Lesson 7. The In and Out of Vignetting

By "vignetting" we mean any property of a lens that reduces the size of the beam relative to the size that is sent in. It is a topic that different programs handle in different ways. Of course, one would often prefer that the beam size remain constant everywhere in the field, since the effective transmission then does not fall off with field on that account. But sometimes the best tradeoff is to accept some vignetting in order to avoid the cost and weight of a more complex lens. In such cases one must know how to manage the varying beam size during optimization and how to set up the lens apertures to model that amount of vignetting for image analysis when the design is finished. That is the subject of this lesson.

Here is a sample lens, a triplet with substantial vignetting.
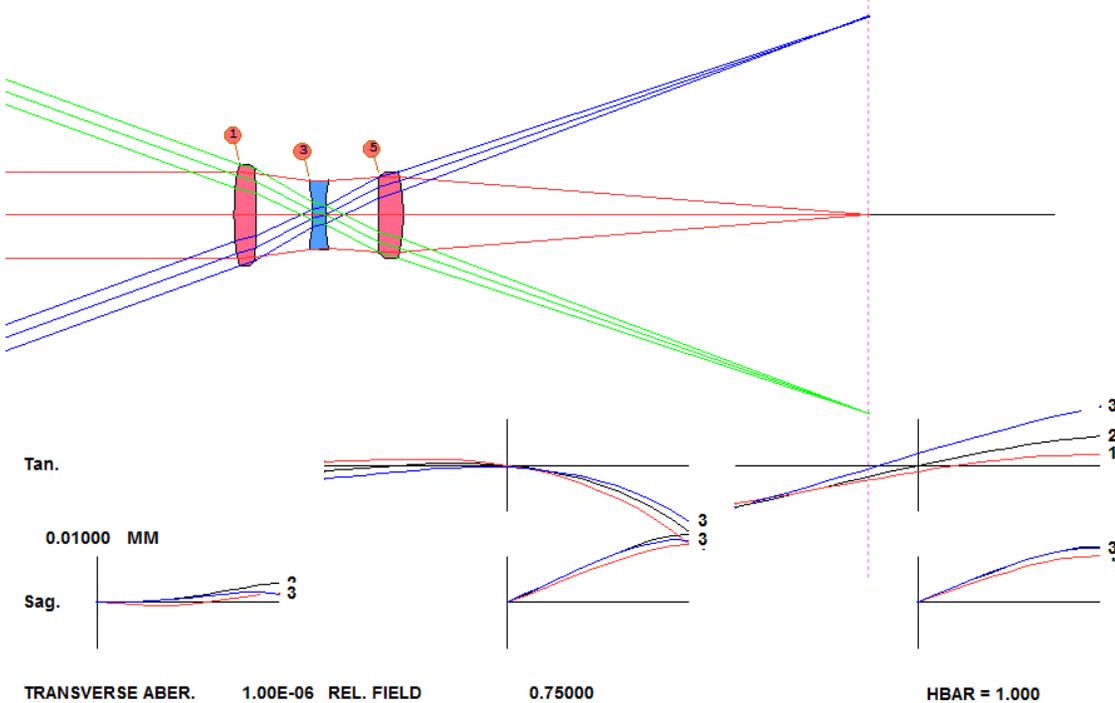
Notice how the size of the beam at the upper and lower field points (in blue and green) is much smaller than the on-axis beam (in red). A look at the RLE file for this lens shows a real stop on surface 3 (to activate ray aiming for the chief ray) and wide-angle pupil option 3 in effect (WAP 3). Here is the RLE file:

```
RLE
ID COOKE TRIPLET F/4.5                          747
 WAVL .6562700 .5875600 .4861300
 APS             -3
 WAP              3
 UNITS MM
 OBB  0.000000     20.00000     5.55500    -2.98488    20.00000    -2.98488     5.55500
   0 AIR
   1 CAO       4.69068139       0.00000000       0.00000000
   1 RAD     21.4939500000000    TH       2.00000000
   1 N1 1.61726800 N2 1.62040602 N3 1.62755182
   1 CTE    0.630000E-05
   1 GTB S     'SK16             '
   2 CAO       4.25560632       0.00000000       0.00000000
   2 RAD    -124.0387000000000    TH       5.25509000 AIR
   3 CAO       3.19251725       0.00000000       0.00000000
   3 RAD     -19.1051800000000    TH       1.25000000
   3 N1 1.61163844 N2 1.61658424 N3 1.62846980
   3 CTE    0.830000E-05
```

```
  3 GTB S    'F4              '
  4 CAO      3.15978037      0.00000000       0.00000000
  4 RAD     21.9794700000000   TH    4.93473000 AIR
  5 CAO      3.48158127      0.00000000       0.00000000
  5 RAD    328.3317499999989   TH    2.25000000
  5 N1 1.61726800 N2 1.62040602 N3 1.62755182
  5 CTE   0.630000E-05
  5 GID 'SK16               '
  5 PIN    1
  6 CAO      4.00000022      0.00000000       0.00000000
  6 RAD    -16.7537700000000   TH   43.24303731 AIR
  6 TH      43.24303731
  6 YMT      0.00000000
  7 CV       0.0000000000000   TH    0.00000000 AIR
END
```

The WAP 3 option reduces the entrance pupil size so that the marginal rays at each field point just clear *all* defined lens apertures. Every surface except the image (at surface 7) has been assigned a hard clear aperture (with the CAO data). This is one way to implement a desired amount of vignetting. Note how the size of the entering beam changes with field angle. This is most easily shown with the PAD Scan button ↑.
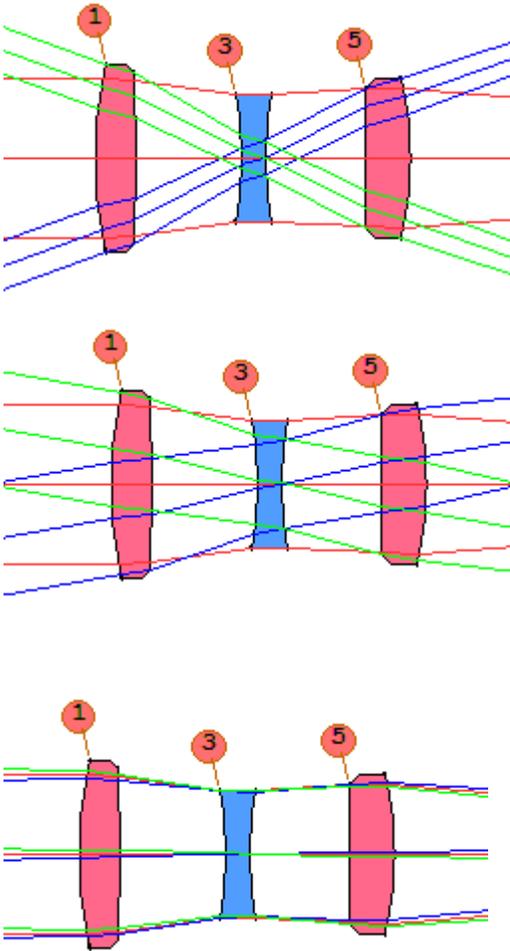
But the WAP 3 option is not the only way, and often is not the best way, to deal with vignetting. During optimization, the size of the beam can change at each surface as the lens changes, and it makes no sense to assign a hard CAO to a surface when you don't even know what that size will be when you are finished. So we never use the WAP 3 option during optimization and seldom use it afterwards.

Instead, we approach vignetting in stages. To illustrate, let us first remove all of the CAOs and the WAP declaration as well.

```
CHG
CFREE
WAP 0
END
```

Ouch! The image is awful and edges are feathered. Maybe that's why the starting design used those options. Perhaps it made sense under the original conditions. What can we do with this lens?
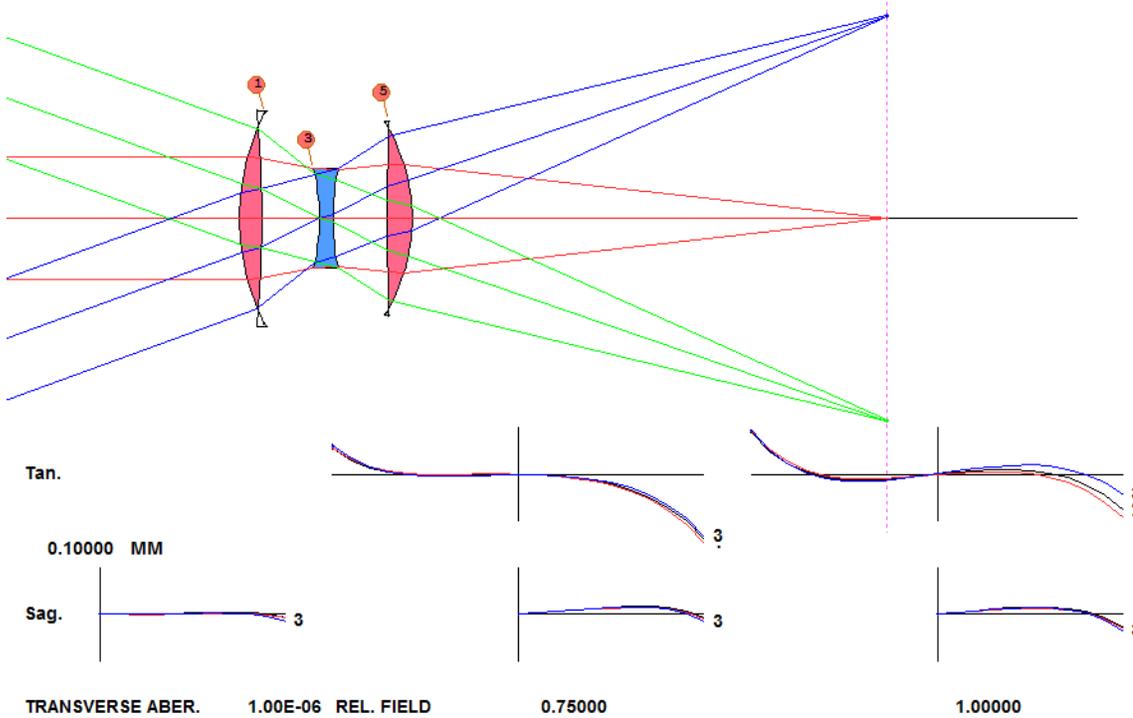
A look at the current options (with POP) shows a YMT solve on 6 but no curvature solve. Let us add one. The lens works at F/4.5 paraxially, and the value of a UMC solve is therefore 0.5/4.5 or -.1111, the minus sign because the marginal ray is going *down* at the image. We change the lens and store a copy in the lens library for future reference.

```
CHG
6 UMC -.1111
END
STORE 3
```

Now we open a new editor with the command AEE and create an optimization MACro.

```
LOG
PANT
VLIST RAD ALL
VLIST TH ALL
END

AANT
AEC
ACC
GSR .5 10 5 M 0
```

```
    GNR .5 2 3 M .7
    GNR .5 1 3 M 1
    END

    SNAP
    SYNO 30
```

Here we have used merit function 6 from the Ready Made Raysets button, .

We make a checkpoint and run this MACro.
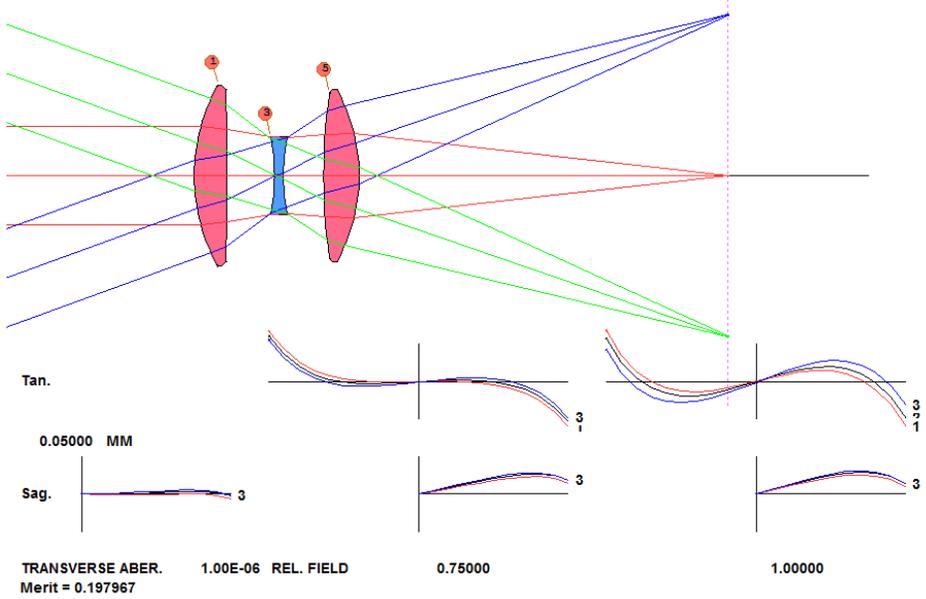
Not too good. The edge feathering has been fixed because of the AEC monitor, but aberrations are out of control at full field. We have to do better. Let's tell the optimization program to vignette the beam size to 40% of the on-axis value at full field. We add to the AANT file a VSET directive,

```
    AANT
    AEC
    ACC
    VSET .4
    GSR .5 10 5 M 0
    GNR .5 2 3 M .7
    GNR .5 1 3 M 1
    END
```

and run it again.

Tan.

0.10000 MM

Sag.

TRANSVERSE ABER.   1.00E-06 REL. FIELD   0.75000   1.00000
Merit = 0.0713521
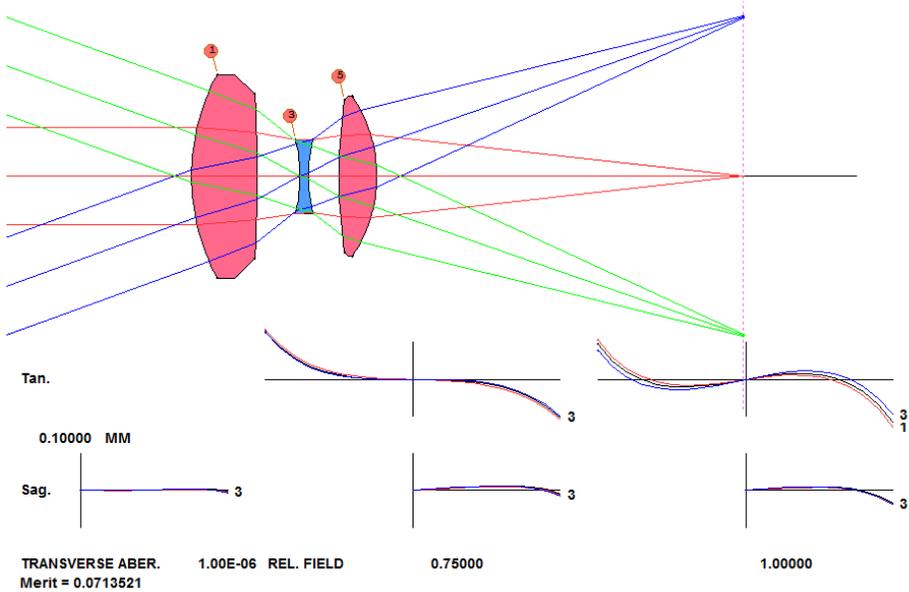
The *edges* of the TFAN got worse, which is not surprising, since they are not corrected anymore. But if we were to vignette this beam at 40% of the size shown in the PAD display, the image would look much better. Let's assume this is the solution we are after. Now we have to model the elements so that amount of vignetting actually happens.

And that's easy to do. Open the WorkSheet (**WS**), enter into the edit pane the directive **CFIX**, and click Update. All surfaces now have a hard CAO assigned, at the same aperture as the default CAO that is currently in effect. Now, click on surface 6. The data for that surface are shown in the edit pane. *Select the CAO radius* with the mouse and then click the **SEL** button. That assigns the top slider to that aperture radius. Move the slider thumb to the left, reducing the aperture. Watch the TFAN at full field, and stop when the unvignetted portion on the left side appears to be about at the 40% position.

This is about where the beam should be vignetted on that side. Do the same at surface 1. Now the beam is vignetted on both sides.
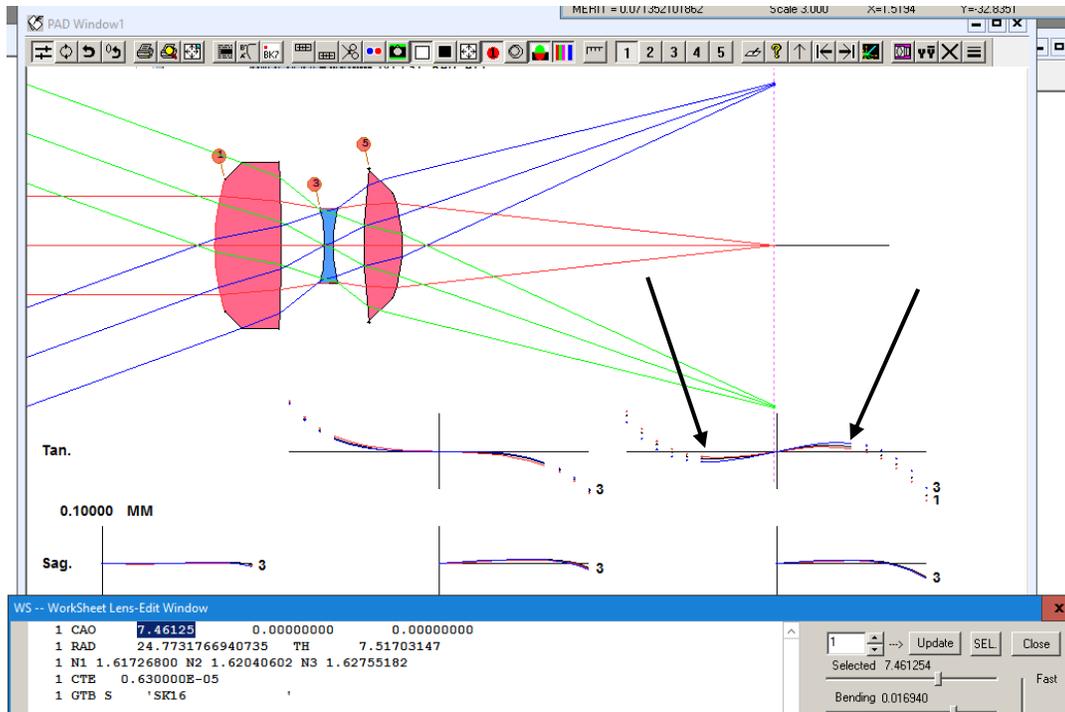
Why does the drawing on top of the PAD display still show the original, unvignetted beam? Well, there is an option for that too, which you can activate by turning off mode switch 65. But since this makes the picture look like what would happen with the WAP 3 option turned on – and it is *not* turned on at the moment, we prefer to leave this switch turned on so we don't get confused later.

Okay, we are almost done. We can activate WAP 3 by adding that directive in the edit pane. That's one way to proceed.

But here's another: We can declare a set of VFIELD parameters. We type into the CW:

    FVF 0 .5 .8 .9 1

That makes the program calculate the vignetting factors that just clear the apertures we have entered at four points in the field. Now the display shows the vignetted beam as it should.

6

Tan.

0.05000   MM

Sag.

TRANSVERSE ABER.   1.00E-06   REL. FIELD   0.75000   1.00000
Merit = 0.0713521
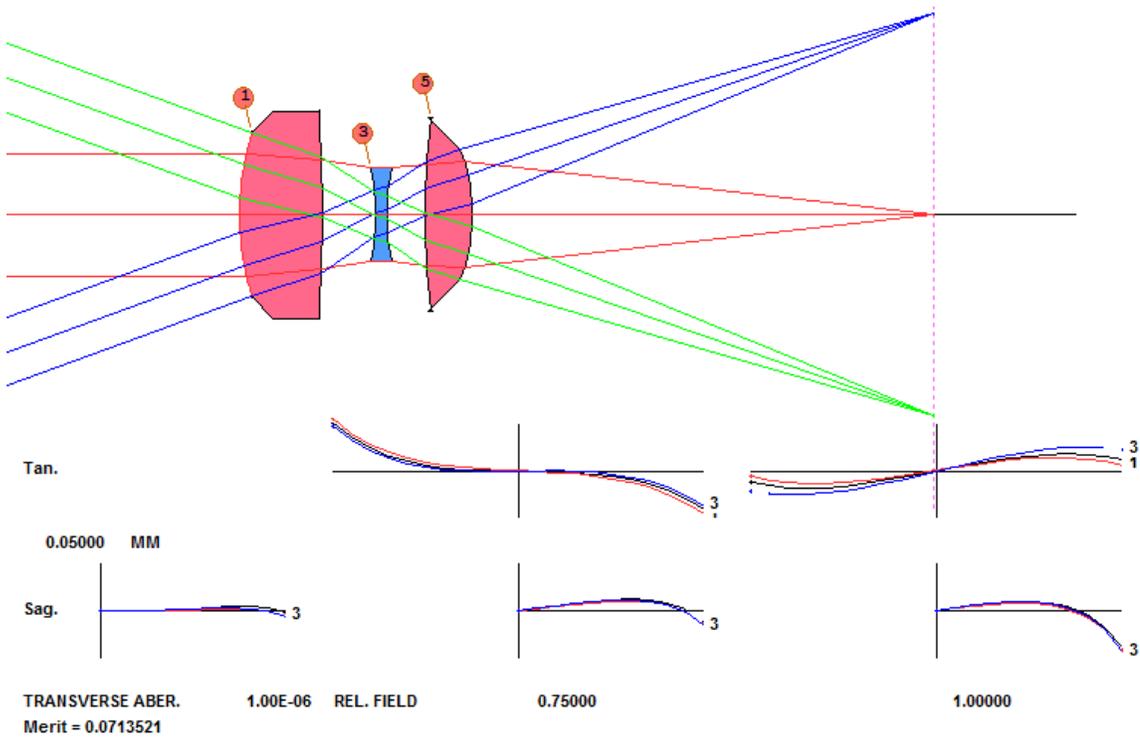
Except the other apertures are still where we left them earlier, because they are all hard apertures. In the edit pane, type **CFREE** and click Update. Now the lens has default apertures – calculated to just pass the beams as adjusted by VFIELD.
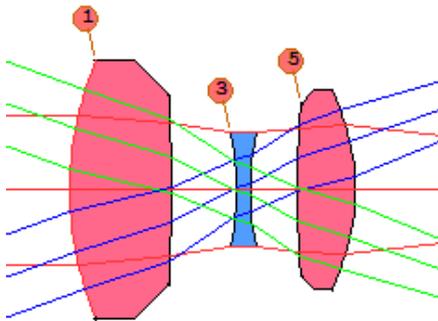
You see the pattern. If we now go back to our optimization MACro, remove the VSET directive, and reoptimize, the edge-control monitor will look at the rays shown above, not the nominal rays. It will therefore reduce the thicknesses if the merit function no longer benefits from the larger values. Then we can adjust the edge geometry with the Edge Wizard (**MEW**) just as we like them.
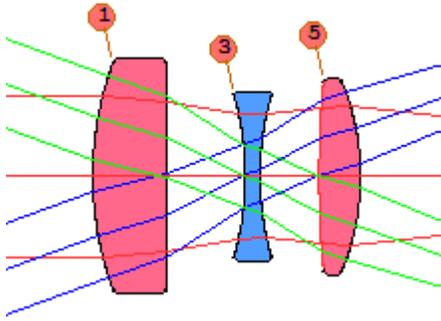
What's the difference between the WAP 3 option and VFIELD, since they look much the same? I'm glad you asked. WAP 3 requires aiming five rays every time you do something that requires ray tracing. That is a rather slow option. VFIELD, on the other hand, has already done that calculation and only needs to aim the chief ray afterwards, interpolating at the requested field. Fast.

And that's what vignetting is all about.