CHAPTER 9

The Optical Design Process

The optical design process includes a myriad of tasks that the designer must perform and consider in the process of optimizing the performance of an imaging optical system. While we often think primarily of the robustness of the optimization algorithm, reduction of aberrations, and the like, there is much more to do. The designer must be at what we sometimes call "mental and technical equilibrium with the task at hand." This means that he or she needs to be fully confident that all of the following are understood and under control:

- All first-order parameters and specifications such as magnification, focal length, f/number, full field of view, spectral band and relative weightings, and others.
- Assure that the optical performance is being met, including image quality, distortion, vignetting, and others.
- Assure that the packaging and other physical requirements, including the thermal environment, is being taken into account.
- Assure that the design is manufacturable at a reasonable cost based on a fabrication, assembly, and alignment tolerance analysis and performance error budget.
- Consider all possible problems such as polarization effects, including birefringence, coating feasibility, ghost images and stray light, and any other possible problems.

Once every one of these items has been addressed and is at least recognized and understood, we start with the sketch of the system. First, the system is divided into subsystems if possible, and the first-order parameters are determined for each subsystem. For example, if we are to design a telescope with a given magnification, the entrance pupil diameter should be chosen such that the exit pupil size matches the eye pupil. A focal length of the objective and the eyepiece should be chosen such that the eyepiece can have a sufficiently large eye relief. Now, when the specs for each subsystem are defined, it is time to use the computer-aided design algorithms and associated software to optimize the system, which will be discussed in the rest of this chapter. Each subsystem can be designed and optimized individually, and the modules joined together or, more often, some subsystems are optimized separately and some as an integral part of the whole system.

What Do We Do When We Optimize a Lens System?

Present-day computer hardware and software have significantly changed the process of lens design. A simple lens with several elements has nearly an infinite number of possible solutions. Each surface can take on an infinite number of specific radii, ranging from steeply curved concave, through flat, and on to steeply curved convex. There are a near infinite number of possible design permutations for even the simplest lenses. How does one optimize the performance with so many possible permutations? Computers have made what was once a tedious and time-consuming task at least manageable.

The essence of most lens design computer programs is as follows:

- First, the designer has to enter in the program the starting optical system. Then, each variable is changed a small amount, called an *increment,* and the effect to performance is then computed. For example, the first thickness may be changed by 0.05 mm as its increment. Once this increment in thickness is made, the overall performance, including image quality as well as physical constraints, are computed. The results are stored, and the second thickness is now changed by 0.05 mm and so on for all variables that the user has designated. Variables include radii, airspaces, element thicknesses, glass refractive index, and Abbe number. If you are using aspheric or diffractive surfaces, then the appropriate coefficients are also variables.

- The measure of performance as used here is a quantitative characterization of the optical performance combined with a measure of how well the system meets its first-order constraints set by the user such as focal length, packaging constraints, center and edge thickness violations, and others. The result of the computation is a single number called an *error function* or *merit function.* The lower the number, the better the performance. One typical error function criteria is the rms blur radius, which, in effect, is the radius of a circle containing 68% of the energy. Other criteria include optical path difference, and even MTF, as described in Chap. 15.
- The result is a series of derivatives relating the change in performance (P) versus the change in the first variable (V_1), the change in performance (P) versus the change in the second variable (V_2), and so on. This takes on the following form:

$$\frac{\partial P}{\partial V_1}, \frac{\partial P}{\partial V_2}, \frac{\partial P}{\partial V_3}, \frac{\partial P}{\partial V_4} \ldots$$

- This set of partial derivatives tells in which direction each parameter has to change to reduce the value of the sum of the squares of the performance residuals. This process of simultaneous parameter changes is repeated until an optimum solution is reached.

A lens system consists of a nearly infinite number of possible solutions in a highly multidimensional space, and it is the job of the designer to determine the optimum solution.

Designers have used the following analogy to describe just how a lens design program works:

- Assume that you cannot see and you are placed in a three-dimensional terrain with randomly changing hills and valleys. Your goal is to locate the lowest elevation or altitude, which in our analogy equates to the lowest error function or merit function. The lower the error function, the better the image quality, with the “goodness” of performance being inversely proportional to the elevation.
- You are given a stick about 2 m long, and you first stand in place and turn around tapping the stick on the ground trying to find which direction to walk so as to go down in elevation.

- Once you determine the azimuth resulting in the greatest drop in elevation, you step forward in that direction by 2 m.
- You now repeat this process until in every direction the elevation goes up or is level, in which case you have located the lowest elevation.
- But what if just over a nearby hill is an even lower valley than you are now in? How can you find this region of solution? You could use a longer stick, or you could step forward a distance several times as long as the length of your stick. If you knew that the derivative or slope downward is linear or at least will continue to proceed downward, this may be a viable approach. This is clearly a nontrivial mathematical problem for which many complex and innovative algorithms have been derived over the years. But the problem is so nontrivial as well as nonlinear that software algorithms to locate the so-called global minimum in the error function are still elusive. Needless to say, the true global minimum in the error function may be quite different or distant from the current location in our *n*-dimensional terrain.

Figure 9.1 shows a two-dimensional representation of solution space as discussed previously. The *ordinate* is the error function or merit function, which is a measure of image quality, and the *abscissa* is, in effect, solution space. We may initiate a design on the left and the initial

Figure 9.1
Illustration of Solution Space in Lens Design

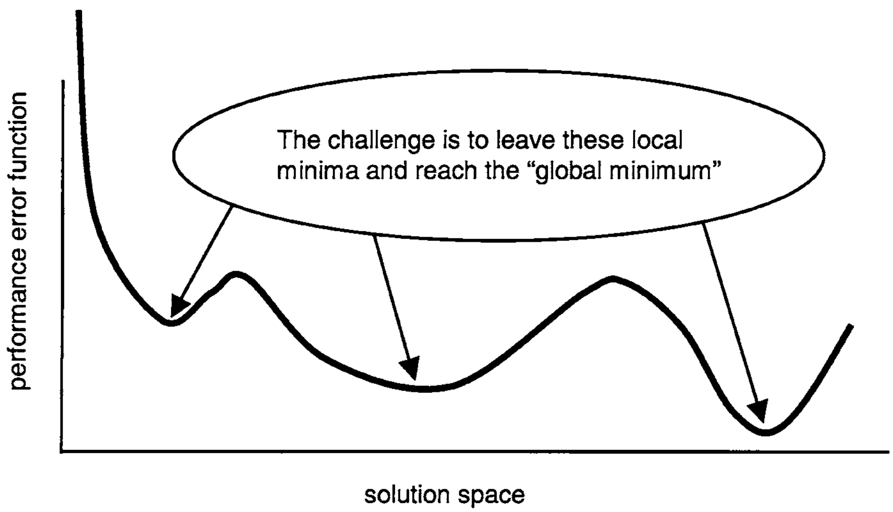

optimization brings the error function to the first minimum called a *local minimum* in the error function. We then change glasses and/or make other changes to the design and ultimately are able to move the design to the next lower local minimum. Finally, we add additional elements and make other changes and we may be able to reach the local minimum on the right. But how do we know that we are at, or even close to, a global minimum? Here lies the challenge as well as the excitement of lens design!

It is important here to note that reaching global minimum in the error function is not necessarily the end goal for a design. Factors including tolerance sensitivity, packaging, viability of materials, number of elements, and many other factors influence the overall assessment or "goodness" of a design. Learning how to optimize a lens system is, of course, quite critical to the overall effort, and learning how to reach a viable local or near-global minimum in the error function is very important to the overall success of a project.

How Does the Designer Approach the Optical Design Task?

The following are the basic steps generally followed by an experienced optical designer in performing a given design task. Needless to say, due to the inherent complexity of optical design, the processes often become far more involved and time consuming. Figure 9.2 outlines these basic steps:

1. The first step in the design process is to *acquire and review all of the specifications.* This includes all optical specifications including focal length, *f*/number, full field of view, packaging constraints, performance goal, environmental requirements, and others.
2. Then we *select a representative viable starting point.* The starting point should, wherever possible, be a configuration which is inherently capable of meeting the specifications for the design. For example, if the specifications are for an *f*/10 monochromatic lens covering a very small field of view and having an entrance pupil diameter of 5 mm, then the lens may very well be a single element. However, if the requirements call for an *f*/1.2 lens over a wide spectral band covering a 40° full field of view, then the solution may very well

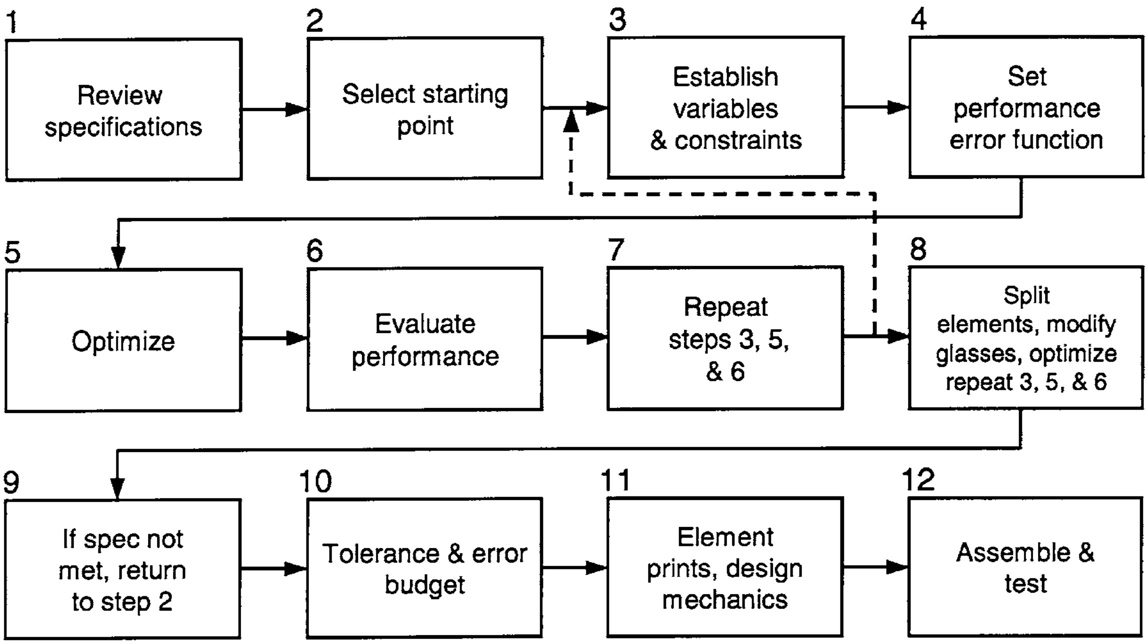

Figure 9.2
Lens Design and Optimization Procedure

be a very complex six- to seven-element double Gauss lens form. If we were to use a single element for this latter starting point, there would be no hope for a viable solution. Finding a good starting point is very important in obtaining a viable solution. The following are viable sources for starting points:

You can use a *patent* as a starting point. There are many sources for lens patents including Warren Smith's excellent book *Modern Lens Design*. There is also a CD-ROM called "LensView," which contains over 20,000 designs from patents. These are all searchable by a host of key parameters. While the authors of this book are not patent attorneys, we can say with confidence that you may legally enter design data from a patent into your computer and work with it in any way that you would like to. If your resulting design is sold on the market, and if the design infringes on the patent you used (or any other for that matter), you could be cited for patent infringement. It is interesting to note that the purpose of our patent system in this country is to promote inventions and innovation. This is done by offering an inventor a 17-year exclusive right to his or her invention in exchange for *teaching* in the patent how to implement the invention. Thus, you are, in effect, invited to use the design data and work with it with the goal of coming up with a better design, which you can then go out and patent. By this philosophy, inventors are constantly challenged to improve upon an invention, which, in effect, advances technology, which is what the patent process is all about. Needless to say, we urge you to be careful in your use of patents.

You could use a so-called hybrid design. We mean a hybrid to be the combining of two or more otherwise viable design approaches so as to yield a new system configuration. For example, a moderate field-of-view Tessar lens design form can be combined with one or more strongly negatively powered elements in the front to create an extremely wide-angle lens. In effect, the Tessar is now used over a field of view similar to its designed field, and the negative element or elements bend or "horse" the rays around to cover the wider field of view. An *original design* can, of course, be a viable starting point. As your experience continues to mature, you will eventually become comfortable with "starting from scratch." With today's computer-aided design software, this works most of the time with simple systems such as doublets and triplets; however, with more complex systems, you may have problems and will likely be better off resorting to a patent or other source for a starting point.

3. Once you have entered your starting point into the software package you are using, it is time to *establish the variables and constraints.* The system variables include the following: radii, thicknesses, airspaces, surface tilts and decenters, glass characteristics (refractive index and Abbe number), and aspheric and/or other surface variables, including aspheric coefficients. The constraints include items such as focal length, f/number, packaging-related parameters (length, diameter, etc.), specific airspaces, specific ray angles, and virtually any other system requirement. Wavelength and field weights are also required to be input. It is important to note that it is not imperative (nor is it advisable) to vary every conceivable variable in a lens, especially early in the design phase. For example, your initial design optimization should

probably be done using the glasses from the starting point, in other words do not vary glass characteristics initially. This will come later once the design begins to take shape and becomes viable. You may also want to restrict the radii or thicknesses you vary as well, at least initially. For example, if adjacent elements have a very small airspace in the starting design, this may be for a good reason, and you should probably leave them fixed. Also, element thicknesses are very often not of great value as variables, at least initially, in a design task, so it is usually best to keep element thicknesses set to values which will be viable for the manufacturer.

4. You now will *set the performance error function and enter the constraints.* Most programs allow the user to define a fully "canned" or automatically generated error function, which, as discussed earlier, may be the rms blur radius weighted over the input wavelengths and the fields of view. In the Zemax program the user selects the number of rings and arms for which rays will be traced into the entrance pupil (rays are traced at the respective intersection points of the designated number of rings and arms). Chapter 22 shows a detailed example of how we work with the error function.

5. It is now time to *initiate the optimization.* The optimization will run anywhere from a few seconds for simple systems to many hours, depending on just how complex your system is and how many rays, fields of view, wavelengths, and other criteria are in the system. Today, a state-of-the-art PC optimizing a six- to seven-element double Gauss lens with five fields of view will take in the order of 5 to 10 s per optimization cycle. Once the computer has done as much as it can and reaches a local minimum in the error function, it stops and you are automatically exited from the optimization routine.

6. You now *evaluate the performance* using whatever criteria were specified for the lens. This may include MTF, encircled energy, rms spot radius, distortion, and others.

7. You now *repeat steps 3 and 5 until the desired performance is met.* Step 3 was to establish the variables and constraints, and step 5 was to run the optimization, and these steps are repeated as many times as necessary to meet the performance goals. You will often reach a solution that simply does not meet your performance requirements. This is very common during the design evolution, so do not be surprised, depressed, or embarrassed if it happens

to you… it happens to the best of us. When it does happen, you may need to *add or split the optical power of one or more of the lens elements* and/or to *modify glass characteristics.* As we have discussed previously, splitting optical power is extremely valuable in minimizing the aberrations of a lens.

8. There is a really simple way of splitting an element in two, and while it is not "technically robust," it does work most of the time. What you do is insert two surfaces in the middle of the current element, the first of which will be air and the second is the material of your initial element. The thickness of each "new" element is one-half of the initial element and the airspace should be small, like 0.1, for example. Now simply enter twice the radius of the original element for both *s*1 and *s*2 of the new elements. You will end up with two elements whose net power sum is nearly the same as your initial element. You can now proceed and vary their radii, the airspace, and, as required, the thicknesses.
9. If you still cannot reach a viable design, then at this point you will need to *return to step 2 and select a new starting point.*
10. Your final task in the design process is to *perform a tolerance analysis and performance error budget.* We will be discussing tolerancing in more depth in Chap. 16. In reality, you should be monitoring your tolerance sensitivities throughout the design process so that if the tolerances appear too tight, you can take action early in the design phase and perhaps select a less sensitive design form.
11. Finally, you will need to *generate optical element prints, contact a viable lens manufacturer, and have your elements produced.* You will also need to work with a qualified mechanical designer who will *design the cell or housing* as well as any required interfaces. It is important to note that while we list the mechanical design as taking place at this point after the lens design is complete, it is extremely important to work with your mechanical designer throughout the lens design process so as to reach an optimum for both the optics as well as the mechanics. Similarly, you should establish a dialog with the optical shop prior to completing the design so as to have time to modify parameters which the shop feels needs attention such as element thicknesses, glass types, and other parameters.
12. Once the components are in house, you will need to *have the lens assembled and tested.* Assembly should be done to a level of precision and cleanliness commensurate with the overall

performance goals. Similarly, testing should be to a criterion which matches or can be correlated with your system specifications and requirements. We discuss testing in Chap. 15.

Sample Lens Design Problem

There was a very interesting sample lens design problem presented at the 1980 International Lens Design Conference. The optimized design for an $f/2.0$, 100-mm focal length, 30° full field-of-view double Gauss lens similar to a 35-mm camera lens was sent out to the lens design community. One of the tasks was to redesign the lens to be $f/5$ covering a 55° full field with 50% vignetting permitted. Figure 9.3 shows the original starting design, as well as the design after changing the $f/$number and field of view, without any optimization.

Sixteen designers submitted their results, and they spent from 2 to 80 h working on the problem. We will present here three representative solutions in Fig. 9.4. The design in Fig. 9.4*a* is what we often call *a happy lens.* What we mean is that the lens is quite well behaved with no steep bending or severe angles of incidence. The rays seem to "meander" nicely through the lens. It is a comfortable design. We show to the right of the layout a plot of the MTF. MTF will be discussed in detail in Chap. 10. For the purpose of this discussion, consider the MTF to be contrast plotted in the ordinate as a function of the number of line pairs per millimeter

Figure 9.3
Starting Design for Sample Lens Design Problem

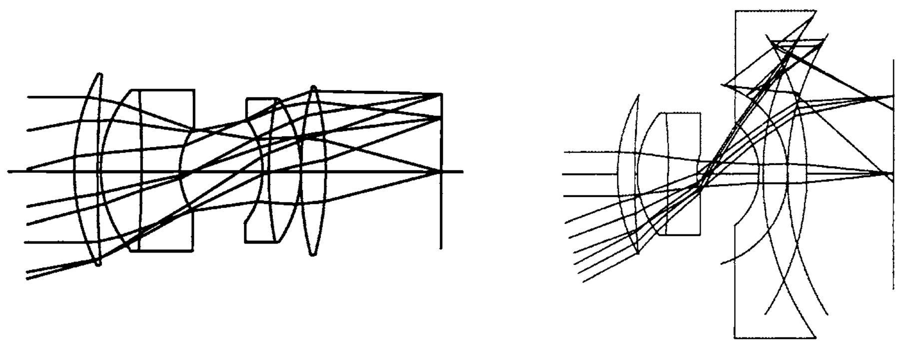

(a) (b)

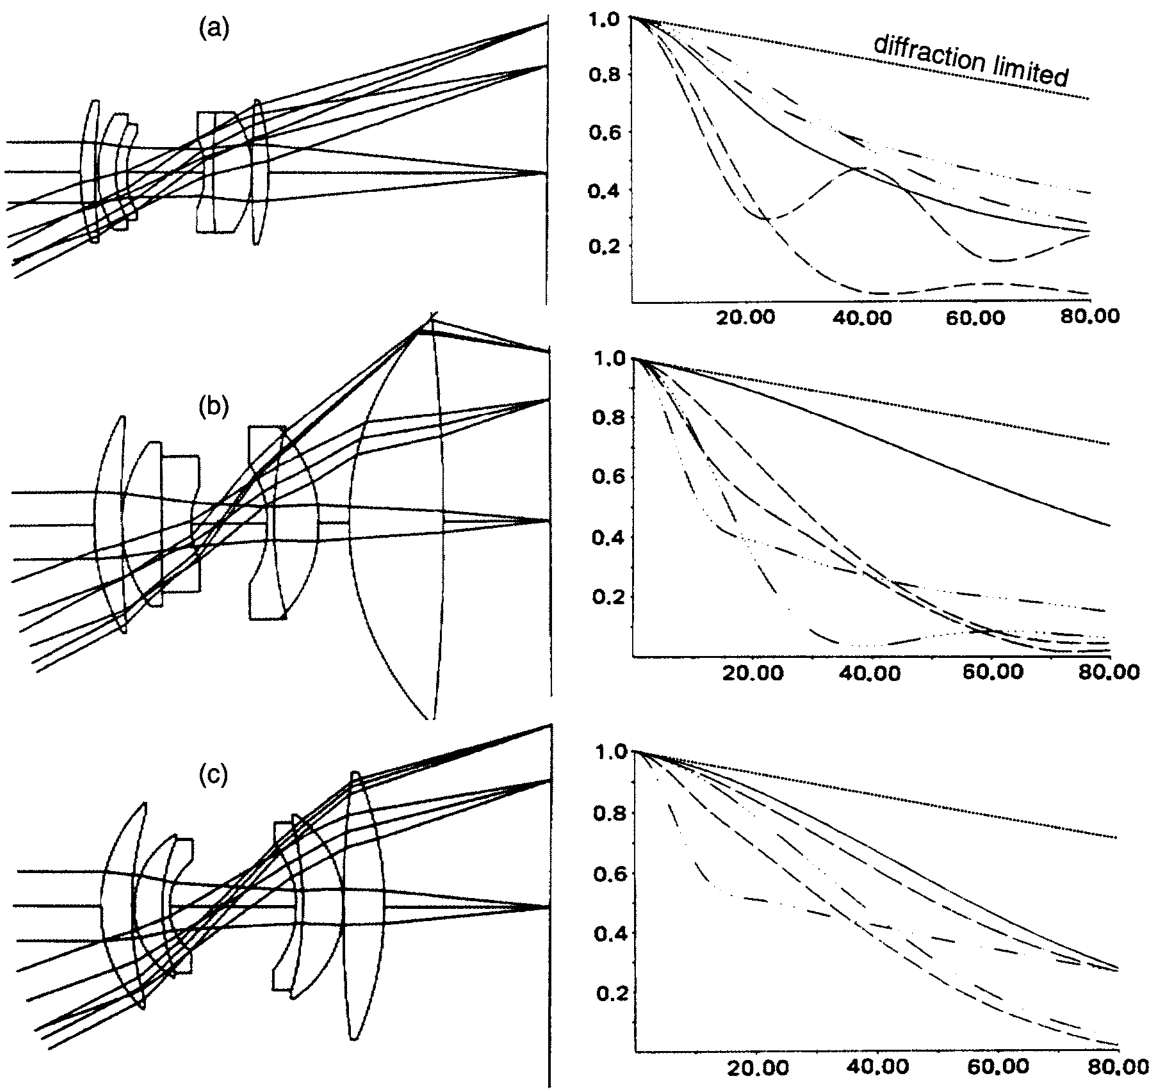

Figure 9.4
Representative Solutions to Sample Lens Design Problem

in the abscissa. The different curves represent different positions in the field of view and different orientations of the resolution patterns. The higher the curves, the better the contrast and the overall performance. The MTF is reasonable for most of the field positions. As will be discussed in Chap. 22, a good rule of thumb for the MTF of a 35-mm camera lens is an MTF of 0.3 at 50 line pairs/mm and 0.5 at 30 line pairs/mm.

The design in Fig. 9.4*b* has a serious problem; the rays entering the last element are at near-grazing angles of incidence. Notice that the exit pupil at full field is to the right of the lens (since the ray cone is descending toward the axis to the right), and at 70% of the field the exit pupil is to the left of the lens (since the ray cone is ascending to the right and therefore appears to have crossed the axis to the left of the lens). This is a direct result of the steep angles of incidence of rays entering the last element. The variation in exit pupil location described here would not itself be an issue unless this lens were used in conjunction with another optical system following it to the right; however, it does indicate clearly the presence of the severe ray bending which will inevitably lead to tight manufacturing and assembly tolerances. Further, the last element has a near-zero edge thickness which would need to be increased. The lens is large, bulky, and heavy. And finally, the MTF of this design is the lowest of the three designs presented.

Finally, the design in Fig. 9.4*c* is somewhat of a compromise of the two prior designs in that it is somewhat spread out from the design in Fig. 9.4*a* but does not have the problems of the design in Fig. 9.4*b*. The MTF of the design in Fig. 9.4*c* is the best of the three designs.

Comparing the three designs is very instructive as it shows the extreme variability of results to the same problem by three designers. The question to ask yourself is what would you do if you subcontracted the design for such a lens, and after a week or two the designer brought you a stack of paper 200-mm thick with the results of the design in Fig. 9.4*b*. And what if he or she said "wow, what a difficult design! But I have this fabulous solution for you!" Prior to reading this book, you might have been inclined to congratulate the designer on a job well done, only to have problems later on during manufacturing and assembly. Now, however, you know that there may be alternate solutions offering superior performance with looser tolerances and improved packaging. Remember that even a simple lens has a near infinite number of possible solutions in a multidimensional space.