Proc. of SPIE Vol. 0181, Contemporary Optical Systems & Components Specifications, ed. Robert E. Fischer (Apr 1979) Copyright SPIE

Keynote Address

Specifications as a driver of performance and cost

Warren J. Smith
Infrared Industries, Incorporated
Santa Barbara, California 93102

Good afternoon -

I see that, according to the program, I am scheduled to deliver what is billed as "The Keynote Address". Normally, I would assume that the syllable "key" in keynote rather implies an operation analagous to the opening of a door, or perhaps providing a guide, to the program delights which are to follow.

You may well wonder, then, at the timing of this talk. A keynote address given after ninety percent of the papers have been presented is at least an anomaly, and I suspect that it establishes some sort of record, perhaps even worthy of inclusion in the Guinness Book of Records. The reason for my unusual location on the program is this: Before Bob Fischer asked me to speak at this seminar, I had already scheduled two Optical Society presidential traveling lectures, one in Pittsburgh on Wednesday, the other in Boston last night, and this afternoon is the earliest that I could get to Washington.

Thus, in the course of doing my presidential duty to the O.S.A., I have had to miss the preceding twenty-five papers of this seminar. This I find highly unfortunate - on two counts - the first being that the program looks like a fine collection of interesting and informative papers which I certainly regret having missed. The second is that I am forced to forfeit the keynoter's traditional advantage, namely that, when he speaks before the program starts, as is customary, no one knows what wonders are yet to be unfolded by the speakers to follow. And the keynoter is free to seem prophetic, wise, all-knowing, and to enjoy a truly magnificent ego trip. On the other hand, someone in my position, which is what you might call that of a "ten percent keynoter", if you were to take my rank from the number of papers remaining, is at a terrible disadvantage. Literally everyone here now knows more about the program which I am keynoting than I do! At least, if I had heard the preceding papers on the program, I could use baseball parlance and bill myself as the clean-up speaker.

Under the circumstances, perhaps I may be pardoned if I regard the "key" in "keynote" as being that which fits the lock to the proverbial barn door.

The title of this talk was originally suggested by Bob Fischer, and I most assuredly must agree: Specifications do drive performance and cost -- usually upward in both cases. The real question is, "how?"; or perhaps "how much?" might be a better choice of phrase.

Actually, the performance required of a system is usually determined by the application for that system, and to a certain extent, the cost is determined by the application as well. So perhaps my title is inverted. It is performance which is, or should be, the driver of the specifications. And, the the best of our ability as engineers, the system should be toleranced so as to produce the minimum acceptable performance with which the application will still function. In this sense, then, performance also drives cost.

In these days of ever-improving optical designs and higher and higher system performance, one can sympathize with a manufacturer of optics if he occasionally feels that the optical designer is dedicated to testing the state of the art in precision optical fabrication.

If we assume that the optical prescription has been established, then the biggest driver of system performance is, of course, the fabrication tolerance budget, that is, the allowable production variation of the constructional parameters of the lens: the value and uniformity of the surface curvatures, the thicknesses and spacings, the refractive indices and the like. The question is, how does one rationally arrive at the tolerance budget, with some assurance that it will produce a useful system, and with the additional hope that this successful system performance will be accomplished at something reasonably approximating a minimal cost?

At best, this is not a minor task. The tolerancing of an optical system has much in common with optical design. The tools tend to be the same, the criteria used to judge success are similar and, as we shall see, the methods and techniques have paralleled each other in the history of their development, although tolerancing has lagged behind design by about a decade. Let's take a look at this development by considering three periods, and the techniques used in each of them.

In the early days (and for the benefit of the younger members of the audience, this means a period well back into the history of optical engineering, say about twenty-five or thirty years ago), a lens designer might, or might not, have what was called a "change table". This was a tabulation of the changes in a few of the aberrations produced by corresponding changes in the constructional parameters of the lens system. If he did have such a table, it was probably rather rudimentary, it was very likely incomplete, and it was almost certainly obsolete, in the sense that it had probably been calculated several design iterations before the final one.

With such a change table, the designer could very crudely estimate the effects of parameter variations on his system, at least to the extent that he understood the somewhat devious relationships between the system performance and the primary aberrations, which were all that he computed in those days. However, given the slow computation techniques of the period, there was rarely enough time to perform a complete analysis and, let's face it, most tolerancing was actually done on a seat-of-the-pants basis. A tolerance budget would simply be based on other, previous budgets which had proven reasonably successful in past applications of a similar nature. Perhaps they would be touched up by a brief mathematical investigation of one or two especially critical parameters, which were typically the more sensitive spacings or thicknesses.

The advent of the Optical Transfer Function as a means of evaluating the performance of an optical system was a great step forward in many ways. In the field of specifications, its greatest value was perhaps in the fact that the relationship between the simpler forms of image aberrations and the MTF of the system could be readily visualized and computed. Thus, for example, graphs became available which related MTF and spatial frequency and the amount of an aberration, and it became quite apparent to even the most crusty old ray benders that the most universal common denominator for the effect of an aberration was when its magnitude was expressed in terms of the wavefront deformation. Thus one could express the spherical aberration as being X wavelengths of OPD, or Optical Path Difference, and have a reasonable basis for comparing its effect with that produced by Y wavelengths of coma. Further, one could say that a half wavelength of aberration would reduce the MTF at some frequency of interest by a known amount, provided of course that the aberrations were of a low order and that one were willing to be reasonably approximate about the whole thing. And, of course, one could also express surface fabrication errors in the same way.

Some ten or twelve years ago, I wrote a set of optical design programs for a time sharing computer.(1) The completed set included a tolerance analysis program which calculated the partial derivatives of the third order aberrations with respect to the constructional parameters and decentrations of an optical system. The program output was given in both conventional transverse aberrations and in wavefront deformations. Thus, one could readily determine that, for example, thickness number three introduced so many hundredths of a wave of OPD of, say, spherical aberration per millimeter change of thickness from the design nominal. If one multiplied these differential coefficients by the tolerance assigned to each constructional parameter, one could then determine, in a simple-minded, low order sort of way, the effect of the tolerances on the wavefront aberration.

And, after a little study, I convinced myself that one could reasonably take the square root of the sum of the squares of all the calculated tolerance effects, a process now called, inelegantly enough, "RSS-ing", and arrive at a statistical prediction that a very large percentage of all the optical assemblies constructed according to a given set of tolerances would have image defects less than the RSS number. And, by doing the work in wavelengths of OPD, one could easily arrive at a straightforward correlation with the modulation transfer function.

The establishment of a tolerance budget from this sort of data was totally manual, and at least partially intuitive. You simply assigned the tightest tolerances to those parameters with the largest OPD derivatives, and by carrying out an elaborate balancing act, with much trial and error, arrived at a budget which would hopefully permit a functional system to be fabricated at a bearable cost. You kept the individual tolerances on the more sensitive parameters as far as possible above the tightest level that your shop could hold, that is, as far above as the lens system performance requirements would permit, and you also held the tolerances on the insensitive dimensions at or below those levels where any further increase in the tolerance ceased to reduce cost.

The notion of a bearable cost means of course that you must have some idea of how the various tolerances relate to the cost of fabricating the optical system. This is a matter that, even given the specific environment of a particular optical shop, is subject to an unbelievable number of interrelationships, so that it is difficult to be exact. However, for better or worse, and with absolutely no pretensions of infallibility, here are a few of the more common relationships as they exist for one particular production-oriented optical shop.

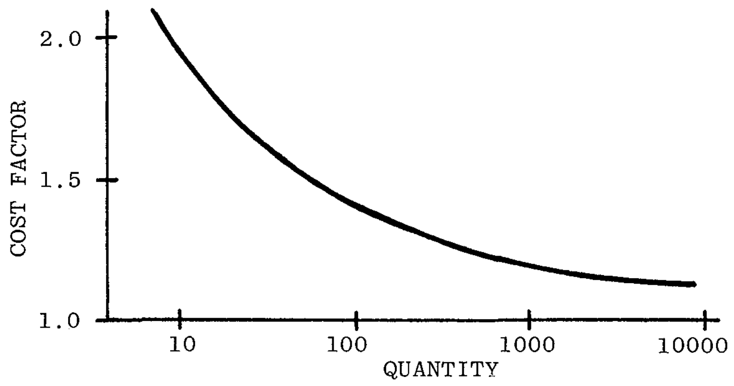

Fig. 1. Relationship between production cost and size of production run.

Fig. 1 and Eqn. (1) illustrate the relationship between the cost of production and the quantity (Q) of pieces made in a production run. This cost factor is frequently expressed as a scrap allowance, a set-up charge, and a learning curve.

$$\text{Quantity Cost Factor} = 1.07 + 2.26\ Q^{-0.42} \qquad (1)$$

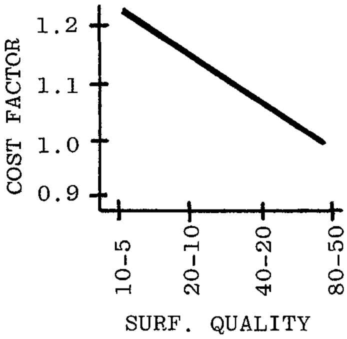

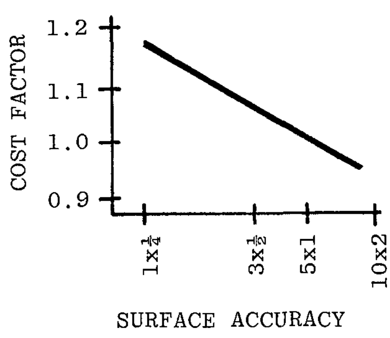

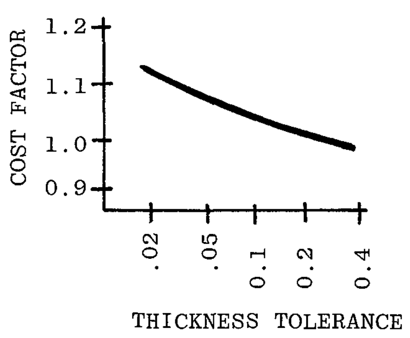

Fig. 2. Relationship between production cost and surface (cosmetic) quality.

Fig. 3. Relationship between production cost and surface accuracy of figure.

Fig. 4. Relationship between production cost and thickness tolerance.

Equation (2) indicates the cost effect of several quality and accuracy factors.

$$\text{Quality Cost Factor} = \frac{1.5}{\sqrt[20]{S \times D \times P \times R \times T}} \qquad (2)$$

where S and D are the conventional scratch and dig specification numbers, P and R are the surface accuracy (test glass reading) specifications, P being the power tolerance in fringes and R being the regularity or asphericity tolerance in fringes, and T is the center thickness tolerance in millimeters. The individual factors are broken out and shown in Figs. 2, 3 and 4.

In addition to these, one might note that a concentricity requirement that necessitates the use of equipment or techniques more precise than conventional "bell" or "cup" centering can raise the cost by about twenty-five percent, and that multilayer (high efficiency antireflection) coatings have a similar cost effect. And finally, the effect of governmental procurement and quality control techniques (e.g., the resident source inspector) is to raise costs by a factor that frequently approximates two.

Among the limitations on the validity of this sort of data is the fact that there are significant cross products, in that one factor will have a different effect, depending on the status of some other factor. Bear in mind that these relationships vary with many outside factors as well, and they may be completely different for another shop, especially a shop with a different operating philosophy or orientation.

Now, to bring us up to the present era, let me briefly decscribe a tolerancing program which is now offered by Optical Research Associates in Pasadena. Key elements of this program were discussed by Rimmer(2) and Koch(3) at an S.P.I.E. seminar last year. The program directly computes the change in the MTF which is caused by a dimensional change, or tolerance. It calculates this without recourse to the low order approximations, and without recourse to the MTF vs. O.P.D. graphs which I mentioned previously. It goes one step beyond the assumption that the relation between a tolerance and MTF is linear; it calculates the relationship to the second order. And it also calculates the cross products -- that is, the change in the value of one partial differential produced by changing some other parameter. The calculation is not based on finite differences -- that is, it does not change a parameter by a small amount to see what happens. It is based on a Taylor series expansion about the nominal system.

As if all this were not enough, this scheme has now been built into a truly automatic tolerancing program. In simplified terms, (a) a target MTF level at some frequency is given, (b) bounds are established for tolerances and, (c) any dimensions which may be adjusted in assembly are indicated. The tolerance bounds are essentially those which I mentioned a moment ago, the lower limit being the smallest tolerance to which your shop can work, and the upper limit being that tolerance value beyond which there is no cost saving. The assembly adjustments almost always include refocussing, and may also include one or two spaces or other features which can be adjusted.

The program presents you with a set of tolerances for the standard constructional parameters, accurately rounded to one or two places:- radius, thickness, index, decentering, irregularity, inhomogeniety, and the like. You also get the statistics of the effects on the Modulation Transfer Function: The standard deviation and a prediction that such and such a percentage of the assemblies will meet or exceed a certain MTF specification. This process has a great similarity to automatic optical design, and the final result is arrived at by a sort of simulated Monte Carlo analysis, but at considerably less cost than a true Monte Carlo, which hardly anybody can afford.

We have certainly come a long way in a couple of decades. Probably, if I hadn't heard other speakers say it so many times in those decades, I would at this point utter some foolish nonsense about "everything in the optical design field having been discovered!". But each time some pundit intones those words, someone else seems to come along and prove him wrong, so I shall restrain myself.

But I certainly can't leave the subject of specifications without at least a passing swipe at my favorite - the Scratch and Dig Specs. We have been discussing specifications as drivers of performance and cost. The constructional, or dimensional, parameter tolerances of a system most assuredly do impact both. But, for most systems, the Scratch and Dig Specification represents a true singularity; it is the only common specification which drives cost, but not performance.

In conclusion, may I leave you with two wishes: If this has been one of the latest keynote addresses on record, I hope that it has also been one of the shortest and least painful. And, secondly, may all your tolerances be functional. Thank you.

References

(1) Smith, W. J., "Optical Computations Using a Time Sharing Computer", Applied Optics, Vol. 6, p. 585 (L) 1967

(2) Rimmer, M. P., "A Tolerancing Procedure Based on MTF", in SPIE Proc. Computer Aided Optical Design, 147-11, 1978.

(3) Koch, D. G., "A Statistical Approach to Lens Tolerancing", in SPIE Proc. Computer Aided Optical Design, 147-12, 1978.