



MICROTHRUSTER-BASED CONTROL FOR PRECISION POINTING OF NEXT GENERATION SPACE TELESCOPES

Thibault L. B. Flinois,^{*} Cameron Haag,^{*} Milan Mandic,^{*} William Sanchez,[†]
Oscar Alvarez-Salazar,[‡] John Ziemer[§]

The next generation of space telescopes, such as the Habitable Exoplanet Observatory mission concept (HabEx), are expected to have milli-arcsecond-level pointing requirements at the spacecraft level. The pointing performance of such telescopes tends to be driven by internal rather than external disturbances, in particular reaction wheel jitter. This paper considers an architecture that relies on microthrusters instead of reaction wheels for spacecraft fine pointing, focusing on the HabEx mission concept as a case study. The proposed implementation uses micro-Newton-class colloidal thrusters that were tested in orbit as part of the ST7 payload on the LISA Pathfinder mission to reach Technical Readiness Level (TRL) 7. Key characteristics of these thrusters and the associated system trade-offs are first discussed. The resulting pointing control loop design is then described. Its ability to meet applicable pointing requirements is shown using three-degree-of-freedom simulations, where relevant disturbances are modelled. The control system is shown to reduce the pointing error by several orders of magnitude in about two hours, starting from coarse reaction control system (RCS) handover conditions. Once settled, the system meets the steady-state milli-arcsecond-level performance requirements with significant margin. Overall, the results show that microthruster-based pointing control architectures offer a promising alternative to traditional reaction-wheel-based designs for space telescope applications.

INTRODUCTION

The ambitious scientific objectives prioritized by the astronomy and astrophysics community in the National Academy of Sciences' 2020 Decadal Survey¹ continue to drive the demand for increasingly stable space-based observation platforms. The objective to conduct direct exoplanet imagery and spectroscopy using high-contrast coronagraphs dictates the stringent pointing stability requirements of future planet-hunting observatories such as the Nancy Grace Roman Space Telescope,² as well as the Large UV/Optical/Infrared (LUVVOIR) Surveyor³ and Habitable Exoplanet Observatory (HabEx)⁴ mission concepts.

Figure 1 is a survey of RMS pointing stability over a wide range of space-based observatory missions, plotted against launch mass (modified from Ref. 5 with Roman, LUVVOIR, and HabEx performance requirements added in bold). It shows that the exoplanet imaging observatories are targeting an entirely new regime of pointing stability, with requirements of sub-milli-arcsecond RMS stability at the instrument level.

It is common for space telescopes to use reaction wheels as the main pointing actuator of the spacecraft attitude control system. However, reaction wheels inherently impart structural vibrations on the spacecraft body, which deteriorate pointing performance at the instrument level. In addition, detailed modelling of the spacecraft's jitter response is a burden on the project. It typically only yields reliable results at relatively low frequencies, so large margins must be carried until the system can be tested.⁶

^{*}Payload/Instrument Pointing Control Analysis Group, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109, USA.

[†]Spacecraft Guidance & Control Systems Engineering Group, Ibid.

[‡]Guidance & Control Section, Ibid.

[§]Astronomy & Physics Strategic Missions & Technology Section, Ibid.

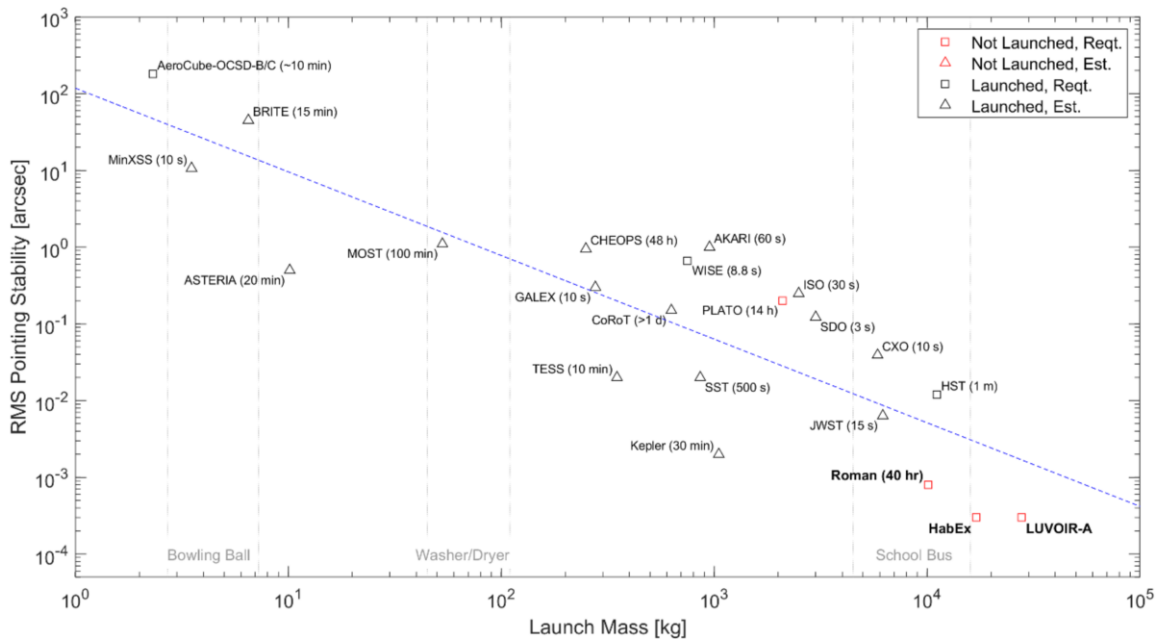


Figure 1. Instrument pointing stability vs. mass of various space observatories. A distinction is made between missions that have not launched and those that have launched, and between pointing stability numbers that are requirement values and those that are estimated performance values. The number in parentheses for each telescope is the required stability duration (if specified).

Reaction-wheel-induced jitter can be mitigated by adding vibration isolators, but these must be carefully tuned and also add mass. For instance, LUVOIR employs a novel non-contact vibration isolation system that mechanically separates the observatory instrumentation from ACS actuator disturbances in the spacecraft bus. Another approach, used for instance by the Roman Space Telescope, is to only operate the wheels at speeds that do not cause significant instrument-level jitter. However, such an operational constraint reduces the maneuverability of the spacecraft and also leads to more frequent desaturation maneuvers. Finally, a fast steering mirror (FSM) can be used in combination with a sensor (such as the low-order wavefront sensing camera on Roman's coronagraph instrument) to reduce the instrument-level jitter, especially at low frequencies.

A different approach is to entirely remove reaction wheel jitter by baselining thrusters for precision spacecraft pointing instead. Traditional thrusters have been shown to be beneficial for modern precision pointing applications where power is limited, as was the case for New Horizons.⁷ However, designs that rely on traditional pulsed thrusters also face several challenges: they consume propellant and impart unintended Δv on the spacecraft. Additionally, while reaction wheels can be placed virtually anywhere on the spacecraft, thrusters must be positioned and oriented in a way that provides the required control authority. Finally, if pulsed thrusters as used, pointing accuracy is driven by their finite minimum impulse. For this reason, some thrusters are specifically designed to have a low minimum impulse to mitigate this limitation.

Solar-electric propulsion (SEP) microthrusters, baselined for instance by HabEx, offer an alternative to reaction wheels and traditional pulsed thrusters. First, they may have a finite minimum thrust level but because they are throttled rather than pulsed, this characteristic does not necessarily deteriorate pointing accuracy, as discussed below. Second, as for any throttled thruster, noise in the output force should be expected. However, in the case of microthrusters, the noise can be at the $0.1 \mu\text{N}/\sqrt{\text{Hz}}$ level,⁸ resulting in instrument-level jitter that can be orders of magnitude smaller than with reaction wheels or traditional pulsed thrusters. Third, SEP is significantly more fuel efficient than chemical propulsion, mitigating the fuel consumption constraint. Fourth, with a typical force in the micro-Newton range, unintended Δv can also be expected to be lower.

In this paper, we introduce a SEP-microthruster-based fine pointing control architecture, based on the HabEx mission concept as a case study. We show that micro-arcsecond-level steady-state pointing stability can be achieved within approximately two hours, when starting from typical pulsed-thruster handover conditions. In the next section, relevant features of the HabEx spacecraft and its pointing architecture are introduced. Next, the proposed control design methodology, the simulation models, and initial conditions used to test the pointing performance are described. The results obtained with four test cases are then shown and key findings are discussed.

HABEX CASE STUDY

HabEx Spacecraft and Environment

The HabEx telescope is in a quasi-halo orbit around the Sun-Earth second Lagrange point (SEL2). Gravity gradients are relatively benign at SEL2 and the main external pointing disturbance that the spacecraft is subject to is solar radiation pressure (SRP) torque. The torque arises because the spacecraft's center of pressure (CP), where the SRP force acts, is shifted with respect to its center of mass (CM). One way to reduce the SRP torque is to reduce the CP-to-CM lever arm by deploying a solar sail during observations, as illustrated in Fig. 2. However, adding a deployable solar sail adds mass and complexity to the architecture. In the sections that follow, performance with and without a solar sail are compared. In addition to SRP torque, thermal settling after a telescope slew to a new target leads to a slow drift of the pointing reference direction that must be tracked by the pointing system.

In order to reject these disturbances, HabEx is equipped with inertial measurement units (IMUs), star trackers, and a Fine Guidance Sensor (FGS) that uses stars in its 2 arcmin field of view to measure the telescope pointing direction at sub milli-arcsecond level. It also has two sets of actuators. The first is an impulsive 1 N monopropellant Reaction Control System (RCS). The nominal RCS thrusters have similar performance to the Aerojet Rocketdyne MR-103H,⁷ which have been used in many missions including Voyager, Cassini, and New Horizons. As an alternative, thrusters designed to have a small minimum impulse, such as the MR-103M⁹ may provide better pointing control. However, the MR-103H thrusters have more flight heritage than the MR-103M. The performance of these two RCS thruster options is compared below. The second actuator system is a set of Busek SEP microthrusters, described in the next section.

The spacecraft uses its RCS thrusters in combination with its IMUs and star trackers to slew to the selected target star. The pointing performance of this configuration is sufficient to hand over from the star trackers to the FGS. Next, the RCS is handed over to the microthrusters. After transients, the spacecraft-level pointing accuracy requirement is 2 mas per axis (1σ). It is expected that this steady-state pointing accuracy will be reached within two hours of the RCS-to-microthrusters handover. Once steady-state is reached, further instrument-level pointing refinement brings HabEx within its science pointing requirement of 0.3 mas per axis. This paper focuses on the handover from RCS to microthrusters and the subsequent control system settling behavior.

Busek Colloid Micro-Newton Thrusters

Colloid micro-Newton thrusters (CMNTs), such as the design developed by Busek Co., Inc. in partnership with NASA's Jet Propulsion Laboratory (JPL), are electrostatic accelerators of electrospray droplets that provide precision spacecraft position and pointing control. Detailed information about the thruster design, heritage, and performance is included for instance in Refs. 8, 10, 11. The microthruster-based control architecture proposed in this paper takes advantage of the Busek thruster characteristics and of its heritage. In this section, after briefly introducing CMNTs, the key trades that have guided the design choices are outlined.

The Busek CMNTs were demonstrated as part of NASA's Space Technology 7 (ST7) payload hosted by the European Space Agency's (ESA's) LISA Pathfinder (LPF) technology demonstration mission in 2016¹¹ and reached Technical Readiness Level (TRL) 7. The CMNT technology met performance requirements operating at 5-30 μN of thrust with 0.1 μN resolution and under 0.1 $\mu\text{N}/\sqrt{\text{Hz}}$ thrust noise. A single cluster of four thruster heads from ST7 is shown in Fig. 3. For ST7, each thruster head contained $n_e = 9$ emitters. The range of thrust provided by a given thruster head scales with the number of emitters in the thruster. The thrust




Figure 2. Telescope geometry, showing the notional microthruster configuration. In the figure, H is 8.23 m, W is 5.25 m. In Cases 1 and 3, which do not have solar sail as described below, the full body length $L_{1,3}$ is 17.22 m. For cases 2 and 4, the body length $L_{2,4}$ is 22.60 m.

from each head can be throttled by changing the beam voltage and/or the propellant flow rate that determines the beam current. By precisely and independently controlling both the beam voltage and the beam current, the CMNT's output thrust can be finely tuned.

Noise. The noise level directly affects the achievable steady-state pointing accuracy. It scales with the number of thrusters firing simultaneously, but it is not affected by n_e to first order. Thus, for a given force command, firing a single large thruster results in less noise than firing several smaller thrusters in parallel.

Time response. Turning thrusters on or suddenly changing the commanded thrust level by more than about 20% leads to transients and overshoots of several seconds before the mean output thrust settles at the commanded value. When turning thrusters off on the other hand, the output force goes to zero within tens of milli-seconds. The duration of the transients and percentage force overshoots do not depend on n_e or the number of thrusters. As a result, frequently changing which thrusters are firing (or pulsing the thrusters) is likely to deteriorate the pointing performance.

Thrust resolution. The size of the output force resolution increment grows with n_e , so the smallest possible number of emitters should be used to improve pointing performance.

Heritage. It is desirable for the number of emitters n_e in each thruster head to be comparable to the 9-emitter ST7 thrusters. Similarly, no more than four thruster heads should be used in parallel to remain within the ST7 heritage.

Lifetime. At 25°C, the smallest thrust achievable by a given thruster is $F_{min} \approx 0.535n_e \mu\text{N}$. Commanded thrust between F_{min} and $F_{DLE} \approx 3.45n_e \mu\text{N}$, does not significantly affect lifetime. Conversely, commanding a thrust force above F_{DLE} can deteriorate life expectancy (hence "DLE") somewhat, and a command above $F_{max} \approx 4.85n_e \mu\text{N}$ significantly reduces lifetime and should be avoided.

Power and mass. Mass and power are not primary drivers of the pointing architecture because the thruster needs are relatively small compared to the total telescope budgets. For power, increasing the number of emitters per head rather than the number of thruster heads is a more power-effective way to increase control authority. The hardware mass of a given thruster head increases linearly with n_e and each new head comes




Figure 3. Picture of the Busek Colloid Micro-Newton Thruster Flight Cluster from ST7, including four thruster heads, electronics, and cathode neutralizer (visible) in thermal-vacuum environmental test setup.

with further mass overhead, suggesting the total number of thrusters should be minimized if possible.

Summarizing, it is desirable to allocate a given force command to a single thruster at a time, at a level between F_{min} and F_{DLE} , and to avoid sudden changes to the command level of a given thruster, or switching which thrusters are firing, as much as possible. An allocation algorithm that behaves in this manner is proposed for steady state, when disturbances to be rejected are quasi-constant. During the comparably short handover periods from coarser control configurations however, the commanded force is time varying and cannot be predicted ahead of time. As a result, a separate thrust allocation strategy is introduced for handovers. It retains the desired allocation characteristics, but is slightly less fuel efficient.

PROPOSED CONTROL SYSTEM DESIGN

In this section, the thruster configuration and two thrust allocation strategies are described: one for handovers and one for steady state.

Thruster Configuration

The thruster configuration is designed to be both simple and consistent with the desired thrust allocation behavior described above. To achieve these goals, 12 clusters of thrusters are mounted on the spacecraft body, at the locations shown in Fig. 2. Each thruster head in a given cluster has a different number of emitters n_e . All clusters have the same number of heads and the same n_e distribution between the heads. For instance, in Case 1 introduced below, each cluster has 4 thrusters, which have 1, 4, 20, and 127 emitters respectively. Additionally, all thrusters in a given cluster are oriented in the same direction. At any given time, thrust is allocated to at most one thruster in any given cluster.

The n_e distribution shared by all clusters is chosen as follows. First, one of the heads in the cluster is always a one-emitter thruster, as it provides the finest possible resolution and the smallest possible F_{min} value. Second, the F_{DLE} value of the largest thruster is sized by the largest expected force command over the full mission. Third, a non-zero overlap is required between the maximum thrust $F_{DLE,i}$ provided by each thruster and the minimum thrust $F_{min,i+1}$ provided by the next thruster. If the thrust ranges provided by smallest and largest thrusters do not overlap, at least one intermediate thruster is needed. The intermediate n_e values are chosen to maximize the smallest overlap ($F_{min,i+1} - F_{DLE,i}$) between any two consecutive thrusters. When the commanded thrust is not in an overlap region, only one thruster can provide the desired force. Within an overlap region, a choice must be made between two thrusters. As illustrated in Fig. 4, hysteresis is included in the thruster selection logic to limit how often thrusters are switched on and off. The approach is most effective if the overlap region is significantly larger than the thrust noise level and any short-term variation in the thrust command.

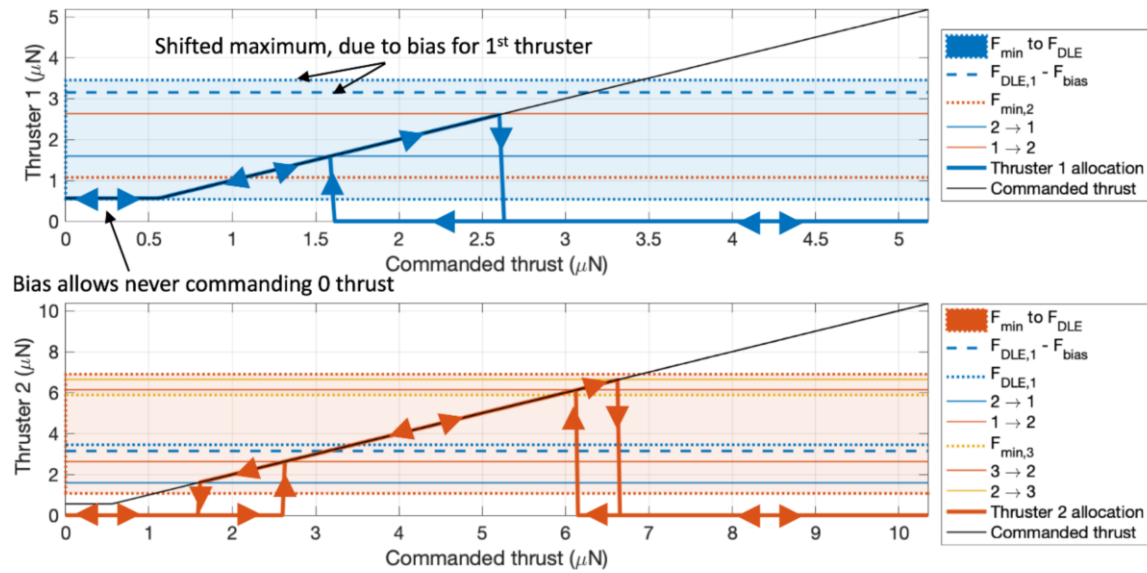


Figure 4. Thrust allocation illustration, showing overlap in achievable thrust between consecutive thruster heads in a given cluster and showing hysteresis.

Steady-State Thrust Allocation

In steady state, the dominant disturbance is the SRP torque, which is quasi-constant at the time scale of an observation. Thermal settling also leads to a slow pointing drift. The steady-state allocation strategy described in this section is tailored to this quasi-steady disturbance torque.

The torque command that the controller outputs is allocated to clusters using a standard pseudo-inverse approach. The resulting raw allocation can have some commands that are larger than the maximum achievable thrust and other that are negative. In addition, with the present configuration, positive but near-zero commands may also be problematic because all heads have a nonzero minimum thrust value. Consequently, commands close to or smaller than $F_{min,1}$ would lead to erratic on-off switching transients, which would deteriorate pointing performance. Instead, a small bias is added by setting the minimum force allocation to a value slightly larger than $F_{min,1}$ rather than zero. Biasing the allocation in this manner ensures the smallest thruster remains on for zero or near-zero commands and thus avoids on-off transients. The fuel penalty associated with this approach is small, as discussed below.

In practice, the transformation from the pseudo-inverse allocations $F_{0,i}$ to the final allocations F_i is of the

form

$$F_i = aF_{0,i} + b. \quad (1)$$

If the spacecraft thruster configuration is symmetric such that the torque resulting from all clusters firing with the same force simultaneously is nominally zero, then its direction is unaffected by the transformation. It can be shown that $F_{bias} \leq F_i \leq F_{lim}$ holds for all clusters i if the parameters a and b are calculated as follows:

$$a = (F_{lim} - F_{bias}) / (F_{0,max} - F_{0,min}), \quad (2)$$

$$b = (F_{bias}F_{0,max} - F_{lim}F_{0,min}) / (F_{0,max} - F_{0,min}), \quad (3)$$

$$F_{0,max} = \max_i (F_{0,i}), \quad F_{0,min} = \min_i (F_{0,i}). \quad (4)$$

In practice, the applied torque does drift in response to thermal settling and slow SRP torque variations. In some cases, it is therefore necessary to switch from one head to the next in a given cluster and the resulting on-off transients are expected to temporarily deteriorate the pointing performance. However, if the distribution of heads is chosen to provide sufficient overlap between thrusters, then the hysteretic switching logic (Fig. 4) ensures that these transients are rare.

Handover Thrust Allocation

When the RCS loop hands over to the microthrusters loop, the system needs to settle to steady state conditions. During this phase, the steady-state allocation strategy described above is not viable because of the time-varying and unpredictable nature of the torque commands. Instead, a simpler allocation approach is adopted, whereby the largest thruster is used in all clusters until the system settles. This strategy ensures the maximum torque authority is available during the handover and also avoids on-off transients. However, the largest thrusters also have the largest F_{min} value, so to avoid large on-off transients the same biasing approach [i.e., Eq. (1)] is used during handovers. The fuel penalty of the handover bias is larger but because handovers are comparatively short, the overall propellant mass impact remains small. The system autonomously switches from handover to steady-state allocation by monitoring a windowed pointing stability metric: when the standard deviation of the pointing error within a sliding time window drops below a threshold, steady-state allocation is enabled. Conversely, if the performance deteriorates for any reason and the windowed pointing stability metric grows beyond another threshold, handover allocation is enabled again to allow the system to settle again.

Controller Design

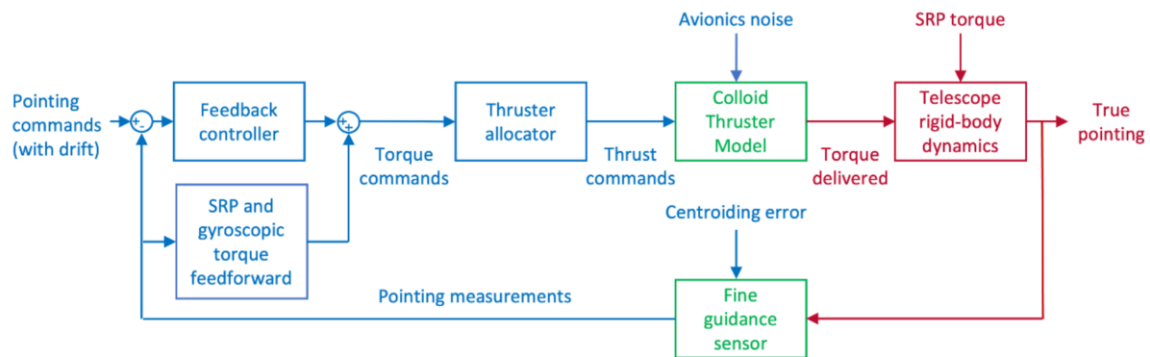


Figure 5. Control loop block diagram. Discrete signals and blocks are shown in blue, continuous/discrete hybrid blocks are shown in green, and continuous signals and blocks are shown in red.

In this section, the design of the control loop, depicted in Fig. 5 is described. Relying on the high performance of the FGS, pointing measurements are fed to the controller without further filtering or estimation. The controller output torque command is the sum of a feedforward and a feedback term. The feedforward term compensates the SRP torque predicted by an on-board model, which is subject to modelling errors compared to the “true” SRP torque experienced by the spacecraft. The feedback term is the output of a controller which comprises of a lead-lag compensator and an integrator. Its purpose is to reject initial errors from the handover, SRP modelling errors, thermal alignment drift, thrust errors, and noise due to the microthruster system. A fixed controller structure is chosen such that it can be tuned with a single parameter that sets the closed-loop bandwidth, while maintaining lower bounds on the gain and phase margins of 6 dB and 30 deg.

Increasing the controller bandwidth improves disturbance rejection and reduces handover settling time. However, it also increases the maximum torque commanded by the controller during the handover. The overall peak force allocation in worst-case handover conditions directly drives the number of emitters in the largest thruster. It can potentially also determine the number of thrusters needed in each cluster. These trade-offs are further discussed in the sections below.

For each case described in this paper, it was found that a single controller was able to settle to steady state with the required pointing accuracy, within an acceptable time and without excessive actuator commands. For more challenging handover conditions, it is possible to design a lower-bandwidth controller that ensures the peak handover commanded force remains below a desired value and a separate higher-bandwidth controller to provide the desired pointing performance in steady state. The transitions between the two controllers and between the two allocations must then be done in sequence to avoid compounding the two sets of transients.

SIMULATION SETUP

Simulation Models

In this section, models used in the three-degree-of-freedom (3-DOF) simulation are described. Thermal drift is modelled as an exponential decay with an amplitude of 1 arcsec and a time constant of 15 hours. An SRP N-plate force model¹² is used and the spacecraft shape is approximated by a box of dimensions shown in Fig. 2. The surface is assumed to be 85% specular, 5% diffusive, and 10% absorptive. In all simulations, the worst case where the sun direction is along the telescope Y axis is considered. The force is applied at the center of pressure, normal to the Z axis. This orientation produces a torque about the X axis with the maximum spacecraft frontal area. On board, the model used for SRP feedforward torque compensation has a fixed error with respect to the true dynamics, applied in the direction of the nominal SRP. The error has a mean of 3 μNm and a standard deviation of 0.5 μNm . In cases where a solar sail is deployed, it is assumed to be capable of rotating about the the telescope Z axis, such that it is always normal to the sun vector.

The thruster cluster position standard deviation is 2 mm/axis and the thrust direction standard deviation of each cluster is 0.67°/axis. The thrusters are also subject to quantization. The center of mass and center of pressure positions are also known to 2 mm/axis (1σ). The uncertainty on all moments and products of inertia is 5% (1σ).

The FGS centroiding accuracy standard deviation is modelled as

$$\sigma_{FGS} = FWHM / (c SNR \sqrt{n_{stars}}) \quad (5)$$

$$FWHM = \sqrt{(1.028\lambda/D)^2 + (2.35\sigma_{smear})^2} \quad (6)$$

$$SNR = \sqrt{n_{photons}} \quad (7)$$

$$n_{photons} = f_{photons} (\pi D^4/4) (t_{integ})(QE)(T), \quad (8)$$

where $FWHM$ is the full-width half maximum of the point spread function, SNR is the signal-to-noise ratio, and $n_{photons}$ is the number of photons collected in a given exposure. This SNR equation assumes read noise, dark current, and background noise are negligible. The parameter c is set to 2, the number of stars is $n_{stars} = 8$, the chosen wavelength is $\lambda = 1000$ nm, the telescope diameter is $D = 4$ m, the assumed smear due to camera motion during the integration is set to the requirement value of 2 mas, the photon flux is $f_{photons}$

= 85000 ph/m²/s, the quantum efficiency is $QE = 80\%$, and the throughput is $T = 74\%$. The FGS accuracy is thus a function of the integration time t_{integ} . With this model, the measurement accuracy is on the order of 0.04 mas for 1 Hz measurements, so its contribution to the total 2 mas requirement is small. The FGS measurements are delayed due to the finite integration time. The avionics add a further delay of 1 second until the measurements can be processed.

Commands to thrusters are processed through an inner 10 Hz avionics loop, which controls the voltage and current commands to the individual thruster heads to track the desired thrust level. This avionics loop is the source of some noise, lag, and sometimes overshoots, as the thrust level settles to the right value. It is also constrained by quantization of the measured voltage and current used in an internal control loop to adjust the output thrust. The avionics loop tuning from Ref. 8 was used and not adjusted in this study. When thrust commands go to zero, however, the avionics loop is bypassed and the thrust level goes to zero in two time steps.

Initial Conditions: RCS Handover

The settling time of the microthruster control system is dependent to the residual angle and rate error of the RCS deadband limit cycle just prior to actuator handover. The rate error of the RCS limit cycle is bounded by the smallest change in angular rate that an opposing pair of RCS thrusters can impart on the spacecraft, as described by Eq. (9):

$$|\omega_{LimitCycle}| < \frac{I_{Bit}d_{ThrusterPair}}{J_{Spacecraft}} \quad (9)$$

For a given spacecraft inertia $J_{Spacecraft}$ and thruster pair moment arm $d_{ThrusterPair}$, the maximum of the limit cycle rate error $\omega_{LimitCycle}$ is computed from the minimum impulse bit I_{Bit} of the chosen RCS thrusters. As mentioned above, this HabEx case study examines two 1-N monopropellant RCS thruster options that achieve minimum impulse bit performance potentially suitable for implementation alongside microthrusters.

The first is the Aerojet MR-103H, which has extensive flight heritage (e.g., Voyager, Cassini, New Horizons) and a flight-proven 5 mN s minimum impulse bit. The second is the ‘‘Minimum Impulse Thruster’’ MR-103M, developed by pairing a MR-103H thruster with a new fast response solenoid valve. The thruster-valve combination was put through an extensive Δ -qualification test program to achieve a fivefold reduction in minimum impulse bit, with $I_{Bit} = 1$ mN s. Table 1 lists the HabEx parameter values used in Eq. (9) and the computed limit cycle maximum rate error for each RCS thruster option considered. The values of $d_{ThrusterPair}$ and $J_{Spacecraft}$ correspond to the RCS thruster geometry and spacecraft inertia for motion about the telescope X axis.

Table 1. Handover initial conditions: key parameters for the two RCS thruster options.

Thruster	MR-103H	MR-103M
I_{Bit}	5 mN s	1 mN s
$d_{ThrusterPair}$	5.25 m	5.25 m
$J_{Spacecraft}$	280,000 kg m ²	280,000 kg m ²
$\omega_{LimitCycle}$	19.3 mas/sec	3.9 mas/sec
$\omega_{IniCond}$	38.6 mas/sec	7.8 mas/sec
K	3	3
$\theta_{Deadband}$	174 mas	35 mas
$\theta_{LimitCycle}$	174 mas	35 mas
$\theta_{IniCond}$	348 mas	70 mas

The maximum angle error of the RCS deadband limit cycle $\theta_{LimitCycle}$ is determined by the RCS thruster controller tuning. For this case study, a phase plane switching logic is implemented to model RCS control about the telescope X axis, based on the Cassini RCS controller, and is illustrated in Fig. 6(a). The RCS

deadband geometry is constrained by two user-defined parameters (shown in green). The first is the RCS angle deadband $\theta_{Deadband}$, which bounds $\theta_{LimitCycle}$ and determines the thruster pulse frequency and fuel consumption during the RCS deadband limit cycle. The second is the parameter K , which tunes the damping behavior of the closed loop as the limit cycle is approached. The overall behavior of the RCS limit cycle is insensitive to K over a range of reasonable values. In this study, K is set to 3 based on Cassini heritage. The values used in this study are listed in Table 1, along with the corresponding max angle error $\theta_{LimitCycle}$.

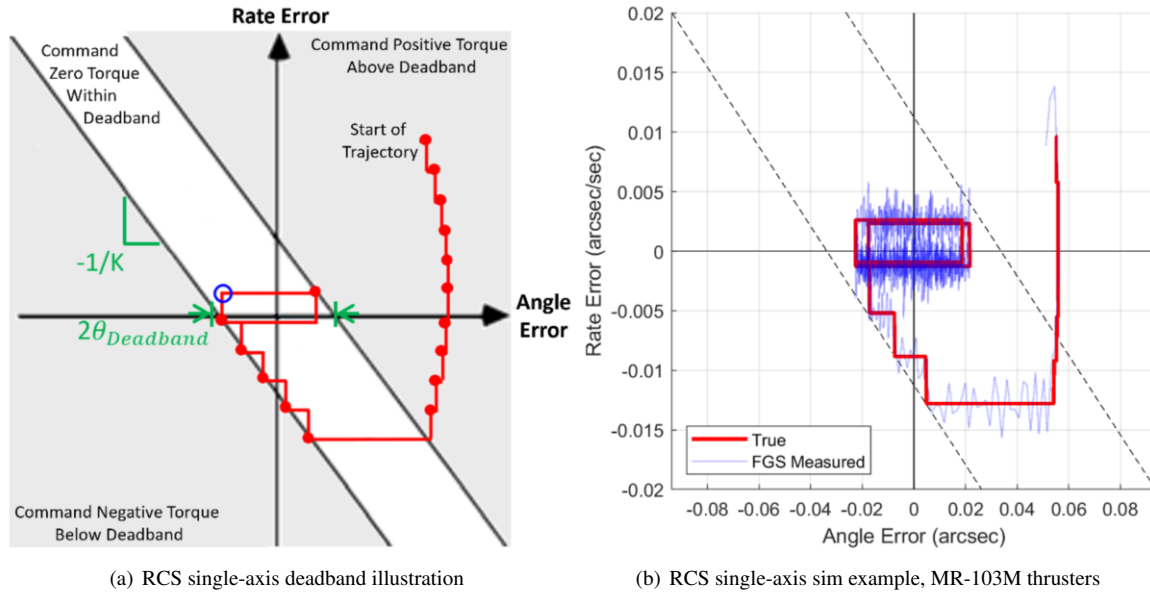


Figure 6. RCS single-axis deadband and limit cycle. Subfigure (a) illustrates a notional RCS deadband with design parameters (green), error trajectory and limit cycle (red), and max limit cycle errors (blue). Subfigure (b) shows a single-axis RCS sim example with MR-103M thrusters. Both figures show the case with solar sail included.

A separate 1-DOF time domain simulation was constructed to model the closed-loop RCS thruster performance and confirm the limit cycle behavior on the telescope X axis. The simulation consists of a single-axis RCS deadband controller running at 8 Hz, coupled with a rigid body spacecraft plant. The RCS thrusters are set to operate at their minimum impulse bit firing times, and also include a 20% impulse variation and 5ms fixed pulse delay. FGS integration time is set to the 8 Hz control rate, which corresponds to a centroiding RMS error of 0.1 mas. SRP torque is set to 0 or 3 mNm to represent cases with and without a solar sail, respectively. Both MR-103H and MR-103M thrusters are simulated, with corresponding deadband parameters and spacecraft inertia listed in Table 1. Figure 6(b) shows the time domain results of the MR-103M thruster case with solar sail included. This scenario confirms that, for either RCS thruster option, the size of the RCS limit cycle is determined by the thruster minimum impulse bit and deadband parameters, rather than being limited by FGS sensor noise or thruster errors. Since a solar sail is included to remove SRP torque, the RCS displays a two-sided limit cycle, bouncing between a positive and negative angle error bounded by $\theta_{Deadband}$. With SRP torque included, the RCS displays a one-sided limit cycle, bouncing off of either a positive or negative angle error limit. In this case, the thruster pulse frequency and fuel usage of the RCS limit cycle increases, but the values of $\omega_{LimitCycle}$ and $\theta_{LimitCycle}$ are unchanged.

The settling time performance of the microthruster control architecture is evaluated by assuming the RCS handover occurs at both the max angle error $\theta_{LimitCycle}$ and max rate error $\omega_{LimitCycle}$ of the RCS deadband limit cycle. In practice, the RCS deadband limit cycle produces a negative $\theta_{LimitCycle}$ when $\omega_{LimitCycle}$ is positive and vice versa, as indicated by the blue circle in Figure 6(a). For additional conservatism, $\theta_{LimitCycle}$ and $\omega_{LimitCycle}$ are assumed to occur simultaneously and in the same direction as the solar pressure torque disturbance. Finally, to calculate the initial condition values $\theta_{IniCond}$ and $\omega_{IniCond}$ used in the 3-DOF

simulations, $\theta_{LimitCycle}$ and $\omega_{LimitCycle}$ are multiplied by a conservative model uncertainty factor (MUF) of 2 to encompass variations for instance in RCS thruster performance and spacecraft inertia.

RESULTS AND DISCUSSION

Table 2. Distribution of number of emitters n_e per thruster head.

Case	n_e	F_{min} (μN)	F_{DLE} (μN)	$F_{Overlap}$ (μN)
1	1	0.54	3.45	-
	4	2.14	13.80	1.31
	20	10.70	69.00	3.10
	127	67.95	438.15	1.06
2	1	0.54	3.45	-
	3	1.61	10.35	1.85
	18	9.63	62.10	0.72
3	1	0.54	3.45	-
	3	1.61	10.35	1.85
	18	9.63	62.10	0.72
	115	61.53	396.75	0.58
4	1	0.54	3.45	-
	2	1.07	6.90	2.38
	6	3.21	20.70	3.69

Using the simulation setup and initial conditions described above, four cases were investigated:

1. MR-103H, no solar sail,
2. MR-103H, with solar sail,
3. MR-103M, no solar sail,
4. MR-103M, with solar sail.

In each case, the controller was tuned to obtain approximately a two-hour settling time, using a standalone simplified controller design tool. The settling time is defined as the instant at which the pointing standard deviation drops below 2 mas over a 10-minute sliding window. The tool is fully linear, single axis, assumes nominal and ideal parameters. Based on the maximum force command predicted by the tool, the thruster head distribution at each cluster was defined using the procedure outlined above. The four thruster distributions are shown in Table 2. The top set of plots in Figs. 7 to 10 shows the pointing error and the applied torque time history predicted by the 1-DOF tuning tool, where each contributor is identified.

The controller designed with the 1-DOF tool is then used in the main 3-DOF simulation. The same controller, with coefficients scaled by each axis' principal moment of inertia, is used for each axis. The bottom set of plots in Figs. 7 to 10 shows the response from a representative (but not necessarily a worst-case) 3-DOF simulation for each considered case. Key results from these simulations are summarized in Table 3.

During the transition from handover to steady-state thrust allocation, the thruster head used in each cluster may change. The resulting thruster on-off response appears as a slight hump in the pointing error and as a small impulse in the thrust response in all cases. Although the controller tuning tool does not account for this transition, setting the transition threshold to a relatively high value ensures the true settling time remains relatively unaffected. In all cases, the sensor noise contribution is negligible during the handover. Once settled, the pointing accuracy meets the 2 mas/axis requirement by a factor of over 50 in all cases. The

Case	1	2	3	4
Closed-loop bandwidth (mHz)	2.0	2.0	1.6	1.6
Settle time (hour)	2.11	2.10	2.44	2.25
Steady-state accuracy (mas/axis RMS)	0.035	0.036	0.035	0.037
Handover propellant consumption (10^{-3} kg)	1.62	0.25	1.20	0.08
Steady-state propellant consumption rate (10^{-3} kg/hour)	1.18	0.07	1.21	0.09
Peak pointing error and associated axis (arcsec/axis)	5.35 (X)	5.42 (X)	2.79 (Z)	1.16 (X)
Peak torque and associated axis (mNm/axis)	2.92 (X)	0.41 (X)	2.62 (X)	0.10 (X)

Table 3. Simulation results summary from representative responses of Cases 1 to 4.

fuel requirements are small: the handover consumption is under 2 grams in all cases and the steady-state consumption corresponds to less than 11 kg of propellant if run continuously for a year in all cases.

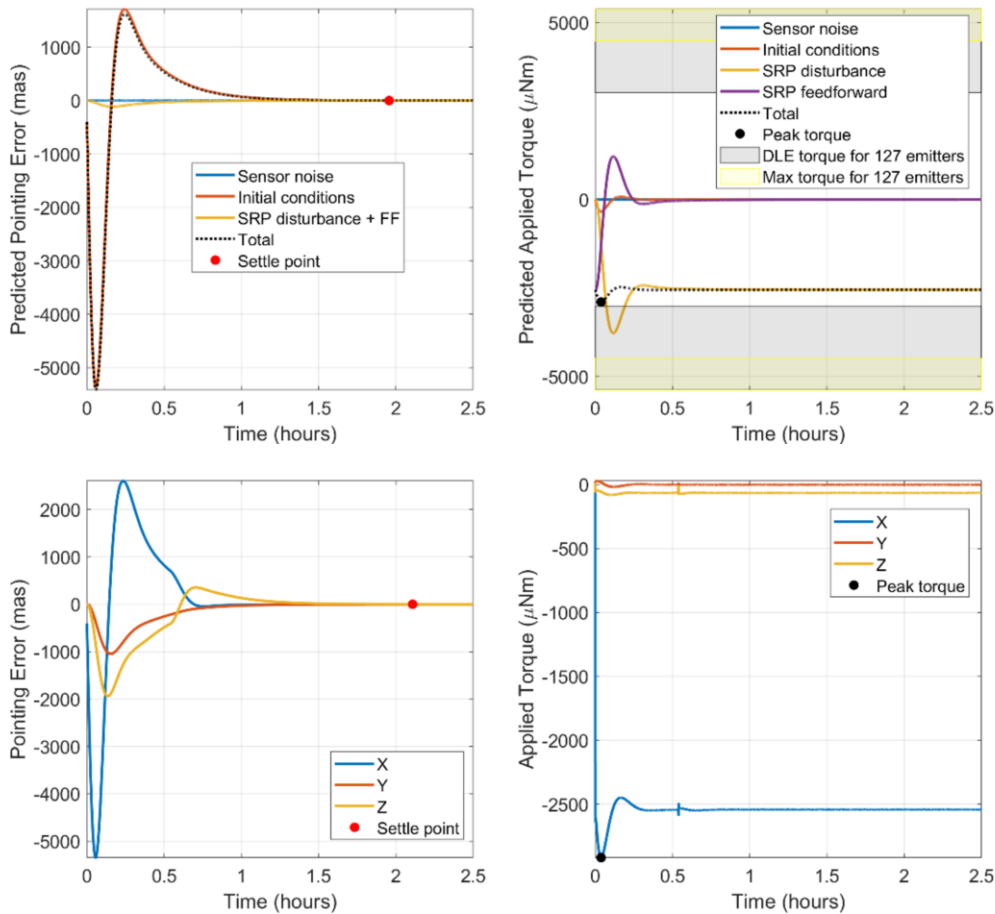


Figure 7. Case 1: MR-103H, no solar sail. Top: Pointing error and applied torque responses, as predicted by the 1-DOF controller design tool, showing individual contributions. Bottom: Pointing error and applied torque responses, from a representative 3-DOF simulation, showing contributions from each axis.

Case 1 is the reference case, where neither mitigation is used. Most of the thrust control authority is needed to reject the SRP torque, while most of the pointing error response is due to initial transients from the RCS handover. The SRP modelling error does not have a dominant impact on the pointing error. Because of the

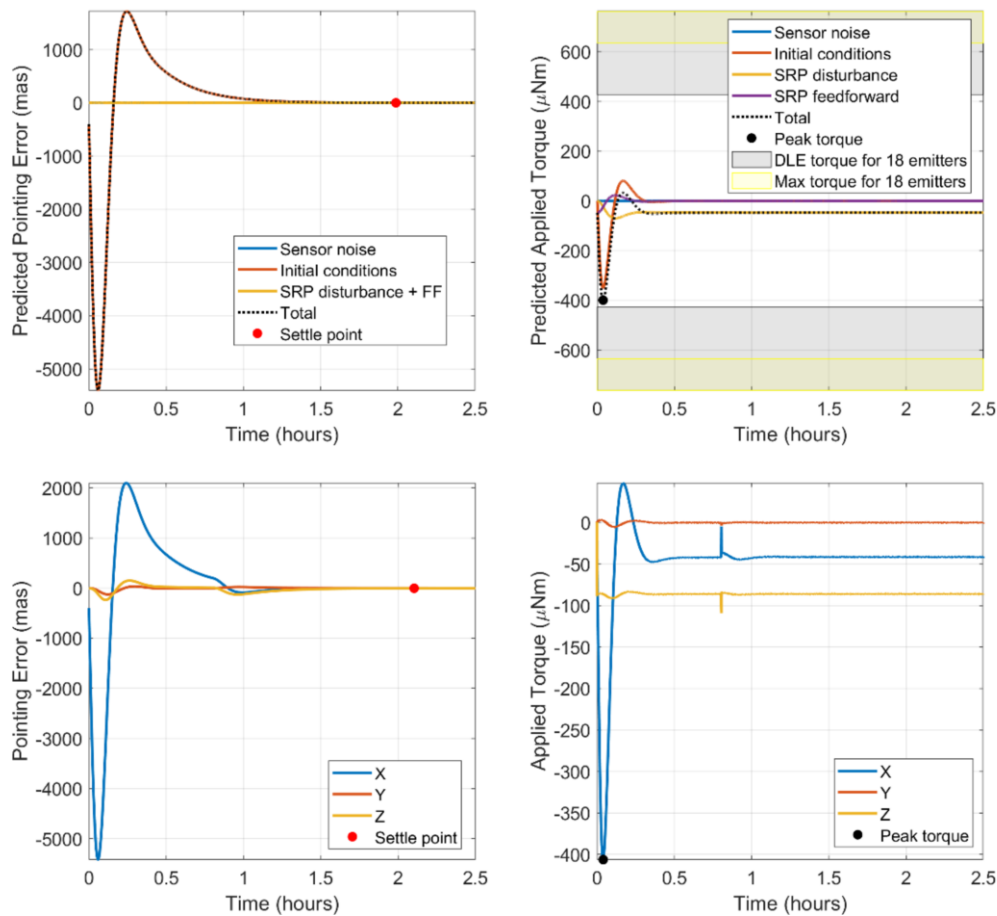


Figure 8. Case 2: MR-103H, with solar sail. Top: Pointing error and applied torque responses, as predicted by the 1-DOF controller design tool, showing individual contributions. Bottom: Pointing error and applied torque responses, from a representative 3-DOF simulation, showing contributions from each axis.

large SRP torque, the thruster distribution has 4 heads and the largest thruster has 127 emitters, i.e., far from the ST7 heritage. The full X-axis response in the 3-DOF simulation is similar to the response predicted by the 1-DOF tool, with a peak error of about 5 arcseconds. The Y-axis and Z-axis pointing errors also reach close to 1 and 2 arcseconds respectively, due to the non-ideal mass properties and thruster response.

In Case 2, a solar sail is added, significantly reducing the SRP torque. The microthruster control authority is therefore mostly used to reject handover initial conditions. As a result, the thruster distribution has only 3 heads and the largest thruster has 18 emitters, i.e., only a factor of two larger than the ST7 9-emitter configuration. Additionally, reducing the absolute SRP torque leads to a correspondingly small X-axis total torque response. In turn, the Y-axis and Z-axis pointing errors are also reduced, since uncertainties in mass properties and thruster alignments lead to smaller absolute torques in these axes. However, because the SRP torque is mostly compensated in feedforward, reducing it does not lead to a significant improvement in the X-axis pointing error response.

When using the minimum impulse thrusters without the solar sail in Case 3, the peak of the RCS-handover contribution and therefore total pointing error is considerably reduced. Thus, the contribution of the SRP torque modelling error is more dominant than in Cases 1 and 2. As there is no solar sail however, the control

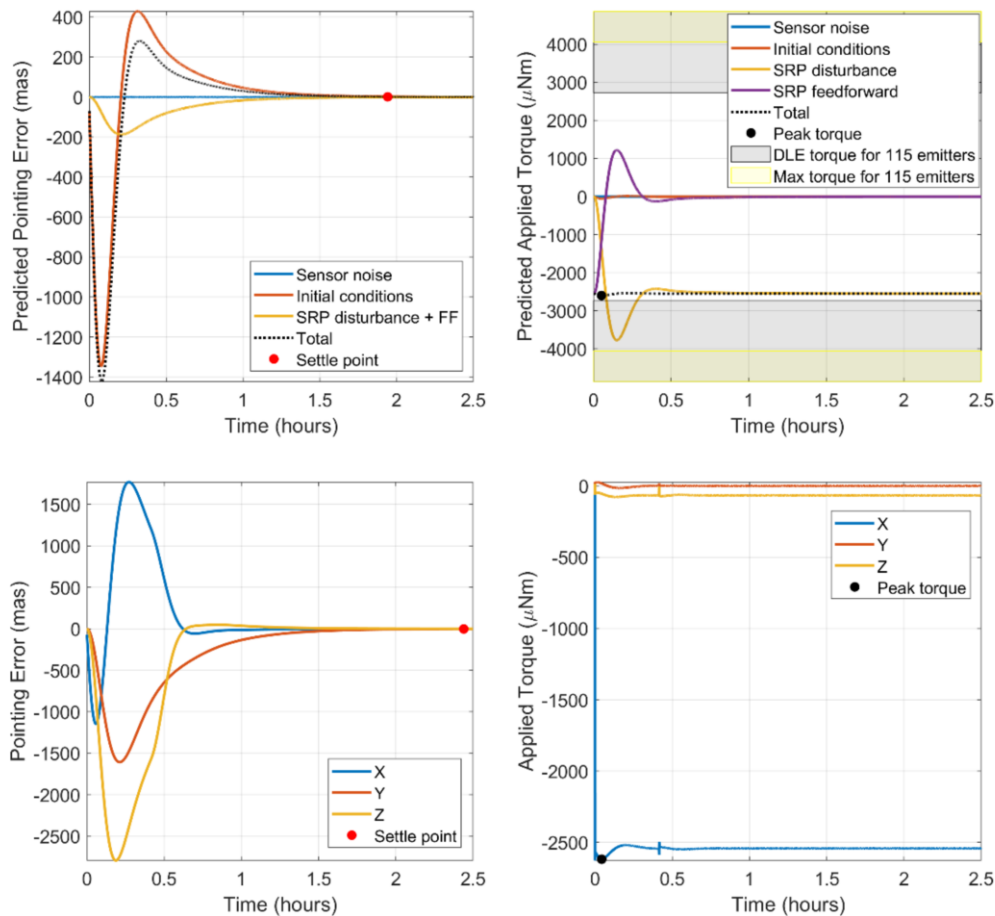


Figure 9. Case 3: MR-103M, no solar sail. Top: Pointing error and applied torque responses, as predicted by the 1-DOF controller design tool, showing individual contributions. Bottom: Pointing error and applied torque responses, from a representative 3-DOF simulation, showing contributions from each axis.

authority is driven almost entirely by the need to compensate the unmitigated SRP torque. Consequently, the thruster distribution needs 4 heads as in Case 1, but with a slightly reduced largest thruster of 115 emitters – still far from ST7 heritage. Unlike in Cases 1 and 2, the full 3-DOF response has a longer settling time of around 2.4 hours, and differs more markedly from the predicted 1-DOF response. The discrepancy is such that the worst per-axis pointing error is in the Z axis and not in the X axis. Thanks to the minimum-impulse thrusters, the pointing error response is not as strongly dominated by the (fixed) worst-case initial conditions, and is more dependent on distributed parameters, such as thruster misalignments and mass properties, and their effect on the Y-axis and Z-axis response. These effects are not included in the simplified 1-DOF tool, which predicts that a less aggressive controller tuning (1.6 vs 2.0 mHz closed-loop bandwidth) is sufficient to obtain a two-hour settling time. Increasing the bandwidth back to 2.0 mHz does bring the true settling time back within 2 hours. Overall, the value of the minimum-impulse thruster lies mostly in the reduction of the pointing error transient following the handover.

Finally, bringing together the solar sail and MR-103M in Case 4, the pointing error is driven by the handover initial conditions, as in Cases 1 and 2, but its magnitude is reduced, as in Case 3. Both initial conditions and SRP modelling error are significant contributors to the peak applied torque, and the thruster distribution has only 3 heads and the largest thruster has only 6 emitters. This configuration is thus within the demon-

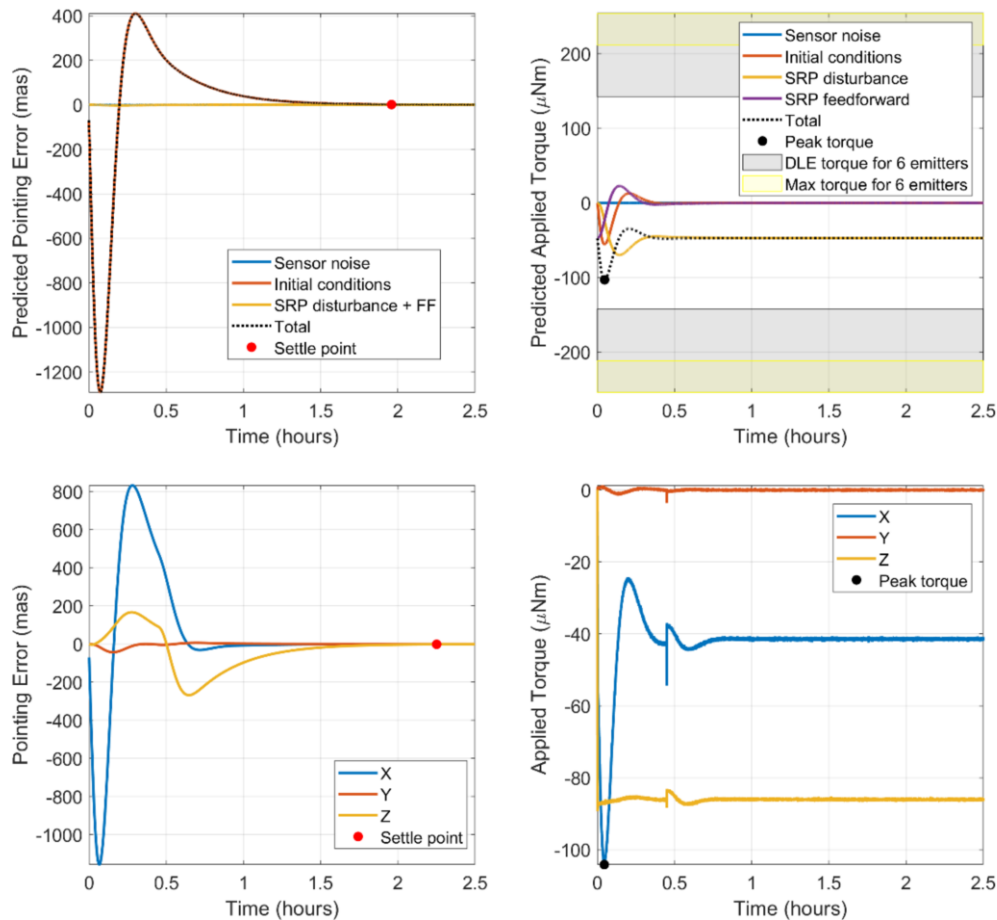


Figure 10. Case 4: MR-103M, with solar sail. Top: Pointing error and applied torque responses, as predicted by the 1-DOF controller design tool, showing individual contributions. Bottom: Pointing error and applied torque responses, from a representative 3-DOF simulation, showing contributions from each axis.

strated ST7 heritage. As in Case 3, thanks to the MR-103M performance, the initial conditions are less of a driver for the pointing error response than in Cases 1 and 2 and a slightly increased settling time of 2.25 hours is observed. As before, the settling time is reduced back to two hours if the controller bandwidth is increased to 2 mHz. Using both mitigations together is synergistic: while MR-103Ms provide favorable hand-over conditions, reducing the X-axis pointing error, the solar sail reduces the total control torque magnitude and therefore its leakage into the other axes.

CONCLUSIONS

In this paper, a control architecture that relies on microthrusters instead of reaction wheels for fine pointing is presented. It relies on the Busek colloidal microthrusters demonstrated in orbit as part of the LISA Pathfinder mission. The architecture includes a controller, with both feedback and feedforward terms, that rejects SRP torque as well as initial pointing errors and rates. It also tracks residual thermo-mechanical alignment drift. Two thrust allocation algorithms are used. They assume a symmetric configuration, with clusters of several thruster heads, chosen to span the needed range of thrust. The first algorithm allows smooth handovers from a coarser RCS-based control loop. The second maintains the system in steady state with high pointing accuracy. This approach takes advantage of the properties of the Busek microthrusters, by minimiz-

ing the occurrence of transients caused by thrusters being switched on or off. Both allocation algorithms also impose a small thrust bias in the commands, which circumvents the performance limitations imposed by the microthrusters' non-zero minimum thrust level.

The architecture was tested in a 3-DOF simulation based on HabEx spacecraft parameters, starting from conservative initial conditions and including all relevant disturbances. In all the considered cases, the control architecture was shown to bring the system to a steady state pointing accuracy under 0.04 mas, i.e., more than 50 times better than the 2 mas requirement, within about two hours of handover. Faster handovers may be achieved by using larger microthrusters. The fuel consumption is small: each handover was found to require under 2 grams of propellant and the steady-state fuel consumption was found to consume under 11 kg per year of operation.

In the baseline case, the system is able to settle within two hours with an arcsecond-level peak handover pointing error. However, the microthrusters required are an order of magnitude larger than the ST7 heritage configuration. Two potential mitigations were investigated to further improve the handover and steady-state performance. First, a solar sail was used to reduce the SRP torque caused by the offset between the spacecraft center of mass and center of pressure. By considerably reducing the required control authority, the solar sail brings the size of the required microthrusters close to or within the heritage of ST7 for a two-hour handover. Second, RCS thrusters designed to have a small minimum impulse were used to improve the conditions at handover to the microthruster loop. Minimum-impulse thrusters considerably reduce the peak pointing error experienced during handovers. Used together, these two improvements lead to a handover with small pointing transients, and which can be achieved with a microthruster configuration that is within the ST7 demonstrated heritage.

These results demonstrate that microthruster-based pointing architectures are feasible, fuel-efficient, and can readily provide sub-milli-arcsecond pointing accuracy. Future work will focus on fully assessing performance statistics using Monte Carlo simulations. Performance may also be further improved by tuning the inner thruster avionics loop to reduce overshoots and on-off transients. Higher fidelity may also be achieved by fully integrating the RCS loop in the simulation and adding in the effects of flexible-body dynamics and slosh.

ACKNOWLEDGMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). © 2022. All rights reserved. California Institute of Technology. Government sponsorship acknowledged.

REFERENCES

- [1] National Academies of Sciences Engineering and Medicine, *Pathways to Discovery in Astronomy and Astrophysics for the 2020s*. Washington DC: The National Academies Press, 2021, 10.17226/26141.
- [2] WFIRST Science Definition Team, "Wide-Field Infra Red Survey Telescope - Astrophysics Focused Telescope Assets Final Report," tech. rep., 2013.
- [3] LUNAR Mission Concept Study Team, "LUNAR Final Report," tech. rep., 2019.
- [4] Astromomy Physics and Space Technology Directorate at the Jet Propulsion Laboratory, "HabEx Final Report," tech. rep., 2019.
- [5] C. M. Pong, "On-Orbit Performance & Operation of the Attitude & Pointing Control Subsystems on ASTERIA," *AIAA/USU Conference on Small Satellites*, 2018.
- [6] C. Dennehy and O. S. Alvarez-Salazar, "A Survey of the Spacecraft Line-of-Sight Jitter Problem," *Advances in the Astronautical Sciences*, 2019, pp. 1–18.
- [7] J. Stratton, "The Use of the Aerojet MR-103H Thruster on the New Horizons Mission to Pluto," *International Astronautical Federation - 55th International Astronautical Congress*, Vol. 11, 2004, 10.2514/6.iac-04-s.1.09.
- [8] J. Ziemer, C. Marrese-Reading, C. Cutler, C. Dunn, A. Romero-Wolf, S. Javidnia, T. Le, I. Li, P. Barela, N. R. Demmons, V. J. Hruba, J. Slutsky, J. I. Thorpe, P. Maghami, O. Hsu, and J. O'donnell, "In-Flight Verification and Validation of Colloid Microthruster Performance," *2018 Joint Propulsion Conference*, 2018, pp. 1–23, 10.2514/6.2018-4643.

- [9] J. M. Parker and M. J. Wilson, "The Minimum Impulse Thruster," *JANNAF Conference*, 2006.
- [10] M. Gamero-Castano, V. Hruby, D. Spence, N. Demmons, R. McCormick, C. Gasdaska, and P. Falkos, "Micro Newton Colloid Thruster for ST7-DRS Mission," *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, No. July, 2003, pp. 1–9, 10.2514/6.2003-4543.
- [11] J. Ziemer, C. Marrese-reading, C. Dunn, A. Romero-wolf, C. Cutler, S. Javidnia, T. Le, I. Li, G. Franklin, P. Barela, O. Hsu, P. Maghami, J. O'Donnell, J. Slutsky, J. I. Thorpe, N. Demmons, and V. Hruby, "Colloid Microthruster Flight Performance Results from Space Technology 7 Disturbance Reduction System," *Proceedings of the 35th International Electric Propulsion Conference*, 2017, pp. 1–17.
- [12] A. Farres, D. Folta, and C. Webster, "Using Spherical Harmonics to Model Solar Radiation Pressure Accelerations," *Adv. Astro. Sci.*, Vol. 162, 2018, pp. 3365–3383.