

The Rationale for Alarm Rationalization

Robin W. Brooks, PhD, Alan Mahoney, PhD, John Wilson, PhD, CEng, Na Zhao, PhD.
Curvaceous Software Limited, PO Box 43, Gerrards Cross, Bucks. SL9 8UX, UK

Correspondence to Robin_Brooks@curvaceous.com

Abstract

Quantifying the economic value of an Alarm System, or even the value of rationalising it, has rarely been attempted. Alarm Systems are in that category of things imposed upon a plant either by legislation or by the fear of litigation and backed by bodies such as OSHA and HSE or, in the case of operator alarms, that come over-enthusiastically configured as part of the DCS along with the built-in need for a later rationalisation project to make them usable. Few, if any, plants actually know the value, as opposed to the cost, of their alarm systems hence they can not justify and do not see a need to initiate projects involving additional expenditure on, for instance, alarm rationalisation or on-going continuous improvement of the alarm systems. Requirements for such projects come from outside the plant either compelled by legislation or coerced by the need to be seen to have adhered to Best Practice in case litigation should follow a plant incident. In the absence of a value-understanding, project justification degrades to the need for just-enough compliance with the legal or Best Practice requirement and may be seen as an imposition deserving of less than their full enthusiasm by those most closely involved with the plant and who are already fully and gainfully occupied in meeting production and up-time targets. Unfortunately these are the very people most essential to any rationalisation or improvement of an Alarm System.

In this paper we show that Alarm Limits and Operating Limits should be the same and that this will allow Alarm Limitsⁱ and their rationalisation to benefit from the well-developed economic understanding already in existence for Operating Limits.

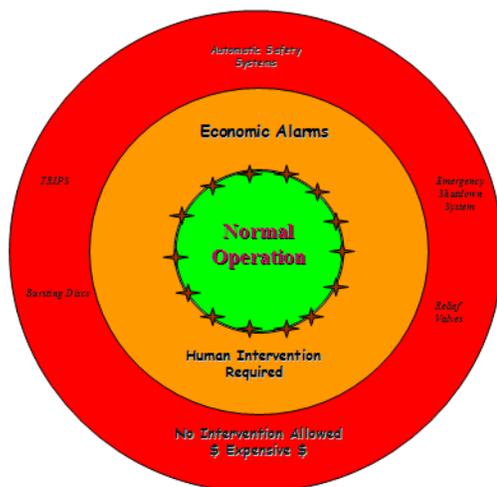


Figure 1 The Vision - Alarm Limits defining the Operating Envelope

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 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 intended to draw the process operator's attention to a situation beyond the capability of the process control system to prevent and requiring 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 these alarms give the operator time to intervene and correct the situation so that they also have an 'insurance premium' value in reducing the demand upon the safety system and thus the small possibility 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 (a) alarm limits are ideally the same as operating limits and (b) the economic cost of violating an alarm limit is the delta cost between the material produced and operating costs of desired and undesired operation.

The 'Envelope' noun has been used by generations of chemical engineers to describe a closed boundary with different properties of something inside and outside the boundary. In two dimensions, such as on a sheet of paper or a computer display, the boundary is a line that separates two areas. In three dimensions the boundary is a surface separating two volumes. So if we are using variables as dimensions and have more than three variables what then is the boundary and what is it separating? The reality has been that a schematic such as Figure 1 was as close as one could get to describing the nebulous concept of an operating envelope involving more than three variables.

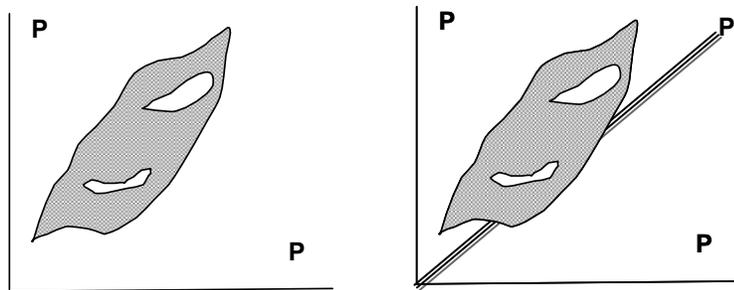


Figure 2 A two-dimensional area; a three-dimensional volume; how can one draw the fourth axis?

coordinate transformation in the 1980's. Instead of trying to draw the axes orthogonally he drew them parallel to each other causing the representation of a point to transform to a poly-line as in Figure 3.

The problem is one of n-dimensional geometry which was fully described by Riemannⁱⁱ in 1853 using equations too complex to be solved except in the simplest cases and without pictures. The problem shown in of how to represent the fourth axis remained as an obstacle to understanding of higher-dimensionality geometry until Inselbergⁱⁱⁱ discovered the parallel

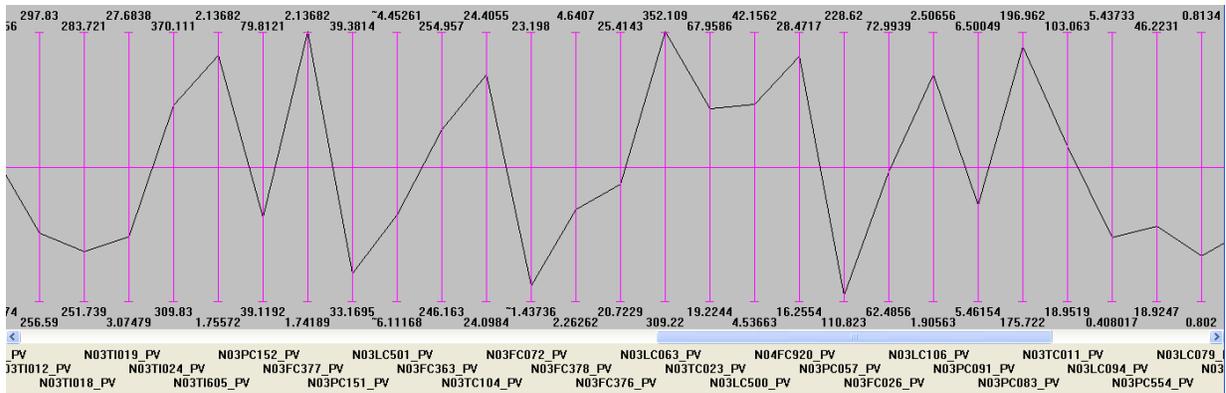


Figure 3 The parallel coordinate representation of a point in 27 dimensions

Adding more points to the graph produces distinctive patterns as in Figure 4, which is the purpose of a graph, 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 gathered by the process historian during three months of unit operation.

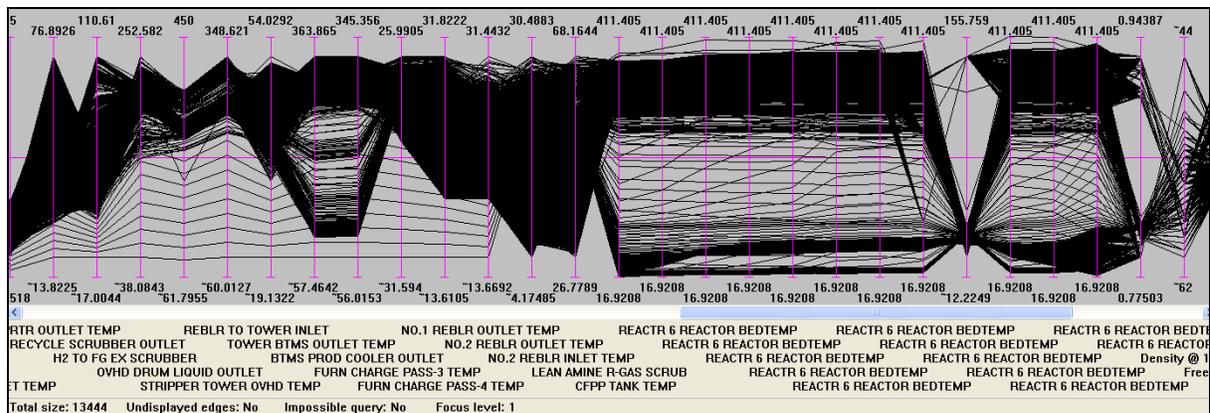


Figure 4 Part of the operating data for an HDS Unit during three months of operation

This plant operates at different times in one of three main ‘Modes’ of Standby, Kerosene desulphurization and Light Gas Oil (LGO) desulphurization. This largely accounts for the bands that are such a prominent visual feature. Like most plants, they have one set of alarm limits to cover all three Modes. When these are superimposed on the graph in Figure 5 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 the alarm system is bad in that the alarm annunciations per hour (Figure 6) and Standing Alarm Count (Figure 7) are both high. 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.

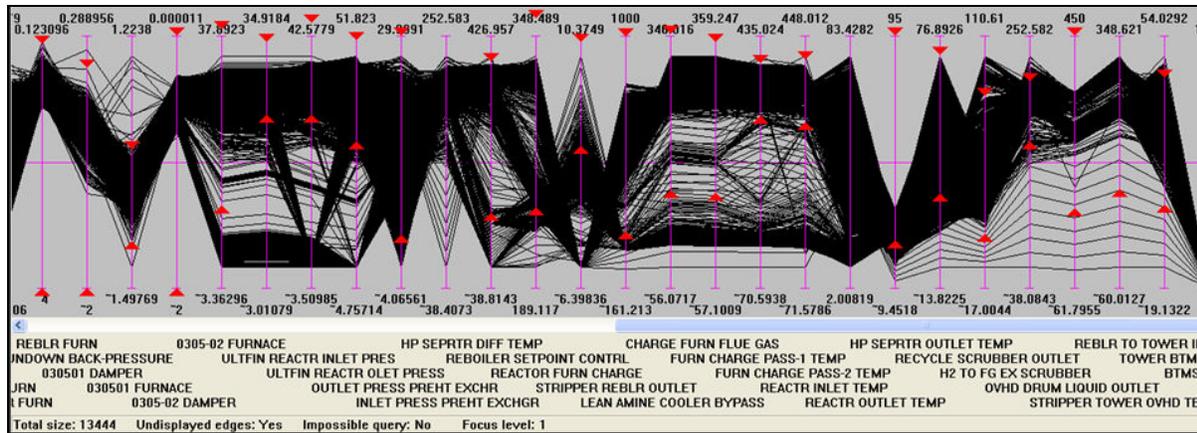


Figure 5 Existing HiLo alarm limits superimposed upon three months of operating data

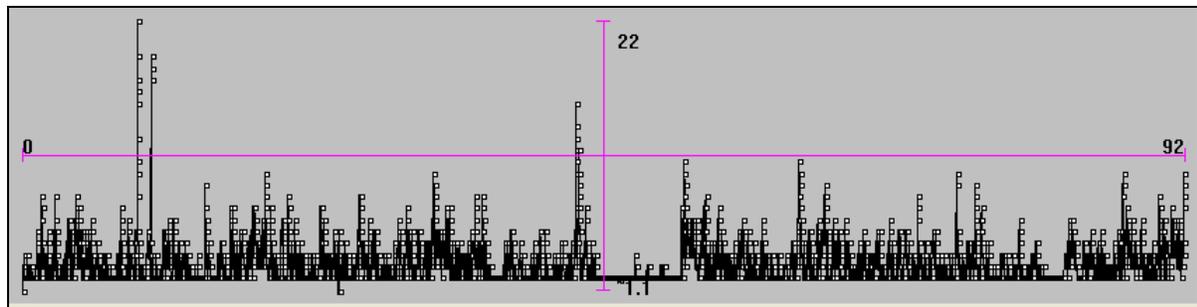


Figure 6 Annunciations per hour peak at 22 during this 92 day period

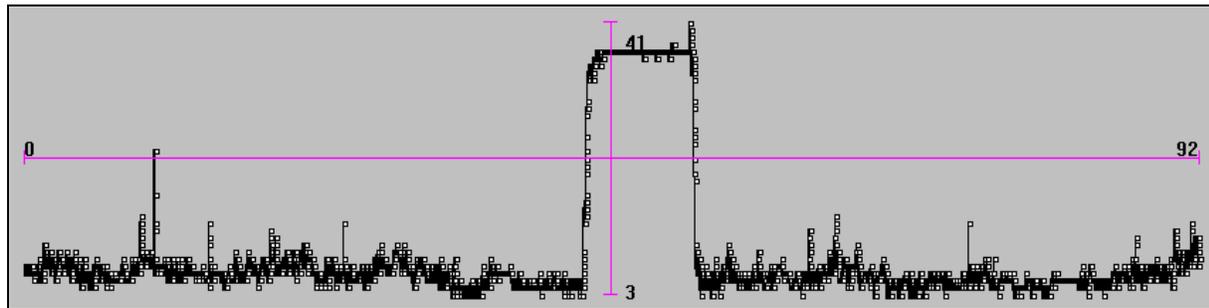


Figure 7 The count of Standing Alarms peaks at 41 during the Standby period and is never less than 3 during the whole 92 days

Moving the alarm limits to the boundaries of where the process has operated safely will cure the problems of the alarm system and most likely allow conformance to the EMUA^{iv}/ISA SP18^v guidelines but will not, in this case, assist the operator in achieving operating objectives (and thus allow the alarm system to demonstrate an economic value) unless we first separate this process into its three Modes of operation. This has been done in Figure 8. With the existing alarm limits superimposed it can be seen that some of them coincide with extremes of the pink Kerosene Mode band and others with the blue LGO Mode band. Perhaps the

operators have been mentally filtering out all but a few alarms depending upon the Mode that they are in.

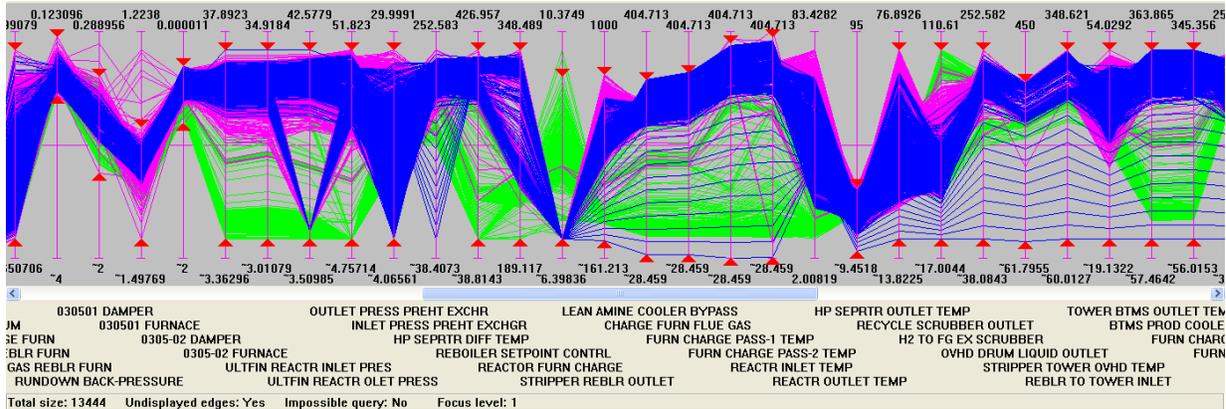


Figure 8 Kerosene Mode is in pink, Gas Oil Mode in blue and Standby Mode in green. One set of alarm limits (the red triangles) set at the boundary of where the plant has actually operated will be used for all three Modes. This is ‘Lumped Mode’ Alarming and is how many plants operate today.

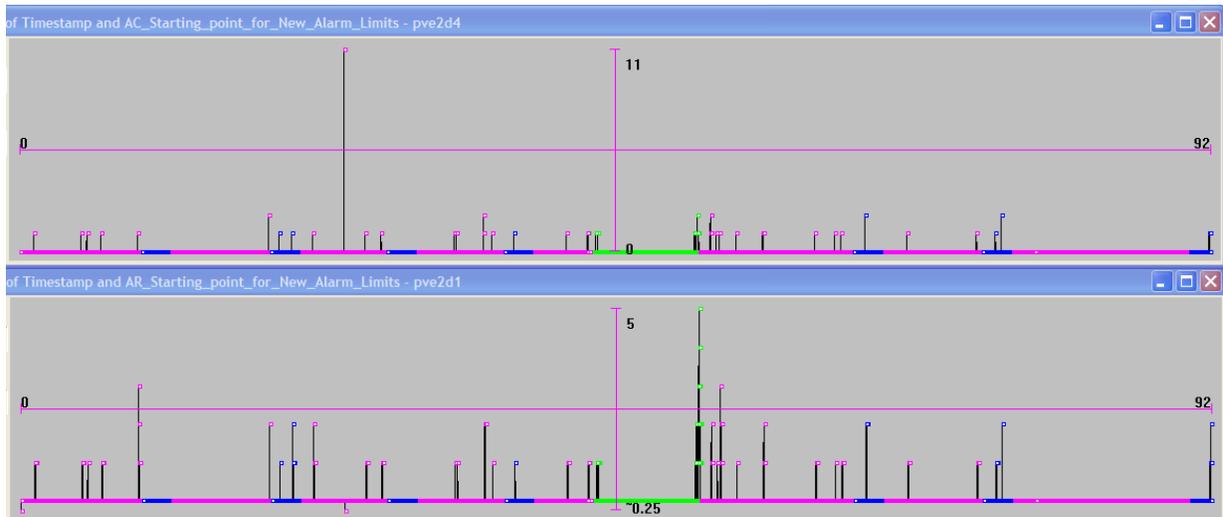


Figure 9 Annunciation Rate per hour and Standing Alarm Count with the 'lumped-mode' alarm limits of Figure 8

The new ‘Lumped-Mode’ Alarm Limits of Figure 8 give the immediate improvement that can be seen by comparing Figure 9 with Figure 6 and Figure 7. The hourly annunciation rate peaks at 5 instead of 22 and the standing alarm count has one peak at 11 instead of 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 give confidence and a realisation that the alarm system can be improved to assist operations.

The next level of improvement is to separate the process Modes and define a set of alarm limits for each Mode separately 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 10 shows the hourly annunciation rate and standing alarm count for Kerosene Mode. The improvement over Figure 9 is clearly visible.

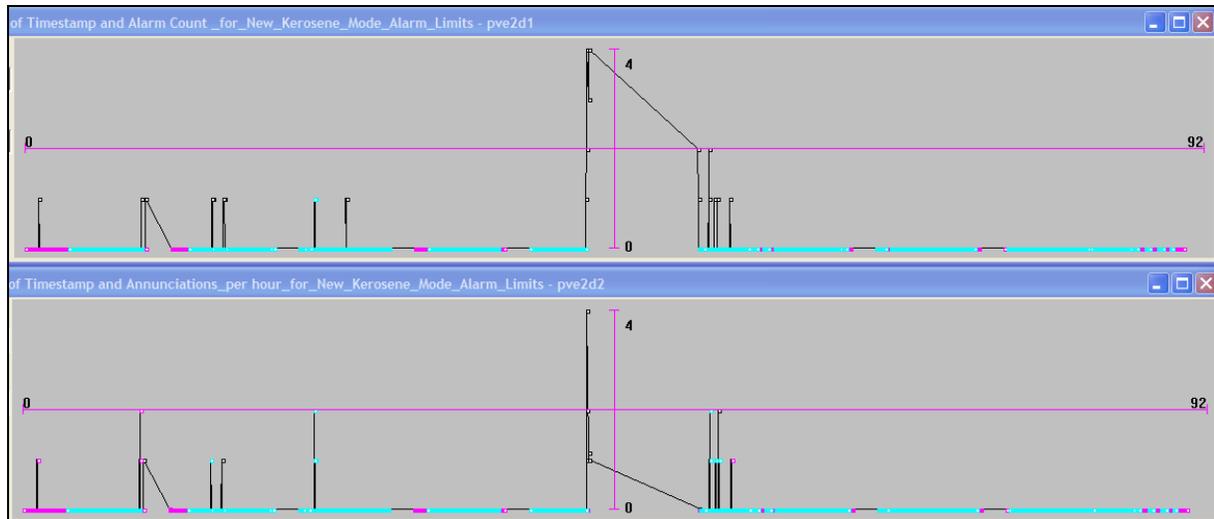


Figure 10 Annunciation Rate and Standing Alarm Count when in Kerosene Mode with Mode-based alarm limits

Alarm monitoring and annunciation will still be performed by the DCS with the addition of a facility to switch between (or download) the appropriate set of alarm limits when the operating Mode changes. It can be seen in Figure 8 that ranges of values of variables used by each Mode often have considerable overlap which 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 11 shows in pink the Kerosene Mode only operations and alarm limits from Figure 8 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 8 or Figure 11 is bad or abnormal operation that should be eliminated with better operation, better process control and better process understanding.

Figure 12 shows what would happen if the Operating Limits of In-Spec Kerosene in Figure 11 were used as alarm limits today with no change in operation.

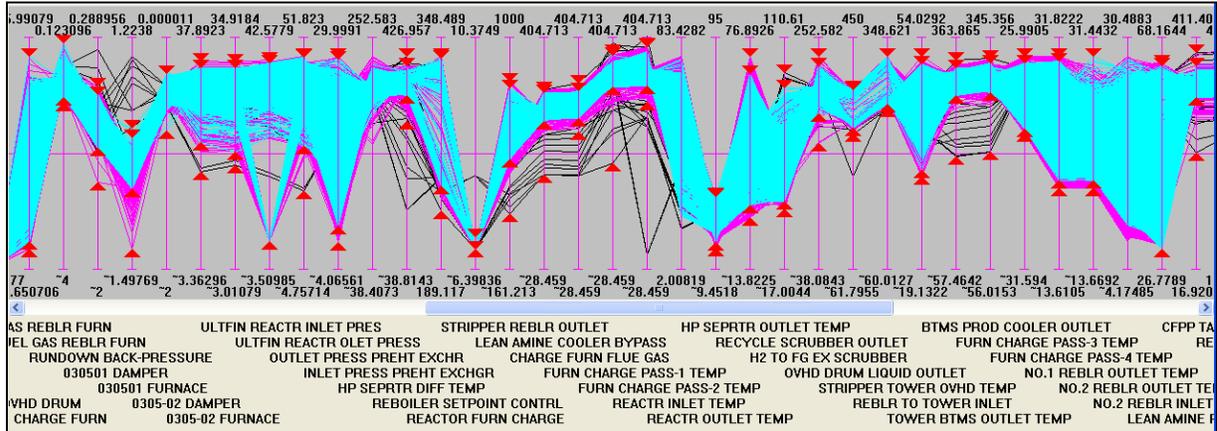


Figure 11 In-Specification Kerosene in turquoise on top of the Starting Alarm Limits for Kerosene Mode

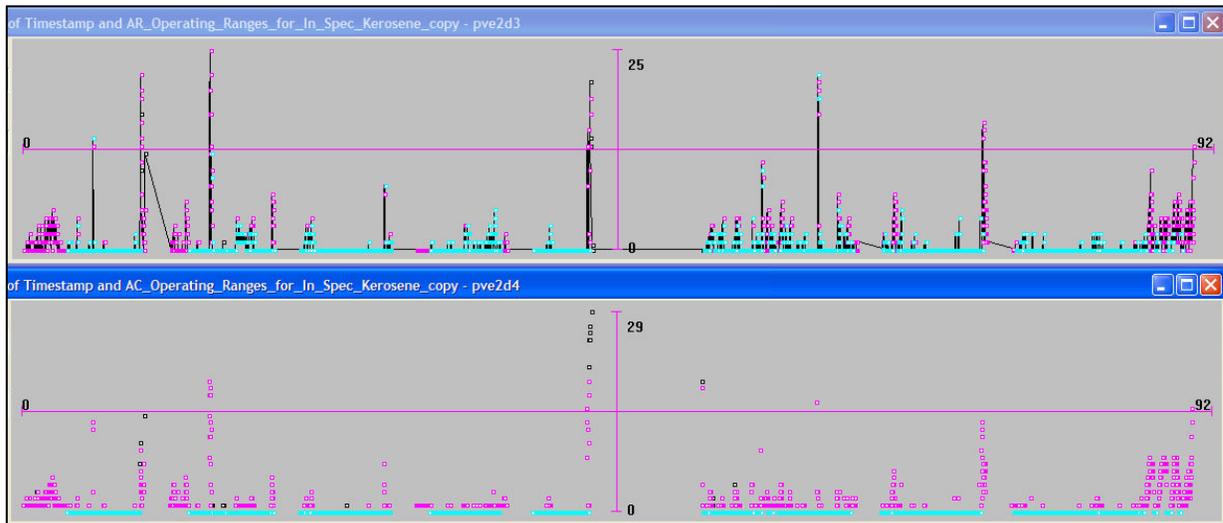


Figure 12 Annunciation Rate per hour and Standing Alarm Count for In-Spec Kerosene Operating Limits

The result in Figure 12 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 envelope? The answer will be to use Figure 11 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 10.

It is a fairly radical concept to set the HiLo alarm limits at the boundary of the economic operating objectives 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. 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 13 shows two years of operation before and two years of operation 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 thus value easily calculated.

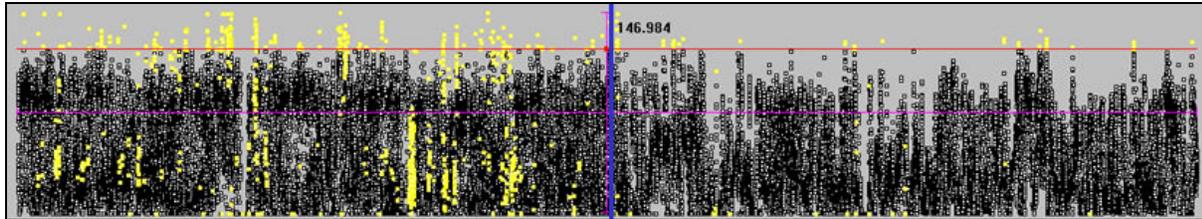


Figure 13 Four years of operation showing the better operation and reduced number of trips (yellow) in the second two years after alarm limits were brought inside the limits of operation of the previous two years

Being able to isolate Modes of operation also allows the actual achievement while in that Mode to be examined and causes for 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.

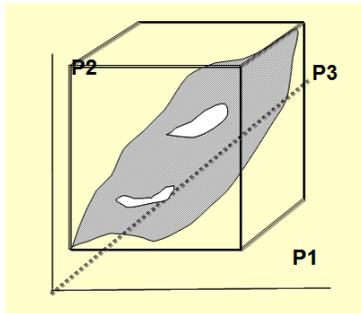


Figure 14 3-variable hypercube enclosing the Operating Envelope

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 14 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. Making them consistent so that they form a hypercube is the first step in improvement.

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 8 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 15. 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' exist already in some plants and are used by Operators for their own individual purposes such as setting a reference value such that they can see some time later whether it was 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. One model can handle multiple Modes of operation by including the Mode number as a variable in the model.

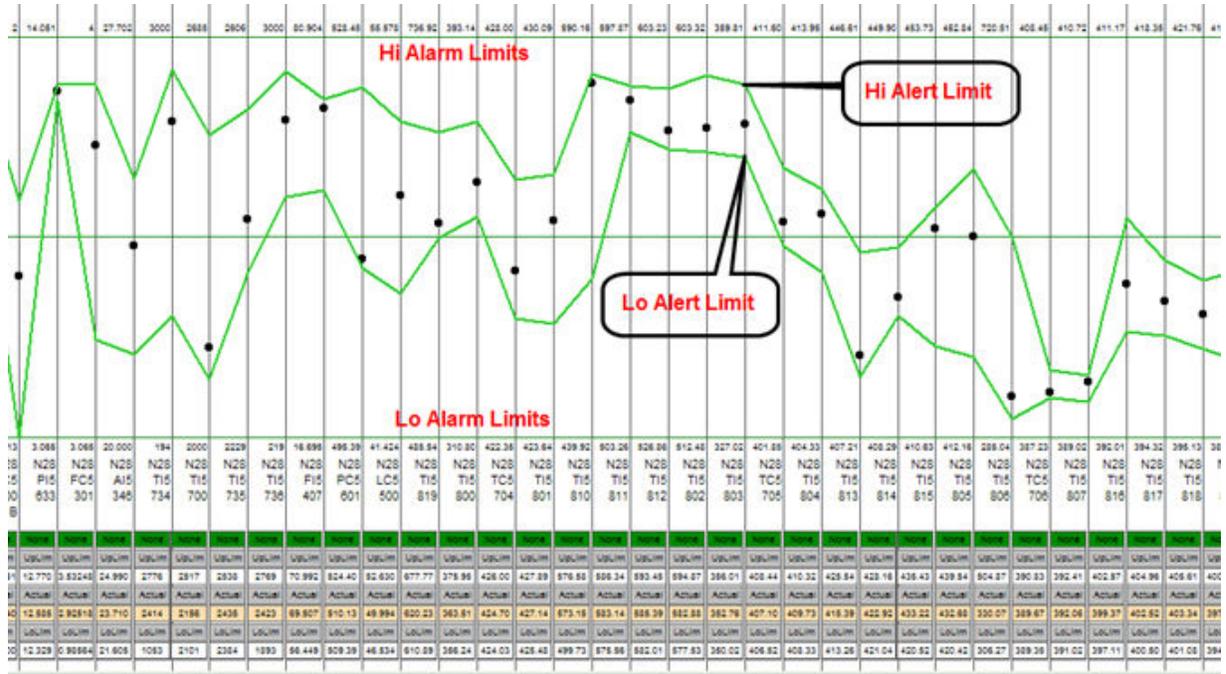


Figure 15 Public Alerts have the objective of keeping the process inside its fixed Alarm/Operating Limits

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 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
Journal of Hazardous Materials 115(2004) 169-174

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

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

^{iv} Alarm Systems. A Guide to Design, Management and Procurement. EEMUA Publication No. 191: 1999 London. ISBN 086931 076 0 www.eemua.co.uk

^v [ISA SP18.02 Management of Alarm Systems for the Process Industries](#)