LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-T2600489-v1** 9/25/2026

Flattening and Construction of the GQuEST End Mirrors

Matthew C. Winkowski[1]
Mentors: Daniel Grass[2], Lee McCuller[2]
[1] *The College of William & Mary,* [2] *California Institute of Technology*

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW17-161
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

1 Introduction

There exists a challenging incompatibility between general relativity and the Standard Model of particle physics. In an attempt to bridge that gap, the study of quantum gravity has emerged to propose theoretical solutions to rectify these differences. A variety of theories have been proposed regarding the nature of quantum gravity, but they are difficult to verify because of the scales at which gravitational interactions occur. One explanation for quantum gravity, proposed by Erik Verlinde and Kathryn Zurek [1], suggests that due to the holographic principle, spacetime can be quantized into Planck-length units of space called pixellons [2]. These individual quantizations of space experience stochastic excitations called geontropic fluctuations.

GQuEST (Gravity from the Quantum Entanglement of Space-Time) is a collaboration involving Fermilab, JPL, and the McCuller Lab at Caltech that aims to measure or constrain these fluctuations. This will be done using a seven-meter-long tabletop Michelson interferometer, shown in Fig. 1.

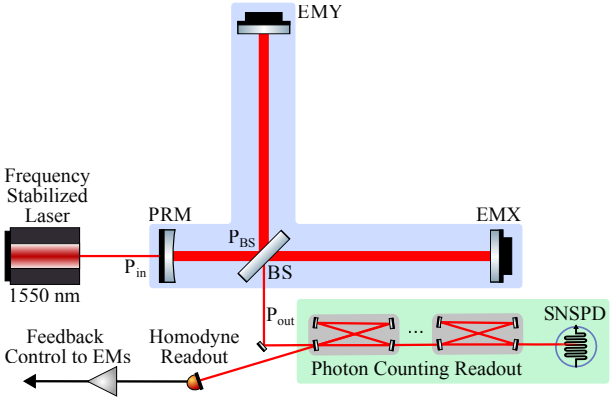


Figure 1: The proposed design for the McCuller Lab Michelson interferometer [3].

Because the proposed quantum gravitational fluctuations are extremely weak, it becomes very challenging to get strong statistical results, such as those of a 5σ significance test. One of the key features of GQuEST is its photon counting method, which filters the output signal of the interferometer through a series of cavities, such that only photons carrying a signal reach the detectors. By heavily filtering the signal, it is possible to reduce the shot noise background, and achieve significant results in a much shorter time span.

The mirrors in this experiment are composed of a bulk silicon substrate with a LIGO-style Bragg coating made up of alternating layers of Ta_2O_5-SiO_2. One of the main sources of noise in this experiment is the mechanical ("Brownian") noise, which has components proportional to both the thickness of the coating and the substrate. Following Grass et al. [4], the coating contribution to the mechanical noise is given by

$$S_{z,\mathrm{CMN\ min}}(\omega) = \frac{4k_\mathrm{B}Th_\mathrm{c}\phi_\mathrm{c}}{\pi M_\mathrm{c} w^2 \omega} \tag{1}$$

Here, k_{B} is the Boltzmann constant, h_{coating} is the coating thickness (10 μm), ϕ_{coating} is the average coating loss angle ($3{\cdot}10^{-4}$), M_{coating} is the average coating P-wave modulus (99 GPa), w is the 2σ beam radius in power (3 mm), and ω is the angular frequency, with $\omega = 2\pi f$ and $f \approx 17$ MHz. The coating thickness is determined by the laser wavelength being used, so disregarding temperature, the coating thickness can be thought of as a constant value. The second source of mechanical noise in the mirror is from the substrate where the coating is applied. The substrate contribution is similarly given by [4]

$$S_{z,\mathrm{SMN\ min}}(\omega) = \frac{2k_{\mathrm{B}}Th\phi}{\pi M w^2 \omega}, \tag{2}$$

where h is the substrate thickness (2 mm), ϕ is the substrate loss angle ($1 \cdot 10^{-6}$), and M is the substrate P-wave modulus (165 GPa). Because the coating noise has a fixed value, in order to minimize noise the component of noise proportional to the substrate needs to be brought below the component proportional to the coating. By using a 2 mm thick silicon mirror, the magnitude of the substrate noise falls below that of the coating.

But using a thinner mirror does not come without tradeoffs. Because the mirror substrate is thinner than is typical (6 mm), they are more prone to deformation under coating stress. This stress is given by the Stoney Equation:

$$\frac{1}{R} = \frac{6(1-\nu_s)\sigma_c h_c}{E_s h_s^2}, \tag{3}$$

where R is the radius of curvature, σ_c is the coating stress, h_c is the coating thickness, h_s is the substrate thickness, E_s is the substrate Young's modulus, and ν_s is the substrate Poisson ratio. The estimated radius of curvature for the unflattened GQuEST mirrors is: $r_{\mathrm{curv}} \approx 7.6\,\mathrm{m}$, or ≈ 0.26 diopters [3]. For a spherical mirror, the radius of curvature is $R = 2f$, and the corresponding optical power is $1/f$ diopters when f is expressed in meters. The initial approach to compensate for this curvature is the application of a thick anti-reflective coating to the rear of the mirrors. This coating reduces the error in diopters by a factor of 10. Even with a coating applied to the rear of the mirrors, there is still expected to be curvature. Past work developing mirror flattening techniques by Rafael Volkamer-Pastor [5], has shown that the GQuEST mirror mount can successfully be used to minimize induced curvature on test mirrors. These mirrors were uncoated, so the only stresses they were subject to were those of the mount, but they were shown to be effectively flattened. Additionally, by pairing interference patterns of flattened mirrors, the losses due to interference can be reduced even further. Ideally, this results in a final mismatch between interferometer end mirrors of ≈ 0.003 diopters.

The key modes that are focused on in the flattening process are the quadratic modes of the Zernike polynomials, shown in the third row of Fig. 2. These can be fitted and calculated for a given mirror surface, allowing for an estimation of the shape of deformation caused by coating stress. Higher-order Zernike teRMS, such as the ones shown, would ideally be minimal. Losses that they cause can be minimized through matching mirrors or an output mode cleaner, but they are not addressable through use of the mount itself.

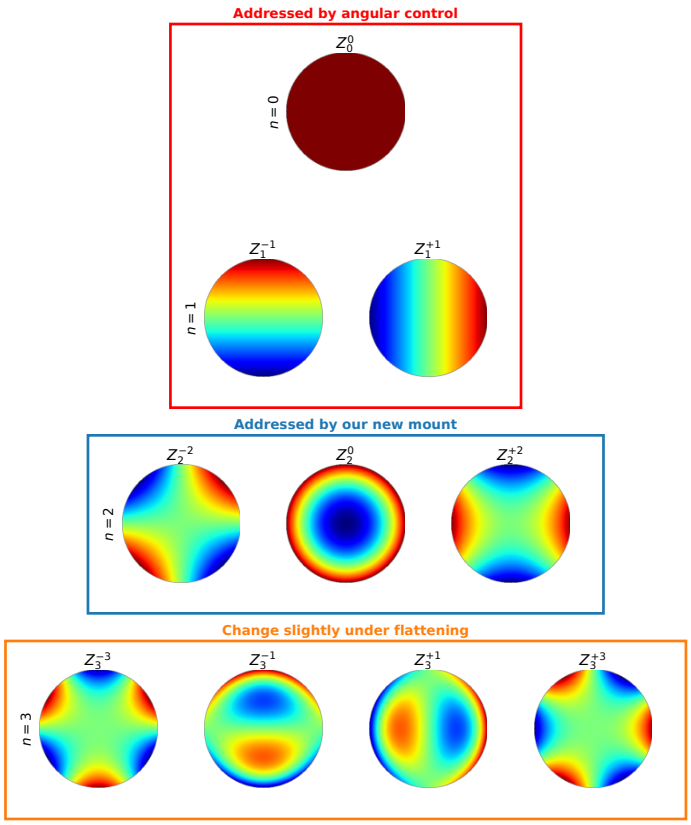


Figure 2: A visualization of Zernike polynomial modes. The Z_{00} mode is the piston, with higher-order teRMS increasing in complexity towards the bottom of the pyramid. The mirror mount can only address up to second-order modes. Losses due to higher-order teRMS will be addressed first by mirror matching and use of an output mode cleaner if necessary.

2 Objective

The main goal of this project was the assembly and testing of the GQuEST end mirrors. This included using the true experimental optics in the actual mounts, as well as implementing the control electronics to match. This consisted of a few key steps, beginning with profiling mirrors in their entirety and under a variety of stresses, such as clamped or completely flattened. After profiling and selecting mirrors, the next part of the process was flattening them. After flattening and profiling the first mirror again, I followed up by assembling a second mount and flattening that mirror. This assembly and flattening process gives an idea not just of how effective the mirror flattening process is, but how repeatable it is with multiple mounts, as four will be necessary for the final experiment.

After completing the mirror flattening process, the next step was to assemble the kinematic components. These parts are necessary to develop an active angular control system to account for minor shifts in end mirror position, and maintain end mirror alignment. The angular control effectively maintains the $Z_{1\pm1}$, modes, or the second row of Fig. 2. The angular control components I used are going to be used on the final GQuEST interferometer,

but they are also a key step in the test setup that I created. The first step towards getting kinematic control working was to assemble the mount itself, by attaching the piezos into the main mount and installing the picomotors. Once those components were installed onto one mirror, the mount could be installed into a test interferometer for angular control testing. This test setup was the desired conclusion of the project, with it being a way to demonstrate control of the end mirrors over a small scale.

The test setup should act as a stand-in for the full GQuEST interferometer, but with only one actively controlled end mirror. In this system, the error will be measured by a quadrant photodetector (QPD), which is connected to the CyMAC, the LIGO experiment's real-time digital signal processing and handling system, via an analog-to-digital converter (ADC). The ADC takes in a voltage reading from the QPD, which is then fed into the CyMAC. On the CyMAC, the software will process the input signal through control loops with the goal of maintaining a constant DC voltage reading at the dark port.

Once the feedback loops have been developed, the computer will output voltage signals corresponding to necessary adjustments to minimize the error via a digital-to-analog converter (DAC). The DAC voltage will need to be amplified to voltages as high as 200 V, to allow for the full range of the piezoelectric actuators. The system will then continually monitor the dark port, and make adjustments in the feedback loop as necessary.

3 Methods

3.1 Mirror Analysis

The tool used to assess mirror flatness was a 4D NanoCam HD surface profiler. With the 2X Michelson objective, it uses instantaneous phase-shifting interferometry and has a specified RMS precision below 0.1 nm for 32-measurement averages. One key issue to address before any scans can be taken is the inherent astigmatism that the NanoCam has. This figure makes it impossible to know if the mirror is truly flat, because at low RMS roughness values, the astigmatism becomes dominant over any other figure. To address this, a reference needs to be calculated and then removed from the figure. This was initially done with a rotational method [5], where a reference flat is rotated 90°, and then anything other than the inherent astigmatism can be removed through a matrix transformation. This produced the original reference flat, shown in Fig. 3.

But this method of reference generation ran into two main issues. First, due to the rotation this method can only generate a square reference, losing information on the edges of the image. Second, due to a limited number of scans, occlusions and imperfections from the reference flat, which are clearly visible in the NanoCam reference, can appear as artifacts in later scans. To address both of these issues, I decided that after having taken many scans with the initial reference, I had a large data set that could be used in an average. By averaging scans, I was able to resolve both the occlusions and edge information losses. The

averages would make outlier pixels in the dust specks irrelevant, and then allow for area outside of the central 3.6 mm-wide square region to be usable. The new reference, shown in Fig. 4, shows improvement of both issues and the smoothness throughout.

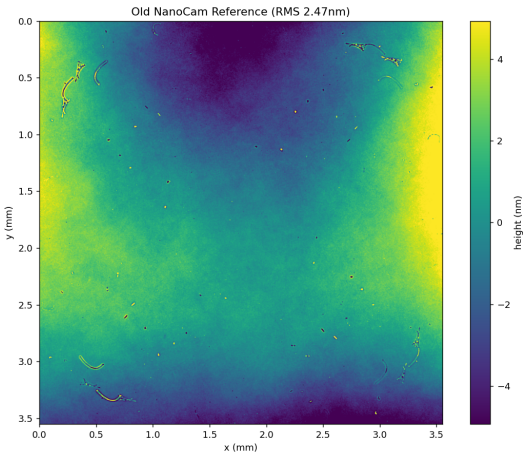


Figure 3: The initial corrected astigmatism for the camera, developed with the rotational method discussed by Volkamer-Pastor [5].

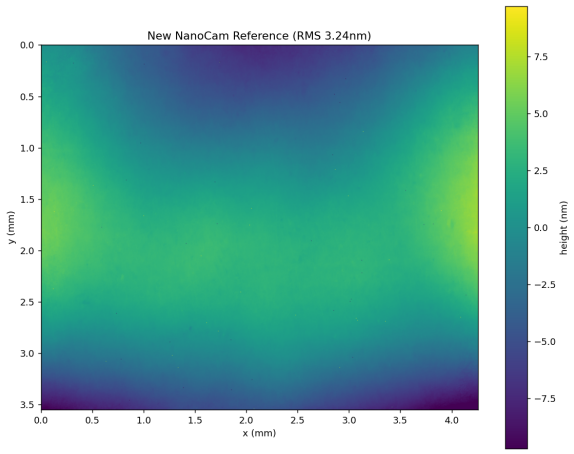


Figure 4: The improved astigmatism correction based on the averaging of ~ 80 scans. The averaging clearly improves the smoothness of the reference and makes the regions outside of the central square from the first reference usable.

One of the key components of this project was improving the mirror-analysis pipeline, including raster scanning, stitching, and Hermite-Gaussian (HG_{nm}) mode analysis. To begin, I improved the centering of each scan by aligning the mirror using the spokes on either end. This required a physical shift of the NanoCam, as well as a re-centering of the mount using the translation stage. Proper alignment allowed larger (5×5) raster scans to cover nearly the entire mirror. In some scans, the corners were slightly clipped by the edge of the measurement region, but this rarely affected the fit and when necessary the affected corner tiles could be removed.

Additionally, centering the mirror properly gave me more room to work with on a typical raster scan. This allowed for greater overlap between tiles, which was a major contributing factor to the improvement of the scans. The stitching was done through a least-squares fit over a constrained boundary between tiles. So a pair of tiles with no overlap between them could not be reconstructed. By increasing the overlap from 200 μm to 500 μm, it allowed for a more tightly constrained and accurate reconstruction of the global mirror map.

With regards to the actual analysis, most of it was done through coupling calculations with the HG_{00} using Eq. 4. This equation is an overlap integral between our incoming and outgoing beam profile. The incoming beam profile is a Gaussian laser input, and the outgoing profile scatters off of the surface into higher-order modes (HOMs).

$$c_{00n'm'} = \iint u^*_{n'm'} e^{2ik\tilde{z}(x,y)} u_{00}(x,y)\, dx\, dy \tag{4}$$

Here, $c_{00n'm'}$ is the coupling coefficient for the n,m mode referred to from here on as $c_{n'm'}$; k is the wavenumber $(2\pi/\lambda)$, (x, y) is the position of each point; u_{00} is the incoming ideal Gaussian mode; and $u_{n'm'}$ is the outgoing mode. The spatial profiles of the Hermite-Gaussian modes u_{nm} are given by:

$$u_{nm}(x,y) = \left(\frac{1}{2^{n+m-1} n! m! \pi}\right)^{1/2} \frac{1}{w_0} H_n\left(\frac{\sqrt{2}x}{w_0}\right) H_m\left(\frac{\sqrt{2}y}{w_0}\right) \exp\left[-\frac{x^2+y^2}{w_0^2}\right]. \tag{5}$$

where n and m refer to the Hermite-Gaussian modes HG_{nm}, H is the respective Hermite Polynomial of order n or m, w_0 is the beam width, x and y are the transverse positions of the beam at the mirror surface.

Reducing the scattering into HOMs in Equation 4 is the main objective of the mirror flattening. This scattering, especially differential scattering (scattering into different modes on either end mirror), needs to be reduced for two key reasons. First, to maintain the darkness at the dark port. Unwanted bright interference can cause noise and increase the uncertainty at signal detection. Scattering into any HOMs also reduces the effectiveness of the power-recycling mirror, making it more challenging to achieve the desired circulating cavity power. For both reasons, it is ideal that the mirror is as flat as possible.

I used a few other calculations, such as a beam-weighted RMS, to provide quick insight into the flatness of the mirror in the relevant region. If instead I had used a total RMS, values would often become contaminated with large edge defects, that do not affect the behavior of the beam, such as major lobes near mounting points. Additionally, beam weighting the RMS almost entirely removes the effect that clipped corners might have on a measurement.

Once the beam weighting was implemented, I also added functionality to test other beam centers. This acted as a point of comparison to tell whether another point on the mirror aside from the center could result in less scattering. Often on my testing of Mirror 2, as shown later in Fig. 9, the raised area in the upper section of the mirror caused the minimal-RMS-center to be shifted downward. Another addition to the analysis pipeline was using a loss calculation using Eq. 6, which follows from the total-integrated-scatter formula for scattering from surface roughness [6]. Here R_{RMS} is the RMS roughness, and k is the wavenumber.

$$\mathrm{Loss} = 4k^2 R^2_{\mathrm{RMS}} \tag{6}$$

All of these equations could give quick insight into the behavior of a global mirror map. The loss and coupling coefficients in particular are a good predictor of final behavior in an actual experiment.

3.2 Mirror Flattening

Before putting any mirror into a mount, I wanted to take simple unmounted scans of the mirrors. Those scans could give a base idea of the mirror's curvature. And by fitting the curvature equation and removing certain modes, I was able to get an idea of the estimated surface roughness, which should match the optimal flatness.

$$\tilde{z}(x, y) = a + b(x - x_0)^2 + c(x - x_0)(y - y_0) + d(y - y_0)^2 + ... \tag{7}$$

Here, the height $\tilde{z}$ is given by a polynomial surface fit, where the series of coefficients $(a, b, c ...)$ give the curvature map of the scan, centered around x_0 and y_0. Although the measured height map contains all the available information about the surface, by fitting it with a smooth polynomial we can isolate and get a good description of large scale deformations. By removing these polynomial modes to second order, we are able to get an estimated flattened height map of the mirror. Examples of these scans, with and without the quadratic modes digitally removed, can be found in Fig. 5 and Fig. 6. The digitally flattened scans are idealized and likely to be unrepresentative of the end mirror outcomes, but they can provide an idea of each mirror's surface roughness, which can help in the mirror selection process.

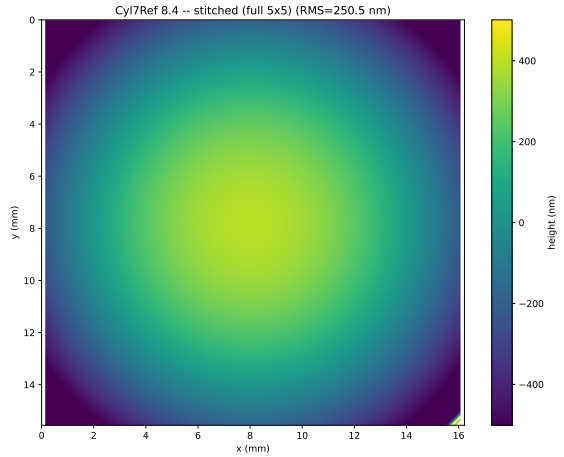


Figure 5: A 5×5 scan of a mirror outside of the mount. Its heavy quadratic bowl curvature can clearly be seen dominating the overall shape of the figure.

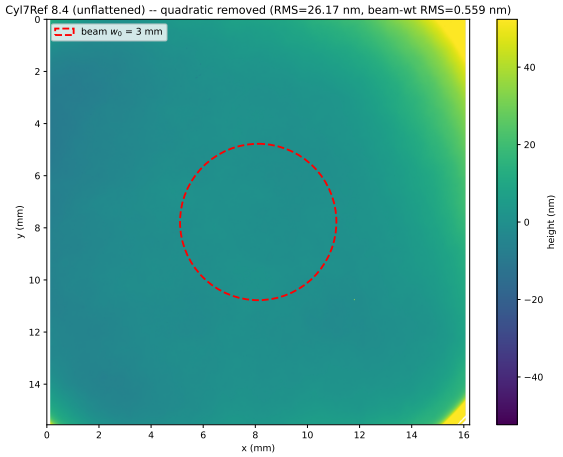


Figure 6: Here, by removing the quadratic figure of the mirror, which is dominant outside of the mount, we were able to get an estimate of the surface roughness. It is not possible to get the mirror this flat in practice.

As outlined earlier, the mirrors have an innate stress due to the coating. Additionally, once they are inserted into the mount itself, there is another stress induced by clamping the mirrors in place. These stresses combined to give a base loss before any flattening of $640,768$ ppm. This is much higher than either the losses before or after flattening, and can be avoided by clamping with less force. Clamping with less force does not appear to have major advantages, as I was still able to flatten the mirror to an acceptable level.

The actual process of flattening followed the explanations provided by Volkamer-Pastor [5]. I began by applying tension to each of the screws until they contacted on the outer "T", after which I would gradually tighten or loosen the screws to cancel out the modes as detailed. An effective heuristic for flattening was to swap back and forth between O-mode and astigmatism. Applying tension in one mode would swap it into the other. Ideally, it would be best to turn each of the screws evenly, flattening consistently through O-mode, but this proved to be very challenging. Additionally, it was also sometimes simpler to just over-tension slightly, and then get the mirror into a flattened state by reducing the tension in the screws, undoing the modes one by one.

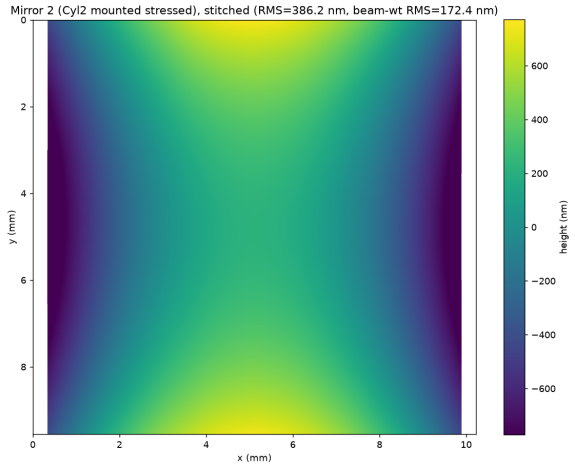


Figure 7: Mirror 2, when placed under clamping stress. This was the state in which it was scattering the most out of the HG_{00} mode, and serves as a point of comparison going forward.

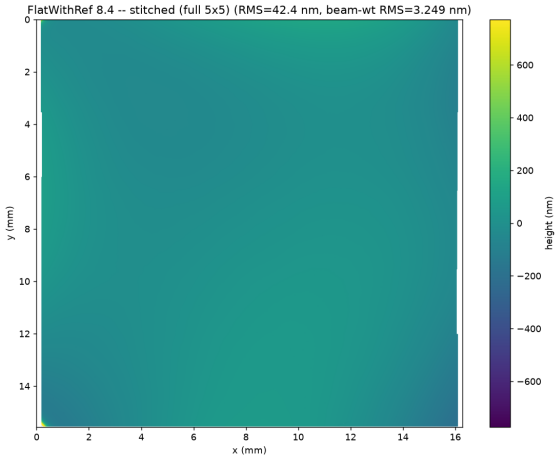


Figure 8: Mirror 2, after flattening, scaled to the same height map as when it was placed under clamping stress. Any global figure is nearly invisible, and it shows the effectiveness of the flattening process.

The flattened mirror, shown in Fig. 8 and Fig. 9, has a clear higher-order structure that cannot be addressed through the standard flattening procedure. The losses from HOMs are generally concentrated in the HG_{30} and HG_{03} modes, indicating that flattening asymmetry is the main contributor to the non-addressable losses.

3.3 Kinematic Mount

A large part of the work done on this project was the assembly and installation of parts to prepare the mirror mounts for use in the experiment. This can be broken up into three key stages. First, the assembly of the mounts and preparation of parts. Second, the installation and assembly of the kinematic mount parts. Finally, the testing and development of the kinematic controls. Each of these is a step along the way towards creating usable kinematic mounts, both for a test interferometer and the final GQuEST experiment.

Because the entire mount is eventually going into vacuum, the parts must be properly cleaned

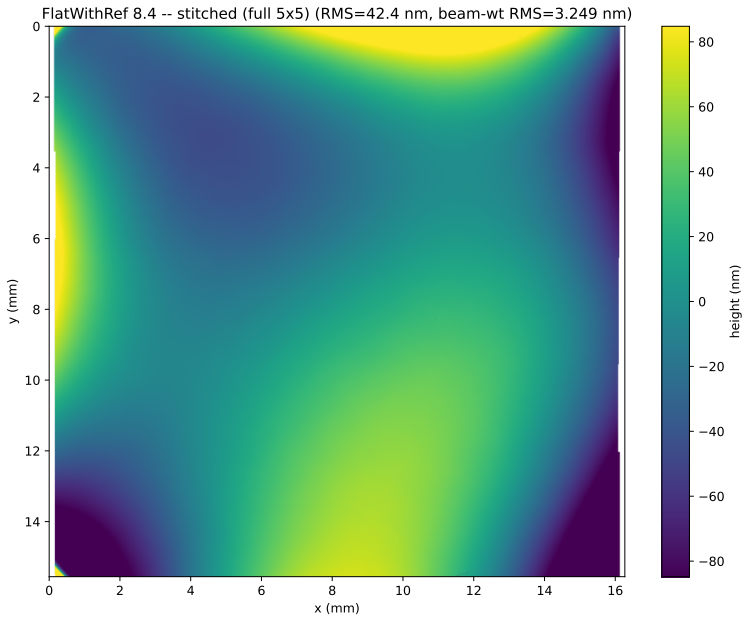


Figure 9: The first flattened mirror, scaled appropriately. The global asymmetric figure is clearly visible, and the lobes near the top and left side are as well. These lobes, while a major deformation of the mirror, have minimal impact on the beam itself.

and baked out before they are put into either a clean room environment or vacuum system. Some of the machined parts may still have grease on them from the machining process, so it is essential to clean off these parts to the best of our ability. The parts that have been prepared are: new bushings, both for the existing mount and a second mount; stainless steel clamps for the active kinematic control system of both the first and second mount; and the main aluminum body and holders for the second mount.

The cleaning is a three-stage process. The parts are cleaned with a solution and then placed into a sonicator for each stage. The first stage is a mix of deionized water and a degreasing agent, with dilution varying depending on how heavy the grease or grime on the parts is. The next step is just deionized water, followed by the final stage, a rinse and bath in isopropanol. After cleaning, the parts are baked in a vacuum oven for temperature and time depending on their material.

After the parts have been prepared, the next step was to assemble them into their respective mounts. One mount already had its base parts assembled and only needed to have the kinematic components installed, the second mount for the other interferometer arm was a fresh assembly. The second mount assembly in progress can be seen in Fig. 10, where the top of the clamps and the flattening screws have yet to be attached.

The other assembly, of the kinematic mount, was a more involved process composed of attaching the piezos to the mount as well as installing the picomotors into the kinematic portion of the mount. The piezos (Noliac NAC2023-H08-UHV-A) were attached to the mount in a two-part process. First, the piezos are attached to the mount with a few drops of EPO-TEK 302-3M Epoxy. They are then clamped into place, as shown in Fig. 11. This is

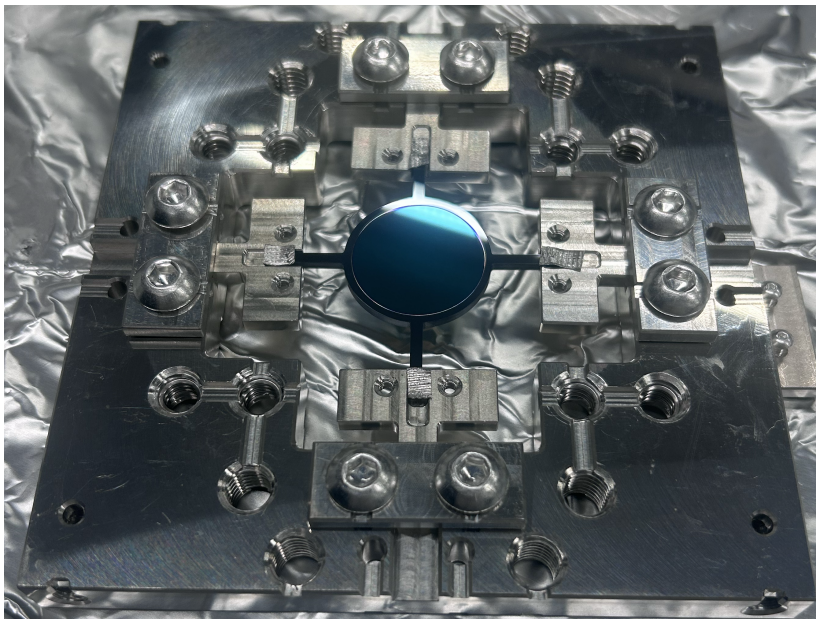


Figure 10: The second mount assembly, with a mirror put into place. The flattening components will be placed on the top and bottom of this central piece of aluminum.

done in all four corners of the mount, and after the epoxy cured we attached the top clamp, letting it cure again afterwards. The top was clamped down by the clamping components of the mount itself, ensuring that what we put together was the final product. The piezos can linearly shift ≈ 10 μm, allowing for actuation of the mount face over a range of ≈ 100 μrad. They are a key part of the angular control setups, and will be used to maintain differential arm length (DARM) and angular control on the final experiment.

After inserting the piezos into the mount, the next step was installation of the picomotors (Newport 8301-V) into the kinematic mount. This was basically a swap, removing existing screws and clips and replacing them with the picomotors. The picomotors are installed into an enlarged version of a typical kinematic mount, allowing for precise two-axis tip/tilt control.

The electronics for the picomotors and piezos are controlled through separate means. The piezos are controlled directly by voltage applied through a CyMAC control loop. The CyMAC DAC outputs ± 10 V, which is then amplified by a set of EVAL-ADHV4702-1 amplifier boards, shown in Fig. 13 to a range of $0-200$ V. The picomotors are controlled over serial by a separate picomotor controller, a New Focus 8752 Picomotor Controller, shown in Fig. 14. I showed that both of these worked to control their respective motors.

4 Results

The key outcome for this project was to demonstrate that the GQuEST end mirror mount could flatten the curvature of the coated mirrors. In effect, it was to show that the mount

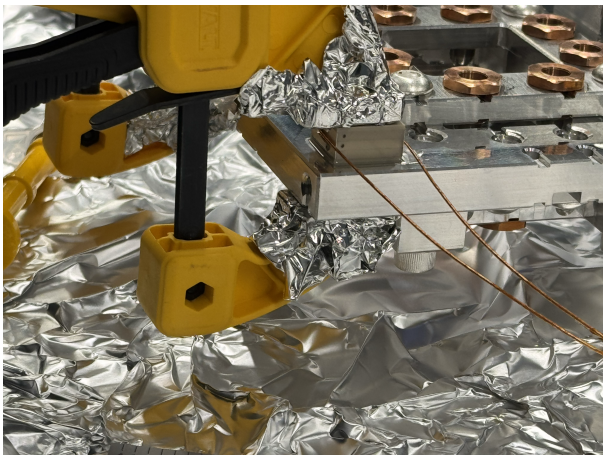


Figure 11: The piezo clamped into place on top of the base mirror mount. The holders were then clamped into place on top of the motors.

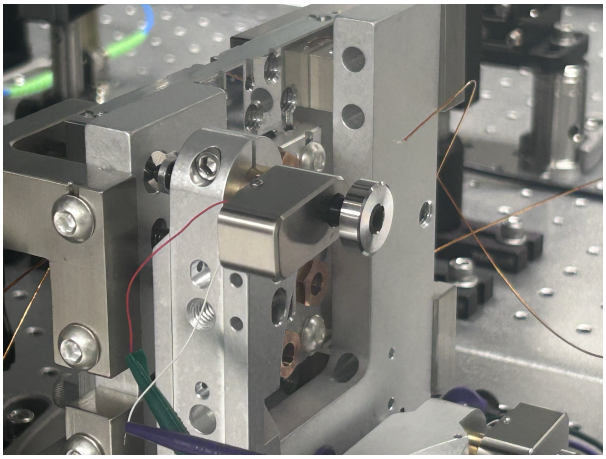


Figure 12: One of the picomotors inserted into the mount. They are placed into opposite corners of the standard kinematic mount, allowing for control over the tip and tilt.

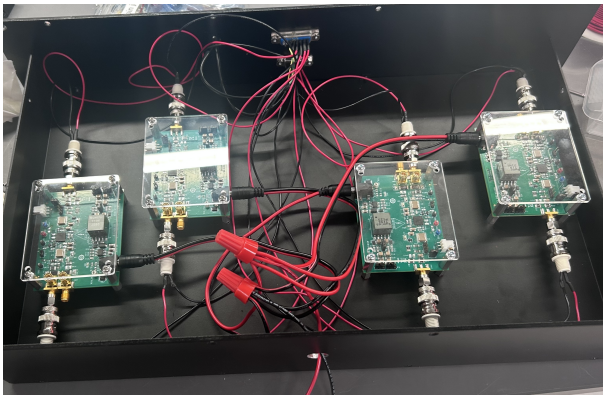


Figure 13: The electronics box for the piezo amplifiers. Each of the circuit boards corresponds to one piezo, amplifying the voltage to the appropriate range.

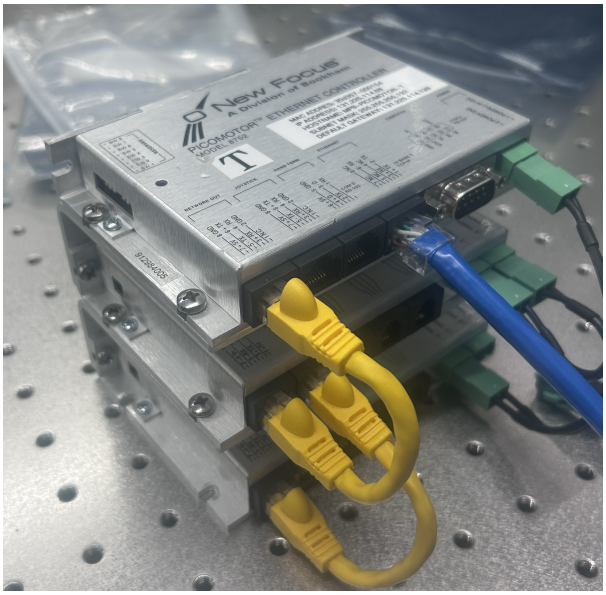


Figure 14: The picomotor controller. Controlled over the serial port and treated as a USB device to handle coarse alignment.

could reduce scattering from the HG_{00} into HOMs. The final scattering comparison for the two end mirrors I selected is shown in Table 1. Overall, this demonstrates a successfully flattened mirror. Reducing the ppm loss to the order of hundreds was the goal, and the quadratic modes and bowl are almost entirely eliminated. Hundreds of thousands of ppm loss is not viable for an experiment that hopes to get 10 kW circulating power in the arm cavity, so this is a good step in the right direction.

These results can be improved further. By matching the surface figures of the two end

mirrors, the differential mode content between the interferometer aRMS can be reduced. Although this does not reduce the scattering loss of either mirror individually, common spatial defects interfere destructively at the dark port. By matching mirrors with similar height maps, and therefore similar mode content, the differential scattering at the interferometer output can be reduced below the individual scattering loss of either mirror.

PPM Couplings

	Cyl2		Cyl7	
Scan	Unflattened	Flattened	Unflattened	Flattened
RMS/TIS Loss Estimate	**142,729**	**693.8**	**126,891**	**1,576**
HG_{00}	874,856	999,309	887,379	998,442
HG_{02}	55,499	3.25	49,055	563.2
HG_{20}	54,038	61.5	50,909	39.1
HG_{03}	0.37	170.2	0.26	606.3
HG_{30}	0.21	247.8	1.75	84.1
HG_{04}	5,160	25.8	4,045	10.5
HG_{40}	5,225	3.44	4,358	3.02

Table 1: Cavity mirror mode content before and after flattening. HG_{nm} coupling is evaluated at $w_0 = 3\,\text{mm}$ and $\lambda = 1550\,\text{nm}$. The flattening is effective at reducing the quadratic (HG_{20} and HG_{02}) modes, but can induce the HG_{03} and HG_{30} modes. Full coupling tables can be viewed in Appendix A.

Take for example, the two mirrors used in this experiment. A pair of difference maps between the two flattened surfaces are shown in Fig. 15. Both orientations produce a differential surface figure that can be evaluated using the same beam-weighted RMS and HG-coupling calculations used for the individual mirrors, which generally improves for the differential maps. But the more interesting result comes from changing the relative mirror orientation. By digitally rotating one of the scans, we can simulate other mirror orientations. In this case, a 90° rotation improves the beam-weighted RMS substantially, and differential loss. By matching up the "lobes" or prominent surface features, their contribution to differential surface figure is reduced, minimizing their individual effects on scattering.

This raises an important question for future work: whether the favorable orientation is determined primarily by the intrinsic mirror figure or by deformation introduced during mounting and flattening. A 90° rotation cannot be done on the GQuEST mount without removing and re-flattening the mirror. If the mounts lead to formation of an inherent figure, then that could cause issues because they can only be rotated 180°. But if the mirrors can be rotated to match their figures, it should lead to a good reduction in losses for the experiment.

There are a few constraints in the current project that make the mirror flattening process challenging. First, there is the large beam waist relative to the diameter of the mirror.

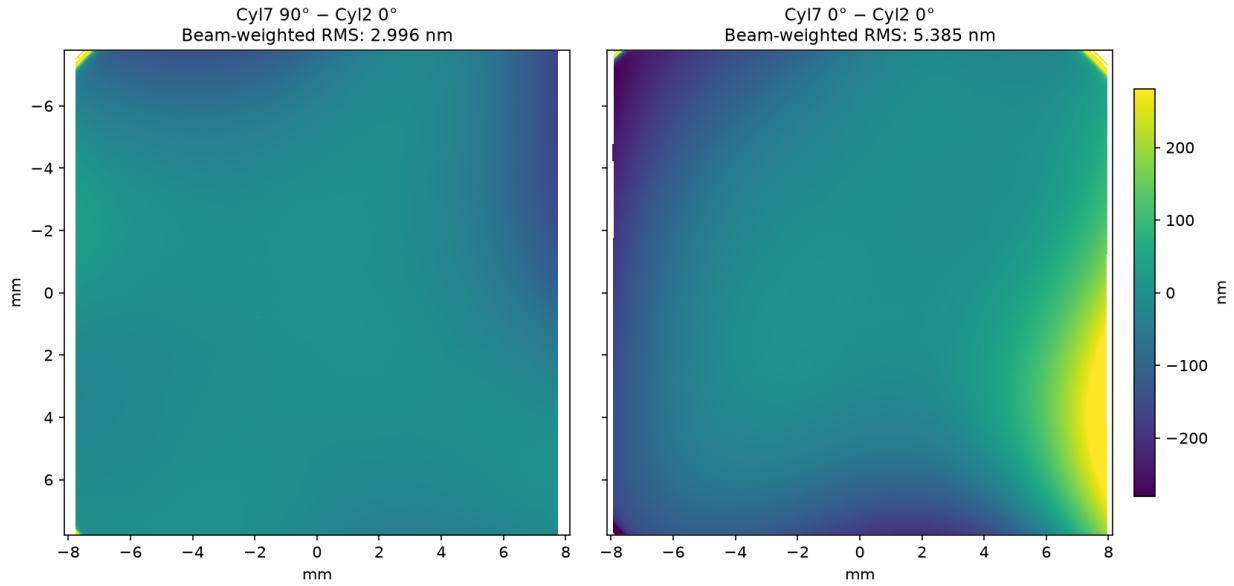


Figure 15: Two interference maps with different rotational orientations. Demonstrates that mirror orientation makes a difference in total scattering loss, and that certain mirror rotations can reduce overall scattering.

GQuEST will use a 3 mm beam waist on a 25.4 mm diameter mirror, which is rather large and leads to larger losses from the outer sections of the beam, which tend to be where the mirror is less flat. For instance, if the beam waist was 1 mm, the loss would be 15.3 ppm (unflattened 1898 ppm), which is over an order of magnitude reduction in loss when compared to the 3 mm beam waist. The majority of those losses are due to surface roughness, which is scattering caused by minor defects in the coating or substrate itself. For our mirrors in particular, a key example of surface roughness would be the lines from polishing, which are visible on a NanoCam scan once the mirror is flattened. The technique addresses the curvature of the mirror at the critical central area, but due to random small effects building up, the large spatial deformations get enhanced near the spokes and edges of the mirror.

There are a few other things to address with regards to the mirror outcomes. First, the mount does effectively reduce the quadratic, and even some higher-order couplings (HG_{04}/HG_{40}). This was the main goal of the project, and the results match closely to the uncoated test mirrors used by Volkamer-Pastor [5]. The main discrepancy is the growth of the asymmetric modes, namely HG_{03}, HG_{30}, and HG_{21}. These modes appear due to flattening in an asymmetric manner, and are not easily addressable by the mount. To get more specific with the process, these modes grow when the screws used to flatten the mirror are used unevenly.

For example, say that the mirror is in astigmatism (Z_2^2 as shown in Fig. 2). It should be flattened by tightening the left and right screws evenly. But if one screw is tightened slightly more or less than the other, it can lead to asymmetry in the mirror, or the growth of the higher-order Zernike modes. This imbalance can accumulate over time, sometimes leading to large losses, such as in Cyl7. These modes are not addressable through normal use of the mount, and will require another attempt from rest to address. It is challenging to quantitatively assess how low one could possibly reduce the magnitude of these modes

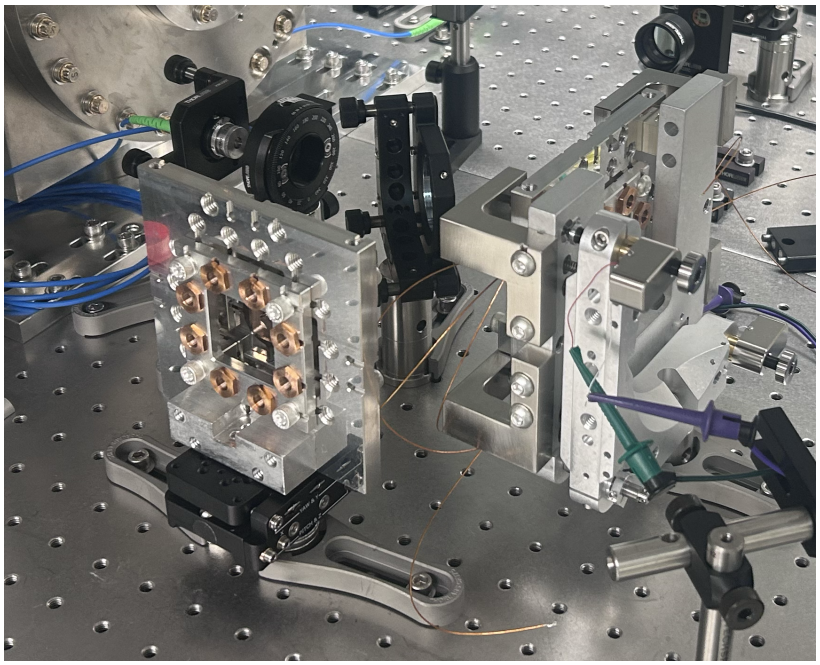


Figure 16: The test interferometer assembly. Rough alignment is handled with picomotors, and then angular and DARM locking should be done with the piezos. The feedback for the system was based on a QPD, seen in the upper right.

without further repetitions, but it is at least worth flattening Cyl7 again. But even under less initial stress, it was shown that the uncoated mirrors were only able to be reduced to a level of 312 ppm loss [5].

The final section of work that I did this summer was the assembly of a test interferometer. This was one of the reasons for the assembly of the second mount, and allowed for a demonstration of the kinematic mount. Ideally, I would have implemented a basic control loop for DARM locking, but I ran into time constraints at the end of the summer and was unable to complete it. But, I was able to successfully demonstrate the actuation and control of both motor types.

The test interferometer, shown in Fig. 16, was a small-scale $\sim$ 100 mm interferometer. It was implemented with one kinematic mirror and one fixed mirror, and all of the optics used, including the beam splitter, were the ones set to be used in the final GQuEST experiment. Both picomotors functioned well, and were used to align onto the center of the QPD. The piezos functioned well individually and together, and were actuated over multiple fringes. With some minor troubleshooting of the CyMAC data acquisition system, it is possible to make this system work, and to DARM lock using the piezos.

5 Future Work

The mirror flattening methods have now been shown to work successfully on the GQuEST end mirrors. Going forward, there are still a few improvements and developments to be made before these end mirror mounts can be implemented into the final setup. The first of these will be improving the flattening of the second mirror. After the second mirror has been brought into the desired end range ($\text{Loss}_{\text{ppm}} \approx 650$), the kinematic parts, including the piezos and picomotors, will need to be attached to the second mount.

When both of the mounts are assembled, the next step will be to implement control loops and continue their development. I attempted to do some basic controls work, but was stalled by technical issues. Continuing to iterate and demonstrate these loops will be a key step towards the final completion of GQuEST as a whole. There are also possible far-future improvements such as inclusion of the mirror correction into feedback loops, or enlarging the mirror and mount so that the edge-based scattering losses are less pronounced.

To conclude, I demonstrated that the GQuEST end mirror design was able to be successfully used to flatten thin-substrate silicon optics. The mirror scattering losses due to curvature were reduced from 142,729 ppm to 693.8 ppm. While not perfectly in line with initial goals, which were closer to 100 ppm, the mirror flattening proves challenging, and it will be difficult to begin to approach that goal using the current mount. Additionally, I assembled a full kinematic end mirror, and demonstrated that the picomotors work well for rough alignment, and that the piezos can be effectively controlled using the CyMAC. Overall, this work acts as another stepping stone towards the eventual completion of the GQuEST experiment.

Acknowledgments

Thank you to Caltech, LIGO Lab, and the NSF for supporting this research. Thank you to GariLynn Billingsley and Camille Makarem for letting me borrow the NanoCam 4D HD. Finally, a huge thank-you to the McCuller Lab for the support and guidance this summer, especially Daniel Grass and Lee McCuller.

References

[1]E. P. Verlinde and K. M. Zurek, "Observational signatures of quantum gravity in interferometers", Physics Letters B **822**, 136663 (2021).

[2]K. M. Zurek, "On vacuum fluctuations in quantum gravity and interferometer arm fluctuations", Physics Letters B **826**, 136910 (2022).

[3]S. M. Vermeulen et al., "Photon-counting interferometry to detect geontropic space-time fluctuations with gquest", Physical Review X **15**, 011034 (2025).

[4]D. Grass et al., "High-frequency thermal noise in michelson interferometers", arXiv preprint arXiv:2605.16710, 10.48550/arXiv.2605.16710 (2026).

[5]R. C. Volkamer-Pastor, *2025 LIGO SURF final report: mapping and correcting the surface of the GQuEST end mirrors*, Technical Note LIGO-T2500248-v1 (LIGO, Aug. 27, 2025).

[6]J. A. Ogilvy and H. M. Merklinger, "Theory of wave scattering from random rough surfaces", The Journal of the Acoustical Society of America **90**, 3382–3382 (1991).

A Complete Coupling Tables

Cyl2 HG coupling

$w_0 = 3\,\mathrm{mm}$ ($1/e^2$ intensity radius),
$\lambda = 1550\,\mathrm{nm}$

Mode	Unflattened Coupling (ppm)	Flattened Coupling (ppm)
HG_{00} retained	874,856.49	999,308.89
HG_{00} loss	125,143.51	691.11
HG_{02}	55,498.84	3.25
HG_{11}	8.71	0.88
HG_{20}	54,037.50	61.48
HG_{03}	0.37	170.18
HG_{12}	0.34	3.29
HG_{21}	0.18	131.06
HG_{30}	0.21	247.84
HG_{04}	5,159.84	25.80
HG_{13}	1.27	2.18
HG_{22}	3,432.13	14.42
HG_{31}	1.08	4.16
HG_{40}	5,224.56	3.44
HG_{06}	724.87	1.11
HG_{24}	313.75	0.04
HG_{42}	327.66	0.11
HG_{60}	410.32	0.50
HG_{26}	43.03	0.02
HG_{44}	29.19	0.02
HG_{62}	25.43	0.03
HG_{46}	4.12	0.00

Cyl7 HG coupling

$w_0 = 3\,\mathrm{mm}$ ($1/e^2$ intensity radius),
$\lambda = 1550\,\mathrm{nm}$

Mode	Unflattened Coupling (ppm)	Flattened Coupling (ppm)
HG_{00} retained	887,378.68	998,441.80
HG_{00} loss	112,621.32	1,558.20
HG_{02}	49,054.67	563.19
HG_{11}	8.17	27.66
HG_{20}	50,908.89	39.05
HG_{03}	0.26	606.30
HG_{12}	0.04	51.78
HG_{21}	0.80	26.39
HG_{30}	1.75	84.07
HG_{04}	4,045.16	10.55
HG_{13}	1.51	31.99
HG_{22}	2,809.46	12.71
HG_{31}	1.18	55.95
HG_{40}	4,357.52	3.02
HG_{06}	391.71	2.27
HG_{24}	243.33	0.06
HG_{42}	255.41	0.11
HG_{60}	418.00	0.43
HG_{26}	25.23	0.02
HG_{44}	23.61	0.05
HG_{62}	26.75	0.04
HG_{46}	2.30	0.01

Flattened mirror difference HG couplings

$w_0 = 3\,\mathrm{mm}$ ($1/e^2$ intensity radius), $\lambda = 1550\,\mathrm{nm}$

Mode	Cyl7 +90°	Cyl7 0°
HG_{00} retained	999,417.65	998,115.94
HG_{00} loss	582.35	1,884.06
HG_{02}	62.78	650.13
HG_{11}	19.10	39.83
HG_{20}	255.62	5.64
HG_{03}	15.15	138.09
HG_{12}	10.84	28.11
HG_{21}	18.52	267.66
HG_{30}	82.86	607.38
HG_{04}	11.67	24.52
HG_{13}	36.15	17.83
HG_{31}	14.97	32.79
HG_{40}	1.67	8.27
HG_{05}	1.19	2.73
HG_{14}	0.25	0.02
HG_{41}	0.32	2.23
HG_{50}	0.53	0.76
HG_{06}	1.25	1.57
HG_{15}	0.79	0.44
HG_{24}	0.18	0.22
HG_{42}	0.21	0.63
HG_{60}	0.27	3.10

Table 2: HG coupling from the full height difference of the flattened mirrors, with Cyl2 fixed at 0° and Cyl7 at the indicated orientation. At a 3 mm beam waist, the +90° difference has lower HG_{00} loss than either mirror alone; the 0° difference does not.