AAS 11-035

ONE-ARCSECOND LINE-OF-SIGHT POINTING CONTROL ON EXOPLANETSAT, A THREE-UNIT CUBESAT

Christopher M. Pong,[*] **Matthew W. Smith,**[*] **Matthew W. Knutson,**[*] **Sungyung Lim,**[†] **David W. Miller,**[‡] **Sara Seager,**[§] **Jesus S. Villaseñor,**[**] **and Shawn D. Murphy**[††]

ExoplanetSat is a proposed 10×10×34-cm space telescope designed to detect down to Earth-sized exoplanets in an orbit out to the habitable zone of bright, Sun-like stars via the transit method. Achieving this science objective requires one-arcsecond line-of-sight pointing control for the science CCD detector, an unprecedented requirement for CubeSats. A two-stage control architecture that coordinates coarse rigid-body attitude control with fine line-of-sight pointing control will be employed to meet this challenging pointing requirement. Detailed testing of the reaction wheels and CMOS detectors has been performed to extract key performance parameters used in simulations. The results of these simulations indicate that a 1.4 arcsecond pointing precision (3σ) is achievable. To meet the 1.0-arcsecond pointing requirement, several options are analyzed. In particular, a new technique to estimate reaction wheel vibrations for feedforward cancellation of reaction wheel vibrations is presented. This estimator adaptively estimates disturbances from noisy sensor measurements and effectively stores disturbance amplitude and phase in memory as a function of wheel speed. In addition to these simulation results, testing results from a hardware-in-the-loop (HWIL) testbed demonstrate the capability of the fine pointing control loop. Future plans for complete HWIL testing of the coarse and fine control loops are presented.

INTRODUCTION

ExoplanetSat, shown in Figure 1, is a concept for a 10×10×34-cm three-unit CubeSat designed to detect exoplanets using the transit method. If an exoplanet's orbital plane is aligned with the line of sight from Earth, the exoplanet will transit in front of the star causing the overall bright-

---

[*] Graduate Research Assistant, Space Systems Laboratory, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139.

[†] Senior Member of the Technical Staff, Aerospace Guidance & Control Group, The Charles Stark Draper Laboratory, 555 Technology Square, Cambridge, MA 02139.

[‡] Professor, Department of Aeronautics and Astronautics, Director, Space Systems Laboratory, Massachusetts Institute of Technology.

[§] Professor, Department of Earth, Atmospheric and Planetary Sciences, Professor, Department of Physics, Massachusetts Institute of Technology.

[**] Research Scientist, CCD Laboratory, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology.

[††] Division Leader, Earth and Space Science Program Office, The Charles Stark Draper Laboratory.

1

ness of its star to dim. Observations of these transit events will provide information on the exoplanet's radius and orbital period. The orbital period determines whether the exoplanet resides in the habitable zone of its star, a region where liquid water can exist on the exoplanet's surface. Detection of these exoplanets will enable other observatories to perform follow-up measurements such as Doppler spectroscopy to determine the exoplanet's mass and atmospheric spectroscopy to determine the composition of the exoplanet's atmosphere. All of these pieces of information form the stepping stones to determining if life exists outside of our solar system.
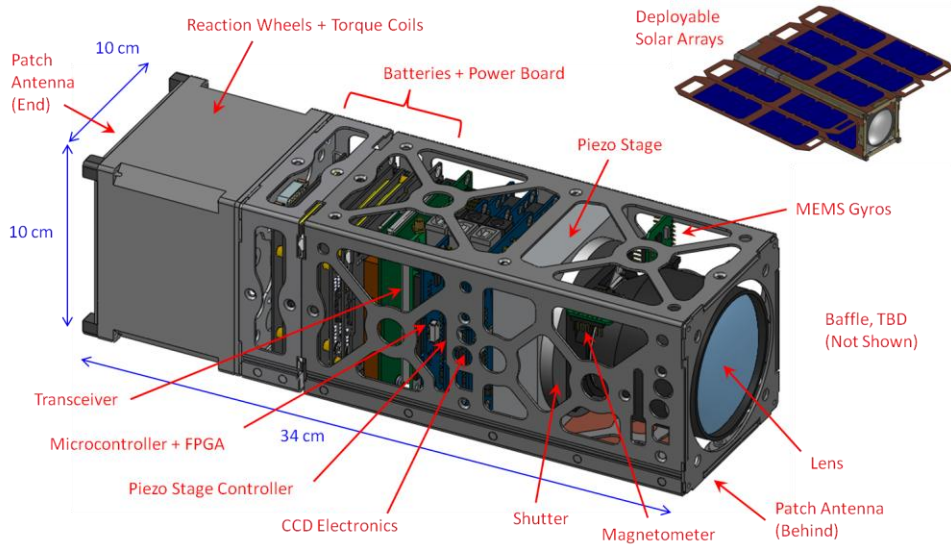

Figure 1. Baseline Design for ExoplanetSat.

These follow-up measurements, however, are only feasible to perform on exoplanets orbiting bright stars. NASA's Kepler, a satellite launched in March 2009, has a mission to provide a census of exoplanets. To observe enough stars (approximately 150,000), Kepler is imaging a field of faint stars. While this will provide valuable information on the frequency of Earth-sized exoplanets orbiting Sun-like stars, follow-up doppler and atmospheric spectroscopy measurements are not currently feasible for low-mass exoplanets around most of these faint stars. Due to the sparse distribution of bright stars across the sky and the low probability of necessary orbital plane alignment, the discovery of Earth-sized exoplanets orbiting these bright stars will rely on dedicated, inexpensive missions continually monitoring a single target. By adhering to the CubeSat standard, ExoplanetSat can fill this need and become an inexpensive platform with multiple launch opportunities. It is envisioned that a constellation of ExoplanetSats, each with a different target star, will enable the possibility of detecting down to an Earth-sized exoplanet in an orbit out to the habitable zone of these bright stars.

To detect down to an Earth-sized exoplanet, the spacecraft must achieve a photometric precision of approximately 10 ppm. One of the largest challenges with obtaining this level of precision is that the sensitivity on the main science instrument, a back-illuminated CCD, may vary on the order of 50% within a pixel.[1] Changes in the line-of-sight pointing of the spacecraft will cause the light from the target star to illuminate different parts of the CCD, which changes the measured intensity of the target star. A model of this phenomenon, known as jitter noise, has been developed and used to determine a pointing requirement for the satellite.[2] Reference 2 presented initial analyses, which only considered sensitivity variations between pixels. Subsequent analyses which

included a subpixel sensitivity model have resulted in the development of a nominal pointing requirement of 1.0 arcsecond (3σ).[*]

This tight pointing requirement has served as the main driving requirement for the satellite design. The current conceptual design for ExoplanetSat is shown in Figure 1. Coarse rigid-body attitude control to within 120 arcseconds (3σ) is provided by Maryland Aerospace Inc.'s MAI-200, a set of three orthogonal reaction wheels, which are desaturated with torque coils. Fine line-of-sight pointing control down to the 1.0-arcsecond (3σ) requirement is provided by a two-axis piezoelectric stage which translates the focal plane. An exploded view of the optical layout is shown in Figure 2. The focal plane consists of comounted CMOS detectors for star tracking and the science CCD, shown in Figure 3. Other sensors include a three-axis MEMS gyro and magnetometer. Further details on the other subsystems in the satellite can be found in References 2 and 3.
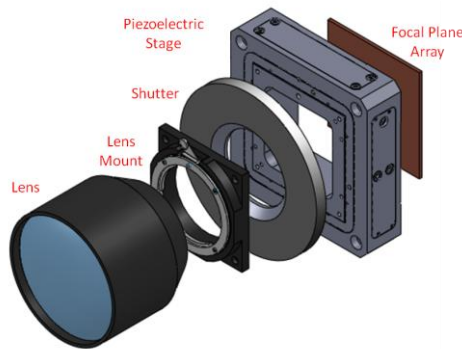
Figure 2. Baseline Optical Layout.
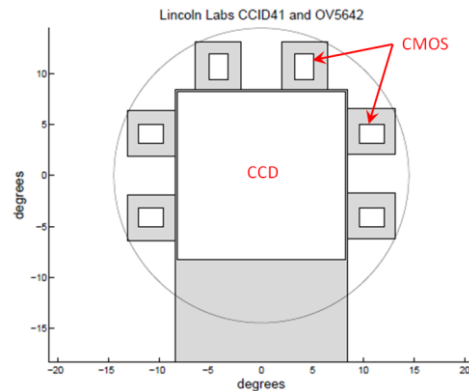
Figure 3. Baseline Focal Plane Array Layout.

A high-fidelity simulation of the satellite has been developed to determine the feasibility of achieving this 1.0-arcsecond pointing requirement.[3] The results presented in Reference 3 have been further refined through more accurate models of the actuators, sensors and flight software. The extensions of this simulation, hardware testing of key components, and initial hardware-in-the-loop testing will be presented in this paper.

SOFTWARE DESCRIPTION

The attitude determination and control subsystem (ADCS) operates in many different control modes, shown in Figure 4. This diagram may also serve as the satellite's general concept of operations. The satellite is inserted in a 600-km, low-inclination orbit[†] and begins in the "Damp Rates" mode. The angular rate of the satellite is slowed down with the torque coils until the reaction wheels can capture the satellite. The satellite then enters the "Sun Acquisition" mode where the satellite performs an attitude maneuver to acquire the sun and point the solar panels directly at

---

[*] Article is in preparation by M. W. Smith, et al.
[†] This is an ideal orbit for low radiation, low aerodynamic drag, and communication with equatorial ground stations. Since the satellite has no propulsion or ability to change its orbit, other orbits are being considered to take advantage of any launch opportunity that may arise.

the sun to recharge the batteries. Then the "Calibration" mode is entered, which initializes the Extended Kalman Filter and allows the estimates to converge to the true attitude. Calibration of any other instruments, such as the CCD, may be performed in this mode.

Once the necessary calibrations are complete, the satellite begins nominal operations. During orbit day, the satellite points the solar array toward the sun, charges its batteries, and desaturates the reaction wheels to their nominal speeds in the "Sun Pointing" mode. As the satellite enters orbit night, the satellite enters "Night Slew" where the satellite performs a smooth spinup-coast-spindown eigenaxis slew to the target star. Because the satellite relies on the gyros and magneto-meter during orbit day, attitude estimate may have accumulated an error on the order of a few degrees. Therefore when the slew is finished, the stars observed with the CCD and CMOS detectors must be matched against a limited star catalog to determine the current attitude in the "Star Survey" mode. A minor "Target Acquisition Slew" is then performed to place the target star on the same orbit-to-orbit position on the CCD and the optimal guide stars on the CMOS detectors. Additional time is allowed after the slew to damp any residual angular rates of the spacecraft during the "Target Acquisition Slew." Finally, the "Science Data Collection" mode is entered where the CCD collects science data and the stringent pointing requirement must be met.
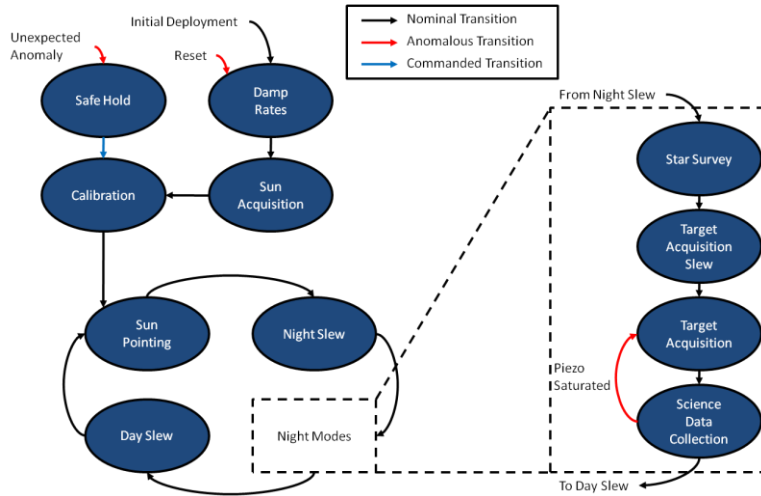
Figure 4. Control Modes for the Attitude Determination and Control Subsystem.

Precise pointing during the "Science Data Collection" mode is achieved through a two-stage control architecture, shown in Figure 5. The outer coarse attitude control loop commands the reaction wheels at rate of 4 Hz. The inner fine pointing control loop commands the piezoelectric stage at 12 Hz.

The 12 Hz output of the CMOS detectors must first be processed to be used for control commands. Image frames from the CMOS detectors are windowed to small regions of interest around the brightest guide stars. These windows are then processed to determine each guide star's centroid using a simple center of mass calculation.[4] Since these centroids contain some amount of effective delay,* this delay can be compensated through a small amount of processing as follows. The

---

* Delays can be caused by the finite integration time of the detector, processing time of the sensor electronics, and processing time of the avionics.

centroid measurements are in the moving frame of the piezoelectric actuator so, to simplify the calculations, the centroids are converted to the focal plane frame, which is fixed relative to the lens, by adding the position of the piezoelectric stage. Through back differencing, the centroid location in the focal plane frame can be propagated forward in time to compensate for fixed, deterministic time delays. It has been determined through simulation that the benefit of compensating these time delays outweighs the sensor noise amplification introduced by this process.

For the coarse attitude control loop, these centroids (in the fixed focal plane frame) from the CMOS detectors are used in a star tracking algorithm. These centroids are matched to known stars in a star catalog stored on board. With this information, the attitude quaternion is estimated with the QUaternion ESTimation (QUEST) algorithm.[5] In addition to the CMOS detectors, gyros are also used. The gyro outputs are filtered and downsampled to match the 12 Hz CMOS output and biases are subtracted out. These biases are estimated by an Extended Kalman Filter (EKF), which combines the attitude quaternion measurements (from the processed CMOS detectors' output) and the processed gyro measurements.[6] The attitude quaternion estimate from the EKF and the processed gyro measurements are fed into a nonlinear proportional-derivative reaction wheel control law.[2,7] An integrator with a small gain can be added to this controller to remove any long-term biases on the piezoelectric stage and help reduce the probability of saturation. To further ensure that the piezoelectric stage is not saturated, any remaining long-term biases in the stage can change the quaternion command given to the RW control block (not shown in Figure 5).

For the fine pointing control loop, the processed centroids (in the fixed focal plane frame) are converted back to the moving piezoelectric stage frame by subtracting the piezoelectric stage position. Since these stars are matched to known stars in a catalog, their desired positions can be calculated knowing the desired attitude of the satellite. The weighted average[*] of the difference between the measured and desired centroid positions gives an error signal. This error signal is fed into a proportional-integral control law to give a relative position command. To convert this into an absolute position command, the current position of the piezoelectric stage is added to the control law output. The output of the CCD also feeds into this control law to mitigate any changes in relative position of the CCD and CMOS sensors (e.g., thermal effects) to ensure the target star remains on the exact same position from orbit to orbit.
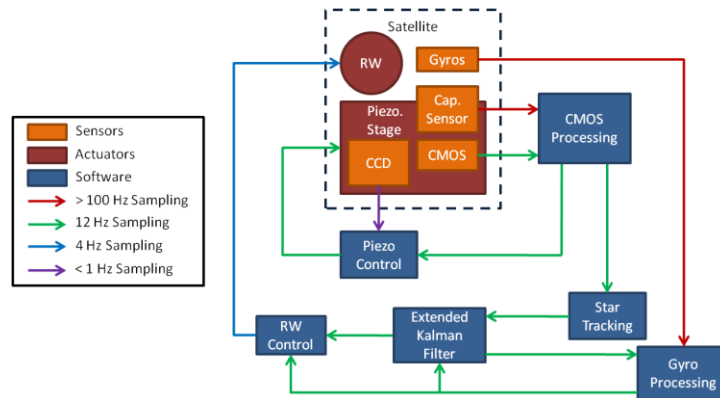
Figure 5. Science Data Collection Block Diagram.

---

[*] The weights can be determined by the brightness of the star since centroid error depends on brightness.

It is important to ensure that the selected gains for the two-stage control architecture have some stability margins. Figure 6 shows a simplified single-axis block diagram of this two-stage control architecture to analyze the gain and phase margins analytically. Some of the important simplifications are that the delay compensation and EKF are not represented and the attitude of the satellite is directly measured. It is also interesting to note that this figure shows that the control loops are only one-way coupled since they are two loops that are in series—the output of the coarse rigid-body attitude control loop is the "reference" input to the fine line-of-sight pointing control loop. This greatly simplifies the control law design since the gains can be selected independently for stability.
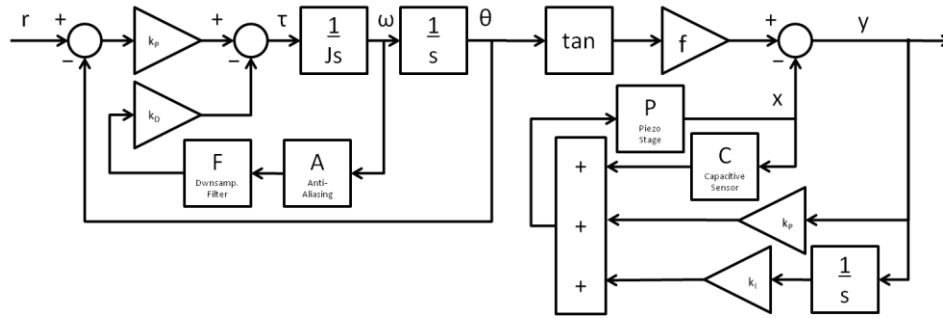
Figure 6. Simplified Single-Axis Two-Stage Control Block Diagram.

Gain and phase margins were calculated analytically using the simplified block diagram shown in Figure 6 and verified with the nonlinear simulation described in the following section. Verification was performed by adding a gain or time delay into the feedback loop until the system went unstable. The crossover frequencies were extracted with a Fourier analysis of the unstable outputs. Table 1 shows the resulting gain and phase margins. Phase margin is also expressed in terms of an equivalent time delay,[*] which gives a requirement on allowable processing delay to maintain stability. Close agreement of the linear analysis and nonlinear simulation gives confidence that these are the true margins. In addition, these margins meet typical gain and phase margins levied on ADCS.

Table 1. Coarse and Fine Loop Gain and Phase Margins.

		Gain Margin	Phase Margin	
Coarse Loop	Lin. Analysis	30 dB @ 0.23 Hz	75° @ 0.020 Hz	10.9 s
	Nonlin. Sim.	28 dB @ 0.23 Hz	73° @ 0.075 Hz	2.7 s
Fine Loop	Lin. Analysis	6.3 dB @ 6.0 Hz	61° @ 1.9 Hz	0.089 s
	Nonlin. Sim.	6.3 dB @ 6.0 Hz	44° @ 1.5 Hz	0.083 s
Typical Requirement		$\geq$ 6 dB	$\geq$ 30°	

SIMULATION DESCRIPTION & COMPONENT TESTING RESULTS

The ADCS software described in the preceding section has been incorporated into a MATLAB Simulink simulation of the satellite to determine the achievable pointing precision.

---

[*] This can be calculated by dividing the phase margin by the crossover frequency.

This simulation has high-fidelity models of the sensors, actuators, satellite dynamics, and space environment. The sensors are modeled with the appropriate bandwidth, quantization, discretization, saturation, and noise characteristics. The actuators contain also models of physical disturbances such as the reactive forces of the piezoelectric stage movement and vibrations produced by the reaction wheels. Satellite dynamics are modeled including gyroscopic effects and external disturbances from the space environment. Further details on the simulation models can be found in Reference 3.

While many hardware specifications can be found from manufacturer's datasheets, some of the components needed to be tested for important performance characteristics that are unavailable from the manufacturers. The two hardware tests that were deemed to be the most critical for this simulation are measuring the reaction wheel vibration characteristics and measuring the CMOS centroiding error.

Reaction Wheel Disturbance Measurements

It is important to quantify the magnitude of the reaction wheel vibrations to determine how much these vibrations will affect the pointing of the satellite. Vibrations that may normally be considered "small" can have large effects at the arcsecond level. Tonal disturbances (both forces and torques) tied to the reaction wheel speed are modeled as sinusoids,[8]

$$d_i(t) = c_i \omega_{RW}^2 \sin\left(h_i \theta_{RW}(t) + \varphi_i\right) \tag{1}$$

where c_i is the coefficient of the i^{th} harmonic, h_i is the harmonic number of the i^{th} harmonic, ω_{RW} is the speed of the reaction wheel, θ_{RW} is the phase of the reaction wheel, and φ_i is a uniformly distributed random variable over the domain $[0, 2\pi)$. Each reaction wheel produces radial disturbances (forces and torques orthogonal to the spin axis) that are offset by 90° in phase between the two axes as well as axial disturbances (forces and torques in the direction of the spin axis) for each of the harmonics. The harmonics and their corresponding coefficients are wheel-dependent and are therefore determined experimentally.

Tests of the MAI-200, the reaction wheel set used on ExoplanetSat, have been performed at the Massachusetts Institute of Technology Space Systems Laboratory (SSL) as well as the Kistler Load table at the Goddard Space Flight Center (GSFC). Figure 7 shows the MAI-200 mounted on the Kistler table and Figure 8 shows a sample of the data produced by this testing. The reaction wheel was run at a set of discrete speeds up to the maximum wheel speed. These time-domain data sets were converted to PSDs and plotted as a waterfall plot. The radial spokes coming out the origin are the tonal disturbances tied to the reaction wheel spin rate.
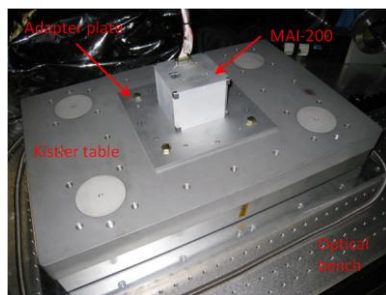

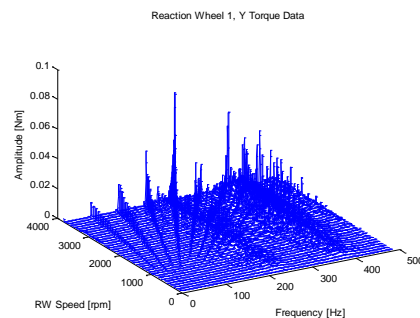**Figure 7. MAI-200 Test Setup at GSFC.**

Figure 8. Example Waterfall Plot.

The waterfall plots are further processed to determine the coefficients, c_i, and the corresponding harmonic number, h_i, for the force and torque disturbances, axial and radial directions, and all of the harmonics. Table 2 shows the results of this testing. These numbers are used along with Equation (1) to simulate the tonal reaction wheel disturbances.

Table 2. MAI-200 Reaction Wheel Disturbance Data.

Harmonic Number	Force [mg·mm]		Torque [mg·mm^2]	
	Axial	Radial	Axial	Radial
1	510	170	9,300	57,000
2	810	520	27,000	100,000
2.5	230	370	4,700	38,000
3	980	540	36,000	140,000
3.5	330	310	10,000	43,000
4	2,400	1,300	190,000	310,000
4.5	340	390	37,000	50,000
5	1,100	1,600	260,000	140,000
5.5	240	250	39,000	31,000
6	1,200	1,300	110,000	180,000

In addition to these tonal disturbances, there are also broadband disturbances caused by the reaction wheels, which change with wheel speed.[9] A model of these broadband disturbances was created and the effect of these broadband disturbances was found to be insignificant.

CMOS Detector Centroiding Error

Another important parameter in the simulation is the error in the measured centroid of the guide stars. This is important because the CMOS detectors are the most precise sensor and they feed into both the coarse and fine control loops. Any change in this parameter greatly affects the simulation results.

Figure 9 shows the experimental setup that was created to measure the centroiding error as a function of guide star magnitude and integration time. An LED behind a pinhole simulates a single guide star. The guide star brightness was adjusted by varying the current powering the LED and calibrated using an optical power meter placed in the optical path. By using neutral density filters, the brightness was brought down to known V magnitudes.[*] The light from this guide star was focused using a Zeiss 85 mm f/1.4 lens onto a custom B/W OmniVision OV5642 detector.[†] Centroids were calculated using a center of mass calculation.[4] Figure 10 shows an example of a picture taken of a 4[th] magnitude star with an 80 ms integration time. Note that the image is not saturated; the maximum is 107 out of 1023 counts, given the 10-bit resolution imager.

V magnitude was varied from 0 to 6 and integration times were varied from 10 to 80 ms. For each setting, the standard deviation of the centroids of 1000 images of the stationary guide star was calculated. Figure 11 shows the results. Note that for 4[th] magnitude and lower, some of the data points for low integration time were excluded because there were not enough counts on the detector to distinguish the star from the noise. Also, paired with each integration time is a maxi-

---

[*] The optical power meter used was not able to accurately measure the brightness of the star with the neutral density filters in place due to its resolution.
[†] Both the lens and CMOS imager used in the test are the models that will be used in the flight unit.

mum readout rate for the detector. Since the CMOS detectors are assumed to be operating at 12 Hz, an integration time of 83 ms or less must be used. The tested integration times comply with this requirement. The results agree with intuition: lower magnitude stars and lower integration times cause the centroid error to increase.
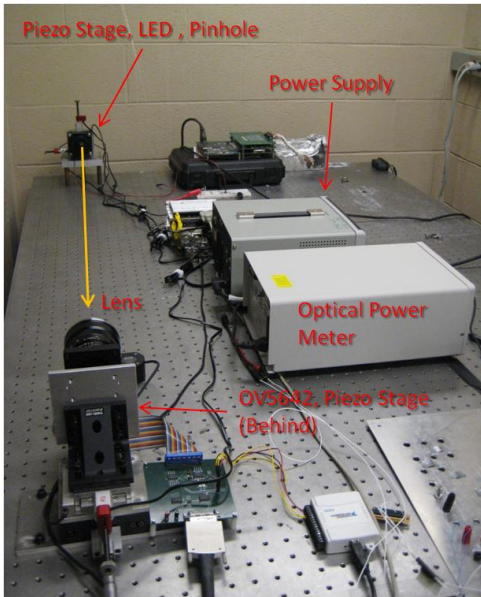

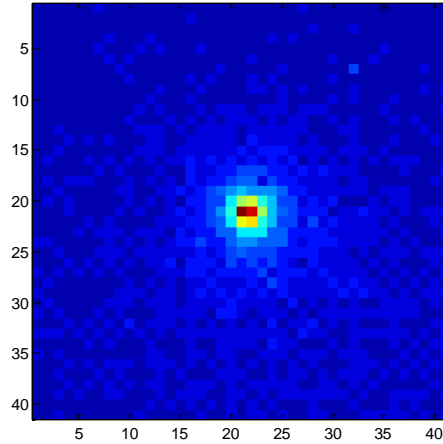**Figure 9. Experimental Setup of Guide Star Centroid Error.**

Figure 10. Example Image of Guide Star: 4th Magnitude, 80 ms Integration Time.

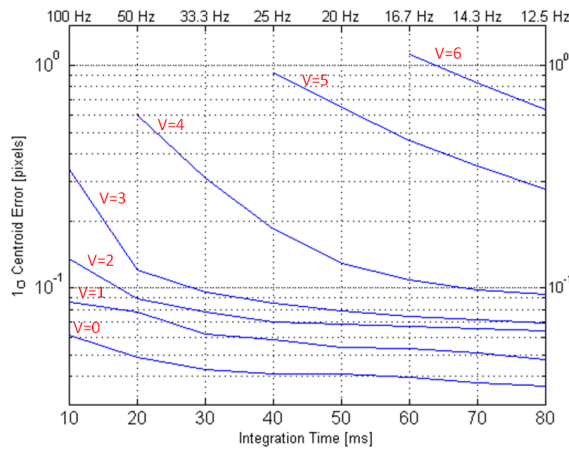

Figure 11. Guide Star Centroid Error as a Function of V Magnitude and Integration Time.

Further testing can be done to improve these results. Two additional parameters (defocus and analog gain in the imager) can be adjusted to determine if the centroid error can be improved. In addition, other centroiding algorithms such as a Gaussian fit, may improve the accuracy of the centroid measurements. For the simulation, a centroid error 0.1 pixels (1σ) was assumed. This is slightly optimistic since 6th magnitude stars are much more common than 4th magnitude stars and

brighter. A more detailed analysis of candidate target stars, their nearby guide stars, and the focal plane array layout will be performed to see exactly what guide stars, and therefore centroiding error can be expected.

SIMULATION RESULTS

Parameters from the previous section were used in a simulation of the closed-loop satellite dynamics to determine the achievable pointing precision. Figure 12 shows the results of the baseline simulation. The plot on the left shows the motion of the target star on the CCD over a period of 5 minutes. It can be seen that there is a large-amplitude, low-frequency component that causes the target star to wander around the CCD. This is due to the low-bandwidth, closed-loop response of the rigid-body attitude dynamics in the presence of disturbances and sensor noise. Superimposed on this is a high-frequency, low-amplitude jitter due to the reaction wheel disturbances. The plot on the right shows that when the piezoelectric stage is activated, most of the low frequency error is removed. The achieved line-of-sight pointing precision from this baseline simulation was 1.4 arcseconds (3σ), which does not meet the requirement of 1.0 arcseconds (3σ).
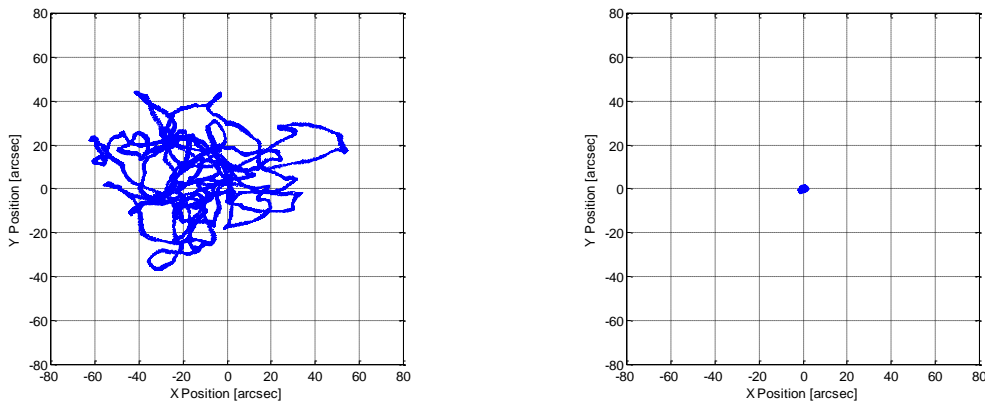

Figure 12. Target Star Motion on the CCD without the Piezoelectric Stage Activated (Left) and with the Piezoelectric Stage Activated (Right).

To see the various contributions to this 1.4 arcsecond pointing precision, it is useful to plot the power spectral density (PSD) and the cumulative mean-square value (MSV) of the pointing error, which is shown in Figure 13. The PSD shows the frequency content of the pointing error; large peaks correspond to tonal pointing errors. The cumulative MSV shows the actual contribution of each of these peaks to the pointing error. There are three peaks near 16 Hz that correspond to large jumps in the MSV at the same frequencies, and therefore contribute greatly to the pointing error. These tonal disturbances correspond to the vibrations of the reaction wheels operating near 1000 rpm. This visualization tool allows the various contributions to the pointing error to be quantified.

Analysis of the MSV plot in Figure 13 shows that 0.7 arcseconds of the 1.4 arcsecond error is due to the closed-loop response of the satellite. This is a result of sensor noise and process noise acting on the closed-loop dynamics of the system. This contribution can be seen visually as the slow rise of the MSV for frequencies less than 10 Hz. The other 0.7 arcseconds of pointing error is due to the reaction wheel vibrations. Of this 0.7 arcsecond pointing error, around 0.4 arcseconds is from the once-per-rev harmonic disturbance (~16 Hz) and 0.3 arcseconds is from the higher harmonics (~33 Hz, ~50 Hz, etc.). It was also determined that the reaction forces due to

the acceleration of the focal plane on the piezoelectric stage was insignificant—it contributed less than 5% of the pointing error.
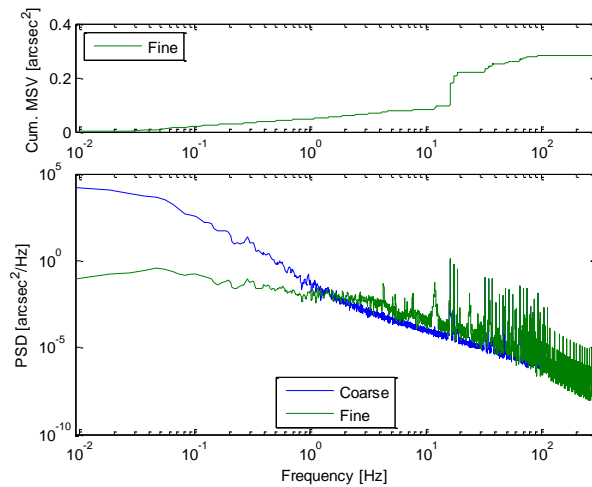
Figure 13. Pointing Error Power Spectral Density (PSD) and Cumulative Mean-Square Value (MSV).

Since reaction wheel vibrations contribute almost half of the pointing error, several options for reducing the effect of these vibrations were analyzed. The first option considered was adding an angular rate sensor to directly measure the high-frequency (> ~10 Hz) angular rate of the spacecraft and feedforward a command to the piezoelectric stage to compensate for this movement. The second option considered was to add a vibration isolator with a cutoff frequency of 5 Hz to filter out the vibrations. The angular rate sensor option achieved a pointing precision of 1.1 arcseconds while the vibration isolator achieved a pointing precision of 0.9 arcseconds. While these results are promising, the addition of this extra hardware, considering the tight volume, mass, and power constraints is hardly feasible.

There are many options, including the ones above, that may be explored to meet the 1.0-arcsecond pointing requirement. A prioritized list of these options and a brief discussion of each is provided below:

- Dynamically balance the reaction wheels. The wheels that were tested at GSFC were only statically balanced, which reduces the force disturbances in the radial direction. Dynamic balancing can also reduce the torque disturbances in the radial direction at the cost of a more complicated balancing procedure. Note that this will only affect the once-per-rev harmonic and not the higher harmonics, which are due to other sources such as bearing noise.

- Develop software algorithms using existing sensors and actuators to measure and cancel the reaction wheel vibrations. This is a very attractive option because it would not require additional hardware. Using gyro measurements, the reaction wheel disturbances can be estimated. This can then be fed forward to the piezoelectric stage to cancel the effect of these tonal vibrations. This method will be discussed further in the "Reaction Wheel Disturbance Estimation" section.

- Add isolation to the individual reaction wheels. While addition isolation was explored above, it still remains to be determined where this isolation can or should be placed. Isolating each wheel is the simplest place to put isolation on the satellite. The satellite is tightly packed into a small volume, making the addition of isolation mounts difficult to place. This introduces additional challenges such as accommodating the larger size of the reaction wheel unit and providing a sufficient thermal path from the motor block to the case.

- Add isolation to entire reaction wheel unit. This would be an attractive option since the reaction wheel unit comes as a single block. However, as seen in Figure 1, the current design of the solar panels provides a mechanical short around the isolator. The solar panel design would have to be heavily customized to accommodate this option.

- Incorporate additional actuators and sensors into the satellite. This option is less desirable than the passive isolation because it would require additional power for these devices. In addition to the angular rate sensor option that has been discussed above, other devices could be used such as a piezoelectric transducer between the reaction wheels and the rest of the satellite. This could be used as both a sensor and actuator: forces and torques generated by the reaction wheels can be measured and also generated with this transducer. This option faces similar integration challenges with the solar panels.

All of these options are being pursued to determine the optimal design choice for Exoplanet-Sat. One option in particular, the estimation of reaction wheel disturbances for feedforward control, will be discussed further in the "Reaction Wheel Disturbance Estimation" section. If the pointing requirement cannot be met with these options and given the project budgets, it may be necessary to relax the requirement. A reduction in pointing precision would increase the jitter noise, causing a corresponding increase in the detectable exoplanet size.

HARDWARE-IN-THE-LOOP TESTING

In addition to the simulation, a hardware-in-the-loop (HWIL) testbed is also being developed as a risk mitigation technique to demonstrate the pointing control architecture. A phased approach is being taken: the fine pointing control alone will be implemented and tested in hardware first, and then the two-stage control will be implemented and tested on a full three-degree-of-freedom air bearing.

The fine pointing control testbed is the same testbed shown in Figure 9. Both the guide star emulator and the CMOS detector are equipped with a piezoelectric stage. Satellite attitude movements can be simulated by moving the guide star with one piezoelectric stage. This motion can then be compensated with the piezoelectric stage with the CMOS detector. The HWIL testbed is controlled entirely by LabVIEW.

A screenshot of the HWIL user interface during testing is shown in Figure 14. The image on the left shows the image of the guide star. The guide star begins by moving around in the image according to the simulated residual attitude errors from the coarse control loop. This can be seen in the plot of the guide star centroid in the top right. About halfway through the run (~150 seconds), the piezoelectric stage is enabled. It can be seen that the centroid immediately snaps into place at the origin and the guide star motion is compensated by moving the piezoelectric stage around.
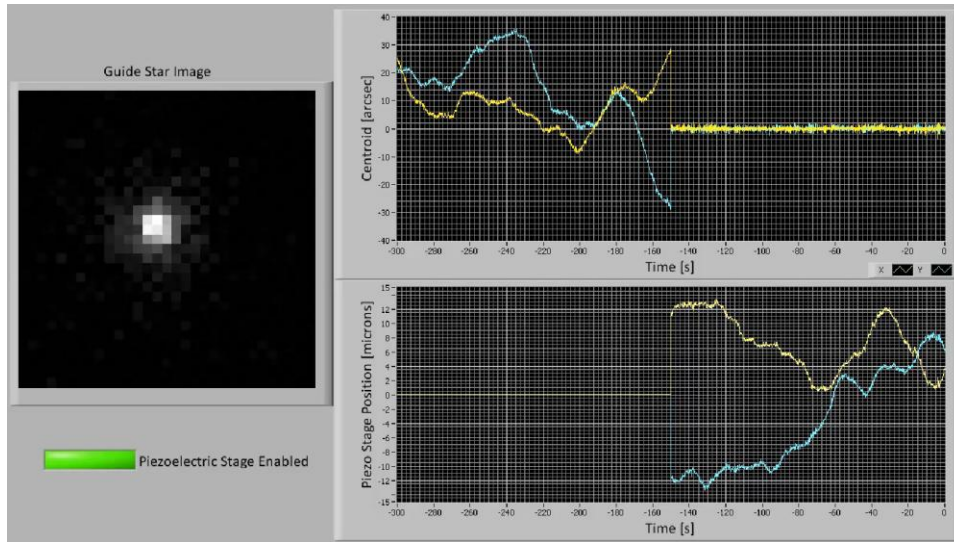
Figure 14. Screenshot of the HWIL Testbed During a Test Run.

This demonstrates the fine line-of-sight pointing control with the hardware running in the loop. While matching the performance of the simulation was not the primary goal of the testbed, it should be possible to achieve the same pointing precision. However, a pointing precision of 2.3 arcseconds (3σ) was achieved compared to the expected 0.7 arcseconds (3σ) from the simulation.[*] Several improvements to this HWIL are planned to improve this pointing precision. First, a more realistic star field will be used. The simulation assumes that there are 6 guide stars while the HWIL testbed only has a single guide star to use. Centroid error improves approximately by the square root of the number of stars in the field of view. A more realistic star field, with multiple guide stars that match the angular position and magnitude of a candidate target star (such as Alpha Centauri) will greatly improve the fidelity of the HWIL simulation. In addition, the use of LabVIEW MathScript nodes necessitated slowing down the sampling rate of the fine control loop. The desktop computer running LabVIEW will be replaced with embedded avionics to process the images and run the fine control loop at the specified 12 Hz.

REACTION WHEEL DISTURBANCE ESTIMATION

The baseline design achieves a pointing precision of 1.4 arcseconds (3σ) in simulation, which does not meet the 1.0-arcsecond requirement, as mentioned in the "Simulation Results" section. There are many options that can be explored to achieve this requirement, but the most attractive one at the moment is one that uses advanced software algorithms and existing hardware. The idea is to estimate the reaction wheel disturbances and feed this forward to command the piezoelectric stage and cancel the effect of the vibrations. This is very similar to the approach taken with the angular rate sensor, but instead using the gyros, which are already in the baseline design. Using the measurements from the gyros directly is not feasible because there is too much noise that would be fed into the control loop. By taking advantage of the fact that the reaction wheel disturbances are at a known frequency (a multiple of the reaction wheel speed, which can be measured

[*] Since the reaction wheel vibrations are not being simulated in this HWIL testbed, the simulation shows that a pointing precision of 0.7 arcseconds should be feasible.

by tachometers and/or encoders) and the fact that the amplitude varies with the wheel speed, the signal can be filtered from the noise.

A Kalman filter can be developed to estimate these disturbances. However, to achieve accurate estimation, process noise weights must be relatively low, which increases the convergence time of the estimates. If the reaction wheel speeds are changing, the Kalman filter may never converge as its settling time may slower than the rate of change of the reaction wheel speeds. Instead, an adaptive estimator can be employed. This adaptive estimator can be equipped with "memory," which allows the estimator to use adaptation information from previously visited wheel speeds. The reaction wheel vibrations can therefore be learned slowly over time, without having to converge faster than the wheel speed changes.

A preliminary overview of this technique will be presented, which develops an estimator for the linear, single-axis attitude dynamics of a satellite:

$$\dot{\theta}_{SC} = \omega_{SC}$$
$$\dot{\omega}_{SC} = \frac{1}{J}(u + d) \tag{2}$$

where θ_{SC} is the attitude of the satellite, ω_{SC} is the angular rate of the satellite, J is the satellite's inertia, u is the control input torque, and d is the disturbance torque. This technique will need to be extended to a three-degree-of-freedom to be used on ExoplanetSat.

The design of the adaptive estimator can be roughly split into three major tasks. The first task is to determine a set of basis functions that can accurately represent the disturbance. It is necessary that this approximation have the element of memory to store disturbance amplitude and phase information as a function of wheel speed. The second task is to develop an estimation law that provides stable estimation dynamics. The third task is to develop an adaptation law that adjusts the weights of the basis functions that ensures stability of the estimator and drives the estimation error to zero.

Disturbance Approximation

A general scalar function, $f(p)$, can be approximated as

$$f(p) = \mathbf{a}^T \mathbf{g}(p) + \varepsilon(p) \tag{3}$$

where $\mathbf{a}$ is a vector of weights, $\mathbf{g}(p)$ is a vector of basis functions, and $\varepsilon(p)$ represents the residual error of the approximation.

A common choice of a set of basis functions or a basis function "network" is the General Regression Neural Network (GRNN). For this particular choice of basis functions, a single basis function is given by

$$\mathbf{g}_i(p) = \frac{\exp\left(-\frac{(p - \mu_i)^2}{2\sigma^2}\right)}{\sum_{j=1}^{N} \exp\left(-\frac{(p - \mu_j)^2}{2\sigma^2}\right)} \tag{4}$$

14

where μ_i is the mean of the Gaussian function and σ is the variance. This is arranged into a network of basis functions:

$$\mathbf{g}(p) = [\mathbf{g}_1(p) \quad \mathbf{g}_2(p) \quad \cdots \quad \mathbf{g}_N(p)]^T \qquad (5)$$

The value of the means and the common variance for this network can be chosen *a priori* to approximate a given function to a given accuracy.[10] With an increasing number of basis functions, the approximation error, ε tends to zero.[11] Figure 15 shows an example of eleven basis functions in a GRNN and Figure 16 shows an example of a weighted output of this GRNN. This shows how an arbitrary function can be approximated by determining the appropriate weights, **a**, for each basis function.
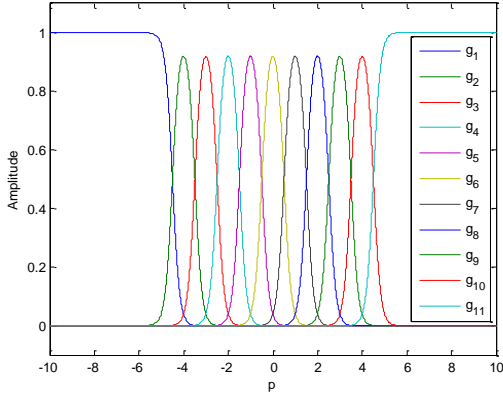

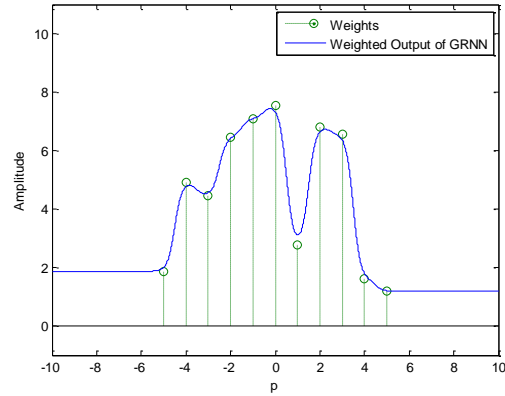**Figure 15. Basis Functions Forming a GRNN.**   **Figure 16. Weighted Output of a GRNN.**

These networks cannot approximate the reaction wheel vibrations because they are missing a sinusoidally varying time dependence. The GRNN can easily be modified to be

$$\mathbf{g}(\omega_{RW}, \theta_{RW}) = \begin{bmatrix} \mathbf{g}_1(\omega_{RW})\sin(\theta_{RW}) \\ \mathbf{g}_1(\omega_{RW})\cos(\theta_{RW}) \\ \vdots \\ \mathbf{g}_N(\omega_{RW})\sin(\theta_{RW}) \\ \mathbf{g}_N(\omega_{RW})\cos(\theta_{RW}) \end{bmatrix}^T . \qquad (6)$$

By adding in sine and cosine terms, the magnitude and phase of these sinusoids can effectively be estimated through adaptation of the weights of these basis functions.

Estimation Law

An estimation law can now be developed. Considering only a single-axis model of spacecraft attitude, the dynamics are linear and can therefore be written as

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}_u u + \mathbf{B}_d d. \qquad (7)$$

where **x** is the state of the system, **A** is the dynamics matrix, $\mathbf{B}_u$ is the control input matrix, and $\mathbf{B}_d$ is the disturbance-input matrix. The sensor output, *y*, and adaptation output, *z*, is given by

15

$$\mathbf{y} = \mathbf{C}_y \mathbf{x}$$
$$z = \mathbf{C}_z \mathbf{x}.$$

(8)

For the satellite model, $\mathbf{C}_y$ was chosen to be the identity matrix, since there are sensors to measure both the attitude and angular rate of the spacecraft. The $\mathbf{C}_z$ matrix, however, was chosen to select only the spacecraft angular rate. This is because this scalar z variable will be used in the adaptation law and only the gyros will have a high enough sampling frequency to capture the reaction wheel vibrations. The CMOS detectors that measure the spacecraft attitude have a sampling frequency that is too low to capture the reaction wheel vibrations.

Given this linear system, and assuming that the disturbance input is known, the estimator can be designed with the following dynamics

$$\dot{\hat{\mathbf{x}}} = \left(\mathbf{A} - \mathbf{L}\mathbf{C}_y\right)\hat{\mathbf{x}} + \mathbf{B}_u u + \mathbf{B}_d \hat{\mathbf{a}}^T \mathbf{g}\left(\omega_{RW}, \theta_{RW}\right) + \mathbf{L}\mathbf{y}$$
$$\hat{z} = \mathbf{C}_z \hat{\mathbf{x}}$$

(9)

The estimator gain matrix, $\mathbf{L}$, can be selected with standard linear techniques such as pole placement or linear quadratic estimator techniques. This would ensure that the estimator dynamics, given by the matrix $\mathbf{A}$-$\mathbf{LC}$, is stable. If the disturbance was known exactly, the estimator would be stable since the disturbance is known to be a bounded input. However, the disturbance is not exactly known. The disturbance is $d \approx \mathbf{a}^T \mathbf{g}(\omega_{RW}, \theta_{RW})$. The true values of $\mathbf{a}$ are not known exactly and are therefore an estimate, $\hat{\mathbf{a}}$, needs to be adapted over time.

Adaptation Law

An adaptation law, which updates the estimates of the weights of the GRNN, $\hat{\mathbf{a}}$, must be developed to ensure stability of the estimator. To do this, first an error signal is defined as the difference between the adaptation output of the system and the adaptation output of the estimator:

$$e = \hat{z} - z.$$

(10)

It can be shown that this error signal can be written as[12]

$$e = H(s)\boldsymbol{\delta}^T \mathbf{g}\left(\omega_{RW}, \theta_{RW}\right)$$

(11)

where $H(s)$ is the error transfer function defined as

$$H(s) = \mathbf{C}_z \left(s\mathbf{I} - \mathbf{A} + \mathbf{L}\mathbf{C}_y\right)\mathbf{B}_d$$

(12)

and $\boldsymbol{\delta} = \hat{\mathbf{a}} - \mathbf{a}$. Equation (10) defines a well known error signal, which has associated provably stable adaptation laws.[12,13] If the error transfer function is strictly positive real (SPR), the following adaptation law can be used

$$\dot{\hat{\mathbf{a}}} = -\eta e \mathbf{g}\left(\omega_{RW}, \theta_{RW}\right)$$

(13)

where $\eta > 0$ is the learning rate. If the error transfer function is not SPR, the following adaptation can be used

$$\dot{\hat{\mathbf{a}}} = -\eta \varepsilon H(s)\mathbf{g}\left(\omega_{RW}, \theta_{RW}\right)$$

(14)

where

$$\varepsilon = e + \hat{\mathbf{a}}^{\mathbf{T}} H(s) \mathbf{g}(\omega_{RW}, \theta_{RW}) - H(s) \hat{\mathbf{a}}^{T} \mathbf{g}(\omega_{RW}, \theta_{RW}). \qquad (15)$$

This adaptation law results in a slower adaptation since the signals are filtered through $H(s)$. The second adaptation law was chosen since the error transfer function was not SPR. The stability proof is a Lyapunov-based proof that shows that a transformed error signal goes to zero due to a positive definite Lyapunov-like function and an associated negative semidefinite time derivative.[12]

Results

A simulation of the spacecraft with the adaptive estimator was developed. Realistic sensor noise (e.g., angular random walk of the gyros) was included. The simulation was run with the reaction wheel speed being swept linearly from 900 to 1100 back to 900 rpm with a period of 200 seconds. With a total simulation time of 1000 seconds, the reaction wheel speed range is swept through 10 times. In addition, a lightly damped resonance was added at 16 Hz (~1000 rpm) to make the problem a little more difficult.

Figure 17 shows the estimate of the disturbance over the length of the simulation. During the first sweep, the disturbance amplitude estimate grows but does not quite reaction the true amplitude. During the second wheel speed sweep (100 to 200 seconds), however, the disturbance estimate grows again to better match the actual disturbance. This is because it is able to use the previously adapted reaction wheel disturbance estimates stored in the GRNN weights. With each of the successive sweeps, the disturbance estimate is refined and by the end of the simulation, any further adaptation is due to sensor noise—steady state has been reached.
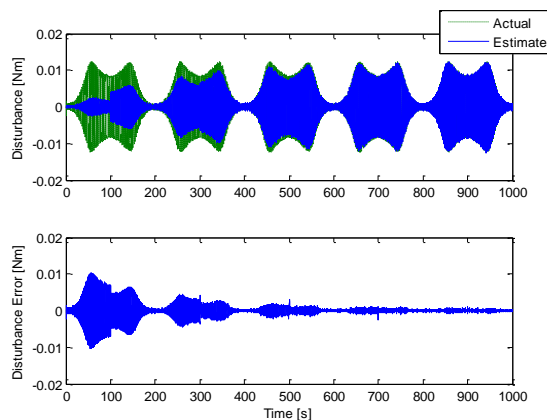
Figure 17. Estimated Disturbance and Disturbance Error Versus Time.

As an alternative way to view this disturbance estimation, several plots were generated, shown in Figure 18, which shows the estimated disturbance amplitudes as a function of wheel speed for various points in the simulation. Note that this disturbance amplitude estimate corresponds to the output of an unmodified GRNN where each weight is determined by the root-sum-square (RSS) of the two corresponding sine and cosine weights of the modifed GRNN. The first plot shows the estimated disturbance amplitude at 25 seconds into the simulation. Since only wheel speeds in the range 900 to 950 rpm have been visited, only estimated amplitudes in that range have been altered due to the spatial locality of GRNN basis functions. Also, as shown previously, the esti-

mated amplitudes are not quite as large as the actual amplitudes. The second and third plots show the estimated amplitudes as the reaction wheel speed increases up to 1000 rpm. The fourth plot shows the estimated amplitudes after the first complete reaction wheel speed sweep. It can be seen that all amplitudes within the range 900 to 1100 rpm have been updated now. The subsequent plots show the incremental convergence of the amplitude estimates to the true values.
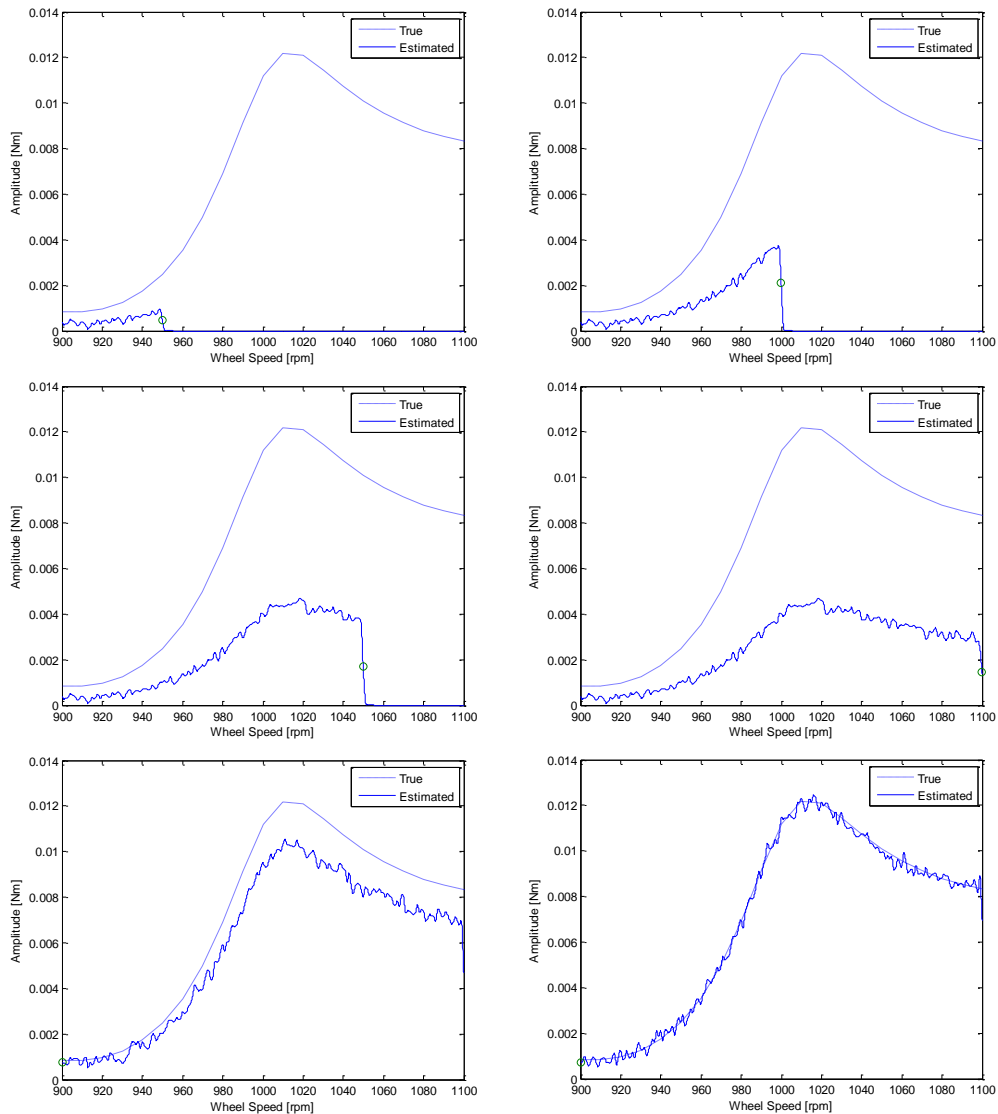
Figure 18. Estimated Disturbance Amplitude as a Function of Reaction Wheel Speed for Various Simulation Times (25 s, 50 s, 75 s, 100 s, 400 s, 1000 s).

This adaptive estimator was compared against a tuned disturbance-estimating Kalman filter.[*] The steady-state RMS estimation error for the adaptive estimator was 2.0×10^{-4} Nm compared to

---

[*] See, e.g., Section 7.1 of Reference 14 for a standard continuous Kalman filter implementation. The state includes the attitude, angular rate, disturbance magnitude for the sine component, and disturbance magnitude for the cosine component.

4.5×10^{-4} Nm for the Kalman filter. This demonstrates a 55% improvement of the adaptive estimator against a standard Kalman filter.

While these results are promising, the improved performance comes at the cost of additional processing and memory requirements. This simulation used on the order of one hundred basis functions in the network. For additional reaction wheels and harmonics, this number can grow quite large.[*] However, since the basis functions are very local only a few weights need to be adapted at any point in time. In addition, the memory requirements are fairly modest. Assuming that 1000 basis functions are used to cover a reaction speed range of interest for three harmonics of three reaction wheels with single-precision weights, 108 kilobytes of memory is required.

The results presented are preliminary and several extensions need to be made to implement this in ExoplanetSat. First, this technique needs to be extended to a three-axis satellite model with nonlinear dynamics and multiple inputs and outputs. Second, this estimator implementation is continuous, while measurements are discrete at different rates. The estimator needs to be modified to handle discrete multi-rate measurements, especially since the attitude measurements (CMOS detectors) cannot measure the reaction wheel vibrations while the gyroscopes have the bandwidth to make these measurements. It is possible that the previously developed EKF can be used for this purpose—the only complication is including this in the transfer function given by Equation (12). Third, the estimator can be extended to handle additional harmonics beyond the once-per-rev harmonic. One interesting issue that occurs with three reaction wheels is what happens when reaction wheel speeds cross. Since the disturbances produced by both wheels will be at the same frequency, it is unclear how the adaptation law will handle this case. With these extensions, it will be possible to say how much of an improvement this new technique will provide and whether or not the 1.0-arcsecond pointing requirement can be met.

CONCLUSION

Achieving a pointing precision of 1.0 arcseconds (3σ) is a very challenging requirement, especially when adhering to the CubeSat standard. To determine the feasibility of achieving this requirement, detailed testing, analyses, and simulations have been performed. Key hardware specifications, reaction wheel vibration characteristics and CMOS detector centroiding errors, were measured. These specifications were fed into a high-fidelity simulation of the satellite to determine the performance. The baseline simulation achieved a pointing precision of 1.4 arcseconds (3σ).

Several options were analyzed, which suggest that meeting the requirement is feasible, however with additional hardware. In an effort to see if meeting the requirement without additional hardware is possible, an adaptive estimator with memory that estimates the reaction wheel vibrations from relatively noisy gyro measurements is developed. It is shown to estimate reaction wheel vibrations 55% better than a standard Kalman filter. Further work is necessary to incorporate this into the full simulation of the satellite to determine how much the pointing precision will be improved.

In addition to the simulation work, a HWIL testbed has been built to demonstrate the fine pointing control loop. A pointing precision of 2.7 arcseconds (3σ) has been demonstrated and further improvements to the setup have been proposed to have the testbed match expected point-

---

[*] The number is calculated as (number of reaction wheels)*(3 axes per wheel)*(number of harmonics per wheel)*(number of basis functions per harmonic).

ing precisions achieved by simulation. This fine pointing control loop is a major milestone in a phased approach to demonstrating the two-stage control architecture in hardware.

ACKNOWLEDGMENTS

REFERENCES

[1] Piterman, A. and Ninkov, Z., "Subpixel sensitivity maps for a back-illuminated charge-coupled device and the effects of nonuniform response on measurement accuracy," *Optical Engineering*, Vol. 41, No. 6, June 2002, pp. 1192–1202.

[2] Smith, M. W., Seager, S., Pong, C. M., Villaseñor, J. S., Ricker, G. R., Miller, D. W., Knapp, M. E., Farmer, G. T., and Jensen-Clem, R., "ExoplanetSat: Detecting and monitoring exoplanets using a low-cost, CubeSat platform," *Proceedings of SPIE*, Vol. 7731, No. 773127, Aug. 2010.

[3] Pong, C. M., Lim, S., Smith, M. W., Miller, D. W., Villaseñor, J. S., and Seager, S., "Achieving high-precision pointing on ExoplanetSat: Initial feasibility analysis," *Proceedings of SPIE*, Vol. 7731, No. 77311V, Aug. 2010.

[4] Liebe, C. C., "Accuracy Performance of Star Trackers—A Tutorial," IEEE Transactions on Aerospace and Electronic Systems, Vol. 38, No. 2, Apr. 2002, pp. 587–599.

[5] Shuster, M. D. and Oh, S. D., "Three-Axis Attitude Determination from Vector Observations," *Journal of Guidance and Control*, Vol. 4, No. 1, Jan.–Feb. 1981, pp. 70–77.

[6] Lefferts, E. J., Markley, F. L., and Shuster, M. D., "Kalman Filtering for Spacecraft Attitude Estimation," *Journal of Guidance, Control, and Dynamics*, Vol. 5, No. 5, Sept.–Oct. 1982, pp. 417–429.

[7] Wie, B., Weiss, H., and Arapostathis, A., "A Quaternion Feedback Regulator for Spacecraft Eigenaxis Rotations," *Journal of Guidance, Control, and Dynamics*, Vol. 12, No. 3, May–June 1989, pp. 375–380.

[8] Masterson, R. A., Miller, D. W., and Grogan, R. L., "Development and Validation of Reaction Wheel Disturbance Models: Empirical Model," *Journal of Sound and Vibration*, Vol. 249, No. 3, Jan. 2002, pp. 575–598.

[9] Liu, K.-C., Maghami, P., Blaurock, C., "Reaction Wheel Disturbance Modeling, Jitter Analysis, and Validation Tests for Solar Dynamics Observatory," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Honolulu, HI, Aug. 2008.

[10] Sanner, R. M. and Slotine, J.-J. E., "Gaussian Networks for Direct Adaptive Control," IEEE Transactions on Neural Networks, Vol. 3, No. 6, Nov. 1992, pp. 837–863.

[11] Beuschel, M., "Nonlinear Function Approximators," *Intelligent Observer and Control Design*, Springer-Verlag, 2000, pp. 83–103.

[12] Lenz, U., "Systematic Intelligent Observer Design," *Intelligent Observer and Control Design*, Springer-Verlag, 2000, pp. 105–134.

[13] Narendra, K. S. and Annaswamy, A. M., *Stable Adaptive Systems*, Prentice Hall Information and System Sciences Series, Prentice Hall, 1989.

[14] Brown, R. G. and Hwang, P. Y. C., *Introduction to Random Signals and Applied Kalman Filtering*, 3rd ed., John Wiley & Sons, 1997.