Analysis of Torus Mirror 18-64

Stephen C. Koehler
steve_koehler@securecomputing.com

January 2004

1 Introduction

This report documents my ongoing efforts to understand the optical quality of my 18" f/4.5 Torus mirror, based on a variety of optical tests. My general approach is to use the tests to cross check each other, and to use the unique strengths of each test for maximum effect. The end result is a rather complete picture of the optical characteristics of this mirror from different viewpoints, not all of them in agreement. I hope eventually to resolve the discrepancies. Until then, I will refrain from drawing final conclusions.

The report is organized by test method. Each quantitative method is compared against interferometry, which provides the most detailed model of the mirror's surface. I begin with a summary of results, which provides a road map to the rest of the paper. The individual sections explain my methodology and justify my conclusions. Some of the sections (especially Interferometry) are rather detailed. If the analysis goes beyond your patience or interest you may skip ahead to sections that are of more interest.

I have tried to make the figures in this report readable both on the screen and in a printed copy. However, there are some graphics that will always appear better on the screen. This includes a few graphs where I use color to make subtle distinctions and some gray images where the contrast is better when viewed on the screen. If you are having trouble understanding a graphic in a printed copy, I recommend looking at that graphic on the screen, instead.

1.1 Background

In the summer of 2003 I took delivery of Torus mirror 18-64. This mirror replaced my original Torus mirror (18-27), which unfortunately was misplaced by the coater after I had sent it in for some rework. As a favor to me, because of my interest in optics, James Mulherin (President/Master Optician at Torus/OMI) supplied me with an unusually large amount of documentation on 18-64, including seven Ronchi images, six Foucault (knife test) images, and twenty interferograms and associated analysis. Table 1 shows the overall specifications reported by OMI. These specifications

Table 1: Mirror Specifications

P-V Wavefront Error	RMS Wavefront Error	Strehl Ratio
.122	.017	.989

indicate a very good mirror (provided the figure is smooth and there is no turned edge), but they don't tell the nature of the remaining errors on the mirror. I wanted to know more.

One of the first things I did after I received the mirror was to perform a Foucault test to confirm the basic specifications. My early Foucault results were much worse than the interferometry, having a Strehl ratio in the range .80 to .85. In subsequent tests I have gotten a Strehl ratio as high as .92, but not anywhere near the .99 of interferometry. This discrepancy remains the biggest mystery in my analysis.

Next, in looking for some way to perform an independent test, I found a local amateur telescope maker, Dale Eason, who has an automated Foucault platform (Robo-Foucault). Dale achieved repeatable results on my mirror with a Strehl ratio in the .99 range, which agrees very well with interferometry. This is very encouraging, but does not, by itself, prove that the interferometry is accurate, because Dale's tester and software have not been calibrated against other optics with known specifications.

In November 2003 I had the opportunity to spend a day at James Mulherin's shop in Iowa City with my mirror. We performed a null Ronchi test, null knife edge test, interferometry on the coated mirror, and a Foucault test. The 20 interferograms of the coated mirror present a rare opportunity to evaluate the change in mirror figure due to coating. In the null knife-edge test we saw some fine details of the mirror shown by interferometry.

Testing on my mirror is ongoing. I would especially like to understand why the Foucault test does not yet agree with interferometry. I will also strive to get more information out of the star test and the null Ronchi test (with star source) when I take the scope outside.

1.2 Zernike Polynomials

In this analysis I make extensive use of a Zernike polynomial representation of surface or wavefront errors. This is a mathematical way to represent a 2-dimensional circular surface (ideal for optics), and provides a basis for comparison of data from different optical tests. Understanding the basic concepts of Zernike polynomials will help in understanding my analysis.

The Zernike polynomials are an infinite set of functions over the unit circle with the property that any well-behaved surface can be approximated by a weighted sum of these functions (i.e., by adding together varying amounts of the Zernike polynomials to approximate the target function). The weights are called the Zernike coefficients.

A Zernike representation works well for modeling optical aberrations for the following reasons:

- The low order polynomials correspond closely to the typical aberrations seen in lenses and mirrors (e.g., coma, astigmatism, and various orders of spherical aberration).
- The magnitude of individual aberrations, as well as RMS error and Strehl ratio, can be calculated directly from the Zernike coefficients.
- Because the contribution of various aberrations are separated out, individual aberrations can be removed from the modeled surface without affecting the other aberrations. (This is useful in removing alignment errors and test stand astigmatism from interferometry results.)

The Zernike representation of a surface is found by performing a least squares fit to find the coefficients that provide the best approximation of the target surface. For practical reasons, only a finite set of Zernike polynomials are used, the specific number depending on the application—the more data available to perform the fit, the more terms can be used.

Table 2: Low-Order Zernike Polynomials

n	m	No.	polynomial	description
0	0	0	1	piston
1	1	1	$\rho\cos\theta$	tilt
		2	$\rho\sin\theta$	
	0	3	$2\rho^2-1$	defocus
2	2	4	$\rho^2\cos 2\theta$	astigmatism
		5	$\rho^2\sin 2\theta$	
	1	6	$(3\rho^2-2)\rho\cos\theta$	coma
		7	$(3\rho^2-2)\rho\sin\theta$	
	0	8	$6\rho^4-6\rho^2+1$	spherical aberration
3	3	9	$\rho^3\cos 3\theta$	trefoil
		10	$\rho^3\sin 3\theta$	

There is no general agreement about the numbering of Zernike polynomials. Table 2 shows the low order Zernike polynomials using the numbering defined in *Basic Wavefront Aberration Theory for Optical Metrology*[9]. In this analysis I identify sets of Zernike polynomials of various sizes by the "order" of polynomials, which corresponds to the value of n in Table 2. For example, the polynomials up to order 2 consist of the first 9 rows in the table. In this analysis I use Zernike polynomials of orders 7, 9, and 12, depending on the amount of available data.

The polynomials with $m = 0$ are called the symmetric polynomials, because they have a radial symmetry. Spherical aberration is an example of a radially-symmetric aberration. Other polynomials (with $m > 0$) come in sin/cos pairs, do not have radial symmetry, and are known as the asymmetric polynomials. Primary astigmatism and coma are examples of asymmetric aberrations.

The symmetrical Zernike polynomials correspond to the various orders of spherical aberration. Notice that there is only one symmetric polynomial for each order. Therefore, I identify a specific kind of spherical aberration by using the order. Thus, low-order spherical aberration is order 2, and the next higher level of spherical aberration is order 3. This notation is non-standard, but it keeps the analysis self-consistent.

A surface modeled by Zernike polynomials usually behaves well (i.e., corresponds to reality) in the interior of the data, but may behave strangely at the periphery. This doesn't much affect statistical metrics, such as RMS error or the Strehl ratio, but can have a larger effect on P-V error. Please keep this in mind when viewing the graphs plotted from Zernike models. I have left the edges in the surface models, but I do not trust that they are accurate. This is especially apparent in graphs of multiple surface plots, some of which turn up, rather than down, at the edge.

There is a lot of information on Zernike polynomials available on the Web. For a general discussion, see [9] or [10]. For an interactive demonstration of Zernike polynomials, see [11].

1.3 Method of Comparing Test Results

This analysis compares quantitative results from very different measurement techniques. Interferometry maps the entire 2-dimensional surface of a mirror, so it makes use of symmetric and asymmetric Zernike terms. In contrast, the Foucault test measures only a single diameter of the mirror, so only the symmetric terms are applicable. (The Foucault test does not directly "see" asymmetric aberrations, such as astigmatism, although it is certainly affected by them to some degree.)

The common ground for comparison is the symmetric terms, i.e., the various orders of spherical aberration. For this reason, all test results that can be quantified are reduced to the symmetric Zernike terms for comparison.

My primary method for comparing symmetric terms is by plotting the surface profile across a radius of the mirror. I annotate the horizontal scale of these plots in millimeters of radius and the vertical scale using both nanometers and fractions of a wavelength. See Figure 1 for an example. If you are trying to read wavefront deviations off these plots, remember that surface error is half of wavefront error, so you would need to multiply the result by two.

1.4 Test Equipment

It is useful to consider which tests share which equipment, so that independence (or lack of independence) of the tests can be ascertained. It is also necessary to understand which pieces of test equipment contribute errors to the results, and how large these contributions might be.

Robo-Foucault, Foucault, Ronchi and the star test are completely independent of themselves and other tests, because they share no common test equipment. The null Ronchi test and interferometry performed at OMI share a large optical flat. Interferometry on the coated mirror and uncoated mirror used different interferometer reference elements (the heart of the interferometer that produces the comparison wave). The most likely source of common error is from the optical flat.

The interferometry on the uncoated mirror used a reference element from Tucson Optical Research Corporation, rated by independent interferometry to $1/20\lambda$ accuracy across its f/2 surface. Mulherin estimates it is accurate accurate to about $1/100\lambda$ P-V within it's f/4.5 central portion. The second reference element has similar specifications.

The optical flat used in interferometry and the null Ronchi test is rated to $1/10\lambda$ P-V across its 22" diameter surface. According to Mulherin, this is a sphere of very long radius of curvature, which primarily affects the focus position of the test being performed. Figuring an optic against this flat imparts a very small amount of spherical under-correction. I have not yet proved to myself the size of the SA introduced by the flat, but I believe it to be very small.

1.5 How the Figures Were Made

All the figures in this report (except the photos and the null knife-edge images) were created using software I wrote in the R language (http://www.r-project.org), based on previous work by Mike Peck. R a is statistics and graphical reporting language that provides a very good foundation for performing optical calculations and presenting the results. I have plans to make my core software routines available when I get time to properly document them. The null knife edge images (Figures 22 and 23) were generated by Jim Burrows' program Diffract from data prepared by my R programs.

2 Summary of Results

I present a brief summary of the conclusions of this paper up front, both to serve as a road map for the rest of the report and to save time for anyone who does not care about the details. The rest of the report explains my methods and justifies my conclusions.

There are four main characteristics of a mirror that describe the optical properties we care about.

1. The mirror should have an overall shape close to the ideal paraboloid.
2. The mirror should not have sudden changes in slope (i.e., it should be smooth and free of zones).
3. The mirror should not have medium or small-scale roughness.
4. The mirror should not have a significant turned edge.

This report makes use of both qualitative and quantitative test methods. Only characteristic 1 is amenable to quantitative analysis with the tools I have available, so 2-4 must be checked with qualitative tests. However, the qualitative tests are able to support the conclusions of the quantitative tests for 1.

Figure 1: Comparison of Surface Profiles

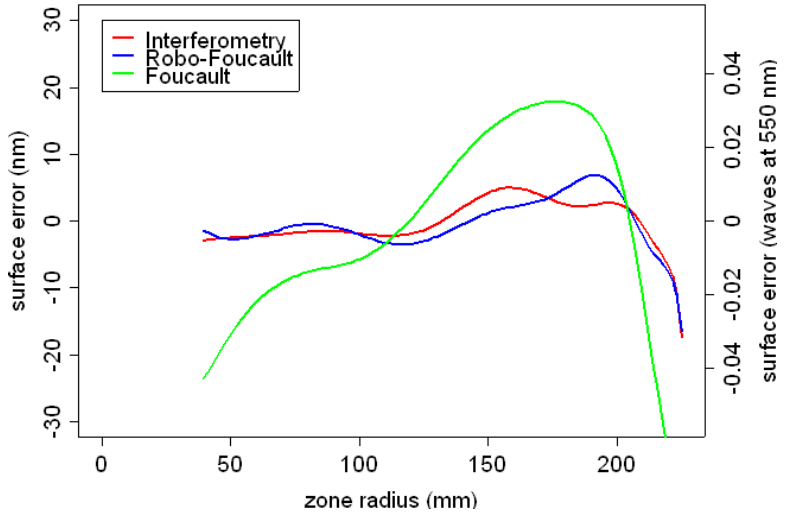

Figure 2: Comparison of Spherical Aberration

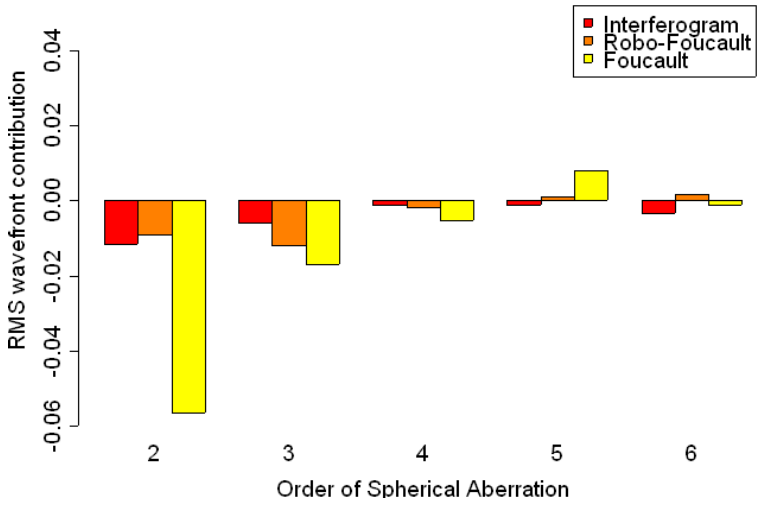

Let's start with the qualitative tests. Foucault, Ronchi and Null Ronchi tests show the mirror to have a smooth figure, to be relatively free of roughness, and to have a good edge. I have not done enough star testing, yet, to draw firm conclusions. I have noticed that turbulence spikes are fairly even on both sides of focus, and are not very prominent, which suggests the edge is good and the mirror is relatively free of steep-slope errors. I have no way to detect small-scale roughness (i.e., micro-ripple).

Of the quantitative tests, interferometry and Robo-Foucault agree to a remarkable degree. The null knife-edge test also confirms some of the fine surface details (mild zones) that are shown by interferometry. The Foucault test doesn't agree with these very well, but I believe the problem

Table 3: Comparison of Specifications

Test	P-V Error	RMS Error	Strehl Ratio
Original specifications	.122	.017	.989
Interferometry	.081	.013	.993
Robo-Foucault	.085	.016	.990
Foucault	.206	.057	.871

Figure 3: Comparison of Relative Transverse Aberration

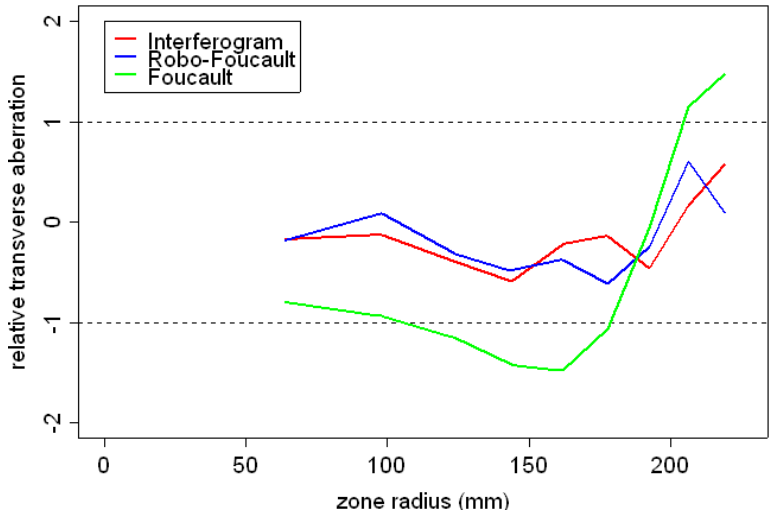

is my environment or technique. Figure 1 compares the quantitative results in surface profile. Figure 2 shows another view of the same thing by breaking out the individual contributions of each order of spherical aberration to the total RMS error.

Table 3 compares the three basic mirror metrics (P-V error, RMS error, and Strehl ratio) of all quantitative tests. At this point in my investigations, I trust the interferometry and Robo-Foucault more than the Foucault test, especially because of the details shown by the null knife-edge test.

In the surface profile comparison, you can see that there are mild zones on the mirror, particularly toward the outer edge. How severe are those zones? One way to compare them is to graph the relative transverse aberration (RTA), which uses ray optics to quantify how far light can be thrown out of the Airy disk. Figure 3 shows this comparison. The y-axis region between -1 and 1 delimits the edge of the Airy disk. Lines within this region keep light within the Airy disk. Lines outside this region show which parts of the surface have a slope that is too steep. For this graph, I reduced all data sets to the nine data points typical of a Foucault test. (It is also possible to compute a continuous RTA curve, but the results do not compare well to a discrete RTA curve.)

Robo-Foucault and interferometry RTA plots look very good, staying well within the Airy disk diameter. This is a good indication that the overall figure of the mirror is smooth, and that the slope of the surface is well under control. The Foucault test RTA plot shows more problems, with the curve deviating out of the Airy disk in two places.

Table 4: Foucault Test Images

image	description	image	description
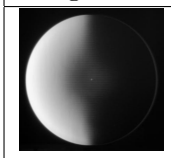	Center zone	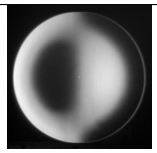	70% zone
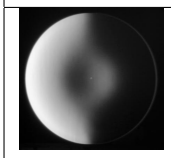	30% zone	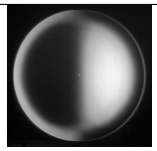	85% zone
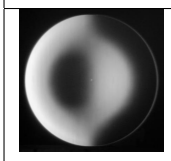	50%zone	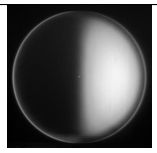	Edge

Table 5: Ronchi Test Images

image	description	image	description
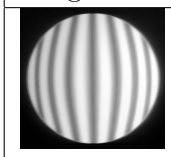	Seven lines, inside focus	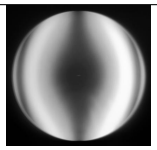	One line, inside focus
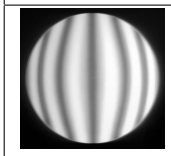	Five lines, inside focus	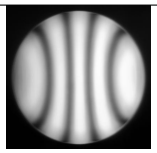	Four lines, outside focus
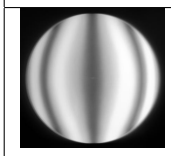	Three lines, inside focus		

One interesting result from interferometry is the effect of coating on the mirror. Even though the coating is relatively thick (about $1/4\lambda$), the before and after interferograms show no significant change in figure due to coating.

3 Ronchi and Foucault Images

I will start out the detailed analysis with an examination of the Ronchi and Foucault images that came with my mirror. Compared to most of the other tests, these are relatively straightforward to interpret without complex analysis.

James Mulherin sent me nine test images of my mirror, including six Foucault (i.e., knife test)

Figure 4: Detail of Five-Line Ronchi Image

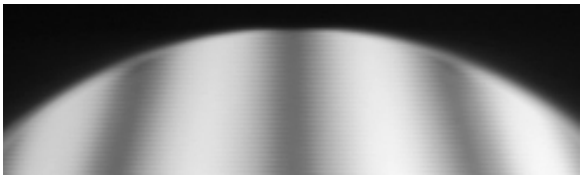

images at a range of zones and five Ronchi images. These images are very useful to assess some qualities of the mirror that do not show up very well in the quantitative tests, such as turned down edge (TDE) roughness, and zones. The Ronchi images were taken using a 60 line per inch grating.

Table 4 and Table 5 show small copies of the images. If you are viewing this in an Acrobat reader, you can click on the thumbnail images to see them full size in your web browser. The images are located at http://www.visi.com/~mkoehler/18-64/images/.

3.1 Turned Down Edge

Both Foucault and Ronchi images are used to assess turned edge. A moderately turned edge shows up in the Foucault images as an uneven brightness in the diffraction ring at the edge of the mirror when the knife positioned at the edge zone. The edge zone image shows an evenly bright ring. Visually, this ring is even brighter and narrower than in the photo.

In the Ronchi images, TDE shows up as a hooking of the lines at the edge of the mirror. The five-line inside focus image is especially good for assessing TDE. In analyzing Ronchi images for TDE, you have to be aware of diffraction effects, which cause the edge of the mirror to be show as a double or triple image offset to the side. This causes the dark Ronchi lines to end before the apparent edge of the mirror.

Figure 4 shows a close up of the top edge of the five-line Ronchi image. The lines show a dark thickening as they approach the edge. From my experience, this is a good sign—mirrors with clear turned edge do not show this. I do not see any hooking of the line at the edge. I believe the slightly wider gap at the edge toward the outside is due to diffraction effects. I would be happy to hear from anyone with a different opinion on this.

3.2 Roughness

Both Foucault and Ronchi images are capable of showing roughness. The center zone Foucault image is good for showing roughness in the center of the mirror. On Foucault images, you have to be careful to ignore vertically aligned striations, which are most likely artifacts of the vertically oriented knife edge used to make the shadows. You also have to ignore momentary turbulence, which shows up in some of the images (e.g., the one-line Ronchi image). Roughness shows up in Ronchi images as a mottling of the lines, and is more apparent when fewer lines are visible.

Overall, my mirror looks very smooth. I think I can see a bit of roughness here and there in the images, but it's hard to be sure. Visually, the mirror looks very smooth under the knife test.

3.3 Zones

The Ronchi images are very good at showing a smooth overall figure of the mirror. Zones, or quick changes in slope, show up as deviations in the lines from a smooth curve. The fewer the lines, the more sensitive the test. I have tried to see some the surface features I detected in interferometry, such as the broad hill at the 75% zone or the fall off at the edge, but I am not able to detect them in the Ronchi images. Overall, the mirror appears to be zone free.

4 Interferometry

This section analyzes the interferograms of my mirror independently from the results reported by OMI. This is an interesting exercise, because there is a wealth of information in the interferograms, beyond the basic mirror specifications. I particularly wanted to answer the following questions.

1. Are any significant symmetric or asymmetric features apparent?
2. Did the coating change the figure of the mirror?
3. Are the specifications reported by OMI (see Table 1) accurate?

I identify individual interferograms as follows:

- A1-A15 are the fifteen interferograms taken in orientation A.
- B1-B5 are the five interferograms taken in orientation B (rotated 90 degrees clockwise with respect to A).
- C1-20 are the twenty interferograms taken of the coated mirror in orientation A.

Please note that the 2D plots derived from interferograms are left in the same orientation as the original interferograms with the support sling at the top (i.e., the images are upside down). Also, I annotate many of these plots with the P-V wavefront error and Strehl ratio at the lower right of the plot. The P-V error is the smaller of the two numbers.

4.1 Derivation of the Asymmetric Model

There are a number of decisions to make in creating a Zernike polynomial model from interferograms:

1. What wavelength should the analysis use?
2. Is it justified to cancel primary astigmatism from the asymmetric terms to reduce the effects of test stand astigmatism?
3. Do you need to use all the asymmetric terms, or can you get by with just the symmetric terms (depends on symmetry of the mirror)?
4. What order of polynomials is appropriate for the asymmetric and symmetric models?
5. How much averaging is needed to reduce the random noise of air currents and vibration (i.e., how many interferograms do you need to take)?

Figure 5: Interferogram A1 Traced on Light and Dark Fringes

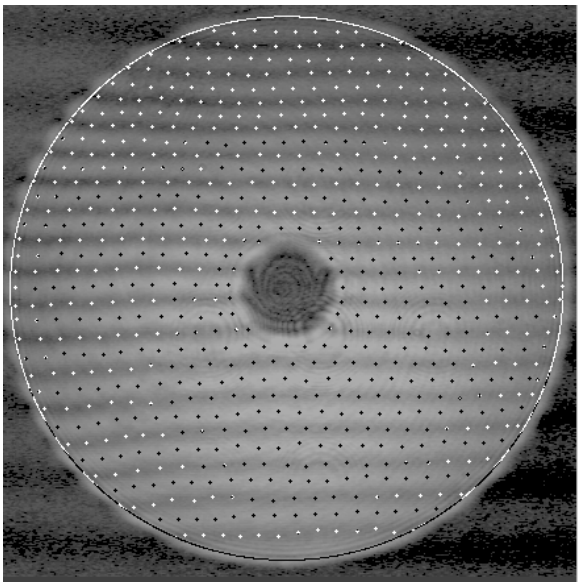

Figure 6: Direction and Magnitude of Astigmatism

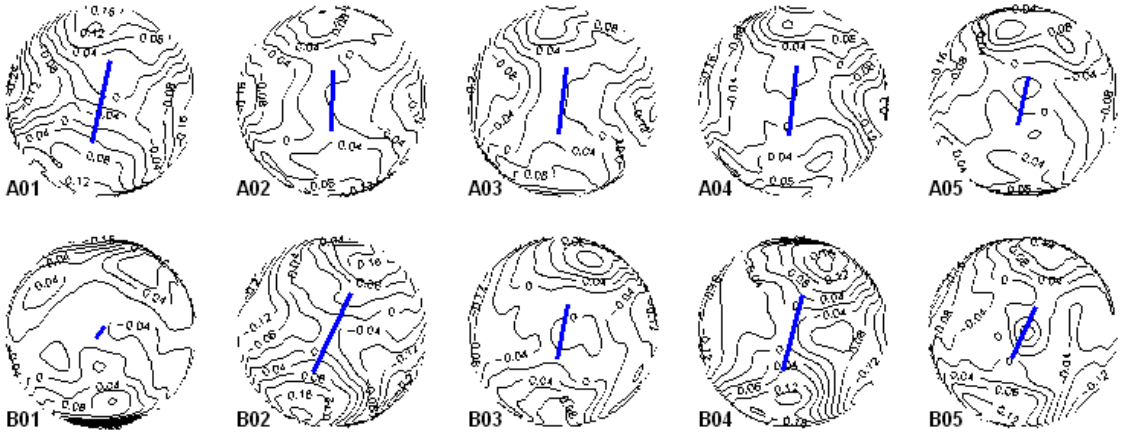

Starting with the raw interferogram images, I used Dave Rowe's freeware fringe analysis program FringeXP[6] to trace the interferogram fringes. I traced both light and dark fringes, which approximately doubles the number of data points taken on each image compared to tracing just dark fringes. Figure 5 shows an example fringe trace from interferogram A1, which resulted in 739 data points.

The interferograms were taken using a HeNe laser having a wavelength of 633nm (red light). However, it is usual when talking about wavefront error of an optic to use a wavelength of 550nm (yellow light). OMI analyzes interferograms at the original 633nm source wavelength. The difference between the two is slight (for example, a Strehl ratio of .980 might drop to .976). I chose to use use the more stringent 550nm wavelength in my analysis.

The result of fringe tracing is a list of coordinate sets for points along the fringes, one set for each fringe. You can think of the fringes as a contour map of a sloped surface. The fringe spacing indicates the degree of slope, which has to do with the tilt of the interferometer's reference element compared to the wavefront from the mirror. Once this tilt is mathematically removed, small undulations in the lines show aberrations on the optical surface (and lots of noise due to air turbulence and vibration).

The next step is to perform a fit of Zernike polynomials to the fringe data, and check whether any astigmatism present is due to test stand deformation. Figure 6 shows the angle and relative magnitude of astigmatism present on 10 interferograms (5 from set A and 5 from set B). As you can see, the astigmatism wanders around a bit in angle and magnitude, but stays oriented in the vertical direction when the mirror is rotated 90 degrees. The magnitude of astigmatism is about 1/4 wave P-V error, which is typical for this size and thickness of mirror. This amount of astigmatism swamps all other errors, and so it must be removed from the asymmetric Zernike terms before further analysis can be done.

Besides astigmatism, it is necessary to cancel tilt, defocus, and coma from the Zernike terms. Tilt is caused by the inclination of the interferometer's reference element with respect to the wavefront. Defocus and coma are due to alignment errors in the test setup.

Now that astigmatism can safely be removed, a decision can be made on the appropriate order of Zernike polynomials to use for an asymmetric model of the surface. I found that order 9 polynomials

Figure 7: Order of Asymmetric Zernike Fit

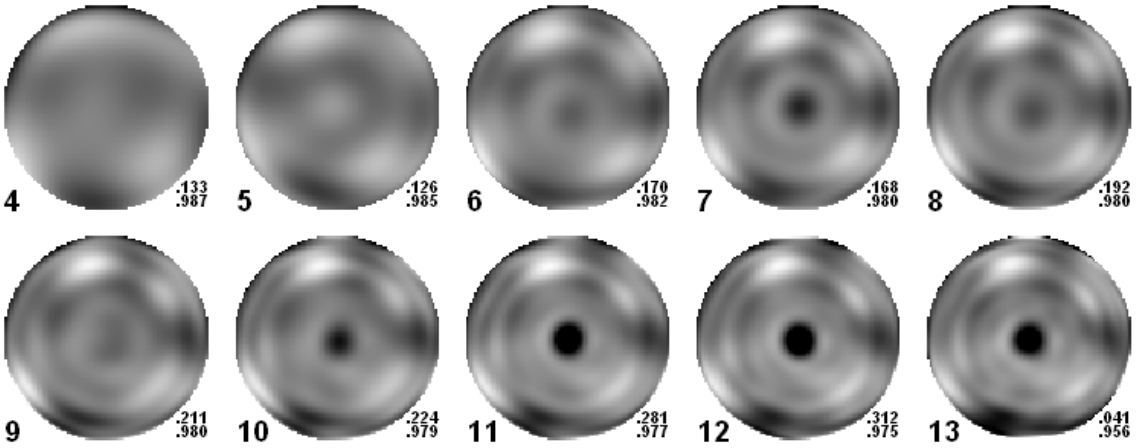

Figure 8: Wavefront Plots of Individual Interferograms

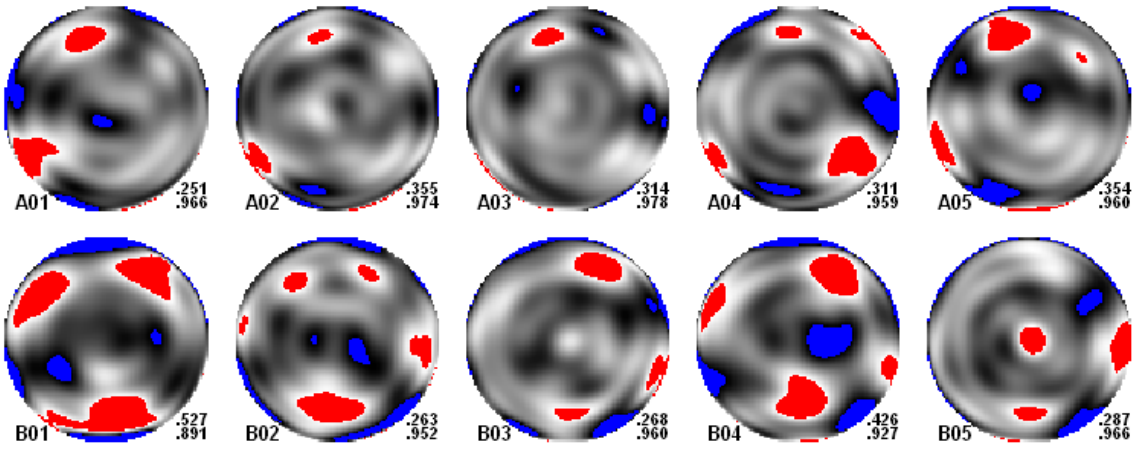

provides a good balance between bringing out more fine surface detail and over-fitting the data. To justify my choice, Figure 7 shows the result of fitting 4th order through 13th order Zernike polynomials to an average of 5 interferograms (A1-5). The individual images are annotated with the order of polynomials fit and the P-V error and Strehl ratio calculated from the wavefront model.

Note the following:

- A fit of only low order terms does not show much detail.
- The Strehl ratio drops as the surface is better fit by higher order terms.
- Above order 10, the central hole was allowed to "blip" down where no fringe data ties it down.
- Starting at order 7, two dark rings start to appear. I believe these are actual surface features, because they persist in the higher order fits.
- The Strehl ratio starts to drop rapidly after order 12, which probably means that the fit is being applied to too few data points.
- QuickFringe (used by OMI) uses a fit equivalent to my order 5 with one additional spherical aberration term from order 6 thrown in. Thus, it would look most similar to my order 5 image.
- The use of 9th order Zernike terms is probably a conservative choice: some surface details are seen, but the fit is below the point where Strehl ratio starts to degrade dramatically.

4.2 Averaging to Reduce Noise

It is interesting to observe how much noise is present in individual interferograms. Figure 8 shows a plot of wavefront error with astigmatism removed. Red and blue areas highlight portions of the surface where the wavefront deviates by more than 1/20 wave from the average position. In other words, the gray areas of the figures are contained within a 1/10 wave envelope. This diagram is

Figure 9: Effect of Averaging Sets of Six Interferograms

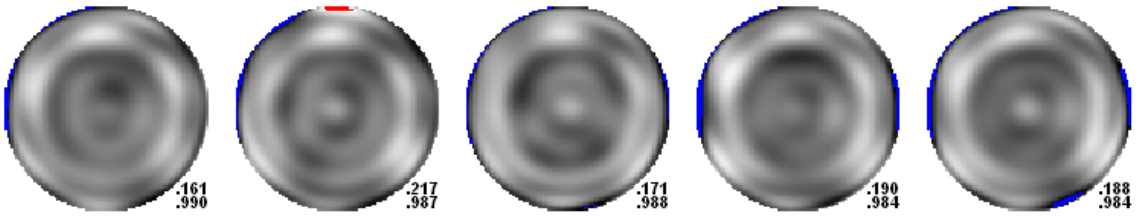

useful to understand the main contributors to P-V error (i.e., the highest peak and the lowest valley on the wavefront) for individual interferograms.

Note the following:

- There is a lot of variation from one interferogram to another with regard to the high and low spots.
- Most images have a high spot toward the top of the figure, which corresponds to the support sling. This is a left-over effect of test stand astigmatism that is not removed by cancelling the primary astigmatism term.
- The left over astigmatism often shows a 3-way symmetry (i.e., three hills), which is called "trefoil" aberration.
- Most low (blue) lips are at the top and bottom of the figures. This is either left over astigmatism, or could be due to the fact that the interferogram fringes in A1-5 were mostly horizontal, so the very top and bottom edges of the interferograms lacked data points to tie down the surface fit.

Clearly, the bulk of the P-V error in individual interferograms is not due to actual surface features, but due to test stand astigmatism or to random effects of turbulence or vibration. Further analysis requires averaging the Zernike terms from multiple interferograms to reduce these random effects. Figure 9 shows the effect of averaging random subsets of 6 interferograms, three from series A and three from series B, rotated to match the mirror orientation of series A. Averaging the same number from the A and B series has the benefit of minimizing the secondary effects of test stand astigmatism.

Note the following:

- The aberrations are much more under control—very little of the surface is now outside the 1/10 wave envelope.
- Some prominent surface features are starting to become apparent, in particular the wide raised ring and the narrow trough just to the outside of the raised ring.
- The Strehl ratio has crept up into the .986 range, probably due to cancelling some of the remaining effects of test stand astigmatism.

Figure 10: Interferograms Averages by Series

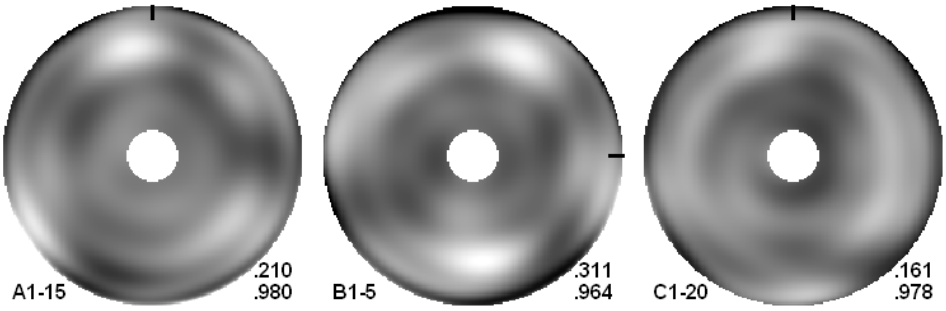

In the analysis that follows, I average much larger sets than the six shown in Figure 9, further reducing noise. Nevertheless, it is useful to see that an average of just six interferograms does so much to reduce the noise of individual interferograms. (OMI typically averages 5 interferograms, which is more than enough.)

4.3 Identification of Asymmetric Features

The next question to answer is whether any significant asymmetric features are apparent in the interferograms. These could show up as a different profile across perpendicular diameters of the mirror, or as localized raised or lowered areas.

Figure 10 shows averages of the series A, B, and C interferograms. The index mark shows how the B series is rotated 90 degrees clockwise.

Note the following:

- The lump at the top of A and C (and the three-way deformation) is probably residual test stand astigmatism (actually, trefoil). It does not show up the same in B.
- There are no clear features of A and C that are rotated 90 degrees clockwise in B.
- The lower Strehl ratio in the B series is probably due to the smaller sample size (only 5, compared to 15 and 20) or due to a "bad hang" on the test stand.

It does not appear that there is any benefit to treating different diameters of the mirror differently in the comparison with one-dimensional tests, such as the Foucault test. Also, it appears that the asymmetric hills and valleys are transient aberrations, not actual surface features. Therefore, it makes sense to ignore the asymmetric terms when performing comparisons with other tests that cannot easily report asymmetric mirror features. As in introduction to this idea, Figure 11 compares an asymmetric average of all 40 interferograms (left) to an extraction of only the symmetric Zernike terms (right). Clearly, use of just the symmetric terms does not change the Strehl ratio by a large amount, while maintaining the primary symmetric surface features.

The middle pane (flip asymm) shows another sort of average where the asymmetric surface is combined with itself rotated by 180 degrees. This technique tends to cancel asymmetric Zernike terms that have an odd rotational order, such as trefoil. As you can see, this technique does such a good job of eliminating the remaining "noise" that its result is not significantly different from the symmetric plot.

Figure 11: Asymmetric versus Symmetric Zernike Fit

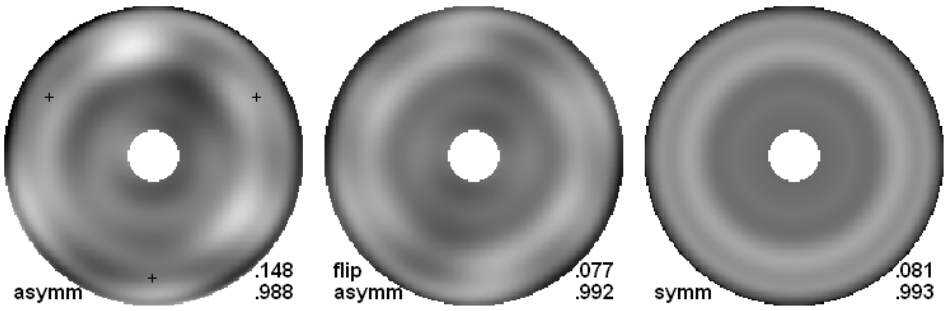

Figure 12: Order of Symmetric Zernike Fit

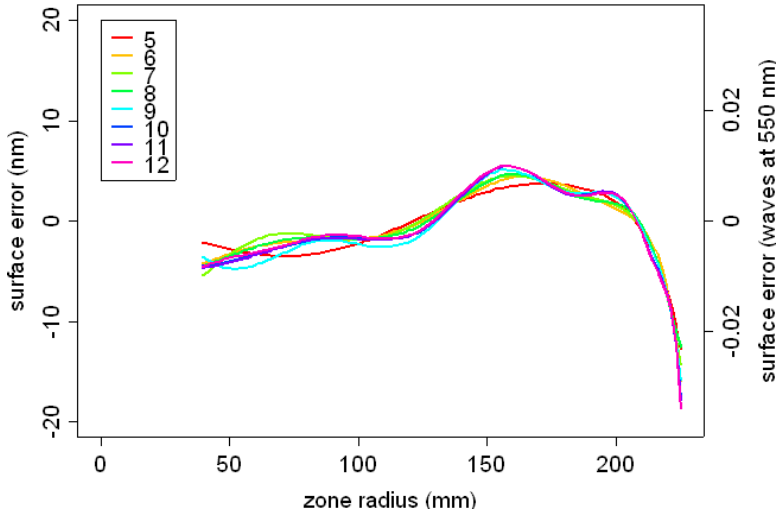

The crosses in the left pane (asymm) show the location of the back support pins on the test stand. They do not appear to line up with specific surface features, so I conclude that the surface figure is not affected by the location of the support pins.

It is interesting to note that the P-V wavefront error is reduced by almost a factor of 2 by using the symmetric terms. This indicates that the major contributor to the P-V error in the asymmetric case is the localized hills and valleys that are probably not actual surface features. This makes me believe that the P-V error derived from the symmetric fit is more accurate.

4.4 Derivation of Symmetric Model

From the above discussion, it looks like a fit of symmetric terms will make a better model of the mirror than the asymmetric terms. This required determining what order of terms to use. I found that order 12 terms are not too many for the symmetric fit, and show slightly more detail than the order 9 terms used in the asymmetric fit.

To demonstrate this, Figure 12 shows surface profiles derived from fits of order 5 through order 12 symmetric terms. By order 10 the curves are superimposed, with a high degree of precision. Clearly, order 12 is not too many terms for the symmetric fit.

A second demonstration will give you an idea of how much the noise is reduced by not using the asymmetric terms. Figure 13 shows individual symmetric profiles for all 40 interferograms. All but three fall within a fairly narrow halo. This gives me confidence that the symmetric model is capable of showing actual surface features, i.e., that the details apparent in the symmetric plot are not due to chance.

Note that the curves are least well behaved toward the edges. In particular, the outside edge sometimes curves up, rather than down. (This can also be seen in Figure 20. This makes me think that the very edge of a Zernike fit cannot be trusted.

Figure 13: Symmetric Interferogram Profiles

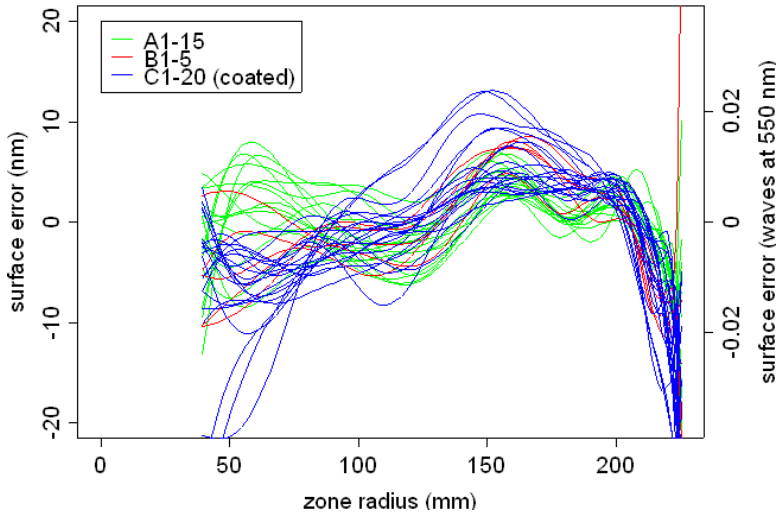

Figure 14: Comparison of Interferograms By Series

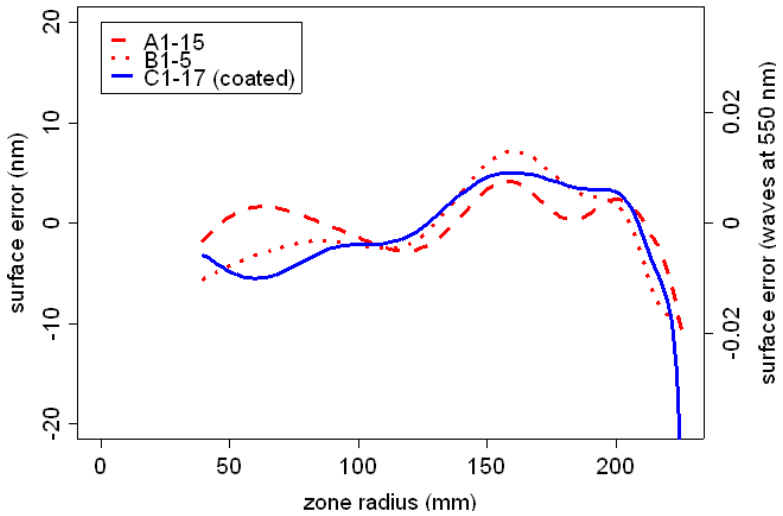

Figure 15: Average Interferogram Profile

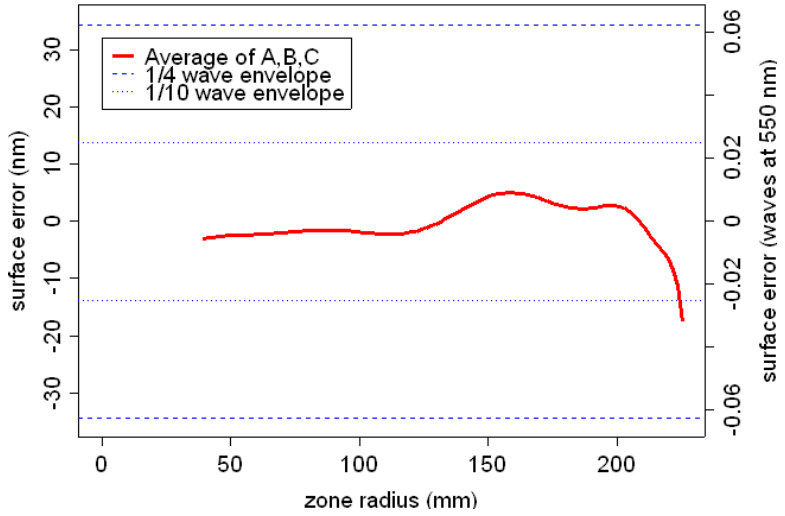

4.5 Effect of Coating

Having interferograms of the coated and uncoated mirror provide an unusual opportunity to assess the effect of coating on the figure of the final mirror. Mirror specifications from the manufacturer are usually on the uncoated mirror, so there is always a question of whether the coating has made the mirror worse. My mirror was coated with a standard aluminum coating with quartz overcoat by Spectrum Coatings. In a phone conversation with Paul Zacharias of Spectrum, he mentioned that a typical coating thickness is 1500Å to 2000Å of aluminum, and 2600Å of quartz overcoat. Thus, the total coating thickness is about $3/4\lambda$, with the reflective aluminum having a thickness of $1/4\lambda$ to $1/3\lambda$. A coating must be applied very uniformly for it not to affect the figure of the mirror! Zacharias also mentioned that any nonuniformity in a coating tends to show up as a thicker coating toward the center of the mirror.

What do the interferograms show? Figure 14 compares the surface profiles derived from averages of series A and B (uncoated), and an average of series C (coated). There are slight differences in the profiles, but not enough to conclude that the coating has changed the figure significantly. The maximum deviation of the curves is about $1/40\lambda$, with the majority within $1/100\lambda$ (on the wavefront).

4.6 Identification of Symmetric Features

Given that the coated and uncoated interferograms do not differ significantly, it is probably safe to use an average of all the interferograms to identify symmetric surface features and to derive final mirror specifications. Figure 15 shows the surface profile from this average. I have scaled the figure to include 1/4 wave (on the wavefront) and 1/10 wave envelopes for comparison.

Here are the primary features apparent in this profile:

- The mirror is relatively flat out 120mm.
- There is a broad hill from from 120mm to around 210mm with a slight depression in the

Figure 16: Foucault/Interferometry Comparison

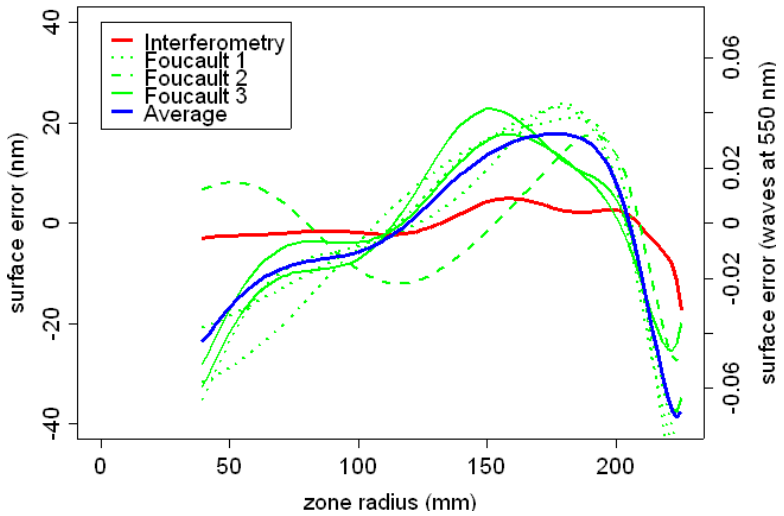

middle of it.

- The edge of the mirror falls off outside 200mm.
- Most of the mirror (all but the outer 1/4") is within a $1/20\lambda$ envelope.

As I will point out in the section on the null Ronchi test, I was able to see the broad hill and the slight depression visually with a null knife edge test. The steep fall-off at the edge was not as obvious as on this profile. That makes me think that the actual surface may not fall off as sharply as the model suggests. Assuming the basic profile is accurate, the extent of this drop is clearly the biggest problem with this mirror, even though it is contained easily within $1/10\lambda$ P-V of wavefront error.

The profile shown in Figure 15 will be used in comparisons with other quantitative tests, such as Foucault and Robo-Foucault. It will also be used, where possible, in comparison with qualitative tests that may show some of the surface details apparent in the profile.

4.7 Derivation of Specifications

Using the results in this section, Table 3 shows my best effort at deriving mirror specifications from the 40 interferograms. Compared to the OMI specs, the Strehl ratio is marginally higher (.993, versus .989), and the P-V wavefront error is significantly lower (.081, versus .122). The reduction in the P-V value is primarily due to eliminating the effect of noise in the interferograms by using only symmetric Zernike terms.

5 Foucault Test

In my first efforts to confirm the interferometry on my mirror, I performed a Foucault test with a 9-zone Couder mask. My first results were significantly different than the interferometry specs

(Strehl ratio of .80 to .84). I was concerned that my tester or test stand were affecting the results, so I replaced the micrometer, upgraded my light source from an incandescent bulb to a green LED, and tried testing with the mirror supported in its mirror cell. The results were all very similar.

I had a chance to redo my Foucault test inside a test tunnel at James Mulherin's shop. I did not have the time that day to perform the test as carefully as I would have liked, but I got different results than before (Strehl=.92). It was also much easier to read the outer three zones of the mirror than in my first tests. This makes me think that air currents were throwing off my previous tests, especially in the outer zones.

I tried a third time with a home-built 12' long test tunnel, but I still had difficulty with air currents, and not very good repeatability of measurements.

Figure 16 compares the surface plots derived from the Foucault test with interferometry. The results are labeled Foucault 1 through Foucault 3 corresponding to the three test sessions. I also show the average of all sets. For this plot I used order 7 Zernike polynomials to fit the 9 data points for each test run. The average profile follows a similar shape to the interferogram profile, but with a much higher peek and more fall off at the edge.

I was hoping to have a good set of Foucault results for this report, but it doesn't look like that is going to happen. I expect I will get a believable set of data some day, but I may have to wait for a time when the temperature in my house is more stable. At this point, I do not trust my Foucault results, because of poor repeatability. In the meantime, I should be able to confirm whether Foucault or interferometry is more accurate using the star test, or possibly the null Ronchi test with star source.

6 Dale Eason's Robo-Foucault

Dale Eason built an automated Foucault test rig and software to test an 8-inch mirror he refigured and a 16-inch mirror he is currently working on. The test platform is a traditional moving-source Foucault tester with the eye replaced by a video camera and longitudinal motion and depth of knife cut controlled by stepper motors. The platform uses high precision linear stages and a micrometer screw (in the longitudinal direction) to precisely control stage and knife motion. The light source is a bright blue LED stopped down to a slit. The knife, light source, and camera ride on the longitudinal stage. When cutting with the knife, only the knife moves, not the source or the camera.

The control software captures video input in real time, controls motion of the stage, and analyzes the gathered test data to produce a set of Foucault test readings suitable for input into analysis programs, such as Sixtests[1]. Robo-Foucault is capable of employing an unusually large number of zones (up to one zone per image pixel), which is the main feature which distinguishes it from a traditional manual Foucault test.

Dale's Robo-Foucault is fairly new, and has not been calibrated against a trusted reference optic. Nevertheless, Robo-Foucault results are repeatable with a very small variance in readings, and the results show remarkable similarity to the interferometric results for my mirror.

6.1 How Robo-Foucault Works

I assume the reader has basic knowledge of how the Foucault test works. In the usual Foucault test using a Couder mask, the tester attempts to find the position of the knife edge for which each zone "nulls", that is, both mask openings for a zone turn gray at the same time as the knife edge is

Figure 17: Foucault Image Intensity Profile

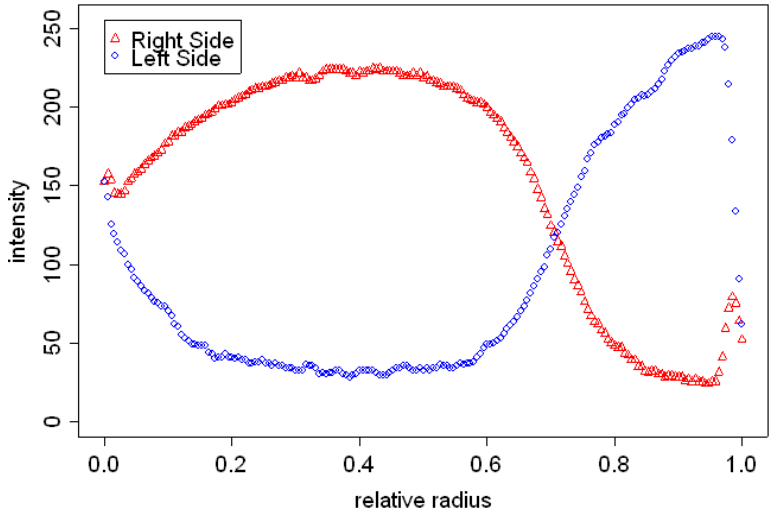

cut across the beam. This results in a table of knife edge positions for each zone radius. Normally, a number of test runs are made, and the results are averaged. A repeatability of readings in the range of .001 inch is considered very good for a moving-source tester.

An automated Foucault tester does away with the Couder mask, and reads the brightness of zones directly off the image of the full mirror. Dick Suiter wrote an article on automated Foucault analysis for the Amateur Telescope Making Journal[7] in which he describes how an automated Foucault test can reverse the dependent variable from a traditional Foucault test by setting knife position, then find the zone radius from the image of the mirror (at the radius where the shadow is equally gray on both sides of the mirror).

For example, look at the Foucault image for the 70% zone in Table 4. It shows the knife edge nulling the 70% zone on the mirror (the radius 7/10 of the way from the center to the edge). You can see a gray circle at this radius. If you look at the light intensity across a horizontal diameter of the mirror, you see that it goes very bright at the left edge, then fades to dark in the center of the mirror. As you continue to the right the image brightens again, then darkens to the edge. Finally, there is a bright diffraction ring at the very right edge of the mirror. If you overlay the brightness profiles of the left and right sides of the mirror, you get the graph in Figure 17. Where the two lines cross at a relative radius of .7 is the null for that knife position.

Dale uses a novel approach to turn Robo-Foucault data into traditional Foucault knife-edge readings. He fixes the zone radii, then interpolates to find the knife edge positions that null those zones. To accomplish this, he steps the knife position in small increments and captures the brightness difference between the right and left sides of the mirror for all zones. A brightness difference of zero for a zone would indicate that the knife is at the null for that zone. However, in a range of knife positions on either side of the true null position, there is a linear relationship between the knife position and the brightness difference. You can see this in Figure 17—in the vicinity where the curves cross, they are relatively straight. This relationship can be used in a linear fit algorithm to accurately determine the knife position to null a zone, even if the knife was never exactly at that position.

Figure 18 shows how this works. The X axis represents brightness difference between right and

Figure 18: Linear Fit of Robo-Foucault Readings

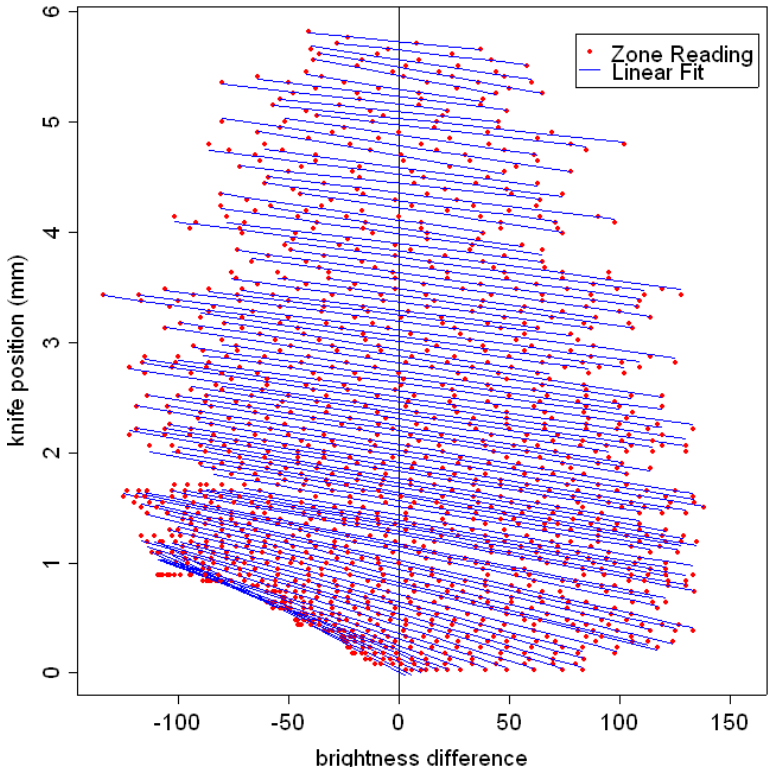

Figure 19: Order of Zernike Fit of Robo-Foucault Data

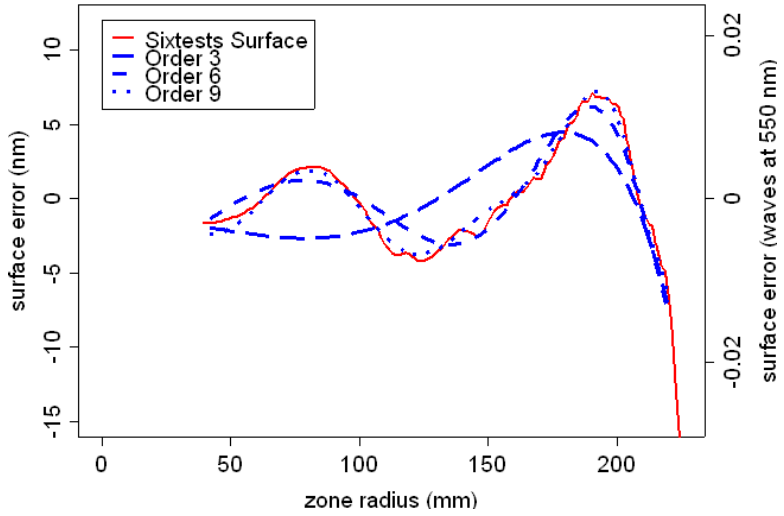

left sides for a zone. Brightness levels are discrete values between 0 and 255, so the brightness difference ranges from -255 to 255. The Y axis represents knife edge position in millimeters. The dots plot brightness difference against knife edge position. The blue lines are a best linear fit of the points for an individual mirror zone. The Y intercept of each of those lines is the knife edge position for that zone.

As you can see, a lot of averaging goes into the determination of each zone position. Even so, the zone position predicted by each data point is already very good. Standard deviation of zone position is often in the range of .0002 inches, which is very, very good by Foucault test standards.

6.2 Modal Analysis

There are some choices in how to process Foucault knife edge readings into wavefront or surface deviations for comparison to other methods, especially with the large number of zones in Robo-Foucault results. The knife position for each zone is related to the slope of the optical surface. Traditional Foucault analysis assumes all knife edge readings are exact and that the surface slope changes smoothly between the zones. The smoothed slope curve is integrated to predict the optical surface deviations [Burrows]. This presents a problem with the many-zone Robo-Foucault test. When using the traditional method, the resulting surface profile is not very smooth, which is unrealistic. It would be better to use a method that finds a smooth curve that best fits the slope data.

Mike Peck wrote a paper describing a modal (i.e., Zernike polynomial fit) method of Foucault test reduction [Peck]. This method produces almost identical results to traditional integration on a normal Foucault test. On a many-zone Foucault test, it allows you to select the amount of smoothing you want, based on the number of Zernike terms used. Also, this is exactly the format we need for comparing Robo-Foucault with interferometry.

Figure 19 compares a modal fit of Foucault data to traditional analysis. The red line was calculated by the program Sixtests from a sample Robo-Foucault run. This line has small-scale wiggles that change from run to run, and are clearly "noise". The three blue lines show the result

Figure 20: Robo-Foucault Surface Profiles

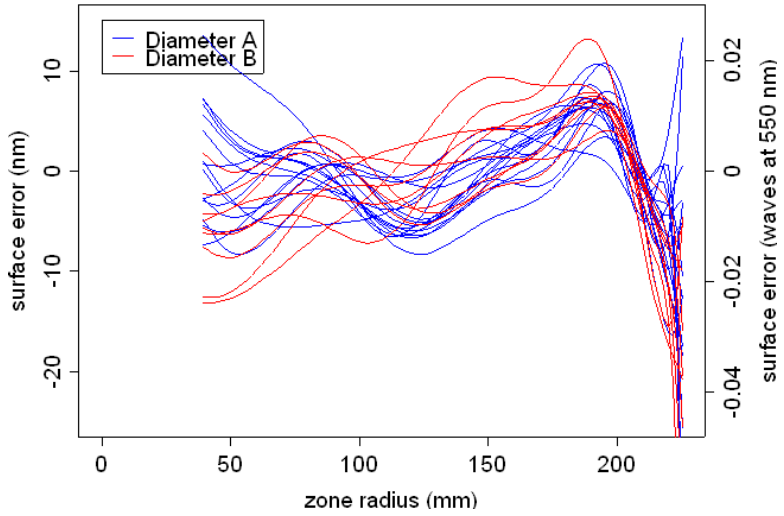

of a modal fit using order 3, order 6, and order 9 Zernike terms. The order 9 curve fits the data very well, but smoothes out all the small-scale undulations. Assuming other Robo-Foucault results have a generally similar appearance, it seems that order 9 Zernike terms are sufficient to model the surface.

6.3 Robo-Foucault Results

Dale analyzed two diameters of my mirror, corresponding to horizontal diameters across the orientation for series A and series B of the interferograms. Overall, 40 clean many-zone data sets were captured. I eliminated about half the results as "outliers" leaving 24 results. (The outliers were probably due to air currents at the time the test was performed.) The number of zones for the tests varied from 47 to 155, with an average of 115 zones.

Figure 20 shows the results. The spread of individual results is remarkably tight. There do not appear to be significant differences between the profiles on diameters A and B.

Figure 21 shows how an average of the Robo-Foucualt results compares to the surface profile from interferometry. Robo shows the main features of the interferometry profile, except that the top of the broad hill is inclined the other way. Within the error tolerances of the two tests, the Robo-Foucault results are indistinguishable from interferometry.

7 Null Ronchi and Null Knife-Edge Tests

During my visit to Mulherin's shop, I was able to look at my mirror under the null Ronchi test. The lines on the mirror with a 60 line grating were perfectly straight, even when only a few lines were showing on the mirror (i.e., the grating was close to the mirror's center of curvature).

We also used the Ronchi grating to perform a null knife-edge test by putting the grating very near the center of curvature so that one Ronchi line was filling the mirror. When cutting across the returning beam with this "knife edge", we were able to see three subtle shadows concentric with

Figure 21: Robo-Foucault/Interferometry Comparison

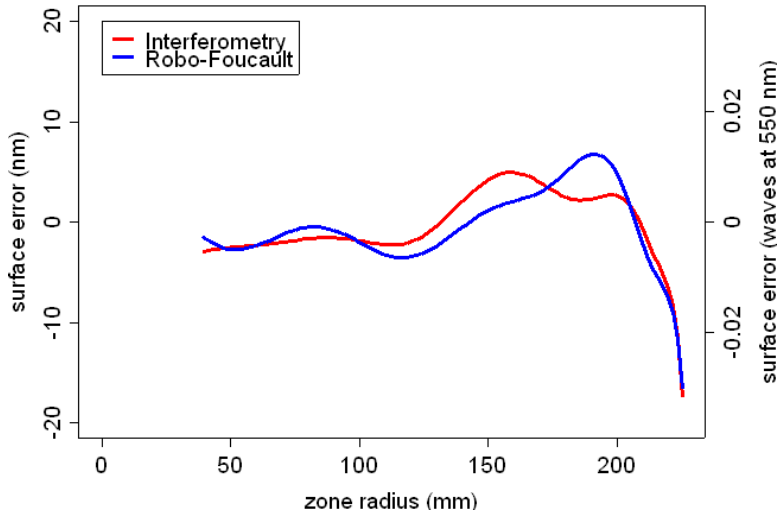

Figure 22: Simulated Null Knife-Edge Test from Interferometry

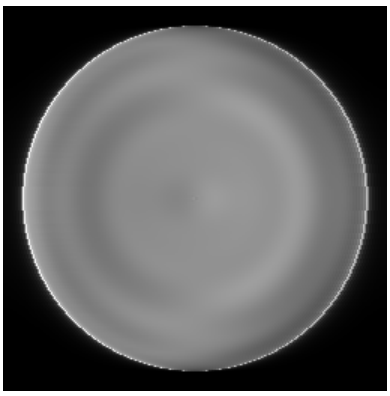

Figure 23: Simulated Null Knife-Edge Test from Foucault Test

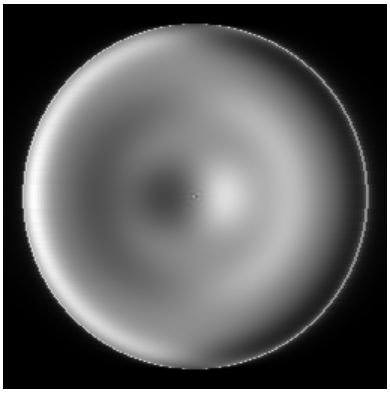

the mirror. These were the falloff at the edge, the faint ring inside the edge, and the wider ring closer to the center of the mirror.

Figure 22 and Figure 23 show simulated null knife edge tests using the surface data derived from interferometry and the Foucault test. These images can help to check whether interferometry or my Foucault results are closer to reality. Visually, I saw very faint annular zones, nearly identical to Figure 22. The null knife-edge test looked nothing like Figure 23. The only way this test could produce erroneous results (i.e., making the Foucault surface optic look like the simulated null knife-edge test from the interferogram surface) is if the optical flat used in the null knife-edge test has a lot of spherical aberration on it, which is highly unlikely.

A similar null knife-edge test can be performed in the field with a star source. This eliminates the optical flat, but is more difficult to perform because it requires steady seeing. I have not performed this test to my satisfaction, yet, but I will use it to confirm the null knife edge test performed on the bench.

8 Star Test

The star test (see Suiter[8]) is very useful, because it requires no additional test equipment, and it is extremely sensitive. However, the star test is not usually quantitative, so its results cannot easily be compared to the other tests described in this paper. In this section, I will describe the examination of the in focus and out-of-focus image of a bright star. A related test, the null Ronchi test with a star source is described in a separate section.

The star test can be used to answer the following questions:

1. Does the objective show any significant aberrations that will affect its usefulness?
2. Assuming other quantitative tests disagree, which quantitative test is more accurate.

8.1 Simulated Star Test

To help answer question 2 (which quantitative test is the most accurate?) I have simulated a star test for each quantitative test in this report that can be compared to the star test in the field. Figure 24 shows star test simulations for ideal objective, interferogram results, robo-foucault results, and foucault results. These simulations were done directly from the wavefront models derived from the test data. They show a somewhat ideal star test with no turbulence and no diffraction due to the spider. The 17% obstruction of my telescope has been included.

Please ignore the overall brightness of each image. There is a problem in the software that generated these images which prevents me from properly coordinating the brightness scale. For example, the 10 waves defocused images of Interferometry and Robo-Foucault should be the same overall brightness as the corresponding inside focus images, but they are dimmer.

I notice the following:

- There is very little to distinguish between the Interferometry and Robo-Foucault patterns.
- The differences between the ideal aperture and Interferometry/Robo-Foucault patterns are subtle. The most obvious difference the relative softness of the outer ring inside focus.
- In the Foucault pattern the soft outer ring is more pronounced, and the entire diffraction pattern inside focus lacks clarity.

Figure 24: Simulated Star Test

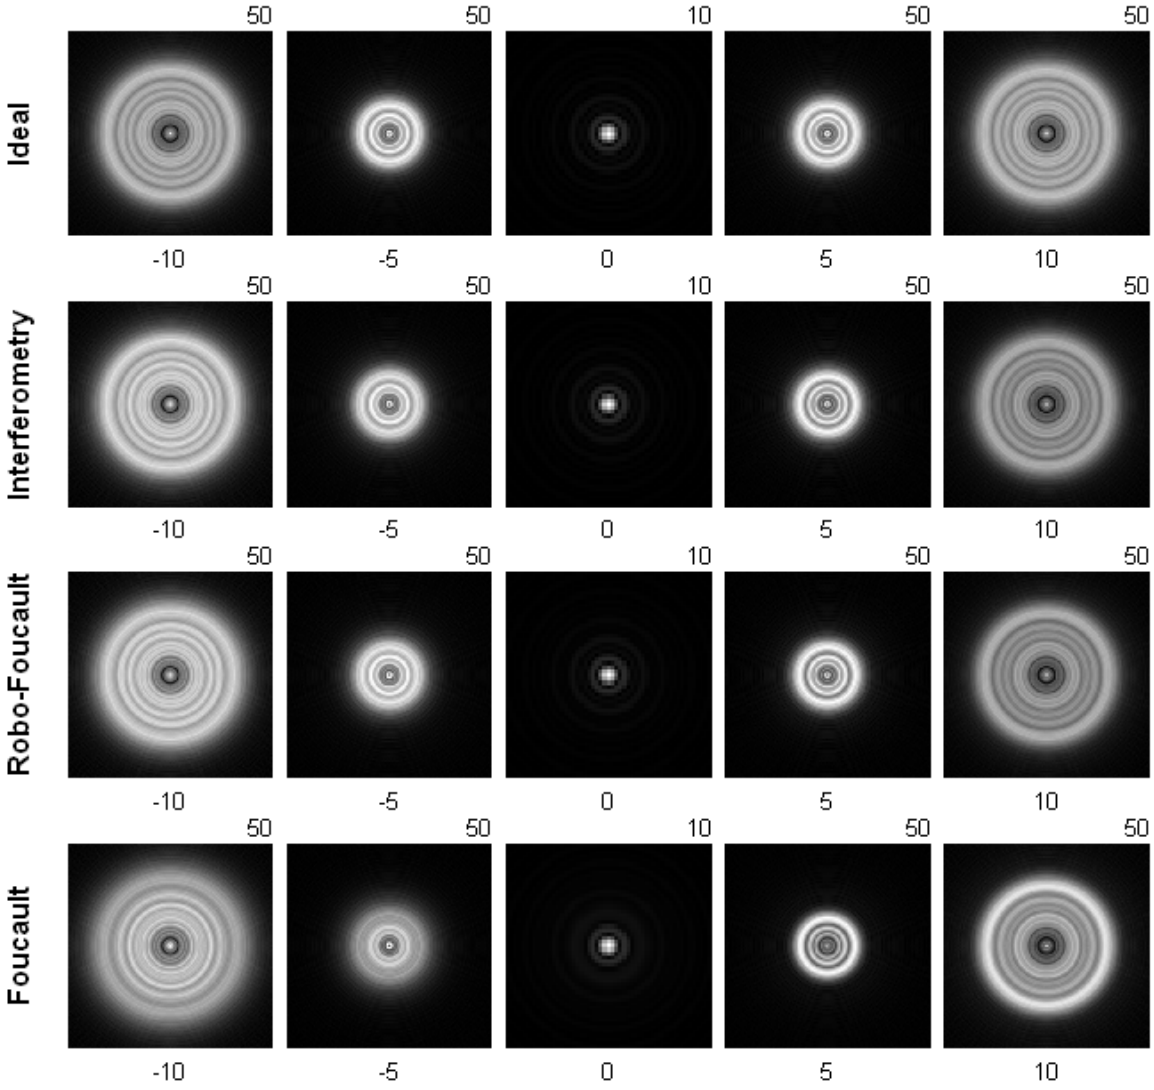

8.2 Star Test Results

I have had only a few opportunities to star test this mirror, so far, but the results are encouraging. One problem I have with the star test is that my location (Minnesota) often has very bad seeing conditions (i.e., turbulence) which makes it difficult to compare the diffraction images.

One feature of the star test I have seen in testing other mirrors is an uneven distribution of turbulence spikes inside and outside of focus. This asymmetry can indicate TDE or the presence of steep zones on the mirror which throw light outside the Airy disk. This mirror shows minimal spikes, and they are evenly distributed on both sides of focus. The spikes grow less as the mirror cools and as the seeing gets better. In focus, there are very few spikes. One time with very steady seeing and a cool mirror I saw the remaining spikes disintegrate into lines of fine points. I believe I was on the verge of seeing the first diffraction ring.

In relatively good seeing I have seen a very condensed Airy disk, but not any evidence of the first diffraction ring.

I have not yet had the opportunity to match the simulated star test in Figure 24 with what I see at the eyepiece.

8.3 Conclusion

Where does all this testing leave me? I am very encouraged that Robo-Foucault and interferometry agree, because they are completely independent tests—no overlapping equipment and different operator. The details I saw in the null Ronchi test are also encouraging, although that shares a flat with interferometry, so it is not a completely independent test. The fact that the coated and uncoated interferograms agree is shows that the two interferometer reference elements are similar. However, I am concerned that I have been unable to get a manual Foucault test to agree with interferometry. For now, I am assuming that it is some remaining problem with my test environment or technique. It's also possible that I am running into the inherent difficulty of performing the Foucault test on a fast mirror. It will take further testing to determine the source of the error. I will continue to use the star test to cross check the other tests. If I solve the Foucault test mystery, acquire additional test results, or find mistakes in my analysis I will publish updates to this paper.

References

[1] Burrows, James. *SIXTESTS: Foucault/Caustic/Poorman's Data Reduction (Windows version)*, http://home.earthlink.net/~burrjaw/atm/odyframe.htm.

[2] Burrows, James. *ATM Mirror Math.* http://home.earthlink.net/~burrjaw/atm/atm_math.lwp/atm_math.htm.

[3] Ceravolo, Peter. *Interferometry and Telescopes: A practical guide to building and using your own interferometer.* http://www.ceravolo.com/testing/interferometry.pdf.

[4] Mulherin, James C. *Averaged Results for Twenty Interferograms of a 15" F/4.5 Newtonian Primary Mirror.* http://www.opticalmechanics.com/interferometric_test_results.htm.

[5] Peck, Michael. *Modal Analysis of Foucault Test Data*, 2002. http://home.netcom.com/~mpeck1/astro/modal/modal.pdf.

[6] Rowe, David, *FringeXP* (fringe analysis software). http://www.ceravolo.com/fringe/FringeXP/FringeXP.htm.

[7] Suiter, Harold R. *Digital Knife-Edge Test Reduction*. ATM Journal, 1998.

[8] Suiter, Harold R. *Star Testing Astronomical Telescopes*. Willman-Bell, 1994.

[9] Wyant, James C. and Katherine Creath. *Basic Wavefront Aberration Theory for Optical Metrology*. Applied Optics and Optical Engineering, Vol XI. Academic Press, 1992. http://www.optics.arizona.edu/jcwyant/Zernikes/Zernikes.pdf.

[10] Weisstein, Eric W. *Zernike Polynomial*. http://mathworld.wolfram.com/ZernikePolynomial.html.

[11] Wyant, James C. *Zernike Polynomials*. http://www.optics.arizona.edu/jcwyant/Zernikes/ZernikePolynomialsForTheWeb.pdf.