STABLE study for the Habitable Worlds Observatory

Tiffany M. Glassman*[a], Jonathan Arenberg[a], Andre Aroyan[a], Nick Burst[a], Suzanne Casement[a], Scott Colby[a], Todd Lawton[a], Gordon Reese[a], Blake Shepherd[a], James Tucker[b]
[a] Northrop Grumman, [b] Kratos SRE
*tiffany.glassman@ngc.com; phone 1-310-812-3789

ABSTRACT

We report on progress on the STABLE study, which was awarded on the NASA 2023 ROSES D.19 opportunity. In the study, we are developing technologies for the Habitable Worlds Observatory (HWO) in four areas: Micrometeoroid Impact Mitigation, Telescope Temperature Control, Ultra-Stable Telescope Structure Materials and Modeling, and Improvement of Thermal Expansion Test Capability. This study is a collaboration between Northrop Grumman, Kratos, NeXolve, JPL, STScI, NASA's White Sands Test Facility, and NASA Goddard's Ultra-Stable Test Facility. We will also discuss our plans for our follow-on study as part of NASA's 2024 ROSES D.19. In this study, we will address additional technologies as well as perform analysis of HWO architecture options.

Keywords: Habitable Worlds Observatory, Telescope, Technology, Industry, D.19

1. INTRODUCTION

Our first study (STABLE) consists of four focus areas with the potential to significantly affect the Habitable Worlds Observatory (HWO) architecture. We focus on two key HWO architectural features: the deployable barrel around the telescope that will mitigate the effects of the external environment (e.g., solar heating variations with attitude, micrometeoroids, and stray light) and the telescope structure that will support the optical train while mitigating the impact of system-induced disturbances. We chose to focus on these areas to both gain timely understanding of how the required capabilities will drive the HWO architecture (in terms of complexity, stowed volume, mass, etc.) and to mature the technologies (thus reducing risk).

Technologies, modeling, and model-validation testing for a barrel that can deploy, maintain thermally stable telescope optics (inclusive of the support structure), and protect against sporadic micrometeoroids of various energies range from very mature (e.g. deployment mechanisms) to very immature (e.g. modeling of low energy micrometeoroid propagation). Before designing a barrel for HWO, the requirements for how many layers, of what materials, etc. must be developed. Our first two study efforts are designed to provide the information necessary to define these parameters.

Similarly, there are technologies, modeling capabilities, and model-validation testing for the telescope support structure that must be extended beyond the current state of the art to reach the ultra-stability needed for HWO. For standard engineering structural parameters such as stiffness, damping, and coefficient of thermal expansion (CTE), this means modeling and testing material properties to levels of precision beyond what has been achieved previously. There are also subtle effects that must be examined more closely to ensure they don't affect observatory stability, such as susceptibility to micro-dynamics.

In our latest study (STABLE-S), we will continue to develop several of these technology areas. For example, developing methods of shielding the telescope from MM and exploring new aspects of the ultra-stable structural design. We will add a technology development effort to study impact of stray light on the observatory and quantify the effects at the level needed for coronagraphy. In addition to these technology development efforts, we have started a study of the HWO design space to identify robust, executable architectures with minimum regrets and to inform the project in time to influence NASA's Mission Confirmation Review and ensure success at that key review.

In this paper, we summarize our progress on each of the STABLE study topics and describe our plans for the STABLE-S study.

2. MICROMETEOROID IMPACT MITIGATION (STABLE AND STABLE-S)

The goal of this study topic is to develop our understanding of how micrometeoroids (MM) behave when impacting thin membranes that we expect to use to protect the HWO telescope. We have developed a physics-based model of their behavior and a series of tests to validate and refine the model. So far, we have completed two of the three planned test campaigns under the STABLE study and are completing analysis of the data. Initial indications are that facture toughness of the impactor (MM) is the most influential variable in determining if an impactor will pass through a potential shield. Lower velocity impacts (which cause less shock to the impactor) result in lower or no fragmentation of tough, metallic, impactors. Our analysis shows that a small fraction of the natural MM spectrum is expected to be in this penetrating class of collisions.

Our tests are conducted at NASA's White Sands Test Facility (WSTF) using the Hypervelocity Test Laboratory. Test series #1, conducted at WSTF in August 2025, consisted of 24 shots covering a range of parameters. We used impactors with different materials (Ruby, Soda Lime Glass, or Steel) and test speeds (4 and 7 km/s). Our test membrane samples, manufactured by NeXolve, are made of 0.5 mil or 1 mil CP-1 thin-film material. Some of the samples were laminated for extra strength, referred to as toughened. A test fixture (Figure 2-1) holds ten single layers of a single material type, with the first five layers separated by ~6 cm and the last five by ~ 3 cm. The layer separation was designed to minimize overlap that complicates hole counting and analysis while also limiting cloud spread to allow capture of the entire impact region in the microscope scan.

Figure 2-1. A test article showing the 10 membrane layers held in the test fixture and placed in the chamber at WSTF.

After each shot, the target layers were removed and photographed under a microscope. These images were then analyzed using a method we developed to count holes (Figure 2-2). This code also measures the size of each hole which is related to, but not identical to, the size of the fragment that caused it. To increase the accuracy of our analysis we must translate hole sizes into fragment sizes (which depend on impactor velocity and film properties). In dense regions, we must also account for overlapping of holes and for "fellow travelers" (small particles that go through larger holes). We have improved our analysis to account for these effects and have therefore increased our statistical understanding of the results.

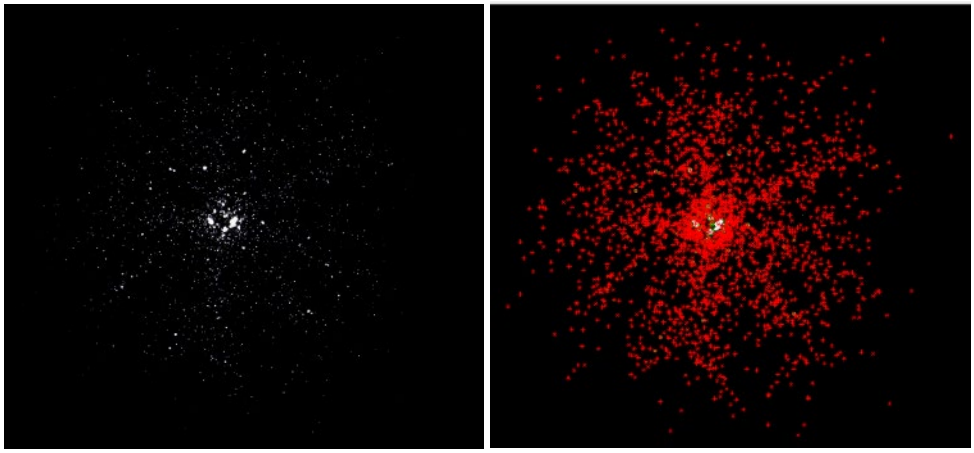

Figure 2-2: Example of a membrane sample with holes created by impactor fragments. On the left is the raw image and on the right is the image with all holes found by our software marked.

Our preliminary conclusions from the results of Test #1 are based on the last damaged layer from each shot, summarized in Figure 2-3. There was very little variation in results depending on the membrane properties, indicating that this is of lesser importance. The toughest impactor type (steel) penetrated all 10 layers in every case and didn't fragment at all for the slower collisions. These results show that the most damaging case is for a tough impactor (metallics) moving at relatively low velocity (~4 km/sec). A benefit of this finding is that it allows for better testability of the MM impact effects, since the most dangerous range of impact velocities overlap with what's achievable on the ground.

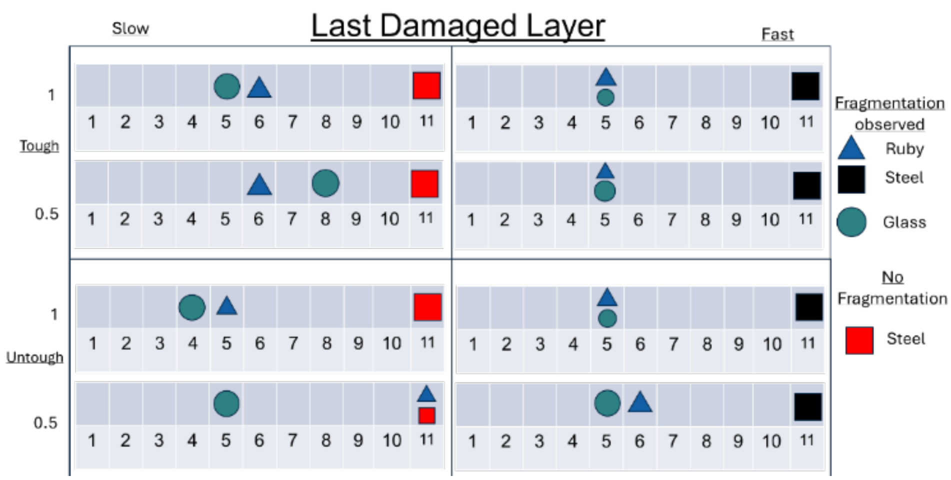

Figure 2-3: Results of test series #1 showing the last layer damaged in each case. The shapes of the points indicate the impactor type, and the rows indicate the membrane properties (toughened or un-toughened and 0.5mil or 1mil thickness). The left side shows shots at 4 km/s and the right at 7 km/s. The toughest impactor type (steel) penetrated all 10 layers in every case and didn't fragment at all for the slower collisions.

Our initial physics-based model of the behavior of impactors moving through multiple layers of material[1] predicted the particle shattering into smaller pieces with each collision. These pieces continue to propagate and fragment until they reach the ballistic limit and are stopped by the final layer. The current data suggests that another important factor at high speeds is the particle being fully ionized the collision shock. The result of this is that faster particles appear to be less damaging than slower particles. We're in the process of updating our model to account for this effect.

Test series #2 was completed in February 2026 using the same set of four membrane types. For this test series, we focused on the steel impactors and did a range of different velocities (2, 3, 5, and 6 km/s) to establish the threshold where fragmentation starts to occur. Our initial analysis of these results shows that the steel impactor penetrated all 10 layers for all tested shots. Fragmentation begins in the middle of the stack because previous layers do not deliver enough energy to fragment the particle but do add to its defect density. The defect density increases until a final shock begins the catastrophic fragmentation process. The starting layer of (catastrophic) fragmentation is a function of impact velocity and layer material properties. However, impact speed is the stronger contributor to this function; layer toughness of the materials spans an order of magnitude and has minimal impact on where fragmentation starts.

We are currently planning test series #3 for July 2026. The results of all three tests will inform the design of sub-scale shields representing flight-like designs for the HWO barrel. These will be tested in our follow-on study under the STABLE-S program.

3. TELESCOPE TEMPERATURE CONTROL (STABLE)

To evaluate our thermal model predictions of Primary Mirror (PM) temperature stability provided by a multi-layer barrel, we conducted a subscale barrel test. The objective of the test was to show that a 5-layer, passive, SLI barrel would meet the HWO goals of keeping the stability of the PM at milli-Kelvin levels at all observation angles to the Sun. In the test, we assessed the temperature stability of the PM Simulator (ΔT_{PM}) in responses to changes in the barrel outer-layer temperature, which simulate the changes in Sun angle in flight. The test utilized a 1/24th scale square barrel configuration (Figure 3-1), with external thermal conditions informed by flight model analyses. The temperatures imposed on the outside of the barrel were matched to predicted temperatures on the outside of the barrel in flight conditions with the barrel protected by a simple shade that's open to the sides[2].

The test was conducted in August 2025 with Sun Simulator balance points of ~ -123 °C (with no power to the Sun Simulator heater), -93 °C, -63 °C, and -33 °C. The measured PM Simulator temperatures were between 17.689 °C and 17.741 °C.

In the process of interpreting the measured temperatures, we realized that the temperature of the outer layer of the barrel (layer 5) was very sensitive to details of the test chamber. In particular, we added details to the model to represent areas in the chamber of bare stainless steel that was near room temperature and with much higher emissivity than assumed. This improved the match between predictions and measurements of the outer barrel temperatures.

The predicted temperatures of the inner barrel layers (layer 1) depend strongly on the details of the model. We created a detailed model (Figure 3-2) with significantly more nodes on the barrel, which enabled us to directly model the spacers between layers. The simplified model with fewer nodes uses an average conductance over the entire barrel area to represent the spacers. The comparison of the two models shows that this average conductance must be much higher than the expected conductance of the individual spacers to produce the same overall heat flow through the barrel.

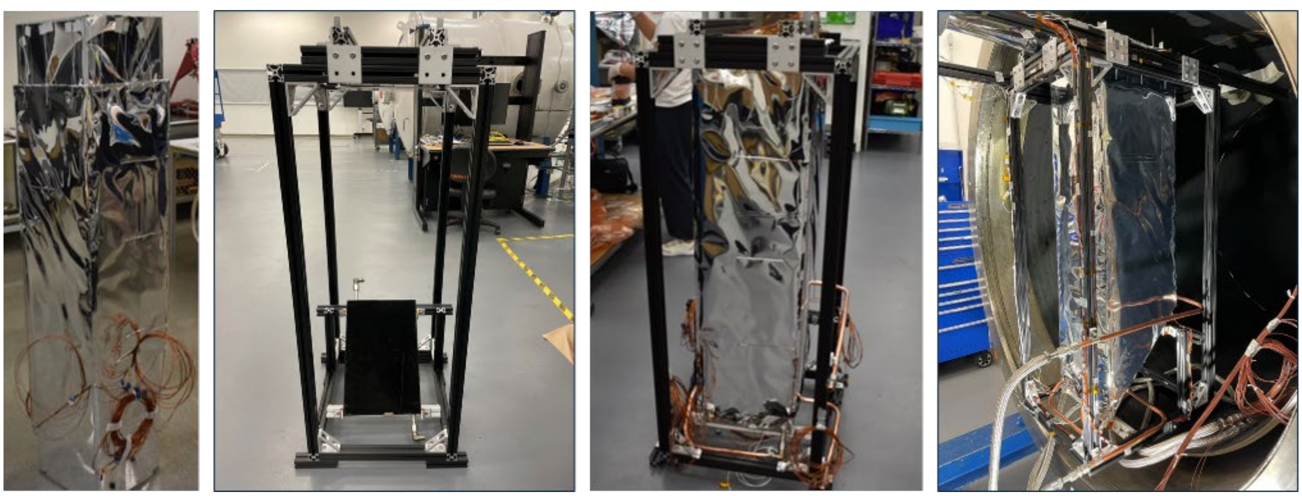

Figure 3-1. The subscale barrel test article being built (left), the test frame including the Deep Space Simulator (cold plate), the barrel installed into the frame, and being loaded into the thermal vacuum chamber (right).

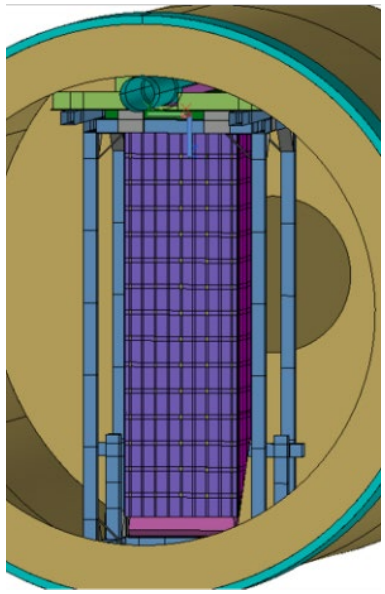

Figure 3-2. Detailed thermal model of the barrel showing the extra nodes added so that the individual spacers between layers could be included in the model.

At each balance point, the PM simulator temperatures were examined and the most stable ~12-hour period was extracted. Though there is an offset of about 150 mK between the two PM Simulator temperature measurements (because of uncalibrated equipment), the two sensors show the same temperature profile (Figure 3-3). We used the average of the data from the two sensors over the balance point period to determine the temperature of the PM. We looked at the delta temperature between balance points to eliminate the uncertainty in absolute temperature. The measured ΔT_{PM} between balance points was ~15-20 milliKelvin. Using model parameters fit to the test setup (chamber properties, barrel layers, etc.) predictions match well to the measurements within the limits of the thermal drift (which makes defining each balance point tricky). The test showed that the PM temperature variation was significantly reduced compared to the imposed temperature variation on the outside of the 5-layer barrel.

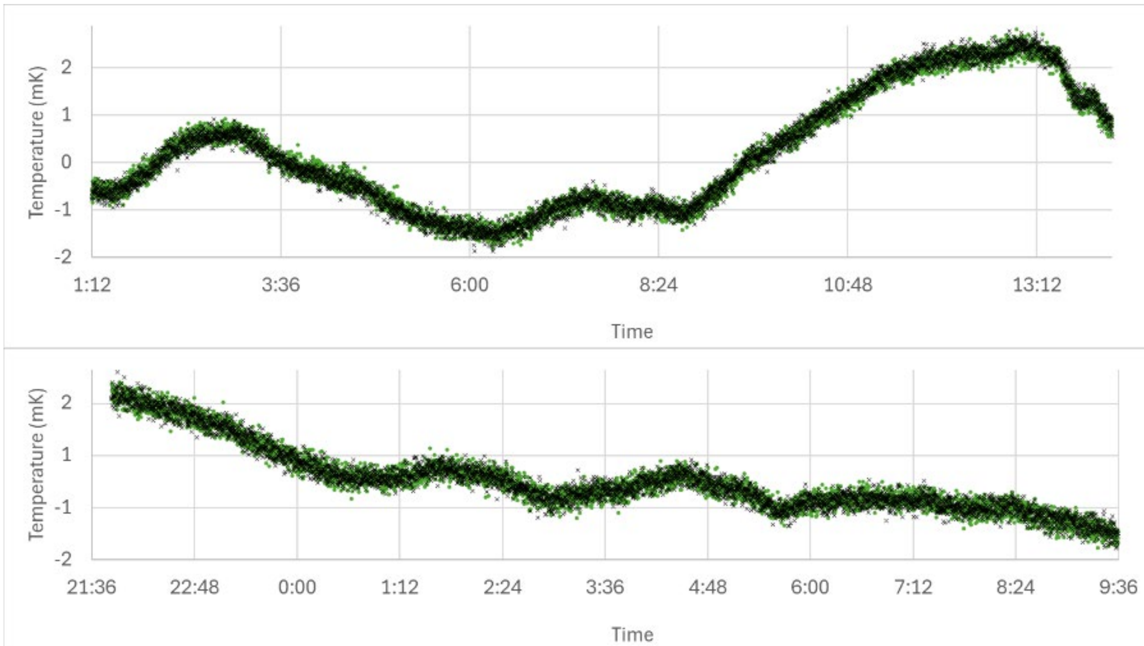

Figure 3-3. Examples of the measured data from the PM simulator (from balance points 1&2). The two sensors on the PM simulator are shown, with their average values over the ~12-hour period subtracted. The profile measured by the two sensors agrees extremely well and temperature drifts over this period of time are ~4mK.

4. ULTRA-STABLE TELESCOPE STRUCTURE MATERIALS & MODELING (STABLE AND STABLE-S)

The focus of this study area is to develop basic tools and testing approaches to design and validate a picometer level of stability in a telescope structure [3]. We have multiple approaches to address this challenge under both the STABLE and STABLE-S studies, and we report here on our progress and plans in a few of these areas.

Damping Foil Testing

We have performed a series of structural damping studies under the ULTRA, ULTRA-TM, and STABLE programs, aimed at increasing damping in a complex structure using 3M damping foil. Early coupon tests attached damping foil to composite laminates (quasi-isotropic and biased) to characterize performance at very low strain levels. Damping ratio increased with foil coverage up to about 50% of the coupon length, after which additional foil provided little benefit. These tests showed that the damping tape can be effective at low strain and across the 20–250 Hz frequency band. An initial sub-element damping test used a latch plane article originally designed for thermal distortion, with a first natural frequency around 306 Hz and damping foil wrapped around round tubes. Several issues limited the effectiveness of this test. For example, the damping tape was located in regions with less than 30% of the strain energy, so it was poorly positioned to dissipate vibration. The structure included bolted joints whose frictional behavior may have dominated or masked the damping tape effects.

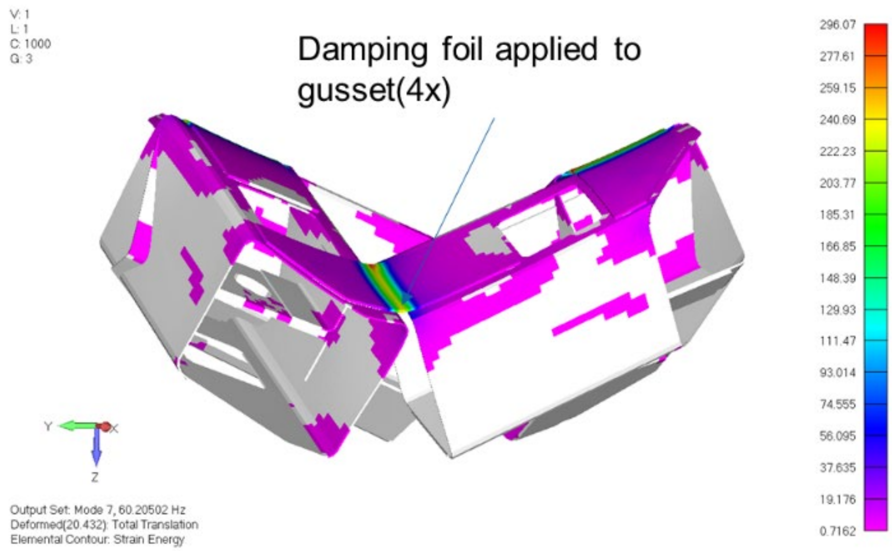

Figure 4-1. The damping test article was modified from previous test to reduce stiffness and eliminate the bolted joint. The figure shows the predicted strain energy in the test and the highest-strain area where the damping foil was applied.

To address these shortcomings, we defined and tested a new test article under STABLE (Figure 4-1). The new article uses half of the original latch-plane structure, retaining composite and metallic components and multiple bonded joints, but eliminating the bolted joint so that damping is driven mainly by the base materials and adhesive (aluminum, graphite-epoxy, and EA9394). Gussets on one side were cut between supports to reduce stiffness and bring the first natural frequency into the 20–250 Hz target band (prediction ~60 Hz). Damping tape was applied over full gussets in regions with high strain energy, improving strain energy participation from below 30% to about 85%.

Finite element analysis was used to predict mode shapes, strain energy distributions, and to choose accelerometer locations. Displacement predictions guided placement of accelerometers on the top or bottom surfaces where good modal response is expected. Analysis indicates a significant increase in damping for the first mode, and potentially for higher modes, when the foil is applied in high-strain regions. The same tap-testing method was used as in the previous tests to measure the actual damping ratio provided by the tape.

The damping test was completed in April 2026 (Figure 4-2) and the test results match pre-test predictions. The predicted increase in damping ratio was 1.5x and the test results showed an actual increase of 1.6x. This new latchplane damping test demonstrates that the damping properties of the tape continue to increase damping performance in a sub-element test article.

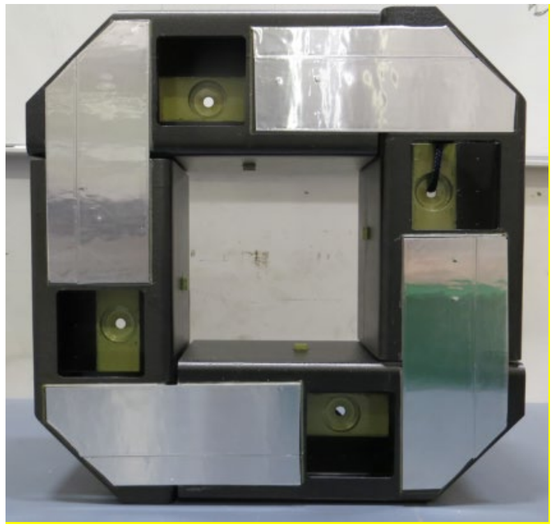

Figure 4-2. The damping test article used in the STABLE sub-element test. The silver areas are the locations where the damping foil was applied.

Picometer Stability Analysis and Testing

We previously reported on our analysis of picometer level structural modeling[4]. In that analysis we identified regimes where the model-predicted displacements are linear. Recent analysis has shown that known limitations of structural models will be even more critical to account for when using them to predict picometer displacements. For example, it has long been understood that the free edges of a composite structure will impact the CTE measurement. We have been successful in correcting for this effect, both through changes to the mirror-tube interface and through analytical approaches, at least to the accuracy required for prior programs. A significant finding from this study is that the axial CTE for the material in a rectangular composite tube does not match the axial CTE in the material itself, even when edge effects are removed using this approach (Figure 4-2). Exploring this effect in the model led us to the theory that this is a Poisson ratio effect due to the mismatch in CTEs in the hoop and through-thickness directions of the structure. This finding means that CTE measurements of such structures will need to be corrected before being used in a finite element model used to predict small motions.

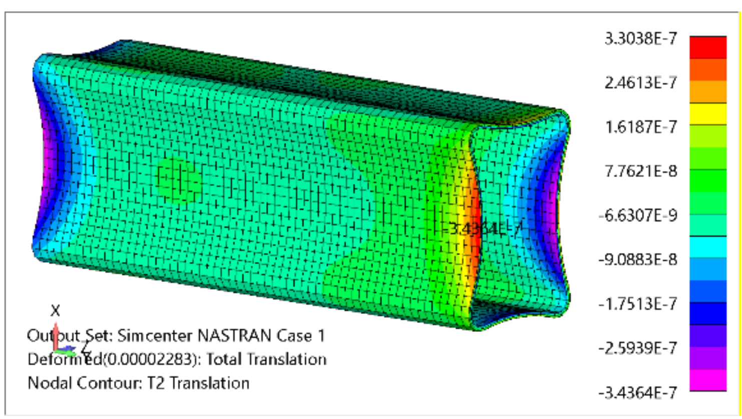

Figure 4-3. Modeling the effective CTE of a composite tube results in values that do not match the material CTE that was the input to the analysis. The CTE values vary along the length of the tube with the largest deviations near the ends.

To further explore the predictability of picometer level displacements, we have designed a coupon-level test to fulfill two purposes: verify that the structure behaves linearly over small temperature ranges and verify that the analytical model bounds the measured deformations. The test article is made of low CTE material (Invar) combined with EA9394 adhesive, with the intent of simulating a bonded joint that is likely to be used on HWO. The joint in this case (Figure 4-3) is based on the design used for the primary mirror interfaces on the JWST backplane (a “napkin” fitting made of invar bonded to the base structure on 2 sides, with the bond perpendicular to the interface surface). This test article will be placed in the pico-chamber at the GSFC Ultra Stable Lab and the tilt of the top of each fitting will be measured over a series of temperature steps. The results will then be compared to the predictions.

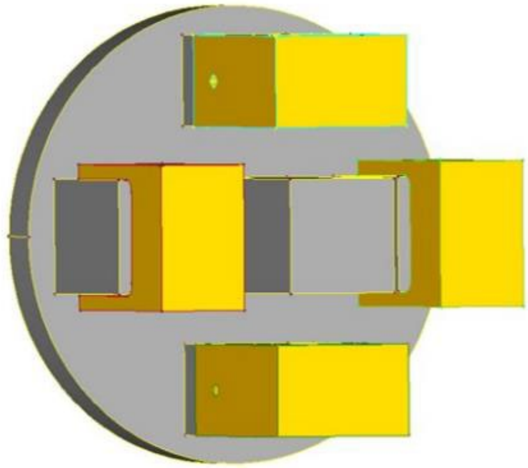

Figure 4-4. Drawing of the test article for our picometer-level stability test. The napkin fittings (yellow) are bonded to the base, and two of the four fittings have un-even bond thickness on the two sides. This leads to a predicted tilt of the fittings over temperature changes that will be compared to the measurements.

Lurching and Micro-Instability Study

This effort has focused on identifying any effect that could cause deformations in the HWO backplane not covered by the typical linear thermo-elastic response and quantifying whether they would exceed the capability of HWO control systems (unpredictable motion exceeding a few picometers over a timeframe less than a few seconds). One basis for this concern is measured jumps in wavefront error (WFE) seen on JWST[5] that is attributed to strain relief in the structure and mechanisms. Though the strain due to cooldown seen on JWST won't be present for HWO, the level of displacements that HWO is sensitive to are ~1000 times smaller. We're using the WFE data and other telemetry from JWST to gain insight into these types of phenomena.

During this study, we have compiled sources of strain in the system and examined each to estimate their potential impact on HWO performance in two ways: the estimated maximum displacement and the time duration over which these effects will last from their initiation. The characteristics of each motion are also examined for the possibility of the control system pre-emptively controlling for the effect. This approach is used to determine the potential impact, whether or not the timing of that impact can be predicted. Effects where past research has not been sufficient to quantify the impact are identified for further study.

For example, during operation the telescope structure is expected to stay thermally stable to within a few milliKelvin, but even this small temperature change causes some change in the stresses in invar fittings. A change in stress in invar fittings means that they will creep to relieve those stresses. A rough order of magnitude calculation shows that the creep strain could be in the range of 0.021 psi per milli-kelvin. A linear displacement due to this level of strain would be ~10 picometers. However, data to verify strain rates at these low levels has not been found and testing will need to be done to determine actual performance impacts.

Once we have identified areas of concern, we will plan a set of tests under our follow-on STABLE-S study to get data needed to better quantify these effects. The outcome of this work will enable us to better predict the micro-dynamics of the HWO structure and design the system in a way to mitigate problematic effects.

Telescope Structure Optimization

Large, stable structures to support segmented mirror systems must be demonstrated to meet HWO's stability needs while also being affordable to design and manufacture. On our STABLE-S study, we will develop structural design concepts using AI and optimization, build examples of these precision structures using JWST-legacy and additive-manufacturing techniques, and test their stability at picometer levels in GSFC's Ultra Stable Lab. Such designs will improve stiffness and stability while balancing other considerations such as manufacturability, CTE control, and mass efficiency.

Structural Latches Study

HWO will have latches between structural elements containing optics that require picometer stability. Micro-dynamically stable latches will be needed for deployable features (e.g., primary mirror wings and secondary mirror support structure) and for serviceable instruments. Latches have not been demonstrated to be stable to this level when subjected to nominal loading induced by thermal changes or other low-level operational loads. This study builds on the work we have done on STABLE, where we demonstrated some of the properties of an "HWO-class latch". We will design and build an HWO telescope latch configuration, evolved from the JWST latches. Unlike JWST, HWO does not have a severe cooldown environment that constrained the latch designs. We will demonstrate their micro-dynamic stability to picometer levels via testing at GSFC's Ultra Stable Lab.

5. IMPROVEMENT OF THERMAL EXPANSION TEST CAPABILITY (STABLE)

Our goal in this part of the study is to improve our capabilities in CTE measurement accuracy to ~0.1 ppb over a 2K ΔT, which is about two orders of magnitude better than the legacy JWST capability. There are two areas of improvement we've targeted: interferometer electronics and testbed environmental stability.

The new FPGA-based heterodyne interferometer reader has been implemented and is now the baseline for all interferometer data acquisition. The phase calculation is based entirely on statistical counts. We achieve nearly single-digit picometer resolution with the double-pass Distance Measuring Interferometer (DMI). This setup also provides lower read noise and elimination of the 180° deadband error. Sampling rate dithering mitigates aliasing issues with the heterodyne beat frequency.

A new vacuum chamber has been designed and built and is in the process of preliminary checkout testing. This setup was designed to have single millikelvin control capability. It also has a rigid connection from the DMI to the specimen support and positive-force specimen constraint to minimize mechanical disturbances.

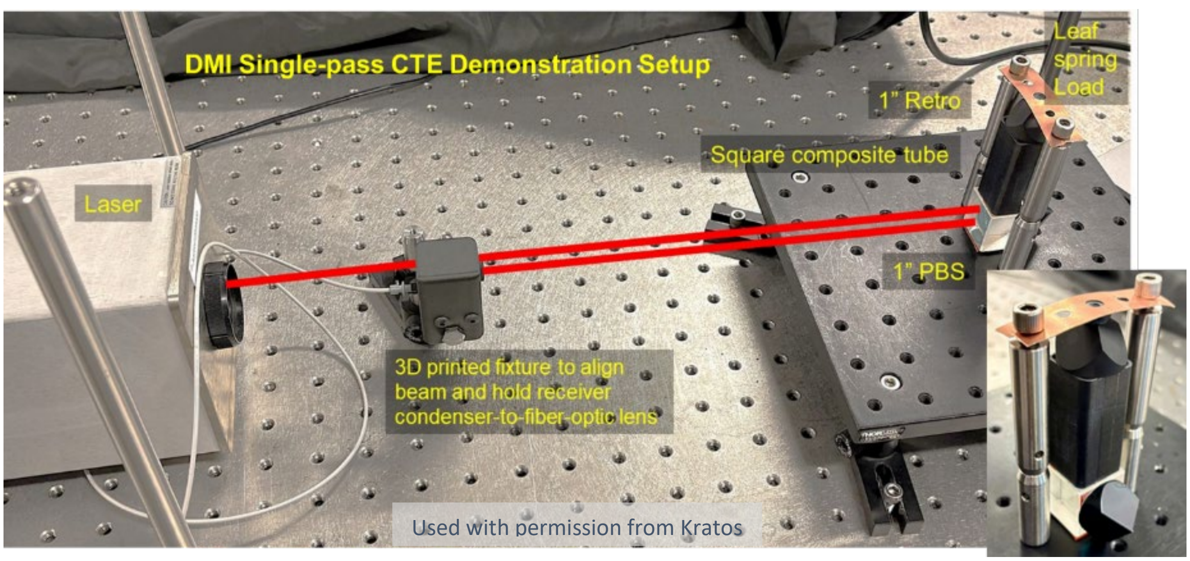

Figure 5-1..The improved testbed setup includes a leaf spring to stabilize the sample in the interferometer.

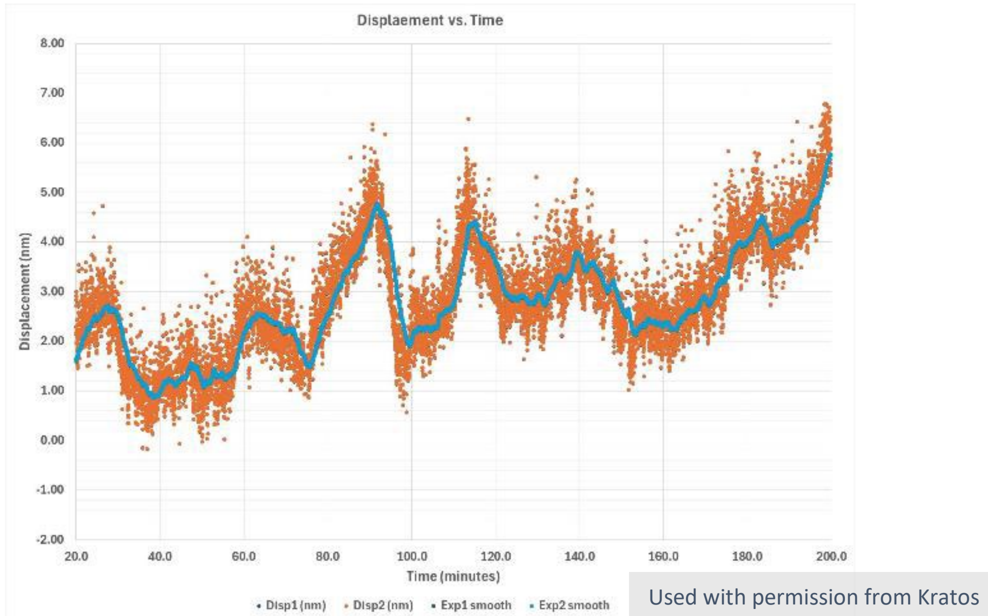

Figure 5-2.Initial measurements of a BK7 PBS, BK7 retro, and a 2.5" long composite with improved electronics shows a reduction in measurement noise. A 3-hour window of data is shown with ~5nm of stability achieved. Exponential smoothing is used to dampen the scatter in the data (blue line). These data were taken with the sample in air, so the larger variations are most likely due to temperature changes.

6. STRAY LIGHT CONTROL (STABLE-S)

Stray light control at the level required for HWO and on the scale of HWO has not been done before. This study will take a holistic approach and attempt to identify all potential sources of stray light and develop a process to assess and allocate values to each. There are many contributors to stray light, some stable and some a function of environment and thus variable over time. Stray light in general is a difficult effect to accurately model – the high contrast ratio required for HWO magnifies this problem.

Mission architecture decisions greatly impact the overall stray light levels achieved: optical design, structure, and barrel architecture are all critical. Different instruments will also have different sensitivities to stray light. We will develop an understanding of relevant weighting functions for each instrument type. Sources of stray light that impact all instruments will be the first priority to assess.

Our first step will be to compile a list of contributors to the stray light error budget and initial estimates for each term. Current stray light budgets for HWO are not detailed enough to identify the key areas where there will be challenges in achieving mission goals. We recognize that some contributors may have no data to develop an initial estimate, but we will include placeholders. We will develop a concept of how the individual contributors roll up into overall stray light performance, along with how stray light contributes to instrument performance and how these answers depend on the observatory and instrument designs.

To probe the effects of optical surface figure, edges, and coatings on the stray light, our team will build mirror samples and measure their scattering properties. Mirrors are being built by L3Harris and NG AOX and scattered light measurements are being performed by our NG Azusa site in collaboration with The Scatter Works. In coordination with the STABLE-S architecture study, we will evaluate if there are other materials that need to be tested for their stray light contribution or if available data is sufficient. With these data, we will update our error budget to reflect the current state of the art.

7. MULTI-DISCIPLINARY ARCHITECTURE STUDY (STABLE-S)

As part of our STABLE-S study, we will study a wide range of architectures for the HWO mission. We will take a multi-disciplinary approach to the study to minimize the number of unconsidered aspects of the design. This will help us ensure that the final architecture meets all mission requirements and constraints.

To achieve our goals on this study, we'll take a wide and shallow approach to begin with. We'll consider a wide design space and evaluate the options against a broad range of criteria. These evaluations will start with flexible, independent models and tools and focus on criteria that provide the opportunity to quickly narrow the option space. Through this process, we will eliminate areas of design space that are not acceptable.

The next step will be to take the narrowed parameter space and start looking at coupling between parameters. We'll add the next level of detail to the modeling to increase fidelity and continue to down-select options to narrow design space.

Finally, we will select a few architectures to continue to a deeper level of evaluation. This will be our initial integrated look at complete, self-consistent design options and allow us to compare architecture performance at mission level.

More details about our approach to this architecture study can be found in a paper by Arenberg et al. in this proceeding.

REFERENCES

[1] Michaela N. Villarreal, Jonathan. W. Arenberg, and Lauren Halvonik Harris, "Predictive model for micrometeoroid damage to gossamer structures," Journal of Astronomical Telescopes, Instruments, and Systems 8(4) 044006 (2022)

[2] Andre Aroyan, "Environmental shielding and Thermal stability of primary mirror for Habitable Worlds Observatory," Proc. SPIE 13623, UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts XII, 13623-9 (2025).

[3] Charles B. Atkinson, Todd Lawton, Robert Hellekson "Structures considerations for Habitable Worlds Observatory (HWO) telescope", Proc. SPIE PC13623, UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts XII, PC1362304 (18 Sept 2025)

[4] Tiffany M. Glassman, Jonathan W. Arenberg, Todd Lawton, et al. "Progress on STABLE: development of technologies for ultra-stable systems", Proc. SPIE 13623, UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts XII, 1362303 (18 Sept 2025)

[5] Randal C. Telfer, Marcio Meléndez, Nicolas Flagey, et al. "Empirical characterization of JWST wavefront error variations", Proc. SPIE 13092, Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave, 1309210 (23 Aug 2024)