Modeling interferometers with lens design software

Bryan D. Stone, MEMBER SPIE
Tropel Corporation
60 O'Connor Road
Fairport, New York 14450

Abstract. Lens design software has been designed primarily to model conventional imaging systems. While interferometers do not generally fall into this category, lens design software nonetheless is well suited for analyzing a variety of aspects of interferometric systems. The general requirements for a ray-based model of interferometers are discussed, and a variety of examples are presented. The examples are designed to demonstrate both the power and flexibility of the approach proposed here for modeling interferometers. [S0091-3286(00)01307-6]

Subject terms: lens design; interferometers; imaging systems.

Paper ENV-13 received Oct. 26, 1999; revised manuscript received Feb. 7, 2000; accepted for publication Feb. 7, 2000.

1 Introduction

Geometrical optics has been used successfully to design conventional imaging systems for hundreds of years. However, simple spot diagrams are insufficient for modeling systems whose aberrations are small—the wave nature of light must be taken into account to accurately predict the behavior of systems as they approach the diffraction limit. A simple ray-based analysis does a good job of predicting the shape of wavefronts away from caustics, but the model breaks down near caustics. Since in an imaging system there is a caustic at the image plane, it is common to trace rays only to estimate the shape of the wavefront in the exit pupil. If quantities like the point spread function are desired, this wavefront is then propagated according to the laws of wave optics to the image plane.

Contrast this with interferometers. By their very nature, they rely explicitly on the wave nature of light for their operation. However, in interferometers, caustics at the detector are typically avoided. Therefore, for any given wavefront that is input into the interferometer, geometrical optics generally does a good job of predicting the output wavefront. And it is the path difference between wavefronts that interferometers measure.

Despite the obvious utility of the relatively simple tools of geometrical optics for modeling interferometers, there are studies in which the propagation of fields is considered.[1,2] However, the bulk of the published work on modeling interferometers centers on the geometric wavefront. Huang discusses propagation errors within an algebraic framework for Fizeau interferometry.[3] A number of papers describing numerical modeling of individual interferometers have been published. For example, Michaloski et al. develop a method for modeling interferometers and apply it to a grazing-incidence interferometer for testing cylindrical parts.[4] Lowman and Greivenkamp[5,6] consider a Twyman-Green interferometer for nonnull testing of aspheres. In another paper, misalignments of optical assemblies in an interferometer for astrometric measurement are studied.[7] All of these papers use ray tracing to estimate geometric wavefronts. An alternative approach that involves tracing a large number of randomly chosen rays through an interferometer has also been described.[8]

A general, geometric framework for modeling the optical performance of interferometers is described in this paper. This framework allows for modeling a variety of aspects of interferometers. These include (but are not limited to) (i) the sensitivity of an interferometer to manufacturing errors in its components; (ii) the effects associated with aberrated input wavefronts; (iii) the reduction in visibility due to an extended source, a polychromatic source, or both; (iv) determining the range of amplitudes and spatial frequencies over which an interferometer can make accurate measurements. This framework is applied to the testing of optical surfaces (interferometers that measure rough parts, for example, are not modeled well with this approach*). Note that not all aspects of interferometer performance can be modeled with this simple geometric framework. For example, if the viewing optics in an interferometer do not image a part under test, say, onto the detector, then *defocus fringes* are observed at the edge of the part. In this case, diffraction effects must be considered to predict the interference pattern seen at the detector.

The basic premise for the proposed approach is to model interferometers in the manner in which they are used. This is discussed further in Sec. 2. Sections 3, 4, and 5 primarily consist of a variety of examples showing the ways in which this basic approach can be applied. The example of Sec. 3 involves a Twyman-Green interferometer. The measurement error associated with an aberrated input beam is modeled, and a simulation of the measurement of a parabolic surface is described. This same interferometer appears in examples presented in Sec. 4 and 5. In Sec. 4, a discussion of modeling extended sources is given, along with examples for a Michelson and for the Twyman-Green interferometer. In Sec. 5, a method to map out the useful mea-

**Rough parts* are those with features on the surface with dimensions on the order of the wavelength.

 0091-3286/2000/$15.00

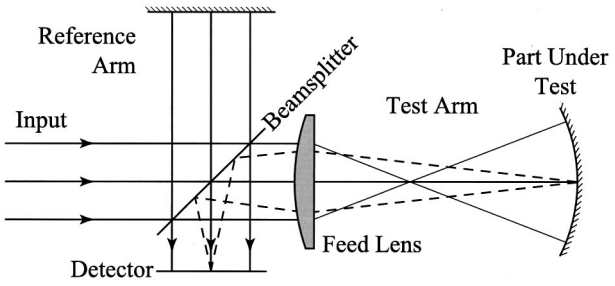

Fig. 1 A schematic of a simple Twyman-Green interferometer.

surement range of an interferometer is described [this covers item (iv) in the above list]. Concluding remarks are offered in Sec. 6.

2 Basic Approach

2.1 *General Description*

Consider the Tyman-Green interferometer shown schematically in Fig. 1. An input beam is incident on a beamsplitter, which sends the light into the two arms of the interferometer. The reference arm is shown with a plane return mirror. The test arm has a feed lens to generate spherical wavefronts. The detector is placed conjugate to the part under test (viewing optics can be added between the beamsplitter and the detector to image the test part onto the detector). If a plane wave enters the system and all elements are ideal (both in design and manufacture), then a plane wave exits each arm of the interferometer. In practice, however, nothing is perfect, and there are a variety of effects that one might want to model: aberrated input wavefronts, extended sources, polychromatic sources, aberrations inherent in the design of the feed or viewing optics, and manufacturing errors or misalignments of elements in the interferometer.

By developing an approach that matches, as closely as possible, the conditions under which the interferometer is used, it is possible to model all these effects. To understand this approach, consider Fig. 2, where the two arms of the interferometer are shown separately. A wavefront (not necessarily plane) is input into the system, and a particular point on the detector is labeled *A*. Now consider a ray through the test arm that (i) is normal to the wavefront and (ii) passes through the point *A* on the detector. This ray intersects the wavefront at the point labeled *B*. A similar ray through the reference arm (that passes through *A* and is normal to the input wavefront) also is shown. The point where this ray intersects the wavefront is labeled *C*. At the point *A*, the interferometer measures the optical path difference (OPD) between the rays *BA* and *CA*. The OPD is obtained by measuring the relative phase between the two interfering wavefronts at *A* and then unwrapping that phase. There are a variety of methods for measuring absolute phase and performing the phase unwrapping. While the choices made for these methods affect the interferometer's performance, the focus of this paper is on modeling the effects of the optics on interferometer performance. Therefore, the assumption is made that phase detection and unwrapping are performed perfectly.

Based on the above discussion, if lens design software is used to model the interferometer, the software must be able

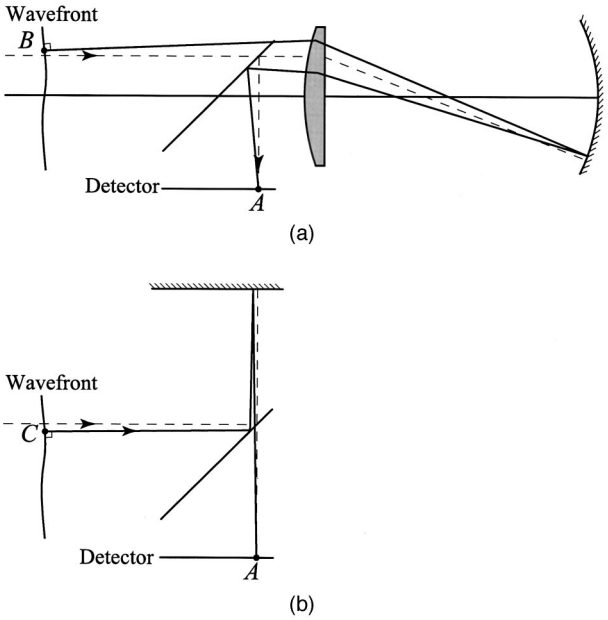

Fig. 2 (a) The test arm and (b) the reference arm in a Twyman-Green interferometer. A ray through each arm of the interferometer is shown that (i) passes through a given point on the detector (labeled *A*) and (ii) is normal to the input wavefront.

to determine the ray that is normal to any given input wavefront and passes through any given point on the detector. This process I call *ray aiming*, because it is related to the ray aiming that commonly takes place in lens design software when determining the initial direction of a ray, say, that starts at a given object point and passes through a given point in the exit pupil. Proper ray aiming is critical for the ultimate success of the approach taken here, and some considerations for performing such ray aiming are discussed in the next subsection.

Before presenting that discussion, a few thoughts are given on the advantages of this procedure. Since it is the interference between wavefronts that gives information about the OPD between arms of an interferometer, one could argue that the step of ray aiming is not essential. Instead, rays normal to the input wavefront can be launched from, say, points on a regular grid across the wavefront. Once these rays are traced to the detector, the wavefronts for the different arms of the interferometer can be fitted and then subtracted. However, this procedure is not always as clean as it sounds. Some of the issues (which are discussed first) relate to the operation of current lens design software, but others are more fundamental. If one wants to use built-in features of lens design software, this typically limits the input wavefronts to planes and spheres. In fitting wavefronts, Zernike polynomials are typically used. However, if the part does not have a circular clear aperture, Zernike polynomials are not the best choice of fitting functions: a set of functions that are orthogonal over the aperture shape is required. Further, mapping errors between the wavefronts from different arms have to be allowed for separately. In general, fitting the output wavefronts is a process fraught with pitfalls, and the user must always be on the lookout for poor fits.

The problems listed above could be considered mere inconveniences—they can all be overcome in one way or another. However, there is a more fundamental reason to avoid the above procedure: when modeling the limits to the performance of an interferometer (as is done in Sec. 5), fitting is inadequate. For example, the ray aiming that is used here sometimes fails. This can occur for large aspheric departures with a nonnull test, for example, or when high-spatial-frequency errors are placed on parts, as is done in Sec. 5. However, if a robust ray-aiming procedure is used, these failures provide useful information. Failures of the ray-aiming algorithm signal that a caustic of the associated wavefront has moved onto the detector. When this occurs, the measurement that is made is not simply related to the surface figure of the part under test. Any robust model of interferometer performance should be vigilant for such occurrences and alert the user to them. If a grid of rays through the input wavefront is sent into the system and the optical path of the output rays is used to fit a wavefront, the result is meaningless once a caustic has moved onto the detector. However, one might never know that this has occurred.

The only drawback of the method proposed here is that it is iterative, and therefore has the potential to be slow. However, such an iterative procedure is not fundamentally different (and no more costly) from determining the initial direction of a ray on the object that passes through a given point on the exit pupil. Computers are fast, and tracing the extra rays is generally of little consequence these days. The advantage is that this method provides a general framework for analyzing a variety of interferometers. By incorporating the ray aiming with conventional lens design software, any component that can be included in a conventional imaging system can be included in an interferometer model. This includes gratings, diffractive elements, gradient-index elements, aspheres, etc. In the next subsection, some thoughts on ray aiming are given.

2.2 *Considerations for Ray Aiming*

There are two properties that are important for any ray-aiming procedure: it must be robust, and it must converge quickly. In assessing the tradeoffs between these two requirements, it is generally desirable to sacrifice speed for the sake of robustness. For quick convergence, I use the Newton-Raphson method (see, e.g., Ref. 9), which is described briefly as follows. Place a Cartesian coordinate system at the detector and one at the input wavefront. Quantities associated with the detector's coordinate system are distinguished by the addition of a prime, and the axes are chosen such that the X' and Y' coordinates of a point correspond to transverse positions on the detector. The coordinate system associated with the input wavefront is similarly aligned so that the X and Y coordinates uniquely identify a point on the wavefront. Given a point of interest on the detector with coordinates $(\overline{x}',\overline{y}')$, the goal is to determine the position on the input wavefront for the ray that (i) starts normal to the wavefront and (ii) passes through the point of interest on the detector. Given some guess for this initial position [call the coordinates of this guess (x_0,y_0)], the coordinates of the ray from this point at the detector can be determined [call them (x_0',y_0')]. Then according to the Newton-Raphson method, an ''improved estimate'' for the desired position on the input wavefront, (x_1,y_1) say, is given by

$$\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} = \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} + \begin{bmatrix} \partial x_0'/\partial x_0 & \partial x_0'/\partial y_0 \\ \partial y_0'/\partial x_0 & \partial y_0'/\partial y_0 \end{bmatrix}^{-1} \begin{pmatrix} \overline{x}' - x_0' \\ \overline{y}' - y_0' \end{pmatrix}, \tag{1}$$

where the derivatives represent the changes in ray position at the detector with respect to changes in ray position on the wavefront. (These derivatives must take into account the change in initial ray direction for non-plane-wave inputs.)

This method, on its own, fails the robustness requirement (see, e.g., Ref. 10). If, for example, the initial guess is not close enough to the actual point on the wavefront, the ''improved estimate'' may actually be worse than the original estimate. To allow for such a possibility, the method can be supplemented with a check of whether the step specified by Eq. (1) (in the second term on the right) gives a point further from the desired point on the detector than the initial guess did. If it does, a smaller step can be searched for that gives an output ray closer to the specified point on the detector. (The method I use for this is similar to the one described in Sec. 9.7 of the book cited in Ref. 9.)

Another way this method can fail is that the output points can oscillate around the desired point, converging only slowly. To fix this problem, an estimate of the rate of convergence can be determined from one iteration to the next. If the convergence after a few iterations is too slow, a smaller step that gives an output ray that is closer to the specified point on the detector can be searched for.

Finally, a means to determine the initial estimate is needed. Since the coordinates of points at the detector tend to be approximately linearly related to the coordinates of points at the input, taking the initial ray to start at the axial point on the wavefront works well. However, in cases of unconventional geometries,[†] more specialized methods for finding initial estimates may be needed.

The process described above is very robust, and I have found that it fails only in cases where there is a caustic at the detector.

3 A Basic Twyman-Green Interferometer

3.1 *Example 1: Measuring a Sphere*

For this example, a Twyman-Green interferometer used to test a spherical part is modeled. The interferometer is illustrated schematically in Fig. 3. The input beam is incident on a cube beamsplitter. A feed lens and the part under test constitute the test arm, and a plane mirror the reference arm. Viewing optics (in the form of an afocal pair) have been placed between the beamsplitter and the detector. This pair images the part under test onto the detector. The prescriptions for the feed lens and viewing optics are given in an appendix. The part under test is taken to have a 100-mm radius of curvature and a 50-mm clear aperture. The center

[†]An embodiment of an interferometer that I consider possesses unconventional geometry is the CylinderMaster™. This interferometer is designed to measure cylindrical parts at near-grazing incidence. CylinderMaster is a trademark of Tropel Corporation, Fairport, NY.

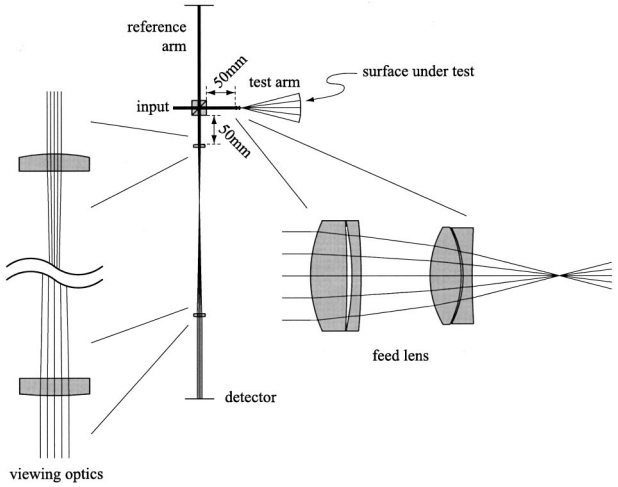

Fig. 3 Illustration of the Twyman-Green interferometer appearing in examples in Secs. 3, 4, and 5.

of curvature of the part is placed at the rear focal point of the feed lens. (The separation from the last surface of the feed to the rear focal point is given in the last row of Table 1 in the appendix.) The input beam diameter that fills the clear aperture of the part is 4.0 mm. The measurement wavelength in all examples presented here is that of a HeNe laser (0.6328 μm).

For the first example, consider the measurement of a perfect spherical part. The points of interest are equally spaced along a line through the center of the detector. For each point, the initial position of two rays (one through the test arm and the other through the reference) are determined. These rays (i) pass through the point on the detector and (ii) are normal to the (in this case plane) input wavefront. Note that because of the nature of interferometers, any constant can be added to the OPD between the two arms. This constant is chosen such that the measured OPD between the rays at the center of the part under test is zero. The resultant OPD is plotted in Fig. 4. The measured OPD is simply twice the wavefront error of the feed lens (twice because the feed lens is double-passed). When high accuracy is desired, the OPD shown in Fig. 4 can be subtracted from any measurements (i.e., it can be calibrated out).

Now consider adding half a wave of third-order spherical aberration to the input beam. The measurement that the interferometer gives for various separations between the beamsplitter cube endface and the reference mirror is determined according to the model, and the measured OPD

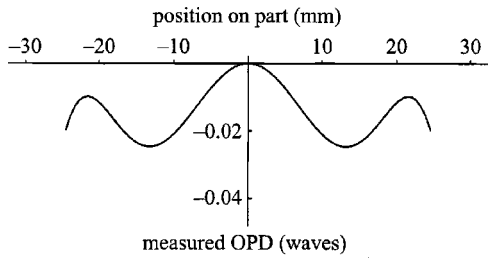

Fig. 4 Measured OPD (in waves) as a function of position on the part for the Twyman-Green interferometer illustrated in Fig. 3.

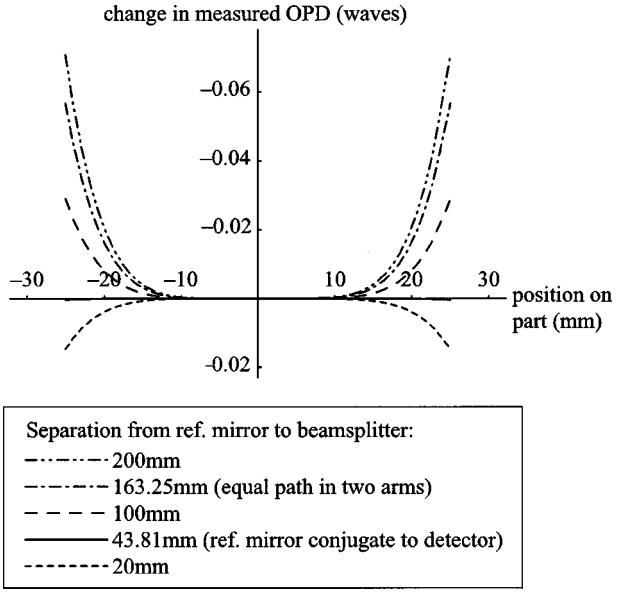

Fig. 5 This figure shows the change in measured OPD when the input beam possesses 0.5 waves of third-order spherical aberration. The different curves are for different positions of the reference mirror.

for a plane wave input is subtracted from this. The result is plotted in Fig. 5 for five different locations of the reference mirror. When the reference mirror is placed 43.81 mm from the beamsplitter, it is imaged by the viewing optics onto the detector. Since the surface under test is also conjugate to the detector, the effect of an aberrated input beam is negligible for this position of the reference mirror. However, the propagation effects as the reference mirror is moved away from this position become noticeable. At a distance of 163.25 mm, the two arms have approximately equal optical paths along the axis (this condition needs to be satisfied, for example, when a temporally incoherent source is used). In this case, the difference between the measured OPD and the OPD found with a plane-wave input is about 0.05 waves at the edge of the part.

3.2 *Example 2: Measuring a Paraboloid*

The Twyman-Green interferometer is now used to measure a paraboloid with the same base radius of curvature as the sphere from the previous subsection. [A paraboloid with a clear aperture of 50 mm and a base radius of 100 mm has about 0.05 mm (or 80 waves at HeNe) of departure from the sphere with the same base radius.] The input wavefront is taken to be plane. In this case, I focus solely on the propagation errors in the system, and therefore again subtract the OPD due to the feed lens from the OPD of the measurement. Say the part is positioned precisely 100 mm from the focus of the feed. Then as one moves out on the part, the test gets further from null, and the measurement error increases. This is shown in Fig. 6 as the curve labeled 0.0 mm. The other curves in Fig. 6 represent the measurement error as the part is shifted axially away from the focus of the feed lens. The label of each curve shows how far the base center of curvature of the part has moved from the focus of the feed lens. When the part is shifted, it can be measured well in two zones: one near the axis and another

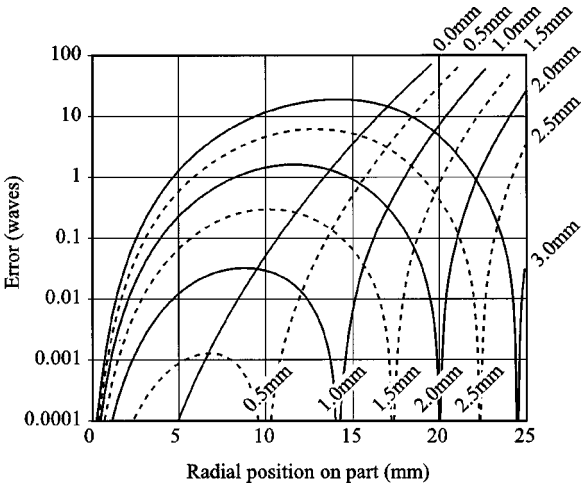

Fig. 6 A logarithmic plot of the measurement error as a function of radial position on the part under test. The different curves are for various positions of the part. The label associated with each curve shows the amount the center of curvature of the part is shifted from the focus of the feed lens.

near a null of constant radius. As the part moves away from the feed, the radius of the second null zone moves out on the part, and the regions (both near the axis and in the second null zone) over which the interferometer gives good data get narrower. This model can determine the size and number of zones required for accurate subaperture stitching. Of course, the fringe densities (which get higher away from the null) also have to be taken into account for detectors with discrete pixels.

4 Extended Sources

4.1 *Discussion*

The approach taken here allows for straightforward modeling of the effects of extended sources. Extended sources tend to decrease fringe visibility. To model this effect, an extended source is approximated by a series of discrete sources that add incoherently. Each discrete source generates an independent input wavefront. Consider the irradiance I at a point on the detector resulting from these discrete wavefronts. In a two-beam interferometer (such as a Fizeau or Twyman-Green), the irradiance is given by

$$I(\varphi_z)=\sum_{i=1}^{N} w_i\left[1+v_i\cos\left(\frac{2\pi}{\lambda}\delta_i+\varphi_z\right)\right], \tag{2}$$

where N is the number of discrete sources, δ_i is the OPD (associated with the point on the detector) for wavefront i, v_i is the visibility associated with source i alone, w_i is a weighting for the different input wavefronts, and φ_z is a phase term. As φ_z is varied from 0 to 2π, the irradiance goes through its maximum and minimum values. Physically, this phase shift can be thought of as resulting from an axial shift of a reference mirror (in a Twyman-Green interferometer, say) or of the part (in a Fizeau interferometer, say). Each v_i must be between zero and one (the v_i are less than unity, for example, when different amounts of energy exit the source and reference arms).

As an aside, there are a variety of reasons for the w_i not to be equal. For example, the w_i can be adjusted for an uneven energy distribution across the source (so that each input wavefront gets a different weighting according to its energy). As another example, Eq. (2) represents a discrete approximation to a continuous integral. Different schemes for numerical integration, such as Gaussian quadrature,[‡] require different weightings to achieve the highest accuracy from the numerical approximation. Therefore, even for a uniform source, the w_i are not necessarily equal.

Equation (2) is now put into a form that makes it obvious how the framework introduced earlier is used to model extended sources. Begin by rewriting it as

$$I(\varphi_z)=\sum_{i=1}^{N} w_i+\cos\varphi_z\left[\sum_{i=1}^{N} w_i v_i\cos\left(\frac{2\pi}{\lambda}\delta_i\right)\right]-\sin\varphi_z\left[\sum_{i=1}^{N} w_i v_i\sin\left(\frac{2\pi}{\lambda}\delta_i\right)\right]. \tag{3}$$

The visibility (denoted by V) is defined as

$$V=\frac{I_{\max}-I_{\min}}{I_{\max}+I_{\min}}, \tag{4}$$

where $I_{\max}$ is the maximum value of the irradiance as φ_z varies, and $I_{\min}$ the minimum value. Let $\varphi_{z,\max}$ and $\varphi_{z,\min}$ be the values of φ_z that maximize and minimize I, respectively. These are found by differentiating Eq. (3) with respect to φ_z and setting the result equal to zero:

$$\varphi_{z,\max}=\tan^{-1}\left[\frac{\sum_{i=1}^{N} w_i v_i\sin\left(\frac{2\pi}{\lambda}\delta_i\right)}{\sum_{i=1}^{N} w_i v_i\cos\left(\frac{2\pi}{\lambda}\delta_i\right)}\right], \tag{5}$$

$$\varphi_{z,\min}=\varphi_{z,\min}+\pi. \tag{6}$$

Equations (3), (4), (5), and (6) can be combined to give the following expression for the visibility:

$$V=\frac{\left\{\left[\sum_{i=1}^{N} w_i v_i\sin\left(\frac{2\pi}{\lambda}\delta_i\right)\right]^2+\left[\sum_{i=1}^{N} w_i v_i\cos\left(\frac{2\pi}{\lambda}\delta_i\right)\right]^2\right\}^{1/2}}{\sum_{i=1}^{N} w_i v_i}. \tag{7}$$

While Eq. (7) can be manipulated further,[§] this form is ideal for numerical implementation. For each point on the source (i.e., for each input wavefront), the OPD between

[‡]For example, the quadrature schemes discussed in Ref. 11 can be used to obtain the most accuracy with the fewest input wavefronts.

[§]For example, Eq. (7) is equivalent to the following equation:

$$V=\frac{\left\{\left[\sum_{i=1}^{N}(w_i v_i)^2\right]+2\sum_{i=1}^{N}\sum_{j=i+1}^{N} w_i w_j v_i v_j\cos\left[\frac{2\pi}{\lambda}(\delta_i-\delta_j)\right]\right\}^{1/2}}{\sum_{i=1}^{N} w_i v_i}.$$

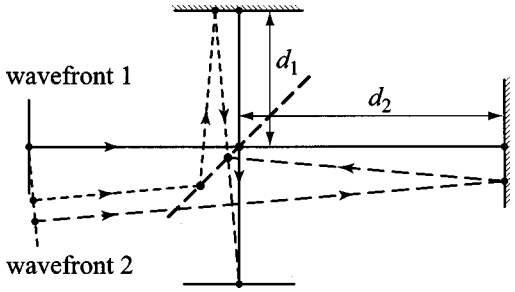

Fig. 7 A schematic illustration of a Michelson interferometer with two plane-wave inputs.

the arms of the interferometer is determined (i.e., the δ_i are found). Running totals are kept of the individual sums appearing in Eq. (7). Once the sums have been computed, it is straightforward to determine the visibility according to Eq. (7).

4.2 *Example 1: Michelson Interferometer*

As an example, consider a Michelson interferometer. As illustrated in Fig. 7, the two arms of the interferometer are taken to have different lengths. The source for this example has a uniform energy distribution, and the visibility at a single point on the detector is modeled. Since there are two arms to consider for each input wavefront, there are four rays of interest for the two wavefronts shown in Fig. 7. These rays are shown individually in Fig. 8. The OPD for wavefront 1 (call this δ_1) is the difference between paths *BCK* [Fig. 8(a)] and *BDK* [Fig. 8(b)]. For wavefront 2, δ_2 is the difference between the paths *EFGK* [Fig. 8(c)] and *HIJK* [Fig. 8(d)]. The values of the OPD for any other wavefronts of interest can be found similarly. These are used in Eq. (7) to determine the fringe visibility.

Before proceeding with the results of this example, the exact visibility is discussed for this case. This allows a comparison between the model and exact results. Consider a plane wavefront propagating roughly in the *Z* direction. Say the *x* and *y* direction cosines of the direction of propagation are α and β (which are assumed small). If this wavefront is input into a Michelson interferometer whose (unfolded) mirrors are precisely normal to the *Z* axis, the path difference between the two arms is given by

$$\delta(\alpha,\beta)=\frac{2\,\Delta d}{(1-\alpha^2-\beta^2)^{1/2}}\approx 2\,\Delta d+\Delta d(\alpha^2+\beta^2), \tag{8}$$

where Δd is defined as

$$\Delta d := d_1 - d_2 . \tag{9}$$

The second (approximate) equality of Eq. (8) holds for small values of α and β. For a continuous source and path differences given according to Eq. (8), the discrete sum of Eq. (2) generalizes to an integral:

$$I(\varphi_z)=\frac{\displaystyle\int_{\theta=0}^{2\pi}\int_{\rho=0}^{\rho_0} w(\rho,\theta)\left[1+v(\rho,\theta)\cos\left(\frac{2\pi}{\lambda}\Delta d\,\rho^2+\varphi_z\right)\right]\rho\,\mathrm{d}\rho\,\mathrm{d}\theta}{\displaystyle\int_{\theta=0}^{2\pi}\int_{\rho=0}^{\rho_0}\rho\,\mathrm{d}\rho\,\mathrm{d}\theta}, \tag{10}$$

where ρ and θ are defined as

$$\rho := (\alpha^2+\beta^2)^{1/2}, \tag{11}$$

$$\theta := \tan^{-1}(\beta/\alpha), \tag{12}$$

and ρ_0 is the sine of the half angular extent of the source. In Eq. (10), $w(\rho,\theta)$ and $v(\rho,\theta)$ are the continuous analogs of the w_i and v_i appearing in Eq. (2). For this example the source is uniform and temporally coherent, so that $w(\rho,\theta)=v(\rho,\theta)=1$. Equation (10) can be evaluated in closed form and the resulting maximum and minimum values for $I(\varphi_z)$ substituted into Eq. (4):

$$V=\left|\frac{\sin(\pi\,\Delta d\,\rho_0^2/\lambda)}{\pi\,\Delta d\,\rho_0^2/\lambda}\right|. \tag{13}$$

This represents the exact result with which the numerical estimates for this example are compared.

The set of discrete wavefronts used in this example to approximate the continuous source are now described. The *Z* axis of a Cartesian coordinate system is aligned with the direction of propagation of the central wavefront in the source (this central wavefront is parallel to the unfolded mirrors of the interferometer). The source is taken to be a series of plane waves arranged on a square grid in optical-direction-cosine space. Only wavefronts whose normals satisfy $(\alpha^2+\beta^2)^{1/2}<\rho_0$ are considered (where α, β are the *X* and *Y* direction cosines of the normal, and ρ_0 is the half angular extent of the source). For this example, Δd is taken to be 30 mm, and in Fig. 9, the fringe visibility is plotted as ρ_0 varies. This is done for grids with 11 points, 21 points, and 41 points across the center of the source. The exact visibility found according to Eq. (13) is also shown in Fig. 9. Note that using 11 points gives good agreement up to the first zero of the visibility. Beyond that, it does not work very well. This is because the phase difference between successive points on the source (i.e., the difference between $[2\pi(\delta_i-\delta_j)/\lambda]$ for *i* and *j* that correspond to neighboring

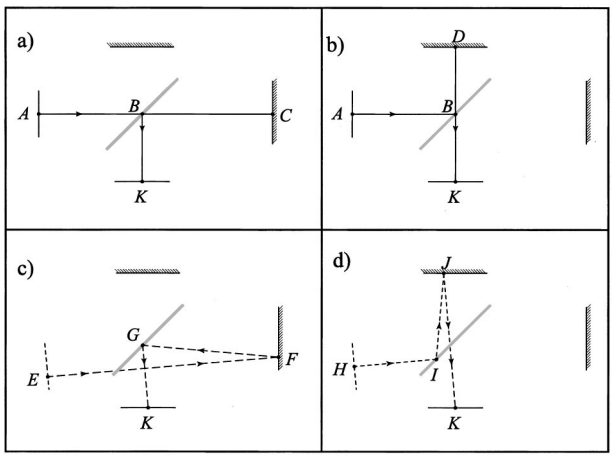

Fig. 8 Paths for the four rays from the two wavefronts that pass through the image point *K*.

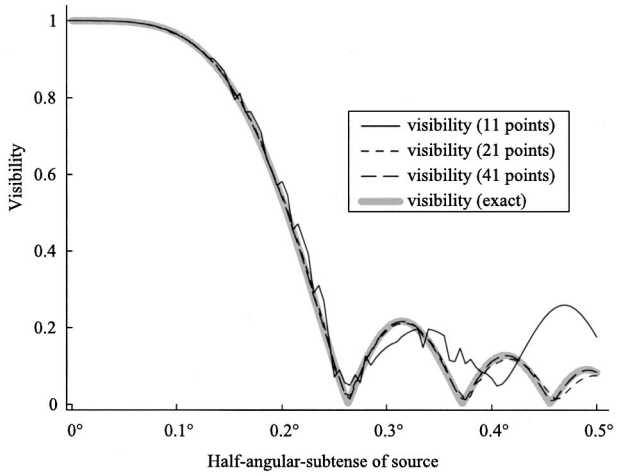

Fig. 9 The fringe visibility as a function of the angular extent of the source for three different densities of sampling the extended source. The exact visibility is also given.

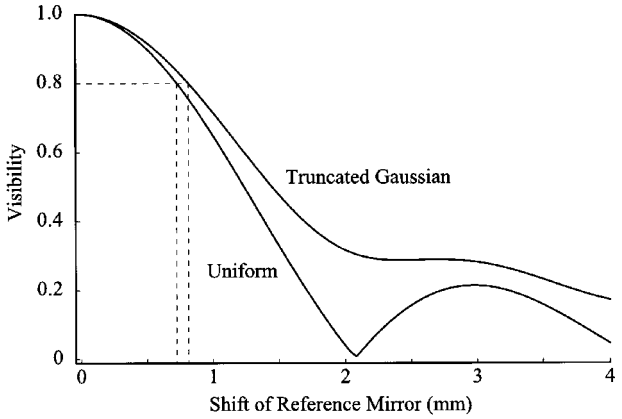

Fig. 10 Fringe visibility as a function of the shift of the reference mirror for two different energy distributions. A shift of zero represents the location where the reference mirror is conjugate to the detector.

wavefronts) is not small compared with 2π. Going to a denser grid alleviates this problem: the curve associated with 41 points across the source is indistinguishable from the exact curve.

4.3 *Example 2: Twyman-Green Interferometer*

As a second example, consider the Twyman-Green interferometer introduced in Sec. 3.1. Recall that when both the reference mirror and the part under test are conjugate to the detector, the measurement error that results from a nonplane-wave input is negligible. A similar situation occurs with an extended source. Namely, the fringe visibility in a Twyman-Green interferometer is unity when the reference mirror and part are both conjugate to the detector. This is analogous to the Michelson interferometer when the optical paths in the two arms are the same [i.e., Eq. (13) shows that $V=1$ when $\Delta d=0$ regardless of the size of the source]. Suppose an extended source is used with the Twyman-Green interferometer. The distance the reference mirror can move from its conjugate position while maintaining good contrast fringes is now investigated. Two energy distributions for the source are considered: a uniform distribution with a 1-deg half angle, and a Gaussian distribution. The Gaussian is truncated at the $1/e^2$ point, which also corresponds to 1 deg for this example. The resulting fringe visibility as the reference mirror shifts from its conjugate position is plotted in Fig. 10 for the axial point at the detector. (In cases where one suspects the fringe visibility might vary, it is straightforward to repeat the calculation for a variety of points across the detector.) This plot indicates that good contrast fringes (with visibility of greater than 0.8, say) result when the reference mirror is placed no more than 0.7 to 0.8 mm from the position where it is conjugate to the detector.

4.4 *Discussion*

I emphasize the importance of the ray aiming used here for generating the numerical results in this section. If some form of random ray tracing were used, such as that proposed in Ref. 8, a very large number of rays would have to be traced to achieve reasonable accuracy. Similarly, launching enough rays from each input wavefront of interest to fit the output wavefronts at the detector also requires many more rays than are necessary simply to estimate the fringe visibility from an extended source.

While only plane-wave inputs are considered in the examples of this section, it is straightforward to add aberrations to the input wavefronts. Also, the measurement error associated with an extended source can be modeled. For this, the average (weighted by the w_i) of the OPDs for the rays coming from the different wavefronts is determined. From these average OPDs for a variety of object points, plots similar to those shown in Fig. 4 and Fig. 5 can be generated for an extended source.

Finally, note that polychromatic sources can be modeled similarly to extended sources. Start with an equation similar to Eq. (2), but now, instead of summing over a set of discrete input wavefronts, sum over a set of discrete wavelengths. In the most general case of an extended source with polychromatic light, the analog of Eq. (7) can be written as

$$V=\frac{\left\{\left[\sum_{j=1}^{M}\sum_{i=1}^{N} w_{ij}v_{ij}\sin\left(\frac{2\pi}{\lambda_j}\delta_{ij}\right)\right]^2+\left[\sum_{j=1}^{M}\sum_{i=1}^{N} w_{ij}v_{ij}\cos\left(\frac{2\pi}{\lambda_j}\delta_{ij}\right)\right]^2\right\}^{1/2}}{\sum_{j=1}^{M}\sum_{i=1}^{N} w_{ij}v_{ij}}, \tag{14}$$

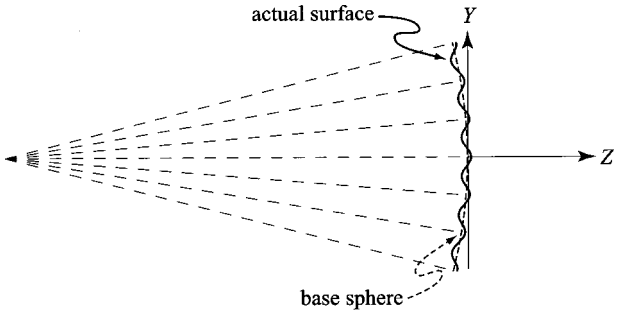

Fig. 11 An illustration of the types of parts used to test the useful measurement range of an interferometer.

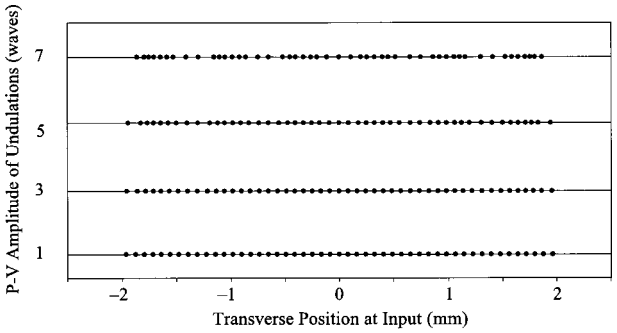

Fig. 12 The locations of points at the input that correspond to equally spaced points at the detector for a test part with undulations of 1, 3, 5, and 7 waves *P-V*.

where M is the total number of discrete wavelengths considered, the λ_j $(j=1,2,3,...,M)$ represent the set of discrete wavelengths, the w_{ij} and v_{ij} are the analogs of the w_i and v_i appearing in Eq. (2), and δ_{ij} is the OPD associated with the detector point for wavefront i and wavelength j.

5 Mapping the Useful Measurement Range of an Interferometer

5.1 *Statement of Problem*

For some applications, one might want to measure more than just low-spatial-frequency surface figure errors. For example, there are manufacturing processes that introduce mid-spatial-frequency errors into a surface. When such processes are used, it is logical to determine the range of spatial frequencies and amplitudes that can be measured effectively by the interferometer used to test the part. To investigate this question, various sinusoidal undulations are placed on a part. By varying the amplitude and spatial frequency of the undulation and looking at the measurement error, one can determine the range of amplitudes and spatial frequencies that can be measured. The interferometer considered in this example is the Twyman-Green introduced in Sec. 3.1.

Interferometers measure the height of a surface not in a direction parallel to the axis, but in a direction normal to a (typically) spherical wavefront. Therefore, it is reasonable to add a radial sinusoid** to the part under test. For this, consider a base sphere with the Z axis of a Cartesian coordinate system aligned with the axis of the interferometer (the coordinate origin is taken to lie on the sphere). The equation for the test part used in this example is of the form

$$z=\frac{c(x^2+y^2)}{1+[1-c^2(x^2+y^2)]^{1/2}}+\frac{1}{2}a\frac{\cos(\eta\sin^{-1}|cy|)}{(1-c^2y^2)^{1/2}}. \qquad (15)$$

The first term represents the base sphere, and the second term the sinusoidal undulation. Here c is the curvature of the part, a is the peak-to-valley (P-V) amplitude of the undulation, and η controls the number of undulations across the part (a higher value of η gives more undulations). With this form of error, arc lengths along the base sphere between successive minima of the undulation are all equal. This is illustrated in Fig. 11, where the base sphere is shown along with the perturbed surface. The amplitude is not precisely that required for a pure radial sinusoid, but the difference between Eq. (15) and a pure radial sinusoid is of order a^2, so that for small-amplitude undulations, this gives a good approximation to a radial sinusoid. Note that the form of Eq. (15) gives undulations in the Y direction. This suffices for a rotationally symmetric interferometer, but if there are any asymmetries (either due to misaligned elements or from asymmetries in the design), Eq. (15) can be modified trivially to give a sinusoid in any direction across the part.

5.2 *Discussion of Ray-Aiming Failures—Caustics at the Detector*

Consider a set of undulations with P-V amplitudes ranging from just under 0.005 waves to 10 waves, and for each amplitude, start with one undulation and increase the number of undulations until ray-aiming failures occur. This signals that a caustic has moved onto the detector. To support this claim, considered an array of equally spaced points at the detector. The initial positions of the rays that (i) are normal to the input wavefront (which is taken to be plane) and (ii) pass through this array of points at the detector are determined. The results are presented in Fig. 12 for parts with six undulations. As the P-V amplitude of the undulations increases from 1 wave to 7 waves (in steps of 2 waves), Fig. 12 shows the initial positions for the rays that satisfy (i) and (ii) listed above. For the smaller amplitude undulations, an equally spaced array on the detector is mapped to a (nearly) equally spaced array at the input. For 5 waves P-V, the rays are seen to be bunched together in some regions and spread apart in others, and by 7 waves P-V this effect is quite pronounced. If the amplitude of the undulations becomes just a bit larger, then the caustic moves onto the detector and there are some points on the detector for which there are no initial rays, and other points

**A radial sinusoid is defined here as follows: Consider some arc along a base sphere. The difference between the base sphere and a sphere with a radial sinusoid is precisely sinusoidal when measured normal to the base sphere and taken as a function of the length along the arc.

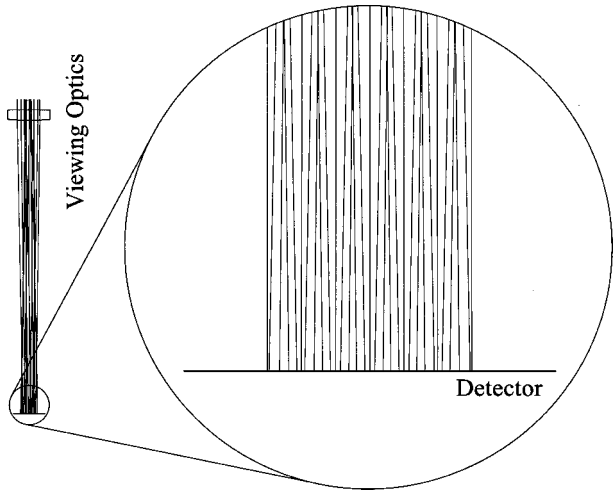

Fig. 13 Rays at the detector for a set of equally spaced points across the input wavefront for a test part with undulations of 7 waves *P-V*.

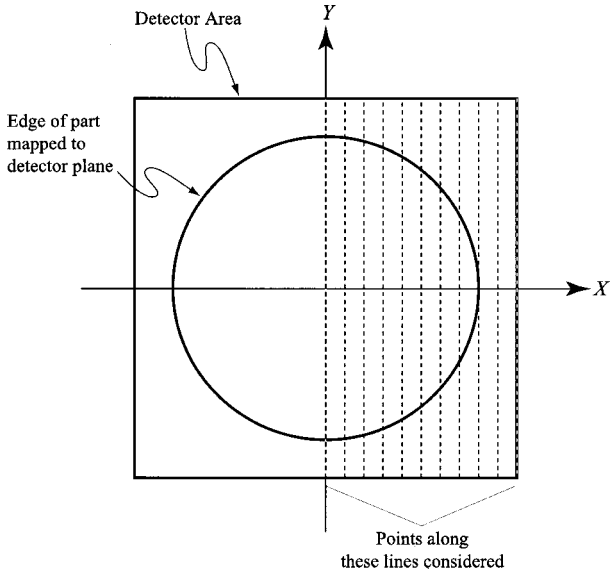

Fig. 14 Lines that contain the points on the detector that were used for evaluating the useful measurement range of the interferometer.

that get mapped to multiple input points.[††] This effect is viewed another way in Fig. 13. In that figure, rays from an equally spaced set of points at the input are traced through the system, and their locations at the detector are shown for the case of six undulations with an amplitude of 7 waves P-V. If the amplitude of the undulations is increased only slightly, the rays begin to intersect at the detector. That is, the caustic moves onto the detector, and the interferometer no longer directly measures the height of the part. (Note that with six undulations at 8 waves P-V, the caustic has moved onto the detector and the ray aiming fails.)

5.3 *Measurement Error*

Before presenting the results of this study, a discussion of the determination of the measurement error is given. Consider a single point on the detector. First, the exact height of the surface being measured is needed. For this, a perfect spherical part is used and the ray (normal to the input wavefront) that passes through the detector point is determined. The actual height of the surface is taken to be the distance between where this ray intersects the perfect spherical part and where it intersects the part with the undulations (with the sign chosen according to which side of the perfect sphere the actual surface lies on). To determine the measured height, the OPD for a perfect spherical part is subtracted from the OPD for the part under test (this simulates calibrating out the errors introduced by the feed lens). Half of this difference in OPD is the measured height. The magnitude of the difference between the measured height and the actual height is the absolute measurement error (for that point on the detector). The relative error is found by dividing the absolute error by the P-V amplitude of the undulations.

For the interferometer considered in this example, the measurement error is determined for points that lie on 11 lines across the detector. These lines are parallel to the Y axis, but are equally distributed from the line $x=0$ to $x=5$ mm (for perfect mapping between the part and detector, the edge of the part is 4 mm from the axis). This is shown in Fig. 14. The number of points on each line is taken to be the greater of (i) 21 and (ii) 8 times the number of undulations on the part. The reported measurement error is the maximum error for all object points in which rays fall within the clear aperture of the part under test. The measured height and the actual height (as found with the procedure described above) are plotted in Fig. 15 for the case of six undulations with an amplitude of 7 waves P-V for points on a line through the center of the detector. Notice that the maximum errors occur in the regions of largest slope (the maximum measurement error for this part is about 65%), but that the interferometer does a reasonable

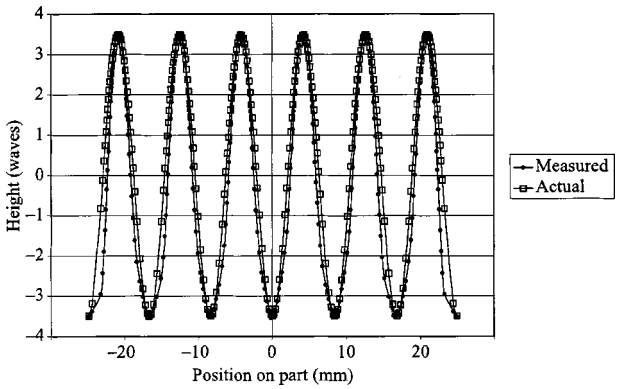

Fig. 15 Plot showing the actual and measured height on a test part with undulations of 7 waves *P-V*.

[††]Even before the caustic moves onto the detector, there are implications for detecting the phase that may be important. The energy where the input points are bunched together gets spread out at the detector (resulting in lower energy densities there). If the test and reference arms have equal total energy, the lower energy densities at some points (and higher energy densities at others) coming from the test arm result in lower-contrast fringes, which makes detecting the phase more difficult. This effect needs to be taken into account in a complete model of an interferometric system, but is not discussed further in this paper.

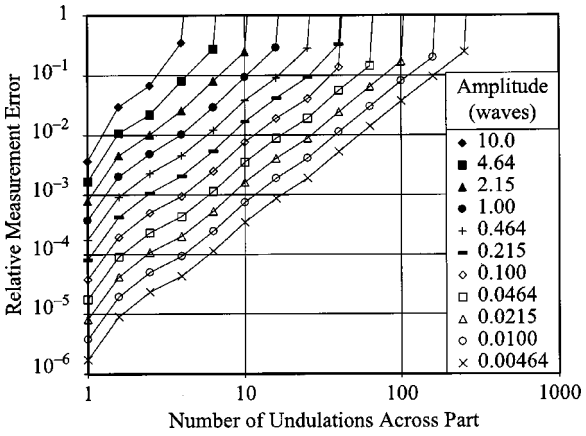

Fig. 16 Relative error plotted as a function of the number of undulations across the part for various values of the amplitude of the undulations.

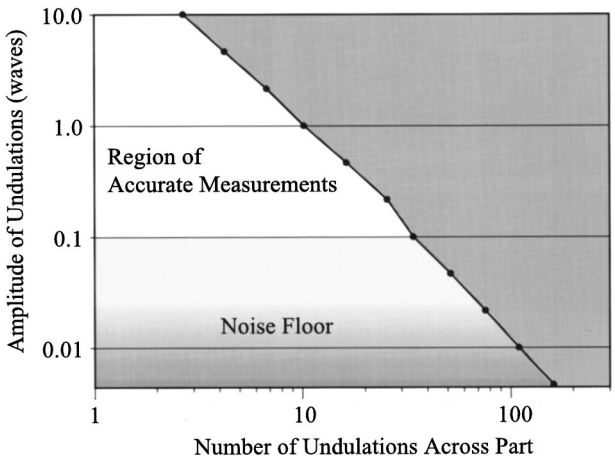

Fig. 17 Plot of the amplitude of undulations that give 10% relative measurement error as a function of the number of undulations across the part. The white area represents the region in which accurate measurements can be made by the interferometer.

job near the extrema of the sinusoid (where the slope is low). Therefore, for a part with a sinusoidal ripple, say, where one wants to measure only the amplitude of the ripples, the measurement error reported here is pessimistic. But for a part that rolls off at the edge, where the slope of the rolloff is the same as the slope of this sinusoid, the absolute measurement error reported here is accurate. This claim is further justified towards the end of Sec. 5.4.

5.4 *Results*

The measurement error is presented now for varying amplitudes and spatial frequencies of the sinusoidal undulations discussed in Sec. 5.1. Figure 16 shows the relative measurement error as a function of the number of undulations across the part for undulations of varying amplitude. For each curve, the data point with the highest relative error is the last one taken before ray-aiming failures occur. Note that in this model, the discrete nature of the pixels in a typical detector is not taken into account—only errors due to the optics that compose the interferometer are modeled. However, for the sake of argument, say that the detector is a square that is 10 mm on a side with 1000×1000 pixel array. Then for undulations with a P-V amplitude as small as 0.005 waves, the transverse size of features on the part that can be resolved is still limited by the optics rather than by the discrete nature of the pixels. (If we take a 10% measurement error to be the maximum acceptable, then from Fig. 16 it can be seen that up to about 200 undulations across the part at an amplitude of 0.005 waves can be resolved. If there are 1000 pixels across the detector, this gives five pixels per undulation, so the discrete nature of the pixels should not limit such a measurement.)

The results presented in Fig. 16 can be used to determine the region in amplitude–spatial-frequency space over which the interferometer makes accurate measurements. For example, say a relative measurement error of at most 10% is acceptable, then Fig. 16 can be used to determine the maximum number of undulations that can be accurately measured for any given amplitude. The results are presented pictorially in Fig. 17. In that figure, the amplitude of the undulation that gives a relative error of 10% is plotted as a function of the number of undulations across the part. Undulations that fall into the gray region to the right of the curve cannot be measured accurately by the interferometer. Typically there is some level of noise in the interferometer that prevents features with too small an amplitude from being measured. The precise level of this *noise floor* depends on the characteristics of the source and detector and on the method used to determine the phase of the fringes. Therefore, it is not shown as a strict boundary in Fig. 17. The area shown in white above the noise floor represents the region in which the optics in the interferometer allow accurate measurements.

One last result is interesting to consider. Rather than the relative error shown in Fig. 16, consider the absolute measurement error. In addition, for each combination of amplitude and spatial frequency, the maximum slope of the part is determined. [By slope of the part, I mean the slope of the difference (measured radially) between the actual part and the base sphere.] Now consider Fig. 18, which shows the absolute measurement error as a function of the maximum slope across the part for undulations with varying ampli-

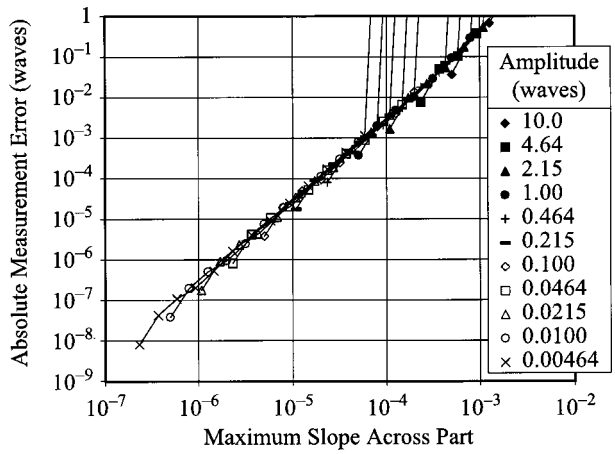

Fig. 18 Absolute error plotted as a function of the maximum slope across the part for various values of the amplitude of the undulations.

tudes. Apparently it is the maximum slope that determines the absolute error. Therefore, to experimentally map out the useful measurement range of this interferometer, it suffices to measure the error in the interferometer as a function of tilt: there is no need to fabricate a variety of parts with various spatial frequencies and amplitudes.

6 Concluding Remarks

In this paper I started with the simple premise that interferometers should be modeled in the way in which they are used. This led to the requirement for a particular ray-aiming scheme. Such a scheme has been implemented as macros in commercial lens design software, which was used to generate all of the data presented in the examples. These tools give the user the capability to model easily a wide variety of interferometers. As a matter of practical interest, it is useful to place a single interferometer within a single lens file. The option to create multiple configurations within a single lens file can be used to model the individual arms in the interferometer. In cases where a known part is used to calibrate a measurement, the calibration setup of the interferometer can be placed in a separate configuration within the lens file. This makes switching between arms in the interferometer (and modeling the calibration of an interferometer) quite convenient.

Note that the examples presented here do not exhaust the types of modeling that can be done, but are intended to demonstrate the power and flexibility of these ideas. For example, while I did not present such examples, any number of system parameters can be perturbed to model the effects of manufacturing errors on interferometer performance. Also, no examples of shearing interferometers were presented, but I claim that any class of interferometer (including shearing) can be modeled within the framework presented here. By staying within the confines of commercial lens design software, any element that can be modeled in the lens design package (such as gratings and diffractive elements) can be incorporated easily into interferometer models. As discussed in Sec. 4, extended, polychromatic sources also can be modeled. While optimization of interferometer performance is not explicitly discussed here, it is possible to use the proposed approach as part of a routine to optimize interferometer components (such as viewing optics). The objective function used in the optimization should emphasize those features of the interferometer that are most important for its intended application. For example, this could include increasing the range of accurate measurement for nonnull testing of aspheres, decreasing the measurement error for high-frequency undulations, decreasing the sensitivity of the interferometer to misalignment of elements, etc.

Table 1 Specifications of the feed objective.

Surface number	Radius (mm)	Separation (mm)	Material
1	6.20077	1.62779	BK7
2	−49.45843	0.21412	Air
3	−12.02225	0.41738	F2
4	−27.45294	3.09281	Air
5	4.41018	1.41910	BK7
6	−5.04456	0.07847	Air
7	−4.80785	0.41738	F2
8	36.50909	3.89484	Air

Table 2 Specifications of the viewing optics.

Surface number	Radius (mm)	Separation (mm)	Material
1	50.000	5.000	BK7
2	0.0	287.91156	Air
3	0.0	5.000	BK7
4	−101.69986	141.55961	Air

The emphasis in this paper is solely on the effects of the optics within an interferometer. However, it is possible to incorporate source characteristics, detector characteristics, and phase detection and unwrapping schemes into the model. That is, the approach described here can form the basis for complete system-level models of interferometers. So, the apparent contradiction notwithstanding, the laws of geometrical optics have a lot to say about the performance of instruments that would not exist but for the wave nature of light.

7 Appendix

The specifications for the lenses in the Twyman-Green interferometer introduced in Sec. 3.1 appear in Tables 1 and 2. This information is required for those wishing to reproduce (perhaps as part of a check of a similar interferometer model) results presented here. In the model, a BK7 beamsplitter cube (made of BK7 that has faces 25.4 mm wide) is placed between the feed and viewing optics. The separation from the beamsplitter cube to both the feed objective and viewing optics is 50 mm (see Fig. 3).

Acknowledgments

The author thanks John Bruning, Paul Michaloski, Andrew Kulawiec, Richard Youngworth, and Yuling Hsu for their thoughtful reading of and comments on this work.

References

1. R. Jozwicki, ''Influence of the field truncation by the aperture stop of interferometers upon the intensity distribution in the observation field,'' *Opt. Appl.* **24**, 439–450 (1989).
2. R. Jozwicki, ''Influence of spherical aberration of an interferometric system on the measurement error in the case of a finite fringe observation field,'' *Appl. Opt.* **30**, 3119–3125 (1991).
3. C. Huang, ''Propagation errors in precision Fizeau interferometry,'' *Appl. Opt.* **32**, 7016–7021 (1992).
4. P. Michaloski, A. Kulawiec, and J. Fleig, ''A method of evaluating and tolerancing interferometer designs,'' in Int. Optical Design Conf., L. R. Gardner and K. P. Thompson, eds., *Proc. SPIE* **3482**, 499–507 (1998).
5. A. E. Lowman and J. E. Greivenkamp, ''Interferometer induced wavefront errors when testing in a non-null configuration,'' *Proc. SPIE* **2004**, 173–181 (1993).
6. A. E. Lowman and J. E. Greivenkamp, ''Modeling an interferometer for non-null testing of aspheres,'' *Proc. SPIE* **2536**, 139–147 (1995).
7. M. C. Noecker, M. A. Murison, and R. D. Reasenberg, ''Optic-misalignment tolerances for the POINTS interferometers,'' *Proc. SPIE* **1947**, 218–231 (1993).
8. N. G. Douglas, A. R. Jones, and F. J. van Hoesel, ''Ray-based simulation of an optical interferometer,'' *J. Opt. Soc. Am. A* **12**, 124–131 (1995).
9. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed., Sec. 9.6, Cambridge Univ. Press, Cambridge (1992).

10. P. A. Stark, *Introduction to Numerical Methods*, Sec. 3.7, Macmillan, New York (1970).
11. G. W. Forbes, ''Optical system assessment for design: numerical ray tracing in the Gaussian pupil,'' *J. Opt. Soc. Am. A* **5**, 1943–1956 (1988).

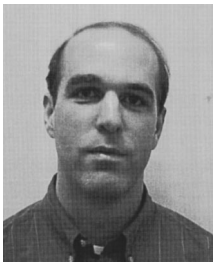

Bryan D. Stone received his BS in optics with high distinction from The Institute of Optics, University of Rochester, in 1985. After working for a year at Amp Incorporated, he returned to the University of Rochester to pursue a PhD degree in optics, which was awarded in 1992. He spent two years in a postdoctoral position at the University of Rochester before joining the faculty of The Institute of Optics full time. In 1998, he joined Tropel Corporation full time as a staff scientist. He is currently the associate editor for lens design of *Optical Engineering*. His professional interests include aberration theory and lens design, asymmetric system design, fabrication and testing of optical components, and assembly and testing of optical systems.