending on the settings, the fine steering system could target a planet's moon, a background star near a comet, or one or the other star in a binary pair. In ground testing this parameter tuning was done manually by downlinking an image and varying the parameters until the desired target was being centroided. This was done with a C++/Python interface using C++ centroiding functions identical to the flight code.

FSM Command Determination and Predictive Algorithm

With a minimum control law frequency of 20Hz, achieving <0.050" error and a 5Hz bandwidth through typical PID control is impossible. Assuming 5" peak to peak amplitude, 0.5Hz residual motion from the course control system and a 20Hz control law frequency the problem of delay becomes apparent in Figure 3 below.

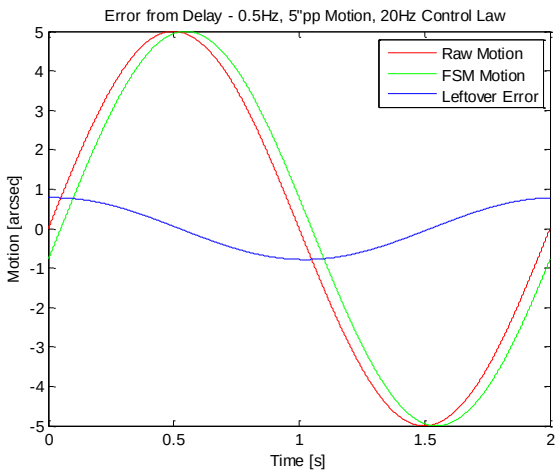

Figure 3: Comparison of ideal PID control at 20Hz for a 0.5Hz, 5" peak to peak input.

In red, idealized raw motion left over from course control is shown. Green depicts ideal FSM motion in response to the raw motion with the control law running at 20Hz. In reality the FSM motion would be in the opposite direction of the raw motion but it is plotted in the same orientation to better show the delay. The 50ms it takes to detect an error and move the FSM in the opposite direction proportionally leaves a leftover one sigma error of 0.55", about a factor of 10 larger than the goal error. To achieve better performance the control law frequency would have to be increase many times over but that was not an option here.

The above plot is idealized because in reality, if the control algorithm ran at 20Hz and only updated the FSM to an opposite proportional position, the FSM motion would instead look steppy as it does in Figure 4 below.

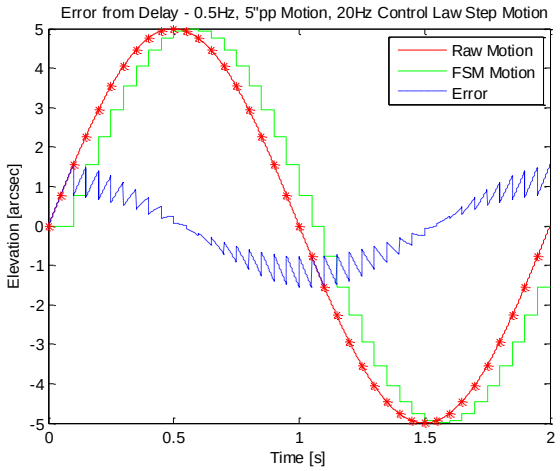

Figure 4: Steppy response of FSM due to 20Hz update rate.

The red asterisks show when the raw motion was sampled. And in greed the delayed FSM position is shown. The blue error now includes the error between the steps and the raw motion. What Figures 3 and 4 really show is that to achieve 0.05" level error an element of prediction is necessary. The control law can not simply react to its input, it has to anticipate it. Figures 3 and 4 are further idealized and just used to communicate the problem because they do not actually include feedback. They are the open loop examples. Figure 5 below shows the closed loop response to the same input.

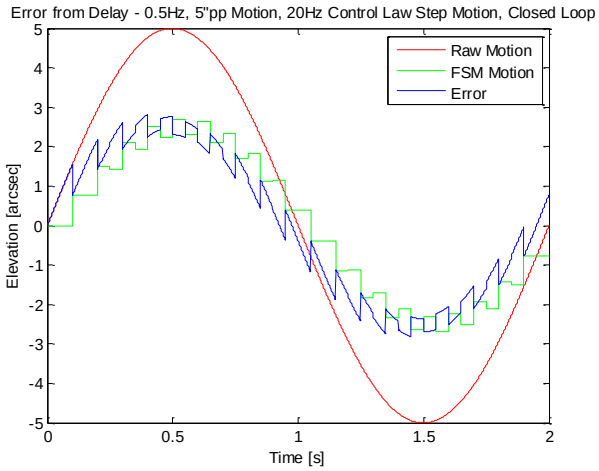

Figure 5: Closed loop response with 20Hz proportional FSM steps.

In this case the one sigma error is much worse, being 1.81". Analysis of the same proportional step response shows that even if the control law were to run at 1kHz, the error would still be greater than 1". This analysis made it quite clear that prediction was needed.

The easiest prediction to implement is a linear fit using a past and current position and adding the difference between the two (a velocity) to the current position, giving a position in the future. When executed at the end of a loop iteration, a future point calculated from current information is actually the appropriate current point. The tricky part of doing this sort of prediction continuously is keeping track of the mirrors actual position. After the first iteration in which the mirror moves, the shape of the raw motion can no longer be seen, rather the only information available is the current error. To make predictions the shape of the motion has to be preserved so a variable dubbed "actual position" in this implementation keeps track of that using the error and the known previous command. The algorithm for this first order prediction is as follows.

$$caa = ce + ppaa1$$

$$v1 = (caa - paa1)/dt$$

$$cpaa1 = caa + v1*dt$$

Where caa is the current actual angle, ce is the current error, ppaa1 is the past predicted actual angle, cv is the current velocity, paa1 is the past actual angle, dt is the current

iteration period, and cpaa1 is the current predicted actual angle. After cpaa1 is calculated, caa and cpaa1 are stored and become paa1 and ppaa1 in the next iteration respectively. This algorithm assumes that cpaa1 output to the mirror is exactly where the mirror will be on the next iteration. Because the UVVis FSM is such high performance, settling to any position in the range in around 1ms, this was considered true. In a system with a lower performance mirror, ppaa1 may have to be measured via a position reference signal from the FSM just before getting a new error measurement. The control law response for the same input using this first order prediction is shown in Figure 6 below.

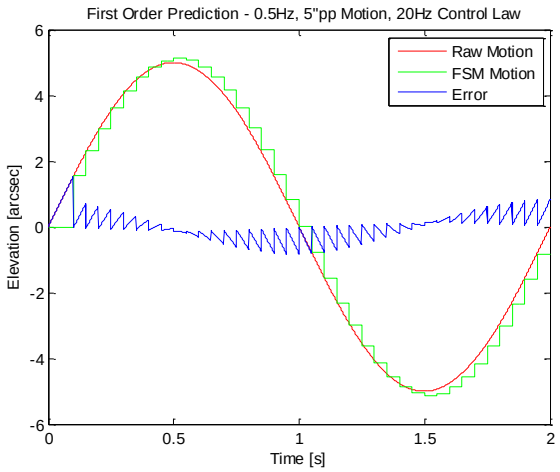

Figure 6: First order prediction response.

The first order prediction response resulted in a one sigma error of 0.33" ignoring the first three control law iterations. These have to be ignored because the mirror does not move until the third iteration. The first iteration gets the original centroid, the second iteration finds the first error, and the third iteration needs the second iteration error for prediction.

With prediction active, the steppy response from the figures before has shifted to start nearly where the raw position is. It can also be seen that where the raw motion changes direction the quickest, the beginning of the steps are the most off. Interestingly, at points where the steps start are the farthest off, the error is the smallest. This reveals a fundamental problem with only updating the FSM position at 20Hz. Even if the prediction were perfect, in places where the slope of the raw motion is large, the error will be large. For the expected motion shown above, to reach below 0.050" in error, this would require that the control loop update near 200Hz rather than 20Hz.

Since that was not thought possible during the design phase, a way around this problem had to be devised. In between prediction points, the mirror has to be updated at >200Hz and keep following some predicted actual motion until the next error calculation is made. Based on Bode plot data of the FSM, the settling time is known to be about 1ms. This meant the FSM's position could be updated with steps at 1kHz, well above the 200Hz minimum needed.

To decide what the FSM should do between control law iterations, more prediction was used. Using the current predicted actual position and the current actual position, a slope was determined and the FSM moved along that slope at 1kHz after starting at the current predicted actual angle. Each iteration the FSM restarts these slope following runs at the current predicted actual angle. This secondary prediction, prediction based on already predicted points, is perhaps the most novel part of the UVVis control law. Figure 7 below show this implementation in practice.

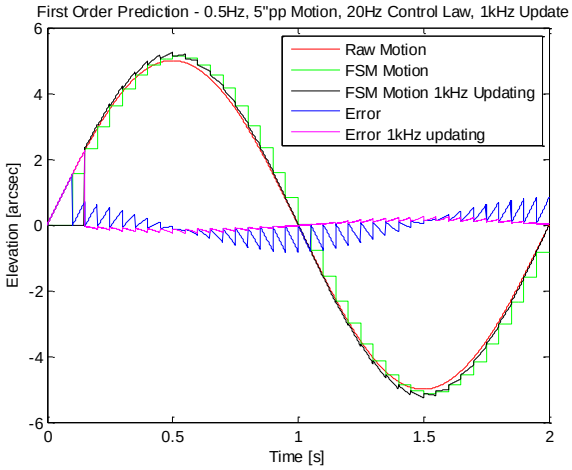

Figure 7: First order prediction with 1Khz updating of FSM position based on a predicted slope.

In black the FSM motion including the 1kHz FSM updating along a predicted slope is shown and magenta is the resultant error. The 1kHz updating technique has the least error where the slope changes the least and vice versa. With this second level of prediction and the FSM position being updated at 1kHz, the one sigma error is decreases to 0.13". A zoom in on the first peak of the raw motion better exemplifies the prediction and 1kHz updating in action. The fuzziness of the 1Khz updating motion is due to there being 50 FSM position steps per 20Hz control law iteration.

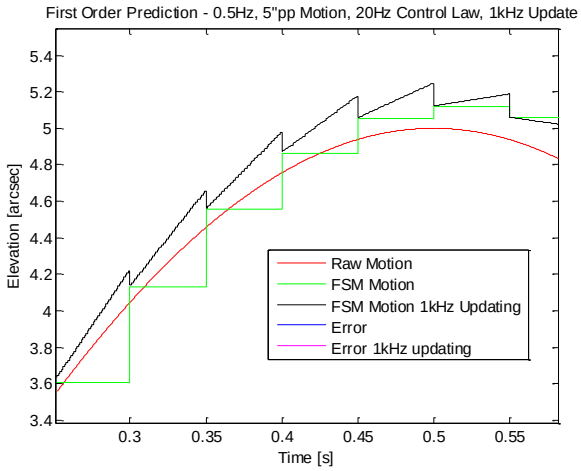

Figure 8: Zoom in of 1kHz FSM position updating for first order prediction at 20Hz.

To increase performance, the next step is to use higher order prediction to get a more accurate current predicted actual angle. Whereas the first order prediction only required one past actual angle point, seconds and third order prediction require the current point and two and three past actual angles respectively. The second and third order algorithms are as follows:

$$caa = ce + ppaa1$$

$$v1 = (caa - paa1)/dt$$

$$v2 = (paa1\text{-}paa2)/dt2$$

$$dt3 = (dt+dt2)/2$$

$$a1 = (v1\text{-}v2)/dt3$$

$$cpaa2 = caa + v1*dt + a1*dt\text{^}2$$

$$v3 = (paa1\text{-}paa2)/dt4$$

$$dt5 = (dt2+dt4)/2$$

$$a2 = (v2\text{-}v3)/dt5$$

$$dt6 = (dt3+dt5)/2$$

$$j1 = (a1\text{-}a2)/dt6$$

$$cpaa3 = caa + v1*dt + a1*dt\text{^}2 + j1*dt\text{^}3$$

Where v2 and dt2 are the velocity and time between the past actual angle one iteration in the past, paa1, and the past actual angle two iterations in the past, paa2. From v1 and v2, an acceleration term, a1 can be found. The second order current predicted actual angle, cpaa2 is then calculated by adding position from the velocity and acceleration terms to the current actual angle. Similarly for third order prediction, another velocity term is found from paa2 and the past actual angle three iterations in the past, paa3. A second acceleration term can be found using v2 and v3 and a jitter term, j1, can be found using a1 and a2. The third order current predicted actual angle uses the jitter term in addition to the acceleration and velocity.

In the same manner, second order prediction was added to the 1kHz mirror updating. Rather than the mirror following a slope from the current predicted actually angle, the FSM would follow a second order line. This was a configurable option and during testing, first order prediction of the 1kHz updating was used. Figure 9 below shows third order prediction with first order 1kHz FSM stepping.

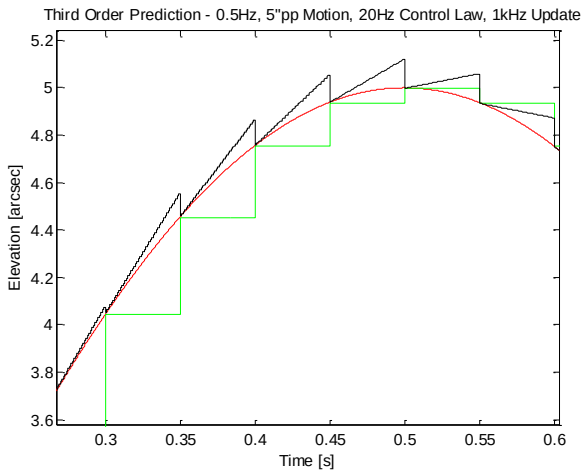

Figure 9: Third order prediction zoom in comparison to first order prediction.

Using third order prediction the FSM position steps start much closer to the actual motion than for the first order prediction in Figure 8. The one sigma error for third order prediction is 0.041", under the 0.50" goal. Using the second order prediction for the 1kHz mirror stepping further decreases this error to 0.014". This prediction seemed to work well in the model but in the real world, motion is not perfectly sinusoidal like above. Higher frequency noise than can be corrected for will be present. This problems leads to how this prediction algorithm was implemented in practice.

Actual Implementation

Because the position errors found between centroids in each iteration would include high frequency motion that could not be compensated for by the control law, a way to smooth the incoming data was needed. The control law would otherwise be reacting to noise and making that noise worse rather than removing slower motion. To minimize this effect in determining a predicted actual angle, a polynomial fit was done on a buffer of the current and past actual angles. Using a least squares fit, a given number of points, and the number of coefficients desired, the current predicted actual position was determined. The first, second, and third order prediction shown in the prior section were instead achieved in practice by using this polynomial fitting. During testing first and seconds order prediction was used most often. Higher order

polyfitting did not yield better results and was more computationally heavy.

Determining how many points to use for a fit depended on the motion observed and the frequency of the control law. To determine the number of points quickly during flight, two rules of thumb were derived. The first is that the number of points used should encapsulate the fitting order minus one inflection points of the motion being corrected. The second rules was that for first order fitting $1/8^{th}$ of the motion's period was desired and for second order $1/4^{th}$ of the period was desired. During flight and ground testing, the polyfitting order was determined by downlinking files containing control law data that were then parsed and graphed. After examining the motion left over by the course control system, the order number and number of points were chosen and retested.

In the software architecture the main control loop that did the star identification, centroiding, and point prediction was contained in one process. It was iterated using interrupt timers at a variable frequency or in a mode where it ran as fast as possible dubbed, "bashee mode." The control law was designed to run at >20Hz. In practice, achieving 50 and 100Hz was routine. As mentioned before, the speed of the control law is primarily dominated by image acquisition. The 60" field of view of the telescoped only covered about a 700x700pixel portion of the detector. Subwindowing this region resulted in a practical minimum frequency of 50Hz. During testing with a static source a control loop frequency of 655Hz was reached with a 40x40 pixel subwindow around an artificial star.

Another process that ran at 1kHz using the same timing was used to convert the predicted values from the main loop to the 16 bit, ±10V range of the DAC. While the plots in the prior section show the motion in arcseconds, the actual implementation dealt with angles, velocities, etc. in pixels. The pixel values were then converted to DAC values based on the FSM's know angular range and voltage command range. The relationship from pixel space to mirror motion had to be found during integration. This calibration resulted in find a relationship of 75.4 DAC values/pixel in FSM Azimuth and 73.9 DAC values/pixel in FMS Elevation.

6. TESTING

Testing of the fine steering system was difficult, and opportunities to do so were few and far between. The FSM was only delivered two weeks before the instrument was shipped from CASA to APL.

Pre-integration Bench Testing

Before the bench was integrated with the telescope and the rest of the gondola, a linear stage, small collimator, and light source were used to simulate a single star. During this testing the calibration between FSM motion and the resultant centroid travel on the guide detector was accomplished. This testing also revealed any directional sign errors in the control law.

High Bay Testing

After integrating with the telescope, testing the control law took some creativity. The setup up include an optics table, a large collimator, two large fold mirrors (one high up on a theodolite), and careful alignment of the telescope and gondola itself. The image below shows the setup.

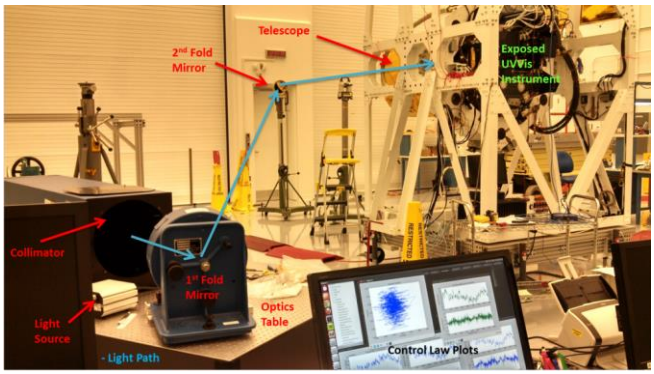

Figure 10: Control law testing setup at APL after UVVis instrument was integrated with telescope.

Blue represents the light path and the different pieces of hardware are labeled in red. On the bottom, control law plots generated by parsing logged control law information are shown. These plots provided all the information used in debugging the control law during development. It proved extremely difficult to adjust either fold mirror even the smallest amount by hand without moving the simulated star outside of the field of view or moving the star at frequencies well above what was expected during flight. Motion for testing ended up having to be induced by having someone lean on the optics table that the collimator light source and first fold mirror were placed on.

This high bay testing also revealed that the cryo-pump used to keep the IR instrument cold created lots of optical noise for the UVVis bench. Running the control law at 220 Hz revealed that the cyro pump induced nearly 5" of motion with peak power at 43 and 86Hz. While this vibration was not detectable to the IR instrument due to its detector's pixel size, this added a complication to the flight mission concept of operations. During UVVis operation it was desired that the cryo-pump be off but for every minute the pump was off, it took 7 minutes to cool back down to operational temperature. This trade off made it so UVVis operation had to take place during brief pump interruptions, deemed "cyrointerruptis", or after the IR instrument had finished all operations.

Hangar Test

On September 22^{nd}, a few days before launch and anomaly, the UVVis bench got its first opportunity for extended on-sky observation and testing in a flight like configuration. The gondola was hanging from a crane inside the Fort Sumner CSBF hangar and looking out open hangar doors at Alpharatz. Unfortunately, this particular evening was rather

windy as a storm rolled in. The gondola control system had to work hard to keep pointed as the entire gondola was visible being blown around. Certain gusts of wind would cause the target to be lost from the UVVis's 60" FOV entirely. In addition to the wind the seeing was particularly poor, on the 3" level, due to atmospheric turbulence over a still warm tarmac. Because no useful data was collected during flight due to the anomaly, the results from this windy night testing are presented in the results section.

Figure 11: The BRRISON gondola looking out the hangar door during Sept 22nd testing.

6. RESULTS

Quantifying the fine pointing systems error can easily be done but looking at the error that the control law used each iteration. This error is the difference between the first centroid in a control law run and the current one. Figure 12 below shows the current error and current actual angle.

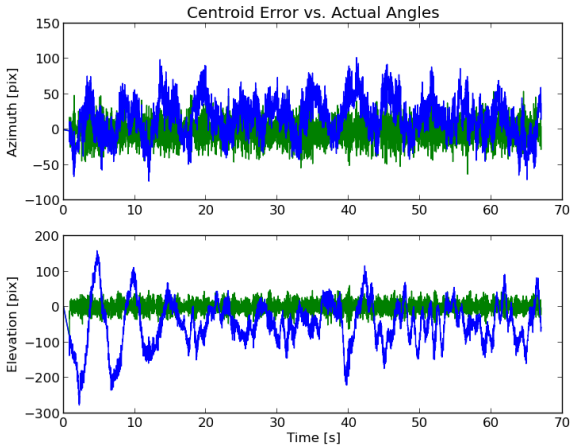

Figure 12: Control law data from Sept 22nd testing.

The current error in shown in green. In blue the current actual angle shows how the FSM had to move to result in that error. The one sigma error, given a plate scale of 0.09"/pixel, is 1.36" in FSM elevation and 1.38" in FSM azimuth.

Another metric of performance involves looking at the guide camera images themselves. Comparing the full width half max (FWHM) of a single image without correction, an image at beast focus, and a co-added stack of images with correction on gives us a sense for the fine steering performance. Figure 13 shows three consecutive images of Alpharatz taken at 100Hz.

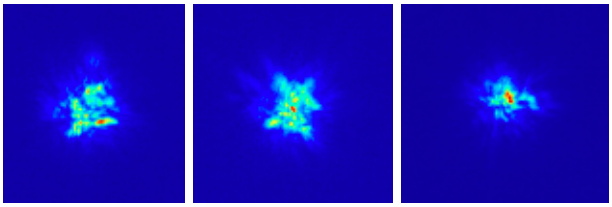

Figure 13: Consecutive Alpharatz images from Sept 22nd ground tests taken at 100Hz.

The effect of the turbulence from the tarmac and wind is immediately obvious. The stars change shape drastically between frames. A composite image of 50 consecutive frame with the fine steering system active is shown in Figure 14 below.

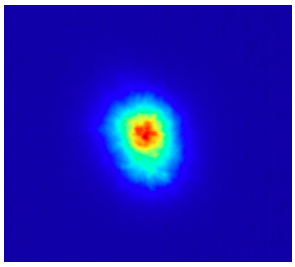

Figure 14: 50 co-added guide camera images with the control law running.

The FWHM of the 50 image stack is 1.5". This is about 10X less than a 50 images stack without fine steering active and about half of the FWHM at beast focus.

Images from the science detector tell the same story. Below, 5 seconds exposures of Alpharatz are shown with the control law off and on.

Figure 15: Science detector images during testing.

7. SUMMARY

While no flight data was collected due to the anomaly, a capable fine correction system for balloon borne payloads was built and tested. Here the design and implementation of that fine correction system were overviewed. Using a predictive control algorithm it was shown that 0.05" performance could be theoretically be achieved for 0.5Hz motion at a control law of 20Hz. While 20Hz was the minimum control law frequency, the UVVis control law was routinely tested at 50 and 100Hz and up to 655Hz. It is expected that the error achieved during ground testing would decrease by a factor of ~100X in the stratosphere environment.

REFERENCES

I need to reference some of the image and stats quoted but I am out of time.

BIOGRAPHY

Jed Diller *received a B.S. in Aerospace Engineering Sciences from the University of Colorado at Boulder in 2012. He is currently lead engineer on the SwRI Solar Instrument Pointing Platform (SSIPP) for the Southwest Research Institute (SwRI) in Boulder, CO. He was controls lead on the UVVis Instrument of the Balloon Rapid Response for Comet ISON (BRRISION) project. His role on the DayStar project was primarily electrical engineering but this shifted during the flight readiness phase to include systems and test engineering as well as project management. His former professional work includes software engineering for geophysical software companies and the University Corporation for Atmospheric Research.*

Zach Dischner *graduated with a BS in Aerospace Engineering in 2012 from the University of Colorado at Boulder. Currently, he is a graduate student pursuing a MS in Aerospace Engineering Systems. He is a graduate data systems technician for the Laboratory for Atmospheric and Space Physics (LASP). Zach is a fervent outdoorsman, whose main interests include skiing, kiteboarding, rock climbing, as well as freelance photography.*

Kevin Dinkel *is an Aerospace Engineering graduate student and a software engineer at the University of Colorado at Boulder. He has been a lead software systems engineer for 5 high-altitude balloon payloads, including BRRISON, DayStar, and BOWSER. He has also developed flight software for the DANDE nano-satellite and for an experimental payload aboard the Surrey Orbital Test Bed. Kevin previously worked on ground and thermal systems at Lockheed Martin Space Systems and BITSystems Co., and currently works as a software and systems engineer at the Southwest Research Institute (SwRI) and Laboratory for Atmospheric and Space Physics (LASP) in Boulder, CO*

Nick Truesdale *is pursuing an M.S. in Aerospace Engineering and a B.S. in Applied Mathematics at the University of Colorado. He was the Electrical Systems Lead for the BOWSER and DayStar balloon payloads, and the Power Lead for the ALL-STAR CubeSat. Through his work on these projects, Nick has gained experience in power systems, microcontrollers, PCB design and layout, and general electrical testing. He is also a skilled programmer, and combines his wide array of skills in system design. Nick has worked with Lockheed Martin, MIT Lincoln Laboratory, and is now supplements his studies as an independent contractor in electrical and systems engineering. He is currently managing the OTCCD and ISON projects for Southwest Research Institute*

MISSING ELIOT'S BIO