UCRL-MA-137950-REV-1

NIF Large Optics Metrology Software: Description and Algorithms

Wade H. Williams

15 October 2002

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

Approved for public release; further dissemination unlimited

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Work performed under the auspices of the U. S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

NIF Large Optics Metrology Software: Description and Algorithms

3/16/00	**Original distribution**
10/15/02	**Rev 1. Change to reflect PSD1 range change; add influence function subtractions on LM1 finished and coated; add CPHASE description**

Summary

Several software packages have been developed for use by NIF large optics vendors during production of NIF optics. These packages allow specific comparison of the interferometer measurements done on optics against the wavefront requirements for those optics, as given on relevant drawings

This document outlines the various packages, and their specific applications, and describes in some detail the calculational algorithms used. It is intended as the primary reference document for the codes (aside from the source codes themselves).

Introduction

In order to ensure good laser beam quality, NIF requires that all large optics be measured with an interferometer to monitor how that optic will affect beam wavefront quality and focusability. Various specifications for transmitted wavefront (TWF) and reflected wavefront (RWF) for the full-aperture parts, and for various sub-apertures, are given on the large optics drawings. For reference, a summary of the various specifications for the NIF large optics is given in Figure 1. Each large optic in production will be measured against some of these specifications. Other specifications will be monitored in a "process validation" fashion by measuring a representative sample of parts to assure the process is yielding parts which pass specification. This document will focus on the specifications requiring measurements on every part. This document will not go into detail concerning the procedures and limitations of the measurements themselves.

The interferometers currently planned for use in NIF optics production are manufactured by Wyko Corp., Zygo Corp., and Phase Shift Technologies, Inc. All manufacturers provide software to analyze the interferograms obtained with their instrument. However, because NIF requires specifications which the interferometer vendors have not had to normally support, the software is not capable of analyzing the results against all the NIF specifications. The Zygo MetroPro® software is able to evaluate fused silica optics blanks against their specifications, so is being used for that purpose at the fused silica manufacturer. For all other optics, rather than split the analysis between different pieces of software, it was decided to build complete packages for these optics which would contain all relevant specifications. Similarly, it was decided to build custom software for each optic type (or process step), rather than combine all optics under one package. This was done to minimize confusion for the vendors: they have only the analysis capability they need for the optics they are fabricating. A summary of the names and applications of these custom packages is given in Table 1. It is the purpose of this document to describe the operation of these codes, including inputs, analysis algorithms, and outputs.

Table 1. Titles and summary purpose of the various large-optic metrology analysis packages.

Software Title	Purpose
PHOMGLASS	evaluate homogeneity of laser glass blanks
FSFIN	evaluate finished fused silica parts
GMPFIN	evaluate finished laser glass, mirrors, and polarizers
MPCOAT	evaluate coated mirrors and polarizers
CPHASE	evaluate crystals, blanks and finished

Specifications covered by these packages are those dealing with spatial scale lengths >= 2 mm. Shorter-scale specifications will be monitored as process control parameters, as discussed above, and will require measurements on small patches of the optics with a different interferometer than used for the other specifications. They will only be done an occasional basis. It is currently anticipated that all such measurements will be done at LLNL. A code to analyze for these short-wavelength specifications (WAV2) has been prepared by Janice Lawson for use at LLNL.

All codes are written in the commercial analysis language, IDL®, sold by Research Systems, Inc. Using a commercial package minimized programming time by allowing use of pre-programmed elements (such as graphics and numerical routines). This has required, however, that each vendor using this software also purchase an IDL® license.

In the next sections we describe the code inputs, calculational algorithms used in the codes, and code outputs.

Code Inputs

As inputs, the codes use a binary file stored from the interferometer, which contains the phase-unwrapped data of the part measured. In most case this will be an image of the entire part with these large-aperture interferometers. It most cases the codes will also request a background (i.e. empty cavity) file for use in background subtraction. The part measurement file should have been stored without any background subtracted.

The code is compiled with the specification values for each optic type as taken from drawings. It also carries information about the required effective ("clear") aperture for each part from the drawings. This allows the analysis to provide "pass / fail" output, rather than just numerical results. It will require, however, that the software be updated if specifications change. The codes use "point and click" interfaces for user input.

Overview of Algorithms

A variety of analysis steps must occur for each optic under evaluation. While there is much overlap in the analysis required for the various optic types, there are some steps unique to certain optics. Some examples include: laser glass homogeneity specifically

allows for certain levels of power and astigmatism; the French-CEA specifications on laser glass homogeneity require somewhat different analyses than LLNL-NIF specifications; the KDP crystals must be evaluated for phase jumps and index homogeneity; background subtraction is handled differently for TWF and normal-incidence and non-normal incidence RWF measurements.

The various analysis steps used in the codes are summarized below, with the more generally-used steps described first, and the special-case steps at the end. Fuller descriptions of many of the steps are given in the comments following. Table 2 lists which steps are used for which optics.

Table 2. Analysis steps used in various optic types.

Part type	Analysis code	Analysis steps
laser glass blanks, PHOM, LLNL	PHOMGLASS	1, 2, 3, 12, 8, 6, 7, 13
laser glass blanks, PHOM, CEA (French)	PHOMGLASS	1, 2, 3, 14
laser glass, windows, polarizer-trans. finished	GMPFIN, FSFIN	1, 2, 3, 4, 6, 7, 8, 9, 10, 11
lenses finished	FSFIN	1, 2, 3, 6, 7, 8, 9, 10, 11
mirrors, polarizer-refl. finished	GMPFIN	1, 2, 3, 4, 5, 6, 7, 8, 9, 10,15
mirrors, polarizer-refl. coated	MPCOAT	1, 2, 3, 4, 5, 6, 7, 8, 9, 10,15
crystals, blank	CPHASE	1, 2, 3, 4, 6, 7, 16, 17
crystals, finished	CPHASE	1, 2, 3, 4, 6, 7, 9, 10, 17, 18

1. The user identifies files and analysis type / mode. This includes measurement and background files, optic type being analyzed (e.g. finished laser glass), and analysis mode (process control or final pass / fail). (See comments A and B below.)

2. Remove piston and tilt in measurement and background (if present) files. (See comment C below.)

3. Determine "effective" aperture. This is the portion of the optic over which the specifications apply, as given on the drawing, and over which the analysis is performed. It includes rounded corners for most optics (30 mm radius). This aperture mask is applied to background and measurement files. (See comment D below.)

4. Subtract background file from measurement file in TWF measurements, pixel-by-pixel. This is the analysis phase array. No background file is used in normal-incidence RWF, lens, and PHOM measurements. For non-normal incidence RWF measurements the background file is subtracted statistically – see steps 7 and 9. Also for non-normal

incidence RWF measurements determine the power in the background file, and subtract this from both background and measurement files. (See comments E and F below.)

5. Correct horizontal spatial dimension and phase amplitude for any difference between the measurement and use angles for the part. In both cases, this is a simple cosine-angle correction. (See comment G below.)

6. Perform a 33 mm low-pass filter of the phase. Fourier transforms are done with a quad-flip technique and an error function edge filter. They also require patching invalid pixels. (See comments H-J below.)

7. For TWF and normal-incidence RWF measurements, calculate the rms gradient of the filtered analysis phase array. For non-normal incidence RWF measurements, calculate the filtered rms gradient of the measurement and background files separately, and then statistically subtract the background rms gradient from the measurement rms gradient. (See comments E and P)

8. Calculate the peak-to-valley phase. For TWF measurements, we use the filtered analysis phase array. For RWF measurements, since the background cannot be directly subtracted, we use the filtered, measured phase array, and take no credit for background subtraction. (See comment N.)

9. For TWF and normal-incidence RWF measurements, perform a band-pass filter (2-33 mm for most optics, 6-33 mm for crystals) on the analysis phase array, and calculate the rms. This value is labeled "PSD1". For non-normal incidence RWF measurements, calculate PSD1 of the background and measurement files separately, and then subtract the background PSD1 from the measurement PSD1 statistically.

These calculations require patching invalid pixels. Fourier transforms are done with a quad-flip technique and an error-function edge filter. (See comments E and H-K below.)

10. Calculate the 1D collapse lines of the analysis phase array (TWF or normal incidence measurements) or the measurement phase array (non-normal incidence RWF measurements), and compare against the Not-to-Exceed line given on drawings. This requires patching invalid pixels. Fourier transforms are done with a Hanning window. (See comments E, H, L and M below.)

11. For surface measurements of the final focus lens and the diagnostic beamsplitter determine the best-fit power of the RWF and the residual peak-to-valley after subtracting this power. (See comments C, F and N below.)

12. For LLNL PHOM measurements, subtract power and astigmatism. (See comments F and O below.)

13. For LLNL PHOM measurements, determine p_v (99%) of surface contour. (See note N.)

14. For CEA (French) PHOM measurements the processing steps are summarized here: subtract power, perform 10 mm low-pass filter, calculate rms and peak (99%) gradients. Determine homogeneity by performing 10 mm high-pass filter on phase with tilt, piston, and power removed. Then look for highest (99%) phase p_v over any 1 x 1 cm patch and divide by part thickness to convert to Δn. Remove piston and tilt in surface measurement, and determine p_v (99%). Additionally, remove power from the surface measurement, and determine the worst p_v (99%) in the ~ sixteen 100 diameter subaperture patches. Each of these patches have tilt removed. (See comments F, J, N, O, P.)

15. For LM1 only, subtract a best fit of measured adaptive optic influence functions from the analysis phase array. The filtered rms gradient calculations are done on this adjusted array. This subtraction is not done in calculations of the filtered peak-to-valley specification or PSD1 specification. (See comment Q.)

16. For crystals (blank and finished) perform the "ramped phase jump criterion" calculation, wherein the phase difference across small distances everywhere on the part is compared against drawing values. This specification guards against large index jumps across crystal phase boundaries which could cause excessive beam modulation in the laser downstream of the crystal. (See comment R.)

17. Perform Orthogonal Polarization Interferometry (OPI) calculations. This requires h and v-polarized measurements on the part, as well as an empty cavity file, and allows determination of local crystal axis wander (away from nominal) in the part, which is converted into a local crystal detuning angle wander. Lineouts through the resulting detuning map are compared against absolute detuning angle scans made in the Crystal Alignment Test System (CATS). OPI is not done on the switch crystal. (See comment S.)

18. Determine the unfiltered P-V surface contour on one face of the crystal. (See comments C and N.)

Comments:

A) Initialize measured phase maps. The phase maps are stored in either Zygo (.dat), Wyko (.opd), or PhaseShift (.MAP) binary file formats. These are read in as a floating point array of the phase values at each camera pixel, along with x and y vectors. Units of waves (633 or 1053 nm) and cm are used. Other information, such as measurement date, part number, part serial number, part type, operator comments, and part center fiducial location are also extracted from the file. Zygo files store only a rectangular phase array, which may be smaller than the full camera array if the part underfilled the camera, so the reader also backfills the camera-size array as invalid data points. This allows subtraction of one file (e.g. background) from another.

Files of size > 1024 x 1024 are interpolated onto a grid of this size. The quad-flip Fourier transform technique uses an array of twice larger dimensions, and is excessively slow if the arrays are larger than this. This is only an issue with the 24" Zygo, which can store data at up to ~2048 x 2048 resolution. Analysis of sample data files shows that the resolution lost from 2048 x 2048 to 1024 x 1024 is not important to the specifications under consideration.

The integrity of the files are checked by assuring the background (if present) and measurement files have the same dimensions (to allow background subtraction). The "wedge" factor in the files is also compared against that expected for the optic type specified. For example, for a finished piece of laser glass the specifications are written as transmitted wavefront. Since the interferometer laser double-passes the optic in measurement, a wedge factor of 0.5 should be used. The code, however, does not attempt to rescale the data using a correct factor. The mismatch is simply reported in the output.

B) Determine optic type and analysis conditions. The user selects menu items to specify which part is being measured (e.g. finished laser glass), allowing the code to compare the measurement against the proper part specifications from the drawings, which are contained in the code. The user also selects between "process control" mode or "final pass/fail" mode, the difference being that in the "process control" mode the user has the opportunity to print out the various graphics as they are calculated (including changing the graph scales to allow more visual part-to-part comparisons), and to store the 1D PSD collapse lines in an Exce®l-compatible (.csv) file.

C) Remove piston and tilt. "Piston" is any average offset from zero phase in the interferogram. Piston is added to the interferogram as an arbitrary factor during the phase unwrapping. As it does not affect beam propagation, but will affect specifications such as the rms phase, it is removed by subtracting the average from the phase array.

Tilt (in the x and / or y directions) has the effect of a wedge on the beam. Because tilt can be introduced in the measurement by simple tilts in the interferometer optics, the amount of wedge in the part cannot be deduced from the interferogram. As with piston, tilt will, however, effect such parameters as the peak-to-valley, so must be removed. This is performing a least-squares fit to the data to two orthogonal planes (x and y), using only valid data points, and removing the resulting fit. An example of the effect of tilt removal on the rms wavefront is shown in Figure 2. Because of the invalid data points in the array (shown as zero's in the figure), the least-squares fit gives a better result than a simple overlap integral.

D) Determine working analysis aperture. For most measurements in the large (24") Wyko and Zygo interferometers the entire part will be viewable in the camera array, leaving a strip of invalid pixels around all or some sides of the part. The code scans the camera array in from the left, right, top and bottom sides to find the first valid data point in each row or column, and calls these the edges of the part.

From these edges, it then attempts to find the "effective" (or "clear") aperture over which the specifications apply, as given on the drawings. This is done by either coming in a fixed distance from each part edge, depending on the part type; or by using fiducial information in the measurement file header which indicates the center of the part, and then tracing the effective aperture from this. In either case, the resulting calculated effective aperture is displayed with the interferogram, and the code user has the opportunity to manually change (drag and click the mouse) this aperture before doing calculations. All subsequent calculations are done only on this specified aperture. The ratio of this aperture area to the effective aperture area on the drawing is also reported by the code as a double-check for coordinators.

E) *Background Subtraction.* Background subtraction is handled differently, depending on whether the part is measured in transmission, normal-incidence reflection, or non-normal incidence reflection. In all cases, the background file is simply an empty cavity (TF and RF) which was measured with a "wedge" factor of 0.5, representing on pass through the TF plus the surface contour of the RF.

For transmission measurements of finished laser glass, windows, and polarizers-trans, the background file is subtracted pixel-by-pixel, and then all calculations are done on the resulting array. Because of this, care must be taken that the part measurement aperture is "seeing" the same portion of the RF/TF as in the background measurement. To this end, for Brewster's angle parts (laser glass and polarizer), where there is beam walk-off, the RF must be shifted by the amount of the walk-off during the part measurement relative to where it was during background measurement. This shift must be accurate to within about +-0.5 mm (as shown in LLNL tests).

For transmission measurements on blank laser glass (PHOM) there is no separate background file (this is removed as part of the PHOM algorithm).

For transmission measurements on lenses it is not possible to do an "empty cavity" measurement (just TF and cal sphere). Hence, to be conservative, we do not attempt to subtract out any background in lens measurements.

For RWF measurements on non-normal incidence optics (the polarizer and all mirrors except M1and M2) we attempt to take as much credit for the background measurement as possible. Though the RF and TF are both present in part measurement and in background measurement, the RF image is flipped around the vertical axis in the part measurement (because of the beam reflection off the part), so the background cannot be directly subtracted. Instead, we attempt to subtract the background in other ways. First, in order to directly subtract out as much background as possible, and since power is symmetric around a vertical axis flip, we fit a sphere to the background measurement, and then subtract this fit from both the background and measurement files. (We could consider higher-order symmetric terms.) Second, we calculate the filtered rms gradient and PSD1 of the power-subtracted background and part measurement files. We assume that the background in both cases adds statistically to the part phase noise, in a root-sum-square

fashion. So we deduce the part rms gradient as the square root of the difference of the squares of the part measurement rms gradient and the background file rms gradient.

$$rms_p = \sqrt{\left(rms_m^2 - rms_b^2\right)},$$

where rms_p, rms_m, and rms_b are the filtered rms gradients of the part alone, part in measurement, and the background, respectively. PSD1 is calculated the similarly. Theoretically, the basis for this statistical differencing is stronger for PSD1, since it is for short-wavelength aberrations. The rms gradient covers > 33 mm aberrations, which will be less random over the 40 cm parts.

For RWF measurements on normal incidence mirrors (M1 and M2), we cannot do background subtraction. The surface finishes on the TF and on the RF have equal weightings in the empty cavity measurement. Until a calibration is done of the TF, we cannot determine what fraction of the rms gradient or PSD1 in the empty cavity come from the TF (which will be present in the part measurement). Hence, to be conservative, we do not attempt to subtract out any background in normal-incidence measurements.

F) Determine Zernike best fit to remove power and/or astigmatism. Zenike polynomials form an orthogonal basis set over a circular aperture, and can be determined in such a situation using the overlap integrals of the functions on the array of interest. They are not orthogonal on a rectangular aperture, as seen in these measurements, however, so are not well suited for this purpose. Nonetheless, because of their familiarity, the low-order Zernike terms of astigmatism and power are used on drawings, so are calculated in this code. Because of the rectangular aperture, the fitting coefficients are determined using a least-squares fit of the data to the polynomials, excluding invalid data points, using Singular Value Decomposition to solve the matrices. An example of a 15-term Zernike fit is shown in Figure 6. To increase code speed, the data array is interpolated onto a 256 x 256 grid when doing these fits. Also, because the Zernikes are not orthogonal on the rectangular grid there is cross-talk between the terms (i.e. the magnitude of any given fit coefficient depends on the number of terms calculated in the fit). Consequently, when removing power we only calculate the first four terms (piston, tilt-x, tilt-y, and power). For astigmatism calculations we add the next two terms, astigmatism (0, 90 degrees), and astigmatism (+-45 degrees).

G) Note on cosine factor correction. The correction factor for adjusting between different measurement and specification angles is done with a cosine correction factor:

$$\phi_{spec} = \phi_{meas} \frac{\cos\theta_{spec}}{\cos\theta_{meas}}$$

and

$$x_{spec} = x_{meas} \frac{\cos\theta_{spec}}{\cos\theta_{meas}},$$

where ϕ_{meas} is the phase as measured, and x_{meas} is the horizontal axis as measured. (All angle mismatches under the current measurement plans are in the horizontal axis.) The rationale for this correction is shown in Figure 13.

Ray trace calculations indicate this correction should be correct unless large aberrations have spatial wavelengths on the order of the camera pixel spacing. The obliqueness of the measurement dictates how much surface area each pixel sees, which becomes significant for short-wavelength aberrations. As the short end of the PSD1 band-pass is 2.5 mm, and the camera pixel spacing is ~0.5 mm, this should not be an issue in these measurements.

H) Patch non-valid points. Invalid points in the binary file are carried as NaN's in the IDL routine. For many calculations they can be ignored. For spectral calculations, however, the Fourier transform requires real numbers in the entire array, so these points must be "patched". Simply replacing them with zero's is not sufficient, as these areas in the phase map will appear to have discrete discontinuities in phase, which adds high-spatial frequency power to the spectrum. Even a few such points can significantly change the shape of the spectrum, and the value of the PSD1 specification (rms phase in the 2.5-33 mm band), as shown in Figure 3.

Interpolation is performed as shown in Figure 4, using the four valid pixels around the invalid point (in the x and y directions). If there are only three valid pixels around the point the calculation is similar, but only the three points are used.

If there are only two valid points then the code attempts to extrapolate the data, as shown in Figure 5. This is difficult to do in a general case. There are a number of different approaches possible, none of them without problems. The approach implemented here adds data to the invalid sections which mostly retains the spectral structure of the valid areas without adding edge discontinuities at the boundary between valid and invalid data. There is sufficient concern about this approach, however, that the code warns the user if more than 5% of the data in the effective aperture is invalid, indicating the PSD calculations may be suspect. Only careful analysis on a case-by-case basis can quantify such uncertainties. Fortunately, it is anticipated that this will rarely be an issue. All measurements where PSD calculations are done are of either rectangular parts or rectangular subsections of parts, which should normally be placed squarely (i.e. with little rotation) in the aperture, minimizing the area needing extrapolation. More importantly, the use of an effective aperture smaller than the part being measured means no extrapolation will be needed for most analyses.

I) Note on quad-flip technique. If the left and right edges of an array (or the top and bottom) are discontinuous one with another, as is usually the case with real data, then these edges will introduce spurious noise in the spectrum when performing a Fourier

transform (by adding the frequency information necessary to form the sharp edges of the discontinuities). To lessen this effect, when performing a band-pass filter we utilize a technique suggested by Mark Henesian wherein the data array is placed in one quadrant of a 4x larger array, and the array flipped three times to fill the other quadrants, as shown in Figure 7. This approach makes the edges of the large array value-continuous (though not slope-continuous). This larger array is then transformed, filtered, and inverse-transformed. The original quadrant is then cut back out, giving the original array, now filtered.

Because the array being Fourier-transformed is larger code runtime is adversely affected. The offsetting benefit is the increased accuracy, due to reduced Gibb's noise, and the ability to better believe the information close to the edges of the array being transformed.

That is, the alternative method of reducing Gibb's noise from edge discontinuities is to use a Hanning window (see comment L). By its nature, however the Hanning window (and its cousins) is a weighting function, whereby the edges of the array are suppressed relative to the center. The resulting value for PSD1 carries this bias. If the array is homogeneous, this is not important. If, however, there is significant inhomogeneity across the part in the PSD1 band, then the difference in these two filtering approaches becomes evident.

Two examples are shown in Figures 8 and 9. Figure 8 shows a measurement where the phase is more homogeneous. The PSD1 value for the two approaches is the same to within ~ 10%. Figure 9, however, is of a measurement where there was relatively more phase noise along the edge of the part. In this case the two approaches show a ~100% discrepancy.

The quad-flip technique can add inaccuracies of its own, however, especially in the farfield (spectral domain). Figure 10 is an example of such a case. Using a Hanning window the periodic structures running at an angle through the phase array show up at isolated locations in the 2D spectral plot, with v_x and v_y having either both positive or negative values (see top graphics in Figure 10). Performing a quad-flip before the transform, however, now places this signal in all four quadrants of the 2D spectrum (see bottom graphics in Figure 10). Because of this possible confusion the code does not show the 2D spectrum from the quad-flip routine. Rather, the 2D and 1D spectrums reported by the code use a Hanning window, as discussed in comment L below.

J) Note on error-function-edge filter. The spectral filters used in calculating both the filtered rms gradient and in the band-pass filtering for PSD1 use a soft edge to minimize Gibb's ringing. The edge is constructed using the error function, though other filter shapes would likely work similarly well for these purposes.

For x and y spatial frequencies of k_x and k_y, and a filter cut-off frequency of k_c, the circular filter is defined as follows:

$$k_r = \sqrt{k_x^2 + k_y^2},$$

$$k_0 = 2./0.1 = 20. \quad \text{(error function} \sim 1.0 \text{ at argument of } 2.0)$$

$$f = 0.5\left(1 - erf\left(k_0\left|\frac{k_r}{k_c} - 1\right|\right)\right) \text{ for } k_r > k_c, \text{ and}$$

$$f = 0.5\left(1 + erf\left(k_0\left|\frac{k_r}{k_c} - 1\right|\right)\right) \text{ for } k_r < k_c.$$

This is shown in Figure 11.

K) Note on calculating rms from band-passed nearfield. The specification PSD1 is the rms phase in the 2.5-33 mm band. It is determined by performing a quad-flip band-pass, as described in comment K, and then simply calculating the rms of the filtered phase. For specification purposes, the rms of the entire effective aperture is used and reported as PSD1. However, because a local defect can heavily influence the overall PSD1, the rms in ~16 patches of ~100 x 100 mm size are also determined and reported. The rms in the worst of these subapertures is reported in the summary statistics. While not a specification, this information may prove useful in identifying parts with local finishing marks or inhomogeneities which may be detrimental to the laser, but which still pass the PSD1 specification over the entire optic.

In the case of normal-incidence RWF or TWF measurements, background subtraction, if present, is done before filtering. In the case of non-normal incidence RWF measurements, background subtraction is performed by determining the rms in the worst subaperture of the measurement array, and statistically subtracting PSD1 of the entire background file from this subaperture value.

L) Note on Hanning windows. To reduce noise from edge discontinuities, as discussed in comment I above, the band-pass filtering for PSD1 calculations uses a quad-flip technique. That technique, however, can distort structure in the frequency domain (see Figure 10). Consequently, for calculations of the 2D PSD map, running rms^2 plot, and the 1D collapse lineouts, the code uses a Hanning window to reduce edge effects.

The Hanning window is simply a weighting function applied to the data before performing a Fourier transform which drives the edges of the data to or near zero (thus eliminating discontinuities). The filter is defined as

$$h(x,y) = \frac{1}{4}\left(1 + \cos(\nu_{0x})\right)\left(1 + \cos(\nu_{0y})\right),$$

where ν_0 is one cycle across the aperture. The use of this window is shown in Figure 12.

Use of the Hanning window, which weights the center of the array more than the edges, assumes that parts are relatively uniform for purposes of the above calculations.

Because the Hanning window decreases the phase magnitude everywhere (except at the array center, where its value is 1.0), it represents a signal loss. Obviously, with such a windowing function, the effect on a parameter such as the peak-to-valley phase will vary from measurement to measurement, and is not derivable. The change in the phase rms, however, being a statistical quantity, is approximately constant for typical interferometer measurements, *if the rms only includes spatial wavelengths relatively short compared to the aperture size*. The PSD1 specification falls in this category, where the low-pass cut-off of 3.3 cm is short compared to the approximate aperture size of 40 cm. Specifically, this loss in a 2D transform is 8/3.

Numerically, the validity of this factor can be seen by considering a 40 x 40 cm surface constructed of a single frequency in x and y, and by simply comparing the rms of the surface before and after the Hanning window is applied. We do this as a function of the aberration wavelength, λ, and the phase shifts in x and y (ϕ_x and ϕ_y):

$$f(x,y) = \cos\left(\frac{2\pi x}{\lambda} + \phi_x\right)\cos\left(\frac{2\pi y}{\lambda} + \phi_y\right).$$

$$\begin{aligned} G(\nu_x,\nu_y) = & \frac{1}{4}F(\nu_x,\nu_y) + \\ & \frac{1}{8}\left(F(\nu_x-\nu_{0x},\nu_y) + F(\nu_x+\nu_{0x},\nu_y) + F(\nu_x,\nu_y+\nu_{0y}) + F(\nu_x,\nu_y-\nu_{0y})\right) + \\ & \frac{1}{16}\begin{pmatrix} F(\nu_x-\nu_{0x},\nu_y+\nu_{0y}) + F(\nu_x+\nu_{0x},\nu_y+\nu_{0y}) + \\ F(\nu_x+\nu_{0x},\nu_y-\nu_{0y}) + F(\nu_x-\nu_{0x},\nu_y-\nu_{0y}) \end{pmatrix} \end{aligned} \qquad \text{L (1)}$$

The ratio of the rms of this surface multiplied by a Hanning window with the rms of the surface alone is shown in Figure 17. The vertical lines at each wavelength represent 10 randomly chosen values for ϕ_x and ϕ_y, At longer wavelengths there is significant spread in the rms ratio both as a function of wavelength, and of phase. These long-wavelength features are not well-resolved in the aperture, and depending on the relation between the phase of the cosine function and the Hanning window, the resulting rms can vary significantly. At short wavelengths, however, this ratio asymptotes to 3/8.

Some intuition for the analytical basis for this factor can be gained by considering a single frequency ripple. If the Fourier transform of a function $f(x)$ is $F(\nu_x)$, then the Fourier transform of the $f(x)*\cos(\nu_{0x})$ is $(F(\nu_x - \nu_{0x}) + F(\nu_x + \nu_{0x}))/2$. A similar splitting occurs in the ν_y dimension, such that a single frequency ripple (ν_x, ν_y) will, when multiplied by a Hanning window, be split into nine frequencies of reduced magnitude. Defining $f(x,y)$ as the phase, $F(\nu_x, \nu_y)$ as the Fourier transform of the phase, $h(x,y)$ as the Hanning window, and $G(\nu_x, \nu_y)$ as the Fourier transform of the phase with the Hanning window applied (f*h):

For an array consisting of many frequencies, as in interferometer measurements, this splitting occurs for each individual frequency in the spectrum.

The rms over a frequency band is defined as

$$rms_{f\cdot h} \equiv \sqrt{\iint_{band} G(\nu_x,\nu_y) * G(\nu_x,\nu_y) d\nu_x d\nu_y} \qquad \text{(L2)}$$

for the windowed function, and

$$rms_f \equiv \sqrt{\iint_{band} F(\nu_x,\nu_y) * F(\nu_x,\nu_y) d\nu_x d\nu_y}$$

for the non-windowed function.

We put (L1) into (L2), and observe that all identically-shifted product terms all have the same integral:

$$\begin{aligned}
&\iint_{band} F(\nu_x - \nu_{x0},\nu_y) * F(\nu_x - \nu_{x0},\nu_y) d\nu_x d\nu_y \\
&= \iint_{band} F(\nu_x + \nu_{x0},\nu_y) * F(\nu_x + \nu_{x0},\nu_y) d\nu_x d\nu_y \\
&= \iint_{band} F(\nu_x,\nu_y) * F(\nu_x,\nu_y) d\nu_x d\nu_y .
\end{aligned}$$

The remaining terms are those containing the integral over F multiplied by F shifted. For a single-frequency ripple these are all zero. For example,

$$\iint_{band} F(\nu_x,\nu_y) * F(\nu_x - \nu_{x0},\nu_y) d\nu_x d\nu_y = 0.$$

For spectra such as encountered in the phase measurements these terms are not zero. They are, however, typically seen to be much less than the integral of $F*F$,

$$\iint_{band} F(\nu_x,\nu_y) * F(\nu_x - \nu_{x0},\nu_y) d\nu_x d\nu_y << \iint_{band} F(\nu_x,\nu_y) * F(\nu_x,\nu_y) d\nu_x d\nu_y,$$

allowing us to ignore them. With this assumption,

$$rms_{f\cdot h} = rms_f \sqrt{\left(\left(\frac{1}{4}\right)^2 + 4\left(\frac{1}{8}\right)^2 + 4\left(\frac{1}{16}\right)^2\right)} = 0.375\, rms_f, \text{or}$$

$$\frac{rms_f}{rms_{f\cdot h}} = 2.667.$$

In use, the code multiplies the windowed array by this correction factor before performing the Fourier transform and calculating the 2D PSD, running rms^2, and 1D collapse graphics.

M) Note on 1D collapse of the PSD. The 1D collapse of the 2D PSD is accomplished using the inverse Riemann sum (a discrete approximation to the Radon transform). This operation is the inverse of that used in tomography where several projections through a 2D image from different angles can be used to reconstruct the image. In this application we construct the projections themselves from the image, which is the 2D PSD. The projections are constructed by viewing the image from six different angles of 0, 30, 60, 90, 120, and 150 degrees, measured with respect to the positive x (horizontal) frequency axis, and them summing along these projection angles. This is shown graphically in Figure 14a. For example, the 1D collapse for an angle θ=0 of the 2D PSD of an aperture of width W and height H having *npw* x *nph* pixels is:

$$PSD_{1D} = \int_{-nph/2H}^{nph/2H} PSD_{2D}\, d\nu = \sum_{nph} PSD_{2D} \frac{1}{H} . \qquad \text{(M1)}$$

For other angles the limits of integration depend on the frequency bin into which the collapse is being done (see Figure 14a). The frequency increment $d\nu$ becomes

$$d\nu = \sqrt{\left(\frac{1}{W}\sin\theta\right)^2 + \left(\frac{1}{H}\cos\theta\right)^2} .$$

The utility of the 1D collapse is the ability to easily identify strong frequency components in the spectrum. For instance, in the 2D PSD of Figure 14a, a strong high-frequency component is present PSD at approximately at ν_x=0 and ν_y =+-1 mm^{-1}. In forming the 1D collapse this component becomes very evident as peaks in the collapse lines (Figure 14b). This example, however, also shows one aspect of the 1D collapse which can be confusing. In the 90 degree 1D collapse, this information collapses close to 1 mm^{-1} in frequency. However, as the angle of projection changes the information collapses into different frequency bins, until at the 0 degree projection it is near 0.02 mm^{-1}. This "shifting" of spectral information is inherent in the collapse, and calls for care in interpreting the collapse lines.

Most optics carry a "Not-to-Exceed" (NTE) specification, wherein none of the 1D collapse lines can exceed a certain line (typically 1.0v$^{-1.55}$) over the PSD1 frequency band. This specification is set to minimize the risk that the spectral noise as measured by PSD1 is not all contained in just one or a few frequency bands. Because the growth of phase noise can be a non-linear process, depending on the noise spectrum, it can be much worse for the laser to have a given rms noise compressed into one band, rather than spread across a spectrum. Consequently, the integral under the NTE line is larger than the PSD1 specification over that band.

Because of the "shifting" of spectral information described above, the 2D PSD is band-pass filtered to contain only frequencies between 2.5 and 33 mm before doing the 1D collapse. This collapse is then compared against the NTE line (Figure 14c). (The 1D

collapse of the unfiltered PSD is also shown in code output for information – Figure 14b.)

Figure 14b/c also show "data valid" lines. These are considered as guides to indicate the regime over which the spectral calculation is felt to be valid. The lower limit is set as three frequency cycles across the measurement aperture. Since the phase array is finite in size, lower-frequency terms in the transform (one and two cycles) are not considered to be accurately represented. The upper limit is set as four pixels per cycle. The transform contains frequency bins as high as two pixels per cycle ("Nyquist"), but we consider the sampling too poor to use this data. The measurements for which this software is intended have appropriate aperture and pixel count to put the PSD1 regime (2.5-33 mm) completely within the data-valid region.

The units of the 2D PSD are nm^2mm^2. The power spectrum (simply the Fourier transform of the phase array multiplied by its complex conjugate) has units of nm^2. The Power Spectral Density (PSD) array is the power spectrum per unit frequency. Since we use frequency units of mm^{-1}, the 2D PSD units become nm^2mm^2. The units for the 1D collapse PSD lines are nm^2mm, as can be seen from eqn. M1.

Since the fundamental frequency unit is $1/L_x$ or $1/L_y$, where L_x and L_y are the physical dimensions of the analysis aperture, both the 2D PSD and the 1D collapse would appear to scale with the aperture size of the measurement. For measurements of interest, and assuming the aperture width and height scale equally (as would happen when changing magnifications in an instrument), this is, in fact, not usually the case, as can be seen as follows.

Considering two different apertures: aperture 1 with width and height of L and np x np pixels, and aperture 2 having width and height of $L/2$ and $np/2$ x $np/2$ pixels. In the special case of a single frequency ripple that fits exactly (i.e. an integral number of cycles) on both of the apertures, we find that both the 2D PSD and the 1D collapse lines do scale with aperture. We define the two spectra as:

$$F1 = F2 = \delta(\nu_x, \nu_y)$$

and so

$$PSD_{2D,1} = (F1 * F1)L^2$$

$$PSD_{2D,2} = (F2 * F2)\left(\frac{L}{2}\right)^2,$$

so

$$\max(\mathrm{PSD}_{2D,1}) = 4 \cdot \max(PSD_{2D,2}).$$

The single-frequency 2D PSD of the larger aperture is 4x that of the smaller aperture.

The 1D collapse becomes:

$$PSD_{1D,1} = \sum_{np} PSD_{2D,1} \cdot \frac{1}{L} = \sum_{np} (F1 * F1)L$$

$$PSD_{1D,2} = \sum_{np/2} PSD_{2D,2} \cdot \frac{2}{L} = \sum_{np2} (F2 * F2)\frac{L}{2},$$

so

$$\max(PSD_{1D,1}) = 2 \cdot \max(PSD_{1D,2}).$$

The single-frequency 1D PSD of the larger array is 2x larger than that of the smaller aperture. Note that because this is a single-frequency ripple, only one frequency bin (+ and -) in each aperture is > 0, so the sum is really over two points in each case.

In the case of a spectrum with much content, as in the interferometer measurements, the spectral information is not confined to one frequency bin, and the power in the measurement is spread out over more bins in the larger aperture. For instance, in the above example, there are four times as many frequency bins in $PSD_{2D,1}$ as in $PSD_{2D,2}$. This means that the average amplitude in each bin is smaller in $PSD_{2D,1}$ than in $PSD_{2D,2}$, assuming the two apertures have equal rms values:

$$rms = \sqrt{\sum_{np}\sum_{np}(F1 * F1)} = \sqrt{\sum_{np/2}\sum_{np/2}(F2 * F2)}$$

or

$$\sum_{np}\sum_{np}(F1 * F1) = \sum_{np/2}\sum_{np/2}(F2 * F2).$$

To good approximation

$$\langle F1 * F1 \rangle = \frac{\langle F2 * F2 \rangle}{4}.$$

Again using the above example,

$$PSD_{2D,1} = (F1 * F1)L^2$$

and

$$PSD_{2D,2} = (F2 * F2)\left(\frac{L}{2}\right)^2,$$

so the mean value of both 2D PSD's is the same

$$\langle PSD_{2D,1} \rangle = \langle PSD_{2D,2} \rangle.$$

The mean values for the two 1D PSD's are also the same

$$PSD_{1D,1} = \sum_{np} PSD_{2D,1} \cdot \frac{1}{L} = \sum_{np} (F1 * F1)L = np\langle F1 * F1\rangle L$$

and

$$PSD_{1D,2} = \sum_{np/2} PSD_{2D,2} \cdot \frac{2}{L} = \sum_{np/2} (F2 * F2)\frac{L}{2} = \frac{np}{2}\langle F2 * F2\rangle \frac{L}{2} = np\langle F1 * F1\rangle L.$$

N) Note on peak-to-valley calculations. Some specifications require a calculation of the phase peak-to-valley, either before or after low-pass filtering the phase. If the phase is unfiltered, it typically contains spurious noise which can dramatically effect the peak-to-valley. Consequently, instead of the true peak-to-valley, we find a "peak" which has 0.5% of the measurement pixels above this value, and a "valley" which has 0.5% of the pixels below this value, and report the difference as the "peak-to-valley (99%)". An example histogram is shown in Figure 15. The use of the 99% value, rather than 98% or 95%, was arbitrary. We report the true (100%) peak-to-valley in the case of low-pass-filtered data because the filtering removes the spurious noise.

O) Note on extrapolating power and astigmatism to full-slab values in PHOM. Two of the PHOM specifications for laser glass blanks dictate the amount of power and astigmatism in the entire part, reflecting the fact that a certain amount of such long-wavelength aberration can be polished out at the finishing vendor. Because the PHOM measurement is done at normal incidence, measurements on the laser glass blanks are done on the left side and right side of the slabs separately. The slab clear aperture is 752 x 400 mm, so the square measurement on each side covers approximately half the slab. Because "stitching" together interferograms is problematic, the code extrapolates from the half-part measurement in each case to infer what the full-part power and astigmatism would be. It does this separately with both measurements.

The power is extrapolated as the ratio of the diagonals of the measured and full apertures:

$$P_{p-v,full-ap} = P_{p-v,meas}\left(\frac{426}{283}\right)^2.$$

The astigmatism is extrapolated by determining the magnitude of the 0/90 and +-45 degree astigmatism terms, and then forming a fictitious surface with these terms over the measured aperture. Following Figure 16, the astigmatism is then calculated over an expanded square aperture, using the quadratic dependence of astigmatism. The inferred aperture portion of this area can then be extracted, and piston and tilt removed. The peak-to-valley of this final array is then compared against specification.

P) Note on gradient calculation. The gradient at each pixel *(i,j)* is a magnitude, which is always positive:

$$g_{i,j} = \sqrt{(g_{x,i,j})^2 + (g_{y,i,j})^2}$$

where $g_{x,i,j}$ is the slope in the x (horizontal) direction, and $g_{y,i,j}$ is the slope in the y (vertical) direction. These horizontal and vertical slopes are calculated using a 5-point formula suggested by M. Henesian (see NBS Circular 55, pg. 914), with correction for points near the array boundary:

for points away from edge:

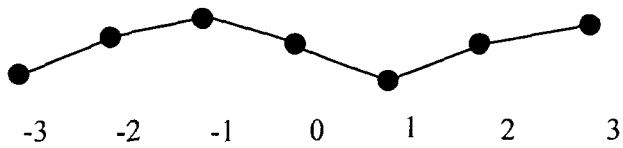

$$g_o = \frac{y_{-2} - y_2 + 8y_1 - 8y_{-1}}{3(x_2 - x_{-2})}$$

for points near edge:

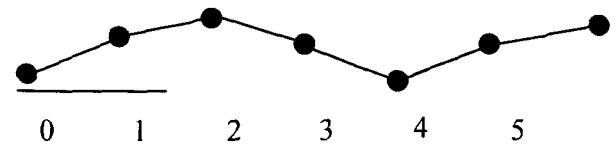

$$g_o = \frac{-25y_0 + 48y_1 - 36y_2 + 16y_3 - 3y_4}{3(x_4 - x_0)}$$

Numerical studies indicate that this 5-point derivative gives equivalent answers to 7, 9 and 19 points derivatives for the arrays of interest in these interferometer measurements.

Note also that after determining the gradient at each pixel, the code ignores a strip of 5 pixels around the edge of the array before calculating a histogram and rms. This reflects the doubt that the gradients near the edges (both from the gradient technique and from the low-pass filtering) are accurate.

Q) Note on LM1 influence function removal. At both the finishing and coating stages the rms gradient of the filtered wavefront for LM1 is calculated by first removing a best-fit of the adaptive optic influence functions from the wavefront, then filtering, and then performing the gradient calculations. This special allowance for this optic is made because it is thinner than other mirrors, and so more difficult to polish flat, and because LM1 will actually be used only in this self-correcting manner. The other specifications of filtered P-V and PSD1 do not have these influence functions removed.

Code to calculate this influence function fit was provided by Rick Sacks. It is accomplished by performing a least-squares fit between the wavefront and a set of 39 independent, height-adjustable "bumps" representing the 39 actuators of the NIF adaptive optic. The shape of each of these bumps is stored in a file, and was derived from a smoothed fit to measured influence functions from the NIF prototype mirror. An example of one such influence function is shown in Figure 18. The entire influence function is simply scaled up or down (i.e., multiplied by an adjustable constant).

The least squares fitting is accomplished as follows. Given

$a(x,y)$: measured wavefront (interferogram)

$f_j(x,y)$: influence function shapes for 39 actuators ($j = 1:39$), and

h_j : height coefficieint of each influence function

we want the minimum of

$$\sum_x \sum_y \left[a(x,y) - \sum_j h_j \cdot f_j(x,y) \right]^2 .$$

This is found through a simple matrix inversion:

$$h_j = \sum_k S_{j,k}^{-1} \cdot B_k ,$$

where

$$S_{j,k} = \sum_x \sum_y f_j(x,y) \cdot f_k(x,y), \text{ and}$$

$$B_j = \sum_x \sum_y a(x,y) \cdot f_j(x,y).$$

The residual wavefront after the best-fit subtraction then becomes:

$$r(x,y) = a(x,y) - \sum_j h_j \cdot f_j(x,y).$$

This approach is more computationally efficient than the alternative of a Hartman sensor loop calculation, which would be more representative of the actual operation of the NIF mirror. Tests using such a code, however, (provided by Mark Henesian and Rick Sacks) have shown the two approaches to give similar results, with the best-fit algorithm consistently giving slightly lower residual filtered rms gradient.

R) Note on ramped phase jump criteria for crystals.

Crystals manifest the unique characteristic of phase boundaries due to abrupt index changes in the material. When a beam is propagated through this phase boundary it acquires a jump in its phase distribution. As the beam is propagated further the phase jump will turn into intensity modulation. The magnitude of this phase change and its transitional width determine how severe will be the resultant intensity modulation. To mitigate against this intensity modulation a "ramped phase jump criterion" was developed by Jerry Auerbach and Janice Lawson. It was based on a calculation of the intensity modulation due to ramped phase jumps of various magnitudes and transition widths. Two situations were considered:

(a) The propagation of a 3w beam with a ramped phase jump from the tripler crystal through the 3w optics. The ramped phase jump width and height (see figure) were adjusted so that the modulation downstream from the tripler crystal would be <= 1.3:1.

(b) The propagation of a 1w beam with a ramped phase jump during the 4th pass through the cavity amplifier. The propagation extended from the PEPC crystal to the input lens of the transport spatial filter. The ramped phase jump width and height were adjusted so that the modulation at the transport spatial filter input would be <= 1.2:1.

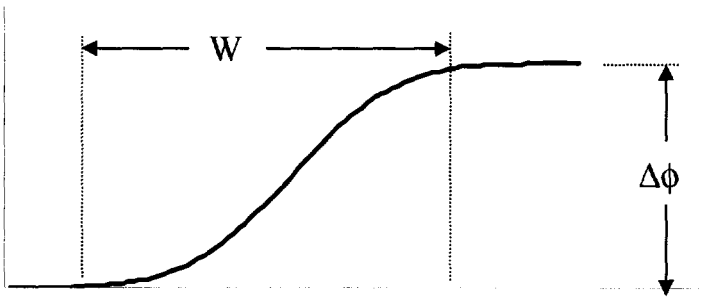

The results are incorporated into the interferometer measurements as a limit on the phase difference between any two closely-separated points on the phase (which roughly represents a local phase slope). Specifically, for the doubler and tripler crystals this limit is:

$$\Delta\phi < 0.029 + 0.016W + 0.022W^2 \text{ for } 0 < W < 3$$

and

$$\Delta\phi < 0.28 \text{ for } 3 < W < 33,$$

where W is the distance between the two points in (mm) and $\Delta\phi$ is the phase difference between the points (633 nm waves).

The specification for the switch crystal is slightly tighter:

$$\Delta\phi < 0.025 + 0.015W^2 \text{ for } 0 < W < 3$$

and

$$\Delta\phi < 0.16 \text{ for } 3 < W < 33,$$

Analysis against this criterion is done by considering two annuli around every valid point in the aperture: 0.5 – 3 mm, and 3 – 33 mm, which correspond to the ranges above. (The lower limit on the smaller annulus is done to lesson calculational noise due to the ~0.4 mm pixel spacing in the Wyko interferometers.) For the smaller annulus, every point in the annulus is differenced with the center point and compared against the spec. For the larger annulus, only every 5th point is used (to reduce computational time). The results are displayed as graphics shown in Appendix D. 1% of the valid pixels are allowed to fail the criterion, per the drawings.

S) Note on Orthogonal Polarization Interferometry (OPI) for crystals.

OPI is an interferometeric method of characterizing frequency conversion crystals, whereby vertically and horizontally polarized transmitted wavefront measurements, along with an empty cavity measurement, are used to estimate the spatial variations of the extraordinary index of refraction in the crystal, and, relatedly, variations in the local phase-matching angle [S1,S2]. The theory behind this technique will only be sketched here, sufficient to understand the workings of the code.

The crystals used for frequency doubling and tripling the NIF laser are potassium dihydrogen phosphate (KDP, or its partially deuterated analog KD*P), machined in the form of plates such that the principle ("c") axis of the crystal is at an angle with respect to the plate normal. KDP is a birefringent material, which means that an incident light wave of arbitrary direction and polarization propagates in the material as two orthogonally polarized waves- one polarized perpendicular to the plane formed by the "c" axis and the crystal surface normal, referred to as the "ordinary" wave, and one polarized in this plane, referred to as the "extraordinary" wave. The index of refraction seen by the ordinary wave depends on wavelength and temperature in the usual sense. The index of refraction of the extraordinary wave has an additional dependence on beam propagation direction *k*, thus the designation extraordinary. For the NIF plate geometry, the "ordinary" and "extraordinary" polarization directions are as shown below.

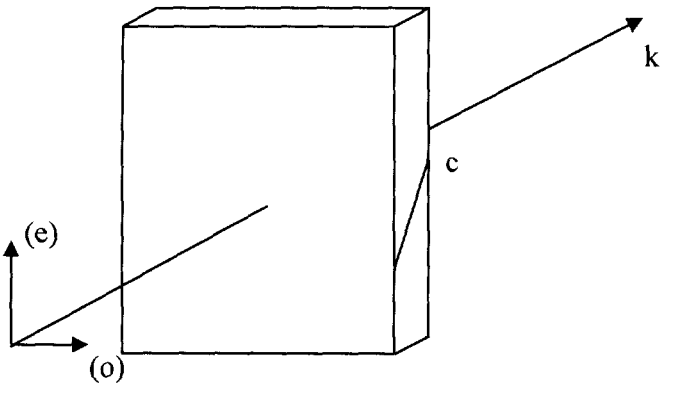

The dependence of the extraordinary index of refraction on the angle θ between k and the c axis is given by

$$\frac{1}{n_e(\theta)^2} = \frac{\sin^2(\theta)}{n_c(\theta)^2} + \frac{\cos^2(\theta)}{n_o(\theta)^2}.$$

Values of n_o and n_c for KDP and 92% KD*P are tabulated at the wavelengths of interest using the Sellmeier formula. Values for intermediate deuteration fractions are interpolated. This equation is differentiated to find $dn_e/d\theta$, which will be used below.

Two interferometer measurements are made with the interferometer laser having an "e" and then an "o" polarization. The part is not adjusted between measurements. A scaled difference between these two measurements gives the local difference in the extraordinary index of refraction at each point in the crystal, as follows.

For both polarizations, and at each point in the crystal, the measured TWF is given by

$$\phi(x,y) = (n(x,y)-1)t(x,y) + C(x,y) + T(x,y) + P(x,y)$$

where

$$n(x,y) = \bar{n} + \delta n(x,y) \text{ and}$$
$$t(x,y) = \bar{t} + \delta t(x,y)$$

represent the local index of refraction and thickness, respectfully. The parameters C, T, and P represent contributions to the local TWF due to the empty cavity, tilt and piston in the measurement. The mean values of n_e and n_o are tabulated (or calculated) as described above. The mean thickness must also be known.

Subtracting the "o" measurement from the "e" measurement, and removing piston (the mean phase) we find (dropping the x,y subscripts):

$$\begin{aligned}\Delta\phi &= \phi_e - \phi_o \\ &= (\bar{n}_e + \delta n_e - 1)(\bar{t} + \delta t) + C + T_e + P_e - (\bar{n}_o + \delta n_o - 1)(\bar{t} + \delta t) + C + T_o + P_o \\ &\approx (\bar{n}_e - \bar{n}_o)\delta t + (\delta n_e - \delta n_o)\bar{t} + \Delta T_{eo},\end{aligned}$$

where we have ignored all terms constant across the aperture, and $\Delta T_{eo}=T_e-T_o$.

The net tilt term, ΔT_{eo}, is the change in tilt of the part between the two measurements. Since the part is not moved between measurements this term is close to zero. It can,

however, be non-zero for the phase-shifting interferometers we use. The phase-shifting process involves physically shifting the transmission flat back and forth a small distance rapidly in time. This movement may not be entirely uniform across the surface of the flat, with one edge shifting slightly less or more than another. This results in a tilt being added to the phase. (Higher order vibrational modes would, presumably, also cause errors. They are believed to be small, however.) Because there is typically an arbitrary tilt in any phase measurement, this error is usually not important. For these difference measurements, however, a tilt difference from one measurement to the next would appear as a linear variation in birefringence across the part. Consequently, an extra step is required to eliminate this tilt error.

The information needed to correct tilt error is obtained by making sure that a portion of the reflection flat is visible to the camera during the TWF measurements. That is, the crystal under measurement must under-fill the aperture somewhat. The subregion of the image where the reflection flat is visible may actually consist of several disconnected subregions. The tilt error (ΔT_{eo}) is then determined by computing a best-fit plane to the difference wavefront in these subregions.

A minimum size for the subregion has not been determined, but it must contain enough pixels to give a reasonable plane fit. The wavefront difference, corrected for tilt error is then $\Delta\phi' \equiv \Delta\phi - \Delta T_{eo}$.

The thickness variation term δt can be estimated from the difference between the ordinary-polarized measurement and the empty cavity measurement:

$$\begin{aligned}\phi_o' &= \phi_o - C \\ &= (\bar{n}_o + \delta n_o - 1)(\bar{t} + \delta t) + T_o - T_c + P \\ &\approx (\bar{n}_o - 1)\delta t + \Delta T_{co},\end{aligned}$$

where it has been assumed that δn_o is small. This assumption is consistent with the OPI approximation that the measured variations in birefringence are due to directional wander of the c axis and therefore variations in n_e. The tilt change between the two measurements, ΔT_{co}, is calculated as above, using the clear area around the part.

Combining the above relations, and, again, assuming δn_o to be small

$$\Delta\phi' - \frac{\bar{n}_e - \bar{n}_o}{\bar{n}_o - 1}\phi_o' = \bar{t}\cdot\left(\delta n_e - \frac{\bar{n}_e - \bar{n}_o}{\bar{n}_o - 1}\delta n_o\right) \approx \delta n_e \bar{t},$$

or

$$\delta n_e = \left(\phi_e - \phi_o - \Delta T_{eo} - \left(\frac{\bar{n}_e - \bar{n}_o}{\bar{n}_o - 1}\right)(\phi_o - C - \Delta T_{co})\right)\left(\frac{1}{\bar{t}}\right) + const.$$

This is the spatially-resolved non-uniformity in extraordinary index. It is converted into a non-uniformity in detuning angle using $dn_e / d\theta$ from above

$$\Delta\theta = \delta n_e \left(\frac{dn_e}{d\theta}\right)^{-1}.$$

The zero-mean rms of the detuning angle across the part is converted back to inhomogeneity according to the above equation and compared against the specification on the drawing.

This OPI information is also used to do final surface cuts on the crystal, in conjunction with the Crystal Alignment Test System (CATS). The CATS system employs harmonic generation and the angular-dependence of phase matching to measure and report Δc, the difference between the actual c-axis orientation and the desired (nominal) orientation c_0 inferred from a co-aligned proof crystal

$$\Delta c_{CATS} = c - c_0.$$

The c-axis is defined such that $c=0$ is the normal to the surface. A correctly cut, homogeneous crystal would have $\Delta c_{CATS} = 0$ at all measurement points. Because there is inhomogeneity, however, for a correctly cut crystal we would have a spread in Δc_{CATS} with $\langle\Delta c_{CATS}\rangle = 0$, given a large number of measurement points across the crystal surface. For an incorrectly cut, inhomogeneous crystal $\langle\Delta c_{CATS}\rangle$ would indicate the mean surface tilt, again assuming enough measurement points.

Because we only take a limited number of measurements in CATS, however, it is useful to use OPI data. OPI uses a full-aperture measurement of birefringence to infer and report a variation in Δc:

$$\Delta c_{OPI} = c - c_0 + P\,(arb.).$$

The piston term, P, is arbitrary in magnitude, and is an artifact of the interferometer measurements in OPI. We can remove the piston term, however, by normalizing against the CATS results. Notionally, the two measurements might look as follows:

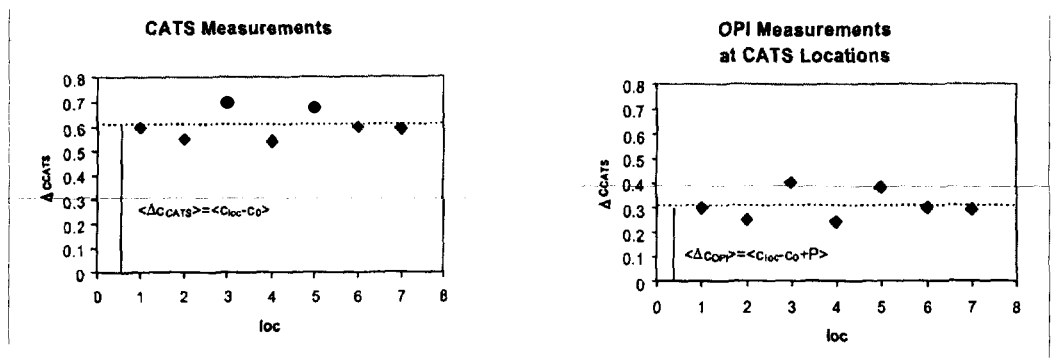

The only difference is the piston term, P. It is equal to $<\Delta c_{OPI}>$ - $<\Delta c_{CATS}>$, calculated at the CATS measurement points. Consequently, the OPI results can be re-defined as

$$\Delta c_{OPI} = c - c_0 + \left(\langle \Delta c_{OPI} \rangle_{CATS} - \langle \Delta c_{CATS} \rangle \right).$$

The desired correction is the average of c-c_0 over the part, which can be calculated as

$$\langle \Delta c \rangle = \langle c - c_0 \rangle = \langle \Delta c_{OPI} \rangle_{OPI} - \left(\langle \Delta c_{OPI} \rangle_{CATS} - \langle \Delta c_{CATS} \rangle\right) = \langle \Delta c_{OPI} \rangle + \left(\langle \Delta c_{CATS} \rangle - \langle \Delta c_{OPI} \rangle_{CATS}\right).$$

Note that $<\Delta c>$ will be in units of crystal *internal* angle.

The required machining correction for the plate is therefore:

$<\Delta c> > 0$: Machine crystal to reduce angle between surface normal and c-axis by $<\Delta c>$ (internal angle).
$<\Delta c> < 0$: Machine crystal to increase angle between surface normal and c-axis by $<\Delta c>$ (internal angle).

[S1] P. Wegner, "Reorientation of proof plates RG8B and RGA8- Round 2," NIF-0037314, Oct 6, 1999.

[S2] J. Auerbach et al., "Modeling of frequency doubling and tripling with measured crystal spatial refractive index non-uniformities," UCRL-JC-139333, 2000.

Code Outputs

Results from calculations are displayed on the screen as graphics and text. During execution the codes allow the compilated results to be sent directly to a printer. Additionally, the codes automatically generate a Postscript (.ps) graphics file containing the same results as printed. Because these files can be rather large (~36 MB for a two-page output), these files can be converted to Adobe Acrobat® format (.pdf), using the commercial code Adobe Acrobat Distiller®. This results in a file of a little over 1 MB in size, which can be more conveniently e-mailed, etc. Sample printouts for each of the codes are given in the appendices, along with descriptions of the information contained in the printouts.

The codes also produce a short summary of the code results in the form of an Excel®-compatible (.csv) file. This file contains the names of the relevant specifications, and the calculational results for the part under test. This file can be "cut-and-pasted" into a spreadsheet, so aids in the transfer of results.

Finally, in "process control" mode, the codes allow the user to print graphics one at a time, and to adjust the numerical scales. This enables a better visual comparison between different parts, or during sequential finishing steps on one part. In this mode the 1D PSD lineouts can also be saved to an Excel®–compatible (.csv) file, so the user can separately plot these results.

Acknowledgements

These codes were a follow-on to the code CVOS, written by Janice Lawson, which has many similar functionalities, and which is used extensively by NIF vendors and at LLNL. Mark Henesian and Janice contributed routines used in these codes, as well as many helpful discussions. Rick Sacks provided the adaptive optic routines influence function routines. The OPI analysis routine for crystals was taken from theory developed by Jerry Auerbach, et al. (ref. UCRL-JC-139333), and codes provided by Paul Wegner.

Summary of Wavefront Specifications for NIF Large Optics

red: Phase or Surface Contour Specifications
blue: Phase Gradient Specifications
green: PSD Specifications

Blanks									
Laser glass blank	PHOM power (.425 waves p-v)	PHOM astigmatism (.22 waves. p-v)	PHOM residual after subracting power and astig (.142 waves, p-v)	surface contour (5 waves p-v)	PHOM gradient after removing power and astig, and after 33 mm low-pass filter (.008333 waves/cm rms)				
Switch window blank	inhomogeniety after removal of up to 0.25 p-v waves of power and 1/12th p-v waves of irregularity (2 ppm p-v)	No discrete inhomogeneities are acceptable within the effective aperture: No 100mm subsection shall exceed PV=0.2waves; No 50mm subsection shall exceed a peak slope of 0.075 waves/cm.	TWF phase gradient after removal of up to 0.25 p-v waves of power, 1/12th wave of irregularity, and 33 mm low-pass filter (0.009 wave/cm rms)						
Target chamber vacuum window blank	inhomogeniety after removal of up to 0.25 p-v waves of power and 1/15th p-v waves of irregularity (2 ppm p-v)	No discrete inhomogeneities are acceptable within the effective aperture: No 100mm subsection shall exceed PV=0.2waves; No 50mm subsection shall exceed a peak slope of 0.075 waves/cm.	TWF phase gradient after removal of up to 0.25 p-v waves of power, 1/15th wave of irregularity, and 33 mm low-pass filter (0.009 wave/cm rms)						

Beamsplitter, DOP, debris shield blank	inhomogeniety after removal of up to 0.25 p-v waves of power (2 ppm p-v)	No discrete inhomogeneities are acceptable within the effective aperture: No 100mm subsection shall exceed PV=0.2waves; No 50mm subsection shall exceed a peak slope of 0.075 waves/cm.	TWF phase gradient after removal of up to 0.25 p-v waves of power and 33 mm low-pass filter (0.009 wave/cm rms)						
All lenses blank	inhomogeniety after removal of up to 0.25 p-v waves of power and up to 0.142 p-v waves of irregularity (2 ppm)	No discrete inhomogeneities are acceptable within the effective aperture: No 100mm subsection shall exceed PV=0.2waves; No 50mm subsection shall exceed a peak slope of 0.075 waves/cm.	TWF phase gradient after removal of up to 0.25 p-v waves of power and up to 0.142 waves of irregularity and 33 mm low-pass filter (0.009 wave/cm rms)						
Polarizer blank	none								
All mirrors blank	none								
Switch crystals, blank	TWF phase after 33 mm low-pass filter (0.17 waves p-v)	TWF phase gradient after 33 mm low-pass filter (11 nm/cm rms)	TWF ramped phase jump for small phase discontinuities, 0% of part to exceed (see note 9)						
Doubler crystals, blank	birefringence variation (1e-6 rms)	TWF phase gradient after 33 mm low-pass filter (11 nm/cm rms)	ramped phase jump for small phase discontinuities, 0% of part to exceed (see note 9)						
Tripler crystals, blank	birefringence variation ((2.99 - 0.5*D)e-6 rms, where D is deuteration fraction)	TWF phase gradient after 33 mm low-pass filter (11 nm/cm rms)	ramped phase jump for small phase discontinuities, <0% of part to exceed (see note 9)						
Finished									

Finished									
Laser glass, finished	TWF phase after 33 mm low-pass filter (0.3333 waves p-v)	TWF phase gradient after 33 mm low-pass filter (.01111 waves/cm rms)	TWF phase in 2.5-33 mm band (1.8 nm rms)	TWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	surface contour in 0.01 to 0.12 mm band (0.4 nm rms)	TWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band			
All lenses, all windows, beamsplitter, DOP, finished (except final focus lens and diagnostic beamsplitter)	TWF phase after 33 mm low-pass filter (0.3333 waves p-v)	TWF phase gradient after 33 mm low-pass filter (.01111 waves/cm rms)	TWF phase in 2.5-33 mm band (1.8 nm rms)	TWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	surface contour in 0.01 to 0.12 mm band (1.0 nm rms)	TWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band			
Final focus lens and diagnostic beamsplitter, finished	TWF phase after 33 mm low-pass filter (0.3333 waves p-v)	RWF surface power (5 waves p-v)	RWF surface p-v, after power subtraction (1 waves p-v)	TWF phase gradient after 33 mm low-pass filter (.01111 waves/cm rms)	TWF phase in 2.5-33 mm band (1.8 nm rms)	TWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	surface contour in 0.01 to 0.12 mm band (1.0 nm rms)	TWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band	
Debris shield, finished	TWF phase after 33 mm low-pass filter (0.3333 waves p-v)	TWF phase gradient after 33 mm low-pass filter (.01111 waves/cm rms)	TWF phase in 2.5-33 mm band (2.5 nm rms)	TWF phase in 0.12 to 2.5 mm band (1.6 nm rms)	surface contour in 0.01 to 0.12 mm band (1.0 nm rms)	TWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band			
Polarizer, finished	TWF and RWF phase after 33 mm low-pass filter (0.3333 waves p-v)	TWF and RWF phase gradient after 33 mm low-pass filter (.01111 waves/cm rms)	TWF and RWF phase in 2.5-33 mm band (1.8 nm rms)	TWF and RWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	surface contour in 0.01 to 0.12 mm band (1.0 nm rms)	TWF and RWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band			
All mirrors except Mirror 1, finished	RWF phase after 33 mm low-pass filter (0.3333 waves p-v)	RWF phase gradient (after 33 mm low-pass filter) (.01111 waves/cm rms)	RWF phase in 2.5-33 mm band (1.8 nm rms)	RWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	surface contour in 0.01 to 0.12 mm band (0.4 nm rms)	RWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band			
Mirror 1 (adaptive optic), finished	RWF phase after 33 mm low-pass filter (0.5 waves p-v)	RWF phase gradient after 33 mm low-pass filter (0.0111 waves/cm rms)	RWF phase in 2.5-33 mm band (1.8 nm rms)	RWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	surface contour in 0.01 to 0.12 mm band (0.4 nm rms)	RWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band			
Switch Crystals, finished	Surface contour (5 waves p-v)	TWF phase gradient after 33 mm low-pass filter (11 nm/cm rms)	TWF ramped phase jump for small phase discontinuities, 0% of part to exceed (see note 9)	TWF phase in 6 - 33 mm band (6.4 nm rms, 3.5 nm rms goal)	TWF 1-D PSD collapse < 0.791v^-2.701 nm^2mm in 6 - 33 mm band	Surface contour in 0.5 - 6 mm band (4.2 nm rms); 1-D collapse < 19.008v^-0.927 nm^2mm	Surface contour in 0.12 - 0.5 mm band (2.0 nm rms); 1-D collapse < 54.911v^-2.457 nm^2mm	Surface contour in 0.0025 - 0.12 mm band (3.0 nm rms)	
Doubler crystals, finished	birefringence variation (1e-6 rms)	Surface contour (5 waves p-v)	TWF phase gradient after 33 mm low-pass filter (11 nm/cm rms)	TWF ramped phase jump for small phase discontinuities, 0% of part to exceed (see note 9)	TWF phase in 6 - 33 mm band (6.4 nm rms)	TWF 1-D PSD collapse < 0.791v^-2.701 nm^2mm in 6 - 33 mm band	Surface contour in 0.5 - 6 mm band (4.2 nm rms); 1-D collapse < 19.008v^-0.927 nm^2mm	Surface contour in 0.12 - 0.5 mm band (2.0 nm rms); 1-D collapse < 54.911v^-2.457 nm^2mm	Surface contour in 0.0025 - 0.12 mm band (3.0 nm rms)
Tripler crystals, finished	birefringence variation ((2.99 - 0.5*D)e-6 rms, where D is deuteration fraction)	Surface contour (5 waves p-v)	TWF phase gradient after 33 mm low-pass filter (11 nm/cm rms)	TWF ramped phase jump for small phase discontinuities, 0% of part to exceed (see note 9)	TWF phase in 6 - 33 mm band (6.4 nm rms)	TWF 1-D PSD collapse < 0.791v^-2.701 nm^2mm in 6 - 33 mm band	Surface contour in 0.5 - 6 mm band (4.2 nm rms); 1-D collapse < 19.008v^-0.927 nm^2mm	Surface contour in 0.12 - 0.5 mm band (2.0 nm rms); 1-D collapse < 54.911v^-2.457 nm^2mm	Surface contour in 0.0025 - 0.12 mm band (3.0 nm rms)
Coated									

Coated									
Polarizer, finished and HR coated	RWF phase after 33 mm low-pass filter (0.5 waves p-v)	RWF phase gradient (after 33 mm low-pass filter) (.006667 waves/cm rms)	RWF phase in 2.5-33 mm band (1.8 nm rms)	RWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	RWF 1-D PSD collapse < 1v^-1.55 in 0.12 to 33 mm band				
All mirrors except Mirror 1, finished and HR coated	RWF phase after 33 mm low-pass filter (0.3333 waves p-v)	RWF phase gradient (after 33 mm low-pass filter) (.006667 waves/cm rms)	RWF phase in 2.5-33 mm band (1.8 nm rms)	RWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	RWF 1-D PSD collapse < 1v^-1.55 in 0.12 to 33 mm band				
All crystals, coated	-none-								
Mirror 1 (adaptive optic), finished and HR coated	RWF phase after 33 mm low-pass filter (3 waves p-v)	RWF phase gradient after 33 mm low-pass filter (0.006667 waves/cm rms)	RWF phase in 2.5-33 mm band (1.8 nm rms)	RWF phase in 0.12 to 2.5 mm band (1.1 nm rms)	RWF 1-D PSD collapse < 1v^-1.55 nm^2mm in 0.12 to 33 mm band				

notes:

1) One wave=633 nm for all optics except for coated mirrors and polarizers (1064 nm), and for birefrengence variation in crystals (657.4 nm). For 633 nm specifications, measurements are taken with a laser diode near 633 nm, and simply converted to a 633 nm equivalent value not allowing for any wavelength variation in index. The 1064 nm specifications are taken with a 1064 nm interferometer. The crystal birefringence specifications are taken with a laser diode of that wavelength, as index control in the converter crystals is essential.
2) PHOM: homogeneity. Involves weighted combinations of TWF, RWF's, and cavity measurements.
3) TWF: transmitted wavefront.
4) RWF: reflected wavefront.
5) Irregularity: p-v after 100 mm low-pass filter.
6) Birefringence for crystal blanks = (OPD(e) - OPD(o)) / thickness, where OPD is measured at 657.4 nm.
7) Wavefront specifications for finished optics apply before AR coating, if applicable, for transmissive oiptics.
8) All specifications except those for rms roughness in the 0.01 to 0.12 mm band will be measured on large aperture interferometers, and, normally, for each optic in production. The latter will be measured on a high-magnification instrument, and will normally be done only occasionally for process verification.
9) ramped phase jump for small phase discontinuities: del_phi < 0.029+.016*w+.022*w*2 for 0<w<3 mm, and del_phi<0.28 for 3 < w < 33 mm, where w is linear distance across crystal, and del_phi is phase difference
10) These specifications are for reference only, and do not supercede the optics drawings.

Figure 1. Summary of large optics metrology specifications. These are for reference only, and do not supercede the optics drawings.

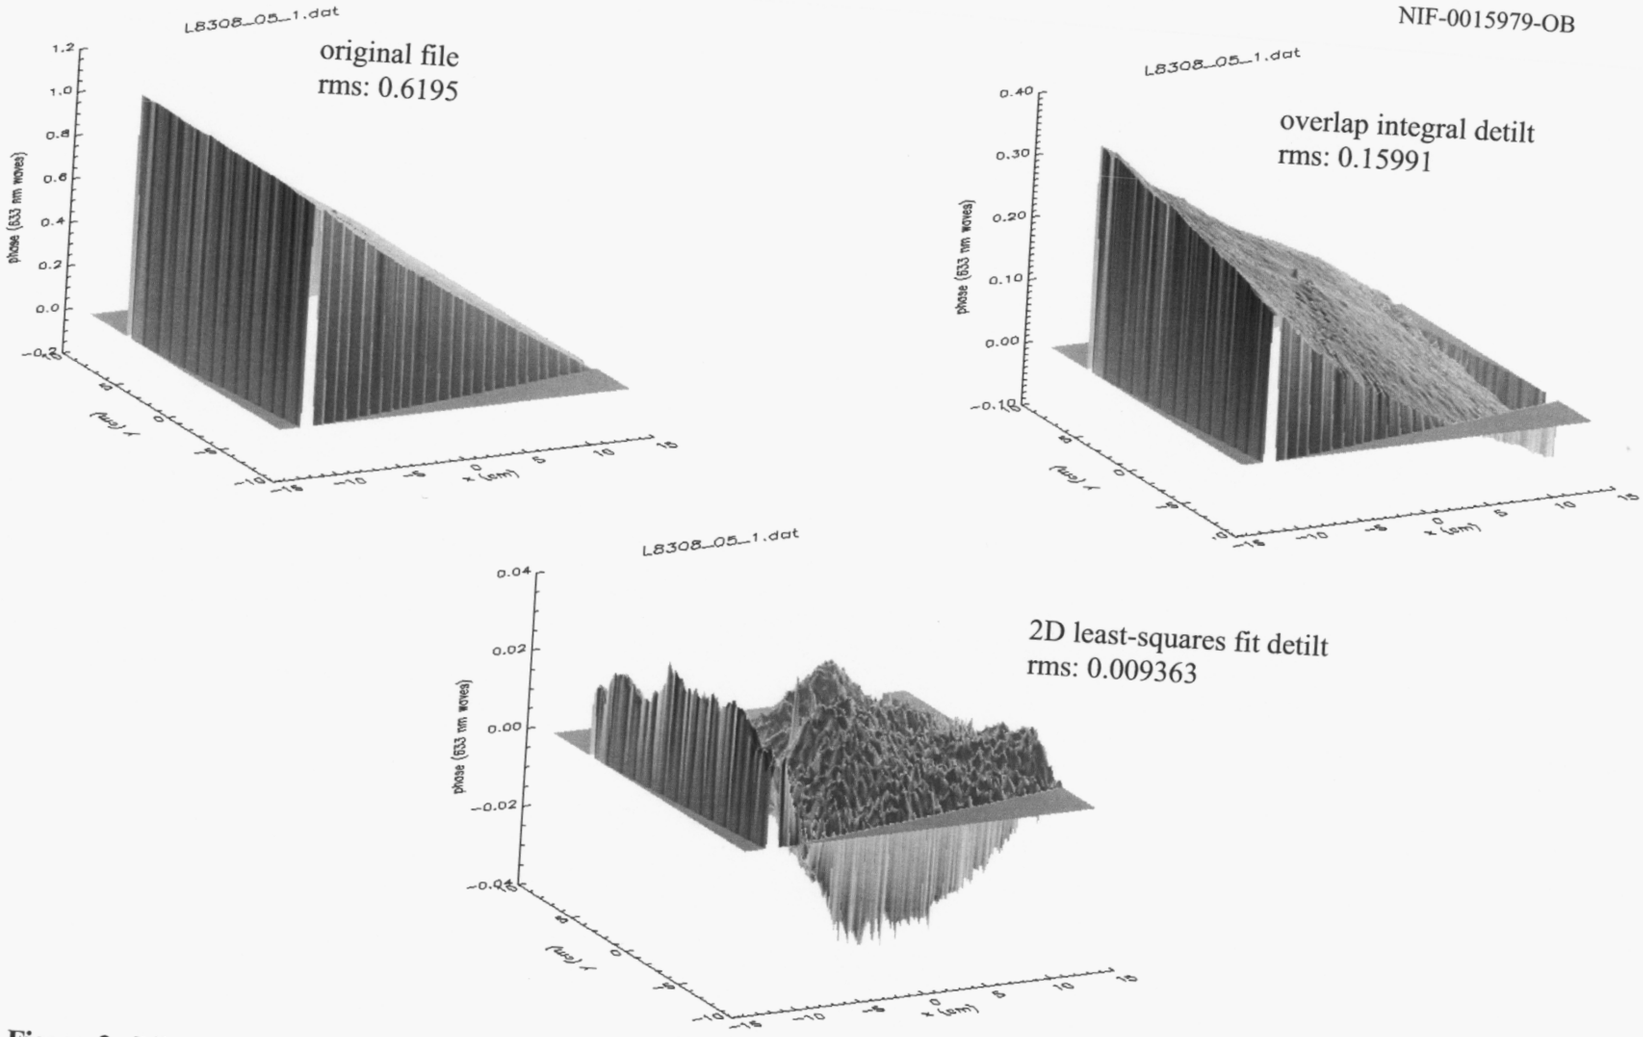

Figure 2. Tilt removal is done by performing a least-squares fit of x and y planes to the data. The least-squares calculation gives a much better fit than a simple overlap integral.

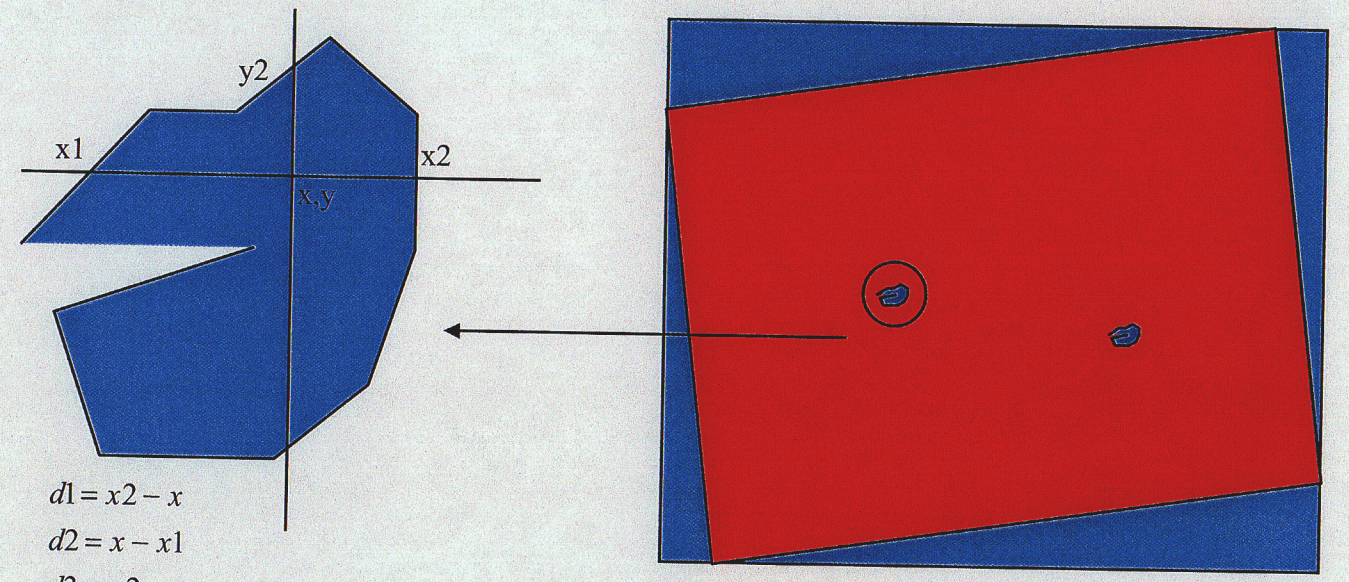

$$d1 = x2 - x$$

$$d2 = x - x1$$

$$d3 = y2 - y$$

$$d4 = y - y1$$

$$z(x,y) = \frac{z(x1,y) * d1 + z(x2,y) * d2 + z(x,y1) * d3 + z(x,y2) * d4}{d1 + d2 + d3 + d4}$$

Figure 4. Interpolation of invalid pixels is done using nearest four valid pixels in x and y directions.

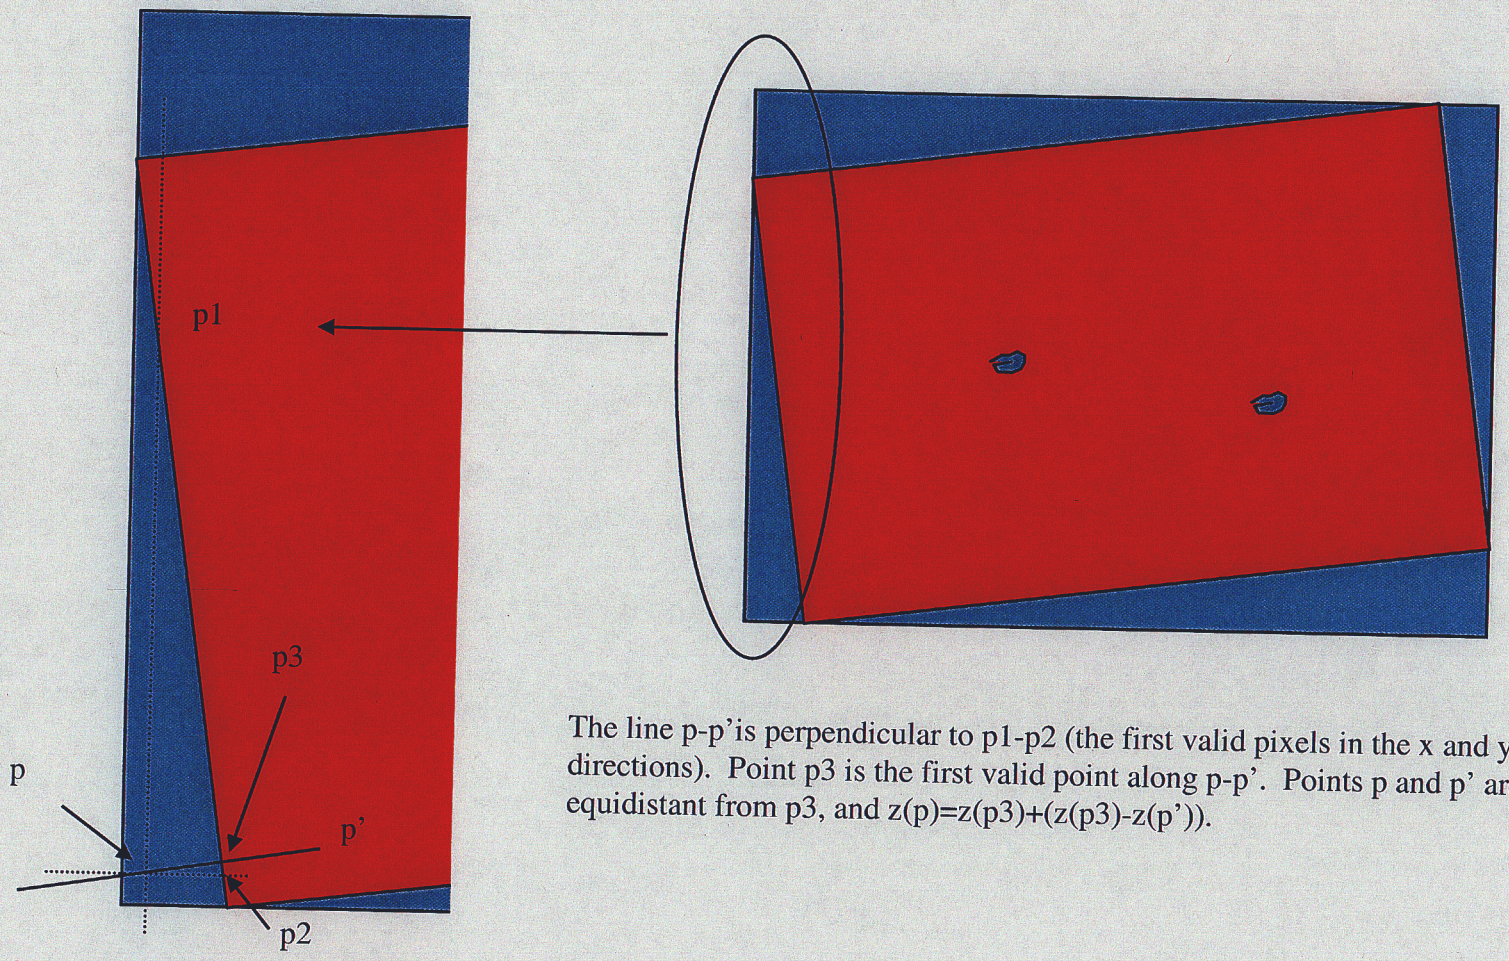

The line p-p'is perpendicular to p1-p2 (the first valid pixels in the x and y directions). Point p3 is the first valid point along p-p'. Points p and p' are equidistant from p3, and z(p)=z(p3)+(z(p3)-z(p')).

Figure 5. Extrapolation is done by attempting to reflect the value of a valid point across the boundary.

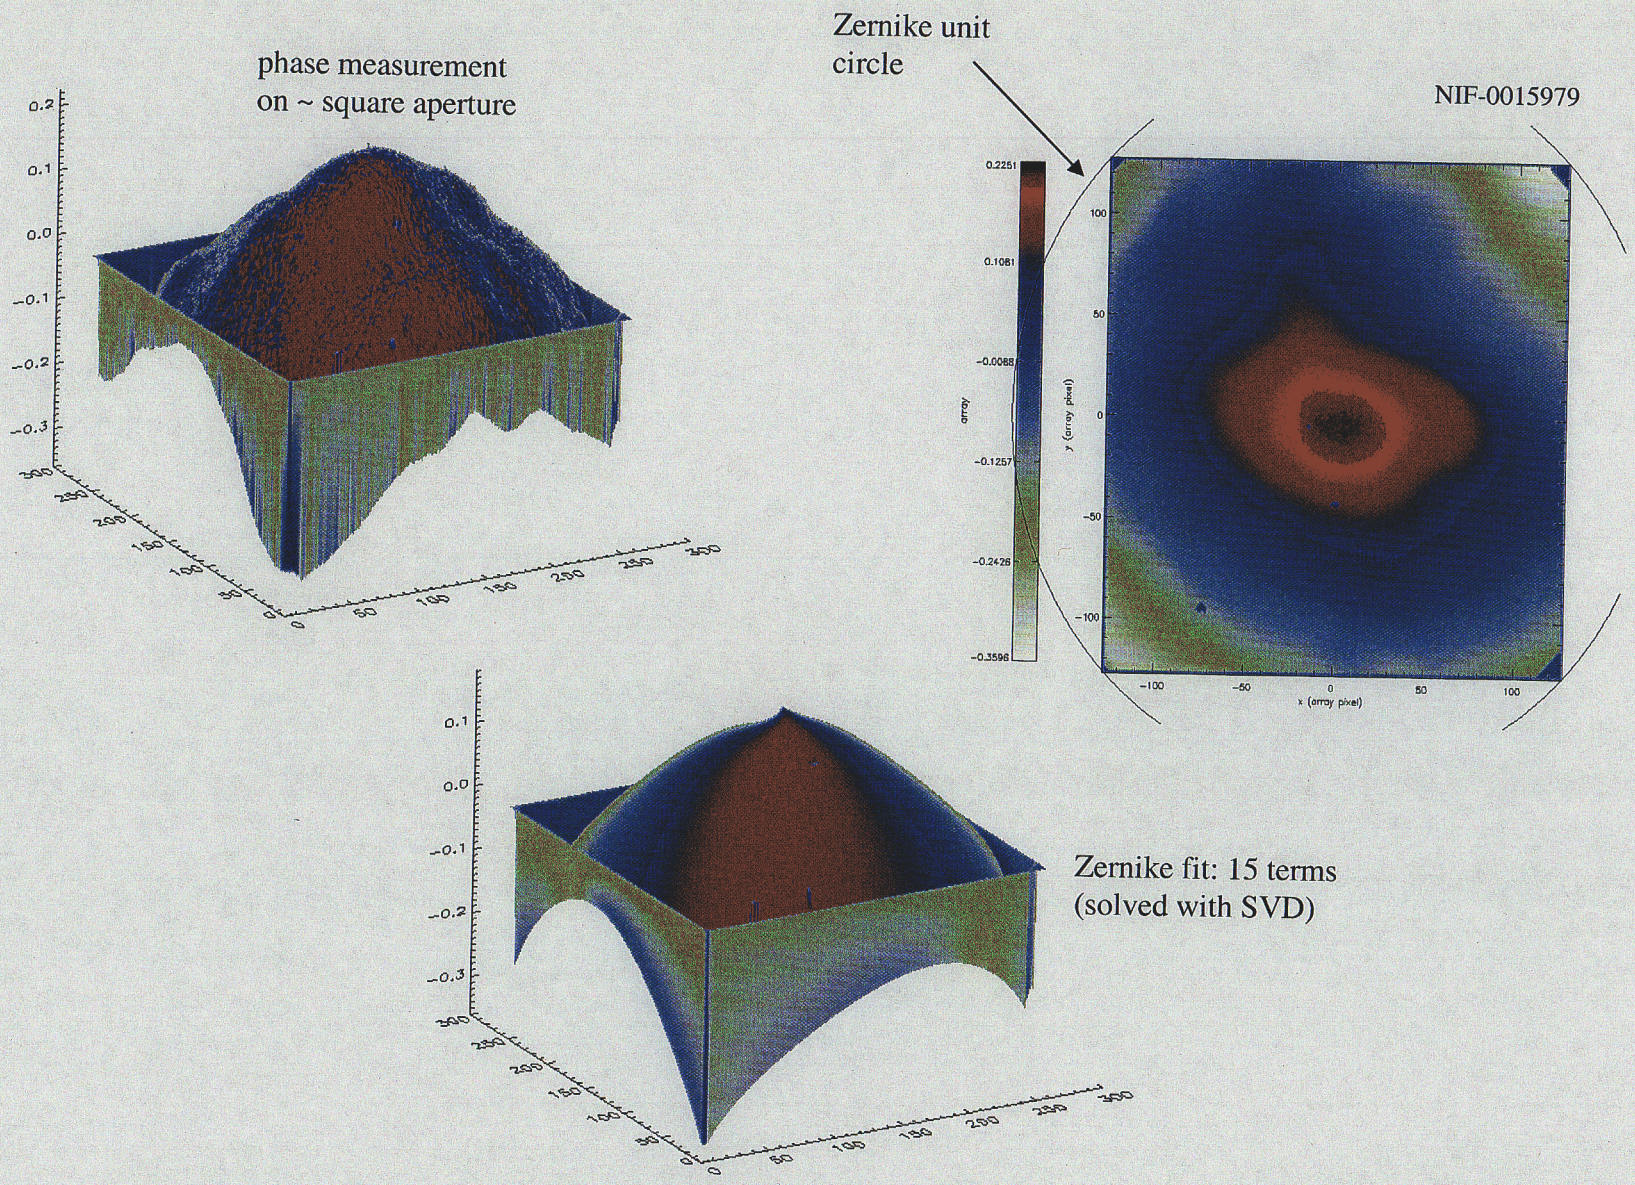

Figure 6. 15-term Zernike fit to a wavefront, using least-squares fitting.

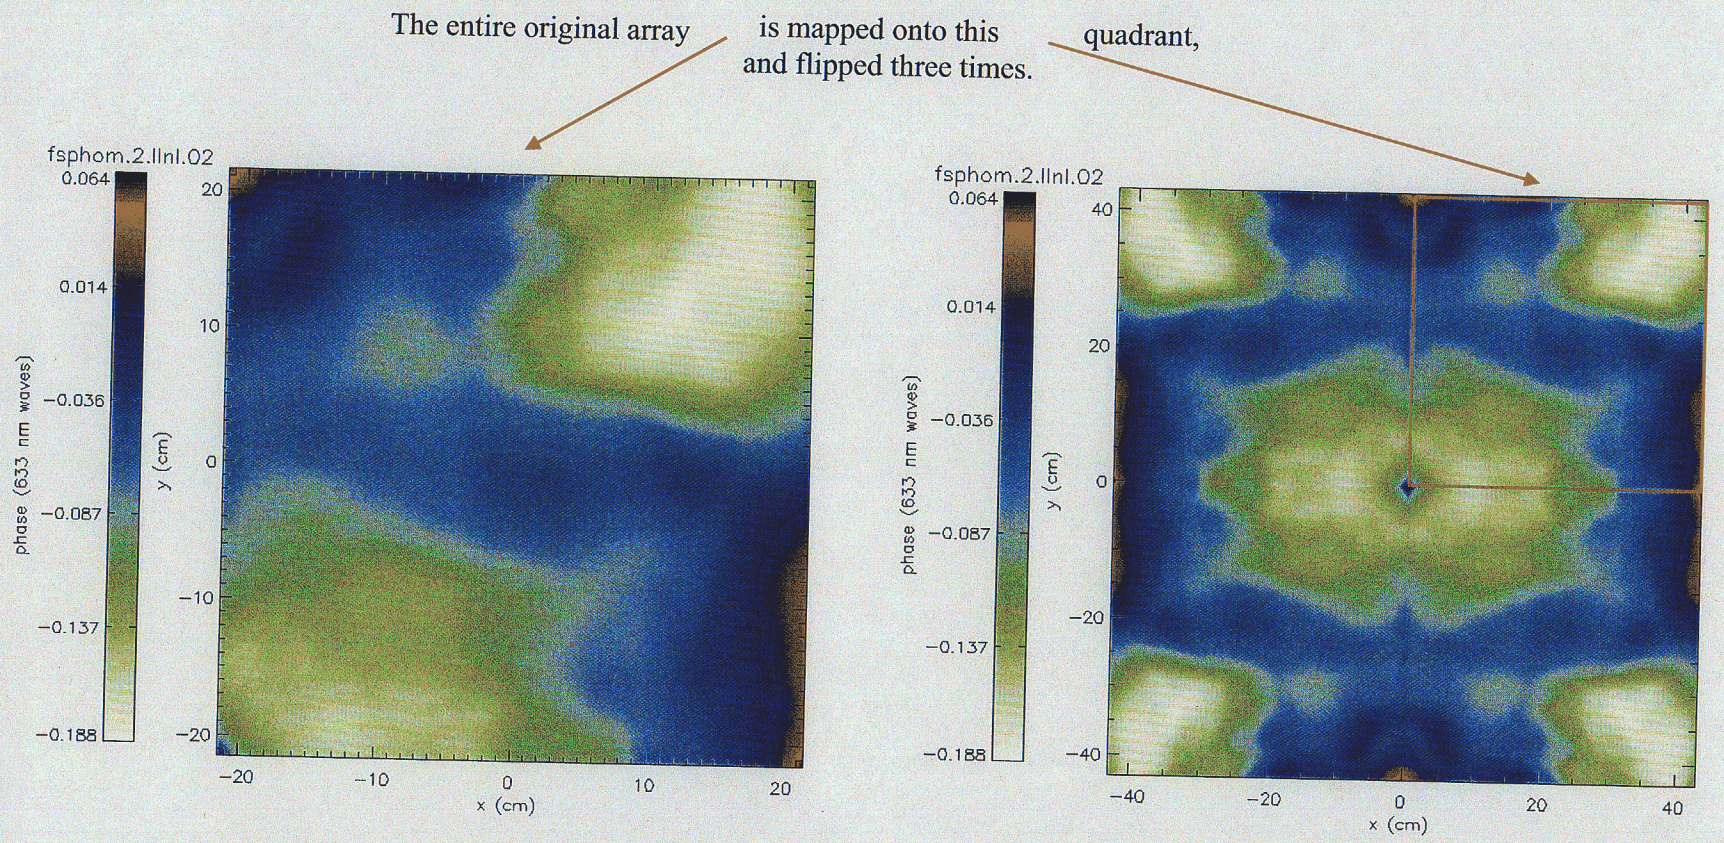

Figure 7. A "quad-flip" technique (suggested by Mark Henesian) is used in performing band-pass filtering of the phase. This approach lessens Gibb's noise from edge discontinuities. The original phase is flipped three times to make an edge-value-continuous array of 4x larger size, which is then Fourier transformed. After filtering, the original quadrant is extracted.

example 1:
KDP crystal TWF

Hanning window:
PSD1=4.0 nm

quad-flip:
PSD1=4.4 nm

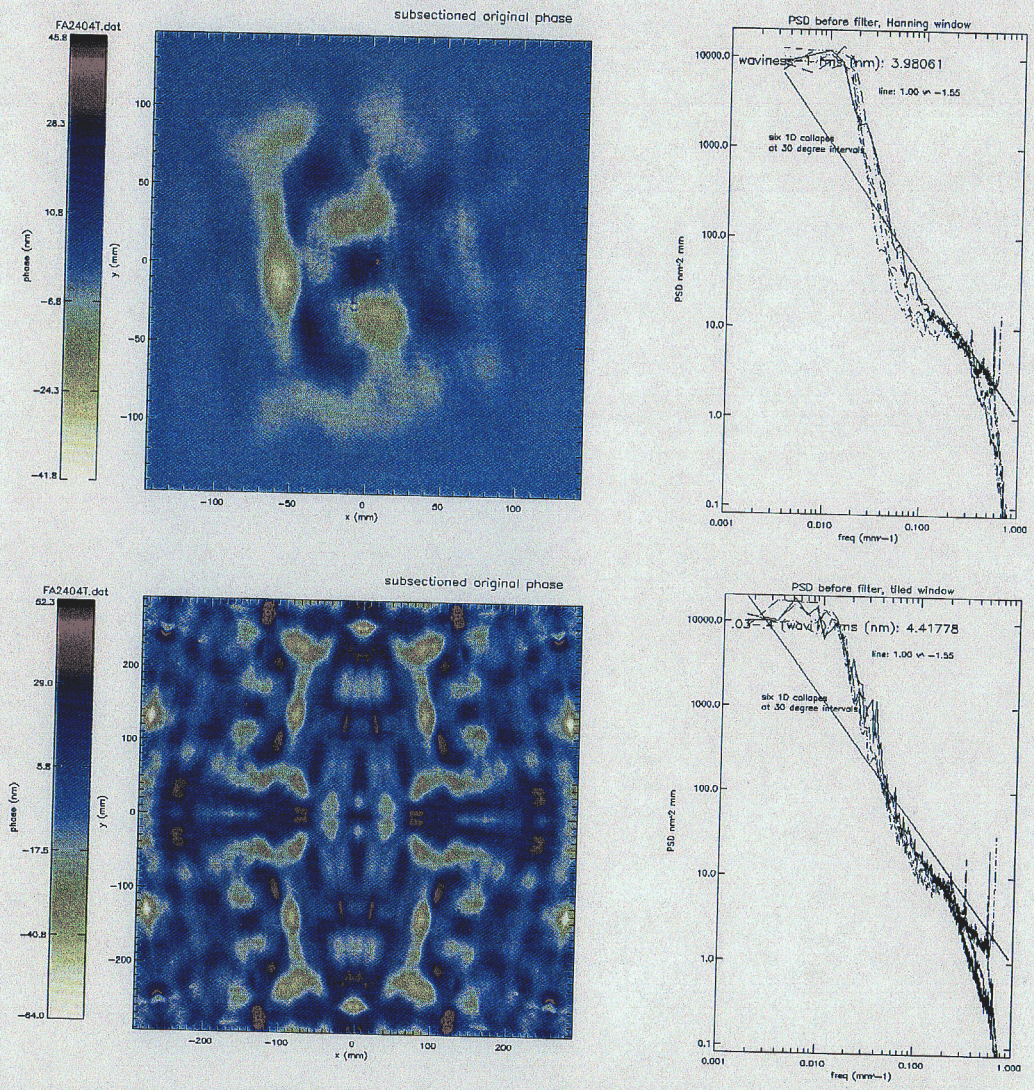

Figure 8. Comparison between the PSD1 determined using a quad-flip technique and using a Hanning window for a more homogeneous phase array. The relative weighting of the array center by the Hanning window has little impact.

example 2:
laser glass PHOM

Hanning window:
PSD1=3.0 nm

quad-flip:
PSD1=5.8 nm

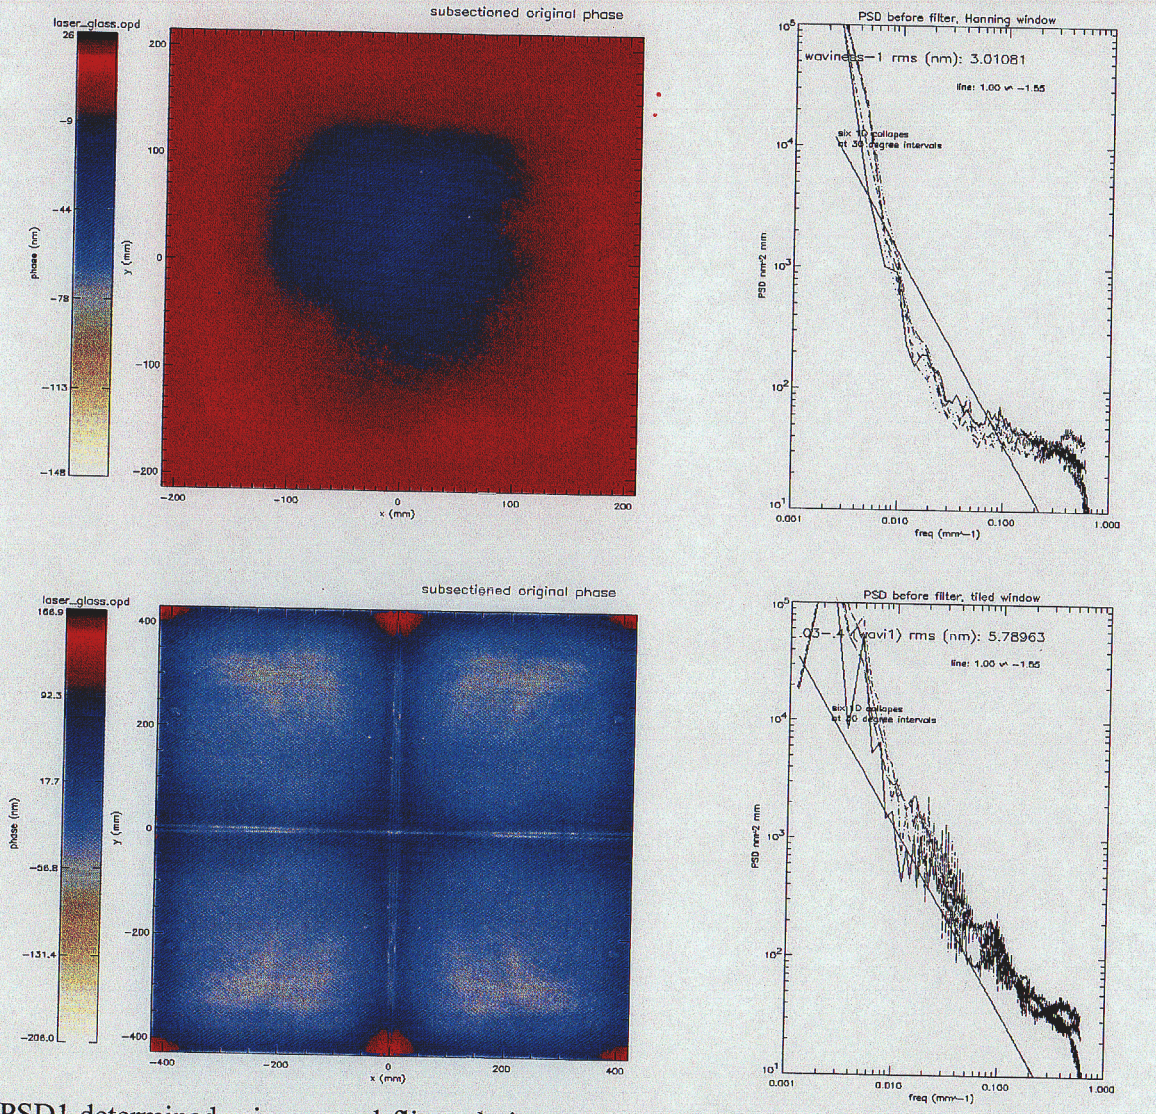

Figure 9. Comparison between the PSD1 determined using a quad-flip technique and using a Hanning window for a less homogeneous phase array. The relative weighting of the array center by the Hanning window has significant impact.

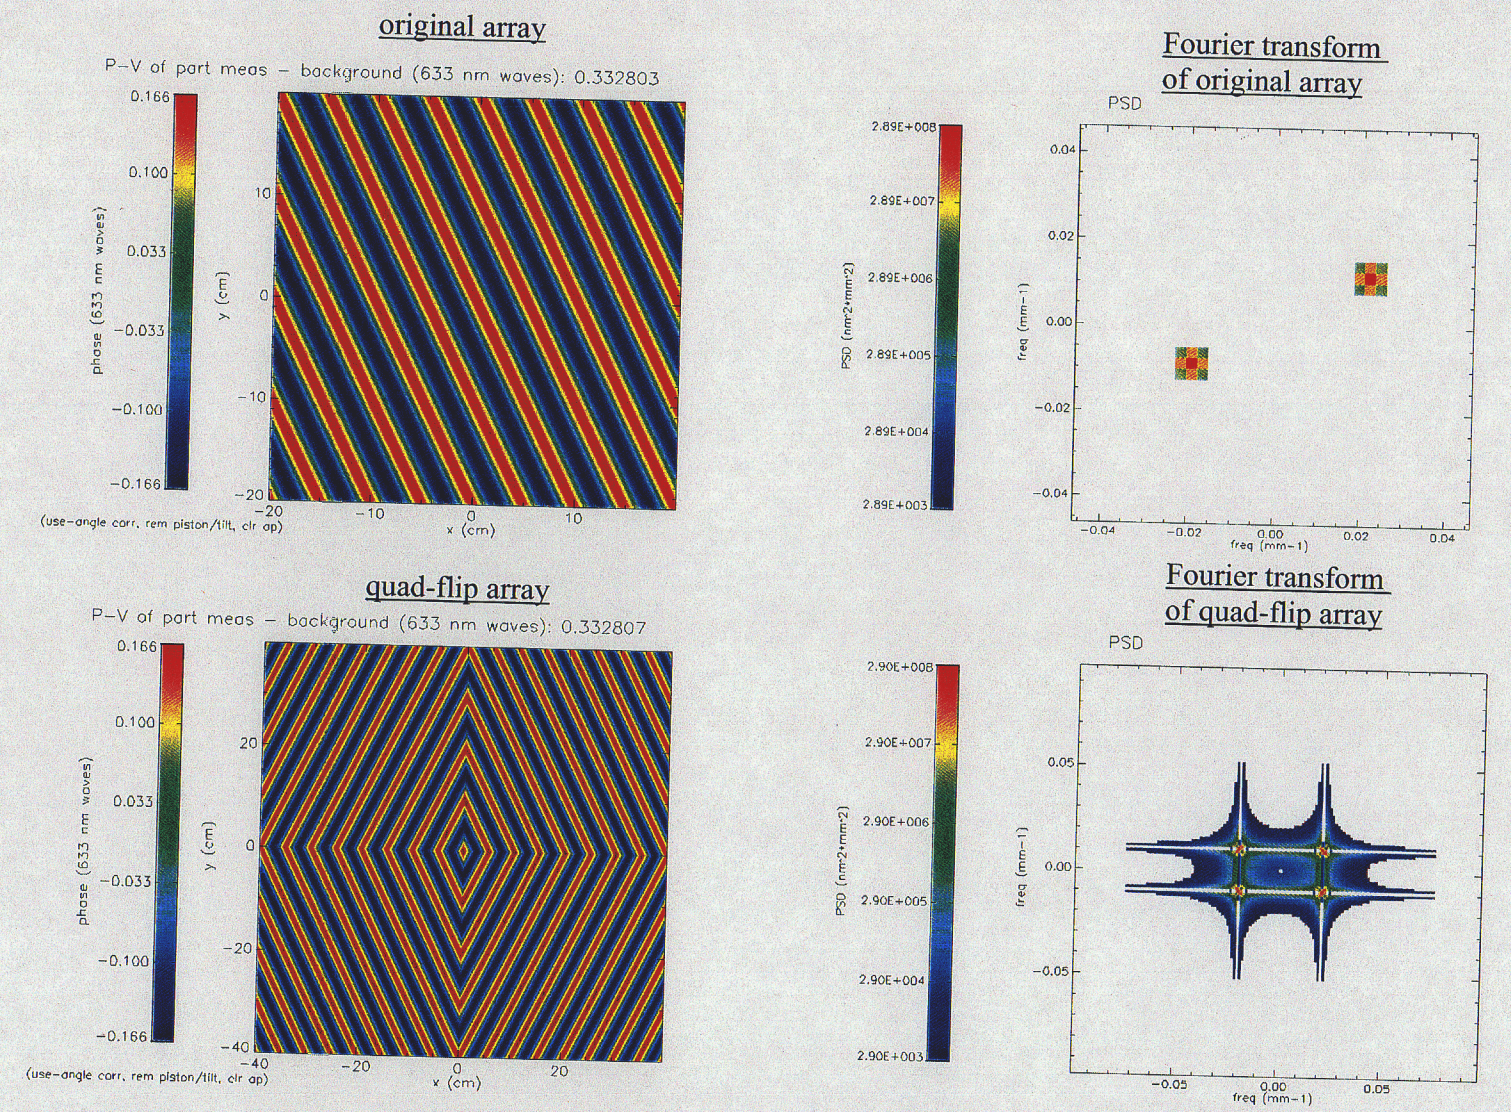

Figure 10. The quad-flip technique can place spectral information in the wrong bins.

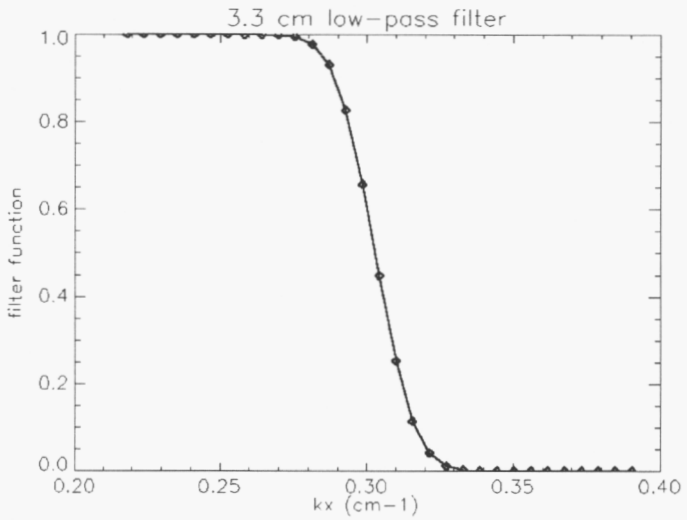

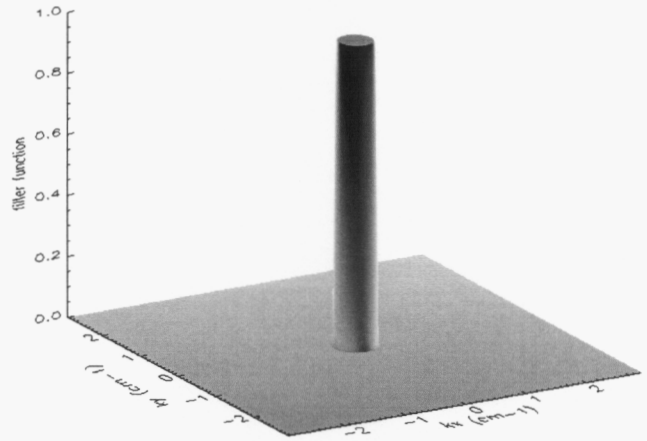

Figure 11. The round spectral filter has an error function shaped roll-off, with the 0-100% width being about 20% of the filter radius.

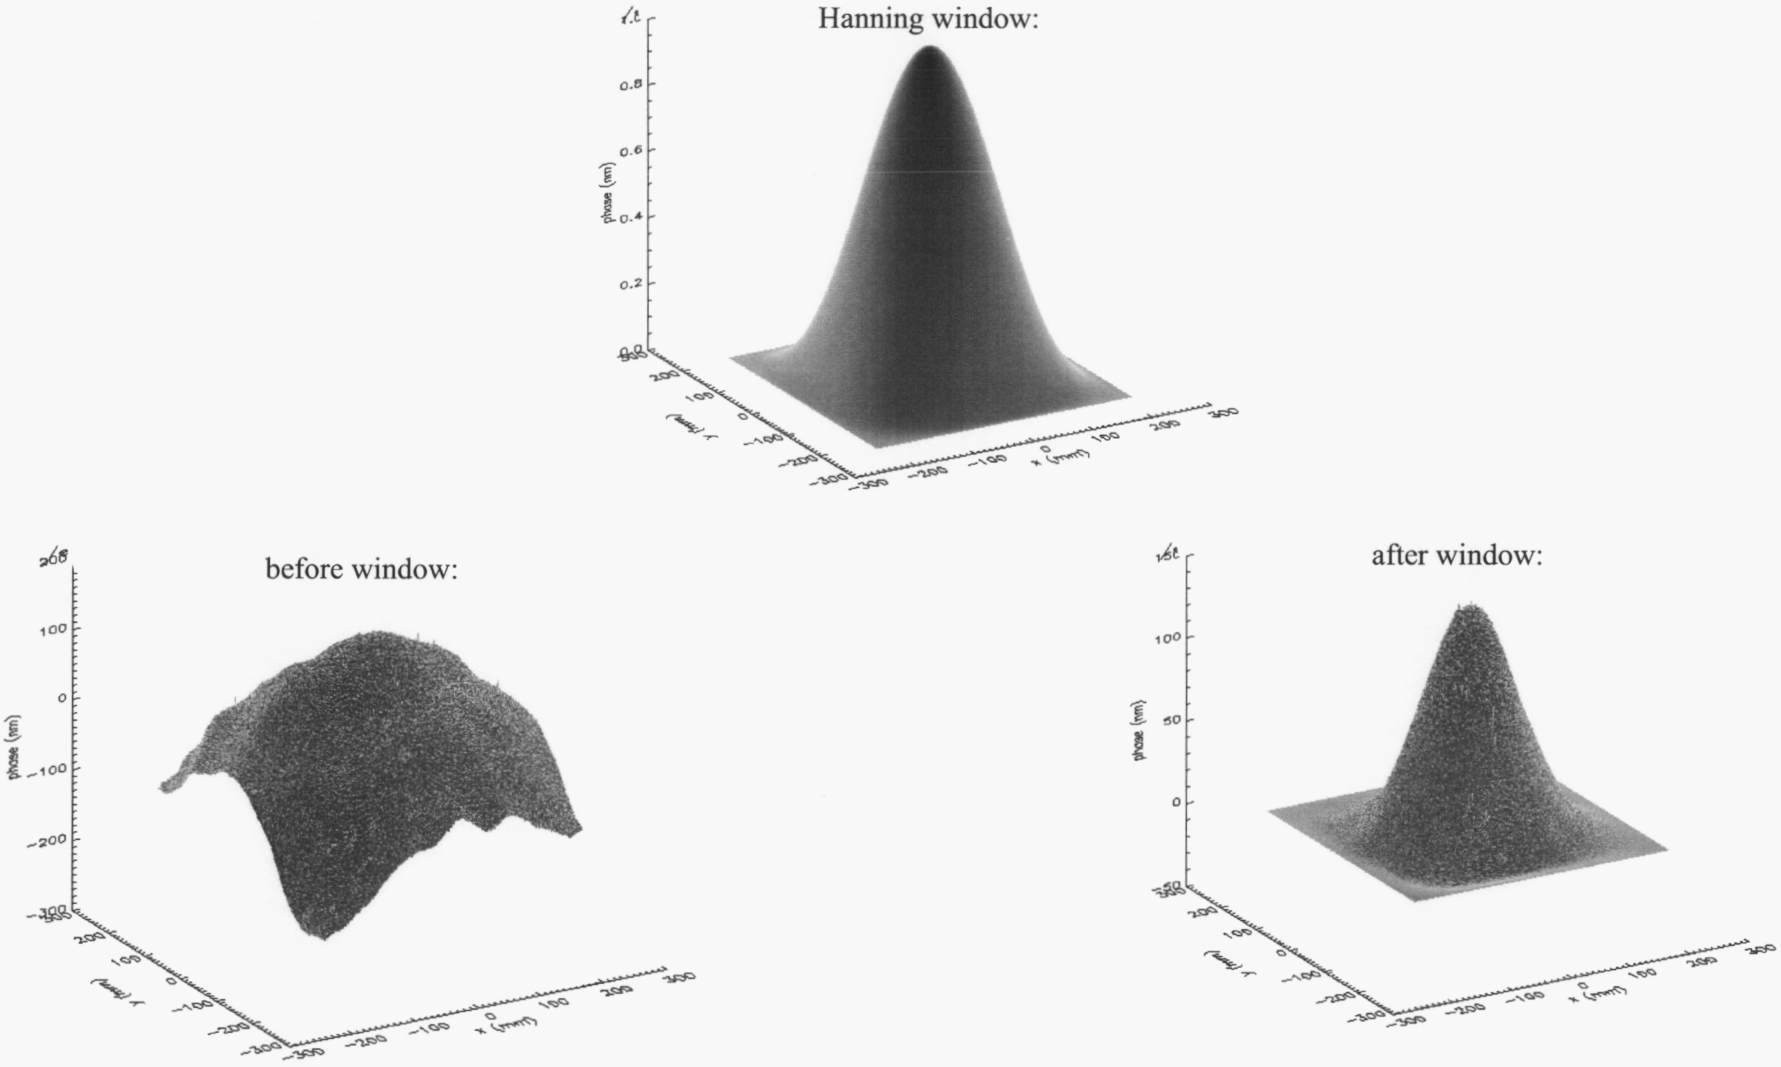

Figure 12. The Hanning window serves to remove edge discontinuities before performing a Fourier transform. It is used in these codes when calculating the 2D PSD, rms^2 running integral, and the 1D collapse of the PSD.

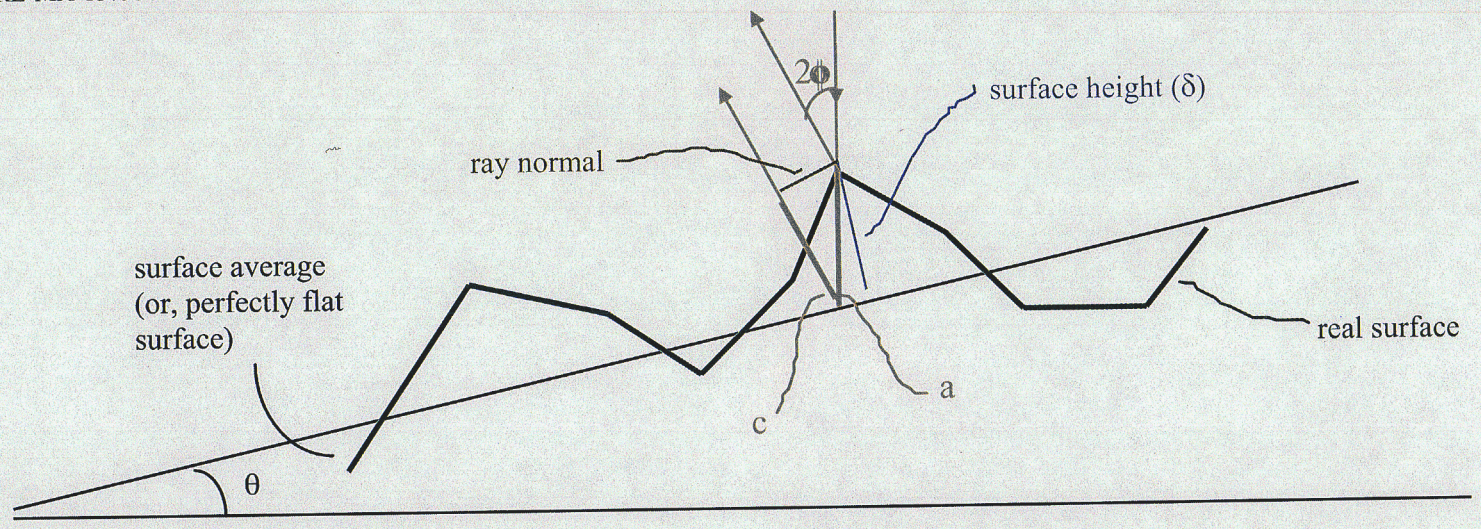

With the surface at angle θ_1, the measured change in phase between ray1 and ray2 (the difference between the real surface and a perfectly flat surface) is:

$$\phi_1=a+c=a+a\cos 2\theta_1=(\delta/\cos\theta_1)(1+\cos 2\theta_1)=\delta 2\cos\theta_1,$$

or

$$\delta=\phi_1/(2\cos\theta_1).$$

At another angle, ϕ_2, the change in phase is

$$\phi_2=2\delta\cos\theta_2=\phi_1\cos\theta_2/\cos\theta_1.$$

example:
phase measured at Brewster's angle (θ_1 = 56.31, θ_2=0): ϕ_B
inferred phase at normal incidence: 1.8 ϕ_B.

Figure 13. The phase amplitude is corrected by the difference between measurement angle and specification angle for the part.

(a)

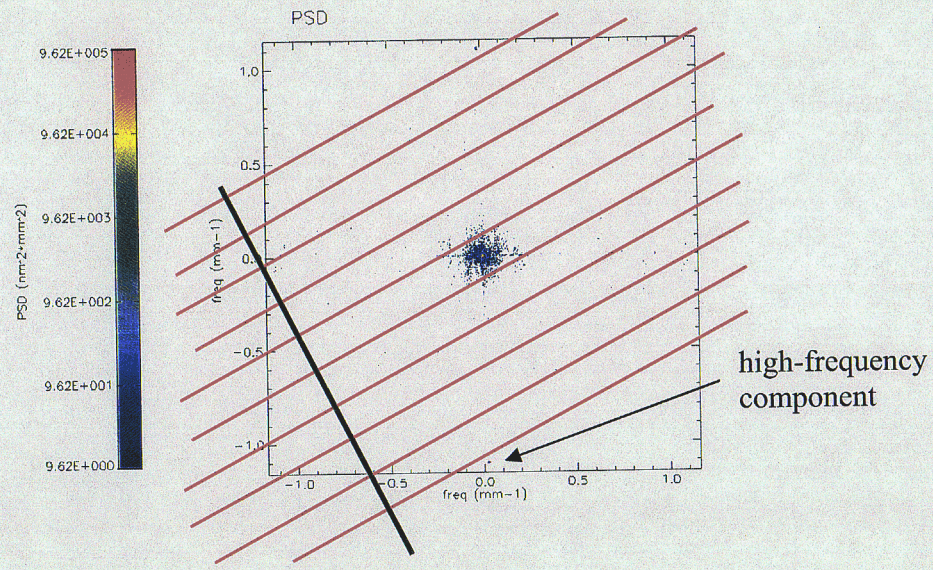

(b)

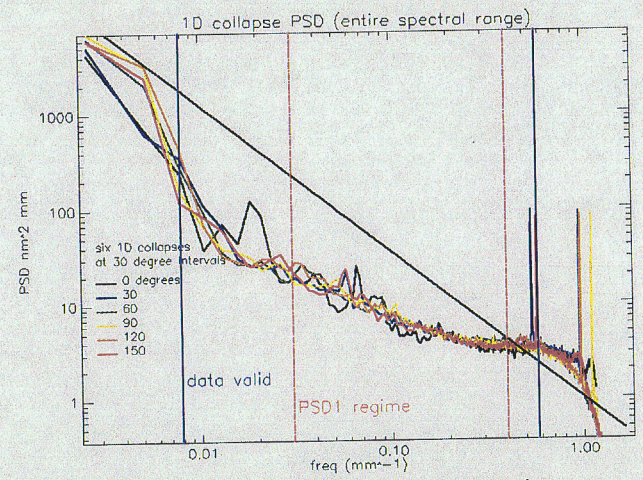

(c)

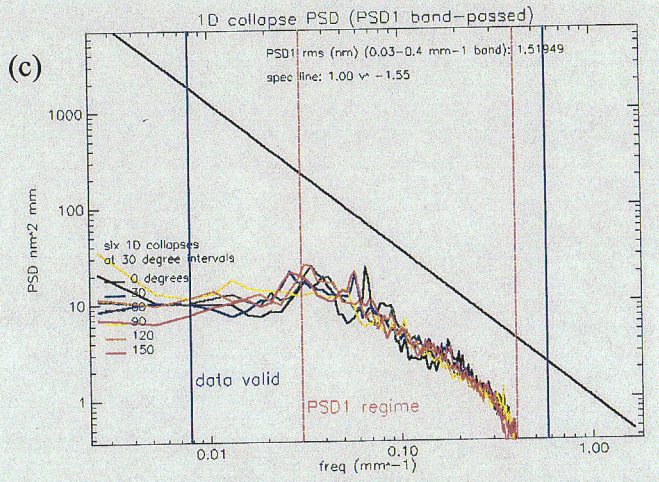

Each 1D collapse line (e.g. green line) is calulcated as a sum along the orthogonal direction (e.g. red lines) :

$$PSD1D(\nu_x) = \int PSD2D(\nu_x, \nu_y) d\nu_y$$

where ν_x and ν_y are arbitrary orthogonal axes.

Figure 14. PSD information is shown both as a 2D PSD (a), and as 1D collapse lineouts, with (c) and without (b) band-pass filtering (2.5-33 mm) before performing the 1D collapse.

(a) before 33 mm low-pass filter

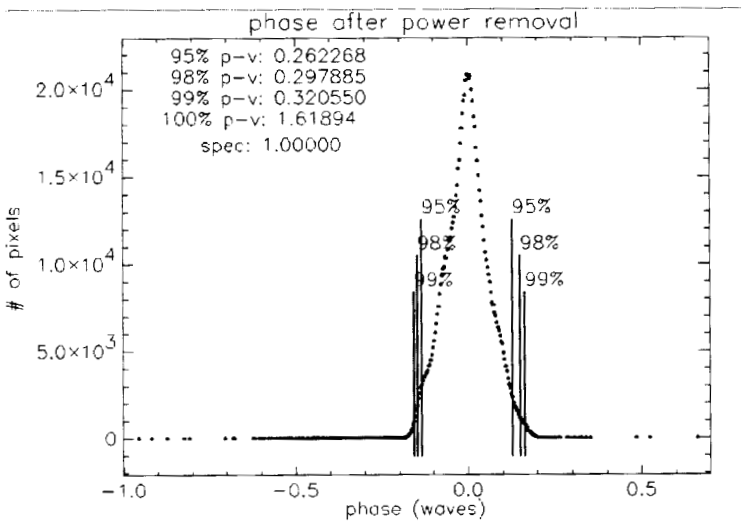

(b) after 33 mm low-pass filter

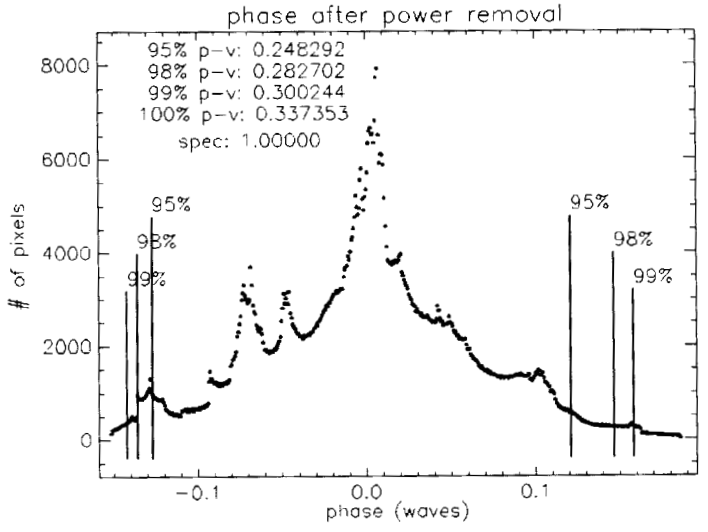

Figure 15. Histograms of a surface measurement, showing how spurious pixels can significantly effect the peak-to-valley, before (a) and after (b) low-pass (33 mm) filtering. For specifications on unfiltered phase we used the 99% value. For specifications on filtered phase we use the 100% value.

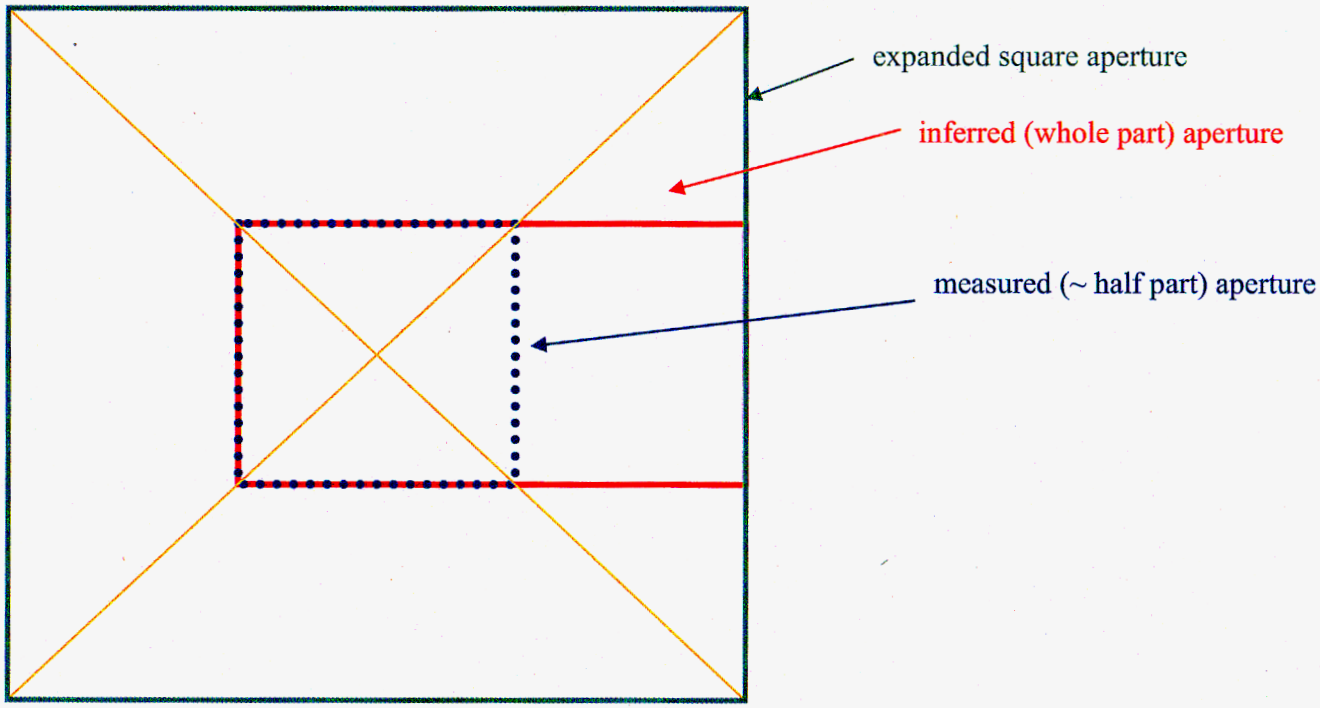

Figure 16. For PHOM measurements, astigmatism is inferred from the measured aperture to the part full aperture (see text).

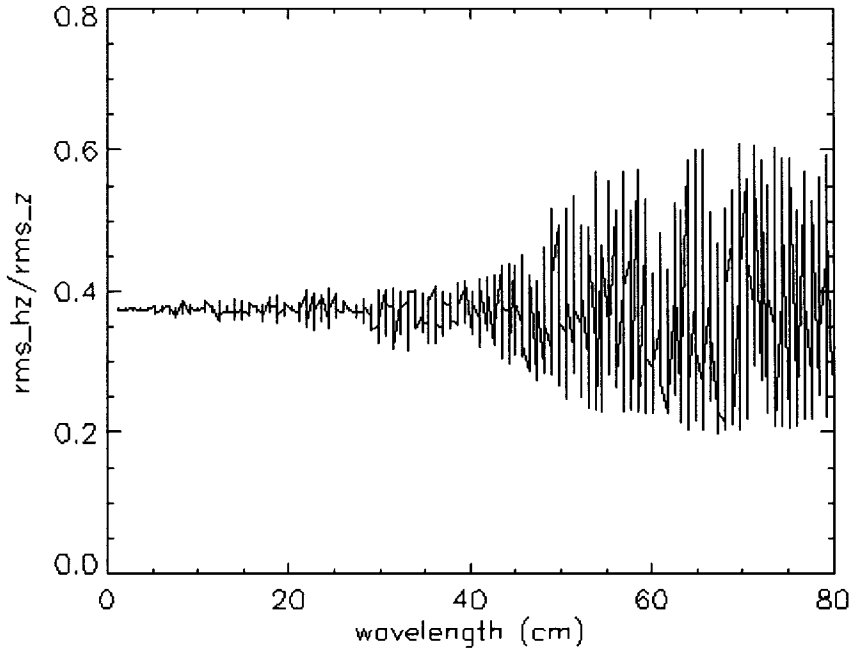

Figure 17. The ratio of the rms of a single-frequency surface with a Hanning window applied to that without a Hanning window on a 40 x 40 cm aperture. The spread at any given wavelength represents 10 randomly chosen phase shifts in x and y. The ratio is seen to approach 3/8 for short wavelengths.

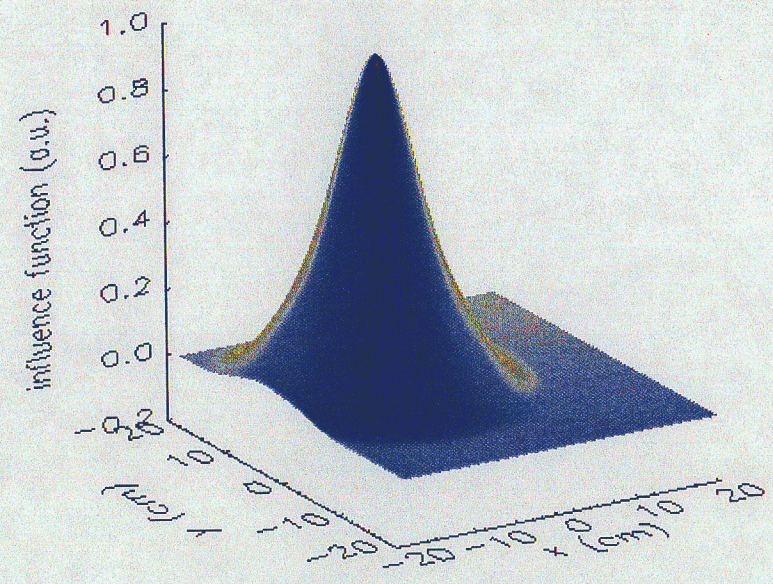

Figure 18. Example of an influence function from one adaptive optic actuator used in the influence function removal analysis.

Appendix A: GMPFIN and FSFIN Printouts

GMPFIN is designed for inspection of laser glass, polarizers, and mirrors at the finishing (i.e. polishing) process stage. FSFIN is designed for inspection of fused silica optics (lenses and windows), also at the finishing process stage.

Figures A-1 and A-2 show an example of the 2-page printout from GMPFIN. FSFIN has identical output. The example analysis here was done for a non-normal incidence RWF measurement (LM3). All optics analyzed with GMPFIN, and most analyzed with FSFIN, have two types of specifications: "nearfield" phase (P-V and rms gradient), and "farfield" (or Fourier-transformed) phase (PSD1 and Not-to-Exceed line). Page 1 of the printout covers the former and page 2 the latter.

Text information is duplicated on both pages giving the title of the part and background (if present) files. (In the attached figures the file titles have been deleted.) Also reported are measurement and analysis dates, analysis type (i.e. "process control" or "pass/fail"), measurement and analysis comments (if any), and part type (e.g. "mirror, LM3, full aperture RWF"). Note that measurement date information for Zygo files is reported as Greenwich Mean Time (GMT) because of the way the date information is stored in the file.

For clarification, the definition of one wave (633 nm) is listed. The array size in the part file and in the analysis aperture are reported for reference. Using the cosine of the measurement angle, the height and width of the clear (of "effective") aperture, as listed on the drawing and as used in the analysis, are reported. Information is also listed on the expected versus actual values for the "wedge" factor (or "interferometer scale factor") in the background and part files to help assure that the measurement represents, for example, a RWF, as opposed to a surface contour. A **Warning** flag is printed if there is a mis-match. The wedge factor in the file is used, regardless, *so this warning flag is the only indication that the file may be reporting, for example, surface contour, when it should have been RWF, and that the results are, correspondingly, inaccurate.*

Results text include the name of each specification type for this optic, the specification value as given on the drawing, the results of the code analysis, and a "pass / fail" flag. In addition, the second page lists two quantities which are not specifications, but which may be useful in later data analysis: PSD1 in the worst sub-aperture (~100 x 100 mm) of the part (see description of graphics H and I), and the ratio of the analysis area to the specification clear aperture area for that part.

The codes print 12 graphics. These are described as follows (see Figures A-1 and A-2):

A) Part measurement phase. This is the phase as retrieved from the part measurement file, converted into 633 nm waves, with piston and tilt removed. The entire camera array is shown, so there may be white space (i.e. "invalid" pixels) around the image. Invalid pixels within the image are also shown in white. The peak-to-valley over the image is

reported on the top of the graphic. Note that this may be a larger number than the difference between the max and min of the color bar scale. This is because extremes of the color bar are set so that 0.25% of the pixels have values higher than the maximum, and 0.25% have values lower than the minimum. This compression yields a more useful image of the data. Noise in the data will often cause a few pixels in the raw interferogram to be abnormally high or low. If these values are used to set the color bar range then the bulk of the image is bleached out. This compression is used in all image representations.

The (0,0) point in the x and y scales is set at the center of the camera.

Two dashed boxes appear also in this example, which relate to the clear aperture to be used in later analysis. The heavier dashed box is the code's attempt to find the clear aperture from the edges of the data set. This will often be wrong, so the user has the ability to manually draw a new clear aperture, shown as the light dashed box. A person viewing the output can tell if this has been done because both boxes will be shown, and the text "Click and drag ..." is superimposed on the image. To aid in this aperturing, information on with width / height of the aperture selected (the light dashed line in this example) and the clear aperture as specified on the drawing are printed in the lower left-hand corner of the graphic. The specification values have been adjusted by the measurement angle. The ratio between the areas represented by these dimensions is reported in the text output at the bottom of the second page, as indicated above. (Note that invalid pixels in the clear aperture are not included in this ratio.) For most parts, the "clear ap spec" and "clear ap used" should have the same values. The large mirrors are an exception because they cannot completely fit into the interferometer field of view.

B) Background file. This graphic shows the background file with piston and tilt removed. If no background is selected (e.g. for normal-incidence RWF measurements) the image is white.

C) Analysis array: measurement minus background over clear aperture. This graphic shows the part file with the background subtracted over the clear aperture. This is the array which is used for all further analysis. The horizontal (x) dimension and the amplitude have been adjusted for any difference between the measurement and use angles of the part. The 30 mm radius clear aperture corners have been added.

As described in the body of this manual, background subtraction is handled differently for different parts. In this example, which is a non-normal incidence RWF, the background file is not directly subtracted. Rather, a power fit is done to the background, and this power is then subtracted from both the measurement and background arrays. The analysis array shown in this graphic has that power removed.

D) Filtered analysis array. This shows the results of low-pass filtering (33 mm) the analysis array in graphic C. The peak-to-valley of this array is compared against specification in the text.

E) Gradient of filtered analysis array. This shows the gradient (i.e. slope) magnitude at all points in the filtered analysis array. If the background has been directly subtracted (TWF), or if there is no background file (normal incidence RWF), the rms of this array is the rms gradient compared against specification in the text at the bottom of Figure A-1. In the case of non-normal incidence measurements, such as in this example, however, we statistically subtract the background. This is done by performing a similar filtered rms gradient calculation on the background file, and reporting the square root of the difference of the squares of the analysis array rms gradient and the background file rms gradient.

F) Histogram of gradient array. This graphic shows the histogram of the array in graphic E. The rms is shown, along with 95, 98, and 99% values (for information only). Because this is a non-normal incidence measurement, the rms gradient of the background is also reported. (As indicated above, the rms gradient compared against specification is the statistical difference between these, i.e. $0.00784^2=0.00902^2-0.00447^2$).

G) RMS^2 running integral. This graphic is for information only. It is determined by patching invalid points in the analysis array, multiplying the result by a Hanning filter, and reporting the running integral of the power spectrum as the running integral of the rms^2. Said differently, it shows the rms^2 of the analysis array as a function of low-pass filter cut-off. The value at the highest frequency is the rms^2 of the unfiltered array.

This graphic can be used to *estimate* the rms phase in an arbitrary band (as opposed to the 2.5 – 33 mm band specified in PSD1) by reading the rms^2 at each of the band limits, and taking the square root of the difference. It also serves to graphically show how the rms in the unfiltered array is distributed. For example, in Figure A-2 the first point on the curve is at ~765 nm^2, out of a total of ~778 nm^2, indicating that most of the rms in the unfiltered array is contributed by frequencies longer than ~0.008 mm^{-1} ~= 12 cm.

Note that the absolute accuracy of this curve is affected by the factor of 8/3 used to correct the power loss from the Hanning filter. As shown in Figure 17 in the body of the manual, this correction is increasingly less accurate if long-wavelength components are included in the rms. *Hence, this figure should be used with caution.*

H) Band-passed analysis array. This graphic shows the analysis array after it has been band-pass filtered over the PSD1 regime (2.5-33 mm) using the quad-flip technique. The rms of the entire array (excluding invalid data points) is reported at the top. In addition, the array is sub-sectioned into ~16 sub-arrays of ~100x100 mm, and the rms in each of these sub-apertures reported. Note that both the rms for the entire array, and those for the sub-apertures, have been corrected for background subtraction. In this case of non-normal incidence reflection, this is done statistically by separately determining the PSD1 of the background file, and subtracting this (as the root of the difference of the squares) from the PSD1 of the analysis array. The rms for the entire array, and that for the worst sub-aperture, are reported in the text at the bottom of the page.

I) RMS in sub-apertures. This graphic plots the sub-aperture PSD1 information from graphic H versus sub-aperture number. The value at 0 is the full aperture. This allows quick identification of bad sub-apertures. It is intended that this information will aid in identifying poorly finished areas or local inhomogeneities. The PSD1 for the background file is reported at the top. Lines for the full-aperture PSD1 specification and twice that are shown for reference.

J) 2D PSD. This graphic shows the 2D PSD of the analysis array. It is generated by taking the Fourier transform of the analysis array, multiplied by a Hanning window, and corrected by the 8/3rd factor. For non-normal incidence RWF measurements (as in this example), no background is subtracted. This graphic is intended to aid in determination of strong spectral components in the measurement.

K) Full-spectrum 1D Collapse. This graphic shows the 2D PSD in the previous graphic, collapsed into its 1D projections, along with a not-to-exceed line. Because it covers the full spectrum, this graphic is for reference only.

L) Band-passed 1D Collapse. This graphic shows the 2D PSD collapsed into its 1D projections, after having first been band-passed filtered to cover only the PSD1 regime. The filtering is abrupt (i.e. does not use an error-function edge) because we are not doing an inverse transform. If any of the lines cross above the not-to-exceed line in the PSD1 regime, this constitutes a failure of the Do Not Exceed specification. Note again that for the non-normal incidence RWF measurement in this example that no background has been subtracted out.

Figure A-3 is a surface form analysis from FSFIN, required for the final focus lens and diagnostic beam splitter. The text and graphics output are very similar to that shown in Figures A-1 and A-2. There is no background subtraction. Because the part measurement file is assumed to represent surface contour the file values are multiplied by a factor of two to convert to RWF. The required specifications are for the peak-to-valley of the best-fit power to the measurement (saggita tolerance), and the residual peak-to-valley after removal of this power (irregularity). No filtering is done, so the irregularity is the central 99% of the histogram (top and bottom 0.5% of pixels are ignored) to reduce the effect of measurement noise.

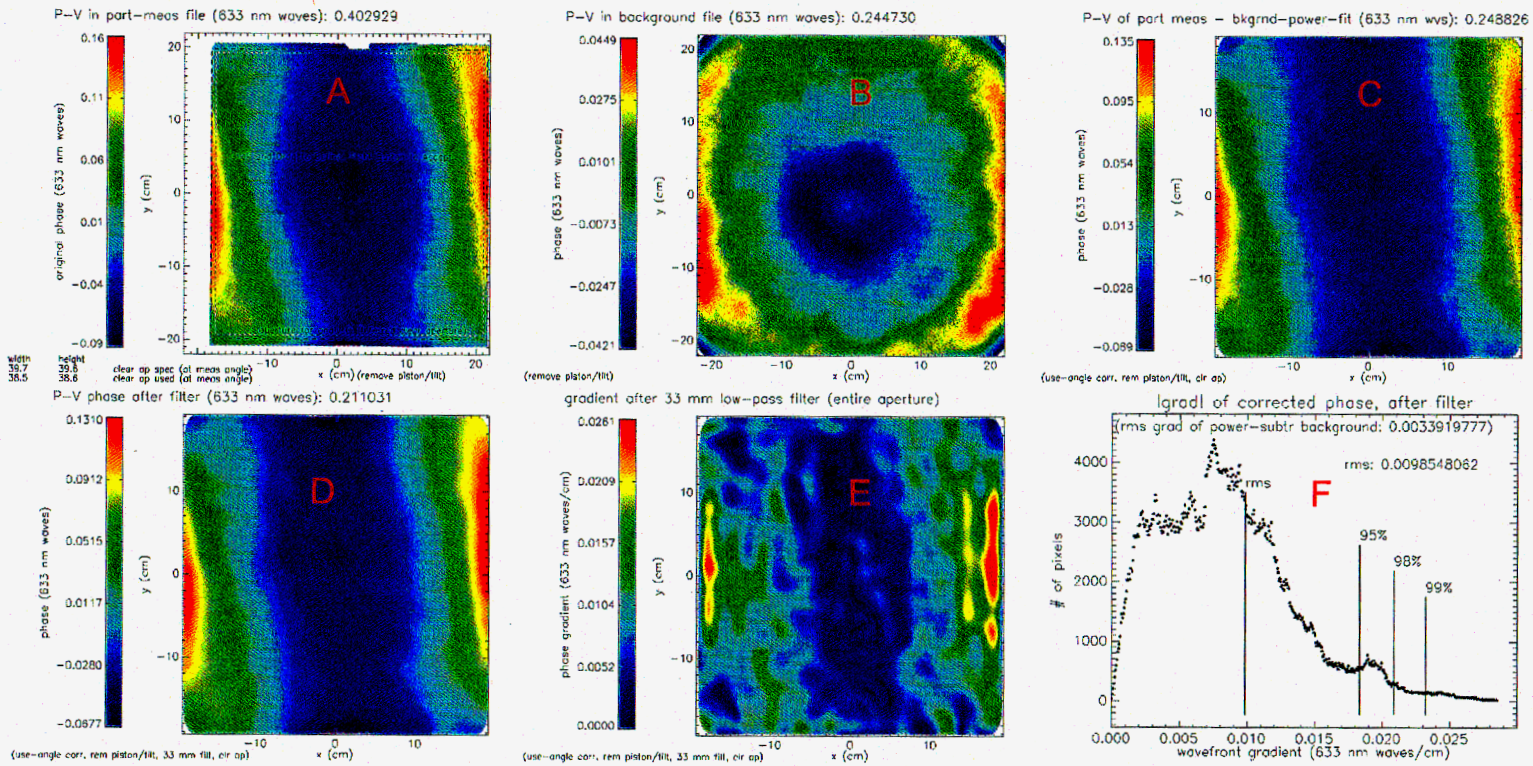

Part file:
Background file:
File part / serial #: /
Part file comments:
Background file comments: 2200 mm cavity length
Analysis comments: test

Part Measured: 1999 Nov 4, 02:44, GMT
Background Measured: 2000 Feb 11, 15:13, GMT
Analysis Performed: Thu Sep 05 10:28:35 2002
Analysis type: Final pass/fail
Part type: mirror, LM3, full aperture RWF

(one wave: 633 nm; file array size: 1024x1024; clear aperture array size: 897x 899; use / meas angle: 56.4 / 56.5; antic. / actual part file 'wedge' factor: 0.5 / 0.5; antic. / actual bkgrnd file 'wedge' factor: 0.5 / 0.5; normal-inc clear ap width spec / used: 71.9 / 69.8; normal-inc clear ap height spec / used: 39.6 / 38.6)

	Specification	Measured	Pass/Fail
Wavefront P-V (33 mm low-pass filtered RWF), waves	0.333333	0.2110306	PASS
RMS Gradient (33 mm low-pass filtered RWF gradient), waves/cm	0.0111111	0.0092526586	PASS

GMPFIN, rev1.7

page 1 of 2

Figure A-1. GMPFIN output, 1st page.

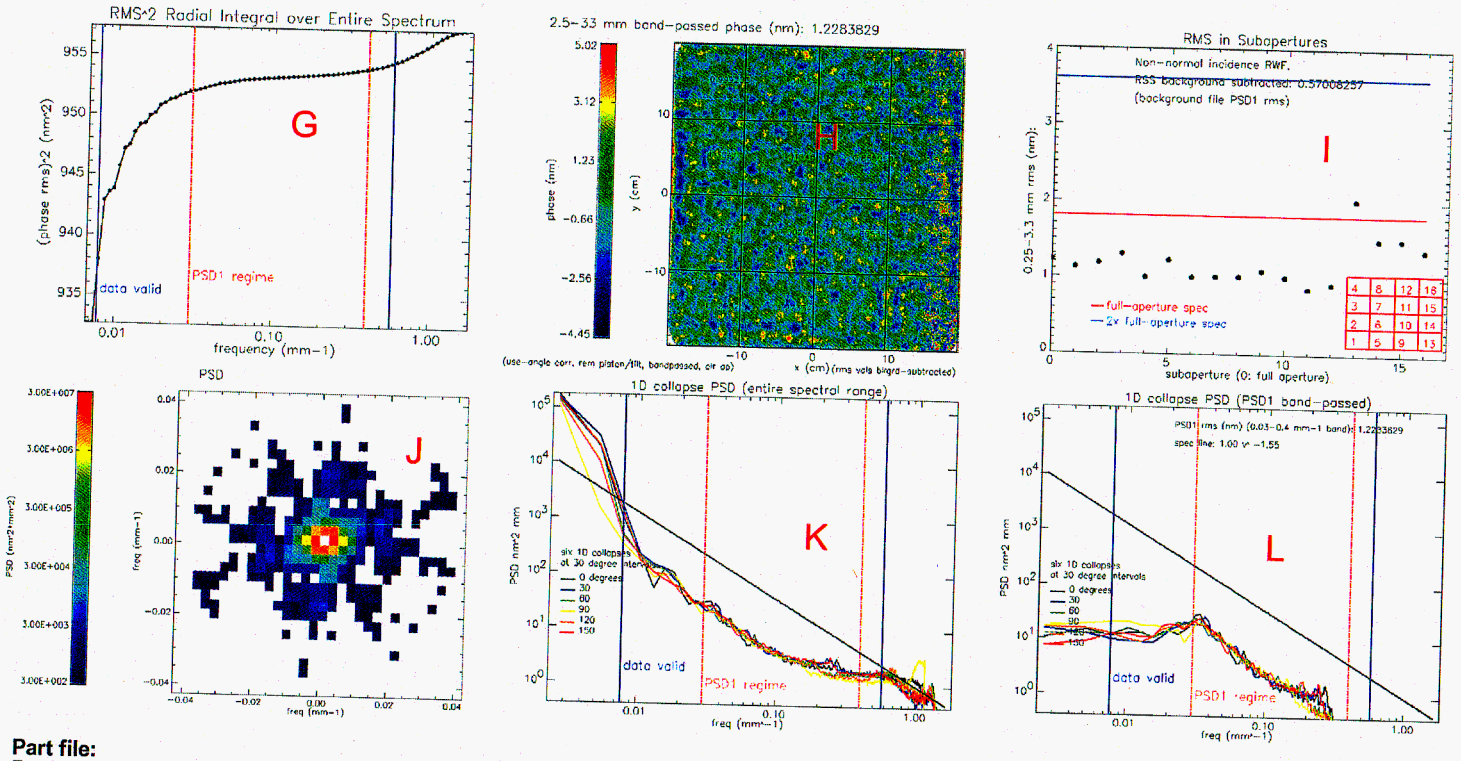

Part file:
Background file:
File part / serial #: /
Part file comments:
Background file comments: 2200 mm cavity length
Analysis comments: test
(one wave: 633 nm; file array size: 1024x1024; clear aperture array size: 897x 899; use / meas angle: 56.4 / 56.5; antic. / actual part file 'wedge' factor: 0.5 / 0.5; antic. / actual bkgrnd file 'wedge' factor: 0.5 / 0.5; normal-inc clear ap width spec / used: 71.9 / 69.8; normal-inc clear ap height spec / used: 39.6 / 38.6)

Part Measured: 1999 Nov 4, 02:44, GMT
Background Measured: 2000 Feb 11, 15:13, GMT
Analysis Performed: Thu Sep 05 10:28:35 2002
Analysis type: Final pass/fail
Part type: mirror, LM3, full aperture RWF

	Specification	Measured	Pass/Fail
PSD1 over effective aperture (RMS RWF in 2.5 - 33 mm band), nm	1.8	1.2283829	PASS
PSD1 Do Not Exceed (1D Collapse RWF in 2.5 - 33 mm band)	1v^-1.55	n/a	PASS
PSD1 in worst subaperture (RMS RWF in 2.5 - 33 mm band), nm	n/a	2.01446	n/a
Ratio of analysis area to spec effective aperture area, %:	n/a	94.5659	n/a

GMPFIN, rev1.7

page 2 of 2

Figure A-2. GMPFIN output, 2nd page.

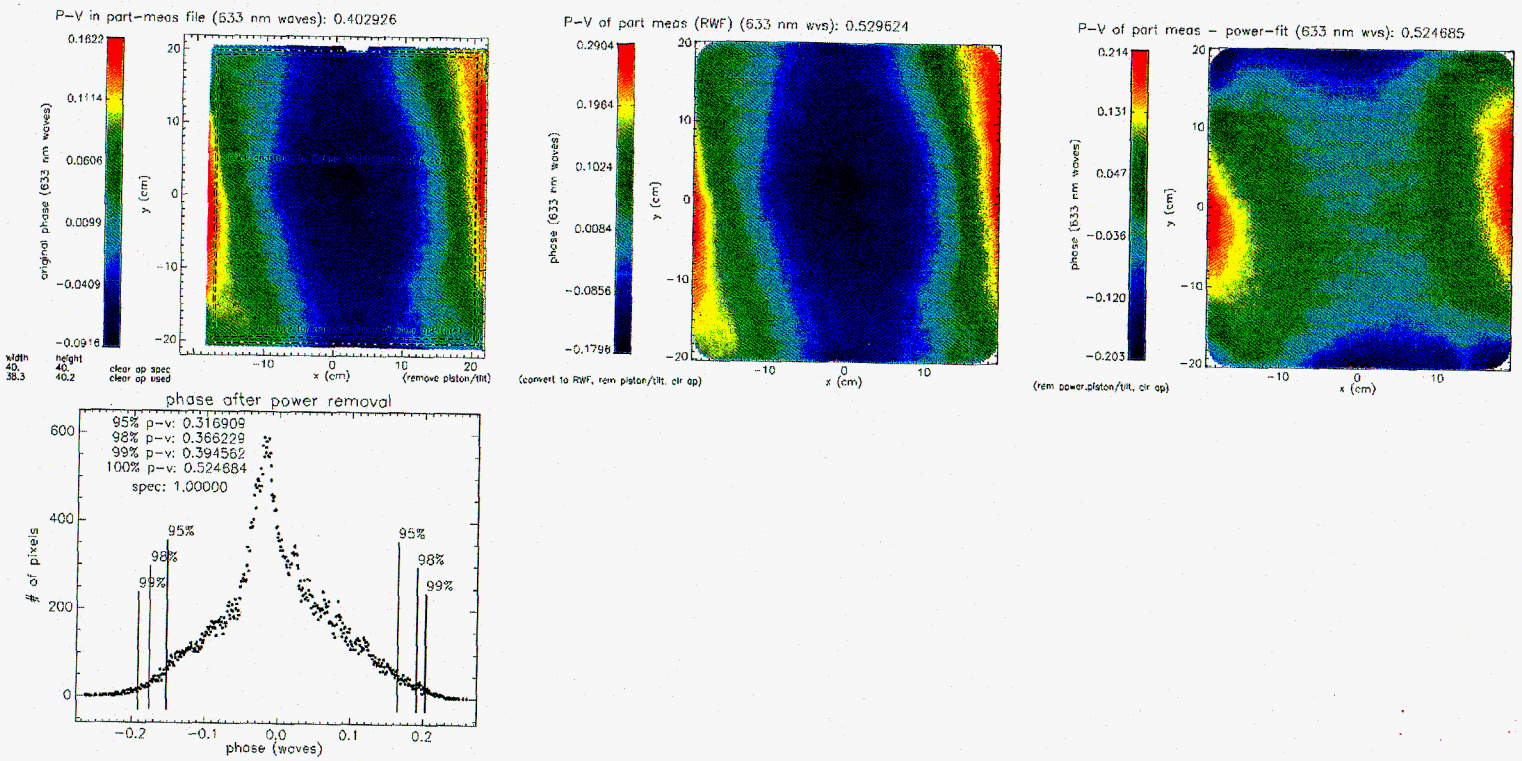

Part file:
Background file:
File part / serial #: /
Part file comments:
Background file comments:
Analysis comments: test
(one wave: 633 nm; file array size: 1024x1024; clear aperture array size: 256x 256;
antic. / actual part file 'wedge' factor: 0.5 / 0.5; ** WARNING ** antic. / actual bkgrnd file 'wedge' factor: 0.5 / 0.;
normal-inc clear ap width spec / used: 40. / 38.3; normal-inc clear ap height spec / used: 40. / 40.2)

Part Measured: 1999 Nov 4, 02:44, GMT
Background Measured:
Analysis Performed: Tue Feb 22 13:15:45 2000
Analysis type: Final pass/fail
Part type: Focus lens SURFACE, full aperture surf contour

	Specification	Measured	Pass/Fail
Saggita tolerance (P-V of best-fit-power to RWF), in waves:	5.	0.426227	PASS
Irregularity (99% P-V of RWF after best-fit-power removal), in waves:	1.	0.394562	PASS
Ratio of analysis area to spec clear aperture area (%):	n/a	96.2288	n/a

FSFIN, rev1.5

page 1 of 1

Figure A-3. FSFIN output for surface form analysis.

Appendix B: MPCOAT Printout

MPCOAT is designed for inspection of polarizers and mirrors at the coating process stage.

Figures B-1 shows an example of the 1-page printout from MPCOAT. The example analysis here was done for a non-normal incidence RWF measurement (LM3). All optics analyzed with MPCOAT have two types of specifications: "nearfield" phase (P-V and rms gradient), and "farfield" (or Fourier-transformed) phase (PSD1 and Not-to-Exceed line). The nearfield specifications are discussed in this appendix. Because it is currently not clear how the farfield specifications will be monitored (due to interferometer limitations) discussion of relevant farfield output will be postponed until those issues are resolved.

Text information includes the name of the measurement file (which has been deleted in the figure). Also reported are measurement and analysis dates, analysis type (i.e. "process control" or "pass/fail"), measurement and analysis comments (if any), and part type (e.g. "mirror, LM3, full aperture RWF"). Note that measurement date information for Zygo files is reported as Greenwich Mean Time (GMT) because of the way the date information is stored in the file.

For clarification, the definition of one wave (1064 nm) is listed. The array size in the part file and in the analysis aperture are reported for reference. Using the cosine of the measurement angle, the height and width of the clear (of "effective") aperture, as listed on the drawing and as used in the analysis, are reported. Information is also listed on the expected versus actual values for the "wedge" factor (or "interferometer scale factor") in the measurement file to help assure that the measurement represents a RWF, as opposed to a surface contour. A **Warning** flag is printed if there is a mis-match. The wedge factor in the file is used, regardless, *so this warning flag is the only indication that the file may be reporting, for example, surface contour, when it should have been RWF, and that the results are, correspondingly, inaccurate.*

Results text include the name of each specification type for this optic, the specification value as given on the drawing, the results of the code analysis, and a "pass / fail" flag. In addition, the ratio of the analysis area to the specification clear aperture area for that part is included for reference.

The codes print 5 graphics. These are described as follows (see Figure B-1):

A) Part measurement phase. This is the phase as retrieved from the part measurement file, converted into 1064 nm waves, with piston and tilt removed. The entire camera array is shown, so there may be white space (i.e. "invalid" pixels) around the image. Invalid pixels within the image are also shown in white. The peak-to-valley over the image is reported on the top of the graphic. Note that this may be a larger number than the difference between the max and min of the color bar scale. This is because extremes of the color bar are set so that 0.25% of the pixels have values higher than the maximum, and 0.25% have values lower than the minimum. This compression yields a more useful

image of the data. Noise in the data will often cause a few pixels in the raw interferogram to be abnormally high or low. If these values are used to set the color bar range then the bulk of the image is bleached out. This compression is used in all image representations.

The (0,0) point in the x and y scales is set at the center of the camera.

Two dashed boxes appear also in this example, which relate to the clear aperture to be used in later analysis. The heavier dashed box is the code's attempt to find the clear aperture from the edges of the data set. This will often be wrong, so the user has the ability to manually draw a new clear aperture, shown as the light dashed box. A person viewing the output can tell if this has been done because both boxes will be shown, and the text "Click and drag ..." is superimposed on the image. To aid in this aperturing, information on with width / height of the aperture selected (the light dashed line in this example) and the clear aperture as specified on the drawing are printed in the lower left-hand corner of the graphic. The specification values have been adjusted by the measurement angle. The ratio between the areas represented by these dimensions is reported in the text output at the bottom of the second page, as indicated above. (Note that invalid pixels in the clear aperture are not included in this ratio.) For most parts, the "clear ap spec" and "clear ap used" should have the same values. The large mirrors are an exception because they cannot completely fit into the interferometer field of view.

B) Analysis array: measurement over clear aperture. This graphic shows the part file over the clear aperture. This is the array which is used for all further analysis. The horizontal (x) dimension and the amplitude have been adjusted for any difference between the measurement and use angles of the part.

C) Gradient of filtered analysis array. This shows the gradient (i.e. slope) magnitude at all points in the filtered analysis array (graphic D). The rms of this array is compared against specification in the text.

D) Filtered analysis array. This shows the results of low-pass filtering (33 mm) the analysis array in graphic B. The peak-to-valley of this array is compared against specification in the text.

E) Histogram of gradient array. This graphic shows the histogram of the array in graphic C. The rms is shown, along with 95, 98, and 99% values (for information only).

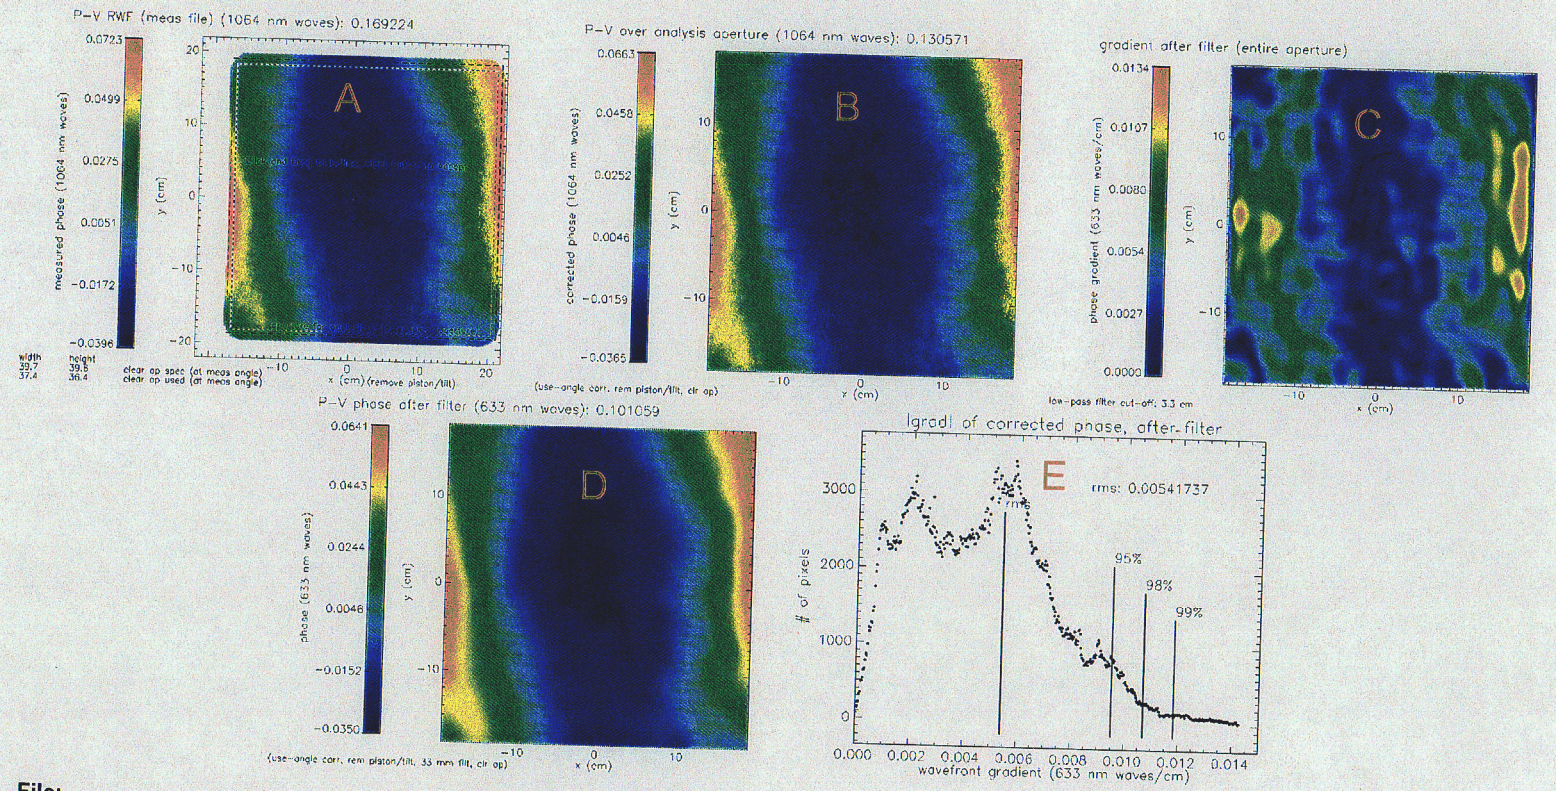

File:

Measurement type: mirror, LM3, full aperture RWF
Analysis comments: test
File part / serial #: ,
File comments:

Analysis mode: Final pass/fail
Analyzed: Tue Feb 22 13:43:34 2000
Measured: 1999 Nov 4, 02:16, GMT

(one wave: 1064 nm; file array size: 1024x1024; clear aperture array size: 871x849; use / meas angle: 56.4 / 56.5;
antic. / actual part file 'wedge' factor: 0.5 / 0.5; normal-inc clear ap width spec / used: 71.9 / 67.8; normal-inc clear ap height spec / used: 39.6 / 36.4)

	Specification	Measured	Pass/Fail
Wavefront P-V (33 mm low-pass filtered RWF), waves:	0.333300	0.101059	PASS
RMS Gradient (33 mm low-pass filtered RWF gradient), waves/cm:	0.00666667	0.00541737	PASS
Ratio of analysis area to spec effective aperture area, %:	n/a	86.8557	n/a

mpcoat, rev. 1.1

Figure B-1. MPCOAT output for nearfield specification.

Appendix C: PHOMGLASS Printouts

PHOMGLASS is designed for inspection of laser glass at the blank process stage, where we are interested in material homogeneity. It analyzes for both NIF and LMJ (CEA, French) specifications. Measurements are done normal to the optic. Because the interferometer aperture is not large enough to see the full rectangular slab, two measurements are done, one on the left side and one on the right. There is no attempt to "stitch" the measurements together. The same specifications are applied independently to both sides.

Figures C-1 and C-2 show an example of the 2-page printout from PHOMGLASS. The first page applies to NIF specifications, and the second to LMJ specifications. (The analysis can also be done for the NIF and LMJ specifications separately, resulting in only 1-page printouts.) The analysis requires a "PHOM" file and a measurement file of the front surface of the glass. The PHOM file contains an array of the material homogeneity at the camera resolution, and is derived from a combination of front surface, rear surface, empty cavity, and transmitted wavefront measurements. This combining is done by the interferometer software, and only the results are used by PHOMGLASS. The laser glass specifications mostly concern the glass homogeneity. However, because the PHOM extraction technique becomes inaccurate if the part surface finish is too poor, there is also a specification that the surfaces of the part be less than 5 waves peak-to-valley over the measurement aperture. We only check one side in the analysis, and assume the other side is similar.

Some text information is duplicated on both pages, giving the title of the PHOM and surface files, the measurement and analysis dates, analysis mode (i.e. "process control" or "pass/fail"), measurement and analysis comments (if any), assumed part thickness and mean index of refraction. Note that the part thickness and index of refraction are taken from the measurement file, if input by the operator. If not there, the code prompts the user for these values. The part thickness is used in the index variation calculation. The mean index of refraction is not used in calculations. For clarification, the definition of one wave (633 nm) is listed, as are the array sizes in the PHOM file and the surface file.

Results text include the name of each specification type for this optic, the specification value as given on the drawing, the results of the code analysis, and a "pass / fail" flag. Because the power and astigmatism specifications are for the entire slab, the measured power and astigmatism must be extrapolated from the square measured aperture to the rectangular part dimensions. The text information includes both the inferred values for the rectangular aperture, and the actual measured values over square aperture.

The codes print 12 graphics. The first six pertain especially to NIF specifications, and the second six to LMJ specifications. These are described as follows (see Figures C-1 and C-2):

A) Part measurement phase. This is the phase as retrieved from the part measurement file, converted into 633 nm waves, with piston and tilt removed. The entire camera array is shown, so there may be white space (i.e. "invalid" pixels) around the image. Invalid pixels within the image are also shown in white. The peak-to-valley over the image is reported on the top of the graphic. Note that this may be a larger number than the difference between the max and min of the color bar scale. This is because extremes of the color bar are set so that 0.25% of the pixels have values higher than the maximum, and 0.25% have values lower than the minimum. This compression yields a more useful image of the data. Noise in the data will often cause a few pixels in the raw interferogram to be abnormally high or low. If these values are used to set the color bar range then the bulk of the image is bleached out. This compression is used in all image representations.

The (0,0) point in the x and y scales is set at the center of the camera.

The dashed box is the clear aperture which will be used for further calculations. This is set to be 40 x 40 cm. The operator can change this box in size or location, in which case two dashed boxes will appear in this graphic, with the words, "Click and drag to define, right button to accept" printed in the center of the picture. The lighter dashed box will be used in further calculations. This user-defined clear aperturing should not be needed in routine analysis.

B) Analysis array. This graphic shows the PHOM file over the analysis aperture. This is the array which is used for all further analysis. Tilt, piston, power and astigmatism have been removed. The magnitude of the removed power and astigmatism are reported in the text at the bottom of the page, and compared against specification.

C) Analysis array histogram. This shows the histogram of the data in graphic B. The peak-to-valley from this histogram is compared against specification at the bottom of the page. This is actually a "99%" peak-to-valley, meaning that 0.5% of the pixels in the array have values above the "peak" and 0.5% have values below the "valley". This definition is used to minimize the effect of measurement noise on this parameter. In this example, this noise is seen as the few stray pixels way out in the tails of the distribution.

D) Filtered analysis array. This shows the results of low-pass filtering (33 mm) the analysis array in graphic B.

E) Gradient of filtered analysis array. This shows the gradient (i.e. slope) magnitude at all points in the filtered analysis array. The rms of this array is compared against specification in the text at the bottom of the page.

F) Histogram of gradient array. This graphic shows the histogram of the array in graphic E. The rms is shown, along with 95, 98, and 99% values (for information only).

G) Part measurement phase. This graphic is identical to graphic A, less the clear aperture indications. The values are also reported in nm instead of waves.

H) Analysis array. This graphic is the same as graphic B, except astigmatism has not been removed.

I) Filtered analysis array. This shows the results of low-pass filtering (10 mm) the analysis array in graphic H.

J) Gradient of filtered analysis array. This shows the gradient (i.e. slope) magnitude at all points in the filtered analysis array (graphic I).

K) Histogram of gradient array. This graphic shows the histogram of the array in graphic J. The rms is shown, along with 95, 98, and 99% values. The rms and peak of this array are compared against specifications in the text at the bottom of the page. We use a "99%" peak, defined as 1% of the array pixels having a value above this to reduce the effects of noise from spurious sources such as dust diffraction rings, etc. In this example the true peak gradient is almost twice that of the 99% value. (The example in this figure is dominated by a quasi-oval vibration pattern, which is probably influencing the true peak gradient, and even the 99% gradient values. Consequently this particular measurement would need to be used with caution.)

L) Histogram of Δn over 1 cm patches. This graphic shows the peak-to-valley (99%) in the index of refraction over the ~ 1600 1 x 1 cm patches in the clear aperture. In each aperture, the peak-to-valley is defined such that 0.5% of the pixels have values above the peak and 0.5% have values below the valley. The index is calculated as the peak-to-valley phase in waves (99%), divided by the part thickness, multiplied by λ. The p_v of the worst patch is compared against specification in the graphic and in the text at the bottom of the page.

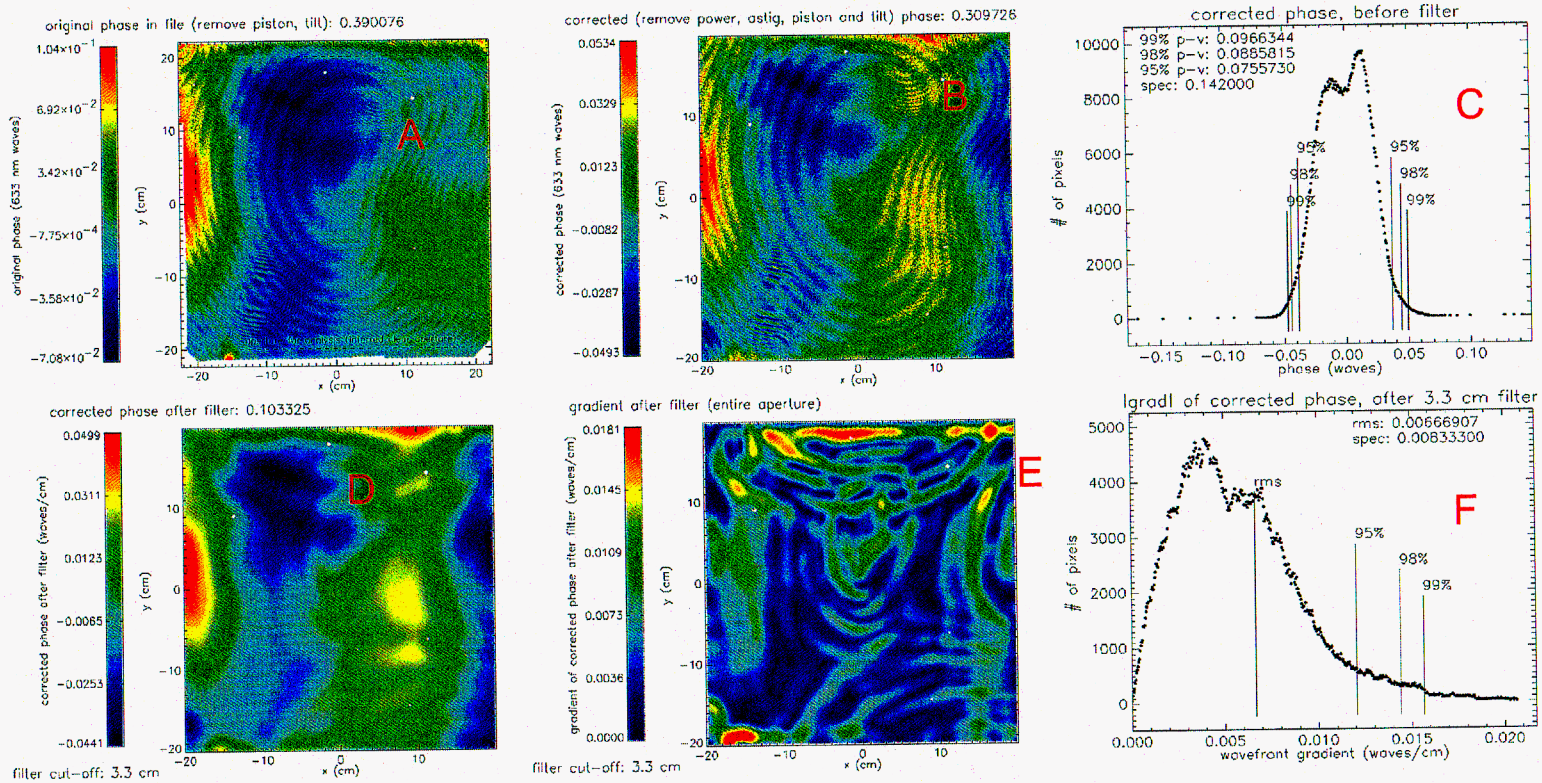

PHOM File:
Surface File:
Files Analyzed: Thu Feb 24 13:16:19 2000
PHOM file comments:
PHOM file title:
Analysis comments: test
(one wave: 633 nm; PHOM file array size: 1024x1024; PHOM calc. array size: 1024x 1024)

PHOM Measured: 04/26/19 10:07:05
Surface Measured: 04/26/19 10:20:16
Analysis Mode: Final pass/fail
Refractive index: 1.52000
Part thickness (cm): 4.63000

	Specification	Measured	Pass/Fail
LLNL, PHOM – Inferred full–slab power / sign (p–v, waves):	0.425000	0.105431 / +	PASS
LLNL, PHOM – Inferred full–slab astigmatism (p–v, waves):	0.220000	0.185600	PASS
LLNL, PHOM – Residual phase, unfiltered (99%) (p–v, waves):	0.142000	0.0966344	PASS
LLNL, PHOM – RMS filtered (33 mm) phase gradient (waves/cm):	0.00833300	0.00666907	PASS
LLNL, Surface – Measured clear aperture P_V (99%, waves):	5.00000	1.24298	PASS
Measured PHOM power / sign (p–v waves):	n/a	0.273602 / +	n/a
Measured PHOM astig. mag. / angle (p–v waves / deg):	n/a	0.0820447 / –5	n/a

phomglass, rev. 1.3d

Figure C-1. PHOMGLASS output, 1st page.

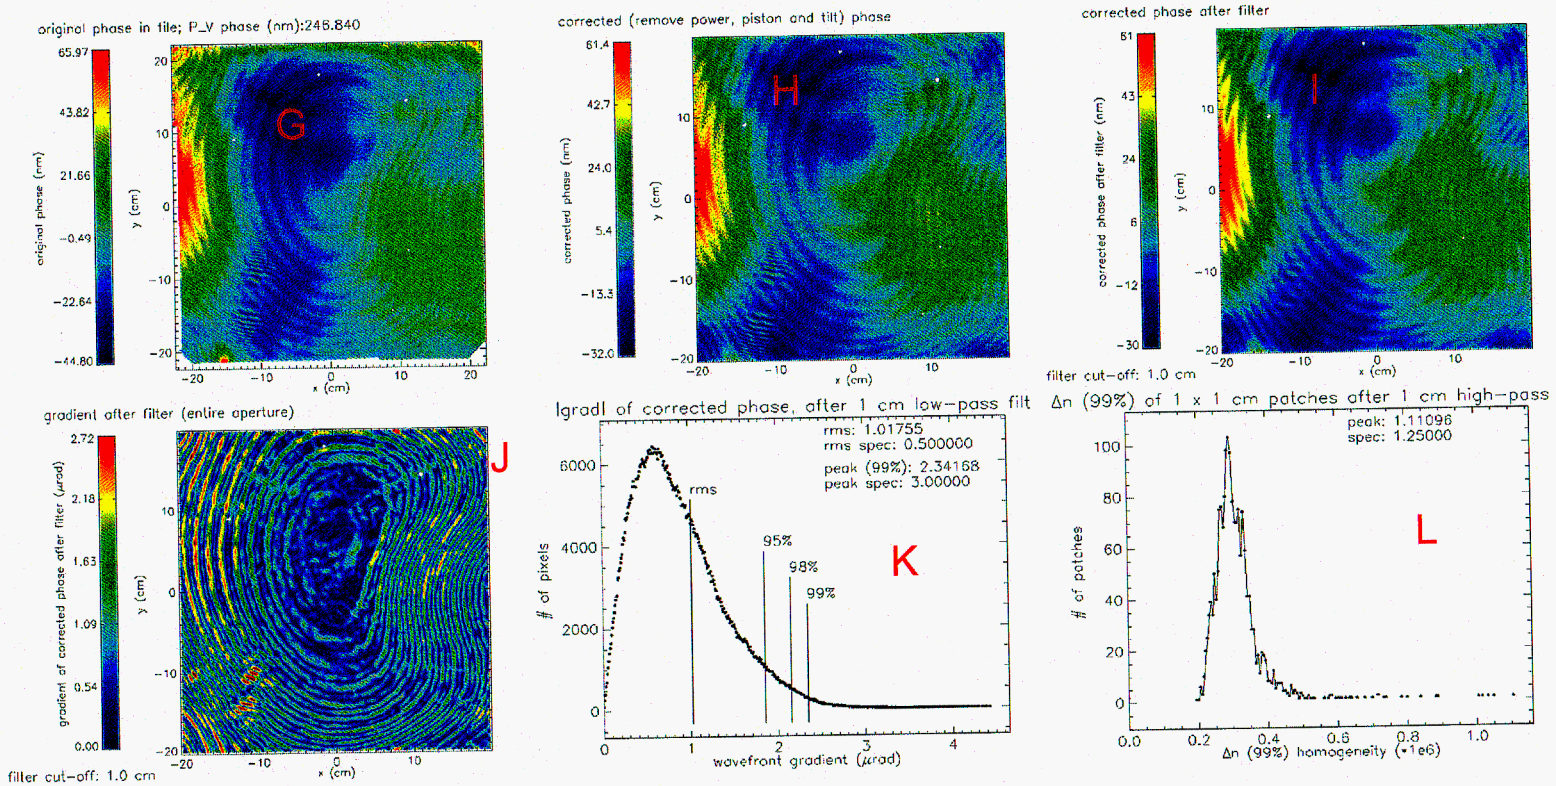

PHOM File:
Surface File:
Files Analyzed: Thu Feb 24 13:16:19 2000
PHOM file comments:
PHOM file title:
Analysis comments: test
(one wave: 633 nm; PHOM file array size: 1024x1024; PHOM calc. array size: 1024x 1024)

PHOM Measured: 04/26/19 10:07:05
Surface Measured: 04/26/19 10:20:16
Analysis Mode: Final pass/fail

	Specification	Measured	Pass/Fail
CEA, PHOM – Inferred full-slab power / sign (p-v, nm):	500	66.7169 / +	PASS
CEA, PHOM – RMS filtered (10 mm) phase gradient (μrad):	0.500000	1.01755	* FAIL *
CEA, PHOM – Peak (99%) filtered (10 mm) phase gradient (μrad):	3.00000	2.34168	PASS
CEA, PHOM – Δn (99%) over 10 mm patches (*1e6):	1.25000	1.11096	PASS
CEA, Surface – Measured clear aperture P_V (99%, waves):	5.00000	1.24298	PASS
CEA, Surface – Max. P_V phase over 100 mm patches (99%, waves):	0.500000	0.177318	PASS
Measured power / sign (p-v, nm):	n/a	173.135 / +	n/a

phomglass, rev. 1.3d

Figure C-2. PHOMGLASS output, 2nd page.

Appendix D: CPHASE Printouts

CPHASE is designed for inspection of switch and converter crystals. There are several specifications for these crystals, applied variously at the blank and finishing stages (see Figure 1 in the main text). This code is structured somewhat different than the others in that it is menu-driven based on the specification of interest, rather than the part. Consequently, complete analysis of a crystal may take several executions of the code, one for each of the specifications. The code was structured this way because some calculations are lengthy, and some require multiple files which may not be available for in-process evaluations.

The various analysis options are as follows:

- RMS wavefront in the PSD1 (33 – 2.5 mm) band.
- RMS gradient of the wavefront (low-pass filtered).
- Ramped phase jump for crystal boundaries.
- Crystal axis wander from Orthogonal Polarization Interferometry (OPI) measurements.
- Crystal axis wander from OPI, compared to CATS data lineouts, to give crystal cutting information.
- Crystal axis wander from a previously completed evaluation, compared to CATS data lineouts.

Figures D-1 thru D-6 show printouts for the various specification evaluations. The required files for these calculations, and the various graphics in the printouts are detailed below. *As in all code outputs, the specification values at the bottom of the sheets are for reference only, and do not supercede the optics drawings.*

RMS Wavefront in the PSD1 Band (33 – 2.5 mm) – Figure D-1.

This evaluation requires a TWF file (and empty cavity, if available). It performs a band-pass filter at 2.5-33 mm, and evaluates the resulting rms. The printout graphics are as follows:

A) Part measurement phase. This is the phase as retrieved from the part measurement file (cavity, if any, subtracted), converted into 633 nm waves, with piston and tilt removed. The entire camera array is shown, so there may be white space (i.e. "invalid" pixels) around the image. Invalid pixels within the image are also shown in white. The peak-to-valley over the image is reported on the top of the graphic. Note that this may be a larger number than the difference between the max and min of the color bar scale. This is because extremes of the color bar are set so that 0.25% of the pixels have values higher than the maximum, and 0.25% have values lower than the minimum. This compression yields a more useful image of the data. Noise in the data will often cause a few pixels in the raw interferogram to be abnormally high or low. If these values are used to set the color bar range then the bulk of the image is bleached out. This compression is used in all image representations.

The (0,0) point in the x and y scales is set at the center of the camera.

Two dashed boxes appear also in this example, which relate to the clear aperture to be used in later analysis. The heavier dashed box is the code's attempt to find the clear aperture from the edges of the data set. This will often be wrong, so the user has the ability to manually draw a new clear aperture, shown as the light dashed box. A person viewing the output can tell if this has been done because both boxes will be shown, and the text "Click and drag ..." is superimposed on the image. To aid in this aperturing, information on with width / height of the aperture selected (the light dashed line in this example) and the clear aperture as specified on the drawing are printed in the lower left-hand corner of the graphic.

B) Band-passed analysis array. This graphic shows the analysis array after it has been band-pass filtered over the PSD1 regime (2.5-33 mm) using the quad-flip technique. The rms of the entire array (excluding invalid data points) is reported at the top. In addition, the array is sub-sectioned into ~16 sub-arrays of ~100x100 mm, and the rms in each of these sub-apertures reported. The rms for the entire array is reported in the text at the bottom of the page.

C) RMS in sub-apertures. This graphic plots the sub-aperture PSD1 information from graphic B versus sub-aperture number. The value at 0 is the full aperture. This allows quick identification of bad sub-apertures. It is intended that this information will aid in identifying poorly finished areas or local inhomogeneities. Lines for the full-aperture PSD1 specification and twice that are shown for reference.

D) 2D PSD. This graphic shows the 2D PSD of the analysis array. It is generated by taking the Fourier transform of the analysis array, multiplied by a Hanning window, and corrected by the $8/3^{rd}$ factor. This graphic is intended to aid in determination of strong spectral components in the measurement.

E) Full-spectrum 1D Collapse. This graphic shows the 2D PSD in the previous graphic, collapsed into its 1D projections, along with a not-to-exceed line. Because it covers the full spectrum, this graphic is for reference only.

F) Band-passed 1D Collapse. This graphic shows the 2D PSD collapsed into its 1D projections, after having first been band-passed filtered to cover only the PSD1 regime. The filtering is abrupt (i.e. does not use an error-function edge) because we are not doing an inverse transform. If any of the lines cross above the not-to-exceed line in the PSD1 regime, this constitutes a failure of the Do Not Exceed specification.

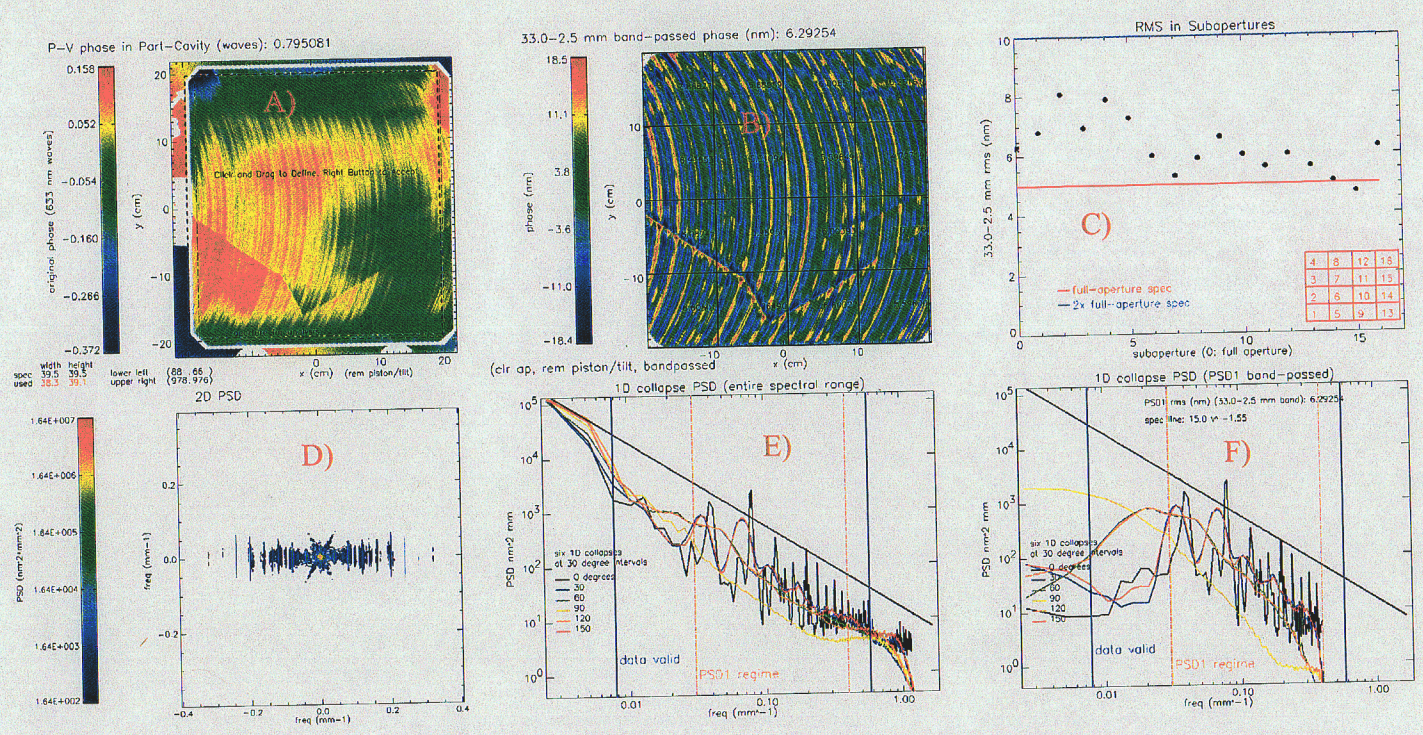

Analysis type: TWF, rms phase in 33-2.5 mm band; 1D collapse (requires h-pol TWF file)

File:
C:\metr_test_data\crystal files\opi tests\620511_SHG\L_SHG_620511_TWFHPOL_AT_02.dat

Measured:
02/07/2001 07:27:16

File Title:
SHG 620511 42cm

File Comments:
KDP Doubler, arrow to TF, 2nd Null OPI, H pol

Analyzed: Thu Aug 08 08:33:38 2002

Analysis comments: test

(one wave: 633 nm; file array size: 1024x1024; clear aperture array size: 890x 910; crystal reference frame)

Drawing	rms phase in 33-2.5 mm band (nm)	not-to-exceed line
switch_finished	5.0(3.5goal)	pass
doubler_finished	5.0	pass
tripler_finished	5.0	pass
(this measurement)	6.29254	fail

cphase, rev. 1.4

Figure D-1. CPHASE output for PSD1 evaluation.

RMS Gradient Wavefront (>33 mm) – Figure D-2.

This evaluation requires a TWF file (and empty cavity, if available). It performs a low-pass filter at 33 mm, and evaluates the resulting rms gradient. The printout graphics are as follows:

A) Part measurement phase. This is the phase as retrieved from the part measurement file (cavity, if any, subtracted), converted into 633 nm waves, with piston and tilt removed. The entire camera array is shown, so there may be white space (i.e. "invalid" pixels) around the image. Invalid pixels within the image are also shown in white. The peak-to-valley over the image is reported on the top of the graphic. Note that this may be a larger number than the difference between the max and min of the color bar scale. This is because extremes of the color bar are set so that 0.25% of the pixels have values higher than the maximum, and 0.25% have values lower than the minimum. This compression yields a more useful image of the data. Noise in the data will often cause a few pixels in the raw interferogram to be abnormally high or low. If these values are used to set the color bar range then the bulk of the image is bleached out. This compression is used in all image representations.

The (0,0) point in the x and y scales is set at the center of the camera.

Two dashed boxes appear also in this example, which relate to the clear aperture to be used in later analysis. The heavier dashed box is the code's attempt to find the clear aperture from the edges of the data set. This will often be wrong, so the user has the ability to manually draw a new clear aperture, shown as the light dashed box. A person viewing the output can tell if this has been done because both boxes will be shown, and the text "Click and drag ..." is superimposed on the image. To aid in this aperturing, information on with width / height of the aperture selected (the light dashed line in this example) and the clear aperture as specified on the drawing are printed in the lower left-hand corner of the graphic.

B) Analysis array: measurement over clear aperture. This graphic shows the part file over the clear aperture. This is the array which is used for all further analysis.

C) Filtered analysis array. This shows the results of low-pass filtering (33 mm) the analysis array in graphic B. The peak-to-valley of this array is compared against specification in the text.

D) Gradient of filtered analysis array. This shows the gradient (i.e. slope) magnitude at all points in the filtered analysis array (graphic D). The rms of this array is compared against specification in the text.

E) Histogram of gradient array. This graphic shows the histogram of the array in graphic D. The rms is shown, along with 95, 98, and 99% values (for information only).

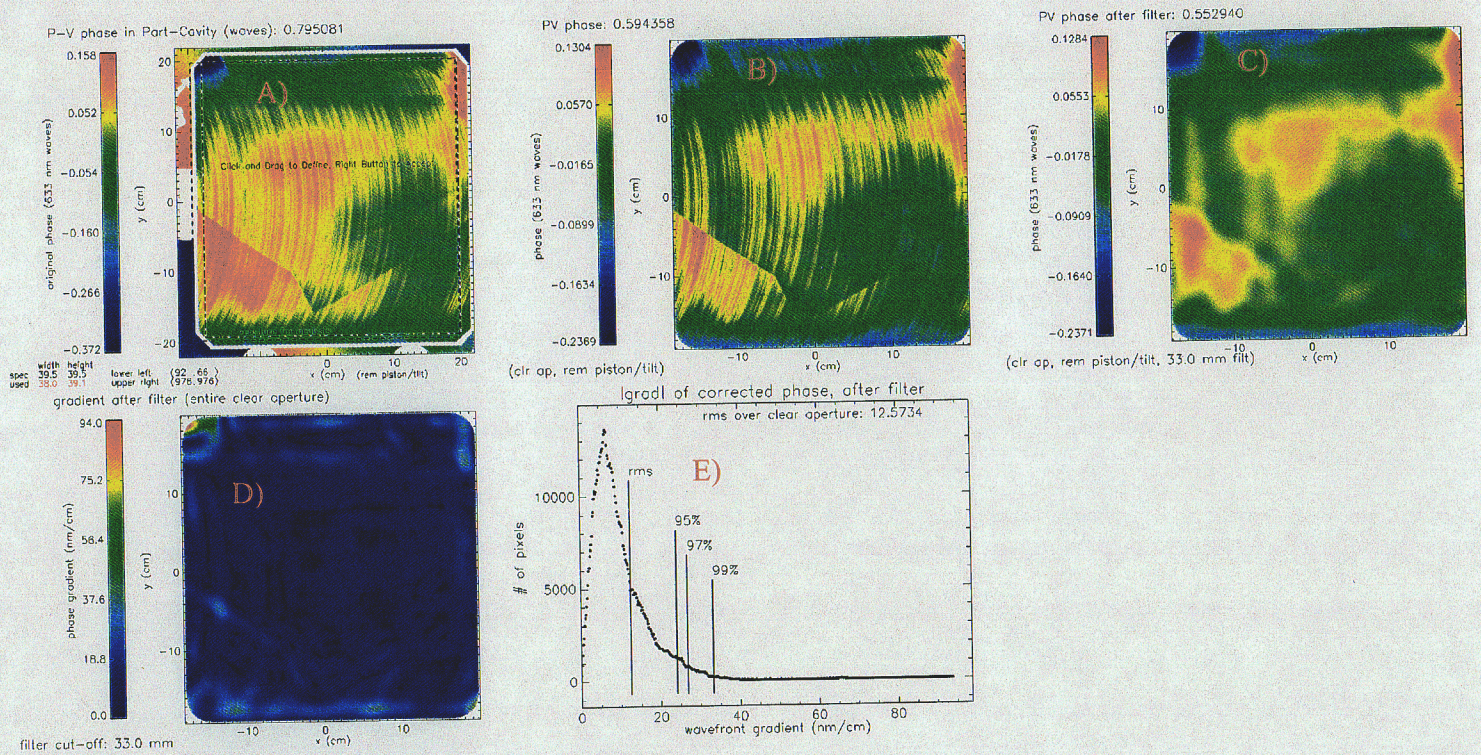

Analysis type: TWF, low-pass filtered phase gradient (requires h-pol TWF file)

File:
C:\metr_test_data\crystal files\opi tests\620511_SHG\L_SHG_620511_TWFHPOL_AT_02.txt

Measured: 02/07/2001 07:27:16

File Title: SHG 620511 42cm

File Comments:
KDP Doubler, arrow to TF, 2nd Null OPI, H pol

Analyzed: Thu Aug 08 08:32:31 2002

(one wave: 633 nm; file array size: 1024x1024; clear aperture array size: 885x 911; crystal reference frame)

Analysis comments: test

Drawing	rms gradient (nm/cm)
switch_finished	11
doubler_finished	11
tripler_finished	11
(this measurement)	12.5734

cphase, rev. 1.4

Figure D-2. CPHASE output for rms gradient evaluation.

Ramped Phase Jump – Figure D-3.

This evaluation requires a TWF file (and empty cavity, if available). It performs the various calculations necessary to look for ramped phase jumps. The printout graphics are as follows:

A) Part measurement phase. This is the phase as retrieved from the part measurement file (cavity, if any, subtracted), converted into 633 nm waves, with piston and tilt removed. The entire camera array is shown, so there may be white space (i.e. "invalid" pixels) around the image. Invalid pixels within the image are also shown in white. The peak-to-valley over the image is reported on the top of the graphic. Note that this may be a larger number than the difference between the max and min of the color bar scale. This is because extremes of the color bar are set so that 0.25% of the pixels have values higher than the maximum, and 0.25% have values lower than the minimum. This compression yields a more useful image of the data. Noise in the data will often cause a few pixels in the raw interferogram to be abnormally high or low. If these values are used to set the color bar range then the bulk of the image is bleached out. This compression is used in all image representations.

The (0,0) point in the x and y scales is set at the center of the camera.

Two dashed boxes appear also in this example, which relate to the clear aperture to be used in later analysis. The heavier dashed box is the code's attempt to find the clear aperture from the edges of the data set. This will often be wrong, so the user has the ability to manually draw a new clear aperture, shown as the light dashed box. A person viewing the output can tell if this has been done because both boxes will be shown, and the text "Click and drag ..." is superimposed on the image. To aid in this aperturing, information on with width / height of the aperture selected (the light dashed line in this example) and the clear aperture as specified on the drawing are printed in the lower left-hand corner of the graphic.

B) Analysis array: measurement over clear aperture. This graphic shows the part file over the clear aperture. This is the array which is used for all further analysis. Phase boundaries show better in black and white.

C) Crystal regions failing the 0.5 – 3 mm specification. This graphic shows regions which fail the small-region spec. Regions in blue fail the switch crystal spec. Those in red fail both the switch and (both) converter crystal specs. (The switch crystal specification is tighter than the converter crystals.) The area failure fraction for both crystal types is shown at the top.

As obvious in this and subsequent graphics, this specification suffers from the weakness that non-crystal boundaries are also found (most notably, crystal finishing marks), so the area failure fraction should be used carefully when looking for crystal boundaries. The spec does, however, have some utility in finding those boundaries visually and evaluating their magnitude.

D) Crystal regions failing a half-actual 0.5 – 3 mm specification. This graphic shows regions which would fail the small-region spec if it were set at half its value. It is for reference only, and is intended as an aid in determining how close visual boundaries are to failing spec. It also show clearly how finishing marks can contribute to the failed-fraction evaluation. Color coding is the same as in graphic C).

E) Crystal regions failing the 0.3 – 33 mm specification. This graphic shows regions which fail the large-region spec. Color coding is the same as in graphic C). Again, because of the way the spec is written, features not associated with crystal boundaries can show up as failures. Most dramatically, the circles evident in the figure come from isolated pixels of especially large (or small) magnitude which cause much of the region in a 33 mm diameter circle to fail. Obviously, this has nothing to do with crystal boundaries.

F) Crystal regions failing a half-actual 3 - 33 mm specification. This graphic shows regions which would fail the large-region spec if it were set at half its value. It is for reference only, and is intended as an aid in determining how close visual boundaries are to failing spec. Color coding is the same as in graphic C).

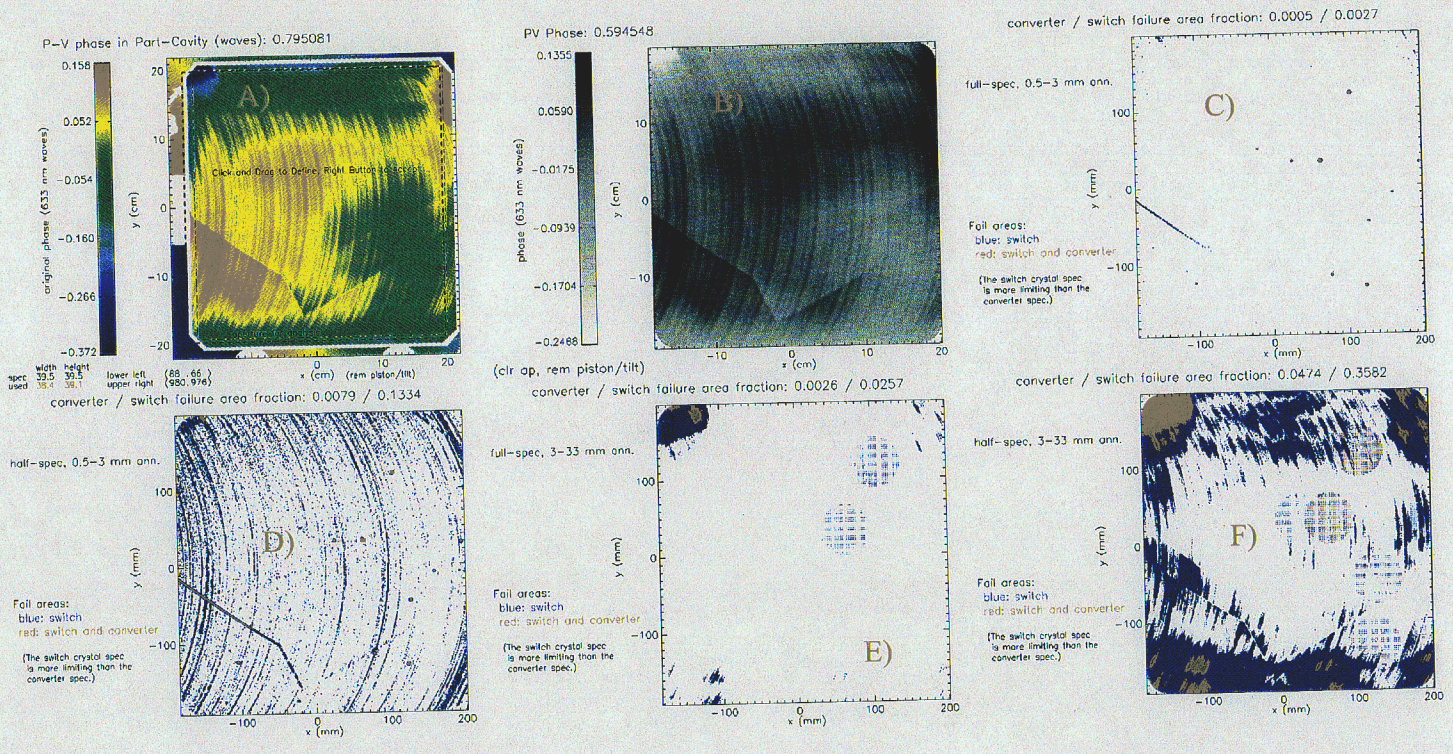

Analysis type: TWF, ramped phase jump (requires h-pol TWF file)

File:
C:\metr_test_data\crystal files\opi tests\620511_SHG\L_SHG_620511_TWFHPOL_AT_02.dat

Measured:
02/04/2001 07:27:16

File Title:
SHG 620511 42cm

File Comments:
KDP Doubler, arrow to TF, 2nd Null OPI, H pol

Analyzed: Thu Aug 08 08:35:48 2002

Analysis comments: test

(one wave: 633 nm; file array size: 1024x1024; clear aperture array size: 893x 911; crystal reference frame)

Drawing	switch ramped phase jump	converter ramped phase jump
switch_finished	0.0100000	0.000000
doubler_finished	0.000000	0.0100000
tripler_finished	0.000000	0.0100000
(this measurement)	0.00267602	0.0257276

cphase, rev. 1.4

Figure D-3. CPHASE output for ramped phase jump specification evaluation.

Crystal Axis Wander from OPI – Figure D-4.

This evaluation requires two TWF files, one at h polarization and one at v polarization. It also requires an empty cavity file. It yields an inhomogeneity map of the crystal. The printout graphics are as follows:

A) V-polarization TWF. This is the phase as retrieved from the part measurement file in v-polarization with no piston or tilt removed. The entire camera array is shown. Since the part underfills the camera, several regions of white space can be seen. Referencing the figure, the two half-moons on the bottom indicate parts of the mount, as does the large white region on the right. A fiducial arrow shows the axis orientation. Most important, the edge of the crystal is masked (in the measurement software), seen as a narrow white strip, so the region inside this boundary is crystal and that outside is the RF behind. It is necessary to be able to view some of the RF so absolute tilt angles can be calculated. The (0,0) point in the x and y scales is set at the center of the camera.

B) Empty Cavity File. This is the empty cavity with no piston or tilt removed.

C) H-polarization TWF. This is the part as measured in h-polarization with no piston or tilt removed. In addition, the RF regions are blacked-out, which is done during the analysis. Input information on the indices of refraction (extraordinary and ordinary), thickness, deuteration fraction, and use angle are reported here for reference, as is the calculated $dn/d\theta$. The text is added during the analysis to aid the user in identifying the part and RF regions.

D) Calculated crystal axis variation. This is the resulting map of variation in the crystal axis, in units of microrad. Since there is no absolute reference (from CATS), *piston is removed in this image.* Consequently, it represents variation in homogeneity around an average. (That is, the average crystal axis orientation relative to the surface normal is unknown.)

E) Histogram of homogeneity map. This is a histogram of graphic D). Again, piston is removed.

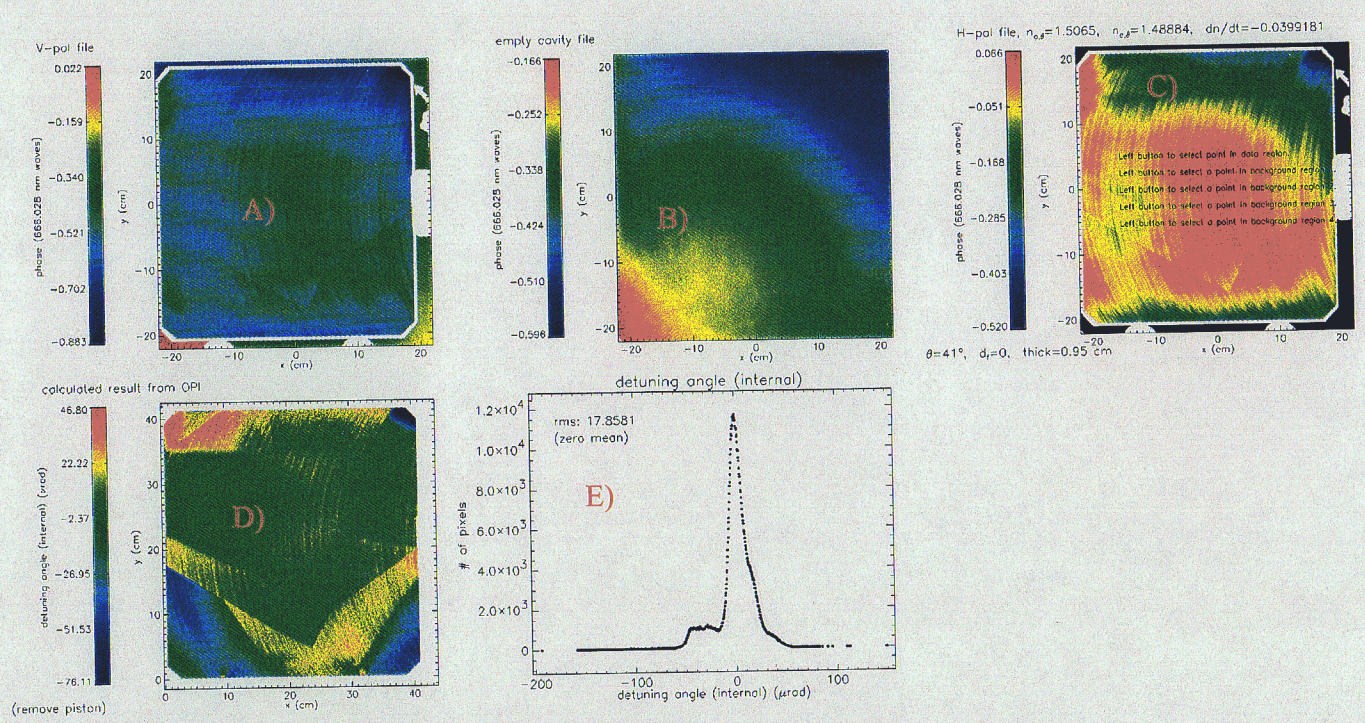

Analysis type: birefringence variation (OPI), analysis 3 (requires h-pol and v-pol TWF, empty cavity files -- outputs OPI detuning file)

File:	Measured:	File Title:	File Comments:
C:\metr_test_data\crystal files\opi tests\620511_SHG\L_SHG_620511_TWFHPOL_AT_01.crd	02/27/2001 07:27:16	SHG 620511 42cm	KDP Doubler, arrow to TF, 2nd Null OPI, H pol
C:\metr_test_data\crystal files\opi tests\620511_SHG\L_SHG_620511_TWFVPOL_AT_01.crd	02/27/2001 07:39:30	SHG 620511 42cm	KDP Doubler, arrow to TF, 2nd Null OPI, V pol
C:\metr_test_data\crystal files\opi tests\620511_SHG\620511_MTcavity_Hpol_OPI_01.crd	02/27/2001 09:47:05	SHG 620511 MTcav	Empty Cavity for SHG 620511, OPI, H pol
no_file			

Analyzed: Fri Mar 23 10:47:56 2001

Analysis comments:

(one wave: 666.028 nm; file array size: 1024x1024; clear aperture array size: 1024x 1024; crystal reference frame)

Spec0: spatial birefringence variation from OPI (delta-n) for doubler
Spec1: spatial birefringence variation from OPI (delta-n) for tripler

Drawing	Spec0	Spec1
doubler_finished	1.00000e-006	0.000000
tripler_finished	0.000000	2.99000e-006
doubler_blank	1.00000e-006	0.000000
tripler_blank	0.000000	2.99000e-006
(this measurement)	7.12860e-007	7.12860e-007

cphase, rev. 1.1

Figure D-4. CPHASE output for evaluation of crystal axis variation.

Crystal Axis Wander from OPI, Compared to CATS Lineouts – Figure D-5.

This evaluation requires two TWF files, one at h polarization and one at v polarization, and an empty cavity file. It also requires a CATS measurement file. It yields a crystal axis variation map of the crystal, with the variation along the CATS measurement lineout compared to the CATS measurement. The difference in the average of these two lineouts is used to estimate the difference between the average crystal axis and the surface normal, which can then be used to re-cut the crystal. The printout graphics are as follows:

A) V-polarization TWF. This is the phase as retrieved from the part measurement file in v-polarization with no piston or tilt removed. The entire camera array is shown. Since the part underfills the camera, several regions of white space can be seen. Referencing the figure, the two half-moons on the bottom indicate parts of the mount, as does the large white region on the right. A fiducial arrow shows the axis orientation. Most important, the edge of the crystal is masked (in the measurement software), seen as a narrow white strip, so the region inside this boundary is crystal and that outside is the RF behind. It is necessary to be able to view some of the RF so absolute tilt angles can be calculated. The (0,0) point in the x and y scales is set at the center of the camera.

B) Empty Cavity File. This is the empty cavity with no piston or tilt removed.

C) H-polarization TWF. This is the part as measured in h-polarization with no piston or tilt removed. In addition, the RF regions are blacked-out, which is done during the analysis. Input information on the indices of refraction (extraordinary and ordinary), thickness, deuteration fraction, and use angle are reported here for reference, as is the calculated $dn/d\theta$. The text is added during the analysis to aid the user in identifying the part and RF regions.

D) Calculated crystal axis variation. This is the resulting map of variation in the crystal axis, in units of microrad. Since there is no absolute reference (from CATS), *piston is removed in this image.* Consequently, it represents variation in homogeneity around an average. (That is, the average crystal axis orientation relative to the surface normal is unknown.) The location of the CATS measurement points is also shown here. The square location indicators are also intended to represent the approximate size of the CATS beam.

E) Histogram of homogeneity map. This is a histogram of graphic D). Again, piston is removed.

F) Comparison of the CATS and OPI lineouts. This shows the results from the OPI analysis in graphic D) along the same points as the CATS data, compared to the CATS data. The difference of the averages of these two lines represents an average tilt in the part surface normal compared to the crystal axis. This information can then be used to re-cut the crystal.

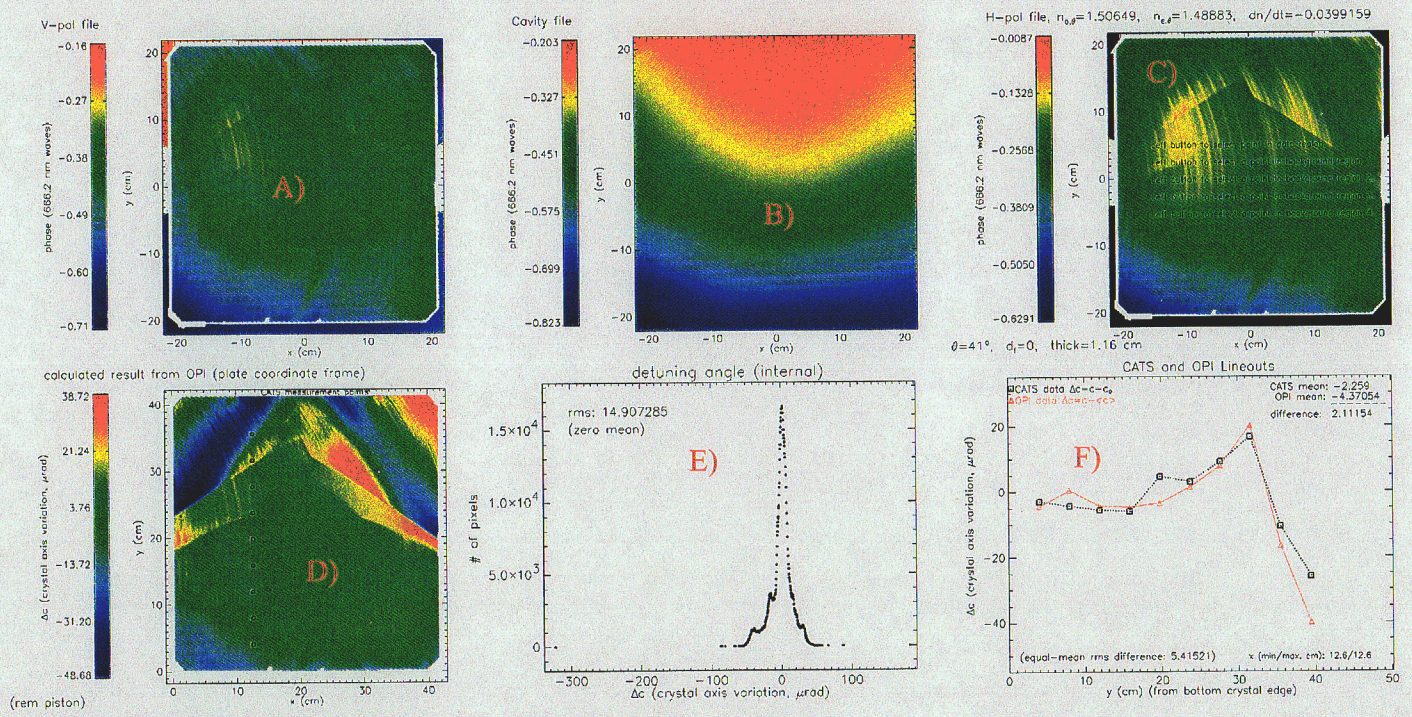

Analysis type: birefringence variation (OPI), analysis 1 (requires h-pol and v-pol TWF, empty cavity, and CATS lineout files)

File:	Measured:	File Title:	File Comments:
E:\Crystals\D11-21\041002\PSI\I_SHG_D11-21_TWFHPOL.opd	04/10/2002 16:46:38	D11-21	TWF, HORIZONTAL
E:\Crystals\D11-21\041002\PSI\I_SHG_D11-21_TWFVPOL.opd	04/10/2002 16:41:52	D11-21	TWF, VERTICAL
E:\Crystals\D11-21\041002\PSI\I_SHG_D11-21_HCAV.opd	04/10/2002 17:12:17	D11-21	CAVITY, VERTICAL
E:\Crystals\D11-21\042302\CATS\D11-21 SHG042302132730OPI.dat	Tue Apr 23 13:58:03 2002		(CATS file)

Analyzed: Tue Apr 23 13:58:43 2002

Analysis comments: AA, pointing right

(one wave: 666.2 nm; file array size: 1024x1024; clear aperture array size: 1024x 1024; crystal reference frame)

if c-axis error is positive, decrease angle between crystal axis and surface normal
if c-axis error is negative, increase angle between crystal axis and surface normal

Drawing	OPI (dn rms)	mean c-axis error (microradians)
doubler_finished	1.00000e-006	10.0000
tripler_finished	2.99000e-006	20.0000
(this measurement)	5.9503775e-007	2.1115446

cphase, rev. 1.4

Figure D-5. CPHASE output for evaluation of crystal axis variation, compared to CATS measurement.

Crystal Axis Wander from Stored OPI File, Compared to CATS Lineouts – Figure D-6.

This evaluation requires a crystal axis wander map from a previous OPI evaluation and a CATS measurement file. It yields a crystal axis variation map of the crystal, with the variation along the CATS measurement lineout compared to the CATS measurement. The difference in the average of these two lineouts is used to estimate the difference between the average crystal axis and the surface normal, which can then be used to re-cut the crystal.

This analysis routine is intended to be used where the crystal may be re-cut, then retested on CATS, then re-cut again. It removes the need of conducting the OPI measurement and analysis, which produces the axis wander map, each time.

The stored crystal axis wander map used as input contains text information about the files used in that analysis, along with the wander map. It also contains, for reference, the intensity frame map from the h-polarization measurement. The printout graphics are as follows:

A) Information related to crystal axis wander map generation. This region of the output lists the various files and other information used in the evaluation which yielded the crystal axis wander map used as input here.

B) Crystal axis variation. This is the input map of variation in the crystal axis, in units of microrad. Since there is no absolute reference (from CATS), *piston is removed in this image.* Consequently, it represents variation in homogeneity around an average. (That is, the average crystal axis orientation relative to the surface normal is unknown.) The location of the CATS measurement points is also shown here. The square location indicators are also intended to represent the approximate size of the CATS beam.

C) Histogram of homogeneity map. This is a histogram of graphic D). Again, piston is removed.

D) Comparison of the CATS and OPI lineouts. This shows the results from the OPI analysis in graphic D) along the same points as the CATS data, compared to the CATS data. The difference of the averages of these two lines represents an average tilt in the part surface normal compared to the crystal axis. This information can then be used to re-cut the crystal.

E) Summary information. A reiteration of the results shown in graphic D).

F) Intensity image. The interferometer stores an intensity image, as well as the resulting phase wavefront estimate. That intensity image was read in during the analysis which generates the axis wander map. It was stored in that file and re-displayed here for reference. Occasionally it is useful to be able to see this map. The dashed line indicates the code's opinion of the part edges. This graphic is for reference only.

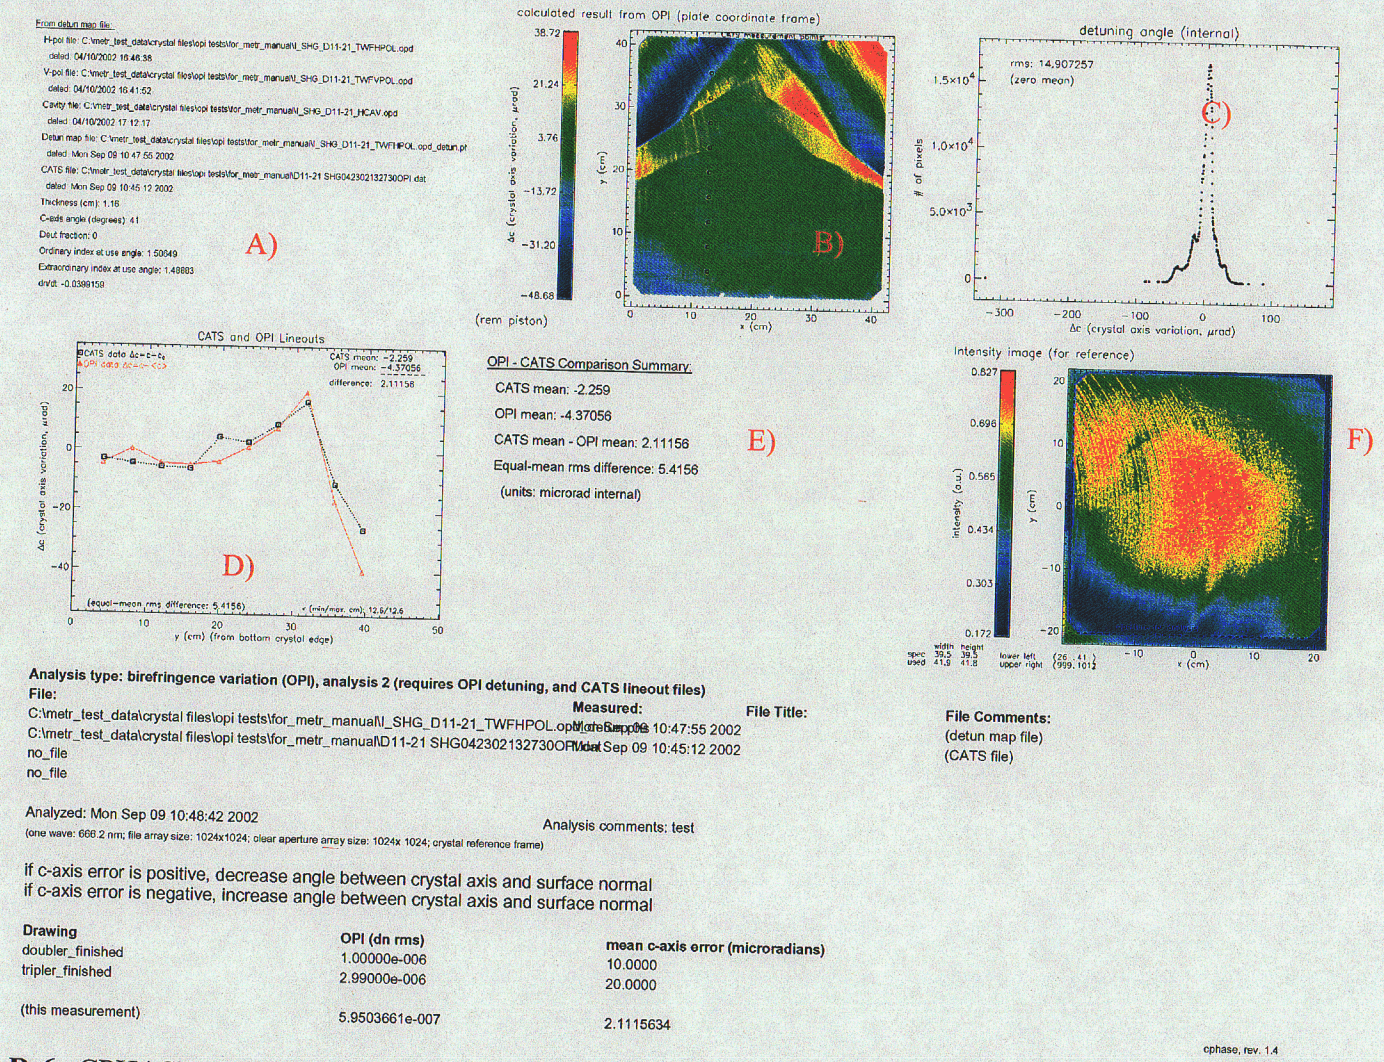

Figure D-6. CPHASE output for evaluation of crystal axis variation, compared to CATS measurement, using stored axis wander map.