

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Man versus machine: a lens design challenge

Donald Dilworth, David Shafer

Donald C. Dilworth, David Shafer, "Man versus machine: a lens design challenge," Proc. SPIE 8841, Current Developments in Lens Design and Optical Engineering XIV, 88410G (25 September 2013); doi: 10.1117/12.2022871

SPIE.

Event: SPIE Optical Engineering + Applications, 2013, San Diego, California, United States

Man versus Machine: a Lens Design Challenge

Donald C. Dilworth

Optical Systems Design, Inc.

David Shafer

David Shafer Optical Design

ABSTRACT

After a generation of writing and improving lens design software, it is time to assess where we are. Specifically, can a modern program compete with, or surpass, the best human designers? Here we describe a friendly contest between two leaders in the field.

Keywords: Lens design, global optimization, optimization algorithms, aberration theory.

1. INTRODUCTION

In an American folk tale the railroad pile driver John Henry pits his prowess with a hammer against a steam-powered drill. Man versus machine. John Henry won the contest, but it killed him. A more modern version is the computer Big Blue's chess victory over human grandmasters. We naturally wonder where we are today in the field of lens design in this continuing man-versus-machine drama. Of course human designers are heavily dependent on computers, but we like to think that our many years of experience doing design and our understanding of aberration theory ought to outclass a completely automatic design program such as those offered by software vendors.

A very remarkable feature of this type is found in Don Dilworth's superb lens design program SYNOPSISTM*, which he will describe in some detail below. It generates many suitable design candidates with no human intervention, once the initial conditions and requirements are specified. Here we will show how that program performed on a particular design problem that was formulated for this man-versus-machine contest, where I (Shafer) was the man. Don and I wanted a problem that would be difficult but which would not require more than about 10 or so lens elements. Ideally it would also be one where my human solution to the problem would have some theoretical basis in aberration theory and design principles that would presumably give me an edge over some blind computer search algorithm.

2. THE HUMAN DESIGNER (SHAFER)

I chose a design problem that I had already written about over 30 years ago, in US patent 4,770,477 which involves an interesting aspect of aberration theory. It is to design a lens that is well corrected at two widely separated laser wavelengths but where there is essentially zero spectral bandwidth at each of the two lines. Chromatic variation in spherical aberration, coma, and astigmatism must be well corrected, and the design must be diffraction-limited over the field at both wavelengths at the same focus location. More recently, in 1995, I wrote again about this design problem¹, and the picture in Figure 1 shows the relevant design. The top drawing from that article shows a monochromatic design for a single wavelength that is diffraction limited over the field.

This design is a 5X relay lens with a .35 NA and a 24 mm field diameter for use at the mercury i-line—but it must also be well corrected at .6328 μ for testing over the whole field at that visible wavelength. That design was to be

* The SYNOPSISTM program is a trademark of Optical Systems Design, Inc. See www.osdoptics.com.

achromatized over the bandwidth of the mercury line, and the $.6328\mu$ image could be at a different focus setting as required by the secondary color of the design (a different focus setting is not allowed in the problem we will do in this paper, but the two lines each have essentially no bandwidth).




Figure 1. Early design problem and solution

The interesting aberration theory here concerns the difference between intrinsic and induced aberrations. The strong negative lens that is 4th from the left end of the design above has a lot of astigmatism because of its power and its location very close to the aperture stop. It also has a lot of chromatic variation of astigmatism because its power changes with wavelength due to the glass dispersion. That is an intrinsic aberration that does not depend on the nature of the surrounding lenses but only on the conjugates and pupil position of that lens. But in the system it has another component to its astigmatism, coma, etc. The lenses before and after that negative lens have a lot of chromatic aberration that causes a wavelength-dependent shift in the input conjugates and pupil position of that negative lens. For example, the chief ray angle coming into the negative lens is different for different wavelengths due to lateral color of the preceding lenses, and that makes the astigmatism of the negative lens change with wavelength, even if that negative lens had no dispersion of its own. This is called *induced* aberration, due to incoming aberrations to the lens and is often more significant than the intrinsic aberration of a lens.

In this design I found that by achromatizing the large positive lenses I could greatly reduce the amount of axial and lateral color going into that strong negative lens (imagine the light going from the right side of the design to the left), and that greatly reduced the induced aberrations at that lens. The result was a complete design that was corrected for chromatic variation of spherical aberration, coma, and astigmatism—almost without effort. Both the $.365\mu$ i-line and $.6328\mu$ were diffraction limited over the field, although at different focus settings. Let us all agree not to think about the explanation just given if the light goes from left to right in the design.

For the new contest problem we altered the goals, requiring the design to be well corrected at two widely separated laser lines, each assumed to have zero bandwidth, and at the *same* focus setting. The two laser lines chosen were $.351\mu$ and $.4461\mu$. To keep the short wavelength transmission high, only two glasses were allowed: BK7 and LF5. The design specs were, .35 NA, a 10 X 10 mm square field, 100 mm focal length, system length less than 250 mm, including the distance to the aperture stop if it was external, back focus greater than 10 mm, no vignetting, a telecentric image, and distortion to be essentially zero at the edge of the field. Also, no cemented lenses are allowed.

Some exploratory design work showed that this is a much more difficult problem than if we allowed the two wavelengths to have different focus positions. The goal of this man-versus-machine contest was to not only find a good design that meets the requirements, but also one with the fewest elements. My starting point was a monochromatic design similar to Figure 1, shown above. I found that I could remove some lenses and get a good monochromatic design with only 9 lenses, all BK7. I then added two LF5 lenses to get the design in Figure 2, which has 11 lenses and good polychromatic MTF. This meets all the specs, and I then told Don that at least one good design exists with 11 lenses and asked if his automatic search program could find it or some other design solution. Don gave his program 11 elements to work with and it came up with a design that is very similar to mine in the back half. I found this to be quite astonishing and also humbling, since I had spent a lot of time on my design and have 47 years of full-time design experience.




Figure 2. 11-element design of David Shafer.

My aperture stop kept wanting to move out in front of the lens, which hurts the system length, so I froze it at the front surface. The 2nd lens from the left side is BK, while the other two negative lenses are LF5.

After studying my 11-lens design for a while I thought I would see if a lens could be removed, namely the next-to-last lens. This worked out quite well and the result polychromatic MTF was almost as good as before. The 10-lens design is shown in Figure 4.




Figure 3. MTF of the 11-element design of David Shafer.




Figure 4. A 10-element solution found by David Shafer.

It then occurred to me to try combining the 2nd and 3rd lenses in this design, which are both BK7, to make a 9-element lens. This worked out quite well, and the result is shown in Figure 5. The polychromatic MTF is still good and seems unchanged from the 11-lens design.




Figure 5. A 9-element solution found by David Shafer.

All three negative lenses are LF5 glass. The polychromatic MTF is a little worse than before but it is still clearly a diffraction-limited design.

Surely, it seemed, it would not be possible to remove another lens. But the 3rd lens from the left is weaker than the other lenses and it was worth a try to remove it. After a whole lot of optimization and the use of some of the optimization “tricks” that I have written about, the result was a very good 8-lens design, shown in Figures 6 and 7. All three negative lenses are of LF5 glass.

The final stage came next. With some difficulty I was able to thicken the 2nd and 3rd lenses and combine them into one thick LF5 lens, making the 7-element design shown in Figure 8, and still get good performance. But now the polychromatic MTF is barely diffraction-limited, as shown in Figure 9.

This is definitely the end of the line, and quite a surprise when one considers the 11-lens design that we started with. In all of my designs here the stop wanted to move out in front, while Don’s solutions sometimes have it inside.

Considering the amount of work all of this required from me and the very impressive results that Don’s search program produced, I am both excited and appalled that there is such a great design tool available. Of course Don did some massaging of his designs after they were found by his search program, but the design forms and almost the final version were all found automatically. This is really quite amazing!




Figure 6. The 8-element solution of David Shafer.




Figure 7. MTF of the lens in Figure 5.




Figure 8. A 7-element design. This is as far as the problem can be pushed, it seems.




Figure 9. MTF of the 7-element solution.

3. THE STEAM DRILL (DILWORTH)

I was delighted when David Shafer replied to my request for a really hard lens design problem. As part of the continuing development of my optics program SYNOPSISTM, I have been studying global optimization—and I wanted to test my algorithm in a way that would challenge and possibly break it.

Here we really come face to face with the John Henry dilemma: can an algorithm running on a stupid but very fast PC do as well as the best human experts? Of course the answer is not simple, since any lens design problem can be divided into separate phases, which start with describing the problem, defining tradeoffs, setting boundary conditions, and the like. Only a human can make those choices, since the desires of the final user determine the answers, and no algorithm can guess what they will be. Once the overall goals and constraints are set, however, just finding a lens construction that minimizes a merit function seems like a proper job for a fast idiot. After that construction is found, the human again takes over, balancing aberrations in a manner suitable for the application. So the contest involves only one of many steps in this process—but a very important one.

The global optimization problem is conceptually simple. One assumes there is a great lens somewhere out there in design space, and your job is to find it. Since there is no closed-form solution, we are poking around in the dark in a universe of perhaps 30 or 40 dimensions. One cannot evaluate every possible design, at least not in a human lifetime, so we simplify.

Algorithms to be run on a computer cannot apply sophisticated knowledge of aberration theory, as do experts like Shafer. That approach works well for those with sufficient skills and experience, but a process that works for a given lens is not likely to apply wholesale to any other problem—and a commercial code like SYNOPSISTM has to work for every lens you throw at it. And not all commercial users have senior-level skills anyway. Another approach is needed.

A common practice is to create an enormous mesh of nodes, where each design variable takes successive values, the number of combinations to be evaluated often exceeding 200,000. Of course that procedure can find any design, in principle, but paraphrasing Mark Twain, “I want to still be young when we get there”. It seems pointless to give the customer a program that will run all weekend—and then tell him that there is no good solution for those requirements and he has to run it again, perhaps with different weights. Is there a faster way? One hopes so.

This quandry lead me to try to envision the geography of the lens design landscape, hoping to gain insights that might lead to a faster procedure. I noted in a previous paper² how, starting with plane-parallel surfaces, one can generate a wide variety of lens forms. Indeed, that procedure works so well that we are tempted to interpret the design space as a mountain range. If you want to find the lowest valley for miles around, it's easy to see if you stand on the highest peak. That peak, we conjecture, is a lens of plane-parallel plates. One can then, according to our engaging metaphor, simply slide downhill and find a minimum somewhere. Changing the direction in which you start out will usually result in a different minimum, and by trying a great many directions one might find them all. Although tempting, this metaphor clearly is too simple, since we cannot visualize a mountain in dozens of dimensions. Nonetheless, we programmed it in a feature we call DSEARCHTM³, which we desired to test and, we hoped, to break.

When Dave gave me his 11-element design problem, I asked DSEARCH to find a solution, and it came back with the lens in Figures 10 and 11. So far, so good.




Figure 10. 11-element design from DSEARCH




Figure 11. MTF of the 11-element design found by DSEARCH.

DSEARCH has a variety of optional modes, some of which execute very quickly, while others are more thorough. In random mode, the design space is sampled in a random manner for a requested number of cycles. This can be slow if that number is large, and for this test we preferred instead a binary search option, described more fully below. In order to make the program run as quickly as possible, there is also an option to either optimize every one of the combinations completely or first perform a quick screening pass where the merit function consists only of 3rd and 5th-order aberrations plus three real rays. A selected number of the lenses that survive this pass are then subjected to a more complete optimization with transverse ray intercepts. This greatly speeds things up, but of course one might miss a good configuration if it does not happen to survive the initial screening pass. Indeed, forcing the program to favor designs with those aberrations corrected can steer it away from better configurations because the required balance of higher-order

aberrations may then not be attainable. If directed, the program can submit all of the finalists to the simulated annealing program, to better evaluate the potentials of the winners.

Lastly, to maximize speed, we can run on more than a single core. A job that runs in, say 40 minutes on a single core, will run in just over five minutes on an 8-core machine. For purposes of comparison, the timings reported here are for a single-core run.

For the 11-element design we selected the following options:

- Binary search method
- Quick screening pass
- Simulated annealing of best candidates.

The options selected took about 25 minutes for this task on our 3.1 GHz PC.

The binary-search option of DSEARCH assigns an initial power to the otherwise flat surfaces according to a binary number, where, for an 11-element lens, there are 2048 combinations of positive and negative power. This can be viewed as a very tiny sub-space of the global design mesh mentioned above. The powers that are initially assigned depend on a user-selected radius of curvature, applied so the bendings are all initially zero. We allowed only two glass types, following Dave's rules, and the algorithm assigned the lower dispersion glass U-BK7 to positive elements, and LF5 to the negative. (The algorithm normally assigns a glass model, which is itself varied in optimization.)

Our very simple algorithm came back with a variety of design forms, the best of which are shown in Figure 12; the top three are all pretty good. We really don't like jobs that run all weekend, and our algorithm seems to be very effective in that regard. Perhaps this was too easy. We wanted to break the program, but it survived. We needed a harder problem.

Dave then pointed the way by observing that one can get by with fewer elements. Indeed, an expert might guess just that by looking at the element at surface 15 in Figure 10. When we tried DSEARCH with 10 and then nine elements, we got equally encouraging results. Still too easy.

Not until we tried the 8-element problem did things start breaking—which was exactly what we were looking for. This is indeed a hard problem. Not to be outdone by a mere human, we found a very good design after running DSEARCH in its slowest mode: random search over 5000 cases, full optimization of every candidate, and some annealing at the end. The design, shown in Figures 13 and 14, is very good, but the run took over eight hours. The statistics of that run are shown in Figure 15, where we see a great many local minima, most not very good, and only one or two great designs over at the left side. Of course the machine will win if you run it long enough—but we want a faster return on our investment. If we stopped here, we would declare the human the winner because we are not willing to wait that long.

So we tried several of the quicker modes—and got nowhere. The lenses that came back were of the “not bad” quality, but not as good as Shafer's design. As Rudolph Kingslake has said, designing a good lens is not difficult—but designing a *very* good lens is much harder. This was the test we were after!

It was time to examine our metaphor again. The direction in which one starts out from the mountaintop is set by a binary number, as noted above, but the *distance* from the top where one starts is given by a user-defined starting radius of curvature. Could a poor choice of that parameter be impeding things?




Figure 12. Example lenses returned by DSEARCH for the 11-element problem. The top three are all pretty good.




Figure 13. Eight-element lens found by a very slow procedure. This run took about eight hours.




Figure 14. MTF of the 8-element lens in Figure 13.

It was time for some research. Dave's excellent design was of the form (P N N P P P N P), so we investigated how the best merit function value that came back (with that construction) depended on the initial value of the radius. This was interesting, and a typical evaluation is shown in Figure 16.




Figure 15. Statistics of the lengthy DSEARCH run. Most of the peaks identify local minima, found by several trials. At the extreme left side of the curve is a good design.




Figure 16. Effect of initial radius selection on best merit function.

In this plot we see that starting out with an initial radius that is either too short or too long results in poor performance. That makes sense: if the radius is too short, many combinations will cause ray failures in the first few elements, while if too long, rays are likely to fail at the last surface, which was assigned a curvature solve to control the F/number. The program can recover from ray failures as noted in our earlier paper, but the design then often moves away from the minimum we are after. There seems to be a sweet spot centered around a radius of 600 mm.

So we tested DSEARCH again , specifying that initial value, and learned another lesson: a seemingly small change to the search parameters often produced very different results. The option to perform simulated annealing, and the value assigned to the relative aperture weighting of targetted rays, all seemed to strongly affect the direction the process took. Some lenses were excellent, as shown in Figures 17 and 18.




Figure 17. Very good 8-element design found by DSEARCH. This run took about 11 minutes.




Figure 18. MTF of the design in Figure 17.

The annealing option inserts some randomness into the process, and we were able to find several other excellent solutions when it was employed, one of which is shown in Figures 19 and 20.




Figure 19. One of several very good 8-element designs found by DSEARCH. This run took about one hour.




Figure 20. MTF of lens shown in Figure 19.

We should point out that DSEARCH is not intended to produce a finished lens design; its purpose is to identify attractive starting points, and each of the designs shown here was then subjected to further optimization, first with transverse intercepts, then with OPD targets, and lastly with MTF targets in the merit function. That is not cheating, since the steam drill found the configuration and we merely improved it somewhat.

Here we have learned something important: perhaps the geography looks more like a battlefield from WWI—with trenches snaking all over and shell craters here and there—than a mountain. If the best design is at the bottom of a crater with raised edges, then of course one cannot get there by sliding downhill. In that case, is there still a fast way to find it? Apparently, one has to allow some randomness, which allows one to jump out of the groove one is in and land in a better one.

There seem to be quite a few very desirable craters to explore. The lens in Figures 21 and 22 shows yet another construction—and we note that the program came up with a 7-element version that is essentially identical to Shafer’s design, after a run of 7.5 minutes. But here we ran into a deeper problem: to get the 7-element result we had to alter the aperture-dependent weighting to favor the center rays even more strongly than before. A naïve user might not know that, so in a purely man-machine contest, we could say that the machine (on its own) lost this round.




Figure 21. One of several very good 8-element designs found by DSEARCH, when some randomness is allowed. This run took about 16 minutes.




Figure 22. MTF of the lens in Figure 21.

Proceeding in this manner, we found that there are in fact a great many configurations that meet the requirements and whose MTF is as good or better than that of the human design. Figure 23 shows some examples. One lesson from this exercise is that, for this problem as for many others, one should not think of the design space as having only one “global optimum”, because there are usually many others. In practice, a human designer would likely stop when he had reached

any one of these solutions, whereas a design search program can produce many others, some of which might better fit the needs than the first one that was found.




Figure 23. Some of the many solutions found by DSEARCH that were as good as or better than the human-designed lens.

We learned a second lesson too: there appears to be some conservation-of-effort law at work, at least on this problem. A run that took 20 minutes usually found a good solution—while a shorter run, using some of the acceleration methods described above and requiring only five minutes, worked about one time out of four. We learn why by comparing the convergence curves for two different trial examples. Figure 24 shows the curve for a well-behaved lens on the left, while on the right is a more chaotic case. If one were evaluating these two cases with only a few cycles, in the interest of speed, then after 20 passes the leftmost would seem the better and the rightmost lens might well be eliminated—even though it turns out to be the superior design when optimized for a greater number of passes.




Figure 24. Comparison of the convergence curves for a well-behaved lens (on the left) and a chaotic lens. A quick screening pass would likely rule out the rightmost of these two, even though it is the better lens when fully optimized.

4. CONCLUSIONS

This study has been very interesting. Dave is impressed that a mere computer can, in some cases, come up with a design equal to or better than what an expert can produce. I am impressed that a mere human can come up with a design that is a real challenge for even the best algorithms. Perhaps we both win.

But every algorithm depends on cogent selection of input options by the user, so in that sense, if John Henry had used his head instead of his hammer, he would have applied for the job of driving the steam drill and would have enjoyed a comfortable retirement.

[¹] Shafer, D. R., "Doing more with less," Proc. SPIE 2537, 2-12 (1995).

[²] Dilworth, D. C., "New tools for the lens designer," Proc. SPIE **7060**, 70600B (2008).

[³] Dilworth, D. C., "Novel global optimization algorithms: binary construction and the saddle-point method," Proc. SPIE **8486**, 84860A (2012).