

PVr—a robust amplitude parameter for optical surface specification

Chris J. Evans, MEMBER SPIE
Zygo Corporation
Laurel Brook Road
Middlefield, Connecticut 06455
E-mail: cevans@zygo.com

Abstract. Peak-to-valley departure (PV) is entrenched in optics design and manufacture as a characterization of an optical figure; modern interferometers commonly use $1k \times 1k$ detectors, the output of which may not be well represented by two points. PVr is a newly proposed robust amplitude parameter that combines the PV of a 36-term Zernike fit and the root mean square of the residual. This provides automatic filtering, is insensitive to system resolution, and relates directly to imaging performance via the Marechal criterion. Use of PVr in place of PV is recommended. © 2009 Society of Photo-Optical Instrumentation Engineers. [DOI: 10.1117/1.3119307]

Subject terms: metrology; interferometry; optical testing.

Paper 080347RR received May 5, 2008; revised manuscript received Feb. 12, 2009; accepted for publication Feb. 18, 2009; published online Apr. 27, 2009.

1 Introduction

Traditionally, optical surfaces have been specified by the peak-to-valley departure [(PV), height difference between highest and lowest points on the surface after removal of piston and tilt for flats and best fit sphere for spherical surfaces] from the nominal surface shape. In some cases, specifications that relate directly to the particular optical function in a given application have been developed. Such cases are in the minority; the majority of optics sold today are still specified as λ/N (e.g., a quarter-wave, tenth-wave, etc.). Such specifications can cause difficulties due, for example, to defects and dirt on the test surfaces or in the interferometer, especially as spatial resolution increases. In addition, it is not uncommon to find discrepancies between outgoing quality assurance at an optic's vendor and incoming inspection by its customer. Interferometers with different spatial resolutions and different noise characteristics report different values of PV. Hence, a robust amplitude parameter (PVr, where "r" represents "robust") is proposed. Section 2 discusses the old connection between PV and optical performance, and Section 3 describes the problems with modern use of PV as an optical surface specification. Section 4 introduces the new parameter (PVr) and outlines its characteristics, before addressing in more detail (Section 4.5) the reasons for its robustness. Standardization is briefly discussed in Section 5.

Optical surface specifications using such characteristics as power spectral density functions, structure function, slopes, or scatter clearly have a direct connection to surface function. Use of such specifications is encouraged. In those cases, however, where PV is currently used, replacing it with PVr may reduce manufacturing costs and the number of disputes between buyers and sellers over part conformance.

2 PV in the Traditional Optical Shop

Classical opticians finished surfaces on pitch laps and assessed their work using a light box and test plate. Experienced opticians automatically filtered out "noise"—dirt or a pit in the test plate, for example—and used the shape of the fringes to assess the PV. The slightly more rigorous approach using, for example, a Polaroid print of the fringes and a parallelogram implies similar filtering. Fringe center finding methods using early laser Fizeau interferometers did much the same, although with slightly higher spatial resolution.

How does this filtered PV relate to functional performance? In an imaging system, one simple criterion for performance is the diffraction limit. Marechal's criterion states that the system will be diffraction limited if the wavefront is better than $\lambda/14$ root mean square (rms). Traditional optical manufacturing processes typically "polish out" quite quickly, and the bulk of the time is spent driving down the low-order surface aberrations. Considering these low-order aberrations as an expansion of Zernike terms (ignoring piston and tilt), we calculate from the definition of the polynomials^{1,2} the ratio of PV to rms term by term (Table 1).

Table 1 PV to rms ratio for low order aberrations.

Aberration	PV:rms ratio
Power	3.5
Astigmatism	4.8
Coma	5.6
Third-order spherical	3.4
Trefoil	6.3




Fig. 1 System wavefront as a function of number of elements made with $\lambda/4$ surfaces in perfectly homogeneous material. Uncorrelated errors add in quadrature (RSS); correlated errors add arithmetically.

It is no surprise, therefore, that the optician “knows” that the ratio of PV to rms is approximately 5:1. Now we can easily see (Fig. 1) that a transmissive system with up to four elements (eight surfaces) made from perfect material ($n=1.5$) with quarter-wave uncorrelated (from surface to surface) surface errors will meet the Marechal criterion when the PV-to-rms ratio is 5. Effects such as material inhomogeneity, some degree of correlation between the errors on different surfaces, alignment, and higher index glasses, all reduce the number of elements that can be used before Marechal’s criterion is reached. The key point, however, is that traditional manufacturing processes combined with traditional measurement methods have embedded within them some functional relationship between imaging performance and PV.

3 Problem with PV in Modern Manufacturing

Dramatic changes in optical fabrication and metrology methods over the last two decades have made PV a troublesome parameter to use when buying and selling optics:

1. Small tool fabrication processes (magnetorheological finishing, computer-controlled polishing, diamond turning, etc) produce surfaces that do not necessarily approximate the 5:1 rule.
2. High spatial resolution interferometers (essential, for example, for generating hit maps for small tool polishers and controlling midspatial frequencies) “see everything.” Small defects may be resolved that will have little or no impact on imaging performance of an optical system containing that surface.
3. There are no standards for frequency cutoffs, filters, or outlier rejection.
4. Different interferometers can give significantly different results when measuring the same wavefront, depending on the instrument resolution and the frequency content of the surface.




Fig. 2 Simulation of measured PV for normally distributed noise with unit standard deviation on a perfect surface (with PV=0).

3.1 PV is Biased

In the presence of noise, PV will give a value that is systematically larger than the “true” form of the surface. The size of the bias depends on the type of surface and the detector resolution. Consider, for a moment, the holy grail—a perfect surface. The reported PV of that perfect surface measured with a nearly perfect interferometer—one that only has Gaussian noise that is uncorrelated between pixels—depends on the sampling of the noise at each pixel.³ For a fixed noise level per pixel, the error in the reported PV increases with detector resolution (i.e., number of pixels). Figure 2 shows the result of a simple, dimensionless, simulation in which the PV was evaluated for 200 random samples of Gaussian noise, with a standard deviation of 1, on matrices sized to match commonly used detectors. The PV plotted is the multiple of the standard deviation of the input Gaussian. Explicitly, this simulation suggests that if there is 1-nm rms noise* in a single phase acquisition, the surface PV would be overestimated by 9–11 nm when measured with a 1k×1k detector. A more detailed discussion is given elsewhere.³

For real surfaces, only some fraction of the surface area is at heights that would be affected by sampling the extremes of the noise distribution. For pure sinusoidal surfaces, the simulation for a 1k×1k detector shows that the error in PV decreases to $\sim 7\times$ noise and is a weak function of frequency and of surface amplitude to noise ratio. On a $\lambda/10$ (approximately) surface, 1 nm Gaussian noise will increase the PV by ~ 7 nm ($>10\%$ error). The bias varies with the details of the surface and is well correlated with skewness and kurtosis.³

For comparison, consider rms as the surface specification. The 1-nm standard deviation noise discussed above would cause the perfect surface to be measured as 1-nm

*In practice, the CCD noise level depends on a number of factors, including electron well depth, dark current, amplifier noise, illumination, and contrast; pixel-to-pixel noise also arises from air turbulence, stray light, etc. Subnanometer noise should be expected in a well-designed, modern interferometer used under good conditions.




Fig. 3 Original data (a) residual after fitting 36 Zernikes (b) and surface generated from the fit Zernikes (c). PVr=11.5 nm (Section 4.1).

rms and a $\lambda/10$ PV (approximately) surface (12 nm rms) would be measured as 12.04 nm (<0.5% error).

3.2 Uncertainty

Another difficulty with PV arises from requirements to evaluate uncertainty in measurement results. This evaluation is complicated by the need to convolve systematic and time varying sources of uncertainty with the shape of the surface under test to find extreme values (evaluation of the uncertainty in the rms of a surface or wavefront is more straightforward and does not involve extreme points.) Uncertainty may vary as a function of position in the aperture, so that the “peak” in the measured data may not be at the same place as the peak uncertainty; the uncertainty in the PV should express the range of reasonably likely estimates of the PV of the surface, and those estimates may come from a range of different point pairs. This difficulty is discussed in more detail elsewhere;³ suffice it to say here that use of PVr as a surface specification avoids this problem.

4 PVr—An Improved Amplitude Parameter

As indicated in the previous sections, optical performance may not be well correlated with PV as measured on modern optical instruments. Other specifications are generally more appropriate for many applications. Despite this, the use of PV, or λ/N , is so entrenched in the culture of optics that it seemed appropriate to derive a robust parameter that retains that tradition. The goal was to provide a parameter that characterizes the high-resolution data from modern instruments while being largely insensitive to differences in detector resolution from instrument to instrument, being robust with respect to noise, requiring no “standardized” filtering, and being easy to implement on any metrology system. Use of such a parameter will reduce disagreements over whether a surface meets specification.

4.1 Definition of PVr

Most optics manufactured today still have circular apertures. Over these apertures, fitting to Zernike polynomials is extremely robust, even on relatively sparse data sets. There is no ambiguity about the definition of the Zernike polynomials,² but significant variation in the literature over their ordering. The 36-term set originally chosen by Loomis⁴ for “Fringe” (includes azimuthal orders up to 5θ

and radial polynomials to r^{12}) has been widely adopted in commercially available interferometers and optical design programs over the last 20 years. This set efficiently captures the low-order aberrations, which dominate the figure of most optics. The rms of the residual after fitting to 36 Zernikes is equally robust and adds a characterization of the midspatial frequencies. Combining these two contributions, the proposed definition of PVr for circular apertures is

$$PVr = PV_{36 \text{ Zernikes}} + 3 \times \sigma_{36 \text{ ZernikeResid}} \quad (1)$$

where the first term is the PV of the surface generated using the 36-term Zernike fit to the data and the second term is three times the rms of the residual after fitting and removing the 36 terms. (Other orthogonal polynomial sets could be used for other aperture shapes, for example Legendres for square apertures, Zernike–Tatian for annular apertures, etc. For such apertures, the polynomial set and order must be included in the specification.) Two constraints on the values of PVr are introduced in Section 5.

Figure 3 shows an example (full area calibration of a 150 mm reference surface). The PV of the surface generated from the 36-term Zernike fit is 10.0 nm, and the rms of the residual is 0.5 nm, hence,

$$PVr = 10.0 + 3 \times (0.5) = 11.5 \text{ nm.}$$




Fig. 4 305 mm TF. PV=53.8 nm, PVr=39.5 mm.

Table 2 Parameter change with detector resolution for data of Fig. 4.

Detector resolution	PV (nm)	PVr (nm)
1024 × 1024	53.81	39.46
512 × 512	38.27	39.41
256 × 256	36.88	39.33
128 × 128	33.90	39.19

Another example is given in Fig. 4. Here, the measurement of a 12-in. flat shows a PV of 53.8 nm. The PV of the Zernike fit was 36.39, and the rms of the Zernike residual was 0.99 nm, giving

$$\text{PVr} = 36.39 + 3 \times (0.99) = 39.46 \text{ nm.}$$

Note that the difference is large compared to the noise estimate of Fig. 2; that simulation considered only Gaussian noise, not edge diffraction, dirt, defects, etc. PVr addresses all such contributions.

An example of PVr's robustness as a function of resolution is shown in Table 2 (calculated by averaging appropriate pixels in raw data, i.e., assuming that system resolution is always optically limited and ignoring detector noise considerations). As noted earlier, the measured PV is sensitive to the pixel density of the interferometer for fixed noise per pixel. In addition, lower resolution systems effectively filter higher spatial resolution information. As shown in Table 2, the use of PVr gives consistent results that are largely independent of interferometer resolution.

For comparison, the data of Fig. 4 were low-pass filtered with a median filter. The resulting PV over approximately

the same range as the binning (3×3 to 9×9 compared to 2×2 to 8×8) ranged from 38.2 to 35.4 nm.

4.2 PVr to rms Ratio

As noted above, every optician “knows” that the PV-to-rms ratio is ~ 5 . For many measurements made in modern interferometers, even of classically produced optics, that is no longer true. Noise, edge diffraction, ghost reflections, small defects in the test optics, and dirt can all generate high-frequency spikes, which have no impact on imaging performance yet can drive up the reported PV. Table 3 shows some sample data from a range of test setups with detector resolutions from 640×480 to 2000×2000 . Clearly, the PVr-to-rms ratio is much closer to the opticians expectations—and to what you would expect from surfaces dominated by low-order aberrations.

Part No. 5 in Table 3 is a particularly good example of the advantage of PVr; the calibration flat has a number of fiducials on it, with associated edge diffraction. PVr is a good representation of the underlying surface shape, which would otherwise be obtained by a tedious and arbitrary masking operation. Direct comparison of the result of masking by an experienced optician to the use of PVr is given elsewhere.⁵

4.3 Implicit Filtering and Clipping

The filtering used in PVr is implicit in the definition, due to the 36-term Zernike polynomial fit, and is tied directly to the test aperture. It does not require filtering to a specified physical length on the part (which is sometimes difficult to implement, depending on part-to-detector mapping), nor does it require definition of the filter characteristics.

PVr should be less than PV; hence, some data points are implicitly truncated to the newly defined amplitude or may be considered to be clipped. In essence, this is a bandpass filter based on amplitude. It is instructive to consider this physically. For the data in Fig. 4, we can choose to clip

Table 3 PV, PVr and rms for a range of parts

No.	Part	PV (nm)	rms (nm)	PV:rms	PVr (nm)	rms (nm)	PVr:rms
1	f/0.75, 65-mm aperture	30.1	1.03	29	8.5	1.03	8
2	f/45 cavity	68.2	10.01	7	63.1	10.01	6
3	Zerodur test sphere	25.9	1.99	13	13.6	1.99	7
4	800-mm aperture flat	162.9	20.68	8	147.9	20.68	7
5	300-mm calibration flat	147.9	11.42	13	51.5	11.42	5
6	305-mm TF	26.4	5.21	5	24.0	5.21	5
7	101.4-mm TF	13.7	1.85	7	10.6	1.85	6
8	300-mm long cavity	45.4	4.9	9	31.1	4.9	6
9	200-mm spherical cavity	23.8	2.64	9	16.1	2.64	6
10	Steep asphere	62.8	6.31	10	43.9	6.31	7




Fig. 5 (a) PV=29.8 nm, PVr=18.1 nm, rms=3.4 nm and (b) Data of (a) clipped uniformly.

such that the peak is $1.5 \times$ rms residual greater than the peak of the Zernike-generated surface and the valley defined analogously. In this case, the difference in pixel count between PV and PVr is 46 pixels out of 785,400 (0.006%).

Figure 5(a) shows data taken in calibrating an $f/1.5$ sphere, whereas Fig. 5(b) shows the result of clipping in exactly the same manner as described in the paragraph above. At first sight, it would appear that a significant contiguous “island” has been removed near the center of the optic, raising the fear that “real” data are being excluded from the range covered by the calculated PVr. Closer examination [Fig. 6(a)] shows that the “real” peak at the center of the part has superimposed upon it at least two on-axis narcissi (ghost interference patterns). Clipping removes some of the peaks of these ghost interference patterns. Thus, the part’s optical performance will be significantly better than would be expected from its PV. Put another way, the optician has expended extra time to make a surface that is actually $\lambda/30$ (<20 nm) simply to appear to be $\lambda/20$ (<30 nm) when measured in the presence of the ghost.




Fig. 6 (a) Close up of the center of the data of Fig. 5 and (b) close up of the center of the data of Fig. 5(b) Truncated pixels are shown here as black.

4.4 Optimized Clipping

Maps showing the data clipped to PVr are helpful in visualizing what data have been excluded by the definition of PVr. Simple, equal bands, as discussed in the preceding section, are easy to implement but may tell an overly pessimistic story, particularly with skewed data.

Figure 7 shows raw data from a “ $\lambda/46$ ” $f/0.75$ spherical surface (PV=13.7 nm) during production at Zygo. PVr for




Fig. 7 $f/0.75$ TS (S/N 109). PV=13.7 nm, PVr=5.4 nm.




Fig. 8 $f/0.75$ TS. Equal clipping (<3900 points). PVr=5.4 nm.




Fig. 9 $f/0.75$ TS. Optimized clipping (<1900 points). PVr=5.4 nm.

this surface is 5.44 nm. Figure 8 shows the surface clipped to PVr with equal clipping bands, a process which truncates nearly 3900 points (0.7%).

Clearly, there is no physical reason that requires equal clipping. Figure 9 shows that optimizing the clipping range (while maintaining the range of PVr) can substantially change the number of points clipped.




Fig. 10 150-mm plano cavity measured with camera mode set at $1\text{ k} \times 1\text{ k}$. PV=85.5 nm, PVr=22.4 nm.




Fig. 11 Same cavity as the data of Fig. 10, measured moments later with the camera mode set to 500×500 . PV=57.7 nm, PVr=21.6 nm.




Fig. 12 Conceptual model of separation of PSD into low-order figure and midspatial frequencies.

Note that, in this case, a visually objectionable ghost does not coincide with either the peak or the valley and hence is not clipped. PVr is not a smoothing filter. Full resolution is maintained. PVr defines an amplitude range that relates to the contribution to imaging system performance of the surface under test.

4.5 Why is PVr Robust?

Table 2 showed the effect of reducing pixel resolution computationally on PV and PVr. Figures 10 and 11 show another example, in this case raw measurement data from a 150-mm aperture flat. Consecutive measurements (within a matter of seconds) using different detector resolutions ($1\text{ k} \times 1\text{ k}$ and 500×500) show a change in PV of nearly 40%, while PVr changes by <4%. Why is this?

The change in PV results from reduced “noise,” specifically from small defects. Removing 16 points (0.002%) from the data of Figure 10 reduces the PV to <57 nm i.e., to less than the PV of Fig. 11. Removing 188 points reduces the data of Fig. 10 to the PVr range (with equal clip). In the absence of standardized filtering or spike removal in Fig. 10, two different competent metrologists could spend significant effort finding “outliers” and declare, in good faith, the part to be anywhere between a $\lambda/10$ and a $\lambda/30$ surface.

Why, by contrast, is PVr robust? One explanation starts from a description of the surface via its power spectral density function (PSD). ISO 10110 notes⁶ that the PSDs of typical optical surfaces have the following form

$$\text{PSD} = \frac{a}{f^B},$$

where B is usually between 1 and 3. This is shown conceptually in Fig. 12, where B is the slope of the fit (dashed) to the PSD.

The 36-term Zernike fit captures the underlying, low-order, optical surface figure,[†] which represents the bulk of the surface amplitude (note that this is a log scale). The Zernike fit (Fringe) represents radial frequencies up to six

[†]The form of the Fringe Zernikes suggests that the cutoff between regimes shown at f_2 is not sharp (particularly in the 2-D representation of Fig. 12); use of terms from both sides of an imprecise cutoff implies, conceptually, no loss of fidelity in representing the functional effect of the spatial frequencies concerned in PVr.

cycles/aperture (r^{12}) and azimuthal frequencies to 5θ it is easy to understand that for even the highest frequencies represented (r^{12}), the fit is robust for ~ 10 points/cycle (64×64) or > 100 points/cycle ($1k \times k$).

The rms of the residual is equally robust. The rms for any bandwidth limited distribution can be computed from the integral of the PSD between the frequencies of interest. Note that changing the detector resolution from $1k \times 1k$ to 500×500 over, say, a 100-mm aperture changes the Nyquist frequency from 5 to 2.5 mm^{-1} , while the low-frequency limit of the integration remains unchanged at $\sim 0.08 \text{ mm}^{-1}$. Given that the axes are logarithmic, it is obvious that the integral between f_2 (in Fig. 12) and f_3 is insignificantly different from the integral between f_2 and f_4 (where f_3 and f_4 represent the resolution limits of different detectors).

Expressed differently, as spatial frequency (f) goes up, the contribution to the surface topography drops exponentially. The rms of the residual is proportional to the sum in quadrature (rss) of these higher frequency components; thus, increasing resolution has an insignificant effect on this sum.

5 PVr Standardization

PVr is a robust amplitude specification that is largely insensitive to interferometer spatial resolution. It is easily computed using functions readily available in commercial interferometer software. For circular apertures, it can easily be integrated into the ISO 10110 drawing formalism; PVr may be defined by default as discussed thus far in this document ($N=3$, circular aperture, 36-term Fringe Zernike polynomial set.) Nonstandard callouts could be covered in a relatively small set of annotations, although it can be argued that allowing such variants will cause more confusion than it is worth.

In addition, the basic definition (Section 4.1), two constraints are appropriate:

1. $PVr < PV$: By definition PVr should be less than PV; pathological cases have been found where large and spatially varying noise gives a large residual rms without affecting the PV.
2. $PVr > 6 \times \sigma_{36\text{ZernikeResid}}$. In the limiting case of extremely good optics (i.e., where the 36-term Zernike fit is essentially zero), then the defined factor of three times the residual will (may) lead to an unacceptably large number of points falling outside the range characterized by PVr.

5.1 Why 3? Why 36?

In general, low-order aberrations tend to dominate the form error of optical surfaces, and the higher spatial frequency errors are typically not correlated (spatially) with the low-order terms. As indicated in Section 5, a surface with no low-order contribution would be well characterized by 6σ . The factor of 3 is shown to give reasonable PVr: rms ratios (see Section 4.2) and acceptable clipping levels for a broad range of optical surfaces.

The choice of the Fringe Zernike set is entirely pragmatic. It appears to be the most widely used set, particularly in commercial interferometer software and lens design

packages. Listings of the Zernike polynomials sometimes are normalized such that the coefficients represent the rms contribution of each term rather than the more common representation where the coefficients are related to the amplitude. Because the Zernikes are used to derive a fit surface, this has no effect on the computation of PVr.

There are a number of different listings of the Zernike polynomials, for example, Mahajan,¹ Wang and Silva,⁷ Kim and Shannon,⁸ and Malacara and DeVore.⁹ Hence, the commonly heard claim of “fitting to the first 36 Zernikes” can be ambiguous. Since the Fringe set appears to be the most widely used set, it has been adopted in this definition of PVr. It is, explicitly, terms Z1–Z35 plus Z48 in ISO/TR 14999-2.¹⁰

It is worth noting, however, that fitting to a different set will likely produce minor differences in PVr. Suppose, for example, that the set sometimes referred to as the “standard” set that includes radial terms to r^8 and azimuthal terms to 7θ is used instead; the surface frequencies represented by the r^{10} and r^{12} terms of the Fringe set will become part of the residual, while the PV contribution of the 6θ and 7θ components will be represented explicitly. The effects of this balancing will depend on the specific surface although, in most cases, the contributions of these higher order Zernike terms are small. Detailed evaluation of this effect is left for future work.

6 Concluding Remarks

PVr is a robust amplitude specification for optical surfaces, which maintains—for surfaces manufactured by a broad range of methods—the implicit linkage to function embodied in manual evaluations of PV of traditionally manufactured surfaces. It is suited to simple specification of mainstream optical surfaces and can reduce disagreements about surface characterizations made on different instruments. PVr is easy to compute with the software in commercially available interferometers; it should be adopted in cases where PV is currently used.

PVr is not a “silver bullet;” it does not capture the detailed aspects of optical function needed when specifying optics for particular applications. Other functional specifications (rms, slope, PSD, structure function, etc.) are likely to have better correlation with specific surface functionality. These more complex descriptors should be used when they are appropriate. It may be useful to supplement PVr with limits on the maximum contiguous area of the surface that fall outside the range that PVr defines or the maximum fractional area.

Acknowledgments

A number of colleagues at Zygo Corp. have made helpful comments and suggestions. Specifically, the author acknowledges the helpful comments of D. Stephenson, D. Sykora, and J. Soobitsky, and the experience acquired by A. Estrin’s extensive use of PVr in Zygo’s optical shop.

References

1. V. N. Mahajan, “Zernike circle polynomials and optical aberrations of systems with circular pupils” OPN, Engineering and Lab. Notes, pp21–23, (August 1994).
2. M. Born and E. Wolf, *Principles of Optics*, 7th ed., pp. 905–910, Cambridge University Press (1999).
3. C. J. Evans, “Uncertainty evaluation for measurements of peak-to-

- valley surface form errors,” *CIRP Ann.* **57** (1) 509–512 (2008).
4. J. S. Loomis, “Fringe users manual,” in D. Anderson, *Fringe Manual Version 3*, University of Arizona (1982).
 5. C. J. Evans, “Robust estimation of PV for optical surface specification and testing,” OSA, *Opt. Fabricat. Test. Workshop Digest*, (October 2008).
 6. ISO “*Optics and photonics—Preparation of drawings for optical elements and systems—Part 5: Surface form tolerances*,” ISO 10110 Part 5, 2nd ed. (2007).
 7. J. Y. Wang and D. E. Silva, “Wavefront interpretation with Zernike polynomials,” *Appl. Opt.* **19**, 1510–1518 (1980).
 8. C.-J. Kim, and R. R. Shannon, “Catalog of Zernike polynomials,” in *Applied Optics and Optical Engineering*, Vol. **10**, Academic Press, New York (1992).
 9. D. Malacara and S. L. DeVore, “Interferogram evaluation and wavefront fitting,” in *Optical Shop Testing*, 2nd ed. (ed D. Malacara), Wiley, Hoboken, NJ (1992).
 10. ISO, “*Optics and Photonics—Interferometric measurement of optical elements and optical systems—Part 2: Measurement and evaluation techniques*,” ISO/TR 14999-2, p. 45 (2005).

Chris Evans is Chief Metrologist and Senior Research Scientist at Zygo Corp. Prior to joining Zygo, he spent over 16 years working at the National Institute of Standards and Technology. His research interests include interferometric testing of optics, optical fabrication and precision engineering.