GAGING FOR SPC: KEEPING IT SIMPLE

After all that has been written about Statistical Process Control (SPC) in the past few years, I am continually surprised at how often shop owners tell me they would like to do SPC, but they're "just too small," or they "can't afford all that fancy equipment." There seems to be a mistaken yet growing perception out there: 1) that SPC is a lot more complicated than it is, and 2) that you need a computerized system on line just to get into it.

This is unfortunate because small shops can often benefit the most from SPC and there really is no reason they should not. The simple fact is, you don't need a computer and you don't need to be a statistical genius to do SPC. I know this latter for a fact, because I understand it. All you need, really, is an indicating gage, a pad and a pencil. So if you're one of those who is still hesitant about charting, try this easy-to-follow recipe for small shop SPC:

Before you start, you do have to understand a little bit about the process itself. But this is not difficult. Bear in mind that the basic principles of SPC were developed back in the '30s -- long before computers were invented -- and have not really changed since. So don't be intimidated by all the bells and whistles. An X/R chart then was the same as an X/R chart now. You don't need to understand all the theory behind the process, just the basics of frequency distribution and charting will do for a start, and there are many guides to help do this. (In fact, Federal Products published one back in 1945 which is currently in its 14th printing and is still in use!) What you're basically looking at is taking a few simple measurements, averaging them, then recording the results.

Next, get the process in place. This involves three steps. First, look at the part you're going to measure. What are its important dimensions and what part of the process controls them? Here is what you want to measure and where you want to start the process. But keep it simple. Start with a single measurement until you get the feel for it.

Second, look at the gaging equipment you are going to use. Make sure you follow all the gaging basics we've talked about in this column. Make sure it is the right gage for that type of measurement, and that you follow the ten-to-one rule for resolution (to measure 0.001" tolerance, you need a gage with at least 0.0001" resolution). You should also run a series of Gage Repeat and Reliability (GR&R) studies to make sure your measurements are as accurate as possible. Remember, whatever "analysis" you do can only be as accurate as the measurements you start with.

The final step in setting up is to benchmark your machining process. You need to find out, simply, if your machine is capable of holding the tolerances or control limits you require. Your SPC guide can help you do this.

Now you're ready to gather data. There are many sophisticated systems available to facilitate this process. However, most trainers agree you are much better off starting manually. With automated data collection systems, the operator often feels left out of the process: it happens without him. He doesn't need to think about it, and therefore makes no effort to understand it. This is bad because, ultimately, it's not SPC that controls your quality, it's the operator. If he doesn't understand the process, he is unable to use it.

Charting manually, on the other hand, puts him immediately in contact with the process. He is able to see that there is a relationship between what he's doing and what the charts show. He sees that his process changes and that the charts provide a prediction of how it changes. He is then able to control it. He sees his wheel is wearing, for example, and knows when, and when not, to compensate. Charting manually -- or at least visually -- empowers the operator, and that, ultimately, is what SPC is all about.

The other advantage of manual charting is that you can use just about any type of indicating gage, so long as it is accurate. Since you probably already have some in your shop, your SPC investment cost is reduced to the Manual, the paper and the pencil. Who knows, you may even already have the pencil.

On the downside, manual charting is tedious and time consuming, and subject to error. Whenever an operator measures and records manually, he has two chances for error: he can observe wrong and he can record wrong. This is where you may want to consider a digital indicator, which can give you several advantages. First, you can eliminate error by gathering and recording data electronically -- as well as interfacing with whatever computerized system you may want to put on line. Second, it makes it quicker and easier to collect the readings: all the operator has to do is measure and press a foot switch. Battery operated gages provide a back-up in case your computer system does go down and, most importantly, these gages still provide the operator with an on-site reading so he can monitor his process.

That's it: simplified SPC for the small shop. Start small, take it one step at a time, and keep your operators involved. And remember, SPC is not a thing you buy, it's a thing you do.

BEDROCK S.Q.C.

Statistical Quality Control has been in use since the 1930s, when it was performed with paper, pencil, and maybe a slide rule. As such, it was originally somewhat time consuming, and it required a fairly high level of care and understanding on the part of its practitioners. SQC got a big boost in the 1970s and '80s, when electronic gages, data loggers, and PCs began to proliferate. Suddenly, untrained line inspectors could easily perform the necessary calculations, without really understanding the process.

Many instructors still believe, however, that inspectors and machine operators should learn to do SQC manually before they plug in their data loggers, on the general principle that people who understand what they're doing tend to do a better job. Thus, we'll be looking at some of the basics of SQC for the next month or two. The rules can apply to any dimensional gaging procedure in a high-volume production application.

All operations that produce features to dimensional tolerances involve variation. Variation cannot be eliminated, but it can be controlled so that it remains within acceptable limits. SQC uses the laws of probability to reliably monitor and control a process. By inspecting variation in a small sample of production, it is possible to draw inferences for the entire lot.

Let's use the following simple example: the feature to be inspected is an OD on a rod, and the specification is 0.375" ±0.005". For our sample, we select 35 parts at random from a certain segment of a production run. (Let's not worry, for the moment, about how we arrive at that figure. Suffice it to say that samples need rarely be greater than 50 parts.)

The first "statistical" procedure is to find the smallest and the largest measurements, then subtract one from the other. This is the "range," which we might express as R = 0.008". Next, we find the average of all the measurements (the sum of the measurements divided by the number of measurements in the sample), which we express as the "X-bar" value, as in X = 0.377".

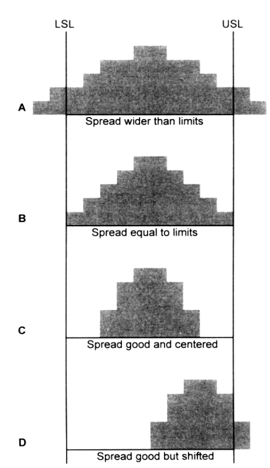

These very simple procedures provide important information. First, they tell us whether all the parts in the sample are in tolerance, and how much of the tolerance range (0.010") the variation (0.008") consumes. And secondly, they show the relationship between the average value and the specification. The average of a sample may lie exactly on the part specification, but the range of the sample may be broader than the tolerance range, so that many parts lie beyond both the upper and lower tolerance limits. On the other hand, the range of the sample may be smaller than the range of the tolerance limits, but the average may be skewed so far from the specification that the entire sample falls outside one of the tolerance limits.

The next step is to chart the data on a histogram. A range of measurement values is divided into equally spaced categories, and each measurement is placed in the appropriate category. If distribution is "normal," the resulting Frequency Distribution Curve will have the familiar bell shape. The "mode"—the category containing the largest number of data points—will be the same as the average in a normal distribution. Distribution curves that are not bell-shaped may indicate a problem in the manufacturing process. For example, if the curve shows a dip where the mode normally appears, it might indicate looseness in the setup or the machine tool.

The histograms show four examples of normal distribution. In A, the range is too wide, indicating that a large percentage of production falls outside the tolerance limits. In B, the range is equal to the tolerance limits: all the items in the sample pass inspection, but there is a statistical probability that some parts in the production run will exceed the limits. In both cases, one would want to reduce the range of variation.

The range of curve C is significantly less than the tolerance range, and falls entirely within the specifications: there is a probability of very few bad parts in the run. In D, the range is acceptable, but because the average is displaced to one side, a significant number of sample parts fall outside of tolerances. Some means must be found to shift the average while maintaining the range.

We'll continue with SQC next month, when we look at control limits and standard deviations.

GAGING STATISTICS—PART DEUX

We have seen how a histogram, showing dimensional variation in a random sample of parts, can be used to determine whether a process is under adequate control. Both the range (R) and the arithmetic mean (X) of the sample must be monitored to make sure that parts consistently remain in tolerance. Let's continue on this topic.

Deviation, as a statistical term, tells how much a given piece of data diverges from the mean. (For example, if the mean value of a sample is 0.010", and the part in question measures 0.012", deviation is +0.002".) *Standard* deviation, when designated as *S*, is a measure of how much the values of *all* the individual items in the sample diverge from the mean value of the sample. Standard deviation can also be designated as s (a small Greek sigma), which is an estimate, based on a sample, of how much the values of the individual items in the *total population from which the sample was drawn* will diverge from the mean value of the population.

We won't go into how to calculate *S* or s here (the methods appear in any basic textbook of statistics, or you can just punch up the values on a scientific calculator). Suffice it to say that once you've found the value of s, further standard deviations are, by definition, simply multiples of s. (e.g., two standard deviations = 2 x s.) Some of the data in a sample will fall beyond one standard deviation from the mean. The second standard deviation tells us how much additional deviation exists among those parts that do not fall within the first standard deviation; the third standard deviation tells how much deviation exists among parts that don't fall within the first two deviations, etc.

Standard deviations can be displayed on a distribution curve as pairs of positive and negative bands on either side of the mean value, as shown in the figure. In any random sample showing a normal, bell-curve distribution, one standard deviation (i.e., one positive and one negative band) will encompass roughly 68% of the values in the population; two standard deviations will encompass roughly 95.5%; and three standard deviations will encompass roughly 99.75% of the workpieces. It always works out this way, because of the laws of random distribution and statistics: wider bell curves naturally exhibit larger standard deviations, and narrower bell curves exhibit smaller ones, and the two are always in proportion to one another.

By comparing the width of the standard deviation bands with the specified tolerance limits, it is possible to calculate, from the sample, how many bad parts will be produced for the production lot from which the sample was drawn. For example, if the tolerance limits are equal to plus or minus three standard deviations, and the mean of the sample is perfectly centered on the specification, one could expect 25 bad parts out of every 10,000 produced (100% - 99.75%). It's simply a matter of calculating the mean and three standard deviations, and comparing the results against the tolerance specification. If the "_3 Sigma" spread falls entirely within the tolerance limits, as in the figure, then the process appears to be under control. If part of the spread falls above or below the tolerance limits, then the process must be adjusted to make the bands narrower (i.e., reduce variation), change the location of the mean, or both.

In any precision manufacturing operation, dimensional variation is very tightly controlled, and the sigma limits are monitored to keep them within the tolerance limits. The span can be compared to the tolerances in various ways. This is known as process capability, and can be discussed in future columns.

In summary: it is possible, based on the laws of statistical probability, to effectively monitor a process and maintain high levels of quality control using a sampling method. In most applications, this tends to be far more cost effective than 100% inspection. Rather than drawing histograms, one can calculate the "control limits" — i.e., the upper and lower boundaries of the standard deviation bands, and the acceptable range of variation of the mean — mathematically.

MORE ON GAGING STATISTIC

A histogram, showing dimensional variation in a random sample of parts, can be used to determine whether a process is under adequate control. Both_the range (R) and the arithmetic mean (X) of the sample must be monitored to make sure that parts consistently remain in tolerance. Let’s continue on this topic. Deviation, as a statistical term, tells how much a given piece of data diverges from the mean. (For example, if the mean value of a sample is .010”, and the part in question measures 0.012”, deviation is +.002”.) Standard deviation, when designated as S, is a measure of how much the values of all the individual items in the sample diverge from the mean value of the sample. Standard deviation can also be designated as σ (the small Greek letter sigma), which is an estimate, based on a sample, of how much the values of the individual items in the total population from which the sample was drawn will diverge from the mean value of the population (see Figure 1).

We won’t go into how to calculate S or σ here (the methods appear in any basic textbook of statistics, or you can just punch up the values on a scientific calculator). Suffice it to say that once you have found the value of σ, further standard deviations are, by definition, simply multiples of σ. (For example, two standard deviations = 2 x σ.) Some of the data in a sample will fall beyond one standard deviation from the mean. The second standard deviation tells us how much additional deviation exists among those parts that do not fall within the first standard deviation; the third standard deviation tells how much deviation exists among parts that don't fall within the first two deviations, and so on.

Standard deviations can be displayed on a distribution curve as pairs of positive and negative bands on either side of the mean value, as shown in the figure.

In any random sample showing a normal, bell-curve distribution, one standard deviation (that is, one positive and one negative band) will encompass roughly 68 percent of the values in the population; two standard deviations will encompass roughly 95.5 percent; and three standard deviations will encompass roughly 99.75 percent of the workpieces. It always works out this way, because of the laws of random distribution and statistics: wider bell curves naturally exhibit larger standard deviations, and narrower bell curves exhibit smaller ones, and the two are always in proportion to one another.

By comparing the width of the standard deviation bands with the specified tolerance limits, it is possible to calculate, from the sample, how many bad parts will be produced for the production lot from which the sample was drawn. For example, if the tolerance limits are equal to plus or minus three standard deviations, and the mean of the sample is perfectly centered on the specification, one could expect 25 bad parts out of every 10,000 produced (100 percent - 99.75 percent). It is simply a matter of calculating the mean and three standard deviations, and comparing the results against the tolerance specification. If the "$\pm$ 3 sigma" spread falls entirely within the tolerance limits, as in the figure, then the process appears to be under control. If part of the spread falls above or below the tolerance limits, then the process must be adjusted to make the bands narrower (that is, reduce variation), change the location of the mean, or both.

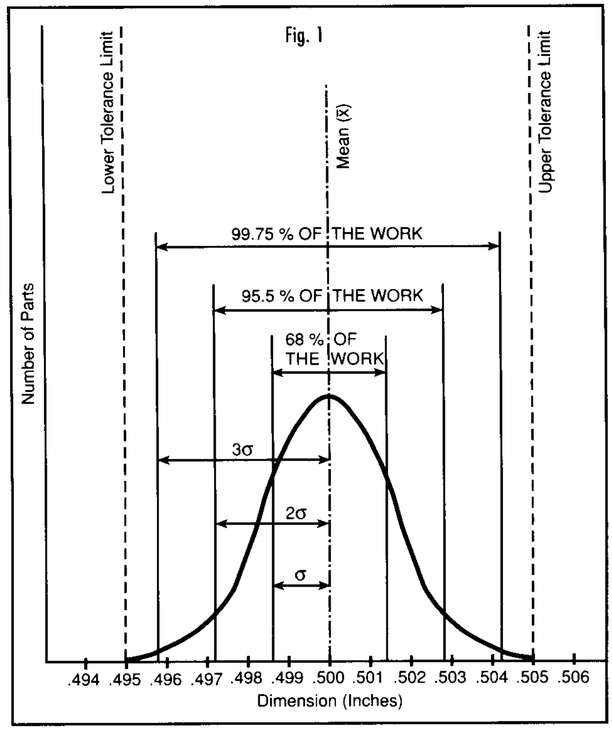

Fig. 1

In any precision manufacturing operation, dimensional variation is very tightly controlled, and the sigma limits are monitored to keep them within the tolerance limits. The span can be compared to the tolerances in various ways. This is known as process capability, and can be discussed in future columns.

In summary: it is possible, based on the laws of statistical probability, to effectively monitor a process and maintain high levels of quality control using a sampling method. In most applications, this tends to be far more cost effective than 100 percent inspection. Rather than drawing histograms, one can calculate the "control limits"-- that is, the upper and lower boundaries of the standard deviation bands, and the acceptable range of variation of the mean-- mathematically.