Investigating a Xinξtics Inc. Deformable Mirror

B. R. Oppenheimer[1]
Palomar Observatory
105-24 California Institute of Technology
Pasadena, CA 91125
e-mail: bro@astro.caltech.edu
Member of the Palomar Adaptive Optics Development Team at the Jet Propulsion Laboratory

[1]With assistance from from Gary Brack (JPL), Richard Dekany (JPL), Brad Hines (JPL), Phil Irwin(JPL), Dean Palmer (JPL), Anand Sivaramakrishnan (Palomar) and Kent Wallace (JPL).

1. Introduction

In 1994 the Jet Propulsion Laboratory (JPL) and the Palomar Observatory entered into an agreement allocating 25% of the 200 inch Hale Telescope observing time to JPL. In partial payment for this observing time, JPL would produce a major new instrument for the telescope. The decision was made that this would be an adaptive optics (AO) instrument, and that Cornell University (the third partner in time sharing of the 200 inch telescope) would construct a near infrared camera to be used in tandem with the AO instrument.

In 1995 the Jet Propulsion Laboratory began the AO project in earnest. The goal of the project is to produce diffraction limited images in the 1 to 2.5 μm wavelength range using natural guide stars. The AO system will mount at the f/15.6 Cassegrain focus of the telescope and includes a 16 $\times$ 16 sub-aperture Hartmann-Shack wavefront sensor and electronics to enable closed-loop operation at 500 Hz. Wavefront correction will be achieved with a tip-tilt mirror and a 349 actuator deformable mirror. Ultimately the system is designed to accomodate easily the addition of a laser beacon guide star. The system should be fully operational at the close of 1998.

This report details specific tests of the deformable mirror (DM) which I conducted between September 1996 and February 1997. The tests were designed to reveal characteristics of the mirror in actual operation, which might affect the quality of wavefront correction. The tests can be broken into two groups: those dealing with the static figure of the mirror surface under the influence of the 349 actuators and those resolving the temporal behavior of the mirror.

The general finding was that the commonly-used simple picture of deformable mirrors—in which the mirror surface is assumed to be linear slopes between adjacent actuators—may not be a completely accurate model when compared to the actual mirrors, at least as evidenced by the one example studied here. The principal idea behind this study was that if we can understand how the mirror works in practice, we can adjust our control loop accordingly and improve wavefront correction.

First, we examined the specifications of the mirror and compared them with those specified by the manufacturer, Xinξtics Inc.

2. The Specifications of the DM

2.1. Xinξtics's Acceptance Tests

The mirror is shown in Figure 1 as it appeared during our tests at JPL. The mirror surface is circular with a 5.850 inch clear aperture and a controllable surface with a 5.520 inch diamter. The casing is about 7.25 inches wide and the base is 8 inches deep. The optical axis of the mirror sits 6.125 inches above the bottom surface of the base, and the mirror is protected by a clear plastic cover. The mirror itself is a 2mm thick facesheet of utralow expansion (ULE) glass silver coated on one side and attached to the actuators on the other. The actuators are attached to a base inside the casing. The mirror surface was extremely clean when it arrived at JPL, and it has remained so through this study.

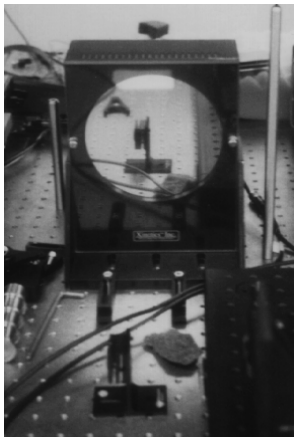

Fig. 1.— The DM as we had it set up during tests of its dynamic capabilities. The plastic cover protects it and causes the secondary reflection visible in this picture. The cover of this report also shows the DM.

The 349 lead magnesium niobate (PMN) actuators are arranged in a

square grid pattern with a spacing of 7 mm. These actuators are meant to accept 0 to 100V to deliver a mechanical stroke of at least 4 μm, with a maximum permissible difference of 2 μm between adjacent actuators. The influence of an actuator on its neighbor was found by Xinξtics to be 9 to 11% of the stroke of that actuator, and the hysteresis in actuator positioning should be less than 1%. Xinξtics states that the actuator frequency response is 4 kHz before it is attached to the mirror facesheet. This is the -3dB frequency for an actuator connected to the Xinξtics driver electronics and moving through half of the full stroke of the actuator. The speed at which the actuators move is dependent upon their capacitance. In the original specifications before the mirror was manufactured this was meant to be 0.25 μF. However, the mirror that was delivered has actuators with an average capacitance of 1.6 μF. This will adversely affect the actuator movement speeds. Fortunately, though, the electronics were designed to exceed expectations, so this loss of a factor of 6 is accounted for and the actuators whould respond fast enough for our application. (See §4.1.)

The mirror surface is coated with "protected silver," an industry standard coating which gives greater than 95% reflectance from about 0.41μm redward. This is needed because although the final result is a corrected near infrared wavefront, the wavefront sensing is done in the optical in the Palomar AO system (0.5 to 0.9 μm).

The acceptance tests run by Xinξtics Inc. found an average stroke of 5.31 μm with a standard deviation of 0.31 μm before they attached the mirror substrate. After attaching the mirror they measured the optical quality of the surface with no voltage applied to any of the actuators. "At a wavelength of 633 nm over the full aperture the peak-to-valley value of the surface is 0.54 λ (surface) with an rms value of 0.072 λ (surface) or $\lambda/14$ rms. Over the control aperture the peak-to-valley of the surface is 0.26λ (surface) with an rms value of 0.063λ or $\lambda/16$ rms."[1] All of these measurments are well within the specifications requested by JPL. The measurements we made are explained in the next two subsections. We find good agreement with the values quoted by Xinξtics.

[1]From *349 Channel Deformable Mirror for the Palomar Telescope: User's Manual, Program No. 5204* (Xinξtics Inc., 410 Great Road # A6, Littleton, MA 01460, (508) 486-0181), p. 7.

2.2. Palomar AO Configuration of the DM

In the configuration used for the Palomar AO system, only the inner 241 actuators are addressed. This assigns every square of four adjacent actuators to one of the sub-apertures on the primary mirror. The actuators outside the inner 241 are "slaved" to those on the edge of the controllable region, so that they mimic exactly the voltages on the outer edge of actuators in the controllable region. (See Figure 2.) Actuators within the shadow of the central obscuration of the primary telescope mirror are "slaved" in software.

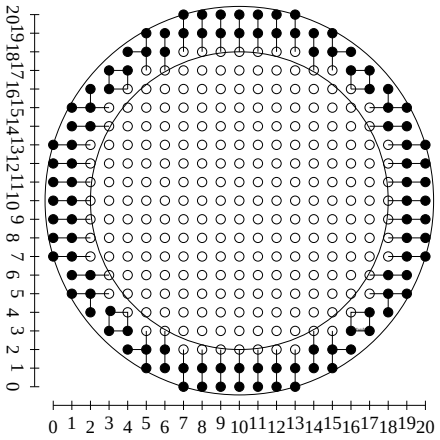

Fig. 2.— A Map of the DM. Each actuator is indicated by a circle. The open circles are actuators addressable by the electronics. The closed circles are "slaved" actuators. They are connected to controllable actuators as indicated by the lines on the plot so that they always have the same voltage as the open circle actuator they are "slaved" to. The addressing scheme is indicated and represents a face on view of the front of the DM. The size of the telescope pupil image is indicated by the inner circle and the edge of the usable mirror is indicated by the outer circle.

The electronics that control the actuators were designed and built by Dean Palmer of JPL. They include a set of zener diodes arranged so that

the difference in voltage between two adjacent actuators can be no more than 27.5V. This prevents the accidental application of a large voltage difference between two actuators, which could result in ripping an actuator off the back surface of the mirror. The epoxy bond between the actuator and the mirror facesheet is apparently weaker than the facesheet itself, so this sort of damage will not render the mirror useless. Only the separated actuator would become useless. The electronics are connected to a high voltage power supply and to a computer which issues the commands to change the voltages on the actuators. Computer commands consist of the address of an actuator followed by an integer from 0 to 4095. Setting an actuator to 4095 results in the maximum voltage being applied to that actuator. This voltage is determined by the voltage limit potentiometer on the high voltage power supply and is nominally set at 100V.

2.3. The Zero Voltage Surface of the DM

The most rudimentary test we conducted was to measure the optical quality of the mirror surface with no voltage applied to any of the actuators. Using a Zygo Corporation interferometer we measured the surface images shown in Figure 3.

The peak-to-valley value of the surface is 0.5728 μm with an rms of 0.0753 μm in the first image taken on 21 September 1996. The second image, taken on 26 September 1996, after five days of moving the actuators in the tests described in §3, has a peak-to-valley value of 0.5747 μm and an rms of 0.0729 μm. These values differ curiously from those reported by Xinξtics, who reported similar numbers but with units of λ rather than μm. We do not understand this discrepancy. (See §2.1)

There was a minor change in the zero voltage shape of the DM over this five day period of testing. Figure 4 illustrates this best: It shows the difference of the two images in Figure 3. The peak-to-valley in this difference image is 0.0957 μm, with an rms of 0.0441 μm. The banded structure in this difference is an artifact of the Zygo interferometer and should contribute fairly negligibly to the peak-to-valley value, though it corrupts the rms value. This change in the zero voltage offsets of the actuators may be due to "creep" or hysteresis in the actuators.[2]

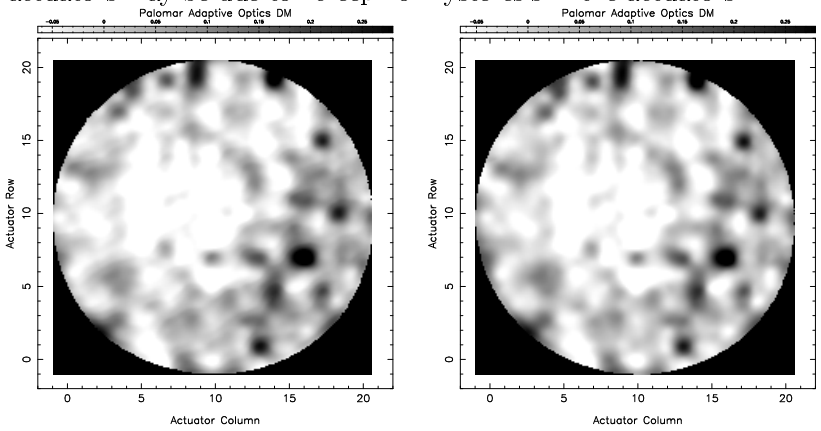

Fig. 3.— The image on the left shows the surface of the DM with no voltages applied to any of the actuators. The image on the right shows the same such surface, but it was made five days later after various tests described in §3. The bar at the top of each plot explains the greyscale and the units are in μm. The x and y axes indicate the positions of the actuators behind the mirror surface.

Hysteresis refers to the effect by which an actuator will move a slightly different amount for the same change in voltage depending upon whether it is approaching the final voltage from above or below. It is a non trivial task to adjust the control loop to account for this non-ideal behavior because it would require an examination of the recent movement history of the actuator to determine the next movement. In principle, though, one could take the hysteresis into account when controlling the DM. We found in our tests described in §3.3 that the actuators in this DM have hysteresis below 0.5% at room temperature. That is, the deviation from the ideal behavior is less than 0.5%.

Creep is a different effect whereby the actuator will continue to move

[2]The phenomenon of creep and hysteresis in PZT and PMN actuators is well-known. See, for example, "Effect of PZT and PMN Actuator Hysteresis and Creep on Nanoindentation Microscopy" by S. M. Hues *et al.* in *Review of Scientific Instruments*, Vol. 65, pg. 1561 (May 1994).

slowly after its initial very fast movement. We have measured the time constant for creep in this DM's actuators. This is discussed further in §4.1.

One of the primary advantages of using PMN actuators as opposed to the more common lead zirconate titanate (PZT) actuators is that PMNs typically have far less hysteresis. In addition they tend to have fairly constant 0V positions over significant aging periods.[3] It has, however, been known for quite some time that the response of PMN actuators is rather temperature dependent.[4] We will have to test this later, once the AO system is operational and on the telescope where the temperature varies significantly.

3. Tests of the Static DM

The tests described in this section were all conducted on time scales much greater than a few seconds. We believe that this timescale is much greater than the timescale in which the DM settles into a "semi-permanent" state after it has been given a command. Confirmation of this belief came only after we had carried out the dynamic tests described in §4.

3.1. Surface Control

The first set of active tests (meaning that voltages were applied to some or all of the actuators) on the DM demonstrated that we could control all of the 241 actuators. These tests were meant to show that the electronics were functioning as intended and that all of the actuators worked, at least on long timescales. This was accomplished by applying patterns to the actuator voltages so that we could see with the Zygo interferometer that every actuator was functioning. We also applied tilts in x and y over the whole mirror surface, manipulated columns and rows and introduced a focus term to the actuator voltages, to produce a roughly parabolic shape

[3]L. E. Cross et al. in *Ferroelectrics*, Vol. 23, p. 647 (1980)

[4]K. Uchino in *Ceramics Bulletins*, Vol. 65, p. 647 (1986)

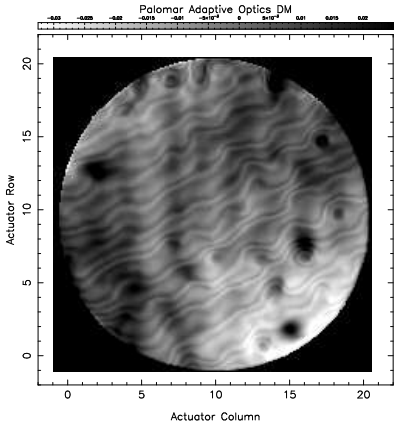

Fig. 4.— The difference of the two images in Figure 3. The scale at the top is in μm.

in the DM.[5] All of these tests were successful, and they were repeated at various times between September 1996 and February 1997 to confirm that all the actuators and the electronics were working properly. At the end of February 1997 all of the actuators were functioning and were all still attached to the back of the mirror.

3.2. Influence Functions

Knowledge of how an actuator deforms the mirror surface is critical for good wavefront correction. Most correction algorithms currently in use assume that the shape of the mirror between two actuators is simply the slope that is a result of the difference in position of each of the actuators. In practice, however, there is an appreciable influence that a given actuator has on the mirror surface at and beyond the adjacent actuators. We performed tests to determine this influence. The influence function of an actuator is the actual surface deformation caused by the movement of that actuator.

[5]The shape was rough because the gains and offsets of the actuators were unknown, so the voltages were set as though all of the actuators had zero offset and identical gain. See §3.4 for a more detailed discussion of gains and offsets.

To determine the influence function of a given actuator we applied a voltage to that actuator. Using the Zygo interferometer we then recorded the surface map and subtracted the zero volt image taken most recently. This provides an image of the actuator pushed up as though the rest of the mirror were perfectly flat. An example of this is shown in Figure 5. The image in this figure is an expanded view of the center of the mirror. The actuator in the 10th column and 10th row has been pushed up to a setting of 1024 or $^1/_4$ of the maximum voltage. Figure 6 shows the same data but in horizontal line cuts through the image.

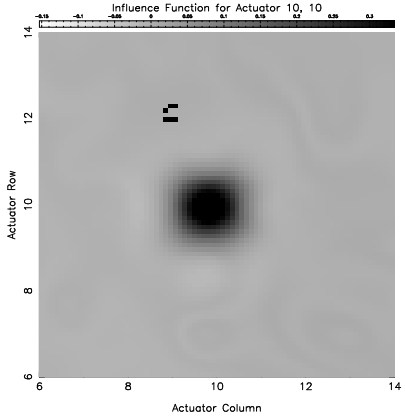

Fig. 5.— Surface map of the central region of the mirror with a setting of 1024 applied to actuator 10, 10 (using the coordinate system defined in Figure 2). The zero voltage offsets (as shown in Figure 3) have also been subtracted from this map to show the movement of the actuator alone. The peak-to-valley value for the mirror in this case is 0.443 μm. There are a few bad data points near the top of the plot.

Upon examining many plots such as those in Figures 5 and 6, it was determined that the influence functions do not vary in shape from actuator to actuator by more than ∼5%. (The gains of the actuators do, of course, vary, but that is discussed below in §3.3.) According to Xinξtics, this variation is probably due to small differences in the placement of the actuators behind the mirror facesheet, or perhaps variations in the thickness of the facesheet.

The most apparent feature of the influence function is that it has a rounded square appearance when viewing the mirror face-on (see Figure 5). This is because the actuators are arranged in a square grid pattern. In the following table is a quantitative assesment of the influence function. This represents an influence function matrix oversampled so that the spacing between successive matrix elements is half the interactuator distance. The upper left corner of the matrix represents the actuator whose influence function is represented. Every other row corresponds to the next row of actuators on the face of the mirror. The influence function is symmetric about the central actuator.

The meaning of the numbers in the "Average Influence Function" is that the actual difference in position between nearby actuators will be adjusted by the fraction in the matrix. For example, if actuator (x, y) is commanded to move 0.700 μm further than actuator (x + 1, y) the actual difference in the positions will be $(1.000 - 0.110) \times 0.700$ or 0.623 μm. Similarly, if actuator (x, y) is commanded to move 0.700 μm farther than actuator (x + 2, y), and actuator (x + 1, y) is commanded to be at the same height as (x + 2, y), then the actual difference in height between actuator (x, y) and (x + 2, y) will be $(1.000 + 0.014) \times 0.700$ or 0.710 μm.

The Average Influence Function Measured every half interactuator distance						
1.000	0.602	*0.110*	-0.024	*-0.014*	0.001	*0.000*
0.598	0.389	-	-	-	-	-
0.110	-	*0.018*	-	-	-	-
-0.025	-	-	-0.001	-	-	-
-0.015	-	-	-	*0.000*	-	-
-0.001	-	-	-	-	0.000	-
0.000	-	-	-	-	-	*0.000*
Values at neighboring actuators are italicized						

3.3. Linearity Considerations

An important aspect of the DM performance is whether the actuators behave linearly and whether the influence functions of neighboring actuators combine linearly.

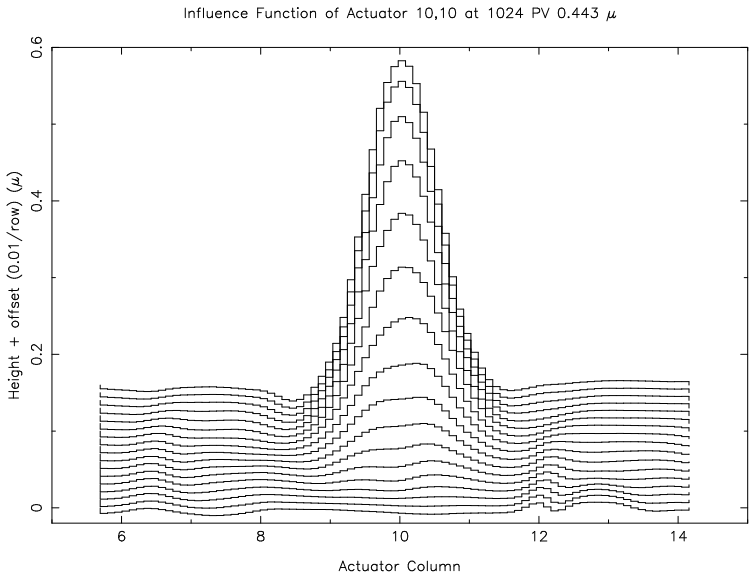

Fig. 6.— A stack of horizontal cuts through the bottom half of the surface map in Figure 5. Each horizontal cut is approximately one tenth of the interactuator distance, so the top plot corresponds to the tenth row of actuators on the mirror, while the bottom plot corresponds roughly to the eigth row. Each successive line plot is shifted up by 0.01 μm to make all of the horizontal cuts simultaneously visible.

Figure 7 shows the stroke of each of the actuators in the mirror as a function of actuator setting. (Only $^3/_4$ of the full range of settings are shown in this plot.) Several pieces of important information can be derived from this plot. First one can easily recognize that for actuator settings above 1300, the gain for any given actuator is constant. To make this more quantitative, we represent the actuator height as a function of the setting in the following form:[6]

$$H(s) = H(0) + as + bs^2$$

Here s is the setting, H is the height of the actuator, a is the so-called "linear gain" of the actuator and b is the "quadratic gain" of the actuator. One finds then by least squares fits to the data that the ratio b/a is on average 0.006, with an rms of 0.001. This shows the large relative importance of the so-called "linear" behavior over the quadratic behavior in the part of the setting range above 1300.

This linearity permits controlling the DM accurately with a linear algorithm in the range of settings above 1300. In that range, the total stroke has an average value of 4.2 μm, which corresponds to 8.4 μm of wavefront. The hysteresis we measured for the actuators was less than 0.5%.

The effect of moving adjacent actuators was investigated to determine whether the influence functions would simply add to produce the final surface on the DM. The assumption that the influence functions simply add seems to be correct to better than the 1% level.

3.4. Calibration and Precise Control of the Surface

Accurate operation of the deformable mirror requires a knowledge of the gain of each individual actuator. By gain, I mean the multiplicative constant relating the actuator setting to its actual stroke in physical units (in this case μm), assuming the linearity discussed above. With this knowledge the wavefront correction will be made more efficient.

Anand Sivaramakrishnan and I have created a scheme by which one can determine the gains of all of the actuators in an automated manner in a

[6]PMN actuators are known to behave quadratically in the sense described here over their full range, while PZTs behave more linearly. See, for example, J. A. Gallego-Juarez in *Journal of Physics E*, vol. 22, p. 804 (1989).

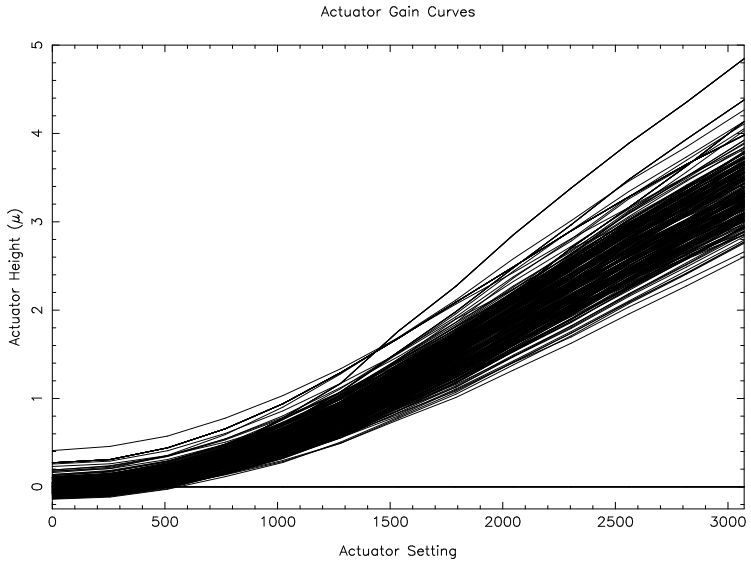

Fig. 7.— This plot shows the gain curves for all of the actuators. Plotted on the ordinate is the stroke of the actuator in μm and on the abscissa is the digital setting of the actuator in the electronics. Only $^3/_4$ of the full range of settings are shown here.

very short period of time, estimated to be less than two minutes under full automation. The utility of this approach is that one can quickly re-calibrate the DM at any time during an observing run on the telescope. Although one would not expect the gains to change markedly during an individual observing run,[7] with this procedure implemented, one could in theory calibrate the DM many times a night. This allows for a quantitative assessment of how much the gains actually do change on many timescales. The scheme is fully described in a separate report entitled "DM Calibration." This report can be found at http://astro.caltech.edu/~bro and is part of Sivaramakrishnan & Oppenheimer (1997, in preparation). Here I simply present the results of the scheme.

The scheme is iterative and attempts to flatten the mirror surface as accurately as possible. It converges extremely rapidly, especially if the initial guesses for the gains are nearly correct. In six iterations during the first trial of this scheme, we were able to flatten the mirror's controllable surface to an rms figure of 19 nm with a peak-to-valley value of 171 nm. This requires a knowledge of the gains to 3 significant figures. Figures 8 through 11 demonstrate this result by showing the mirror surface at each iteration. Once this is fully automated, we believe we can do significantly better than this value of 19 nm rms. The accuracy of the scheme is only determined by the precision of the DM surface maps and how accurately the actuators can be controlled, which is nominally on the order of 5 to 10 nm.

3.5. The Effect of Zener Diode Protection of the DM

The electronics rack which Dean Palmer designed and constructed for controlling the DM has a built in safety feature. The electronics will not permit a voltage difference of greater than 27.5V between two adjacent actuators, or 55V between two actuators separated by one in between. This is achieved with a grid of double zener diodes connecting all the actuators.

[7]The AO team for the Keck observatory, at CARA, who are using an almost identical Xinetics mirror, have reported that the actuators used in this type of mirror are rather sensitive to temperature variations. This might prove to be an important reason why one would want to calibrate the DM several times in a given night. See §5.1.

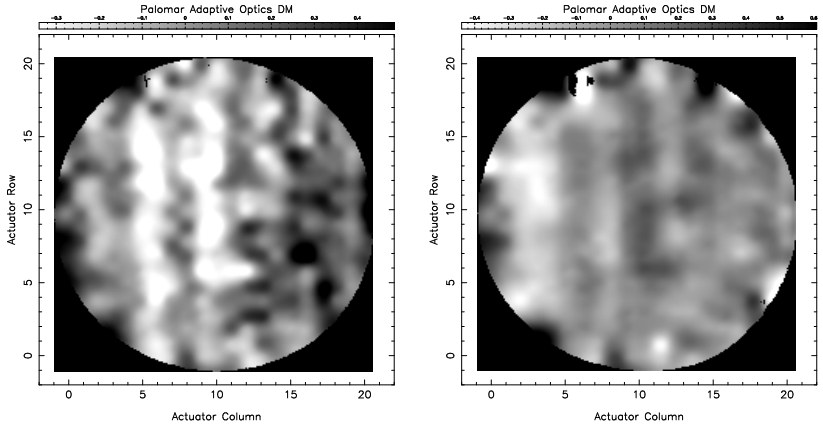

Fig. 8.— The image on the left shows the surface of the DM in the first step of the flattening scheme. The peak-to-valley (PV) value in the inner controllable part of the mirror is 1.083 μm and the rms of that region is 0.189 μm. The right hand figure shows the DM surface after the first iteration (PV = 0.656 μm, rms = 0.124 μm).

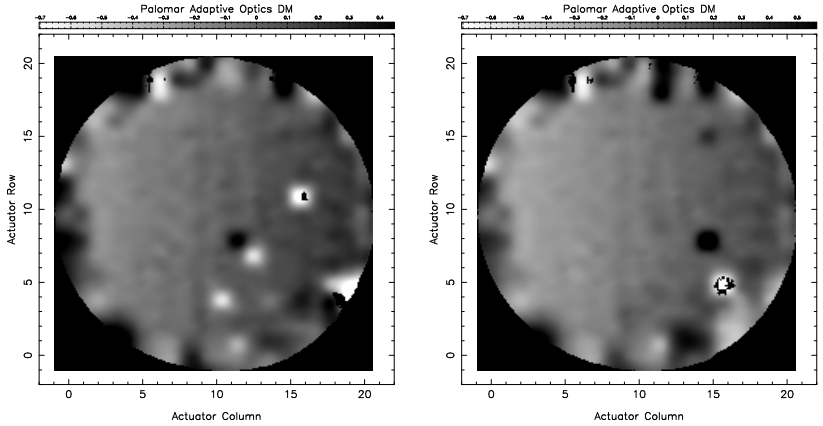

Fig. 9.— The image on the left shows the surface of the DM after the second iteration of the flattening scheme (PV = 1.518 μm, rms = 0.093 μm). The right hand figure shows the DM surface after the third iteration (PV = 2.159 μm, rms = 0.092 μm). In these images one can easily see that there are several actuators with far greater errors than the majority of the mirror. These are actuators whose gains were not well determined at these iterations.

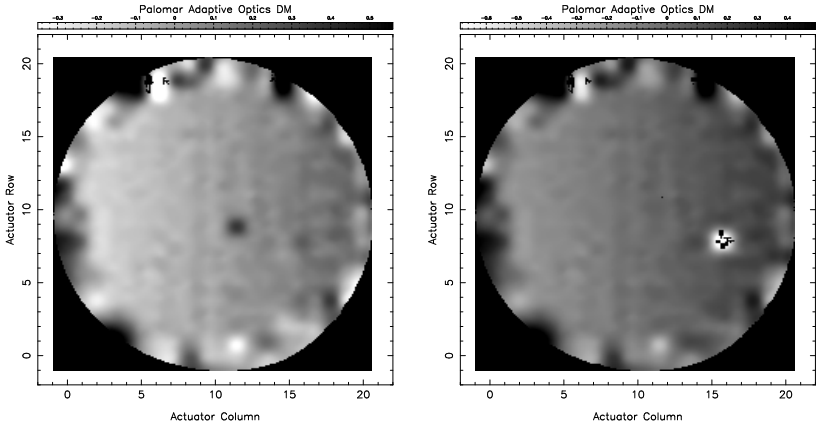

Fig. 10.— The image on the left shows the surface of the DM in the fourth iteration of the flattening scheme (PV = 1.282 μm, rms = 0.066 μm). The right hand figure shows the DM surface after the fifth iteration (PV = 0.436 μm, rms = 0.031 μm).

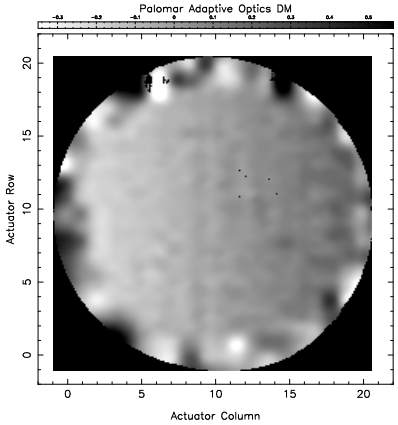

Fig. 11.— The final surface of the DM after 6 iterations of the flattening scheme. The inner controllable surface has PV = 0.171 μm and rms = 0.019 μm. The edges of the mirror surface are not flat. This is because that portion of the surface is controlled by the slaved actuators. These actuators simply recieve 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 fits easily within the flattened part of the surface.

The documentation from Xinξtics Inc. states that the maximum difference in height between two actuators should not exceed 2 μm, to avoid risking damage to the bonds holding the actuator heads to the mirror facesheet. Setting the maximum voltage difference to 27.5V places a more stringent restriction on the actuator behavior. This should only permit roughly 1 μm of position difference between adjacent actuators.

4. Dynamic Tests

In a dynamic test one moves the mirror while measuring its configuration many times per second in an attempt to resolve the temporal response of the actuators. The purpose of such tests is to determine how quickly a given actuator actually reaches its final position and whether one can operate the mirror at the desired closed loop frequency of 500 Hz.

The dynamic tests were conducted with an optical heterodyne interferometer. This system, shown in Figure 12, is capable of measuring a single optical path length at a rate of 10 kHz. Since it can only measure one path length, we can only measure the movement of one point on the mirror at any given time. During these tests the laser beam was positioned at a number of locations on the mirror to ascertain whether different actuators respond differently. There are no measurable differences in the temporal characteristics of the ten different actuators measured.

4.1. Time Resolved Motion of Individual Actuators

Figure 13 shows the mirror displacement as a function of time for a commanded motion of 1536 units. All of the actuators were commanded to move this amount and then back to the starting position at about a 10Hz rate. The position of the mirror was determined at a rate of 5 kHz. The initial and final positions of the actuators were chosen to be well within the linear part of the gain curves of the actuators (see Figure 7). In later tests we ran the mirror at actuator update rates up to 1.5 kHz, to determine whether the response curve changes based on this frequency. It does not.

The plot in Figure 13 has four important components. These are labelled with letters A through D. A corresponds to the movement of the actuator from the moment it receives a new voltage from the electronics. The rise time here is proportional to the distance the actuator has to move. It turns out that the slope of part A is the same regardless of the movement length. The slope is found to be the following (an average for several actuators):

$$v_{act} = 2.67 \times 10^{-4} m/s$$

This is the speed of the actuator in meters per second. What this immediately reveals is that at a 500 Hz update rate, the maximum distance that an actuator can move is 0.53 μm. Commanding the actuator to move further than this is possible at a 500 Hz update rate, but the actuator will only have time enough to move 0.53 μm before it is commanded to move again. This actuator speed is determined by the capacitance of the actuator and the current used to drive it. The voltage on the actuator will obey the simple relation $dV/dt = I/C$. In this case I is 6.4 mA, and C is 1.6 μF.

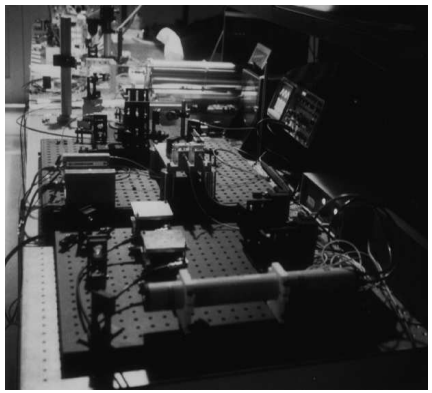

Fig. 12.— The heterodyne interferometer used in the dynamic tests of the DM. The laser is the long tube at the bottom of the photograph. Its beam is split into the two acousto-optical mixers, which are the two square boxes just beyond the laser. The light is shifted by several MHz and then recombined in the optics just beyond the mixers. The interferometer is located at the end of the black optical pad, just in front of the silver tube near the top of the picture. The DM is to the right, in front of the silver tube. Its plastic cover is on in this picture. The two rectangular boxes are the sensors, which feed analog output to a computer (not in the photograph).

(We did not actually remeasure the capacitances of the actuators. This number comes from the Xinξtics report.) This results in a slope of 3989 V/s which translates roughly to 2.11×10^{-4} m/s. Thus, the expectation is slightly slower than the measured value, but this is a rough calculation. In the first set of specifications, the actuator capacitance was set to be 0.25 μF. With the present electronics rack, this capacitance would have yielded a speed of 1.3×10^{-3} m/s, and a maximum stroke at the 500 Hz update rate of 2.5 μm. This value was chosen in the design of the electronics to have quite a margin beyond the needs of the AO system. As a result of this, the fact that the actuators have much larger capacitance has not had a terrible impact on the operability of the system: We expect that moves of 0.53 μm should be relatively rare except in very bad seeing conditions.

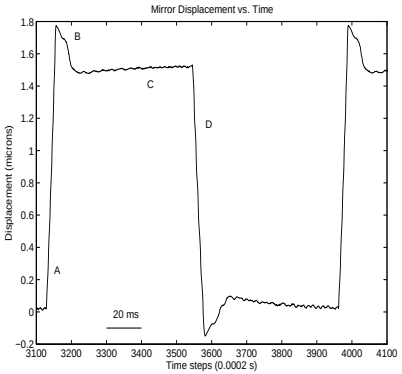

Fig. 13.— The displacement of a single point on the mirror as a function of time. See the text for a full discussion.

Parts B and C of the plot represent the compensation of the overshoot of the actuator. The actuator reaches its commanded position and then overshoots in part A, by approximately 13 to 14 % of the commanded distance. The reversal of the overshoot takes place in two separate phases. The first is a short phase which overcompensates for the overshoot, but only to within 3 % of the final position. This is designated B in the plot. This movement always seems to take the same amount of time, no matter what distance is commanded to the actuator. B has a time constant of 9.0 ms. Part C represents the final settling of the actuator to its commanded position. This is the "creep" effect mentioned at the end of §2.3. It appears to be an exponential decay with a time constant on the order of half a second. However, within the 60 ms of part C shown here, the actuator reaches the final position to within 0.5%.

Finally part D corresponds to the next voltage impulse from the electronics. It has the same slope magnitude as part A.

An important consideration is how long it takes for the settings to take effect. How much time elapses between the sending of a command from the computer to the actual physical motion of the actuators on the mirror? In the test set up, we issued commands from the VME crate to the DM electronics rack via a coaxial taxi line. Determining the time between the issuing of the command and the movement of the first actuator is difficult, but appeared to be approximately half a ms. In contrast, it was easy to determine the time between setting the first actuator and setting the last actuator on the mirror. This is 1.0 ms. These measurements were made using an oscilloscope attached to a diagnostic feed on the DM electronics and the output of the heterodyne interferometer's D/A. Clearly the dominating timescale is the rise time of the actuator, which will average 3.1 ms for a 1 micron average movement. (That corresponds to fairly bad seeing conditions at the telescope, so one can expect better performance, of course, during calmer atmospheres.) The accuracy of the measurement presented here is ± 0.1 ms.

5. Unresolved Issues

5.1. Temperature Dependence of DM Performance

One issue which these tests were unable to resolve is whether the DM performace changes as a function of temperature. The CARA Adaptive Optics group, who are building the system for the Keck telescope, report that their 37 actuator Xinξtics DM, which is identical to our DM except in size, has greatly degraded performance in the lower temperatures characteristic of the Mauna Kea summit. All of our tests were done at the lab temperature which is approximately 20° C. Some of these tests can be

redone once the entire optics train is complete. We can use the stimulus to fill the capacity of the Zygo interferometer and redo all of the static tests at the range of temperatures at the top of the mountain. Unfortunately the dynamic tests will not be repeatable on the mountain because of the lack of availability of the heterodyne interferometer.

More specifically, the CARA group report that the hysteresis and linearity of the actuator stroke functions degrade considerably at low temperatures. They also report that the equivalent of part C in Figure 13—the creep of the actuators—has a much longer timescale at low temperatures. Unfortunately we will not be able to remeasure the plot in Figure 13 at lower temperatures.

6. Operation Recommendations

6.1. Procedures for Turning the DM On and Off

I present a recommended procedure for turning the DM on and off. Although the DM is protected by the zener diode grid at all times, it is possible that a zener diode could go bad. In that case, using this procedure will prevent any damages from happening during the beginning and end of an observing session.

Power up procedure:

0. Set the voltage limiter on the high voltage power supply (HVPS) to zero, if it is not already there.
1. Boot the VME computer (the interface computer).
2. Power up the DM electronics rack.
3. Start the dm_srv software on the VME computer.
4. Send a "zero all actuators" command to the electronics.
5. Turn on the HVPS and increase the voltage limiter to 115V.

Power down procedure:

5. Send a "zero all actuators" command to the electronics.
4. Reduce the voltage limiter to zero and turn off the HVPS.
3. Terminate the dm_srv software on the VME computer.
2. Turn off the DM electronics rack.
1. Turn off the VME computer.