The Role of Process History in Reducing False Alarms

Alan Mahoney, Technical & Operations Director; Oksana Koltsova, R&D Manager; Robin Brooks, Managing Director; PO Box 43, Gerrards Cross, Buckinghamshire, SL9 8UX, UK

It is essential to consider process history when reviewing operator alarm limits in the context of alarm stewardship and formal rationalization. For far too long limits have been implemented on a 'try-it-and-see' approach that leads to higher operator load weakening the operators' trust in the alarm system, potentially leading to delays in acting, and adding extra work later in re-reviewing the limits. Reasons for avoiding such consideration in current alarm reviews may include the complexity of the data and overhead of review. In this paper we demonstrate techniques of data analysis based on the parallel coordinate plot that streamline and improve this, enabling the inclusion and reference to process operating envelopes in all alarm reviews.

Operator alarms are essential for the economic operation of process plants, avoiding process downtime and contributing toward increased process safety. There has been much recent attention on these systems and the introduction of the EEMUA 191 guidelines and IEC62682 standard for the management systems concerned with alarms. However, of necessity, little detail is provided for the practice of setting these alarm limits. In most approaches to alarm limit setting and philosophies, while attention is paid to consequences and consequence threshold, current process performance and capability is rarely considered in this process. This leads to alarm sets that cause unacceptable operator performance and do not contribute to improved operation or safety.

Including historic process data in the alarm rationalization process is required to avoid these pitfalls. The size of the required datasets, and the difficulty of visualizing, let alone interrogating, this data using traditional methods, has led to adapting the parallel coordinate projection as the enabling technique for visualizing sets of alarm limits and their relationship with operating history, operating envelopes, and operator response.

Using interrogative visualization of process history in the alarm review context increases effectiveness, producing limits that already consider process operation, and identifying early in the process issues that are usually only seen after the new limits have been put in place, allowing necessary operational and engineering changes to be investigates months or more earlier than now, while producing a set of limits consistent with this operation. These methods also increase the speed of the review, allowing smaller teams to perform most of the work independently and providing a common framework for communication.

Pitfalls and issues that can be identified by using the historic data include mis-sized equipment, poor control, lack of capability and failed equipment. We demonstrate how these are identified in the context of alarm review.

Introduction

Operator alarms are systems that visually and audibly alert the operator to process operating conditions and are one of the main systems in an operating console. The operator alarms are the first line of defence in process plants, designed to alert the operator to deviations, perhaps dangerous, from the normal carefully designed operation. Operator alarms are not the only defence from process excursions. Outside the operator alarm, automatic safety systems such as trip systems, interlocks, bursting disks, pressure relief valves, etc., should bring the plant to a stable, safe, shutdown. Use of these systems have an economic cost, though. Triggering them causes process disturbances with effects ranging from degraded operation for minutes to a full plant trip, requiring hours or days to come back online. Operator alarms help avoid these systems activating, alerting the operator to take more benign action to keep the process from activating these safety systems or to bring the process down in a controlled fashion. Thus, one measure of the goodness of an alarm system is the number incidents and duration of downtime due to activation of these automatic systems.

Often the tools and information taken to alarm review meetings are design limits, design operating limits and regions, equipment limits, the results of HAZOP studies and the like. These are essential for alarm review but if they were sufficient, the alarms as initially designed or reviewed for plant changes would be sufficient and there would be no need for current review. Some of the things that may have changed since the alarm setpoints were set are deliberate equipment modification/plant change, equipment wear, changing feedstocks either discrete as to a different supplier or continuous as with reservoir aging or changes due to production rate, changing plant operating targets, equipment/catalyst aging and degradation, etc. Some of these things are deliberately planned and should include alarm review as part of the implementation process, but some are unplanned, or even unknown, and can have an effect on plant operations and therefore alarm systems. Other things may also impact alarm performance, like equipment faults, operator action, or errors in setting limits in the past.

Operator Alarm Objectives

Figure 1 shows our conceptual idea of an alarm system. On the left we see alarm limtis positioned at the edge of normal operation in green. These alarms signify deviations from normal expected operation, giving the operator maximum reaction time to take action, with three possible interim goals. First, actions returning the unit to normal operation by adjusting process inputs or requesting diagnosis or maintenance. Especially in the case of maintenance, this may be a slow process and is often insufficient, and even not the most immediate required action. Second is maintaining the system in an abnormal, degraded state, for instance by relying on a field operator to monitor a level while waiting for maintenance. If that is not feasible, the third class of actions is a more benign shutdown or unit isolation than would result if the automatic safety systems were activated.

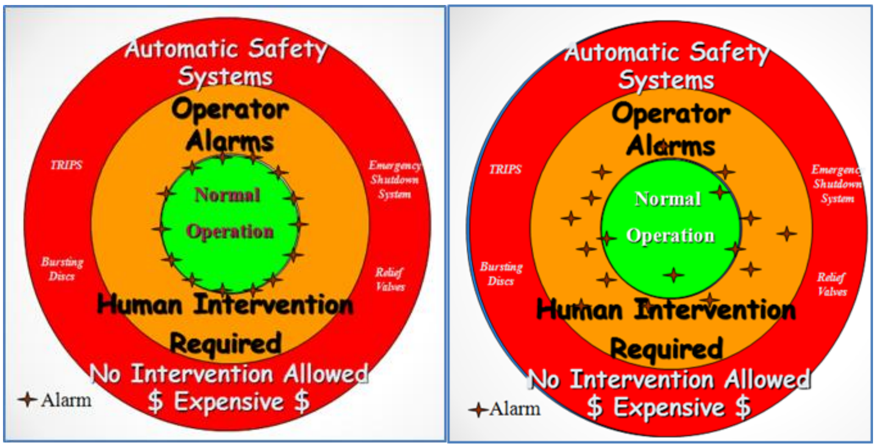

Figure 1: Schematic depiction of alarm limits. On the left, the expected configuration of alarm limits at the edge of normal operation. On the right, a more realistic picture of what is commonly seen in practice.

The primary goal of operator alarms is to increase performance of a unit by minimizing reduced capacity resulting from trips, downtime, maintenance issues, and improving on-spec and efficient operation. The key measures of this are economic: frequency of trips and periods of resulting recovery as well as overall plant efficiency. Evaluating normal human factor KPIs (annunciation rate, standing alarm count) are a necessary minimum for operators to be able to respond, but alone do not capture the alarm system quality. As alarms must direct behaviour, alarms that don't or can't direct action when it is required are also poorly set, but will never appear from a review of historic alarm events.

Many sources of information feed into alarm review including design, HAZOP studies, LOPA analysis, operation and alarm philosophy, and control system capability, but process history is one that is often overlooked. Sometimes spot values or a brief recent period of process information are all that are examined. This data is nearly always examined in isolation, individually for each variable as it is considered. What is often needed is to consider process history, including controller outputs and setpoints, for a year or more (to include seasonal variation) across a piece of equipment or even a unit.

Interrogating Process History with Parallel Coordinates

In order to set the alarm limits at the boundary of normal operation, it helps to have an idea of what constitutes normal operation. For a plant that has been running some time, this is implicit in the data in the plant historian, and this data constitutes the space of historical operation. There is a corresponding 'operating envelope' that corresponds to this historical operation. The term 'envelope' is used to describe a closed boundary between an inside and an outside space. In two dimensions the boundary is a closed curve that separates two areas. In three dimensions the boundary is a surface. For use in alarms, the data needs to be filtered to remove periods of trips, shutdowns, and other abnormal operation to arrive at a normal operating envelope. This is not the same as ideal operation which occurs much more rarely. In a process unit of any reasonable size, it is likely that there are instrument or operational issues ongoing at any point in time. In these cases, one would expect alarms, if not shelved or suppressed, on variables around the faults, so normal operation also doesn't necessarily correspond to alarm-free operation.

Many of the process constraints are relationships between two or more variables, and if these are plotted for a single pair of variables at a single point in time, would look something like the left image in Figure 2. The constraint relationship between these two variables will change as other process variables change, resulting in a sequence of two-dimensional spaces as seen in the right image in Figure 2. Conceptually, all of these constraints together constitute our operating envelope. While some of these, such as production targets, product qualities, and equipment operating limits are known and understood, other constraints are hidden deep in the thermodynamics, chemistry, and fluid dynamics of the process and are not explicitly known in detail. They are implicit in our space of historic normal operation, though.

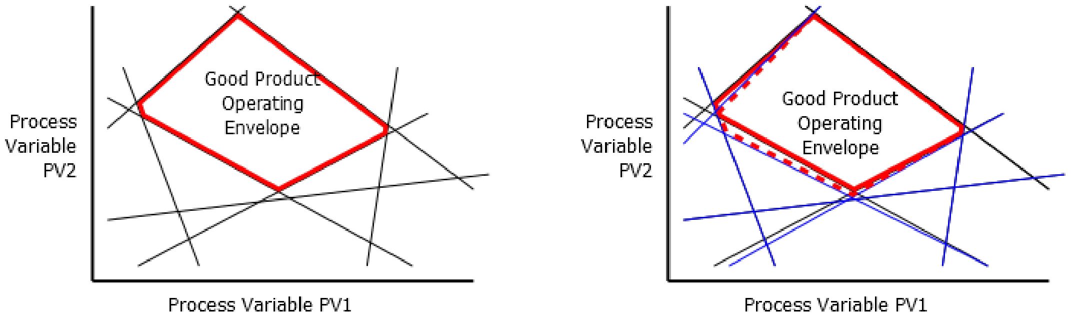

Figure 2: Schematic of two-dimensional operating envelope. Black lines are initial constraints on these variables. Blue lines show the constraints sometime later, with a shifted operating envelope.

Alarm limits and operating limits are nearly always defined as high and low values on a single variable so are by definition independent of any other process variable. Where more sophisticated relationships are used, calculated variables are introduced so that the alarms are still high or low values. They appear as vertical or horizontal lines when drawn onto our picture of a simple operating envelope. Figure 3 shows high and low limits for both PV1 and PV2 in green and two high safety system limits in brown.

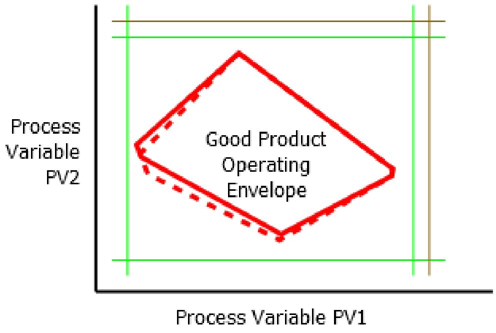

Figure 3: The relationship of operating envelope (red) with alarm limits (green) and safety system limits (brown).

This 2-d picture is conceptually convenient, but any interesting process has many more variables in our envelope. It would be useful to have a picture or graph that allows this extension to many more dimensions.

The solution is Inselberg's parallel coordinate transformation where the axes are drawn parallel to each other [Inselberg, 2009], introduced in Figure 4. Each variable has its own axis just as in a conventional graph. A set of related variable values, most commonly related by a moment in time, are plotted on their own axes and the values are joined with straight lines to give a zigzag line which indicates that these values 'belong' together. The zigzag line is a representation of a point in time. The individual variable values are the coordinates of the point.

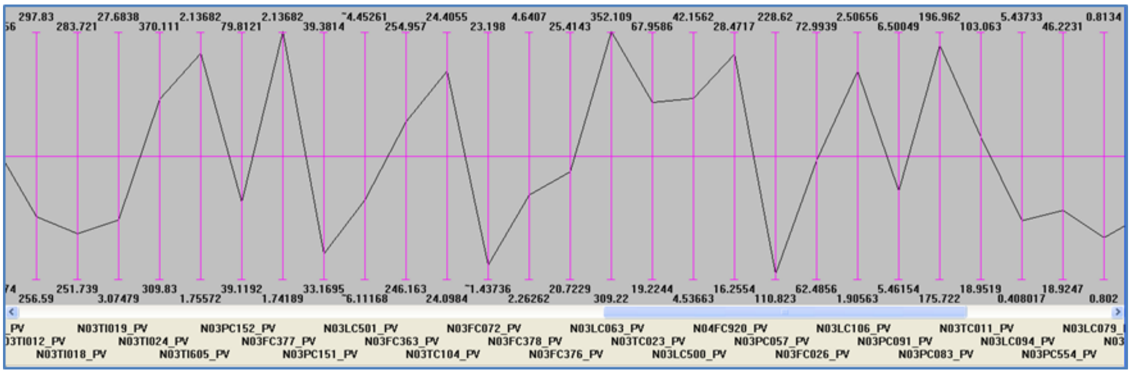

Figure 4: a parallel coordinate representation of 27 process variables at a single point in time.

The power of graphs is exploiting humans' visual pattern recognition capabilities. To exploit this, any useful graph must show many points, and the parallel coordinate graph produces its own set of patterns. This data in Figure 5 represents 3 months of operation of an hydro-desulphurisation (HDS) unit and was taken from the process historian database at 10-minute intervals to give a total of 13,444 time points, each plotted as a zigzag lines. The dataset has a total of 175 variables. Patterns among the lines are apparent and can be correlated to process behavior. Conditions under which the process has often operated are darkly coloured, those with rare operation such as transitions between modes have sparse lines. Additionally, relationships between the axes are apparent, and banding shows correlation between operations across different variables in the unit.

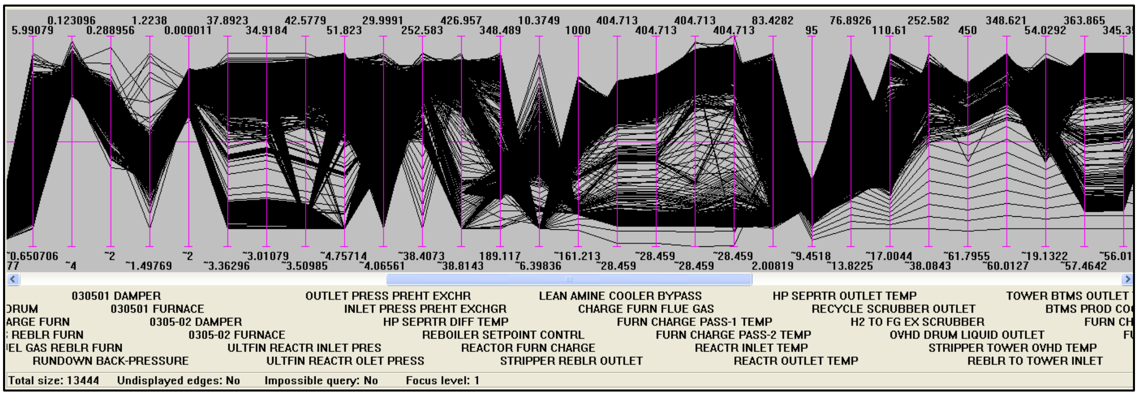

Total size: 13444 Undisplayed edges: No Impossible query: No Focus level: 1

Figure 5: three months of operation at 10-minute intervals. A total of 13,344 points in time (zigzag lines) are shown. This is the space of 'all operation.'

Superimposing the existing alarm limits on the graph as red triangles as seen in Figure 6 makes the relationship between alarms and the operation immediately clear. Some alarm limits are inside the solid black area so give alarms at least some of the time; others are so far outside the black area that they may never annunciate at all.

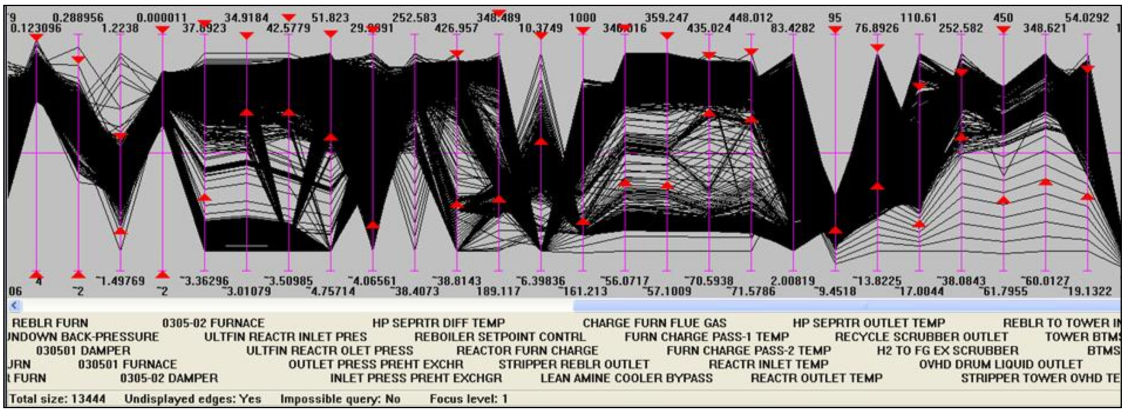

Total size: 13444 Undisplayed edges: Yes Impossible query: No Focus level: 1

Figure 6: existing high and low alarm limits imposed on three months of operating data.

The parallel plot allows a way to quickly evaluate existing and proposed alarm limits as well as set operating alarm limits at the edge of normal operation. The basics of setting limits with attention to human factors KPIs can be seen in earlier papers [Brooks 2012].

Alarm Limits in the Context of Operating History

Discrepancy between alarm settings and historic operation can arise from many causes, have many effects on the alarm system, and show different behaviour in the parallel plot. There can be alarm limits in the middle of operation, alarm limits set so far outside operation they will never be violated, and operation that is in standing alarm. Other issues such as poor operation, redundancy, poor design may also contribute. Each of these are considered and some practical examples from real practice are given in this section.

Figure 7 has a year of process operation for a process unit with superimposed high and low alarm limits in cyan. These limits are represented as the edges of a selection. That none of the data rows are coloured cyan indicates that at least one variable was in alarm at each point in time. Here there are at least four variables visible that were violating their alarm limits the entire year. For TI97, second from the left, the operation was over the high (only) limit of -25 for the entire year. This alarm has been effectively disabled. Whatever its design intention, it will never annunciate again unless the process moves significantly cooler and then warms again. Additionally, these alarms sit in the standing alarm list, making it difficult for the operator to use. The standing alarm list should ideally have zero alarms (a clean board), and the operator should be able to explain any alarms appearing there. This makes it very easy to check this list at the beginning and periodically through the shift to ascertain the current operating situation. If there are dozens of alarms permanently in this list, spotting small changes becomes impossible and operators tend to ignore the list completely. These may arise from maintenance issues, for instance sensor replacement, and the alarms should be temporarily disabled with procedures to ensure they are enabled again when the maintenance is performed. More serious are when they are set in relationship to an idea of normal process operation that has not recently, or even ever, been attained. In this case, examining the process nearby seems to show no suspicious operation. In fact, the alarms had been designed for a mode of heat integration that was used only briefly during plant commissioning and the alarms were no longer required. Another variable, FC118, in the same figure, shows a flow always at zero, a candidate for state-based suppression.

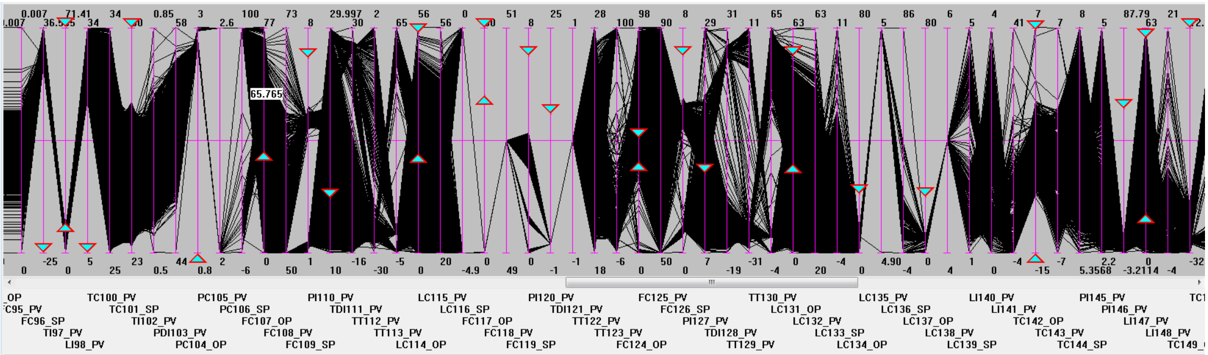

Figure 7: A subset of alarm limits (cyan) superimposed on a year of normal process operation. Note evidence of variables always in alarm (TI97, TI98, TC100 at left), alarm limits buried in operations and alarms far outside operating history.

Alarm limits in the middle of normal operation are usually the target of the most attention. These include frequently annunciating alarms, or ‘bad actors’, but also cases where the alarm is positioned between or among various operation modes where the process variable stays in alarm for hours or day at a time, effectively disabling it for those periods. The ‘bad actor’ terminology presupposes the fault is with the alarm limit, but it frequently that the alarm limit is reasonable and correct, but that operation cannot, or will not, operate within the limits. Addressing this operation often requires maintenance, engineering or other support and there may be a multi-step action plan for resolving the alarm issue. Where there are no faults, moving alarm limits to the edge of normal operation is the simplest and often appropriate action, but there are many cases that require consideration.

One is actual physical constraints. These are a large component of limits in non-process areas, such as storage. Figure 8 shows a temperature in a reaction bypass that is routinely in alarm, and only goes briefly out of alarm when the bypass is used. During normal operations, it is clear that the alarm limit is not, or cannot be achieved by the operators. Inhibiting this alarm when the bypass isn’t in use might be considered, however, one intent is to keep reactants from condensing when the bypass is activated. Currently, the alarm annunciates going from abnormal to ‘normal’ operation and has less urgency and has been ignored for years, hiding the true issues. In this case, the alarm review brought attention to a vital maintenance issue. The alarm is useful and required where it is, engineering must be done to address why operation cannot currently meet it. These instances are relatively rare, but of very high operation and safety importance when discovered.

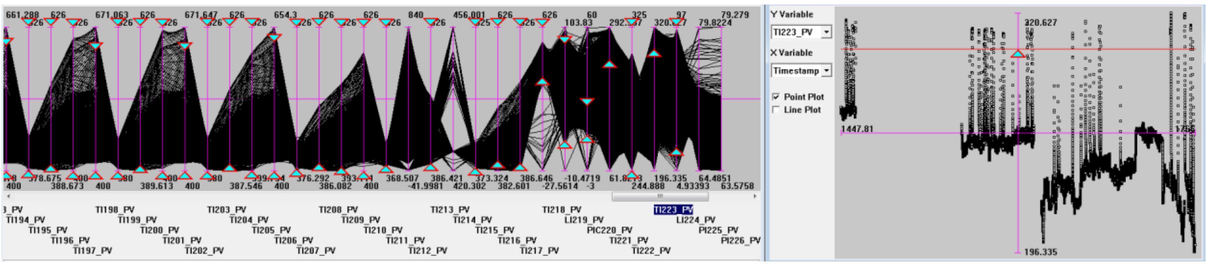

Figure 8: Alarm limit of TI223 is normally in alarm. Time-trend plot on right shows only brief periods when it is not in alarm.

Figure 9 shows a variable that has a very distinctive time trend. The parallel plot shows the upper limit buried in normal operation, suggesting that the alarm may be set too low, but the time trend illustrates a slow fill over days, followed by a quick drain (less than the 10 minute data frequency). Here, the level is under manual control and the high alarm is used as an action trigger for the operator. This is contrary to most definitions of an alarm, indicating an abnormal situation, which this (at least currently) is the normal operation of this tank. If the capability exists, moving the categorization of this to an alert and adding an alarm outside the normal response range would be appropriate, but that capability is not commonplace, so these sit in the alarm system alongside the other alarms. In this case, attempting to move the alarm limit outside normal operation would be unproductive and impossible.

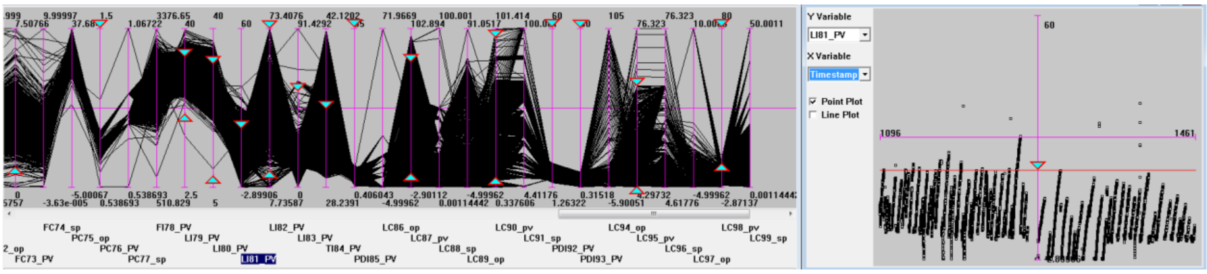

Figure 9: LI181 is frequently violated in normal operation, the level going above the high limit.

Figure 10 shows a number of alarm limits far outside normal operation. These never show up on alarm reports, because they never annunciate. These have varying frequency, but are often 10-30% of the configured alarm limits. Occasionally, these limits are reasonable, as there is normally no action until the limit is hit. For instance, a filter that is changed at every shutdown may have a limit on the pressure differential that would require a specific process interruption to change. There is no reason to bring the limits nearer the normal operation as no action should be taken for smaller values. In most cases, however, the limit can and should be brought in. This provides the operator earlier warning that the process is deviating from normal and allows earlier action to address the issue. TI67 is one of these where the control is very good and stable over a year, and the alarm range is so far out that significant process disturbances would have to occur. By bringing the limits in the operator has additional time to respond and a broader range of options before automatic systems are engaged.

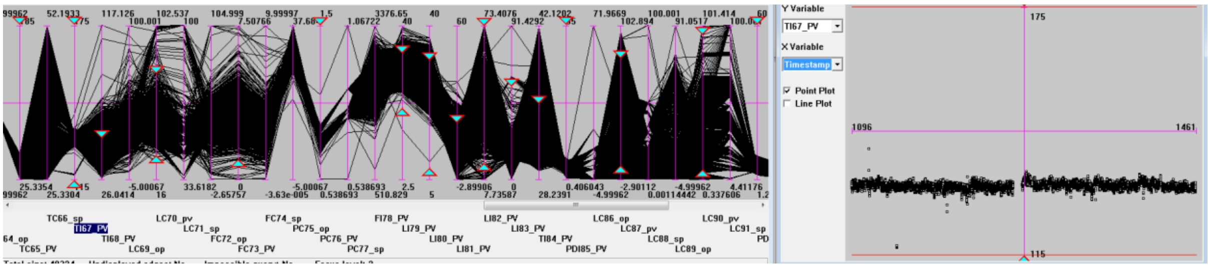

Y Variable TI67_PV

X Variable Timestamp

☑ Point Plot

☐ Line Plot

Figure 10: TI67 is among the variables with alarm limits set far outside normal variation.

Redundant/correlated alarms are also easy to see from the history. Figure 11 shows two correlated levels that seem to occasionally vary in calibration. The high limits are seen to be at comparable levels, but the low will ring in at different times. Normally replacing these with a single configured alarm with a potential second deviation alarm would be preferable. Additionally, here, the high limit is well within normal operation and should be addressed. In this case, the process setpoint had shifted and the alarms were not adjusted.

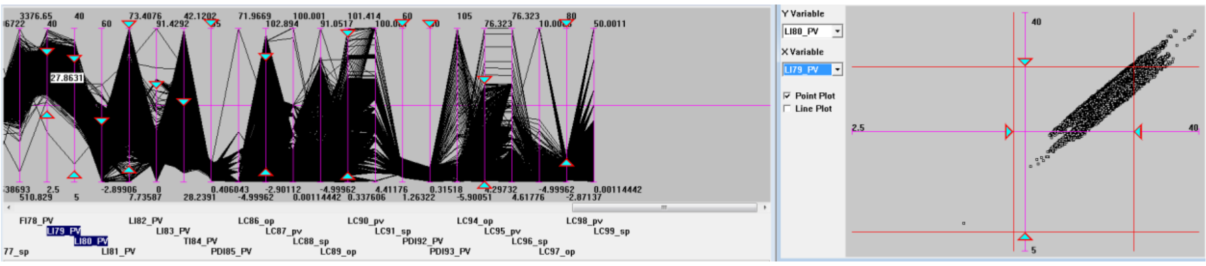

Y Variable LI80_PV

X Variable LI79_PV

☑ Point Plot

☐ Line Plot

Figure 11: Redundant alarms. High alarm settings on LI79 and LI80 are at a similar level, and ring together. Low alarm settings are at two different levels.

Figure 12 shows another relationship between two related variables. Here we have the output temperature of two parallel fin-fan coolers. The alarm limits here are buried in normal operation and cause frequent alarms. In looking at the appropriate level for these limits a new problem is encountered. There are three main populations in the scatter plot between the two variables. The bottom left shows the normal correlation between these two variables, going up and down together as process temperatures or ambient temperature changes. The two populations above and to the right of this population consist conditions where one of the two fans has tripped and stopped running. This leads to an increase in the outlet temperature of the corresponding cooler. The alarms are meant to detect this condition. Two different strategies have been used in placing the alarm limits. On the x-axis, TI35, the alarm limit has been positioned just above the maximum normal temperature. As can be seen that it intersects the tripped population, it will miss alarming on trips when the temperature is cool. On the other hand, on the y-axis, TI37 has an alarm limit positioned below the minimum tripped temperature but well inside the normal operation, causing the operator to doubt whether the alarm is ever accurate. In this case the alarm must be redesigned, looking at heat across the exchanger, or a deviation from process or ambient temperatures as no alarm limit can properly do the job on either of these variables. This also illustrates the general case that the limits and process constraints we are often interested in have multi-dimensional and even complex constraints. This can be approached with higher dimensionality models and better approximations of the operating envelope than simple single-range alarms [Brooks 2004].

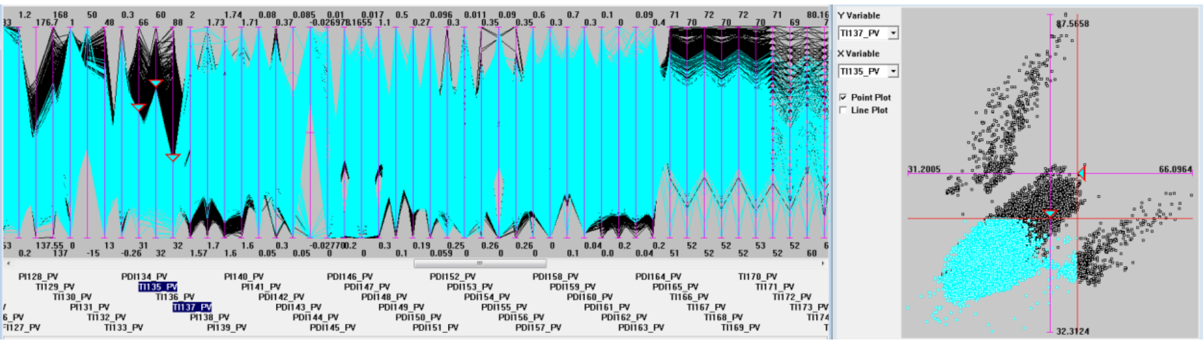
Y Variable
TI137_PV
X Variable
TI135_PV
☑ Point Plot
□ Line Plot

Figure 12: Alarm limits on two cooler output temperatures. Note three distributions of the temperature relationship.

Another use of process history is quantifying process rates-of-change (ROC). Maximum rates of process movement are dependent on system operation, and the developing fault/situation. Ideally, response times could be calculated for all possible faults. In practice, this is infeasible for most process upset scenarios and so common or arbitrary rates of change are chosen for broad classes of variables in a plant. While still of smaller magnitude than what is expected during and actual process disturbance, 'normal' rate of change will provide a better estimate of time to respond. This can be taken from process data, avoiding noise issues by considering, for instance, 10 minute data to capture rate-of-change. This will necessarily give a lower bound on the maximum expected rate-of-change, alternatively an upper bound on the required stand-off between alarm and safety system set points. This can also highlight undersized equipment or lack of capability, for instance where the normal rate of change can be 6% of range a minute and a ten minute response time is desired for both high and low alarms. Regardless of process operation and capability, such a set of alarm limits is impossible without process changes.

The images so far have focused on a single set of alarm limits. Often it is instructive to bring in multiple sets on the same graph. As a more complete example, Figure 13 shows four sets of limits and their relationship to process history, quickly illustrating the concentric arrangement of normal operation, operator alarms pre-trip and safety limits. Additional sets, including integrity limits can easily be accommodated.

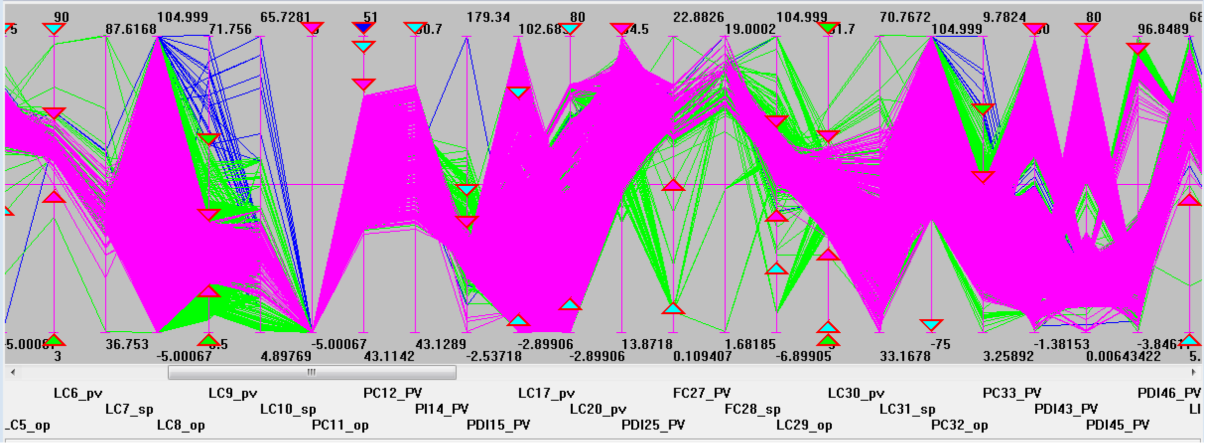

Figure 13: four sets of alarm limits with process history. Safety trip levels are shown in dark blue, hihi/lolo limits in green, hi/lo limits in red, and pvervious hi/lo limits in light blue. Coloration shows periods where no alarms of that class were active.

Conclusion

Determining the appropriate values for alarm limits is a key part of rationalization and ongoing alarm maintenance. Today, this is often done with very limited reference to process history, but the process history is essential in driving improvements to the plant operation. Both in determining the context of alarms in ways that alarm logs, alarm statistics and even single-variable histories cannot capture, but also driving process improvement as often alarm problems are symptomatic of deeper process problems. Accessing and visualizing this information easily and quickly during the review process is key to utilizing it efficiently. The alarms can only be judged for consistent with regard to their created enclosing hypercube with access to the process history. With the proper tools, adding process history to the alarm review process doesn't slow the process, but tends to speed discussions and conclusions.

References

Brooks, R.W., Thorpe R., and Wilson, J.W., 2004, A New Method for Defining and Managing Process Alarms and for Correcting Process Operation when an Alarm Occurs. *Journal of Hazardous Materials* 115:169-174.

Brooks, R.W., Mahoney, A, Wilson, J., and Zhao, N. 2010, The Rationale for Alarm Rationalization. *TAMU Instrumentation Symposium.*

Brooks R.W., Mahoney, A., Wilson, J., and Zhao, N. 2012, Operator Alarms are the First Line of Defence. *Hazards XXIII Proceedings*.

Inselberg, A, 2009, *Parallel Coordinates*, Springer-Verlag, New York.