New Approaches to SPC

Implementing control charts appears easy, but appearances can be deceiving.

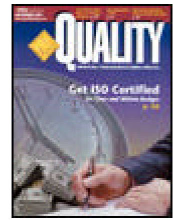

By Steve Wise and Doug Fair

The meaningful and correct application of control charts can be deceptively difficult. At first glance, introducing a control chart to the shop floor seems easy enough. Pick an important parameter to measure, determine a sampling strategy, teach someone how to plot the points and voilá--a control chart is born. The results: audit scores are higher, C_{pk} values get reported and customers are happy to see statistical process control (SPC) activity. Using control charts might even result in cost savings. However, do the alarms generated by control charts provide justification to perform process interrogation? Are the sources of variation captured within each subgroup useful for identifying true process changes? These are just some of the questions that must be answered before the chart can be considered a successful process improvement tool.

Not to diminish the importance of passing audits and making customers happy, but a control chart's true function is to provide real-time feedback needed to control and improve processes. The enthusiasm enjoyed from the benefits listed above will fade if the data displayed on the charts do not help front line operators make better process decisions. Operators will reject SPC if they are required to feed data into a control chart that provides no useful information. Introducing useless statistical tools on the shop floor can negate improvement activities and jeopardize an SPC program's effectiveness.

Of the many SPC implementation gaffs that occur in the workplace, two mistakes are widespread. Unfortunately, these mistakes are not only the most common, but they are the most severe and pose the greatest risks to an SPC program.

SPC application mistake #1

The first mistake is incorrect subgrouping. A fundamental assumption for any traditional Shewhart control chart is that the individual data values within a subgroup are independent, meaning that one measurement in a subgroup should not be influenced by another. For example, an aerosol can whose height is a critical dimension is measured in three locations (A, B, and C) that are 120 degrees apart. These multiple checks are performed to ascertain a can's height uniformity. Because the three values come from the same can, each value within the subgroup size of three is dependent on, or related to, the other heights. That is, if the can is extra tall, all three height measurements will be high; if the can is short, all three will be low.

A common approach for monitoring can height is to measure the three heights, then plot the average and range of these three readings on an Xbar and R chart. The average and range values, provided in the table, Can Height Dataset (p. 35), contain a sample data set of ten subgroups gathered with this strategy, including the Xbar and range calculations. These are graphed in the Xbar and range chart, Incorrect Method (p. 35), and are in a state of statistical control, which indicates a consistent amount of within-subgroup variability. Because the range chart is in control, the C_p value can be calculated, which is 3.83. However, the Xbar chart is out-of-control, with seven of the ten plot points exceeding its control limits. One explanation for this is that the control limits are calculated incorrectly, but checking the math proves that the control limits are calculated correctly. Therefore, what could cause so many of the subgroup averages to be out of control? In addition, the high C_p value of 3.83 is questionable, but a check of the math verifies that it too has been calculated correctly. What is the problem?

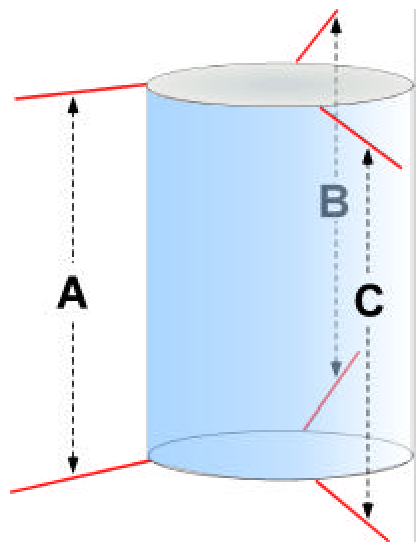

Aerosol Can Height Measurements - Three height measurements are made on a can.

	Can Height Dataset Subgroup Number									
Location	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
A	160.60	160.32	160.89	161.61	161.70	161.74	159.90	160.58	161.43	160.45
B	160.59	160.23	161.42	161.79	161.99	161.55	159.83	160.83	161.48	160.61
C	161.08	160.16	161.28	161.56	161.58	161.86	159.80	160.93	161.64	160.53
Average	160.76	160.24	161.20	161.65	161.76	161.72	159.84	160.78	161.52	160.53
Range	0.49	0.16	0.53	0.23	0.41	0.31	0.10	0.35	0.21	0.16
A sample data set of ten subgroups are gathered with this strategy, including the average and range calculations.										

	Can Height Data for 3-D Chart. Subgroups									
Location	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
A	160.60	160.32	160.89	161.61	161.70	161.74	159.90	160.58	161.43	160.45
B	160.59	160.23	161.42	161.79	161.99	161.55	159.83	160.83	161.48	160.61
C	161.08	160.16	161.28	161.56	161.58	161.86	159.80	160.93	161.64	160.53
Average	160.76	160.24	161.20	161.65	161.76	161.72	159.84	160.78	161.52	160.53
Moving R	-	0.52	0.96	0.45	0.11	0.04	1.88	0.94	0.74	0.99
Range/w	0.49	0.16	0.53	0.23	0.41	0.31	0.10	0.35	0.21	0.16
Calculations for plotting on a 3-D chart. This table is the same as above, with the addition of a moving range calculation.										

Let's explore the facts.

- The ranges represent within-can variation.
- The Xbar chart is based on can-to-can variation. Each Xbar value represents the average height of an individual can.
- The control limits are computed from the average range (Xbar ± A_2 Rbar).

The sampling procedure violates the assumption of data independence because the three values making up the subgroup are from the same part. However, appropriate subgrouping practice requires each measurement within a subgroup to be independent. If the subgroup were composed of one height measurement from three different cans, each individual measurement would not be influenced by any of the others and would be independent.

In the current subgrouping scheme, the Xbar control limits reflect within-can height variation, not can-to-can variation. Because the within-can variation is less than can-to-can variation, the Xbar chart appears to be out of control. But, the Xbar chart's state-of-control is incorrectly assessed because the wrong component of variation is used to calculate the control limits. The common "work around" in this situation is for operators to ignore the control limits and compare the plot points to the specification limits, completely disregarding normal process variation. This approach is also wrong because only individual measurements should be compared to the specification limits. Even if the process were in control, the C_p and C_{pk} values would be reflective of only within-can variation, while the can-to-can capability is unknown. After hearing this, a common response is to "just make it to spec" or give up on SPC altogether.

Incorrect Method

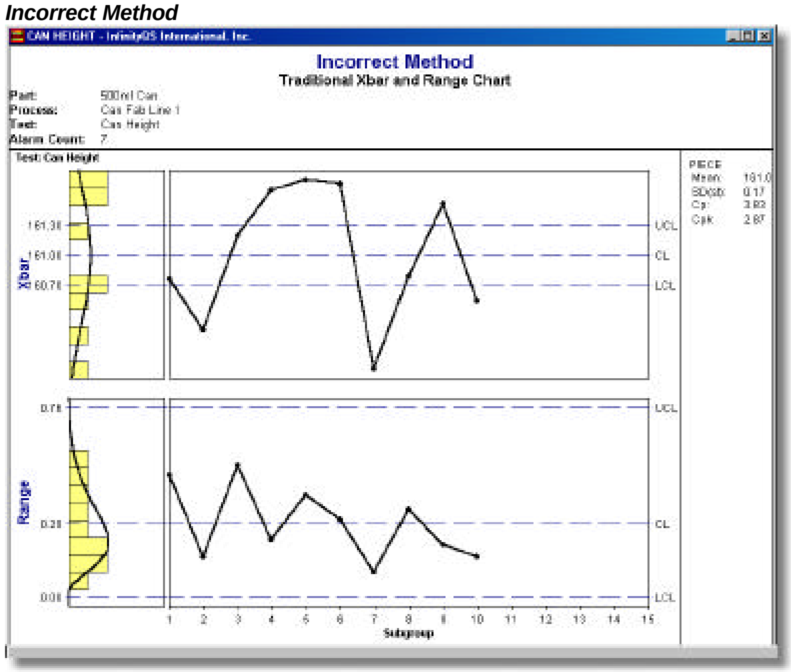

The X-bar and range chart graphs the Can Height Dataset. While the range chart is in control, the X-bar has seven of 10 points exceeding control limits.

Solving mistake #1

In addition to calculating a range, calculate the moving range of the averages. Each subgroup now has an average, a moving range and a within-subgroup range. These three "dimensions" of variation can be correctly analyzed with a 3-D chart, the solution to mistake #1.

A 3-D chart consists of an Xbar chart for displaying time-to-time variation, a moving-range chart for can-to-can variation, and a range chart for within-can variation. Even though the Xbar chart plot points are subgroup averages, they are treated as individuals on this chart because each average is the best one number representation of height for an *individual* can.

The 3-D chart, titled the Correct Method (p. 37), is constructed for the can height measurements. The R(within) chart is identical to the Range chart, titled the Incorrect Method, but it is more correctly called a Range chart for within-piece variation. The control limits are calculated in the conventional manner, D_3_R and D_4_R, where _R is the average of the within-can ranges. The subgroup size used to determine the D_3 and D_4 factors is based on the number of measurements per piece, which is three in this example. Control limits for the moving range chart are computed using the standard method where the D_3 and D_4 factors are based on n=2, minus the consecutive averages used to calculate the moving range.

The Xbar portion of this 3-D chart is treated as an Individual X (IX) chart, with each plot point representing the average height of a single can. However, the control limits for this special version of an IX chart are computed using the centerline of the moving-range chart (Xbar ± A_2 MRbar). Note that the within piece variability from the R(within) chart is not mistakenly used for calculating the control limits for the IX chart. Instead, the correct estimate of can-to-can variability is used, which is indicated by the MRbar.

With the correct control limits, the 3-D chart now indicates stability for three different statistics, time-to-time (IX chart), can-to-can (MR chart) and within-can (R(wi thin) chart). Now that control has been demonstrated, process capability can be estimated. The within-can C_{pk} is 2.87. The can-to-can capability, 0.77. Finally, the combined, or pooled, C_{pk} is 0.74.

The original, and incorrect, analysis led to the belief that the process C_p was 3.83, but this 3-D analysis correctly reports the C_p as 0.77.

Correct Method

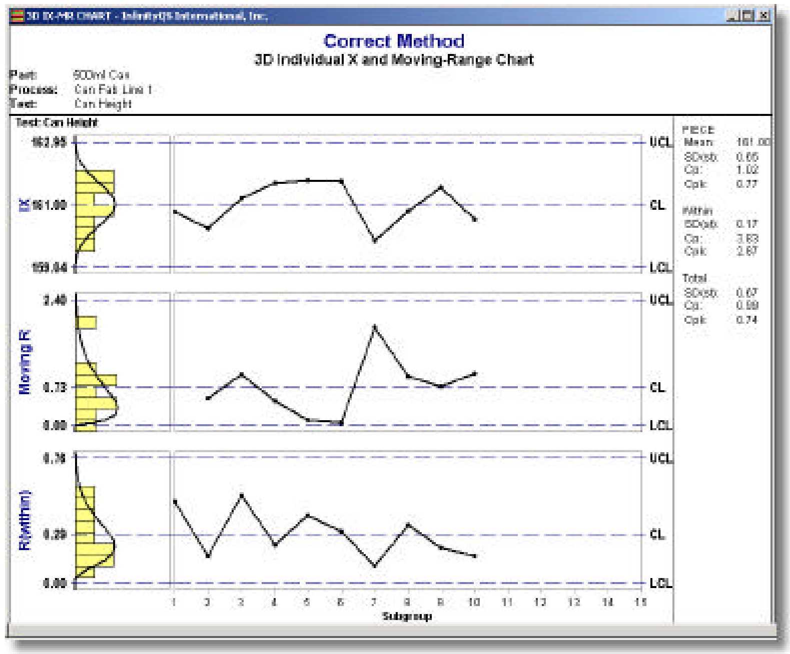

The 3-D Individual X and Moving Range chart for can height.

SPC application mistake #2

The second mistake is also related to incorrect subgrouping. Sometimes several process streams are combined in the same subgroup. For example, one process stream might be considered a machine or assembly line. Mixing process streams is not advisable because the standard deviations among the processes are more than likely different. The end result is control limits that do not correctly represent any individual process stream, and, thus, there is an inability to suitably interpret process signals on the control chart. This error can mask statistical signals and seriously degrade problem solving and process improvement.

Another example of this mistake might include mixing data from multiple spindles within the same subgroup or creating a subgroup whose injection mold data includes measurements across multiple cavities. To illustrate this mistake, consider a machine with eight filling heads, where each head is controlled by a separate pump. Because the heads are controlled separately, each can be considered a separate process worthy of its own individual control chart.

The data in the table, Bottle Fill Volume (p. 39), represents fill volumes for bottles filled by each head. A common strategy is to combine one fill volume from each head into a subgroup of size eight, then calculate the average and sample standard deviation, *s*.

Because of its incorrect subgrouping strategy, the resulting Xbar and *s* chart, titled Incorrect Method-Traditional Xbar/Sigma Chart (p. 37), appears to be in

Incorrect Method

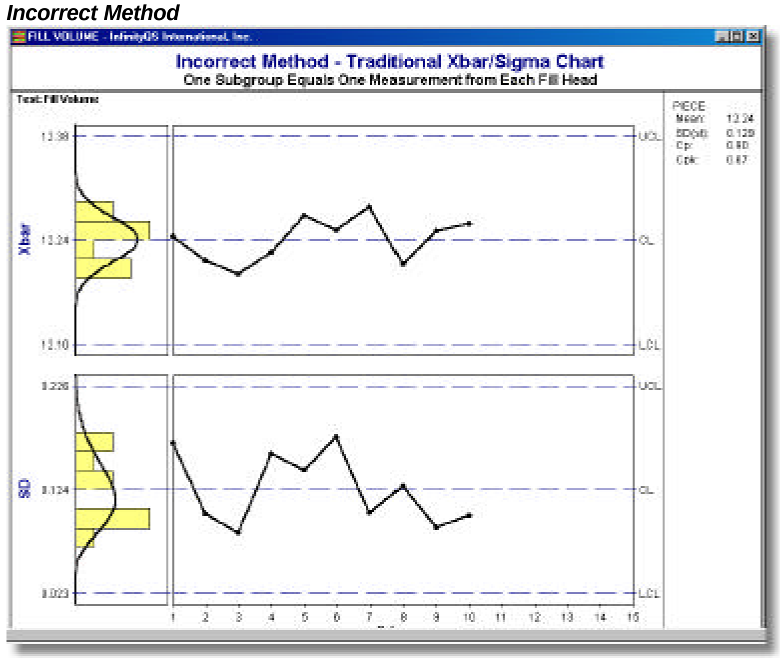

Single Xbar and s chart confounding all eight fill heads.

Xbar/Sigma Chart (p. 37), appears to be in control. All that can be said about this process is that bottles are being consistently filled with liquid, but this conclusion may be invalid because all eight heads are averaged into one value. If any of the heads were "misbehaving," the Xbar chart would probably mask the problem. Although combining the eight fill heads into a single subgroup seems convenient, it will seriously undermine the ability to problem solve and improve filling performance. Why? Because the differences between the heads are lost in the average and range.

To avoid the subgrouping mistakes associated with combining all eight heads into a single subgroup, plot each head on its own chart. This is statistically correct, but not very practical because this solution requires eight control charts to manage this one filling machine, as is seen in the graphic, One Chart Per Head (p. 38). Imagine trying to control a filling machine with 20 or 30 heads.

Another option might be to plot all process streams on a single chart, resulting in 8 separate data lines. However, displaying all data stream lines results in a "spaghetti chart" where the highs and lows are indistinct and confusing, as is seen in the graphic, Spaghetti Chart (p. 38). The confusion would grow as the number of heads increases.

One Chart Per Head

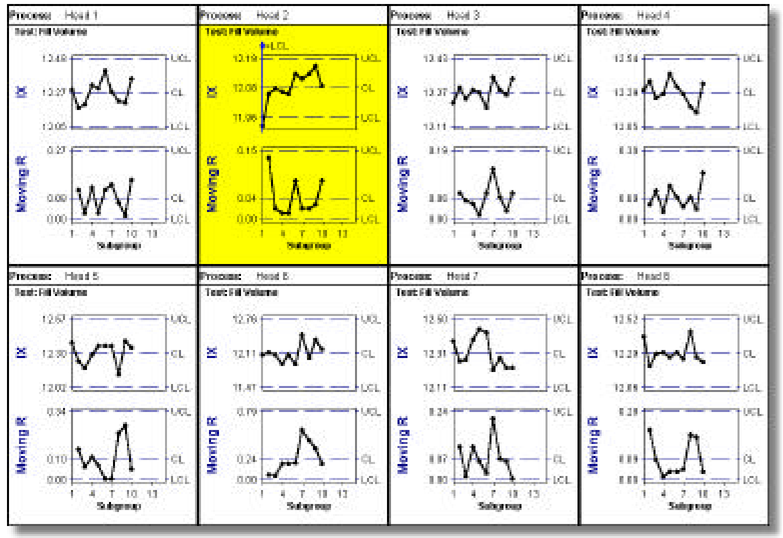

Each fill head plotted on its own chart.

Spaghetti Chart

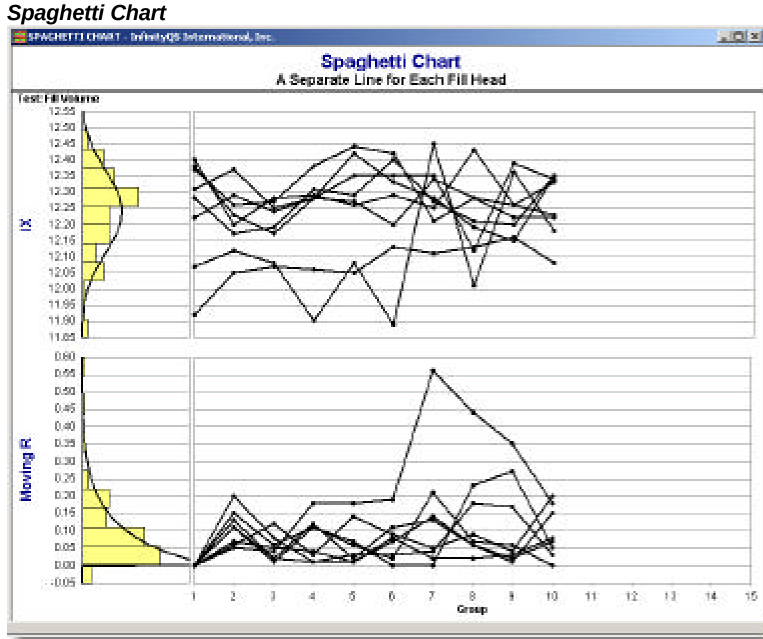

The Group IX-MR chart shows the MIN and MAX fill volume values (IX chart) and the MIN and MAX moving ranges from each sample (Moving R chart).

Solving mistake #2

The solution to this problem is a group chart that provides the simplicity of one chart, but maintains a separate analysis of each fill head. Instead of averaging the fill heads, or plotting each on its own chart, treat all eight individual values as a group representing fill volumes at a given period of time. From this group, only the minimum (min) and maximum (max) fill volumes are highlighted and plotted, as seen in the table, Minimum and Maximum Fill Volumes (p. 39). Then each min and max plot point on the chart is labeled with a number representing the associated fill head. The resulting chart is called a Group IX-MR chart.

The group chart is intentionally absent of control limits because sometimes group charts are used for plotting dependent data streams. As illustrated in the mistake, subgrouping dependent data results in incorrect control limits. The power of a group chart is two-fold: one, to clearly and distinctly illustrate the extremes in a data set; and two, to present the data to users so that opportunities for improvement are clearly indicated.

A group chart can be considered a running

Pareto chart of variables data--the most important information is displayed to the user, which facilitates the identification of improvement actions.

Correct Method-Group IX-MR Chart

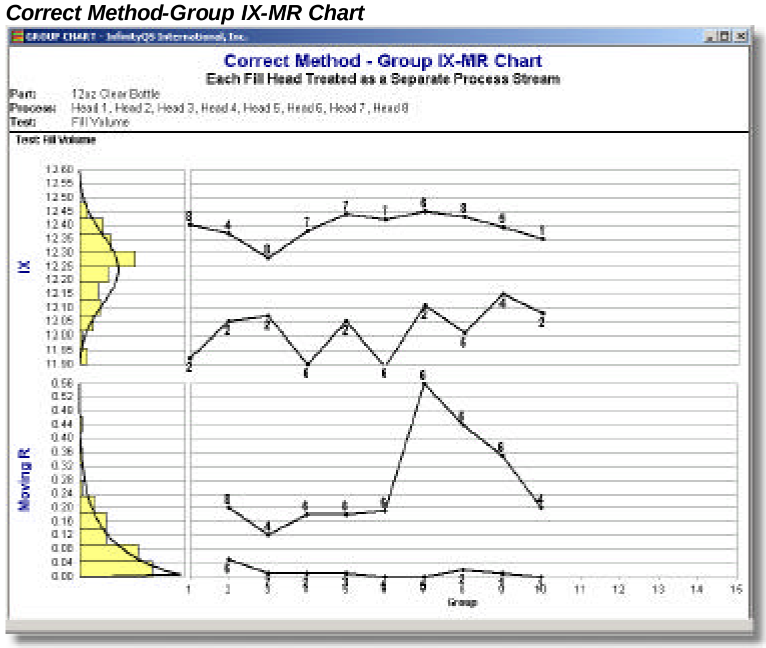

The Group IX-MR Chart shows the MIN and MAX fill volume values (IX chart) and the MIN and MAX moving ranges (Moving R chart) from each sample.

Interpreting the group chart

When interpreting group charts, look for runs in the min and max positions, as shown in the Correct Method Group IX-MR chart (p. 40). Notice that the min plot points on the IX chart are predominately represented by heads 2 and 6, which are consistently filling less volume than the other heads. Looking at the moving range chart, the majority of the max plot points are from head 6. Therefore, the variability from head 6 is consistently greater than the other heads. The distance between the IX chart's min and max lines displays the variation within a single rotation through all eight heads. That is, the closer the lines, the more consistent the fill volume across all heads.

There are several different variations of a group chart, but they all are designed to plot multiple process or data streams on a single chart. With just a single group chart, one can answer many questions about a process and promptly uncover tremendous amounts of information. Conversely, users armed only with traditional SPC methods would need to closely examine each of several different charts, while simultaneously juggling and comparing those charts and their center lines. They would also likely require the aid of off-line analysis by a statistician.

	Bottle Vill Volume Subgroup Number									
Head #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
1	12.28	12.17	12.19	12.31	12.29	12.40	12.27	12.21	12.20	12.35
2	11.92	12.05	12.07	12.06	12.05	12.13	12.11	12.13	12.16	12.08
3	12.22	12.29	12.24	12.28	12.27	12.20	12.34	12.28	12.26	12.33
4	12.31	12.37	12.25	12.28	12.42	12.33	12.28	12.30	12.15	12.35
5	12.38	12.23	12.17	12.28	12.35	12.35	12.35	12.30	12.39	12.34
6	11.37	12.12	12.08	11.90	12.08	11.89	12.45	12.01	12.36	12.18
7	12.37	12.26	12.27	12.38	12.44	12.42	12.21	12.28	12.22	12.22
8	12.40	12.20	12.28	12.29	12.26	12.29	12.25	12.43	12.26	12.23
Average	12.24	12.21	12.19	12.22	12.27	12.25	12.28	12.21	12.25	12.26
s	0.169	0.100	0.082	0.159	0.143	0.176	0.102	0.127	0.087	0.099
The bottle fill volume is organized with one subgroup consisting of one measurement from each of eight fill heads.										

	Minimum and Maximum Fill Volumes Group Number									
Head #	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
1	12.28	12.17	12.19	12.31	12.29	12.40	12.27	12.21	12.20	**12.35**
2	**11.92**	**12.05**	**12.07**	12.06	**12.05**	12.13	**12.11**	12.13	12.16	**12.08**
3	12.22	12.29	12.24	12.28	12.27	12.20	12.34	12.28	12.26	12.33
4	12.31	**12.37**	12.25	12.28	12.42	12.33	12.28	12.30	**12.15**	12.35
5	12.38	12.23	12.17	12.28	12.35	12.35	12.35	12.30	**12.39**	12.34
6	12.07	12.12	12.08	**11.90**	12.08	**11.89**	**12.45**	**12.01**	12.36	12.18
7	12.37	12.26	12.27	**12.38**	**12.44**	**12.42**	12.21	12.28	12.22	12.22
8	**12.40**	12.20	**12.28**	12.29	12.26	12.29	12.25	**12.43**	12.26	12.23
Values in red represent the **MAX** volume from the group and blue represents the **MIN** volume.										

SPC solves problems

With SPC's ability to put tremendous amounts of information and problem-solving knowledge into the hands of operators, it is surprising how many companies have not adopted SPC or have abandoned its use. Surprising, that is, unless the situations that bedevil companies trying to leverage the power of SPC are analyzed. Upon closer inspection, it is likely that there are frustrated operators, misinformed managers and a lack of support for statistical methods on the manufacturing shop floor. Why? Because there is a general lack of understanding of correct subgrouping conventions and a belief that SPC "just doesn't work" in their particular situations. More likely, the lack of success could be attributed to an incorrect subgrouping strategy as well as the use of the wrong statistical tool.

For the most part, users are familiar with the standard Xbar and R chart. However, the 3-D chart is a superior tool for evaluating three important sources of variation: time-to-time, piece-to-piece, and within-piece. Additionally, the Group Chart is a clever analysis tool that allows users to visualize the individual performances of multiple process streams on a single chart. With these tools in the hands of operators and support personnel, SPC has the opportunity to become not only less troublesome, but infinitely more capable of providing the critical real-time information required for sustaining a successful SPC system.

VIEWING MULTIPLE PROCESS STREAMS ON A BOX PLOT

Group charts are great for time-ordered analyses of multiple process streams. Another tool used to compare multiple process streams is the box plot. Like a histogram, the box plot illustrates a data set's distribution. A histogram can be thought of as the distribution's front view, while a box plot represents its top view.

In the chart, Box Plot, the yellow box plot represents the combined distribution of sampled measurements across all fill heads. Each blue box represents the distribution of fill volumes from each head. The blue boxes, when combined, make up the yellow box and whisker plot. A quick scan of the blue boxes shows that heads 2 and 6 are filling less volume than the others. By looking at the width of their boxes, Head 6 appears to exhibit most variability while head 3 is exhibiting the least. This interpretation agrees with the conclusions derived about process performance from the previous analysis of the group chart.

Box Plot View

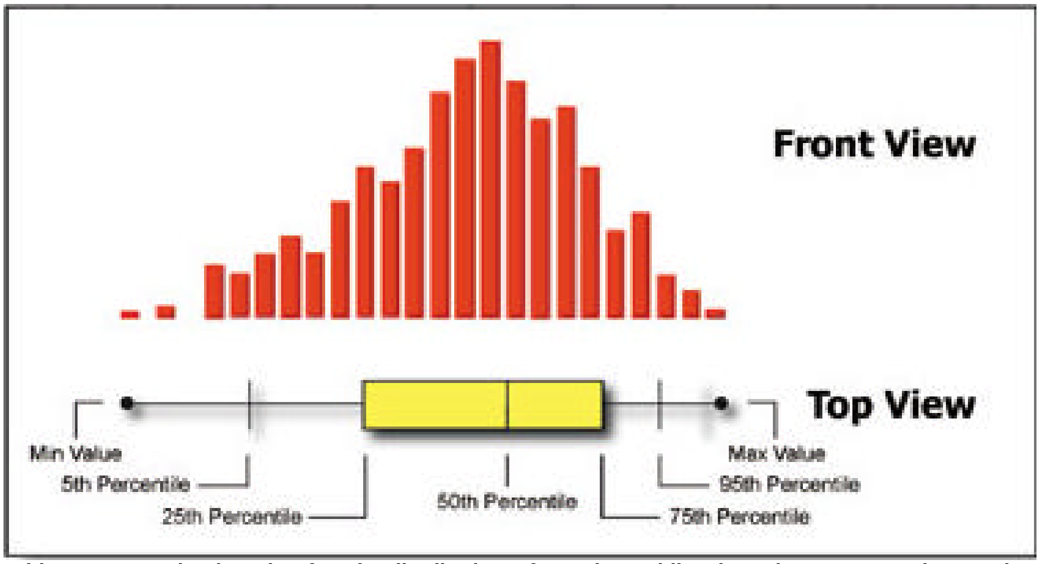

A histogram can be thought of as the distribution's front view, while a box plot represents its top view.

Box Plot

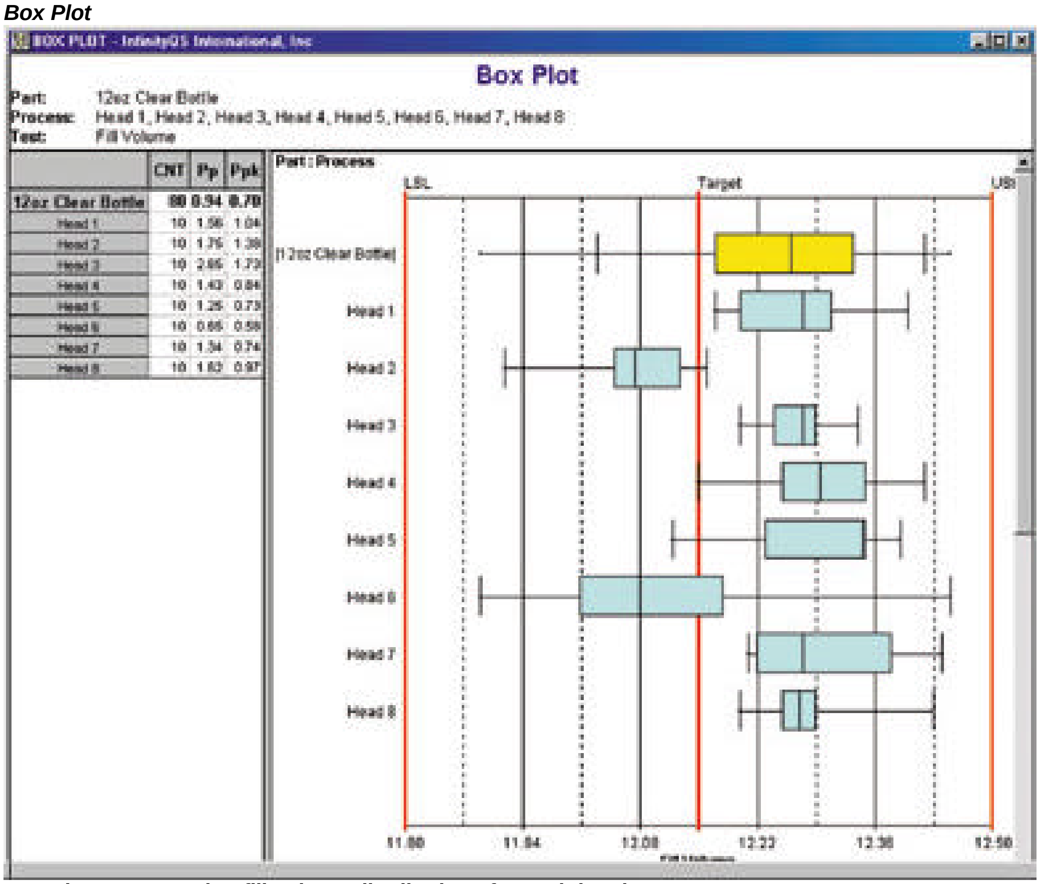

Box Plots representing fill volume distributions for each head

Steve Wise and Doug Fair are the Directors of Statistical Applications for InfinityQS International Inc. (Manassas, VA), a quality systems software company. They are co-authors of ***Innovative Control Charting: Practical SPC Solutions for Today's Manufacturing Environment*** published by ASQ Quality Press, 1998. Fair is also the co-author of ***Principles and Methods of Quality Management in Healthcare***, Aspen Publishing, 2000. For more information, e-mail swise@infinityqs.com or dfair@infinityqs.com.