

ZEMAX

Optical Design Program

User's Guide

Version 10.0

IMMERVISION Ex. 2011
LG v. ImmerVision
IPR2020-00179

SIMON TRIBUIT

INV

ZEMAX[®]

Optical Design Program

User's Guide

Version 10.0

Focus Software, Incorporated
P. O. Box 18228
Tucson, AZ 85731-8228 USA
Tel. (520) 733-0130
Fax. (520) 733-0135
support@focus-software.com
www.focus-software.com

Important notice

ZEMAX® is a registered trademark of Focus Software, Inc. Copyright © Focus Software, Incorporated 1990-2001. All rights reserved.

All other product names or trademarks are property of their respective owners.

Information in this manual is subject to change without notice and does not represent a commitment on the part of the vendor. The software described in this manual is furnished under a license agreement and may be used or copied only in accordance with the terms of the agreement.

The material in this manual is copyrighted and may not be reproduced or electronically distributed without written permission from the copyright holder.

Focus Software, Incorporated (FSI) provides this publication "as is" without warranty of any kind, either express or implied, including but not limited to the implied warranties or conditions of merchantability or fitness for a particular purpose. In no event shall FSI be liable for any loss of profits, loss of business, loss of use or data, interruption of business, or for indirect, special, incidental, or consequential damages of any kind, even if FSI has been advised of the possibility of such damages arising from any defect or error in this publication or in the Software.

Published in April, 2001.

Focus Software, Incorporated
10001 West 13th
Tulsa, Oklahoma 74117
Phone: 918-439-1111
Fax: 918-439-1112
E-mail: info@focussoftware.com
www.focussoftware.com

Introduction

This chapter describes conventions and defines terminology used throughout this manual. Most of the conventions and terms ZEMAX uses are common in the optics industry, however there may be some important differences.

Active configuration

The active configuration is the configuration currently being displayed in the lens data editor. For details see the Chapter "Multi-Configurations".

Angular magnification

The ratio of the paraxial image space chief ray angle to the paraxial object space chief ray angle. The angles are measured with respect to the paraxial entrance and exit pupil locations.

Apodization

Apodization refers to the uniformity of illumination in the entrance pupil of the system. By default, the pupil is always illuminated uniformly. However, there are times when the pupil should have a non-uniform illumination. For this purpose, ZEMAX supports pupil apodization, which is a variation of amplitude over the pupil.

Three types of pupil apodization are supported: uniform, Gaussian, and tangential. For each type (except uniform), an apodization factor determines the rate of variation of amplitude in the pupil. See the discussion on apodization types and factors in the chapter "System Menu".

ZEMAX also supports user defined apodizations, which may be placed on any surface. Surface apodizations behave differently than pupil apodizations, because surfaces need not be located at a pupil. For more information on surface apodizations, see the "Surface Types" chapter under "The User Defined Surface" section.

Back focal length

ZEMAX defines the back focal length as the distance along the Z axis from the last surface made of glass to the paraxial image plane. If no surfaces are made of glass, the back focal length is the distance from surface 1 to the paraxial image plane.

Cardinal planes

The term cardinal planes (sometimes called cardinal points) refers to those special conjugate positions where the object and image surfaces have a specific magnification. The cardinal planes include the principal planes, where the magnification is +1, the nodal planes, where the angular magnification is +1, the anti-nodal planes, where the angular magnification is -1, and the focal planes, where the magnification is 0 for the image space focal plane and infinite for the object space focal plane.

Except for the focal planes, the cardinal planes are conjugates with each other, that is, the image space principal plane is conjugate with the object space principal plane, etc. If the lens has the same index in both object space and image space, the nodal planes are identical to the principal planes.

ZEMAX lists the distance from the image surface to the various image space planes, and lists the distance from the first surface to the various object space planes.

Chief ray

If there is no vignetting, and there are no aberrations, the chief ray is defined to be the ray that travels from a specific field point, through the center of the entrance pupil, and on to the image plane. Note that without vignetting or aberrations, any ray passing through the center of the entrance pupil will also pass through the center of the stop and the exit pupil.

When vignetting factors are used, the chief ray is then considered to be the ray that passes through the center of the vignetted pupil, which means the chief ray may not necessarily pass through the center of the stop.

If there are pupil aberrations (and there virtually always are), then the chief ray may pass through the center of the paraxial entrance pupil (if ray aiming is not used) or the center of the stop (if ray aiming is used), but generally, not both.

If there are vignetting factors which decenter the pupil, then the chief ray will pass through the center of the vignetted entrance pupil (if ray aiming is not used) or the vignetted stop surface (if ray aiming is used).

The common convention used is that the chief ray passes through the center of the vignetted pupil, while the principal ray passes through the center of the unvignetted stop. ZEMAX never uses the principal ray. Most calculations are referenced to the chief ray or the centroid. Note the centroid reference is generally superior because it is based upon the aggregate effect of all the rays that actually illuminate the image plane, and not on the arbitrary selection of one ray which is "special".

Coordinate axes

The optical axis is the Z axis, with the initial direction of propagation from the object being the positive Z direction. Mirrors can subsequently reverse the direction of propagation. The coordinate system is right handed, with the sagittal X axis being oriented "into" the monitor on a standard layout diagram. The tangential Y axis is vertical.

The direction of propagation is initially left-to-right, down the positive Z axis. After an odd number of mirrors the beam physically propagates in a negative Z direction. Therefore, all thicknesses after an odd number of mirrors should be negative.

Diffraction limited

The term diffraction limited implies that the performance of an optical system is limited by the physical effects of diffraction rather than imperfections in either the design or fabrication. A common means of determining if a system is diffraction limited is to compute or measure the optical path difference. If the peak to valley OPD is less than one quarter wave, then the system is said to be diffraction limited.

There are many other ways of determining if a system is diffraction limited, such as Strehl ratio, RMS OPD, standard deviation, maximum slope error, and others. It is possible for a system to be considered diffraction limited by one method and not diffraction limited by another method.

On some ZEMAX plots, such as the MTF or Diffraction Encircled Energy, the diffraction limited response is optionally shown. This data is usually computed by tracing rays from a reference point in the field of view. Pupil apodization, vignetting, F/#'s, surface apertures, and transmission may be accounted for, but the optical path difference is set to zero regardless of the actual (aberrated) optical path.

For systems which include a field point at 0.0 in both x and y field specifications (such as 0.0 x angle and 0.0 y angle), the reference field position is this axial field point. If no (0, 0) field point is defined, then the field coordinates of field position 1 are used as the reference coordinates instead.

Edge thickness

ZEMAX uses two different definitions for the term "edge thickness". Usually, the edge thickness is computed for a specific surface by:

$$E_i = Z_{i+1} - Z_i + T_i$$

where Z_i is the sag of the surface at the +y semi-diameter of the surface, Z_{i+1} is the sag of the next surface at the +y semi-diameter of the next surface, and T_i is the axial thickness of the surface. Note that the edge thicknesses are computed accounting for the sag at the respective semi-diameter of each surface, which in general are different.

Note also that edge thickness is normally computed for the +y radial aperture, which may be inadequate if the surface is not rotationally symmetric, or if surface apertures have been placed upon either of the surfaces.

The exception to this rule is when computing edge thickness solves. Because the edge thickness solve can change the center thickness, the edge thickness solve can change where rays strike the following surface, which in turn means the semi-diameter of the next surface may change. If the semi-diameter of the next surface is used in the edge thickness computation, an "infinite loop" or circular definition may occur.

For this reason, edge thickness solves compute the edge thickness strictly at the semi-diameter of the first surface, for both surfaces. The semi-diameter of the second surface is never used, although the curvature or shape of the surface is used.

Introduction

This chapter provides detailed descriptions of each of the analysis features ZEMAX supports. Analysis in this context means any graphical or text data computed from data defining the lens. This includes aberrations, MTF, spot diagrams, and many other computations. Program features which modify the lens data or which manipulate other data (such as glass catalog data) are described in the chapter "Tools Menu".

Selecting a menu option will immediately perform the requested calculation. Once the graph or text window is displayed, you may select the Settings menu option to modify the default settings for that window. Once you have made the appropriate changes, click on "OK" and the program will recalculate and redisplay the data presented in the window. If you prefer to change the settings before the graphic or text data is displayed, use the "Show Options First" checkbox on the Graphics tab of the File, Preferences dialog box.

For a description of the OK, Cancel, Save, Load, Reset, and Help buttons present on most of the "Settings" dialog windows, see the chapter "User Interface".

Each analysis window has an "Update" menu item. The update function forces ZEMAX to recompute and redisplay the data presented in the window. This is useful if the lens data has changed and the graph now displays obsolete data. Double clicking within the window has the same effect as selecting Update. Clicking with the right mouse button is equivalent to clicking on "Settings". For more information, see the chapter "User Interface".

Layout

2D Layout

Purpose:

Layout diagram. This is a simple YZ cross section through the lens.

Settings:

Item	Description
First Surface	The first surface to be drawn.
Last Surface	The last surface to be drawn.
Wavelength	Either any one or all wavelengths may be shown.
Field	Either any one or all field positions may be shown.
Number of Rays	The number of rays specifies the number of tangential rays to be drawn for each defined field. The rays will be evenly distributed along the fan of the pupil, unless apodization has been specified. This parameter may be set to zero.
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the range of surfaces drawn to fill the graphic page. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the lens actual size on the printer (not the screen). A factor of 0.5 will plot the lens at half scale.
Upper Pupil Limit	The maximum pupil coordinate to draw rays to.
Lower Pupil Limit	The minimum pupil coordinate to draw rays to.
Marginal and Chief Only	Draws only the marginal and chief rays, overriding the other ray settings.
Square Edges	Draws flat faces and edges if selected, otherwise uses the semi-diameter value to draw lens edges.

Item	Description
DXF File	In this text cell enter the name of the DXF format file to use. This option is only used if the button "Export As DXF File" is subsequently pressed. The file will be stored in the default directory for output.
Export As DXF File	If pressed, this button will cause the generation of a DXF formatted file which will contain the same data that is displayed in the graphic window. The file name is given by the "DXF File" option. The DXF file generated is a 2D model of the lens system suitable for importing into CAD programs that can read DXF files. See the discussion section for details.
Color Rays By	Select "Fields" to use color to distinguish between each field position, or "Waves" to distinguish between each wavelength.
Suppress Frame	Suppresses drawing of the frame on the bottom of the screen, which leaves more room for the layout plot itself. No scale bar, address block, or other data will be displayed.
Delete Vignetted	If checked, rays are not drawn if they will be vignetted by any surface.
Fletch Rays	If checked, small arrows are drawn on each ray to indicate the direction of propagation.

Discussion:

This feature is not available if you use coordinate breaks, spider obscurations, obscuration decenters, X-angles, holograms, or other attributes which spoil the rotational symmetry of the lens. Use the 3D layout instead.

The "Export As DXF File" button will generate a 2D DXF file and store it in the file name provided in the "DXF File" data field. The DXF file will consist of arcs and lines. The arcs are used to show the shape of curved lens faces. If only spherical (or plane) lenses are used, then the arc is a perfect representation of the lens. However, the arcs are only approximations to non-spherical surfaces. If the surface sag is described by an asphere, then the arc is correct at only three points: the vertex, and the top and bottom points. ZEMAX uses the exact surface sag at these three points to fit the arc. See also "Export IGES/STEP Solid" on page 151.

If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail by using the Ray Trace calculation described later in this chapter.

3D Layout

Purpose:

Draws 3D layout plots of the lens system. The algorithm draws a wireframe style representation of the lens.

Settings:

Item	Description
First Surface	The first surface to be drawn.
Last Surface	The last surface to be drawn.
Wavelength	Either any one or all wavelengths may be shown.
Field	Either any one or all field positions may be shown.
Number of Rays	The number of rays specifies the number of rays to be drawn for each selected field and wavelength. The rays will be evenly distributed along the fan of the pupil, or around the perimeter if Ring is the selected Ray Pattern, or unless apodization has been specified. This parameter may be set to zero. It is ignored if the Ray Pattern is set to List.
Ray Pattern	Choose XY Fan, X Fan, Y Fan, Ring, List, or Random to indicate what the pattern of rays to be traced should be. The List option indicates that the rays to be traced are user defined and listed in a file, see the discussion below for information on the ray list format. If List is selected the Number of Rays setting is ignored.

Item	Description
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the range of surfaces drawn to fill the graphic page. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the lens actual size on the printer (not the screen). A factor of 0.5 will plot the lens at half scale.
Hide Lens Faces	If checked, this option will suppress drawing of the lens faces, and only the lens edges will be drawn. This is useful because some complicated systems look cluttered with the faces drawn.
Hide Lens Edges	If checked, this option will suppress drawing of the outer aperture of the lens. This is useful for giving the 3D layout a 2D "cross section" appearance.
Hide X Bars	If checked, this option will suppress drawing of the X portions of the lens faces. This option is useful when "Hide Lens Edges" is checked and "Hide Lens Faces" is not checked.
Rotation About X	The angle in degrees by which the lens appears to be rotated about the X axis.
Rotation About Y	The angle in degrees by which the lens appears to be rotated about the Y axis.
Rotation About Z	The angle in degrees by which the lens appears to be rotated about the Z axis.
Color Rays By	Select "Fields" to use color to distinguish between each field position, or "Waves" to distinguish between each wavelength, or "Config" to distinguish between configurations.
Suppress Frame	Suppresses drawing of the frame on the bottom of the screen, which leaves more room for the layout plot itself. No scale bar, address block, or other data will be displayed.
Delete Vignetted	If checked, rays are not drawn if they will be vignetted by any surface.
Configuration	Select "All" to draw all configurations at once, or select the one configuration to draw, or select "Current" to show the active configuration.
Offset X, Y, Z	The X, Y, and Z direction offset between configurations in lens units. Only has an affect on the drawing if "All" configurations are being drawn.
Fletch Rays	If checked, small arrows are drawn on each ray to indicate the direction of propagation.
Split NSC Rays	If checked, rays from NSC sources will be statistically split at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Scatter NSC Rays	If checked, rays from NSC sources will be statistically scattered at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Square Edges	Draws flat faces and edges if selected, otherwise uses the semi-diameter value to draw lens edges.

Discussion:

Pressing the left, right, up, down, Page Up, or Page Down keys will rotate the displayed image for a different perspective.

For rays from the sequential entry port only: If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail by using the Ray Trace calculation described later in this chapter.

When drawing all configurations, an offset may be added to each configuration in the x, y, and z directions independently. The offsets may all be zero if desired. If the offsets are zero, then all the configurations are superimposed; otherwise, the configurations are all displaced from one another by the specified amount. Note that all offsets are defined from the global coordinate reference surface position. The global coordinate reference surface is defined on the Miscellaneous tab of the System, General dialog box. If all offsets are zero, the multiple configurations are all overlapped at the global coordinate reference surface.

Raylist file format

If List is chosen for the ray pattern, the rays to be traced are defined in a file. The file must be called RAYLIST.TXT and be placed in the main ZEMAX directory. The file format is ASCII, with two distinct methods for defining the rays supported, implicit and explicit. The implicit format file consists of two numbers on each line, one for the px and one for the py normalized pupil coordinates. The specified rays are traced at each defined field and wavelength selected.

Example: Four marginal rays are defined by:

```
0.0 -1.0
0.0 1.0
-1.0 0.0
1.0 0.0
```

The explicit format file consists of the word EXPLICIT followed by the values x, y, z, l, m, n, and wavenumber where x, y, and z are the ray starting coordinates, l, m, and n are the direction cosines, and wavenumber is an integer indicating the wavelength to use. All coordinates are in object space. If the object thickness is infinity, then the spatial coordinates are relative to surface 1. If the object is not at infinity, then the coordinates are relative to surface 0. In both cases the ray itself is in the object space medium, prior to refraction into surface 1. If explicit format is used, then the field and wavelength settings are ignored, and only those rays listed in the file are traced.

Example: Three rays at wavelengths 1, 2, and 3 along the Y axis parallel to the Z axis are defined as follows:

```
EXPLICIT
0.0 -5.0 0.0 0.0 0.0 1.0 1
0.0 +0.0 0.0 0.0 0.0 1.0 2
0.0 +5.0 0.0 0.0 0.0 1.0 3
```

Solid Model

Purpose:

Draws a hidden-line representation of the lens.

Settings:

The options are similar to those available for the 3D layout feature. The "Hide Lens Edges" and "Hide X Bars" checkboxes are not available, and several new controls are added as described below.

Item	Description
DXF File	In this text cell enter the name of the DXF format file to use. This option is only used if the button "Export As DXF File" is subsequently pressed. The file will be stored in the default directory for output.
Export As DXF File	If pressed, this button will cause the generation of a DXF formatted file which will contain the same data that is displayed in the graphic window. The file name is given by the "DXF File" option. The DXF file generated is a 3D faceted model of the lens system suitable for importing into CAD programs that can read DXF files.
Draw Section	Select "Full" to draw each lens element completely. The 3/4, 1/2, and 1/4 options draw just that much of the element, yielding a cut away perspective of the lens interior.
Radial Segments	The number of radial segments used to approximate the lens shapes. Larger numbers require more processing time.
Angular Segments	The number of angular segments used to approximate the lens shapes. Larger numbers require more processing time.

Discussion:

The solid model algorithm describes the lens as a collection of polygon facets. The lines and facets which are hidden from view are removed, which gives the lens a solid appearance. This algorithm is slower than the other layout plots, but produces the best looking results. The number of facets used to display the lens elements can be modified using the radial and angular segment options.

The "Export As DXF File" button will generate a 3D DXF file and store it in the file name provided in the "DXF File" data field. The DXF file will consist of small faceted surfaces in a fully 3D orientation. The facets are used to show the shape of curved lens faces. However, the facets are nearly flat segments which only approximate the surface contour. The corners of each facet always lie exactly on the real optical surface, but arbitrary points within the facet do not follow the contour of the surface. ZEMAX uses the exact surface sag at the corners of each facet to define the facet shape. See also "Export IGES/STEP Solid" on page 151.

Pressing the left, right, up, down, Page Up, or Page Down keys will rotate the displayed image for a different perspective.

If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail, see "Ray Trace" on page 119.

Wireframe

Purpose:

Draws a wireframe representation of the lens.

Settings:

The settings are identical to those available for the Solid Model feature, including support of DXF export.

Discussion:

The wireframe model is identical to the solid model, except hidden lines are not removed. This representation may cause the screen to become cluttered with lines. The "Hide Lens Faces" option can be used to clean up the display. The advantage to this display method is speed; it is faster than the solid model.

If rays miss a surface, then the rays will not be drawn to the surface where the error occurred. If the ray is total internal reflected, then the ray will be drawn up to but not past the surface where the error occurred. Ray failures can be evaluated in detail by using the Ray Trace calculation described later in this chapter.

Shaded Model

Purpose:

Draws a shaded solid model representation of the lens using OpenGL graphics.

Settings:

The options are almost identical to those available for the Solid Model feature, except there are additional controls for setting the lighting level and background color.

ZEMAX Element Drawing

Purpose:

This feature creates a mechanical drawing of surface, singlet, or doublet elements suitable for use in optical shop fabrication.

Settings:

Item	Description
Surface	The first surface of the element to be drawn.
Show As	Select either "Surface", "Singlet" or "Doublet".
Note File Name	The name of the ASCII file which contains the notes to be appended to the notes section of the element drawing. Notes should always start at number 2, since number 1 is reserved for the units specification.
Edit Note File	Clicking on this button will invoke the Windows NOTEPAD.EXE editor, which can then be used to modify the selected note file.
Rad n Tol	The radius (1,2,or 3) tolerance box value.
Pow/Irr n	The power/irregularity (1,2, or 3) tolerance box value.
Clear Ap n	The clear aperture of the lens on surface n. The default value is twice the semi-diameter value.

Item	Description
Thick n Tol	The center thickness tolerance of surface n.
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the element to fill the right half of the element drawing. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the element at actual size on the printer (not the screen). A factor of 0.5 will plot the element at half scale.
Title	This field is for any user defined text. The default is the lens title.
Drawing Name	All of these fields are for user defined text. Any text may be entered. No default value.
Approved	
Revision	
Drawn By	
Project	
Note Font Size	
Reset all but titles	If selected, this button will reset all the default tolerances and apertures for the specified surfaces, but the current text titles will remain as they are.

Discussion:

The element drawing settings may be stored for the specific lens file by pressing the Save button. Unlike most analysis features, the element drawing feature saves all the settings for each surface separately. For example, the notes and tolerances for surface 1 may be saved, and then new notes and tolerances for surface 3 may be entered and then saved. To recall the settings for any specific surface, change the surface number to the desired surface, and then press the Load button. If a match is found with a previously saved surface, the settings for that surface will be displayed. This feature makes it easy to regenerate complex drawings for multiple element systems.

An important feature of the element drawing capability is the ability to load different note files and place them on the drawing. The default note file "DEFAULT.NOT" is a generic set of notes which will rarely be useful as is. However, the user can modify the note files (they are ASCII files which any word processor or text editor can modify) and store them under different names. For example, you may want to have a .NOT file for each type of optic you design, and then load the most appropriate note file when the element drawing is generated.

The note files should always start at note number 2. Note number 1 is reserved by ZEMAX for the line "1) All dimensions in millimeters" or whatever the current lens units are. The line breaks and spacings in the note file will be replicated exactly on the element drawing.

Whenever a new element drawing is generated, or the "Reset" button is pressed, the default settings will be regenerated. The default tolerances are taken from the tolerance data editor. The maximum of the min/max tolerance range is used as the default. For example, if the TTHI thickness tolerance is -.03, +.05, the tolerance generated, the tolerance is set to zero. Note all tolerance fields are text; and may be edited to suit any requirement.

A handy conversion between radius tolerance and the power tolerance in fringes for a Newton's rings type optical test against a test plate is given by

$$\#fringes = \frac{\Delta R \rho^2}{\lambda R^2}$$

where ΔR is the radius error, λ is the test wavelength, p is the radial aperture, and R is the radius of curvature. This formula is an approximation for shallow curvatures. For more information, see Malacara, Optical Shop Testing, J. Wiley & Sons, Inc.

ISO Element Drawing

Purpose:

This feature creates an ISO 10110 type drawing of surface, singlet, or doublet elements suitable for use in optical shop fabrication.

Settings:

Item	Description
Surface	The first surface of the element to be drawn.

Discussion:

The ISO 10110 Element Drawing is an interpretation of the drawing specification "ISO 10110 Optics and Optical Instruments -- Preparation of drawings for optical elements and systems: A User's Guide", by Ronald K. Kimmel and Robert E. Parks, eds., published by the Optical Society of America. For more information see OSA's web site at www.osa.org.

NSC 3D Layout

Purpose:

Draws 3D layout plots of the sources and objects in a single NSC group.

Settings:

Item	Description
Fletch Rays	If checked, small arrows are drawn on each ray to indicate the direction of propagation.
Split Rays	If checked, rays from NSC sources will be statistically split at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Scatter Rays	If checked, rays from NSC sources will be statistically scattered at ray-surface intercepts. Rays entering from the entry port are not affected by this setting.
Suppress Frame	Suppresses drawing of the frame on the bottom of the screen, which leaves more room for the layout plot itself. No scale bar, address block, or other data will be displayed.
Configuration	Select "All" to draw all configurations at once, or select the one configuration to draw, or select "Current" to show the active configuration.
Color Rays By	Select "Sources" to use color to distinguish rays traced from each source, or "Waves" to distinguish between each wavelength, or "Config" to distinguish between configurations.
Scale Factor	If the scale factor is set to zero, then "Fill Frame" will be selected, which will scale the range of surfaces drawn to fill the graphic page. If a numeric value is entered, then the plot will be drawn in "real" scale, times the scale factor. For example, a scale factor of 1.0 will plot the lens actual size on the printer (not the screen). A factor of 0.5 will plot the lens at half scale.
Rotation About X	The angle in degrees by which the lens appears to be rotated about the X axis.
Rotation About Y	The angle in degrees by which the lens appears to be rotated about the Y axis.
Rotation About Z	The angle in degrees by which the lens appears to be rotated about the Z axis.
Offset X, Y, Z	The X, Y, and Z direction offset between configurations in lens units. Only has an affect on the drawing if "All" configurations are being drawn.

Item	Description
Filter	If blank, all rays are drawn. Otherwise, only rays which meet the criteria defined by the filter string will be drawn. See the discussion for information about the filter string syntax.
Ray Database	If "none" is selected, new rays will be traced and displayed. If a ZRD data base file is selected, then rays contained within the database will be displayed. In either case the filter, if any, is applied. Generally, reading rays from a large database is faster than retracing them. The other advantage to using a database of rays is that the rays are always the same, until the database is replaced. ZEMAX cannot tell if the selected database is for rays from the current lens being displayed; so care should be taken when selecting the ZRD file that it corresponds to the lens data being displayed. For more information on the ZRD file, see "Ray database files" on page 267.

Discussion:

The settings in the above table are very much like their counterparts on the 3D Layout feature. However, the feature only draws objects in a single NSC group, and only traces and draws rays from sources defined within the group.

The filter string

It is frequently convenient to draw only rays which have certain properties. For example, when scattering and splitting are turned on, the layout diagram will become very confusing if many rays are traced. The filter string allows definition of a "test" rays must pass before they are drawn. The filter string syntax consists of logical operations between flags that indicate if a ray hit, miss, reflected, or refracted from an object within the NSC group. The supported flags are:

Hn: Ray hit object. To test whether a ray hit an object, the flag is of the form Hn. For example, to test if a ray hit object 5, the flag would be H5.

Mn: Ray missed object. To test whether a ray missed an object, the flag is of the form Mn. For example, to test if a ray hit missed object 15, the flag would be M15.

Rn: Ray reflected from object. The flag R7 would test if the ray reflected at object 7.

Tn: Ray transmitted (refracted) in to or out of object. The flag T4 would test if the ray refracted in or out of object 4.

Sn: Ray scattered from object n.

Dn: Ray diffracted from object n.

Each flag is evaluated for each ray traced, and the flag is assigned a status of TRUE or FALSE. The flags may be used alone, or may be combined using logical operands. Logical operands generally act on two other logical flags (the exception is the NOT operand which acts only on the flag to the right of the operand). The supported logical operands are:

&: Logical AND. Both flags on either side of the & symbol must be TRUE for the AND operation to return TRUE.

|: Logical OR. If either of the flags are TRUE, OR returns TRUE.

^ Exclusive OR (XOR). If either of the flags are TRUE, but not both, XOR returns TRUE.

! Logical NOT. Returns TRUE if the flag to the right was FALSE and vice-a-versa.

The parentheses symbols may also be used to set operator precedence.

Examples

For example, if the only rays to be draw are those which hit object 7, the filter string would be:

H7

If the only rays to be draw are those which hit both objects 7 and 9, the filter string would be

H7 & H9

If the rays to be drawn must have either a) hit object 7 and object 9, but did not reflect off object 6, or b) miss

(H7 & H9 & !R6) | M15

The filter string is checked for basic syntax errors, such as mismatched parentheses, but not all possible syntax errors are checked and reported.

The number of rays which meet the filter string test and are thus drawn may be very small, and perhaps even zero. ZEMAX still traces the number of rays defined for the source (by the # of layout rays parameter), but only that fraction of these rays which pass the filter will actually be drawn.

NSC Shaded Model

Purpose:

Draws a shaded solid model representation of the lens using OpenGL graphics.

Settings:

The options are very similar to those available for the Shaded Model feature, except this feature only draws components and rays from sources in a single NSC group.

One additional feature is the option to color detector objects by either the energy incident on the detector in the last analysis or by just the rays traced in the layout view. The detector shows false color or black and white displays, using either coherent intensity or any of the other options supported by the detector viewer. See "Detector" on page 244 for details.

Fans

Ray Aberration

Purpose:

Shows ray aberrations as a function of pupil coordinate.

Settings:

Item	Description
Plot Scale	Sets the maximum vertical scale for the plots. The maximum scale is in microns for ray fans, waves for OPD plots, or percent for entrance pupil aberration plots. This overrides the automatic selection of scale for the plots. Enter zero for automatic scaling.
Number of Rays	This is the number of rays traced on each side of the origin of the plot.
Wavelength	The wavelength number for which the calculation should be performed.
Field	The field number for which the calculation should be performed.
Tangential	Selects which aberration component to plot for the tangential fan. Since tangential fans are functions of the y pupil coordinate, the default is to plot the y component of the aberration.
Sagittal	Selects which aberration component to plot for the sagittal fan. Since sagittal fans are functions of the x pupil coordinate, the default is to plot the x component of the aberration.
Use Dashes	Selects either color or dashes.
Check Apertures	Specifies whether or not to check if rays pass all surface apertures. If this is selected, rays which do not pass surface apertures will not be drawn.
Vignetted Pupil	If checked, the pupil axis will be scaled to the unvignetted pupil, in which case the data will reflect the vignetting in the system. If unchecked, the pupil axis will be scaled to fit the vignetted pupil.

Discussion:

The tangential fans show either the x or the y component of the transverse ray aberration as a function of the y pupil coordinate of the ray. The default option is to plot the y component of the aberration. However, since transverse ray aberrations are vectors, this is an incomplete description of the aberration. When ZEMAX plots the y component, the plot is labeled EY, when plotting the x component of the aberration, the plot is labeled EX.

The vertical axis scale is given at the bottom of the graph. The data being plotted is the difference between the ray intercept coordinate and the chief ray intercept coordinate. The tangential fan is the plot of the difference between the ray x or y coordinate and the chief ray x or y coordinate at the primary wavelength, as a function of the y pupil coordinate. The sagittal plot is the difference between the ray x or y coordinate and the chief ray x or y coordinate as a function of the x pupil coordinate. The horizontal scale for each graph is the normalized entrance pupil coordinate, either PX or PY.

If "All" wavelengths are shown, then the plot is referenced to the primary wavelength chief ray. If monochromatic, then the chief ray for the selected wavelength is used as a reference. For this reason, the data for non-primary wavelengths will in general change when switching between monochromatic and polychromatic displays.

Because ray aberrations are vectors, with both an x and a y component, the ray aberration fan is an incomplete description of the aberrations, especially when the image plane is rotated or the system is otherwise non-rotationally symmetric. Also, the fans only indicate aberrations along two "slices" through the pupil, rather than over the entire entrance pupil. The primary purpose of the ray fan plot is to determine what aberrations are present in the system; it is not a complete description of the system performance, especially for systems without rotational symmetry.

Optical Path

Purpose:

Shows optical path difference as a function of pupil coordinate.

Settings:

The options are identical to those for ray aberration fans, except the only option for "Tangential Fan" and "Sagittal Fan" is OPD, since OPD is a scalar quantity.

Discussion:

The vertical axis scale is given at the bottom of the graph. The data being plotted is the optical path difference or OPD, which is the difference between the optical path length of the ray and the optical path length of the chief ray. Usually, the calculation is referenced back to the difference between the ray path lengths at the system entrance pupil. The horizontal scale for each graph is the normalized entrance pupil coordinate.

If "All" wavelengths are shown, then the plot is referenced to the primary wavelength based reference sphere and chief ray. If monochromatic, then the reference sphere and chief ray for the selected wavelength is used as a reference. For this reason, the data for non-primary wavelengths will in general change when switching between monochromatic and polychromatic displays.

Pupil Aberration

Purpose:

Shows entrance pupil distortion as a function of pupil coordinate.

Settings:

The options are identical to those for ray aberration fans, except the only option for "Tangential Fan" and "Sagittal Fan" is pupil aberration, since pupil aberration is a scalar quantity.

Discussion:

Entrance pupil aberration is defined as the difference between the real ray intercept on the stop surface and the on axis primary wavelength paraxial ray intercept as a percentage of the paraxial stop radius. If the maximum aberration exceeds a few percent, ray aiming (see the chapter "System Menu") may be required to get the rays in object space to correctly fill the stop surface. If ray aiming is switched on, the entrance pupil aberration will appear to be zero (or a very small residual value), because the distortion is accounted for by the ray tracing algorithm. This is used as a check that ray aiming is working correctly. The definition used here for pupil aberration is not intended to be complete or in agreement with other definitions. The sole purpose of this feature is to provide guidance as to whether or not ray aiming is required.

Spot Diagrams

Standard

Purpose:

Show spot diagrams.

Settings:

Item	Description
Pattern	The pattern may be either hexapolar, square, or dithered. These terms refer to the pattern of rays as they appear in a pupil plane. Defocus the lens significantly to see the pattern, if desired. Dithered spot diagrams are generated by pseudo-random rays which eliminate the symmetrical artifacts in the spot diagram typical of rectangular or hexapolar patterns. The pattern is distorted to give the correct distribution of rays if pupil apodization is specified. There is no "best" pattern to use, each shows a different character of the spot diagram.
Refer To	The spot diagrams by default are referenced to the real chief ray. The RMS and GEO spot sizes listed at the bottom of the diagram (and defined in the discussion section) are calculated assuming the chief ray is the "zero aberration" point. However, this option allows selection of two other reference points: the centroid or the middle. The centroid is defined by the distribution of rays traced. The middle is defined so that the maximum ray errors are equal in the x and y directions.
Show Scale	Scale bar is the default. Selecting "Airy Disk" will draw an elliptical ring around each spot showing the size of the Airy ellipse. The Airy disk radius is 1.22 times the wavelength (primary wavelength is used if polychromatic) times the F/# of the system; which in general depends upon field position and pupil orientation. If the Airy disk is larger than the spot, the Airy disk will set the scale size, otherwise the spot size will set the scale. Selecting "Square" will draw a box, centered on the reference point, whose width is twice the distance from the reference point to the outermost ray. Selecting "Cross" will draw a cross through the reference point location. The "Circle" setting will draw a circle centered on the reference point.
Wavelength	The wavelength number for which the calculation should be performed.
Field	The field number for which the calculation should be performed.
Surface Number	Selects the surface at which the spot diagram is to be evaluated. This is useful for evaluating intermediate images or vignetting.
Plot Scale	Sets the maximum scale size of the display in microns. A setting of zero will generate an appropriate scale.
Delta Focus	The delta focus option is only used if through-focus spot diagrams are selected. This is the Z-axis spacing between spot diagram planes. Five spot diagrams will be shown for each field angle. The defocus will be -2, -1, 0, 1, and 2 times the delta focus provided. The delta focus units are in microns. For some systems the default delta focus may be too small to cause a change in the spot diagram structure.
Ray Density	The ray density specifies the number of hexapolar rings to be traced if a hexapolar or dithered pattern is selected, or the number of rays across the width and height if a rectangular pattern is selected. The more rays traced, the greater the accuracy of the RMS spot size, although the computation time increases. There are 6 rays in the first hexapolar ring, 12 in the second, 18 in the third, and so on.
Use Symbols	If checked, this option will draw different symbols rather than dots for each wavelength. This helps distinguish the various wavelengths.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. Only ZEMAX-EE supports this capability.

Item	Description
Direction Cosines	If checked, the data presented will be the direction cosines of the rays rather than the spatial coordinates of the rays. The x direction data will be the x direction cosine of the ray, the y direction data will be the y direction cosine. The image coordinates will also be given as the reference point direction cosines. Direction cosines are dimensionless.
Configuration	Select "All" to draw all configurations at once, or select the one configuration to draw, or select "Current" to show the active configuration.
Color Rays By	Select "Fields" to use color to distinguish between each field position, or "Waves" to distinguish between each wavelength, or "Config" to distinguish between configurations.

Discussion:

The ray density has a maximum value based upon the number of fields displayed, the number of wavelengths defined, and available memory. Through-focus spot diagrams will trace half of the maximum number of rays possible on standard spot diagrams.

The GEO spot size listed on the plot for each field point is the distance from the reference point (which is either the chief ray at the primary wavelength, the centroid of all the rays traced, or the middle of the spot cluster) to the ray which is farthest away from the reference point. In other words, the GEO spot size is the radius of the circle centered at the reference point which encloses all the rays.

The RMS spot size is the root-mean-square radial size. The distance between each ray and the reference point is squared, and averaged over all the rays, and then the square root is taken. The RMS spot size gives a rough idea of the spread of the rays, since it depends upon every ray. The GEO spot size only gives information about the one ray which is farthest from the reference point.

For information on the X and Y RMS spot sizes; see the "text" listing for the spot diagram.

The Airy disk radius is given by 1.22 times the wavelength (primary wavelength is used if polychromatic) times the F/# of the beam, which in general depends upon field position and pupil orientation. This is the radius to the first dark ring of the Airy disk for a circular, uniformly illuminated entrance pupil. The Airy disk may be optionally drawn to give an idea of the scale of the plot. For example, if all the rays are well within the Airy disk, then the system is often said to be "diffraction limited". If the RMS spot size is significantly larger than the Airy disk radius, then the system is not diffraction limited. The threshold for diffraction limited performance depends upon which criterion is used. There is no absolute boundary at which the system becomes diffraction limited. The Airy disk shown is not an accurate representation of the diffraction dark ring shape or size if the system does not have uniform illumination or if vignetting is used to eliminate some of the rays. ZEMAX does not plot vignettted rays on spot diagrams, nor are they used in computing the RMS or GEO spot sizes.

ZEMAX generates grids of rays based upon the wavelength weighting factors and the pupil apodization, if any. The wavelength with the largest weight uses the maximum grid size set by the "Ray Density" option. Wavelengths with lower weights use grids with fewer rays to maintain the correct representation in the diagram. Ray grids are also distorted to maintain the correct ray distribution, if apodization is specified. The RMS spot size stated on the spot diagram considers the wavelength weighting and apodization factors. However, it is only an estimate of the RMS spot size based on the rays actually traced. It is not a very accurate estimate for some systems.

The image plane intercept coordinates of the reference point are shown underneath each spot diagram. If a surface other than the image plane is specified, then the coordinates are the intercept coordinates of the reference point on that surface. Since the reference point may be selected to be the centroid, this provides a convenient way of determining the centroid coordinates.

Through Focus

Purpose:

Show spot diagrams as they change through focal plane shifts.

Settings:

The options are identical to the standard spot diagram.

Discussion:

The through focus spot diagrams are useful for estimating astigmatism, or for analyzing best focus or depth of focus.

Full Field

Purpose:

Shows spot diagram with all field points on a common scale.

Settings:

The options are identical to the standard spot diagram.

Discussion:

The "Full Field" spot diagram type is similar to the "Standard" type, except all of the spots are plotted with respect to the same reference point, as opposed to a separate reference point for each field position. This provides some idea of how the spot would look relative to the other field points. For example, this can be used to determine if two closely spaced image points can be resolved. The "Full Field" spot diagram type is useless if the spot size is small compared to the total field size, because in this case the spots for each field will appear as simple "dots". If "chief ray" is selected as the reference point, then the chief ray for field position 1 will be used.

Matrix

Purpose:

Show spot diagram as a matrix of individual diagrams, with each field along a row and each wavelength down a column.

Settings:

The options for the matrix spot diagram are similar to those for the standard spot diagram, except for the addition of the following option.

Item	Description
Ignore Lateral Color	If checked, this option will reference each spot diagram to the reference point for each field and wavelength independently. This in effect ignores the effects of lateral color which can displace the reference points for each wavelength.

Discussion:

The matrix representation is a convenient way of distinguishing the components of wavelength dependent aberrations.

Configuration Matrix

Purpose:

Show spot diagram as a matrix of individual diagrams, with each field along a row and each configuration down a column.

Settings:

The options for the configuration matrix spot diagram are similar to those for the standard spot diagram.

Discussion:

The configuration matrix representation is a convenient way of distinguishing the components of configuration dependent aberrations. Along the left side of the plot the field positions are listed; only the field positions for the last configuration are listed if more than one configuration is shown and the field definitions are part of the multi-configuration data which changes.

MTF

FFT MTF


Purpose:

Computes the diffraction modulation transfer function (MTF) data for all field positions using an FFT algorithm.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Show Diffraction Limit	Select whether or not the diffraction limited data should be displayed.
Max Frequency	Select the maximum spatial frequency (in cycles per millimeter) plotted.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number to be used in the calculation.
Type	Select either modulation, real, imaginary, phase, or square wave response.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:

 See the discussion sections of the FFT and Huygens Point Spread Functions. Those comments also apply to this feature.

The diffraction MTF computation is based upon an FFT of the pupil data. The resulting MTF is the modulation as a function of spatial frequency for a sine wave object, although optionally the real, imaginary, phase, or square wave response is available. The Square wave MTF is the modulation response for a square wave target of the specific frequency, as opposed to the response to a sine wave target for the other plots. The square wave response is computed from the MTF data using the following formula:

$$S(v) = \frac{4}{\pi} \left[\frac{M(v)}{1} - \frac{M(3v)}{3} + \frac{M(5v)}{5} - \frac{M(7v)}{7} + \dots \right]$$

where $S(v)$ is the square wave response, $M(v)$ is the sinusoidal modulation response, and v is the spatial frequency.

The cutoff frequency at any one wavelength is given by one over the wavelength times the working F/#. ZEMAX computes the working F/# at each wavelength for each field for the sagittal and tangential response separately. This yields accurate MTF data even for systems with anamorphic and chromatic distortion, such as those incorporating cylinders or gratings.

The diffraction calculations are more accurate as the sampling increases, the peak-to-valley and maximum slope of the OPD decrease, and the transverse ray aberrations decrease. If the peak-to-valley OPD in the pupil is too great, then the wavefront sampling is too coarse and aliasing occurs. Aliasing will result in inaccurate data. ZEMAX will attempt to detect when aliasing occurs, and issue an appropriate error message. However, ZEMAX cannot automatically detect when the sampling is too low in all cases, especially in the presence of very steep slopes on the wavefront phase.

The FFT based MTF assumes a (reasonably) uniform distribution of rays on the exit pupil in cosine space to be accurate. Some systems, such as extremely fast off-axis reflectors, have dramatic stretching of the exit pupil and the FFT based MTF will thus be inaccurate. For these systems, the Huygens MTF should be used instead. For more information, see "Huygens MTF" on page 83.

When the OPD in waves is large, such as more than 10 waves, it is probably a good idea to switch to geometric MTF instead of diffraction MTF. For these highly aberrated systems, the geometric MTF is very accurate, especially at low spatial frequencies (the higher frequency MTF falls off rapidly when aberrations are large).

Because ZEMAX does not account for vector diffraction, the MTF data may not be accurate for systems faster than about F/1.5 (the deterioration in accuracy is gradual). For these systems, the OPD fan data are more fundamental and therefore more reliable indicators of performance. If the system is not too close to the diffraction limit, the geometric MTF may prove useful.

If shown, the diffraction limit curve is for the aberration free response at the reference field position (see "Diffraction Limited" in the chapter "Conventions and Definitions").

The spatial frequency scale of the MTF plot is always in cycles per mm in image space, which is the correct term for sinusoidal MTF response. The term line pairs per mm is often used, but strictly speaking line pairs per mm only applies to bar, as opposed to sinusoidal targets. ZEMAX uses these terms interchangeably, as is common in the industry. MTF is always measured in image space, so any magnification of the system needs to be considered when determining spatial frequency response for object space.

FFT Through Focus MTF

Purpose:

Computes the FFT modulation transfer function data as a function of focus shift at a specific spatial frequency.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Delta Focus	The range of defocus used.
Frequency	The spatial frequency (in cycles per millimeter) for which data is plotted.
# Steps	The number of focal planes at which the data is computed. A smooth curve is drawn through the computed points. More steps yield higher accuracy and longer computation times.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select either modulation, real, imaginary, phase, or square wave response.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:

See "FFT MTF" on page 79 for details.

FFT Surface MTF

Purpose:

Displays the FFT computed MTF data as a 3D surface, contour, grey scale or false color map. This plot is useful for visualizing the MTF response for object orientations other than purely sagittal or tangential.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Rotation	Rotation specifies how the surface plots are rotated for viewing; either 0, 90, 180, or 270 degrees.

Item	Description
Scale	The scale factor overrides the automatic vertical scaling set by the program on the surface plots. Generally, this value should be set to unity. The scale factor can be greater than unity to vertically stretch the plot, or less than unity to compress it.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number determines for which defined field position the calculation should be performed.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Discussion:

The regular MTF plot is just two orthogonal cross sections through the surface MTF plot. This plot is primarily qualitative. See "FFT MTF" on page 79 for details.

FFT MTF vs. Field

Purpose:

Computes the FFT MTF data as a function of field position, and displays the data in a graph.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64 etc. Although higher sampling yields more accurate data, calculation times increase.
Frequency 1, 2, 3	The spatial frequencies (in cycles per millimeter) for which data is plotted.
Wavelength	The wavelength number to be used in the calculation.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various spatial frequencies.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See the comments about vignetting factors in the discussion section below.
Field Density	The field density is the number of points between zero degrees and the maximum field at which the MTF is calculated, intermediate values are interpolated. A maximum of 100 field points is allowed.

Discussion:

See "FFT MTF" on page 79 for details. This feature plots MTF vs. Y field height up to the maximum defined radial field coordinate.

Comment about vignetting factors

Vignetting factors determine the size and shape of the pupil as seen from different field points (see "Vignetting factors" on page 65 for a full discussion). Because this analysis feature needs to trace rays at arbitrary intermediate field points where no specific vignetting factors are defined, the use of vignetting factors is not recommended. If "Remove Vignetting Factors" is checked on (the default), any defined vignetting factors will automatically be replaced with surface apertures for this computation. The surface aperture method is generally more accurate than the vignetting factor method when the pupil is overfilled with light. The resulting data may be

different between the two methods. In some cases, particularly where the vignetting factors are being used to define the shape of the source beam rather than the apertures of the optics, it may be required to use the defined vignetting factors. In this case, check the "Remove Vignetting Factors" box off. ZEMAX will then use the closest defined field to determine the vignetting factors to use for an arbitrary field point.

FFT MTF Map

Purpose:

Computes the FFT MTF as a function of field position, and displays the data over a rectangular region of field.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
X or Y Field Width	The X or Y field width in field units. This is the total width or height, not the half width or height. Field units are degrees in object space if field angle is used, otherwise field units are the same as lens units.
Frequency	The spatial frequency at which to compute the MTF.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Wavelength	The wavelength number to be used in the calculation, or All for polychromatic MTF.
X or Y Pixels	The number of pixels at which to compute the MTF in each respective direction. Note the size of the pixels is determined by both the number of pixels and the width of the field; the pixels are not required to be square. The MTF is computed at the center of the pixel and the MTF is assumed to have that value over the entire region of the pixel for display purposes.
MTF Data	Choose Tangential, Sagittal, or average MTF to be displayed.
Reference Field	This control selects the field number that corresponds to the center of the map. If zero is selected the (0, 0) field coordinate is used as the center of the map.
Show As	Choose grey scale or false color map as the display option.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

See "FFT MTF" on page 79 for details.

This feature computes the MTF at each field point on a 2D grid. If the total number of points is large, the computation time may become quite large. See also the Geometric MTF Map feature on page 86.

Huygens MTF


Purpose:

Computes the diffraction modulation transfer function (MTF) data using a Huygens direct integration algorithm.

Settings:

Item	Description
Pupil Sampling	Selects the size of the grid of rays to trace to perform the computation. Higher sampling densities yield more accurate results at the expense of longer computation times.
Image Sampling	The size of the grid of points on which to compute the diffraction image intensity. This number, combined with the image delta, determine the size of the area displayed.
Image Delta	The distance in microns between points in the image grid.
Zero Padding	Adds a zero-value guard band around the computed PSF before performing the Fourier transform.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select the data to display, currently modulation is the only option.
Max Frequency	The maximum spatial frequency in cycles per millimeter to display.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:

 See the discussion sections of the Huygens Point Spread Functions. Those comments also apply to this feature.

The Huygens MTF computes an FFT of the Huygens Point Spread Function. The initial PSF sampling settings for Image Sampling and Image Delta settings are the same as for the Huygens PSF, therefore it is instructive to do a Huygens PSF first (see "Huygens Point Spread Function" on page 90). The zero padding setting adds zero intensity values around the PSF and has the visual effect of increasing the pixel density in the transform. Since the transform is done on the PSF in image space coordinates, the tangential response corresponds to spatial frequencies in the y direction in local image surface coordinates, and the sagittal response corresponds to spatial frequencies in the x direction. The Huygens MTF also has no dependence on the location of rays in the paraxial pupils. The MTF can therefore be computed for any system that the Huygens PSF can be computed for including many non-sequential systems using ports where reference rays required by other diffraction algorithms would not make it through, or for systems where pupils or images formed by multiple non-sequential sub-apertures are overlapped. Systems with extreme exit pupil distortion, such as very fast off-axis reflectors, are also handled correctly with the Huygens technique.

The spatial frequency scale of the MTF is always in cycles per mm in image space.

Huygens Surface MTF

Purpose:

Computes the diffraction modulation transfer function (MTF) data using a Huygens direct integration algorithm and displays the data as a surface, grey scale, false color, or contour plot.

Discussion:

This feature is very similar to the Huygens MTF feature, as described on page 83.

Geometric MTF

Purpose:

Computes the geometric MTF, which is an approximation to the diffraction MTF based upon ray aberration data.
Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Max Frequency	The maximum spatial frequency (in cycles per millimeter) for which data is plotted.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Multiply by Diffraction Limit	When checked, will scale the geometric MTF by the diffraction limited MTF to yield a more realistic result for systems with small aberrations. Should always be used.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. Only ZEMAX-EE supports this capability.

Discussion:

The geometric MTF is a useful approximation to the diffraction MTF if the system is not close to the diffraction limit. The primary advantage to using the geometric MTF is for systems which have too many waves of aberration to permit accurate calculation of the diffraction MTF. The geometric MTF is also very accurate at low spatial frequencies for systems with large aberrations.

Geometric Through Focus MTF

Purpose:

Computes the geometric MTF data through focus at a specific spatial frequency.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Delta Focus	The range of defocus used.
Frequency	The spatial frequency (in cycles per millimeter) for which data is plotted.
# Steps	The number of focal planes at which the data is computed. A smooth curve is drawn through the computed points. More steps yield higher accuracy and longer computation times.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Multiply by Diffraction Limit	When checked, will scale the geometric MTF by the diffraction limited MTF to yield a more realistic result for systems with small aberrations. Should always be used.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Item	Description
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. Only ZEMAX-EE supports this capability.
Use Dashes	Selects either solid lines or dashed lines to differentiate the various curves.

Discussion:

See the Geometric Transfer Function section on page 84 for details.

Geometric MTF vs. Field

Purpose:

Computes the geometric modulation transfer function data as a function of field position.

Settings:

The settings are identical to those for the (diffraction) MTF vs. Field feature, with the added ability to scatter rays.

Discussion:

This feature is nearly identical to the (diffraction) MTF vs. Field feature, except the geometric MTF is used rather than the diffraction based MTF. The GMTF is always scaled by the diffraction limit for this analysis.

Geometric MTF Map

Purpose:

Computes the geometric modulation transfer function data as a function of field position, and displays the data over a rectangular region of field.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
X or Y Field Width	The X or Y field width in field units. This is the total width or height, not the half width or height. Field units are degrees in object space if field angle is used, otherwise field units are the same as lens units.
Frequency	The spatial frequency at which to compute the GMTF.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Wavelength	The wavelength number to be used in the calculation, or All for polychromatic GMTF.
X or Y Pixels	The number of pixels at which to compute the GMTF in each respective direction. Note the size of the pixels is determined by both the number of pixels and the width of the field; the pixels are not required to be square. The GMTF is computed at the center of the pixel and the GMTF is assumed to have that value over the entire region of the pixel for display purposes.
MTF Data	Choose Tangential, Sagittal, or average GMTF to be displayed.
Reference Field	This control selects the field number that corresponds to the center of the map. The value zero refers to the center of the object ($h_x = 0$, $h_y = 0$), even if no field point is defined there.
Show As	Choose grey scale or false color map as the display option.

Item	Description
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. Only ZEMAX-EE supports this capability.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

See the Geometric Transfer Function on page 84 for details. This feature computes the GMTF at each field point on a 2D grid. If the total number of points is large, the computation time may become quite large. See also the (Diffraction) MTF Map feature on page 83.

PSF

FFT Point Spread Function

Purpose:

Computes the diffraction point spread function using the Fast Fourier Transform (FFT) method.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Display	The display size indicates what portion of the computed data will be drawn. The display grid can be any size from 32 x 32 up to twice the sampling grid size. Smaller display sizes will show less data, but at higher magnification for better visibility.
Rotation	Rotation specifies how the surface plots are rotated for viewing; either 0, 90, 180, or 270 degrees.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select linear (intensity), logarithmic (intensity), or phase.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Image Delta	The delta distance between points in image space, measured in microns. If zero, a default spacing is used. See the discussion for details.
Normalize	If checked, the peak intensity will be normalized to unity. Otherwise, the peak intensity is normalized to the peak of the unaberrated point spread function (the Strehl ratio).

Discussion:

The FFT method of computing the PSF is very fast, however, a few assumptions are made which are not always valid. The slower, but more general Huygens method makes no such assumptions, and is described in the next section.

Assumptions used in the FFT PSF calculation

The FFT PSF computes the intensity of the diffraction image formed by the optical system for a single point source in the field. The intensity is computed on an imaginary plane which lies perpendicular to the incident chief ray at the reference wavelength. The reference wavelength is the primary wavelength for polychromatic computations, or the wavelength being used for monochromatic calculations. Because the imaginary plane lies normal to the

chief ray, and not the image surface, the FFT PSF computes overly optimistic (a smaller PSF) results when the chief ray angle of incidence is not zero. This is often the case for systems with tilted image planes, wide angle systems, systems with aberrated exit pupils, or systems far from the telecentric condition.

The other main assumption the FFT method makes is that the image surface lies in the far field of the optical beam. This means the computed PSF is only accurate if the image surface is fairly close to the geometric focus for all rays; or put another way, that the transverse ray aberrations are not too large. There is no hard and fast limit, however if the transverse aberrations exceed a few hundred wavelengths, the computation is likely not accurate. Note that even systems with very little wavefront aberration can have large transverse aberrations; for example, a cylinder lens which only focuses rays along one direction. In this case, the transverse aberrations along the unfocused direction will be on the order of the beam diameter. The Huygens PSF method may provide more accurate results in these cases as well.

For most lenses, a less important assumption is that scalar diffraction theory applies. The vectorial nature of light is not accounted for. This is significant in systems that are very fast, around F/1.5 (in air) or faster. The scalar theory predicts overly optimistic (a smaller PSF) results when the F/# is very fast.

For systems where the chief ray is nearly normal (less than perhaps 20 degrees), the exit pupil aberrations are negligible, and the transverse ray aberrations are reasonable, then the FFT PSF is accurate and generally much faster than the Huygens PSF method.

When in doubt, both PSF methods should be employed for comparison. A solid understanding on the part of the user of these assumptions and the method of computation is essential to recognize cases where the accuracy may be compromised.

Discussion of the FFT method and sampling issues

The FFT PSF algorithm exploits the fact that the diffraction PSF is related to the Fourier transform of the complex amplitude of the wavefront in the exit pupil of the optical system. The amplitude and phase in the exit pupil are computed for a grid of rays, an FFT is performed, and the diffraction image intensity is computed.

There is a tradeoff between the sampling grid size in the pupil, and the sampling period in the diffraction image. For example, to decrease the sampling period in the diffraction image, the sampling period in the pupil must increase. This is done by "stretching" the pupil sampling grid so that it overfills the pupil. This process means fewer points actually lie within the pupil.

As the sampling grid size is increased, ZEMAX scales the grid on the pupil to yield an increase in the number of points that lie on the pupil, while simultaneously yielding closer sampling in the diffraction image. Each time the grid size is doubled, the pupil sampling period (the distance between points in the pupil) decreases by the square root of 2 in each dimension, the image plane sampling period also decreases by the square root of 2 in each dimension, and the width of the diffraction image grid increases by a factor of the square root of 2 (since there are twice as many points in each dimension). All ratios are approximate, and asymptotically correct for large grids.

The stretching is referenced to a grid size of 16 x 16. The 16 x 16 grid of points is placed over the pupil, and the points that lie within the pupil are actually traced. For this grid size, the default distance between points in the diffraction image plane is given by

$$\Delta X = \lambda F \frac{n-2}{2n}$$

where F is the working F/# (not the same as the image space F/#), λ is the shortest defined wavelength, and n is the number of points across the grid, in this case 16 (the sampling grid is of size n x n). The -2 factor is due to the fact the pupil is not centered on the grid (since n is even), but is offset at n/2 + 1. The 2n in the denominator is due to the zero-padding described later.

For grids larger than 16 x 16 (and ZEMAX starts at 32 x 32), the grid is by default stretched in pupil space by a factor of $\sqrt{2}$ each time the sampling density doubles. The general formula for the sampling in image space is the

$$\Delta X = \lambda F \frac{n-2}{2n} \left[\frac{16}{n} \right]^{1/2}$$

and the total width of the image data grid is

$$W = \Delta x(2n - 1).$$

Since the stretching of the pupil grid decreases the number of sample points in the pupil, the effective grid size (the size of the grid that actually represents traced rays) is smaller than the sampling grid. The effective grid size increases as the sampling increases, but not as quickly. The following table summarizes the approximate effective grid size for various sampling density values.

DEFAULT EFFECTIVE GRID SIZES FOR PSF CALCULATIONS

Sampling Grid Size	Approximate Effective Pupil Sampling
32 x 32	23 x 23
64 x 64	32 x 32
128 x 128	45 x 45
256 x 256	64 x 64
512 x 512	90 x 90
1024 x 1024	128 x 128
2048 x 2048	181 x 181

The sampling is also a function of wavelength. The discussion above is only valid for the shortest wavelength used in the calculation. If the computation is polychromatic, then the longer wavelengths will be scaled to have smaller effective grids. The scale factor used is the ratio of the wavelengths. This should be considered when selecting sampling grids for systems with broad wavelength bands. For polychromatic computations, the data for shorter wavelengths is more accurate than for longer wavelengths.

The default image delta, Δx , can be selected manually if a different sampling distance is required. If the image delta is zero, ZEMAX uses the default spacing and sampling grids described above. If the image delta is greater than zero, then ZEMAX scales the pupil sampling to yield the desired image delta size. The actual amount of stretching depends upon the grid size, the image delta, the defined wavelengths, the F/#s at each field and wavelength, and the aspect ratio of the exit pupil. If the image delta is set too small, then the not enough points will be left to sample the pupil; if the image delta is too big, then the pupil grid will not extend over the full width of the exit pupil. Both of these cases are trapped by ZEMAX and an error message will be issued if they occur.

Once the sampling is specified, ZEMAX doubles the array size in a process called "zero padding". This means for a 32 x 32 sampling, ZEMAX uses the center portion of a 64 x 64 grid. Therefore, the diffraction point spread function will be distributed over a 64 x 64 size grid. The sampling in the image space is always twice the pupil sampling. Zero padding is performed to reduce aliasing.

FFT PSF Cross Section

Purpose:

This feature plots cross sections through the diffraction Point Spread Function.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Row/Col	The row or column to display. For a sampling of 32 x 32, there are 64 rows and 64 columns (see the discussion section of the Point Spread Function feature). Whether a row or a column is used depends upon the "Type" setting.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select X or Y cross sections, either linear, logarithmic, or phase. X cross sections are called rows, and Y cross sections are called columns, but this is arbitrary

Item	Description
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Normalize	If checked, the peak intensity will be normalized to unity. Otherwise, the peak intensity is normalized to the peak of the unaberrated point spread function (the Strehl ratio).

Discussion:

☞ See the discussion section of the FFT Point Spread Function. Those comments also apply to this feature.

The cross sections are taken directly from the PSF data. Because the PSF is computed directly from the phase in the exit pupil, the orientation of the coordinate system may not be correct in all cases. What constitutes positive x or y may not agree with data presented in image space coordinates such as the spot diagram.

Huygens Point Spread Function

Purpose:

Computes the diffraction point spread function using direct integration of Huygens wavelets method.

Settings:

Item	Description
Pupil Sampling	Selects the size of the grid of rays to trace to perform the computation. Higher sampling densities yield more accurate results at the expense of longer computation times.
Image Sampling	The size of the grid of points on which to compute the diffraction image intensity. This number, combined with the image delta, determine the size of the area displayed.
Image Delta	The distance in microns between points in the image grid.
Rotation	Rotation specifies how the surface plots are rotated; either 0, 90, 180, or 270 degrees.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Type	Select linear (intensity), or logarithmic (intensity).
Normalize	If checked, the peak intensity will be normalized to unity. Otherwise, the peak intensity is normalized to the peak of the unaberrated point spread function (the Strehl ratio).
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Use Centroid	If checked, the plot will be centered on the geometric image centroid. If unchecked, the plot will be centered on the chief ray.

Discussion:

One way of considering the effects of diffraction is to imagine each point on a wavefront as a perfect point source with an amplitude and phase. Each of these point sources radiates a spherical "wavelet", sometimes called a "Huygens wavelet" after Huygens, who first proposed the model. The diffraction of the wavefront as it propagates through space is given by the interference, or complex sum, of all the spherical wavelets radiated.

To compute the Huygens PSF, a grid of rays is launched through the optical system, and each ray represents a particular amplitude and phase wavelet. The diffraction intensity at any point on the image surface is the complex sum of all these wavelets, squared. Unlike the FFT PSF, ZEMAX computes the Huygens PSF on an imaginary plane tangent to the image surface at the chief ray intercept. Note the imaginary plane is normal to the normal of the surface, not the chief ray. Therefore, the Huygens PSF accounts for any local tilt in the image surface caused by either the image plane slope, the chief ray incidence angle, or both.

The Huygens method accounts for the evolving shape of the diffraction image as the beam propagates along the image surface. This is an important effect if the image surface is tilted with respect to the incoming beam. Another advantage to the Huygens PSF method is that any grid size and spacing, may be selected by the user. This allows direct comparison between PSF's from two different lenses, even if the F/#'s or wavelengths are different.

The only disadvantage of the Huygens PSF is speed. Direct integration is slow when compared to the FFT method (see the previous section for details). The computation time depends upon the pupil grid size squared times the image grid size squared, times the number of wavelengths. ZEMAX accounts for any symmetry the system has.

Wavefront

Wavefront Map

Purpose:

Displays the wavefront aberration.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Rotation	Rotation specifies how the surface plots are rotated for viewing; either 0, 90, 180, or 270 degrees.
Scale	The scale factor is used to override the automatic vertical scaling set by the program on the surface plots. The scale factor can be greater than unity to vertically stretch the plot, or less than unity to compress it.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Reference To Primary	By default, the wavefront aberration is referenced to the reference sphere for the wavelength being used. If this box is checked, then the primary wavelength reference sphere will be used instead. In other words, checking this box will cause the data to exhibit the effects of lateral color.
Use Exit Pupil Shape	By default, the shape of the pupil is distorted to show the approximate shape of the exit pupil as seen from the on axis chief ray image point. If this box is unchecked, then instead the plot will be scaled to circular entrance pupil coordinates, no matter how distorted the exit pupil may actually be.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.

Discussion:

See also the Interferogram feature described below.

Interferogram

Purpose:

Generates and displays interferograms.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64 etc. Although higher sampling yields more accurate data, calculation times increase.
Scale Factor	Determines the number of fringes per wave of OPD. Useful for modeling double pass interferometers (i.e. use a scale factor of two).
Show As	Choose contour map, grey scale, or false color map as the display option.
Wavelength	The wavelength is the wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
X-Tilt	The number of waves of tilt to add in the x-direction after applying the scale factor.
Y-Tilt	The number of waves of tilt to add in the y-direction after applying the scale factor.
Beam 1	Selects the first beam for the interferogram. See the discussion.
Beam 2	Selects the second beam for the interferogram. See the discussion.

Discussion:

This feature works by computing two pupil maps, one each from beams 1 and 2. The phase (or OPD) of these two pupil maps is subtracted, and then optionally some linear phase is added as a function of the x and y pupil coordinate to simulate tilt fringes. The individual beams may be OPD as computed for any one configuration, or a "reference" beam which has identically zero OPD may be selected.

Interferometers may be simulated by modeling the two paths through the system using two configurations, and then computing the interferogram of the two resulting beams. The accuracy of this approach is limited by some simplifying assumptions:

- Any lateral shift or magnification difference between the two beams is ignored; it is assumed that the pupils perfectly overlap at the exit pupil.
- Any differences in transmission are ignored; so the two OPD values at any one point in the pupil are assumed to be of equal value in intensity and the phase can be subtracted to yield the net phase difference.

Foucault Analysis

Purpose:

Generates and displays Foucault knife-edge shadowgrams.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64 etc. Although higher sampling yields more accurate data, calculation times increase.
Type	Selects either linear or logarithmic display of the data.
Show As	Choose surface, contour map, grey scale, false color map, or cross section as the display option.
Wavelength	The wavelength is the wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Row/Col	When "Show As" is selected to be a cross section, this control defines the row or column number to display.

Item	Description
Knife	Choose Horizontal Above, Horizontal Below, Vertical Left, or Vertical Right. The Vertical Left knife blocks out all the light near focus from the knife position coordinate left; that is, towards negative x coordinates. The Vertical Right blocks all light from the knife position rightward. The Horizontal Above blocks all light from the knife position up, and Horizontal Below blocks all light from the knife position down. The terms left, right, up and down refer to the -x, +x, +y, and -y directions in the local coordinate system of the image surface.
Position	The position in microns relative to the chief ray of the knife. The coordinate is assumed to be in X or Y depending upon whether an X or Y knife is selected.
Data	The computed shadowgram, the reference shadowgram, or the difference between the two may be selected. See the discussion for details.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Reference	The name of the bitmap reference image file.
Decenter X/Y	The decenter in X or Y of the reference shadowgram image relative to the computed shadowgram image. The units are relative to the full width or height of the reference shadowgram image. For example, an X Decenter of 0.25 will shift the reference image relative to the computed image by 25% of the full width of the reference image.
Scale X/Y	The scale factor in X or Y of the reference shadowgram image pixels relative to the computed shadowgram image pixels.

Discussion:

This feature simulates the placement of either an X- or Y- oriented knife edge at any position near focus; then computes the resulting shadowgram after propagating the vignetted beam back to the near field. The method of calculation involves computing the diffraction based complex amplitude point spread function at focus via the FFT method; then a portion of the complex amplitude is vignetted by the simulated knife edge, and the remaining complex amplitude is propagated back to the near field. The shadowgram calculated this way is called the "computed" shadowgram for this feature.

This feature also allows the import of either a BMP or JPG bitmap file of a reference or measured shadowgram. The reference shadowgram may be displayed for convenient check of orientation.

The difference between the computed and reference shadowgram may be displayed. ZEMAX computes the RMS difference between the computed and reference shadowgrams, and this RMS difference may be optimized using the FOUC operand described in the chapter "Optimization". Optimizing the RMS difference permits quantitative determination of the aberrations present in the beam that created the measured shadowgram. When calculating the difference between the computed and reference shadowgram, the two images must be registered together to overlap correctly. The decenter x/y and scale x/y controls are used for registering the two images.

Surface

Surface Sag

Purpose:

Displays the sag of a surface as a 2D color or contour map, or as a 3D surface plot.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64 etc. Although higher sampling yields more accurate data, calculation times increase. ZEMAX actually increases the sampling by 1, so that an odd sampling number is used. This allows for a more symmetric display of the data.
Contours	The contour increment to use if the data is to be displayed as a contour map. This setting has no affect on the other types of displays. The contours are defined in lens units.
Surface	The surface number to compute the sag display for.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.

Discussion:

This feature accounts for the size and shape of any aperture present on the surface; even if the aperture is decentered. The sag is computed on a uniform grid of points in XY plane, and the Z value of the sag is the displayed data.

See also the Surface Phase feature described below.

Surface Phase

Purpose:

Displays the phase of a surface as a 2D color or contour map, or as a 3D surface plot.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase. ZEMAX actually increases the sampling by 1, so that an odd sampling number is used. This allows for a more symmetric display of the data.
Contours	The contour increment to use if the data is to be displayed as a contour map. This setting has no affect on the other types of displays. The contours are defined in periods; each period is a phase change of 2π .
Surface	The surface number to compute the phase display for.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.

Discussion:

This feature accounts for the size and shape of any aperture present on the surface; even if the aperture is decentered. The phase is computed on a uniform grid of points in XY plane, and the phase value is the displayed data. This feature defines phase in units of periods; so one period represents a phase change of 2π . Surfaces which do not impart a phase change to the ray, such as the Standard surface, will display a phase of zero everywhere on the surface phase display. See also the Surface Sag feature described above.

RMS

RMS vs. Field

Purpose:

Plots RMS radial, x, and y spot size, RMS wavefront error, or Strehl ratio as a function of field angle; either monochromatic or polychromatic.

Settings:

Item	Description
Ray Density	If the method is Gaussian quadrature, then the ray density specifies the number of radial rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. The maximum number is 20 which is sufficient for pupil aberrations up to order 40. If the method is rectangular array, then the ray density indicates the grid size. Rays outside the circular entrance pupil are ignored. See the "Discussion" section for details.
Field Density	The field density is the number of points between zero degrees and the maximum field angle specified at which the RMS/Strehl ratio is calculated, intermediate values are interpolated. A maximum of 100 field points is allowed.
Plot Scale	Sets the maximum vertical scale for the plot. Zero results in automatic scaling.
Method	Selects either Gaussian quadrature or rectangular array. The Gaussian quadrature method is very fast and accurate, but only works if there is no vignetting. If any rays are to be vignetted, then rectangular array is more accurate.
Data	Selects either wavefront error, spot radius, spot x-direction, spot y-direction, or Strehl ratio.
Refer To	Select either chief ray or centroid. For monochromatic calculations, the specified wavelength is used for reference. For polychromatic calculations, the primary wavelength is used for reference. Both reference points subtract out wavefront piston. The centroid reference mode also subtracts out the tilt of the wavefront, which yields smaller RMS values.
Orientation	Select +y, -y, +x, or -x field direction. Note the data will only be computed to the limits of the defined fields in the selected direction.
Use Dashes	Selects either color or dashes.
Wavelength	Select "All" to display data for each wavelength and a polychromatic computation, select any one wavelength to plot monochromatic data, or select "Poly Only" which only plots the polychromatic data.
Show Diffraction Limit	If checked, then a horizontal line indicating the diffraction limited response will be drawn on the plot. For RMS radius, x, or y; the diffraction limit is assumed to be 1.22 times the working F/# on axis times the wavelength (primary wavelength if polychromatic). The change in the diffraction limit due to changes in F/# with field are ignored; a single value is used across the range of the plot. For Strehl ratio, 0.8 is used, and for RMS wavefront, 0.072 waves is used. These are all approximate indicators for convenience only; the actual meaning of "diffraction limited" may be open to interpretation.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

This feature calculates the RMS error or Strehl ratio as a function of field angle for each wavelength and gives a wavelength-weighted polychromatic result as well.

Two different methods of calculation are used; either a Gaussian quadrature method, or a rectangular array of rays. For the Gaussian quadrature method, the rays traced are arranged in a radial pattern with an optimal weighting to estimate the RMS with a minimum number of rays. The method is described in a paper by G. W. Forbes (JOSA A 5 P1943). Although the method is very efficient, the algorithm is not accurate if some of the rays are clipped due to surface apertures.

For wavefront RMS computations, ZEMAX automatically subtracts out the mean OPD, which yields in effect the standard deviation rather than the true RMS. However, ZEMAX uses the term RMS in this case to comply with common conventions in the optics industry. To compute the RMS wavefront in systems with surface apertures requires the use of the rectangular array method, and a larger number of rays for sufficient accuracy.

RMS vs. Wavelength

Purpose:

Plots RMS radial, x, and y spot size, RMS wavefront error, or Strehl ratio as a function of wavelength.

Settings:

Item	Description
Ray Density	If the method is Gaussian quadrature, then the ray density specifies the number of radial rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. The maximum number is 20 which is sufficient for pupil aberrations up to order 40. If the method is rectangular array, then the ray density indicates the grid size. Rays outside the circular entrance pupil are ignored. See the "Discussion" section for details.
Wave Density	The wave density is the number of points between the minimum and maximum defined wavelength at which the RMS/Strehl ratio is calculated, intermediate values are interpolated. A maximum of 100 points is allowed.
Plot Scale	Sets the maximum vertical scale for the plot. Zero results in automatic scaling.
Method	Selects either Gaussian quadrature or rectangular array. The Gaussian quadrature method is very fast and accurate, but only works if there is no vignetting. If any rays are to be vignetted, then rectangular array is more accurate.
Data	Selects wavefront error, spot radius, spot x-direction, spot y-direction, or Strehl ratio.
Refer To	Select either chief ray or centroid. Both reference points subtract out wavefront piston. The centroid reference mode also subtracts out the tilt of the wavefront, which yields smaller RMS values.
Use Dashes	Selects color or dashes.
Field	Select "All" to display data for all fields, select one field to display data for a single field.
Show Diffraction Limit	If checked, then a curve indicating the diffraction limited response will be drawn on the plot. For RMS radius, x, or y; the diffraction limit is assumed to be 1.22 times the working F/# on axis times the wavelength. The change in the diffraction limit due to changes in F/# with field are ignored. For Strehl ratio, 0.8 is used, and for RMS wavefront, 0.072 waves is used. These are all approximate indicators for convenience only; the actual meaning of "diffraction limited" may be open to interpretation.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Discussion:

This feature calculates the RMS error or Strehl ratio as a function of wavelength for each field position. See the "RMS vs. Field" feature for more details.

RMS vs. Focus

Purpose:

Plots RMS radial, x, and y spot size, RMS wavefront error, or Strehl ratio as a function of focus change.

Settings:

Item	Description
Ray Density	If the method is Gaussian quadrature, then the ray density specifies the number of radial rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. The maximum number is 20 which is sufficient for pupil aberrations up to order 40. If the method is rectangular array, then the ray density indicates the grid size. Rays outside the circular entrance pupil are ignored. See the "Discussion" section for details.
Focus Density	The focus density is the number of points between the minimum and maximum focus shift specified at which the RMS/Strehl ratio is calculated, intermediate values are interpolated. A maximum of 100 points is allowed.
Plot Scale	Sets the maximum vertical scale for the plot. Zero results in automatic scaling.
Method	Selects either Gaussian quadrature or rectangular array. The Gaussian quadrature method is very fast and accurate, but only works if there is no vignetting. If any rays are to be vignetted, then rectangular array is more accurate.
Data	Selects either RMS wavefront error, RMS spot size (radial), RMS X-direction, RMS Y-direction, or Strehl ratio.
Use Dashes	Selects color or dashes.
Refer To	Select either chief ray or centroid. For monochromatic calculations, the specified wavelength is used for reference. For polychromatic calculations, the primary wavelength is used for reference. Both reference points subtract out wavefront piston. The centroid reference mode also subtracts out the tilt of the wavefront, which yields smaller RMS values.
Wavelength	Select "All" to display data for a polychromatic computation, select any one wavelength to plot monochromatic data.
Min Focus	The minimum value of the defocus to plot. The units are lens units.
Max Focus	The maximum value of the defocus to plot. The units are lens units.
Show Diffraction Limit	If checked, then a horizontal line indicating the diffraction limited response will be drawn on the plot. For RMS radius, x, or y; the diffraction limit is assumed to be 1.22 times the working F/# times the wavelength (primary wavelength if polychromatic) on axis. The change in the diffraction limit due to changes in F/# with field are ignored; a single value is used across the range of the plot. For Strehl ratio, 0.8 is used, and for RMS wavefront, 0.072 waves is used. These are all approximate indicators for convenience only; the actual meaning of "diffraction limited" may be open to interpretation.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Discussion:

This feature calculates the RMS error or Strehl ratio as a function of a change in focus position for each field position. The method of calculation is identical to that described in the previous section on RMS vs. Field; see that section for a detailed discussion. ZEMAX adds the specified focus shift to the value of the thickness of the surface prior to the image plane. If the optical system has an odd number of mirrors, this surface normally would have a negative thickness, and therefore negative focus values move the image plane farther away from the last component. For systems with an even number of mirrors, then negative focus values move the image plane closer to the last component.

Encircled energy

Diffraction

Purpose:

Encircled energy diagram. This is the percentage of total energy enclosed as a function of distance from either the chief ray or the image centroid at the image of a point object.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Type	The analysis type option specifies how the encircled energy is calculated; either encircled (radial), X-only, Y-only, or ensquared.
Maximum Distance	This setting overrides the default scaling. The units are microns. To choose the default scaling option, enter zero.
Use Dashes	Selects either color or dashes.
Refer To	Select chief ray, centroid, or vertex as the reference point. Vertex refers to the coordinates (0, 0) on the image surface. This option will only return meaningful data if the diffraction image at all selected fields is within the maximum distance of the vertex. When vertex is selected, ZEMAX is unable to detect if the sampling is sufficient, so some care should be taken to set the sampling high enough for accurate results.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Show Diffraction Limit	If checked, the diffraction limited results are computed and displayed. See the discussion below.
Use Huygens PSF	If checked, the more accurate but slower Huygens PSF method is used to compute the PSF. This option should always be used if the image surface is tilted, or if the chief ray is not close to normal to the image surface.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Discussion:



See the discussion sections of the FFT and Huygens Point Spread Functions. Those comments also apply to this feature.

The accuracy of the diffraction encircled energy calculation is limited by the magnitude and slope of the OPD error and the sampling density used. If the sampling density is insufficient, ZEMAX will issue an error message indicating that the data is inaccurate. To increase the accuracy, increase the sampling density or decrease the OPD error. If shown, the diffraction limit curve is for the aberration free response at the reference field position (see "Diffraction Limited" in the chapter "Conventions and Definitions").

Geometric

Purpose:

Computes encircled energy using ray-image plane intercepts.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Type	The analysis type option specifies how the encircled energy is calculated; either encircled (radial), X-only, Y-only, or ensquared.
Maximum Distance	This setting overrides the default scaling. The units are microns. To choose the default scaling option, enter zero.
Use Dashes	Selects either color or dashes.
Refer To	Select either chief ray or centroid as the reference point.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number for which the calculation should be performed.
Multiply by Diffraction Limit	If checked, ZEMAX approximates the diffraction encircled energy by scaling the geometric data by the theoretical diffraction limit curve computed for a rotationally symmetric Airy disk. The only way to compute an obscured or asymmetric pupil diffraction limit function would be to perform an exact diffraction calculation, in which case the diffraction encircled energy feature should be used instead. The diffraction limit approximation is only useful for systems with unobscured pupils, reasonably rotationally symmetric images, and modest field angles since the approximation ignores the change in F/# with field.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. Only ZEMAX-EE supports this capability.

Discussion:

The X- and Y-only options will compute the fraction of rays which are contained with plus or minus the specified distance from either the chief ray or the image centroid. If a scale of 10 microns is shown, then the region enclosed is 20 microns across (and infinite in the other direction). The geometric encircled energy is not a good indicator of performance if the system is close to diffraction limited.

Line/Edge Response

Purpose:

Computes the geometric response to a line object and an edge object.

Settings:

Item	Description
Sampling	The size of the ray grid used to sample the pupil. The sampling may be 32x32, 64x64, etc. Although higher sampling yields more accurate data, calculation times increase.
Maximum Radius	The maximum radius setting overrides the default scaling. The units are microns. To choose the default scaling option, enter zero.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number to be used in the calculation.

Item	Description
Type	The type option specifies which data is displayed on the graph; line and edge, line only or edge only.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Discussion:

The line response function (or line spread function, LSF) is the cross section of the intensity pattern of the image of a line object. The edge response function (ERF) is the cross section of the intensity pattern of the image of an edge (a semi-infinite plane). The tangential and sagittal data refer to the orientation of the line or edge. This is a geometric calculation.

Extended Source

Purpose:

Computes encircled energy using an extended source similar to the geometric image analysis feature.

Settings:

Item	Description
Field Size	This value defines the full width of the square image file in field coordinates, which may be either lens units or degrees, depending upon the current field definition (heights or angles, respectively).
Rays x 1000	This setting determines approximately how many rays will be traced. The number of rays traced is approximately 1000 times the specified value. The reason the number of rays is only approximate is because the distribution of rays over the pixels in the image must be uniform. For example, if there are 1500 pixels in an image file, then at least 1500 rays will be traced, even if a value of 1 is selected. The distribution of rays at each wavelength is in proportion to the wavelength weights.
Type	The analysis type option specifies how the encircled energy is calculated; either encircled (radial), X-only, Y-only, or ensquared. The X-only and Y-only options are sometimes called "enslitted" and correspond to the total fraction of energy contained within an expanding slit. There are also options for X- or Y-distributions, which show the energy distribution in either x or y directions. These latter two options also report the geometric full width at half max. The X- or Y- distributions are the amount of energy falling on a pixel which is narrow in one direction and infinite in the other direction.
Refer To	Select either chief ray or centroid as the reference point.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Multiply by Diffraction Limit	If checked, ZEMAX approximates the diffraction encircled energy by scaling the geometric data by the theoretical diffraction limit curve. The diffraction limit curve ZEMAX uses is based upon the unobscured circular pupil. The only way to compute the obscured pupil diffraction function would be to perform the exact diffraction calculation, in which case the diffraction limit approximation would be used instead. The diffraction limit approximation encircled energy feature should be used for pupils and modest field angles, since the approximation ignores the change in F/# with field.
Wavelength	The wavelength number to be used in the calculation.

Item	Description
Field	The image file may be centered on any defined field position. This permits a small target such as a bar chart to be moved to any location in the field of view.
File	The name of the .IMA image file. This file must reside in the \ImaFiles directory. See the discussion section in the Geometric Image Analysis feature for a full description of the IMA file format.
Max Distance	This setting overrides the default scaling. The units are microns. To choose the default scaling option, enter zero.
Use Dashes	Selects either color or dashes.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

The X- and Y-only options will compute the fraction of rays which are contained with plus or minus the specified distance from either the chief ray or the image centroid. If a scale of 10 microns is shown, then the region enclosed is 20 microns across (and infinite in the other direction). The geometric encircled energy is not a good indicator of performance if the system is close to diffraction limited.

See the Geometric Image Analysis feature discussion for details about extended source modeling and the IMA file format.

Illumination

Relative Illumination

Purpose:

Computes the relative illumination as a function of radial field coordinate for a uniform Lambertian scene.

Settings:

Item	Description
Ray Density	The number of rays on one side of an array of rays used to integrate the illumination of the exit pupil. A value of 10 will trace about $10 \times 10 \times \pi / 4$ or 78 rays. Higher ray densities yield more accurate results at the expense of longer computation times.
Field Density	The number of points along the radial field coordinate to compute the relative illumination for. Larger field densities yield smoother curves.
Wavelength	Selects the wavelength for computation. Relative illumination is a monochromatic entity.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Log Scale	If checked, a logarithmic rather than linear scale will be shown.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

This feature computes the relative illumination (RI) as a function of radial y field coordinate. RI is defined as the intensity of illumination per unit area of image surface normalized to the illumination at the point in the field that has maximum illumination (which may not be on axis). The computation considers apodization, vignetting, apertures, aberrations of both the image and pupils, variations in F/#, chromatic aberrations, image surface shape, angle of incidence, and optionally, polarization effects assuming unpolarized light. The method is based upon one described in "Relative illumination calculations" by M. Rimmer, SPIE Vol. 655, pp. 99 (1986). The

published method was extended to include apodization, transmission, polarization, and non-planar image surface effects. The computation method assumes the following are all true:

1. The object scene is plane, uniform, and Lambertian.
2. The image surface is a reasonably good conjugate (that is, an image) of the object surface, so that rays coming from small patches of light on the object surface are imaged to patches of light on the image surface. Aberrations are fine, but the rays should be reasonably localized on the image surface.
3. The exit pupil is not too close to the image surface. This condition will be satisfied if the F/# is larger than about 0.1 and the ray aberrations are small compared to the exit pupil distance.

The relative illumination is computed by integration of the effective area of the exit pupil as seen from the image point(s). The integration is carried out in direction cosine space using a uniform grid in image cosine space.

Note that the RI computation will not in general yield a cosine-fourth law curve, because the cosine-fourth law is in fact a rough approximation based upon a thin, aberration free lens with the stop at the lens illuminating a plane image surface. For general systems including telecentric systems, systems with pupil or image aberrations, or vignetting, the RI can be computed using an integration of the solid angle or effective area of the exit pupil as seen from the image location.

For systems with very high amounts of vignetting, or for systems with non-linear cosine space aberrations that would violate the assumptions of the computation, the relative illumination cannot be calculated, and an error message will be displayed. Cosine space aberrations may be displayed using the spot diagram feature.

Vignetting Plot

Purpose:

Calculates fractional vignetting as a function of field angle.

Settings:

Item	Description
Ray Density	The ray density specifies the number of rays to be traced. The more rays traced, the greater the accuracy, although the computation time increases. For a ray density of n, ZEMAX traces a grid of (2n+1) x (2n + 1) rays at each field point.
Field Density	The field density is the number of points between zero degrees and the maximum field angle specified at which vignetting is calculated, intermediate values are interpolated.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

Fractional vignetting is the percentage of rays incident upon the entrance pupil which pass all obscurations and apertures in the system and survive to the image plane, normalized to relative pupil area. The graphic generated by this function shows fractional vignetting as a function of field position. If too few rays are used, the results may be inaccurate. This is especially true in systems with many apertures and large field angles.

Only the primary wavelength is used in this calculation. This is a geometric calculation. Only positive y field positions are used, therefore this feature is only appropriate for rotationally symmetric lenses and fields. Rays which cause errors such as missing a surface or those which are TIR are considered vignetted.

See also the relative illumination feature.

Illumination XY Scan

Purpose:

Computes relative illumination for an extended source along a line across the image surface.

Settings:

Item	Description
Sampling	Selects the grid size used to sum the illumination on the image surface. The sampling determines how many pixels are used to collect the ray data.
Rays x 1000	Determines the approximate total number of rays to be traced in computing the illumination due to the extended source.
File	The name of the IMA file used to define the shape of the extended source. See the Geometric Image Analysis feature for details.
Source Size	The full width in field units of the extended source.
Rotation	The angle to rotate the extended source in object space about the normal to the center of the extended source.
Show As	Select either X or Y direction scan.
Smoothing	Smoothing helps remove jagged lines due to low ray sampling by applying a smoothing operator to average the data in adjacent pixels.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number indicates which field position is used as the center reference point for the extended source.
Surface	The scan may be computed at any surface, however, the relative illumination computation is only accurate at the image surface.
Detector Size	The total width of the detector in lens units. The detector size is divided into pixels according to the "Sampling" setting.
Use Relative Illumination	If checked, then the RI computation described in the feature "Relative illumination" is used to weight the rays from various points in the field of view to accurately account for the effects of exit pupil radiance and solid angle. The computation is generally more accurate, but slower, if this feature is used. RI may only be used if the selected surface is the image surface, and ZEMAX assumes the image surface is a reasonable conjugate of the object surface. See the RI feature for details.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

The illumination XY scan is similar to the relative illumination (RI) feature, with the added capability to estimate the RI for non-uniform extended sources. For uniform extended Lambertian sources, the RI feature is faster and more accurate. However, for systems with complex source properties, the illumination XY scan can estimate illumination by Monte Carlo ray tracing combined with the conventional RI computation.

The extended sources are defined in the same way as described in the Image Analysis feature.

Illumination 2D Surface

Purpose:

Computes relative illumination for an extended source over a 2D surface.

Settings:

The options are identical to the XY scan, except that a 2D surface is drawn as either an isometric surface or contour plot, or as a grey scale or false color map.

Discussion:

See the illumination XY scan feature.

Image Analysis

Geometric Image Analysis

Purpose:

The geometric image analysis feature has many applications. It can be used to model extended sources, analyze useful resolution, represent the appearance of imaged objects, and provide intuition as to image rotation. Image analysis is also useful for estimating multi-mode fiber coupling efficiency.

This feature is based strictly upon geometrical ray tracing; see the Diffraction Image Analysis for a diffraction based feature with similar functionality.

The ability of this feature to render color images is limited, for analyzing color bitmap images see "Geometric Bitmap Image Analysis" on page 107.

The image analysis feature uses special IMA or BIM files to describe the object to be imaged. The IMA and BIM file formats are described in the discussion section which follows.

Settings:

Item	Description
Field Size	This value defines the full width of the square image file in field coordinates, which may be either lens units or degrees, depending upon the current field definition (heights or angles, respectively).
Image Size	This value sets the size of the scale bar which is superimposed on the image diagram. It has no effect on the actual size of the image. The image size is set by the object scale and the magnification and aberrations of the system. The default may not be acceptable to see the desired portion of the image.
Parity	The "Even" setting leaves the object as it would appear when viewed looking down the negative Z axis in object space. The parity can be set to "Odd" which reverses the object from top to bottom.
Rotation	The rotation can be set to any angle in degrees. The algorithm actually rotates the object before tracing the rays, so this feature can be used to switch from tangential to sagittal orientation of bar targets, for example.
Rays x 1000	This setting determines approximately how many rays will be traced. The number of rays traced is approximately 1000 times the specified value. The reason the number of rays is only approximate is because the distribution of rays over the pixels in the image must be uniform. For example, if there are 1500 pixels in an image file, then at least 1500 rays will be traced, even if a value of 1 is selected. The distribution of rays at each wavelength is in proportion to the wavelength weights.
Show	Choose surface plot, contour map, grey scale, false color map, or spot diagram as the display option.
Source	The source may be uniform or Lambertian. The uniform setting weights all rays equally. Lambertian weights all rays by the cosine of the angle the ray makes with the axis of the object surface.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Scatter Rays	If checked, rays will be statistically scattered at ray-surface intercepts that have defined scattering properties. Only ZEMAX-EE supports this capability.
Wavelength	The wavelength number to be used in the calculation.

Item	Description
Field	The image file may be centered on any defined field position. This permits a small target such as a bar chart to be moved to any location in the field of view. The resulting image is then centered on the chief ray coordinate of this field position.
File	The name of the .IMA or .BIM image file. This file must reside in the \ImaFiles directory. See the discussion section for a full description of the IMA and BIM file formats.
Edit IMA File	Pressing this button will invoke the Windows Notepad editor which allows modification of the currently selected IMA file. This button is disabled if the file type is BIM.
Surface	The surface number at which to evaluate the rays. The default is the image surface. Other surfaces may be selected, for example, to visualize the beam footprint on an optical surface.
# Pixels	The number of pixels across the width of the selected image size. This value is not used if "spot diagram" is the method of displaying the image data.
NA	The numerical aperture (NA) cut-off. If zero, this feature is ignored. If a number greater than 0 is entered, then all rays with a numerical aperture greater than the specified number are ignored.
Total Watts	The total power in watts radiated by the source into the entrance pupil of the optical system. This flux is then used to normalize the detected power according to the relative pixel values and the total efficiency.
Use Symbols	If checked, this option will draw different symbols rather than dots for each wavelength. This helps distinguish the various wavelengths. This value is only used if "spot diagram" is the method of displaying the image data.
Configuration	Select "All" to draw all configurations at once, or select the one configuration to draw, or select "Current" to show the active configuration.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

ZEMAX supports three different file formats. Two of these formats end in the IMA extension, one in the BIM extension.

The IMA format

There are two different IMA file formats, one ASCII and one binary. Whichever file format is used, the file must end in the extension IMA. ZEMAX will distinguish between the two types of file formats automatically.

The ASCII image file is a text file which ends in the extension .IMA. At the top of the file is one number which indicates the size of the file in pixels. The remaining rows and columns contain the pixel data, with one character to each pixel. All IMA files must be square, with $n \times n$ pixels defined. For example, a 7 x 7 representation of the letter "F" could be described by the following IMA file:

```
7
0111110
0100000
0100000
0111100
0100000
0100000
0100000
```

Note that the single entry "7" starts the file, and it is followed by a carriage return. Then there are 7 rows of 7 columns each, and each row is terminated with a carriage return. The columns are not separated by a space or any other character. The image file must be square. ZEMAX will attempt to allocate enough memory to hold the image file and will report an error if there is not enough memory.

The "intensity" at each pixel can be any digit between 0 and 9. The number of rays each pixel will generate is proportional to this value. Pixels with a value of 0 do not radiate any rays.

The binary IMA file format is more complicated than the ASCII format, and binary IMA files cannot be edited with a text editor. However, the binary IMA files are dramatically more powerful. Each pixel in the binary IMA file is represented by an unsigned byte, which means there are 256 "gray-scale" levels of intensity. Furthermore, each wavelength can be assigned a separate pixel map. Therefore, very realistic photograph like extended sources can be modeled.

The binary IMA file format requires 3 16-bit header values. The first 16-bit value is a signed integer that must be equal to zero. The second 16-bit signed integer is the width of the pixel map in pixels, which can be any number from 1 to 4000. The third 16-bit signed integer is the number of pixel maps, which correspond to the number of colors (or wavelengths) represented in the file.

For example, a 3-color binary pixel map of a 50 by 50 image would have 6 bytes of header (0, 50, and 3), followed by 2500 bytes for color 2, then 2500 bytes for color 1, then 2500 bytes for color 3, for a total of 7506 bytes. The data for each color is stored by columns for each row (the column index changes faster than the row index).

The BIM format

The drawback to the IMA format is that a maximum of 256 grey scale levels are supported. The BIM format is a binary double precision floating point file format which effectively makes the number of grey scales many trillions. The BIM format consists of the following binary values:

1 32 bit integer representing the number of x pixels, n_x .

1 32 bit integer representing the number of y pixels, n_y .

followed by $n_x \times n_y$ 64 bit double precision floating point values representing the relative intensity.

Currently, the n_x and n_y values must be identical or an error message will be issued.

How rays are chosen for analysis

The rays generated by each pixel are chosen randomly from coordinates within the pixel cell. The entrance pupil coordinates are also chosen randomly for each ray. The distribution of rays is uniform over the pixel and over the circular paraxial entrance pupil (if ray aiming is used, then there may be some pupil distortion). For the ASCII IMA files, the number of rays generated by each pixel is equal to the pixel intensity times the number of wavelengths times the ray density. The wavelength used for each ray is selected randomly in proportion to the wavelength weights provided on the wavelength data screen. For the binary IMA files, the number of rays generated from each pixel is proportional to the ray density times the fractional intensity relative to 256.

The field size determines the physical size of the image file as seen by the optical system. For example, if a 30 x 30 pixel size image file is used, and the field size is 2.0 mm (this assumes the fields are defined in terms of object or image heights), then each pixel represents a 66.67 micron region. If the same image file is later used with a system with a 40 degree full field of view (using field angles), then the field size can be set to 40 to cover the entire field. Each pixel will now represent a 1.33 degree square. The difficulty in using field angle for defining the object field of view is that field angle units are inherently anamorphic. X-direction angles represent a different subtended angle at a Y angle of 80 degrees than at an Y angle of 10 degrees. If field angles are being used, and the field of view is fairly wide (more than about 40 degrees in any direction) then great care should be taken in interpreting the results for an extended object. For a precise definition of the field angles ZEMAX uses, see "Field angles and heights" on page 27.

By separating the form of the object from the scale, the same image file can be used for many applications. For example, the sample image file "letterf.ima" contains a 7x7 grid of pixels defining the capital letter F. The object scale can be set to 1 mm, then 0.1 mm, then 0.01 mm to get a feel for how small a character F the optical system can resolve, without the need to change the IMA file.

Note that if fields are defined by image height, then field size determines the size of the object in image space, not object space. The field size is always in whatever units the fields are defined in, and so for image height the field size determines image height. The size of the object is then determined by the field size divided by the magnification of the lens.

The choice of field position also permits great flexibility in analyzing image quality. For example, the letter F image file can be tested at several field points to see if the resolution is strongly affected by field aberrations. The object scale is set to the height of the letter, but the image will be centered about the chief ray intercept of the selected field point.

The source is by default a uniform radiator of rays. Uniform here means uniform in the entrance pupil. All rays generated fall within the entrance pupil, and they are all weighted equally. Since ray wavelengths are selected randomly in proportion to the wavelength weights, no explicit wavelength weighting is required. The uniform setting is usually preferred for large object distance systems with small fields of view. The source may also be defined to be Lambertian, which weights all rays by a cosine factor.

The percent efficiency is defined by

$$\%E = \frac{\sum W_i}{\sum W_j}$$

where the sum i is over all rays which were unvignetted, and the sum over j is over all rays which were launched. The efficiency calculation considers vignetting, source distribution, wavelength weights, and reflection and transmission losses in the optical system if the "Use Polarization" checkbox is selected.

Calculating efficiency of multi-mode fibers

ZEMAX has a algorithm for accurately computing fiber coupling into single-mode fibers; for details see "Fiber Coupling Efficiency" on page 128.

To estimate the coupling efficiency for multi-mode fibers, a geometric approach may be used. Place a circular aperture at or just before the image surface with the appropriate maximum radial aperture representing the core size. Then set the NA (see the table above) to the maximum acceptance NA of the fiber. The percent efficiency will then be calculated by summing all the rays that pass the core aperture within the specified NA. The NA of a typical multi-mode fiber with an inner core of index n_i and an outer cladding of index n_o is given by

$$NA = \sqrt{n_i^2 - n_o^2}$$

Text output

Selecting the "Text" option on the image analysis window menu bar will generate and display an ASCII file listing the ray data. If the "Show" option is set to "Spot Diagram", the file will have 9 columns. The first column is the sequential ray number. The second and third columns are the x and y field coordinates (either degrees or object height). The fourth and fifth columns are the normalized pupil coordinates, P_x and P_y . The sixth column is the integer wavelength number. The seventh column is the weight of the ray, which depends upon the source properties. The eighth and ninth columns are the image coordinates in lens units, relative to the reference ray.

If the "Show" option is not set to "Spot Diagram", then the text display will list the weighted ray count in each pixel. Use the "Escape" key to terminate a lengthy image analysis computation.

Geometric Bitmap Image Analysis

Purpose:

This feature creates an RGB color image based upon ray tracing data using an RGB bitmap file as the source. This feature has many applications. It can be used to model extended sources, analyze useful resolution, display distortion, represent the appearance of imaged objects, provide intuition as to image rotation, display beam footprints, and show surface plots of illumination on any surface, to name just a few.

This feature is based strictly upon geometrical ray tracing. The bitmap image analysis feature uses standard Windows BMP and JPG files as the source image, see the discussion for details.

Settings:

Item	Description
Field Y Height	This value defines the full height of the source bitmap in field coordinates, which may be either lens units or degrees, depending upon the current field definition (heights or angles, respectively).
Parity	The "Even" setting leaves the object as it would appear when viewed looking down the negative Z axis in object space. The parity can be set to "Odd" which reverses the object from top to bottom.
Rotation	The rotation can be set to any angle in degrees. The algorithm actually rotates the object before tracing the rays, so this feature can be used to switch from tangential to sagittal orientation of bar targets, for example.
Rays x 1000	This setting determines how many rays will be traced. The distribution of rays at each wavelength is in proportion to the relative RGB intensities of each pixel.

Item	Description
X-Pixels	The number of pixels across the X direction on the detector.
Y-Pixels	The number of pixels across the Y direction on the detector.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.
Grey Scale	If checked, the RGB intensities will be averaged at each detector pixel to yield a grey scale detected image. Rays will still be traced according to the relative RGB intensities of the source bitmap, but all color information will be lost when the detected image is displayed.
Wavelength	If "RGB" is selected, then 3 wavelengths will be defined, 0.656, 0.587, and 0.486 microns for red green, and blue, respectively; no matter what the current wavelength definitions are. If "1+2+3" is selected, then wavelengths 1, 2, and 3 as currently defined on the wavelength data box will be used. The red channel of the source bitmap will be used for wavelength 3, green channel for wavelength 2, and the blue channel for wavelength 1. The displayed image will be in RGB format no matter what wavelengths are defined using this option. For selection of a specific wavelength, such as 1, 2, 3, etc... the B, G, or R channel image will be used; for wavelengths higher than 3 the B channel is always used, for wavelengths selected which are not defined the highest defined wavelength will be used.
Field	The source image may be centered on any defined field position. This permits a small target such as a bar chart to be moved to any location in the field of view. The resulting image is then centered on the chief ray coordinate of this field position.
Input	The name of the .BMP or .JPG source image file. This file must reside in the \lmaFiles directory.
Surface	The surface number at which to evaluate the rays.
X-pixel size	The size in lens units of each pixel measured in the X direction.
Y-pixel size	The size in lens units of each pixel measured in the Y direction.
Show	Select either "object" or "image". If object is selected, then the source bitmap will be drawn. If image is selected then rays will be traced and the detected image will be displayed. Note the number of rays, pixels, and pixel sizes are ignored when drawing the object bitmap.
Output	The name of the .BMP or .JPG to write the detected bitmap to. The detected bitmap size is determined by the number of x and y pixels defined; but the pixels size must be the same in x and y for the aspect ratio to be correct in the output bitmap file. The file name must end in either a BMP or JPG extension, with no path name supplied. This file will be created or overwritten without warning and will be placed in the \lmaFiles directory.
Remove Vignetting Factors	If checked, vignetting factors are automatically removed. See "Comment about vignetting factors" on page 82.

Discussion:

See also the Geometric Image Analysis feature discussion, as this feature is very similar.

The BMP files used as source bitmaps must be standard Windows format, uncompressed, 24-bit color RGB bitmaps. Not all files ending in the .BMP extension meet this definition. Contact Focus Software Technical Support if you have a BMP file that does not work correctly for advice on converting the file to the proper format. Any properly formatted JPG file will work.

The rays generated by each pixel are chosen randomly from coordinates within the source pixel cell. The entrance pupil coordinates are also chosen randomly for each ray. The distribution of rays is uniform over the pixel, the ray distribution over the entrance pupil is uniform unless pupil apodization is defined.

The Field Y Height determines the physical size of the source file as seen by the optical system. For example, if a 50H x 100W pixel source file is used, and the Field Y Height is 2.0 mm (this assumes the fields are defined in terms of object or image heights), then each pixel represents a 0.040mm x 0.040mm region, and the source bitmap covers a 2.0mm high x 4.0 mm wide area.

Note that if fields are defined by image height, then Field Y Height determines the size of the object in image space, not object space. The Field Y Height is always in whatever units the fields are defined in, and so when using image height as a field type the Field Y Height determines the source bitmap height in image space. The size of the source in object space is then determined by the Field Y Height divided by the magnification of the lens.

The pixels from which rays are traced are chosen randomly according to the RGB intensity of each pixel in each color channel. Rays are generated randomly from anywhere within the bitmap until the specified number of rays have been launched, although not all these rays may make it through the lens.

Once a ray is generated, it is traced through to the selected surface. If the ray is vignetted or an error occurs, the ray is ignored. Otherwise, the pixel on the receiving detector that the ray struck is determined, and the intensity of the ray is added to the detector's bin count in the appropriate color channel.

After all the rays are traced, an RGB image is created from the normalized counts in each detector pixel. The percent efficiency is defined by

$$\%E = \frac{\sum W_i}{\sum W_j}$$

where the sum i is over all rays which were unvignetted, and the sum over j is over all rays which were launched. The efficiency calculation also includes consideration of reflection and transmission losses in the optical system if the "Use Polarization" checkbox is selected, as well as ray errors.

Selecting the "Text" option on the menu bar will generate and display an ASCII file listing the detected pixel data. Use the "Escape" key to terminate a lengthy image analysis computation.

Diffraction Image Analysis

Purpose:

The diffraction image analysis feature is similar to the geometric image analysis feature, except the complex system Optical Transfer Function (OTF) is used to compute the image appearance. This method accounts for the finite pass band and other diffraction related effects real optical systems have on image formation.

The image analysis feature uses special IMA files to describe the object to be imaged. The IMA file is described in "The IMA format" on page 105.

Settings:

Item	Description
Image Size	The full width in lens units of the region defined by the IMA file. Note IMA files are always square.
Oversampling	Sets the factor by which the IMA file pixels are oversampled. This increases the effective resolution of the IMA file without the need to define a new IMA file.
Zero Padding	Determines the actual size of the region to compute the diffraction image of by adding zero intensity values around the IMA file pixels. This increases the size of the displayed diffraction image without changing the size of the unaberrated image; this allows study of the energy diffracted well away from the perfect image location.
OTF Sampling	The grid size of the sampling in the pupil; larger grids yield a more accurate representation of the system OTF. This has no affect on the size of the diffracted image, just the accuracy of the predicted frequency response.
Show As	Choose surface plot, contour map, grey scale, or false color map as the display option.

Item	Description
Data Type	Choose either the incoherent image, coherent image, raw image, incoherent transfer function, coherent transfer function, or the transform of the raw image. See the discussion below for information about the coherent transfer function and the limitations of this computation.
Diffraction Limited	If checked, aberrations are ignored. Apertures are still considered.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number indicates for which field position the optical transfer function is computed.
File	The name of the .IMA image file. This file must reside in the \ImaFiles directory. See the discussion section for a full description of the IMA file format.
Edit IMA File	Pressing this button will invoke the Windows Notepad editor which allows modification of the currently selected IMA file.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Discussion:

This feature can compute complex diffraction image properties from extended sources. The method involved in the computation is based upon Fourier Optics, which is described in clear and insightful detail in "Introduction to Fourier Optics" by Joseph Goodman, McGraw-Hill 1968. See that reference for more information on coherent vs. incoherent imaging and other Fourier optics theory. There are several important assumptions in the method that must be understood by the user before this feature may be used to draw important conclusions.

The IMA file defines the relative intensity at each of an arbitrary number of pixels arranged on a square grid. For example, the letter "F" may be described by a 7 x 7 grid such as:

```
0111110
0100000
0100000
0111110
0100000
0100000
0100000
```

The image size parameter determines how big each pixel is in *the image space of the optical system*. Note this is different from the geometric image analysis feature, where the IMA file defines the size and shape in "field" space, which maybe either object or image space. For the diffraction image analysis, the IMA file defines the ideal image shape in image space. If the image size is 0.1 mm, then each pixel is 14.286 microns wide for this 7 x 7 pixel image.

Although the IMA file is convenient for defining simple shapes, the resolution of files created by hand is generally too low to see sufficient detail in the diffracted image. The oversampling option remedies this problem by increasing the resolution; the number of pixels is increased, and the data from each pixel is replicated as required to yield the same shape at higher resolution. With an oversampling of 2X; the letter "F" file IMA image becomes a 14 x 14 grid:

```

0011111111100
0011111111100
001100000000000
001100000000000
001100000000000
001100000000000
00111111110000
00111111110000
001100000000000
001100000000000
001100000000000
001100000000000
001100000000000
001100000000000
001100000000000
001100000000000

```

Note the shape is the same, there is just twice the sampling in each direction. The width of the image is unaffected; the pixels are half the width and there are twice as many in each direction now.

Because the effects of diffraction tend to blur and extend the ideal image, it is desirable to increase the size of the displayed area beyond the limits of the IMA file. This can be done with the zero padding option. This option increases the IMA file size by adding zero intensity values around the defined IMA file pixels. The letter "F" file would look like this with a zero padding of 2X:

```

000000000000000
000000000000000
000000000000000
000000000000000
00000111110000
000001000000000
000001000000000
00000111100000
000001000000000
000001000000000
000001000000000
000001000000000
000001000000000
000000000000000
000000000000000
000000000000000

```

The size and shape of the image is the same, but additional area has been defined that some energy may now diffract into. Note the width of the new image has increased by the zero padding factor, but the width of the "F" part of the image is the same. Zero padding is not available when using the binary IMA file format.

The oversampling and zero padding may be used together, however the array sizes quickly become quite large. For an oversampling of 8X, a zero padding of 4X, and an original pixel size of 12 x 12, the resulting array becomes 384 x 384, which ZEMAX internally further zero pads up to 512 x 512 for transform purposes.

Once the input image is defined, the image is transformed into frequency space, multiplied by the OTF, and transformed back into position space. The resulting image is filtered by the complex OTF to yield the diffraction image.

The primary assumption made by this implementation of the Fourier method is that the OTF does not change over the extent of the image size. This means the field of view defined by the size of the image is small enough so that the OTF is the same over all points in the image. The user must take care to be sure the image size defined is small with respect to the rate of change of the field aberrations. ZEMAX computes the OTF for the field point selected, and assumes this OTF is valid over the entire region covered by the image.

Because of this assumption, distortion will not be visible in the predicted image, since only the variation in OTF over the field will introduce distortion. To see the effects of distortion or other "large field" effects, use the geometric image analysis feature.

Note that the diffraction image analysis feature is good at computing detailed image data for small images, while the geometric image analysis feature is good at computing image data for larger scale images.

The other assumption this calculation makes is in the method used to compute the coherent optical transfer function. The coherent optical transfer function is assumed to be the complex pupil function:

$$H(f_x, f_y) = P(\lambda d f_x, \lambda d f_y) e^{iKW(\lambda d f_x, \lambda d f_y)}$$

where H is the complex OTF, d_i is the pupil to image distance, f_x and f_y are the spatial frequencies, P is the pupil function (which determines the relative transmission over the pupil, and is zero outside the pupil), and W is the wavefront aberration function. This is an approximation which is only exact for the special case of the system stop being located at the Fourier transform plane of the image. This condition is met in general for systems which are telecentric in object space. For other systems, this is only an approximation to the coherent transfer function. For more information, see the reference by Goodman.

Miscellaneous

Field Curvature/Distortion

Purpose:

Displays the field curvature and distortion plots.

Settings:

Item	Description
Max Curvature	Defines the maximum scale for the field curvature plot in lens units. Enter zero for automatic.
Max Distortion	Defines the maximum scale for the distortion plot in percent. Enter zero for automatic.
Wavelength	The wavelength number to be used in the calculation.
Use Dashes	Selects either color or dashes.
Ignore Vignetting Factors	See the discussion section.
Distortion	Select Standard, F-Theta, or Calibrated. See the discussion section for details.
Do X-Scan	If checked, then the data is computed based upon a scan along the positive X field, otherwise, the scan is along the positive Y field.

Discussion:

The field curvature plot shows the distance from the currently defined focal (image) plane to the paraxial focal plane as a function of field coordinate. The tangential data are the distances measured along the Z-axis from the currently defined focal plane to the paraxial focal plane measured in the tangential (YZ) plane. The sagittal data are the distances measured in the plane orthogonal to the tangential plane. The base of the plot is on axis, and the top of the plot represents the maximum field (angle or height). There are no units on the vertical scale because the plot is always normalized to the maximum radial field.

The field curvature for the tangential and sagittal rays is defined as the distance from the defined image plane to the paraxial focus for that ray. In non-rotational systems, the real ray and the chief ray may never intersect, and so the value presented is at the point of closest approach.

The field scan is along the positive Y field by default. If "Do X-Scan" is selected, then the maximum field is along the positive X field, in which case the tangential curve is for the XZ plane, and the sagittal curve is in the YZ plane.

Users often ask why the field curvature plot does not always start at zero for zero field. The reason is that the plot shows the distance from the currently defined image plane to the paraxial focal plane, and the currently defined image plane need not be coincident with the paraxial focal plane. If there is any defocus, then the two planes are offset, and so is the field curvature data.

"Standard" distortion is defined as the real chief ray height, minus the reference ray height, divided by the paraxial chief ray height, times 100:

$$Distortion = 100 \times \frac{y_{chief} - y_{ref}}{y_{ref}}$$

where all heights are taken to be the image plane radial coordinate, at whatever image plane is defined (the data is not referred to the paraxial image plane). The paraxial reference ray height is computed by tracing a real ray

from a very small field height, and then scaling the results as required. This generalization permits the computation of reasonable distortion values even for systems not well described by paraxial ray tracing. The reference height for the undistorted ray in a rotationally symmetric system at paraxial focus is given by

$$y_{ref} = f \tan \theta,$$

where f is the focal length and θ is the angle in object space. ZEMAX actually uses the generalization described above, rather than this equation, however the idea is the undistorted height goes as the tangent of the field angle. "F-Theta" distortion does not use the tangent relationship, but instead uses the height given by the focal length of the lens multiplied by the angle the chief ray makes in object space. This so called "F-Theta" height is only meaningful in systems with the object at infinity, when field heights are measured in angles. The reference height for the undistorted ray is given by

$$y_{ref} = f\theta,$$

where f is the focal length and θ is the angle in object space. Generally, F-Theta is used only in scanning systems where the image height needs to be linear with scan angle.

"Calibrated" distortion is similar to "F-Theta" distortion, except that the "best-fit" focal length is used rather than the system focal length. Calibrated distortion measures the deviation from linearity between the image height and the field angle, without the restriction that the proportion of linearity be defined by the focal length of the system. A focal length is chosen that best fits the data, rather than the system focal length, although in general the best fit focal length is close to the system focal length. The calibrated focal length used is given on the text listing for this feature. The reference height for the undistorted ray is given by

$$y_{ref} = f'\theta,$$

where f' is the calibrated focal length and θ is the angle in object space. One somewhat surprising effect of this definition of calibrated distortion is the non-zero distortion at zero field angle. The reason for this non-zero distortion is best explained by looking at the limiting behavior of the definition of calibrated distortion. Calibrated distortion is defined as

$$Distortion = 100 \times \frac{y_{chief} - y_{ref}}{y_{ref}}.$$

For small angles, the real y chief ray coordinate in any reasonable optical system is well described by $y_{chief} = f\theta$, and the reference ray coordinate by $y_{ref} = f'\theta$, so the distortion is

$$Distortion = 100 \times \frac{y_{chief} - y_{ref}}{y_{ref}} = 100 \times \frac{f\theta - f'\theta}{f'\theta} = 100 \times \frac{f - f'}{f'}.$$

which is not zero, unless $f = f'$. Therefore, the calibrated distortion is not generally zero on axis. This does not mean the image height is not zero, it is simply an artifact of the different choice of focal lengths for the reference and actual ray coordinates near the optical axis. Note that a percentage distortion near the axis is not significant because the field itself approaches zero at zero field angle.

For non-rotationally symmetric systems, or systems with curved image planes, distortion is poorly defined and the data presented is probably meaningless. The reason is that no single number adequately describes distortion at a single field point if the system is not rotationally symmetric. Instead, see "Grid Distortion" on page 114.

Strictly speaking, the field curvature and distortion plots are only valid for rotationally symmetric systems with plane image surfaces. However, ZEMAX uses a generalization of the field curvature and distortion concepts to give reasonable results for some, but not all, non-rotationally symmetric systems. Caution should be used when interpreting these graphs for non-rotationally symmetric systems.

By default, ZEMAX ignores vignetting factors when computing the field curvature and distortion plots. Vignetting factors can change the chief ray location on the stop surface, such that the chief ray no longer goes through the center of the stop.

Grid Distortion

Purpose:

Displays a grid of chief ray intercept points to indicate distortion.

Settings:

Item	Description
Display	Select either "Cross" to mark each chief ray intercept with a cross, or "Vector" to plot a vector from the ideal image point to the actual chief ray image point.
Grid Size	The size of the grid.
Wavelength	The wavelength number to be used in the calculation.
Field	The reference field position. See the "Discussion".
Scale	If the scale is different than 1.0, then the "x" points on the distortion grid will be exaggerated by the selected scale factor. The scale factor may be negative to change positive to negative distortion on the plot.
H/W Aspect	If unity, then a square field will be selected. The output image may not be square if the system is not symmetric, but the object field will be square. If the H/W aspect is greater than 1, then the "height" or "y" field will be expanded by the aspect ratio. If the aspect is less than 1, then the "height" or "y" field will be compressed by the aspect ratio. The aspect ratio is the "y" field height divided by the "x" field width. The aspect ratio only affects the input field; the image aspect ratio is determined by the optical system imaging properties.
Symmetric Magnification	If checked, then the X magnification is required to be identical to the Y magnification. This causes distortion to be referenced to a symmetric predicted grid rather than an anamorphic predicted grid.
Field Width	The width, in field units, of the "x" field of view.

Discussion:

This feature displays or computes the coordinates of a grid of chief rays. In a system without distortion, the chief ray coordinates on the image plane follow a linear relationship with the field coordinate:

$$\begin{bmatrix} x_p \\ y_p \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} f_x \\ f_y \end{bmatrix},$$

where x_p and y_p are the predicted image coordinates relative to a reference image point and f_x and f_y are the linear coordinates on the object surface relative to a reference point. For optical systems using angles as a field definition, f_x and f_y are the tangents of the field angles (the field coordinates must be linear, therefore, tangents of angles rather than angles are used). To compute the ABCD matrix, ZEMAX traces rays over a very small region centered upon the reference field position. Usually, this is the center of the field of view. ZEMAX allows selection of which field position to use for reference.

By default, ZEMAX sets the corner of the field grid in object space to be at the maximum radial field distance. Because the object height is linear with field angle and not angle directly, the full width of the field when angles are used to define the field is given by

$$\theta_{wide} = 2 \tan^{-1} \left[\frac{\sqrt{2}}{2} \tan \theta_r \right],$$

where θ_r is the maximum radial field angle at the corner of the field.

The ray coordinates in image space for the very small field of view are used to determine the ABCD matrix components. The use of an ABCD matrix allows for coordinate rotations. If the image plane is rotated, such that a y object coordinate images to both an x and a y image coordinate, the ABCD matrix will automatically account for the rotation. The grid distortion plot shows the linear grid, and then marks the actual chief ray intercept for a ray with the same linear field coordinates with an "X" for each point on the grid.

The text listing available tabulates the predicted image coordinate, the actual image coordinate, and the "percent distortion" defined by

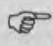
$$P = 100\% \frac{R_{\text{distorted}}}{R_{\text{predicted}}}, \text{ where}$$

$$R_{\text{distorted}} = \sqrt{(x_p - x_r)^2 + (y_p - y_r)^2}, \text{ and}$$

$$R_{\text{predicted}} = \sqrt{(x_p)^2 + (y_p)^2},$$

and the subscripts r and p refer to the real and predicted coordinates on the image plane relative to the reference field position image location, respectively. This definition may not be applicable in all cases, and the results should be used with caution.

ZEMAX cannot compute the grid distortion in exactly the manner described here if the field type is real image height. The reason is that when using real image height, ZEMAX iterates each ray trace to find the exact field coordinate to hit the desired image coordinate; implicitly removing the distortion. As a workaround, ZEMAX automatically changes the field type from real to paraxial image height for the purposes of this computation, and this may affect the accuracy of the results for some systems. Generally the results will be accurate if the distortion is small.

 **ZEMAX cannot compute the grid distortion in exactly the manner described here if the field type is real image height.**

Footprint Diagram

Purpose:

Displays the footprint of the beam superimposed on any surface. Used for showing the effects of vignetting and for checking surface apertures.

Settings:

Item	Description
Ray Density	Determines the number of rays traced across the half pupil; a setting of 10 will trace a grid of 21 x 21 rays.
Surface	The surface to show the beam footprint on.
Wavelength	The wavelength number to be used in the calculation.
Field	The field number to be used in the calculation.
Delete Vignetted	If checked, then rays which are vignetted by subsequent surfaces will not be drawn. Rays which are vignetted by prior surfaces are never drawn.

Item	Description
Use Symbols	If checked, this option will draw different symbols rather than dots for each wavelength. This helps distinguish the various wavelengths.
Configuration	Select "All" to draw all configurations at once, or select the one configuration to draw. Or select "Current" to show the active configuration.

Discussion:

This feature will draw the shape of the surface, and then superimpose on that surface a grid of rays. If there is no aperture on the surface, then a circular surface shape with a radial clear aperture of the semi-diameter value is shown. Otherwise, the shape of the aperture is shown. The surface aperture is always shown as centered in the frame; even if the aperture is decentered on the actual surface. If there is an obscuration on the surface, then the obscuration will be drawn along with the circular aperture defined by the semi-diameter.

The ray grid size is specified by the ray density parameter, and rays may be from any or all fields, at any or all wavelengths. Rays which are vignetted by surfaces prior to the surface shown are not drawn. Rays which are vignetted by the surface or subsequent surfaces are not drawn if "delete vignetted" is checked, otherwise, they are drawn. The ray set is apodized if any system pupil apodization is selected. The number of rays shown divided by the total number of rays launched is shown as a percentage.

Longitudinal Aberration

Purpose:

Displays the longitudinal aberration as a function of pupil height at each wavelength.

Settings:

Item	Description
Scale	Defines the maximum scale for the plot in lens units. Enter zero for automatic.
Wavelength	The wavelength number to be used in the calculation.
Use Dashes	Selects either color or dashes.

Discussion:

This feature computes the distance from the image surface to where a zonal marginal ray "focuses", or crosses the optical axis. The computation is performed only for the on axis field point, and only for zonal marginal tangential rays as a function of pupil zone. The base of the plot is on axis, and the top of the plot represents the maximum entrance pupil radius. There are no units on the vertical scale because the plot is always normalized to the maximum entrance pupil radius. The horizontal scale is in lens units, and represents the distance from the image surface to the point where the ray crosses the optical axis.

Because longitudinal aberration is defined in terms of the distance to the ray-axis crossing point, this feature may produce meaningless data for non-rotationally symmetric systems. Great care should be exercised in interpreting this plot for non-rotationally symmetric systems.

Lateral Color

Purpose:

Displays the lateral color as a function of field height.

Settings:

Item	Description
Scale	Defines the maximum scale for the plot in lens units. Enter zero for automatic.
Use Real Rays	Either real or paraxial rays may be used. The default is paraxial rays. To use real rays check this box.

Item	Description
All Wavelengths	If checked, then data for all defined wavelengths will be displayed. Each wavelength will be referenced to the primary wavelength. If not checked, then the difference between the shortest and longest wavelength rays will be used. See the discussion.
Show Airy Disk	If checked, then the Airy disk radius at the primary wavelength will be plotted on either side of the reference line to indicate the extent of the Airy disk.

Discussion:

This feature computes the lateral color either of two ways:

If "All wavelengths" is checked: The data computed is the distance on the image surface from the intercept of the chief ray at each wavelength to the primary wavelength chief ray intercept.

If "All Wavelengths" is not checked: The data computed is the distance on the image surface from the shortest wavelength chief ray intercept to the longest wavelength chief ray intercept.

The base of the plot is on axis, and the top of the plot represents the maximum field radius. Only positive field angles or heights in the Y direction are used. The vertical scale is always normalized to the maximum field angle or height. The horizontal scale is in microns. Either real or paraxial rays may be used.

If "Show Airy Disk" is check on, then the Airy disk radius will be approximated as the square root of the sum of the squares of the X and Y direction Airy disk radii.

This feature may produce meaningless data for non-rotationally symmetric systems. Great care should be exercised in interpreting this plot for non-rotationally symmetric systems.

Y-Ybar Diagram

Purpose:

Y-Ybar diagram.

Settings:

Item	Description
First Surface	The first surface for which data will be plotted.
Last Surface	The last surface for which data will be plotted.
Wavelength	The wavelength number to be used in the calculation.
Plot Scale	Sets the maximum scale for the plot. The plot is always shown in a square box; the default scale is the maximum transverse ray coordinate. Enter zero for automatic.

Discussion:

The Y-Ybar diagram is a plot of marginal ray height as a function of chief ray height for a paraxial skew ray at every surface in the lens.

Chromatic Focal Shift

Purpose:

Chromatic focal shift plot.

Settings:

Item	Description
Maximum Shift	The maximum extent in lens units for the horizontal axis. The vertical axis scale is set by the range of wavelengths defined. Enter zero for automatic.
Pupil Zone	The radial zone in the pupil used to compute the back focus. The default value of zero means a paraxial ray will be used. Values between 0 and 1 mean real marginal rays are used in the appropriate zone in the entrance pupil. A 1 is at the edge of the pupil, or full aperture.

Discussion:

This is a plot of the shift in back focal length with respect to the primary wavelength. At each plotted wavelength the shift in the image plane required to reach focus for that color marginal ray is computed. This plot may not be meaningful for non-paraxial systems.

The maximum shift setting overrides the default scaling. Units are in lens units. The entire plot is always referenced to the primary wavelength paraxial focus. The diffraction limited depth of focus listed is given by the formula $4\lambda F^2$, where F is the working F/# and λ is the primary wavelength.

Dispersion Diagram

Purpose:

Plots index of refraction as a function of wavelength for any material in the glass catalog.

Settings:

Item	Description
Min Wavelength	Defines the left X- axis of the dispersion plot.
Max Wavelength	Defines the right X- axis of the dispersion plot.
Minimum Index	Defines the bottom Y- axis of the dispersion plot. Enter zero for automatic scaling.
Maximum Index	Defines the top Y- axis of the dispersion plot. Enter zero for automatic scaling.
Glass	The name of the material.
Use Temp, Pres	If checked, then the change in index due to temperature and pressure effects will be considered.

Discussion:

This is useful for checking if constants of dispersion or other formula data were entered correctly.

Glass Map

Purpose:

Plots names of glasses on the glass map according to the index of refraction and Abbe V-number. The index and Abbe number are taken directly from the entries in the glass catalog, and are not computed based upon the wavelength data or dispersion coefficients. All currently loaded glass catalogs are searched for glasses within the boundary values specified in the following table.

Settings:

Item	Description
Min Abbe	Defines the left X- axis of the plot.
Max Abbe	Defines the right X- axis of the plot.
Min Index	Defines the bottom Y- axis of the plot. Enter zero for automatic scaling.
Max Index	Defines the top Y- axis of the plot. Enter zero for automatic scaling.

Discussion:

This is useful for locating a glass with particular refractive and dispersive properties. By convention, the glass map is shown with the Abbe number decreasing while going from left to right, which is why the default min and max Abbe numbers seem to be reversed.

Internal Transmittance vs. Wavelength

Purpose:

Plots the internal transmittance for any thickness as a function of wavelength for any material in the glass catalog.

Settings:

Item	Description
Min Wavelength	Defines the left X- axis of the plot.
Max Wavelength	Defines the right X- axis of the plot.
Minimum Transmission	Defines the bottom Y- axis of the plot.
Maximum Transmission	Defines the top Y- axis of the plot. Enter zero for automatic scaling.
Glass	The name of the material.
Thickness	The thickness of the glass in millimeters.

Discussion:

This is useful for checking the transmittance of a particular glass. See the Chapter "Polarization Analysis".

System Summary Graphic

Purpose:

Displays in a graphic window a summary of the system data, similar to the text based system data report.

Settings:

None.

Discussion:

This graphic is primarily used to display a summary of system data within the printed page of the report graphics 4 or 6 feature; described in the Reports Menu chapter.

Calculations

Ray Trace

Purpose:

Paraxial and real trace of a single ray.

Settings:

Item	Description
Hx	Normalized x-field coordinate. The value should be between -1 and 1.
Hy	Normalized y-field coordinate. The value should be between -1 and 1.
Field	Select either a specific field or "Arbitrary" to enter in Hx and Hy. If a specific field is selected; the Hx and Hy controls are grayed out.
Wavelength	The wavelength number of the ray to trace.
Px	Normalized x-pupil coordinate. The value should be between -1 and 1.
Py	Normalized y-pupil coordinate. The value should be between -1 and 1.
Global Coordinates	If checked, all ray trace data is given in global coordinates rather than local coordinates, except tangent angles.
Type	Select "Dir Cosines" to display the ray direction cosines at each surface, "Tangent Ang" to display the tangent of the angle the ray makes at each surface, or "Ym,Um,Yc,Uc" to display the paraxial marginal and chief ray intercept/tangent angle display. The tangent angle is the ratio of the x (or y) direction cosine to the z direction cosine.

Discussion:

The Hx, Hy, Px, Py, and Global Coordinates settings are ignored if the "Ym, Um, Yc, Uc" option is selected. For a description of normalized coordinates, see the chapter "Conventions and Definitions". For the other options, this feature allows the user to specify the normalized object coordinates, the normalized pupil coordinates, and wavelength number and then view the real and paraxial ray coordinates at every surface. The first set of data presented is for a real ray. The values presented are the coordinates (in the surface local or the global coordinate system) of the ray intercept point. The direction cosines (or tangent angles) are the values of the ray in the surface after refraction. The direction cosine value is the cosine of the angle the ray makes with respect to the specified axis (the x-direction cosine is the cosine of the angle the ray makes with respect to the x-axis, etc.) The second set of data is just like the first set, except it is for a paraxial ray. Tangent angles are always computed relative to the local Z axis; regardless of the Global Coordinates setting.

Gaussian Beam

Purpose:

Computes Gaussian beam parameters.

Settings:

Item	Description
Wavelength	The wavelength number to use for the calculation.
M ² Factor	The M ² quality factor used to simulate mixed mode beams. See the Discussion.
Waist Size	The radial size of the embedded (perfect TEM00 mode) beam waist in object space in lens units.
Surf 1 to Waist	The distance from surface 1 (NOT the object surface) to the beam waist location. This parameter will generally be negative if the waist lies to the left of surface 1.
Update	See below.
Orient	
Surface	

Discussion:

This feature computes ideal and mixed mode M² Gaussian beam data, such as beam size, beam divergence, and waist locations, as a given input beam propagates through the lens system. This discussion is not meant to be a complete tutorial on laser beam propagation theory. For more information on Gaussian beam propagation, see one of the following references: "Lasers", A. E. Siegman, University Science Books (1986), "Gaussian beam ray-equivalent modeling and optical design", R. Herloski, S. Marshall, and R. Antos, Applied Optics Vol. 22, No. 8 pp. 1168 (1983), "Beam characterization and measurement of propagation attributes", M. W. Sasnett and T. F. Johnston, Jr., Proc. SPIE Vol. 1414, pp 21 (1991), and "New developments in laser resonators", A. E. Siegman, Proc. SPIE Vol. 1224, pp 2 (1990).

A Gaussian laser beam is described by a beam waist size, a wavelength, and a location in object space. The Gaussian beam is an idealization that can be approached but never attained in practice. However, real laser beams can be well described by an "embedded" Gaussian beam with ideal characteristics, and a quality factor called M², which defines the relative beam size and divergence with respect to the embedded Gaussian mode. The ideal M² value is unity, but real lasers will always have an M² value greater than unity.

This feature requires the definition of the initial input embedded beam properties, and the M² value. The input embedded beam is defined by the location of the input beam waist relative to surface 1 (note this is not the object surface, but the first surface after the object) and the waist radial size at this location. ZEMAX then propagates this embedded beam through the lens system, and at each surface the beam data is computed and displayed in the output window. ZEMAX computes the Gaussian beam parameters for both X and Y orientations.

Default beam parameters

ZEMAX defaults to a waist size of 0.05 lens units (no matter what the lens units are) and a surface 1 to waist distance of zero; this of course means the waist is at surface 1. The default values may be reset by clicking on

the "Reset" button. After the default values are computed and displayed, any alternate beam waist size and location may be entered and that Gaussian beam will be traced instead.

Propagating the embedded beam

Once the initial beam waist and location parameters are established, ZEMAX traces the embedded beam through the system and computes the radial beam size, the narrowest radial waist, the surface coordinate relative to the beam waist, the phase radius of curvature of the beam, the semi divergence angle, and the Rayleigh range for every surface in the system. ZEMAX calls these parameters the Size, Waist, Waist Z, Radius, Divergence, and Rayleigh, respectively, on the text listing that is generated.

The dimensions for all parameters are lens units except for the semi divergence angle, which is in units of radians. The embedded Gaussian beam parameters can be computed using the following standard formulas. The Rayleigh range, z_r , is defined as

$$z_r = \frac{\pi \omega_0^2}{\lambda}$$

The phase radius of curvature is given by

$$R(z) = z + \frac{z_r^2}{z}$$

where z is the distance from the beam waist. The radial beam size at any z is computed from

$$\omega(z) = \omega_0 \left[1 + \left(\frac{z}{z_r} \right)^2 \right]^{1/2}$$

where ω_0 is the beam waist. The divergence angle of the beam is given by

$$\theta = \tan^{-1} \frac{\omega_0}{z_r}$$

The quality factor

All of the preceding results are correct for the ideal embedded Gaussian beam. For aberrated, mixed-mode beams, a simple extension to the fundamental Gaussian beam model has been developed by Siegman. The method uses a term called the beam quality factor, usually denoted by M^2 . The factor M^2 can be thought of as "times diffraction limited" number, and is always greater than unity. The M^2 factor determines the size of the real, aberrated Gaussian beam by scaling the size and divergence of the embedded Gaussian mode by M . It is common practice to specify M^2 for a laser beam, rather than M , although the factor M is used to scale the beam size. The M^2 factor must be measured to be determined correctly. If the M^2 factor is set to unity, the default value, ZEMAX computes the TEM00 data described above. If M^2 is greater than unity, then ZEMAX computes both the embedded Gaussian beam parameters as well as the scaled data.

Because the embedded Gaussian beam parameters are based upon paraxial ray data the results cannot be trusted for systems which have large aberrations, or those poorly described by paraxial optics, such as non-rotationally symmetric systems. This feature ignores all apertures, and assumes the Gaussian beam propagates well within the apertures of all the lenses in the system.

Interactive analysis

The Settings dialog box for this feature also supports an interactive mode. After defining the various input beam parameters, clicking on "Update" will immediately trace the specified Gaussian beam, and display the usual results in the dialog box. The parameters for any surface may be viewed, and the surface number selected from the drop down list. The orientation may also be selected using the provided control.

The interactive feature does not in any way modify the lens or the system data; it is a handy "calculator" for displaying Gaussian beam data.

Seidel Coefficients

Purpose:

Displays Seidel (unconverted, transverse, and longitudinal), and wavefront aberration coefficients.

Settings:

Item	Description
Wavelength	The wavelength number to use for the calculation.

Discussion:

ZEMAX will compute the unconverted Seidel, transverse, longitudinal, and some wavefront coefficients. The Seidel coefficients are listed surface by surface, as well as a sum for the entire system. The coefficients listed are for spherical aberration (SPHA, S1), coma (COMA, S2), astigmatism (ASTI, S3), field curvature (FCUR, S4), distortion (DIST, S5), longitudinal color (CLA, CL), and transverse color (CTR, CT). The units are always the same as the system lens units, except of course for the coefficients measured in waves.

These calculations are only valid and accurate for systems consisting entirely of standard surfaces. Any systems which contain coordinate breaks, gratings, paraxial, or other non-standard surfaces are not adequately described by the paraxial rays which are used to compute the coefficients.

Transverse aberration coefficients are also listed by surface with a system sum. The coefficients given are transverse spherical (TSPH), transverse sagittal coma (TSCO), transverse tangential coma (TTCO), transverse astigmatism (TAST), transverse Petzval field curvature (TPFC), transverse sagittal field curvature (TSFC), transverse tangential field curvature (TTFC), transverse distortion (TDIS), transverse axial color (TAXC), and transverse lateral color (TLAC). The transverse aberrations are in the system lens units. The transverse aberration coefficients may be very large in optical spaces where the light is nearly collimated, and have little meaning in these optical spaces.

Longitudinal aberration coefficients are computed for longitudinal spherical aberration (LSPH), longitudinal sagittal coma (LSCO), longitudinal tangential coma (LTCO), longitudinal astigmatism (LAST), longitudinal Petzval field curvature (LPFC), longitudinal sagittal field curvature (LSFC), longitudinal tangential field curvature (LTFC), longitudinal distortion (LDIS), longitudinal axial color (LAXC), and longitudinal lateral color (LLAC). The longitudinal aberrations are in the system lens units. The longitudinal aberration coefficients may be very large in optical spaces where the light is nearly collimated, and have little meaning in these optical spaces.

The wavefront coefficients given are spherical aberration (W040), coma (W131), astigmatism (W222), field curvature Petzval (W220P), distortion (W311), axial color defocus term (W020), lateral color tilt term (W111), field curvature sagittal (W220S), field curvature medial (W220M), and field curvature tangential (W220T). All the wavefront coefficients are in units of wavelengths at the edge of the exit pupil. The various aberration coefficients are interrelated according to the following table. The symbols n and u refer to the index of refraction and the paraxial marginal ray angle in the object space side of each surface. The primes above the n and u symbols indicate that these are the values on the image space side of the surface. For a discussion of the meaning and derivation of the Seidel aberration coefficients, see Welford, *Aberrations of Optical Systems*, Smith, *Modern Lens Design*, or O' Shea, *Elements of Modern Optical Design*. Complete references can be found in the chapter "Introduction".

INTERRELATIONSHIP OF ABERRATION COEFFICIENTS

Name	Seidel	Wave	Description	Transverse	Longitudinal
Spherical	S_1	$\frac{S_1}{8}$	Spherical	$-\frac{S_1}{2n'u'}$	$\frac{S_1}{2n'u'^2}$

Name	Seidel	Wave	Description	Transverse	Longitudinal
Coma	S_2	$\frac{S_2}{2}$	Sagittal	$-\frac{S_2}{2n'u'}$	$\frac{S_2}{2n'u'^2}$
			Tangential	$-\frac{3S_2}{2n'u'}$	$\frac{3S_2}{2n'u'^2}$
Astigmatism	S_3	$\frac{S_3}{2}$	From tangential to sagittal foci	$-\frac{S_3}{n'u'}$	$\frac{S_3}{n'u'^2}$
Field Curvature	S_4	$\frac{S_4}{4}$	Gaussian to Petzval	$-\frac{S_4}{2n'u'}$	$\frac{S_4}{2n'u'^2}$
	$S_3 + S_4$	$\frac{S_3 + S_4}{4}$	Gaussian to sagittal	$-\frac{(S_3 + S_4)}{2n'u'}$	$\frac{S_3 + S_4}{2n'u'^2}$
	$2S_3 + S_4$	$\frac{2S_3 + S_4}{4}$	Gaussian to medial	$-\frac{(2S_3 + S_4)}{2n'u'}$	$\frac{2S_3 + S_4}{2n'u'^2}$
	$3S_3 + S_4$	$\frac{3S_3 + S_4}{4}$	Gaussian to tangential	$-\frac{(3S_3 + S_4)}{2n'u'}$	$\frac{3S_3 + S_4}{2n'u'^2}$
Distortion	S_5	$\frac{S_5}{2}$	Distortion	$-\frac{S_5}{2n'u'}$	$\frac{S_5}{2n'u'^2}$
Axial Color	C_L	$\frac{C_L}{2}$	Chromatic aberrations are measured between the extreme defined wavelengths, referenced to the selected wavelength.	$\frac{C_L}{n'u'}$	$\frac{C_L}{n'u'^2}$
Lateral Color	C_T	C_T		$\frac{C_T}{n'u'}$	$\frac{C_T}{n'u'^2}$

The Petzval radius of curvature shown on the Seidel screen is in system lens units, as is the Lagrange invariant.

Zernike Fringe Coefficients

Purpose:

Calculates the Zernike coefficients using the Fringe or University of Arizona polynomial set.

Settings:

Item	Description
Sampling	Specify the density in the pupil to use for coefficient fitting. Larger grid sizes are more accurate, although the computation time increases.
Max Term	Specify the maximum Zernike coefficient to compute. Any value up to 37 may be specified.

Item	Description
Wavelength	The wavelength number to use for the calculation.
Field	The field number to use for the calculation.

Discussion:

The Zernike screen displays the individual coefficients as well as the peak-to-valley, RMS, variance, Strehl ratio, residual RMS fit error, and maximum fit error.

The RMS of the wavefront error, σ , when referenced to the mean, is defined as

$$\sigma^2 = \overline{W^2} - \overline{W}^2$$

where W is the wavefront error, $\overline{W^2}$ is the mean square error, and \overline{W} is the average wavefront error. The RMS can actually be computed several different ways. If the mean wavefront term is ignored, then the RMS "referenced to zero" results. This computation yields the square root of $\overline{W^2}$ directly, and is rarely used.

If the mean wavefront is subtracted from all the wavefront phase values (the absolute phase reference has no physical meaning), then the RMS is "referenced to the mean".

Typically, the RMS is further referenced to the tilted and shifted reference sphere which minimizes the RMS. This is equivalent to subtracting out not only the mean (which is piston) but the average tilt in x and y as well. This is justified because tilt shifts the location of the diffraction image centroid, but otherwise has no effect on image quality. For brevity, ZEMAX calls this reference point the "centroid", although it is a reference point which is usually close to but not exactly at the diffraction image centroid. Most of the time, the RMS is taken to mean the RMS referenced to the centroid, which is always the lowest of the three numbers.

The Strehl ratio is computed using the RMS referenced to the centroid by the following approximation:

$$S = e^{-(2\pi\sigma)^2}$$

The approximation is only valid for Strehl ratios higher than about 0.10.

This feature computes a maximum of 37 Zernike terms. The particular Zernike terms used are not orthonormal but are instead all normalized to have unity magnitude at the edge of the pupil. Some of the higher order terms in the expansion were dropped to keep the total number of terms small, and the terms remaining were selected to favor accurate fitting of higher order spherical aberration. This particular set of Zernike polynomials is sometimes called the "Fringe" or "University of Arizona" notation. The more formal, and more general polynomial set is the Standard notation, sometimes called the "Born & Wolf" or the similar "Noll" notation, which is described under the "Zernike Standard Coefficients" feature.

The Zernike Fringe polynomials are defined in the following table.

ZERNIKE FRINGE POLYNOMIALS

Term	Z(ρ, φ)
1	1
2	$\rho \cos \varphi$
3	$\rho \sin \varphi$
4	$2\rho^2 - 1$
5	$\rho^2 \cos 2\varphi$

DEFOCUS

ASTIGMATISM

Term	$Z(\rho, \varphi)$	
6	$\rho^2 \sin 2\varphi$	
7	$(3\rho^2 - 2)\rho \cos \varphi$	COMA (PRIMARY)
8	$(3\rho^2 - 2)\rho \sin \varphi$	
9	$6\rho^4 - 6\rho^2 + 1$	PRIMARY SPHERICAL
10	$\rho^3 \cos 3\varphi$	
11	$\rho^3 \sin 3\varphi$	TRIANGULAR COMA AST.
12	$(4\rho^2 - 3)\rho^2 \cos 2\varphi$	
13	$(4\rho^2 - 3)\rho^2 \sin 2\varphi$	SECONDARY ASTIG.
14	$(10\rho^4 - 12\rho^2 + 3)\rho \cos \varphi$	
15	$(10\rho^4 - 12\rho^2 + 3)\rho \sin \varphi$	
16	$20\rho^6 - 30\rho^4 + 12\rho^2 - 1$	
17	$\rho^4 \cos 4\varphi$	SQUARE ASTIG
18	$\rho^4 \sin 4\varphi$	
19	$(5\rho^2 - 4)\rho^3 \cos 3\varphi$	TRIANGULAR COMA (7)
20	$(5\rho^2 - 4)\rho^3 \sin 3\varphi$	
21	$(15\rho^4 - 20\rho^2 + 6)\rho^2 \cos 2\varphi$	
22	$(15\rho^4 - 20\rho^2 + 6)\rho^2 \sin 2\varphi$	
23	$(35\rho^6 - 60\rho^4 + 30\rho^2 - 4)\rho \cos \varphi$	
24	$(35\rho^6 - 60\rho^4 + 30\rho^2 - 4)\rho \sin \varphi$	
25	$70\rho^8 - 140\rho^6 + 90\rho^4 - 20\rho^2 + 1$	
26	$\rho^5 \cos 5\varphi$	
27	$\rho^5 \sin 5\varphi$	
28	$(6\rho^2 - 5)\rho^4 \cos 4\varphi$	

Term	$Z(\rho, \varphi)$
29	$(6\rho^2 - 5)\rho^4 \sin 4\varphi$
30	$(21\rho^4 - 30\rho^2 + 10)\rho^3 \cos 3\varphi$
31	$(21\rho^4 - 30\rho^2 + 10)\rho^3 \sin 3\varphi$
32	$(56\rho^6 - 105\rho^4 + 60\rho^2 - 10)\rho^2 \cos 2\varphi$
33	$(56\rho^6 - 105\rho^4 + 60\rho^2 - 10)\rho^2 \sin 2\varphi$
34	$(126\rho^8 - 280\rho^6 + 210\rho^4 - 60\rho^2 + 5)\rho \cos \varphi$
35	$(126\rho^8 - 280\rho^6 + 210\rho^4 - 60\rho^2 + 5)\rho \sin \varphi$
36	$252\rho^{10} - 630\rho^8 + 560\rho^6 - 210\rho^4 + 30\rho^2 - 1$
37	$924\rho^{12} - 2772\rho^{10} + 3150\rho^8 - 1680\rho^6 + 420\rho^4 - 42\rho^2 + 1$

Zernike Standard Coefficients

Purpose:

Calculates the orthonormal Zernike coefficients computed using the notation defined in "Zernike polynomials and atmospheric turbulence", R. Noll, J. Opt. Soc. Am., Vol. 66, No. 3, March 1976. This notation allows terms to be consistently defined up to (almost) arbitrary order, rather than the limited 37 term Fringe polynomials.

Settings:

Item	Description
Sampling	Specify the density in the pupil to use for coefficient fitting. Larger grid sizes are more accurate, although the computation time increases.
Max Term	Specify the maximum Zernike coefficient to compute. Any value up to 231 may be specified.
Wavelength	The wavelength number to use for the calculation.
Field	The field number to use for the calculation.

Discussion:

This feature is almost identical to the "Zernike Fringe Coefficients" feature, except a somewhat different numbering scheme is used, more terms in the expansion are available for fitting, the terms are orthogonal and normalized, and none of the terms are skipped. See that description for some important details. The most general way to express the Zernike polynomials is in the form

$$R_n^m(\rho)e^{im\theta} = \begin{cases} R_n^m(\rho)\cos m\theta \\ R_n^m(\rho)\sin m\theta \end{cases}$$

The radial portion of the polynomial is defined by two indices, n and m . The n index defines the order of the radial power; so an n value of 5 would indicate all polynomials whose maximum radial power was ρ^5 . Only certain values for m are allowed once n is chosen; $n + m$ must be even, and $0 \leq m \leq n$. For details on Zernike Standard polynomials, see "Principles of Optics", by Born & Wolf, referenced in Chapter 1, or the JOSA reference listed above. The terms are orthonormal such that the magnitude of the coefficient of each term is the RMS contribution of the term.


This feature lists the formulas next to the fitted coefficients; the entire 231 term table is too long to include here. The first 28 terms are given below.

ZERNIKE STANDARD POLYNOMIALS

Term	$Z(\rho, \varphi)$
1	1
2	$\sqrt{4}\rho \cos \varphi$
3	$\sqrt{4}\rho \sin \varphi$
4	$\sqrt{3}(2\rho^2 - 1)$
5	$\sqrt{6}(\rho^2 \sin 2\varphi)$
6	$\sqrt{6}(\rho^2 \cos 2\varphi)$
7	$\sqrt{8}(3\rho^3 - 2\rho) \sin \varphi$
8	$\sqrt{8}(3\rho^3 - 2\rho) \cos \varphi$
9	$\sqrt{8}\rho^3 \sin 3\varphi$
10	$\sqrt{8}\rho^3 \cos 3\varphi$
11	$\sqrt{5}(6\rho^4 - 6\rho^2 + 1)$
12	$\sqrt{10}(4\rho^4 - 3\rho^2) \cos 2\varphi$
13	$\sqrt{10}(4\rho^4 - 3\rho^2) \sin 2\varphi$
14	$\sqrt{10}\rho^4 \cos 4\varphi$
15	$\sqrt{10}\rho^4 \sin 4\varphi$
16	$\sqrt{12}(10\rho^5 - 12\rho^3 + 3\rho) \cos \varphi$
17	$\sqrt{12}(10\rho^5 - 12\rho^3 + 3\rho) \sin \varphi$
18	$\sqrt{12}(5\rho^5 - 4\rho^3) \cos 3\varphi$
19	$\sqrt{12}(5\rho^5 - 4\rho^3) \sin 3\varphi$

Term	Z(ρ , φ)
20	$\sqrt{12}\rho^5 \cos 5\varphi$
21	$\sqrt{12}\rho^5 \sin 5\varphi$
22	$\sqrt{7}(20\rho^6 - 30\rho^4 + 12\rho^2 - 1)$
23	$\sqrt{14}(15\rho^6 - 20\rho^4 + 6\rho^2) \sin 2\varphi$
24	$\sqrt{14}(15\rho^6 - 20\rho^4 + 6\rho^2) \cos 2\varphi$
25	$\sqrt{14}(6\rho^6 - 5\rho^4) \sin 4\varphi$
26	$\sqrt{14}(6\rho^6 - 5\rho^4) \cos 4\varphi$
27	$\sqrt{14}\rho^6 \sin 6\varphi$
28	$\sqrt{14}\rho^6 \cos 6\varphi$

Fiber Coupling Efficiency

 This feature is only available in the XE and EE editions of ZEMAX.

Purpose:

This feature computes the coupling efficiency for single mode fiber coupling systems. For multi-mode fiber coupling, see "Calculating efficiency of multi-mode fibers" on page 107.

Settings:

Item	Description
Source Fiber	
NA x/y	Numerical Aperture of the source fiber in object space in the xz and yz planes respectively. This is n times the sin of the half angle to the 1 over e squared intensity point.
X Angle, Y Angle	The angular rotation of the source fiber in object space, in degrees, measured from the nominal orientation of the source fiber. The X angle is the angle measured in the XZ plane; the Y angle is measured in the YZ plane.
Field	The field number to use for the calculation.
Wavelength	The wavelength number to use for the calculation.
Sampling	The grid size to use for the numerical integration.
Align Source to Chief Ray	If unchecked, the fiber is oriented along the object z axis, centered on the field point. If checked, the center of the fiber points along the chief ray for that field point.
Ignore Source Fiber	If checked, then illumination of the pupil as specified by the System, General, Apodization setting will be used. All computations are then referenced to the total energy incident upon the entrance pupil.
Receiving Fiber	

Item	Description
NA	Numerical Aperture of the receiving fiber in image space. This is n times the sin of the half angle to the 1 over e squared intensity point.
Tilt About X, Y	The angular rotation of the receiving fiber in image space, in degrees, measured from the nominal orientation of the receiving fiber. If both X and Y rotations are specified, the rotation is first around X, then around Y. Note the rotations are about the specified axes; so Tilt About X rotates the fiber toward or away from the Y axis. Rotations are done <u>after</u> the XYZ decenters below.
Decenter XYZ	The linear position shift in the x/y/z direction in lens units of the receiving fiber, measured from the nominal position of the receiving fiber. Shifts are done <u>before</u> the XY tilts above.
Align Receiver to Chief Ray	If unchecked, the fiber is positioned at XY coordinates (0.0, 0.0) on the image surface. If checked, the fiber is repositioned in x and y to be centered at the point the chief ray intercepts the image surface. In either case the receiving fiber is oriented along the image space z axis, unless rotated using the controls above.
Use Polarization	If checked, polarized rays will be traced for each ray required, and the resulting transmitted intensity through the system will be accounted for. See the "System Menu" chapter under "Polarization" for information on defining the polarization state and other details. Only ZEMAX-EE supports this capability.

Discussion:

This feature computes fiber coupling for single-mode fibers with a Gaussian shaped mode. For multi-mode fiber coupling, see "Calculating efficiency of multi-mode fibers" on page 107.

Fiber coupling efficiency is computed based upon a two fiber or a one fiber model. In the two fiber model, light emerges from a source fiber to fill (or partially overfill) the entrance pupil of an optical system. Energy not collected by the entrance pupil is lost, reducing the overall efficiency. If desired, the source fiber may be ignored for a one fiber model; and in this case the efficiency is computed relative to the energy entering the entrance pupil; which in turn is a function of the system apodization (see "Apodization Type" on page 57).

The system efficiency (S) is the sum of the energy collected by the entrance pupil which passes through the optical system, accounting for both the vignetting and transmission of the optics, divided by the sum of all the energy which radiates from the source fiber:

$$S = \left[\frac{\iint t(x, y) F_s(x, y) dx dy}{\iint F_s(x, y) dx dy} \right]^2$$

where F_s is the source fiber amplitude function and the integral in the numerator is only done over the entrance pupil of the optical system, and $t(x, y)$ is the amplitude transmission function of the optics. The transmission is affected by bulk absorption and optical coatings (if use polarization is checked on). If the source fiber is ignored, then the integral in the denominator is only done over the entrance pupil, and the F_s function is determined by the system apodization, if any.

Aberrations in the optical system introduce phase errors which will affect the coupling into the fiber. Maximum coupling efficiency is achieved when the mode of the wavefront converging towards the receiving fiber perfectly matches the mode of the fiber in both amplitude and phase at all points in the wavefront. This is defined mathematically as a normalized overlap integral between the fiber and wavefront amplitude:

$$T = \frac{\left| \iint F_r(x, y) W'(x, y) dx dy \right|^2}{\iint F_r(x, y) F_r'(x, y) dx dy \iint W(x, y) W'(x, y) dx dy}$$

where $F_r(x, y)$ is the function describing the receiving fiber amplitude, $W(x, y)$ is the function describing the wavefront from the exit pupil of the optical system, and the $*$ symbol represents complex conjugate. Note that these functions are complex valued, so this is a coherent overlap integral.

Large phase errors in the wavefront function are difficult to numerically integrate accurately. Fortunately, rays with huge phase aberrations are not going to contribute to the fiber coupling anyway, as these rays move through 2π phase cycles so quickly that they tend to all average to zero. To improve the accuracy of the fiber coupling computation for highly aberrated systems, the integration algorithm will ignore rays with more than $n/8$ waves of OPD error, where n is the number of rays across the pupil, which is determined by the sampling setting.

T has a maximum possible value of 1.0, and will decrease if there is any mismatch between the fiber amplitude and phase and the wavefront amplitude and phase.

ZEMAX computes the values S and T . The total power coupling efficiency is the product of these numbers. A theoretical maximum coupling efficiency is also computed; this value is based upon ignoring the aberrations but accounting for all vignetting, transmission, and other amplitude mismatches between the modes.

To model the reflection loss from the front face of the receiving fiber, set the image space glass to the material the fiber is made of, and check the box "use polarization". The reflectivity of the receiving fiber (and other surfaces) will be accounted for.

YNI Contributions

Purpose:

This feature lists for each surface the paraxial YNI value, which is proportional to the Narcissus contribution of that surface.

Settings:

Item	Description
Wavelength	The wavelength number to use for the calculation.

Discussion:

See "Narcissus: reflections on retroreflections in thermal imaging systems", Applied Optics, Vol. 21, #18, p3393 (1982) for a discussion.

Sag Table

Purpose:

This feature lists for the selected surface the surface sag (z-coordinate) at various distances from the vertex. The best fit spherical radius is computed, and the sag of the best fit sphere and the difference is also tabulated. Only the Y-coordinate of the surface is considered, therefore the data may not be useful for non-rotationally symmetric surfaces.

Settings:

Item	Description
Surface	The surface number to use for the calculation.
Step Size	The distance between steps measured from the vertex at which the sag is computed. The default value of zero will automatically choose a reasonable step size.
Negative Remove	See "Discussion".

Discussion:

This feature can be used to determine the maximum aspheric deviation of a surface. The sag table is also useful for lens fabrication.

This feature lists out five columns:

Y-Coord: The y coordinate of the point on the surface being computed.
 Sag: The sag of the surface at the y coordinate.

BFS Sag: The sag of the best fit sphere. The best fit sphere is determined by finding the radius of the sphere that minimizes the RMS deviation between the surface sag and the sag of the best fit sphere. ZEMAX computes this best fit sphere radius using an iterative bisection technique.

Deviation: The difference between the BFS Sag and the Sag, i.e. (BFS Sag) - (Sag).

Remove: The amount of material to remove assuming that the surface was first generated to the best fit sphere. By default, this number is given to be positive. Checking the "Negative Remove" box will yield all negative values for the remove column. Note that the remove column values are identical to the deviation values, but a constant offset value has been added. The remove values can be selected to be either all positive or all negative. Which is correct depends upon whether the surface is convex or concave. ZEMAX cannot always make this determination correctly, and so it is left as a user selectable option.

It is always advisable to check the sign convention carefully before any fabrication decisions are made.

Cardinal Points

Purpose:

This feature lists for the selected range of surfaces and wavelengths the locations of the principal, nodal, anti-nodal, and focal planes. The calculation is done for any defined wavelength and either the X-Z or Y-Z orientation.

Settings:

Item	Description
First Surface	The starting surface number of the group to compute the cardinal points for.
Last Surface	The ending surface number of the group to compute the cardinal points for.
Wavelength	The wavelength number to use for the computation.
Orientation	The orientation to use for computing the cardinal plane locations.

Discussion:

This feature may not return reliable results if coordinate breaks or non-centered optics are included within the specified surface range.

Gradient Index

Gradium™ Profile

Purpose:

Plots the axial index profile of a Gradium surface.


Settings:

Item	Description
Profile	The name of the profile. See the chapter on surface types for details.
Surface	The surface number to plot.
Wavelength	The wavelength number to use for the calculation.

Discussion:

If the surface is selected to "None", then the reference wavelength for the profile glass family is used, no matter what wavelength is selected. If a Gradium surface is selected, then any defined wavelength, or the reference wavelength may be selected. Also, if a surface number is selected, then the starting and ending points for the glass blank will be indicated using an "X" on the plot. The starting and ending positions include consideration of the sag of the surface at the defined semi-diameter.

Universal Plot

 This feature is only available in the EE edition of ZEMAX.

New Universal Plot...

Purpose:

Displays as either a plot or as a text listing the value of any optimization operand as a function of some other parameter.

Settings:

Item	Description
Independent variable settings	
	Select either surface, system, or multi-configuration data as the independent variable.
	Select the parameter to use as the independent variable. The available selections include radius, curvature, thickness, conic, parameter, and extra data values if surface parameters are chosen. For system parameters, the selections include the system aperture, field and wavelength data, apodization factor, temperature, and pressure. For configuration data, all multi-configuration operands are listed.
	The surface number for the independent variable parameter if the parameter is associated with a surface. This field is used for the configuration number if configuration data is selected.
Start/Stop Value	The beginning and ending range of the independent variable.
# of Steps	The number of values between the start and stop values, inclusive, to compute the dependent variable function value.
Dependent variable settings	
Operand	The optimization operand to use as the dependent variable function. The whole system merit function is also included at the top of the list. If the system merit function is selected, then either the entire merit function or any specific operand's value may be selected using the "Line" control below.
Line	The optimization operand # to use as the dependent variable function. The whole system merit function is also included at the top of the list.
Int1/Int2	The Int1 and Int2 values for the selected operand. The dialog box will display the short name for the Int1 and Int2 values to aid in identifying the data that needs to be supplied. If the edit control is grayed out; then the selected operand does not use that value.
Hx/Hy/Px/Py	The Hx, Hy, Px, Py values for the selected operand. The dialog box will display the short name for these values to aid in identifying the data that needs to be supplied. If the edit control is grayed out; then the selected operand does not use that value.
Min/Max Plot Value	The min and max value to plot for the dependent variable (Y) function. If both of these values are zero, then a default scale will be chosen. If either one is not zero, then the y range of the plot will extend from the minimum to the maximum. If the data does not fit within the specified range, then the data may plot beyond the borders of the data box.
Plot Title	Any text may be entered which will be used as the title for the plot or text listing.

Item	Description
Save As	This will save the current settings in a file specified by the name in the edit control to the right of this button. Any name which may be used as a valid file name may be entered. ZEMAX will append the appropriate extension to the given name and use this as the file name for storing the settings. The settings from any previously saved settings file may be loaded using the "Load From" button described below. Once the settings are saved, the name will appear as a menu option for quick regeneration of the plot.
Load From	Load settings previously saved with the "Save As" button described above.

Discussion:

This feature will either plot as a graph or create a text listing of the value of any optimization operand as a function of any number of system, surface, or multi-configuration data.

For example, suppose a plot of the sagittal MTF at 30 lp/mm was needed as a function of the decenter of a lens group (which may be a useful diagnostic for tolerancing analysis). Since the MTFS operand computes the sagittal MTF, the Universal Plot can generate such a graphic or text listing. See the Chapter "Optimization" for a list of available optimization operands. A decenter of a lens group is defined by Parameter 1 or 2 on the relevant coordinate break surface, and Parameter 1 and 2 are both listed in the available surface group parameters.

Because of the number of different plots this feature can generate, there are no intelligent defaults for either the independent or dependent settings. These values need to be carefully provided in the settings dialog box. If the optimization operand cannot be computed, an error message will be displayed and the plot will not be generated.

Because many optimization operands accept Hx and Hy values to define the field point for the calculation, these operands may require that the number of fields be set to 1, then set Hx = 0 and Hy = 1, then finally choose Y Field 1 as the independent variable for a plot of the operand as a function of field. A similar trick works for getting plots as a function of wavelength.

Some individual operands, like DIFF and SUMM, have no meaning if selected from the "Operand" column because these operands are only defined when used as part of a larger merit function. If results from a calculation like DIFF are required, then the merit function should be selected from the Operand list and the individual DIFF's within the merit function selected for the "Line".

The escape key will terminate the analysis if the computation is taking too long.

Polarization



These features are only available in the EE editions of ZEMAX.

For more information on polarization calculations, see "POLARIZATION ANALYSIS" on page 369.

Polarization Ray Trace

Purpose:

The polarization ray trace feature generates a text window which displays all the polarization data for a single ray.

Settings:

Item	Description
Ex	Electric field amplitude in X direction.
Ey	Electric field amplitude in Y direction.
X-Phase	Phase of the X component of the electric field in degrees.
Y-Phase	Phase of the Y component of the electric field in degrees.
Hx	Normalized x-field coordinate. The value should be between -1 and 1.
Hy	Normalized y-field coordinate. The value should be between -1 and 1.

Item	Description
Px	Normalized x-pupil coordinate. The value should be between -1 and 1.
Py	Normalized y-pupil coordinate. The value should be between -1 and 1.
Wavelength	The wavelength number of the ray to trace.

Discussion:

This feature tabulates all of the data computed by ZEMAX to perform polarization ray tracing. See the chapter "Polarization Analysis".

Polarization Pupil Map

Purpose:

Generates a graph of the polarization ellipse as a function of pupil position. This aids in visualizing the change in polarization over the pupil.

Settings:

Item	Description
Ex	Electric field amplitude in X direction.
Ey	Electric field amplitude in Y direction.
X-Phase	Phase of the X component of the electric field in degrees.
Y-Phase	Phase of the Y component of the electric field in degrees.
Wavelength	The wavelength number of the rays to trace.
Field	The field position number of the rays to trace.

Discussion:

The polarization ellipse is a representation of the figure traced out by the electric field vector as the wave propagates during one cycle. The magnitude of the ellipse is determined by the transmission of the ray, which is generally a function of pupil position. See the chapter "Polarization Analysis".

Transmission

Purpose:

Computes the integrated and surface by surface transmission through the optical system considering polarization effects.

Settings:

Item	Description
Ex	Electric field amplitude in X direction.
Ey	Electric field amplitude in Y direction.
X-Phase	Phase of the X component of the electric field in degrees.
Y-Phase	Phase of the Y component of the electric field in degrees.
Unpolarized	If checked, the electric field definition is ignored and the unpolarized computation is performed.

Discussion:

This feature tabulates for each field position and wavelength the integrated transmission of the optical system for the specified polarization. The transmission is computed as a fraction of 100%, with 100% being that transmitted.

Phase Aberration
 Purpose:
 Computes the polariz
 Settings:

Item
Ex
Ey
X-Phase
Y-Phase
Wavelength
Field

Discussion:
 This feature comp
 and Y orientation
 media, and by refl
 Like the usual Of
 cases where the
 system, if the inci
 ray, and therefore
 the polarization,
 this discontinuity
 ray phase. The p
 cases, the phas
 zero.

Transmission
 Purpose:
 Generates a gr
 sagittal pupil fa
 Settings:

Item
Ex
Ey
X-Phase
Y-Phase
Wavelength

If there were no absorption, reflection, or vignetting losses. The transmission calculation accounts for vignetting factors, vignetting due to apertures or obscurations, ray clipping due to ray trace errors, surface Fresnel or coating losses, and bulk internal transmittance due to absorption. Also tabulated for each field and wavelength is the relative and total transmission of the chief ray. This allows identification of where significant surface losses originate. See also the relative illumination feature.

Phase Aberration

Purpose:

Computes the polarization induced phase aberration of an optical system.

Settings:

Item	Description
Ex	Electric field amplitude in X direction.
Ey	Electric field amplitude in Y direction.
X-Phase	Phase of the X component of the electric field in degrees.
Y-Phase	Phase of the Y component of the electric field in degrees.
Wavelength	The wavelength number of the rays to trace.
Field	The field position number of the rays to trace.

Discussion:

This feature computes for the specified field position and wavelength the phase aberration in the image space X and Y orientations. Polarization phase aberrations are induced by the effects of refraction through dielectric media, and by reflection from metallic or dielectric mirrors.

Like the usual OPD plots, the polarization phase aberration is referenced to the chief ray. However, there are cases where the chief ray phase cannot be determined in both orientations. For example, in an axial symmetric system, if the incident polarization is linear in the Y direction, there is zero intensity in the X direction for the chief ray, and therefore the X phase is indeterminate. For other rays in the pupil, there is generally a slight rotation of the polarization, and therefore the resulting electric field in the X direction yields a valid phase angle. To avoid this discontinuity in the phase, ZEMAX averages two rays on either side of the chief ray to interpolate the chief ray phase. The problem may still appear under certain circumstances, even with this averaging technique. In all cases, the phase data is still valid, because the phase aberration has no effect on image quality if the intensity is zero.

Transmission Fan

Purpose:

Generates a graph of the transmitted intensity for each field and wavelength as a function of either tangential or sagittal pupil fans.

Settings:


Item	Description
Ex	Electric field amplitude in X direction.
Ey	Electric field amplitude in Y direction.
X-Phase	Phase of the X component of the electric field in degrees.
Y-Phase	Phase of the Y component of the electric field in degrees.
Wavelength	The wavelength number of the rays to trace.

Item	Description
Field	The field position number of the rays to trace.
Unpolarized	If checked, the electric field definition is ignored and the unpolarized computation is performed.

Discussion:

The polarization transmission fan is useful for determining the transmission variation over the pupil as a function of field and wavelength. See the chapter "Polarization Analysis".

Coatings

 *These features are only available in the EE editions of ZEMAX.*

For more information on coating calculations, see the chapter "Polarization Analysis".

Reflection vs. Angle

Purpose:

Computes the S, P, and average polarization intensity coefficients for reflection for the specified surface as a function of incident angle.

Settings:

Item	Description
Min Angle	The minimum angle of incidence to plot. This defines the left edge of the plot.
Max Angle	The maximum angle of incidence to plot. This defines the right edge of the plot.
Min Y	The minimum y value to plot. This defines the bottom edge of the plot.
Max Y	The maximum y value to plot. This defines the top edge of the plot.
Surface	The surface number which defines the interface to use.
Wavelength	The wavelength number to use.

Discussion:

The incident angle is measured in the medium prior to the specified surface. See the chapter "Polarization Analysis" for details on the computation.

Transmission vs. Angle

Purpose:

Computes the S, P and average polarization intensity coefficients for transmission for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface. See the chapter "Polarization Analysis" for details on the computation.

Absorption vs. Angle

Purpose:

Computes the S, P, and average polarization intensity coefficients for absorption for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface. See the chapter "Polarization Analysis" for details on the computation.

Diattenuation vs. Angle

Purpose:

Computes the R (reflected) and T (transmitted) diattenuation for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface. See the chapter "Polarization Analysis" for details on the computation.

Phase vs. Angle

Purpose:

Computes the S and P polarization phase for reflection (if the surface is a mirror) or for transmission (if the surface is not a mirror) for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface. See the chapter "Polarization Analysis" for details on the computation.

Retardance vs. Angle

Purpose:

Computes the retardance for the specified surface as a function of incident angle.

Settings:

See "Reflectivity vs. Angle" above.

Discussion:

The incident angle is measured in the medium prior to the specified surface. See the chapter "Polarization Analysis" for details on the computation.

Reflection vs. Wavelength

Purpose:

Computes the S, P, and average polarization intensity coefficients for reflection for the specified surface as a function of incident wavelength.

Settings:

Item	Description
Min Wave	The minimum wavelength to plot. This defines the left edge of the plot.
Max Wave	The maximum wavelength to plot. This defines the right edge of the plot.
Min Y	The minimum y value to plot. This defines the bottom edge of the plot.
Max Y	The maximum y value to plot. This defines the top edge of the plot.
Surface	The surface number which defines the interface to use.
Angle	The angle of incidence, in degrees, to use.

Discussion:

See the chapter "Polarization Analysis".

Transmission vs. Wavelength

Purpose:

Computes the S, P, and average polarization intensity coefficients for transmission for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Discussion:

See the chapter "Polarization Analysis".

Absorption vs. Wavelength

Purpose:

Computes the S, P, and average polarization intensity coefficients for absorption for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Discussion:

See the chapter "Polarization Analysis".

Diattenuation vs. Wavelength

Purpose:

Computes the R (reflected) and T (transmitted) diattenuation for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Discussion:

See the chapter "Polarization Analysis".

Phase vs. Wavelength

Purpose:

Computes the S and P polarization phase for reflection (if the surface is a mirror) or for transmission (if the surface is not a mirror) for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Discussion:

See the chapter "Polarization Analysis".

Retardance vs. Wavelength

Purpose:

Computes the retardance for the specified surface as a function of incident wavelength.

Settings:

See "Reflectivity vs. Wavelength" above.

Discussion:

See the chapter "Polarization Analysis".

Introduction

The optimization feature provided by ZEMAX is quite powerful, and is capable of improving lens designs given a reasonable starting point and a set of variable parameters. Variables can be curvatures, thicknesses, glasses, conics, parameter data, extra data, and any of the numeric multi-configuration data. ZEMAX uses an actively damped least squares method. The algorithm is capable of optimizing a merit function composed of weighted target values; these target values are called "operands". ZEMAX has several different default merit functions, described in a subsequent section. These merit functions can be changed easily using the Merit Function Editor. For details on this procedure see the section "Modifying the merit function".

Optimization requires three steps: 1) a reasonable system which can be traced, 2) specification of the variables, and 3) specification of the merit function. A reasonable system is a rather loose concept which simply means that poorly conceived designs are not likely to be transformed into exceptional designs by the optimization algorithm (although there are exceptions). The variables, and there must be at least one for the optimization algorithm to be able to make any progress, are specified on the various editors, as described in the next section. To reach the optimization screen, select Tools, Optimization. You must specify all variables before using optimization.

The algorithm used by the optimization feature described in this chapter is designed to find the "local" minimum of the specified merit function. However, ZEMAX-XE and EE also have a capability to search for a "global" minimum of the merit function. The global minimum is the lowest possible value for the merit function, and if the optimization feature is not for novice users, and is not appropriate for interactive designing. For details see the chapter "Global Optimization".

Selecting variables

Variables for optimization are specified by pressing Ctrl-Z when the highlighted bar is on the parameter to be varied in the Lens Data Editor. Note that Ctrl-Z is a toggle. The Multi-Configuration and Extra Data Editors also contain numeric data that may be made variable by using Ctrl-Z. Glasses cannot be made variable directly because they are discrete. To optimize glasses, see "Optimizing glass selection" later in this chapter.

Defining the default merit function

The merit function is a numerical representation of how closely an optical system meets a specified set of goals. ZEMAX uses a list of operands which individually represent different constraints or goals for the system. Operands represent goals such as image quality, focal length, magnification, and many others.

The merit function is proportional to the square root of the weighted sum of the squares of the difference between the actual and target value of each operand in the list. The merit function is defined this way so a value of zero is ideal. The optimization algorithm will attempt to make the value of this function as small as possible, and so the merit function should be a representation of what you want the system to achieve. You do not have to use the default merit function, you may construct your own as described in a later section.

The easiest way to define a merit function is to select the Tools, Default Merit Function option on the Merit Function Editor menu bar. A dialog box will appear which allows selection of options for the default merit function. Each option is explained in the following paragraphs.

Selecting the type of optimization

Several different types of merit functions are available. The default merit function is constructed using four key choices: The optimization type, data type, reference point, and integration method. The choices are described in the following tables.

DEFAULT OPTIMIZATION TYPES

Name	Description
RMS	RMS is an abbreviation for Root-Mean-Square. This type is by far the most commonly used. The RMS is the square root of the average value of the squares of all the individual errors.
PTV	PTV is an abbreviation for Peak-To-Valley. There are rare cases where the RMS is not as important as the maximum extent of the aberrations. For example, if all the rays need to land within a circular region on a detector or fiber. In these cases, the Peak-To-Valley (PTV) may be a better indicator of performance. This merit function type attempts to minimize the PTV extent of the errors.

DEFAULT OPTIMIZATION DATA

Name	Description
Wavefront	Wavefront is the aberration measured in waves.
Spot Radius	The radial extent of the transverse ray aberrations in the image plane.
Spot X	The x extent of the transverse ray aberrations in the image plane.
Spot Y	The y extent of the transverse ray aberrations in the image plane.
Spot X and Y	Both the x and y extent of the transverse ray aberrations in the image plane. The x and y components are considered separately, and both are optimized together. This is similar to Spot Radius, except the signs of the aberrations are retained, which yields better derivatives. Note that in computing the radius of an aberration, the sign information is lost.

DEFAULT OPTIMIZATION REFERENCE POINTS

Name	Description
Centroid	The RMS or PTV computation of the data is referenced to the centroid of all the data coming from that field point. Centroid reference is generally preferred, especially for wavefront optimization. For wavefront optimization, reference to the centroid subtracts out piston, x-tilt, and y-tilt of the wavefront, none of which degrade image quality. Centroid reference also yields more meaningful results when coma is present, since coma shifts the image centroid away from the chief ray location. Historically, the chief ray has been used because it is simpler to compute, but this complexity is dealt with by ZEMAX with virtually no performance loss.
Chief	The RMS or PTV computation of the "data" is referenced to the chief ray at the primary wavelength.
Mean	This option is only available if the selected optimization data is wavefront. The mean reference is very similar to the centroid reference, except only the piston (mean wavefront) is subtracted out; not the x- or y-tilt. Since the exact point at which the OPD is defined to be zero is arbitrary; referencing to the mean is generally preferred to the chief ray reference for those cases where the centroid is not preferred.

Physically significant merit functions

Note that the numerical value of the merit function is physically significant. If the merit function is RMS-Wavefront-Centroid, then the numerical value of the merit function is the RMS wavefront error in waves. If the merit function is RMS-Spot Radius-Chief, then a value of 0.145 means the RMS spot size is 0.145 lens units. If the lens units were millimeters, this would correspond to 145 microns RMS.

Novice designers often ask why the RMS spot radius merit function should yield a different optimum design than the RMS wavefront merit function. The basic reason they are different is that ray aberrations are proportional to the derivative of the wave aberrations. Therefore, it is unreasonable to expect that the minimum of one corresponds to the minimum of the other. A general rule of thumb to use is that if the system is close to diffraction limited (say a PTV wavefront error of less than two waves) then use wavefront error. Otherwise, use spot radius.

Generally speaking, the merit functions which use the centroid as a reference are superior to those that use the chief ray. Most diffraction based performance measures, such as MTF or encircled energy, improve when the RMS wavefront error referenced to the centroid decreases. However, it is always best to reoptimize a final design with the various merit functions to verify which one provides the best performance for the system being designed. For example, the RMS wavefront centroid reference often yields better low frequency MTF response, but worse medial frequency response, than the RMS chief ray reference optimization does.

Selecting the pupil integration method

There are two different pupil integration methods used to construct the merit function: Gaussian quadrature (GQ) or rectangular array (RA). The GQ algorithm is VASTLY superior for almost all cases of practical interest. The GQ algorithm uses a carefully selected and weighted ray set to accurately compute the RMS or PTV error over the entrance pupil (strictly speaking, the PTV algorithm is not a GQ algorithm, but it is very similar). The weighting for all rays is applied according to the weights set on the wavelength and field data dialog boxes, any pupil apodization function, and by the GQ merit function algorithm. For RMS merit functions, the weighting and ray set selection used is based on a method described in a paper by G. W. Forbes, J. Opt. Soc. Am. A, Vol. 5, No. 11, November 1988, p1943. For the PTV merit functions, the ray set is based on solutions to the Chebyshev polynomials, described in Numerical Recipes, Cambridge University Press (1989). If you are interested in detailed information on the basis and accuracy of these methods, see these references. GQ is much, much more accurate than any other known method, and requires fewer rays. Therefore, you get the best of both worlds: greater speed and greater accuracy. The GQ algorithm requires specification of the number of "Rings" and the number of "Arms", and these terms are defined in subsequent paragraphs. The only drawback to GQ is that it does not work if there are surface apertures in the optical system. For these systems, GQ is probably no longer a better choice than RA. GQ does work fine when used with vignetting factors, since the ray pattern is simply redistributed.

The RA algorithm traces a grid of rays through the pupil. The "Grid" size determines the number of rays traced, and is described in a subsequent paragraph. The "Delete Vignetted" option (also described later) allows the vignetted rays to be deleted from the ray set. Vignetted rays in this context are those rays clipped by surface apertures, not rays which have been altered by the use of vignetting factors (see the chapter "Conventions and Definitions"). The advantage to the RA algorithm is the ability to accurately account for the effects of vignetting in the merit function. This is useful in systems such as obscured telescopes and camera lenses which intentionally clip troublesome rays. The disadvantage to the RA algorithm is speed and accuracy. Usually, more rays are required to achieve a given degree of accuracy than the GQ algorithm. The bottom line: don't use RA unless you are using surface apertures.

Rings


The "Rings" setting is only used in the GQ algorithm. It determines how many rays are traced at each field and at each wavelength. For on-axis fields (zero degrees field angle in a rotationally symmetric system), the number of rays is equal to the number of rings. For all other fields in symmetric systems, the number of rays traced per ring is equal to half the number of "arms" (defined in the next paragraph). Only half the rays are traced because of the left-right symmetry of the system is exploited. Each set of rays is traced for each defined wavelength. For example, if you have one on-axis field, two off-axis fields, three wavelengths, and four rings selected, the number of rays traced is $3 * (4 + 4*3 + 4*3) = 84$. For systems without rotational symmetry, the number of rays per ring is the number of "arms" independent of field. In the prior example, this means $3 * 3 * 4 * 6 = 216$ rays. ZEMAX automatically calculates these numbers for you; the only reason it is described here is so you will understand how the default merit function is defined. Optimization runs are longer if more rays are traced.

Arms

The "Arms" setting is also only used in the GQ algorithm. It determines how many radial arms of rays in the pupil are traced. By default six equally spaced (in angle) arms are traced (or three if the system is rotationally symmetric). This number may be changed to eight, ten, or twelve. For most common optical systems, six is sufficient.

You should select the number of rings and the number of arms according to the order of aberrations present in your system. A simple way of determining the correct number of rings is to select the minimum number, one. Then go to the optimization dialog box and note the merit function. Now go back to the default merit function tool, and select two rings. If the merit function changes by more than a few percent, go back and select three, and so on until the merit function does not change significantly (perhaps 1%). Repeat the procedure for the number of arms

(six arms is almost always plenty). Selecting more rings or arms than required will not improve the optimization performance, it will only slow the algorithm down needlessly. Tracing more rays than required will not help you find better solutions!

 **Selecting more rings or arms than required will not improve the optimization performance, it will only slow the algorithm down needlessly.**

Grid

The "Grid" is only used by the RA algorithm, and the value determines the number of rays to be used. The grid size can be 4x4 (16 rays per field per wavelength), 6x6 (36 rays per field per wavelength) etc. Rays on the grid are automatically omitted if they fall outside the entrance pupil, so the actual number of rays used will be lower than the grid size squared. Selecting a larger grid size generally yields more accurate results at the expense of slower execution. However, there may be an advantage in choosing a large grid density, and then selecting the "Delete Vignetted" checkbox (described in the next paragraph). The reason is that a large grid density will fill the pupil with rays, and then the operands which are vignetted will be deleted. The result is a reasonable number of rays which accurately reflect the aperture of the system.

Delete Vignetted

The "Delete Vignetted" checkbox option is only used by the RA algorithm. If selected, then each ray in the merit function will be traced through the system, and if it is vignetted by a surface aperture, if it misses any surface, or if it is total internal reflected at any surface, the ray is deleted from the merit function. This keeps the total number of rays in the merit function to a minimum. The disadvantage is that if the vignetting changes as the design is optimized, then the merit function may have to be regenerated. It is always a better choice to use the vignetting factors and then use the GQ algorithm than to delete vignetted rays if possible. Vignetting factors can be adjusted, if required, during optimization using SVIG in the merit function.

Note that ZEMAX will attempt to trace any ray defined in the merit function regardless if that ray is vignetted or not. For example, if the chief ray height is targeted using REAY, and there is a central obscuration that vignettes the chief ray, ZEMAX will still trace the ray and use the operand results as long as the ray can be traced. ZEMAX does not check to see if defined rays are vignetted, because this introduces substantial overhead during optimization.

In general, avoid vignetting of rays by surface apertures, and use vignetting factors to shape the beam size when possible. To optimize on the fraction of unvignetted rays, a macro must be defined to perform the required computations. However, this method is very prone to stagnation during optimization because small changes in lens parameters lead to discrete changes in the merit function as rays abruptly jump from being vignetted to not being vignetted.

Setting thickness boundary values

Boundary constraints may be automatically generated and included in the default merit function by checking the air and/or glass boundary values on. If selected, then MNCG, MXCG, and MNEG operands will be added to the merit function to constrain the minimum center thickness, maximum center thickness, and minimum edge thickness for glass surfaces, respectively. MNCA, MXCA, and MNEA operands will be added to the merit function to constrain the minimum center thickness, maximum center thickness, and minimum edge thickness for air spaces, respectively.

The automatic boundary constraint feature is meant to save some manual entry of routine boundary constraints on optical systems with or without mirrors. More complex lenses, such as those with complex coordinate breaks, or multi-configurations usually require additional boundary constraints to be added to the merit function manually.

Start At

The "Start At" option is used to add the default merit function at a specific position within the Merit Function Editor operand list. ZEMAX will attempt to place the default merit function after the existing targets. The algorithm used to determine the starting pointing may be ineffective if the default merit function has been edited. To control where the default start at value will be, see the DFMS operand definition.

Assume Axial Symmetry

If selected, then the default merit function will exploit the left-right and rotational symmetry of the lens when constructing and evaluating the merit function. Fewer rays will be traced, accelerating the optimization with no loss of accuracy. In systems with coordinate breaks or non-rotationally symmetric systems, the default is unselected, which means symmetry will not be exploited. Overriding the default system symmetry is useful, however, if you are designing a lens that ZEMAX thinks is non-symmetric, but the lack of symmetry does not affect

the aberrations. For example, if tilted but flat fold mirrors are present, these mirrors do not eliminate the left-right symmetry of the system, but ZEMAX will by default assume symmetry does not exist. Some gradient index surfaces also use non-symmetric index variation terms which are often zero (they are used only for tolerancing). Select the checkbox to accelerate the optimization in these cases. See also the "USYM" operand description. If you do not understand this feature, just use the default.

Ignore Lateral Color

By default, ZEMAX references all RMS or PTV computations to a common reference point for each field. All of the rays are traced for all wavelengths for each field point, and the primary wavelength chief ray or the centroid of all the rays is used as the reference point. If "Ignore Lateral color" is selected, then ZEMAX computes an independent reference point for each wavelength as well. This is useful for designing systems that intentionally divide the beam by wavelength, such as a prism or spectrometer system. This option will cause the merit function to optimize each color spot independently.

Relative X Weight

The relative X weight is the additional weighting to be placed on the X component of the transverse aberrations when computing the PTV or RMS SPOT X + Y merit function. This setting has no effect on the other merit functions. If the relative X weight is less than unity, then the Y components are weighted more heavily; if the relative weight is greater than unity, then the X components are weighted more heavily. If left at the default value of unity, then the components are equally weighted. This control is useful for systems which intentionally form slit images, such as spectrometers.

Pitfalls with the default merit function

The default merit function is easy to set up, numerically efficient, and suitable for a large number of optimization problems. However, most optical designs require extensions or modifications to the default as the design progresses. ZEMAX offers significant flexibility in the definition of the merit function, as described in the following sections.

Note that if the field or wavelength values or weights are changed, you must reconstruct the default merit function. If you are using the RA algorithm, reconstruct the default merit function if the vignetting influence changes appreciably during optimization.



If field or wavelength values or weights are changed, reconstruct the merit function.

Optimization with apodized beams

If no pupil apodization has been specified (see the "System" chapter for details on specifying the pupil apodization) then ZEMAX assumes uniform illumination when constructing the default merit function. If the illumination is not uniform, then the rays in the default merit function are weighted according to the apodization factor. Since the rays selected may be insufficient to adequately represent an apodized beam, use a larger number of rays (described previously) when using apodization factors. See the chapter "Conventions and Definitions" for more information about apodization.

Modifying the merit function

The merit function can be modified by the user. To change the merit function, select Editors, Merit Function from the main menu bar. New operands can be added to the list, or others deleted, using the insert and delete keys. The current merit function value and the value of each operand can be updated by selecting Tools, Update.

Operands are set by typing the name in the first column and then filling in the remaining data fields. There are eight fields that may be required to define an operand: Int1, Int2, Hx, Hy, Px, Py, target, and weight. The Int values are integer parameters whose meaning depends upon the operand selected. Usually, Int1 is the surface indicator, and Int2 is the wavelength indicator, but not always. Not all of the operands use all of the fields provided.

Int1 and Int2

For those operands that use Int1 to indicate the surface number, the parameter specifies at which surface the target should be evaluated. Similarly, the Int2 value, when used as a wavelength specifier, describes which wavelength to use. Int2 must be an integer value equal to the wavelength number. The Int1 and Int2 parameters have other uses as described later.

Hx, Hy, Px, Py

Many of the operands use Hx, Hy, Px, and Py; these are the normalized field and pupil coordinates (see "Normalized field and pupil coordinates" in the chapter "Conventions and definitions"). Note that ZEMAX does not check to see if the specified Hx, Hy, Px, and Py coordinates are within the unit circle. For example a pupil coordinate of (1, 1) is actually outside the entrance pupil, but you will not get an error message when tracing these rays unless the rays cannot physically be traced. Some operands use the Hx, Hy, Px, and Py data fields for other values, depending upon the operand type.

The target is the desired value of the specified parameter. The difference between the target and the value of the operand is squared, and summed over all operands to yield the merit function. The value of the target and the operand itself is unimportant in optimization, only the difference between the two. The larger the difference, the greater the contribution to the merit function.

The weight is the relative importance of that parameter. The weight can be any number, positive or negative. However, the optimizer will act somewhat differently if the weight is negative, zero, or positive.

Operand weights less than zero

When the weight is negative, the operand will be treated as a Lagrangian multiplier. The Lagrangian multipliers force the optimization algorithm to find a solution which exactly meets the specified constraint, regardless of the effect on the other operands. This is sometimes useful to exactly meet an optimization target, such as focal length or magnification. In some respects, this is similar to a weight of "infinity", however it is implemented in a way that is numerically more stable.

Because there is generally a non-linear relationship between the variables and the operand targets, ZEMAX may not converge to the exact target value in a single optimization cycle; however, multiple cycles will usually converge to the Lagrangian targets with extremely high precision in a few cycles if a solution exists. It is possible to define Lagrangian targets that cannot be met with the variables provided, especially if there is more than one Lagrangian target defined.

Note that when using Lagrangian multipliers, the merit function may increase after optimization, as ZEMAX modifies the system to meet the exact constraints. For purposes of computing the overall merit function value, ZEMAX will use the absolute value of the weight.

For best results, use Lagrangian multipliers sparingly, if at all. Better optimization and adequate accuracy is usually just as easily achieved using heavier weights on those operands which require exact (or nearly so) values.

Operand weights equal to zero

When the weight is zero, the optimization algorithm calculates but ignores the operand. This is very useful for computing a result that does not have a specific target, but might be used elsewhere in the merit function; or if the value is simply used as a check or monitored parameter.

Operand weights greater than zero

If the weight is greater than zero, then the operand will be treated as an "aberration" to be minimized along with the merit function. The vast majority of operands should have positive weights.

Merit function definition

The merit function is defined as:

$$MF^2 = \frac{\sum W_i (V_i - T_i)^2 + \sum W_j (V_j - T_j)^2}{\sum W_i}$$

The summation over i includes only positive weighted operands, while the summation over j includes only Lagrangian multiplier operands. The absolute value of the Lagrangian multiplier weight is used. This convention is chosen so that adding Lagrangian multipliers to control boundary conditions has no effect on the merit function if the Lagrangian targets are all met.

Optimization operands

The following tables describe the available operands. The first table is a "quick reference" guide which categorizes the operands by general subject. The second table provides a detailed description of each operand (listed alphabetically) and states which operands use which data fields. Note in particular that some of the operands

(such as SUMM) use the Int1 and Int2 values to indicate something other than surface and wavelength. If an operand does not use a data field, a "-" is shown.

OPTIMIZATION OPERANDS BY CATEGORY

Category	Related Operands
First-order optical properties	EFFL, PIMH, PMAG, AMAG, ENPP, EXPP, LINV, WFNO, POWR, EPDI, ISFN, OBSN, EFLX, EFLY, SFNO, TFNO
Aberrations	SPHA, COMA, ASTI, FCUR, DIST, DIMX, AXCL, LACL, TRAR, TRAX, TRAY, TRAI, OPDC, PETZ, PETC, RSCH, RSCE, RWCH, RWCE, ANAR, ZERN, TRAC, OPDX, RSRE, RSRH, RWRE, TRAD, TRAE, TRCX, TRCY, DISG, FCGS, FCGT, DISC, OPDM, RWRH, BSER
MTF data	MTFT, MTFS, MTFA, MSWT, MSWS, MSWA, GMTA, GMTS, GMTT
Encircled energy	DENC, DENF, GENC, XENC
Constraints on lens data	TOTR, CVVA, CVGT, CVLT, CTVA, CTGT, CTLT, ETVA, ETGT, ETLT, COVA, COGT, COLT, DMVA, DMGT, DMLT, TTHI, VOLU, MNCT, MNET, MXCT, MXET, MNCG, MNEG, MXCG, MXEG, MNCA, MNEA, MXCA, MXEA, ZTHI, SAGX, SAGY, CVOL, MNSD, MXSD, XXET, XxEA, XXEG, XNET, XNEA, XNEG, TTGT, TTLT, TTVA, TMAS, MNCV, MXCV, MNDT, MXDT
Constraints on parameter data	PMVA, PMGT, PMLT
Constraints on extra data	XDVA, XDGT, XDLT
Constraints on glass data	MNIN, MXIN, MNAB, MXAB, MNPD, MXPD, RGLA, GCOS, GTCE, INDX
Constraints on paraxial ray data	PARX, PARY, PARZ, PARR, PARA, PARB, PARC, PANA, PANB, PANC, PATX, PATY, YNIP
Constraints on real ray data	REAX, REAY, REAZ, REAR, REAA, REAB, REAC, RENA, RENB, RENC, RANG, OPTH, DXDX, DXDY, DYDX, DYDY, RETX, RETY, RAGX, RAGY, RAGZ, RAGA, RAGB, RAGC, RAIN, PLEN, HHCN, RAID, RAEN, RAED, IMAE
Constraints on element positions	GLCX, GLCY, GLCZ, GLCA, GLCB, GLCC
Changing system data	CONF, PRIM, SVIG
General math operands	ABSO, SUMM, OSUM, DIFF, PROD, DIVI, SQRT, OPGT, OPLT, CONS, QSUM, EQUA, MINN, MAXX, ACOS, ASIN, ATAN, COSI, SINE, TANG
Multi-configuration (zoom) data	CONF, ZTHI, MCOV, MCOL, MCOG
Gaussian beam data	GBWA, GBW0, GBWZ, GBWR, GBWD
Gradient index control operands	I1GT, I2GT, I3GT, I4GT, I5GT, I6GT, I1LT, I2LT, I3LT, I4LT, I5LT, I6LT, I1VA, I2VA, I3VA, I4VA, I5VA, I6VA, GRMN, GRMX, LPTD, DLTN
Ghost focus control	GPIM
Fiber coupling operands	FICL
Optimization with ZPL macros	ZPLM
User defined operands	UDOP
Merit function control operands	BLNK, ENDX, USYM, DMFS, SKIS, SKIN

Category	Related Operands
Constraints on non-sequential object data.	NPXG, NPXL, NPXV, NPYG, NPYL, NPYV, NPZG, NPZL, NPZV, NTXG, NTXL, NTXV, NTYG, NTYL, NTYV, NTZG, NTZL, NTZV, NPGT, NPLT, NPVA
Non-sequential ray tracing and detector operands.	NSDD, NSTR
Constraints on construction optics for optically fabricated holograms	CMFV

OPTIMIZATION OPERANDS AND DATA FIELD USAGE

NAME	Description	Int1	Int2	Hxy, Pxy
		Op #	-	-
ABSO	Absolute value.	Op #	Flag	-
ACOS	Arccosine of the value of the specified operand number. If flag is 0, then the units are radians, otherwise, degrees.	-	Wave	-
AMAG	Angular magnification. This is the ratio of the image to object space paraxial chief ray angles. Not valid for non-paraxial systems.	-	Wave	-
ANAR	Angular aberration radius measured in the image plane with respect to the primary wavelength chief ray. This quantity is defined as $1 - \cos\theta$, where θ is the angle between the chief ray and the ray being traced. See TRAR.	-	Wave	Yes
ASIN	Arcsine of the value of the specified operand number. If flag is 0, then the units are radians, otherwise, degrees.	Op #	Flag	-
ASTI	Astigmatism in waves contributed by the specified surface. If the surface value is zero, the sum for the entire system is used. This is the third order astigmatism calculated from the Seidel coefficients, and is not valid for non-paraxial systems.	Surf	Wave	-
ATAN	Arctangent of the value of the specified operand number. If flag is 0, then the units are radians, otherwise, degrees.	Op #	Flag	-
AXCL	Axial color in lens units. This is the paraxial focal plane separation between the two most extreme wavelengths defined. The distance is measured along the Z axis. Not valid for non-paraxial systems.	-	-	-
BLNK	Does nothing. Used for separating portions of the operand list. A comment line may optionally be typed in the space to the right of the operand name; this comment will be displayed in the editor as well as in the merit function listing.	-	-	-
BSER	Boresight error. Boresight error is defined as the radial chief ray coordinate traced for the on axis field divided by the effective focal length. This definition yields a measure of the angular deviation of the image.	-	Wave	-

NAME	Description	Int1	Int2	Hxy, Pxy
CMFV	Construction merit function value. This operand calls the merit functions defined in either of the two construction systems used to define an optically fabricated hologram. The Con# is either 1 or 2, for the first or second construction system, respectively. The Opr# is either 0, which will return the entire merit function value from the construction system; or it is an integer which defines the operand row # from which to extract the value from. For example, if Con# is 2 and Opr# is 7, CMFV will return the value of merit function operand 7 in construction file 2. If there are more than one optically fabricated hologram surfaces in the playback system being optimized, the Con# may be incremented by 2 to specify the second surface parameters be used, or by 4 to indicate the third hologram surface construction optics be used, and so on. For example, a Con# of 7 would indicate construction system 1 on the fourth optically fabricated hologram surface present.	Con#	Opr#	-
COGT	Boundary operand that constrains the conic of surface "Surf" to be greater than the specified target value.	Surf	-	-
COLT	Boundary operand that constrains the conic of surface "Surf" to be less than the specified target value.	Surf	-	-
COMA	Coma in waves contributed by the specified surface. If the surface value is zero, the sum for the entire system is used. This is the third order coma calculated from the Seidel coefficients, and is not valid for non-paraxial systems.	Surf	Wave	-
CONF	Configuration. This operand is used to change the configuration number during merit function evaluation, which permits optimization across multiple configurations. This operand does not use the target or weight columns.	New #	-	-
CONS	Constant value. This is used to enter in constant values for use in other operand computations. The value will be identical to the target value.	-	-	-
COSI	Cosine of the value of the specified operand number. If flag is 0, then the units are radians, otherwise, degrees.	Op #	Flag	-
COVA	Conic value. Returns the conic constant of a surface.	Surf	-	-
CTGT	Center thickness greater than. This boundary operand constrains the center thickness of surface "Surf" to be greater than the specified target value. See also "MNCT".	Surf	-	-
CTLT	Center thickness less than. This boundary operand constrains the center thickness of surface "Surf" to be less than the specified target value. See also "MXCT".	Surf	-	-
CTVA	Center thickness value. Constrains the center thickness of surface "Surf" to be equal to the specified target value.	Surf	-	-
CVGT	Curvature greater than. This boundary operand constrains the curvature of surface "Surf" to be greater than the target value.	Surf	-	-
CVLT	Curvature less than. This boundary operand constrains the curvature of surface "Surf" to be less than the target value.	Surf	-	-

NAME	Description	Int1	Int2	Hxy, Pxy
CVOL	Cylinder volume. This operand computes the volume in cubic lens units of the smallest cylinder that will contain the specified range of surfaces. Only the vertex positions and semi-diameters are used in the calculation, not the sag. The range of surfaces should not include any coordinate breaks.	First surf	Last surf	-
CVVA	Curvature value. This operand constrains the curvature of surface "Surf" to be equal to the specified target value.	Surf	-	-
DENC	Diffraction Encircled Energy (radius). This operand computes the radius in microns to the specified fraction of diffraction encircled, ensquared, x only, or y only (enslitted) energy. Int1 specifies the pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc. Int2 is the integer wavelength number; 0 for polychromatic. Hx specifies the field number. Hy is the fraction of energy desired, and must be between 0.0 and 1.0. Px is the type: 1 for encircled, 2 for x only, 3 for y only, and 4 for ensquared. Py is the reference point: 0 for centroid, 1 for chief ray, 2 for vertex. If the sampling is too low, the radius returned is a 1e+10. See also DENS, GENC and XENC. This operand is supported in XE and EE Only.	Sam- pling	Wave	(see left)
DENS	Diffraction Encircled Energy (fraction). This operand computes the fraction of diffraction encircled, ensquared, x only, or y only (enslitted) energy at a given distance from the reference point. The options and settings are identical to DENC, except Hy, which here is used as the distance at which the fraction of energy is desired. See also DENC, GENC and XENC. This operand is supported in XE and EE Only.	Sam- pling	Wave	(see left)
DIFF	Difference of two operands (OP#1 - OP#2). The two arguments are the row numbers of the operands to subtract.	Op #1	Op #2	-
DIMX	Distortion maximum. This is similar to DIST, except it specifies only an upper bound for the absolute value of the distortion. The Field integer can be zero, which specifies the maximum field coordinate be used, or any valid field number. Note the maximum distortion does not always occur at the maximum field coordinate. The value returned is always in units of percentage, for the system as a whole. This operand may not be valid for non-rotationally symmetric systems.	Field	Wave	-
DISC	Distortion, calibrated. This operand computes the calibrated distortion across the field of view, and returns the absolute value of the maximum deviation from linearity of the f-theta condition. This operand is extremely useful for designing f-theta lenses.	-	Wave	-
DISG	Generalized distortion in percent. This operand computes the distortion for any ray in the pupil, from anywhere in the field, at any wavelength, using any field point as a reference. The method used and assumptions made are the same as for the grid distortion plot described in the Analysis Menu Chapter.	Ref Fld	Wave	Yes

NAME	Description	Int1	Int2	Hxy, Pxy
DIST	Distortion in waves contributed by surface "Surf". If the surface value is zero, the sum for the entire system is used. Also, if the surface number is zero, the distortion is given in percent instead. This is the third order distortion calculated from the Seidel coefficients, and is not valid for non-paraxial systems.	Surf	Wave	-
DIVI	Division of first by second operand. The two arguments are the row numbers of the operands to divide.	Op #1	Op #2	-
DLTN	Delta N. Computes the difference between the maximum and minimum index of refraction on axis for a gradient index surface. The minimum and maximum z coordinates used are computed by accounting for the sag of both ends of the surface. See the section "Using gradient index operands".	Surf	Wave	-
DMFS	Default merit function start. This operand is simply a flag to indicate where the default merit function should be appended, if one is subsequently created. The row number after this operand will appear as the default "Start At" row on the default merit function dialog box.	-	-	-
DMGT	Diameter greater than. This boundary operand constrains the diameter of surface "Surf" to be greater than the specified target value. The diameter is twice the semi-diameter value displayed on the main spreadsheet.	Surf	-	-
DMLT	Diameter less than. This boundary operand constrains the diameter of surface "Surf" to be less than the specified target value. The diameter is twice the semi-diameter value displayed on the main spreadsheet.	Surf	-	-
DMVA	Diameter value. This operand constrains the diameter of surface "Surf" to be equal to the specified target value. The diameter is twice the semi-diameter value displayed on the main spreadsheet.	Surf	-	-
DXDX	Derivative of transverse x-aberration with respect to x-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate.	-	Wave	Yes
DXDY	Derivative of transverse x-aberration with respect to y-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate.	-	Wave	Yes
DYDX	Derivative of transverse y-aberration with respect to x-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate.	-	Wave	Yes
DYDY	Derivative of transverse y-aberration with respect to y-pupil coordinate. This is the slope of the ray fan plot at the specified pupil coordinate.	-	Wave	Yes
EFFL	Effective focal length in lens units. This is paraxial, and may not be accurate for non-paraxial systems.	-	Wave	-
EFLX	Effective focal length in the local x plane of the specified range of surfaces at the primary wavelength.	First surf	Last surf	-
EFLY	Effective focal length in the local y plane of the specified range of surfaces at the primary wavelength.	First surf	Last surf	-
ENDX	End execution. Terminates the computation of the merit function. All remaining operands are ignored.	-	-	-

NAME	Description	Int1	Int2	Hxy, Pxy
ENPP	Entrance pupil position in lens units, with respect to the first surface. This is the paraxial pupil position, valid only for centered systems.	-	-	-
EPDI	Entrance pupil diameter in lens units.	-	-	-
EQUA	Equal operand. This operand constrains all operands within the specified range of operands to have the same value within the tolerance specified by the target. The value of this operand is computed by finding the average of the range of values, and then summing the absolute value of the errors between each operand and the average if the error exceeds the target value. See SUMM and OSUM.	First	Last	-
ETGT	Edge thickness greater than. This boundary operand constrains the edge thickness of surface "Surf" to be greater than the specified target value. The edge thickness is calculated at the semi-diameter radius along the +y axis if code is zero, the +x axis if code is 1, the -y axis if code is 2, and the -x axis if code is 3. See also "MNET".	Surf	Code	-
ETLT	Edge thickness less than. This boundary operand constrains the edge thickness of surface "Surf" to be less than the specified target value. The edge thickness is always calculated at the semi-diameter radius along the +y axis if code is zero, the +x axis if code is 1, the -y axis if code is 2, and the -x axis if code is 3. See also "MXET".	Surf	Code	-
ETVA	Edge thickness value. Constrains the edge thickness of surface "Surf" to be equal to the specified target value. The edge thickness is always calculated at the semi-diameter radius along the +y axis if code is zero, the +x axis if code is 1, the -y axis if code is 2, and the -x axis if code is 3. See also "MNET".	Surf	Code	-
EXPP	Exit pupil position in lens units, with respect to the image surface. This is the paraxial pupil position, valid only for centered systems.	-	-	-
FCGS	Generalized field curvature, sagittal. The field curvature value for any field point, at any wavelength. The value is generalized to return reasonable results even for non-rotationally symmetric systems; see the Field Curvature feature in the Analysis Menu Chapter.	-	Wave	Hx,Hy
FCGT	Generalized field curvature, tangential; see FCGS.	-	Wave	Hx,Hy
FCUR	Field curvature in waves contributed by the specified surface. If the surface value is zero, the sum for the entire system is used. This is the third order field curvature calculated from the Seidel coefficients, and is not valid for non-paraxial systems.	Surf	Wave	-
FICL	Fiber coupling efficiency for single mode fibers. The sampling defines the grid size used for the integration; with 1 being 32 x 32, 2 being 64 x 64, etc. The wavelength must be monochromatic, and the wavelength number specified in the Int2 column. The Hx value is the integer field position number. If Hy is zero, then the object source fiber is considered; if Hy is non zero, the object source fiber is ignored. Px and Py are used to define the source and receiver fiber NA's respectively. The calculated value is the total coupled energy efficiency, relative to unity. See "Fiber Coupling Efficiency" on page 128 for details. This operand is only supported in the XE and EE editions of ZEMAX.	Sam- pling	Wave	See left

NAME	Description	Int1	Int2	Hxy, Pxy
FOUC	Foucault analysis. This operand returns the RMS difference between the computed and reference shadowgram as computed by the Foucault analysis feature, using whatever the current default settings are. To use this operand, first define the settings on the Foucault analysis feature as desired, then press Save on the settings box. The data option "difference" MUST be selected to return valid data. The operand FOUC will return the RMS difference between the computed and reference shadowgrams. Using this operand, the optical system wavefront aberrations may be optimized to produce the reference shadowgram.	-	-	-
GBWO	Minimum Gaussian beam waist in image space of specified surface. If Hx is non-zero, then the computation is for the x-direction beam, otherwise, it is for the y-direction. The Hy value is used to define the input beam waist, and Px is used to define the distance from Surface 1 to the waist location. See the Gaussian beam feature for details.	Surf	Wave	See left
GBWA	Gaussian beam size on surface "Surf". If Hx is non-zero, then the computation is for the x-direction beam, otherwise, it is for the y-direction. The Hy value is used to define the input beam waist, and Px is used to define the distance from Surface 1 to the waist location. See the Gaussian beam feature for details.	Surf	Wave	See left
GBWD	Gaussian beam divergence in optical space following surface "Surf". If Hx is non-zero, then the computation is for the x-direction beam, otherwise, it is for the y-direction. The Hy value is used to define the input beam waist, and Px is used to define the distance from Surface 1 to the waist location. See the Gaussian beam feature for details.	Surf	Wave	See left
GBWR	Gaussian beam radius of curvature at the specified surface. If Hx is non-zero, then the computation is for the x-direction beam, otherwise, it is for the y-direction. The Hy value is used to define the input beam waist, and Px is used to define the distance from Surface 1 to the waist location. See the Gaussian beam feature for details.	Surf	Wave	See left
GBWZ	The distance from the image space Gaussian beam waist to the surface. If Hx is non-zero, then the computation is for the x-direction beam, otherwise, it is for the y-direction. The Hy value is used to define the input beam waist, and Px is used to define the distance from Surface 1 to the waist location. See the Gaussian beam feature for details.	Surf	Wave	See left
GCOS	Glass cost. This operand returns the relative cost factor as listed in the glass catalog for the glass on the specified surface.	Surf	-	-
GENC	Geometric Encircled Energy. This operand computes the radius in microns to the specified fraction of geometric encircled, ensquared, x only, or y only (enslitted) energy. Int1 specifies the pupil sampling, where 1 yields 32 x 32, 2 yields 64 x 64 etc. Int2 is the integer wavelength number; 0 for polychromatic. Hx specifies the field number. Hy is the fraction of energy desired, and must be between zero and 1, exclusive. Px is the type: 1 for encircled, 2 for x only, 3 for y only, and 4 for ensquared. See also DENC and XENC. This operand is supported in XE and EE Only.	Sam- pling	Wave	(see left)

NAME	Description	Int1	Int2	Hxy, Pxy
GLCA	Global x-direction orientation vector component of surface "Surf".	Surf	-	-
GLCB	Global y-direction orientation vector component of surface "Surf".	Surf	-	-
GLCC	Global z-direction orientation vector component of surface "Surf".	Surf	-	-
GLCX	Global vertex x-coordinate of surface "Surf".	Surf	-	-
GLCY	Global vertex y-coordinate of surface "Surf".	Surf	-	-
GLCZ	Global vertex z-coordinate of surface "Surf".	Surf	-	-
GMTA	Geometric MTF average of sagittal and tangential response. The Int1 parameter must be an integer (1, 2, ...) where 1 yields 32 x 32 sampling, 2 yields 64 x 64 sampling, etc. The Int2 can be a valid wavelength number, or 0 for polychromatic. The Hx value must be a valid field number (1, 2, ...). Hy is the spatial frequency in cycles per mm. Px is a flag; if zero, then the diffraction limit will be used to scale the results (recommended) otherwise, no scaling is done. See the discussion "Using MTF operands" in this chapter for details.	See left.	+	+
GMTS	Geometric MTF sagittal response. See GMTA for details.	See left.	+	+
GMTT	Geometric MTF tangential response. See GMTA for details.	See left.	+	+
GPIM	Ghost pupil image. GPIM controls the location of ghost pupils (and optionally ghost images) relative to the image plane. Double-bounce ghosts form images of the pupil, and if these images are formed near the focal plane will contaminate the image with unwanted light. This is the cause of the familiar "sun flare" images of the pupil seen through camera lenses pointed near the sun. The operand computes any one specific or all possible ghost pupil image locations and returns one over the absolute value of the distance from the image plane to the closest pupil ghost. The operand is defined in this manner so it can simply be targeted to zero and weighted and optimized to reduce ghost pupil affects. If the int1 and int2 parameters are set to specific surface numbers, that specific ghost path is computed, if either or both of the int values are -1, then all possible surface combinations are considered. For example, if Int1 is 12 and int2 is -1, then all double bounces that first bounce off surface 12 and then 11, 10, 9, etc. are considered, if both numbers are negative, all possible ghosts are considered. This same operand also can be used for detecting and controlling image ghosts (which are distinct from pupil ghosts) by changing the "mode" flag in the Hx column from 0 to 1, or to control ghost pupil magnification, by setting the mode to 2. The WFB and WSB columns will list the worst combination found for reference and possible further analysis. Only surfaces with index changes are considered as possible ghost generators. First bounces off mirrors are ignored.	First surface	Second surface	See left
GRMN	Gradient index minimum index. This boundary operand sets the minimum allowable index of refraction value for the gradient index surface "Surf" at the wavelength number "Wave". The index is checked at six places: the front vertex, front +y top, front +x side, rear vertex, rear +y top, and the rear +x side. See also "InLT", "InLT" and "GRMX".	Surf	Wave	-

NAME	Description	Int1	Int2	Hxy, Pxy
GRMX	Gradient index maximum index. This boundary operand sets the maximum allowable index of refraction value for the gradient index surface "Surf" at the wavelength number "Wave". The index is checked at six places: the front vertex, front +y top, front +x side, rear vertex, rear +y top, and the rear +x side. See also "InGT", "InLT" and "GRMN".	Surf	Wave	-
GTCE	Glass TCE. This operand returns the Thermal Coefficient of Expansion Alpha1 as listed in the glass catalog for the glass on the specified surface.	Surf	-	-
HHCN	Test for the hyperhemisphere condition. ZEMAX traces the specified ray to the specified surface, and computes the x, y, and z intercept coordinates. Then, the x and y coordinates only are used in the sag expression for that surface to see what z coordinate results. If the z coordinates are not the same, then HHCN returns 1, otherwise it returns zero. This operand can be used to prevent optimizations from reaching solutions that require hyperhemispheric surface shapes.	Surf	Wave	Yes
IMAE	Image analysis efficiency. This operand returns the fractional efficiency as computed by the geometric image analysis feature, using whatever the current default settings are (except for "show" which is always set to spot diagram for this computation). To use this operand, first define the settings on the geometric image analysis feature as desired, then press Save on the settings box. The operand IMAE will return the efficiency (normalized to unity) as computed by the image analysis feature. See the discussion "Optimizing with the IMAE operand" on page 306.	-	-	-
INDX	Index of refraction. Returns the current index at any surface and defined wavelength.	Surf	Wave	
InGT	Index "n" greater than. This boundary operand constrains the index of refraction at wavelength number "Wave" of gradient index surface "Surf" at one of six points inside the gradient index lens. For n=1, the point is the front vertex, n=2 is the front +y top, n=3 is the front +x side, n=4 is the rear vertex, n=5 is the rear +y top, and n=6 is the rear +x side. In all cases the operand bounds the index at the specified point to be greater than the specified target value. For example, "I4GT" constrains the minimum index at the rear vertex of surface (i.e. the vertex of the next surface) of the gradient index lens. In all cases the +y top and +x side distance is defined by the larger of the front and rear semi-diameters set on the main spreadsheet. See also "GRMN" and "GRMX", which are similar operands that are easier to use.	Surf	Wave	-
InLT	Index "n" less than. This operand is similar to "InGT" except it constrains the maximum value of the index of refraction rather than the minimum. See "InGT" for a complete description of the parameter "n".	Surf	Wave	-
InVA	This operand is similar to "InGT" except it simply constrains the current value of the index of refraction. See "InGT" for a complete description of the parameter "n".	Surf	Wave	-
ISFN	Image space F/#. This operand is the paraxial infinite conjugate F/#. See "WFNO".	-	-	-
LACL	Lateral color. This is the y-distance between the paraxial chief ray intercepts of the two extreme wavelengths defined. Not valid for non-paraxial systems.	-	-	-

NAME	Description	Int1	Int2	Hxy, Pxy
LINV	Lagrange invariant of system in lens units. The paraxial marginal and chief ray data are used to compute this value.	-	Wave	-
LPTD	This boundary operand constrains the slope of the axial gradient index profile from changing signs within a gradient index component. See the section "Using gradient index operands".	Surf	-	-
MAXX	Returns the largest value within the indicated range of operands. See MINN.	First	Last	-
MCOG	Multi-configuration operand greater than. This is used to constrain values in the multi-configuration editor.	MC operand #	Config #	-
MCOL	Multi-configuration operand less than. This is used to constrain values in the multi-configuration editor.	MC operand #	Config #	-
MCOV	Multi-configuration operand value. This is used to directly target or compute values in the multi-configuration editor.	MC operand #	Config #	-
MINN	Returns the smallest value within the indicated range of operands. See MAXX.	First	Last	-
MNAB	Minimum Abbe number. This boundary operand constrains the Abbe number of surfaces between "First Surf" and "Last Surf" to be greater than the specified target value. See also "MXAB". This operand only considers surfaces using model or substitute status catalog glasses.	First surf	Last surf	-
MNCA	Minimum center thickness air. This boundary operand constrains each of the center thicknesses of surfaces from "First surf" to "Last surf" which have air (i.e. no glass) as the glass type to be greater than the specified target value. See also "MNCT" and "MNCG". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MNCG	Minimum center thickness glass. This boundary operand constrains each of the thicknesses of surfaces from "First surf" to "Last surf" which have a non-air glass type to be greater than the target value. See also "MNCT" and "MNCA". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MNCT	Minimum center thickness. This boundary operand constrains each of the center thicknesses of surfaces from "First surf" to "Last surf" to be greater than the specified target value. See also "MNCG" and "MNCA". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MNCV	Minimum curvature. This boundary operand constrains each of the curvatures of surfaces from "First surf" to "Last surf" to be greater than the specified target value. See also "MXCV". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MNDT	Minimum diameter to thickness ratio. Controls the minimum allowable value on the ratio of surface diameter to center thickness. Only surfaces with non-unity index of refraction are considered. See also "MXDT". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-

Hxy, Pxy

NAME	Description	Int1	Int2	Hxy, Pxy
MNEA	Minimum edge thickness air. This boundary operand constrains each of the edge thicknesses of surfaces from "First surf" to "Last surf" which have air (i.e. no glass) as the glass type to be greater than the specified target value. See also "MNET", "MNEG", "ETGT", and "XNEA". This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XNEA for constraining non-rotationally symmetric surfaces.	First surf	Last surf	-
MNEG	Minimum edge thickness glass. This boundary operand constrains each of the edge thicknesses of surfaces from "First surf" to "Last surf" which have a non-air glass type to be greater than the specified target value. See also "MNET", "MNEA", "ETGT", and "XNEG". This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XNEG for constraining non-rotationally symmetric surfaces.	First surf	Last surf	-
MNET	Minimum edge thickness. This boundary operand constrains each of the edge thicknesses of surfaces from "First surf" to "Last surf" to be greater than the specified target value. See also "MNEG", "MNEA", "ETGT", and "XNET". This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XNET for constraining non-rotationally symmetric surfaces.	First surf	Last surf	-
MNIN	Minimum index at d-light. This boundary operand constrains the Nd value of surfaces between "First Surf" and "Last Surf" to be greater than the specified target value. See also "MXIN". This operand only considers surfaces using model or substitute status catalog glasses.	First surf	Last surf	-
MNPD	Minimum $\Delta P_{g,F}$. This boundary operand constrains the deviation of the partial dispersion of surfaces between "First Surf" and "Last Surf" to be greater than the specified target value. See also "MXPDP". This operand only considers surfaces using model or substitute status catalog glasses.	First surf	Last surf	-
MNSD	Minimum semi-diameter. Constrains the semi-diameter to be larger than the specified target over the surface range.	First surf	Last surf	-
MSWA	Modulation square-wave transfer function, average of sagittal and tangential. See "MTFT" for details.	Sampling	Wave	(see left)
MSWS	Modulation square-wave transfer function, sagittal. See "MTFT" for details.	Sampling	Wave	(see left)
MSWT	Modulation square-wave transfer function, tangential. See "MTFT" for details.	Sampling	Wave	(see left)
MTFA	Modulation transfer function, average of sagittal and tangential. See "MTFT" for details.	Sampling	Wave	(see left)
MTFS	Modulation transfer function, sagittal. See "MTFT" for details.	Sampling	Wave	(see left)

NAME	Description	Int1	Int2	Hxy, Pxy (see left)
MTFT	Modulation transfer function, tangential. This computes the diffraction MTF. The Int1 parameter must be an integer (1, 2, ...) where 1 yields 32 x 32 sampling, 2 yields 64 x 64 sampling, etc. The Int2 can be a valid wavelength number, or 0 for polychromatic. The Hx value must be a valid field number (1, 2, ...). Hy is the spatial frequency in cycles per mm. If the sampling is set too low for accurate computation of the MTF, then the MTF operands all return zero. If both the tangential and sagittal MTF are needed; place the MTFT and MTFS operands on adjacent lines and they will be computed simultaneously. See the discussion "Using MTF operands" in this chapter for details. This operand and the related MTFA, MTFS, MSWA, MSWT, and MSWS are supported in XE and EE Only.	Sampling	Wave	(see left)
MXAB	Maximum Abbe number. This boundary operand constrains the Abbe number of surfaces between "First Surf" and "Last Surf" to be less than the specified target value. See also "MNAB". This operand only considers surfaces using model or substitute status catalog glasses.	First surf	Last surf	-
MXCA	Maximum center thickness air. This boundary operand constrains each of the thicknesses of surfaces from "First surf" to "Last surf" which have air (i.e. no glass) as the glass type to be less than the target value. See also "MXCT" and "MXCG". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MXCG	Maximum center thickness glass. This boundary operand constrains each of the center thicknesses of surfaces from "First surf" to "Last surf" which have a non-air glass type to be less than the target value. See also "MXCT" and "MXCA". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MXCT	Maximum center thickness. This boundary operand constrains each of the center thicknesses of surfaces from "First surf" to "Last surf" to be less than the specified target value. See also "MXCG" and "MXCA". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MXCV	Maximum curvature. This boundary operand constrains each of the curvatures of surfaces from "First surf" to "Last surf" to be less than the specified target value. See also "MNCV". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MXDT	Maximum diameter to thickness ratio. Controls the maximum allowable value on the ratio of surface diameter to center thickness. Only surfaces with non-unity index of refraction are considered. See also "MNDT". This operand controls multiple surfaces simultaneously.	First surf	Last surf	-
MXEA	Maximum edge thickness air. This boundary operand constrains each of the edge thicknesses of surfaces from "First surf" to "Last surf" which have air (i.e. no glass) as the glass type to be less than the specified target value. See also "MXET", "MXEG", "ETLT", and "XXEA". This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XXEA for constraining non-rotationally symmetric surfaces.	First surf	Last surf	-

NAME	Description	Int1	Int2	Hxy, Pxy
MXEG	Maximum edge thickness glass. This boundary operand constrains each of the edge thicknesses of surfaces from "First surf" to "Last surf" which have a non-air glass type to be less than the target value. See also "MXET", "MXEA", "ETLT", and "XXEG". This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XXEG for constraining non-rotationally symmetric surfaces.	First surf	Last surf	-
MXET	Maximum edge thickness. This boundary operand constrains each of the edge thicknesses of surfaces from "First surf" to "Last surf" to be less than the specified target value. See also "MXEG", "MXEA", "ETLT", and "XXET". This operand controls multiple surfaces simultaneously. The boundary applies to the top "+y" edge of the surface only; see XXET for constraining non-rotationally symmetric surfaces.	First surf	Last surf	-
MXIN	Maximum index at d-light. This boundary operand constrains the Nd value of surfaces between "First Surf" and "Last Surf" to be less than the specified target value. See also "MNIN". This operand only considers surfaces using model or substitute status catalog glasses.	First surf	Last surf	-
MXPD	Maximum $\Delta P_{g,F}$. This boundary operand constrains the deviation of the partial dispersion of surfaces between "First Surf" and "Last Surf" to be less than the specified target value. See also "MNPD". This operand only considers surfaces using model or substitute status catalog glasses.	First surf	Last surf	-
MXSD	Maximum semi-diameter. Constrains the semi-diameter to be less than the specified target over the surface range.	First surf	Last surf	-
NPXG	Non-sequential object position x greater than.	Surf	Object	-
NPXL	Non-sequential object position x less than.	Surf	Object	-
NPXV	Non-sequential object position x value.	Surf	Object	-
NPYG	Non-sequential object position y greater than.	Surf	Object	-
NPYL	Non-sequential object position y less than.	Surf	Object	-
NPYV	Non-sequential object position y value.	Surf	Object	-
NPZG	Non-sequential object position z greater than.	Surf	Object	-
NPZL	Non-sequential object position z less than.	Surf	Object	-
NPZV	Non-sequential object position z value.	Surf	Object	-
NSDD	Non-sequential detector data. Detector refers to the object number of the desired detector. If the pixel number is zero, the detector is cleared. Otherwise, the data from the specified pixel is returned. Data is 0 for flux, 1 for flux/area, and 2 for flux/solid angle. See "Optimizing with sources and detectors in non-sequential mode" on page 305 for complete details.	Surf	Detector	See left
NSTR	Non-sequential trace. Source refers to the object number of the desired source. If source is zero, all sources will be traced. See "Optimizing with sources and detectors in non-sequential mode" on page 305 for complete details.	Surf	Source	See left
NTXG	Non-sequential object tilt about x greater than.	Surf	Object	-

NAME	Description	Int1	Int2	Hxy, Pxy
NTXL	Non-sequential object tilt about x less than.	Surf	Object	-
NTXV	Non-sequential object tilt about x value.	Surf	Object	-
NTYG	Non-sequential object tilt about y greater than.	Surf	Object	-
NTYL	Non-sequential object tilt about y less than.	Surf	Object	-
NTYV	Non-sequential object tilt about y value.	Surf	Object	-
NTZG	Non-sequential object tilt about z greater than.	Surf	Object	-
NTZL	Non-sequential object tilt about z less than.	Surf	Object	-
NTZV	Non-sequential object tilt about z value.	Surf	Object	-
NPGT	Non-sequential parameter greater than. The Hx value is used to define the parameter number.	Surf	Object	See left.
NPLT	Non-sequential parameter less than. The Hx value is used to define the parameter number.	Surf	Object	See left.
NPVA	Non-sequential parameter value. The Hx value is used to define the parameter number.	Surf	Object	See left.
OBSN	Object space numerical aperture. This is only useful for finite conjugate systems, and is calculated on axis at the primary wavelength.	-	-	-
OFF	This operand indicates an unused entry in the operand list. OFF operands are automatically converted to BLNK operands upon evaluation of the merit function. OFF is only used to indicate that the merit function operand type was not recognized.	-	-	-
OPDC	Optical path difference with respect to chief ray in waves.	-	Wave	Yes
OPDM	Optical path difference with respect to the mean OPD; this operand computes the OPD referenced to the mean OPD of all rays in the pupil. OPDM has the same restrictions that TRAC does; see TRAC for a detailed discussion.	-	Wave	Yes
OPDX	Optical path difference with respect to the shifted and tilted reference sphere that minimizes the RMS wavefront error; which ZEMAX calls the centroid reference. OPDX has the same restrictions that TRAC does; see TRAC for a detailed discussion.	-	Wave	Yes
OPGT	Operand greater than. This is used to make any operand an inequality constraint.	Op #	-	-
OPLT	Operand less than. This is used to make any operand an inequality constraint.	Op #	-	-
OPTH	Optical path length. This is the distance, in lens units, the specified ray travels to the surface "Surf". The distance is measured from the object for finite conjugates; otherwise the distance is referenced to the first surface. The optical path accounts for the index of refraction of the media, and for phase adding surfaces such as gratings and binary optics. See PLEN.	Surf	Wave	Yes
OSUM	Sums the values of all operands between the two specified operands. See SUMM.	First	Last	-

NAME	Description	Int1	Int2	Hxy, Pxy
PMGT	Parameter greater than. This boundary operand constrains the value of a parameter to be greater than the target value. The parameter values have different meanings depending upon the surface type. See the Chapter "Surface Types" for a description of the parameter values.	Surf	Param	-
PMLT	Parameter less than. This boundary operand constrains the value of a parameter to be less than the target value. The parameter values have different meanings depending upon the surface type. See the Chapter "Surface Types" for a description of the parameter values.	Surf	Param	-
PMVA	Parameter value. This boundary operand constrains the value of a parameter to be equal to the target value. The parameter values have different meanings depending upon the surface type. See the Chapter "Surface Types" for a description of the parameter values.	Surf	Param	-
PnGT	This operand is obsolete, use PMGT instead.	Surf	-	-
PnLT	This operand is obsolete, use PMLT instead.	Surf	-	-
PnVA	This operand is obsolete, use PMVA instead.	Surf	-	-
PANA	Paraxial ray x-direction surface normal at the ray-surface intercept. This is the x component of the surface normal vector at the intersection point of the specified paraxial ray and the surface "Surf", in the local coordinate system.	Surf	Wave	Yes
PANB	Paraxial ray y-direction surface normal at the ray-surface intercept. This is the y component of the surface normal vector at the intersection point of the specified paraxial ray and the surface "Surf", in the local coordinate system.	Surf	Wave	Yes
PANC	Paraxial ray z-direction surface normal at the ray-surface intercept. This is the z component of the surface normal vector at the intersection point of the specified paraxial ray and the surface "Surf", in the local coordinate system.	Surf	Wave	Yes
PARA	Paraxial ray x-direction cosine of the ray after refraction from the surface "Surf".	Surf	Wave	Yes
PARB	Paraxial ray y-direction cosine of the ray after refraction from the surface "Surf".	Surf	Wave	Yes
PARC	Paraxial ray z-direction cosine of the ray after refraction from the specified surface.	Surf	Wave	Yes
PARR	Paraxial ray radial coordinate in lens units at the specified surface. This is the radial distance from the local axis to the intersection of the surface "Surf" and the specified paraxial ray, in the local coordinate system.	Surf	Wave	Yes
PARX	Paraxial ray x-coordinate in lens units at the surface "Surf".	Surf	Wave	Yes
PARY	Paraxial ray y-coordinate in lens units at the surface "Surf".	Surf	Wave	Yes
PARZ	Paraxial ray z-coordinate in lens units at the surface "Surf".	Surf	Wave	Yes
PATX	Paraxial ray x-direction ray tangent. This is the tangent of the angle the paraxial ray makes in the X-Z plane after refraction from surface "Surf".	Surf	Wave	Yes

NAME	Description	Int1	Int2	Hxy, Pxy
		Surf	Wave	Yes
PATY	Paraxial ray y-direction ray tangent. This is the tangent of the angle the paraxial ray makes in the Y-Z plane after refraction from surface "Surf".	-	Wave	-
PETC	Petzval curvature in inverse lens units. Not valid for non-paraxial systems.	-	Wave	-
PETZ	Petzval radius of curvature in lens units. Not valid for non-paraxial systems.	-	Wave	-
PIMH	Paraxial image height at the paraxial image plane for the specified wavelength. Not valid for non-paraxial systems.	Surf 1	Surf 2	Yes
PLEN	Path length. This operand computes the total optical path length (including index of refraction) between surfaces 1 and 2 for the specified ray, which is always traced at the primary wavelength. See OPTH.	-	Wave	-
PMAG	Paraxial magnification. This is the ratio of the paraxial chief ray height on the paraxial image plane to the object height. Only useful for finite conjugate systems. Note the paraxial image plane is used even if the system is not at paraxial focus.	Surf	Wave	-
POWR	The surface power (in inverse lens units) of the surface "Surf". This operand only works for standard surfaces.	-	Wav #	-
PRIM	Primary wavelength. This is used to change the primary wavelength number during merit function evaluation. This operand does not use the target or weight columns.	Op #1	Op #2	-
PROD	Product of two operands. The two arguments are the row numbers of the operands to multiply.	First	Last	-
QSUM	Quadratic sum. This operand squares and then adds all operands between the first and last operand (inclusive), then takes the square root of the sum. See also SUMM, OSUM, EQUA.	Surf	Wave	Yes
RAED _x	Real ray angle of exitance. This is the angle in degrees between the surface normal and the ray after refraction or reflection. See also RAID.	Surf	Wave	Yes
RAEN	Real ray angle of exitance. This is the cosine of the angle between the surface normal and the ray after refraction or reflection at that surface. Does not return correct results if the glass prior to the surface is a gradient index medium. See also RAIN.	Surf	Wave	Yes
RAGA	Global ray x-direction cosine. This is the direction cosine of the ray in the global coordinate system. The origin of the global coordinate system is at the global reference surface.	Surf	Wave	Yes
RAGB	Global ray y-direction cosine. See "RAGA".	Surf	Wave	Yes
RAGC	Global ray z-direction cosine. See "RAGA".	Surf	Wave	Yes
RAGX	Global ray x-coordinate. This is the coordinate in lens units in the global coordinate system. The origin of the global coordinate system is at the global reference surface.	Surf	Wave	Yes
RAGY	Global ray y-coordinate. See "RAGX".	Surf	Wave	Yes
RAGZ	Global ray z-coordinate. See "RAGX".	Surf	Wave	Yes

NAME	Description	Int1	Int2	Hxy, Pxy
RAID	Real ray angle of incidence. This is the angle in degrees between the surface normal and the incident ray. See also RAED.	Surf	Wave	Yes
RAIN	Real ray angle of incidence. This is the cosine of the angle between the surface normal and the ray before refraction at that surface. Does not return correct results if the surface prior to surf is a gradient index surface. See also RAEN.	Surf	Wave	Yes
RANG	Ray angle in radians with respect to z axis. The angle is measured with respect to the local Z axis.	Surf	Wave	Yes
REAA	Real ray x-direction cosine of the ray after refraction from the surface "Surf".	Surf	Wave	Yes
REAB	Real ray y-direction cosine of the ray after refraction from the surface "Surf".	Surf	Wave	Yes
REAC	Real ray z-direction cosine of the ray after refraction from the surface "Surf".	Surf	Wave	Yes
REAR	Real ray radial coordinate in lens units at the surface "Surf".	Surf	Wave	Yes
REAX	Real ray x-coordinate in lens units at the surface "Surf".	Surf	Wave	Yes
REAY	Real ray y-coordinate in lens units at the surface "Surf".	Surf	Wave	Yes
REAZ	Real ray z-coordinate in lens units at the surface "Surf".	Surf	Wave	Yes
RENA	Real ray x-direction surface normal at the ray-surface intercept.	Surf	Wave	Yes
RENB	Real ray y-direction surface normal at the ray-surface intercept.	Surf	Wave	Yes
RENC	Real ray z-direction surface normal at the ray-surface intercept.	Surf	Wave	Yes
RETX	Real ray x-direction ray tangent (slope).	Surf	Wave	Yes
RETY	Real ray y-direction ray tangent (slope).	Surf	Wave	Yes
RGLA	Reasonable glass. This operand restricts the deviation the index, Abbe, and deviation of the partial dispersion values may take from actual glasses in the currently loaded glass catalogs. See "Optimizing glass selection" on page 303 for a complete discussion. The constraint is active over the surface range specified.	First Surf	Last Surf	See Text
RSCE	RMS spot size (ray aberrations) with respect to the geometric image centroid, measured in lens units. This operand is similar to RSCH, except the reference point is the image centroid instead of the chief ray. See RSCH for details.	Rings	Wave	Hx, Hy
RSCH	RMS spot size (ray aberrations) with respect to chief ray. This operand uses a Gaussian quadrature method to estimate the RMS spot size at a specified field coordinate and wavelength. The number returned is in lens units. The method used is only accurate for systems with circular pupils. The Int1 column is used to specify the number of rings of rays traced (use no more than required to converge the result). Only Hx and Hy are used to define the field point, Px and Py are not used. If the "wave" value is zero, then a wavelength weighted polychromatic calculation is performed.	Rings	Wave	Hx, Hy

NAME	Description	Int1	Int2	Hxy, Pxy Hx, Hy
RSRE	RMS spot size (ray aberrations) with respect to the geometric image centroid, measured in lens units. This operand is similar to RSCE, except a rectangular grid of rays is used instead of the Gaussian quadrature method. This operand always considers vignetting. A grid value of 1 will trace 4 rays, 2 will trace a 2 x 2 grid per quadrant (16 rays), 3 will trace 3 x 3 rays per quadrant (36 rays), and so on. Symmetry is considered.	Grid	Wave	Hx, Hy
RSRH	Similar to RSRE, except the reference point is the chief ray.	Grid	Wave	Hx, Hy
RWCE	RMS wavefront error with respect to the diffraction centroid. This operand is useful for minimizing the wavefront variance, which is proportional the Strehl ratio and the area under the MTF curve. The units are waves. See RWCH. See RSCH for details.	Rings	Wave	Hx, Hy
RWCH	RMS wavefront error with respect to the chief ray. The units are waves. This RMS really refers to the standard deviation of the wavefront, since the mean OPD is subtracted out. See RWCE. See RSCH for details.	Rings	Wave	Hx, Hy
RWRE	Like RSRE, except for wavefront error rather than spot size.	Grid	Wave	Hx, Hy
RWRH	Like RSRH, except for wavefront error rather than spot size.	Grid	Wave	Hx, Hy
SAGX	The sag in lens units of the surface "Surf" in the XZ plane at the semi-diameter distance.	Surf	-	-
SAGY	The sag in lens units of the surface "Surf" in the YZ plane at the semi-diameter distance.	Surf	-	-
SFNO	Sagittal working F/#, computed at any defined field and wavelength. See TFNO.	Field	Wave	-
SINE	Sine of the value of the specified operand number. If flag is 0, then the units are radians, otherwise, degrees.	Op #	Flag	-
SKIS	Skip if symmetric. If the lens is rotationally symmetric, then computation of the merit function continues at the specified operand number.	Op #		
SKIN	Skip if not symmetric. See SKIS.	Op #		
SPHA	Spherical aberration in waves contributed by the specified surface. If the surface value is zero, the sum for the entire system is used.	Surf	Wave	-
SQRT	Square root of operand. The argument is the row number of the operand value to take the square root of.	Op #	-	-
SUMM	Sum of two operands. The two arguments are the row numbers of the operands to add. See OSUM.	Op #1	Op #2	-
SVIG	Sets the vignetting factors. When included, updates the vignetting factors for the current configuration.	-	-	-
TANG	Tangent of the value of the specified operand number. If flag is 0, then the units are radians, otherwise, degrees.	Op #	Flag	-
TFNO	Tangential working F/#, computed at any defined field and wavelength. See SFNO.	Field	Wave	-

NAME	Description	Int1	Int2	Hxy, Pxy
TMAS	Total mass. Computes the mass of the glass lenses within the specified range of surfaces. Only valid for plane and spherical standard surfaces with circular edges. See the chapter "Reports Menu" for a discussion of how element masses and volumes are computed.	First surf	Last surf	-
TOTR	Total track (length) of lens in lens units.	-	-	-
TRAC	Transverse aberration radial direction measured in the image plane with respect to the centroid. Unlike most other operands, TRAC critically depends upon the placement of other TRAC operands within the Merit Function Editor to work correctly. TRAC operands must be grouped together by field position and wavelength. ZEMAX traces all TRAC rays with a common field point together, and then uses the collective data to compute the centroid of all the rays. Each ray individually is then referenced to the computed centroid. This operand should only be entered into the Merit Function Editor by the Default Merit Function tool, and is not recommended for use directly by the user.	-	Wave	Yes
TRAD	The x component of the TRAR only. TRAD has the same restrictions that TRAC does; see TRAC for a detailed discussion.		Wave	Yes
TRAE	The y component of the TRAR only. TRAE has the same restrictions that TRAC does; see TRAC for a detailed discussion.		Wave	Yes
TRAI	Transverse aberration radius measured at the specified surface with respect to the chief ray. Similar to TRAR, except a surface other than the image surface may be specified.	Surf	Wave	Yes
TRAR	Transverse aberration radial direction measured in the image plane with respect to the chief ray. See ANAR.	-	Wave	Yes
TRAX	Transverse aberration x direction measured in the image plane with respect to the chief ray.	-	Wave	Yes
TRAY	Transverse aberration y direction measured in the image plane with respect to the chief ray.	-	Wave	Yes
TRCX	Transverse aberration x direction measured in the image plane with respect to the centroid. See TRAC. This operand should only be entered into the Merit Function Editor by the Default Merit Function tool, and is not recommended for use directly by the user.	-	Wave	Yes
TRCY	Transverse aberration y direction measured in the image plane with respect to the centroid. See TRAC. This operand should only be entered into the Merit Function Editor by the Default Merit Function tool, and is not recommended for use directly by the user.	-	Wave	Yes
TTGT	Total thickness greater than. This boundary operand constrains the total thickness (including both front and back surface sags) of surface "Surf" to be greater than the specified target value. The thickness is calculated at the semi-diameter radius along the +y axis if code is zero, the +x axis if code is 1, the -y axis if code is 2, and the -x axis if code is 3. This operand automatically changes the sign on thicknesses in mirror spaces to always yield a positive value for physically possible lenses. See TTLT and TTVA.	Surf	Code	-
TTHI	Sum of thicknesses from first to last specified surface. Note that the sum is inclusive, it is not the thickness between the two surfaces.	First surf	Last surf	-

NAME	Description	Int1	Int2	Hxy, Pxy
		Surf	Code	-
TTLT	Total thickness less than. See TTGT.	Surf	Code	-
TTVA	Total thickness value. See TTGT.	Macro #	Data#	Yes
UDOP	User defined operand. Used for optimizing numerical results computed in externally compiled programs. See the section "User defined operands". See also ZPLM.	-	-	-
USYM	If present in the merit function, this operand instructs ZEMAX to assume radial symmetry exists in the lens even if ZEMAX detects symmetry does not exist. This speeds execution of the merit function in some special cases. See "Assume Axial Symmetry" in the default merit function description earlier in this chapter.	-	-	-
VOLU	Volume of element(s) in cubic cm. The resulting value is only valid for regions bounded by plane and spherical surfaces. Total volume. Computes the volume of the lenses and air spaces within the specified range of surfaces. Only valid for plane and spherical standard surfaces with circular edges. See the chapter "Reports Menu" for a discussion of how element masses and volumes are computed.	First surf	Last surf	-
WFNO	Working F/#. This is computed from the angle the real marginal ray makes with respect to the chief ray in image space.	-	-	-
XDGT	Extra data value greater than. The Int2 number must be between 1 and 200 to indicate which of the extra data values is selected.	Surf	Number	-
XDLT	Extra data value less. The Int2 number must be between 1 and 200 to indicate which of the extra data values is selected.	Surf	Number	-
XDVA	Extra data value. The Int2 number must be between 1 and 200 to indicate which of the extra data values is selected.	Surf	Number	-
XENC	Extended source encircled energy. This operand computes the radius in microns to the specified fraction of extended source geometric encircled energy, using whatever the current default settings are. To use this operand, first define the settings on the extended source encircled energy feature as desired, then press Save on the settings box. The only setting that is overwritten is the type, which is 1 for encircled, 2 for x only, 3 for y only, and 4 for ensquared. Hx is the fraction of energy desired, and must be between zero and 1, exclusive. See also DENC and GENC.	Type	-	See left
XNEA	Minimum edge thickness for air surfaces. This operand checks the edge thickness at numerous points around the perimeter of the surface, and insures all points are at least the minimum specified thickness. See "MNEA".	First surf	Last surf	-
XNEG	Minimum edge thickness for glass surfaces. This operand checks the edge thickness at numerous points around the perimeter of the surface, and insures all points are at least the minimum specified thickness. See "MNEG".	First surf	Last surf	-
XNET	Minimum edge thickness. This operand checks the edge thickness at numerous points around the perimeter of the surface, and insures all points are at least the minimum specified thickness. See "MNET".	First surf	Last surf	-

NAME
XXEA
XXEG
XXET
YNIP
ZERN
ZPLM
ZTHI

The operati
CVLT, CTG
in the secti
Because o
(more) a
units. How
the weight
variable
then yield
target valu
After maki
ydata. T
bestest c

NAME	Description	Int1	Int2	Hxy, Pxy
XXEA	Maximum edge thickness for air surfaces. This operand checks the edge thickness at numerous points around the perimeter of the surface, and insures all points are no more than the maximum specified thickness. See "MXEA".	First surf	Last surf	-
XXEG	Maximum edge thickness for glass surfaces. This operand checks the edge thickness at numerous points around the perimeter of the surface, and insures all points are no more than the maximum specified thickness. See "MXEG".	First surf	Last surf	-
XXET	Maximum edge thickness. This operand checks the edge thickness at numerous points around the perimeter of the surface, and insures all points are no more than the maximum specified thickness. See "MXET".	First surf	Last surf	-
YNIP	YNI-paraxial. This number is the product of the paraxial marginal ray height times the index times the angle of incidence. This quantity is proportional to the narcissus contribution of the specified surface. See Applied Optics, Vol. 21, 18, p3393.	Surf	Wave	-
ZERN	Zernike Fringe coefficient. The Int1, Int2, Hx, and Hy data values are used to specify the Zernike term number (1-37), the wavelength number, the sampling density (1 = 32 x 32, 2 = 64 x 64, etc.), and the field position, respectively. Note that if you use multiple ZERN operands which only differ in the term number, they should be placed on adjacent lines in the editor; otherwise, the computation is slower.	Term	Wave	See left
ZPLM	Used for optimizing numerical results computed in ZPL macros. See the section "User defined operands". See also UDOP.	Macro #	Data#	Yes
ZTHI	This operand controls the variation in the total thickness of a range of surfaces over multiple configurations. It is similar to the TTHI operand, except it is an inequality operator. The target value specified is the maximum allowed difference between the TTHI at each defined configuration. For example, if there are 3 configurations where TTHI 3 8 would evaluate to 17, 19, and 18.5, respectively, ZTHI will return 2 (i.e. 19-17) if the target is less than 2. Otherwise, ZTHI returns the target value. To keep all zoom configurations the same length, use a target of 0.	First surf	Last surf	-

The operational operands (SUMM, OSUM, DIFF, PROD, DIVI, SQRT) along with the parametric operands (CVGT, CVLT, CTGT, CTLT, etc..) can be used to define very general and complex optimization operands, as discussed in the section "Defining complex operands", which can be found later in this chapter.

Because of the dimensional differences between parameters such as effective focal length (tens of millimeters or more) and RMS spot size (microns), usually a weighting of one is sufficient for quantities measured in lens units. However, the residual value of the effective focal length with this weighting is not likely to be zero. Increasing the weighting will bring the resulting system closer to the desired effective focal length. This effect is often noticeable when defining ETGT (edge thickness greater than) operands. Usually ETGT with a target of zero will often yield a value just slightly less than zero. Rather than increase the weight, it is much simpler to provide a target value of .1, or some such number.

After making changes to the operand list, the current values of each operand can be updated by selecting Tools, Update. This is also useful for checking to see what the current values of each operand are, and which has the greatest contribution to the merit function. The percent contribution is defined as

$$\% \text{ contrib}_i = 100 \times \frac{W_i(V_i - T_i)^2}{\sum_j W_j(V_j - T_j)^2}$$

where the index j indicates the sum over all operands.
The merit function is automatically saved with the lens file.

Understanding boundary operands

The boundary operands such as MNCT, CTGT, DIMX, and others behave somewhat differently than specific target operands such as TRAR and REAY. When you specify a boundary on a parameter, you specify the target value as the definition of the boundary. For example, to maintain a minimum center thickness on surface 5 of 10 mm, you might use a command such as CTGT 5 10 (where the 5 is in the Int1 column and 10 is in the target column). If you update the merit function and then observe the "value" column of that operand, there are two possibilities for the value: 1) if the boundary is violated, that is, the center thickness is less than 10, then the actual value of the thickness will be displayed, or 2) if the boundary is not violated, that is, the center thickness is greater than 10, then the value 10 will be displayed.

The rule is simple- if the boundary is violated, the actual value is shown; if it is not violated, the value is set to the target and is therefore ignored by the optimization algorithm. If during optimization the boundary becomes violated, then the value will automatically be updated and the optimization algorithm will attempt to correct the offending parameter.

The boundary operands which constrain a range of surfaces are slightly more complicated. These multiple surface operands return values which represent the total effect of all violated boundaries within the specified surface range. For example, the operand MNCT 1 10 will constrain the minimum center thickness of surfaces 1 through 10. If the target is 3.0, which defines the boundary, then the difference between the value of the operand and the target is the sum of the difference between 3.0 and the thicknesses of all surfaces between 1 and 10 whose center thickness is less than 3.0. If only one surface in the range has a center thickness less than 3.0, say 2.2, then the operand has a value of 2.5. If a second surface is added that has a thickness less than 3.0, say 2.2, then the operand will have a value of 1.7 (2.5 minus .8; the .8 is 3.0 - 2.2). The total difference between the target of the operand and the value is 3.0 - 1.7 or 1.3. This difference of 1.3 is due to the violation of 0.5 by the first surface and another 0.8 by the second surface.

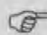
If calculating the value of these boundary operands seems confusing, don't worry; ZEMAX does all the calculations for you. All you need to do is to specify the boundary type (such as MNCT or MNET) the boundary range (surface 1 through 10, or whatever) and the desired value (3 mm or whatever). If all the boundary constraints are met, then the operand value is equal to the target, otherwise, the value will be different and the merit function will increase. The increased merit function will cause the optimization algorithm to seek a reduction of the operand contribution.

If a boundary operand does not seem to work, there are several things to check:

- 1) Make sure the variables you have defined can have some effect on the boundary operands. A common mistake is to specify MNCT and have some "frozen" thickness within the surface range. If the thickness violates the boundary and it is not variable, ZEMAX can't fix it. The operands DO NOT ignore violated but frozen boundaries.
- 2) If there is a small residual error, try increasing the boundary value. For example, if MNCT is used with a target of 0.0, and the value is a small number (like -.001) the problem is not that the operand doesn't work, it is simply that the residual error is too small to increase the merit function significantly. It is usually better to increase the target to 0.1, or some other number, rather than to increase the weight. Increasing the weight will only lead to a smaller violation (like -.0000001) rather than meeting the boundary.
- 3) Check to see if there is a reasonable contribution to the merit function. You can easily check this with the percent contribution column. By looking at the percent contribution column, verify that the operand in question has enough influence on the total merit function. If it does not, increase the weight, or see the preceding paragraph for advice on changing the target.

Understanding the boundary operands is a crucial part of mastering ZEMAX optimization, and with a little practice you will find them to offer excellent control and flexibility.

Using MTF operands

 **These operands are only supported in the XE and EE editions of ZEMAX.**

The MTF operands such as MTFT, MTF5, and MTF6 provide a capability to directly optimize the diffraction MTF. This is a powerful capability, however, using the MTF operands requires some care on the part of the user.

For systems not close to the diffraction limit, geometric equivalent MTF operands are provided: GMTT, GMTS, and GMTA. These should be used in place of the diffraction operands for systems with greater than about 2-5 waves of aberration.

The MTF operands calculate the complete diffraction or geometric MTF exactly like the plots available on the Analysis, Diffraction menu option. Therefore, any optical system for which the MTF plot produces invalid data (due to excessive OPD in the pupil, see the chapter "Analysis" for details) will also produce meaningless data during optimization. For example, it is unrealistic to optimize a lens for MTF starting from plane parallel plates, because of course the MTF generally cannot be accurately computed for such a system. Also, MTF optimization is considerably slower than RMS spot size or RMS wavefront error optimization, typically 5 to 50 times slower. Note that if you use both a MTFT and a MTFS operand for the same field and wavelength data, they should be placed on adjacent lines in the editor; otherwise, the MTF is computed twice. If the sampling is too low for accurate computation of the MTF, then the MTF operands return zero rather than a meaningless number.

The slow execution speed may be noticeable when updating the Merit Function Editor display, and entering and exiting the optimization dialog box. At these times the merit function is updated. If several MTF operands have been entered, ZEMAX may take several minutes on a slow computer just to refresh the screen.

A good approach is to interactively design your system using RMS wavefront error. Generally speaking, systems with low RMS wavefront error will have reasonable MTF performance. After the design is very nearly in the final form, then try switching over to MTF optimization for a "touch-up". It is also a good idea to completely erase the default merit function before entering any MTF operands, except of course for any boundary constraints needed (such as MNCT).

The MTF operands use the operand data columns such as Int1, Int2, and Hx differently than most operands. The Int1 column determines the sampling density to be used. A value of 1 specifies a 32 x 32 grid should be used, 2 specifies a 64 x 64 grid, and so on. Use the smallest grid size for which the data are computed accurately; again see the "Analysis" chapter for details.

The Int2 column selects the wavelength, like most operands. However, a value of zero may be entered to specify a polychromatic calculation. In this case, the wavelength-weighted polychromatic MTF is computed. Obviously, this is slower than a monochromatic MTF calculation.

The Hx column is used to specify the field position, and it must be an integer between 1 and the number of defined fields.

The Hy column is the spatial frequency in cycles per millimeter, independent of the lens system units. Any value may be entered; the operand will return 0 if it is past the incoherent cut-off frequency. The value does not need to be an integral multiple of some fraction; the exact value of the MTF is computed using a cubic spline fit to the neighboring data points, just like the MTF data plots.

The target and weight columns are used just like any other operand. Of course, if the target is set to 1.0, and the frequency is non-zero, then the value of the operand can never reach the target.

Performing an optimization

To begin optimization, choose Tools, Optimization from the main menu bar. The optimization control dialog box will appear with the following options.

OPTIMIZATION OPTIONS

Item	Description
Automatic	Executes until ZEMAX determines the system is no longer improving significantly.
1 Cycle	Executes a single optimization cycle.
5 Cycles	Executes 5 optimization cycles.
10 Cycles	Executes 10 optimization cycles.
50 Cycles	Executes 50 optimization cycles.
Inf. Cycles	Executes optimization cycles in an infinite, continuous loop until "Terminate" is pressed.
Terminate	Terminates a running optimization, and returns control back to the dialog box.
Exit	Closes the optimization dialog box.

Auto Update	If checked, ZEMAX will automatically update and redraw all open windows at the end of each optimization cycle. This allows monitoring of the optimization progress using any analysis feature.
# CPU's	Selects the number of CPU's over which to spread the optimization task. More than 1 may be selected, even on a single CPU computer, in which case the single CPU will time share the multiple simultaneous tasks. The default is the number of processors detected by the operating system.

Selecting automatic will cause the optimizer to run until no progress is being made. The other options will run the specified number of cycles. Automatic mode is highly recommended. The time required to run a given optimization cycle varies enormously with the number of variables, the complexity of the system, the number of solves, the number of operands, and of course the computer speed. If the cycle is taking too long, or if it appears to be hung up, or if you feel the design is not progressing adequately, click on Terminate to end the optimization run.

When the optimization begins, ZEMAX first updates the system merit function. If any of the operands cannot be computed, the optimization cannot begin, and an error message will be displayed. Operands cannot be computed if they require the tracing of rays which miss surfaces or which undergo total internal reflection (TIR) at an index boundary. If such an error message appears, usually the starting lens prescription is in error, or the ray targets are incorrectly defined (this will not happen with the default merit functions, but might happen with user defined rays). ZEMAX can automatically recover if the merit function cannot be evaluated during the course of optimization; only the starting system need be adequate to compute all operands in the merit function.

Defining complex operands

Although the default merit function, coupled with a few predefined operands, is perfectly suitable for the majority of optical designs, there are times when an unusual constraint needs to be added to the merit function. Rather than define a very large number of very specific operands, ZEMAX allows you to build your own operands out of simple building blocks.

ZEMAX allows very general operand definitions. There are two tricks to creating these operands. First, use certain operands with zero-weighting to define the parameters you need, and second, use the operational operands to define relationships between them. For example, suppose you require that the thickness of surface 3 and the thickness of surface 4 sum to 10. There is an operand that does this, TTHI. The command structure would look like this:

Number	Type	Int1	Int2	Target	Weight
1	TTHI	3	4	10	1

However, for illustration only, note that there is an alternate way of calculating the same thing:

Number	Type	Int1	Int2	Target	Weight
1	CTVA	3		0	0
2	CTVA	4		0	0
3	SUMM	1	2	10	1

Operand 1 uses the Center Thickness Value (CTVA) command to extract the value of the thickness of surface 3. Similarly, operand 2 is used to extract the thickness of surface 4. The zero weighting on both operands insures that the optimization algorithm ignores the constraint; it is only used as an intermediate step. Operand 3 now sums two operands: number 1 and number 2. The result is the sum of the thicknesses of surfaces 3 and 4 is the value of operand 3, and this has a non-zero weight. The optimization algorithm will attempt to drive the sum to 10. Why go to all the trouble of this three-step process if a single TTHI command would do the same thing? The reason is that this approach can be extended to develop very general operands. For example, suppose you

wanted the radius of curvature of surface 5 to be centered on the vertex of surface 8. Study the following commands to see if you understand how this is done:

Number	Type	Int1	Int2	Target	Weight
1	CVVA	5			
2	TTHI	5	7	0	0
3	PROD	1	2	1	1

The CVVA command extracts the curvature of surface 5, the curvature we want to control. TTHI 5 7 calculates the distance from surface 5 to surface 8 (note we only sum to surface 7 to get to surface 8, since the thickness of surface 8 gives the distance to surface 9). Since the curvature of the surface is the reciprocal of the radius, the product of the curvature and the distance must be one; hence the target is 1 for operand 3. Operand 3 is also the only weighted operand in the sequence.

Now consider the requirement that the thickness of surface 5 must be greater than twice the radius of curvature of surface 4 plus the conic constant of surface 2 (this is nonsensical, but illustrative of the flexibility in the approach):

Number	Type	Int1	Int2	Target	Weight
1	CTVA	5		0	0
2	CVVA	4		0	0
3	CONS			2	0
4	DIVI	3	2	0	0
5	COVA	2		0	0
6	SUMM	4	5	0	0
7	DIFF	1	6	0	0
8	OPGT	7		0	1

Operand 1 extracts the (center) thickness of surface 5. Operand 2 extracts the value of the curvature of surface 4. Operand 3 sets a constant of two, and operand 4 divides the value 2 by the curvature (yielding twice the radius of curvature). COVA extracts the conic, and SUMM adds operands 2 and 4. Operand 7 takes the difference of the thickness and twice the radius plus the conic. Since we want the former to exceed the latter, we set an operand greater than constraint; the only one to have a non-zero weighting.

Optimizing glass selection

Optimization of glasses is handled somewhat differently than other data. Optimizing the glass choice directly is a difficult and unpredictable process because there does not exist a continuum of glasses on the glass map. There are two methods for dealing with this problem: by using model glasses or by using glass substitution. Glass substitution is usually far superior, but is only supported in the XE and EE editions of ZEMAX.

Using model glasses

The model glass method is to idealize the glass dispersion using a few simple numerical parameters, and then optimize these parameters while constraining either the parameter values or the computed index values to be similar to available glasses. This is the "model" glass method. Model glasses are described in detail in the chapter "Using Glass Catalogs". One disadvantage of the model glass method is that the optimized parameters and resulting index values may not correspond to any physically existing glass. Another disadvantage is that model glasses are only sufficiently accurate in the visible spectrum. This method is used by the conventional optimizer described in this chapter.

To optimize glasses then requires several steps. First, change the glass of the appropriate surface to a "Model" glass using the glass solve dialog box in the Lens Data Editor. For information on model glasses, see the Chapter

"Using Glass Catalogs". When you change the glass from "Fixed" to "Model", ZEMAX will make a suitable guess for the index, Abbe number, and partial dispersion; you only need to change the values if you wish. The three values can be made variable by clicking on the "Vary" box next to each entry.

The Ctrl-Z shortcut on the glass column will also work; it makes the index, Abbe, and partial dispersion variable automatically. The model glass data values can now be optimized using the optimization feature in the usual way.

Unconstrained glass optimization usually will lead to very high index materials being selected. This is because surfaces with high refractivity (a large difference in index across the boundary) need less curvature than low refractivity surfaces to have the same optical power. Lower curvature surfaces introduce less aberration.

Unfortunately, high index materials are expensive, heavy, harder to fabricate, and may be brittle, delicate, or susceptible to stains and scratches. Also, very high index materials do not always exist; there are few glasses (for the visible spectrum) available with an N_d higher than about 1.9. The V_d value also is limited to the range of roughly 20 to 80. Therefore, it is essential to limit the N_d and V_d values to reasonable ranges during optimization. The partial dispersion deviation also must be limited in range.

There are two ways to limit the N_d , V_d , and $\Delta P_{g,F}$ values. The simplest way is to add the RGLA operand somewhere in the operand list. The RGLA operand measures the "distance" on the glass map between the index, Abbe number, and partial dispersion of the model glass to the closest glass in the currently loaded catalogs. For example, if you are optimizing the index and Abbe and you have specified that the Schott and Hoya catalogs are used (these are specified on the general data window), the RGLA operand computes the "distance" to each glass in these catalogs. If the smallest "distance" is less than the target value specified for the RGLA operand, then the boundary condition is met, and the value of the operand is equal to the target. If the closest glass is farther than the target value, then the RGLA value is the actual "distance". The "distance" is defined by the square root of the weighted sum of the squares of the difference between the index, Abbe, and partial dispersion terms for two glasses. The "distance" between any two glasses is given by

$$d = [W_n(N_{d1} - N_{d2})^2 + W_a(V_{d1} - V_{d2})^2 + W_p(\Delta P_{g,F1} - \Delta P_{g,F2})^2]^{\frac{1}{2}}$$

where the factors W_n , W_a , and W_p weight the various terms. The weighting factors may be user defined on the RGLA operand parameter list, or if left at zero, will default to 1.0, 1E-04, and 1E+02, respectively.

The best way to use RGLA is to specify the surface range that covers all of the surfaces you are optimizing. For a target value, start off with 0.05. This will allow the glasses to easily move all over the glass map, since the spacing between various glasses is usually less than 0.05. After optimization, decrease the target to roughly 0.02 and reoptimize. This will encourage the optimized system to choose index and Abbe numbers reasonably near actual glasses.

The other method for constraining index and Abbe values is to use the MNIN, MXIN, MNAB, and MXAB controls. These operands are mnemonics for Minimum and Maximum Index and Abbe values, and they are documented in the preceding tables. These operands can be used to restrict the optimization to specific rectangles on the glass map. It may be useful to use RGLA along with MXIN, for example, to restrict the glass selection to existing glasses with an index lower than some value.

At some point you will want to convert your variable index data back to a real glass. There will usually not be a perfect match between the optimized N_d and Abbe values and those of an actual glass in the current catalog. However, ZEMAX will search through the catalog and find the "best fit" glass using a least-squares criterion similar to the RGLA definition above (the partial dispersion term is omitted). The glass in the catalog which differs the least from the variable index parameters is the glass selected. This glass is also reported on the "Surface Data Summary" feature (select Reports, Surface Data). The index of refraction data shown is that calculated from the N_d and Abbe values, not the best fit glass. After converting from a model glass to a real glass, another optimization run is generally required. For systems with delicate chromatic aberration balancing, the best glass choice might never be found using variable glasses, simply because the model glass dispersion is never identical to the dispersion of a real glass.

Using glass substitution


The glass substitution method is to directly alter the glass types, and then reoptimize to see if the new glasses yield a better solution. This method can be used manually, by simply changing the glass type and then reoptimizing, or the process can be automated using the global optimization techniques described in the next chapter, "Global Optimization". The global optimization method uses only actual glass catalog materials, and in this sense is the superior method.

See the discussion "Optimizing Glass Selection" in the next chapter, "Global Optimization" for a discussion on this method for optimizing glass selection.

Optimizing zoom and multi-configuration lenses

Optimizing zoom lenses is virtually identical to optimizing conventional single-configuration lenses. See the Chapter "Multi-Configurations" for details.

Optimizing extra data

 *This discussion is only relevant for users of ZEMAX-EE.*

Certain surface types supported by ZEMAX-EE, such as the Zernike, Zernike phase, extended polynomial, and binary optic surface use the extra data values. These extra data values can be edited, loaded from ASCII files, and may be used as variables for optimization. For details on the editor, see the chapter "Editors menu".

To make an extra data value a variable, open up the Extra Data Editor. When the extra data editor appears, move the cursor to the row and column with the value you want to optimize, and press Ctrl-Z (the same command used to set variables on the main screen). The variable will now be optimized when the optimization is run.

There are also several boundary constraints for use with the extra data values. XDVA, XDGT, and XDLT are extra data value, greater than, or less than, respectively. The Int1 column on the spreadsheet display indicates the surface number to which the operand applies, and the Int2 column is used to specify which of the extra data values is to be used.

Optimizing objects in a non-sequential group with sequential rays

Optimizing variables within a non-sequential group is fundamentally no different from optimizing other numerical parameters. Variables are set in the same way as for parameters in the Lens Data Editor.

The difficulty is optimizing non-sequential object properties is the unpredictable way in which rays may (or may not) propagate through a non-sequential group. For non-sequential objects such as prisms, usually a small change in position or size of the prism does not dramatically affect the ray path. However, for objects such as light pipes, a small change in the object definition can dramatically affect the ray path. Rays that once propagated through an object may miss the object completely if the object position or angle changes slightly. This usually causes severe errors in the computation of derivatives, and the optimization performs either poorly or not at all.

Another problem with some non-sequential systems is the exit pupil may not be a reasonable image of the entrance pupil. For this reason, rectangular array rather than Gaussian Quadrature should be used if the system is a non-imaging system that does not form an image of the entrance pupil at the exit pupil.

For these systems, optimization may proceed more effectively using the global optimization algorithms, which do not rely exclusively on derivative computation.

Optimizing with sources and detectors in non-sequential mode

Optimizing illumination systems or other optical systems that use non-sequential sources and detectors is supported using the NSDD and NSTR operands.

A typical merit function would consist of three groups of operands:

First, NSDD operands would be used to clear the data in the current detectors. NSDD 0 0 clears all energy in all detectors. Usually a single NSDD 0 0 at the top of the merit function is all that is needed. NSDD returns a value of zero and has no effect on the merit function value when used to clear detectors.

Second, NSTR operands are used to trace rays from NSC sources. NSTR *i* traces analysis rays from source *i*; NSTR 0 traces all analysis rays from all sources. Note the number of analysis rays on the NSC editor determines how many rays are traced... and how long the evaluation of the NSTR operand will take. NSTR always returns a value of zero and has no effect on the merit function value.

Third, a new group of NSDD operands are used to read out the detector data. NSDD has four arguments: surface, detector, pixel, and data. Surface is the surface number of the NSC group (use 1 if the program mode is non-sequential). Detector is the object number of the detector. Both detector objects and faceted detectors may be used as detectors. If pixel is zero, the sum of all the flux in the detector is returned. If pixel is an integer greater than zero, the flux, flux/area, or flux/solid angle is returned for that pixel. Which of the three is determined by the data argument, which should be 0, 1, or 2 for flux, irradiance, or intensity, respectively. The

units of the returned data is determined by the system units, see "Irradiance/Illuminance Units" on page 59. If the object is a faceted detector, only data options 0 and 1 are supported.

The practical difficulty in optimizing these systems is the difficulty of computing derivatives of detected energy with respect to variable parameters because of the relatively large uncertainty in computing detected energy. Many rays must be traced to determine illumination patterns approximately.

Optimizing with the IMAE operand

The IMAE operand estimates the efficiency of an optical system by launching many rays into the entrance pupil, computing the fraction of rays that pass through all surface apertures to any surface. Optimization with this operand may not proceed smoothly if only hard-edged surface apertures, such as the circular aperture, are used. This is because ZEMAX estimates the derivative of operand values by making very small differential changes in the value of each variable, then computes a finite difference of the operand value. For the IMAE operand, a small change in the value of a variable may not change the efficiency estimate; if no rays are close enough to an aperture to change from being vignettted to unvignettted or vice-a-versa.

The solution is to replace hard edged apertures with soft-edged apertures placed on a user-defined surface. A soft-edged aperture has a transmission that is unity over most of the clear aperture, but near the edge the transmission drops to zero gradually over a small region, rather than abruptly.

Filter functions for doing this are included with ZEMAX as sample DLL files; see the Chapter "Surface Types" under "User Defined Surface" for details. See in particular the discussion of the US_FILTER4.DLL sample.

IMAE uses the current saved settings for the image analysis feature; except for "show" which is always set to spot diagram for this computation. See "Geometric Image Analysis" on page 104.

Using gradient index operands

There are several optimization operands which are used to control the properties of gradient index materials during optimization. Some of them are described below.

DLTN

DLTN is used to control the maximum total change in index within a gradient index lens. Int1 is used to define the surface number, and Int2 is used to define the wavelength number. DLTN is defined as:

$$DLTN = n_{max} - n_{min}$$

The min and max index values are computed at the extreme z coordinates, Z min and Z max. Z min and Z max are the Z coordinates of the minimum and maximum axial positions of the blank used to make the lens, before the shaping begins. For a convex surface, they correspond to the vertex. For a concave surface, they correspond to the maximum sag at that surface.


LPTD

LPTD is used to control the profile of the gradient within the material. Only the Int1 value is used to define the surface number of the gradient index surface. LPTD is an acronym for LightPath Technology Delta, and the constraint is used to keep a nonlinear profile monotonically increasing or decreasing. It only needs to be used when the quadratic or cubic term of the axial gradient is variable. This operand only affects GRIN 5 surface types. The LPTD operand should be used with a target of 0. The boundary constraint enforces the following conditions:

$$\frac{\partial n}{\partial z_{min}} > 0 \text{ and } \frac{\partial n}{\partial z_{max}} > 0, \text{ or } \frac{\partial n}{\partial z_{min}} < 0 \text{ and } \frac{\partial n}{\partial z_{max}} < 0.$$

Z min and Z max are the Z coordinates of the minimum and maximum axial positions of the blank used to make the lens, before the shaping begins. For a convex surface, they correspond to the vertex. For a concave surface, they correspond to the maximum sag at that surface. If the residual value of the operand is less than zero, then the target may be decreased slightly (try 0.1). Changing the target is usually more effective than increasing the weight. The value of the LPTD operand must be zero for the blank to be fabricated. Always check the gradient profile to make sure the slope does not change sign.

User defined operands

 This discussion is only relevant for users of ZEMAX-EE.

There are times when very complex calculations need to be performed, and the results of the computation need to be optimized. ZEMAX supports some of these calculations already, such as the MTF operand which traces many rays, computes the MTF, and then returns a single resulting number to the Merit Function Editor "value" column. Some limited calculations can be performed within the merit function itself; see for example the discussion in the section "Defining complex operands" earlier in this chapter.

However, there are problems for which only the flexibility of a user defined program is sufficient for defining the data computed by an operand. There are two ways of achieving this:

- 1) Through the use of a ZPL macro
- 2) Through the use of an externally defined and compiled program

The use of ZPL macros is simpler, well integrated with ZEMAX, and requires very little programming experience. However, it is limited by the capabilities of the ZPL macro language, and ZPL macros are interpreted, which means slow execution for complex calculations. ZPL macro optimization is generally a better choice for simpler macros that execute fairly quickly.

Externally defined programs are more complex to program, require an external C or other language compiler, and at least some programming experience. However, externally defined programs can be vastly more complex than what is supported by the ZPL macro language, and since external programs are compiled, they run significantly faster. The speed difference can be dramatic, and generally more complex calculations will benefit more from being externally compiled. Indeed, externally defined operand programs may be very complex, tracing millions of rays or doing other lengthy calculations before returning control to ZEMAX. Note this interface can be used to optimize lenses based upon data computed by other analysis programs, such as a stray light analysis program.

Both the ZPL and the externally compiled methods of implementing UDO's are described in detail below.

Optimizing with ZPL macros

If the ZPL macro language is sufficient to perform the required computations, then the operand ZPLM may be used to call a ZPL macro from within the merit function. The macro performs the required computations, then returns the result using the ZPL OPTRETURN keyword.

ZPLM is simple to use. The Int1 and Int2 values are used to specify the macro number and data field number, respectively. The macro number is used to indicate which ZPL macro should be executed, while the data field number indicates which value computed by the macro should be optimized.

The macro number must be an integer between 0 and 99. If the Int1 value for a ZPLM operand is set to 17, for example, then the macro number is 17, and the macro to be executed must be named ZPL17.ZPL. The macro name must always use a two digit representation of the macro number. If the macro number was 6, then the macro to be executed would be ZPL06.ZPL. The ZPL macro file must reside in the default directory for ZPL macros; see the chapter "File Menu" for details.

The data field number may be any number between 0 and 50, inclusive. This number refers to a position in a global array associated with the lens in memory. During execution of the macro, the macro keyword OPTRETURN specifies which data field number stores the results of the macro calculation. There are 51 different data fields, so that a single macro call can be used to optimize up to 51 different values simultaneously. For example, suppose you needed a macro which computed the total length of the lens from surface 1 to the image plane (this is in effect a user-defined version of the TOTR operand). The macro might look like this:

```
n = nsur()
x = 0
for i = 1, n, 1
  x = x + thic(i)
next
optreturn 0 = x
```

Note the use of the OPTRETURN keyword. This keyword stores the resulting value for "x" in the global array position 0. Suppose this macro was named ZPL15.ZPL. To optimize the resulting value for x, the ZPLM merit function operand would be added to the Merit Function Editor, with Int1 = 15 and Int2 = 0. After updating the merit function, the "value" would be the same as that returned by TOTR, and it can be optimized in the same way.

ZPLM also permits the use of the Hx, Hy, Px, and Py data fields. These data fields can be read by the ZPL macro using the PVHX, PVHY, PVPX, and PVPY ZPL functions, respectively. "PV" is a mnemonic for "Pass Value". There is one very important thing to know about the data field number. If it is zero, then the macro is executed and the value from OPTRETURN 0 is returned. However, if the data field number is not zero, then the macro is not executed, but any previous value stored from an earlier call to the macro is used instead. The advantage to this convention is substantial. If the macro computes many values, all of which need to be optimized, the macro only needs to be called once, yet multiple ZPLM operands can access the data. This is much more efficient than calling the macro multiple times.

For example, suppose a macro named ZPL11.ZPL computes three values, all of which require optimization. In the macro, the values are stored using OPTRETURN:

```
OPTRETURN 0 = x
OPTRETURN 1 = y
OPTRETURN 2 = z
```

Then three ZPLM operands in the merit function can extract the data and perform the optimization with a single call to the macro:

```
ZPLM 11 0
ZPLM 11 1
ZPLM 11 2
```

The macro ZPL11.ZPL is only called during the evaluation of the ZPLM 11 0 operand. Note the Hx, Hy, Px, and Py values can only be used if the Int2 value is zero, since only in this case is the macro evaluated.

Lastly, it is very important that no changes be made to the lens data during the macro execution. These changes may interfere with the subsequent evaluation of other operands. ZEMAX does not restore the lens being evaluated to the state it was in prior to the evaluation of the ZPLM specified macro. Also, ZPLM should not be used in the middle of a default merit function, but should instead be placed either prior to or after the portion of the merit function that ZEMAX defined by default. If lens data is changed during the operation of the macro, ZEMAX has no way of knowing what data was altered, and cannot restore the lens to its original unaltered state.

This could be avoided by allowing the ZPL macro to execute only on a copy of the lens being optimized, rather than the actual lens, however this is not currently supported. The reason is there may be times where a macro needs to alter the lens data prior to evaluation of subsequent operands. In this case, two macros should be executed. The first should modify the data as required, and the second should restore the data to the original condition. Both macros can be listed in the merit function editor, with the intervening operands executing on the altered lens data.

Optimizing with externally compiled programs

The second method of creating a user defined operand (UDO) is to write an external Windows program which computes the data, then use Dynamic Data Exchange (DDE) to pass the data to and from ZEMAX. The DDE interface in ZEMAX is documented and described in the chapter "ZEMAX Extensions". The material presented there is not duplicated here; this discussion assumes the material in that chapter is understood.

The operand UDOP is used to call an external client program from within the merit function. The client program performs the required computations, possibly by making multiple DDE calls back to the ZEMAX server, then returns the result to ZEMAX using the DDE interface. The computed data is then placed in the "value" column of the Merit Function Editor and thus may be optimized in the usual way.

UDOP is simple to use. The Int1 and Int2 values are used to specify the client program number and data field number, respectively. The client program number is used to indicate which client program should be executed, while the data field number indicates which value computed by the client program should be optimized.

The client program number must be an integer between 0 and 99. If the Int1 value for a UDOP operand is set to 17, for example, then the client program number is 17, and the client program to be executed must be named UDO17.EXE. The client program name must always use a two digit representation of the client program number. If the client program number was 6, then the client program to be executed would be UDO06.EXE. The client program file must reside in the \UDO directory off the main ZEMAX directory.

When reaching a UDOP operand with a data field number of zero (more on the data field number shortly), ZEMAX will call the client program. The client program is called with the following syntax, assuming the client program number was 17:

```
UDO17.EXE buffercode hx hy px py
```

The buffercode is an integer value provided by ZEMAX to the client that uniquely identifies the correct lens. Because ZEMAX is capable of evaluating multiple lenses simultaneously, the buffercode is used as an identifier so that when the client requests or returns data, it is associated with the correct lens. Note that while optimizing

ZEMAX may be evaluating dozens of lenses simultaneously, each slightly different as derivatives are evaluated and the optimization progresses. The client must compute data for the indicated lens.

Once execution in the client has begun, the client program must do the following critical steps:

- 1) Establish a DDE link with the ZEMAX server.
- 2) Load the correct lens in the ZEMAX server's memory.
- 3) Compute the required data.
- 4) Pass the data back to ZEMAX.
- 5) Clear the ZEMAX server's memory.
- 6) Terminate the DDE link and exit.

The DDE link is typically maintained by the ZCLIENT code, which is described in the DDE chapter (the user is of course free to write their own, if preferred). ZCLIENT calls the UserFunction defined by the user to compute the operand data.

To load the correct lens into the ZEMAX server's memory, a single item must be sent to the ZEMAX server, GetUDOSystem. The syntax is "GetUDOSystem, buffercode". This will cause ZEMAX to retrieve the correct lens from system memory, and all subsequent DDE calls will be for actions (such as ray tracing) on this lens.

The data is then computed, optionally using any of the DDE item calls defined in the DDE chapter. Once the data is computed, up to 51 values are placed in a long formatted string. The string is sent back to the server, and ultimately to the optimizer within ZEMAX, with the SetUDOData item. The syntax is "SetUDOData, buffercode, free formatted integer, exponential, or floating point numbers.

It is critical that the client program pass back a string using SetUDOData, even if it only includes the buffercode with the remaining fields blank. ZEMAX will wait for the client program to pass this string back. ZEMAX has no way of knowing how long the computation may take; and so ZEMAX will "hang" until it receives the data. If the client program crashes, or never returns the data, ZEMAX will never complete execution of the operand, and will hang forever. Pressing the escape key from within ZEMAX will "break" the evaluation and cause ZEMAX to skip the operand evaluation.

Terminating the link is achieved by the client program UserFunction returning control to ZCLIENT.

A sample code that computes 3 data items, called a, b, and c, might look like this:

```
void UserFunction(char *szCommandLine)
{
    double a, b, c;
    char szBuffer[5000], szSub[256];
    int buffer_code;


    /* get the buffer code that identifies the lens */
    buffer_code = atoi(GetString(szCommandLine, 0, szSub));

    /* set the correct lens in the server's memory */
    sprintf(szBuffer, "GetUDOSystem,%i", buffer_code);
    PostRequestMessage(szBuffer, szBuffer);

    /* Here is where we compute the data... these lines are omitted */

    /* now return the data for dataflag is 0, 1, and 2 */
    /* DO NOT use SetUDOData if the total line length will exceed 255 characters, */
    /* use SetUDODItem instead. See the chapter on ZEMAX Extensions for details. */
    sprintf(szBuffer, "SetUDOData,%i,%.7f,%.7f,%.7f", buffer_code, a, b, c);
    PostRequestMessage(szBuffer, szBuffer);
}
```

Note that multiple data values may be returned with a single call to the client program. This is where the data field number comes into play. The data field number may be any number between 0 and 50, inclusive. This number refers to a position in a global array associated with the lens in memory. During execution of the client program, the client DDE item "SetUDOData" is used to return a long string of 51 numbers (or less; blank or omitted values are assumed to be zero). These values, which must be numeric, are stored in the lens buffer for later use by UDOP. In practice, the 255 character limitation of DDE item names limits SetUDOData to passing back far fewer than 51 items; but the DDE item "SetUDODItem" works around this limitation by passing back data one at a time.

 *There is a limitation of 255 characters in any DDE item name; and so SetUDOData is limited to passing back a small number of values. To work around this limitation, use SetUDOItem, described in the chapter "ZEMAX Extensions".*

There are 51 different data fields, so that a single client program call can be used to optimize up to 51 different values simultaneously. The data field number indicates which of the returned values should be placed in the "value" column for that UDOP operand.

UDOP also permits the use of the Hx, Hy, Px, and Py data fields. These data fields can be read by the client program, because they are passed on the command line after the buffercode.

There is one very important thing to know about the data field number. If it is zero, then the client program is executed and the value from data position 0 is placed in the value column. However, if the data field number is not zero, then the client program is not executed, but any previous value stored from an earlier call to the client program is used instead. The advantage to this convention is substantial. If the client program computes many values, all of which need to be optimized, the client program only needs to be called once, yet multiple UDOP operands can access the data. This is much more efficient than calling the client program multiple times.

For example, suppose a client program named UDO25.EXE computes three values, all of which require optimization. In the client program, the values are passed back using "SetUDOData,buffercode,x,y,z". Then three UDOP operands in the merit function can extract the data and perform the optimization with a single call to the client program:

```
UDOP 25 0
UDOP 25 1
UDOP 25 2
```

The client program UDO25.EXE is only called during the evaluation of the UDOP 25 0 operand. Note the Hx, Hy, Px, and Py values can only be used if the data field value is zero, since only in this case is the client program evaluated.

Unlike using ZPL macros, UDO's may freely change the lens data during evaluation since all DDE commands are executed on a copy of the lens, not the actual lens being optimized.

There is a sample UDO source code file called UDO_DEMO.C; it may be compiled and linked in with ZCLIENT. The executable needs to be renamed UDOxx.EXE where xx is a two digit integer. The sample UDO returns 6 values: Hx, Hy, Px, Py, and two dummy constants, in data positions 0 through 5, respectively.

Suggestions for use

In preliminary design stages, it is rarely required to trace all of the rays for all of the wavelengths at each field position during optimization. For this reason, execution times may be substantially decreased by limiting the number of fields and wavelengths used during optimization. If the weight of selected fields and wavelengths is set to zero, then the default merit function algorithm will skip the zero weighted fields or wavelengths when constructing the merit function. This results in fewer rays being traced, speeding execution.

For example, if the lens is being evaluated at five field points, it is possible that only the first, third, and fifth field need be included in the merit function. Of course, later in the design process all fields may need to be included and the default merit function reconstructed.

There are a few other tricks to improve performance. Avoid setting boundary operands on variables unless the optimized solution persists on implausible designs. Boundary operands add computational overhead. Use solves instead of explicit operands whenever possible. For example, use a curvature solve to control the focal length rather than an operand if possible.

Optimization is inseparable from the art of modern lens design, and only practice will make a designer a proficient user of optimization algorithms. Users who are expert at other software optimization algorithms will probably find ZEMAX easier to use, and with a little practice, the mechanics of using the interface will slip into the subconscious, and the designer can concentrate upon the design itself. If you are new to computerized optimization of lenses, there is no better way to learn than to practice.

The global optimum

The design that yields the lowest possible value of the merit function is called the "global" optimum and is by definition the best possible design. However, there is no known optimization algorithm that can universally find the global optimum for an arbitrary design problem, unless you consider "direct search" an optimization algorithm (in other words, try all of the infinite number of possible solutions to see which is best). The art of optical design

with computer assistance has two basic components. First, the designer must be able to determine a suitable starting design, and second, he or she must play the role of supervisor during the optimization process. A good supervisor knows when and how to back up and coax the program into a more fruitful direction. Unfortunately, this often requires considerable experience, and even more often, excessive tedium. An experienced designer uses a combination of intuition, analysis, and luck in searching for new, better, design forms. ZEMAX provides an automated capability for performing this global optimum search; the feature is described in the next chapter.

In this special case, ZEMAX needs to use the TCE supplied in the Lens Data Editor to define the mounting material properties, rather than the TCE supplied in the glass catalog. This can be accomplished by setting the "Ignore Thermal Expansion" switch for the material in the glass catalog.

Other glass catalog options


Selecting "Sort Index" sorts the glasses by the d-light (.587 microns) index of refraction, "Sort Abbe" sorts the glasses by the d-light Abbe Vd number, and "Sort Name" sorts by glass name (alphabetically then numerically). Calc Nd/Vd computes the Nd and Vd values based upon the dispersion coefficients. Report provides a text listing for the selected glass.

Finding a glass quickly

The fastest way to view the data for any glass is to click once on the name of the glass in the Lens Data Editor, then select Tools, Glass Catalogs (or select the GLA button if displayed). The correct catalog and glass will be displayed.

Using MIL number glasses

MIL number glasses are those described by a six-digit number, such as 517640 for BK7. The first three digits in the MIL number is the d-light index minus one, without the decimal place. The last three digits is the Abbe V-number times 10. You can enter a six-digit number directly in for the glass rather than the glass name if you desire. ZEMAX uses a formula for computing the index at each defined wavelength based upon the index and Abbe number defined by the MIL number. The formula is based upon a least-squares fit of coefficient data of many typical glasses. Typically, the index data calculated is accurate to roughly .001. At wavelengths in the deep UV or infrared the index value becomes less reliable. MIL number glasses are generally inferior substitutes for the constants of dispersion (or other) models for the glass, however they are useful if no other data is available.

 **Note MIL number glasses are an approximation, although usually a very good approximation in the visible range. Outside the visible wavelength range, such as in the ultraviolet or infrared, the MIL number glass is not accurate and should not be used.**

It is important to note that the indices calculated from the six-digit MIL number are not the same as those calculated from the glass catalog, even if the MIL number you are using corresponds to a glass in the catalog. Index data is calculated directly from the MIL number entered on the main screen; not from the glass catalog data, even if a glass with that MIL number is in the catalog.

Using model glasses

ZEMAX can idealize the dispersion of glass using the index at d-light (.5875618 microns), the Abbe number, and a term which describes the deviation of the partial dispersion from the "Normal Line". The index at d-light is given the symbol Nd. The Abbe number (also called the V-number) is given the symbol Vd and is defined by

$$V_d = \frac{N_d - 1}{N_F - N_C}$$

where N_F and N_C are the indices of refraction at .4861327 and .6562725 microns, respectively. The partial dispersion term is $\Delta P_{g,F}$.

ZEMAX uses a formula based upon the typical dispersion of standard glasses in the visible range to estimate the index at any defined wavelength within the visible range as a function of the Nd and Vd values. This formula is accurate to roughly .0001 for typical glasses.

The Nd, Vd, and $\Delta P_{g,F}$ values are specified on the glass solve dialog box, which can be reached from the Lens Data Editor.

Note model glasses are an approximation, although usually a very good approximation in the visible range. Outside the visible wavelength range, such as in the ultraviolet or infrared, the model glass is not accurate and should not be used.

Glass catalog sources

The data for the glass catalogs included with ZEMAX are generally provided by the manufacturers. For example, the data in the Schott catalog was provided to Focus Software by Schott Glass Technologies. The data has been converted to the ZEMAX .AGF format.

Although the data contained in the glass catalogs is generally reliable, it is always possible for errors to be made during translation or editing of the data. *It is absolutely crucial that all index data be verified for accuracy by the end user! This is especially true where fabrication of the optics is being considered.*

The catalog data may be in error, and it needs to be verified before it can be trusted.

The data contained in the INFRARED.AGF catalog has been compiled from published sources as described in the following table. A few other materials may be included in the catalog which are not referenced in the table. These materials should be especially scrutinized before trusting the data accuracy.

INFRARED CATALOG DATA SOURCES

Material	Source
AG3ASS3	Handbook of Optics Vol. II
AGCL	Handbook of Optics Vol. II
AGGAS2	Handbook of Optics Vol. II
AGGESE2	Handbook of Optics Vol. II
AL2O3	Handbook of Optics Vol. II
ALN	Handbook of Optics Vol. II
ALON	Handbook of Optics Vol. II
AMTIR1	Amtir Spec Sheet
AMTIR3	Amtir Spec Sheet
BAF2	Handbook of Optics Vol. II
BEO	Handbook of Optics Vol. II
CALCITE	Handbook of Optics Vol. II
CAF2	Handbook of Optics Vol. II
CAMOO4	Handbook of Optics Vol. II
CAWO4	Handbook of Optics Vol. II
CDSE	Handbook of Optics Vol. II
CDTE	Handbook of Optics Vol. II
CLEARTRAN	Spec Sheet (Morton)
CSBR	Handbook of Optics Vol. II
CUGAS2	Handbook of Optics Vol. II