Deformable Mirror Calibration for Adaptive Optics Systems

A. Sivaramakrishnan[a] & B. R. Oppenheimer[b]

[a]Space Telescope Science Institute
3700 San Martin Dr.
Baltimore, MD 21218
email: anands@stsci.edu

[b]Palomar Observatory
105-24 California Institute of Technology
Pasadena, CA 91125
e-mail: bro@astro.caltech.edu

ABSTRACT

We present a rapid technique for the simultaneous calibration and flattening of a deformable mirror in an adaptive optics system. We achieve an rms surface figure equivalent to better than $\lambda/100$ in the K band in our first application of the scheme. The technique requires the presence of a phase shifting interferometer within the adaptive optics system. Measurements of the mirror surface are made while each mirror actuator's gain is estimated, in a stable, iterative, convergent sequence. Knowledge of the mirror response function to single actuator activation (i. e. the so-called "influence function") is not needed. The iterative process ends when the mirror surface reaches some preset desired flatness, which can be specified by the rms deviation of the surface. An overall scale correction to the gains is measured and applied after the iterations converge. Our technique uses the mirror under conditions similar to those present when the adaptive optics system is running. This technique provides the user the ability to calibrate the mirror in a matter of a minute or two, enabling re-calibration of the actuator gains many times during a night with minimal impact on observing duty cycle, in order to adapt to changing conditions.

Keywords: Adaptive Optics, Deformable Mirrors

1. INTRODUCTION

Adaptive optics (AO) systems using deformable mirrors with a few hundred channels of sensing and actuation are becoming quite common. Several functioning systems have demonstrated the feasibilty of this approach to ground-based, high-resolution imaging. The key to making an AO system work well is a thorough understanding of the many parts that comprise the instrument. A critical part of most AO systems is the deformable mirror (DM), the component that actually performs the wavefront correction.

During the development of the AO system for the Palomar Hale 5 m telescope, we spent considerable time on understanding the behavior of the 349 actuator deformable mirror made by Xinetics. The mirror and our work on understanding it are described elsewhere.[1] The results of our previous work are an understanding of the DM's dynamics on millisecond timescales and the characterization of the behavior of the mirror surface under the influence of the actuators. During that exercise, we devised a rapid calibration scheme which we present in detail here. This scheme permits DM recalibration many times a night with little impact on the observing duty cycle. This is important because PMN-actuated DMs (such as our Xinetics mirror) exhibit strongly temperature-dependent behavior.[2] In our first test of this calibration scheme, we controlled the mirror surface to an accuracy of 19 nm rms. We anticipate doing much better once the process is fully automated.

It should be noted that it is possible to operate an AO system in closed loop with a fairly rough knowledge of the actuator response to specified voltages, because the AO loop is itself an iteration of a (usually) convergent process. However, such an approach achieves DM control at the expense of precious bandwidth, and the quality of the wavefront correction is degraded, especially for faint targets.[3] Without well-calibrated gains, one cannot assert a specified shape onto the DM surface in one step.

2. THE CALIBRATION SCHEME

In the case of the Palomar Xinetics DM, we need to determine the gains for 241 actuators of the DM—although there are 349 actuators, only the inner 241 are used in the system.

We define the gain of an actuator as that multiplicative constant which converts a digital actuator setting (which comes from the control computer) to a relative position of that actuator in physical units, such as microns. The actuators have been shown to behave linearly over most of their commanded range,[1] so the assumption of a constant gain is a useful approximation. Our calibration scheme flattens the DM and determines the relative gains of all of the commanded actuators simultaneously. The scheme is iterative, and it flattens the mirror surface to an arbitrarily chosen accuracy.* This calibration scheme requires the presence of an interferometer with which one can measure the surface figure of the DM. In our first tests of the scheme, we used a Zygo Corporation interferometer, and analyzed the data it took with our own software. Once the AO system is fully operational, we will be able to do this using the "stimulus" subsystem of the instrument,[4] which is capable of producing interferograms similar to those produced by the Zygo. In the following we refer to "the interferometer" generically. With the stimulus interferometer integrated into the AO system, the calibration and flattening scheme will be fully automated, since the computer that analyzes the interferograms will also be capable of issuing new commands directly to the DM.

The technique consists of the following steps.

1. We make an initial estimate of the gains for all the actuators. This could be either the set of gains from previous measurements, or simply the crude estimate that all the gains are equal to maximum movement (in μm) divided by the maximum digital setting on the actuators's DAC. In the case of the Palomar AO system, this estimate is $5.0\mu m/4096$. We represent the gains as g_i^j, the gain of actuator i at iteration j. Initially j is zero and i ranges from 1 through N, the total number of actuators.

2. We set all actuators to the midrange value

$$s_i^0 = 2048,$$

where s_i^0 is the setting for actuator i at the zeroth iteration. This moves each actuator to a height $h_i(s_i^0)$.

3. We then measure the height of the mirror surface with the interferometer. We call this measurement z_i^j, where i and j have the same meaning as they do in the notation for g and s. Because of the nature of the interferometric measurement, the value of z_i^j is actually the height of the mirror surface at actuator i relative to the *average* height of the mirror surface:

$$z_i^j = h_i(s_i^j) - \textstyle\sum_i h_i(s_i^j)/N$$

4. We attempt to reduce each z_i^j to zero in the next iteration by moving the actuators. To do this, we compute new settings for all the actuators:

$$s_i^{j+\frac{1}{2}} = s_i^j - z_i^j/g_i^j$$

These new settings are given the $j + \frac{1}{2}$ iteration number because they are not the actual settings we should assert on the mirror. An additional constraint on the settings is necessary for convergence of the scheme. Application of this constraint comprises step 5.

5. We constrain the average of the settings at all iterations to be s_i^0, the midrange value of the digital setting:

$$s_i^{j+1} = s_i^0 + s_i^{j+\frac{1}{2}} - (\textstyle\sum_i s_i^{j+\frac{1}{2}})/N$$

6. We apply these new settings, s_i^{j+1}, and make a new measurement with the interferometer, thereby determining a new set of deviations from the average height of the surface, z_i^{j+1}.

*The actual accuracy depends on how well the actuators can be controlled and on how well the mirror surface can be measured. Actuators with significant hysteresis or non-linearities will of course degrade the accuracy of the calibration and control. The final figure of merit, after convergence is achieved, is a measurement of some combination of all such errors. We suspect that the largest source of error is probably hysteresis.

7. With the new measurements of z_i^{j+1}, we can correct the gains, with a proviso on the size of correction being larger than some threshold—we recompute new gains for all settings that have changed by more than a threshold value, which we chose to be some small fraction t of the largest setting s_{max}:

$$\text{If } |s_i^{j+1} - s_i^j| > ts_{max}, \text{ then}$$
$$g_i^{j+1} = (z_i^{j+1} - z_i^j)/(s_i^{j+1} - s_i^j)$$

The purpose of this tolerance is to prevent division by small numbers, and to set a tolerance on the final flatness of the mirror surface. In practice t should be chosen to reflect the level of controllability of the actuators, which in this case will be dictated by their intrinsic hysteresis. Before convergence of the gains is achieved, there may be small differences between the mean heights of the mirror in any two consecutive steps. This is essentially because the constraint in step 5 cannot be applied to h_i^j directly, which is unmeasurable. More detailed discussion of this follows step 8.

We have not proved rigorously why convergence of the gains is possible using the above expression, though we have demonstrated that this process does converge in practice.

8. We loop back to step 4 until the gains and settings no longer get changed during an iteration. After this convergence is achieved (which in practice occurs in less than ten iterations, depending on the value of the tolerance t), it is possible that some of the actuators will have g and s equal to their initial values (i.e. they were never changed during the iterations). Such actuators fortuitously start off within the tolerance setting of the correct position for the final flat surface with the initial midrange setting. The gains of these actuators will not be well determined, since they remain at whatever their initial estimate was. To account for this possibility, and to ensure that the gains are determined for all the actuators, we repeat the whole process at new values of the initial settings, $s_i'^0$. The new values are computed by imposing a 1 μm increase in the heights of all the actuators from the flattened position:

$$s_i'^0 = s_i^{final} + 1\mu m/g_i^{final}$$

This forces the mirror to a new position where most of it will still be as flat as before, but the actuators with poorly determined gains will be at incorrect heights (unless the initial gain estimates were very close to their real gains). Repeating the iterative flattening process at this height will determine accurate gains for all the actuators. If necessary yet one more shift can be executed. Verification of a good gain determination is simply achieved by choosing an arbitrary height and trying to flatten the mirror at that height.

Finally gains have been determined, but an additional correction is necessary.

3. GAIN CORRECTION

The iterative process described above produces the ability to achieve good absolute flatness, but at a height that is known only approximately, unless all the actuator gains are equal. To see this, we note that after convergence the constant height of the flattened surface is related to the harmonic mean of the gains (we have dropped the iteration indices here for convenience):

$$h = h_i = s_i g_i$$

This means that all the heights are equal. Furthermore we know that

$$Ns^0 = \sum s_i,$$

so

$$Ns^0 = h \sum (1/g_i).$$

If the variance of the gains is small compared to their values, this height is approximately (excluding terms or third or higher order)

$$h = \bar{g}s^0(1 + \sigma_g^2/\bar{g}^2),$$

where $\bar{g}$ is the mean of the gains produced by the iterative process, and σ_g^2 is the variance in these gains. Thus the surface is slightly higher than it would have been had all the gains been equal.

This has important ramifications for the gain calculations. In fact, the gains calculated in the iterative process, ending with the completion of step 8 will be scaled slightly from their true values by the amount $(1+\sigma_g^2/\bar{g}^2)$, excluding higher order terms in the binomial expansion. In our case this scale change amounts to about a 4% error (because of the variance of the gains in our particular mirror). This scale change is easily calculated using the expression given above, or it can be ignored at the expense of very little bandwidth in the AO closed loop operation. However, the application of this correction is essentially trivial.

Another way to see this complication is by considering the constraint in step 5 again. By changing the settings we expect to apply after using the gain in step 4, we have slightly different new settings. The difference between $s_i^{j+\frac{1}{2}}$ and s_i^j is the key here and is related to the variance of the gains as described above.

A way to test our gains experimentally is to attempt to assert a flat but *tilted* surface on the DM. We expect the interferometric measurement of heights on the DM surface to be a plane with the slope that we have chosen. Fitting the measured heights with a plane will reveal the accuracy of our gain estimates. If the slope of the fitted surface is constrained to be equal to the asserted slope, spatially correlated errors in the resultant fit would probably reveal where our assumptions of linear gains break down.

4. RESULTS

Here we present the results of this scheme in a series of surface maps. The surface maps were made at each iteration. In this case, the scheme converged in only 6 iterations. The final image shows the flattened mirror. The central region of the mirror, which is the only controllable part of it, has a residual rms figure of 19 nm.

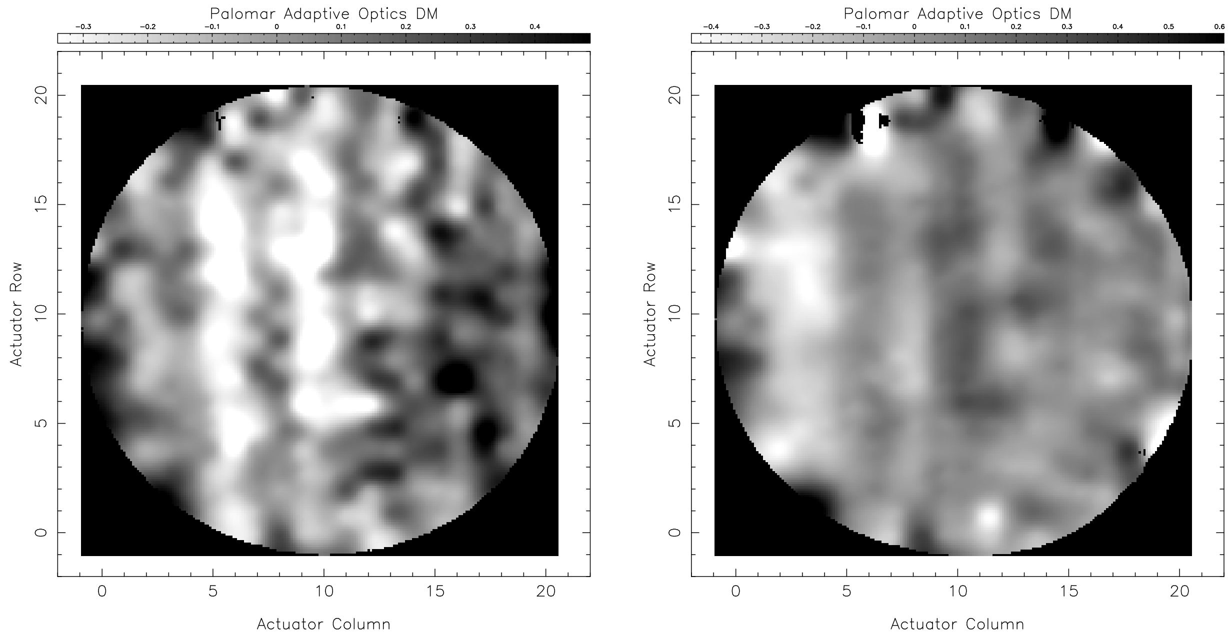

Figure 1. The image on the left shows the surface of the DM in the first step of the flattening scheme ($j = 0$). The peak-to-valley (PV) value in the inner controllable part of the mirror is 1.083 μ and the rms of that region is 0.189 μm. The right hand figure shows the DM surface after the first iteration (PV = 0.656 μm, rms = 0.124 μ).

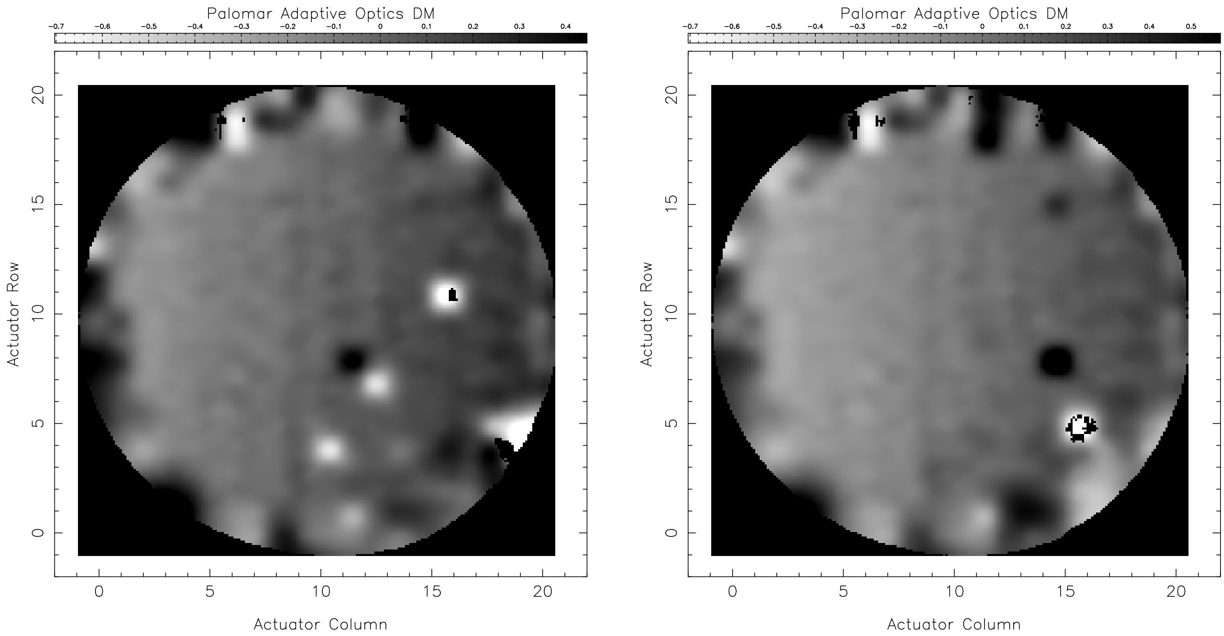

Figure 2. The image on the left shows the surface of the DM after the second iteration of the flattening scheme (PV = 1.518 μm, rms = 0.093 μm). The right hand figure shows the DM surface after the third iteration (PV = 2.159 μm, rms = 0.092 μm).

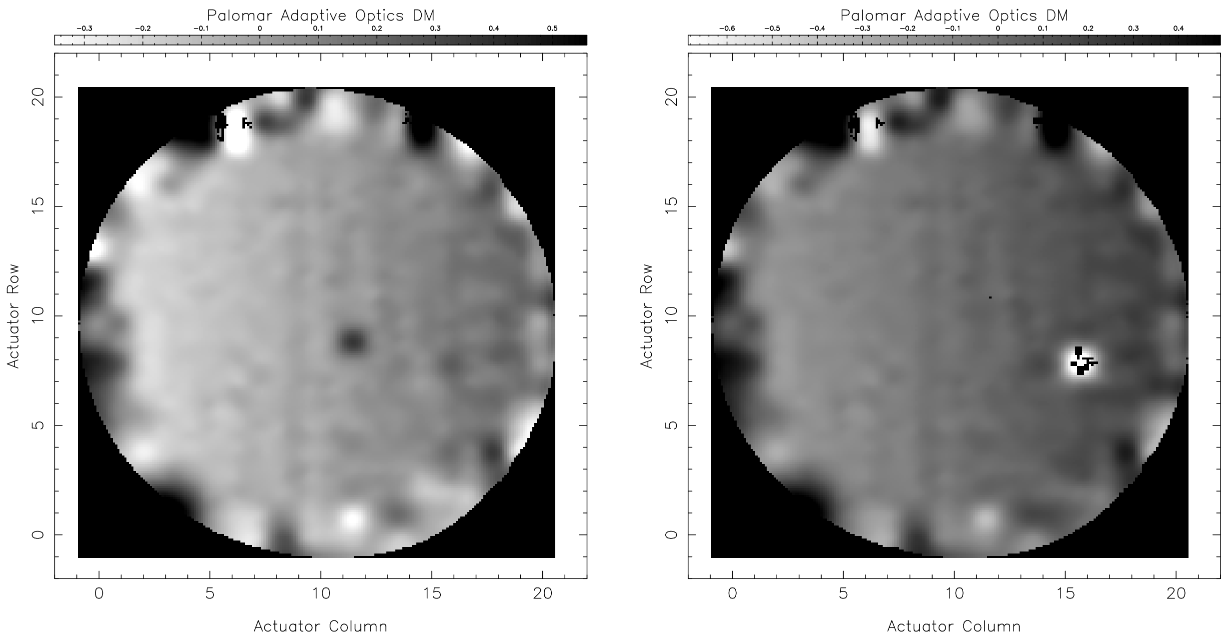

Figure 3. The image on the left shows the surface of the DM in the fourth iteration of the flattening scheme (PV = 1.282 μm, rms = 0.066 μm). The right hand figure shows the DM surface after the fifth iteration (PV = 0.436 μm, rms = 0.031 μm).

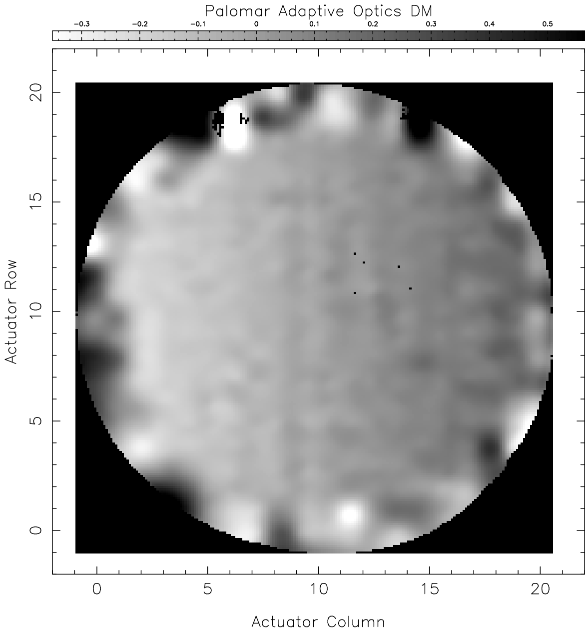

Figure 4. The final surface of the DM after 6 iterations of the flattening scheme. The inner controllable surface has PV = 0.171 μm and rms = 0.019 μm. The edges of the mirror surface are not flat. This is because that portion of the surface is controlled by "slaved" actuators. These actuators simply receive the same voltages as those they are slaved to. Since they have different gains, they will not be at the same height as the actuators inside the controlled surface. The beam used by the AO system fits entirely within the flattened part of the surface.

In practice, the final set of s_i can be used as initial offsets the next time the mirror is calibrated. The mirror can be set flat in the middle of its throw with these settings, and operation of the closed loop AO control can simply commence using the gains determined in this calibration procedure.

During the iterations of the loop occasionally single actuators will poke dramatically above or below the rest of the surface. This can be explained simply, if one considers exactly what the scheme does. Our tolerance requirement in step 7 does not disallow new updates of actuator settings. It only prevents gains from being updated in a given iteration if the new actuator setting is within a certain tolerance of the last actuator setting. In fact, the scheme updates all of the acutator settings at every iteration. As a result, the actual height of the flattened mirror surface can fluctuate by a small amount from iteration to iteration. Thus, for an actuator whose original setting only changed slightly, within the tolerance in step 7, it may take several iterations before the mirror surface moves enough that the next setting will be outside the tolerance. Then since the gain for that actuator is not well known (i. e. it might never have been updated), the new setting will be quite wrong. The gain for this particular actuator then will be corrected in the next time step. Indeed, one can see from the figures that these actuators only poke up or down for a single iteration. Subsequently they have correct gains and are well-controlled for the rest of the iterations.

This particularly demonstrates why one needs step 8, in which the mirror is then flattened again at different heights. The actuators whose gains never get updated will be grossly out of place at the other heights.

5. DISCUSSION

If we do not apply the constraint in step 5, the whole surface (or its average height) will drift away from its starting value. Such drift is a generic instability that must be dealt with in some fashion. If the iterative scheme is viewed as an abstract dynamical system in a phase space of the estimated gains, without the constraint the initial point representing out first guesses escapes in a few iterations to a globally attracting point at the origin. At this stage we would be attempting to assert a surface with infinite values of the settings s_i. Adding the constraint changes the dynamical system qualitatively: it removes this instability, and the resultant dynamics is enacted in the basin of the attracting fixed point which corresponds to the desired gains, which is colinear with both the point representing the true gains and the origin.

A second order analysis of the behavior of the unconstrained dynamical system in the neighborhood of the fixed point of the relative gains shows that the instability is quadratic in the rms of the gain errors at that iteration.

The computation required for these iterations is quite modest, since no gain matrix inversions are required. Estimates of actuator response functions are also not used in this scheme.

6. ACKNOWLEDGMENTS

The authors wish to dedicate this work to Horace Babcock.

REFERENCES

1. B. R. Oppenheimer, D. Palmer, R. Dekany, A. Sivaramakrishnan, M. A. Ealey, and T. R. Price, "Investigating a xinetics inc deformable mirror," in *Adaptive Optics and Applications*, R. K. Tyson and R. Q. Fugate, eds., *Proc. SPIE* **3126**, pp. 569–579, 1997.
2. S. Acton, P. Stomski, P. Wizinowich, J. Maute, and T. Gregory, "Test of a deformable mirror in a freezing environment," in *Optical Society of America Adaptive Optics Conference*, P. Wizinowich, ed., *Journal of the Optical Society of America* **431**, p. 1542, 1996.
3. R. Brigantic, M. Roggeman, B. Welsh, and K. Bauer, "Closed loop bandwidth optimization to maximize adaptive optics system performance," in *Adaptive Optics and Applications*, R. K. Tyson and R. Q. Fugate, eds., *Proc. SPIE* **3126**, p. 151, 1997.
4. R. Dekany, K. Wallace, G. Brack, B. R. Oppenheimer, and D. Palmer, "Initial test results from the palomar 200 inch adaptive optics system," in *Adaptive Optics and Applications*, R. K. Tyson and R. Q. Fugate, eds., *Proc. SPIE* **3126**, pp. 269–276, 1997.