

SIMPLE METHODS FOR ALARM SANITATION

Jan Eric Larsson

Department of Information Technology
Lund Institute of Technology
Box 118, 221 00 Lund, Sweden

Abstract: Large industrial processes are often equipped with many alarms. A significant proportion of these alarms may be badly tuned, and therefore activated too often or too seldom. This means that the alarms in question are not useful, and such a situation can potentially be dangerous. *Alarm sanitation* is the process of finding badly tuned alarms and retuning them. This paper describes an ongoing project, in which we are developing simple methods to detect and retune such alarms.

Keywords: Alarm systems, Automatic process control, Calibration, Complex systems.

1. INTRODUCTION

Complex industrial processes are usually equipped with a large number of alarms. It is also common that a significant proportion of these alarms are more or less badly tuned. Many examples are known in different systems, from steel mills to airplane cockpits, where alarm limits are too narrowly set, so that the alarm is activated incorrectly. In such cases, the alarm will disturb the operators, but will not work as a reliable source of supervision information. Alternatively, the alarm limits may be set too widely, so that the alarm is never activated, which also may have negative consequences.

To improve the alarm system of a plant, one has to detect wrongly tuned alarm limits, and retune these limits to correct values, a process, which we will call *alarm sanitation*. To conduct an alarm sanitation is not difficult in principle, but will demand an effort, which can be large if there are many alarms in the plant.

In an ongoing project, we are developing simple methods for detection and retuning of badly tuned alarms. Through automatic logging of alarm activations and signals, it is possible to detect suspicious

alarms and to suggest new, more realistic alarm limits. If, for example, an alarm has been activated ten times during one hour, with no other response than a reset, the alarm could be presented to operators or service engineers, together with the measured signal's mean value, standard deviation, and upper and lower extreme values, with a suggestion for retuning of the alarm's limits. This will make an alarm sanitation easier, because a large fraction of the troublesome alarms will be automatically detected, and saved knowledge about the signal's actual values will be helpful in the tuning operation.

The project described in the following is concerned with alarm detection and presentation, but not so much with validation and consistency checking of sensor values. Instead, we assume that we already have sensors, an on-line data collection system, and algorithms for measurement validation in place.

2. BACKGROUND

Industrial processes usually have a large alarm system, in which measured signals are monitored. Such processes are conventional and nuclear power plants, pulp and paper plants, petroleum refineries, and

chemical and biochemical plants. However, they also include other systems, such as airplane cockpits, medical intensive-care units, and command centrals for different systems. In all such systems, there is a large number of automated alarms, which are usually set to monitor a measured value, and to activate a warning signal if the value goes outside an interval with a lower and an upper limit.

There can be several problems with complex alarm systems. Some of these are:

- The alarm system is badly designed. Some important alarms may be missing, while other alarms may actually be unnecessary for the operation of the plant.
- The alarm system is designed for one operating state but fails to meet the requirements of another. For example, an alarm system may work well in normal operation and for small problems, but fail in case of a major accident.
- The alarm system has a badly designed interface, so that it fails to present the current fault state to operators in an efficient way.
- The alarm system presents too much information and causes mental overload and stress for the operators.

A good example of such problems is provided by the Three-Mile Island accident, see for example Perrow (1999). The alarm system of the Three-Mile Island nuclear power plant was designed to provide alarms for even very small problems, which was useful during normal operation. During the accident, however, too many alarms were activated, and one of the major problems for the operators was to understand the actual state of the plant. In such cases, one would wish for methods or algorithms that could degrade or suppress less important and purely consequential alarms, so as to avoid overloading the operators.

2.1 Trying to Improve the Alarm Situation

In order to improve the situation concerning the problems mentioned above, several methods and algorithms are under development. For example, in the Swedish research program for Complex Technical Systems, sponsored by NUTEK, the Swedish National Board for Industrial and Technical Development, the following projects have been performed:

- Separation of alarms into *primary* and *consequential* based on an analysis of the logic of the alarm and control system, a project from the De-

partment of Computer Science, Stockholm University.

- Grouping of connected alarms into *super alarms*, based on learning from existing alarm lists, the same project as above.
- Better *visualization* of the process, the process state, and the alarm state through better cognitive design of the interface, a project from the Department of Computer Science, Uppsala University.
- Separation of alarms into *primary* and *consequential* based on models of goals and functions of the plant, see Larsson (1998, 1999), Larsson and Öhman (1998), a project of the author from the Department of Information Technology, Lund Institute of Technology.

These are just a few examples. Several other projects are going on, within different research disciplines, such as computer science, artificial intelligence, control theory, human factors, safety, and user interface design. Hopefully, some of these projects will result in improved alarm system technologies.

2.2 What Are Badly Tuned Alarms?

Before more advanced methods and algorithms can be used, some simpler problems must first be solved. Among other things, the problem of *badly tuned alarms* must be addressed. Before it makes sense to send alarm information to advanced algorithms, the alarms that produced the information must be correctly tuned.

In many industrial processes, a significant number of alarms have *wrongly set alarm intervals*. The limits may be either too tight or too wide. If the limits are too tight, the alarm will be activated unnecessarily, when there are small disturbances or noise, but no real fault or problem. If the limits are too wide, the alarm will be activated too seldom or never, and faults will remain undetected. In both these cases, the alarm in question will not fulfill its function, and in practice, it will be useless, which may potentially mean great risks.

We do not know how large the fraction of badly tuned alarms are in a typical plant. It most certainly varies from branch to branch, between different countries, different technologies, and from plant to plant. From automatic control, it is a well-known fact that a significant fraction of all PID controllers are suboptimally or badly tuned. The figures are usually stated to be in the order of 30%, see Bialkowski (1993), Ender (1993), Hägglund (1995). We would

guess that the fraction of badly tuned alarms in many applications is as large or even larger. That they cause large practical problems is clear, see Perrow (1999). Our contacts with people in industry tell us the same thing, Paavola (1999).

2.3 Why Are There Badly Tuned Alarms?

It is our guess that the causes of badly tuned alarms are not of any complex or basic nature. Rather, we believe that the main cause is that too small resources are put into the design and maintenance of alarm systems. The detection and retuning of an erroneous alarm is most probably a simple matter, but when it comes to handling several thousands of alarms, the task becomes complex and costly. In addition, the gains of performing an alarm sanitation are not always immediately visible, at least not until faults and accidents have occurred.

2.4 What Are the Consequences?

The existence of badly tuned alarms has several more or less serious consequences. First of all, let us see the effects of different errors in the alarm limit tuning:

- If the limits are too wide, this will cause an *absence of alarms*. This situation is potentially dangerous, since it is in principle equivalent to not having any alarm, with the *added risk* of false security, because the operators may believe that the alarm is working properly and the plant is in a normal state.
- If the alarm limits are too tight, this will cause *false alarms*. This situation is as bad as the previous one. The consequence of too many false alarms is that the operators will have to disable or routinely ignore the alarm in question. In essence, the situation is again equivalent to not having an alarm at all, with the added risk of mental overload on behalf of the operators.

In short, a badly tuned alarm is as bad or worse than a non-existing alarm. Now, let us analyze the potential negative consequences of having badly tuned alarms:

- One bad effect of having false alarms is that the operators may suffer from stress and mental overload, which in turn may cause them to commit mistakes.
- Another bad effect of false alarms is that the operators may have to use valuable time taking

care of alarms, while they could have used that time to tune the process instead.

- The most obvious bad effect of any kind of badly tuned alarm is, of course, that faults may be overlooked, and that accidents may happen. In such cases, the consequences may be both costly and dangerous.
- Finally, the most severe (but maybe not most obvious) consequence is the threat of unusual and complex accidents, which may be called *once-in-a-lifetime events*. Even if the operators are skilled in handling “normal” accidents without the support of a working alarm system, unusual and rare circumstances may pose severe threats of disaster. In these cases, a good alarm system is of great importance, and any problem here will mean large risks.

In conclusion, badly tuned alarms are common, and do pose serious risks. Before more advanced methods of alarm handling can be used, the problem of badly tuned alarms must be solved. The main obstacle against alarm sanitation is probably that the large number of alarms makes any simple “hand tuning” a large resource-consuming task.

3. AUTOMATED ALARM SANITATION

In an ongoing project, we are developing simple methods for automated alarm sanitation. The basic ideas are:

- Equip the monitoring and control system with a few simple functions, which monitor the behavior of the alarms over time.
- Trace those alarms that are activated often, and are either ignored by the operators, or followed by no other action than a reset of the alarm, and store information about them in a database.
- Likewise, trace those alarms that are never activated in spite of their signals changing mean value, variance, etc., and store this information too in the database.
- In addition, store samples of the signals monitored by “suspicious” alarms, together with a few statistical parameters, such as mean value, standard deviation, minimum and maximum, etc.
- At regular intervals, or when it is suspected that there is a high frequency of false alarms, present the contents of the database to the operators and to service personnel. Each potential “culprit alarm” could be evaluated, and the signal data would provide a good basis for new alarm limits.

3.1 An Alarm Sanitation Toolbox

Our proposed design incorporates an *automated tool for detection* of candidate culprits, i.e., badly tuned alarms, and *computer-based support for retuning* of alarm limits. The evaluation of whether an alarm is actually badly tuned, whether a retuning should be done, and the actual choice of limits, are all to be performed by a human operator or service engineer. We strongly believe in letting human and machine share the burden of solving problems. The tasks that the proposed system should be able to handle can be grouped as follows:

- *Detection of too active alarms.* Here, the methods may be based on activation frequency during a specified time interval, where both the frequency and time interval must be set to match the general time scale of the process. Simply looking at the most frequent alarms may also give good results. Other methods can be based on an analysis of the actions taken after alarm activations. If the operators seem to ignore the alarm or only to reset it, this may be an indication of a “known” false alarm. Finally, the detection of alarms should also be related to the operating state of the plant, so that it is possible to detect alarms that are tuned for one state but not another.
- *Detection of too passive alarms.* Here, the methods used could analyze the monitored signals to detect a fault state, especially in connection to other alarms indicating a fault situation. A detection of changes in the mean value is an obvious suggestion, but more sophisticated statistical methods could also be used. Simulated test signals could be sent to the alarm system to check that all alarms that should be activated actually do react. Finally, formal methods for consequence propagation could be used to see whether some alarms are too silent. MFM would be one possible basis for this, Öhman (1999).
- *Tuning of alarm limits.* Here, the methods used should provide a human user with enough information to select a new set of alarm limits. Different simple statistical measures from the monitored signal should be presented, such as mean value, standard deviation, and maximal and minimal value. In addition, a visual plot of a part of the signal should obviously be provided. Other algorithms could relate differences in these measures to changes of plant operating state, which could lead to the conclusions that an alarm should only be active during certain states, or that different limits should be used during different operating states.

- *Presentation of data.* Here, it is important to provide the system with a simple but good interface, so that the alarm sanitation becomes simple enough that the operators or service personnel would actually have time to perform it.
- *Deciding to initiate an alarm sanitation.* The proposed system will supervise the activity of the alarm system, and when there are enough potentially bad alarms, or such alarms are firing with a high enough frequency, the alarm sanitation system must activate *an alarm of its own*, meaning that it is now time to consider an alarm sanitation. Here, it is indeed important to make the activation decision behave well, so that the alarm sanitation system is not seen as a nuisance and ignored by the operators.

A selection of available and newly developed methods will be chosen and put together for use in the project, to perform all the tasks described above. Rather than to evaluate and select one specific method for each task, we will construct a toolbox of algorithms, from which several different systems can be assembled and tested. In the long run, some methods may prove to be better than others, but the choice of particular method or algorithm may also proved to be a function of the target process and branch.

4. SOME EXAMPLES

Here are some simplified examples of the different cases, where the alarm limits are either too narrow or too wide.

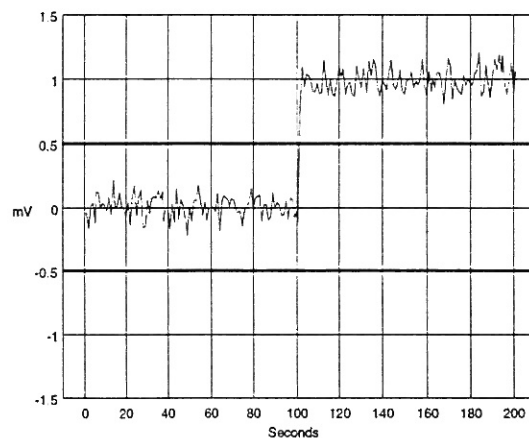


Fig. 1. In this case, the alarm limits are well tuned, and when the signal goes out of bounds at time 100, it can be detected easily and reliably.

In Figure 1, we see a normal situation with well-tuned alarm limits. When a fault occurs, this is seen

clearly in the step change of the signal. An alarm system would have no problems detecting this fault with the used limits.

In Figure 2, we see a situation with too narrow alarm limits. The noise in the signal will cause the alarm to be activated over and over again, even though the signal actually contains no fault information.

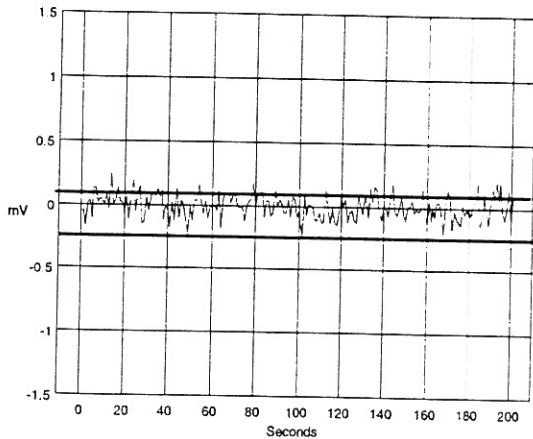


Fig. 2. In this case, the upper alarm limit is set too narrowly, so that the noise in the signal will activate an alarm repeatedly.

Finally, in Figure 3, we see a case of too wide alarm limits. Although a change of level occurs in the signal, this is not detected, because the alarm limits are too widely set. Thus, faults may go undetected in this case.

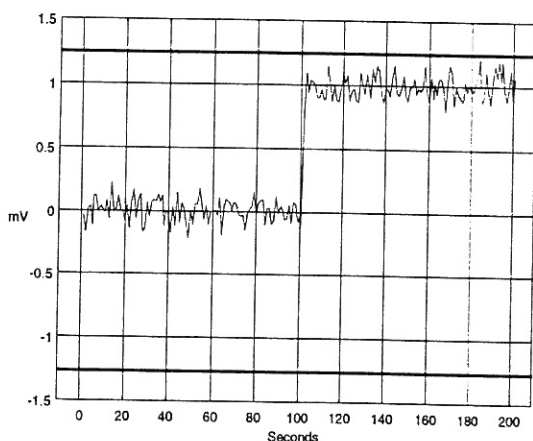


Fig. 3. In this case, the alarm limits are set too widely, and changes in the monitored signal may go undetected.

The examples given here all concern static alarm limits, but similar results would be valid for dynamical alarm limits and methods using model-based fault

detection, see Frank (1996). All the examples are, of course, trivial, but it is hoped that they will help to explain the problems with which alarm sanitation is concerned. These trivial problems do cause large problems in real life, and when they appear in thousands, together they become non-trivial.

5. DISCUSSION

Here are a few further clarifications of the ongoing project and the proposed ideas.

5.1 Design of the Alarm Sanitation Toolbox

It is important to observe that the detection of the badly tuned alarms should be done with full automation. Here, it is of paramount importance to use the good properties of the computer, i.e., that it will not “forget” events or miss them because of stress or tiredness, etc. The retuning of the alarms should make full use of any available human knowledge and judgment, though. Here it is instead important to provide the system with a good interface, so that data can be interpreted quickly and correctly by human operators. For these reasons, we believe that the concept of a toolbox with several methods will be useful, at least until further experiences with the whole task of alarm sanitation have been gained.

5.2 Testing the Methods on Real Cases

The proposed ideas are fairly easy to implement. An implementation and a few test runs are not enough, though, especially since this project addresses a problem that occurs only in practice. Thus, a large part of the project will be to test and evaluate the implemented system in practice, on several different kinds of processes. Here we plan to use our contacts within the research program for Complex Technical Systems, and try our system on conventional and nuclear power plants, on Swedish steel plants, and hopefully other processes too.

5.3 The Decision to Warn about Badly Tuned Alarms

One very important design issue is how the detection system should warn the operators about a situation where an alarm sanitation may be needed. It is obvious that the system must be very careful and not give warnings about potential bad alarms too often. Otherwise, these “meta alarms” may cause more of the very problem that they were meant to relieve. Thus, a

serious effort must be put into the design of the meta alarm decision, and the interface to it in the alarm sanitation toolbox. Obviously, there should be a possibility of clearing an alarm from suspicion for ever, or for ever as long as a set of specific conditions last. Maybe it would also be possible to have a menu of alternatives on how to perform the bad alarm detection.

An interesting observation is that by warning about common alarms, it may also be possible to detect a deteriorating part of the process, where alarms start occurring more often. The operators may ignore the alarms while they are actually an indication of a developing failure. Such cases have been described in, for example, a Swedish steel plant, see Paavola (1999).

6. CONCLUSIONS

Complex industrial processes normally are equipped with a large number of alarms. Most probably, a significant fraction of these alarms are badly tuned. The bad tuning of alarms poses risks of suboptimal production, accidents, and disasters. It also poses problems in using advanced algorithms for consequence analysis of alarms and visualization of the process state. In an ongoing project, we are developing simple methods for automated detection of badly tuned alarms, and for computer-based support for retuning of alarms. We are suggesting automated support for detecting badly tuned alarms, and computer-based support for retuning these alarms. A set of possibly useful methods will be implemented in a toolbox and tested and evaluated on several real processes.

7. ACKNOWLEDGEMENTS

The author would like to thank his students, Fredrik Dahlstrand and Bengt Öhman, for encouragement and inspiration, Jan Tuszyński at Southern Sweden Power Supply Inc. for cooperation and support, and Rutger Gyllenram, Kobilde, Jan Hill, Qualtech, and John Graffman, Bengt Larsson, and Arne Otteblad, NUTEK, for providing inspiring problems and questions, and finally Anu Uus for excellent suggestions on how to improve the paper. The project is sponsored by NUTEK, the Swedish National Board for Industrial and Technical Development, under the project number 1K1P-99-05989.

REFERENCES

- Bialkowski, W. L., "Dreams versus Reality: A View from Both Sides of the Gap," *Pulp and Paper Canada*, vol. 94, no. 11, 1993.
- Ender, D. B., "Process Control Performance: Not as Good as You Think," *Control Engineering*, vol. 40, no. 10, pp. 180-190, 1993.
- Frank, P. M., "Analytical and Qualitative Model-Based Fault Diagnosis - A Survey and Some New Results," *European Journal of Control*, vol. 2, pp. 6-28, 1996.
- Hägglund, T., "A Control-Loop Performance Monitor," *Control Engineering Practice*, vol. 3, no. 11, pp. 1543-1551, 1995.
- Larsson, J. E., "Alarm Analysis for a Nuclear Power Plant Using Multilevel Flow Models," invited paper, Proceedings of the 9th International Symposium on System, Modeling, Control, Zakopane, Poland, 1998.
- Larsson, J. E., "Diagnostic Reasoning Based on Means-End Models: Experiences and Future Prospects," invited paper, Proceedings of the IEEE International Conference on Control Applications and the IEEE International Symposium on Computer-Aided Control System Design, Kohala Coast, Hawaii, 1999.
- Larsson J. E. and B. Öhman, "Model-based Alarm Analysis for Large Plants," invited paper, Proceedings of the International Conference on Systems, Signals, Control, Computers, Durban, South Africa, 1998.
- Öhman, B., "Failure Mode Analysis Using Multilevel Flow Models," Proceedings of the 5th European Control Conference, Karlsruhe, Germany, 1999.
- Paavola, B., "Larmsanering masugnsavsnittet, (Alarm Sanitation on the Melting Oven)," Internal report, SSAB Tunnpå, Swedish Steel Inc., 1999.
- Perrow, C., *Normal Accidents: Living with High-Risk Technologies*, Princeton University Press, Princeton, New Jersey, 1999.