OCTAVE

WHITE PAPER

Making a big dent in nuisance alarms

Contents

Introduction

Nuisance alarms are common in alarm systems. High quantities of these alarms can significantly reduce the effectiveness of the alarm system and render it useless for providing timely notifications of process situations that demand operator attention. Nuisance alarms can result in incidents, lost production, environmental excursions or even serious accidents. Solving the nuisance alarm problem is an important and easy step in alarm system improvement.

The seven-step approach from Octave to alarm system improvement is a well-known and proven methodology based on hundreds of successful projects and many terabytes of real-world data.

Step 1: Develop, adopt and maintain an alarm philosophy

Step 2: Collect data and benchmark your systems

Step 3: Perform bad actor alarm resolution

Step 4: Perform alarm documentation and rationalization (D&R)

Step 5: Implement alarm audit and enforcement technology

Step 6: Implement real-time alarm management

Step 7: Control and maintain your improved system

The first three steps are universally needed for the improvement of an alarm system. Often, these steps are simultaneously implemented at the start of a project and they collectively provide the most improvement for the least expenditure of effort. They provide the best possible start and the fundamental underpinnings for the remainder of steps necessary for effective alarm management. This white paper covers step 3, in which nuisance alarms (bad actors) are identified and addressed.

Several categories of nuisance alarms exist with various methods for dealing with them. With enough bad actors, an alarm system is rendered virtually useless. This may lead to hazardous plant conditions, since important or critical alarms are lost in the “sea” of bad actor alarms. These conditions can result in adverse economic, environmental or safety consequences.

Expected results from nuisance alarm reduction

The top 20 most frequent alarms comprise 25% to 95% of the entire system load. Operators can achieve major system improvement with comparatively little effort if those alarms are dealt with successfully. It is quite amazing that such high numbers of nuisance alarms exist because it is doubtful the best control engineer in a company could intentionally design alarms to behave in the ways we will discuss. Yet, they do exist. All varieties are in almost every system we analyze.

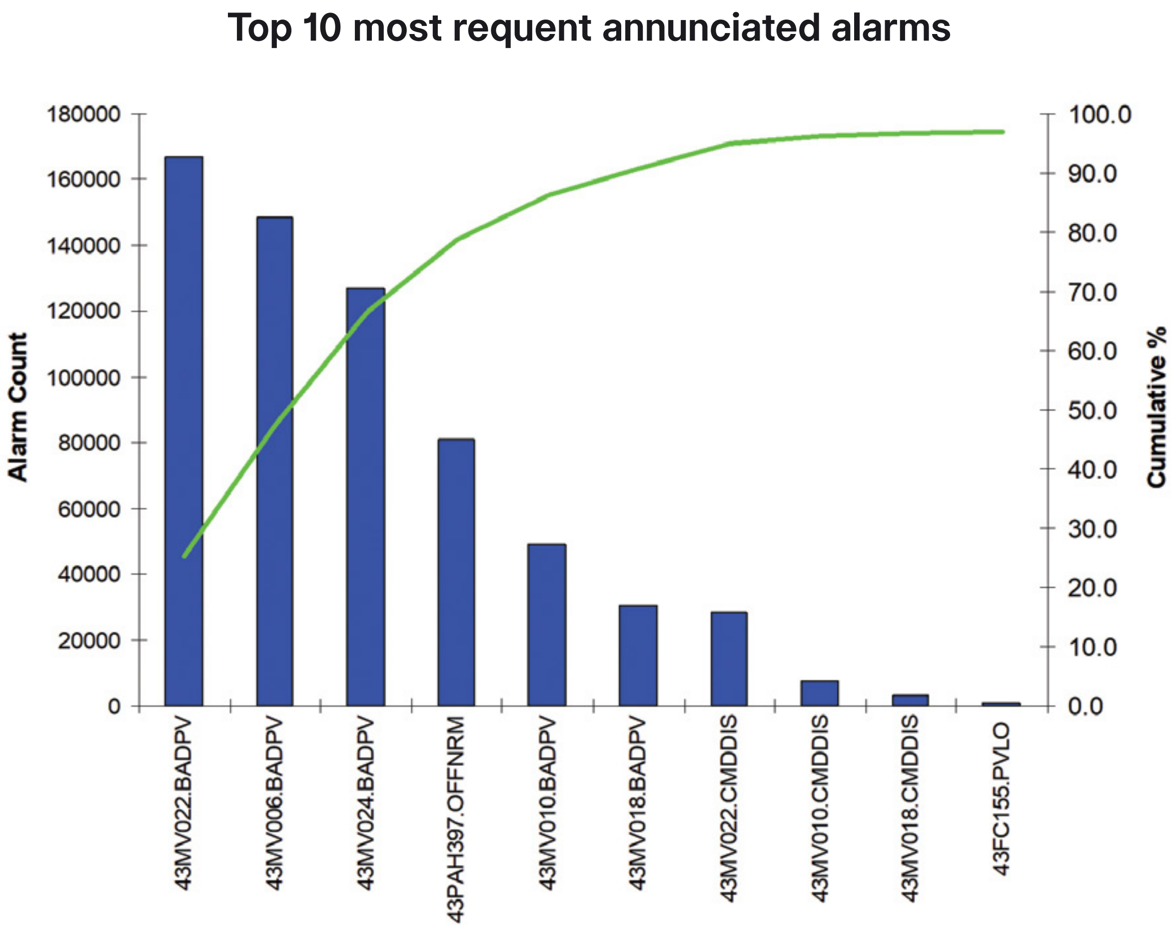
Top 10 most requent annunciated alarms

Figure 1. Top 10 most frequent alarms on a single system

In Figure 1, only 10 alarms comprise 96% of the total alarm load. The chart is based on eight weeks of data and several of the alarms went off more than 100,000 times. This performance was never intentional, and fixing only these 10 alarms would reduce system load by 96%.

Here are some examples from 15 different control systems:

Bad actor alarm work process results	Baseline alarms	Reduction from bad actor recommendations	Reductions %
System 1	339,521	325,423	95.8%
System 2	644,487	593,904	92.2%
System 3	79,434	72,935	91.8%
System 4	58,049	51,782	89.2%
System 5	482,375	413,094	85.6%
System 6	414,887	333,395	80.4%
System 7	93,848	71,372	76.1%
System 8	64,695	46,749	72.3%
System 9	33,115	22,646	68.4%
System 10	225,668	133,307	59.1%
System 11	44,527	24,882	55.9%
System 12	183,312	77,417	42.2%
System 13	106,212	38,566	36.3%
System 14	91,686	29,188	31.8%
System 15	39,305	8,625	21.9%

Figure 2. Improvement amounts from alarm bad actor resolution

In the above systems, fewer than 50 alarms were analyzed each using the techniques in this paper. The average reduction achieved was 66% (based on systems) or 77% (based on total alarms). This is a substantial gain for a small amount of work.

By doing this task near the start of an improvement effort, you immediately achieve a significant improvement and establish the effort's credibility. You are fixing things operators have probably known (or complained) about for years and may have just given up on ever getting fixed.

Wouldn't you be pleased if you analyzed about 30 alarms and cut your alarm rate by more than half?

Chattering and fleeting alarms

Imagine an alarm that cycles between annunciating and clearing three or more times per minute. This is the initial definition we will use for a chattering alarm. This alarm condition is not cleared because an operator detects it, analyzes the situation and makes a change in the process, which changes the process condition and results in the alarm clearing. Such chattering alarms are quite common and are a big nuisance and distraction to the operator. They are relatively easy to fix.

Both analog values and digital (on-off) signals, such as from switches, can and do chatter with digital values typically being the worst case.

There is a sub-category of alarms similar to chattering called fleeting. Fleeting alarms come in and clear very quickly (too quickly for the operator to have been responsible) but do not necessarily repeat. With minor differences, fleeting alarms can be addressed using the same methods for addressing chattering alarms. For chattering analog sensors, operators should first consider the deadband of the alarm. A brief review of the deadband concept follows.

Alarm deadband

Deadband and on-off control

On-off control is the most basic form of control. A certain degree of deadband is placed around the process setpoint. If the process variable is lower or higher than the deadband, the control action is turned fully on or fully off.

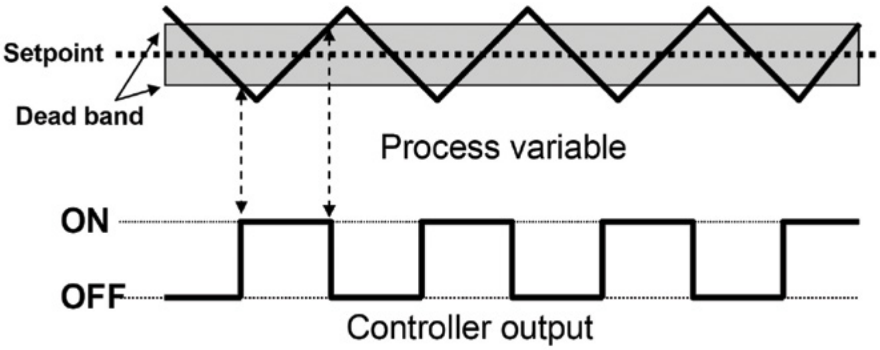

Figure 3. Deadband and on-off control

On-off control is commonly used for regulating the temperature of non-critical processes, such as central heating and air conditioning, lubrication oil temperature and cold-room temperature. It can be used on a pump to fill (or empty) a tank (on at 20%, off at 80%, similar to a toilet tank).

With on-off control, the process variable is always cycling through the deadband. To achieve tighter control, reduce the amount of the deadband. The resulting side effect is the frequency of the oscillations increases, which reduces the life of the final control element such as relays or control values.

Deadband and alarms

Similar to deadband for setpoints and process control, alarms on analog values should also have a deadband specified. As a process value passes through an alarm setpoint, any noise or slight variation of the signal causes multiple alarms if there is too small of an alarm deadband. All process signals have noise.

Figure 4 shows how a proper deadband, larger than the noise in the signal, reduces alarm events as the process value moves above a high alarm setpoint. The alarm deadband should be larger than any expected signal noise. Most distributed control systems (DCSs) allow for deadband but may specify it in measurement units, percent of range or in other ways. The deadband's plus or minus positioning relative to the setpoint may also vary based on the alarm type. There may be individual deadband settings for each different alarm or one applying to all the analog alarms on the point. You should check the DCS documentation to configure the deadband properly. Note that deadband should always be applied to analog signals before the application of the delay time technique.

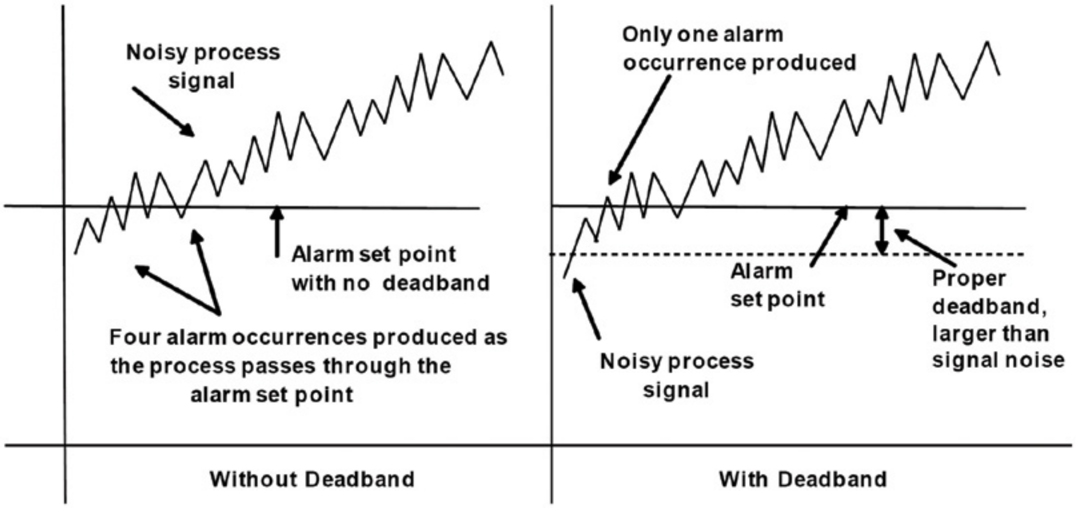

Figure 4. Deadband and alarms

Deadband should be configured on every analog alarm. Rigorous calculation is not usually necessary; the following good starting values can be used. Trial and error can also be used. Pick a small starting point (1% to 2%) and increment based on the results.

Signal type	Deadband %
Flow	5%
Level	5%
Pressure	2%
Temperature	1%

Figure 5. Deadband settings based on sensor type

Process value filtering and alarms

It is possible to filter process variable signals in a DCS, usually in various ways. The primary reason to use these filters has to do with control loop performance, not with alarms. All process signals have some noise. Noisy signals interfere with good control loop performance. An unfiltered noisy input signal to a PID (proportional-integral-derivative) controller produces a noisy output signal. This provides poor control and excessive valve movement and wear.

Generally, control systems have a variety of filter algorithms to improve noisy process variable signals. However, such filtering may have a similar effect on alarm activation as deadband (depending on if the alarm processing is performed on the signal before or after the filter is applied). An optimum filter setting smooths out signal noise but has little effect on the desired response of the system. If a filter is too large, it may obscure process problems from the operator.

In addition, process value filters introduce additional time lag into control loops, which is seen as additional apparent dead time in the loop. This may significantly affect the loop's settling time. A PID controller must be retuned after adding or modifying a filter on the process variable.

Typically, a full chapter in a control engineering textbook discusses how to determine proper filter settings for control. It is more important for process value filters to act correctly for control than for their resulting alarm effect. Therefore, we do not advocate signal filtering as a good way to address chattering analog alarms. They are mentioned here only because they can have an alarm-related effect. If you suspect a control problem related to a noisy signal, the following values are good filter starting points.

Signal type	Filter time constant
Flow	2 seconds
Level	2 seconds
Pressure	1 second
Temperature	none

Figure 6. Filter time constants based on sensor type

Doesn’t everyone already know this?

Deadband and process value filtering are control system engineering 101» concepts. Some readers may think this is so obvious as to be a waste of paper. Well, place yourself as a newly hired consultant to a multibillion-dollar petrochemical project six months from startup. The project manager asks you to “take a look at the DCS configuration and tell me what you think.”

The DCS configuration engineering is done by a team of engineers supplied by the DCS manufacturer’s company. Their individual charge rate is over $200 per hour, and they have been working for months. You talk to the team leader, a person with many years of experience.

You: “So, how are you defaulting your analog alarm deadbands? What values are you using? Are they the same for all sensor types?”

Team Leader: “They are all defaulted to zero.”

You: “Hmmm. So, how are you defaulting the process value filtering to your controllers?”

Team Leader: “Ah, uhh, all the same, zero.” As you walk off, you see him grab the phone …

Later, you report this to the project manager. You explain that these settings guarantee chattering alarms and jittery process control during startup. Long ago, the project manager had taken control systems engineering 101, so he understands.

You also point out that the initial download to the DCS has been made and changes to individual points will be needed to set things right. This will take some time. You advise the project manager not to accept a change order for additional funds from the configuration team because they should have known better. “You are paying top dollar for talent that should produce a configuration that reflects top expertise.” The project manager emphatically agrees.

Moral: Just because we know something does not mean it will be implemented in practice. We challenge every reader to examine your control system configuration. You will be surprised at the number of analog points with alarm deadbands set to zero. We’ve seen it often.

Delay time analysis and alarms

Deadband and process value filtering apply only to analog values. Often, worst-case chattering or fleeting alarms are associated with on-off signals such as pressure and level switches. These devices may have a screwdriver-type, trial-and-error mechanical deadbandadjustment on them, but can you find the manual? There is another powerful method to use that is probably already a capability of your DCS. This method applies to both analog and digital point types. The method and technique require a bit of explaining, but once explained, the technique itself is simple. The results you will get are so powerful it is well worth the effort.

There are two types of alarm delays available in many DCSs, namely the on-delay and the off-delay. The off-delay is sometimes referred to as a “debounce timer.” Some point or alarm types may have either delay available and some only one of them. Again, there is a need to read the DCS documentation. On and off delays work differently and have significantly different implications when used. These settings provide powerful methods for dealing with chattering and fleeting alarms.

A fleeting alarm transitions between the alarm state and the normal state quickly - usually within a few seconds to a minute. They do not immediately repeat. If they repeat, they are called a chattering alarm. In both cases, the alarm durations are too short for the alarm clearing to have been because of operator action in response.

Alarm analysis software is needed to improve an alarm system, but it can also detect and solve problems. To do this, we want to take one of our nuisance-chattering or fleeting alarms and perform two frequency analyses. These are analyses of the times-in-alarm (durations) and times-between-alarms (intervals).

Time-in-alarm and time-between-alarms

DCSs produce time-stamped event records of at least three things: the alarm occurrence itself, the return-to-normal event (created when the condition causing the alarm to occur has cleared) and the operator acknowledgment event (created when the operator hits the acknowledge key for the alarm). Only the first two are of interest to us in this discussion.

In your alarm analysis software, you will have recorded thousands of occurrences from your nuisance alarms. For each specific nuisance alarm, take each pair of alarm occurrences and return-to-normal events and then subtract the timestamps. The result is the time-in-alarm (duration) for the alarm occurrence. In a similar method, subtracting an alarm occurrence timestamp from the prior alarm clear event timestamp produces the time-between-alarms (interval). This is shown in Figure 7.

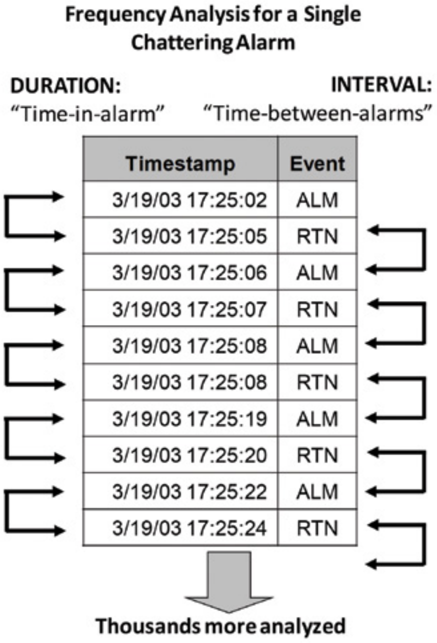

Frequency Analysis for a Single Chattering Alarm

Timestamp	Event
3/19/03 17:25:02	ALM
3/19/03 17:25:05	RTN
3/19/03 17:25:06	ALM
3/19/03 17:25:07	RTN
3/19/03 17:25:08	ALM
3/19/03 17:25:08	RTN
3/19/03 17:25:19	ALM
3/19/03 17:25:20	RTN
3/19/03 17:25:22	ALM
3/19/03 17:25:24	RTN

Figure 7. Chattering and fleeting alarm durations and intervals

Plotting the results for thousands of events from a single alarm may show a graph similar to Figure 8.

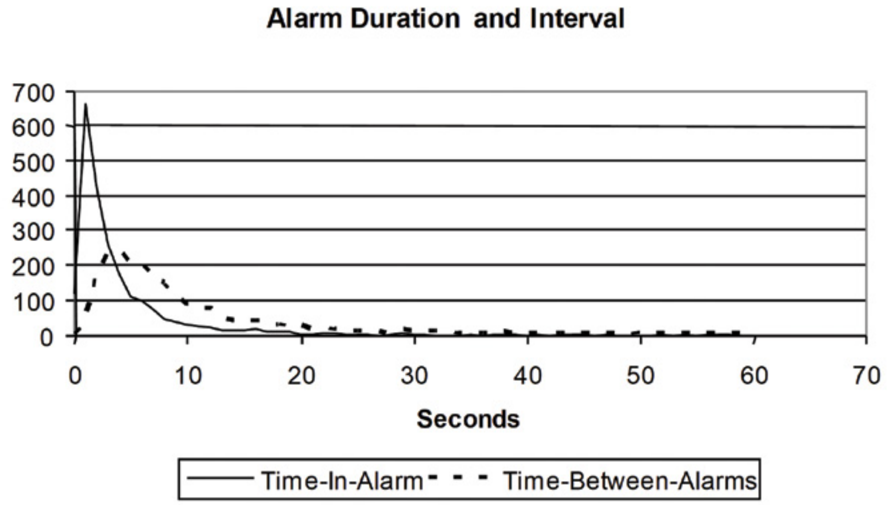

Figure 8. Alarm delay time analysis graph

The two graphs' curves are determined as follows:

- Time-in-alarm (duration) graph: Y = count of alarm occurrences having the duration of X seconds.
- Time-between-alarms (interval) graph: Y = count of alarm occurrences having the interval of X seconds.

In the case shown based on thousands of alarm-return pairs, most of the alarms from this point have durations (solid line) less than 10 seconds and the time-between-alarms (dotted line) is mostly less than 20 seconds. An alarm that comes in, lasts less than 10 seconds and then goes away all by itself does not meet the basic criteria for an alarm – something requiring operator action to resolve.

When you plot these durations, the area under the curve totals 100% of the alarm occurrences from the single alarm being analyzed. Figure 9 further illustrates this principle.

In Figure 9, the alarm in question had thousands of activations lasting 10 seconds or less. In fact, 93% of all activations of the alarm lasted 15 seconds or less. Those alarms did not return to normal because of responsive operator action. They indicated some transient conditions that did not require operator action to resolve. However, some of the alarms did remain valid for several minutes.

This is compelling information to use when coupled with the on-delay and off-delay abilities of the DCS. Here is exactly how those abilities work.

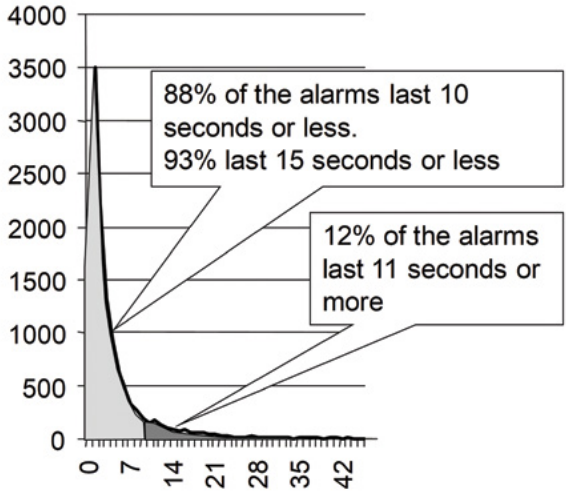

Figure 9. On-delay (duration) histogram percentage determination

On-delay

Using on-delay time parameter can prevent a transient alarm from ever being seen by the operator. The alarm does not immediately annunciate; it must remain in effect and not clear for the on-delay timer's value before it is announced to the operator.

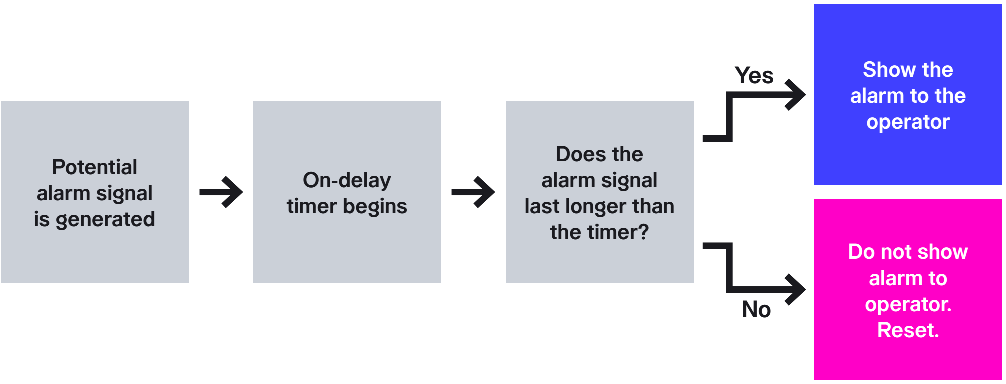

Figure 10. On-delay alarm processing

The correct choice of the on-delay time parameter is important since, if used, even a valid alarm is not immediately presented to the operator. This will increase the overall time it takes for a proper response to be made. Such a delay could be a safety concern on some alarms. On-delays of 30 seconds or less are generally not problematic for priority 3 alarms. On-delays of more than 30 seconds or a minute should be applied carefully, even for priority 3 alarms. On-delays of more than a few seconds are a concern for priority 2 or priority 1 alarms.

Off-delay

This powerful method can turn a string of repetitive, nuisance and chattering alarms into a single, longer-duration alarm event with no initial delay. The alarm is immediately annunciated. When it clears, that clearing is not shown to the operator unless it remains clear for longer than the off-delay timer. The operator perceives an alarm that attempts to re-occur as a single sustained alarm rather than a recurring, short-duration one.

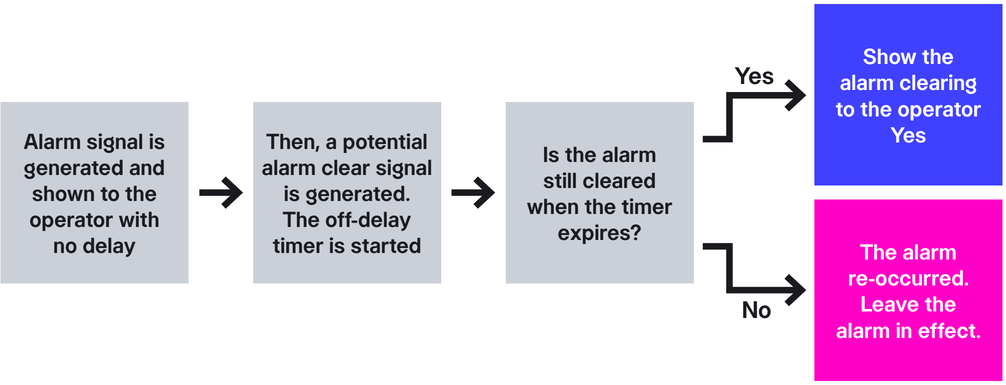

Figure 11. Off-delay alarm processing

Using this technique, hundreds or thousands of nuisance alarm occurrences can become a single, longer-duration alarm occurrence with no initial annunciation delay. The key is the correct choice of the delay time parameter to be greater than the normal time-between-alarms.

The disadvantage to this technique also concerns the delay time. Suppose the operator gets the alarm and takes corrective action to eliminate it. In that case, he will not see a return-to-normal condition until after the delay time has expired, regardless of whether the action was immediately successful. In most cases, this is quite acceptable for off-delays of up to a few minutes. The operator can see (for analogs) that the process value has moved below the alarm setpoint. Another concern for alarms that chatter for long periods is that the resulting "off-delayed" alarm may become stale. Well, this is still a far preferable condition to chattering. The source of the problem is probably that the alarm setpoint is too close to where the process likes to live and isn't indicating a transition to an abnormal range.

Both on-delay and off-delay apply to analog and discrete inputs such as switches. Deadbands may be applicable only for analog signals unless the discrete sensor has some sort of mechanical deadband adjustment.

Delay in seconds	**% Reduction**	**% Reduction**
	Time-in-alarm (on-delay)	Time-between-alarms (off-delay)
5	77.7	19.7
10	87.6	37.8
15	93.0	48.7
20	95.4	58.4
25	96.1	62.4
30	**96.5**	64.1
35	97.6	66.5
40	97.8	68.7
45	97.9	69.6
50	98.2	70.6
55	98.5	71.6
60	98.5	**72.2**
65	98.6	72.4
70	98.7	73.2
75	98.7	73.6
80	98.7	74.1
85	98.7	74.6
90	98.9	75.1
95	99.0	75.7
100	99.0	75.8
105	99.0	76.0
110	99.2	76.4
115	99.2	76.9
120	99.2	77.2

Figure 12. Delay time alarm reduction table

For each alarm, generating a table similar to Figure 12 is straightforward. This numerical analysis yields the exact percentage of how many alarms would be eliminated based on the choice and type of delay time. The table and charts let you locate the diminishing returns and pick your delay correctly.

For this alarm, an on-delay of 30 seconds would eliminate over 96% of the events. An off-delay of one minute would eliminate 72% of the events. It is typical for an off-delay to be less powerful than an on-delay; for the same specified time, an off-delay will generally eliminate fewer alarm occurrences. Depending on the control system, there may be restrictions around the choices of delay types.

These calculations do not determine "why" the chattering alarm behavior is occurring. The process conditions and the sensing hardware that result in chattering and fleeting behavior and a root cause investigation might find installation or hardware problems. The implementation of delay times is more of a highly effective band-aid solution.

Octave Tempo Alarm Management

Octave Tempo Control System Effectiveness (formerly PAS PlantState Integrity) automates this analysis with the alarm mechanic feature. Any alarm in a typical analysis list (such as most frequent, chattering or fleeting) can be selected and a new analysis window opens. Thousands of that alarm's past occurrences are automatically analyzed, and the effect of both on and off-delay times are shown with a graph and a table. The graph shows the point of diminishing returns, and the table provides the exact amount of alarm reduction for any specific delay time choice.

The alarm mechanic analysis can be saved as a PDF for documentation of management of change.

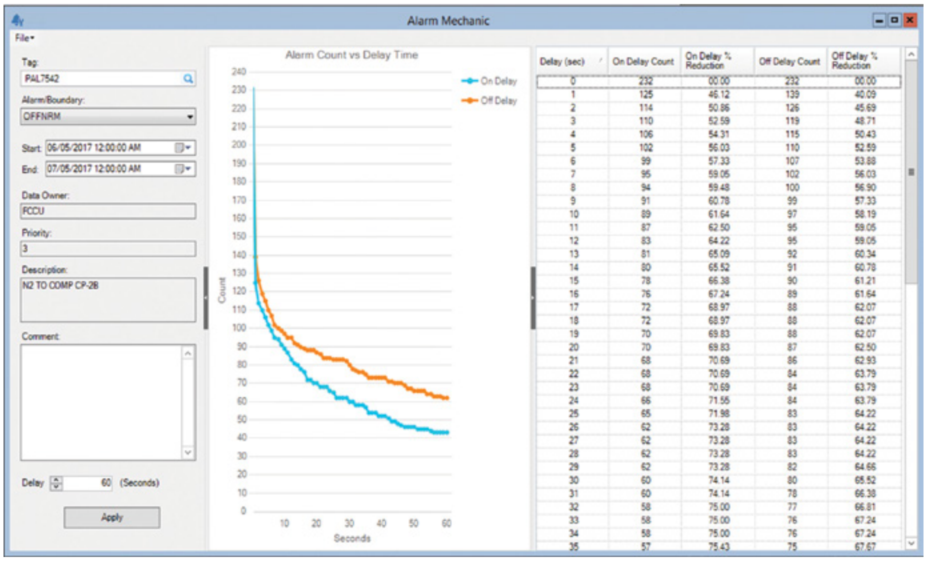

Tag: PAL7542
Alarm/Boundary: OFFNRM
Start: 06/05/2017 12:00:00 AM
End: 07/05/2017 12:00:00 AM
Data Owner: FCCU
Priority: 3
Description: N2 TO COMP CP-2B
Comment:
Delay 60 (Seconds)

Delay (sec)	On Delay Count	On Delay % Reduction	Off Delay Count	Off Delay % Reduction
0	232	00.00	232	00.00
1	125	46.12	139	40.09
2	114	50.86	126	45.69
3	110	52.59	119	48.71
4	106	54.31	115	50.43
5	102	56.03	110	52.59
6	99	57.33	107	53.88
7	95	59.05	102	56.03
8	94	59.48	100	56.90
9	91	60.78	99	57.33
10	89	61.64	97	58.19
11	87	62.50	95	59.05
12	83	64.22	95	59.05
13	81	65.09	92	60.34
14	80	65.52	91	60.78
15	78	66.38	90	61.21
16	76	67.24	89	61.64
17	72	68.97	88	62.07
18	72	68.97	88	62.07
19	70	69.83	88	62.07
20	70	69.83	87	62.50
21	68	70.69	86	62.93
22	68	70.69	84	63.79
23	68	70.69	84	63.79
24	66	71.56	84	63.79
25	65	71.98	83	64.22
26	62	73.28	83	64.22
27	62	73.28	83	64.22
28	62	73.28	83	64.22
29	62	73.28	82	64.66
30	60	74.14	80	65.52
31	60	74.14	78	66.38
32	58	75.00	77	66.81
33	58	75.00	76	67.24
34	58	75.00	76	67.24
35	57	75.43	75	67.67

Figure 13. Alarm Mechanic analysis

This method uses actual occurrence data to determine the proper delay-time value. When implementing new points, you initially have no such data to use. What should be the defaults? The answer requires some explanation.

Implementation of either on-delay or off-delay is different than the implementation of deadband. In specifying deadband, the physics of the situation generally indicates the use of a zero default. A zero on or off delay may be perfectly acceptable for many points. Both EEMUA 191 and ISA-18.2 document some basic warnings about the use of delay time. Here is some more thorough guidance.

Signal type	On-delay time: Default is zero. Use on-delay only on identified problems and priority 3 alarms. Use on priority 1 or 2 should be individually evaluated.	Filter time constant Default and use as shown for priority 3 alarms. Use on priority 1 or 2 should be individually evaluated for acceptability.
Flow	0-15 seconds	15 second default
Level	Use >30 seconds with care	30-60 second default. Use >30 with care, considering tank volume and throughput rates.
Pressure	Use >15 seconds with care	15 seconds default. Upper limit of 60-120 seconds is not usually a concern.
Temperature	Use >30 seconds with care	30-60 second default. Upper limit of 60-120 seconds is not usually a concern.

Figure 14. Recommended delay times based on signal type

Both methods are fixes or workarounds that address the behavior of the alarm without determining the root cause as to why the signal is chattering or fleeting. Each type of input hardware (such as a switch) may have different causes and compensation mechanisms. This technique immediately addresses the chattering behavior without suppressing the view of the alarm entirely from the operator (as alarm suppression would do).

Once you fix all the chattering alarms identified by the three alarms in one-minute criteria, then widen the criteria to three alarms in two minutes and you can find and fix more chattering alarms. While appropriate methods can dramatically improve alarm performance, the underlying process and mechanical causes should also be investigated. This involves a review of sensors and installations. That is, if you have the time, money or people available.

Suppressed alarms

An initial analysis of a system to determine the bad actor resolution list must also identify any configured alarms that are suppressed. Alarm suppression is often uncontrolled. At the end of the bad actor resolution step, no suppressed alarms should remain.

Stale (long-standing) alarms

Stale alarms remain in alarm for extended periods. A value of more than 24 hours is a good starting value to identify them. They distract the operator by filling up the alarm summary screens. We have seen alarms that have been in effect continuously for years. They often reflect stable unit conditions, such as equipment that is intentionally shut down or sensor malfunction, and generally indicate alarms that were not configured in accordance with the principles contained in *The Alarm Management Handbook*.

Stale alarms can be addressed only when you understand the process states and hardware involved. They are usually eliminated by reconfiguring them, so they truly reflect only abnormal, unexpected conditions that require operator action to resolve. This may require imagination or implementation of logic or state-based alarm methodologies.

Duplicate alarms

There are a few types of duplicate alarms.

Dynamic duplicate alarms

Dynamic duplicate alarms consistently occur within a short time period of other specific alarms. If you use your alarm analysis software to list the alarms always occurring within (for example, one second of each other), you will likely find a large list of duplicate alarms to address. Such alarms are highly likely to be multiple annunciations of the same process event in different ways – an undesirable situation. The individual situation will determine which are kept and which are not or what adjustments must be made. A high quantity of potential duplicates shows the need for rationalization to eliminate them.

Configured duplicate alarms

Interconnections between points in a DCS can create cases of duplicate alarm configuration. For example, a process measurement may be sent from a sensor point to a selector point, to a totalizer point, to a logic point, to a controller point and so forth. Often, a bad measurement alarm is configured on each point, and thus, if the sensor point goes into that condition, several simultaneous alarms will result. These distract the operator by annunciating multiple alarms caused from a single event. There should only be one such alarm configured on the point where the operator is most likely to take the action. If a controller is involved, it (and not the sensor point) is the proper place, since the action taken from a bad reading will likely put the controller in manual and adjust the output.

Nuisance instrument diagnostic alarms

It is surprising to see the number of alarm occurrences on most systems, representing a bad measurement or similar instrument problem. These are often in the hundreds or thousands. (See Figures 1 and 14. In Figure 1, half of the top 10 are instrument diagnostic alarms.)

When a loop was designed, did someone tell the control engineer the following? "Oh, and by the way, I want this sensor to go into bad measurement under the following conditions and I want 650 bad measurement alarms per week at a minimum." And, if that had been told to the best control engineer in the company, could they have done it? Probably not. Yet, we find these on almost every system we see.

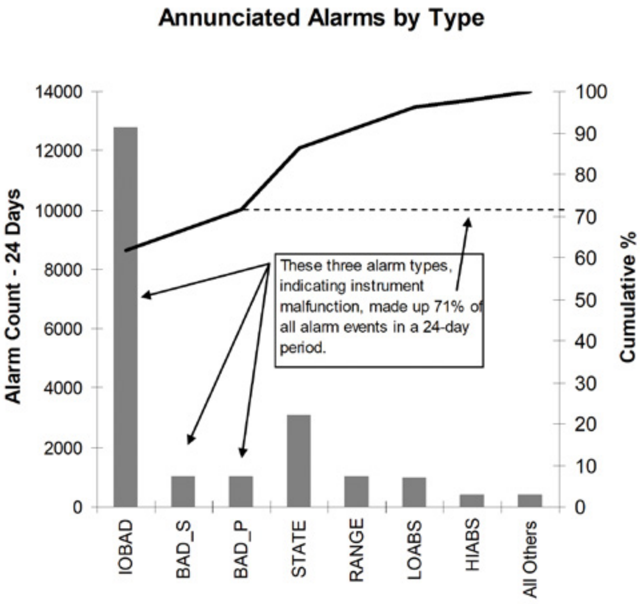

These three alarm types, indicating instrument malfunction, made up 71% of all alarm events in a 24-day period.

Figure 15. Alarm system dominated by instrument diagnostic alarms

Since no instrument was designed to be in such a state, every one of these situations can be fixed, and they should not be tolerated. They are misconfigured in range, in "measurement clamping" or there is an installation problem such as impulse leads filling up. The original justification for installing a flow meter probably did not include a specification that it was okay if it didn't work half the time. If that had been proposed, the money would have never been spent to buy it in the first place.

These situations must be addresses promptly since, often an instrument malfunction removes an identified, rationalized indicator of an abnormal situation from the operator's view. The time operators spend confirming the instrument problem reduces their attention to other operator duties.

Long ago, the available instrument sensors had a significant trade-off between accuracy (significant digits) and range; you could obtain high accuracy only over a small range, probably less than the possible variation of the process. Control engineers knew this trade-off and were accustomed to designing within those constraints. But when such sensors with constrained ranges were implemented in a DCS, the "bad measurement" alarms occur frequently and did not represent abnormality.

The digital electronic revolution that gave us the DCS also gave us much-improved measurement sensors. Those old scaling constraints can usually be thrown out of the window. Modern sensors can generally provide all the accuracy needed over the entire range, but the process will likely vary. However some engineers continue to follow the older configuration practices and do not consider the consequences of generating many bad measurement alarms during conditions such as startup and shutdown.

The default should now be configuring the instrument range for the entire range of possible values the process can have (including ambient) and then seeing if the accuracy you get is enough. If not (rarely, with modern transmitters), buy a better transmitter. But don't configure the range where you know you will get a bad measurement state at ambient or shutdown conditions.

Differential pressure flows are often the worst offenders. If, at zero flow, there is a slight imbalance in the leads, the meter attempts to report a slight backward or negative flow. The flow range might not be configured for a slight negative, so the bad measurement condition and alarm occur. Such points should be configured to handle the zero case. A cutoff can be configured and clamped at a zero value, so a small negative flow number is not used, which could affect some downstream calculations.

Most DCSs can clamp an analog value at the end of the range rather than go into a bad measurement state. This ability should be fully understood and used properly, which means reading the documentation. Controller points using the value will usually have "shed modes." These are predetermined actions to take when a measurement goes bad. These should be chosen with care.

Ongoing work process

A work process must be in place to identify and resolve new nuisance alarms. The process will change or be modified, sensors will age or develop problems and new nuisance alarms will appear. Ongoing alarm analyses can spot and report these, but it must be someone's job to act and correct the situation. We have seen that once nuisance alarms are initially resolved, the operators will notice that, realize it can be done and not be very tolerant of new nuisance alarms. This is a good thing.

Summary

Elimination of nuisance alarms is an essential, early step toward a properly functioning alarm system.

Nuisance alarms can be dealt with in several ways. The methods have been proven in hundreds of projects in all industry segments. Dealing with very few alarms, in the ways we have shown, can create a large and easily calculated improvement in an alarm system with very little time and effort.

About the author

Bill R. Hollifield

Retired principal alarm management and HMI consultant

Bill is a retired principal consultant responsible for the areas of both alarm management and high performance HMI. He is a member of the ISA SP-18 Alarm Management committee, the ISA-SP101 HMI committee, The American Petroleum Institute's API RP-1167 Alarm Management Recommended Practice committee and the Engineering Equipment and Materials Users Association (EEMUA) Industry Review Group.

Bill has multi-company, international experience in all aspects of alarm management and HMI development. He has 28 years of experience in the petrochemical industry in engineering and operations, and an additional 18 years in alarm management and HMI software and services for the petrochemical, power generation, pipeline, pharmaceutical and mining industries.

Bill is co-author of *The Alarm Management Handbook, The High Performance HMI Handbook and The Electric Power Research Institute (EPRI) Guidelines on Alarm Management for both Power Generation and Power Transmission.*

Bill has authored several papers on alarm management and HMI and is a regular presenter on such topics in such venues as API, ISA, and Electric Power symposiums. He has a BSME from Louisiana Tech University and an MBA from the University of Houston.

In 2014, Bill was made an ISA Fellow.

About Octave

Octave is a leader in enterprise software, turning data into decisive action and intelligence into your edge. Our software solves for and simplifies complexity, from the design and build to operations and protection of people, property, and assets– for any scope, at any scale. For decades, we've partnered with customers to sharpen performance, elevate efficiency, and amplify results. From factory floors to entire cities, our solutions are tuned to scale up what's possible from day one onward.

octave.com