

Statistical Process Control (SPC)—The Basics

Robert B. Austenfeld, Jr.

(Received on October 28, 2008)

1. Introduction

The purpose of this paper is to introduce the reader to statistical process control (SPC) by showing where it fits in the area of quality management, why it is important, and covering some of its basics. The main focus of this paper will be the control chart which is the “heart and soul” of SPC.

This paper is organized as follows:

1. Introduction
2. Quality Management, Variation, and SPC
3. Types of Control Charts
4. Control Charts for Variables
5. Control Charts for Attributes
6. Process Capability
7. Other Types of Control Charts
8. Summary and conclusion

2. Quality Management, Variation, and SPC

The ultimate purpose of quality management is to deliver products and services to the customer than not only satisfy but also and ideally delight the customer. That is, products/services¹⁾ that more than meet the customer’s expectations in terms of usability, esthetics, reliability, durability, etc. To do this the

1) From now on the term “product” will stand for both product and service.

process(es) that create(s) the product must be good. A “good” process, in turn, is one with little variation in terms of what it is producing—once an “ideal” product is created it is highly desirable that the “same” product be delivered each and every time. This means that every part of the product needs to be similar. For example if you are producing something like a gearbox where several gears must mesh together, if each gear isn’t almost “perfect” in terms of meeting some specification, the gears either will not work or will not work well and be subject to excessive wearing due to the mismatch. Whereas if your processes consistently produce gears well within the specification’s tolerances, your gearbox will not only work smoothly but also last a long time due to minimal wear. The point here is that although we can never eliminate all variation, we want to minimize it.

The question then is how do we minimize process variation? The first thing to understand is that there are two types of causes of variation: common and special.²⁾ The common causes are those that are inherent in the process itself. They are generally random and nominally form a normal distribution. They are due to those things that influence the process over the long term such as the type of material being used, the capability of the machinery involved, the ability/training of the operators, environmental conditions, etc. By changing these factors it is possible to improve the process if such seems warranted by the cost involved; i.e., is the process good enough or will the benefit be worth the improvement cost?

The other source of variation is that due to special causes. These are due to what is usually a temporary condition such as a machine getting out of adjustment, a tool wearing, an input to a chemical process becoming diluted, an operator who is new and doesn’t know how to properly operate the process, or some

2) Sometimes called “chance” and “assignable” respectively.

temporary environmental condition such as temperature or pressure changes that affect the process.

It is these special causes that are the main target of statistical process control (SPC) and control charts. The purpose of a control chart is to control a process by revealing when some significant change has occurred; i.e., when some special cause of variation has occurred. It does this by showing when the variation in the process causes some value being plotted on the chart—for example the process average (or mean, denoted by the symbol for mu, μ)—to go beyond a reasonable value, usually plus or minus three standard deviations (3 sigma or, symbolically, 3σ) from that average value. Another thing of interest besides a significant change in the process mean is whether the process' spread³) has changed. Since we are usually dealing with values that are essentially normally distributed, a changed in spread will also affect the process output since what was once considered an acceptable spread when the process was under statistical control (i.e., no special causes present) will now cause the value of interest to go beyond those established specification limits. Accordingly, there are usually two charts plotted, one monitoring the process mean and the other the process spread.

3. Types of Control Charts

In general, we are dealing with two types of data: *variable* data from a process that produces something where we can measure some value and *attribute* data where we are concerned with whether the product of the process is defective or not, or has one or more defects. In the section that follows we will describe in detail two types of charts for *variable* data:

- X-bar/R charts
- Individuals and moving range charts

3) Also called “dispersion.”

And in section 5 four types of charts for *attribute* data will be described:

- np-chart
- p-chart
- c-chart
- u-chart

Since these are the most common charts used this will give us a basic understanding of control charting as the key tool for SPC. Some other types of control charts will be very briefly described in section 7.

4. Control Charts for Variables

X-bar and R charts. The most common control charts for variables are the X-bar (\bar{X}) and R charts. \bar{X} is the symbol for the sample average (mean) of the variable X, and R stands for the sample range of the variable. The variable, of course, is some important quality characteristic of the part being produced by the process. As mentioned, when controlling a process we are interested in two things: has the average shifted to some new value and has the dispersion/spread (standard deviation) changed. If either of these things has occurred, it may mean the process is no longer “stable and under control”—that is, producing predictable results. Given that we are usually working with data that is essentially normally distributed when dealing with variables, ideally the process mean will be centered between the specification limits. Furthermore, if our process is to produce almost all of the subject parts “within spec,” the dispersion must be such that chance of finding a sample value beyond the specification limits is almost nil. For control chart purposes the value of three standard deviations is used. Figure 1 shows a normal distribution and how often a randomly selected variable will fall within ± 3 standard deviations, namely 99.7%. Section 6 of this paper will discuss the question of whether the process can “meet the spec,” now we are concerned only with how stable our process is.




Figure 1. A normal distribution (from Oakland, 2008, p. 90)

It is important to understand that we can never know the true process mean, designated by the Greek letter mu (μ), or the true process standard deviation designated by the Greek letter sigma (σ). Accordingly, we must find ways to estimate these two important values. For control chart purposes the process mean (μ) will be estimated by the grand mean, $\bar{\bar{X}}$, of the sample means, \bar{X} . The central limit theory tells us that even if the process (i.e., population) values are not normally distributed, *the sample means will be* as the sample size (n) increases, and $\bar{\bar{X}}$ becomes a very good estimate of the process mean. This holds true even if n is as small as 4 (Oakland, p. 94).

As for the process standard deviation (σ), for sample sizes of say 12 or less \bar{R}/d_2 provides a good estimate. \bar{R} is the average of the sample ranges and d_2 , something called Hartley's constant, is a function of sample size (n). Figure 2 is a table showing Hartley's constant along with other control chart factors.

Now we can use these proxies for μ and σ to set up a control charts that will let us monitor μ and σ . To monitor the process average, μ , we will create an X-bar (\bar{X}) chart with control limits at ± 3 standard deviations from the mean of the

TABLE 45.1 Control Chart Factors

<i>n</i>	<i>d</i> ₂	<i>d</i> ₃	<i>c</i> ₄	<i>c</i> ₅	<i>A</i>	<i>A</i> ₂	<i>A</i> ₃	<i>B</i> ₃	<i>B</i> ₄	<i>B</i> ₅	<i>B</i> ₆	<i>D</i> ₃	<i>D</i> ₄
2	1.13	0.85	0.80	0.60	2.12	1.88	2.66	0	3.27	0	2.61	0	3.27
3	1.69	0.89	0.89	0.46	1.73	1.02	1.95	0	2.57	0	2.28	0	2.58
4	2.06	0.88	0.92	0.39	1.50	0.73	1.63	0	2.27	0	2.09	0	2.28
5	2.33	0.86	0.94	0.34	1.34	0.58	1.43	0	2.09	0	1.96	0	2.11
6	2.53	0.85	0.95	0.30	1.23	0.48	1.29	0.03	1.97	0.03	1.87	0	2.00
7	2.70	0.83	0.96	0.28	1.13	0.42	1.18	0.12	1.88	0.11	1.81	0.08	1.92
8	2.85	0.82	0.97	0.26	1.06	0.37	1.10	0.19	1.82	0.18	1.75	0.14	1.86
9	2.97	0.81	0.97	0.25	1.00	0.34	1.03	0.24	1.76	0.23	1.71	0.18	1.82
10	3.08	0.80	0.97	0.23	0.95	0.31	0.98	0.28	1.72	0.28	1.67	0.22	1.78

Figure 2. Control chart factors (from Wadsworth, 1999, Table 45.1, p. 45.6)

sample means using these two formulas:

$$UCL_{\bar{X}} = \bar{\bar{X}} + 3\bar{R}/(d_2\sqrt{n}) = \bar{\bar{X}} + A_2\bar{R} \text{ for the upper control limit}$$

$$LCL_{\bar{X}} = \bar{\bar{X}} - 3\bar{R}/(d_2\sqrt{n}) = \bar{\bar{X}} - A_2\bar{R} \text{ for the lower control limit}$$

Let us see how these formulas are derived. The reason we are using sample means as our variable of interest is this gives us a much more realistic picture of the actual process. For example if we were to simply plot the individual values taken from some process they would show a much greater dispersion than the sample means. Oakland (p. 84) uses the example of sampling the lengths of steel rods. His first sample ($n = 4$) consists of these values: 144 mm, 146 mm, 154 mm, and 146 mm. The sample mean is 147.5, which is obviously more representative of the true mean of the process than the more widely dispersed individual values. And, since we are using sample means, the standard deviation of these means is not σ but σ/\sqrt{n} , something called the standard error of the means. Therefore the fundamental form of our control limits formula is: $CL_{\bar{X}} = \bar{\bar{X}} \pm 3\sigma/\sqrt{n}$. As mentioned above, a good estimate for σ for small sample sizes ($n \leq 12$) is \bar{R}/d_2 . To simplify calculation of the limits further, the value of $3/(d_2\sqrt{n})$ is shown in the table in Figure 2 for various sample sizes as the “constant” A_2 .

In a similar way, to monitor the process dispersion, σ , we will create a range chart with control limits at ± 3 standard deviations from the mean of the sample ranges using these two formulas:

$$UCL_R = \bar{R} + 3d_3(\bar{R}/d_2) = (1 + 3d_3/d_2)\bar{R} = D_4\bar{R} \text{ for the upper control limit}$$

$$LCL_R = \bar{R} - 3d_3(\bar{R}/d_2) = (1 - 3d_3/d_2)\bar{R} = D_3\bar{R} \text{ for the lower control limit}$$

The fundamental form of these equations is: $CL_{\bar{R}} = \bar{R} \pm 3\sigma_R$ where \bar{R} is our estimate of the process' mean range. To estimate the range standard deviation we use another factor from the table in Figure 2: d_3 , and the relationship $\sigma_R = d_3\sigma$. Again using our estimate of σ , \bar{R}/d_2 , we get $\sigma_R = d_3\sigma = d_3(\bar{R}/d_2)$ and the first form of the above formulas. As with the formulas for the X-bar (\bar{X}) chart control limits these formulas are simplified into a final form using the “constants” D_4 and D_3 , which again are read from the table in Figure 2 according to the sample size.

Borrowing from Griffith, let's walk through the X-bar/R⁴⁾ charts example shown on page 99 (Figure 3). This is a milling process producing a “mount assembly” and the characteristic measured is the “gap dimension.” It is customary to place both charts on the same sheet so that sample to sample changes in both the process average and the process range (dispersion) can be easily observed. To compute our X-bar chart control limits we need to know three things: $\bar{\bar{X}}$ (the mean of the sample means, sometimes called the grand mean), \bar{R} (the mean range of the sample ranges), and n (the sample size). In Griffith's example, 25 samples of size 5 are taken. Therefore $n = 5$. If we add up all the sample means (\bar{X}) and divide by 25 we get the grand mean, $\bar{\bar{X}}$, of 0.716 (see table on the next page). And, if we add up all the sample ranges and divide by 25 we get the mean of the sample ranges, \bar{R} , of 0.178.⁵⁾ As can be seen on the control charts these two values, $\bar{\bar{X}}$ and \bar{R} , are the “center lines” about which we will plot the sample values and are drawn with a solid line.

Using our above formulas we can now compute the control limits for each

4) Also known as an “Average and Range” chart.

5) Although Griffith does not give any measurement units in this example it can be assumed that all values are of some reasonable amount such as millimeters.

chart. For the X-bar chart the upper control limit (UCL) will be:

$$\bar{\bar{X}} + A_2\bar{R} \text{ or } 0.716 + (0.58)(0.178) = 0.819$$

where the value for A2 is taken from the table in Figure 2 for a sample size of five. Similarly the lower control limit (LCL) will be: $\bar{\bar{X}} - A_2\bar{R}$ or $0.716 - (0.58)(0.178) = 0.613$.

For the range chart our UCL will be $D_4\bar{R}$ or $(2.11)(0.178) = 0.376$ where D4 is taken from the table in Figure 2. Similarly the LCL will be $D_3\bar{R}$ or $(0)(0.178) = 0$. To complete our charts these upper and lower control limits are drawn with dashed lines as seen on the example.

Once the centerlines and control limits are drawn the sample means and ranges are plotted on their respective charts.

Now that we know how to construct these control charts, what do we do with them? Recall that the purpose of a control chart is to see if our process is stable and under control which means has anything occurred to cause either the process

mean or process dispersion (standard deviation) to change significantly. According to Wadsworth if the process is under control it will exhibit the following characteristics:

- Most of the plotted points occur near the centerline.
- A few of the points occur near the control limits.
- Only an occasional rare point occurs beyond the control limits.
- The plotted points occur in a random manner with no clustering,

sample no.	sample sums	sample means	sample ranges
1	3.50	0.70	0.20
2	3.85	0.77	0.20
3	3.80	0.76	0.10
4	3.40	0.68	0.15
5	3.75	0.75	0.20
6	3.65	0.73	0.25
7	3.65	0.73	0.15
8	3.60	0.72	0.20
9	3.90	0.78	0.20
10	3.35	0.67	0.20
11	3.75	0.75	0.40
12	3.80	0.76	0.20
13	3.60	0.72	0.05
14	3.55	0.71	0.25
15	4.10	0.82	0.15
16	3.75	0.75	0.15
17	3.80	0.76	0.15
18	3.35	0.67	0.15
19	3.50	0.70	0.20
20	3.10	0.62	0.05
21	3.30	0.66	0.30
22	3.45	0.69	0.20
23	3.50	0.70	0.15
24	3.20	0.64	0.10
25	3.20	0.66	0.10
Taking the averages we get >>		0.716	0.178
		grand mean	range mean




Figure 3. Example of an X-bar/R control chart (from Griffith, 1996, p. 16)

trending, or other departure from a random distribution. (p. 45.7)

Should any of these characteristics not be present it means there is a good chance the process is not in control and we have a special cause of variation. In our Griffith example we can see a typical instance where Wadsworth's fourth characteristic is "violated" starting with sample number 15. There is a downward trend that would suggest that something is causing the process to no longer exhibit randomness. In this case it is traced to "tool wear." Note the comment at sample number 11 of "unknown cause" for the point falling beyond the range chart's UCL. Since this seems to be an isolated instance it may simply be that "occasional rare point" mentioned in Wadsworth's third characteristic. However, since the chance of that occurring is so small it should be investigated. Apparently in this case nothing was found to be amiss.

Griffith states that any indication of an out-of-control condition "deserves some action to find the special cause and eliminate it" and "[i]f the cause is found and eliminated, recalculate the control limits (p. 17)." He further recommends a fairly frequent recalculation of the control limits in any event to be sure they accurately reflect the actual current process that could, in fact, be changing. Griffith also emphasizes the importance of marking anything of significance on the chart so it becomes a running record of the process. Examples of this are shown on the Griffith control chart.

There are many other "rules" for interpreting this type of control chart and the reader is referred to the Griffith, Wadsworth, Oakland and other references. As a matter of interest, Oakland writes from a UK/European perspective and describes control charts in terms of not only upper and lower control limits set at three standard deviations but an additional set of control limits set at two standard deviations. The former are called "action lines" and the latter "warning lines." The idea is to give a more precise way of detecting possible problems with the addition of the warning lines.

Individuals and moving range charts. Another common control chart for variables is the individuals chart. This is used when it is not possible to take a meaningful subgroup⁶⁾ from the process. Wadsworth (p. 45.10) mentions these examples: a chemical or other continuous process, or when measuring such things as pressure, temperature, accounting data, efficiency, ratios, expenditures,

and quality costs. Now we are taking only one observation ($n = 1$) and we will be plotting X , not the sample X -bar. To estimate the process mean, μ , we will use the mean of the observations designated \bar{X}_i . Now our control limits will be set according to this formula: $CL = \bar{X}_i \pm 3\bar{R}/d_2$.

The process standard deviation, σ , is estimated as with the X -bar chart using \bar{R}/d_2 . (Note that here we are using σ not σ/\sqrt{n} since we are dealing with individual process values, X_i , not sample means, \bar{X} .)

Now the question is how do we find \bar{R} since we no longer have two or more values in a subgroup (sample). The answer is the moving range. A moving range is calculated using consecutive X values from the observations. For example if our first observed X value is 8.0 and our next observed X is 8.5, the range value associ-

observation no.	obs. values (X_i)	obs. moving ranges (R_m)
1	8.0	
2	8.5	0.5
3	7.4	1.1
4	10.5	3.1
5	9.3	1.2
6	11.1	1.8
7	10.4	0.7
8	10.4	0.0
9	9.0	1.4
10	10.0	1.0
11	11.7	1.7
12	10.3	1.4
13	16.2	5.9
14	11.6	4.6
15	11.5	0.1
16	11.0	0.5
17	12.0	1.0
18	11.0	1.0
19	10.2	0.8
20	10.1	0.1
21	10.5	0.4
22	10.3	0.2
23	11.5	1.2
24	11.1	0.4
Taking the averages we get >>	10.57	1.31
	mean of observations	mean of moving ranges

6) By “subgroup” we mean a sample of n values where we are dealing with the production of (usually) many individual objects such as a steel rod or tablets or bottles being filled.

ated with the “8.5” observation will be 0.5. Note there will always be $k-1$ range values, where k is the number of observations. Taking the average of the moving ranges, R_m , we get \bar{R}_m and use this for our \bar{R} . This \bar{R} is also used to calculate our control limits for our moving range chart in the same way we did that for the previous range chart; i.e., $UCL = D_4\bar{R}_m$ and $LCL = D_3\bar{R}_m$.

Again drawing on a Griffith example as shown on the next page (Figure 4), let’s calculate the control limits. For the observed values as listed in the table on the previous page we calculate the mean of the individual X values, \bar{X}_i , and the mean of the moving ranges, \bar{R} . These respectively are 10.57 and 1.31. Using the above formulas our UCL is $10.57 + 3(1.31)/1.13$ (for $n = 2$) equals $10.57 + 3.48$ or 14.05 .⁷⁾ And our LCL is $10.57 - 3.48$ or 7.09 . Similarly for the moving range chart: UCL is $3.27(1.31)$ equals 4.28 and the LCL is $0(1.31)$ equals zero. As with Griffith’s X-bar/R chart example this individuals/moving range chart shows examples of problems that a control chart might typically reveal such as contamination at observation 13 and how cleaning the tank improved things from sample 15 on.

5. Control Charts for Attributes

Introduction. So far we have been talking about control charts for variables, that is for processes where we measure some variable quality characteristic such as the diameter of a shaft or the number of tablets being placed in a container or the amount of liquid being put in a bottle. There are also processes where the question is does this product meet some stated standard to make it acceptable or not. Once the standard has been clearly established, we will use our control chart to determine how many of those parts or products actually have met the standard. Oakland (p. 192) gives these examples:

7) The control chart (Figure 4) shows 14.06, perhaps due to “rounding error.”




Figure 4. Example of an Individuals and Moving Range control chart (from Griffith, 1996, p. 26)

...bubbles in a windscreen [windshield], the general appearance of a paint surface, accidents, the particles of contamination in a sample of polymer, clerical errors in an invoice and the number of telephone calls.

We are usually interested in two types of attributes: (1) whether the unit of the product is acceptable or not, i.e., it is conforming or not, and (2) the number of defects or non-conformities in a unit of product.⁸⁾ Furthermore we can break down these two categories in terms of sample size where in one case the sample size is constant and in the other it is not constant. This chart shows the four com-

Type of attribute chart	Name of chart	Measures
Does unit conform or not? Constant sample size.	np-chart	number of non-conforming units
Does unit conform or not? Variable sample size.	p-chart	number of non-conforming units
How many non-conformities are in the unit? Constant sample size.	c-chart	number of non-conformities (defects) per unit.
How many non-conformities are in the unit? Variable sample size.	u-chart	number of non-conformities (defects) per unit.

mon types of attribute control charts according to this classification. In dealing with attribute data we are now concerned with questions such as this: given that the average proportion of defectives (non-conforming) units for a process is such and such (say 0.02) what is the chance that for a random sample of say 100 units drawn from the process the number of defectives will exceed a certain number. In other words we are talking about is a binomial statistic; i.e., a case where the results of each observation can only be one of two situations: good or bad. In very simple terms if we flip a coin two times and ask what is the probability that the number of tails (where a tails equals “bad”) will be more than

8) Note that for attribute data we are not dealing with “ranges” only the number (or proportion) of non-conforming units or non-conformities.

one. For this to happen both flips would have to be tails, and the probability of that is only 0.25 if we have a “fair” coin. If we were to repeat this experiment several times and keep getting two tails we would suspect the coin is no longer fair and that the “average proportion of defectives” for the process has changed from 0.5 to some much higher value. Let’s now go through an example, this time from Oakland, and see how these ideas apply for an np-chart.

The np-chart. In this example Oakland (p.198) uses a process producing ball-point pen cartridges and draws 50 samples of 100 each over a “typical” time frame. The results are shown in the

table to the right. Our formula for the control limits on the np-chart is:

$$CL = n\bar{p} + 3\sqrt{n\bar{p}(1-\bar{p})}$$

where n is the sample size (100 in this case) and \bar{p} is the estimate of the average proportion of defectives in the process. Using the data from our table we can calculate the average number of defectives per sample ($100/50 = 2$). This is also equal to $n\bar{p}$ and \bar{p} must be equal to $2/100$ or 0.02. As with our other charts we will set our control limits at $\pm 3\sigma$. For a binomial distribution: $\sigma = \sqrt{n\bar{p}(1-\bar{p})}$. So now we can setup our np-chart with our “centerline” at $n\bar{p}$ and our control limits at 3σ . We already know $n\bar{p}$ and $\sigma = \sqrt{n\bar{p}(1-\bar{p})} = \sqrt{2(0.98)} = 1.4$. Therefore our UCL = $2 + 3(1.4) = 6.2$ and, theoretically, the LCL = -2.2 which

sample no.	no. of defectives	sample no.	no. of defectives
1	2	26	0
2	4	27	3
3	1	28	1
4	0	29	2
5	0	30	1
6	4	31	2
7	5	32	1
8	3	33	5
9	2	34	3
10	3	35	0
11	2	36	2
12	3	37	2
13	0	38	1
14	3	39	3
15	1	40	1
16	2	41	1
17	3	42	3
18	1	43	0
19	2	44	2
20	1	45	1
21	2	46	2
22	4	47	0
23	2	48	4
24	1	49	2
25	6	50	1
totals >>	57		43
Total no. of defects >>			100

makes no sense so is set to zero. To quote Oakland:

In control charts for attributes it is commonly found that only the upper limits are specified since we wish to detect an increase in defectives. Lower control lines may be useful, however, to indicate when a significant process improvement has occurred, or to indicate when suspicious results have been plotted. (p. 201)

Using the data we've now calculated we get the control chart shown in Figure 5.




Figure 5. An np-chart for number of defectives (from Oakland, 2008, p. 202)

As mentioned Oakland writes from a UK/European perspective so his charts show upper and lower “warning” and “action” lines versus the single upper/lower control lines common on U.S. control charts. His warning line in Figure 5 is set at 2σ and is meant to give a more precise way of detecting when the process is no longer stable, i.e., the process mean has shifted.

Examining the chart in Figure 5 we can conclude that this process is in statistical control since “all samples contain less defectives than the action limit and only 3 out of 50 enter the warning zone, and none of these are consecutive...” (Oakland, p. 201). Had the results of any sample exceeded the action line or exceeded the warning line consecutively twice or more action to find a cause for this would have been called for.

The p-chart. When controlling for nonconforming product (defectives) and

the sample size is *not* constant the p-chart is used. Perhaps one of the best explanations of this chart is also from Oakland and we will use his excellent example. In this example “textile components” are being delivered in varying quantities and for control chart purposes we have sampled 24 deliveries as shown by the table to the right. With a constant sample size (n), the control limits remain the same but with a variable sample size the control limits will change as n changes. Theoretically we should calculate a different set of control limits for each n . However, Oakland tells us (p. 205) that as long as the sample size is within 25 percent of the average n we can use the average n . For those cases where n falls outside this range we must calculate separate control limits. Oakland’s example will ably illustrate this.

sample no.	sample size	no. of defectives	proportion defective
1	1,135	10	0.009
2	1,405	12	0.009
3	805	11	0.014
4	1,240	16	0.013
5	1,060	10	0.009
6	905	7	0.008
7	1,345	22	0.016
8	980	10	0.010
9	1,120	15	0.013
10	540	13	0.024
11	1,130	16	0.014
12	990	9	0.009
13	1,700	16	0.009
14	1,275	14	0.011
15	1,300	16	0.012
16	2,360	12	0.005
17	1,215	14	0.012
18	1,250	5	0.004
19	1,205	8	0.007
20	950	9	0.009
21	405	9	0.022
22	1,080	6	0.006
23	1,475	10	0.007
24	1,060	10	0.009
totals >>	27,930	280	
avgs >>	1,164		

Let’s first see what an acceptable range of n would be and then calculate the control limits for those samples. Since the total of all the sample sizes is 27,930, the average, \bar{n} , would be 1,164 as shown in the table. This give us an acceptable range of $1,164 \pm (0.25 \times 1,164) = 1,164 \pm 291$ or from 873 to 1,455.

Similar to the np-chart the formula for control limits is $\bar{p} \pm 3\sigma$ where, for this type of chart, sigma is: $\sigma = \sqrt{\bar{p}(1-\bar{p})} / \sqrt{n}$. The average proportion of defectives, \bar{p} , is the total number of defective (280 in this case) divided by the total number of items sampled (27,930) giving us a value of 0.01. Using the above formu-




Figure 6. A p-chart for number of defectives (from Oakland, 2008, p. 208)

las and our average sample size, $\bar{n} = 1164$, this results in a σ of 0.003 and control limits at $0.01 \pm 3(0.003)$ or 0.019 and 0.001. These control limits will suffice for all samples where n was within the 25 percent range of 873 to 1455; i.e., all samples except numbers 13 and 16, which were greater than 1,455,⁹⁾ and numbers 3, 10 and 21, which were less than 873. For these samples we must calculate individual control limits using their specific n 's and plot these on our control chart. For example for sample number 10 the calculation would be as follows:

$$\begin{aligned}
 CL &= \bar{p} \pm 3\sigma = \bar{p} \pm 3\sqrt{\bar{p}(1-\bar{p})} / \sqrt{n} = 0.01 \pm 3\sqrt{0.0099} / \sqrt{540} \\
 &= 0.01 \pm 3(0.0043) = 0.01 \pm 0.013
 \end{aligned}$$

which gives us an UCL of 0.023 and a LCL of -0.003 (or, in effect, zero). Figure 6 shows Oakland's p-chart and how those samples whose n fell outside the 25 percent range have had their control limits drawn individually. Note that sample number 10, whose control limits we just calculated, exceeds even the greater UCL of 0.023. Also note that sample number 21, although exceeding Oakland's upper "warning line" of 0.02 did not exceed the UCL of 0.025, had \bar{n} been used for this sample's control limits it would have far exceeded the UCL

9) Actually sample 23 was a little greater than 1,455 but apparently not enough to warrant having individual control limits computed for it on Oakland's p-chart example.

of 0.019 indicating a potentially serious problem.

Analyzing this chart, Oakland (p. 206) notes that all is reasonably well until delivery of sample number 10 which probably resulted in some discussions between the supplier and the customer bringing the quality back to within acceptable levels until another possible problem occurred as indicated by the results of sample number 21.

The c-chart. So far with the np and p charts we've been concerned with whether the unit (item being inspected) conforms or doesn't. Now we take up the situation where we want to assess the number of non-conformities, that is defects, in a unit. As with the np/n charts two possibilities exist: the size of the unit being inspected is either constant or varies. Oakland (p. 209) uses the example of fisheyes (blemishes) in polythene film being produced where the number of fisheyes in randomly selected identical lengths of film is counted. However rather than staying with Oakland let's borrow

from Wadsworth for our c-chart example so the reader can see what a typical control chart might look like from a third perspective.¹⁰⁾ In this example we are interested in the number of pinholes in a paper that is suppose to be impervious to oils. Samples of identical size are taken from the production process and tested using colored ink to detect any pinholes. The data is shown in the table at the right. Whereas for the np/p charts the underlying distribution was binomial, here it is the Poisson distribution whose standard deviation is simply

sample no.	no. of defects	sample no.	no. of defects
1	8	14	6
2	9	15	14
3	5	16	6
4	8	17	4
5	5	18	11
6	9	19	7
7	9	20	8
8	11	21	18
9	8	22	6
10	7	23	9
11	6	24	10
12	4	25	5
13	7		
Tot. no. of defects >>			200
Avg. no. of defects per sample \bar{c} (200/25) >>			8

10) Having already looked at examples from Griffith and Oakland.




Figure 7. A c-chart for number of defectives, uniform sample size (from Wadsworth, 1999, p. 45.16)

the square root of the mean, in this case $\sqrt{\bar{c}} (= \sqrt{8} = 2.83)$. Using our usual formula for the control limits, $CL = \bar{c} \pm 3\sigma$ we get $CL = \bar{c} \pm 3\sqrt{\bar{c}} = 8 \pm 8.5$ for an UCL of 16.5 and a LCL of -0.5 (i.e., zero). Figure 7 shows Wadsworth's c-chart.

Except for sample 21, the process appears to be stable and under control, producing predictable results.

The u-chart. The final control chart this paper describes in some detail is the u-chart, where you are trying assess and control non-conformances (defects) and your sample sizes are not uniform. Returning to Griffith for an example—see next page, Figure 8—we have a forging process producing covers where the number of covers being inspected in each sample varies as shown by the data in the table on page 112. For this type of chart sigma is: $\sigma = \sqrt{\bar{u}} / \sqrt{\bar{n}}$. And, similar to the p-chart, $CL = \bar{u} \pm 3\sigma = \bar{u} \pm 3(\sqrt{\bar{u}} / \sqrt{\bar{n}})$. As with the p-chart, for any sample sizes falling within the range of $\bar{n} \pm 0.25\bar{n}$ (4.2 to 6.9 in this case), the control limits can be calculated using \bar{n} ($111/20 = 5.55$) in the above formula. Therefore, in this example $\sigma = \sqrt{1.8} / \sqrt{5.55} = 0.57$ and the UCL would be $1.8 + 3(0.57) = 3.5$ and the LCL $1.8 - 3(0.57) = 0.09$ or essentially zero. However, a

quick look at the table shows that most of the n 's do not fall within the 25 percent range so individual control limits must be calculated for almost every point on the chart. And it can be seen in Figure 8 that he has calculated control limits for each sample point.

There are several other kinds of control charts and these will be briefly described in section 7 of this paper. Now we turn to the matter of process capability.

sample no.	sample size (n) (units)	no. of defects	no. of defects per unit (u)
1	2	7	3.5
2	4	5	1.3
3	6	12	2.0
4	6	10	1.7
5	8	25	3.1
6	4	16	4.0
7	2	2	1.0
8	2	4	2.0
9	4	6	1.5
10	8	20	2.5
11	8	15	1.9
12	7	16	2.3
13	3	17	5.7
14	10	3	0.3
15	2	6	3.0
16	8	10	1.3
17	12	11	0.9
18	3	6	2.0
19	10	4	0.4
20	2	6	3.0
Totals >>	111	201	
$\bar{u} = 201/111 = 1.8$			

6. Process Capability

Process capability for variables. A process may be stable and under control but is it capable? In other words does the process' mean and standard deviation conform to the requirements of the specification. There are two common indices used to determine this: C_p to see if the process standard deviation (spread) is sufficiently small, and C_{pk} to see if the process mean is sufficiently centered. Again borrowing from Griffith, Figure 9 will be used to help us understand these two indices. Figure 9 shows three process distributions, two whose mean is centered between the upper and lower specification limits (USL and LSL) and one (the top one) whose mean is not centered. *If our process is centered* then all we are

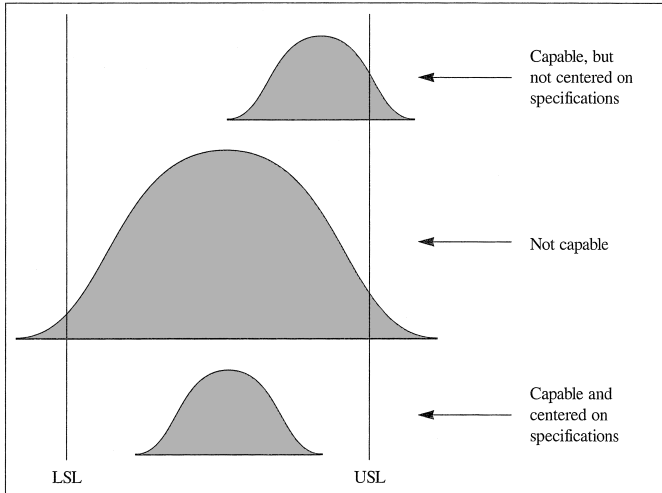


Figure 9. A diagram to illustrate process capability (from Griffith, 1996, p. 126)

concerned about is whether the specification limits exceed 3σ on both sides of the process mean. This is the index C_p and is determined as follows:

$C_p = \frac{USL - LSL}{6\sigma}$. Oakland provides the example of a pharmaceutical company

manufacturing tablets. Twenty samples of a sample size (n) of four are used to create a control chart. Recall that when we calculated the control limits for our X-bar chart we used $\bar{\bar{X}}$ as our estimate of the process mean and \bar{R}/d_2 as our estimate of σ . Given the following data from our control chart calculations, $\bar{\bar{X}} = 2500$ mg and $\bar{R} = 91$ mg, and upper/lower specification limits of 2650/2350 (mg), we can now determine C_p and C_{pk} . First C_p :

$$C_p = \frac{USL - LSL}{6\sigma} = \frac{USL - LSL}{6\sigma(\bar{R}/d_2)} = \frac{2650 - 2350}{6(91/2.059)} = 1.13$$

This tells us most of our process measurements will be within the specification limits since those limits are just outside the $\pm 3\sigma$ points of the distribution. This means (see Figure 1 again) that if our process remains stable and assuming we

do have a normal distribution at least 99.7% of our values will be “within spec.”

The problem with the C_p index is it doesn't account for a process that is not centered on the specification limits, hence, the C_{pk} index. Now it is necessary to calculate two values: one for the USL and one for the LSL: $C_{pk(u)} = \frac{USL - \bar{X}}{3\sigma}$ and $C_{pk(l)} = \frac{\bar{X} - LSL}{3\sigma}$. It should be apparent that if we used the same data as before in Oakland's tablet example we would get the same result of 1.13 since that process was centered. To show how C_{pk} works Oakland changed the above values to $\bar{X} = 2650$ mg and the upper and lower specification limits to 2750 mg and 2250 mg; there is no change in σ with \bar{R} remaining at 91 mg. This new data yields:

$$C_{pk(u)} = \frac{USL - \bar{X}}{3\sigma} = \frac{2750 - 2650}{3(91/2.059)} = 0.75 \text{ and}$$

$$C_{pk(l)} = \frac{\bar{X} - LSL}{3\sigma} = \frac{2650 - 2250}{3(91/2.059)} = 3.02.$$

This tells us that the distance between our process mean and the upper specification limit (USL) is less than 3σ and, therefore, some of our values will be outside the USL and be unacceptable. It also tells us the distance between our process mean and the LSL is very large, in fact, three times 3σ so there is little chance that any values would fall below the LSL. The top distribution in Figure 9 shows this situation.¹¹⁾ This process is definitely not capable. Had we computed only C_p we would have gotten:

$$C_p = \frac{USL - LSL}{6\sigma} = \frac{500}{6(91/2.059)} = 1.89$$

and might have concluded the process was capable.

At what value of C_{pk} should we consider the process as not capable? If the specification and 3σ points coincide it means that on average 99.7% of the

11) Although Figure 9 shows this as a “capable” process it actually would not be since too much of its distribution exceeds the USL.

measurements will be within specification and only 0.3% out of specification if the process is centered ($C_p = 1.0$). If it is not centered¹²⁾ but the 3σ point of the “tail” that is closest to either the upper or lower specification limit coincides with that limit then only about 0.15% of the values will be out of specification ($C_{pk} = 1.0$). Because we probably never have a purely normal process and processes do tend to become unstable we would want our specification limits to be more than just 1.0 and Oakland (p. 285) recommends a C_{pk} of 2.0 for a “high level of confidence in the producer”—for example, like the bottom distribution in Figure 9. Remember, however, that measures of process capability are premised on the process being in statistical control.

If the process is not capable then the alternatives are either to relax the specification if that is feasible or improve the process through whatever means; e.g., better input material, equipment, procedures, training, or, even, some entirely new approach to producing desired product.

Process capability for attributes. Process capability for attribute data is simply the process average given the process remains in statistical control. For example for the p-chart the process capability will be 1.00 minus \bar{p} , the average proportion of defectives. So if $\bar{p} = 0.02$ the process capability will be 0.98 meaning the process is theoretically capable of producing 98% good product.

As can be seen by now control charts and process capability go hand in hand and provide the basis for not only controlling existing processes but for carrying out the most basic dictum of quality management; i.e., continuous improvement of the processes.

12) There are times when the measurement in question is the time it takes to do something, like process an insurance claim or carry out some transaction. Here we are only concerned with one side of the distribution where the process values are less than some desired standard.

7. Other Types of Control Charts

There are many types of control charts and so far we've only discussed the most common types as shown to the right:

To give the reader a feel for what other control charts exist we will now briefly discuss the following:

- X-bar and standard deviation charts
- Moving mean and moving range charts
- Median chart
- Exponentially weighted moving average chart
- Cumulative sum (CUSUM) chart
- Short run charts

Variable control charts:

- X-bar and R charts
- Individuals and Moving Rang charts

Attribute control charts:

- np-chart
- p-chart
- c-chart
- u-chart

X-bar and standard deviation charts. These are similar to the X-bar/R charts except \bar{s} , the average of the sample standard deviations, is used instead of \bar{R} to calculate the control limits. The sample standard deviations are calculated with this formula:

$$s = \sqrt{\frac{\sum (X_i - \bar{X})^2}{n-1}}$$

Using appropriate factors, \bar{s} is used to determine the control limits for both the X-bar and standard deviation charts.

There is not much difference in the results compared to the X-bar/R charts however the X-bar/standard deviation charts are said to be a bit more sensitive to process changes. Also, according to Griffith (p. 19), the X-bar/standard deviation charts are used when the sample size is 10 or more "because the ranges become inefficient" at these sizes.

Moving mean and moving range charts. These charts are similar to the indi-

viduals and moving range charts already discussed. According to Griffith (p. 28) they are better since they “dampen some of the effects of over-control.....and also provide increased ability to defect shifts in the process level.” In an example given in Oakland (starting on p. 164) daily readings are taken from a polymerization process of an important quality characteristic. Control limits have already been established from previous data taken when the process was in control. After the fourth daily reading the moving mean and range are calculated. Using Oakland’s example the table to the right shows the calculations for the first six days. Starting on Day 4 the moving mean and range are plotted as if they were simply individual values on a conventional mean (X-bar) and range chart.

	Value	4-day moving total	4-day moving mean	4-day moving range
Day 1	0.29			
Day 2	0.18			
Day 3	0.16			
Day 4	0.24	0.87	0.218	0.13
Day 5	0.21	0.79	0.198	0.08
Day 6	0.22	0.83	0.208	0.08
etc.	etc.	etc.	etc.	etc.

Median chart. This chart is easier to use than the X-bar chart. This is especially true for odd sample sizes where all that is necessary is to plot each individual value and then circle the middle (median) value. The disadvantage is that it is not as sensitive to variation. However Griffith (p. 32) states they are “very useful to monitor a process that has already had some level of improvements made.”¹³⁾ Once the median values of a series of samples have been determined, the “grand median,” the median of the medians, is used as the centerline of the chart. The “median range,” the median of the ranges, is then used along with appropriate factors¹⁴⁾ to determine the control limits for the median chart. A regular range chart can be used in conjunction with the median chart for monitoring spread or, as Oakland says (p. 161) it may be easier to use the median

13) I believe he is saying that the process is now at a point where it is generally running in a controlled way and only nominal monitoring is required.

14) From readily available tables.

range to calculate those control limits also.

Exponentially weighted moving average (EWMA) chart. According to Wadsworth (p. 45.22), this chart is supposedly better at detecting small shifts in the mean (but not large shifts). The idea is to give more weight to the previous moving mean plotted when calculating the current moving mean. To do this a “smoothing constant” (as Oakland puts it) between 0 and 1 is chosen. According to Oakland, 0.2 is commonly used. The example Oakland uses to illustrate this chart is data measured in centistokes (cSt), a unit related to viscosity. As before when the process is stable a value for the average mean would be determined, in this case 80.00 cSt. The

table to the right shows how each new weighted moving mean is calculated using a smoothing constant, a , equal to 0.2. The control limits are calculated in a similar manner as before using appropriate formulas.¹⁵⁾

	Column 1	Column 2	Column 3	
$a = 0.2$	Viscosity value (cSt)	a X new value (column 1)	$(1-a)$ X previous value	new moving mean*
				80.00
Batch 1	79.10	15.82	64.000	79.82
Batch 2	80.50	16.10	63.856	79.96
Batch 3	72.70	14.54	63.965	87.50
Batch 4	84.10	16.82	62.804	79.62
etc.	etc.	etc.	etc.	etc.

*This is the sum of last two entries in columns 2 and 3.

Cumulative sum (CUSUM) chart. Charts like the X-bar/R and various attribute charts we looked at are meant to trigger action based on the last point plotted, it would be good to if a chart used more of the data available for this purpose. The moving average and moving range charts and the EWMA chart do help in this regard by taking into account part of the previous data. However, according to Oakland, a chart that uses all the information available is the cumulative sum or CUSUM chart. Oakland (p. 225) lauds it as “...one of the most powerful management tools available for the detection of trends and slight

15) See references such as Wadsworth or Oakland for details.




Figure 10. A CUSUM chart compared with a standard X-bar chart (from Wadsworth, 1999, p. 45.19)

changes in data.” The drawback is construction of the chart is rather involved. In fact, Oakland devotes a full chapter to this. Rather than try to fully explain this type of chart Figure 10, from Wadsworth, shows an example of such a chart and how it compares with a standard X-bar chart. As seen in Figure 10, something called a “V-mask” is constructed to provide an objective way to determine if a significant shift in the process has occurred. As might be guessed, the scaling of the CUSUM chart is critical. Another commonly used method for determining if a significant change has occurred is the use of “decision intervals.” The reader is referred to other sources, such as Oakland, for details.

Short run charts. These control charts are used when, as the name implies, the amount of data is not sufficient to take, say, 20 or 25 samples as we do for the traditional X-bar/R charts. This may be due to the fact that only a certain relatively small number of a particular part is to be produced at a time or due to

a the process cycling so fast the production run is over before the data needed can be gathered. Various short run charts have been devised to permit control of these process for both variable data and attribute data. A key feature of these types of charts is the use of coded data based on some target value. The reader is referred to other sources for the details of constructing short run charts. Griffith discusses this type of chart in considerable detail for both variable and attribute data.

Concluding remarks. As mentioned, there are many different types of control charts to accommodate various situations or the need for more sensitivity to process changes. This section has touched on a some of the more common ones. Examples of others not mentioned here include the mid-range, multivariate, and the Box-Jenkins Manual Adjustment charts.

8. Summary and Conclusion

The purpose of this paper has been to provide an introductory look at statistical process control (SPC) concentrating mainly on control charts since they are at the heart of the matter. Several common types of control charts have been described in detail, two for variable data and four for attribute data. Some other types of control charts have also been briefly described in section 7. There is much, much more to SPC than given here such as more detail on how to interpret these charts and, even before that, how to judiciously select the variable/attribute to be charted. Perhaps the most important thing to remember is that control charts are not meant for controlling some quality characteristic but rather for controlling the process and, by virtue of that, making sure the characteristic being measured (and others inherent in the part/product) remains within desired values. In fact, control charts should be considered as a means to go beyond simply maintaining a “status quo” and as a tool for getting to know the process better and better and learning of ways to improve it!

References

- Oakland, J. S. (2008). *Statistical Process Control* (6th edition). Oxford: Elsevier.
- Griffith, G. K. (1996). *Statistical Process Control for Long and Short Runs* (2nd edition). Milwaukee: ASQC Quality Press.
- Wadsworth, H. M. (1999). Section 45, Statistical Process Control. In J. M. Juran (Ed.), *Juran's Quality Control Handbook*, 5th ed. New York: McGraw-Hill.