THE NANCY GRACE ROMAN SPACE TELESCOPE'S ATTITUDE CONTROL SYSTEM

Eric Stoneking[*], Oscar C. Hsu[+], Gary Welter[†], Lia Lewis[+], Blair Carter[+], Bryce Doerr[+], Jean-Etienne Dongmo[+], Matt Heron[+], Jacob Shapiro[+]

The Nancy Grace Roman Space Telescope ("Roman") mission has completed its design phase and is currently in integration and testing, with a launch readiness date of October 2026. The mission is designed to investigate essential questions in the areas of dark energy, exoplanets, and infrared astrophysics. Roman uses a 2.4-meter primary telescope and two instruments: the Wide Field Instrument (WFI) and the Coronagraph Instrument (CGI). In order to address the critical science requirements, the Roman mission will conduct large-scale surveys of the infrared sky, requiring both agility and precision pointing (8 milli-arcsec stability, 14 milli-arcsec jitter). This paper describes some of the challenges this mission profile presents to the attitude control system (ACS), and some of the design elements chosen to accommodate those challenges. The Fine Guidance System (FGS), which is integral to the WFI instrument, is capable of tracking up to 18 guide stars, enabling robust FGS acquisition and precision pointing. Mission requirements drive the FGS acquisition process to be rapid, with a requirement of ten seconds on average. Control and estimator designs that support this capability are described. The wheel nullspace is used to minimize zero-crossing disturbances and jitter excitation. Aside from pointing precision, one of the most challenging ACS requirements was thruster maneuver accuracy. The design for meeting this requirement is discussed. Also, algorithms for pointing the high-gain antenna (HGA) are described. Finally, the ACS hardware suite is presented.

INTRODUCTION

The Nancy Grace Roman Space Telescope ("Roman"), named for NASA's first Chief of Astronomy, is currently in integration and testing, progressing toward a launch readiness date of October 2026. Roman is a Hubble-class infrared observatory, with mission objectives to investigate dark energy, dark matter, exoplanets, and general infrared astrophysics.[1] Its Wide Field Instrument (WFI) has similar angular resolution as the Hubble Space Telescope, but a field of view 200 times larger. This supports a key requirement of Roman, to conduct large-area surveys to map weak gravitational lensing, as well as detect and observe supernovae and exoplanet microlensing events. Typical survey operations feature science exposures of one to three minutes in duration, linked by slews from 0.25 to 3 degrees. Observational efficiency is heavily impacted by slew performance,

[*] Senior Aerospace Engineer, Attitude Control Systems Engineering Branch, NASA Goddard Space Flight Center, Code 591, Greenbelt MD 20771.
[+] Aerospace Engineer, Attitude Control Systems Engineering Branch, NASA Goddard Space Flight Center, Code 591, Greenbelt MD 20771.
[†] FGS Algorithm Lead, Flight Software Branch, NASA Goddard Space Flight Center, Code 582, Greenbelt MD 20771.

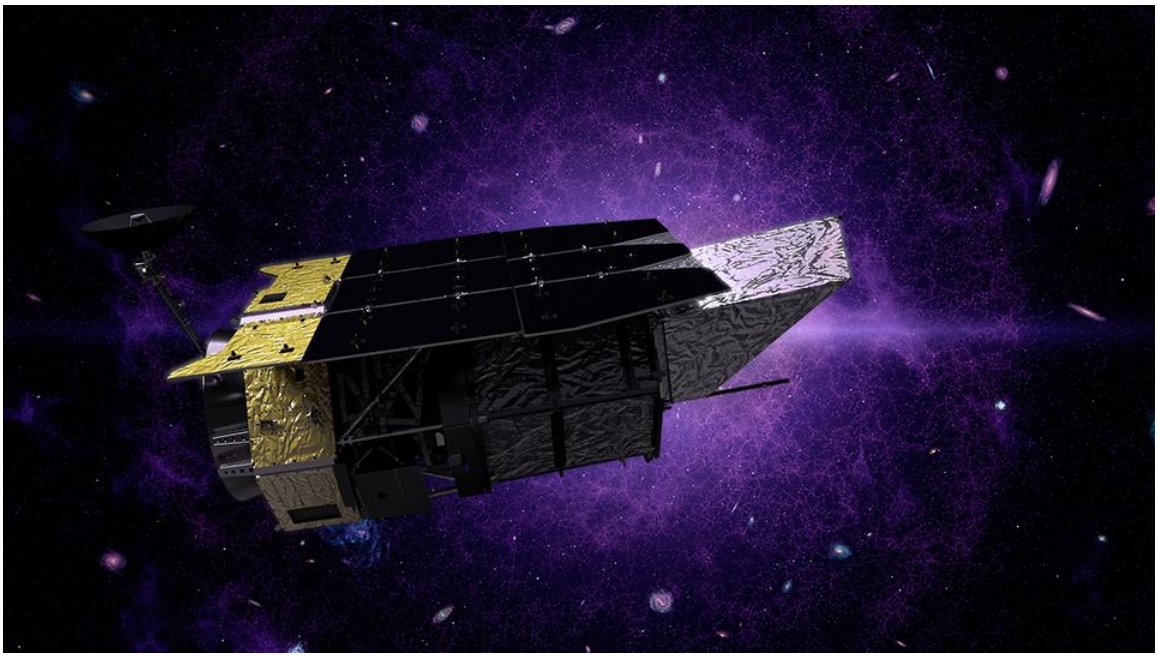

Figure 1. The Nancy Grace Roman Space Telescope.

so the attitude control system has been designed to slew, acquire guide stars in the Fine Guidance System (FGS), and settle to science-quality pointing stability, all as rapidly as possible.

Roman will be placed in a quasi-halo orbit about the Sun-Earth L2 Lagrange point. Solar radiation pressure is the dominant environmental disturbance torque. The angular momentum it imparts must be absorbed by the reaction wheel system and periodically unloaded using thrusters. To maximize the interval between unload events, the wheel system is unloaded to a "negative" extreme and wheel speeds are allowed to pass through zero on their way to the "positive" extreme. Control laws acting in the nullspace of the wheel system act to reduce attitude disturbances when wheels cross through zero speed, as well as to reduce excitation of structural vibration modes due to wheel imbalance torques.

The second instrument, the Coronagraph Instrument (CGI) is a technology demonstration of exoplanet coronagraph observations. Its internal pointing stability requirements are even more stringent than the WFI stability requirements, requiring additional stabilization internal to the CGI instrument. The pointing stability requirements levied onto ACS are the same as for WFI operations. The CGI is, however, especially sensitive to jitter disturbances generated by the spinning reaction wheels, which influences momentum management during CGI operations.

HARDWARE COMPONENTS

The Roman Space Telescope attitude control system hardware suite contains a mixture of commercial hardware and custom-built hardware for the Roman Space Telescope. The Roman mission is designed as a cross-strapped observatory where a single failure will not be mission-ending. The sensor suite consists of twenty-four coarse sun sensors (CSS) from Adcole Space, one Space Scalable Inertial Reference Unit (SSIRU) from Northrop Grumman, three Astro APS Star Sensors from Jena Optronik, one two-axis high gain antenna system using custom-built control electronics with Honeybee actuators/encoders, the Fine Guidance System that was custom-built for Roman, and the Coronagraph Instrument built by NASA Jet Propulsion Lab. The actuator suite

consists of six (modified) HR18 Reaction Wheels from Honeywell, sixteen 5-N class thrusters from Moog, eight 22-N class thrusters from Aerojet Rocketdyne, and one two-axis high gain antenna system.

Sensors

The twenty-four coarse sun sensors are divided in two sets such that there are two CSSs aligned at each location. Roman ACS uses a single set of CSSs with the second set being a cold spare. Each single set of CSS provides 4π steradian coverage. The single SSIRU contains four individual gyros in a pyramid configuration. The ACS uses all four individual gyros within the SSIRU at any given time and can operate with any single-gyro channel failure. The three Astro APS Star Sensors are used all the time by the ACS. The Astro APS Star Sensors are mounted so the fields of view of the star sensors do not overlap. The ACS systems can meet mission requirements using any two of three star sensors. The high gain antenna system consists of 1.7-m dish and a two-axis gimbal system. There are two sets of optical encoders, motor winding, and control electronics. The fine guidance system tracks up to 18 guide stars. The last sensor used by the ACS is the Coronagraph Instrument (CGI) which provides pointing information from their internal cameras when CGI is the prime instrument.

Actuators

The six reaction wheels are based on the Honeywell HR18-250 Constellation series reaction wheels. For Roman, custom motor windings were procured to provide a required higher torque (0.45 Nm), resulting in a lowered minimum momentum storage capacity of 100 Nms. All six wheels are actively used by ACS to rapidly perform slew-settle operations and fine pointing. The sixteen 5-N thrusters and eight 22-N thrusters are broken into two groups of 5-N and 22-N thrusters. The groups are redundant so that any bank of thrusters could become non-operational and still achieve mission success. The last actuator controlled by ACS is the two-axis high gain antenna system. The actuators are designed for steps sizes of 0.00375 degrees.

FINE GUIDANCE SYSTEM

The WFI uses 18 Teledyne HAWAII-4RG 4K×4K detectors[2] (Figure 2) to collect science data, each pixel having a linear pitch of ~110 milli-arcsec on the sky. Each detector has the capability to read data from a small, selected guide window separately from, the science data stream. This capability is used to implement a Fine Guidance System (FGS) integral to the WFI. Up to 18 guide stars are tracked, centroided, and combined to obtain attitude error measurements, which are then forwarded to the Attitude Control System (ACS) for its use.

The WFI has two operating modes for science support: imaging and spectroscopy. The FGS has parallel modes for fine guidance support. Both modes have guide star acquisition and guide star tracking phases. The acquisition phase for each mode has the FGS and ACS working together to transition from Star Tracker (ST) / Inertial Reference Unit (IRU) / Reaction Wheel (RW) based attitude knowledge and control precision (~2 arcsec (3σ) for attitude) to FGS / RW based knowledge and control. During the acquisition phase the initial guide window (GW) sizes are large enough to accommodate ST/IRU level attitude uncertainty and transition to the smaller, track-size GWs when control is sufficiently tight. Initial placement of the GWs is based on the expected location of the guide stars selected by the Guide Star Selection System (GSSS) for the desired observatory pointing.

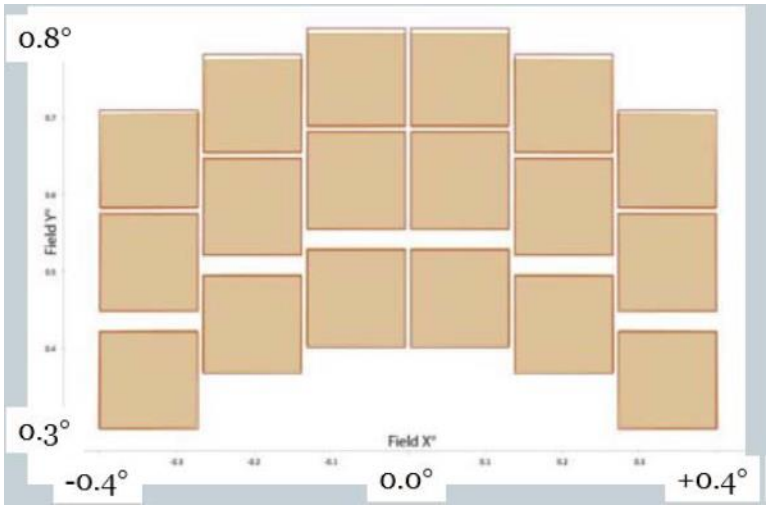

Figure 2. The Roman Field of View is over 200x that of Hubble's Wide Field Camera; Image shows relative size and position of 18 WFI detectors.[3]

FGS Image Mode Track Phase Support

During imaging science, the FGS flight algorithm does centroiding on the images of the guide stars. While the detector readout frame time is approximately five seconds, the nominal 16×16-pixel image mode guide window exposure time is ~0.191 s. Each WFI science image uses one of eight filters ([R: 0.48-0.76], [Z: 0.76-0.98], [Y: 0.93-1.20], [J: 1.13-1.46], [H: 1.38-1.78], [F: 1.68-2.00], [K: 1.86-2.41], and [W: 0.97-2.00] μm). FGS guide star images are acquired with the filter used for science. The associated point spread function (PSF) will have a half-width of ~1.2 λ/D (measured in radians), where λ is the bandwidth mid-region wavelength, and D is the diameter of the primary mirror. The half-widths range between ~0.58 pixel (R) to 2.00 pixels (K). Figure 3 shows sample simulated PSFs for the R and K filters for detector 1, including secondary mirror strut masking and "typical" optical aberration.[*] Regarding the latter, ray propagating wave phase shifts at the detector focal plane, the cause of optical aberration, enter the PSF equations as $2\pi\delta(-)/\lambda$, where $\delta(-)$ is the effective optical path difference projected back to the primary mirror as a function of position on the primary mirror and incoming direction of the ray wavefront. For a given $\delta(-)$ at the primary mirror, $2\pi\delta(-)/\lambda$ is largest, and hence the effect of aberration is greatest, for small λ - if $2\pi\delta(-)/\lambda << 1$. This is true for Roman when its optics are well focused. PSF details change as a function of WFI field-of-view position, and one needs to apply corrections for the effects of both optical aberration and finite pixels size. The size of these corrections is greater for the short wave-length filters than for the long wavelength filters.

Track phase GW data accumulation for each detector for each FGS cycle consists of five steps. First, a GW reset, reading and discarding one frame of GW data to let reset-associated noise fade. Second, reading N (nominally 2) frames of GW "pedestal" data to measure the GW zero level. Third, exposing for ~0.17 s (during which a sub-block of science data is being read). Fourth, reading and discarding N_{TFPN} GW frames (nominally 3) to let GW transient fixed pattern noise (TFPN)[†] noise fade. And last, reading N frames (again, nominally 2) of GW "signal" data.

---

[*] Our thanks to our colleague Guangun Gao of the Roman optics team for providing the aberration model.

[†] GW TFPN is triggered in the WFI detectors by the transition from reading science data to reading GW data. It was discovered and characterized by our Roman detector characterization team colleagues, in particular Dan Kelly and Augustyn Waczynski.

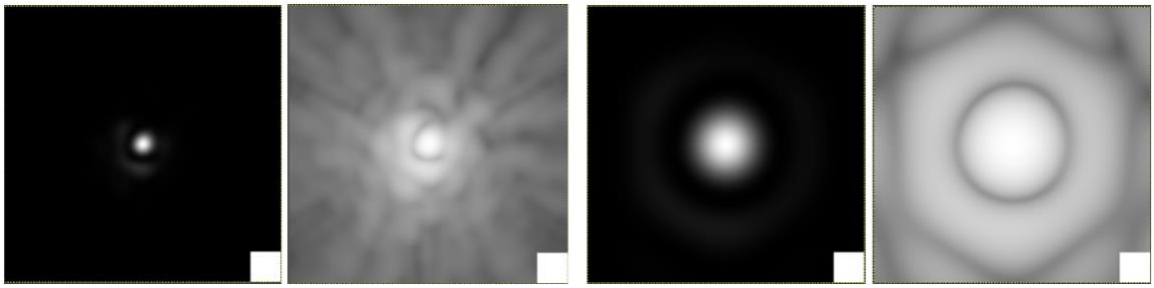

Figure 3. Simulated linear and logarithmic PSF plots for the WFI R (left) and K (right) filters. Plot box size for these images is 9×9 WFI pixels; the white square on the lower right of each image is the size of a single WFI pixel. Image resolution is at a scale of 128 fine pixels per WFI pixel.

The WFI detector hardware averages the N frames of pedestal data, and separately the N frames of signal data, forming "resultant" frames of each. The resultant frames are sent to the flight software (FSW), which subtracts the pedestal from the signal to form the GW correlated double sample (CDS) image. Total CDS exposure time extends from the middle of the pedestal frames to the middle of the signal frames, equal to ~0.19 s.

After forming a GW CDS image, the FSW scans it to find the up to N_B (nominally 5) 3×3-pixel squares with integrated brightness within a tolerance range specified by the GSSS for the selected guide star for the detector of immediate interest. The expectation is that usually there will be only one obvious star in each guide window satisfying the brightness restriction, but false positives can occur – e.g., due to unknown hot bad pixels or cosmic ray hits. For each bright 3×3 square, the FSW computes a 2D centroid using an algorithm based on the discrete version of the Fourier transform (DFT) shift theorem.[*] Each such centroid estimate is then corrected for effects of finite pixel size and optical aberration. These corrections are most important when using short wavelength filters (e.g., R vs. K).

With a set of centroids computed for each of the detectors selected for FGS use for a given observation, the set of sets (i.e., all centroids for each selected detector) are processed with a pattern match algorithm to find the proper guide star for each detector. The proper guide stars should be in nearly the same relative position in their respective GWs, despite motion of the guide stars during attitude settling and resultant motion of the GWs commanded by the FGS to track[†] the guide stars. After the proper set of centroids have been isolated, those centroids and their corresponding GSSS-provided reference coordinates are used in a weighted-for-star-brightness least-squares fit algorithm to determine the WFI frame attitude shift relative to the desired attitude. Optionally, the FGS FSW can be commanded to replace the GSSS-provided reference coordinates with FGS-generated centroids, possibly adjusted for a mean pointing shift relative to the GSSS coordinates, at the start of science observations. The solved-for attitude error is used (a) by the FGS to move the GWs to track the guide stars, and (b) by the ACS to control the observatory attitude. Based on high-fidelity simulations, we estimate that the precision of the horizontal and vertical pointing components of the image mode attitude solution in the WFI frame ranges between 0.3 and 0.7 milli-arcsecond (1σ), depending on which filter is being used. The precision of the effect of the roll component at a lever arm of ~10,000 pixels (~0.31°) from the WFI boresight is comparable.

---

[*] This is essentially the same algorithm as that used for JWST centroiding, the only difference centroiding box size.

[†] The GWs are commanded to move in lock-step, with the same GW horizontal and vertical shifts commanded to all detectors.

FGS Spectroscopy Mode Track Phase Support

The WFI has two dispersing elements to support spectroscopy: a grism and a prism. They have similar band passes (0.96 - 1.98 mm for the grism, and 0.76 - 1.81 mm for the prism), but significantly different dispersions (nearly constant at ~920 pixels/mm for the grism, varying between ~510 and 110 pixels/mm from short-to-long wavelength for the prism). For convenience, the short- and long-wavelength spectrum edges are referred to as the "blue" and "red" edges, respectively. Figure 4 shows sample simulated prism-dispersed PSFs for the blue and red spectrum edges for detector 1, including "typical" optical aberration and spectrum tilt*. Note that the spectrum width near the blue edge is smaller than towards the red edge because of the wavelength differences.

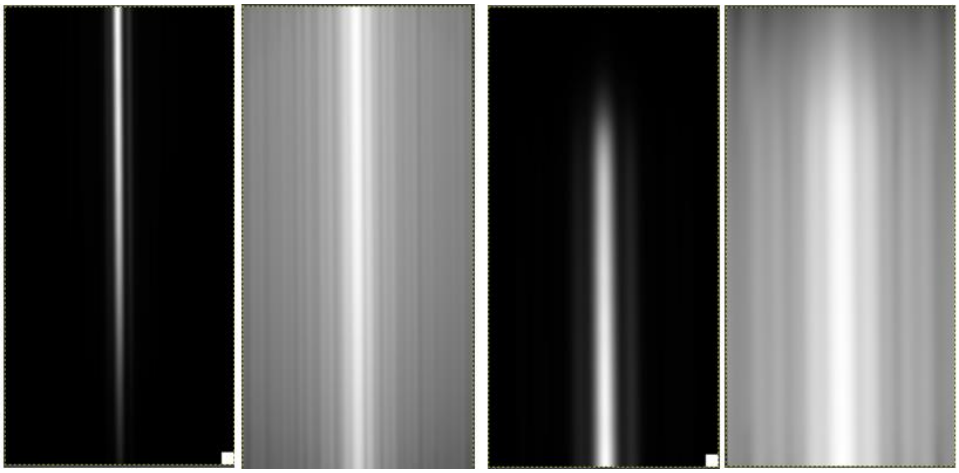

Figure 4. Simulated linear and logarithmic dispersed PSF plots for the WFI prism blue (left) and red (right) spectrum edges. The total box size for these images is 32×16 WFI pixels; the white square on the lower right of the linear images is the size of a single WFI pixel. Image resolution is at a scale of 64 fine pixels per WFI pixel.

At a high level, the procedure for computing a "centroid" for a spectrum edge when in FGS spectroscopy support mode mirrors the procedure for computing a star image centroid when in FGS imaging support mode. The biggest exception is that the FGS is not computing a center of light in the dispersion direction, but rather fitting the observed spectrum to a model that includes a 50% cutoff point for the edge.

Accumulation of FGS spectroscopy mode track phase GW window data is much the same as that for imaging mode, with the following exceptions. The parameter N for pedestal and signal frames is set to 12 to get better statistical reduction of detector read noise and to lengthen exposure time thereby getting more guide star photons. (Dispersing the guide star spectrum implies only a small fraction of total photons will be in each edge region, thus reducing signal relative to detector read noise. Increasing N helps to offset that effect.) The GW area is 32×16 pixels, vs. 16×16 for imaging mode, approximately doubling the time to read a GW once. This allows N_{TFPN} to be reduced from 3 to 2 with still adequate TFPN fading.

The FGS FSW creates a GW CDS image from pedestal and signal resultants. It calculates a first estimate for the horizontal (approximately cross-spectrum) centroid by (a) forming a mean row by averaging the N_{MR} rows (nominally 8) at the brightest edge of the GW, (b) locating the brightest

* Spectrum tilt for both grism and prism ranges over ~ -6° to +6° over the WFI field-of-view.

contiguous N_B pixels (3 for blue edge, 5 for red edge) in the mean row (which also provides a brightness estimate for the spectrum edge region), (c) calculates a 1D centroid using the DFT-based algorithm discussed above for imaging, and (d) optionally applies a correction for finite pixel size. With the spectrum horizontal centroid near the GW bright side known, the FSW constructs a centroid sample box with vertical length of 32 pixels (i.e., the GW length) and width equal to N_B plus padding to allow for the local tilt of the spectrum. Because tilt is more significant farther from the center of the WFI, this padding is wider for GWs nearer the left and right sides of the WFI field-of-view. This box is compressed horizontally to form a 1D, 32-pixel-long, column-matrix, spectrum edge region sample.

With the spectrum sample constructed, the FSW then applies an iterative, unweighted, least-squares algorithm to fit a spectrum model to the spectrum sample. The spectrum model used by the FSW is constructed as a product of three relatively simple components: the optics fast cutoff region, the optics slow decline region immediately brightwards of the edge, and the spectrum of the star through those regions modeled as a black body. The state vector for the fit consists of two elements: the optics fast cutoff region model 50% dropoff location in the GW reference frame, and a "gain" constant to match the model brightness to the data. For testing the algorithm, the stellar spectra used for creating guide window data in our HiFi simulator are taken from the Castelli and Kurucz catalog of stellar spectra at the Space Telescope Science Institute.[4] We have found in testing that use of our simple model, probably in particular use of a black body spectrum, results in a systematic bias in the estimate for the 50% cutoff (our vertical "centroid") as a function of stellar surface temperature. We have calibrated the effect and plan on having the GSSS-provided guide star reference coordinates adjusted for it.

FGS Special Challenges Synopsis

We summarize here some of the special challenges encountered during FGS algorithm development and strategies selected to deal with them.

- Transient fixed pattern noise. This was discussed above. Immediately following the transition from reading science sub-block pixels to GW pixels, the TFPN standard deviation can be upwards of 200 e^-. The FGS team strategy is to wait for it to fade, which typically happens within the time to read ~3 16×16-pixel GW frames.
- Imaging centroid correction for the effects of finite pixel size and optical aberration. (Recall the discussion associated with Figure 3 regarding these effects.)
- Guide star catalog error or detector relative alignment error. Systematic errors of these types may cause the expected positions of the guide stars to have unacceptably large relative errors. As discussed above, the FGS FSW can replace GSSS-specified guide star reference positions with FGS-measured centroids as science begins, optionally adjusted to keep the mean position of the new reference coordinates unchanged from those provided by the GSSS.
- Known bad pixels. Pixels can go bad, registering either cold or hot. Most will be known from ground operations regularly monitoring pixel health, with onboard bad pixel files being kept (mostly) up to date. For FGS imaging mode support, any 3×3-pixel box contaminated by a known bad pixel is deemed unsuitable for centroiding. (Note, for imaging mode, we are mostly relying on the GSSS avoiding selecting guide stars near bad pixels.) For FGS spectroscopy mode support, the signal for any known bad pixel within a guide window is replaced with the weighted mean value of its good neighbors (side neighbors given higher weight than diagonal neighbors).
- Unknown bad pixels. There can be a time gap between the last update of the onboard bad pixel files the time of FGS data processing. For FGS imaging mode support, we allow hot

unknown bad pixels to be centroided like a normal star image. However, there are a number of mechanisms in place to filter out such centroids; a centroid is removed if it (a) is outside the GSSS-specified brightness range for the guide star, (b) is too close to another centroid in the GW, (c) fails the pattern match test for acceptable position relative to the guide star centroids in the other detectors, or (d, unmentioned above) gets rejected during optional editing of post-fit residuals during attitude least-squares fitting. For FGS spectroscopy mode support, hot bad pixels will stand out as localized bright spots within the dispersed image. The strategy here is to identify any pixels with a signal that is unreasonably bright relative to their immediate neighbors and replace the value thereof with a weighted mean of the neighboring signals.

- Cosmic ray hits. Most cosmic ray hits are expected to manifest as localized bright spots. They are treated like unknown hot bad pixels.

CONTROL LAWS AND ESTIMATORS

FGS Acquisition Sliding Mode Control Law

Over the 5-year nominal mission lifetime, Roman will perform hundreds of thousands of science exposures. Before each science exposure, the observatory must slew to the target attitude, acquire the FGS, and settle attitude errors. To achieve the observational efficiency required to meet mission objectives, ACS is required to perform FGS acquisition and settling within 10 seconds, on average, per exposure. This ACS requirement is verified via Monte Carlo simulation campaign.

During Roman's slews to target attitude, ACS attitude estimation is derived from star tracker (ST) and Inertial Reference Unit (IRU) measurements. These slews exit within attitude accuracy errors of 5 arcsec, per axis, in angle error and 3 arcsec/sec, per axis, in rate error.

During Roman's science exposures, ACS attitude estimation is derived from FGS measurements. ACS initiates a science exposure once both FGS acquisition and settling are complete. FGS acquisition is complete once ACS has transitioned from ST/IRU-based control to FGS-based control. Settling is complete once control errors are within 0.02 arcsec, per axis, in angle error and 0.04 arcsec/sec, per axis, in rate error. These settling criteria ensure that attitude stability requirements will be met for the duration of the science exposure.

To meet the 10-second settling time requirement, Roman ACS relies on a nonlinear sliding mode control law that combines elements from linear control and bang-bang control.

Linear control law. Most control laws within Roman ACS are single-axis linear control laws. These controllers are relatively simple to implement, and stability can be proven via classical linear stability analysis techniques. Therefore, they are the preferred controller type for Roman ACS whenever their performance is sufficient to meet ACS requirements. A proportional-derivative (PD) linear control law was considered for FGS acquisition and settling. Eq. (1) defines the control torque $\mathrm{u_{PD}}(x, v)$, a function of angle error x and rate error v. K_x and K_v are controller gains.

$$\mathrm{u_{PD}}(x, v) = -K_x x - K_v v \tag{1}$$

The PD zero-torque line $\mathrm{v_{PD}}(x)$, defined by Eq. (2), is found by solving $\mathrm{u_{PD}}\big(x, \mathrm{v_{PD}}(x)\big) = 0$. Note that for $v > \mathrm{v_{PD}}(x)$, $\mathrm{u_{PD}}(x, v)$ is negative; for $v < \mathrm{v_{PD}}(x)$, $\mathrm{u_{PD}}(x, v)$ is positive.

$$\mathrm{v_{PD}}(x) = -\frac{K_x}{K_v} x \tag{2}$$

As demonstrated within the Roman ACS high-fidelity simulation environment, this linear control law is too slow to meet Roman's stringent settling time requirement.

Bang-bang control law. Bang-bang control is the classical minimum-time optimal control law assuming finite-control authority a_{max}. The term "bang-bang" refers to the control behavior – "bang" to one extreme to reach the switching curve, then "bang" to the other extreme to brake to a stop. Eq. (3) and (4) define the bang-bang switching curve $\mathrm{v_{BB}}(x)$ and control torque $\mathrm{u_{BB}}(x, v)$, respectively. $\mathrm{v_{BB}}(x)$ is derived from kinematic equations for constant accelerations of $\pm a_{max}$.

$$\mathrm{v_{BB}}(\mathrm{x}) = -\operatorname{sign}(x)\sqrt{2a_{max}|x|} \tag{3}$$

$$\mathrm{u_{BB}}(x, v) = \begin{cases} -a_{max} & v > \mathrm{v_{BB}}(x) \\ +a_{max} & v < \mathrm{v_{BB}}(x) \\ -a_{max}\operatorname{sign}(v) & v = \mathrm{v_{BB}}(x) \end{cases} \tag{4}$$

$\mathrm{u_{BB}}(x, v)$ brings the state from any initial condition to $\mathrm{v_{BB}}(x)$ as fast as possible. Along the curve, $\mathrm{u_{BB}}(x, v) = -a_{max}\operatorname{sign}(v)$ which brings the state along $\mathrm{v_{BB}}(x)$ to the origin in minimum time.

While bang-bang control is optimal in an ideal setting, small errors and signal delay lead to excessive actuation and overshoot. Modifying the bang-bang switching curve can result in a control law robust to signal delay. Generalizing $\mathrm{v_{BB}}(x)$ to the solution to a general parabola yields Eq. (5).

$$\mathrm{v_P}(x, \sigma_0, \sigma_1, \sigma_2) = \frac{\operatorname{sign}(x)}{2\sigma_2}\left(\sigma_1 - \sqrt{{\sigma_1}^2 + 4\sigma_2(|x| - \sigma_0)}\right) \tag{5}$$

Note that $\mathrm{v_{BB}}(x) = \mathrm{v_P}(x, 0, 0, (2a_{max})^{-1})$. Eq. (6) defines a modified switching curve.

$$\mathrm{v_{P^*}}(x) = \mathrm{v_P}(x, 0, T, (2a_{max})^{-1}) \tag{6}$$

It can be observed via simulation that using $\mathrm{v_{P^*}}(x)$ as the switching line instead of $\mathrm{v_{BB}}(x)$ results in signal delay robustness equivalent to T. While this control law is robust to delay, it is still subject to excessive actuation, and linear stability analyses are not applicable. For those reasons, it is not the best choice for the Roman ACS settling control law.

Sliding Mode Control Law. The concept of "sliding mode control" is to constrain the state trajectory to a subspace of the total phase space (i.e., to "slide" along a surface). To meet the ACS settling time requirement, Roman ACS uses a sliding mode control law to drive the observatory error states along the following switching curve:

$$\mathrm{v_s}(x) = \frac{|x|}{|x| + \beta}\mathrm{v_{P^*}}(x) + \frac{\beta}{|x| + \beta}\mathrm{v_{PD}}(x) \tag{7}$$

β, defined as the blending angle, is a controller parameter. β defines how $\mathrm{v_{PD}}(x)$ and $\mathrm{v_{P^*}}(x)$ blend together to create the sliding mode switching law. For comparison, Figure 5 shows the various switching curves considered. β is set to 2 arcsec for clarity of illustration. In the implementation, β is much smaller for performance.

Note the following features of $v_s(x)$:

$$\lim_{x\to\infty} \mathrm{v_s}(x) = \mathrm{v_{P^*}}(x) \tag{8}$$

$$\lim_{x\to 0} \mathrm{v_s}(x) = \mathrm{v_{PD}}(x) \tag{9}$$

Eq. (8) ensures that this control law behaves like a bang-bang control law for $|x| \gg \beta$; Eq. (9) ensures that this control law behaves like a linear control law for $|x| \ll \beta$. Therefore, linear stability techniques apply once in steady state. Control torque $\mathrm{u_s}(x, v)$ is defined by Eq. (10), where K_a is the sliding mode controller gain.

$$\mathrm{u_s}(x, v) = K_a(\mathrm{v_s}(x) - v) \tag{10}$$

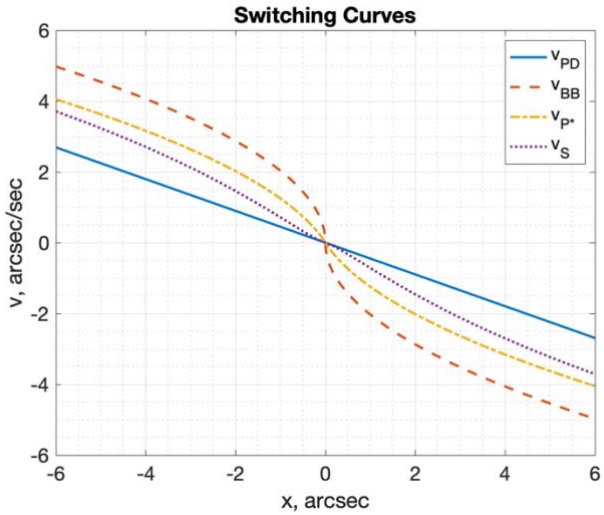

Figure 5. Comparison of PD, bang-bang, and sliding mode switching curves.

Monte Carlo simulation campaigns within the Roman ACS high-fidelity simulation environment verify that the performance of this control law is sufficient to meet the 10-second ACS settling requirement.

Attitude Estimators

Roman has several fixed-gain estimators associated with its control system design. In general outline, they follow the well-known Multiplicative Extended Kalman Filter (MEKF). It is found that acceptable performance is obtained using fixed gains, however, which allows avoidance of the complexity of using a state covariance matrix. Therefore, the estimator form may be more precisely termed a Multiplicative Extended Luenberger Observer (MELO). The first estimator used star trackers as its attitude measurement sensor and SSIRU as rate measurement sensor. The second estimator uses the fine guidance sensor as its only sensor input. The third estimator is used to estimate the misalignment between star trackers and the fine guidance sensor by comparing star tracker and FGS measurements when both are active.

ST/IRU Estimator

The Roman ST/IRU estimator is a fixed-gain estimator that uses the star tracker attitude measurements to update the IRU propagated attitude, IRU rate, and IRU rate bias states. The Roman ST/IRU estimator estimates four quaternion components for attitude, three components for rates and three components of the IRU drift rate or bias. The general state vector associated with this estimator is $\hat{x} = [\hat{q}(\delta\theta) \quad \hat{\omega} \quad \hat{b}]$ and its dynamics are represented by Eq. (11).

$$\dot{\hat{x}} = \mathrm{F}(\hat{x}) + L\big(z - \mathrm{h}(\hat{x})\big) \tag{11}$$

z represents the sensor measurements; ***L*** is the constant estimator gain matrix computed using pole placement; $\hat{\omega}$ is the estimated angular rate; $\delta\theta$ is the attitude update vector; $\hat{q}(\delta\theta)$ is the estimated quaternion computed from $\delta\theta$; $\hat{b}$ is the IRU bias estimate. $\mathrm{F}(\hat{x})$ is the time-varying state transition matrix. $\mathrm{h}(\hat{x})$ is the expected measurement vector. Although Eq. (11) is defined in continuous time, the implementation is done in discrete time. As in the MEKF, attitude update states are represented by a three-vector of small rotations between the measurement and the estimate.[5,6,7,8] During each attitude control cycle, the IRU and the star trackers are sampled, asynchronously with the ACS

cycle, but with timestamps provided to ACS. The previous control cycle's estimated attitude is propagated to the current control cycle time with a first order propagator using the IRU measured rate corrected with the previous cycle's bias estimate. The star tracker data are processed in the sensor data processing algorithms and back-propagated to the IRU sample time using the current body rate measurement. The estimator finally uses the synchronized measurements to update the propagated quaternion, estimated rate and bias states.

Fine Guidance System Estimator

The Roman Fine Guide Estimator has the same structure as the ST/IRU estimator. It uses only the FGS for measurement inputs. The internal states are attitude, rate, and an angular acceleration bias. Since the IRU is not used as a sensor, no gyro bias states are needed, but the acceleration states are useful to facilitate FGS acquisition. Their treatment parallels the rate bias states that they replace.

The FGS attitude error measurements are provided to ACS, transformed into the body frame for attitude control. The state vector associated with FGS estimator is, $\hat{x} = [\hat{q}(\delta\theta) \quad \hat{\omega} \quad \widehat{\alpha_b}]$ where α_b is the angular acceleration bias. The dynamic structure of the estimator is the same as the ST/IRU estimator. The estimator gains are also computed using pole placement.

The FGS measurements are taken at a different sample rate than the ACS cycle time, with the measurement rate varying depending on internal FGS activities. Each measurement is propagated to the current ACS time for the measurement update.

Roman's high-performance estimators are simple, robust and designed to accommodate its agility. Performance has been verified through Monte Carlo runs using high-fidelity simulation of the FGS and ACS systems working together.

Star Tracker Misalignment Estimator

Accurate alignment of the star tracker-derived attitude and the FGS is key to rapid FGS acquisition. This co-alignment will be calibrated during observatory commissioning, but there is concern that changes in thermal loading due to large attitude changes could cause alignment drift due to thermal deformations. To address this possibility, an onboard star tracker misalignment estimator is implemented. For each star tracker there is a misalignment estimator to track the misalignment between that star tracker's attitude estimate and the attitude estimate of the fine guidance sensor. The estimator takes the FGS attitude estimate and the individual star tracker measurement and finds the difference between them. The difference is converted to an angle vector and used as a residual. The current offset is shifted towards the measured offset by a fractional amount. The form of the estimator is a Luenberger observer:

$$\hat{x}_{k+1} = (1 - L)\,\hat{x}_k + L y_k \tag{12}$$

$\hat{x}$ is the estimate at a discrete point in time k, and y is the measurement. Note that there is no term included for dynamics. As the thermal misalignment is very slow (2.5 arcsec over an hour), it allows for this simplification.

WHEEL NULLSPACE MANAGEMENT

The Roman Space Telescope uses six reaction wheels as its primary ACS actuators. In addition to exerting feedback torques for attitude control, the wheel system must absorb the angular momentum that is steadily imparted by solar radiation pressure. Since three degrees of freedom are occupied by providing attitude control torque, there are three degrees of freedom remaining in the wheel nullspace.[9] The nullspace is used to address several, sometimes competing, ancillary

objectives: pushing through zero speed, maintaining wheel speed separations, and avoiding wheel speeds below 5 Hz.

Due to the observatory geometry and pointing constraints, solar radiation pressure exerts a steady torque along the +Y observatory axis. The wheel speeds vary at a steady rate in response. To maximize the intervals between thruster-based momentum unload events, the system momentum is initialized with a large -Y component, and the wheels absorb the solar torque until the system momentum has reached a specified large +Y value. So it is unavoidable that each wheel will cross through zero speed. The zero crossing event generates an attitude disturbance due primarily to the stiction torque. This disturbance may be partially mitigated by commanding a torque to push the wheel through zero. Appropriate torques are also commanded to all the other wheels such that no net torque is exerted on the observatory.

Wheel imbalance torques may excite structural vibration modes as the wheel speed (or a harmonic) crosses through the mode's resonant frequency. As this resonance is amplified if two wheels cross the mode at the same time, one goal of the nullspace control law is to keep wheel speeds separated from each other by a 1 Hz spacing. Because there are only three degrees of freedom in the nullspace, only three speed separations may be specified. This makes four wheel speeds trend in parallel, with the remaining two wheels constrained by the requirement to maintain attitude control. The control law sorts the wheels so that the two "free" wheels are those with the smallest speeds. The rationale is that wheel jitter disturbances are smaller at lower frequency, so that when two wheel speeds do coincide, it is at a less problematic speed.

Finally, there are several structural modes clustered at or below 4 Hz. It isn’t possible to keep all wheel speeds out of this range due to competing requirements, but it is possible to keep up to three wheels out of this range most of the time.

Figure 6 shows a typical wheel speed profile. At the beginning and end of the time interval, four wheel speeds are spaced at 1-Hz intervals. As the wheel speeds trend toward zero, three selected wheels are held at constant speeds to reduce interaction with low-frequency structural modes. And when a wheel approaches zero speed, it is actively pushed through zero, with other wheel speeds being adjusted to maintain momentum.

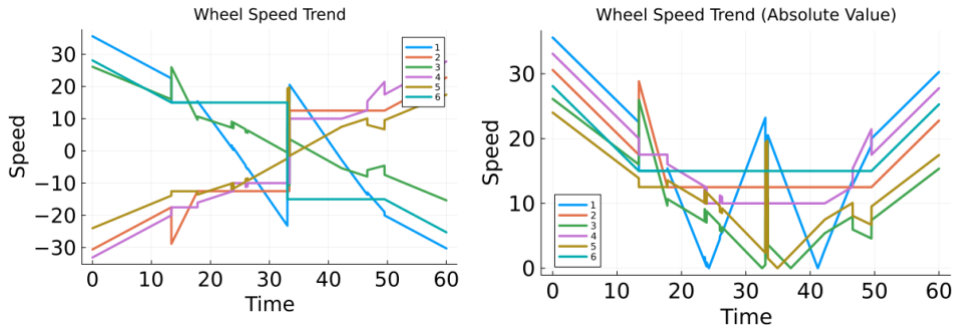

Figure 6. Typical wheel speed profile.

MANEUVER PERFORMANCE

Roman’s ACS incorporates two operational modes for maneuvering: Maneuver Mode for larger propulsive maneuvers such as Mid-Course Correction (MCC), Libration Orbit Insertion (LOI) and End-of-Mission Disposal, and Stationkeeping (SK) Mode for small maneuvers needed for orbit maintenance.

The performance of Roman's ACS in both Maneuver and SK Modes is influenced by several sources of error. Extensive Monte Carlo simulation shows that the driving sources of error are mass knowledge error, blowdown factor knowledge error, and thruster scale factor error.[10] All of these contribute to Delta-V (DV) errors because the onboard algorithm computes the DV achieved by integrating the estimated acceleration exerted by the thrusters during each ACS cycle. Errors in mass knowledge, blowdown knowledge, and individual thruster scale factor knowledge all appear as scale factor errors in the acceleration calculation.

Figure 7 shows DV execution error as a function of commanded maneuver magnitude. The effect of the main error sources is apparent in the growth of error with maneuver magnitude. Operationally, it is expected that mass estimate will need to be tracked and adjusted to maintain required maneuver accuracy over the mission.

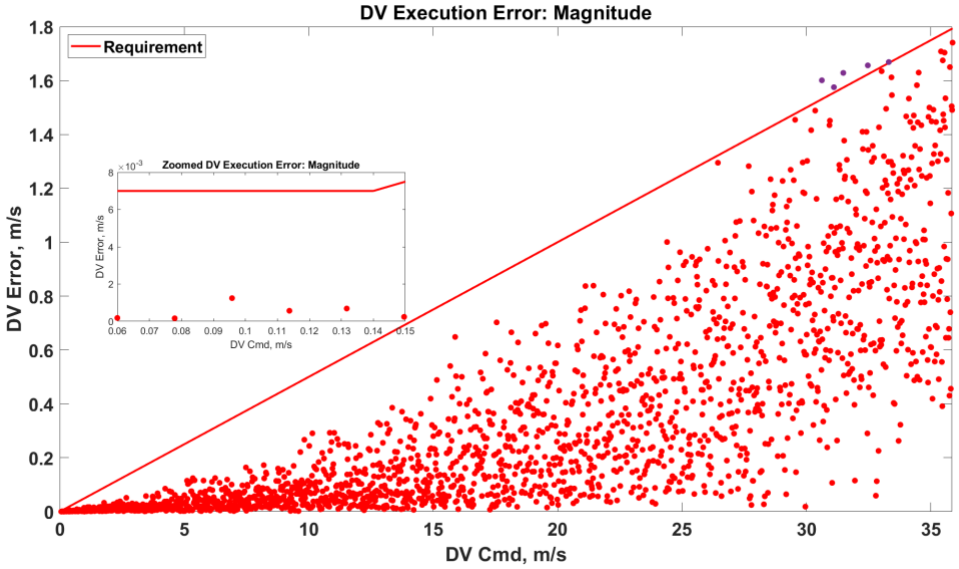

Figure 7. Maneuver Mode DV execution error magnitude.

HIGH-GAIN ANTENNA POINTING

The High Gain Antenna System (HGAS) on Roman consists of a two-axis gimbal system to control the HGAS. The two axes, denoted GA1 and GA2 as shown in Figure 8, receive commands from the ACS system to step to target angles. In the attitude control, the algorithms are computed in the Observatory Coordinate System using elevation and azimuth. The elevation axis coincides with the Observatory +Y axis and the azimuth axis coincides with the Observatory +Z axis when elevation is 0 deg. As shown in Figure 8, the transformation between the gimbal axes and the elevation and azimuth axes is defined by a scale factor and a bias. The GA1 axis has a scale factor of -1 and a bias of -10 degrees. The GA2 axis has a scale factor of +1 and a bias of 0 degrees.

Nominally the target of the HGAS will be directly at the ground station located on the Earth. The ground station position and velocity are computed on board and used by the HGA algorithm to determine the gimbal target angles needed to point the HGAS at the ground station. This determination is made by using inverse kinematics from the target boresight vector. The Attitude Ground System (AGS) will provide parameter updates after the HGAS is calibrated on orbit.

The High Gain Antenna is commanded into one of five main control modes: PARK, STOP, TRACK, SCIENCE, or HOLD. SCIENCE mode is further subdivided into two sub-modes: HOLD

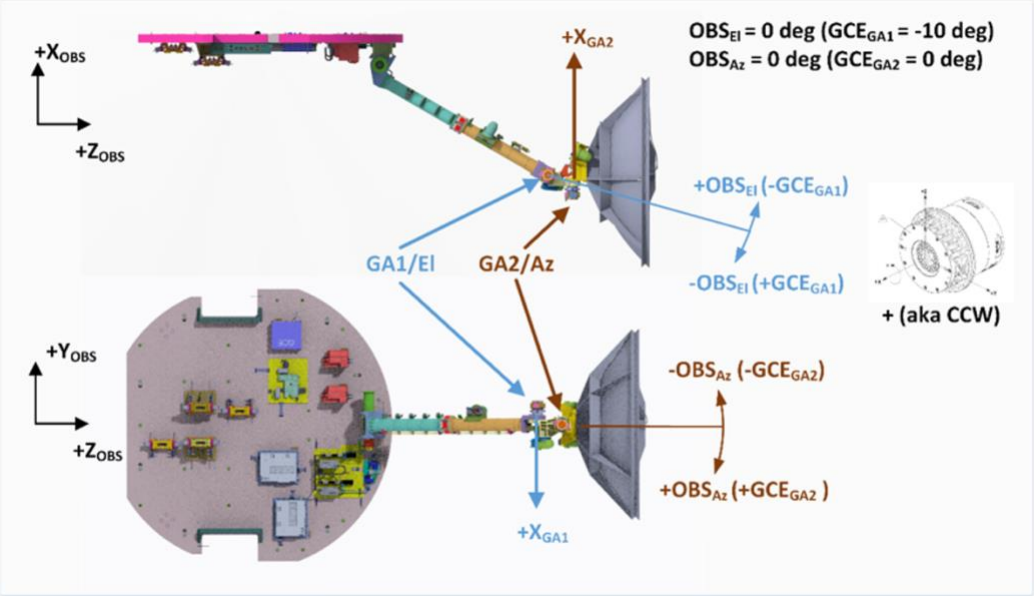

Figure 8. HGA axes definitions.

and STOP. Park mode is used to park the antenna at table-defined gimbal angles for either safing or other activities in which the HGA needs to be held at specified gimbal angles. Stop mode is used to ensure the gimbals are not moving. While in this mode all gimbal rate commands will be set to zero to ensure no movement from the gimbals. Track mode is used to follow the current ground station. Gimbal commands are sent to ensure the error from the current ground station is small. Science mode is used during observatory science activities. While the HGA is in Science mode, there is a further subdivision of modes into Science Hold and Science Stop. The difference between these modes depends on if the observatory is currently in a contact with the ground or not. Communication downlink takes priority over science when in contact with the ground. In this case, Science Hold is used and the HGA will send commands to move the gimbals to a new target if the computed error from the current ground station is larger than a threshold. If the observatory is not in contact with the ground, Science Stop is used and there will be no gimbal commands, to preserve the science

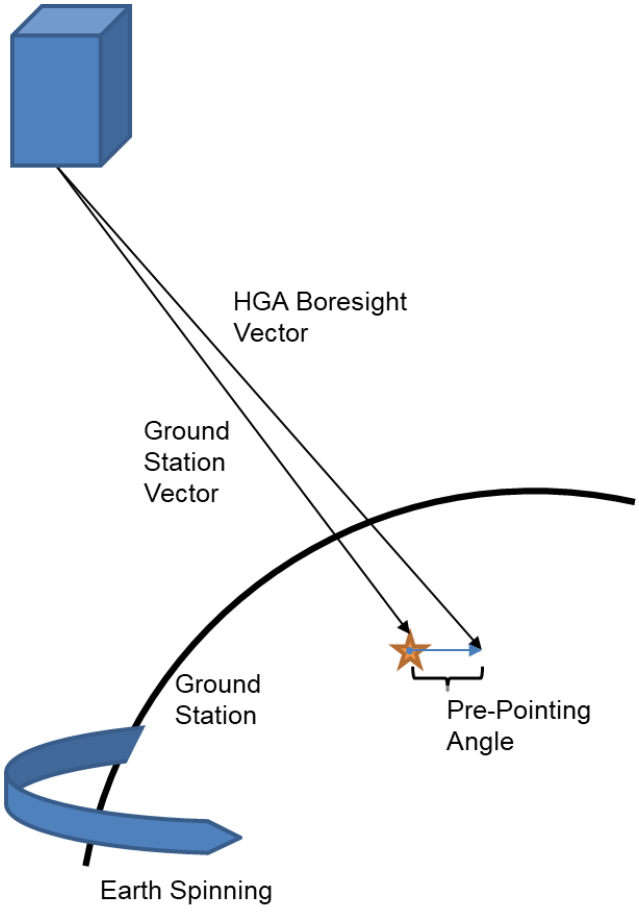

Figure 9. Diagram of HGA pre-pointing.

observation performance. Hold mode can be commanded directly as well and will operate similarly to Science Hold mode. This mode will send commands to move the gimbals only if the error from the current ground station is above a table-defined threshold.

In addition to the HGA modes, there are three separate Target Modes used by the HGA algorithm: None, Offset, and Prepoint. A target mode of None will set the target for the HGA directly at the current ground station. A target mode of Offset will use small offsets from the current ground station as the target to determine gimbal angles. A target mode of Prepoint will intentionally set the HGA boresight target ahead of the current ground station in the direction of the ground station drift vector. This will allow the maximum time for communication downlink and science observations to occur before the HGA would be required to move again.

After the target gimbal angles are determined, a gimbal step rate is computed based on defined maximum step rates and/or accelerations for each axis independently. The step error and step rate are transformed into the gimbal frame to send over to the Command Generation application for the commands to be sent to the gimbal electronics.

CONCLUSION

The Roman Space Telescope's attitude control system has several key design features that enable it to address Roman's challenging attitude accuracy, agility, and jitter requirements. The Fine Guidance System (FGS) is integral with the science instrument, and can guide on star images or on stellar spectra. The sliding mode controller used during FGS acquisition is key to rapid settling and science readiness. The ST-based and FGS-based estimators carry the sophisticated measurement handling of the MEKF, but opt for the simplicity of a constant-gain Luenberger observer to avoid maintaining a state covariance matrix. Potential thermal alignment drift of the star trackers relative to the FGS is addressed by an onboard alignment estimator, preventing such drift from degrading FGS acquisition performance. Excitation of structural vibration by wheel exported disturbances is minimized to the extent possible through wheel nullspace management. Stringent requirements on propulsive maneuver accuracy are verified by thorough Monte Carlo simulation campaigns. A lesson learned is that an onboard accelerometer would be a valuable aid in achieving maneuver accuracy. High-gain antenna control meets stringent pointing requirements while minimizing excitation of structural vibrations. The Roman ACS hardware suite consists of off-the-shelf and some custom components to support all aspects of the Roman mission with a single-fault-tolerant architecture.

REFERENCES

1. https://roman.gsfc.nasa.gov
2. http://www.teledyne-si.com/ps-h4rg.html
3. Pasquale et al., "Optical Design of the WFIRST-AFTA Wide Field Instrument", (SPIE-OSA 929305-1), SPIE-OSA/ Vol. 9293, 2014.
4. The Castelli and Kurucz 2004 Stellar Atmosphere Models, https://www.stsci.edu/hst/instrumentation/reference-data-for-calibration-and-tools/astronomical-catalogs/castelli-and-kurucz-atlas
5. F. Landis Markley, John L Crassidis. "Fundamentals of Spacecraft Attitude Determination and Control." Springer New York, 2014.
6. F. Landis Markley, Lefferts E. J.; 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. A. Gelb. "Applied Optimal Estimation." The MIT Press, 1974.

8. D G Luenberger. “Observers for multivariable systems.” IEEE Transactions on Automatic Control, Vol. AC-11, No. 2, April 1966, pp. 190 – 197.

9. F. Landis Markley, Reid G. Reynolds, Frank X. Liu, Kenneth L. Lebsock. "Maximum Torque and Momentum Envelopes for Reaction Wheel Arrays". https://ntrs.nasa.gov/api/citations/20110015369/downloads/20110015369.pdf

10. John M Hanson and Bernard B Beard. Applying monte carlo simulation to launch vehicle design and requirements verification. Journal of Spacecraft and Rockets, 49(1):136-144, 2012.