# Consistent Alarms Improve Process Economics and Increase Process Safety

Robin W. Brooks PhD, Alan Mahoney PhD, Timothy Triplett BSEE, John Wilson PhD, CEng, Na Zhao PhD.
Timothy Triplett is with Coherent Technologies Inc., Palestine, Texas, USA.
The other authors are with PPCL, PO Box 43, Gerrards Cross, Bucks. SL9 8UX, UK

Correspondence to Robin_Brooks@ppcl.com

## *Abstract*

*There has never been a scientific method for finding process operating alarm limits other than a few strength-of-materials calculations of safety limits which are outside the scope of this paper. Almost all previous work on Alarm Rationalization has started from the Process Alarm Log and proceeded from there based upon qualitative judgements. We present our scientific method for determining safer and more profitable operating alarm limits based upon process historical data and demonstrate this technique with implementation examples.*

*Much tighter alarm limits are usual with our method giving the operator more correction time in the event of a process excursion. A process excursion caught earlier in its development usually requires less dramatic corrective action. The new alarms are geometrically consistent with each other and this approach thus takes the first steps towards recognising that alarm limits interact with each other due to the process variables themselves interacting with each other. Our better alarm limits reduce the proportion of false alarms resulting in increased operator confidence and leading to alarms being used as positive aids to assist the operator in the achievement of operating objectives and in identifying improvement targets for process control.*

*Alarms are shown to contribute economic value to a plant in their own right, leading to a more positive environment for improving or rationalizing alarm systems than exists in today's environment - where the motivating forces are either legislation or avoiding the possibility of a major plant disaster.*

*Implementation for multiple process operating Modes (also known as State-based alarming) is straightforward using this method and is shown.*

*This method requires no user mathematical analysis. As such it not only gives better alarm limits and a safer process but is also being faster and easier to use than existing methods of alarm rationalization.*

*The method extends to the provision of dynamic Alerts which fully include the effects of variable interactions. These Alerts help operating personnel keep the process inside its fixed HiLo alarm limits, thus contributing even further to increased process safety. This extension to Alerts will be introduced and, based upon experience with its use in fault detection and multi-phase batch*

*process control applications, the possibility exists for its use in economically automating time-varying Modes such as Start-Up and Shut-Down of continuous processes.*

## *Alarm Systems Today*


Figure 1 Alarm Limits as they are imagined to be

There are two major alarm systems to be considered. The first is the Safety Alarm System responsible for taking control and shutting down the process in extreme process excursions which both the process control system and the operator have been unable to prevent. The value it provides is as a last line of defense in preventing an extreme excursion from turning into a disaster with liabilities and costs that can run into hundreds and even thousands of millions of dollars. Its costs are viewed as an insurance premium against a disaster that most plants will, thankfully, never experience.

The second is the Operator Alarm system acting as the first line of defense and intended to draw the process operator's attention to a situation beyond the preventive capability of the process control system that requires application of his considerably greater human intelligence to resolve and correct before the safety system intervenes and trips or shuts down the plant. Automatic plant shutdowns are expensive in terms of lost production and possible consequential plant damage, and operator alarms, by giving the operator time to intervene and correct the situation, also have an 'insurance premium' value in reducing the demand upon the safety system and thus the small probability that it will fail when called upon.

More significant though is that these alarms are often known collectively as 'Economic Alarms' because they are also intended to help the operator in the achievement of the plant's economic objectives by assisting him in keeping the plant inside the operating envelope where these objectives can be achieved. Most plants would describe this as 'Normal' operation and imagine that their alarm limits are positioned around, and thus define, the boundary of the Operating Envelope within which desired economic results are achieved similarly to Figure 1. This would suggest that alarm limits are ideally the same as operating limits and the economic cost of violating an alarm limit is the delta of the values of the materials produced and the delta of the operating costs of the desired and undesired operation.


Figure 2 Actual alarm limits

The reality of today's individually-set and managed alarm limits is more like that in Figure 2 than Figure 1. Some alarms are inside the Normal Operation space so create frequent false alarms and in todays terminology are 'bad actors'. They annunciate frequently so are a nuisance to operators. The remedy is to move their limit values outwards towards or beyond where the boundary of the 'Normal' space is believed to be. But since the boundary position is unknown an uncertainty allowance is added to ensure that the alarm ceases to be a nuisance so it ends up somewhere in the orange recovery space where it can never annunciate during normal operation. But it gets worse because alarms are adjusted one at a time so over the course of time alarms are individually moved outwards, 'leap-frogging' each

other and eventually reaching a scenario such as shown in the schematic in Figure 3. The alarm limits are so far out that an extremely large process excursion would be needed to cause any alarm to annunciate and then the operator finds he has very little recovery time before the safety alarms are triggered. It is easy to recognize that the alarms are providing very little contribution to process safety which now relies almost entirely on operator vigilance. It is at this point that a rationalization project is invoked to reset all the alarm limits and start over. But, since the method is unchanged, the same cycle will be repeated with an observed period, in the oil and petrochemicals sectors, of 5-7 years. And since we believed that the alarm limits defined the boundary of Normal Operation, Normal Operation got redefined as well.

Figure 3 Process safety now reliant upon operator vigilance

In reality, only the innermost alarms in Figure 2 can ever annunciate. We have seen several situations since beginning our focus on alarm managementi , some not long after a rationalization project had been performed, where the proportion of alarms able to annunciate has been as low as 10% of the alarmed variables. The 'Bad Actors' identified by previous alarm rationalization methods may actually have been the only performing actors in the troupe and thus the stars of the show. Taking them off-stage won't improve the show!

The practical alternative to moving the alarm limit was to have improved process control capability for the affected variables so that the alarm limit was violated less often but this costs more, takes longer to implement and requires a small project which is often time-consuming to justify economically.

## *Operating Envelopes and The New Alarms*

Many of the poor alarm systems we have examined were initially considered 'good' by their owners because they met EEMUA 191[ii] performance targets (and would probably have met the more recent ISA SP18[iii] targets too). This contradiction arises because EEMUA 191 and ISA SP18 are targets on the econometric or human factors performance of an alarm system and, by themselves, are an incomplete set of targets for the primary roles of an alarm system. What was missing, given the use of alarms to identify Normal Operation as in Figure 1, is a target on the achievement of Normal Operation and the alerting of non-normal operation.

Identification of the Operating Envelope within which operating objectives can be achieved brings together process control, alarms and achievement of process operating objectives. Process control will have the role of assisting the operator to stay inside the Operating Envelope which is an extension from today's objectives of maintaining operation at a target operating point and managing the process dynamics. It is known that operating objectives can be met in the vicinity of the target operating point but generally not known how far the process can move from the target point while still meeting its operating objectives. The availability of an Operating Envelope definition therefore provides additional information to process control that is not generally available today.

The term 'Operating Envelope' has been used by generations of chemical engineers to conceptually describe a closed boundary with different properties of something inside and outside the boundary. The concept is easy to understand when only 2 or 3 variables are involved but becomes nebulous

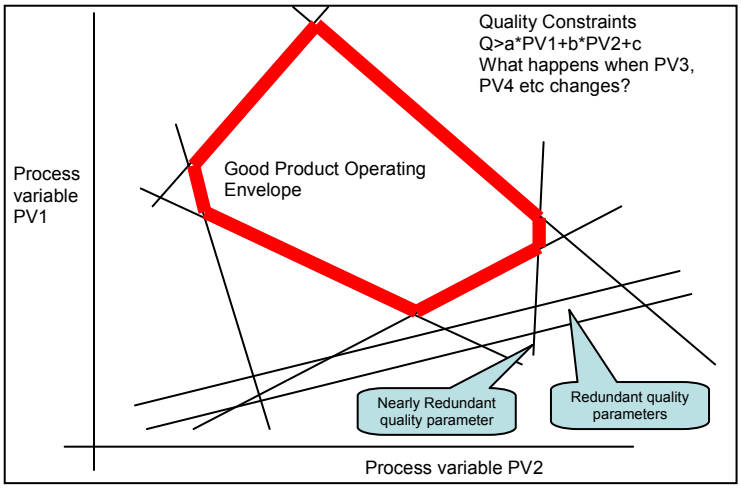


Figure 4 A simple operating envelope defined by its constraints

when extrapolated to the many tens or hundreds of variables that are involved in the operation of a process. Yet it is known to exist in that generations of operators and engineers have known that their process had a 'sweet spot' (which is actually a point in an imaginary multi-dimensional space) where it operated 'best' and have striven to locate the point and the extents of its movement as the many process variables, whose values constitute the points coordinates, changed. This is what we are referring to as an Operating Envelope. There can be infinitely many Operating Envelopes and it is the definition of 'good' that differentiates one from another. Remaining inside the Envelope will achieve the 'good' objective by which the Envelope was chosen and is defined. We refer to the 'good' objective as the primary objective and most commonly and in its simplest form in an operating plant it will be the production of saleable product meeting the customer's specifications. There will be secondary objectives such as operating at least cost that will be sub-spaces within the primary objective envelope.

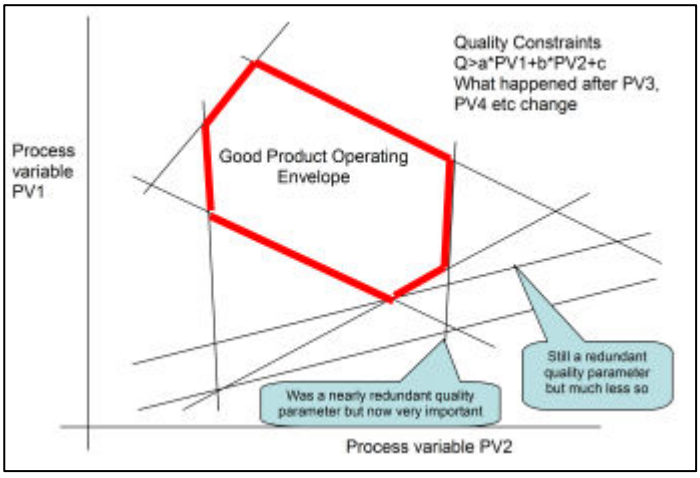


Figure 5 But the constraints are functions of many other variables so can change their position

Figure 4 and Figure 5 show what is meant by the good product operating envelope of a plant. It is bounded by quality constraints whose correlation with two process variables PV1 and PV2 only are shown so that it can be drawn in two dimensions. Such a diagram would highlight quality parameters that are redundant or rarely applicable and might be eliminated. But under different operating conditions the constraints move because the constants a, b and c in their equations are not really constants but functions of other process variables PV3, PV4…etc and we have no way to show those in our 2-d diagram.

The problem of seeing all the tens or hundreds, even thousands, of variables is one of n-dimensional geometry which was fully described by Riemann[iv] in 1853 using equations too complex to be solved except in the simplest cases and with no ability to draw pictures of n-space. The problem shown in Figure 4 and Figure 5 of how to represent the fourth axis remained as an obstacle to understanding of higher-dimensionality geometry until Inselberg[v] discovered the parallel coordinate transformation late in the 1980's. Instead of trying to draw the axes orthogonally he drew them vertically and parallel to each other causing the representation of a point to transform to a poly-line as in Figure 6 where each variable is left-justified to its own individual axis.

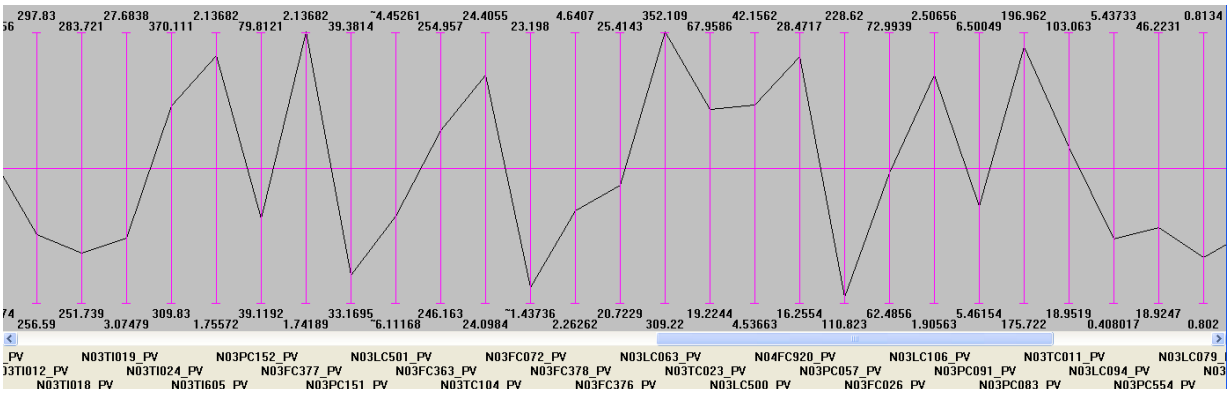

Figure 6 One set of observations of 27 variables

Adding more points to the graph produces distinctive patterns, which is the purpose of a graph, as in Figure 7 and for the first time gives the ability to see with our own eyes where the process has operated and how the variables interact with each other. This data came from a refinery hydro-desulphurisation (HDS) unit and is part of a graph of 178 variables at 5-minute intervals covering three months of unit operation and taken from the process historian.

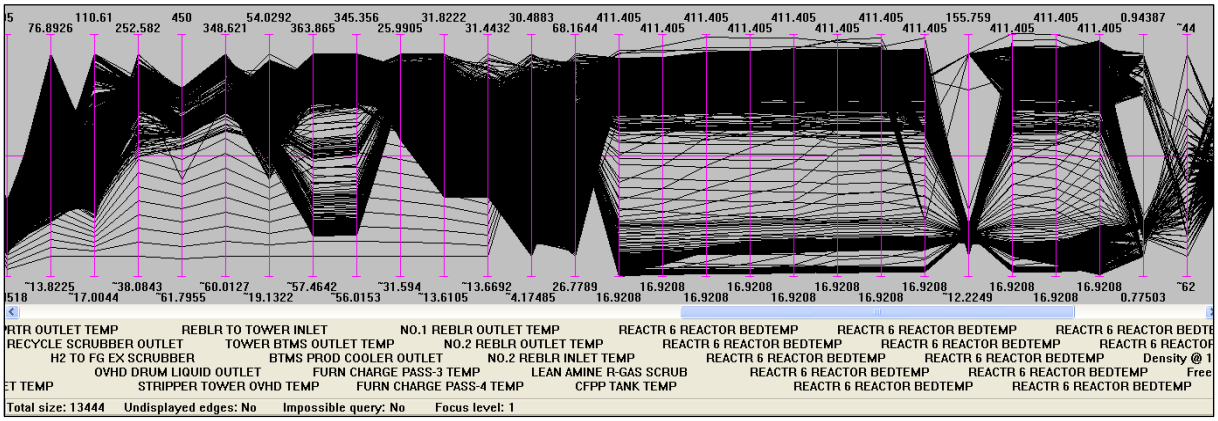


Figure 7 13,444 points from the process historian covering 3 months of operation at 5-minute intervals. 27 of 178 variables are visible

This plant operates at different times in one of three main operating Modes[1] of Standby, Kerosene desulphurization and Light Gas Oil (LGO) desulphurization. This largely accounts for the bands that are such a prominent visual feature. As in most plants today, one set of alarm limits covers all three Modes. When the alarm limits are superimposed on the graph in Figure 8 as red triangles it is immediately apparent that there has been some attempt to move some of the alarm limits inside the black area to equal operating limits, thus alarming undesired operation and so defining the economic Operating Envelope. Other limits have been set so wide that they will never annunciate. They are 'good actors' in today's uni-variate alarm rationalization terminology so will receive no attention and may escape the HazOp scrutiny of the multi-disciplinary Alarm Review Panel. The performance of

[1] Modes refer to the operating intention set by the production planner whereas States are usually taken to refer to the actual State the plant is operating in now. We think of the relation of States to Modes in the same way as that of PV's to SP's.

the alarm system is bad in that the alarm annunciations per hour (Figure 9) and Standing Alarm Count (Figure 10) are both high and in the whole 3 months there was never even one 5-minute sample of operation with no alarms present. In fact during Standby Mode (low values of most variables) the alarm display showing 41 variables in alarm means that any real alarm has a high probability of going unnoticed.

The discussion has moved to alarm limits because it is apparent that the dense coloured areas of the parallel coordinate plot are composed of a cloud of points whose envelope is an Operating Envelope of the process. There are many possible Operating Envelopes corresponding to the many possible Operating Objectives (not all are necessarily desirable) and we can see them in historical process data by applying the operating objective criteria to select and highlight points that met that objective. At last we can see an Operating Envelope composed of as many variables as we wish and shown in as many dimensions as necessary. We will go on to see that alarm limits and operating limits are merely simple ways to get a first approximations of two Operating Envelopes as hypercubes and later in this paper will introduce Alerts to give a much better representation of an Operating Envelope.

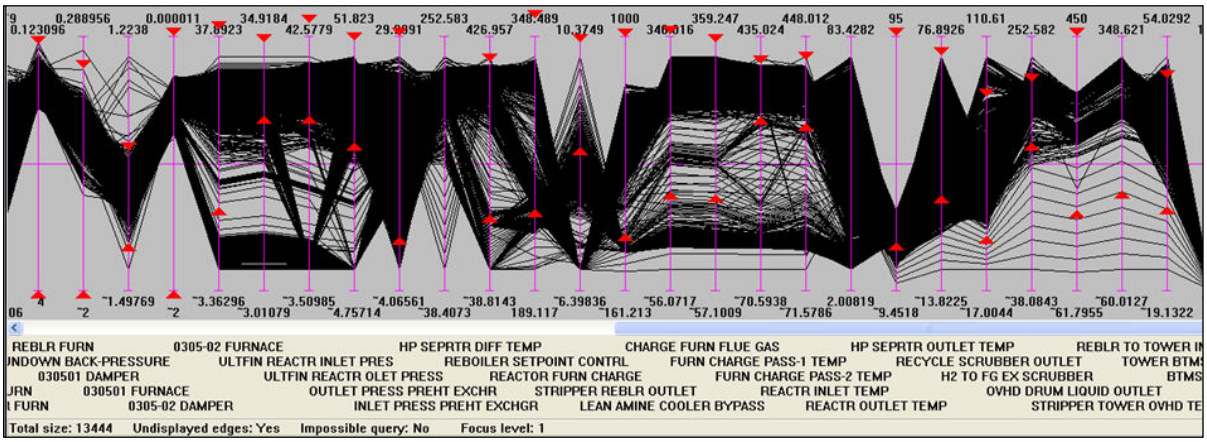

Figure 8 Existing HiLo alarm limits superimposed upon three months of operating data

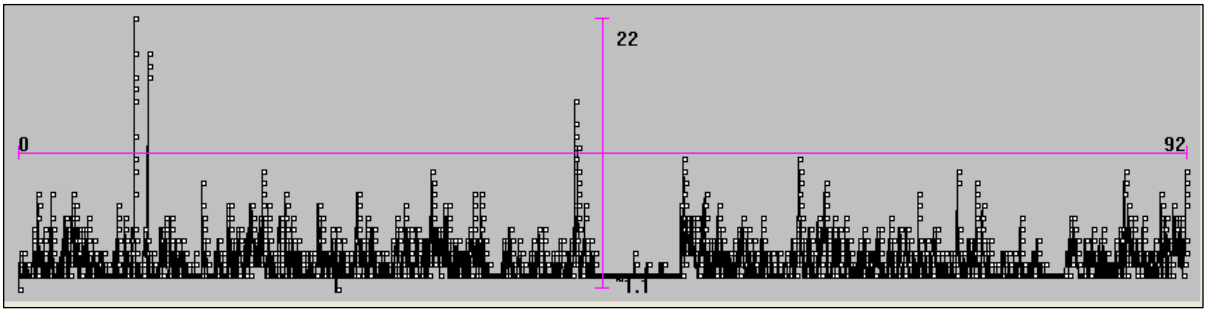

Figure 9 Annunciations per hour peak at 22 during this 92 day period and are typically 3-4

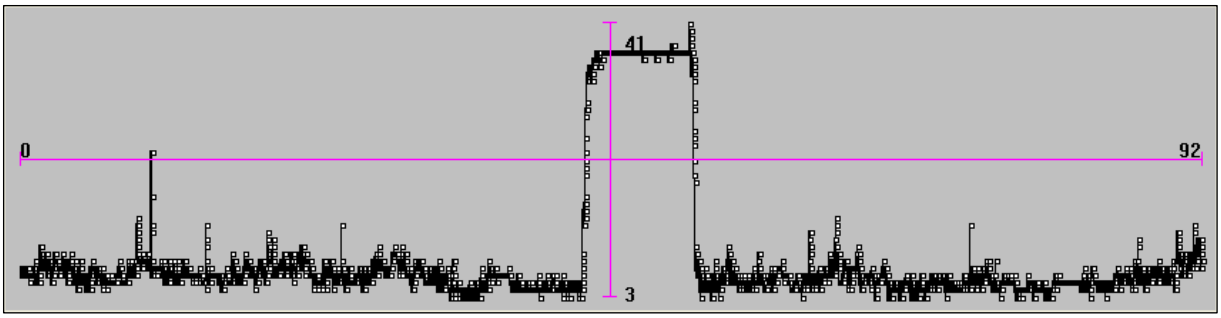

Figure 10 The count of Standing Alarms peaks at 41 during the Standby period and is never less than 3 during the whole 92 days

Moving the alarm limits to the boundaries of where the process has operated safely and then further to remove known past undesirable operation (see, for an example, the reduced upper limits on the 3rd and 4th variables from the left of Figure 11 which eliminate some known process upsets that one would want annunciated if they re-occurred in the future) is the starting point for the Alarm Review meeting where further refinement will take place. One set of alarm limits will be used covering all three Modes of operation of this multi-Mode process.

The new 'Lumped-Mode' Alarm Limits of Figure 11 give the immediate improvement that can be seen by comparing Figure 12 with Figure 9 and Figure 10. The hourly annunciation rate peaks at 5 instead of 22 and the standing alarm count has one peak at 11 instead of at 22 with other infrequent peaks that are rarely greater than 2 and at zero otherwise compared to the 'never less than 3' of the past.

The Lumped-Modes Limits will be further improved during the Alarm Review, perhaps in this case by removing many of the alarm limits altogether, and the much better operating environment that results will build confidence and a realisation that the alarm system can after all be improved to assist operations.

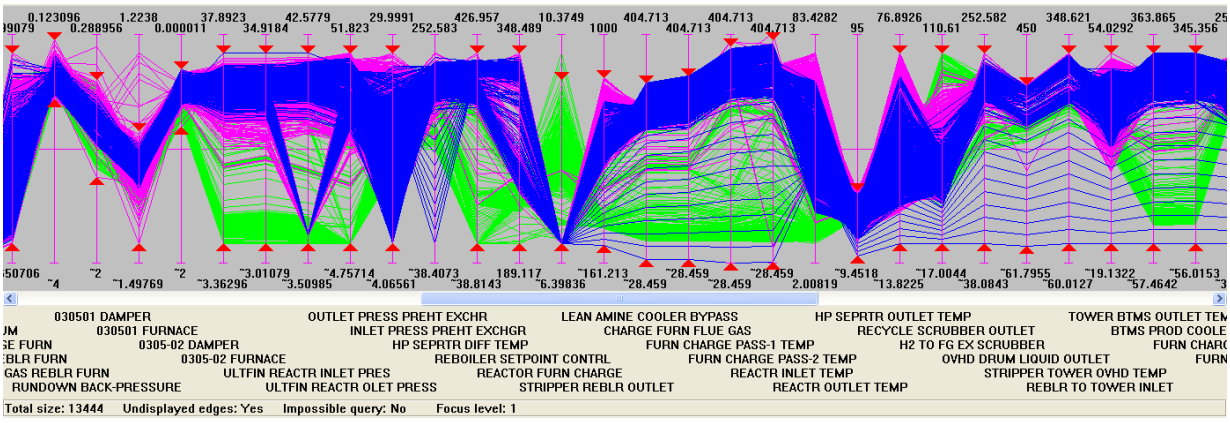

Figure 11 Kerosene Mode in pink, Gas Oil Mode in blue and Standby Mode in green. One set of alarm limits (the red triangles) is shown at the boundary of where the plant has operated and will be used for all three Modes. This is 'Lumped Mode' Alarming and is how most plants operate today.

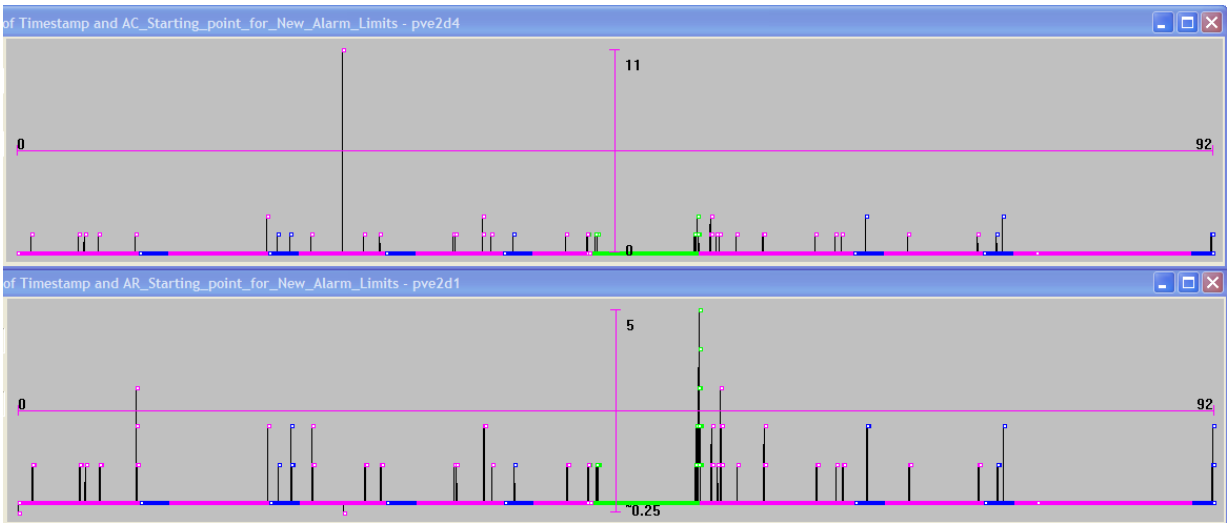


Figure 12 Annunciation Rate per hour and Standing Alarm Count with the 'lumped-mode' alarm limits of Figure 11. The colors on the time-axis show when each Operating Mode was in use

If we superimposed the existing alarm limits from Figure 8 upon Figure 11 we would see that some of the existing alarms coincide closely with boundaries of the Kero or LGO Modes. Perhaps the operators have been mentally filtering out all but a few alarms depending upon the Mode that they are in.

The next level of improvement is to define a separate set of alarm limits for each Mode at the limits of where the plant has operated in that Mode. These values can then be used as the starting point for the alarm review process as before Figure 13 shows the hourly annunciation rate and standing alarm count for Kerosene Mode. The improvement over Figure 12 is clearly visible.

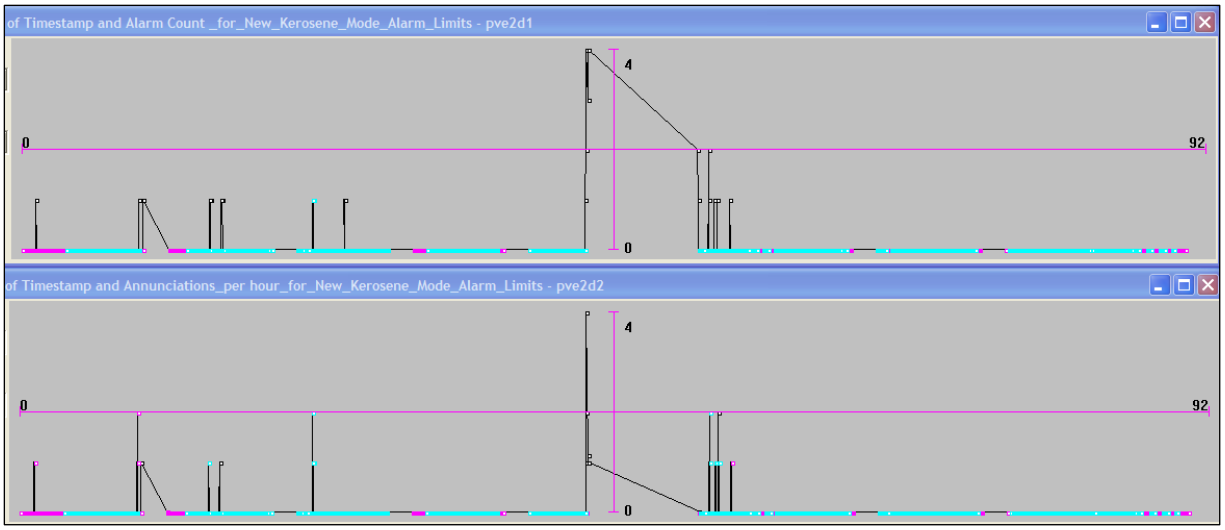


Figure 13 Annunciation Rate and Standing Alarm Count when in Kerosene Mode with Mode-based alarm limits

Alarm monitoring and annunciation would still be performed by the DCS with the additional facility required to switch between (or download) the appropriate set of alarm limits when the operating Mode changes. It can be seen in Figure 11 that ranges of values of variables used by each Mode often have considerable overlap. This will make the construction of an automatic State Detector difficult so it is probably better, at least initially, to have the Operator select the Mode he wishes to operate in.

Figure 14 shows in pink the Kerosene Mode only operations and alarm limits from Figure 11 with, in turquoise, the Operating Limits derived from the subsequent lab analyses when the Kerosene was in specification. The obvious question is why should the Alarm Limits be outside the Operating Limits? The definition of 'Normal' in Figure 1 implies, at the least, making product that is saleable and hence in specification. The conclusion is that Alarm Limits and Operating Limits are and should be two names for the same thing and that wherever pink is visible in Figure 11 or Figure 14 is bad or abnormal operation that should be eliminated with better operation, better process control and better process understanding.

Figure 15 shows what would happen if the Operating Limits of In-Spec Kerosene in Figure 14 were used as alarm limits today with no change in operation. It emphasises that too small an operating envelope for the current operating and control capability will cause violations of the EEMUA and SP18 guidelines and lead to rejection by operations. Migration to the in-spec kerosene Operating Envelope must be an incremental process following improvements in operational and process control capability.

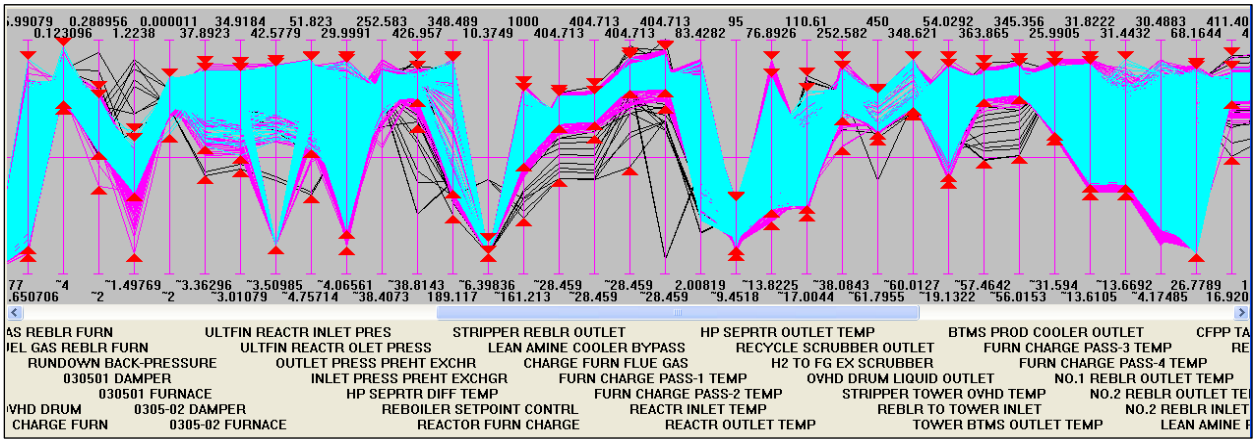

Figure 14 In-Specification Kerosene in turquoise on top of the Starting Alarm Limits for Kerosene Mode

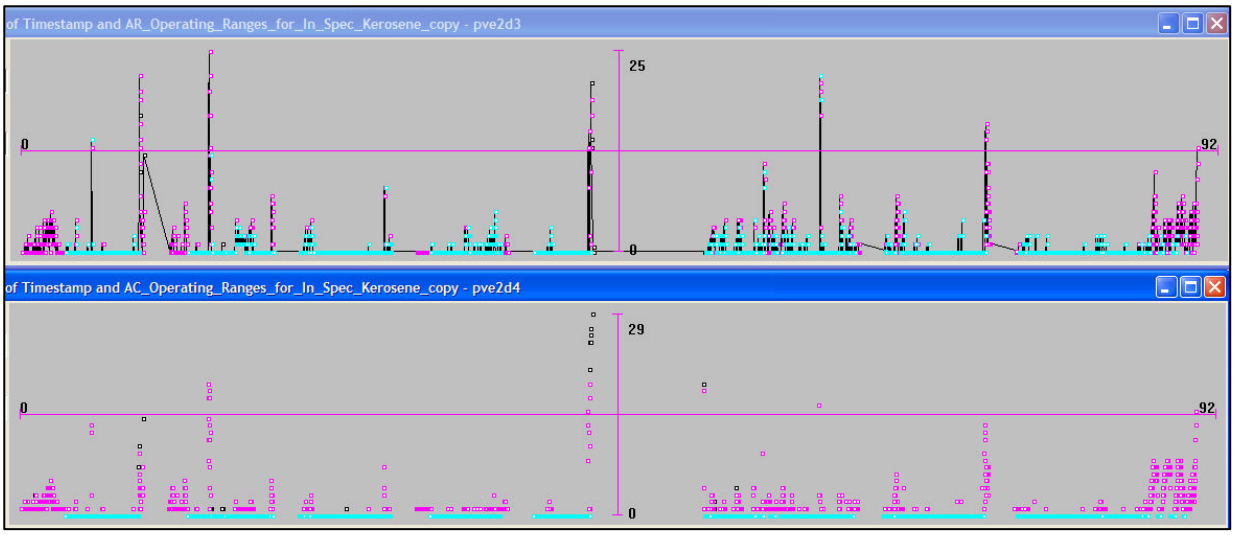

Figure 15 Annunciation Rate per hour and Standing Alarm Count for In-Spec Kerosene Operating Limits

The result in Figure 15 is sufficiently good to indicate an achievable objective. The question to ask repeatedly until the whole site becomes involved in answering it is 'why do we operate outside of our in-spec product Operating Envelope? The answer will be to use Figure 14 as a guide to explaining why pink areas are present while steadily improving operations and/or process control until it is practical to operate there all of the time and the alarm situation for the operator looks no worse than, for instance, that in Figure 13. Why hasn't this been done already? Probably because no one could 'see' the in-spec Operating Envelope so process control improvements were applied without being able to 'see' where improvement was really required.

It is a fairly radical concept to set the HiLo alarm limits at the boundary of the economic operating envelope primarily because process control, economic objectives and process alarms have always been treated as separate topics with only the process operator being concerned with all three. The Operating Envelope is actually the missing unifying root of all three.

Mallinckrodt Chemicals, UK demonstrated that alarm limits can be used to add value to process operation. In their para-aminophenol (PAP) process the alarm system would trip the pumps on some alarm limit violations and bring the process to a graceful halt. They had their alarm limits set at the boundaries of where the plant had previously operated and then tightened them by a small percentage. The operators soon learned to operate within the new tighter limits, encouraged, some said, by the desire to avoid going out into the cold and rain to the field-mounted pump starters. Figure 16 shows two years of operation on one of many process variables to the left of the blue line and two years of operation to its right after this tightening of limits with the horizontal red line indicating the tightened alarm limit and trip incidents shown in yellow. The incidence of trips was reduced from 2% to 0.1% and the gain in production time and value were easily calculated.

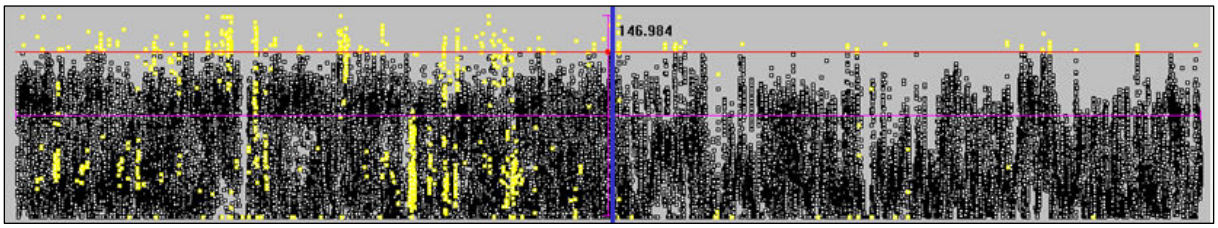

Figure 16 Four years of operation showing the better operation and reduced number of trips (yellow) in the second two years after alarm limits were brought inside the limits of operation of the previous two years

Being able to isolate Modes of operation also allows the actual achievement while in that Mode to be examined and the causes of non-achievement identified. Immediate improvement is obtained by re-setting operating limits/alarm limits to be consistent with the economic objectives. This also provides a way of, first, identifying the variables where process control most needs improvement and, second, continuously tracking improvement progress.

But, delineating the operating envelope with fixed ranges of values on individual variables ignores the richness of variable interactions that occurs in all processes and is geometrically equivalent to constructing a hypercube that encloses the used part of a variables operating envelope as illustrated in the simple 3-variable example in Figure 17 where it is apparent that fixed values for operating limits / alarm limits don't adequately describe the shape of the operating envelope. But, for everyday use they are simple and widely used and ensuring they form a consistent hypercube is the first step in improvement.

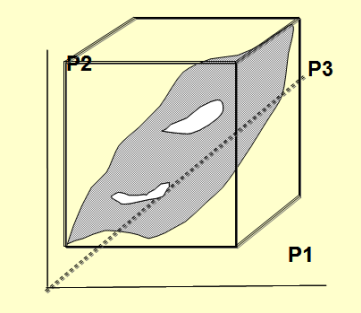

Figure 17 3-variable hypercube enclosing the Operating Envelope

The second step is to model the shape of the operating envelope itself by finding visually the cloud of multi-dimensional points where the desired objective has previously been achieved and then wrapping the cloud in a skin to obtain the operating envelope. This is much easier than it sounds requiring no further effort since, for instance, the cloud of blue Gas Oil-Mode points in Figure 11 has already been found as the cloud of blue points inside the fixed alarm/operating limits. We just take the blue points and then use a wrapping algorithm to find the envelope. The result has proven to be most effectively shown as a real-time 'you are here' display such as that in Figure 18. Here the fixed Alarm/Operating limit values are on the horizontal grey lines at the top and bottom and the black dots indicate the value now of each process variable. These black dots are collectively the current process operating point. The green lines indicate the space available around the current operating point when all variable interactions are taken into account. These green lines move at every time-step as the process operating point changes.

Violations of the green space are multi-variable excursions outside of the Operating Envelope and are called 'Public Alerts'. They are distinguished from Alarms because (a) their limit values are not fixed (b) they are not included in the Change Management requirements that are normally mandatory for fixed Alarm Limits.

'Alerts' are already provided by some DCS systems and are used by Operators for their own individual purposes such as setting a reference value of a variable such that they can compare with the reference some time later to see whether it has been reached or passed. Used this way they are specific to one operator but very valuable to him so we propose they should be re-named 'Private Alerts' to distinguish them from the 'Public Alerts' that we have just introduced to you.

The geometric basis for the calculation of the green lines is remaining interior to the cloud of points so that should the process stray outside the green space it is possible to calculate, using geometry, the smallest distance to move the manipulable process variables in order to bring the operating point back inside the Operating Envelope. This gives the operator intrinsically safe advice to correct the process problem and avoid a violation of the fixed alarm/operating limits. Once the Operating Point is inside the Operating Envelope the primary operating objective is being achieved and the task of process control is just to keep it there since the primary objective is achievable equally well anywhere inside the envelope. If there are secondary objectives there will be sub-spaces inside the primary operating envelope where they are also achieved so 'optimisation' consists of additional process moves inside the primary envelope until the process is also inside a secondary envelope when movement will again cease. This contrasts with traditional process control which attempts to position the process at a point whose coordinates are the values of all the setpoints. Our models expect process dynamics to be handled by the process control system so depend upon there being at least a regulatory PID level of control present. One model can handle multiple Modes of operation by including the Mode number as a variable in the model. For a time-varying process a time variable is included and for a multi-Phase Batch process one includes both the Phase number and time from the start of each Phase. For automating a Start-Up one would proceed similarly as for a Batch Process by dividing the procedure

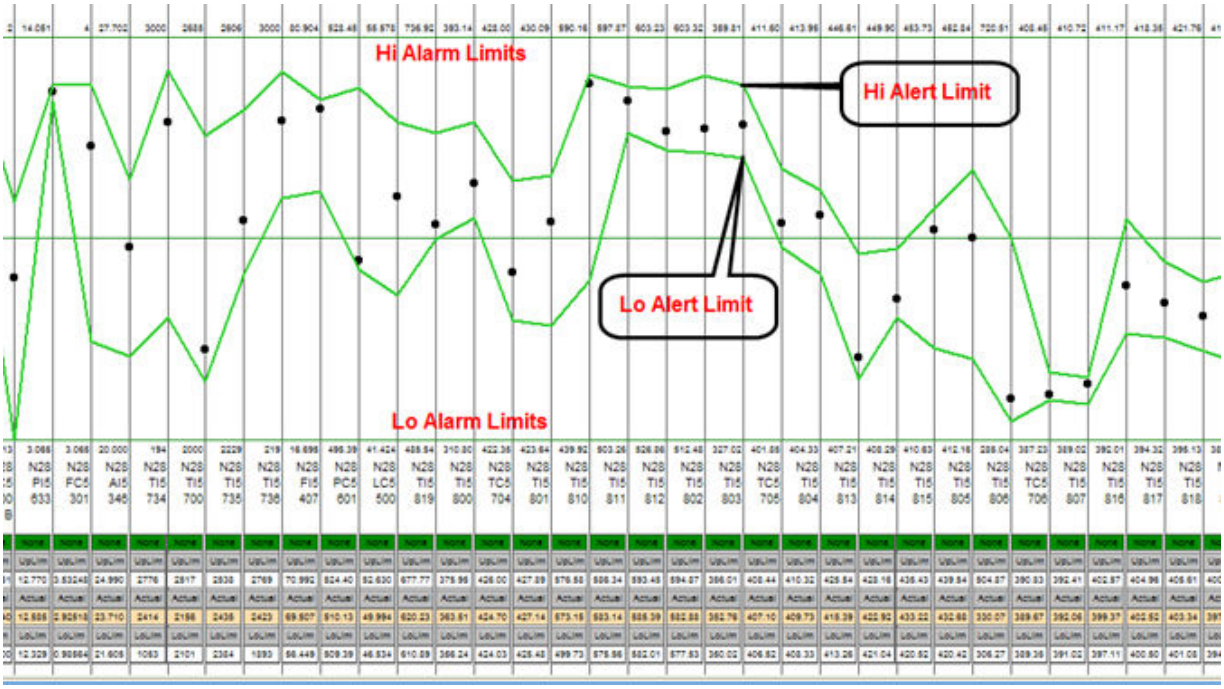


into its Phases and including the phase number and phase time as variables.

Figure 18 Public Alerts have the objective of keeping the process inside its fixed Alarm/Operating Limits

So, starting from process history data instead of alarm log data and using a wholly graphical method we have shown how fixed alarm limits and operating limits are first approximations to Operating Envelopes and should be combined and can be improved with little or no change to existing methods of working. We have shown how a Multi-Mode process (and all processes have at least two Modes viz. Operating and Shutdown) can be treated as a Lumped-Mode process with one set of alarm limits as is usually the situation today and how it can easily be separated into its Modes and separate sets of alarm limits found and implemented for each Mode.

We have also shown how to proceed beyond the limitations of fixed limits with little additional effort to a new dynamic method of operator guidance allowing operation even as tight as the capabilities of modern process control systems will allow. And by showing that Alarm Limits and Operating Limits are, or should be, the same we can use the same well-developed methods of calculating value from the reduction of excursions outside operating limits for calculating value from alarm limits, thus giving an economic Rationale to Alarm Rationalization and to Alarms as a whole.

---

[i] A New Method for Defining and Managing Process Alarms and for Correcting Process Operation when an Alarm Occurs. Brooks RW, Thorpe RJ, Wilson JW. Journal of Hazardous Materials 115(2004) 169-174

[ii] Alarm Systems. A Guide to Design, Management and Procurement. EEMUA Publication No. 191: 1999 London. ISBN 086931 076 0 www.eemua.co.uk

[iii] ISA SP18.02 Management of Alarm Systems for the Process Industries

[iv] Bernhard Riemann's inaugural lecture Nature, Vol. VIII. Nos. 183, 184, pp. 14--17, 36, 37

[v] A.Inselberg, Parallel Coordinates, DOI 10, 1007/978-0-387-68628-8_5, Springer Science+Business Media 2009