NEW ALARMS AND ALERTS FROM OPERATING ENVELOPES DRIVE ECONOMIC BENEFITS AS WELL AS SAFER PROCESSES

Robin W. Brooks, Alan Mahoney, John Wilson, and Na Zhao
Process Plant Computing Limited (PPCL), Gerrards Cross, Bucks, UK

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.

The root cause has been lack of a fundamental understanding of alarms exemplified by the fact that there has never previously been a general method to find values at which to set the alarm limits, although this is where many or most of the problems of alarm systems begin and therefore could end. We show that alarms should be related to the Operating Envelope required to achieve the plant's business objectives and so provide a general and easily implemented method for finding alarm limit values. Operating Envelopes have been poorly understood, although the term is in common use, and no method for finding or using them has existed. We show how they are easily found for various business and process objectives and how they are related to alarm limits. The example used throughout this paper is a hydro-desulphurisation (HDS) unit having three distinct Modes of operation.

In this paper we show that Alarm Limits and Operating Limits are linked by the Operating Envelope. They should in principle be the same, thus allowing Alarm Limits and their rationalisation to benefit from the well-developed economic understanding already in existence for Operating Limits. That they are not the same today is probably because it was not previously possible to see, compare and work with Operating Envelopes.

Alarms in a process plant fall into one of two categories. They are either Operator Alarms (aka. Economic Alarms) which form the first line of defence against process mis-operation or mal-function or they are Safety Alarms which form the second and usually the last line of defence. This main subject of this paper is Operator Alarms.

Safety Alarm Systems are 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 they provide is in preventing a loss-of-control from turning into a disaster with liabilities and costs that can run into hundreds and even thousands of millions of dollars. The costs of a Safety Alarm System are viewed as an insurance premium against the economic consequences of a disaster that most plants will never experience.

OPERATING ENVELOPES

Operator Alarm systems are 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 shut-downs are expensive in terms of lost production and possible consequential plant damage and operator 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 while in alarm and the operating costs of desired and undesired operation.

The "Operating Envelope" noun-phrase has been used by generations of chemical engineers to describe a closed boundary with different properties of something inside and outside the boundary. It is obviously multi-variable or multi-dimensional but that meant that we couldn't draw a picture of it.

Figure 1. The Vision – Alarm limits defining the Operating Envelope

The problem of how to see all of the process variables in one graph is one of n-dimensional geometry which was fully described by Riemann[1] in 1853 using equations too complex to be solved except in the simplest cases and without pictures. The problem shown in Figure 2 and Figure 3 of how to represent the fourth axis remained as an obstacle to understanding of higher-dimensionality geometry until Inselberg[2] discovered the parallel coordinate transformation in the 1980s. 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 5.

Adding more points to the graph produces distinctive patterns as in Figure 6, 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 an oil refinery hydrodesulphurisation (HDS) unit and is part of a graph of 178 variables at 13,444 5-minute intervals gathered by the process historian during three months of unit operation.

This plant operates at different times in one of three main operating Modes[3] 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 today, they have one set of

[1]Bernhard Riemann's inaugural lecture Nature, Vol. VIII. Nos. 183, 184, pp. 14–17, 36, 37.
[2]A. Inselberg, Parallel Coordinates, DOI 10, 1007/978-0-387-68628-8_5, Springer Science + Business Media 2009.
[3]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.
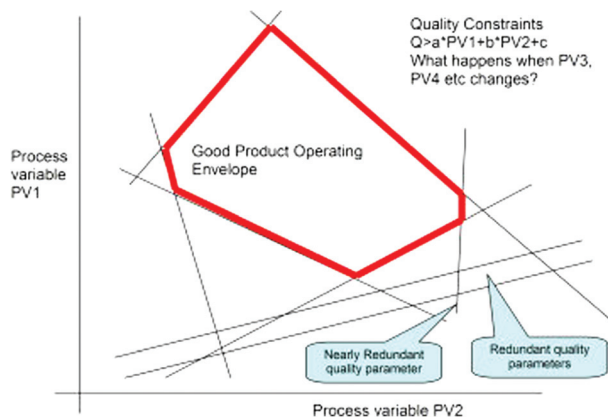

Figure 2. How we think of an Operative Envelope of two process variables and several quality constraints expressed as functions of the two process variables

alarm limits to cover all three Modes. When these are superimposed on the graph in Figure 7 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 previous uni-variate alarm rationalization terminology so would receive no attention and might escape the HazOp scrutiny of the multi-disciplinary Alarm Review Panel.

The performance of the alarm system is poor in that there are typically five–seven alarm annunciations per hour (Figure 8) and a Standing Alarm Count (Figure 9) of four–five during normal operation. During Standby Mode (low values of most variables) the alarm display showing 41 variables in alarm during standby means that any real alarm has a high probability of going unnoticed. This alarm system conforms to the EEMUA 191 and ANSI/ISA SP18 guidelines in its Human Factors performance
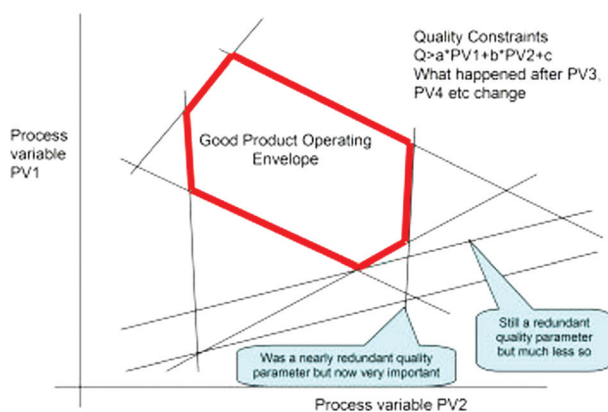

Figure 3. But because the quality constraints are also functions of several other process variables they change when those process variables change
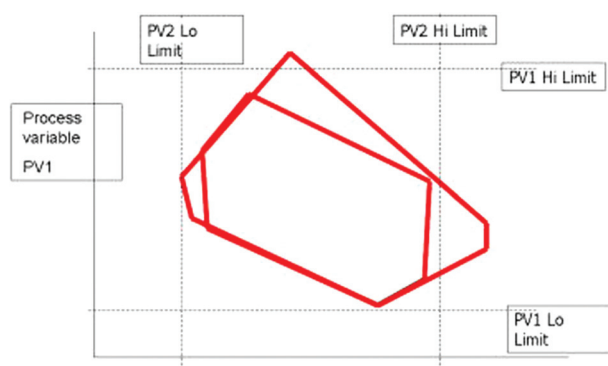
Figure 4. The operating envelope is 3-d and alarm limits form a box

but this anomaly simply emphasises that there is a missing quantitative guide measuring the "goodness" of an alarm system and/or the "quality" of its alarms in protecting the process.

Few have attempted to measure the "goodness" of their alarms. One case that we were associated with during a field trial was performed by Ineos Chlor, Runcorn, UK[4] to prove the quality of the alarms on a real process. They defined alarm quality quantitatively as the proportion of false alarms and defined an alarm as false if subsequent laboratory analysis received several hours later showed that the product had actually been good at the time of the alarm and vice versa thus effectively adopting the viewpoint of Figure 1 that alarms should delineate the boundary of Normal Operation. Ineos Chlor reset their alarm limits using the best experience and knowledge of the plants engineers and the results were compared with

results from alarm limits generated using the methods described in this and an earlier paper.[5] The objective in both cases was to produce product within specifications as measured by subsequent laboratory analyses. Alarms raised with the traditionally set alarm limits were false 49% of the time whereas those raised by the new methods for finding Alarm and Alert Limits described in this paper were false only 10% of the time. The 10% was further reduced by improved choice of variables in the envelope.

Reducing the total number of alarms annunciated by 39% also reduced the Annunciation Rate, in a similar proportion.

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. It is the Envelope of all operation during the three month period so contains process excursions and periods when process faults and problems were known to exist. Removing these from the graph would leave the Envelope of safe process operation.

There are many possible Operating Envelopes corresponding to 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 giving what most would regard as the Envelope of Normal operation. 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, which begins to take account of variable interactions compared to the strictly univariate methods of the past, and
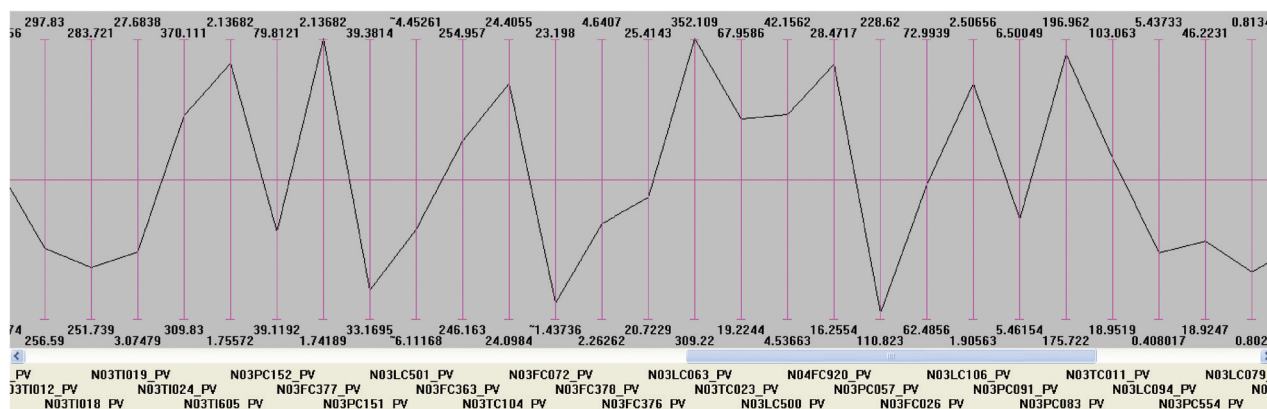

Figure 5. The parallel coordinate representation of a point in 27 dimensions

[4]D. Armstrong, M. Tyrrell, S. Casey, Ineos Chlor Ltd. R. Brooks, R. Thorpe, J. Wilson, Curvaceous Software Limited. First Experiences at Ineos Chlor Ltd. with GPC for Product Quality and Process Operations Improvement. Proceedings of the AspenWorld Conference, November 2002. Also available via www.ppcl.com.

[5]Brooks, R., Thorpe, R., and Wilson J. 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.
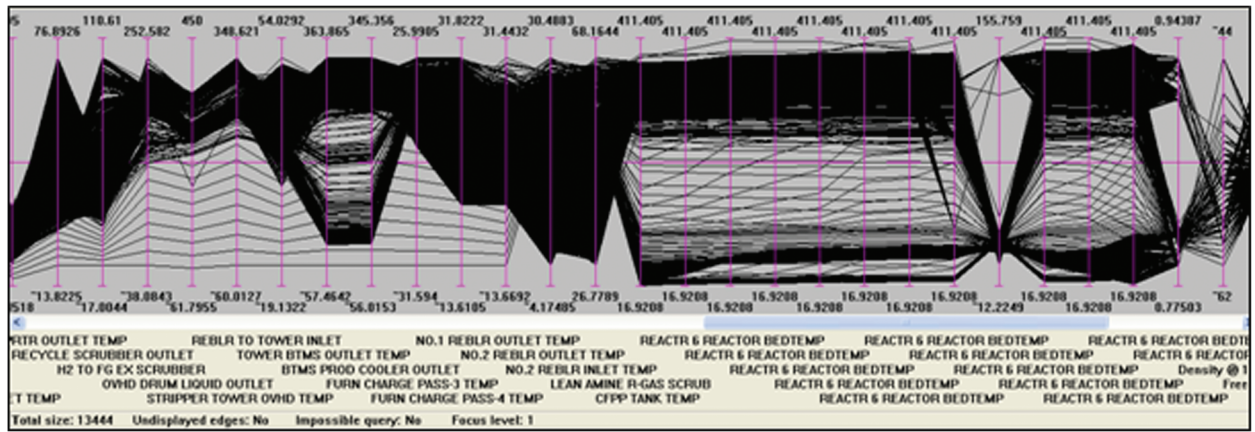
Figure 6. Part of the operating data for an HDS Unit during three months of operation
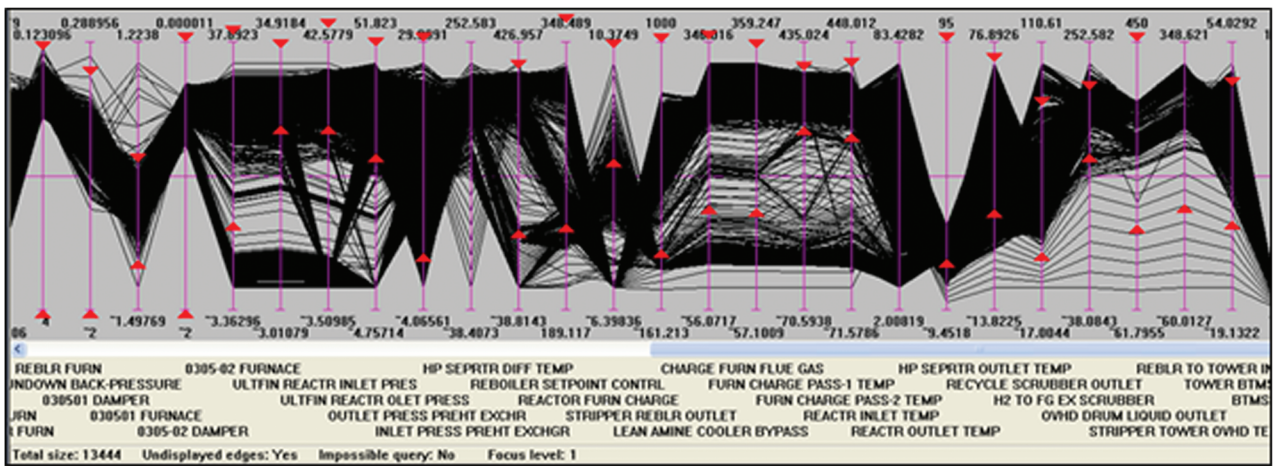

Figure 7. Existing HiLo alarm limits superimposed upon three months of operating data
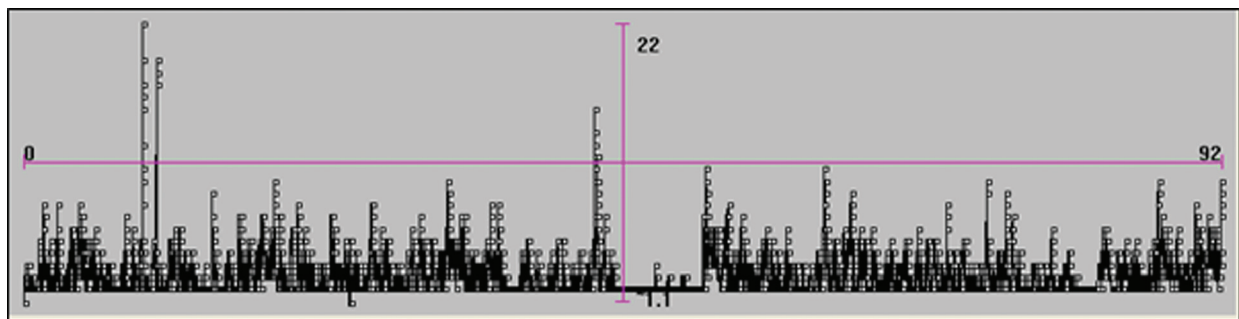

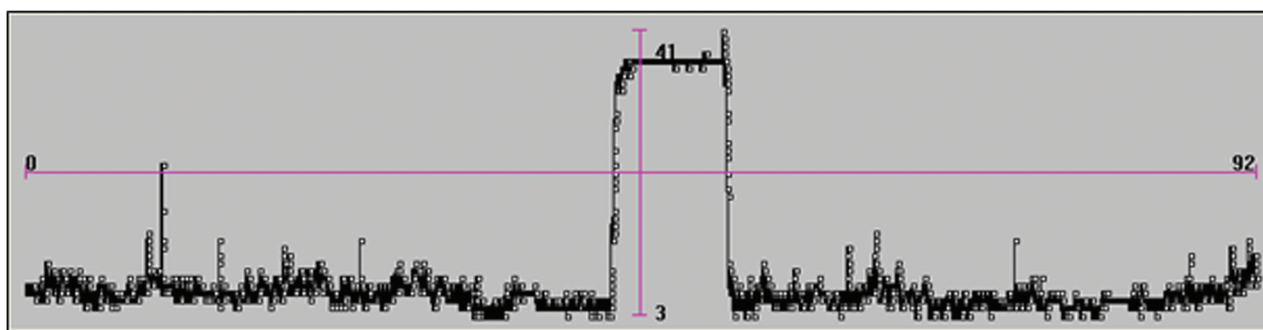**Figure 8.** Annunciations per hour peak at 22 during this 92 day period

Figure 9. The count of Standing Alarms peaks at 41 during the Standby period and is never less than 3 during the whole 92 days
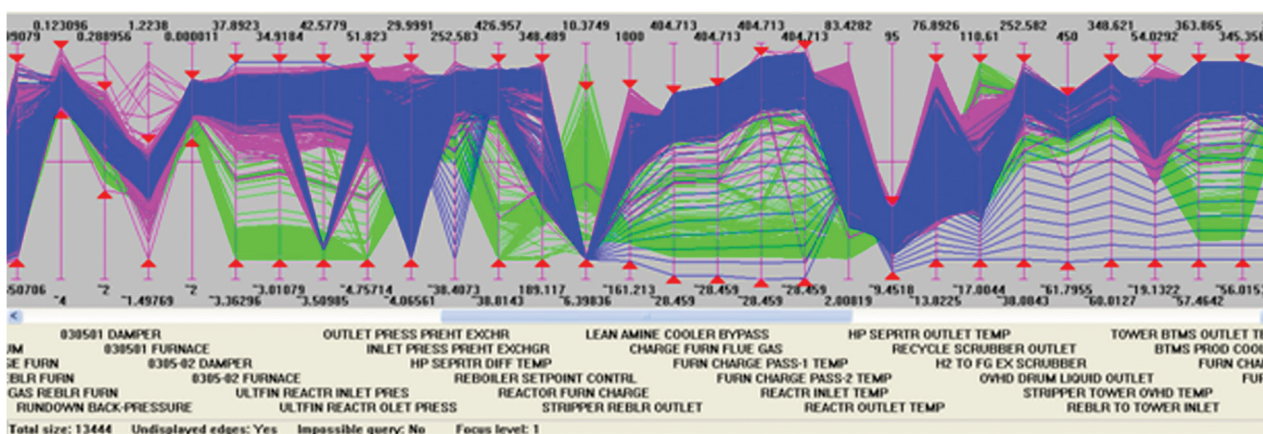

Figure 10. Kerosene Mode is in pink, Gas Oil Mode is blue and Standby Mode is green. One set of alarm limits (the red triangles) set at the boundary of where the plant has actually operated will be used for all three Modes. This is "Lumped Mode" Alarming and is how most plants operate today

later in this paper will introduce Alerts to give a much better representation of an Operating Envelope.

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[6]/ISA SP18[7] 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 10. 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.

The new "Lumped-Mode" Alarm Limits of Figure 10 give the immediate improvement that can be seen by comparing Figure 11 with Figure 8 and Figure 9. 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, which will be considerably assisted by the ability that is lacking today to confidently predict the annunciation rates and standing alarm counts that would have resulted from any set of alarm limits having been in use during the time period of the data. The much better operating environment that results will give confidence and a realisation that the alarm system can be improved to positively assist operators in achieving their operating objectives which it does not do today.

[6]Alarm Systems. A Guide to Design, Management and Procurement. EEMUA Publication No. 191: 1999 London. ISBN 086931 076 0 www.eemua.co.uk.
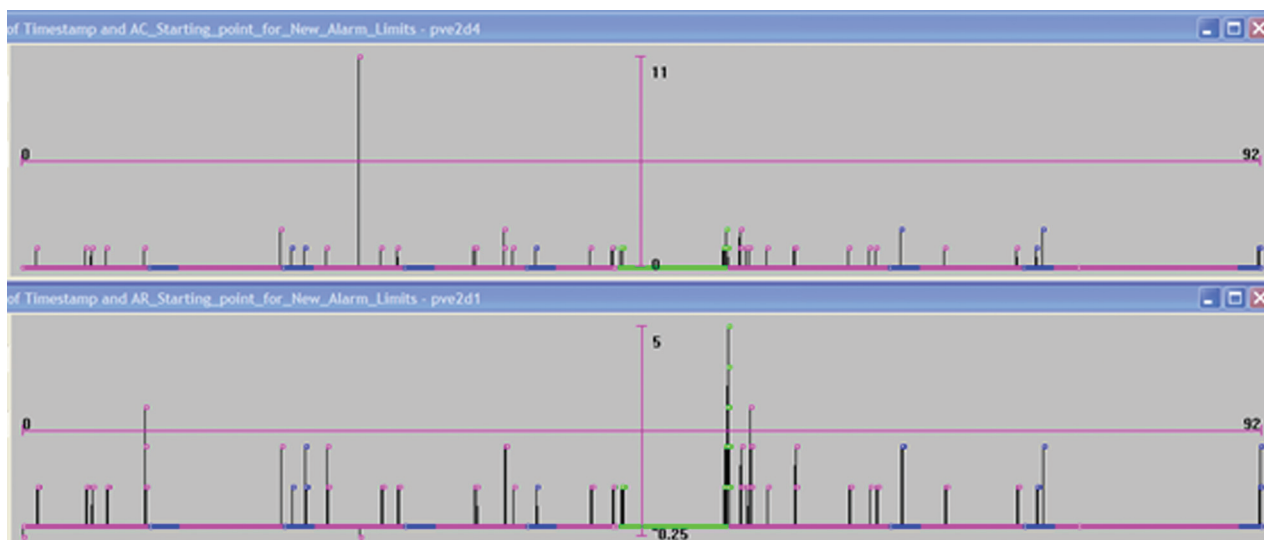[7]ISA SP18.02 Management of Alarm Systems for the Process Industries.

Figure 11. Annunciation Rate per hour and Standing Alarm Count with the "lumped-mode" alarm limits of Figure 10

ALARMS BY MODE OF OPERATION

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 12 shows the hourly annunciation rate and standing alarm count for Kerosene Mode. The improvement over Figure 11 is clearly visible.

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 10 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 13 shows in pink the Kerosene Mode only operations and alarm limits from Figure 10 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
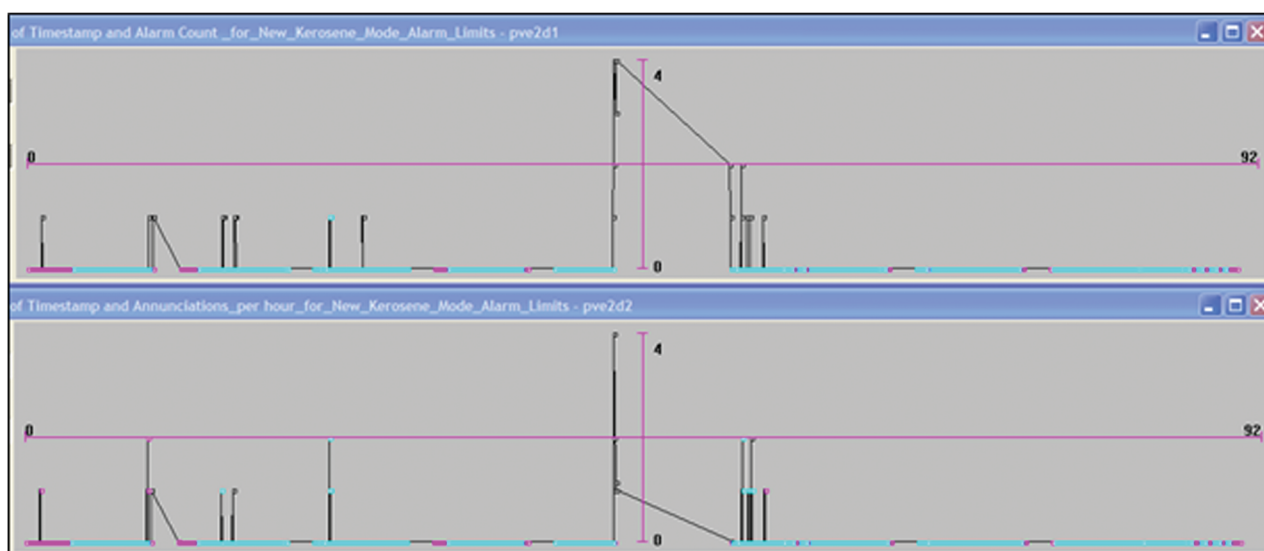

Figure 12. Annunciation Rate and Standing Alarm Count when in Kerosene Mode with Mode-based alarm limits
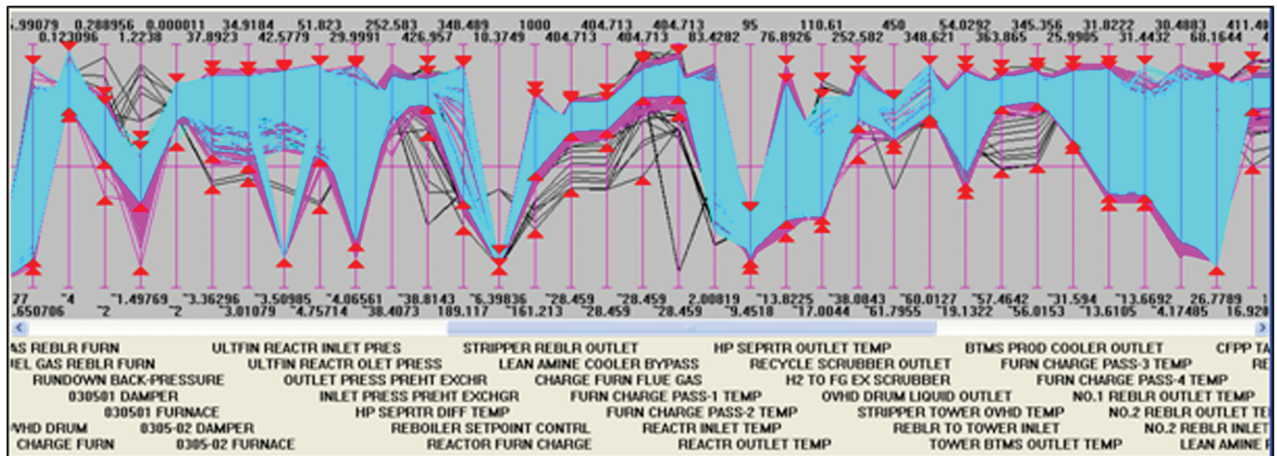
Figure 13. In-Specification Kerosene in turquoise on top of the Starting Alarm Limits for Kerosene Mode

Alarm Limits and Operating Limits are and should be two names for the same thing and that wherever pink is visible in Figure 10 or Figure 13 is bad or abnormal operation that should be eliminated with better operation, better process control and better process understanding.

Figure 14 shows what would happen if the Operating Limits of In-Spec Kerosene in Figure 13 were used as alarm limits today with no change in operation.

The result in Figure 14 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 13 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 12. 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.

To set alarm limits that will achieve the operating objective with fewer alarms than in Figure 14 we first set Operating Limits on the necessary process variables. These will alarm if we attempt to operate outside the normal space but won't necessarily alarm well for abnormal situations involving other variables. We introduce additional alarms on other variables from consideration of unusual conditions that could occur and position these alarms at the edges of the operating space so that they cannot alarm during unless there is an abnormality. We find which variables are necessary for operation using a
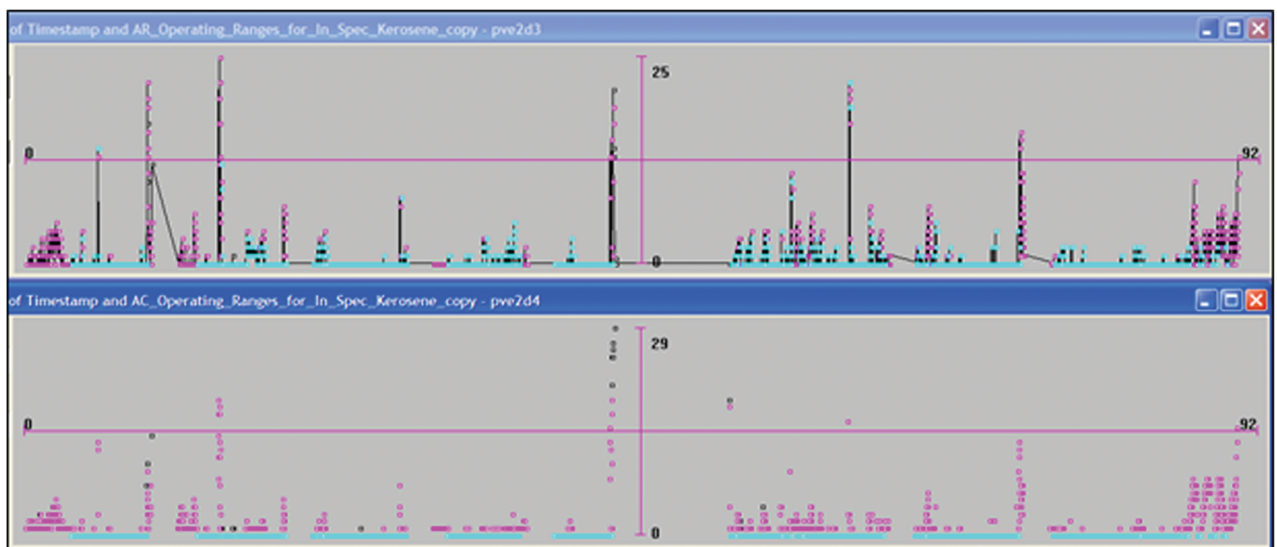

Figure 14. Annunciation Rate per hour and Standing Alarm Count for In-Spec Kerosene Operating Limits
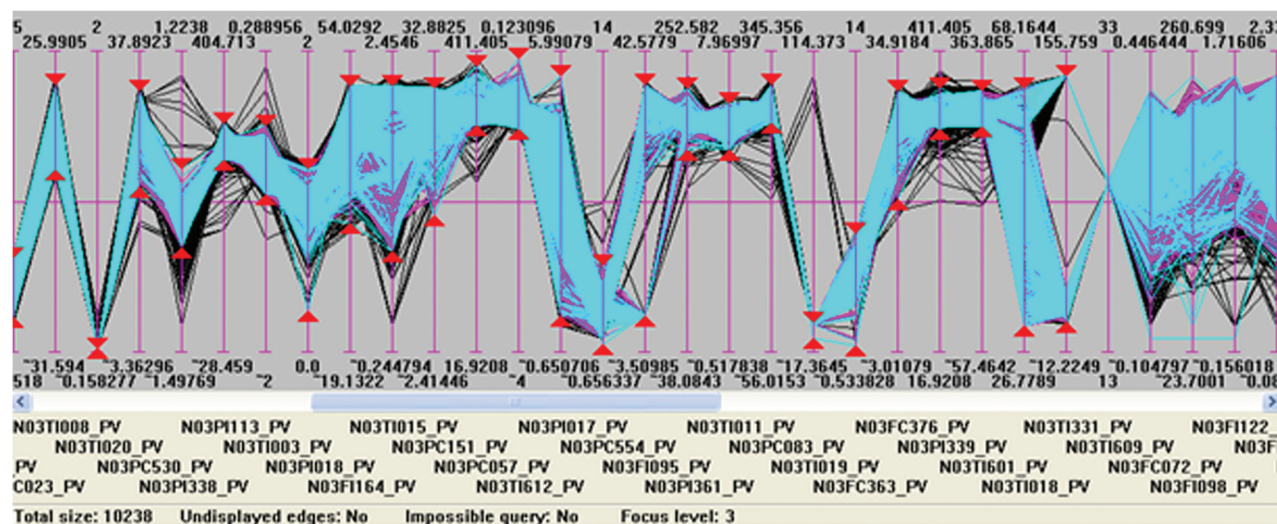
Figure 15. Showing part of the result of the question "how much better would the process perform if it were operated in the turquoise ranges of the in-spec kerosene query". Yield would rise from 24% to 30% even if there were no change in operating practices. Note that operating limits are not imposed on the variables to the right of N03TI609.

patented algorithm called the Box Query which answers the question of (a) how well would the plant perform if we operated in the turquoise ranges projected onto the process variables by the query on the quality specifications in Figure 13 and (b) what is the order of importance of these variables in achieving this result. The algorithm works by finding the lowest dimensionality box that encloses all the turquoise points on the process variables. It defines the Selectivity of each process variable as the ratio of selected to rejected points within the operating range, selects the most selective variable and then repeats the calculation on the remaining variables and points stopping when all turquoise points have been selected. An example is shown in Figure 15 where the order of importance is shown by the left-to-right ordering of the variables with those having no red triangles not contributing further to reducing variability. It is these unselected variables to which we can apply additional alarm limits in anticipation of abnormal or unusual situations that require to be alarmed but by setting them at the bounds of the turquoise ranges they will not cause alarms during normal operation. We have defined the boundary of Normal Operation as in Figure 1 and can now put as

many alarms around that boundary as we wish without fear of the annunciation rates becoming unmanageable (see Figure 16).

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.

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.

ALERTS FROM THE OPERATING ENVELOPE
But, delineating the operating envelope with fixed ranges of values on individual variables ignores the richness of
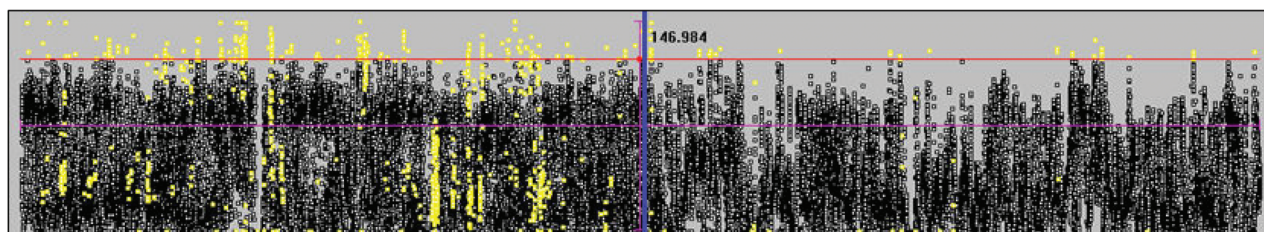

Figure 16. Four years of operation showing the better operation and reduced number of trips (yellow) in the second two years after alarm limits were brought inside the limits of operation of the previous two years
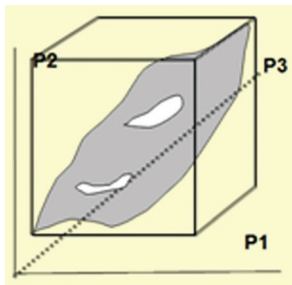
Figure 17. 3-variable hypercube enclosing the Operating Envelope

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. 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 10 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" 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
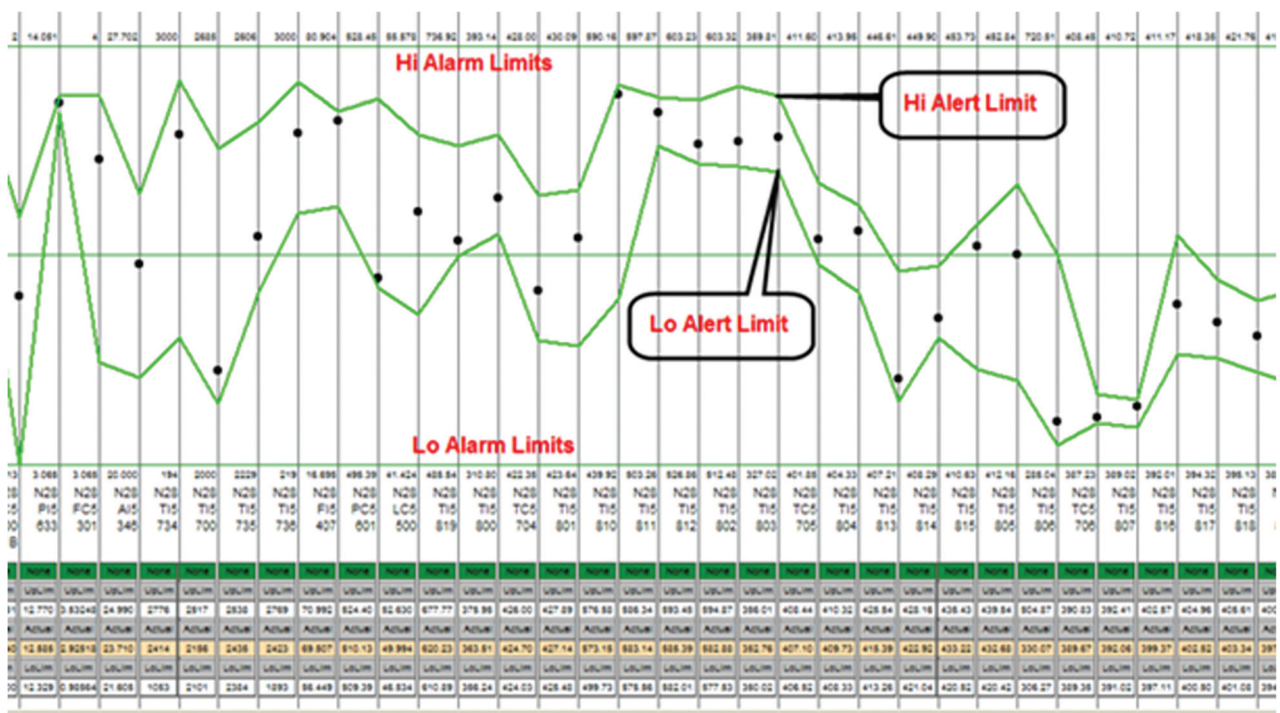

Figure 18. Public Alerts have the objective of keeping the process inside its fixed Alarm/Operating Limits

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.

CONCLUSIONS

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.