Prevent Noise from Degrading Control Performance
A DYP article in CONTROL Magazine
R. Russell Rhinehart 2026-08-21
How to prevent noise from degrading control performance | Control Global

This is a 3-part article related to several types of "noise" in measurements that corrupt knowing the true process value, and how to compensate for the noise features.

PART 1 – Introduction and Conventional Concepts

Process noise confounds our interpretation of 1) when a change starts, 2) when the process variable (PV) reaches steady state, and 3) what is the true PV value. See Figure 1. The dots represent a noisy measurement of a process that moved from one steady value to another, and the line is a first-order filtered value of the data. Looking at the filtered value, one might claim that the transient started at a time of 3, and reached steady state at a time of 20. Alternately, looking at the noisy measurements, the beginning and end of the transient might be considered at times of 2 and 15.

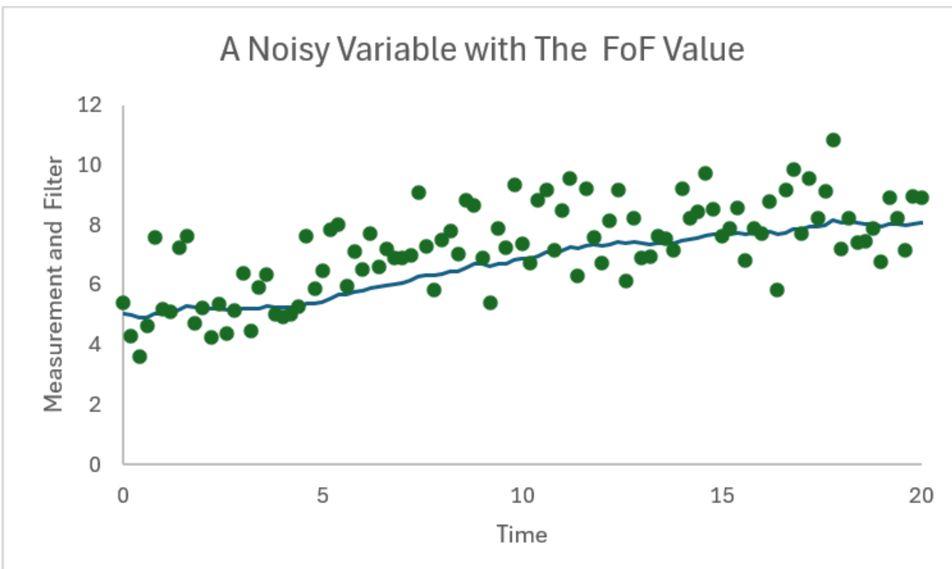

Figure 1 – Illustration of a shift in a noisy variable and its first-order filter

During a transient the filtered value lags behind. Notice that nearly all the measurement values in the time interval of 5 to 10 are above the filtered indication. If the filtered PV value were to be used in on-line process analysis, for example in material balance closure, the filtered value would misrepresent the process value during the transient. However, during the initial and ending steady state values, the first-order filtered value is in the centroid of the data, which does provide a relatively noiseless and valid indication of the true PV value.

Notice that during either the initial or final steady state, the filtered trace is not steady; it wiggles a bit responding to the data perturbations. One could use less filtering to make the filter track the transient faster; but then undesirably, it would not reduce the noise as much during steady periods.

The lag and continual variation are undesirable trade-off features of first-order filtering.

This is simulated data, and Figure 2 also reveals the true (but unknowable) PV trend, which generated the data for Figure 1. The true PV value makes a first-order rise, starting at a time of 4 and reaching 95% of the way to its asymptotic final value at a time of about 10. Neither observation from the data nor from the first-order filtered value are great indications of the transient start or end times.

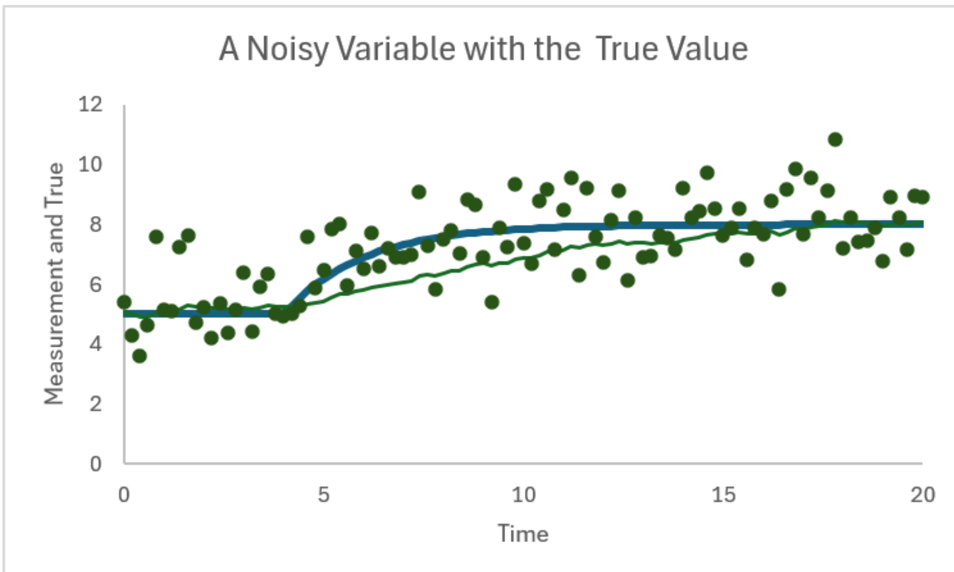

Figure 2 – Figure 1 including the true value of the process variable

Noise can be produced by flow turbulence, which confounds pressure or flow rate measurements. It can be generated by incomplete in-process mixing, confounding composition or temperature measurements as rich or lean fluid packets pass by the sensor. It can be generated by physical vibrations from external equipment or in-process splashing, boiling, cavitation, or bubbling. It can be generated by external electromagnetic fields.

Noise is usually considered to be normally and independently distributed (NID) perturbations that are added to a process variable; and most often, the perturbations are close enough to being Gaussian distributed with a mean of zero and noise amplitude characterized by standard deviation, NID(0, σ).

This ideal characterization of noise has permitted foundational mathematical analysis of many techniques for tempering noise, such as the first-order filter or a moving average. The NID(0, σ) and additive perturbation to a true PV value assumptions have also grounded the development of statistical filters, such as Kalman or SPC-based methods. Legitimized by theoretical analysis of idealized concepts, these methods dominate applications.

Statistical filtering

Contrasting the behavior of a first-order filter, Figure 3 shows a statistically based filter behavior. It holds the prior filtered value until there is adequate evidence to confidently report a change, which is characterized by step-and-hold patterns. Here, filter delay (not lag) is an undesirable feature, but desirably the filtered value does not wiggle during periods of steady state, and it tracks the transient faster.

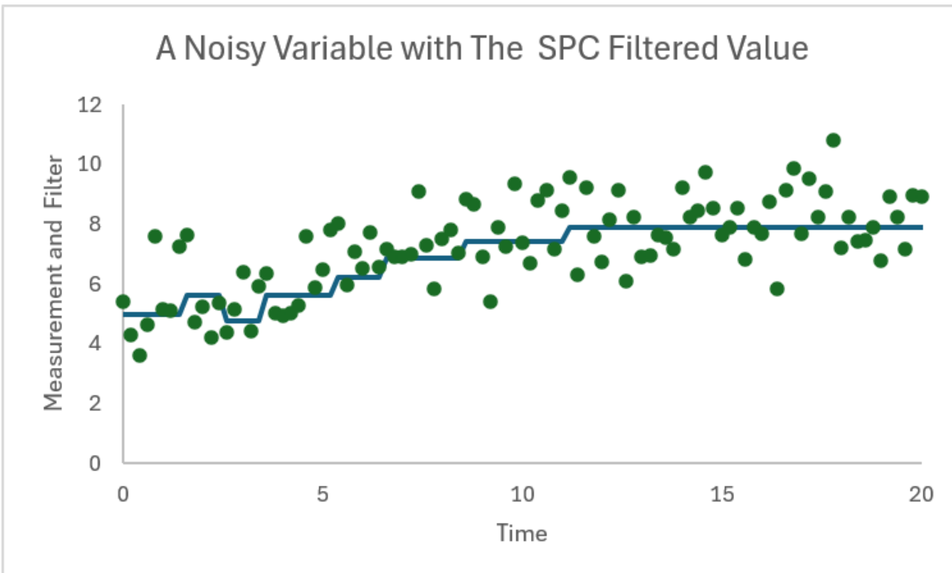

Figure 3 – Characteristics of the statistical filter

Search keywords – first-order filter, SPC filter, Box-Mueller method to generate noise.

Kernel filtering

The first-order and statistical filter methods update the filtered value in real time, as each sampling provides a new noisy PV value. They can only "see" present and past data values. They cannot "see" future values. By contrast, after the data is collected, a kernel filter goes back through the data and determines an "average" at each point in time using a few of both past and future data. An advantage is that it will more closely track transient responses, but the disadvantage for on-line control is that it is always reporting results represented half of a data window behind.

Kernel methods would be good for accuracy and transient analysis in off-line analysis of process behavior.

The "kernel" is the rule for weighting past and future data values. It could simply be an average, or it could be an exponential or radial basis weighting, or a median value. There are many weighting methods, including best fitting a local curve through the data and interpolating values from the curve.

Search keywords – kernel method filtering, median filter

Tempering is not removing

Even so, methods to "remove" noise cannot truly remove it, they temper the perturbations, they reduce the variance, but do not remove uncertainty. Further, on-line filtering methods cause either a lag or delay in the filtered signal catching up to the true signal. If the process just transitioned to a new steady state, filtering takes some time to see the new average (within the residual uncertainty), but during transients or a sequence of real changes, the filtered signal can remain aggravatingly behind, misdirecting action.

Point: Filtering to remove noise, can't. The greater the noise reduction, the greater is the lag or delay. You must choose the filter parameter values to balance precision (at steady state) and speed of response with lag or delay.

Point: If you don't like the noise impact on the controller, only filter on the derivative function (not on the process variable) so that the proportional action can be responsive to data, not the lagged value.

Search keywords – CUSUM filter, SPC filter, Kalman filter

\- - - - - - - -

Russ Rhinehart started his career in the process industry. After 13 years and rising to engineering supervision, he transferred to a 31-year academic career, serving as the ChE Head at Oklahoma State University for 13 years. Russ is a fellow of AIChE and ISA. Now "retired", he enjoys coaching professionals through books, articles, short courses, and postings on his web site www.r3eda.com.

Understanding asymmetric noise and autocorrelation in process measurements | Control Global

PART 2 – Aberrations and Solutions

This second of a three-part article discusses nonideal aberrations in "noise", and solutions to consider.

Nonsymmetric/Non-Gaussian

Noise may not be ideally symmetric (upward perturbations may have a different amplitude from downward perturbations). This might be due to pressure pulses in a flow line (small ones are larger than large ones), or due to a nonlinear feature such as square root extraction of sensor values, or due to a statistical distribution of perturbations that puts a long tail on one side of a distribution (for instance, impurity level is bounded on the low side by zero, but unbounded on the high side). If the perturbations are not symmetric, filtering will shift the reported average toward the larger amplitude side, creating a bias, confounding even steady state analysis.

There are several visual characteristics of asymmetric noise. During a steady state period, the negative deviations will be clearly farther from (or nearer to) the average, and there will be more data on one side of the average than the other. Figure 4 reveals both departures from ideal. During the steady period of time 10 to 20, there are more markers above the filtered value than below, and the below markers are generally farther from the average than the above markers.

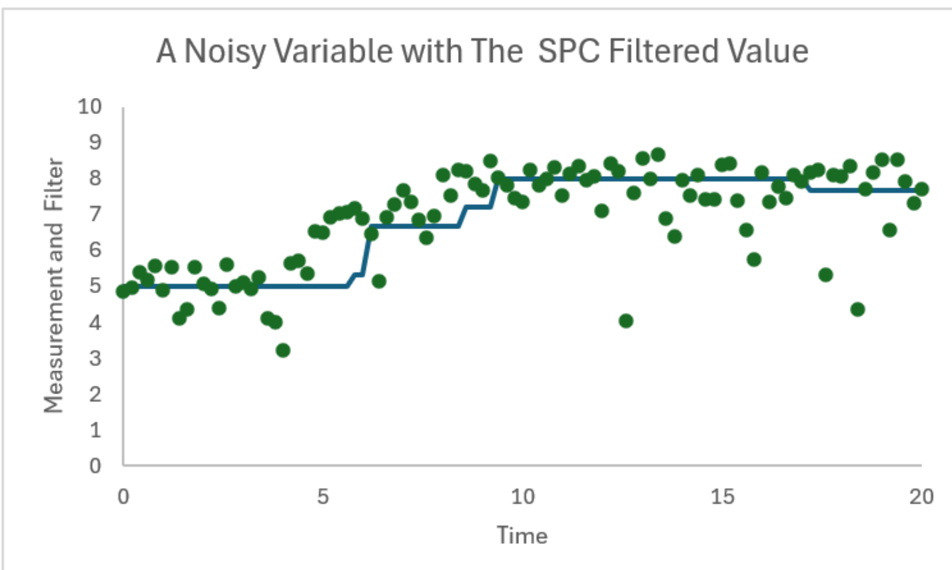

Figure 4 – Illustrating asymmetric noise

The problem is that this creates a bias on the measurement – the average is pulled toward the more extreme data. And the distortion from symmetric data violates the assumptions about normally distributed data that are the basis for statistical analysis (such as t-tests, or variance estimation). But if the noise level is relatively small or the disparity between plus and minus deviations is not excessive, the nonideality can be usually ignored.

Point: If there is copious data at steady state, which does not have symmetric noise, the median (the 50%-ile value) will be a better representation of the true process value than the average.

Search keywords – Median filter

Autocorrelation

Noise is generally considered to be perturbation additions to the true measured value, which are independent at each sampling. Something real causes the perturbation. If that something persists to the next sampling, it will have a similar impact on the next measurement. What this means is that if a prior measurement is on the high side, then the next measurement will likely be high also. Both measurements are contaminated by the same event. These persisting influences cause autocorrelation in the perturbations. Sequential perturbations are not independent.

As an example of a mechanism causing autocorrelation, clouds passing by the sun might be randomly causing changes in solar energy hitting your process, which affect internal temperatures. But the clouds don't blink on and off at the frequency of your data sampling. Both cloud cover and clear periods persist for a brief duration, so the sequence of temperature measurements will progressively rise and fall.

As another example, batch-to-batch raw material feed causes composition changes that persist for a duration.

Finally, at-sensor filtering to temper noise or the sensor lag caused by a thermowell both cause autocorrelation. There are many sources of autocorrelation.

Figure 5 illustrates autocorrelation during a steady period. The dashed connect-the-dots line helps visualize the trend in the data values. If one point is high, then next will likely also be high. If one is low, likely so will the next. The trend is like oscillation, but the above and below periods are not regular in either amplitude or duration. And unlike asymmetric noise, the number and range above the average is about the same as below.

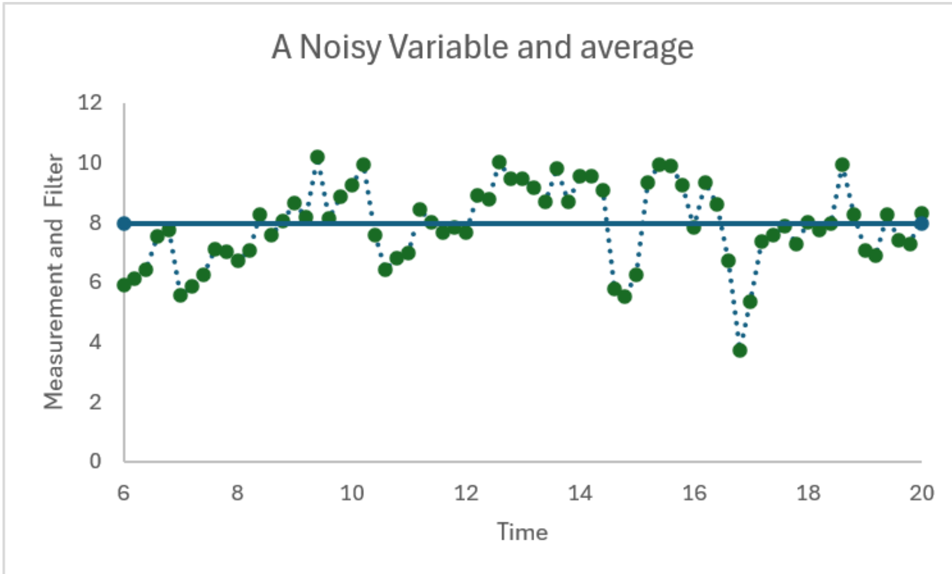

Figure 5 – Illustrating autocorrelated noise at a steady period

Another way to detect autocorrelation is to plot one data value with respect to the prior data value. If there is no autocorrelation, the pattern will be circular. If there is autocorrelation the data pattern will be an oval with an upward or downward trend. The term "Lag-1" means that the data sequence is shifted by one, which is what Figure 6 represents.

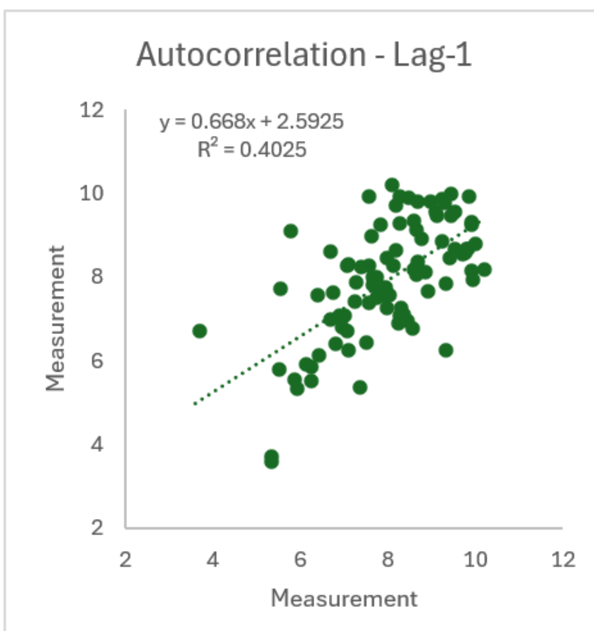

Figure 6 – A Lag-1 autocorrelation plot

Although the correlation trend is visually obvious, one could do a statistical test on the R^2 correlation coefficient to support an autocorrelation claim.

Figure 7 indicates that at Lag-4 (one data point plotted vs. the 4th prior point) no autocorrelation is visible. In this simulation, the impact of influence on the 4th past data point has faded enough so that it does not have a noticeable impact on the most recent measurement.

Note that the axis lengths and measurement range in Figures 6 and 7 are identical. If not, the circular pattern in Figure 7 will appear elongated.

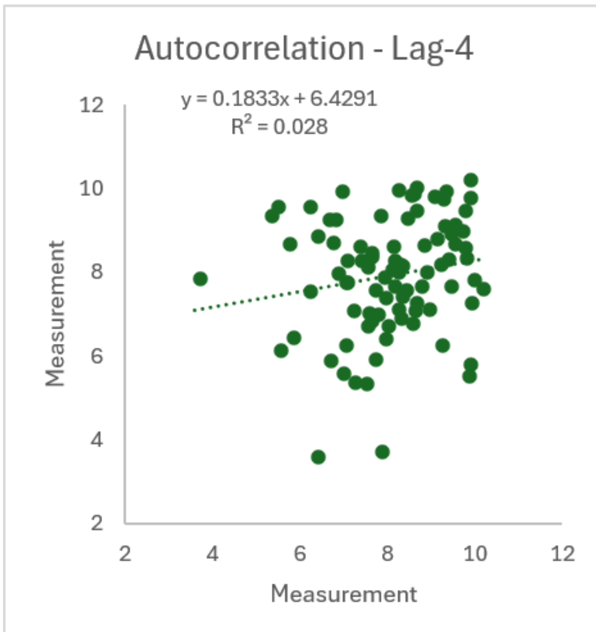

Figure 7 – A Lag-4 autocorrelation plot

Autocorrelation might be the result of pulsing in piston pumps, on-off level or temperature control, unit switching for regeneration, changing properties of raw material feed, BTU content of a shovel-full of coal, or process device issues such as stick-slip (sticktion) in a valve. It may appear to be random noise on a less frequent sampling interval (as Figure 7 reveals), but normal measurement frequency can track the persisting ups and downs making it look like the real trend, which it represents.

Similarly, external environmental influences on the process can cause perturbations that have persistence. Consider wind gusts, rain showers, or passing clouds that create intermittent cooling of external units, which persist for a brief time. Again, the persistence of these on-off disturbances will cause temporary measurement changes with autocorrelation.

Filtering also causes autocorrelation. Filtering is a form of averaging, and each measurement persists in the averaging window, until time passes it out. Or, in the case of a first-order filter, in which the influence is exponentially weighted by time, the influence of all past data remain in the calculation, but after enough samplings the exponential weighting make the past perturbations inconsequential.

The controller can cause autocorrelation. If a step disturbance pushes the process off of setpoint, the controller begins changing the manipulated variable, until the integral action accumulates enough to move the controlled variable to the set point. The deviation persists until fixed.

Point: Instead of using the instrument and control system to cover-up (filter) or to control the persisting perturbation, eliminate the source. For instance, improve feed material blending, use pulse dampers, or put positioners on valves. Alternately, reduce the sampling frequency so that the short-term persistence has faded prior to the next sampling. However, reducing sampling frequency may cause a delay in sensing process changes.

Point: Many statistical tests are based on the no-autocorrelation assumption. These include t-tests on differences, steady state detection, and Statistical Process Control charts. Although you might want a controller to work on a high frequency, you also might want to use less frequent sampling in process analysis to remove autocorrelation from the data.

Point: Just because the manufacturer designed a controller to operate at a 10 Hz frequency, doesn't mean that is better. An old rule of thumb is that there should be 30 actions in a settling time. If your process has a 20-sec time-constant, the 95% settling time is roughly 60 sec, and a control frequency of 0.5 Hz should be adequate. If there is perturbation autocorrelation at the high sampling frequency, it might disappear at the lower, but still fully adequate, rate.

Point: The high value of one sampling does not cause the high value of the next. There is a third variable that persists for a while, which affects sequential measurements similarly.

Point: Filtering at the sensor/transmitter may have been chosen by someone installing or calibrating it to generate a smooth signal. You may want to reconsider the filter tuning.

Search keywords – Autocorrelation, Correlation Coefficient, EWMA filter (exponentially weighted moving average)

Outliers are not noise

Relative to a control system, errors that should be ignored are missed/dropped signals or other spurious signals such as those due to electromagnetic pulses from motor startups. Here, median filters are a sensible solution. In a middle-of-three filter, the middle value of the last three signals is accepted as the representative value. If one signal is absurdly high or low, one of the other two will have the middle value. The mid value might be the most recent, the second past, or the third past. On average this will cause a delay of one sampling as the reported value. If you think a spurious or outlier event might have a persistence of three sequential signals, use the middle of the past 5 measurements.

Point: Methods to remove outliers (infrequent spurious events that should be rejected) are different from those intended to reduce noise by filtering data. Using a first-order filter strong enough to render the outlier inconsequential probably will cause an undesirably significant lag.

Search keywords – Median filter

Resolution

Signal discretization can give an appearance of noise. Discretization can come from many sources such as digital transmission devices or analog-to-digital display. This is like the digital time display when seconds or minutes jump from the past value to the present, then hold that value during a time interval until it jumps to the next. However, for a process signal that is progressively drifting about a value, signal discretization generates either 1) a flatline hold until enough change makes it jump, which is a false indication of a noiseless steady state, or 2) a square wave signal that seems to randomly jump between certain discrete values (with no in between values).

The visible clue to signal discretization is that the signal jumps between certain values. See Figure 8. Note the lines of dots with identical values, and no dots with in between values. If a process is nonlinear, the discretization may have very small intervals (not noticeable) in one range, but large enough to be visible in another range.

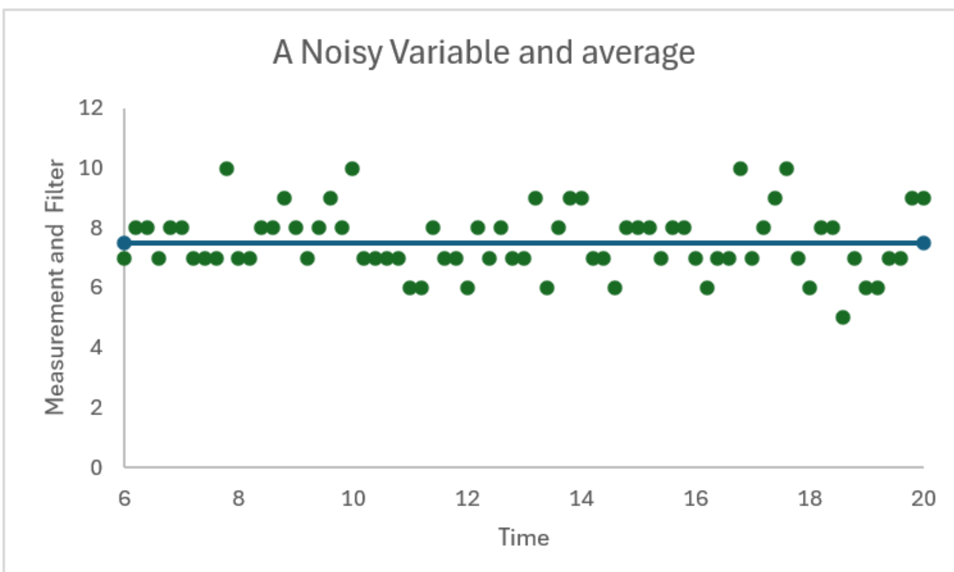

Figure 8 – Characterization of a resolution problem

I would not label this discretization signal as noise; however, filtering could minimize its amplitude (but also introduces a lag).

Point: To eliminate the appearance of noise due to signal discretization, use higher bit converters.

Point: All digital devices have discretization. If an instrument is calibrated for a very wide PV range, the discretization interval on the reported measurement will be large, and possibly visible. Consider calibrating over a smaller PV range.

- - - - - - - -

Russ Rhinehart started his career in the process industry. After 13 years and rising to engineering supervision, he transferred to a 31-year academic career, serving as the ChE Head at Oklahoma State University for 13 years. Russ is a fellow of AIChE and ISA. Now "retired", he enjoys coaching professionals through books, articles, short courses, and postings on his web site www.r3eda.com.

This is part 3 of the 3-part article on coping with noise in process measurements.

PART 3 - Solutions

Noise or Reality?

Are these sorts of influences noise or real process changes? Something caused the measurement to have perturbations. The perturbation is caused by something real. But that something might not mean the process changed. For example, the process flow rate may be steady (unchanged), but turbulence or vibrations made the measured flow rate value change. Or, in an inline blending, the mixing might be incomplete, and a composition or temperature measurement may vary as rich or poor packets pass by the sensor.

However, the noise might represent small and temporary process changes from an uncontrolled influence.

Is the perturbation "noise"? From a data-based view, it depends on the sampling frequency of the observation. "Noise" is uncorrelated. Sampling faster, while the influence of the "something real" persists, will make the perturbations correlated. Making it appear that the process is drifting.

From a process control view, it also depends on the speed with which the control system can correct the deviation. The deviation might reveal a real process change, but because of transport delay or analyzer delay, or batch processing, the material may have passed through the process and controller will not be able to correct it.

If the effect passes before the controller can correct it, or if the perturbation is noise rather than a real persisting deviation, then a controller would simply be tampering – responding to noise that it cannot fix. The cause of the perturbation will be gone before the controller "correction" is implemented, and the controller will just be amplifying the "noise" by adding another source of deviations.

Point: Don't try to control noise. The speed of the control system to correct a PV upset, must be much faster than the duration of the upset.

Search keywords – Tampering, W. Edwards Deming's funnel, Statistical Process Control

Ensure instruments are calibrated

Certainly, an improperly calibrated sensor/transmitter can introduce a systematic error. For instance, time, heat, degradation of internals, or erosion may have made a once-calibrated device to now be out of calibration. Additionally, if the sensor to transmitter response is nonlinear (for example temperature to thermocouple mV, yet it was calibrated assuming to be linear with a 2-point calibration, then between the two calibrated points, the sensor will report a systematic error. Probably, nearly all sensors are nonlinear, but within a reasonable range, the linear assumption is acceptable for most.

Point: Be sure that instruments are still in calibration and that calibration procedures match the sensor-to-measurement reality.

Match the solution to the character of the "noise"

Noise solutions include first-order filtering, statistical filtering, median filtering, improved resolution, nonlinear compensation, less frequent sampling,

Point: The solution must match the "noise" character.

Detecting truth through the noise

As part of a prior discussion about data, Ed Farmer suggests an even broader context, "Sometimes, though, it becomes essential to differentiate the underlying reality from the corrupted measurement data."

Filtering tempers noise, but adds a lag or delay, and does not correct measurement bias.

To detect reality: Seek to minimize sources of the various types of perturbations or misrepresentations in the measurements. Keep instruments calibrated. Perhaps consider data reconciliation to correct measurements that have a systematic bias.

Measurement bias might change on a daily or weekly basis due to sensor device damage, progressive failure, or changes in operating conditions. Data reconciliation adjusts measured values to make material and energy balances about process sections close. It uses variance on each measurement source, to determine uncertainty.

Search keyword – data reconciliation

In Human Communication

Filtering may be desirable to temper control action response to noise. But observing the filtered variable masks features that an operator may want to see, such as when a transient starts or stops, whether there are spurious signals, etc.

Point: It may be beneficial to let operators see the noisy signal on their display, rather than the filtered signal which is used for control or on-line analysis.

Noise is not just a part of process control systems. Erroneous messages and "noise" also refer to distraction from or degradation of a message, background nonsense or confusion; such as that which is promulgated by social media, distortionists, or propagandists. Humans lie to cover-up or distract attention from things they don't want known. Humans preserve traditions as reality, even when the tradition is technical folklore, perhaps a legacy misunderstanding of process mechanisms or a a feature of the process that no longer exists. Humans claim certainty whether they are certain or not, because uncertainty does not engender followership action or their future leadership-based promotion.

Point: Use your understanding of filters to reject human-generated nonsense. Perhaps median filters to reject extremes, fact checking, mechanistic cause-and-effect reasoning, or data reconciliation.

Conclusion

Understand the various characterizations and solutions about "noise" and filters.

\- - - - - - - -

Russ Rhinehart started his career in the process industry. After 13 years and rising to engineering supervision, he transferred to a 31-year academic career, serving as the ChE Head at Oklahoma State University for 13 years. Russ is a fellow of AIChE and ISA. Now "retired", he enjoys coaching professionals through books, articles, short courses, and postings on his web site www.r3eda.com.