Lesson 25. Tips and Tricks of a Power User

If you have gone through the previous lessons in this Online Tutorial, you already know about the unusual power and friendliness of SYNOPSYS™. Let us now go a little deeper and show some rather nice ways you can get just what you are after.

Example MACro

Here is an example of a sophisticated optimization MACro. By now, you should be able to read this kind of file fluently.

```
LOG
AWT: 1.0
CHG
NCOP
END

PANT
VY 0 YP1 50 -50
VY 0 BTH
VLIST RD ALL
VLIST TH ALL
END

AANT P
AEC
ACC
M 100 10 A FOCL
LLL 2 1 1 A BACK
AAC 49.5 .5 5

SKIP
GSR     AWT     6.000000        4   1     0.000000
GNR     AWT     3.000000        4   1     0.100000
GNR     AWT     3.000000        4   1     0.300000
GNR     AWT     3.000000        4   1     0.500000
GNR     AWT     3.000000        4   1     0.70000
GNR     AWT     3.000000        4   1     0.80000
GNR     AWT     3.000000        4   1     0.90000
GNR     AWT     3.000000        4   1     1.000000
EOS

!SKIP
GSO     0       0.8         4   1     0.000000
GNO     0       0.27        4   1     0.100000
GNO     0       0.27        4   1     0.300000
GNO     0       0.27        4   1     0.500000
GNO     0       0.27        5   1     0.70000
GNO     0       0.27        5   1     0.80000
GNO     0       0.37        5   1     0.930000
GNO     0       0.27        5   1     0.950000
GNO     0       0.27        5   1     1.000000
EOS
```

```
        !SKIP

        LUL 29 1 1
        A BLTH 3

        EOS

        END

        !EVAL
        !EDS

        SNAP/DAMP 1
        SYNOPSYS 40
```

Tip 1:

In this example, we vary the paraxial quantity YP1. The lens does not have an explicit stop defined, and this variable will let the chief-ray intercept on surface 1 vary – thus sending it into the lens at the current location and *implying* a stop position wherever it happens to cross the axis. This is a powerful way to find out where the stop should go. If the design looks promising, then it is simple to just assign the stop at that location or close to it and reoptimize.

Tip 2:

Note the monitor AAC 49.5 .5 5. This lens has to fit inside a tube of 100 mm diameter, and this monitor will ensure that none of the clear apertures exceed the value 49.5 mm. The other two arguments give the relative weight and the monitor window. You can adjust them according to the importance of that control. (And you can always *look up* that topic in the User's Manual if you want to see how it works. Type HELP AAC in the Command Window. It pays to know how to use these powerful tools.)

Tip 3:

Notice how we have defined a symbol, AWT: 1.0. That symbol shows up in the AANT file as the aperture weighting parameter on some ray sets. A value of 0 says to weight all rays in the generated grid by the same amount, the value given by the second argument. A weight of 1.0 weights the central rays rather more heavily than the edge, and for this lens that is a good idea. Uniform weights tend to give an image of high contrast, while a higher weight gives better resolution. Here, too, you can experiment to see what works best with your lens. We often start with a value of 0.5. And the point of making this a *symbol* is so you can try a different value just by changing that symbol and reoptimizing. So you don't have to change all the weights on every line of the MACro. Neat.

Tip 4:

The shop making this lens has a curious problem: they already have a lens blank for the element at surface 3, and have measured its thickness at 30 mm. So that lens must be controlled during optimization to be sure it does not require a thicker blank.
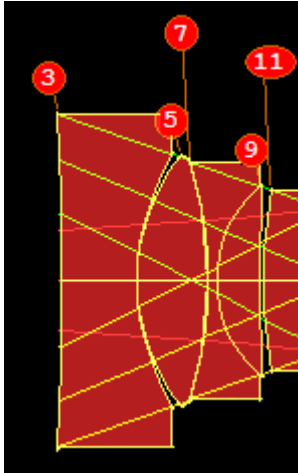
That's where the AANT entries below come in.

```
LUL 29 1 1
A BLTH 3
```

Here we specify a maximum value (with LUL) of 29 mm for the thickness of the blank for this element.   "LUL" means limit, upper limit, and you can read about this useful feature by typing HELP LUL. You can also make your own aberration by combining the TH and the sag at any surface that happens to be concave.   Section 10.3.3 of the User's Manual describes the target SCAO. It's useful to know how to make aberrations like this.

Tip 5:

Notice the use of the SKIP directive in this MACro. Ray grid definitions and weights are easily generated with the buttons on the MACro editor toolbar, and those either generate rays whose transverse intercepts are to be controlled, or rays with OPD targets. But which should you select? The SKIP directive in the above example lets you choose one or the other (or both) just by commenting or uncommenting that directive. As shown above, the first set of ray grids, which target transverse aberrations, will be skipped. When the program reaches the EOS (End Of Skip) line, it stops skipping -- so the ray grids targeting OPDs are then in effect. To see the effect of transverse targets instead of OPD, just uncomment the first SKIP and comment out the other (with "!"). You see how simple it is to switch from one to the other in SYNOPSYS™, which separates the variable list and merit function definition from the lens file itself (unlike other programs.) With a few clicks and keystrokes you can study the effect. Neat.

Tip 6:

More on the subject of TAP targets *vs* OPD targets: Look at the following ray-fan plots:
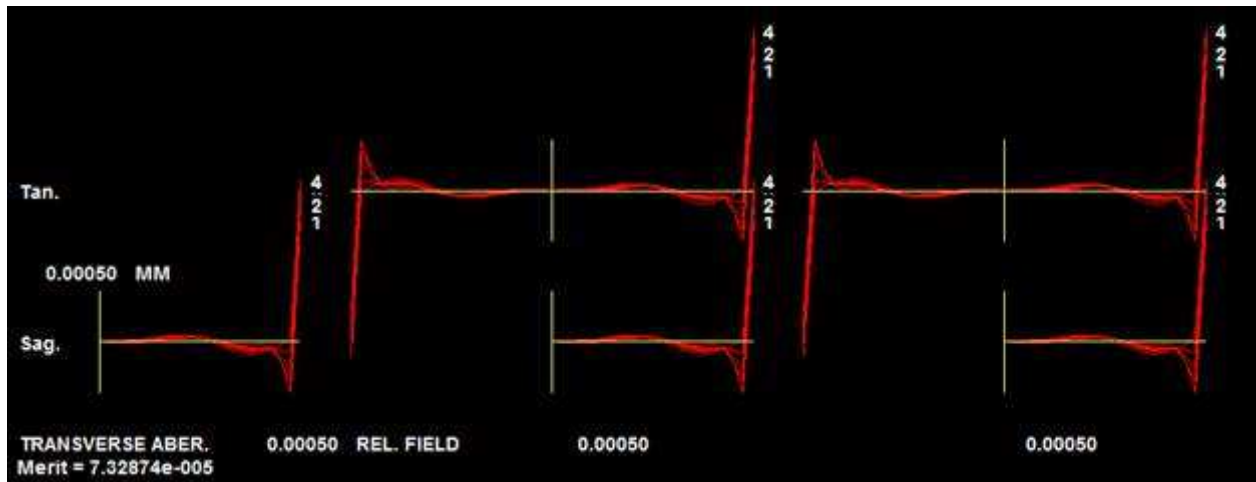
You might think this is a terrible image, with the rays flying off like that at the ends of the fans. But look at the OPD fans:
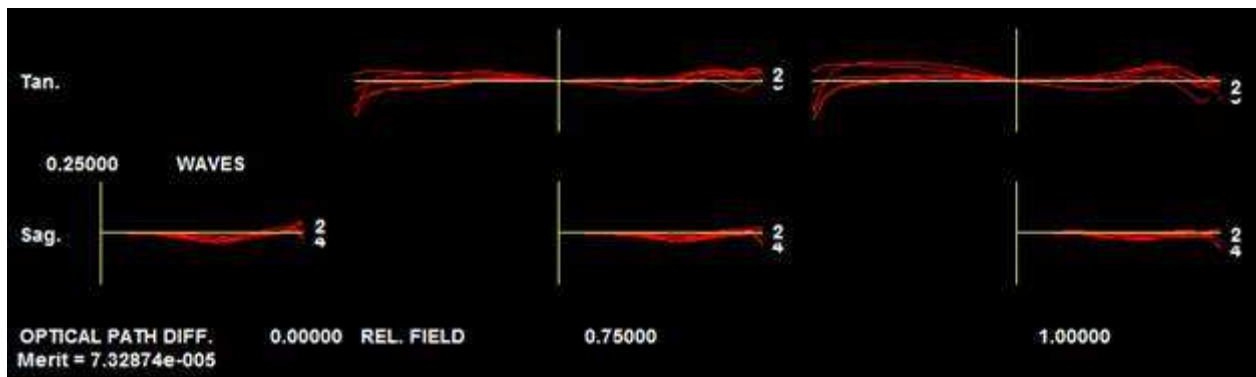
Can't get much better than that! This is a good lesson on why you should switch to OPD targets if your lens is close to the diffraction limit. Someone targeting only ray intercepts would likely look at the top curves and then discard the lens and start over -- or give a heavy weight to the marginal rays and keep optimizing. But that would not work very well. Specifying a large value of the aperture weight parameter (which we called AWT in the above MACro) would work somewhat better, but still, OPD targets are better in cases like this.

Tip 7:

So why use transverse targets, anyway? Well, for the lenses we have studied, those targets can change a lens faster – and thus arrive at a good configuration – more quickly than can OPD targets. So start with transverse and switch to OPD when you are almost there.

Tip 8:

Suppose we only want to evaluate the current lens, to spot the largest aberrations – but don't want to change anything just yet. Simple. Just uncomment the lines

```
!EVAL
!EDS
```

and run the MACro. The program will evaluate the merit function and then end the data set (the MACro) at the EDS line. Later, if you want to optimize the lens, just comment those lines out again and rerun it.

Tip 9:

When you optimize your lens, assuming you have switch 1 turned on (the default), it is a good idea to run the program twice. Thus,

      SYNOPSYS 50

      SYNOPSYS 50

is better than

      SYNOPSYS 100.

This is because if any of the variables in the first run encounters a boundary, that variable is dropped for the remaining iterations. The rule saves time, since usually that variable will continue to try to violate the same boundary, and if it is not dropped it will slow things down considerably. But if the lens changes its form significantly during the run, it often happens that the same variable now wants to move in the other direction. The second run frees all variables so it can do so. This is especially true for glass model variables (GLM), which usually hit a boundary early on.

Tip 10:

In the MACro given at the start of this lesson, we included both transverse aberrations and OPDs. Notice the difference in weights. This reflects the fact that an OPD error of one unit (one wave) usually gives a much better image than a transverse error of one inch (or one mm). Sometimes one can get excellent results from giving targets to *both* kinds of errors -- but now the relative weights are important. And if you have, say, a target for a mechanical property, such as an aperture or spacing somewhere, you want to be sure that your very nice balance is not upset when the program thinks the one-wave OPD is just horrible and tries to bring it down at the expense of the other errors. To make finding a suitable weight easier, we have provided two very nice tools. If you click the button in the MACro editor and choose option 8, the program will create ray grids for both kinds of targets, assigning weights to the OPD errors that reflect the current wavelength and F/number. These weights ensure that the difference in units is accounted for in a reasonable way, and of course you can adjust things from there after you see the effects. The button also lets you select OPD targets, and in this case, you assign relative weights as usual and then click the box "Calculate special OPD weights". The OPD weights will be modified by the same rule when the ray grid requests are added to the MACro.