LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-CONF-400061

MEMS adaptive optics for the Gemini Planet Imager: control methods and validation

L. A. Poyneer, D. Dillon

December 21, 2007

MEMS Adaptive Optics II
San Jose, CA, United States
January 22, 2008 through January 24, 2008

Disclaimer

MEMS adaptive optics for the Gemini Planet Imager: control methods and validation

Lisa A. Poyneer[a] and Daren Dillon[b]

[a]Lawrence Livermore National Lab, 7000 East Ave, Livermore, CA, 94550
[b]UCO Lick Observatory, Laboratory for Adaptive Optics,
University of California, Santa Cruz,
1156 High Street, Santa Cruz, CA, 95064

ABSTRACT

The Gemini Planet Imager (GPI) Adaptive Optics system will use a high-order MEMS deformable mirror for phase compensation. The MEMS mirror will be used in a Woofer-Tweeter configuration, with a frequency-domain based splitting of the phase between the two mirrors. Precise wavefront control depends on the ability to command them MEMS to make the exact phase desired. Non-linearities in the MEMS may prevent this. We determine that influence-function pre-compensation can remove most, but not all, open-loop error. We use simulation and a simulation of a non-linear MEMS to address the issue of how much non-linearity can be tolerated in closed-loop by GPI.

Keywords: Adaptive Optics, wavefront control, MEMS mirror

1. INTRODUCTION

The Gemini Planet Imager (GPI)[1] is an instrument which is designed to directly image extra-solar planets. In order to achieve contrasts of greater than 10^6 between star and planet, GPI combines a high-performance Adaptive Optics (AO) system with an Apodized Pupil Lyot Coronagraph[2] to suppress starlight and a science-path interferometer to provide calibration information.[3]

The GPI AO system will operate at frames rates up to 2 kHz. In order to provide high-Strehl images, thousands of actuators are necessary for the phase correction. GPI's design has an actuator spacing of 18 cm in the pupil, which requires a mirror with at least 45 actuators across its diameter. For GPI the high-spatial frequency phase correction will be accomplished with a microelectricalmechanical system deformable mirror (MEMS mirror, for short). This choice of a MEMS mirror, as opposed to a conventional piezo-electric deformable mirror (DM), was initially driven by both actuator count and size. Though conventional mirrors will soon be available with thousands of actuators (e.g. the high-order mirror for SPHERE[4] which is being manufactured by Cilas) for GPI size is an even bigger driver. The GPI MEMS mirror, with an actuator spacing on the device of 400 μm, will correct a pupil of 1.76 cm, which at f/16 easily fits unfolded in the space-constrained GPI optical design.[5] By contrast, a conventional DM with actuator spacing of 1 cm would require a pupil and an optical path 25 times larger.

The MEMS mirror for GPI is still under development by Boston Micromachines (BMC). The final mirror will have 64×64 actuators and a total stroke of 4 μm. In the meantime, 32×32 devices from BMC have been rigorously used and tested at the University of California, Santa Cruz's Laboratory for Adaptive Optics. Evans et al. have shown that the MEMS can be controlled in closed-loop to sub-nm accuracy with direct phase measurements.[6] This can also be done[7] with a spatially filtered Shack-Hartmann wavefront sensor and the Fourier reconstructor. Morzinski et al. have shown that MEMS actuator position is highly repeatable and temporally stable.[8] Further challenges to using a MEMS mirror for GPI remain, and in this paper we discuss these problems and our progress towards solving them.

Even with 4 μm stroke, the GPI MEMS will not be able to fully compensate for high-order atmospheric turbulence. It simply does not have enough dynamic range to prevent occasional saturations under turbulence.

Send correspondence to Lisa Poyneer: poyneer1@llnl.gov, 1 925 423 3360

GPI therefore uses a "Woofer-Tweeter" strategy where two mirrors are used in tandem. The "Woofer" is a high-stroke, low-order conventional DM. The "Tweeter" is the MEMS. The two mirrors will control separate portions of frequency space, with the Woofer controlling the low-order Fourier modes, and the Tweeter the high-order modes. GPI's Woofer has been specified to ensure that the Tweeter rarely saturates.[9] A detailed discussion of GPI's Woofer-Tweeter strategy is in preparation by our colleagues Lavinge and Véran;[10] this paper will present a summary of the strategy and show end-to-end simulation results which verify its efficacy.

The second major challenge with the MEMS is making the surface of the mirror go exactly where the control system wants it to go. If the MEMS mirror was a linear, shift-invariant (LSI) system, its behavior could be completely compensated for with a filter. We apply such a filter to the problem of forming Kolmogorov-like shapes on the MEMS and conclude that this filtering pre-compensation reduces the open-loop shaping error by a factor of 6, to 10's of nm. This level or error may be too much for GPI, even in closed-loop. We use a simulation of a non-linear MEMS in our GPI end-to-end AO code to address how much non-linearity can be tolerated.

2. WOOFER-TWEETER CONTROL

Woofer-Tweeter control in general describes a system where two devices are used in tandem for wavefront correction. One device, the Woofer, controls aberrations of a low-frequency, large-amplitude nature. The second device, the Tweeter, controls high-frequency, small-amplitude aberrations. This designation can be used for temporal frequency, such as Véran and Herriot's tip-tilt control method as proposed for TMT's NFIRAOS.[11] This strategy is possible for tip-tilt because the theoretical models for atmospheric and windshake tip-tilt follow power laws with a negative exponent of temporal frequency. The same distribution of power holds for Kolmogorov turbulence (see, for example, the spatial frequency description of it used in phase-screen generation[12]) which means that most of the power is in the lower spatial frequencies. Conan et al. are studying spatial Woofer-Tweeter with a testbed.[13] Their setup is significantly different than ours, in that their mirrors have similar ratios of pupil diameter D to actuator spacing d of $D/d = 7$ for the Woofer and $D/d = 9$ for the Tweeter, while for GPI we have 8 for the Woofer and 44 for the Tweeter. They are also are not using a Fourier-mode wavefront control method.

For GPI we will be using wavefront reconstruction and control in the Fourier basis. In particular, we are using Fourier Transform Reconstruction (FTR).[14] FTR provides the computational efficiency that is necessary for GPI; using FTR instead of a matrix multiplication requires 45 times fewer FLOPs. A second major reason to use the Fourier modes is that they provide a useful basis set for sophisticated temporal control of the modes individually. GPI's baseline algorithm is Optimized-gain Fourier Control (OFC).[15] This is a form of modal gain optimization,[16] wherein closed-loop AO telemetry is used to estimate the optimal gain for each Fourier mode's integral controller. A further advancement beyond gain optimization is Predictive Fourier Control (PFC).[17] Under the assumption of frozen flow, each Fourier mode can be independently Kalman-filtered for improved atmospheric correction, while still meeting GPI's computational budget. Given these reasons, the Woofer-Tweeter strategy for GPI must work with the Fourier modes, and not require excessive computation.

These requirements led to a strategy of directly splitting the Fourier modes into two groups. Low spatial frequencies are sent to the Woofer. The remaining Fourier modes are kept for the Tweeter. This extraction is done directly in the Fourier domain, as is illustrated in Fig. 1. The slopes are reconstructed into residual phase, and then the temporal controller is applied. This is all done while still in the Fourier modal domain. Then specific complex-valued Fourier modal coefficients are extracted from the signal and arranged in the vector as real and imaginary parts of the low-order Fourier modes. This vector is multiplied by the Modes-to-Actuators (MTA) matrix. This MTA matrix is pre-determined; for each sine or cosine is calculates the actuator commands which will make that phase shape on the Woofer. As such, MTA includes information about the Woofer influence function.

Those low-order Fourier modes are zeroed-out in the Tweeter signal. Then the MEMS pre-compensation filter is applied and the inverse Fourier Transform is calculated. This produces the actuator commands for the Tweeter. The details of the strategy, including how to determine the appropriate spatial frequencies to offload, and if any feed-back correction is needed, will be discussed in the forthcoming paper by Lavinge and Véran.[10]

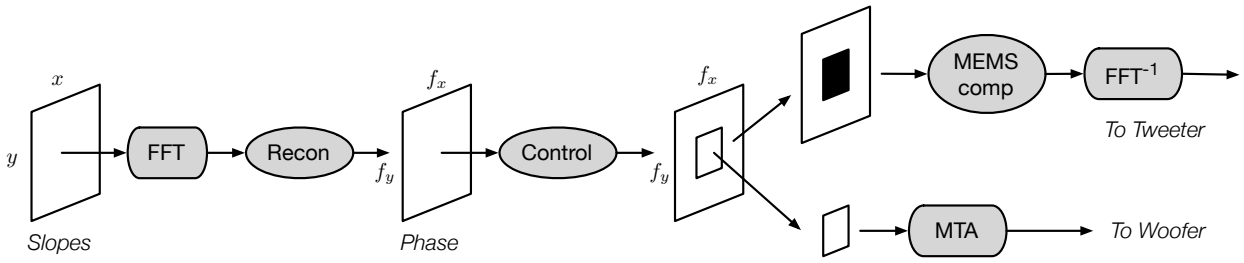

Figure 1. GPI's Woofer-Tweeter control strategy. Specific low-order Fourier modes are extracted in the frequency domain and are converted to Woofer commands via the Modes-to-Actuators (MTA) matrix. The remaining Fourier modes are pre-compensated for the MEMS and then converted into Tweeter commands with an inverse Fourier Transform.

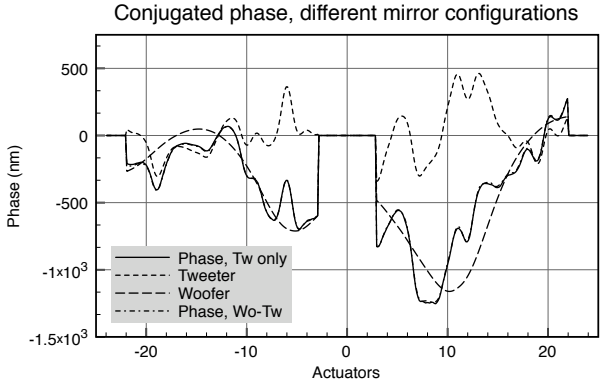

Figure 2. Phases conjugated in a noise-free simulation by either just the Tweeter, or the Woofer-Tweeter configuration. The spatial frequency splitting, the two mirrors act just as if they were one, producing to the same correction.

This splitting strategy has been implemented in the GPI end-to-end simulation (see this paper[18] for a detailed discussion of the simulation). Fig. 2 shows mirror shapes during a closed-loop simulation run. In this case the GPI simulation was run with no WFS noise and either an infinite-stroke tweeter or the Woofer-Tweeter configuration. The phases conjugated by the mirror(s) are shown. The combined phase conjugated by the Woofer-Tweeter configuration differs only slightly from the phase in the Tweeter-only configuration. For this case the Tweeter shape in 665 nm RMS in the controllable band, and the difference between it and the shape achieved in closed-loop with the Woofer-Tweeter algorithm is 4.5 nm RMS tip/tilt on the Woofer and 1.2 nm RMS controllable spatial frequencies. The total difference is 4.7 nm RMS. Note that this is not just a demonstration that the same phase can be made with two mirrors instead of one, but that in closed-loop the two mirrors behave like one.

3. LSI CONTROL STRATEGIES

A linear, shirt-invariant system can be completely described by either its impulse response or its transfer function[19] (which is just the Fourier transform of the impulse response). For a LSI DM, the impulse response is termed the influence function, which is the shape that the surface of the mirror takes when a single actuator is moved. Fourier modes and LSI systems go together, as Fourier mode are eigenfunctions of LSI systems.[19] This means that in Fourier control the entire characterization of the LSI DM is reduced simply to a real-valued DM gain. As shown in Fig. 1, pre-compensation for the DM is accomplished with a filter. This filter is just the division of the phase signal by the transfer function of the DM.

In most cases it is assumed that the DM or MEMS mirror obeys linear superposition, such that the surface is simply the sum of the responses of individual actuators. For example, see Ellerbroek, where a matrix of influence

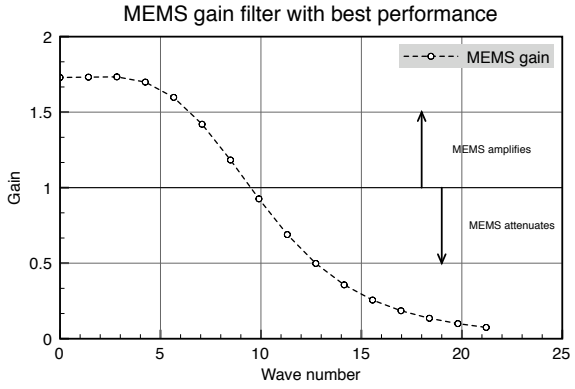

Figure 3. MEMS gains from the filter which produces best open-loop shaping performance. This is a slight scaling of the Fourier transform of the normalized influence function of the MEMS as measured by the PSDI.

functions is used in the minimum-variance unbiased reconstruction approach.[20] In our own work with FTR and Fourier modal control, we assume that the DM influence function is pre-compensated. For Optimized-gain Fourier Control, it was shown that accurate DM pre-compensation is essential (see Fig. 6 of that paper[15]). Otherwise the modal gain optimization will primarily correct for the DM, not the true signal-to-noise in the atmospheric measurement.

In our end-to-end simulation for GPI the Woofer and Tweeter are modeled as LSI and the gains are directly available for each mirror. In an experimental system the MEMS gains must be determined empirically. In our recent work[7] we have shown that a Gaussian fit to the experimentally measured MEMS gains substantially improves the closed-loop convergence of the LAO testbed when run with the spatially-filtered Shack Hartmann and FTR. This MEMS gain filter was determined by placing Fourier modes on the MEMS of a specific magnitude, then reconstructing the phase from the WFS measurement. The scaling between the desired and estimated phase represented the necessary pre-scaling to shape the MEMS correctly. The resulting gains were very well-fit by a Gaussian, so that model was used instead of the actual experimental gains, which had a small amount of noise.

However, because the FTR filter was based on a model, not a measurement of the physical WFS, this pre-scaling will also incorporate any corrections necessary because the reconstruction filter is not exact. As such, it will not be just a MEMS compensation filter, but also have a WFS-compensation component. For our experiment here with open-loop shaping of the MEMS, no WFS or FTR was used. For this application we need the pure MEMS gains.

To develop this filter we began with a measurement of the LAO MEMS influence function offset away from bias. This was then matched to a numerical model (see the discussion in Section 4.1, and in particular Fig. 6). This model was sampled with 8 pixels per actuator in a 256×256 grid. The DFT of this signal was calculated, and the magnitude was taken to exclude any possible linear phase terms due to alignment. The center 32×32 region of frequency space was extracted; this contains the gains for the exact spatial frequencies which the MEMS can make. In this method there was a slight uncertainty in normalization; it turned out that an extra normalization by 1/1.15 was necessary. A slice across the diagonal of the final filter is shown in Fig. 3. Spatial frequencies of $\mathrm{sqrt}(k^2 + l^2) < 9$ are amplified by the MEMS. The highest spatial frequencies are substantially attenuated. The MEMS pre-compensation is executed by dividing the Fourier transform of the phase by this filter.

For this initial try at a direct MEMS pre-compensation filter, we used the experience (discussed later in Section 4.2) that a filter made with a low-stroke influence function performs best. As such, the gains at the lowest-spatial frequencies of the filter (see Fig. 3) are lower than in the case of a maximum stroke influence function. As the shapes tested have very little low-frequency content, it remains to be determined is this choice is the best for the LAO MEMS.

Given this filter, we conducted an experiment to determine to what extent an LSI model can be used to precisely control the MEMS. This is best examined through the open-loop error, in which we attempt to in a

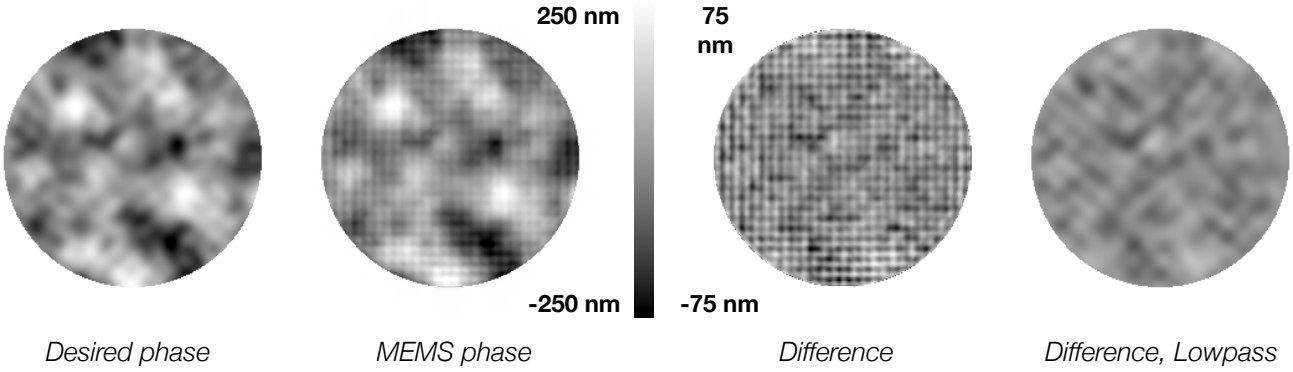

Figure 4. Desired phase (far left) and actual phase made by MEMS, with influence function pre-compensation. The input signal is 530 nm peak-to-valley and 90 nm RMS, and has the spatial frequency content of a GPI Tweeter correcting Kolmogorov turbulence. On the right are the open-loop error signals, at all frequencies (exhibiting strong actuator ripple) and in the controllable band. The controllable open loop error is 11 nm RMS. Note different color scales.

single step drive the MEMS to a specific shape. For the desired phase shape we chose a specific realization of phase on the Tweeter for a median r_0 of 14.5 cm in the GPI simulation. Since Woofer-Tweeter control was used, the Tweeter does not have the few lowest spatial frequencies. This signal, cropped to 32×32, was initially 1885 nm peak-to-valley and 300 nm RMS. Two versions of this signal were used: the raw signal, and the signal pre-compensated by the MEMS filter.

After the MEMS mirror was flattened to remove the inherent shape of the surface, that phase was added to successive scalings of the Tweeter phase and applied to the MEMS. The PSDI was then used to measure the phase on the MEMS. This process should allow us to directly produce the desired phase shape. The data were analyzed using a just a few manipulations. First, the PSDI measurements, which have 10.65 pixels per actuator, were resampled to have 8 pixels per actuator, to better agree with our ideal, sinc-interpolating MEMS simulation. Second, the PSDI measurements had to be shifted to align correctly with our estimate of the desired phase. With just this resampling and shifting, the actual measured phase signals can be directly compared to the ideal phase. The difference between these signals was calculated and evaluated for error in the controllable spatial-frequency band. An example of this is shown in Fig. 4.

For a wide range of input aberration amplitudes, the open-loop error is in the 10's of nm when the pre-compensation is used. The error increases linearly with the amount of input aberration. For the case shown in Fig. 4, where the input is 530 nm peak-to-valley and 90 nm RMS, we achieve 11 nm RMS error. The results for a wide range of input amplitudes are given in Fig. 5. For comparison, the open-loop error when no MEMS compensation is used is also given. It is 6 times larger than the pre-compenation result, because the MEMS amplifies spatial frequencies up to wave number 9 (see Fig. 3). For both methods the open-loop error initially grows linearly. For the raw, uncompensated phase signal, once saturation occurs the error decreases below the trend, as saturation is reducing the amount of excess signal. For the pre-compensated signal, the error increases once saturation begins, as the saturation prevents the MEMS from taking the full shape. When the saturated regions are masked, the open-loop error peaks at 34 nm RMS for an input of 300 nm RMS (which is a reasonable input amount for GPI under median atmospheric conditions). These saturated regions represent 25% of the pupil, however, so some uncertainty about the accuracy of this calculation does exist. The ability to shape a 300 nm RMS aberration to 34 nm (as opposed to 44 nm) will only be verified when a MEMS with adequate stroke is available.

For comparison, Morzinski et al.[21] report 15 nm RMS error for a 500 nm peak-to-valley Kolmogorov-like shape with a non-linear shaping algorithm. Stewart et al.[22] give a result of less than 15 nm RMS open-loop error for input aberrations up to 1.5 μm peak-to-valley, but only simple modes, and not a Kolmogorov shape, were used in those tests.

Further work remains to be done. In particular, the MEMS compensation filter can be refined. The low-

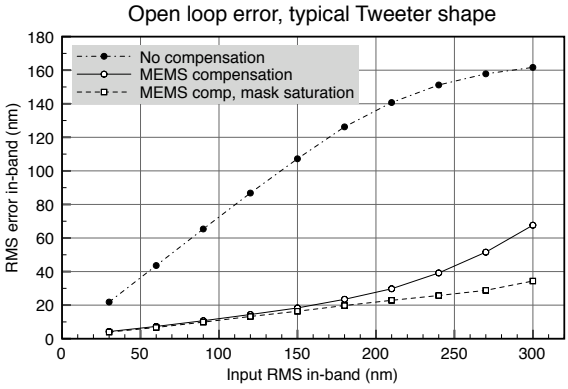

Figure 5. Open-loop error for the phase shape given in Fig. 4, for a variety of input amplitudes. Actuator saturation on the LAO MEMS starts to occur past an input of 150 nm RMS, and this saturation grows to be at least half the total open-loop error. The dashed line gives RMS error when the saturated regions are masked. For comparison, the error with no MEMS compensation is much higher.

spatial frequency gains need to be determined exactly, as this test does not have significant power there. We also can explore if the best filter changes as a function of total input aberration, which is a sign of non-linearity in the MEMS. In addition, the best estimate of the open-loop error for typical Tweeter shapes will come from trying dozens of different inputs shapes, not just the single result given here.

4. NON-LINEAR MEMS IN CLOSED-LOOP

The non-linearity of the MEMS is important when the device is used in open-loop control, as will be necessary for such future AO applications as Multi-Object Adaptive Optics (MOAO).[23] Open-loop control algorithms that account for inherent device linearity have been proposed by Vogel and Yang,[24] Morzinski et al.[21] and Stewart et al.[22] among others. For GPI our primary concern is whether we will need such a non-linear compensation algorithm, or if the non-linearity of the GPI MEMS will be small enough that use on closed-loop will not degrade performance.

Our first step in analysis of this problem has been to implement a non-linear MEMS simulator for use with our end-to-end AO simulation code. At least three different approaches exist for such a modeling and solution of the MEMS surface from input voltages: the same work of Vogel and Yang as above,[24] Baker,[25] and the work of Gavel[26] followed by Morzinski et al.[21] For both its ease of implementation and reasonable computational cost, we have chosen the Vogel and Yang model. Though this model is incomplete (see Section 4.B of that paper for discussion of how), it provides a reasonable starting point for a non-linear MEMS in our simulation.

4.1. Tuning the model to the LAO MEMS

The first step we took before using the non-linear MEMS was to tweak its parameters to make it similar in character to the Boston MEMS at the LAO. It must be emphasized that our non-linear MEMS model is not an accurate simulation of the actual LAO MEMS; it is simply a type of non-linear mirror with some similar characteristics to the LAO MEMS. The model has two major parameters which can be adjusted. The first is σ, which controls the stiffness of the mirror surface. This changes the width of the influence function and the inter-actuator coupling. The second parameter is the range of possible voltages used. As the range of voltages increases, the non-linearity of the model also increases.

As discussed above, the influence function completely characterizes an LSI mirror. Even when the mirror has non-linear characteristics, the influence function accounts for a substantial portion of the mirror response, as was shown in the previous section. The non-linear model parameter σ was adjusted until the shape of the model's influence function matched that of the LAO MEMS, as is illustrated in Fig. 6

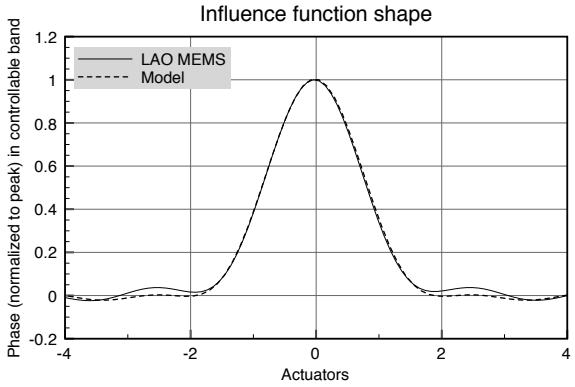

Figure 6. Slice through the influence function of the LAO MEMS in the controllable band of spatial frequencies, as compared to the non-linear model. A parameter in the model was adjusted to match influence functions.

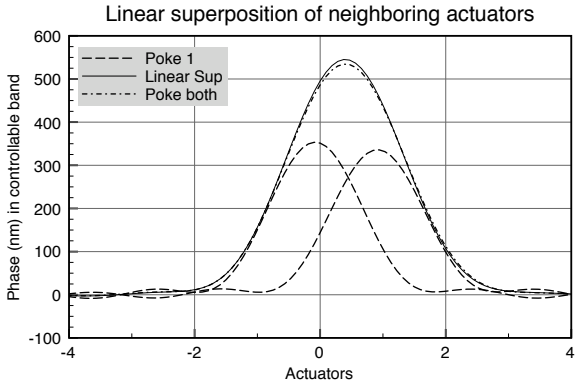

Figure 7. The first non-linearity calibration test, where the linear superposition of two neighboring actuators is assessed. Slices through the controllable phase are shown. The LAO MEMS has 2% less stroke than what is predicted by a linear superposition of the two individual pokes.

The second step in matching the model to the LAO MEMS is characterizing the non-linearity of the mirror. As our measure of non-linearity, we consider the deviation from LSI in the case of two adjacent actuators being poked. We term this "two-poke non-linearity" and we express it in the percent that the actual mirror deviates from the shape predicted by linear superposition. In all cases, it is measured in the controllable spatial frequencies only, to exclude higher-frequency effects such as actuator ripple. Determining the two-poke non-linearity requires three measurements of the MEMS. Two directly adjacent actuators are poked individually, and then together. For the LAO MEMS, in the controllable band the non-linearity causes the mirror shape to be 2% less than what linear superposition predicts, as is shown in Fig. 7 for pokes at maximum stroke away from bias. The non-linear MEMS model was tuned by adjusting the maximum voltage range to produce this same amount of non-linearity

This simple measurement is not sufficient to characterize the total non-linearity of the MEMS. In particular, the Vogel and Yang model neglects factors which will reduce the non-linearity when many actuators are moved, as opposed to the extreme case above where a single actuators is moved to maximum deflection while the rest are held at bias.[27] To test this, a 3×3 region was actuated one column at a time and the displacement of the center actuators was examined. The LAO MEMS exhibited less than 2% non-linearity, while the model exhibited 10%. This means that the non-linear model that we use in the simulation will be pessimistic; it has more non-linearity then the LAO MEMS.

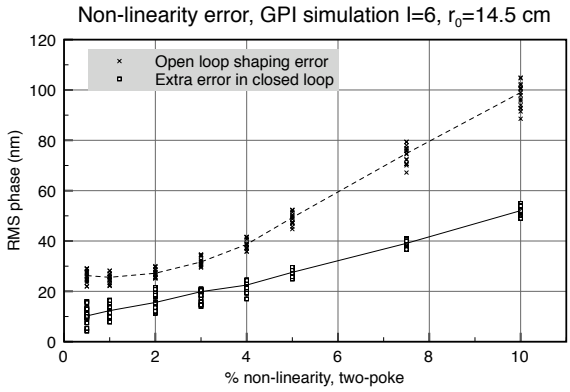

Figure 8. Error from using the simulated non-linear mirror in the GPI simulation, for 16 different random atmospheres for the case of NGS I magnitude 6 and r_0 of 14.5 cm. Extra error in closed loop represents the non-linearity error beyond the temporal lag and WFS noise terms. Open loop shaping error is the phase difference caused by the non-linaer mirror when forming a typical Tweeter shape from each closed-loop run.

4.2. Influence function compensation

As discussed above, pre-compensation of the actuator commands to account for the influence function improves performance. However, the influence function changes shape slightly as the amount of the actuation increases. This is observed both in the LAO MEMS and in the model MEMS. This gives rise to the question of what the best influence function pre-compensation filter is. To test this, filters were constructed based on the pokes at 20%, 40%, 60%, 80% and 100% of the maximum stroke away from bias. These filters differed primarily in the gain on the lowest spatial frequencies. In closed-loop the filter which produced the minimum extra error was the 20% actuation filter. This is consistent with the fact that when correcting atmospheric phase, the Tweeter will form shapes in which neighboring actuators are not at maximum stroke apart from each other.

4.3. Closed-loop simulations

To evaluate the impact of a non-linear mirror in GPI AO, we incorporated the Vogel and Yang simulator into our existing end-to-end AO simulation. We specified several different amounts of non-linearity for the model mirror, ranging from 0.5% to 10% two-poke non-linearity. Each model MEMS was run with the best influence function compensation filter (see above). The simulation was run on a five-layer, median r_0 atmosphere. For a set of 16 different realizations of the five-layer turbulent atmosphere, the simulation was run with the LSI Tweeter, and with each of the various non-linear Tweeters. Each trial was run until convergence, then the RMS error was calculated from the median variance over 25 time steps. Using the reference closed-loop residual from the LSI mirror, the amount of extra error due to non-linearity was calculated with a RSS of the error terms.

This extra error in closed-loop is shown for 16 different random atmospheric inputs in Fig. 8. These trials were all done with an optimized-gain filter (using OFC[15]) which was determined with the LSI mirror. Even the 0.5% non-liner mirror still had a median of an extra 10 nm RMS in-band. The amount of extra error in closed-loop increases as a linear function of the two-poke non-linearity of the MEMS model. The open-loop error was also calculated by taking a sample Tweeter phase from each trial and shaping it on the LSI Tweeter and the non-linear Tweeters. In this case a base level of open-loop error existed at around 25 nm RMS in-band, as is shown in Fig. 8. For models with 3% or more two-poke non-linearity, the open-loop error also increases as a linear function.

The exact relationship between the open-loop error and the closed-loop error is complicated and difficult to model. The shape difference on the non-linear MEMS most likely has a temporal power spectrum similar to the input atmosphere. But the correlation (or lack thereof) of this shape difference with the signal and noise inputs is unknown. The specific temporal controller (uniform modal gains, optimized modal gains, predictive Kalman filter) which is used will also affect the relationship between open-loop and closed-loop error.

5. CONCLUSIONS

We have described and investigated several remaining issues involved in using a MEMS mirror for GPI. The 64×64 MEMS mirror will be the Tweeter of the system, correcting the high-spatial frequency phase aberrations in tandem with a low-order but high-stroke Woofer DM. We have demonstrated with our end-to-end GPI simulation that an LSI Woofer and LSI Tweeter controlled with this frequency-splitting algorithm behave in closed-loop just like a single, infinite stroke mirror, to within a few nm RMS.

In our closed-loop simulations, with a non-linear mirror that is similar to, but not exactly like the LAO MEMS, we found that a 2% two-poke non-linearity mirror will produce 16 nm RMS of extra in-band error. This mirror will have a typical open-loop error of 27 nm RMS. Knowledge of the simulation model assumptions leads us to predict that it has more non-linearity than the LAO MEMS.

Our initial MEMS pre-compensation filtering experiment on a similar amount of Kolmogorov-shape Tweeter phase produced an open-loop error of 34 nm RMS (44 nm RMS without masking the saturated regions). This puts the current LAO MEMS at 21 nm RMS over the simulated open-loop error. This is contradictory with the prediction that the simulator has more non-linearity than the LAO MEMS. Further work on refining the exact MEMS pre-compensation filter must be done to determine the true limitation of the LSI algorithm to shape the MEMS.

Two other uncertainties remain in the open-loop error result. First, tests with dozens of random Kolmogorov-like Tweeter shapes much be conducted to obtain a statistical profile of open-loop errors. Second, a device with adequate stroke, namely the new 64×64 BMC MEMS, must be used to ensure that the saturation masking of the residual error is not biased. It also remains to be verified if the new 64×64 BMC MEMS will have the same amount of non-linearity as the current LAO MEMS. A decision on whether or not a special non-linearity compensation algorithm must be used in GPI cannot be made until full testing of the actual GPI MEMS device is executed.

ACKNOWLEDGMENTS

We would like to thank Bruce Macintosh for his helpful commentary and advice. This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. The document number is LLNL-CONF-400061. This work has been supported by the National Science Foundation Science and Technology Center for Adaptive Optics, managed by the University of California at Santa Cruz under cooperative agreement No. AST - 9876783. This work has been supported by the Gordon and Betty Moore Foundation through its grant to the UCO/Lick Laboratory for Adaptive Optics.

REFERENCES

1. B. Macintosh, J. Graham, D. Palmer, R. Doyon, D. Gavel, J. Larkin, B. Oppenheimer, L. Saddlemyer, J. K. Wallace, B. Bauman, J. Evans, D. Erikson, K. Morzinski, D. Phillion, L. Poyneer, A. Sivaramakrishnan, R. Soummer, S. Thibault, and J.-P. Véran, "The Gemini Planet Imager," in *Advances in Adaptive Optics II*, B. L. Ellerbroek and D. B. Calia, eds., *Proc. SPIE* **6272**, p. 62720L, 2006.
2. R. Soummer, "Apodized Pupil Lyot Coronagraphs for Arbitrary Telescope Apertures," *Ap. J.* **618**, pp. L161–L164, 2005.
3. J. K. Wallace, R. Bartos, S. Rao, R. Samuele, and E. Schmidtlin, "A laboratory experiment for demonstrating post-coronagraph wavefront sensing and control for extreme adaptive optics," in *Advances in Adaptive Optics II*, B. L. Ellerbroek and D. B. Calia, eds., *Proc. SPIE* **6272**, p. 62722L, 2006.
4. T. Fusco, C. Petit, G. Rousset, J.-F. Sauvage, K. Dohlen, D. Mouillet, J. Charton, P. Baudoz, M. Kasper, E. Fedrigo, P. Rabou, P. Feautrier, M. Downing, P. Gigan, J.-M. Conan, J.-L. Beuzit, N. Hubin, F. Wildi, and P. Puget, "Design of the extreme AO system for SPHERE, the planet finder instrument of the VLT," in *Advances in Adaptive Optics II*, B. L. Ellerbroek and D. B. Calia, eds., *Proc. SPIE* **6272**, p. 62720K, 2006.
5. B. Bauman, "Personal communication." Lawrence Livermore National Laboratory, 2007.
6. J. W. Evans, B. Macintosh, L. A. Poyneer, K. Morzinski, S. Severson, D. Dillon, D. Gavel, and L. Reza, "Demonstrating sub-nm closed loop MEMS flattening," *Opt. Exp.* **14**, pp. 5558–5570, 2006.

7. L. A. Poyneer, D. Dillon, S. Thomas, and B. A. Macintosh, "Laboratory demonstration of accurate and efficient nanometer-level wavefront control for extreme adaptive optics," *submitted to Appl. Opt.* , 2007.
8. K. M. Morzinski, J. W. Evans, S. Severson, B. Macintosh, D. Dillon, D. Gavel, C. Max, and D. Palmer, "Characterizing the potential of MEMS deformable mirrors for astronomical adaptive optics," in *Advances in Adaptive Optics II*, B. L. Ellerbroek and D. B. Calia, eds., *Proc. SPIE* **6272**, p. 627221, 2006.
9. B. A. Macintosh, "Gemini planet imager preliminary design document volume 2: Instrument design," tech. rep., Lawrence Livermore National Lab submitted to International Gemini Project Office, 2007.
10. J.-F. Lavinge and J.-P. Véran, "Woofer-tweeter control in an adaptive optics system using a Fourier reconstructor," *in preparation* , 2007.
11. J.-P. Véran and G. Herriot, "Woofer-tweeter tip-tilt control for NFIRAOS on TMT," in *Advances in Adaptive Optics II*, B. L. Ellerbroek and D. B. Calia, eds., *Proc. SPIE* **6272**, p. 62721R, 2006.
12. E. M. Johansson and D. T. Gavel, "Simulation of stellar speckle imaging," in *Amplitude and Intensity Spatial Interferometry II*, J. B. Breckinridge, ed., *Proc. SPIE* **1237**, pp. 372–383, 1994.
13. R. Conan, C. Bradley, P. Hampton, O. Keskin, A. Hilton, and C. Blain, "Distributed modal command for a two-deformable-mirror adaptive optics system," *Appl. Opt.* **46**, pp. 4329–4340, 2007.
14. L. A. Poyneer, D. T. Gavel, and J. M. Brase, "Fast wave-front reconstruction in large adaptive optics systems with use of the Fourier transform," *J. Opt. Soc. Am. A* **19**, pp. 2100–2111, 2002.
15. L. A. Poyneer and J.-P. Véran, "Optimal modal Fourier transform wave-front control," *J. Opt. Soc. Am. A* **22**, pp. 1515–1526, 2005.
16. E. Gendron and P. Léna, "Astronomical adaptive optics I. Modal control optimization," *Astron. Astrophys.* **291**, pp. 337–347, 1994.
17. L. A. Poyneer, B. A. Macintosh, and J.-P. Véran, "Fourier transform wavefront control with adaptive prediction of the atmosphere," *J. Opt. Soc. Am. A* **24**, pp. 2645–2660, 2007.
18. L. A. Poyneer and B. A. Macintosh, "Optimal Fourier Control performance and speckle behavior in high-contrast imaging with adaptive optics," *Opt. Exp.* **14**, pp. 7499–7514, 2006.
19. A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and Systems (2nd Ed.)*, Prentice Hall, New Jersey, 1997.
20. B. L. Ellerbroek, "Efficient computation of minimum-variance wave-front reconstructors with sparse matrix techniques," *J. Opt. Soc. Am. A* **19**, pp. 1803–1816, 2002.
21. K. M. Morzinski, K. B. W. H. e, D. T. Gavel, and S. M. Ammons, "The open-loop control of MEMS: modeling and experimental results," in *MEMS Adaptive Optics*, S. S. Olivier, T. G. Bifano, and J. A. Kubby, eds., *Proc. SPIE* **6467**, p. 64670G, 2007.
22. J. B. Stewart, A. Diouf, Y. Zhou, and T. Bifano, "Open-loop control of a MEMS deformable mirror for large-amplitude wavefront control," *J. Opt. Soc. Am. A* **24**, pp. 3827–3833, 2007.
23. R. G. Dekany, M. C. Britton, D. T. Gavel, B. L. Ellerbroek, G. Herriot, C. E. Max, and J.-P. Veran, "Adaptive optics requirements definition for TMT," in *Advancements in Adaptive Optics*, D. B. Calia, B. L. Ellerbroek, and R. Ragazzoni, eds., *Proc. SPIE* **5490**, pp. 879–890, 2004.
24. C. R. Vogel and Q. Yang, "Modeling, simulation, and open-loop control of a continuous facesheet MEMS deformable mirror," *J. Opt. Soc. Am. A* **23**, pp. 1074–1081, 2006.
25. G. J. Baker, "A fast high-fidelity model for the deformation of continuous facesheet deformable mirrors," in *Advances in Adaptive Optics II*, B. L. Ellerbroek and D. B. Calia, eds., *Proc. SPIE* **6272**, p. 627224, 2006.
26. D. Gavel, "Laboratory for adaptive optics at UC Santa Cruz: project status and plans," in *Advances in Adaptive Optics II*, B. L. Ellerbroek and D. B. Calia, eds., *Proc. SPIE* **6272**, p. 62721U, 2006.
27. P. Bierden, "Personal communication." Boston Micromachines Corporation, 2007.