An SPC-Based Filter

R. Russell Rhinehart – 29 June 2016 – www.r3eda.com

A simple procedure is developed which uses concepts from statistical process control to filter noise from a process variable.

For the original publication, see Rhinehart, R. R., "A CUSUM-type of on-line Filter," Process Control and Quality, Vol 2, 1992, p. 169-176. For applications see Szela, J. T., and R. R. Rhinehart, "A Virtual Employee to Trigger Experimental Conditions", Journal of Process Analytical Chemistry, Vol. 8, No. 1, 2003; Mahuli, S. K., R. R. Rhinehart and J. B. Riggs, "pH Control Using a Statistical Technique for Continuous On-Line Model Adaptation," Computers and Chem. Engr., Vol 17, No 4, 1993, p. 309-317; Rhinehart, R. R. "A Statistically Based Filter", ISA Transactions, Vol. 41, No. 2, April 2002, pp 167-175; and Muthiah, N., and R. Russell Rhinehart, "Evaluation of a Statistically-Based Controller Override on a Pilot-Scale Flow Loop", ISA Transactions, Vol. 49, No. 2, pp 154-166, 2010

Development

Process variables (either measured or calculated) are subject to noise (uncertainty, error, natural vagaries); and since automatic control decisions in any form (coefficient adjustment, process change) are calculated from the measurements, they consequently propagate the noise. This propagation of noise is counter to a principle of statistical process control (SPC), which directs that action should be in response to real process changes, not noise. "Tampering" is the SPC label for control action that responds to noise, which causes unnecessary change and actually increases variability.

Traditional deterministic filters (first-order and more complex) do not remove noise. They reduce it, but some action is still influenced by the noise. Further, in partially reducing noise, they create an undesirable lag in the filtered process variable.

This tutorial reviews the development of a method for statistically-based filtering and summarizes its application to a pilot-scale unit.

Consider Figure 1, an illustrative "Shewhart" individuals chart of a process variable (PV) response with respect to sampling number.

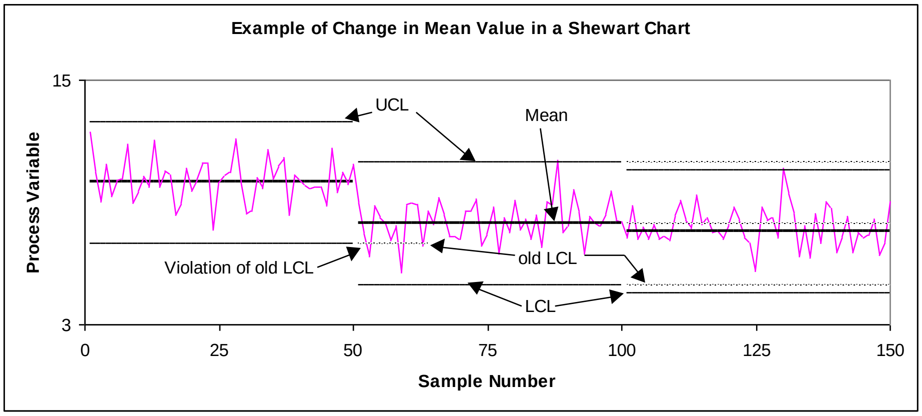

Figure 1. Principles illustrated by a modified Shewhart Chart.

The process is at steady conditions during the 0 to 50 sampling period with a mean value of 10 units and standard deviation of the noise value of 1 unit. $\bar{X}$ is the average (labeled as "Mean" on the figure), and

$\hat{\sigma}_x$ is the estimated population standard deviation from the N data points on the x versus time graph. These values are conventionally calculated as:

$$\bar{x} = \frac{1}{N} \sum_{i=1}^{N} x_i \qquad \textbf{(1)}$$

$$\hat{\sigma}_x = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (x_i - \bar{x})^2} \qquad \textbf{(2)}$$

If the PV is “stationary” its mean and variability are constant (the process is at steady conditions). If “noise” can be considered as independent, random influences added to the mean at each sampling, and if the noise is normally (Gaussian) distributed, there is only a 0.27% chance that data will appear in the tails of the histogram (more than 3 $\hat{\sigma}_x$ from the average). Only 1 out of 370 points will fall in the extreme tails. The upper and lower 99.73% confidence limits (UCL and LCL) are also shown on Figure 1.

The filtering concept of this paper attempts to find the mean within noisy data from a stationary process. The estimate of the mean is unchanged until there is sufficient statistical confidence that the value is incorrect.

If the PV “really” changes, then the average will shift, and there is a good chance that the new process will place points outside of the old 3 $\hat{\sigma}_x$ limits. Again refer to Figure 1. At sampling number 50 the mean shifted to a value of 8 units, and it sustains that value throughout the 50 to 100 sample period. At sample number 52 a measurement violates the old LCL, indicating that a shift in the mean has probably occurred. The SPC philosophy is that if x_i is more than $3\hat{\sigma}_x$ from the old average, then one will accept that the process really has changed. Although one will be wrong 0.27% of the time, the basic SPC philosophy is:

If $\left| x - \bar{x}_{old} \right| > 3\hat{\sigma}_x$, then respond. **(R1)**

But, Rule (R1) waits until one point is in the improbable tail region. If the process shift is small compared to σ_x (see the 100 to 150 sample period of Figure 1 where the mean shifted to a value of 7.8 units), then it may take many points before the first measurement is more than $3\hat{\sigma}_x$ from the average. However, if there is a shift, then the measurements will show a systematic bias from the old average. This is illustrated in Figure 1 where the data from the 100-150 samplings are uniformly scattered about the new average, but show a definite bias about the old average. There is no violation of the old $3\hat{\sigma}_x$ individual limit, however, there is a cumulative bias, averaging 0.2 units at each sampling.

Define

$$\text{CUSUM} = \sum \frac{(x - \bar{x}_{old})}{\hat{\sigma}_x} \qquad \textbf{(3)}$$

If the process is stationary (mean and variance remain constant in time), and noise is independent at each sampling, and the average is the mean, then we expect CUSUM to be a random walk variable starting from an initial value of zero. However, if the process has shifted, then $|\ CUSUM\ |$ will steadily, systematically grow with each sampling. CUSUM is the cumulative number of $\hat{\sigma}_x$'s that the process has deviated from $\overline{x}_{old}$. CUSUM could have a value of 3 if there is a PV violation of the $3\hat{\sigma}_x$ limit. Regardless as to whether it is a single or accumulated value of 3, for this procedure,

If $|\ CUSUM\ | > 3\sqrt{N}$ take action **(R2)**

The variable N is the number of samples for which the variable CUSUM is calculated, and the square root functionality arises from a propagation of variance on Equation (3). The trigger value, 3 (three sigma), representing the traditional SPC 99.73% confidence level in a decision, is not a magic number. If the value was 2, nominally representing the traditional 95% confidence level of economic decisions, then action will occur sooner; but it will be more influenced by noise: 5% of data in a normal steady process will have a $\pm 2\sigma_x$ violation. If the trigger value was 4 representing the 99.99% confidence level, then the filtered value will be less influenced by noise, but the procedure will wait longer to take action.

While 3 is the conventional SPC trigger, values from 2 to 4 are generally chosen to balance responsiveness and false alarms for particular SPC applications. For integrating SPC and automatic control, often termed algorithmic process control, I prefer the traditional economic decision trigger values of about $\pm 2\sigma_x$, representing a 95% confidence in a decision, give "best" results. Choose TRIGGER=2, and

If $|\ \mathrm{CUSUM}\ |$ > TRIGGER * $\sqrt{N}$, then take control action. **(R3)**

The action will be to report a change in the mean of the noisy process variable. How much change should one take? Short of detailed process model inference procedures, primitively credit that $\overline{x}$ had sustained an average offset since the last change in $\overline{x}$. Since the previous CUSUM was about zero, the CUSUM has now accumulated a value of $N \cdot \left(\overline{x}_{new} - \overline{x}_{old}\right) \sigma_x$. Accordingly, if Rule (R3) is true, then reset the filtered value by:

$$\overline{x}_{new} = \overline{x}_{old} + \hat{\sigma}_x \cdot \mathrm{CUSUM} / \mathrm{N} \qquad \textbf{(4)}$$

And reset the counter N and CUSUM to zero.

My preference is to eliminate possible instances of a divide-by-zero in a computer procedure (which may occur is the PV value is "frozen" and $\hat{\sigma}_x$ becomes zero). So, re-define

$$\mathrm{CUSUM}_{new} = \mathrm{CUSUM}_{old} + \left(x - \overline{x}_{old}\right) \qquad \textbf{(5)}$$

Then the filter decision will be executed at each sampling. Coded in VBA it is:

```
N = N +1
CUSUM = CUSUM + X – XSPC
IF ABS(CUSUM) > TRIGGER * SIGMAX * SQR(N) THEN
        XSPC = XSPC + CUSUM/N
        N = 0
```

```
            CUSUM = 0
        END IF
```
(C1)

Here XSPC is the SPC-filtered value of x, the estimated average value of x.

The method requires a value for $\hat{\sigma}_x$, "SIGMAX", and the value should automatically adjust when the process variability changes. Based on a technique detailed in Cao, S., and R. R. Rhinehart, "An Efficient Method for On-Line Identification of Steady-State," Journal of Process Control, Vol. 5, No 6, 1995, pp. 363-374, define

$$\hat{\rho}^2 = \frac{1}{M-1}\sum_{i=1}^{M}\left(x_i - x_{i-1}\right)^2 \tag{6}$$

where ρ^2 is a measure of σ^2, the process variance. Conventionally,

$$\hat{\sigma}^2 = \frac{1}{M-1}\sum_{i=1}^{M}\left(x_i - \bar{x}\right)^2 \tag{7}$$

A bit of algebra reveals that in the limit of large M and for a stationary process, ρ^2 is an unbiased estimate of σ^2:

$$\lim_{M\to\infty}\hat{\rho}^2 = 2\overline{\sigma}^2 \tag{8}$$

Equation (6) prescribes a computationally expensive method to calculate ρ^2. Therefore, as an equivalent approach to estimate ρ^2 this work uses

$$\hat{\rho}^2_{f_{new}} = \left(\frac{M-2}{M-1}\right)\hat{\rho}^2_{f_{old}} + \left(\frac{1}{M-1}\right)\left(x_{new} - x_{old}\right)^2 \tag{9}$$

Equation (9) is an exponentially weighted moving variance (EWMV), or a first-order filter, on the value of ρ^2 requiring one storage location and few arithmetic operations at each update. Alternately, Equation (9) can be written in forms that may be more familiar to the reader.

$$EWMA_{new} = (1 - \lambda)\ EWMA_{old} + (\lambda)\ Signal \tag{9a}$$

$$Filtered_{new} = e^{-T/\tau_f}\ Filtered_{old} + (1 - e^{-T/\tau_f})\ Signal \tag{9b}$$

where $\lambda = \frac{1}{M-1} = 1 - e^{-T/\tau_f}$

Conceptually, the SPC filtering procedure is to first estimate $\hat{\rho}^2$ from Equation (9), then $\hat{\sigma}_x$ from a rearrangement of Equation (8):

$$\hat{\sigma}_x = \sqrt{\hat{\rho}^2_{f_{new}} / 2} \tag{10}$$

Then to use this value of $\hat{\sigma}_x$ in the CUSUM test. Whether one uses Equation (6) or (9) to estimate the variance in the process noise, one has to choose a value for M, the number of past data points that are used in the estimation of variance. How should one choose a value for M [or λ in Equation (9a) or τ_f in Equation (9b)]? Each new value of $\hat{\rho}^2_f$ is influenced by each new value of $(x_{new} - x_{old})^2$. If $\lambda \left[= \frac{1}{M-1} \right]$ is very small, then each random value of $(x_{new} - x_{old})^2$ will have a small influence on $\hat{\rho}^2_f$, and $\hat{\rho}^2_f$ will not vary much from the true value. Therefore, a large value for M seems desired. However, if the process σ_x changes, and M is large, then $\hat{\rho}^2_f$ will be very slow to respond. So, one wants a small M (or large λ, or small τ_f) for responsiveness on σ_x. I find that M ≅ 11, or $\lambda \cong .1$, gives the "best" balance of removing variability from the estimate, yet remaining responsive with a human-convenient λ value.

For computationally simplicity Equations (9) and (10) can be combined.

$$\hat{\sigma}^2_{f_{new}} = \left(\frac{M - 2}{M - 1} \right) \hat{\sigma}^2_{f_{old}} + \left(\frac{1}{M - 1} \right) \left(\frac{1}{2} \right) \left(x_{new} - x_{old} \right)^2 \tag{11}$$

So the elementary VBA algorithm for the SPC filter becomes:

```
N = N + 1
V = 0.9 * V + 0.05 * (X – XOLD) ^ 2
XOLD = X
CUSUM = CUSUM + X - XSPC
If ABS(CUSUM) > TRIGGER * SQR(V * N) THEN
        XSPC = XSPC + CUSUM / N
        N = 0
        CUSUM = 0.0
END IF
```
(C2)

Here V, in Line 2 of Code (C2), represents the noise variance as calculated from Equation (11) using M=11.

There is no need to initialize V, XSPC, XOLD or CUSUM with their true values. If everything is initialized with zero; then, on the first call, N=1, V = 0.05 * X^2, XOLD = X and CUSUM = X. Since $|CUSUM| = |X| > TRIGGER* \sqrt{0.05X^2} = .2236 * TRIGGER * |X|$, on the first call XSPC will automatically be initialized to 0 + X/1 = X, the initial value. (As long as TRIGGER is less than 4.47 (=1/0.2236). Recall that values of 2 to 4 are standard practice.) So, the complete algorithm in VBA is:

```
IF first_call THEN
        N = 0
```

```
        XOLD = 0.0
        XSPC = 0.0
        V = 0.0
        CUSUM = 0.0
        M=11
        FF2 = 1.0 / (M - 1) / 2.0
        FF1 = (M - 2) / (M - 1)
        First_call = "FALSE"
END IF
Obtain X
N = N + 1
V = FF1 * V + FF2 * (X – XOLD) ^ 2
XOLD = X
CUSUM = CUSUM + X - XSPC
IF ABS (CUSUM) > TRIGGER * SQR (V * N) THEN
        XSPC = XSPC + CUSUM / N
        N = 0
        CUSUM = 0.0
END IF                                                                  (C3)
```

Simulation

Available on the www.r3eda.com site is a simulator for testing the algorithm, and comparing the SPC filter to a conventional first-order (FO) filter. It generates a noisy signal that can make either a step or a ramp change after an initial steady-state. You can specify the initial and final values of the PV and its noise, and the point of change. The simulator calculates several statistics that can be used to assess a filter performance. Key ones are root-mean-square (rms) deviation of the filtered value from the true value, and the settling time (time after a step change in level for the filtered value to be within a half-sigma of the true).

Figure 2 compares the SPC and FO filtered responses to one data realization. Initially, the true PV value is 8 units during the 0 to 50 sampling period, and the variance is unity. At a sampling of 50, the true value begins a ramp to a value of zero at the 200th sampling. The standard deviation simultaneously ramps from unity to a value of zero. The trigger value for the SPC filter is 2. The FO filter coefficient is 0.053, which gives the same rms value as the SPC filter during a stationary period. The data are represented by the dots, the solid line is the SPC filtered value and the dashed line is the FO filtered value.

Figure 2. Simulation study.

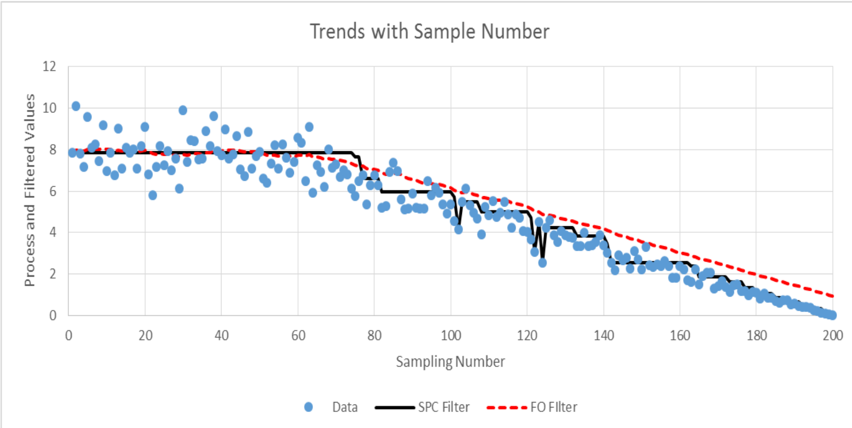

Note several aspects contrasting the FO and SPC filters: First, the SPC filter generally does a better job of representing the true PV value during the ramp period. The FO filter lags behind.

Second, as the process variability reduces (see the 180 to 200 sampling period), the SPC filter acts more aggressively, but with smaller increments, to track the signal. Because of the EWMA estimate of the variance, it is a self-tuning filter. It adapts to the noise level.

Third, during the initial steady period, the SPC filter does not fluctuate, it holds a constant value, while the FO filter fluctuates as each value is influenced by the sampling PV value.

Fourth, when the ramp change started at sampling 51, the SPC filter held its prior value. Until a sampling number of about 75 there was insufficient evidence to implement a change. By contrast, the FO filter begins noticeably changing value at sample number 60. This delay, until there is adequate statistical evidence to make a change is characteristic of the SPC filter.

Experimental

Tests were performed on a pilot-scale two-phase flow apparatus from the Chemical Engineering Laboratory at Oklahoma State University. Air and water flow upward through a vertical tube 14-foot tall with a 1.5 inch diameter. The process is computer controlled with conventional flow control valves, orifice flow meters, and pressure transducers. Instrument signals are either 4-20 mA or 3-15 psig. The SPC filter sampling rate was set at 2 Hz. At faster sampling rates the noise on the differential pressure signal showed autocorrelation, which violates the noise independence assumption which allows ρ_f^2 to be an unbiased estimate of σ^2.

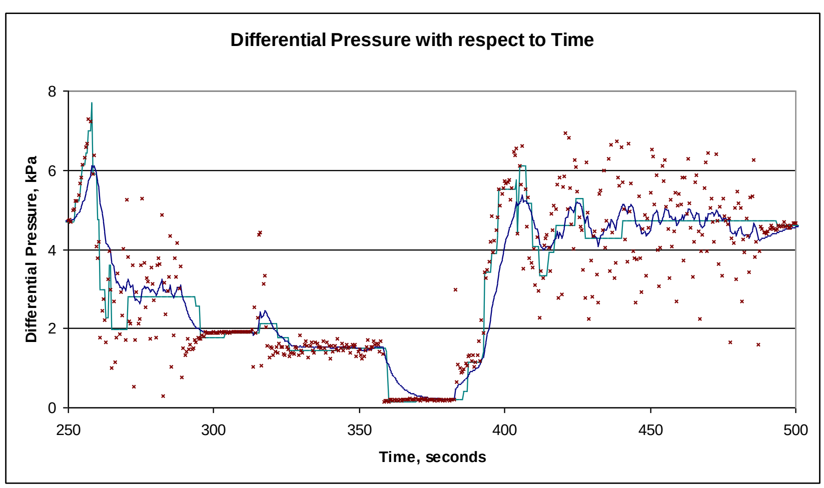

Figure 2 shows a plot of the differential pressure with respect to time during an experimental run that had changes in air and water flow rates. The changes placed the flow in several patterns (stagnant, bubbly, slug, churn, or annular), which had very distinct values for both the pressure drop and the noise patterns.

Figure 3. SPC filter performance compared to a first-order filter. Experimental data.

Observe several features of the signal in Figure 3. The average value of the PV changes significantly over a relatively short time period – of a few samplings. Some changes have the character of steps and others of ramps. Some data events are short-lived spikes (pulses or spurious signals), while others are first-order like. The noise level changes significantly, about 40:1, over short time periods.

The SPC filter responds faster than the first-order filter when the PV changes, and ignores the noise during the hold periods of constant conditions. By contrast, the conventional filter wiggles mildly throughout the same time window.

One might claim that the conventional filter is tuned for too much filtering, and that the lag is an artificial contrivance of the author. However, the noisy data that follows indicates insufficient filtering by the conventional filter.

Discussion

This method assumes that noise is a random, independent (un-correlated) deviation with a zero mean. However, filtering or lags related to process instrumentation, or persistent perturbations of a process will cause autocorrelation in the PV signal. The process/control engineer should choose a sampling interval greater than about one-half of the lifetime of such noise transients for which no filter action is desired, otherwise the PV data will be autocorrelated and the filter will track them. In the experimental work of Figure 3, the 0.5 second sampling interval (2 Hz) was chosen because it eliminated autocorrelation. However, I find that it is simpler to have the filter sampling rate match the controller/system rate and to let the engineer/operator tune the trigger value for the desired SPC-Filter performance. Autocorrelation makes V smaller and CUSUM larger than they "should" be, requiring a larger value of TRIGGER (2.5 perhaps instead of 2).

In this development the CUSUM value is credited to be due to a constant bias over the last N samples and the new SPC-filtered PV value is calculated from that assumption. This may not be true. A large CUSUM value may reflect, for example, a ramp change in PV or a step change in PV level after a long steady period; and, if so, the procedure for adjustment the SPC-filtered value will not make a sufficient change. However, I find that in such an event the next adjustment will be both soon and sufficient, and attempts to model PV change add complexity with little real tracking benefit.

A simple CUSUM formulation and resetting criteria was proposed. Certainly, one can argue for some of the traditional CUSUM techniques of dual "+" and "–" sums, "slack" adjustment, or for time-scheduled XSPC and CUSUM resetting methods. Perhaps, the performance enhancement of such methods would justify the additional complexity and user choices. However, I find the method presented above to be both simple and sufficient.

Here the SPC filter is applied to a PV which was used in process monitoring, not as a controlled variable. But, the technique has been useful for tempering the output of a controller, for batch-to-batch recipe corrections, and for model coefficient updating on-line.