applied optics

Ray-tracing model of a perfect lens compliant with Fermat's principle: the Cardinal Lens

Jeffrey P. Wilde

E. L. Ginzton Laboratory, Stanford University, Stanford, California 94305, USA (jpwilde@stanford.edu)

Received 13 October 2023; revised 5 January 2024; accepted 6 January 2024; posted 9 January 2024; published 29 January 2024

When using ray tracing for optical system design, it is often the case that the designer would like to implement simplified versions of one or more compound lens groups. This could be the case during initial layout when idealized versions of such compound lenses are needed or, perhaps alternatively, to mimic a well-corrected commercially available lens for which the prescription details are unavailable. One option is to use a paraxial thin lens as a proxy for the actual lens group, but doing so will yield a layout that is not consistent with Fermat's principle or the Abbe sine condition. For example, a paraxial lens version of a compound microscope objective typically produces the wrong numerical aperture for a given entrance pupil diameter, and vice versa. A better option is to use a lens model that provides perfect imaging for a specified paraxial magnification and obeys Fermat's principle. A variant of the model can yield a perfect Fourier transform lens. In addition, it is desirable to implement an idealized thick lens in which the principal planes are separated by a user-specified distance. This paper presents such a model, referred to as the Cardinal Lens, with implementation in Zemax OpticStudio via a user-defined surface.

https://doi.org/10.1364/AO.507605

1. INTRODUCTION

In sequential optical system design, it is often useful to begin by using a paraxial version for some, or perhaps all, of the lens elements. Such a first-order layout helps define what the performance of a system looks like when aberrations are neglected for the idealized elements. However, there are instances when the properties of a paraxial lens are insufficient to mimic a real, well-corrected compound lens. The foremost limitation stems from the fact that a paraxial lens, when used outside of the paraxial regime, does not satisfy the Abbe sine condition, whereas many actual lenses do obey the sine condition, at least approximately, and are known as aplanatic lenses. They are free from, or have minimal, spherical aberration and coma (for small off-axis fields). If a lens is free from astigmatism over its field of view, while also having no spherical aberration and no coma, it is referred to as being a stigmatic lens. On a related historical note, the well-known lens designer Paul Rudolph, while working for the Carl Zeiss company in Germany, coined the term "anastigmatic" to describe a lens for which the astigmatism at one off-axis field angle is zero ([1], p. 236). More recently, this term has come to mean a lens that is designed to have reduced astigmatism over its field of view [2]. An interesting example of an actual stigmatic lens is the Luneburg lens based on a spherical gradient of the refractive index (see e.g., [3,4], p. 712). However, for modeling an idealized perfect lens, we can ignore the physical structure of the lens and instead treat it as a black box that has properties consistent with Fermat's principle.

When capturing images with a planar sensor, one would also like to have minimal field curvature. Distortion, on the other hand, may or may not be a problem. For perfect imaging there should be no distortion, but a perfect Fourier transform lens has intrinsic distortion when viewed from an imaging perspective [5]. So, it is beneficial to have access to a generalized perfect lens model that (a) is free from spherical aberration, coma, astigmatism, and field curvature; (b) provides a specified form for the distortion; and (c) accurately captures the constraints associated with Fermat's principle and hence satisfies the sine condition for arbitrary numerical apertures. Additional flexibility arises if this perfect lens model can mimic a thick lens with its two principal planes separated from one another.

Commercial ray-tracing software packages provide various types of idealized lens models. For example, Zemax OpticStudio [6] includes a thin paraxial lens surface but has no version of a perfect lens satisfying the sine condition. Other software such as CODE V [7,8] and OSLO [9] offer versions of perfect lenses that are based on so-called "mock ray tracing" using an eikonal formulation that typically requires an iterative numerical solution when the lens is represented by a specific eikonal function ([10], Ch. 31).

In this paper a relatively straightforward perfect lens ray-tracing model is presented that complies with Fermat's principle and satisfies a generalized sine condition, but does not require eikonal function theory. Two versions are developed, one that has no distortion, characterized by an $f\tan(\theta)$ response, as needed to model perfect imaging, and a second version that has

intrinsic $f\sin(\theta)$ distortion, which is representative of an ideal Fourier transform lens.

We refer to our particular formulation of this model as a "Cardinal Lens" in order to distinguish it from other embodiments of idealized lenses [11]; it has been implemented as a user-defined surface for sequential ray tracing in Zemax OpticStudio, written in C++ and compiled as a 64-bit DLL (dynamic link library) and made available in Code 1, Ref. [12]. Along with the form of the distortion, this model requires the user to specify the effective focal length and the paraxial magnification at which the lens provides optimal performance. The lens can be embedded in media with different refractive indices before and after the lens and can be used at large numerical aperture (NA), thereby allowing, for example, simulation of (a) air or immersion microscope objectives, (b) Fourier transform lenses outside of the paraxial regime, and (c) wide-angle, high-speed imaging lenses.

The remainder of the paper is structured as follows. Section 2 provides an overview of the key features associated with modeling a perfect lens. Section 3 describes the basic formulation of the Cardinal Lens model. The details associated with its implementation in OpticStudio follow in Section 4. A few examples are provided in Section 5, while additional example layouts are included in Supplement 1. Finally, Section 6 offers a brief summary and conclusion.

2. MODELING AN IDEALIZED PERFECT LENS

A. Problem Statement

Consider a ray propagating in a sequential optical system and passing through an axially symmetric perfect lens as shown in Fig. 1. The lens effective focal length is f, and the thickness of the lens, t, corresponds to the separation between the principal planes. The input space to the left of the lens is homogeneous with refractive index n_1, while the output space on the right is homogeneous with index n_2. The ray z coordinates are measured locally relative to the principal planes, with the input and output spaces being referred to the first and second principal planes, respectively. A z coordinate lying to the left of either principal plane is negative, and to the right it is positive. A ray segment is defined by a unit vector pointing along the direction of travel and a point in space through which the ray passes. We denote a ray vector by $\vec{s} = (L, M, N)$, where (L, M, N) are direction cosines such that $N = (1 - L^2 - M^2)^{1/2}$.

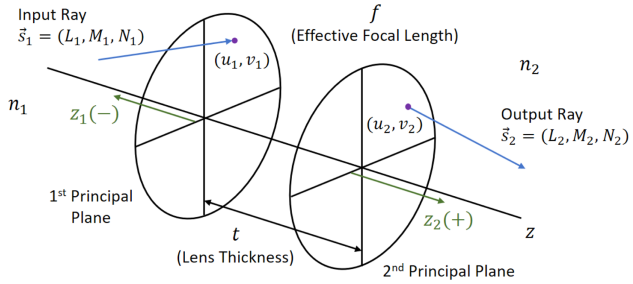

Fig. 1. Local coordinate system for ray tracing through an axially symmetric perfect lens of thickness t. For a given input ray, we seek to find the corresponding output ray that complies with Fermat's principle.

In Fig. 1 an incoming ray segment $\vec{s}_1 = (L_1, M_1, N_1)$ is incident on the first principal plane of the lens at point (u_1, v_1). These values are readily available from the ray-tracing software. Our task is to find the outgoing ray segment coordinates (u_2, v_2) on the second principal plane along with the direction cosines $\vec{s}_2 = (L_2, M_2, N_2)$ in such a way that this outgoing ray complies with Fermat's principle, and then send these results back to the main program so the ray can continue to propagate. Of course proper analysis of a general optical system that incorporates a perfect lens requires tracing many rays, so this process is repeated for each ray launched into the system and arriving at the perfect lens.

B. Localized Imaging Perspective and the Principal Reference Spheres

While the perfect lens may be just one element in a more complicated sequential lens train, for purposes of ray tracing we can analyze it in isolation, based on a localized imaging viewpoint since determination of an outgoing ray from the lens only relies on knowledge of the incoming ray, independent of any other specific optical elements that may come before or after the lens. Therefore, let us consider the imaging properties of a perfect lens taken by itself, as doing so is key to construction of the ray-tracing model. The local object and image planes for the perfect lens need not correspond to the actual physical object or image planes for the overall system unless, of course, the system is constructed that way. Instead, the localized conjugate planes are, in general, virtual constructs used solely to determine the ray path through the lens.

Sharp imaging, also known as stigmatic imaging, implies that all rays leaving any given point on an object surface will arrive at a single corresponding point on a conjugate surface in image space. For homogeneous object and image spaces, it is well-established that a rotationally symmetric perfect lens can stigmatically image just one object surface to one image surface (neglecting certain degenerate cases, such as perfect virtual imaging of the whole object space via reflection by a plane mirror, which are of little practical value) ([1], p. 149; [10], Ch. 22). Here we are interested in object/image surfaces that are planar. Therefore, a perfect lens can only provide stigmatic imaging at one specific transverse magnification (or paraxial magnification when distortion is present). This is distinctly different from a paraxial lens, which provides aberration-free imaging at any magnification and is therefore inconsistent with Fermat's principle. In our perfect lens model, the optimal (paraxial) magnification is a user-specified quantity.

Because we adopt a localized imaging model, we are not concerned here with an overall system stop or the location of the pupils. Instead, while the lens can have a limiting aperture that restricts ray transmission, the principal planes of the lens are of primary interest. More specifically, it is helpful to utilize the concept of principal reference spheres as described next.

When tracing a given ray, we first find the corresponding local object and image points, determined by the locations where the ray pierces the object and image planes of the lens as illustrated in Fig. 2. The first and second principal reference spheres are centered on the object and image points, respectively, and intersect the optical axis at the first and second principal points

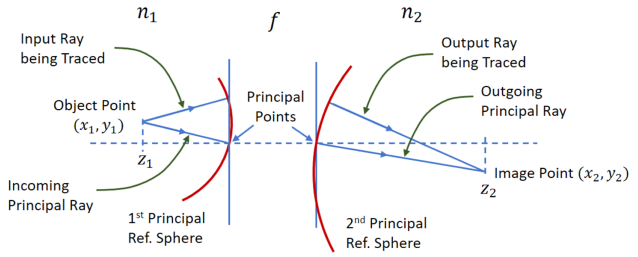

Fig. 2. Schematic of a perfect lens showing the localized object and image points for a given ray that is being traced through the lens. The principal reference spheres are centered on these points. The principal rays, which connect the object/image points to their respective principal points, are utilized internally within the lens model.

(see, e.g., [10], p. 194; [4], p. 430; and [13], p. 15). In some limited respects the principal reference spheres are analogous to the entrance and exit pupil reference spheres commonly used in aberration theory. However, the principal reference spheres pertain solely to the lens of interest and have no connection to the system stop. In the paraxial regime, the principal spheres become conjugate planes with unity magnification, while the pupil spheres become conjugate planes with a magnification that depends on the stop location. So these two sets of reference spheres are fundamentally different entities.

While the principal reference spheres are conceptually very important, for general ray-tracing implementation it is preferable to calculate the ray coordinates at the principal planes and to display the rays in a layout diagram such that they terminate at the first principal plane and resume propagation at the second principal plane (as shown in Fig. 1). If the user would like to draw the principal reference spheres for one or more localized object/image points, and optionally turn off drawing of the ray segments to the principal planes, they can do so manually as illustrated in Section 5.A, Fig. 8.

C. Optical Path Length Considerations and the Cosine Rule

For a perfect lens, the optical path length (OPL) is the same for all rays connecting any given object point to its image point. A subtler, but vitally important, property of a perfect lens derives from Hamiltonian optics, in particular from the cosine rule as first derived by Smith in 1922 [14] and subsequently discussed in other references; for example, see [10,15,16]. For the sake of completeness, we present here a short derivation of the cosine rule, by adopting a slight variation of the description provided by Walther [10], Ch. 6. With reference to Fig. 3, we begin by showing an object point P_1 that is stigmatically imaged to its conjugate point P_2. Therefore the OPL from P_1 to P_2, denoted as $[P_1\,P_2]$, is the same for all rays connecting P_1 and P_2. Now, consider a second object point Q_1 that resides within a small neighborhood of P_1 and is located by an infinitesimal displacement vector $d\vec{r}_1$, and assume that it too is stigmatically imaged to point Q_2 in the neighborhood of P_2 with displacement vector $d\vec{r}_2$. This assumption requires that $[Q_1\,Q_2]$ be the same for all rays connecting Q_1 and Q_2. An equivalent requirement is that the optical path length difference, given by $[Q_1\,Q_2]-[P_1\,P_2]$, be constant.

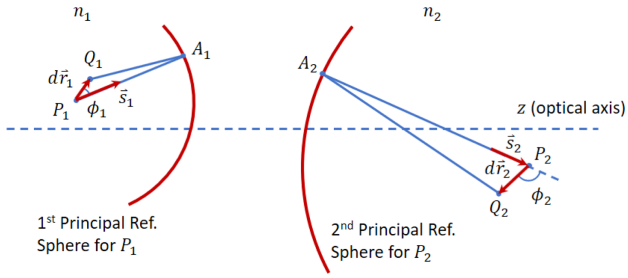

Fig. 3. Geometry for derivation of the cosine rule.

To quantify the situation, consider a single arbitrary ray segment that begins at P_1, pointing along a unit vector $\vec{s}_1=(L_1, M_1, N_1)$ and intersecting the first principal reference sphere at A_1. This ray then goes on to intersect the second principal reference sphere at A_2, finally terminating at P_2 via $\vec{s}_2=(L_2, M_2, N_2)$. Its OPL is $[P_1\,P_2]=[P_1\,A_1\,A_2\,P_2]$. To find $[Q_1\,Q_2]$ we can apply Fermat's principle. Instead of using an actual ray path from Q_1 to Q_2, we take a path that is closely adjacent to the ray path from P_1 to P_2 (going through A_1 and A_2), which is possible because the two object points and their respective image points are infinitesimally close to one another. Fermat's principle of stationarity ensures that the adjacent path OPL from Q_1 to Q_2, denoted $[Q_1\,A_1\,A_2\,Q_2]$, is approximately equal to the actual ray OPL, which we write as $[Q_1\,Q_2]$, at least to first order in the small displacements. We therefore have

$$\begin{aligned}[Q_1\,Q_2]-[P_1\,P_2] &\simeq [Q_1\,A_1\,A_2\,Q_2]-[P_1\,A_1\,A_2\,P_2]\\ &=([A_2\,Q_2]-[A_2\,P_2])-([P_1\,A_1]-[Q_1\,A_1])\\ &=n_2\vec{s}_2\cdot d\vec{r}_2-n_1\vec{s}_1\cdot d\vec{r}_1\\ &=n_2\cos(\phi_2)|d\vec{r}_2|-n_1\cos(\phi_1)|d\vec{r}_1|.\end{aligned}\tag{1}$$

Here we have expressed the vector dot products in terms of the cosines of the angles between the vectors as shown Fig. 3. As noted previously, this expression for the OPD must equate to a constant if Q_1 and Q_2 are to be stigmatically imaged. Dividing by $|d\vec{r}_1|$ and noting that the differential magnification associated with these field points is $|m_d|=|d\vec{r}_2|/|d\vec{r}_1|$, the above expression leads to the cosine rule as follows:

$$|m_d|n_2\cos(\phi_2)-n_1\cos(\phi_1)=C,\tag{2}$$

where C is a constant specific to the given field points. Since the cosine rule must apply for all rays connecting P_1 and P_2, if we change $\vec{s}_1$ and $\vec{s}_2$, the angles ϕ_1 and ϕ_2 will also change, but the constant remains the same. So, we can calculate the constant using any convenient ray of our choice, and then use this constant to find other rays connecting P_1 and P_2 by calculating the outgoing ray segment $\vec{s}_2$ given an incoming ray segment $\vec{s}_1$. In other words, the net consequence of the cosine rule is that the angle of a ray coming to focus at an image point can be found from its angle of departure from the object point.

D. Generalized Sine Condition

The cosine rule applies to displacement vectors $d\vec{r}_1$ and $d\vec{r}_2$ residing in small three-dimensional volumes centered on P_1 and

P_2, respectively. Of primary interest here is imaging from an object plane to an image plane, both planes being normal to the optical axis. In this case we restrict the displacement vectors to lie in infinitesimal two-dimensional circles on these planes.

In general, for a lens with distortion, the displacement vectors need not be parallel to one another. For example, a small displacement of an arbitrary object point in the x direction may yield a shift of the image point having both x and y components. However, for a lens with axial symmetry, as we assume here, this coupling can be avoided by constraining the object point to lie on either the x or y axis. We choose the y axis, so the object and image points both reside in the yz plane, which we refer to as the tangential (or meridional) plane. Doing so simplifies calculation of the differential magnification components. To accommodate an arbitrary off-axis object point, the lens system (including the incoming ray) can first be rotated about the z axis to bring this point into the tangential plane, then the ray of interest can be traced from the new object point location through the lens, and finally the lens coordinate system containing the outgoing ray segment can be rotated back to its original orientation. More detail about this process is discussed in Section 3.A.

Now, assuming the object point lies in the tangential yz plane, if we choose one set of displacement vectors to be oriented along the x direction and a second set along the y direction, the cosine rule yields the following two relations, which constitute the generalized sine condition (GSC) (for related versions see [14], p. 37; [17], Eq. 11):

$$\begin{aligned} m_{dx} n_2 L_2 - n_1 L_1 &= C_x \\ m_{dy} n_2 M_2 - n_1 M_1 &= C_y. \end{aligned} \tag{3}$$

Here $m_{dx} = dx_2/dx_1$ and $m_{dy} = dy_2/dy_1$ are the components of the differential magnification along the x and y axes, respectively, while C_x and C_y are field-dependent constants. For a lens having distortion, such as a perfect Fourier transform lens, m_{dx} and m_{dy} vary with field (i.e., with the y_1 value) and can take on different values, meaning the system is locally anamorphic as discussed further in Appendix A. However, for a perfect imaging lens m_d is isotropic and constant over the field of view, and can be equated to the paraxial magnification m_p (i.e., $m_{dx} = m_{dy} = m_p$). So when modeling a perfect imaging lens it is not necessary to confine the object point to the yz plane, thereby avoiding the rotation operation.

This formulation of the GSC makes use of the ray direction cosines, but these (L, M) values for a ray segment can be found by projecting the ray segment onto the xz and yz planes, respectively, and then computing the *sines* of the angles that these two projected vectors make with respect to the z axis. For this reason, Eq. (3) is considered a "sine" condition. In addition, note that Eq. (3) uses m_{dx} and m_{dy}, which can be positive or negative, whereas Eq. (2) uses $|m_d|$. With a little effort, the interested reader can confirm the signs of m_{dx} and m_{dy} are needed when converting the cosine rule based on the angles ϕ_1, ϕ_2 (now measured with respect to differential displacement vectors confined to the xy plane) to the sine condition based on (L, M) values. Lastly, Eq. (3) is referred to as the "generalized" sine condition because traditionally the Abbe sine condition applies to on-axis points with the constants $C_x = C_y = 0$ and $m_{dx} = m_{dy} = m_p$, while the generalized version allows us to explicitly consider off-axis conjugate points, either with or without distortion. For a perfect lens, the GSC is satisfied for all object/image conjugate points throughout the local field of view of the lens, with each field point having its own set of constants C_x and C_y.

Enforcement of the GSC at any given field point also ensures local stigmatic imaging in the near vicinity of that point. However, we do not rely on the GSC to yield stigmatic imaging; instead, as will be seen in Section 3, we implement a model that by design has no ray or wavefront aberrations, independent of the GSC, when the lens is used at its specified optimal magnification. We do, though, utilize the GSC to calculate the direction cosines of the rays leaving the lens. Doing so ensures that the angular composition of a ray bundle (e.g., an elliptic cone) coming to focus at an imaging point is properly related to the angular composition of the corresponding ray bundle leaving the object point. It is in this sense that the model complies with Fermat's principle because the GSC derives from the cosine rule, which in turn derives from the application of Fermat's principle.

E. Distortion Considerations

The image from a perfect lens may or may not suffer from distortion, but, regardless, it yields stigmatic imaging with a one-to-one mapping of object-plane points to image-plane points. Without distortion, perfect imaging is characterized by a constant transverse magnification that is applicable across the entire image plane. (So again there is no need to constrain the object point to the yz plane via a lens system rotation as discussed in the previous section.) Therefore, when imaging at infinity, the image point location follows an $f_1 \tan(\theta_1)$ relationship in which f_1 is the object-space focal length and θ_1 is the incoming field angle. However, if distortion is present, then the conventional transverse magnification, as well as the differential magnification values, vary with the object point location (i.e., they are field dependent). In this case we can specify the paraxial magnification for object points very near the optical axis, but the change in transverse magnification with field depends on what type of perfect lens the user would like to simulate. For example, a perfect Fourier transform lens imparts $f_1 \sin(\theta_1)$ distortion (see [5]; and [13], p. 36). Our Cardinal Lens model simulates a perfect imaging lens as well as a perfect Fourier transform lens. Additional details related to these two distortion relationships are discussed in Sections 3.D and 3.E.

F. Conjugate Plane Locations

Given the effective focal length (f) of the lens and the paraxial transverse magnification (m_p) at which stigmatic imaging occurs, we can easily determine the locations, z_1 and z_2, of the local object and image conjugate planes (as shown in Fig. 2). Beginning with the lens equation,

$$\frac{1}{f} = \frac{n_2}{z_2} - \frac{n_1}{z_1}, \tag{4}$$

and combining it with an expression for the paraxial magnification,

$$m_p = \frac{z_2/n_2}{z_1/n_1}, \tag{5}$$

leads to

$$z_1 = n_1 f\left(\frac{1}{m_p} - 1\right), \quad z_2 = n_2 f\left(1 - m_p\right). \tag{6}$$

Again, f is the effective focal length, with the front and back focal lengths being $f_1 = n_1 f$ and $f_2 = n_2 f$, respectively.

3. RAY PROPAGATION ANALYSIS: THE CARDINAL LENS MODEL

We now describe how any given ray incident on the lens is traced through the lens; the details comprise our Cardinal Lens model. It is a two-step process in which we first find the principal ray segments and then second use these principal ray segments to facilitate tracing the ray of interest via the GSC. For purposes of discussion, assume the object plane is located to the left of the lens (negative z_1) and the image plane is on the right (positive z_2) as depicted in Fig. 2. In this case the lens is positive, so both the object and image are real and the magnification is negative. In general, though, the lens can be either positive or negative, so the object and image planes can reside on either side of the lens. This means the object, as well as the image, may be either real or virtual in accordance with Eq. (6). The ray-tracing formulation developed here is applicable to any of these combinations (illustrated in Table S1 of Supplement 1).

A. Conjugate Point Locations

To determine the object point location (x_1, y_1), we simply calculate where the incoming ray being traced crosses the object plane. For the perfect imaging lens we can proceed directly to calculation of the corresponding image point coordinates as discussed below. For the perfect Fourier transform lens, as noted in Section 2.D, we need to first apply a rotation to the lens coordinate system to bring the object point onto the y axis unless, of course, the object point already happens to reside in the yz plane. The angle of rotation is given by $\alpha = -\text{atan2}(y_1, x_1) + \pi/2$, where atan2 is the four-quadrant inverse tangent function (yielding the azimuthal angle of the object point as measured from the $+x$ axis). The minus atan2 term takes the object point to the $+x$ axis, and the $\pi/2$ term transfers it to the $+y$ axis. By using the standard 2D rotation matrix, $R(\alpha) = [\cos\alpha, -\sin\alpha; \sin\alpha, \cos\alpha]$, the following vectors are transformed: (x_1, y_1), (L_1, M_1), and (u_1, v_1). For simplicity we use the same notation for the vectors after rotation, with it being implicitly understood that for the Fourier lens the ray analysis is conducted in the rotated coordinate system, and the results are subsequently transformed back into the original coordinate system via $R(-\alpha)$.

To determine the image point location, we make use of the input/output principal ray segments as shown in Fig. 4. These principal rays are not seen by the user; instead, they are only used internally within the model. The incoming principal ray is directed from the object point to the center of the first principal plane; its direction cosines are denoted $\vec{s}_{1p} = (L_{1p}, M_{1p}, N_{1p})$. From Fig. 4 we see $M_{1p} = \sin(\theta_{1y})$. For the perfect imaging lens we similarly have $L_{1p} = \sin(\theta_{1x})$, while for the perfect Fourier lens $L_{1p} = 0$. In either case, $N_{1p} = (1 - L_{1p}^2 - M_{1p}^2)^{1/2}$.

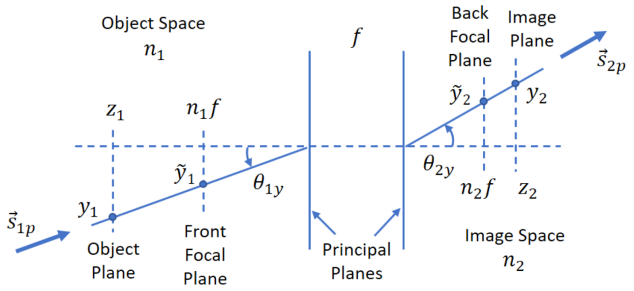

Fig. 4. Cross-sectional diagram of the tangential yz plane showing the principal ray segments in object/image spaces. The object-space principal ray $\vec{s}_{1p}$ connects the object point (associated with a given input ray) to the center of the first principal plane. The image-space principal ray $\vec{s}_{2p}$ propagates from the center of the second principal plane toward a point on the back focal plane determined by the user's choice for the form of the lens distortion. The intersection of the image-space principal ray with the image plane determines the image point location.

Next, the form of the distortion is used to find the outgoing principal ray in image space via its intersection with the back focal plane. For the imaging lens, the back focal plane coordinates are $(\tilde{x}_2, \tilde{y}_2)$, which then allow us to find the image point location (x_2, y_2) by extending the ray segment to the image plane. The details are presented below in Section 3.4. In this case the conjugate points are related via the paraxial magnification, and the differential magnification components are equal to m_p. However, for a Fourier lens (after rotation), the x coordinates in both the back focal plane and image plane are nominally zero. However, to calculate the differential magnification components, we need to examine small x and y displacements of the object point (e.g., using $\delta x_1 = \delta y_1 = z_1/1000$) and find the corresponding image point perturbations $(\delta x_2, \delta y_2)$, from which we can readily calculate the differential magnification values $m_{dx} = \delta x_2/\delta x_1$ and $m_{dy} = \delta y_2/\delta y_1$. So our analysis for the Fourier lens case (Section 3.E) accommodates the necessary δx_1 (and corresponding δx_2) displacements out of the yz plane.

B. Managing Infinite Conjugates

At this point we digress briefly to address the issue of infinite-conjugate imaging. For small values of magnification, the input beam is close to being collimated, while for large magnification the output beam is approximately collimated. The limiting case of $m_p \to 0$ (with $z_1 \to \pm\infty$) corresponds to an infinite front conjugate case, so that all incoming rays are parallel to the principal ray in object space. A similar situation applies to the infinite rear conjugate case in which $m_p \to \pm\infty$ (with $z_2 \to \mp\infty$). Numerically, we assume the infinite front and rear conjugate scenarios correspond to $|m_p| \le 1\text{e}-10$ and $|m_p| \ge 1\text{e}+10$, respectively. In these two limiting cases, we apply a long-conjugate approximation and somewhat arbitrarily set $z_1 = -1\text{e}10$ (infinite front conjugate) and $z_2 = +1\text{e}10$ (infinite rear conjugate), both measured in lens units (typically millimeters). We then internally set the paraxial magnification so that it equals z_2/z_1 and, for the infinite front conjugate case, temporarily modify the incoming ray direction cosines to make them consistent with the finite long-conjugate approximation. Next, as will be discussed in Section 3.F, we apply the GSC to

obtain the output ray direction cosines, from which we calculate the outgoing ray coordinates (u_2, v_2) (see Fig. 1). Because either the incoming or the outgoing ray direction cosines are only approximate, we lastly make the necessary minor adjustments to the direction cosines to ensure the beam is properly collimated in either object or image space. For the infinite front conjugate case, this simply entails restoring the original direction cosines of the incoming ray. For the infinite rear conjugate case, we set the outgoing ray direction cosines equal to those of the outgoing principal ray.

C. General Approach and Role of the Paraxial Magnification

Our general approach utilizes the form of the distortion to propagate the outgoing principal ray from the center of the second principal plane into image space. As mentioned previously, we consider two cases: (a) a perfect imaging lens without distortion, characterized by an $f\tan(\theta)$ response, and (b) a perfect Fourier transform lens with $f\sin(\theta)$ distortion. For each case we consider two separate ranges of magnification, $|m_p| \le 1$ and $|m_p| > 1$, which helps facilitate handling of the front and rear infinite-conjugate scenarios. After deriving equations for the outgoing principal ray direction cosines and corresponding image point coordinates in Section 3.D (imaging lens) and Section 3.E (Fourier lens), we turn our attention to tracing the ray of interest in Section 3.F. The reader who is primarily interested in using the Cardinal Lens model in OpticStudio, and is less concerned with the mathematical details, may go directly to Sections 4 and 5.

D. Perfect Imaging Lens

For a perfect imaging lens without distortion, the output principal ray leaves the center of the second principal plane and intersects the back focal plane at coordinates

$$\tilde{x}_2 = n_1 f \tan(\theta_{1x}) = n_1 f L_{1p}/N_{1p},$$
$$\tilde{y}_2 = n_1 f \tan(\theta_{1y}) = n_1 f M_{1p}/N_{1p}. \quad (7)$$

Likewise, the front focal plane intercept coordinates are related to the direction cosines (L_{2p}, M_{2p}, N_{2p}) of the outgoing principal ray,

$$\tilde{x}_1 = n_2 f \tan(\theta_{2x}) = n_2 f L_{2p}/N_{2p},$$
$$\tilde{y}_1 = n_2 f \tan(\theta_{2y}) = n_2 f M_{2p}/N_{2p}. \quad (8)$$

1. Imaging Lens: $|m_p| \le 1$

Note that $|m_p| \le 1$ includes the case of infinite front conjugate imaging for which $m_p = 0$, so in this regime we use an approach that does not rely on the object plane location. From similar triangles in the image space of Fig. 4 we see $y_2/z_2 = \tilde{y}_2/(n_2 f)$, with a corresponding result for x_2 and $\tilde{x}_2$, which are then combined with Eq. (7) to determine the image point coordinates,

$$x_2 = \frac{z_2\tilde{x}_2}{n_2 f} = z_2\left(\frac{n_1}{n_2}\right)\left(\frac{L_{1p}}{N_{1p}}\right),$$
$$y_2 = \frac{z_2\tilde{y}_2}{n_2 f} = z_2\left(\frac{n_1}{n_2}\right)\left(\frac{M_{1p}}{N_{1p}}\right). \quad (9)$$

The length of the outgoing principal ray segment (from the center of the second principal plane to the image point) is $r_p = (x_2^2 + y_2^2 + z_2^2)^{1/2}$, so the principal ray direction cosines in image space are simply

$$L_{2p} = \sin(\theta_{2x}) = x_2/r_p,$$
$$M_{2p} = \sin(\theta_{2y}) = y_2/r_p. \quad (10)$$

2. Imaging Lens: $|m_p| > 1$

For $|m_p| > 1$ (which includes the case of infinite rear conjugate with $z_2 \to \infty$) we proceed by first calculating the outgoing principal ray direction cosines. To do so we use Eq. (8) and write

$$a = \frac{L_{2p}}{N_{2p}} = \frac{\tilde{x}_1}{n_2 f} = \left(\frac{n_1}{n_2}\right)\left(\frac{x_1}{z_1}\right),$$
$$b = \frac{M_{2p}}{N_{2p}} = \frac{\tilde{y}_1}{n_2 f} = \left(\frac{n_1}{n_2}\right)\left(\frac{y_1}{z_1}\right), \quad (11)$$

where, as seen from similar triangles in object space, we have made use of $\tilde{y}_1/(n_1 f) = y_1/z_1$ (along with a corresponding result for x_1 and $\tilde{x}_1$).

It is straightforward to show $N_{2p} = (1 + a^2 + b^2)^{-1/2}$. So the object point coordinates yield values for a and b, and hence N_{2p}, from which we determine the remaining direction cosines:

$$L_{2p} = a N_{2p},$$
$$M_{2p} = b N_{2p}. \quad (12)$$

The image point coordinates then follow from these outgoing principal ray direction cosines:

$$x_2 = \left(L_{2p}/N_{2p}\right) z_2,$$
$$y_2 = \left(M_{2p}/N_{2p}\right) z_2. \quad (13)$$

Note, because we constrain $|z_2| \le 1e10$, for the infinite rear conjugate case ($z_2 \to \infty$) these image coordinates are only approximations, but they are adequate for determining the (u_2, v_2) coordinates of the outgoing ray at the second principal plane of the lens.

E. Perfect Fourier Transform Lens

For a perfect Fourier transform lens, in order to apply the GSC, the object and image conjugate points are assumed to nominally reside in the yz plane (no x components). However, the following analysis is not subject to this constraint; it includes both x and y coordinates and is generally applicable to arbitrary conjugate point locations. Therefore, it can be used to analyze $(\delta x_1, \delta x_2)$ displacements of the nominal object/image points

out of the yz plane as needed to calculate the differential magnification in the x direction. Bearing this important detail in mind, we can proceed in a fashion similar to that for the perfect imaging lens, and begin by noting the output principal ray intersects the back focal plane at coordinates

$$\tilde{x}_2 = n_1 f \sin(\theta_{1x}) = n_1 f L_{1p},$$
$$\tilde{y}_2 = n_1 f \sin(\theta_{1y}) = n_1 f M_{1p}, \tag{14}$$

while again the front focal plane intercept coordinates are related to the direction cosines of the outgoing principal ray,

$$\tilde{x}_1 = n_2 f \sin(\theta_{2x}) = n_2 f L_{2p},$$
$$\tilde{y}_1 = n_2 f \sin(\theta_{2y}) = n_2 f M_{2p}. \tag{15}$$

1. Fourier Transform Lens: $|m_p| \le 1$

As before, from Fig. 4 we have $y_2/z_2 = \tilde{y}_2/(n_2 f)$, with a corresponding result for x_2 and $\tilde{x}_2$, which are combined with Eq. (14) to determine the image point coordinates as follows:

$$x_2 = \frac{z_2 \tilde{x}_2}{n_2 f} = z_2 \left(\frac{n_1}{n_2}\right) L_{1p},$$
$$y_2 = \frac{z_2 \tilde{y}_2}{n_2 f} = z_2 \left(\frac{n_1}{n_2}\right) M_{1p}. \tag{16}$$

Again, using $r_p = (x_2^2 + y_2^2 + z_2^2)^{1/2}$, the principal ray direction cosines in image space are

$$L_{2p} = x_2/r_p,$$
$$M_{2p} = y_2/r_p. \tag{17}$$

2. Fourier Transform Lens: $|m_p| > 1$

For $|m_p| > 1$ (which includes the case of infinite rear conjugate with $z_2 \to \infty$) we proceed to calculate the outgoing principal ray direction cosines by combining Eq. (15) with $\tilde{y}_1/(n_1 f) = y_1/z_1$ (along with a corresponding result for x_1 and $\tilde{x}_1$), as seen from similar triangles in object space, to yield

$$L_{2p} = \left(\frac{n_1}{n_2}\right)\left(\frac{x_1}{z_1}\right),$$
$$M_{2p} = \left(\frac{n_1}{n_2}\right)\left(\frac{y_1}{z_1}\right). \tag{18}$$

The image point coordinates then follow from these outgoing principal ray direction cosines,

$$x_2 = \left(L_{2p}/N_{2p}\right) z_2,$$
$$y_2 = \left(M_{2p}/N_{2p}\right) z_2. \tag{19}$$

Again we note for the limiting case of infinite rear conjugate (such that numerically $|z_2| = 1e10$) these are approximate values.

F. Output Ray Direction Cosines from the Generalized Sine Condition

Having found the image point coordinates through which the output ray being traced must pass, we can turn our attention to finding the direction cosines (L_2, M_2, N_2) of this output ray segment. This is done by application of the GSC derived in Section 2.D, so the result is consistent with Fermat's principle. Once we have the output direction cosines, we can determine the intersection point (u_2, v_2) of the outgoing ray with the second principal plane of the lens, which completes our task.

Consider first the perfect imaging lens case. With reference to Fig. 5, we restrict attention to object points lying in the z_1 plane and image points in the z_2 plane. We begin with the known principal ray segment direction cosines, (L_{1p}, M_{1p}) and (L_{2p}, M_{2p}), for the specific set of conjugate points associated with the ray that is being traced. These principal ray direction cosines are found as outlined in the previous section. Next we apply the generalized sine condition to these principal ray segments and compute values for C_x and C_y:

$$C_x = m_{dx} n_2 L_{2p} - n_1 L_{1p} = m_{dx} n_2 \sin(\theta_{2x}) - n_1 \sin(\theta_{1x}),$$
$$C_y = m_{dy} n_2 M_{2p} - n_1 M_{1p} = m_{dy} n_2 \sin(\theta_{2y}) - n_1 \sin(\theta_{1y}). \tag{20}$$

Again, we note that for a perfect imaging lens $m_{dx} = m_{dy} = m_p$. The constants of Eq. (20) can then be applied to the ray being traced to find its output direction cosines given its input direction cosines via Eq. (3),

$$L_2 = \frac{n_1 L_1 + C_x}{m_{dx} n_2},$$
$$M_2 = \frac{n_1 M_1 + C_y}{m_{dy} n_2},$$
$$N_2 = \left(1 - L_2^2 - M_2^2\right)^{1/2}. \tag{21}$$

It is now a simple matter to calculate the output ray segment intercept coordinates at the second principal plane via back-propagation from the image point, so

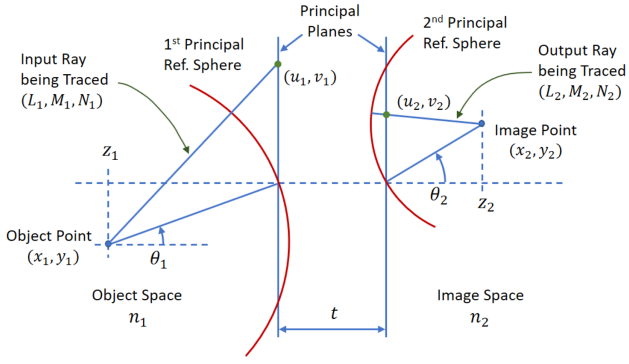

Fig. 5. Cross-sectional diagram showing the input and output ray segments, along with the principal ray segments. The output ray segment being traced is found from application of the generalized sine condition.

$$u_2 = -(L_2/N_2)z_2 + x_2,$$
$$v_2 = -(M_2/N_2)z_2 + y_2. \quad (22)$$

This completes the ray-tracing task. We reiterate that this formulation applies to both real and virtual conjugate planes. (An example of virtual imaging with a negative Cardinal Lens is shown in Supplement 1.)

For the case of a perfect Fourier transform lens, rotation of the lens system to place the conjugate points in the yz plane means $x_1 = x_2 = 0$, and the principal ray segments have $L_{1p} = L_{2p} = 0$, which in turn yields $C_x = 0$. Additionally, the m_d components are found using the approach described in Section 3.A. Otherwise, Eqs. (21) and (22) remain applicable in the rotated coordinate system, and the resulting vectors (L_2, M_2) and (u_2, v_2) can be subsequently transformed back into the original lens orientation via the inverse rotation.

G. Optical Path Length Modeling

Calculating the proper optical path length (OPL) of the ray as it transits the lens (i.e., the OPL assigned to the ray in between the principal planes) requires a few steps. First, the ray OPL is internally computed as the sum of the OPL from the local object point to the first principal plane, plus the OPL from the second principal plane to the local image point—call this sum OPL_1. Second, this process is repeated using the principal ray for the same local object/image points, yielding OPL_0. Last, the final OPL value assigned to the ray being traced equals $\mathrm{OPL}_0 - \mathrm{OPL}_1$. Doing so forces the optical path difference (OPD) between the ray being traced and the principal ray to equal zero. This allows analysis features in OpticStudio that depend on the OPD to function properly (e.g., wavefront map, point spread function, etc.). When a Cardinal Lens is used in isolation at the specified paraxial magnification, no aberrations are introduced, so the OPD is zero and the lens response is diffraction-limited. A similar situation occurs when modeling a more complicated optical train that includes glass elements, but the local object/image planes for the Cardinal Lens happen to correspond to intermediate conjugate planes in the optical system (either real or virtual). In this case the ray aberrations arriving at a Cardinal Lens are simply relayed through the Cardinal Lens with the appropriate magnification—no additional aberrations are generated. However, if a Cardinal Lens is used in a way that deviates from its optimal magnification, then aberrations associated with the lens itself arise; an example is provided in Section 5.D.

4. IMPLEMENTATION IN ZEMAX OPTICSTUDIO

The Cardinal Lens ray-tracing procedure has been implemented in Zemax OpticStudio v23.2 by using a custom user-defined surface. It is written in C++ and compiled as a 64-bit DLL file. The user must supply a value for the lens thickness (zero for a thin lens), along with values for three parameters: (1) the effective focal length, (2) the paraxial magnification for which the lens is aberration-free, and (3) a binary switch value (0/1) to indicate whether the lens is distortion-free with an $f\tan(\theta)$ response, or it operates as a Fourier transform lens with intrinsic $f\sin(\theta)$ distortion. An example of these parameters entered in the Lens Data Editor is shown in Fig. 6.

A. Cardinal Lens Model: Key Features

Here we summarize the key features of the Cardinal Lens DLL implementation:

- The code internally determines the local object and image conjugate plane locations for the Cardinal Lens; it accommodates arbitrary conjugates (i.e., finite conjugate, infinite front conjugate, or infinite rear conjugate—real or virtual).
- The Cardinal Lens may be used following reflection by a mirror (i.e., in mirror space) in which the propagation distances are negative.
- In OpticStudio, the first-order properties of a sequential optical system are found using paraxial ray tracing. Therefore, the Cardinal Lens also supports paraxial ray propagation based on the effective focal length. Using the "Single Ray Trace" analysis feature, the user can see both the real and paraxial ray parameters; various operands available for use in the merit function editor also rely on this paraxial capability.
- As mentioned previously, if the lens is used in an infinite-conjugate arrangement, then appropriate corrections are made to the ray direction cosines (which are initially computed using a long-conjugate approximation) so that the final values are accurate for true infinite-conjugate propagation.

B. Cardinal Lens Model: User Guidelines

In order to use the Cardinal Lens DLL in Zemax OpticStudio, there are a few important guidelines that the user should follow:

- When entering the paraxial magnification, values of $|m_p| \le 1e-10$ are treated as infinite front conjugate, while $|m_p| \ge 1e+10$ corresponds to infinite rear conjugate.
- Proper visualization of the rays in a layout window requires the user to insert a dummy surface immediately prior

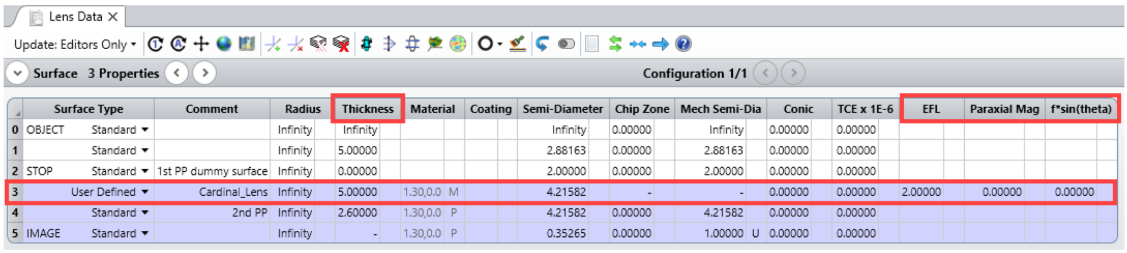

	Surface Type		Comment	Radius	Thickness	Material	Coating	Semi-Diameter	Chip Zone	Mech Semi-Dia	Conic	TCE x 1E-6	EFL	Paraxial Mag	f*sin(theta)
0	OBJECT	Standard		Infinity	Infinity			Infinity	0.00000	Infinity	0.00000	0.00000			
1		Standard		Infinity	5.00000			2.88163	0.00000	2.88163	0.00000	0.00000			
2	STOP	Standard	1st PP dummy surface	Infinity	0.00000			2.00000	0.00000	2.00000	0.00000	0.00000			
3		User Defined	Cardinal_Lens	Infinity	5.00000	1.30,0.0 M		4.21582	-	-	0.00000	0.00000	2.00000	0.00000	0.00000
4		Standard	2nd PP	Infinity	2.60000	1.30,0.0 P		4.21582	0.00000	4.21582	0.00000	0.00000			
5	IMAGE	Standard		Infinity	-	1.30,0.0 P		0.35265	0.00000	1.00000 U	0.00000	0.00000			

Fig. 6. Example of the Cardinal Lens user-defined DLL surface (with its three input parameters) in the OpticStudio Lens Data Editor. The thickness of the Cardinal Lens surface corresponds to the spacing between the principal planes.

to the Cardinal Lens surface. This dummy surface should have zero thickness and use the same material as the surface prior to it. Therefore, the refractive index of the dummy surface corresponds to the index for the front space of the lens.

- The material choice for the Cardinal Lens surface itself establishes the refractive index for the rear space of the lens. Therefore, the material choice for the surface immediately following the lens surface should be the same as that for the lens surface.
- The thickness of the Cardinal Lens surface corresponds to the separation between the principal planes. If the Cardinal Lens resides in mirror space, then this thickness should be negative.
- The user is advised to select the "Hide Rays To This Surface" drawing option for the Cardinal Lens surface in order to suppress the drawing of connecting lines between ray intercepts points on the principal planes. These connecting lines do not represent actual ray paths and are simply artifacts due to the way the OpticStudio layout window works.
- For an optical train containing a Cardinal Lens, the system stop can in principle be located at any physically meaningful location, either before or after the Cardinal Lens. However, if the stop is located at or behind the second principal plane, then paraxial ray aiming should be turned on, and utilization of the pupil compress feature may be needed in order to shrink the paraxial entrance pupil used for the initial condition of the iterative ray aiming algorithm. It appears that the most robust system aperture type is "Float by Stop Size." The stop size may also need to be reduced to ensure there actually exist rays that can propagate through the lens to the edge of the stop.
- Unfortunately we find that the ray aiming algorithm may fail to work in certain circumstances when using the Cardinal Lens. The ray aiming algorithm details are not described in the OpticStudio user documentation, so it is hard to know precisely why failure occurs. The goal of ray aiming is to produce a desired spatial distribution of rays that fills the stop based on the apodization setting (e.g., uniform). Also, we note that for a typical surface in OpticStudio an output ray leaves the surface from the same location that the input ray strikes the surface. However, for the Cardinal Lens this is not the case, and, moreover, the ray intercept mapping between the principal planes is nonlinear. It appears that if this mapping becomes sufficiently nonlinear and the outgoing rays have high NA, then the ray aiming algorithm simply fails to work properly. *Therefore, the user is strongly encouraged to try and construct the optical path so that the stop is placed in front of or on the first principal plane of the Cardinal Lens.* Doing so should yield a robust solution without encountering the aforementioned difficulties.

5. EXAMPLES

In this section, we provide a few examples that illustrate the use of the Cardinal Lens in Zemax OpticStudio via our user-defined surface, Cardinal_Lens.dll. Both the DLL and the example models are provided in Code 1, Ref. [12].

A. Infinite-Conjugate High-NA Immersion Objective

First, consider the case of imaging at infinity ($m_p = 0$) with a high-NA immersion lens. In Fig. 7 we compare a paraxial lens to a thin Cardinal Lens, both having an effective focal length (EFL) of 5 mm and no aberrations. The object space is air ($n_1 = 1.0$), and the image space is water ($n_2 = 1.3$). The field angles are 0° and 20°. The aperture is set to an entrance pupil diameter (EPD) of 10 mm. For the Cardinal Lens, the paraxial magnification parameter is zero as appropriate for this infinite front conjugate case. Even though both lenses have the same EPD and EFL, the focused beam NAs are quite different. For the paraxial lens the NA = 0.80, while for the Cardinal Lens it is 1.0. As a result, diffraction-related quantities such as the Airy disk diameter and the modulation transfer function cutoff frequency are also different for the two models. Clearly, if one were trying to model a well-corrected microscope objective, the Cardinal Lens version would provide a more realistic option. In OpticStudio, the "Image Space NA" is taken to be a paraxial quantity; the actual NA based on a real marginal ray should be calculated as one over two times the "Working F/#." We note for the Cardinal Lens with an object at infinity, the relation EPD = 2 · NA · EFL is generally applicable.

It is also interesting to compare the value of the "Offense against the Sine Condition" (OSC) for the two lenses of Fig. 7. For an actual lens, the OSC is a dimensionless quantity equal to the relative sagittal coma [18]. If spherical aberration is negligible, and the OSC assumes a very small value, then linear coma is also negligible for small off-axis fields. This is typically the

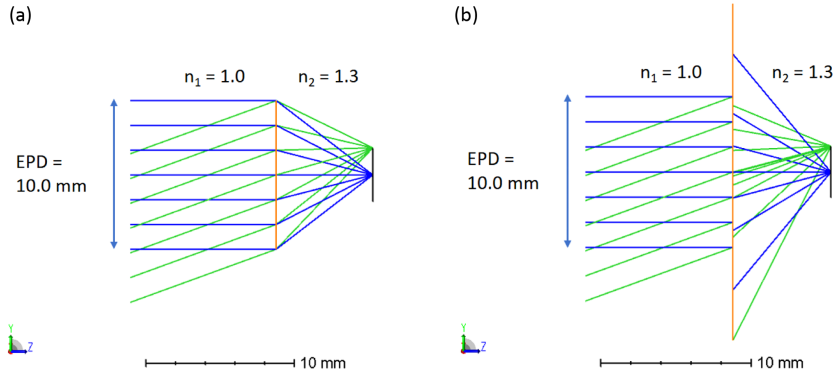

Fig. 7. Ray-trace layouts for imaging at infinity with an immersion objective, showing a comparison between (a) a paraxial thin lens with NA = 0.80 and (b) a thin Cardinal Lens with NA = 1.0. Both lenses have an EFL = 5 mm and an EPD = 10 mm, but yield different numerical apertures. Field angles of 0° and 20° are shown.

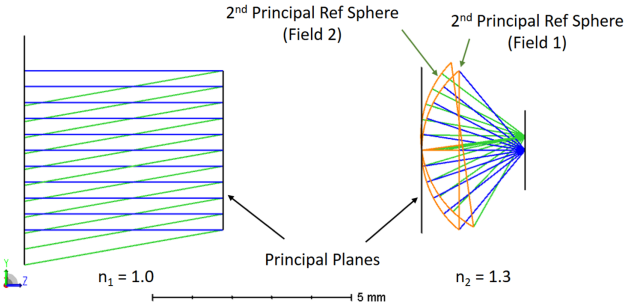

Fig. 8. Modified version of the Cardinal Lens shown in Fig. 7(b) to make it a thick lens. Also, for each of the two incoming field angles (0° and 10°), the corresponding output principal reference spheres have been added to the model. The drawing options available in OpticStudio have been used to suppress rays in between the first principal plane and the second principal reference spheres.

case for well-corrected objectives. In our example here, both lenses yield aberration-free focusing by design, so elimination of linear coma does not come into play, but nonetheless the OSC value provides a measure of how closely a lens follows the sine condition. To compare the paraxial and Cardinal lenses, we use a simplified version of the OSC $= (u_2/\sin U_2) - 1$, where u_2 is the paraxial marginal ray slope and U_2 is the real marginal ray angle, both taken in image space. The paraxial lens has OSC $= 0.261$, while the Cardinal Lens has OSC $= 0$. In passing, we should mention that OpticStudio provides a merit function operand (OSCD) that calculates a more general version of the OSC for imaging an object at infinity (based on [18], Eq. 9.41).

If desired, the user can simulate a thick version of the Cardinal Lens. For example, Fig. 8 shows a 5 mm gap in between the principal planes, obtained by simply setting the thickness of the Cardinal Lens surface to 5 mm. In addition, the second principal reference spheres can be manually added to the model and displayed in the layout as shown in Fig. 8. In this case we use two configurations, one for each of the two fields, and display both configurations simultaneously. For a given field the corresponding principal reference sphere radius and tilt angle, which correspond to the output principal ray segment length and angle, can be found via the merit function and then applied to a corresponding dummy spherical surface via the multi-configuration editor. The ray drawing options can be used to suppress the rays in between the input principal plane and the output principal reference spheres.

B. Microscope Objective with Tube Lens

When combining a Cardinal objective lens with a glass tube lens, both the ray and wave aberrations associated with the tube lens are apparent. Figure 9 shows the layout for an immersion objective with an effective focal length of 5 mm and an object-space index of 1.5, operating at NA $= 1.3$, followed by a simple doublet tube lens having a focal length of 200 mm. This creates a simple microscope system having a magnification of -40. In this case the system aperture stop is placed on the first principal plane of the Cardinal objective lens. The figure shows rays from three field points (0, 0.25, and 0.50 mm). The principal planes of the objective are separated by 10 mm. The objective lens has

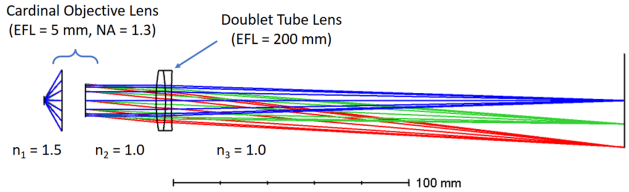

Fig. 9. Simple 40× microscope system based on a 10-mm-thick Cardinal Lens objective (EFL $= 5$ mm, NA $= 1.3$) and a doublet tube lens (EFL $= 200$ mm). Field points with y coordinates of 0, 0.25, and 0.50 mm are shown.

its paraxial magnification parameter set to $-1.0e10$, which the DLL interprets as infinite rear conjugate. The ray aberrations across the visible spectrum for the 0.5-mm field point, arising solely from the tube lens, are clearly observed in the spot diagram of Fig. 10(a). The OPD fans of Fig. 10(b) show the associated wavefront aberrations.

C. Fourier Transform Lens: One- and Two-Element Versions

A Cardinal Lens can function as a perfect Fourier transform (FT) lens. This is done in OpticStudio by setting the $f\sin(\theta)$ distortion switch in the lens data editor (Parameter 3 for the lens surface) to a value of 1, which ensures the incoming angular distribution of rays is mapped to the proper spatial distribution of ray intercept points across the back focal plane. In a coherent optical system, the back focal plane serves as the Fourier plane. The focused spots in the Fourier plane have (x, y) coordinates that are directly proportional to the (L, M) direction cosines of the collimated beams, or angular spectrum components, leaving the front focal plane, with the constant of proportionality being the front focal distance ([13], Ch. 2). Of course an optical FT is a diffraction phenomenon and should therefore be simulated using wave optics (i.e., physical optics propagation, or POP, in OpticStudio). However, from a ray-tracing perspective, an ideal FT lens should provide the proper angle-to-spatial mapping as well as the correct focusing NA, consistent with the generalized sine condition, which is what we demonstrate here.

Figure 11(a) shows a Cardinal Lens configured as a thin Fourier transform (FT) lens in air with f/1.0 and EFL $= 10$ mm. The aperture stop is located in the front focal plane, so this is an image-space telecentric configuration. The plot of Fig. 11(b) illustrates the desired distortion.

A somewhat more complicated version of a compound FT lens in air using two separate Cardinal lenses is shown in Fig. 12. In this case each lens has an EFL $= 50$ mm, is 2 mm thick, and the gap between them is 2 mm. The EFL of this lens assembly is 25.51 mm. The first lens operates in infinite-front-conjugate mode with a paraxial magnification of zero, and its distortion switch is set to 1, corresponding to an $f\sin(\theta)$ mapping as required for the FT operation, where f is the front focal length of this first lens (note: in air this equals the EFL of the lens). The second lens operates in finite-conjugate perfect-imaging mode and applies no additional distortion. The appropriate magnification and back focal distance for the second lens are found by setting those values as variables and optimizing for minimum focused spot size. The distortion associated with the first Fourier lens is scaled by the magnification of the second

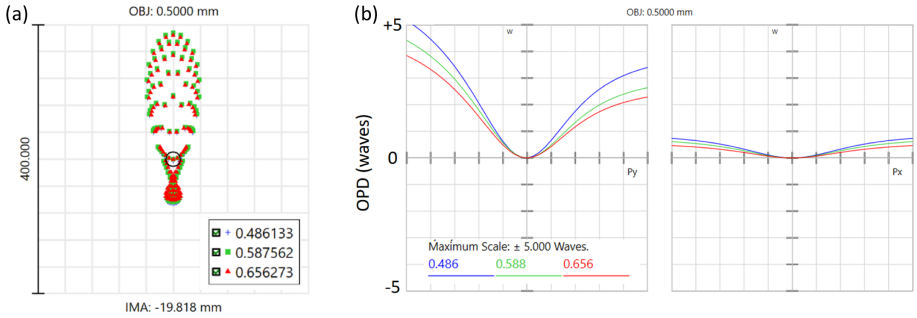

Fig. 10. Aberrations for the 0.5-mm off-axis field point of the microscope system shown in Fig. 9, which are solely attributable to the tube lens. (a) Spot diagram showing transverse ray aberrations, and (b) OPD curves (plotted on a scale from −5 to +5 waves) representing the wavefront aberrations.

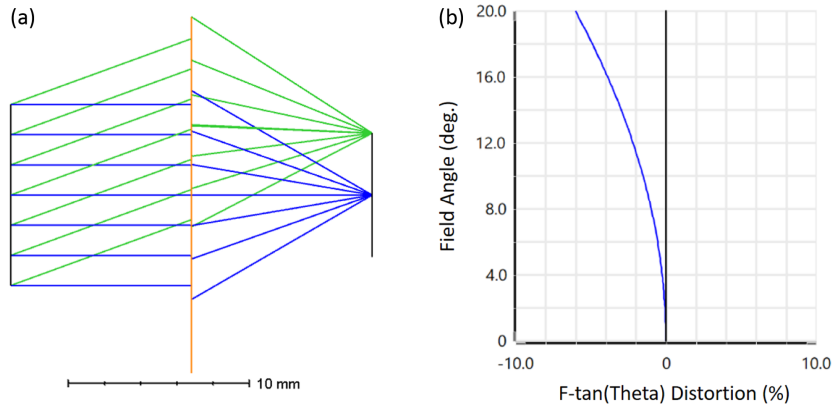

Fig. 11. Example of a thin Cardinal Lens used in Fourier transform mode (in air with f/1.0 and EFL = 10 mm). (a) Ray-tracing layout for field angles of 0° and 20° and (b) corresponding distortion plot showing the $f\sin(\theta)$ response.

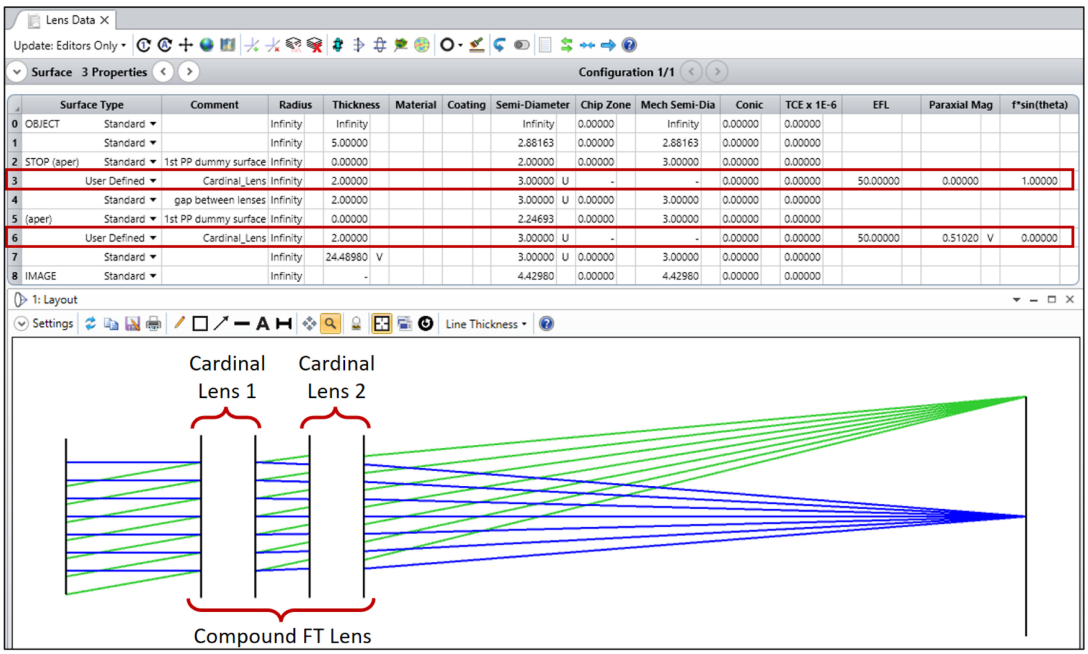

Fig. 12. Example of a compound FT lens using two Cardinal lenses, each having a thickness of 2 mm, with the lenses separated by 2 mm. The optical path as constructed in the OpticStudio lens data editor is shown with the Cardinal Lens surfaces highlighted.

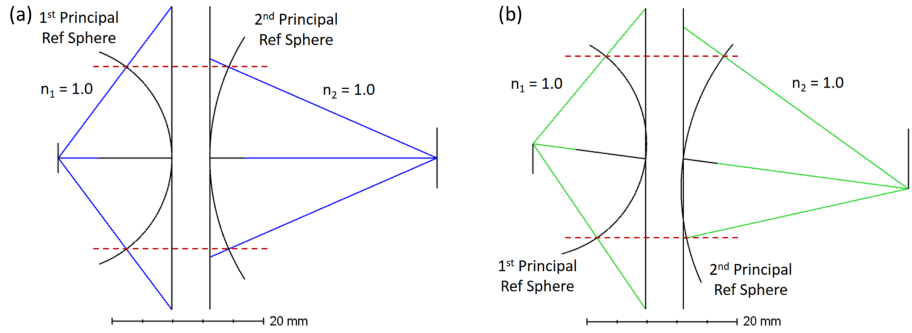

Fig. 13. Ray-trace layouts for finite-conjugate imaging with a Cardinal Lens in air (EFL = 10 mm, $m_p = -2.0$, $\mathrm{NA}_1 = 0.80$, $\mathrm{NA}_2 = 0.40$, $t = 5$ mm). Two fields are shown: (a) on-axis imaging and (b) off-axis imaging with $y_1 = 2$ mm and $y_2 = -4$ mm. Because both the image space and object space have the same refractive index, the upper and lower marginal rays each cross the two principal reference spheres at the same distance from the optical axis as shown by the red dashed lines.

imaging lens such that the compound lens provides the correct FT angle-to-spatial mapping based on the front focal length of the compound lens assembly.

We conclude this section by mentioning that if two Cardinal FT lenses are used for coherent imaging in a 4f configuration, then an object in the form of a grating can be properly imaged ([13], p. 21), whereas using two paraxial lenses leads to an image having the wrong period; an example of this is shown in Supplement 1.

D. Finite-Conjugate Imaging

In this section we illustrate a Cardinal Lens used for finite-conjugate imaging, first at its optimal aberration-free magnification; then we next show how aberrations arise when the magnification deviates from the optimal value. Figure 13 shows the optimal case with a 5-mm-thick lens used by itself, and the layout magnification is the same as that of the user-specified value ($m_p = -2.0$), so no aberrations are present. The diffraction-limited Airy disk radius is 0.76 μm. The ratio of the object-space to image-space numerical apertures equals the paraxial magnification, or $\mathrm{NA}_1/\mathrm{NA}_2 = m_p$. Also, because both the object and image spaces have the same refractive index, the upper and lower marginal rays each individually cross the two principal reference spheres at the same radial distance from the optical axis ([13], p. 15). If the object/image-space refractive indices are different, then this simple geometrical result is no longer exact, but instead only holds approximately.

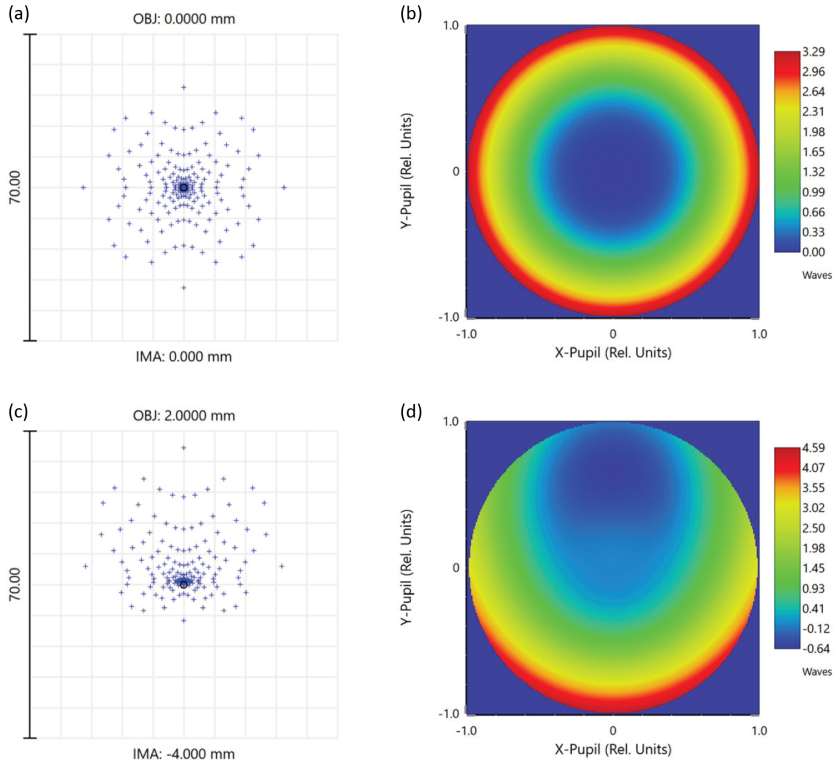

Fig. 14. Aberrations arise when using the Cardinal Lens away from its optimal magnification value. Shown are results for the lens of Fig. 13 when $m_p = -2.01$. (a) On-axis spot diagram, (b) on-axis wavefront map, (c) spot diagram for the 2-mm off-axis field point, and (d) wavefront map for the 2-mm off-axis field point. The wavelength is 0.50 μm.

If the lens of Fig. 13 is used at a magnification other than the designated optimum value, significant aberrations can arise. For example, if the paraxial magnification specification of the lens is changed to $m_p = -2.01$, but we retain the same object and image plane locations, the resulting spot diagrams and wavefront maps are shown in Fig. 14. In this case, for both fields, the RMS spot radius is ~10 μm and the RMS wavefront error is ~1 wave (with tilt removed from the off-axis field), so the lens is now operating far outside of the diffraction limit. From Eq. (6) we have $z_1 = -14.975$ mm and $z_2 = 30.10$ mm, while the layout object and image plane locations remain at −15.0 mm and 30.0 mm, respectively, in accord with a layout magnification of −2.00.

E. Ray and Wavefront Aberrations for the Cardinal Lens

As a final example, it is instructive to consider how both transverse ray aberrations and wavefront aberrations propagate through the Cardinal Lens and to assess their self-consistency with one another, which serves as a check as to whether or not the lens complies with Fermat's principle. Figure 15(a) shows a biconvex lens operating with an object-space NA = 0.23 and a magnification of −1. The first lens surface is selected as the stop surface, and the aperture size is set via the specified NA value. The marginal ray suffers from significant aberration as seen in the associated OPD fan of Fig. 15(b) with a value of approximately 100 waves. Now consider, as shown in Fig. 15(c), a Cardinal Lens used to relay this image with a magnification of −1. Similarly, Fig. 15(e) shows a paraxial lens used for the relay. For both cases the transverse ray aberrations are identical, but the two corresponding OPD fans [Figs. 15(d) and 15(f)] are quite different. The Cardinal Lens version has a maximum OPD of 500 waves, while the paraxial lens imparts only 50 waves. In OpticStudio, various modes exist for computing the OPD through a paraxial lens; here we use Mode = 1, which provides the highest degree of accuracy.

The OPD is measured with respect to a reference sphere that is centered on a point where the chief ray intersects the image plane. In OpticStudio this reference sphere surface is constructed so it contains the point where the chief ray crosses the exit pupil plane. We expect there to be a change in the OPD when adding a relay lens because the location of the exit pupil changes relative to the new image plane; hence the radius of the reference sphere changes. However, for both the Cardinal and paraxial relays the exit pupil location is the same [specifically, −6.6 mm from the final image plane in Figs. 15(c) and 15(e)]. So which OPD curve is correct? To answer this question, we can employ the following well-established relations ([1], p. 206) linking the transverse ray aberrations to the OPD, which is

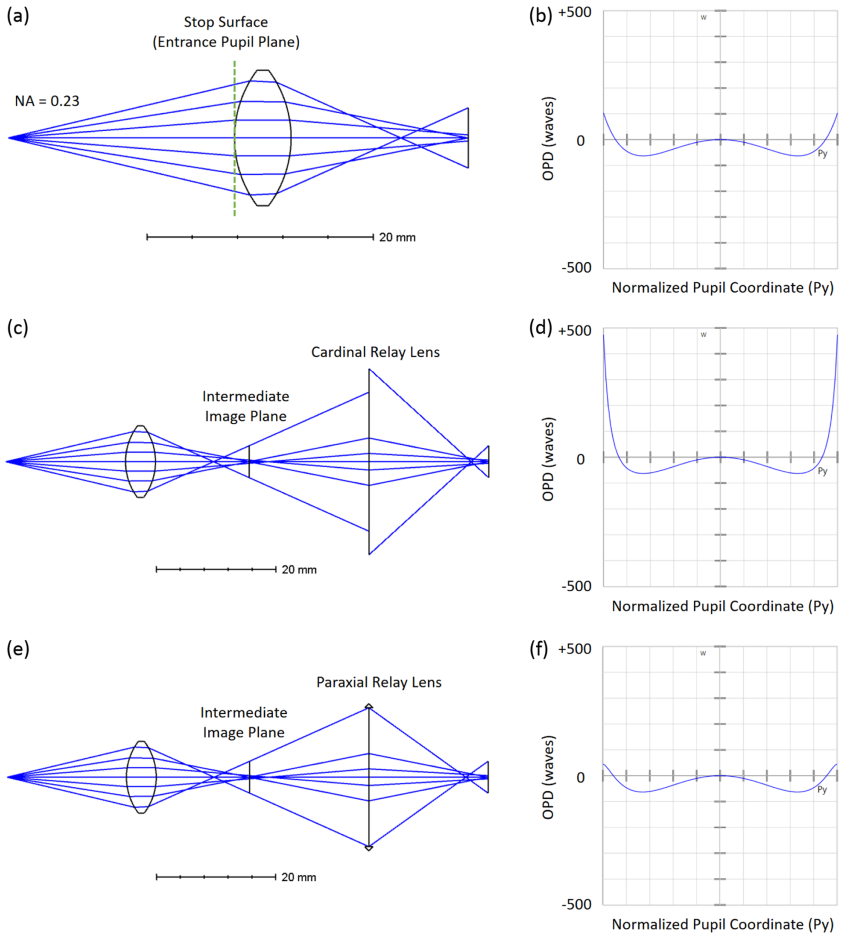

Fig. 15. Biconvex lens OPD and the change in OPD when using a relay lens. (a) Biconvex lens with 1:1 imaging and (b) its OPD fan. (c) Relay imaging with a Cardinal Lens and (d) resulting OPD fan. (e) Relay imaging with a paraxial lens and (f) corresponding OPD fan.

denoted by the function $\Phi(x_p, y_p)$:

$$\Delta x = \frac{R'}{n}\frac{\partial \Phi}{\partial x_p},$$

$$\Delta y = \frac{R'}{n}\frac{\partial \Phi}{\partial y_p}. \tag{23}$$

Here R' is the distance from a ray intercept point on the image plane to its intercept point on the reference sphere, (x_p, y_p) are exit pupil coordinates on the reference sphere, and n is the index in image space. When the aberrations are small, then R' can be well approximated by the radius of the reference sphere. However, in our case the aberrations are significant, so we must use the exact value for R'. Equation (23) is derived using Hamiltonian optics (specifically Hamilton's point characteristic function) and therefore intrinsically applies to rays obeying Fermat's principle.

To proceed, we use a custom macro in OpticStudio to trace a tangential fan of rays spanning the entrance pupil and record the following data for each of the two relay lens options: y_p, Δy, $\Phi(y_p)$, and $R'(y_p)$. Using Eq. (23) we calculate Δy from the derivative of $\Phi(y_p)$ and compare this to the result found directly from ray tracing as shown in Fig. 16. The ray fan computed using the Cardinal Lens relay OPD agrees precisely with the ray fan from direct ray tracing, while the paraxial lens relay OPD leads to ray fan errors at the edges of the pupil (for $|P_y| > 0.9$). Therefore, the OPD in the Cardinal Lens case is correct as it is fully consistent with the transverse ray aberration, which is another indicator that the Cardinal Lens complies with Fermat's principle.

F. Summary of Additional Examples in the Supplemental Document

Other examples involving the Cardinal Lens are provided in Supplement 1, including (a) virtual imaging with a negative Cardinal Lens, (b) its use in mirror space, (c) placing the system stop at or behind the Cardinal Lens output principal plane in which case ray aiming is required, (d) quantitative evaluation of spherical aberration when the lens is not used at the prescribed magnification, including comparison to analytical results as reported by Walther ([10], Sec. 31.2), (e) comparisons of the Cardinal and paraxial lenses via 4f coherent imaging of a grating, and (f) a zoom lens example that includes a discussion of how to convert a paraxial lens to a Cardinal Lens. We also include a table summarizing various configurations in which the Cardinal Lens can be used, along with the associated paraxial magnification values.

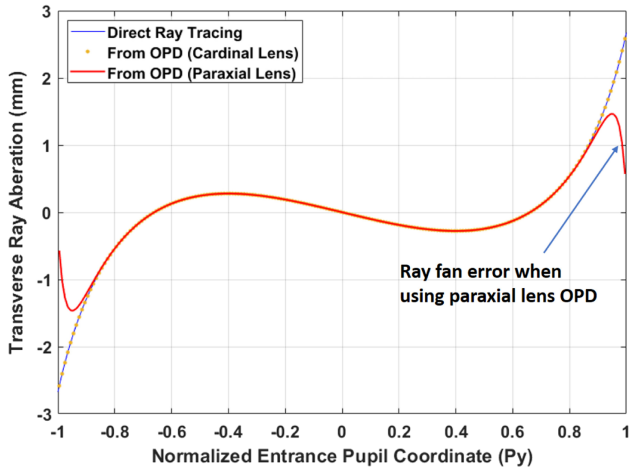

Fig. 16. Comparison of ray fans based on direct ray tracing as well as computed from the OPD fans for the Cardinal Lens and paraxial lens relay cases [per Eq. (23)].

6. CONCLUSION

This paper presents a ray-tracing model for a perfect thick lens, referred to as the Cardinal Lens, that satisfies the generalized sine condition and is therefore consistent with Fermat's principle, but its formulation does not require the use of an eikonal function to represent the lens. When simulating a well-corrected lens, this model provides a more accurate alternative compared to use of a simple paraxial thin lens. Two versions of the model are described, one being a distortion-free $f\tan(\theta)$ lens for imaging applications, and the other having $f\sin(\theta)$ distortion as appropriate for a Fourier transform lens. In both cases, the lens is free of aberration only when used at its prescribed (paraxial) magnification. A version of the Cardinal Lens suitable for use in Zemax OpticStudio has been coded as a user-defined surface DLL, and a variety of ray-tracing examples using the DLL have been provided (see Code 1, Ref. [12]). It is hoped that this model will be of some use to the lens design community as a relatively straightforward means to mimic well-corrected compound lenses for which detailed prescription data are not available.

APPENDIX A: FOURIER LENS DIFFERENTIAL MAGNIFICATION

Consider a finite-conjugate imaging configuration based on a single Cardinal Lens immersed in air, operating in Fourier lens mode. For an object point located on the y axis at $(0, y_1)$, a distance z_1 from the lens, the image point has a y coordinate given by Eq. (16):

$$y_2 = z_2 \sin(\theta_1) = z_2 y_1 \left(y_1^2 + z_1^2\right)^{-1/2}, \tag{A1}$$

where θ_1 is the angle of the incoming principal ray with respect to the optical axis. Due to this intrinsic distortion, the two conjugate points are related by a field-dependent transverse magnification, $m(y_1)$, such that $y_2 = m(y_1)y_1$. Comparison with Eq. (A1) shows

$$m(y_1) = z_2\left(y_1^2 + z_1^2\right)^{-1/2}. \tag{A2}$$

For an infinitesimal y displacement of the object point, dy_1, we have

$$y_2 + dy_2 = m(y_1 + dy_1)(y_1 + dy_1)$$

$$\simeq \left[m(y_1) + \frac{\partial m(y_1)}{\partial y_1} dy_1\right](y_1 + dy_1). \tag{A3}$$

For notational convenience going forward, we drop the explicit field dependence of the magnification. The differential image point displacement becomes, to first order in dy_1,

$$dy_2 = \left[m + \frac{\partial m}{\partial y_1} y_1 \right] dy_1, \tag{A4}$$

from which the y-direction differential magnification follows:

$$m_{dy} \equiv \frac{dy_2}{dy_1} = m + \frac{\partial m}{\partial y_1} y_1. \tag{A5}$$

From Eq. (A2) we have

$$\frac{\partial m}{\partial y_1} = -z_2 y_1 \left(y_1^2 + z_1^2 \right)^{-3/2} = -m y_1 \left(y_1^2 + z_1^2 \right)^{-1}. \tag{A6}$$

Substitution into Eq. (A5) yields

$$m_{dy} = m \left[1 - y_1^2 \left(y_1^2 + z_1^2 \right)^{-1} \right] = m \left[1 - \sin^2 \theta_1 \right]. \tag{A7}$$

For small principal ray angles, $m_{dy}(y_1)$ is approximately equal to the conventional transverse magnification, $m(y_1)$, but as the angle θ_1 increases (i.e., $|y_1|$ increases), the value of $m_{dy}(y_1)$ becomes progressively smaller than $m(y_1)$. Finally, we note that for this configuration, the x-direction differential magnification remains equal to the transverse magnification, or $m_{dx}(y_1) = m(y_1)$. As a result, a Fourier lens is locally anamorphic.

Acknowledgment. The author gratefully acknowledges the valuable assistance of Prof. José Sasián at the University of Arizona for reviewing of draft of this paper and providing very helpful feedback, including, among other things, recommendations to use a unique name for the lens model and to have it support thick lenses, which initially was not the case. The author also thanks one of the anonymous reviewers for pointing out the need to properly treat the locally anamorphic nature of a Fourier lens.

Disclosures. The author declares no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the author upon reasonable request.

Supplemental document. See Supplement 1 for supporting content.

REFERENCES AND NOTES

1. M. Born and E. Wolf, *Principles of Optics*, 6th ed. (Pergamon, 1980).
2. J. Bentley and C. Olson, *Field Guide to Lens Design* (SPIE, 2012), p. 32.
3. R. Luneburg, *Mathematical Theory of Optics* (University of California, 1966), pp. 182–187.
4. H. Gross, ed., *Handbook of Optical Systems* (Wiley, 2005), Vol. **1**.
5. K. von Bieren, "Lens design for optical Fourier transform systems," Appl. Opt. **10**, 2739–2742 (1971).
6. Zemax OpticStudio, Ansys Inc., https://www.ansys.com.
7. "CODE V optical design software," Synopsys, Inc., https://www.synopsys.com/optical-solutions/codev.html.
8. T. G. Kuper and M. P. Rimmer, "Lens modules in optical design," Proc. SPIE **0892**, 140–151 (1988).
9. Lambda Research Corporation, "OSLO optics software for layout and optimization," (Optics Reference, Revised 9-March-2021), pp. 165–169, https://lambdares.com/oslo.
10. A. Walther, *The Ray and Wave Theory of Lenses* (Cambridge University, 1995).
11. The choice for the name Cardinal Lens stems, in part, from the fact that Stanford University is well-known by its school color, which is cardinal. Also, it is a play on the fact that the lens principal planes and principal reference spheres are central to the model, and they are important cardinal surfaces.
12. J. P. Wilde, "Cardinal Lens publication code files (Zemax OpticStudio DLL and models)," figshare, 2024, https://doi.org/10.6084/m9.figshare.24295720.
13. M. Mansuripur, *Classical Optics and its Applications*, 2nd ed. (Cambridge University, 2009).
14. T. Smith, "The optical cosine law," Trans. Opt. Soc. **24**, 31–39 (1922).
15. M. Herzberger, *Modern Geometrical Optics* (R. E. Krieger, 1980), p. 160.
16. H. Buchdahl, *An Introduction to Hamiltonian Optics* (Dover, 1993), pp. 93–94.
17. T. T. Elazhary, P. Zhou, C. Zhao, *et al.*, "Generalized sine condition," Appl. Opt. **54**, 5037–5049 (2015).
18. W. T. Welford, *Aberrations of Optical Systems* (Adam Hilger, 1986), pp. 172–176.