New Solutions for Alarm Problems

Jan Eric Larsson[1], Bengt Öhman[1], Antonio Calzada[1], and Joseph DeBor[2]

1) GoalArt, Scheelevägen 17, 223 70 Lund, Sweden
2) DeBor & Associates, 3630 North 21st Avenue, Arlington, VA 22207, USA
Phone: +46 46 286 4880, Fax: +46 46 286 4882, E-mail: janeric@goalart.com

Abstract – *Alarm problems have been around since the move towards computerized control systems and control rooms. They have been surprisingly hard to solve, and are still here today. In fact, it seems that modern control systems produce more alarm data, and that alarm problems become worse with new systems, not better. We contend that there are two separate reasons for why alarm problems are still largely unsolved today. First, there are different kinds of alarm pro blems, and each kind demands a different kind of solution. Thus, no single technology will suffice to solve all alarm problems. Secondly, some alarm problems have so far lacked viable technical solutions. The hardest problem is that of consequential alarm cascades. However, we can now present an efficient and industrially proven method for handling large alarm cascades through automated root cause analysis.*

I. INTRODUCTION

Today's advanced, low-cost sensors and computerized control systems enable us to include an unprecedented number of measurements and alarms in the alarm systems for new control rooms. This allows a higher level of awareness of plant state and faults than before. However, it also creates problems of information overload, so called alarm problems.

We contend that there are several different types of alarm problems, and each type demands its own solution. One solution cannot solve all alarm problems. Among the different kinds of alarm problems are:

- Badly configured alarm systems, where there are too many unnecessary alarms in the system.

  The solution is to perform an alarm system revision, where the function and use of each alarm is analyzed, and unnecessary alarms are removed from the system forever. To aid this process, we have developed a method for *sensor placement analysis*. This method can calculate which sensor equipment and placement is needed in order to observe all faults in the process. In this way, it is possible to validate that a sensor can indeed be removed without making the operators blind to some faults.

- Badly tuned alarm parameters, such as low and high limits, dead bands, and filters. This may lead to *nuisance alarms*, when limits are too close, and so called *silent alarms*, when limits are too wide.

  The solution is to re tune the alarm parameters. We have developed a method called *alarm cleanup and maintenance*, which uses logged analog trend data to automatically retune alarm parameters.

- Alarms may be irrelevant for certain operating states. Typical examples are alarms tuned for normal operation, which are generated but irrelevant during startups, shutdowns, state changes, and from disabled equipment.

  The solution is to assign dynamic priorities (including suppression) to alarms depending on the plant operating state. We have developed a method called *state-based alarm priority*, with which it is efficient to assign such dynamic priorities, estimate the current state, and perform the prioritization and suppression.

- The most dangerous alarm problem is that of *alarm cascades*. An initiating event usually creates several consequential faults, and instead of a single alarm, the system may produce hundreds of alarms. This can make a fault situation extremely difficult to understand, and the problem typically occurs at the most critical moments. These are correctly tuned alarms and relevant to the current situation, but the amount of information creates a difficult diagnostic situation.

  Our solution is based on a new *root cause analysis* algorithm, which allows efficient modeling and fast execution, so that even large plants can be handled with a reasonable effort. Our method reduces large alarm cascades containing hundreds of alarms to one or two alarms directly connected to the initiating events.

GoalArt is a company that specializes in alarm reduction and automated operator support for industrial processes and complex technical products. Our main experience is in nuclear power plants, power grids, and conventional heat and power.

II. ALARMS AND INFORMATION OVERLOAD

Alarm problems have been known since the introduction of control room technology in the sixties. In spite of this, they have been more difficult to solve than expected, and are still present today. We propose that there are two reasons for this.

- There are several different types of alarm problems, and each type of problem demands its own solution. So far, many efforts to solve alarm problems have used only one method or technology, and therefore failed to solve all different types of problems.

- So far, there has been a lack of methods to solve some of the more difficult problems in a technically and commercially viable way. Notably, the problem of consequential alarm cascades needs an efficient method for root cause analysis, and so far, the proposed solutions have demanded a prohibitively large amount of knowledge engineering or modeling work.

So when there are alarm problems in a control room, they often consist of a mixture of several different kinds of problems. Some common types of alarm-related problems are:

- The alarm system may be badly configured and the alarm measurement points may be badly placed. In today's computer-based control systems, it is quite easy to add another alarm point, and sometimes there are simply too many unnecessary alarms in the system. The solution here is to perform an *alarm system revision* and remove unnecessary alarms.

- Alarm limits and other parameters may be wrongly tuned, so that spurious alarms are generated because of noise, etc. For example, if a limit is too tight, signal or process noise may cause the alarm to be activated a large number of times, when there is really nothing wrong, just because the signal is close to the limit and the noise pushes it over. The solution to these kinds of problems is to tune the alarm parameters, limits, and filters, based on either process knowledge or historical trend data.

- Alarms are often designed for a certain operating state, such as, for example, 100 % production, while they may be irrelevant for other states, for example, stand-by or emergency shutdown. It is a well-known phenomenon that alarm showers tend to appear during state changes, such as startup and shutdown, and often alarms are generated from equipment, which is switched off and not in operation. In themselves, such alarms are expected and easily understood, but the problem is that *they may hide other alarms*, from faults occurring during the state changes. If a fault occurs during a startup, the operators may not see the corresponding alarms, because they were drowned out and lost in the large shower of "usual" alarms. The solution to this problem is to make the alarm generation or presentation state-sensitive, and to suppress those alarms that are irrelevant to the current state. A potentially efficient method for accomplishing this is has been developed by GoalArt and tested in several projects.

- When there is a fault in a process, it usually causes several consequential faults. If all faults are monitored by the alarm system, an original fault is usually followed not by one alarm, but often by tens or hundreds of alarms, a so-called *alarm cascade*. Because the alarms in the cascade seldom arrive in exactly correct time order, it is usually difficult to analyze and understand the fault situation. The solution to this problem is to apply a root cause analysis to the fault situation, to find out the original fault and the causal chain of events. An efficient method for has been developed by GoalArt and tested in several projects.

- No matter how well an alarm system is working, there will, from time to time, be nuisance alarms of two kinds. One kind is alarms that come and go repeatedly in special circumstances. These can be temporarily suppressed, *shelved*, and brought back when the situation calms down. GoalArt has developed an adaptive, automatic method for shelving of nuisance alarms. The other type is alarms that belong to broken equipment, which will remain active until the equipment is repaired. GoalArt has developed an adaptive, automatic method for shelving so-called "long lasters." Once the alarm is fixed, it is automatically de-shelved again.

- Another type of problem relates to the presentation of alarms in the human-machine interface. If too many alarms are presented, the operators will be overloaded by the information and unable to utilize the information provided. If, on the other hand, to few alarms are presented, while others are suppressed, the operators may not have the right information needed to successfully analyze and understand the situation. The solution here is to design the alarm presentation system so that the more important information is easily available and other information can be hidden or shown according to the current needs of the operators.

- There may also be alarm problems related to the organization. For exa mple, there should be clear and well-documented decisions about the alarm policy of the company, about the roles and responsibilities of different people and groups, and well-established routines for operation and handling of the plant in fault situations. The solution here consists in creating an alarm philosophy or policy for the organization, and education, training, and problem handling in the group of people operating and maintaining the plant and control system.

- When suppressing alarms, it is important that alarms should never be permanently removed from the system. They may be sorted into different lists and windows, or hidden from view, but the alarm system must clearly indicate that there are hidden alarms, and it must be quick and easy to retrieve the hidden alarms. GoalArt's philosophy is to never remove an alarm, but to hide it or present it in a different place.

- Finally, there are people who state that it is dangerous to suppress and/or hide alarms, and that this should never be done. We contend that this is wrong. There *is* a potential danger in having an automated system suppress an alarm. The design may have hidden faults in it. However, the situation where all alarms are presented no matter what, is much more dangerous. The risk that the operator will commit an oversight, under stressful conditions, facing a cascade of alarms, and having to do all the reasoning himself, is much larger than the risk of the engineer making mistakes in the calm design situation, going over alarm by alarm in a systematic way, and with time to go back and think again if needed. *Thus, although alarm suppression is associated with risks, not doing it is much more dangerous.*

The main conclusion to be drawn from the above is that there are a number of different types of alarm-related problems, and each type of problem demands a separate solution. For example, an alarm revision and alarm system redesign may improve the average alarm load, but it will not help against consequential alarm cascades. A new alarm presentation may help the operators to notice a problem quicker, but if there are too many spurious or false alarms, there will still be a problem.

MEASURING ALARM PROBLEMS

How do you know whether you have alarm problems? Often, the control room personnel have a vague feeling that the system produces too many alarms, and is difficult to use in complex fault situations.

In order to get more hands-on data, we recommend the use of some simple key numbers. Donald Campbell Brown suggests the following measures, [1], inspired by previous work by EEMUA, [2]. At GoalArt, we call these measures DCB numbers.

- Average number of new alarms per hour. This number gives a general measure of the alarm load during normal operation.

- Worst case over any ten-minute period of average number of alarms per hour. This gives a measure of alarm load during state changes and plant upsets.

- Percentage of hors with more than 30 new alarms. This gives a measure of how often the plant is in a "disturbed" or unclear mode.

These key figures are simplistic, and do not catch all aspects of the tuning on an alarm system. On the other hand, they are easy to measure or calculate, and they will give you a quick but factual indication of the alarm load.

Donald Campbell Brown suggests the following table for interpreting DCB numbers, see Table 1.

Alarms/hour	Max/hour	% hours	Classification
> 600	> 6 000	> 50 %	Overloaded
> 60	> 6 000	> 25 %	Reactive
> 6	> 600	> 5 %	Stable
> 6	> 60	> 1 %	Robust
< 6	< 60	< 1 %	Predictive

Tab. 1. Classification table using DCB numbers.

The table gives the following interpretations of the alarm system measures.

- *Overloaded* means that operators will regularly miss important alarms. There will be so-called "human errors" induced by the general alarm load, simply because it is not humanly possible to keep up with this information rate. The first remedies to take are alarm revision and alarm cleanup.

- *Reactive* means that the operators are busy with acknowledging alarms, rather than running the process. The alarm system is controlling the operators instead of the other way around. The alarm rate is high enough that important alarms will be missed. Remedies are alarm revision, alarm cleanup, and state-based alarm priority.

- *Stable* means that the alarm rate during normal operation is acceptable, but that the operators loose control during upsets and state changes. This may be the most dangerous level, because there are serious alarm problems, but they only show up at incidents. Remedies here are state-based alarm priority and root cause analysis. Note that alarm revision and alarm cleanup will not help. The tuning of single alarms has almost no impact on alarm cascades and state-change induced alarm showers.

- *Robust* means that normal operation is good, while during plant upsets, there are trouble, but no total loss of the alarm system's usefulness. Such a situation can and should be improved by state-based alarm priority and root cause analysis.

- *Predictive* is the ideal state of affairs when the operator team anticipates trouble and takes predictive action to avoid ending up in trouble.

GoalArt's internal statistics show that many conventional power plants in northern Europe are operated under *stable* or *robust* conditions, while sometimes having periods of *reactive* alarm loads. Unofficial figures state that *overloaded* and *reactive* operation occurs quite often in the petrochemical area.

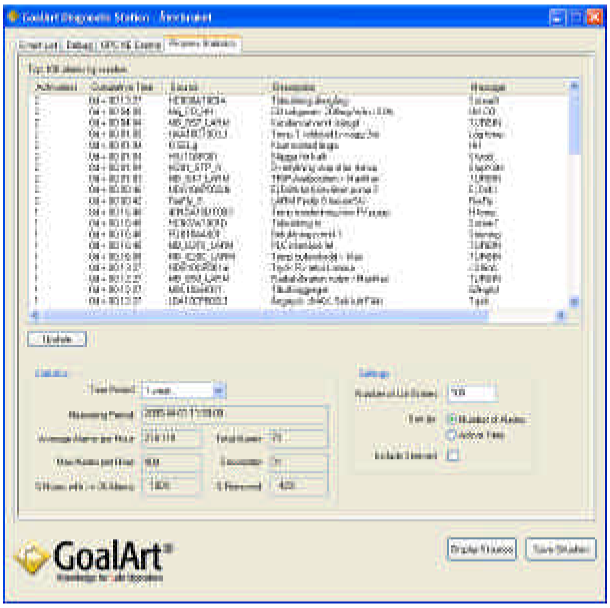

Fig. 1. A top list of worst-offender alarms from GoalArt's alarm statistics and plant classification algorithm.

Several conventional plants operated by current GoalArt customers actually have periods of *predictive* operation. Maybe this is a sign that being interested in advanced alarm management goes hand in hand with having a well-tuned alarm system.

It is our own impression that many nuclear power plants operate in the *stable* category. During normal operation, the alarm system is well-tuned, while most process upsets cause a very large amount of alarms, bringing the plant into the > 600 alarms per hour level. This can be remedied by GoalArt's state-based alarm priority and root cause analysis.

GoalArt provides algorithms for both off-line and on-line measurement of DCB numbers. In addition to this, our system also produces on-line top lists of worst offender alarms, and longest lasting alarms. This is all collected in an algorithm called alarm statistics and plant classification (ASPC), see Figure 1. By running this on top of the existing control system, it is possible to get a short-term and long-term monitoring of the alarm load and "health condition" of the alarm system and operators' working conditions. You will also detect periods when the alarm tuning gets worse, for example, after plant revisions where new problems were accidentally introduced.

III. ALARM CLEANUP

GoalArt has developed an easy-to-use and cost-effective method for validation and retuning of alarm parameters. It works in the following way.

- Analog signals with one or several alarm limits are logged during a period of 1-2 weeks. We recommend using a couple of periods, with full and lower operation, and preferably a period with a plant trip.

- GoalArt performs an automated analysis, with an algorithm optimizing the settings of alarm parameters, such as limits, filters, and hysteresis. There is also a visual analysis for validation and additional conclusions.

- GoalArt produces a report where each signal is described, with current settings, recommended settings, and projected gains form possible re-tuning. Each signal has a separate page in the report, see Figure 2.

- The plant owner and GoalArt go through the report together, and may retune some of the alarm parameters settings in the control system. The alarm cleanup can be repeated regularly, to maintain well-tuned and validated alarm system settings.

The alarm cleanup and maintenance method has several benefits:

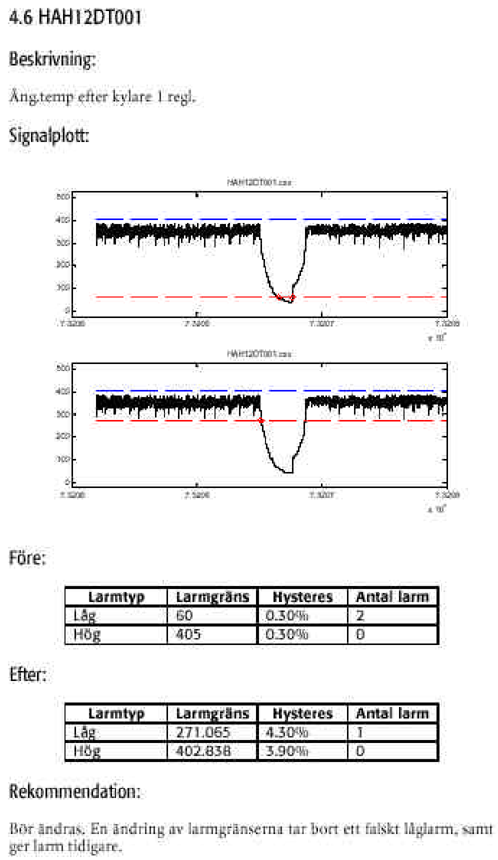

4.6 HAH12DT001

Beskrivning:

Ång.temp efter kylare 1 regl.

Signalplott:

Före:

Larmtyp	Larmgräns	Hysteres	Antal larm
Låg	60	0.30%	2
Hög	405	0.30%	0

Efter:

Larmtyp	Larmgräns	Hysteres	Antal larm
Låg	271.065	4.30%	1
Hög	402.838	3.90%	0

Rekommendation:

Bör ändras. En ändring av larmgränserna tar bort ett falskt låglarm, samt ger larm tidigare.

Fig. 2. A page from an alarm cleanup report.

- The method tests and validates that there are no wrongly tuned alarm parameters in the system, such as bad alarm limits, unstable filters, or erroneous dead-bands.

- If there are problems with nuisance alarms, that depends on too tight alarm limits, the methods gives advice on how to tune these alarms away.

- The method discovers so called silent alarms, where the limits are too wide, so that faults can occur without any alarm being generated.

- The method validates that there are no sensor or measurement faults. These shows up immediately as a difference in the number of alarms calculated by the algorithm and the number of alarms actually generated by the control system.

- The risk of human error in validation and tuning is reduced.

In summary, it is possible to maintain well-tuned alarm parameters by manual efforts, and this is done today in all nuclear plants. However, this does a demand a large work effort of manual and tedious work. GoalArt's method is easy to use and automates the work, making the effort less and reducing the risk of induced faults.

IV. STATE-BASED ALARM PRIORITY

Alarms and events are often tuned for one operational state (typically production), while they may be irrelevant in other states, such as startup, shutdown, and switched off. GoalArt's algorithm state-based alarm priority (SBAP) identifies the current process state and gives each alarm and event a dynamic priority, including suppression in irrelevant states.

GoalArt's tool *SBAP Configurator* lets the user define the different states, the paths between the states, and the priorities for each alarm in each state. There is also a logical definition of how each state should be recognized. All these definitions are performed with a graphical tool much like an interactive table or "Excel sheet," see Figure 3. It is also possible to define that an event should be created when an alarm or event, which is expected to appear in a certain state, does not come.

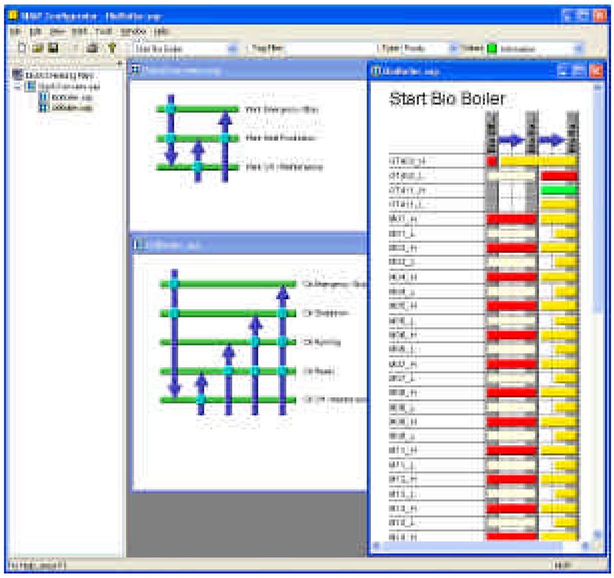

Fig. 3. A screen shot from the state-editor of GoalArt's *SBAP Configurator.*

Once the dynamic alarm priority scheme has been edited and saved, it is used by an on-line algorithm, which reads the alarms and events of the underlying control system, and adds a dynamic priority status number. This is then used by the human-machine interface to sort the alarm into different lists, and to suppress the irrelevant alarms.

The direct benefits of state-based alarm priority are that it gives a large reduction of presented alarms and removes irrelevant alarms and events. The operators are also notified about expected but missing events.

Fewer presented alarms means that it is easier for the operator to understand the current process state, and the risk of missing a critical alarm (or acknowledging it by mistake) is reduced. The presentation of relevant alarms only, increases the trust in the alarm system.

An alternative to SBAP is programming alarm suppression directly in the control system code. With *SBAP Configurator*, the works effort needed is orders of magnitude less, and SBAP allows the operator to see suppressed alarms if needed.

V. ROOT CAUSE ANALYSIS

A fault in a plant usually leads to several consequential faults. When the plant is well equipped with alarms, this causes a large number of alarms. A single fault may lead to hundreds of alarms, a so-called *alarm cascade*. Quite often, the root alarm does not appear first, and alarm cascade situations can be very difficult to analyze.

GoalArt's algorithm alarm analysis (AA) analyzes complex alarm situations and separates root causes from consequential faults. The root causes can be shown in a separate alarm list, while the consequential faults are shown in another list or are suppressed, for presentation when needed.

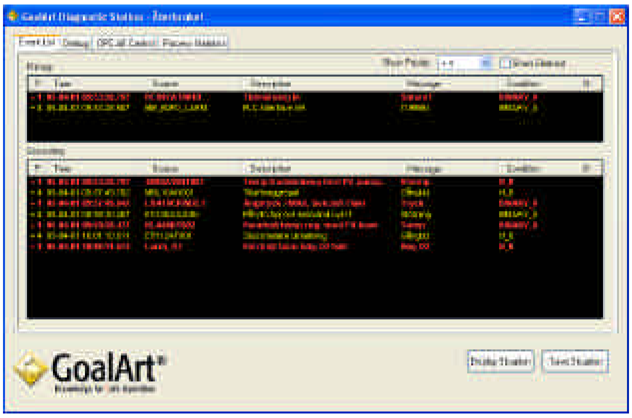

Fig. 4. A screen shot from a GoalArt Diagnostic Station, showing a (simulated) situation with two independent root cause alarms, and several consequences. Root causes are shown in the upper list, consequences in the lower one. The coloring indicates the dynamic priority set by SBAP.

The algorithm uses a simple description of the plant causality, called an MFM model. In order to install an alarm analysis system, one must build an MFM model of the plant. The most important aspect of this is that the construction of MFM models is easier and demands considerably less effort than other modeling, such as, for example, programming in control system code or building simulation models.

The main benefit of alarm analysis is that it helps operators and service personnel to understand complex alarm situations rapidly and correctly. This, in turn, means increased productivity, fewer production stops, shorter re-pair times, avoidance of sub-optimal operation, and decreased risk of accidents and emissions. The operator's trust in the alarm system will also in-crease.

The alternative to alarm analysis is to use fault trees or rule-based expert systems. The effort of using one of these techniques is order of magnitudes higher than with MFM, and usually prohibitive.

MULTILEVEL FLOW MODELS

The root cause analysis algorithm is based on a modeling methodology called Multilevel Flow Models (MFM). These are graphical models of goals and functions of technical systems. The goals describe the purposes of a system or subsystem, and the functions describe the capabilities of the system in terms of flows of mass, energy, and information. MFM also describes the relations between the goals and the functions that achieve those goals, and between functions and the sub-goals.

MFM was invented by Morten Lind at the Technical University of Denmark, [8-10]. Several new algorithms and implementations have been contributed by Jan Eric Larsson at Lund Institute of Technology, [3-6]. MFM development started in the late seventies and has reached industrial application in the beginning of this century, [7, 11]. MFM provides a good basis for diagnostic algorithms.

The algorithms described in [3] are based on discrete logic. The MFM algorithms all operate by searching in fixed graphs. All cases are handled by search methods of linear or sub-linear complexity. Together with the discrete logic, explicit means-end concepts, and graphical nature of MFM, this gives several advantages:

- The graphical representation provides strong support for knowledge base overview and consistency.
- The high level of abstraction makes knowledge acquisition, knowledge engineering, and knowledge base validation and support considerably easier than with standard rule-based systems or fuzzy logic systems.
- The graphical nature of the models allows the algorithms to have good real-time properties, such as an easily computed worst-case time, low memory demands, and high efficiency.

- The high level of abstraction allows the algorithms to be very fast. A worst-case alarm analysis on the full Hambo system takes a second on a standard PC.

These advantages have been observed in practice, during two nuclear projects, [7, 11].

MFM knowledge bases can be constructed with much less effort than needed for other knowledge-based methodologies. In an ongoing project, we have built an MFM model for the Hambo simulator, comprising the majority of the status indicators (around 6 500 signals) for the Forsmark 3 nuclear power plant in Sweden. For this we have used around 4 man-months of modeling, testing, and validation time. Another 4 man-months were used by IFE in testing the system, [11].

MORE MFM DIAGNOSTIC ALGORITHMS

Based on the MFM technology, GoalArt has developed several algorithms that can perform diagnostic reasoning tasks, for use either during plant design and redesign, or on-line during actual operation.

- *Sensor placement analysis* uses an MFM model to calculate whether a certain set of sensors can detect all faults in a process, and if not, where the blind spots are. This can be used to validate that all faults can be detected by the alarm system, to validate that a certain sensor is indeed superfluous and can be removed, and to validate sensor redundancy in a formal way.

- *Probability and safety analysis* (PSA) uses an MFM model and pre-calculated or measured values of reliability and availability to calculate reliability values for an entire sub-system or plant. It automatically obtains reliability and availability values for the plant, and checks that all systems fulfill SIL value demands. The algorithm follows the method describe sin the IEC 61508 standard and can be used off-line during design, and on-line during operation.

- *Single-fault tolerance analysis* uses an MFM model to validate that any single fault cannot cause a stop of the entire plant. Alternatively, the algorithm can be set up to validate that a design fulfills N -way redundancy for all involved systems.

All these design and validation algorithms share the formal property that, if the MFM model is correct, the validation can be *proved* to be correct.

- Root cause analysis uses an MFM model and incoming alarms and events to reduce consequential alarm cascades to single root cause alarms.

- Sensor fault detection uses an MFM model and incoming alarms and events to validate that the alarms are consistent, that is, to detect erroneous measurements, alarm limits, and alarm indications.

- Failure mode analysis uses an MFM model to predict consequences of the current fault situation, as well as of proposed actions taken by the operator. This resembles FMEA, but can be used on-line as a real-time planning support tool.

- Startup planning and ready-to-run validation uses an MFM model to either plan a complex startup procedure, or to validate that each step in a maneuver is allowed. For example, it can be used to check that support systems and activated before operations are commenced.

All these algorithms use *the same MFM model* as a database. Once a model has been built, all algorithms are available at no further knowledge engineering cost.

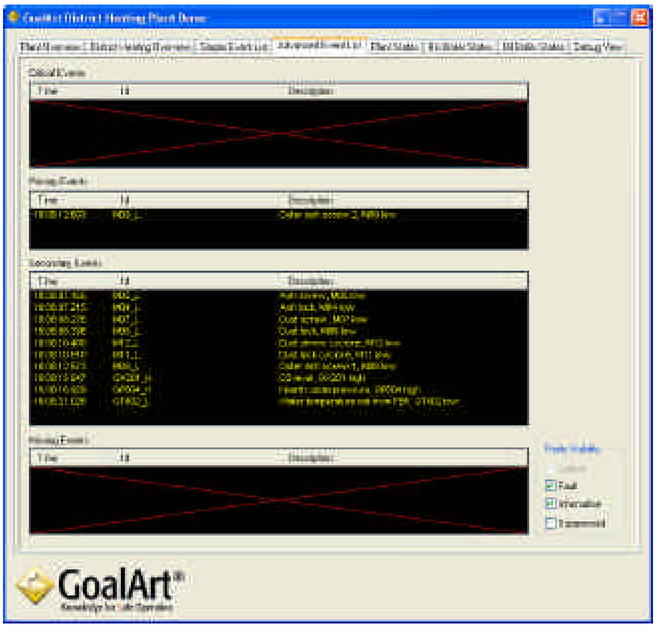

Fig. 5. A screen shot of GoalArt's improved alarm presentation, with separate lists for critical priority alarms, root cause alarms, consequential alarms, and missing alarms.

HUMAN-MACHINE INTERFACE

One very important aspect of an alarm system is the presentation of the alarms. The alarm list is far from ideal. It dates back from the line printer times, and is designed as much for documentation as for real-time use. We strongly believe that the best presentation strategy is to combine classical alarm lists with alarm presentation in process schematics.

The list presentation can easily be improved using GoalArt's algorithms for state-based alarm priority and root cause analysis. Our current standard solution is to use several lists, see Figure 5.

In our presentation, we separate root cause alarms and present them in a special list, while all consequences are shown in another list. This gives the operator an idea of the number of independent root causes at a glance. For example, in Figure 5, there is a single root cause, while in Figure 4 there are two independent fault chains.

Using SBAP, our system assigns a dynamic priority to every alarm. If an alarm has the highest priority (critical), we show it in the top list of critical alarms, independent of whether it is a root cause or a consequence. For example, we recommend giving all major trip signals a critical priority. In this way, a scram signals will appear in the critical list, while the cause of the scram will appear in the root cause list. All 300 – 500 consequence alarms that normally appear during a scram will be shown in the consequential list.

Finally, SBAP knows what alarms to expect in different states, and can calculate so-called missing alarms. These are messages saying that an alarm should have come (and might have been suppressed) but did not. These missing alarm messages are shown in a fourth list, see Figure 5.

In this way, if something goes wrong during, say, a scram operation, and some system fails to work, the operator will see missing actions at a glance.

CONCLUSIONS

This paper gives an overview of new methods for solving alarm problems. In particular, we have introduced new methods for reducing alarm cascades to single root cause alarms. This can be one without permanently removing alarms from the system. In fact, with GoalArt's methods in place, the new availability of large amounts of alarm information can be used without the risk of information overload.

ACKNOWLEDGEMENTS

The authors would like to thank Anu Uus at GoalArt for proofreading, feedback, and valuable advice, and our nuclear partners in Scandinavia for the enthusiasm they have shown for our methods and for testing them in the two nuclear projects.

REFERENCES

[1] D. C. Campbell Brown, "Horses for Courses—A Vision for Alarm Management," Proceedings of the 3rd in the Series of Seminars & Work-shops on Alarm Systems, IS1172, IBC Global Conferences, London, UK, (2002).

[2] EEMUA, "Alarm Systems—A Guide to Design, Management, and Procurement," Publication 191: 1999, The Engineering Equipment and Materials Users Association, London, (1999).

[3] J. E. Larsson, *Knowledge-Based Methods for Control Systems*, Doctor's thesis, TFRT – 1040, Department of Automatic Control, Lund Institute of Technology, Lund, (1992).

[4] J. E. Larsson, "Diagnostic Reasoning Strategies for Means-End Models," *Automatica*, **30**, 5, (1994).

[5] J. E. Larsson, "Diagnosis Based on Explicit Means-End Models," *Artificial Intelligence*, **80**, 1, (1996).

[6] J. E. Larsson, "Diagnosis Reasoning Based on Explicit Means-End Models: Experiences and Future Prospects," *Knowledge-Based Systems*, **15**, 1-2, (2002).

[7] J. E. Larsson, et al. "A Revival of the Alarm System: Making the Alarm List Useful During Incidents," Proceedings of the NPIC&HMIT, Albuquerque, 2006.

[8] M. Lind, "Representing Goals and Functions of Complex Systems — An Introduction to Multilevel Flow Modeling," Technical report, 90–D–38, Institute of Automatic Control Systems, Technical University of Denmark, Lyngby, (1990).

[9] M. Lind, "Modeling Goals and Functions of Complex Industrial Plant," *Applied Artificial Intelligence*, **8**, 2, (1994).

[10] M. Lind, "Plant Modeling for Human Supervisory Control," *Transactions of the Institute of Measurement and Control*, **21**, 4-5, pp. (1999).

[11] J. Tuszynski, J. E. Larsson, C. Nihlwing, B. Öhman, and A. Calzada, "A Pilot Project on Alarm Reduction and Presentation Based on Multilevel Flow Models," Proceedings of the Enlarged Halden Programme Group Meeting, HPR-358, Storefjell, Gol, Norway, (2002).