THE ENGINEERING EQUIPMENT AND MATERIALS USERS' ASSOCIATION

Alarm Systems
A Guide to Design, Management and Procurement

PUBLICATION No 191

Edition 2

ISBN 0 85931 155 4

Imprint reference 07-2007

Registered and Trading Address
10-12 Lovat Lane London EC3R 8DN
Telephone: +44 (0)20 7621 0011
Fax: +44 (0)20 7621 0022
E-mail: **sales@eemua.org**
Website: **www.eemua.org**

ENGINEERING EQUIPMENT AND MATERIALS USERS' ASSOCIATION

The Engineering Equipment and Materials Users' Association, more commonly known as EEMUA, is a European based, non-profit distributing, industry Association run for the benefit of companies that own or operate industrial facilities.

EEMUA aims to improve the safety, environmental and operating performance of industrial facilities in the most cost-effective way.

EEMUA Members pursue these aims by sharing engineering experiences and expertise, and by the promotion of their distinct interests as the *users* of engineering products.

More specifically, the aims of EEMUA Member companies are achieved by:

- providing the organisation within which networking, information sharing and collaboration on non-competitive technical matters can take place;
- influencing the way written regulations are interpreted and applied in practice;
- presenting and promoting Members' views, and encouraging the application of good, sound engineering practices;
- developing and publishing user guides, specifications and training materials;
- facilitating Members' participation in national and international standards making;
- influencing relevant national and European legislation and regulations.

Formed in 1949 as the Engineering Equipment Users Association, and re-named in 1983 (as a result of taking over the materials association, OCMA), **EEMUA** has for more than fifty years given companies that own and operate process plants, power stations and other significant industrial facilities a collaborative voice in addressing technical and engineering related issues that impact on good integrity management and asset management practices. The Association is open to companies of all sizes that meet the 'engineering user' criteria.

The following were the leading plant owners and operators participating in EEMUA at the time of publication: AstraZeneca, BASF, BP, Chevron, ConocoPhillips, Dow Corning, E.ON, ExxonMobil, Flexsys, Hydro, Innospec, Johnson Matthey, RWE npower, SABIC UK Petrochemicals, Shell, Statoil, Syngenta, Total, Vopak.

EEMUA activities often lead to the production of publications. These are prepared primarily for Members' use, but may be offered for sale as well.

A list of EEMUA publications for sale is given at the end of this publication. The full list is also on the Association's website at www.eemua.co.uk/acatalog/shop.html, together with on-line shopping facilities.

To enquire about corporate Membership, write to info@eemua.org or call +44 (0)20 7621 0011.

ABOUT THIS PUBLICATION

Legal Aspects

It has been assumed in the preparation of this publication that: the user will ensure selection of those parts of its contents appropriate to the intended application; and that such selection and application will be correctly carried out by appropriately qualified and competent persons for whose guidance this publication has been prepared. EEMUA does not, and indeed cannot, make any representation or give any warranty or guarantee in connection with material contained in its publications, and expressly disclaims any liability or responsibility for damage or loss resulting from their use. Any recommendations contained herein are based on the most authoritative information available at the time of writing and on current good engineering practice, but it is essential for the user to take account of pertinent subsequent developments and/or legislation.

Any person who encounters an inaccuracy or ambiguity when making use of this publication is asked to notify EEMUA without delay so that the matter may be investigated and appropriate action taken.

Consultation and Feedback

EEMUA encourages constructive comments on this publication from both Members and non-members of the Association.

Comments should be sent on the standard feedback form, a copy of which is provided towards the end of this publication. An electronic version (MS-Word document) of the form is available from EEMUA - e-mail your request for a copy to info@eemua.org, telephone your request to +44 (0)20 7621 0011, or download the form from the EEMUA website at www.eemua.org. Submit comments by e-mail (preferred) or post.

Comments will be considered by the relevant EEMUA Technical Committee and may be incorporated in future editions of the publication. New editions will be publicised on the EEMUA website.

Cover photograph reproduced with kind permission of Shell Photographic Services, Shell International Ltd

Printed and bound by CPI Antony Rowe, Eastbourne

Foreword

Alarm systems are increasingly important in the safe management of plant and machinery. Where major hazards may be present, effective alarm management is an area of continuing concern for the Health & Safety Executive (HSE)[1].

This guide provides clear – and now tried-and-tested – guidance on alarm system design, maintenance and continuous improvement. Although inspection experience has shown that significant progress has been made in industry practices since the first edition of the guide, there is still work to do; and of course maintaining a well-managed system requires continued vigilance.

While more sophisticated control systems can help to optimise plant control and production, and reduce immediate operator exposure to risk, there can be less easily-foreseen, and sometimes major, consequences if these more complex systems are not well designed and managed. In particular, account should always be taken of the human capacity to respond effectively to alarms. There are existing standards for systems with safety-related applications[2] but there are still some stringent requirements for non safety-related control systems: failure to manage these appropriately can be a significant contributory factor to loss of control and accidents.

Following the original HSE survey on alarm systems in 1998, HSE contributed towards the production of the original edition of this guide by EEMUA, and was consulted on this revised 2007 edition. The guide is highly relevant to the process industries, though clearly the principles apply more widely.

Inspectors carrying out assessment and inspection activities may look, when necessary, for evidence that the principles and recommendations in the EEMUA 191 guide (or an equally effective equivalent) are being, or have been, applied to alarm system design and management. HSE experience is that such application produces significant safety improvements as well as other operational benefits.

KJ Allars

Kevin Allars
Deputy Director, Hazardous Installations Directorate
Health & Safety Executive
UK 2007

[1] See e.g. 'Better alarm handling', HSE information sheet at www.hse.gov.uk/pubns/chis6.pdf.
[2] The IEC 61508 series of standards offers authoritative guidance on functional safety for electrical, electronic and programmable electronic safety-related systems. The process industry implementation of this standard, IEC 61511, offers advice tailored to the process sector.

ASM Consortium

The Abnormal Situation Management Consortium (ASM Consortium) fully endorses the comments of the British Health & Safety Executive in the Foreword to this guide.

ASM Consortium members contributed to the first edition of the EEMUA 191 guide and have been consulted in the production of this second edition. Since the publication of the first edition in 1999, our members have carried out assessments of the performance of alarm systems actually achieved in practice. The results of these assessments have indicated that a metrics-focused continuous improvement program is the best way to address the EEMUA alarm system performance recommendations. The ASM Consortium is pleased to see that such an approach is incorporated and explained in the current second edition of the EEMUA 191 guide.

EEMUA 191 has achieved widespread recognition in the USA, Europe, Asia and elsewhere since it was released eight years ago. In the opinion of the ASM Consortium, this guide continues to currently represent the best publicly-available benchmark of accepted industry good practice for alarm system design, management and procurement.

Kevin R. D. Harris
Director, Abnormal Situation Management Consortium
Phoenix
Arizona
USA
March 2007

ASM

Contents

List of Figures

List of Tables

About This Guide

Aim

The aim of this Guide, EEMUA 191, is to assist in the design, development, procurement, operation, maintenance and management of industrial alarm systems. Following the guidance in EEMUA 191 should result in better alarm systems that are more usable and that result in safer and more cost-effective industrial operations.

This Guide has been produced by the *users* of alarm systems in industry. It is based on what some leading companies are doing, but it is also intended to be challenging and to promote continuous improvement in alarm handling practices from every starting level. It is intended to help both in improving existing systems and in developing new facilities during plant construction or during alarm system refurbishments.

Scope

This Guide is primarily concerned with alarm systems provided for people operating industrial processes. These include alarm systems in industries such as chemical manufacture, power generation, oil and gas extraction and refining and others. The Guide has been compiled from the experience gained within such industries, however much of the guidance is generic and with appropriate interpretation can be applied in other sectors. For example, the guide has been used successfully as a basis for training in the rail and transport sectors, in the nuclear industry, and elsewhere.

Guidance is included on:

- the design of alarm processing systems and their functionality;
- the operation of existing alarm systems and performance optimisation;
- the specification and purchase of new alarm systems.

Alarm system functionality is now commonly included in computer-based systems used to control and monitor industrial processes, as well as being found in stand-alone dedicated computer-based alarm systems and the more basic annunciator-based or hard-wired 'light box' systems. This Guide is appropriate to alarm systems implemented by any of these methods, wherever a human response may be needed to an alarm condition. The guidance is designed to be practical and usable. There is overview information, and also much detailed guidance on specific issues.

Why this Guide is Important

Alarm systems form an essential part of the operator interfaces to large modern industrial facilities. They provide vital support to the operators by warning them of situations that need their attention.

Alarm systems thus have an important role in preventing, controlling and mitigating the effects of abnormal situations. If these alarm systems do not work well, the effects can be very serious. The explosion and fires at the Milford Haven Refinery in the UK (23) – which caused £48 million of plant damage plus major production loss - could have been prevented by the operating staff. They failed to do this partly because they faced a continuous barrage of alarms for the whole five hour period leading up to the accident. Alarm overload during an abnormal situation is a common problem: it usually means the alarm system is at its least effective when it is needed most. In the England-France Channel Tunnel fire (23) - which resulted in damage and revenue losses totalling around £200 million - problems with alarms led to slow recognition of the fire and inhibited the effective management of the incident.

In neither of these accidents were people seriously injured or killed. However, there have been serious industrial accidents of a similar nature which have resulted in deaths. There have also been prosecutions and breaches of environmental limits. Whilst the degree to which alarm systems have contributed to these accidents is difficult to establish, alarm information has often been material. Alarm system performance is still being cited in continuing major incidents as a contributory factor (for instance, see the Australian Longford refinery incident (28)).

Alarm system shortcomings cause a multitude of smaller avoidable incidents which increase the risks to people and increase operating costs. A survey carried out for the GB Health & Safety Executive in the chemical and power industry (5) identified many problems with alarm systems and many cases where inadequacies in alarm system performance had led to financial loss or to equipment or environmental damage. Such incidents are often difficult to identify and quantify precisely. However, the total costs incurred in all such incidents are large and significant savings could be made by eliminating them.

Revised 2nd Edition 2007

Thousands of copies of this EEMUA 191 Guide have been sold since it was first launched in 1999. Purchasers have included process operating companies, suppliers, engineering contractors, training establishments, regulatory agencies, inspection bodies and individual engineers.

The principles and guidance in EEMUA 191 have been applied worldwide and for many it has become the *de facto* code of good and best practice. Significant improvements are being seen by those companies investing appropriate time and resources to ensure that their alarm systems meet the requirements of good operational ownership.

From 2007, this revised 2nd Edition has been published. It builds on the original 1st Edition, by taking note of feedback from the large body of users of this Guide, updating items as appropriate, especially where knowledge has moved forward.

Any reader is able to make comments or provide suggestions on the current Edition and a feedback form is provided towards the end of this publication for that purpose.

Key Changes in 2nd Edition

The following sections are new or have been significantly revised:

- 2.3.3 Safety Related Alarms;
- 3.6 Testing of Alarms;
- 4 Measuring Performance;
- A1 Glossary;
- A4.1.1 Identification of Risks;
- A5.3 Priority Distribution of Alarms;
- A5.5.3 Method 3;

- A5.6 Record Keeping;
- A7 Alerts;
- A13 Performance Levels;
- A20 Batch;
- A21 Alarm System Improvement.

Acknowledgements

The 1st Edition of this Guide was prepared by an EEMUA sub-committee comprising representatives of EEMUA member companies and the GB Health & Safety Executive (HSE). Funding for the Guide was provided by EEMUA member companies and by the Abnormal Situation Management Consortium[3] (ASM®). A large number of comments on the draft Guide were received from EEMUA Member Companies, ASM Consortium members and elsewhere. The development was assisted by the late Matthew Bransby. Tragically Matthew died in a climbing accident in 2000.

This 2nd edition has been prepared by a working group of EEMUA's Instrumentation & Control Committee, with input from (in Europe) the HSE and (in America) the ASM Consortium, as well as many comments received from readers around the world.

EEMUA gratefully acknowledges the contributions of all who have been involved, in many capacities, during the preparation and production of this Guide.

Notes on use of this Guide

In order to assist readers, the next two pages highlight the core principles underlying the guidance and give a 'roadmap' providing directions to key information.

Sections 1 to 6 of this Guide concentrate on alarm management philosophy and general principles. Key messages are highlighted in **bold**.

Much supporting detail is contained in the Appendices A1 to A23.

A glossary of terms is given in Appendix 1.

References are listed in Appendix 22. Throughout the text, they are referred to in parentheses, for example (1).

Footnotes are referred to using superscripts, for example[1].

[3] ASM and Abnormal Situation Management are US registered trademarks of Honeywell International, Inc.

Core Principles and the Roadmap

This Guide begins with four core principles and a roadmap. The core principles encapsulate the most important ideas presented in the Guide. The roadmap directs different users to the parts of the Guide relevant to their particular needs. It also provides a quick and easy point of reference and an overview of the Guide's advice.

Four core principles run through this Guide:

First - Usability. Alarm systems should be designed to meet user needs and operate within the operator's capabilities. This means that the information alarm systems present should:

- be relevant to the operator's role at the time;
- indicate clearly what response is required;
- be presented at a rate that the operator can deal with;
- be easy to understand.

Second - Safety. The contribution of the alarm system to protecting the safety of people, the environment and the plant equipment should be clearly identified. Any claims made for operator action in response to alarms should be based upon sound human performance data and principles.

Third - Performance monitoring. The performance of the alarm system should be assessed during design and commissioning to ensure that it is usable and effective under all operating conditions. Regular auditing should be continued throughout plant life to confirm that good performance is maintained. This requires a real and continuing commitment from the senior management of the plant.

Fourth - Investment in engineering. Alarm systems should be engineered to suitably high standards. When new alarm systems are developed (or existing systems are modified), the design should follow a structured methodology in which every alarm is justified and properly engineered. The initial investment in system design should be sufficient to avoid the operational problems and the safety, environmental and financial risks that often arise and which result in overall higher lifetime costs. Contract strategies should be chosen to ensure that alarm systems are engineered to good standards.

Roadmap

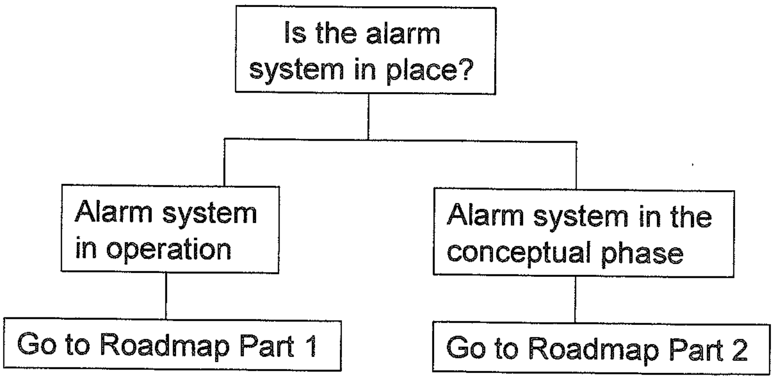
Is the alarm system in place?
Alarm system in operation
Alarm system in the conceptual phase
Go to Roadmap Part 1
Go to Roadmap Part 2

Roadmap Part 1

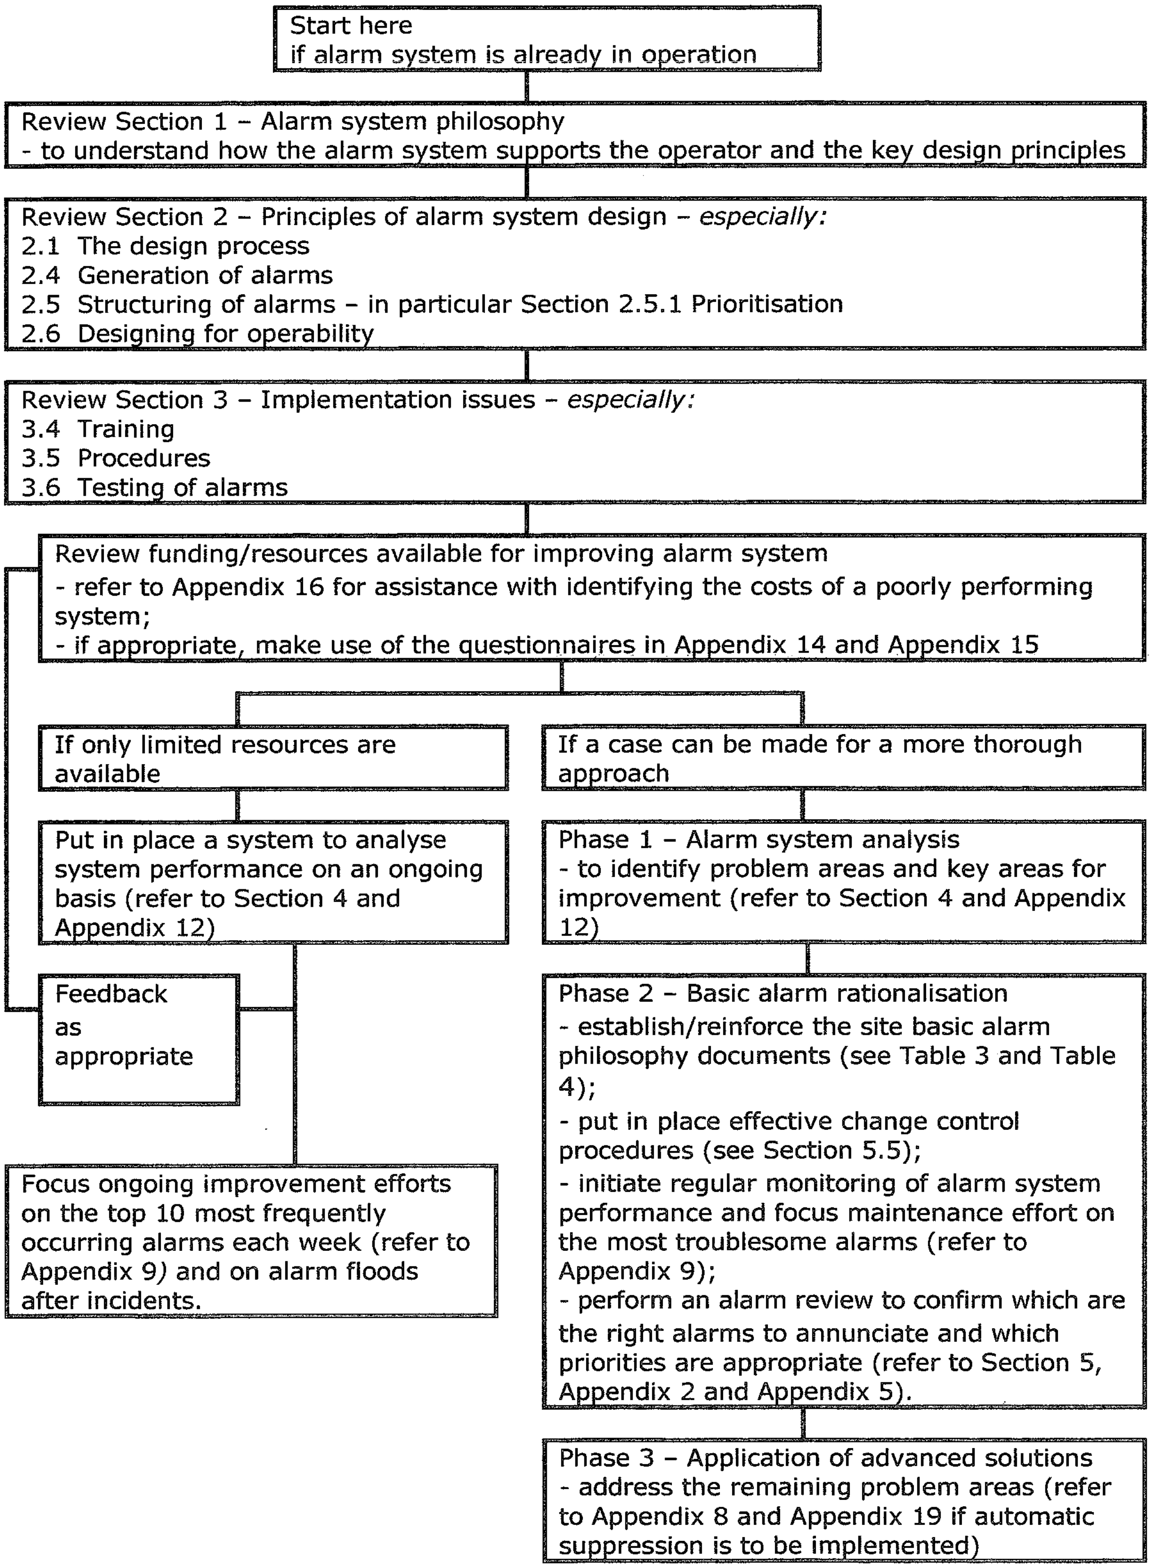

Start here
if alarm system is already in operation

Review Section 1 – Alarm system philosophy
- to understand how the alarm system supports the operator and the key design principles

Review Section 2 – Principles of alarm system design – *especially:*
2.1 The design process
2.4 Generation of alarms
2.5 Structuring of alarms – in particular Section 2.5.1 Prioritisation
2.6 Designing for operability

Review Section 3 – Implementation issues – *especially:*
3.4 Training
3.5 Procedures
3.6 Testing of alarms

Review funding/resources available for improving alarm system
- refer to Appendix 16 for assistance with identifying the costs of a poorly performing system;
- if appropriate, make use of the questionnaires in Appendix 14 and Appendix 15

If only limited resources are available

Put in place a system to analyse system performance on an ongoing basis (refer to Section 4 and Appendix 12)

Feedback as appropriate

Focus ongoing improvement efforts on the top 10 most frequently occurring alarms each week (refer to Appendix 9*)* and on alarm floods after incidents.

If a case can be made for a more thorough approach

Phase 1 – Alarm system analysis
- to identify problem areas and key areas for improvement (refer to Section 4 and Appendix 12)

Phase 2 – Basic alarm rationalisation
- establish/reinforce the site basic alarm philosophy documents (see Table 3 and Table 4);
- put in place effective change control procedures (see Section 5.5);
- initiate regular monitoring of alarm system performance and focus maintenance effort on the most troublesome alarms (refer to Appendix 9);
- perform an alarm review to confirm which are the right alarms to annunciate and which priorities are appropriate (refer to Section 5, Appendix 2 and Appendix 5).

Phase 3 – Application of advanced solutions
- address the remaining problem areas (refer to Appendix 8 and Appendix 19 if automatic suppression is to be implemented)

Roadmap Part 2

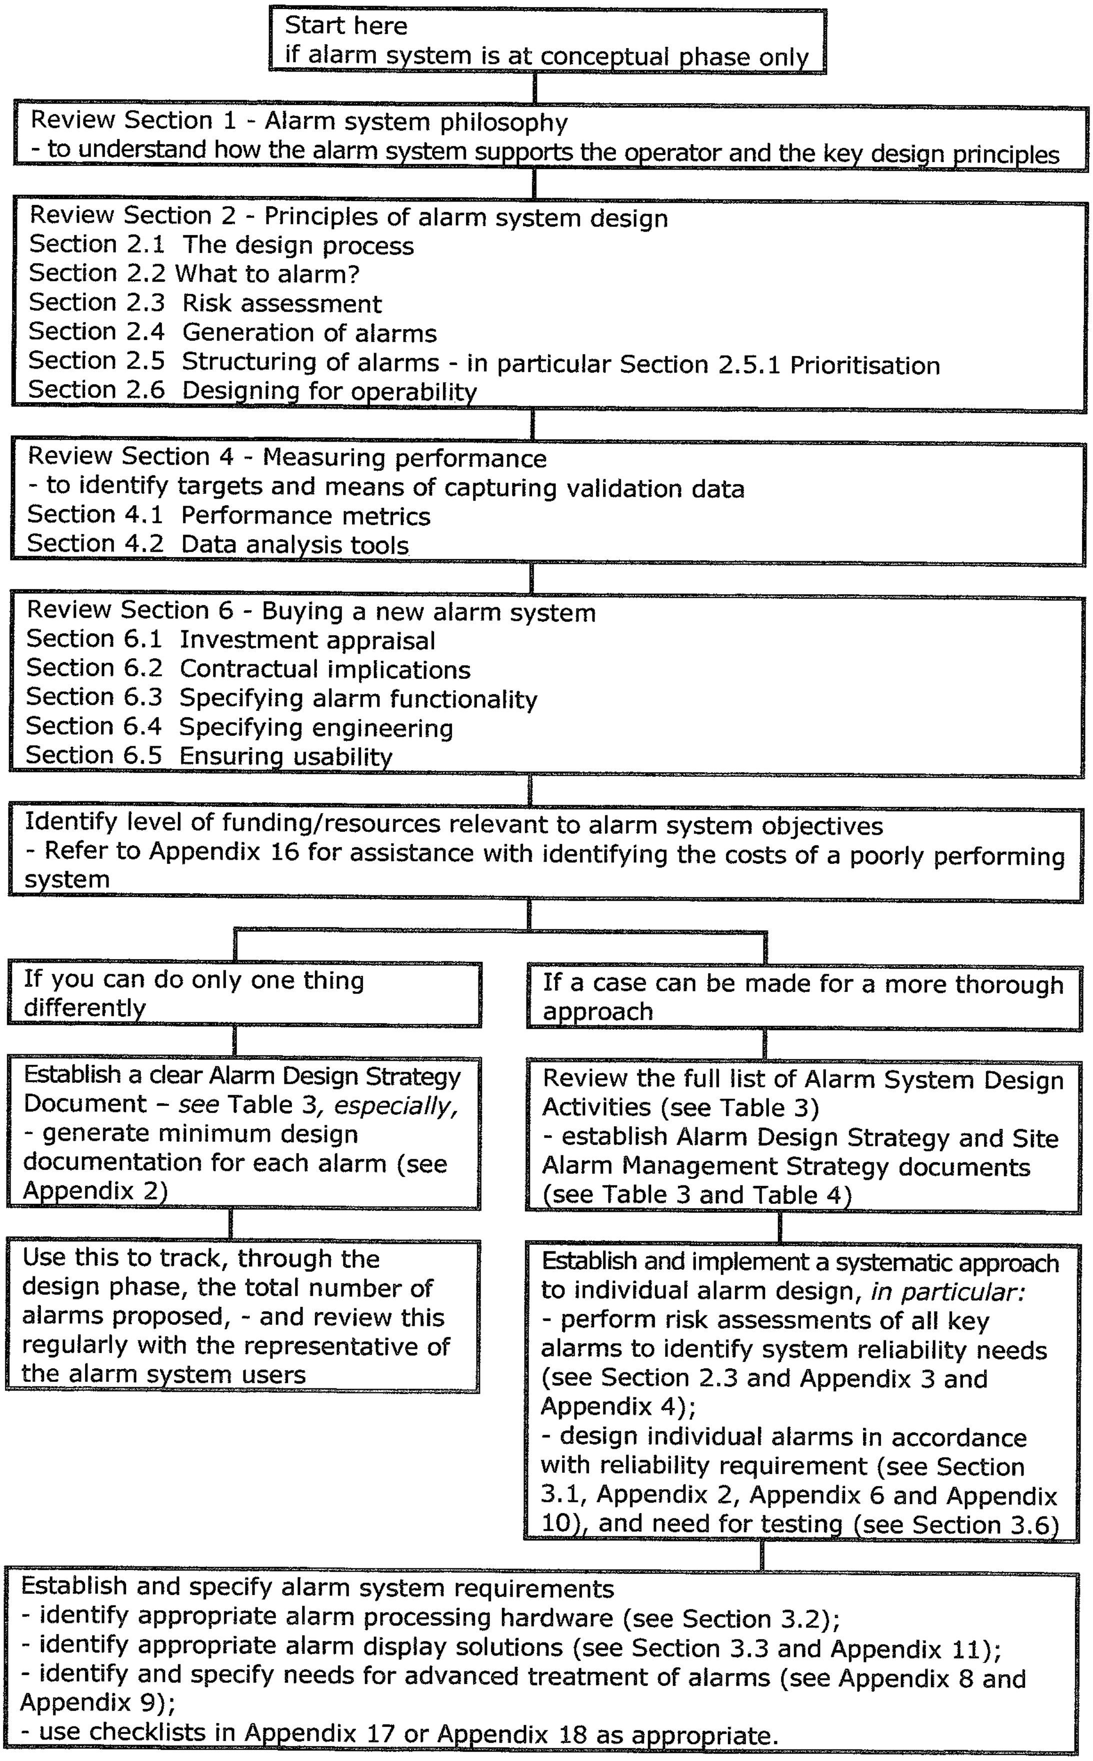

Start here
if alarm system is at conceptual phase only

Review Section 1 - Alarm system philosophy
- to understand how the alarm system supports the operator and the key design principles

Review Section 2 - Principles of alarm system design
Section 2.1 The design process
Section 2.2 What to alarm?
Section 2.3 Risk assessment
Section 2.4 Generation of alarms
Section 2.5 Structuring of alarms - in particular Section 2.5.1 Prioritisation
Section 2.6 Designing for operability

Review Section 4 - Measuring performance
- to identify targets and means of capturing validation data
Section 4.1 Performance metrics
Section 4.2 Data analysis tools

Review Section 6 - Buying a new alarm system
Section 6.1 Investment appraisal
Section 6.2 Contractual implications
Section 6.3 Specifying alarm functionality
Section 6.4 Specifying engineering
Section 6.5 Ensuring usability

Identify level of funding/resources relevant to alarm system objectives
- Refer to Appendix 16 for assistance with identifying the costs of a poorly performing system

If you can do only one thing differently

Establish a clear Alarm Design Strategy Document – *see* Table 3, *especially*,
- generate minimum design documentation for each alarm (see Appendix 2)

Use this to track, through the design phase, the total number of alarms proposed, - and review this regularly with the representative of the alarm system users

If a case can be made for a more thorough approach

Review the full list of Alarm System Design Activities (see Table 3)
- establish Alarm Design Strategy and Site Alarm Management Strategy documents (see Table 3 and Table 4)

Establish and implement a systematic approach to individual alarm design, *in particular:*
- perform risk assessments of all key alarms to identify system reliability needs (see Section 2.3 and Appendix 3 and Appendix 4);
- design individual alarms in accordance with reliability requirement (see Section 3.1, Appendix 2, Appendix 6 and Appendix 10), and need for testing (see Section 3.6)

Establish and specify alarm system requirements
- identify appropriate alarm processing hardware (see Section 3.2);
- identify appropriate alarm display solutions (see Section 3.3 and Appendix 11);
- identify and specify needs for advanced treatment of alarms (see Appendix 8 and Appendix 9);
- use checklists in Appendix 17 or Appendix 18 as appropriate.

1. Alarm System Philosophy

Section overview

This section provides an introduction to alarm systems and comprises:

- a basic explanation of what an alarm system is;
- a discussion of the role of the operator, how this changes according to operating state, and what support the operator needs in the different states;
- a more detailed discussion of the function of an alarm system and of some of the characteristics that a good alarm system should have. Key principles identified are that:
- every alarm should have a defined response;
- adequate time should be allowed for the operator to carry out this response.

These are important principles that underlie much of the thinking in this Guide.

1.1 What is an Alarm System

Alarm systems form a core element of almost all modern operator interfaces to industrial plants including oil refineries, power stations, chemical plants and many others. Traditionally they were based around hard-wired indication lamps and annunciator panels, but more modern systems use computer-driven display devices (e.g. VDUs) to present the operator with graphical representations or text lists of alarms.

Alarms are signals which are annunciated to the operator[4], typically by an audible sound, some form of visual indication, usually flashing, and by the presentation of a message or some other identifier. An alarm will indicate a problem requiring operator attention, and is generally initiated by a process measurement passing a defined alarm setting as it approaches an undesirable or potentially unsafe value.

Alarm systems are a very important way of automatically monitoring the plant condition and attracting the attention of the process plant operator to significant changes that require assessment or action. They help the operator:

- to maintain the plant within a safe operating envelope. A good alarm system helps the operator to correct potentially dangerous situations before the Emergency Shutdown (ESD) system[5] is forced to intervene. This improves plant availability. It also reduces the demand rate on the ESD system and, thus, increases plant safety[6];
- to recognise and act to avoid hazardous situations. It is the role of the ESD system to intervene before a hazard arises. However, there may be cases, hopefully extremely infrequent, where the plant deviates outside its design conditions into a state where the ESD system is no longer capable of acting effectively. Also there may be cases where operator action following an alarm

[4] Note that throughout this Guide the term 'operator' should be taken to include the desk operator plus other primary users (e.g. supervisors) assisting in the control of the plant in the control room.
[5] See Glossary: Appendix 1.
[6] For automatic protection systems the dangerous failure rate equals the demand rate times the average probability of failure on demand. Reducing the demand rate thus reduces the dangerous failure rate.

has been explicitly identified within the plant Safety Case as a measure of protection;
- to identify deviations from desired operating conditions that could lead to financial loss such as off-specification or overly expensive product;
- to better understand complex process conditions. Alarms should be an important diagnostic tool and are one of several sources that an operator uses during an upset.

In this Guide the term 'alarm system' refers to the complete system for generating and handling alarms including field equipment, signal conditioning and transmission, alarm processing and alarm display. It also includes hardware, software and supporting information (e.g. alarm response procedures, management controls). The term 'alarm processor' refers to the part of the system for processing and displaying alarms. Often this function will be carried out within proprietary electronic annunciator boxes, DCSs or SCADA systems.

1.2 The Role of the Operator

The role of an operator on an industrial plant generally encompasses a range of different activities including plant operation, optimisation of production, fault identification, co-ordination of maintenance, etc. The tasks involved change depending on plant state, e.g. whether it is in normal operation, upset operation, emergency shut down, planned shut down, start up, or operating mode change. Some of these changes are illustrated in Figure 1.

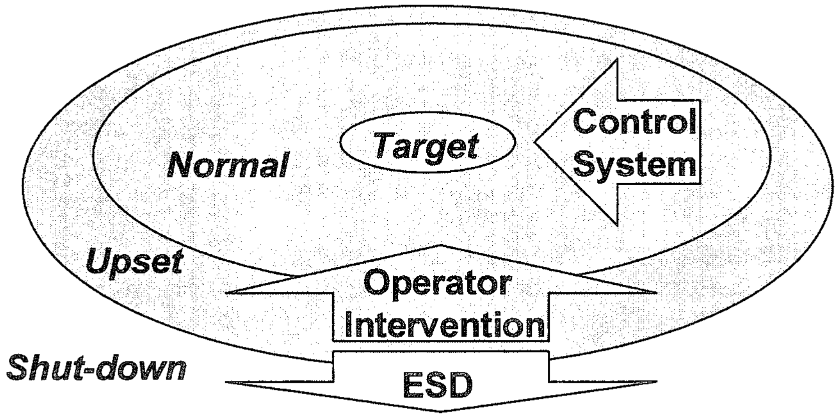

Plant state	Operator's primary role	Key alarm information
Normal	monitoring & optimisation	minor operating adjustments needed
Upset	situation management	operator intervention needed
Shut-down	ensure safe shut down	safety actions needed

Figure 1 The operator's role in the different plant states

Most of the complex systems covered by this Guide do not operate steadily but are subject to constant disturbances (e.g. changes in fuel quality, changes in ambient temperature). Under normal operation the automatic control systems will, typically, act to mitigate these disturbances to keep the plant close to target operating conditions. The operator's primary role is to monitor operation and make fine adjustments, e.g. to control set points or to plant equipment under manual control. Alarms may be provided to draw the operator's attention to the need for adjustments. However, they should be very carefully designed to ensure that they do not become a nuisance in other operating states.

If there are significant disturbances (e.g. mechanical failure of equipment, feedstock change), they may push the plant into an 'upset' state from which the control system is not able to effect a recovery without operator intervention. Alarms should be provided to annunciate this need for operator intervention/action.

If the upset state is not corrected satisfactorily by the operator and the plant condition approaches a state where damage or danger is likely to occur, the ESD system, where provided, should intervene and shut down the plant area affected. The operator role is to check that the automatic shut down takes place safely and take complementary action to minimise the size of upset. If the plant does not shut down safely, or where no automatic shut down system is provided, the operator should take action to bring the plant into a safe state. Alarms should be provided to inform of shut down system failures, or of other unsafe situations requiring operator action.

It should be noted that, in practice, the role of the operator during abnormal situations can be very complex. As shown in Figure 2, the operator response may involve several quite different types of task. Furthermore, the operator response required to one abnormal situation may be quite different from that required to an apparently similar situation at another time.

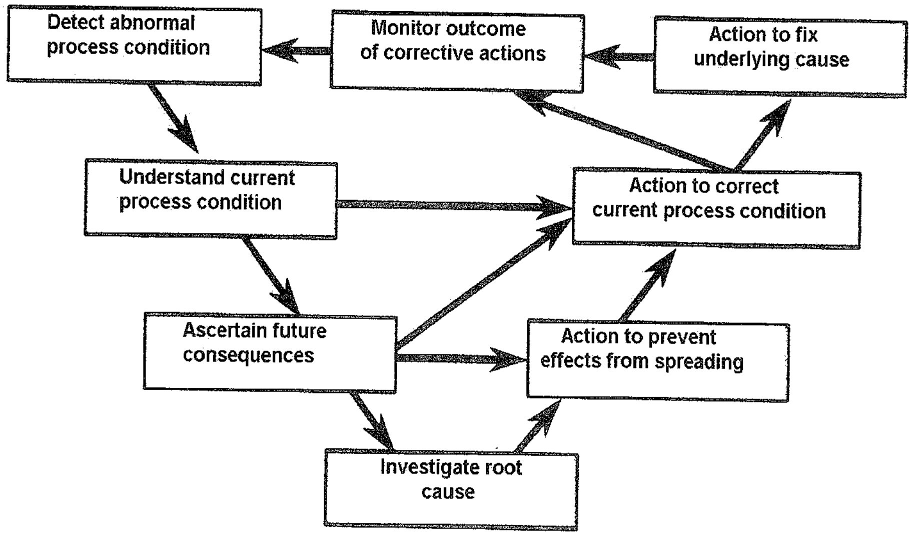
Detect abnormal process condition
Monitor outcome of corrective actions
Action to fix underlying cause
Understand current process condition
Action to correct current process condition
Ascertain future consequences
Action to prevent effects from spreading
Investigate root cause

Figure 2 Operator response to an abnormal situation

The correction of an abnormal situation often requires a number of separate tasks to be carried out, some of them in parallel. Due to the nature of the process, which may have long time delays and slow dynamics, each task may involve a series of short sub-tasks separated by waits of minutes or hours before results of the actions can be seen and the operator can decide what sub-task to perform next. Thus, to efficiently correct an abnormal situation the operator will often have to work under time pressure and stress to string together a series of unrelated sub-tasks. It is clear that whilst alarms are a useful tool to help the operator, the need for more general task management during an abnormal situation should be addressed.

Note that the need for different types of alarm in different plant operational states implies that the alarm system may need to be '**context sensitive**' and take account of operational state in the determination of what signals should be alarms. Some signals that would be alarms in normal operation may not be relevant in other operational states. This implies the need for careful analysis of what signals should be alarms, and in some cases, logical processing to control the presentation of alarms (see Section 2.5.2 and Appendix 8). The types of operational state to be considered include start up/shut down, normal, abnormal/upset, emergency and maintenance (see EEMUA 201 (20) and Appendix 8, Section A8.3.4).

It is important that the role of the operator in maintaining safety and plant integrity should be clearly identified for all operational states. Even on highly automated plants with extensive automatic protection systems there almost certainly will be potential fault scenarios that require operator intervention. These scenarios should be identified and it should be determined whether and how the alarm system will support the operator in carrying out his corrective action.

The role of other people such as engineering and maintenance staff is recognised in this Guide. An alarm system may provide them with important information, but this is often mediated through the plant operator. If so, this should be explicitly stated in the definition of the operator's tasks. By comparison with the operator, the activities of these other people tend to be less plant state or time dependent.

1.3 Key Design Principles

The primary function of an alarm system can be defined as:

> **The purpose of an alarm system is to direct the operator's attention towards plant conditions requiring timely assessment or action.**

Thus an alarm system is saying to the operator "do something about this urgently", "look at this soon", "don't forget this problem", etc. Thus, it should help the operator to manage tasks and resources, and should focus attention on the most important issues.

To achieve the above:

> **Each alarm should alert, inform and guide.**

Thus, when an alarm occurs, the operator should be alerted, which, typically, is done by an audible warning and some form of flashing indication. The alarm should present information defining what the problem is, e.g. by displaying a descriptive text message, lighting a labelled annunciator window, or changing a graphic display. Information should also be provided (e.g. written procedures or operator-selected on-screen help) giving guidance on how to respond to the alarm. Ideally the alarm system should also provide feedback to the operator on the success of actions taken in response to alarms[7].

For an alarm system to be effective in supporting the operator, every alarm presented to the operator should be a help rather than a hindrance. The objective should be to avoid the operator wasting time on deciding whether the

[7] This information will generally have to be provided on supplementary displays, for example, accessed by clicking on an alarm on an alarm list (see Appendix 11).

alarm can be ignored and ensure that the operator does not adopt a mind frame that the alarms can be ignored[8]. Thus:

> **Every alarm presented to the operator should be useful and relevant to the operator[9].**

One way in which the designer can ensure that every alarm is useful is to require that:

> **Every alarm should have a defined response.**

Generally, this response should be an action (e.g. altering a control set point, changing over to a standby pump, asking a fitter to repair a piece of failed plant equipment). Sometimes the response to the alarm will have to be conditional. For example, the operator might select a graphic display, check the plant condition, and only in certain circumstances carry out any control action. In a few cases the defined response to an alarm will be purely mental. For example, in response to a 'plant tripped' or 'start up sequence complete' alarm the operator may need only to change the form of plant monitoring the operator is carrying out. There may not be any immediate control action required, but it is important for the operator to make this cognitive switch.

The key point is that every alarm (or combination of alarms) should have some response which should have been clearly defined by the designer of that alarm. If a response cannot be defined, then the signal should not be an alarm. A common problem is that such event information often gets mixed in with alarms[10].

If a response is defined for every alarm during the design process, this information provides an excellent starting point for the production of alarm response procedures (see Section 3.5) and training material.

Given that the operator is expected to respond to every alarm, it follows that in a usable alarm system:

> **Adequate time should be allowed for the operator to carry out a defined response.**

This implies that:

- the alarm should occur early enough to allow the operator to correct the fault (see Section 2.4.2);
- the alarm rate should not exceed that which the operator is capable of handling. A typical operator's role involves many different activities and responsibilities. It is important that the overall role is manageable, and the

[8] There is experimental evidence (2) which shows that there is a 'cry wolf' effect and operators do adapt their response according to the probability of the alarm being useful.

[9] This implies that a change of plant state relevant to the maintenance staff but not relevant to the operator should not be an 'alarm'. Appropriate means should be provided for presenting this information to these staff, and facilities similar to those provided to the operator for presenting alarms may be appropriate. Note, however, that the operator may need to be made aware through alarms of the operational consequences of faults requiring maintenance.

[10] There will be events that occur on plant which are not alarms - such as changes in mode of control loops - which the operator will need to be able to observe. The designer should consider how these are presented. It may be sufficient to show them as status indications on plant mimic formats. Alternatively, an event list display which is similar to but separate from the alarm list display may be appropriate.

plant, the control systems and the operator interface should be designed to achieve this. Furthermore, the time required for other activities often imposes severe limits on what alarm handling workload is acceptable[11]. As will be seen later in Section 2.6, this has some very important implications in terms of overall alarm system design.

To ensure that the above objective is achieved:

The alarm system should be explicitly designed to take account of human limitations.

Some of the characteristics that an alarm should have are summarised in Table 1

CHARACTERISTICS OF A GOOD ALARM

• Relevant	i.e. not spurious or of low operational value
• Unique	i.e. not duplicating another alarm
• Timely	i.e. not long before any response is needed or too late to do anything
• Prioritised	i.e. indicating the importance that the operator deals with the problem
• Understandable	i.e. having a message which is clear and easy to understand
• Diagnostic	i.e. identifying the problem that has occurred
• Advisory	i.e. indicative of the action to be taken
• Focusing	i.e. drawing attention to the most important issues

Table 1 Characteristics of a good alarm

So far this section has discussed the primary alarm system function of providing warnings to the operator. An alarm system may also have a secondary function of providing an alarm log which can be used for optimising plant operation, for analysis of incidents[12], and for improving the performance of the alarm system (see Section 5). Manufacturers often provide this as part of a broader capability for logging all significant plant events including alarms, plant status changes not included in the alarm list, operator actions, etc. Where this is done it may be necessary to be able to extract alarm logs separate from other events.

[11] The long term average workload, W, that the alarm system imposes on the operator may be expressed by the equation:

$$W = R * T$$

where:

R is the average rate at which alarms are presented to the operator, and
T is the average time taken by the operator to respond to the alarm.

For example, if alarms are presented at an average rate of 1 per minute, and the operator takes on average 30 seconds to deal with each alarm, then the alarm system would on average be consuming 50% of the operator's time.

[12] Note that high speed logging of some parameters may be necessary to analyse certain sorts of incidents. Often separate post-incident recorders are used to do this.

2. Principles of Alarm System Design

Section overview

This section provides an introduction to the process of designing an alarm system. It identifies the stages in the process and lays down the principle that the design process should be clearly defined and formally documented. Furthermore, the design of individual alarms should follow a structured procedure.

A fundamental issue when designing each alarm is to consider how important it is and how reliable it should be. To do this it is necessary to go through some form of qualitative or quantitative risk assessment. Important considerations are:

- *whether the alarm should be classified as safety related according to the definitions given in the international standard IEC 61508 (29);*
- *whether it should be implemented in a stand-alone system independent of the process control system.*

The decision whether an alarm is safety related will be influenced by national legislation and by existing practices within the industry sector. Alarms which are safety related should be given special consideration in terms of the design of the operator interface and the operator support provided.

Risk assessment only provides a starting point in the design process. This section goes on to consider broader design issues. This includes how alarm settings are chosen, how priorities should be defined and set, and how alarms should be processed to make them as meaningful as possible. There is also consideration of how an overall alarm system should be designed for operability.

2.1 The Design Process

The design of an effective operator interface for a modern industrial plant is generally a complex task which should involve consideration of a wide range of human factors issues. It is too broad a topic to be covered in depth in this Guide. Some pointers to the issues involved may be found in (13), (19) and (43) and in Section 6.5.

For safe and efficient plant operation it is necessary that the operator's role is designed to be one that the operator can perform effectively and that the operator is provided with suitable tools to support him in doing this. An alarm system is only one of these support tools and, however good it is, the operator will not be able to run the plant safely and efficiently if an operator is being required to carry out a task that is beyond their capabilities, or if the other parts of the operator interface are inadequate. Thus the alarm system should be designed within a framework of considering the complete operator role and support system.

Because of these factors, the design of an alarm system is a complex task. As indicated in Table 2, it may involve a large number of different design activities.

ALARM SYSTEM DESIGN ACTIVITIES

RISK ASSESSMENT
Development of the plant safety case
Identification of the safety role of the operator
Risk assessment to identify alarms to protect against safety, environmental or economic risks
Review to identify alarms providing significant risk reduction (e.g. safety related alarms)

ERGONOMICS
Identification of number of operators and their roles
Overall design of operator interface (e.g. numbers of screens, colour usage, information aids)
Design of the alarm interface (e.g. display methods, annunciation)

DESIGN OF INDIVIDUAL ALARMS
Review of proposed alarms not deriving from risk assessments
Identification of alarms with special integrity or display requirements
Completion of data form for each alarm (see Appendix 2)
Production of alarm response procedures
Design of plant alarm sensors
Design of hardware for conditioning individual alarm signals
Installation of plant alarm sensors and signal conditioning

DESIGN INTEGRATION
Rationalisation of lists of proposed alarms
Review of overall system design to meet Key Design Principles (see Section 1.3)
Identification of required alarm processing functionality

ALARM SYSTEM CONFIGURATION
Installation of alarm system hardware
Configuration of alarm system hardware/software facilities
Construction of alarm information database
Configuration of hardware/software for individual alarms
Configuration of alarm combinatorial logic

TESTING AND COMMISSIONING
Testing of alarm system facilities
Testing of alarm sensors and signal conditioning
Testing of individual alarm hardware/software configurations
Evaluation of overall ergonomic acceptability
Measurement of alarm system performance
Determination of ongoing testing needs
Optimisation of operational performance

Table 2 Activities in the design of an alarm system

The performance of alarm systems can have a significant safety, environmental and economic impact (5), (14), (23) and (36). Consequently, for all plants where the alarm system is to play a part in preventing or mitigating significant hazards or economic losses, it is important that good design practices are followed in all the above design activities. In order to ensure that good practice is established and sustained, it is recommended that an alarm system design strategy should be defined by a suitably experienced multi-disciplined team and formally documented. Topics that this should cover are given in Table 3.

ALARM DESIGN STRATEGY DOCUMENT

- allocation of roles and responsibilities for design of the alarm system - including what user involvement there is to be
- identification of the alarm system users and their needs
- a definition of what an alarm should be
- a definition of the safety role of the alarm system
- a list of any alarms claimed to contribute to safety cases
- definitions of alarm system performance targets (e.g. maximum rates)
- rules for prioritisation of alarms
- checklist for designers on the information to be recorded for each alarm (see Appendix 2)
- dictionary of terms and abbreviations to be used in alarm messages
- guidance to sub-contractors on the design of alarms (where appropriate)
- guidance on content and structure of alarm response procedures
- guidance on interpreting patterns of alarms, and their grouping, masking and acceptance (where appropriate)
- calculations of alarm equipment test frequencies

Table 3 Content of alarm design strategy document

Before the plant becomes operational a further document should be developed defining the site alarm system management strategy. Topics that this should cover are given in Table 4. Together these two strategy documents should help to ensure that valuable design and maintenance information is maintained throughout the plant lifetime.

SITE ALARM MANAGEMENT STRATEGY

- allocation of roles and responsibilities for maintaining and managing the alarm system
- definition of alarm review procedures
- performance targets for confirming compliance with alarm design strategy (e.g. average alarm rate, number of alarms following incidents)
- routine measurements to be taken of alarm performance
- requirements for logging alarms and for storage of logs
- maintenance and test procedures
- lists of documentation relating to the alarm system that the site should maintain
- description of modification procedures to be followed when changes are made to alarms, when new alarms are introduced, or when documentation is changed
- training and competence

Table 4 Content of site alarm management strategy document

The design of the individual alarms and their configuration in the alarm processor is a complex multi-stage process. To achieve good performance, good design practices should be consistently applied to all alarms. To encourage this:

> **The design of each alarm should follow a systematic structured procedure in which design decisions are documented.**

Appendix 2 provides a list of questions that may be worked through for each alarm to help structure the design process. The answers should be recorded as they provide a useful part of the plant maintenance documentation. In addition:

> **The overall alarm system design should accord with the Key Design Principles given in Section 1.3.**

This will require a review of all proposed alarms, a check that the Key Design Principles are adhered to, and an assessment of whether desired system performance is likely to be met. In practice this is a very iterative process.

2.2 What to Alarm

Proposals for alarms come from a large variety of sources, e.g.:

- custom and practice ("We always specify one of these", "There's one on our existing plant");
- equipment manufacturer's standard provision;
- fault reports from the alarm processing hardware/software;
- requirements of safety authorities, insurers, headquarters departments, etc.;
- designer - especially software designer - whim ("I thought it would be good idea", "It is easy and it is cheap");
- as an alternative to automation ("It was cheaper than fitting a control system");
- operating experience/feedback;
- simulator studies;
- task analysis[13];
- safety reviews, e.g. HAZOPs;
- qualitative or quantitative risk analysis.

The majority of these are informally structured design processes and experience shows that they lead to many poorly justified alarms being proposed. Ideally the whole design processes should be integrated and formally structured, but this may not be completely achievable. However, what is practical is to ensure that every proposed alarm is properly scrutinised. Thus:

> **Whatever its source, every alarm should be justified, properly engineered and be consistent with the overall alarm philosophy and plant risk assessment.**

2.3 Risk Assessment

2.3.1 The Need to Identify and Minimise Risks

Alarms are provided to reduce the likelihood of sub-optimal plant operation or of plant damage, as these may lead to injury to people, environmental damage and/ or economic loss. It follows that the design of the alarm system should involve a

[13] Task analysis (24) is a technique for formally breaking each of the operator's tasks down into a hierarchy of mental and physical actions and identifying the tools and information that the operator needs to perform these. Task analysis can also assist in deciding what tasks should be automated.

consideration of the risks[14] of injury to people, damage to the environment and economic loss, and a decision about which risks the alarm is intended to reduce.

In the European Union, national legislation implementing the Framework Directive 86/391[15] requires employers to identify any hazards to people associated with their work activities and to assess the risks from these hazards. Many of the industrial systems covered by this Guide have potential for causing major accidents possibly resulting in multiple deaths, and for these the assessment of risk will be particularly important.

Having identified potential hazards to people, every employer is also legally obliged to take steps to reduce the risks from these hazards to a tolerable level. In addition, a record of the risk assessment has to be kept, and it has to be reviewed as appropriate.

Risk assessment is also a useful (and sometimes necessary) tool for identifying risks of environmental damage and economic loss, and reducing these to a tolerable level. Appendix 3 discusses the subject of quantitative and qualitative risk assessment in more detail and Appendix 4 provides examples.

2.3.2 Risk Reduction

Alarms are one of many tools which may be used to reduce risks to a tolerable level. Risk reduction should start with the initial plant design by selecting processes and plant configurations which have appropriate inherent safety. For example, when designing a chemical plant, a production process might be chosen which avoided generation of hazardous intermediate products. The risk reduction should follow through into the detailed design of the plant where steps should be taken to eliminate hazards, protect against them, reduce their frequency of occurrence and minimise their consequences.

Alarms are a form of protection against risks, i.e. they sense conditions indicating increases of safety, environmental or economic risk and should result in corrective action being taken. Where they are used, they are often only part of one of several layers of protection against the specified risks. For example, there might also be:

- automatic control systems which act to maintain process variables within safe bounds;
- operator information displays, manual controls and manual trips;
- automatic protection systems;
- mechanical protection, e.g. safety relief valves.

Together these systems should reduce the specified risk to a tolerable level.

[14] When applied to safety issues, risk is a measure of the probable rate of occurrence of a hazard and its severity. For example, if a hazard occurred once per 100 years and involved a 1 in 10 chance of injuring a person, the risk would be 10^{-3} injuries per year. Risk also has a broader everyday meaning, and can be applied in a similar way to environmental and economic losses. Economic risks would be measured in money lost per year.

[15] 'The Council Directive of 12 June 1989 on the introduction of measures to encourage improvements in the safety and health of workers at work' (89/391/EEC) is usually referred to as the Framework Directive. In the UK the Health & Safety at Work Act places general duties on employers, employees, etc. and enables supporting legislation to be produced as needed. The Management of Health and Safety at Work Regulations 1992, which is one of these sets of supporting legislation, was produced to implement the Framework Directive. Regulation 3 specifically requires employers to "make suitable and sufficient assessment of the risks". Additional regulations defining requirements for risk assessment in more detail apply to many of the systems covered by this Guide, e.g. plant containing large quantities of hazardous materials, nuclear power plant, railways, etc.

2.3.3 Safety Related Alarms

Whilst the main focus of this document is concerned with non-safety related alarms, there are some situations where an alarm is classified as safety related. It should be noted that safety related alarms will be the exception – the vast majority of alarms on most plant will not be safety related.

IEC 61508 (29) is the generic international standard that applies to electrical/ electronic/ programmable electronic safety related systems in all industrial sectors. IEC 61511 (30) is the process industry-specific international standard that embodies this.

Many countries recognise such international standards, although there may be no specific legal requirement to apply them. However, in practice, an organisation that can show that it has followed the appropriate standard may have the best defence if an accident occurs that was caused by the failure of the system or systems concerned.

There are stringent requirements for all for alarm systems to be adequately designed, installed, commissioned and maintained. In general, the design, engineering and maintenance practices required for safety related alarm systems are even higher.

An alarm system is an electrical/electronic/programmable electronic system (E/E/PES) under the definitions of the international standard IEC 61508 (29). An alarm system should be considered to be safety related if:

- it is a claimed part of the facilities for reducing the risk from hazards to people to a tolerable level, and;
- the claimed reduction in risk provided by the alarm system is 'significant'.

For a system operating in demand mode, e.g. an alarm system, 'significant' means a claimed Average Probability of Failure on Demand (PFD_{avg}) of less than 0.1.

If any alarm system is safety related then:

- it should be designed, operated and maintained in accordance with requirements set out in the standard;
- it should be independent and separate from the process control system (unless the process control system has itself been identified as safety related and implemented in an appropriate manner).

Often safety related alarms will be implemented in some form of stand-alone alarm system driving individual discrete alarm annunciators. These can provide good reliability and can be designed so that critical alarms are very obvious and easy to recognise (see Section 3.3.2).

It is implicit in the above paragraphs that both the equipment delivering the alarm and **the operator response** are part of the safety related system and hence they both need to be considered within the scope of IEC 61511 (30). At the time of publication, guidance on the use of this standard is being developed, see Bibliography (Appendix 23).

In seeking to provide a safety function that provides risk reduction, it should be recognised that the operator response will usually be the numerically weakest part of the calculation chain. A major issue that has to be addressed is that of the ergonomics of the system (see Section 2.3.4), which can have a huge impact on the likely timely success of the response.

2.3.4 Reliability Claims

The risk reduction achieved by an alarm will depend both on the reliability of the equipment (i.e. field instrumentation plus alarm processing system) and on the reliability of the operator in responding to the alarm with the appropriate action. The reliability of the operator will in turn depend on factors such as:

- the way in which alarms are presented;
- the time available for the operator to decide what to do and to implement a decision;
- the stress level;
- the basic unreliability of humans in performing defined tasks e.g. due to distraction, forgetfulness, negligence.

Where claims are made for risk reduction using alarms, consideration should be given to both the equipment and the human dimensions. However, particular attention should be paid to human factors issues. This is because experience shows that the majority of the failures of alarm systems derive from human failures rather than from hardware failures. Alarm hardware integrity is important, but in practice, risk reduction benefits are generally more easily derived from improving usability than from improving hardware integrity. Thus, in every alarm system:

- the operator should not be overloaded with alarms presented by the chosen display arrangement - either in normal operation or upsets[16];
- alarm system performance should be regularly checked to ensure that alarm overload is not occurring;
- alarms presented by the chosen display arrangement should be operationally very useful with very few spurious or low value annunciations;
- the operators should be trained in the use of the alarm systems;
- alarms should be properly prioritised.

For a safety related alarm, more stringent requirements should be imposed on both equipment and human performance. These are given in Table 5.

A target $\mathrm{PFD_{avg}}$ is generated from a consideration of the risk reduction required from a safety system - this will require hazard and risk assessment activities to be performed.

There is a limit to the amount of risk reduction which can be achieved using alarms even when the equipment is of the highest integrity. This is because of basic human reliability limitations. Consequently, as shown in Table 5, it is recommended that in no circumstances should a $\mathrm{PFD_{avg}}$ of less than 0.01 be claimed for any operator action in response to an alarm, even if there were

[16] An important point about alarm overload is that it is a dependent failure that affects all alarms on the system that is overloaded. If there are several alarms all derived from the same cause, they are likely to all be missed during an overload. Thus duplication of alarms does not reduce this risk, and indeed due to the increased alarm load, may actually increase risk. Note also that overload is most likely to occur when the plant state is changing, e.g. during start up or during an upset. These tend to be the periods of greatest risk when protection from the alarm system would be particularly valuable.

multiple alarms and the response was very simple[17]. This puts a limit on the level of reliability that should be claimed for any alarm function.

Claimed PFD_{avg}	Alarm system integrity/ reliability requirements	Human reliability requirements
1-0.1 (standard alarm)	alarms may be integrated into the process control systems	no special requirements - however the alarm system should be operated engineered and maintained to the good engineering standards identified in this Guide.
0.1-0.01 (safety related alarm)	alarm system should be designated as safety related and categorised as SIL 1 (Safety Integrity Level 1 as defined in IEC 61508 (29)); alarm system should be independent from the process control system (unless this has also been designated as safety related).	the operator should be trained in the management of the specific plant failure that the alarm indicates; the alarm presentation arrangement should make the claimed alarm very obvious to the operator and distinguishable from other alarms (see Recommendation 6. (23)); the alarm should be classified at the highest priority in the system (see Appendix 5); the alarm should remain on view to the operator for the whole of the time it is active; the operator should have a clear written alarm response procedure for the alarm; the required operator response should be simple, obvious and invariant; the operator interface should be designed to make all information relevant to management of the specific plant failure easily accessible; the claimed operator performance should have been audited.
below 0.01	alarm system would have to be designated as safety related and categorised as at least SIL 2.	it is not recommended that claims for a PFD_{avg} below 0.01 are made for any operator action even if it is multiple alarmed and very simple.

Table 5 Reliability requirements for alarms

A general principle expressed in various places in this Guide is that the operator should be able to easily identify alarms and should have adequate time to deal properly with them. This principle is particularly relevant to safety related alarms. Consequently it is recommended that:

> **For all credible accident scenarios, the designer should demonstrate that the total number of safety related alarms and their maximum rate of presentation does not overload the operator.**

[17] Techniques do exist for quantifying human error (37), examples being the THERP (41) and the HEART (47) techniques. When using these it should be noted that dealing with alarms in general (e.g. accepting alarms, moving up and down an alarm list) is a completely familiar and routine task that can be done consistently and reliably. However, diagnosing the cause of a specific alarm, working out an appropriate response and carrying this out successfully is a much more skilled task where the operator performance is much less predictable.

This might be interpreted as requiring that no credible accident generates more than a certain number of safety related alarms within a specified period.

Special efforts should be made to avoid spurious safety related alarms.

All safety related alarms should be tested at a frequency necessary to achieve the claimed PFD_{avg} (see Section 3.6 and Appendix 4).

2.4 Generation of Alarms

2.4.1 Types of Alarms

Alarm detection should provide early warning that there is a problem requiring operator intervention whilst minimising unnecessary or nuisance alarms. To achieve this, the most appropriate alarm detection mechanism should be chosen for each parameter. Some methods used for detecting alarms are:

- absolute alarms
- deviation alarms
- rate of change alarms
- discrepancy alarms
- calculated alarms
- recipe-driven alarms
- bit-pattern alarms
- controls and instrumentation systems alarms
- adjustable alarms
- operator-set alarms
- adaptive alarms
- re-triggering alarms
- statistical alarms
- first-up alarms

Descriptions of these and examples of when they might be applied are given in Appendix 6.

It is also useful to re-iterate what signals should not be alarms, i.e.:

- signals without a defined operator response;
- process variable or plant status changes that do not require operator attention;
- warnings of events which are too fast for the operator to prevent;
- events that are recorded in an alarm/event log but which the operator does not need to see;
- signals which confirm successful operator actions[18];
- signals which corroborate or duplicate another alarm (these signals may need to be logically suppressed - see Appendix 8).

Although these signals are not alarms, they may still need to be displayed to the operator, e.g. on event list displays or as status indications on plant mimic graphics (see Appendix 7 Alerts).

2.4.2 The Selection of Alarm Settings

All of the different types of alarms described in Section 2.4.1 included alarm settings or other parameters which influence when the alarm will be generated. It is important that these should be properly defined during plant design and commissioning.

[18] For example, the operator does not want an alarm from a pump he has just shut down, but does want an alarm if the pump trips unexpectedly. Similarly, he wants an alarm if an automatic changeover from a main to a standby pump suddenly takes place without him initiating it.

This is particularly important on absolute alarms indicating approaches to plant operating limits. All too often conservative alarm settings are chosen - with the consequence that alarms occur during normal operational fluctuations. Nuisance alarms can often be traced to bad alarm settings. Experience shows that such alarms are given little credibility by the operator and are often ignored or disconnected.

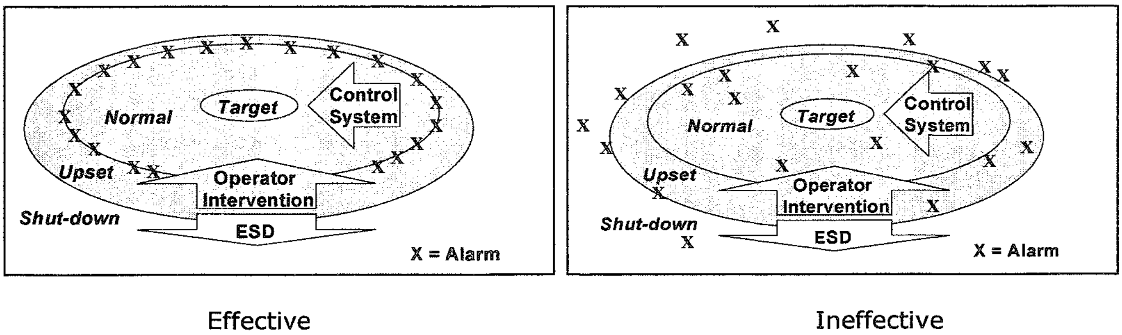

Figure 3 An effective and ineffective alarm system

The key to an effective alarm is that it should mark the point at which the operator has to take action. For example, as shown in Figure 3, alarms on variables reflecting operating limits should be set on the boundary between the normal and the upset state of the plant.

In practice the choice of alarm settings is complicated and must take account of the following factors:

- plant dynamics - i.e. the amplitude and duration of acceptable normal operational fluctuations;
- the limits at which the automatic protection system will operate (or economic loss will start to occur if this is not serious enough to need automatic protection);
- the rate at which the alarmed variable may be changing during a severe upset;
- the time it will take for the operator to respond and correct the problem generating the alarm.

An illustration of a possible relationship between these factors is shown in Figure 4 for a high pre-trip alarm. It is seen that there is a margin, A, between the normal operating limit and the alarm setting and another margin, B, between the point where the operator brings the disturbance under control and the trip limit.

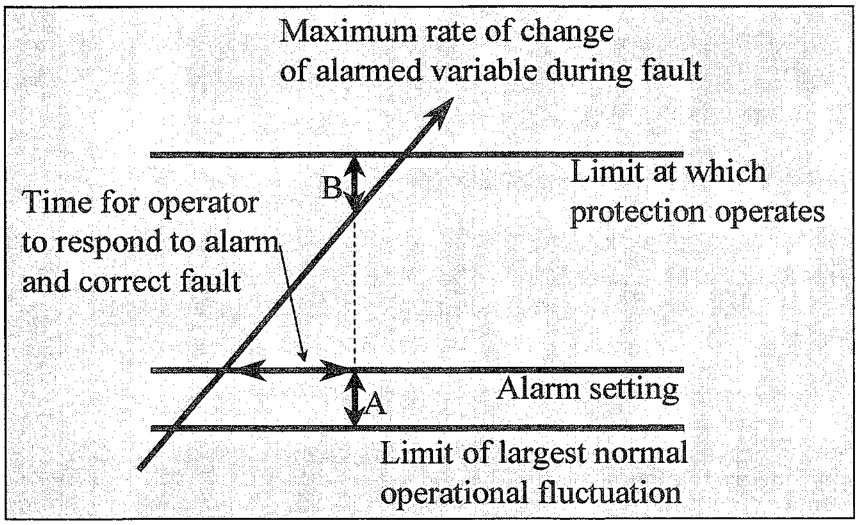

Figure 4 The setting of a high alarm

Sometimes conflicts between the factors will mean that it will be impossible to set an alarm at a value where alarms will not occur during the largest normal operational fluctuations yet will allow the operator time to correct the problem after the alarm occurs. In this case the designer has to choose between:

- **redesign** - i.e. redesigning the plant or its controls to provide greater margin between the normal operating limits and the trip limits. This is the most desirable solution but is often impractical or too expensive;
- **setting within normal operating limits** - i.e. setting the alarm within the limits of normal operating fluctuations and accepting that spurious alarms will occur during large normal disturbances. This is ergonomically very undesirable and will tend to increase alarm rates and reduce the operator confidence in the alarm system. In effect it increases the Average Probability of Failure on Demand (PFD_{avg}) for the alarm system as a whole;
- **setting nearer trip limits** - i.e. setting the alarm closer to the trip limits and accepting that some fast transients will not be corrected by the operator before they reach the trip level. This will increase the production losses due to plant trips and, because there are more demands on the protection system, tend to make the plant less safe. It also implies an increased PFD_{avg} for the alarm system.

Failing to resolve these conflicting requirements often leads to incorrect or conservative setting of alarms - which can result in additional costs when the plant becomes operational. A more detailed engineering study - e.g. modelling process dynamics and reviewing safety margins and cost penalties - may be justified if the plant is to operate with the alarmed variable close to plant constraints. Alternatively some other type of alarm, e.g. a discrepancy or rate of change alarm, may be appropriate.

A further issue is the accuracy of alarm setting that can be achieved in practice during calibration and maintenance. Frequently an absolute setting is defined plus an associated tolerance band.

Due to the operational importance of achieving effective alarm settings it is recommended that:

> **All alarm settings should be systematically determined and documented during plant design, commissioning and operation. All changes should be documented with reasons.**

The above discussion was based around the setting of absolute alarms. The same principles apply to the parameters in other types of alarms.

2.5 Structuring of Alarms

2.5.1 Prioritisation

In an alarm system of any significant size it is extremely useful to prioritise alarms such that the more important alarms at any given time are more obvious to the operator. This helps the operator to decide which alarms to deal with when several occur at the same time. This can be particularly useful during periods of high alarm activity when the operator needs to structure a response so that essential and important operator actions are carried out first. Prioritisation also helps during less busy times, as it brings it clearly to the operator's attention that alarms have occurred which are specially important and should be dealt with urgently.

It is usually appropriate to prioritise alarms according to two factors:

- the **severity of the consequences** (in safety, environmental and economic terms) that the operator could prevent by taking the corrective action associated with the alarm;
- the **time available** compared with the time required for the corrective action to be performed and to have the desired effect.

Both the severity of consequences and the time available may depend on the state of the plant, and the prioritisation system may need to recognise that fact.

In theory it might be desirable to have some dynamic prioritisation system which continuously and automatically ranked alarms taking account of severity of consequences, time available and changing plant state. In practice manufacturers of alarm processors generally take the simple approach of categorising alarms into a small number of priority bands. The priority is then used to vary the conspicuity of the alarm as displayed to the operator and the stridency of the audible warning.

Experience has shown that the use of **three priority bands within any one type of display**[19] is ergonomically effective for the normal presentation of alarms. However, a plant may have more than one alarm system. As an example, there may be alarms implemented within the process control system, alarms (some of which are safety related) on a hard-wired annunciator, plus a separate fire and gas alarm panel. The use of priority should therefore be adapted to suit the particular arrangement chosen. However, it is recommended that definitions of alarm priority should be consistent across systems and in total there should be no more than four normal priority bands in any plant. Some examples of different arrangements are illustrated in Figure 5.

Example C shows a mixed alarm system with some alarms in the process control system and some in a stand-alone system. Here it would be acceptable to categorise alarms into four overlapping priority bands. In this example, the high priority alarms are implemented both in the process control system and the stand-alone system. In addition there is a higher priority band of **'critical alarms'**. This would include the safety related alarms which would have to be implemented to at least integrity level SIL 1. Provided it does not cause confusion, it may also be acceptable to include alarms representing significant environmental or economic risks among the critical alarms. Designers may choose to implement these in the same manner and to the same integrity level as safety related alarms. This may have both reliability and ergonomic benefits.

[19] In this report the three priority categories have been termed high, medium and low (and there is also a critical category as discussed later). Different manufacturers use different terminology, e.g. 1, 2, 3; emergency, urgent, alert, warning, etc. Readers should take care not to confuse the terminology used in this report with that used by particular manufacturers. Note also that the use here of the term 'low priority' should not imply that the alarm can be ignored by the operator. As stated earlier, every alarm should require a response from the operator otherwise it should not be an alarm.

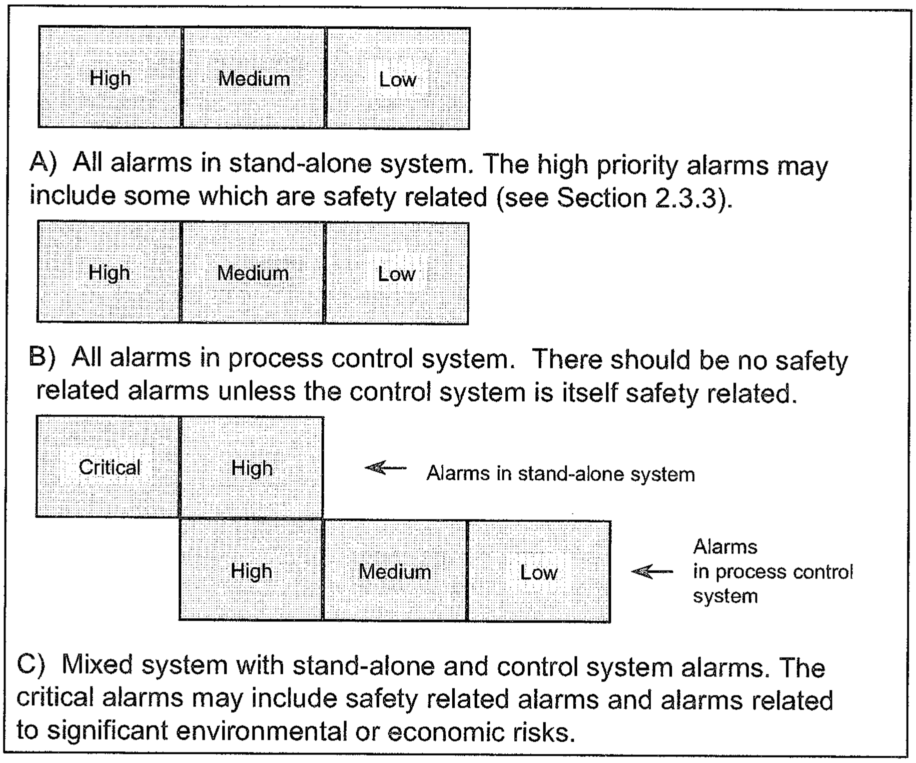
High
Medium
Low
High
Medium
Low
Critical
High
Alarms in stand-alone system
High
Medium
Low
Alarms in process control system

A) All alarms in stand-alone system. The high priority alarms may include some which are safety related (see Section 2.3.3).

B) All alarms in process control system. There should be no safety related alarms unless the control system is itself safety related.

C) Mixed system with stand-alone and control system alarms. The critical alarms may include safety related alarms and alarms related to significant environmental or economic risks.

Figure 5 Use of priority for different types of alarm system

There are other areas of flexibility in the general guidance of having three priorities in any alarm system, e.g.:

- **additional bands** - to simplify the operator interface it may be convenient to assign additional priority bands to signals not normally displayed[20] but which the operator may wish to sometimes examine using the standard alarm display facilities such as the alarm list;
- **sub-division** - it may be desirable to further sub-divide priorities. For example, fire alarms might be categorised as critical safety related alarms but displayed differently from other critical alarms and have a different audible warning. If adopted, such sub-division should be done with great care to avoid the prioritisation becoming confused.

The priority of every individual alarm should not necessarily be fixed for all time. If achievable within the alarm system, it can be very effective to dynamically modify an alarm's priority according to the prevailing plant state.

[20] These signals might include alarms which have been logically suppressed, control loop status changes, etc. If such a facility is provided, care should be taken not to undermine the requirement that every alarm should have a defined response, i.e. the display of information-only events should not come to represent normal usage of the alarm system. Note that it may also be useful to show logically suppressed alarms - in a suitable coding - on schematic displays

In order to properly manage priority:

> **Every site should have written rules on how priorities should be assigned. These should be applied consistently to all alarms in all systems used by the operator.**

Appendix 5 discusses the prioritisation of alarms in more depth. It includes an example procedure for how alarms may be prioritised based on their associated safety, environmental and economic consequences and on the time available for response. An example of the prioritisation of an economic alarm is given in Figure 6. The Appendix also discusses the prioritisation of safety related alarms.

2.5.2 Logical processing of alarms

A general principle that follows from the discussion in Section 1.3 is that:

> **Alarms should be designed so that they are worthy of operator attention in all the plant states and operating conditions in which they are displayed.**

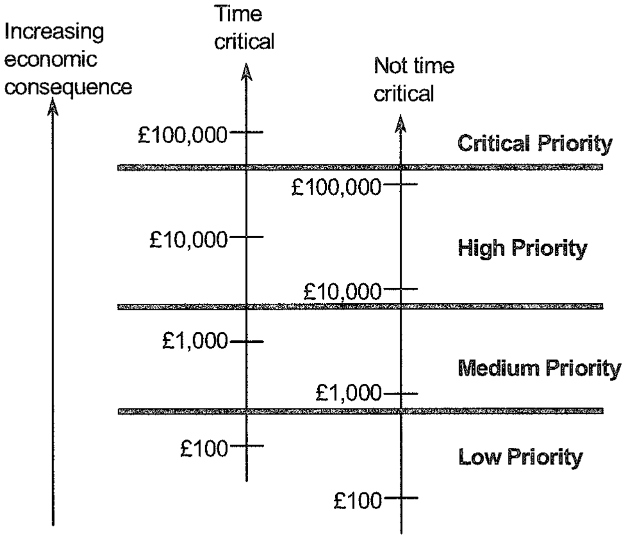

Figure 6 shows an example of how an economic alarm might be prioritised. The priority depends both on the severity of consequence and on the time available. Thus, e.g. if the alarm is not time critical then it is high priority if the consequence is between £6,000 and £150,000 (the figures are purely illustrative). If it is time critical it is high priority if the consequence is between £2,000 and £50,000.

Figure 6 Possible rules for alarm priority

In order to generate alarms that meet this requirement, the alarm signals detected by one of the means described in Section 2.4.1 often need to be logically processed and/or combined with other information in order to generate alarms in a form suitable for display to the operator.

Methods for logically processing signals to generate useful alarms are discussed in Appendix 8. This covers:

- grouping of alarms
- logic for handling redundant alarms
- eclipsing of multiple alarms on the same variable
- alarms from out of service plant
- operating mode suppression
- major event suppression
- intelligent alarm processing
- automatic alarm load shedding
- alarms from equipment under test

Appendix 9 adds additional techniques related to the treatment of repeating alarms since these are important particularly for the effective display of alarms on alarm list graphics (see Appendix 11). There is no general guidance on which of the methods described in these two Appendices will be appropriate to any particular alarm, but the techniques should be applied until the designer is confident that the above general principle is satisfied for all alarms.

2.6 Designing for Operability

The previous sections have considered the detection of an alarm and the logical processing and combination of alarms. However, as discussed in Section 1.3, there is a higher level design issue of matching the alarm system to the human user. Even if each alarm that is generated is operationally useful (in the sense that the correct operator response to it would reduce risk or save money), that does not mean it should be always displayed to or acted on by the operator. This is because the operator has only a limited 'processing power' and there may be more important things that the operator should be doing with their limited time which would have a greater impact on reducing risk or saving money.

If an operator is to be effective, it is essential that the operator can cope with all the tasks demanded of him under both normal and emergency conditions. Given the variety of other activities required from the operator it is crucial that the operator is not overloaded with alarms. As discussed in Section 1.3, this may severely limit the operator's capacity for handling alarms.

To take an example, if there is on average 1 alarm per 2 minute and the operator on average needs 1 minute to deal with each alarm[21], then the alarm system will be demanding 50% of the operator's time. For many operators, such a workload would be unacceptable even during normal operation - and there will be situations where other important tasks reduce the amount of time that the operator can devote to alarm management.

This example illustrates the gulf between the operator capabilities and the design of many existing systems. There can be debate about what the appropriate parameters are for operator workload and response time, and what is meant by 'on average'. However, it is an inescapable fact that there is a limit on what any operator can achieve, and that this implies a limit on the rate at which the operator should be expected to deal with alarms.

There are two parallel approaches that the designer should take to reduce the problems of alarm overload:

Eliminate the alarm overload

Alarm overload represents a significant threat to operator effectiveness, so strenuous efforts should be taken to avoid it. This requires that efforts are made during design to ensure that all alarms are justified and properly configured. Further, once a plant becomes operational, performance should be audited, overload incidents should be identified and steps should be taken to minimise their frequency.

[21] Those used to working with existing alarm systems might consider 1 minute per alarm to be excessive. However, their experience is often based on systems with large number of spurious or low value alarms.

Typically overloads are due to a small number of spurious or badly designed alarms. A review programme should be set up to identify these (see Section 5.3) and to re-engineer them. However, since the alarms that cause overloads tend to vary from incident to incident, the review will ultimately have to cover many different alarms. Techniques to re-engineer alarms range from simple re-tuning through to complex suppression methods and ultimately perhaps to automatic alarm load shedding (see Appendix 8).

Improve the management of alarm overloads

The elimination of overloads should be the first priority. This may however be a difficult and long drawn out exercise and even then it may be very hard to be certain that all potential for alarm overload has been eliminated. Consequently steps should be taken to ensure that the operator performs as effectively as possible during any overload incidents that do occur.

To assist in this, the operator information displays should be designed such that the operator can easily access all key plant information even if the alarm system does become overloaded. This may require a suite of overview displays of critical safety parameters to be developed. In addition the operator should be explicitly instructed and trained in how to respond when the alarm load is too high. Progressively drastic operator responses might be to:

- select the alarm list display to show only high priority alarms and ignore all medium and low priority alarms;
- ignore all standard alarms on the process control system and only look at safety related alarms on the stand-alone system;
- ignore all alarms and operate the plant using overview displays showing critical plant information[22].

A further consideration is that it can be useful to define low priority alarms as ones that may be "turned off" (i.e. not displayed to the operator – but still recorded in a journal) during alarm floods. This must clearly be recognised at design time (and in the Alarm Philosophy) and will limit the opportunities for adjustment of priorities referred to above. If low priority alarms can be turned off, the impact of alarm flooding can be ***substantially*** reduced. It should be clear to the operator when the alarm system is running in this mode.

Whatever strategy the operator is trained to follow, the implications of missing important information should be carefully assessed.

The key requirement relating to alarm overload is that the alarm system designer should recognise the fundamental human usability limitation, and have some explicit strategy to deal with it.

[22] As an illustration, operators at many nuclear power stations are trained so that, following a reactor trip, they work step-by-step through a defined post-trip operating procedure using specified post-trip displays. This systematically checks that the automatic shut down system has successfully carried out all necessary safety actions. This monitoring does not rely on use of the alarm system, so the operator can confirm safe shut down even if the alarm system were overloaded or failed - though of course it is preferable if it does remain fully usable.

3. Implementation Issues

Section overview

This section discusses practical issues of implementing alarm systems. This covers the derivation of alarms from field measurements, the design of alarm processing systems in general, and the display of alarms to the operator. It also covers training, procedures and testing.

3.1 Field Derivation of Alarms

The process sensing field devices used to generate alarm signals are critical elements of an alarm system. In practice they are frequently neglected and given insufficient attention during design, installation and maintenance. If the field device is incorrectly specified, badly calibrated, installed in the wrong location or rarely/never maintained, then the operator may be unaware of risks to people, plant or the environment which require attention.

Appendix 10 highlights some key issues relating to field alarm sensors including:

- the choice of analogue sensors or switches;
- the location of sensors;
- the choice of sensor range;
- the validation of signals from alarm sensors;
- the transmission of alarm signals;
- dependent failures in alarm sensors.

However, instrumentation practice is a large and important topic. The Appendix cannot provide comprehensive guidance and should only be taken as an introduction to the subject.

3.2 Alarm Processing Hardware

3.2.1 The Commercial Environment

There are a variety of off-the-shelf and individually designed systems for handling alarms available from different manufacturers. The commercial scene is continually changing, and this Guide can only deal with general issues.

Typically most large alarm systems, and many smaller ones, are implemented within proprietary DCS or SCADA packages which are also carrying out other control or information processing activities. The next most common approach is to use dedicated alarm systems typically driving annunciators or some dedicated programmable display devices. A checklist of functionality found in a typical dedicated alarm system is given in Appendix 17.

Integrated alarm/control/information systems will generally be constructed around a shared set of input hardware, processors and displays. They will also have standard functionality for processing alarms integrated within their system software. Similarly, dedicated alarm systems will also use standardised alarm handling software.

It will often be costly for users to have this standard system software customised to meet their individual needs. Also, although manufacturers do continually upgrade their offerings, there are often restrictions and high costs on upgrading existing systems. This means (see also Section 6.3 and Appendix 17) that before buying an alarm system:

> **Purchasers should carefully assess any proposed system to ensure that it meets their current and likely future needs.**

To achieve this, the purchaser will need to have developed an alarm design strategy.

Consideration should also be given to limits on expanding or modifying the system hardware. Hardware is generally modular. It may be desirable to purchase both installed spare capacity plus space for expansion.

3.2.2 Reliability Issues

Modern alarm systems are often highly distributed with input facilities distributed around the plant, and multiple processors for alarm processing and display. The implications of failure should be carefully considered.

Often there are a range of possible failures, e.g.:

- failure of an individual alarm input;
- failure of a scanner;
- failure of one processing or display device;
- failure of system communications;
- failure of the total alarm processing system;
- failure due to high alarm generation rates[23].

Alarm systems should be designed so that failures are made obvious to the operator.

Techniques used include watchdogs, check channels on scanners, integrity checks on communication lines, etc. In practice, the detection of failures on the individual alarm signals is often an issue. Duplication of inputs and validation techniques (see Appendix 10) can be used, but these may have a high cost implication if applied to all inputs. The designer should assess the implications of undetected input failures and whether these justify additional investment.

The operational implications of potential failures should be assessed, and consideration should be given to the need for operator procedures to cover them.

More sophisticated systems will include redundancy and segregate components to provide defence against partial failures and reduce the likelihood of total system failure. However, experience shows that total system failures are credible, so consideration should be given to the required operator response. For example, if the operator loses all display screens:

- can the operator be certain that, whatever the state of the control systems, the automatic protection will ensure that the plant remains in a safe state, or should the operator initiate a manual shut down?

[23] An alarm system should continue to function if there is a very high alarm generation rate. However, if this could occur in practice, then the designer should also consider the ergonomics issues. As has been stressed in several places in this Guide, operators have a limited capacity for dealing with alarms. If an alarm system cannot reduce a flood of alarm inputs down to a manageable number of displayed alarms, then it has in effect failed.

- is the loss of all alarms and displays acceptable, or should there be some independent system for providing critical process measurements and alarms (e.g. those relevant post-shut down)?

It is noted that alarm processor failure alarms are often a cause of considerable nuisance to operators - particularly during electrical supply disturbances (e.g. due to lightning strikes). Alarm systems designers should apply stringent criteria to what should be alarmed, similar in character to those applied by plant equipment designers. Also capabilities for logically processing alarms should be as applicable to alarm system alarms as they are to any other alarm.

3.2.3 Functionality

In various places this Guide identifies functions that may be performed by the alarm system (e.g. logical processing of alarms in Appendix 8, handling of repeating alarms in Appendix 9). Purchasers of alarm systems should review which of these functions they need. They should also consider how easy the system is to use, e.g. how logic is configured, how data such as alarm settings and messages are changed.

The capability of the alarm system should also be assessed, e.g.:

- how many alarms can be configured in the system?
- what rates of data capture can be handled?
- is the operator advised of data overload/lost alarms?
- how accurately are inputs time-stamped?
- how does the alarm system handle alarms received out of time sequence? To what extent can alarm list displays get out of time order? Can the operator correct the time ordering?
- what is the maximum time delay between input change and display to the operator?
- what length of alarm list can be stored?
- are there limits on the numbers of alarms that can be configured at each priority?
- how many different alarm lists can be configured?
- are there limits on size of alarm logs?

Any requirements on alarm system capability should be related back to the needs of the operator. For example, a delay of 2 seconds between an alarm occurring and the operator seeing it is unlikely to be of great operational significance - although a loss of order of occurrence of alarms may matter. Similarly, if there is a very large flood of alarms in a short space of time - say, 1000 alarm inputs becoming active then clearing within 10 seconds - it is operationally irrelevant if the system loses some transitions and fails to display them to the operator. The important thing is that the alarm system reduces the flood down to a manageable number of displayed alarms - and in this example it would almost certainly have to discard much fleeting alarm information.

3.2.4 Environmental Requirements

Consideration should be given to the environmental conditions under which the alarm system should operate, some of which include:

- **power supply quality** - e.g. voltage variations, frequency variations, transient over-voltage, interruptions, dips, notches, voltage imbalances, harmonics. An important issue is whether bursts of spurious alarms are generated at system power-up or following supply interruptions;

- **electromagnetic** - both emission and immunity requirements. Alarms installed near high power equipment (e.g. high voltage switchgear, battery rooms) may require special precautions to be taken;
- **electrical isolation** - e.g. channel-to-channel isolation between analogue and digital inputs, isolation between power supplies and retransmitted outputs;
- **explosive atmospheres** – whether equipment has to be certified to operate in zoned areas;
- **operating temperature** - the upper and lower operational temperature limits;
- **dust and moisture** - i.e. the degree of protection against ingress to be provided by enclosures.

3.3 Display Options

3.3.1 Types of Display

Alarm system displays fall into two main categories:

- alarm displays on VDU screens or other programmable display devices;
- alarm annunciator panels, consisting of 'light boxes'.

There may be supplementary alarm displays such as single indicator lamps, LED arrays, etc. for special purposes.

Alarms may be displayed on programmable devices both as:

- dedicated lists of alarm text messages, and;
- symbolic indications of alarms on general graphics such as plant schematics or control faceplates.

In some cases plants may have a mixture of alarm displays on individual annunciators and on programmable devices. This might be because:

- some alarms are safety related or critical priority (see Section 2.3.4);
- some alarms are needed in the event of the failure of the programmable device;
- a staged replacement of annunciators by programmable displays is in progress.

Care should be taken to ensure that any such combined system has an integrated and consistent overall alarm philosophy and is easily usable. For example, if the same alarm is duplicated on two different displays, it should have the same categorisation, prioritisation and coding on both types of display and should be accepted by a single operator action. Furthermore, it should not be possible for an operator to accept an alarm on one system whilst thinking he is accepting an alarm on the other.

Section 1.3 has discussed the function of the alarm system as a whole. Within this, the function of the alarm display is to:

- draw the operator's attention to the occurrence of new alarms, and;
- provide facilities for viewing the state of alarms.

The alarm display should also be integrated into the operator interface in such a way that, when an alarm occurs, the operator can quickly implement the required response. For example, if alarms are normally managed via an alarm list, and this is not permanently on display, then the operator should be able to access this very quickly and easily. In some cases (e.g. offshore), where plant may be supervised from two locations, an alarm acceptance protocol may be required.

Display systems can draw attention to alarms using colour changes, symbol changes and dynamics (e.g. flashing). This is often supported by audible annunciation of the alarm. Wherever colour is used (either on annunciators or alarm graphics), account should be taken of the possibility of operators having impaired colour vision. This requires care to be taken in the choice of colours and for colour-coded information to be replicated in another form such as a shape or intensity change. Colour usage should be consistent across all types of graphic displays and should be consistent with the colours used on other control room indications and on the plant (see also Section 3.3.4). BS EN 60073 (7) provides general guidance on colour coding.

All alarm displays should be easily viewable and provide an easy to use and intuitive interface. Alarm displays based on programmable devices may be able to provide operator access to other important information related to the alarm, such as extra detail on the plant area affected, associated schematics or trends, or operating procedures. These facilities can significantly enhance the usefulness of the alarm system.

3.3.2 Annunciator Display

Traditional alarm panels, made up of arrays of annunciators, provide a basic alarm system interface. These are becoming less common as alarm systems make use of computer technology. However, they can still be useful for small systems or more especially for display of critical or safety-related alarms.

Annunciators can provide:

- excellent spatial pattern recognition;
- large easily visible displays;
- immediate access to information;
- ease of use.

However they do not provide access to associated information about the alarm, nor are they appropriate for systems with potentially large numbers of alarms.

Within this framework the following rules should be applied to make best use of annunciators:

- groups of annunciators should be logically arranged, by plant area, and should reflect the structure of the process;
- within each annunciator, individual alarms should be arranged so that spatial layout consistently reflects functional relationships, e.g.:
  - alarms associated with a single plant item should be together;
  - 'high' alarm indicators should be physically above 'low' alarm indicators;
  - layout patterns should be similar across all annunciators.
- each annunciator should allow easy pattern recognition so should only display a limited number of alarms (e.g. an 8x8 array - although more may be acceptable);
- annunciators should be positioned so as to be easily read from the normal operating position;

- alarm annunciation should accepted by a button push which should also silence any audible annunciator;
- operating buttons should be positioned within reach of normal operating position;
- a new alarm should cause the associated annunciator window to flash in an obvious colour;
- acceptance should cause the window to remain lit but in a steady mode;
- clearance of the alarm condition should be indicated in the display, e.g. by colour or brightness change. It should then be darkened by pressing the reset button;
- a lamp test facility should be provided;
- alarms should be colour-coded according to priority;
- first-up alarm facilities may be valuable for some applications (see Appendix 6).

3.3.3 Alarm Lists

Alarm list displays are the most common form of alarm display provided within software-based alarm processing systems - both stand-alone systems and those implemented within proprietary DCS or SCADA systems. This is because a list display provides a single channel through which very many different alarms can be presented; such large numbers of different alarms cannot be easily handled on annunciators or other types of display. Operators making the transition from annunciator-based systems can initially find alarm lists difficult to use, but provided the implementation has the necessary features, lists can provide a good operator interface.

The performance and usability of alarm list displays can be greatly affected by the underlying sophistication of the alarm manager software. The performance in handling high alarm loads and repeating alarms is particularly important. Thus:

> **Alarm list displays should be designed such that repeating alarms do not cause them to become unusable.**

Different list handling facilities are provided by different suppliers and often the user is very constrained in what changes can be made once the system has been purchased. Appendix 11 gives guidance based on user experience with a range of different systems.

Alarm list displays give the opportunity of providing the operator with a wealth of information that is not accessible using an annunciator system. This information can be used to guide the operator with respect to actions to be taken, plant problems to be investigated, etc.

A useful method for providing such information is to configure the system so that a right mouse click on the alarm text produces a pop-up pick-list. Where the DCS permits this, functionality should be configured in the alarm system. Three main types of information can be provided via this method:

Quick navigation to operational displays

- **schematic** - a pre-configured schematic of the plant area associated with the alarm. (Note that some alarm processing systems provide push-buttons for selection of schematics which light up when there are alarms in the associated plant area);
- **control panel** - a control panel of the associated plant item which 'pops up' so that action may be taken directly on the problem plant item;
- **trend** - a history trend of the signal generating the alarm;

Actual Alarm information – what it is, why it is in the system, what to do etc.

- **alarm point information** - general information on the point in alarm - e.g. settings, normal operating value, database identification, equipment reference;
- **alarm data** - what the alarm is, why it is in the system;
- **action** - text specifying the action to be taken on receipt of the alarm;
- **operating procedures** - more detail on associated alarm response procedure or other pertinent information.

Alarm performance history data

- **alarm and event log** - providing operator access through a screen to a time-ordered list of all alarms and other events such as status changes can be useful to enable him to look back and examine the sequence of events that occurred, say, just prior to a plant item trip;
- **history** - times of the last 'n' occurrences of the alarm, number of occurrences within the last 'n' minutes, or top ten alarms in the last 'n' minutes;
- **trend** - a history trend of the signal generating the alarm.

3.3.4 Alarm Display on Schematics

Alarm lists are the most common way of displaying alarms on programmable devices as they provide a channel through which every alarm on the system may be presented. However, the operator interface can be significantly improved by also communicating alarm information via schematics. In practice on large plants it is generally impossible to have every active alarm on view to the operator on schematics. Therefore, display of alarms on schematics needs to be carefully thought out.

Alarm information on schematics is often associated with:

- highly critical/important alarms;
- generalised plant status information to guide to more detailed investigation.

Schematics containing alarm information are usually:

- overview displays (e.g. plant wide);
- detailed plant area or plant item displays.

Overview displays give the following properties:

- **spatial positioning** - there are benefits if these displays are designed to be on view continuously and in a fixed position. This provides an arrangement similar to an alarm annunciator where the alarm is recognised almost solely from its position. Overviews allow all key plant parameters and status indicators to be easily scanned to determine where problems may be occurring (3). It is noted that overview displays can be very valuable when managing severe upsets, and, for example, the HSE commented in its report on the Milford Haven explosion (23) that the lack of process overviews made it difficult for operators to identify the cause of the upset and to manage the incident;
- **grouped alarms** - the level of detail and amount of space is limited. Hence the alarm indication on the overview tends to be a group alarm, indicating a problem with a specific plant item which needs further investigation;
- **colour/dynamics** - given the space constraints, information needs to be coded by the use of colour and dynamics. For example an orange flashing symbol might indicate an unaccepted urgent alarm;
- **criticality** - overview displays are an excellent place to highlight critical items as they are permanently on view;

- **navigation** - overviews can provide facilities to bring up more detailed information displays (either as a temporary overlaid 'mini-window' or a display on another screen).

Detail plant area or plant item schematics can be used to show the individual alarm indication that contributes to the group alarm on the overview. For instance, the overview display may indicate a problem on a pump, the detail schematic showing the high temperature alarm and its location.

A full discussion of the subject of colour coding on schematic displays is outside the scope of this document. However, in general, the conspicuity of information should be related to its operational importance; background information should be given a low conspicuity, normal plant measurements a medium conspicuity and highly conspicuous colours should be reserved for abnormal states and alarms. Where alarms are colour coded, priority should also be reflected in the colour conspicuity. Thus, high priority alarms should be displayed in highly conspicuous and distinguishable colours, and lower priorities displayed in progressively less conspicuous (but still very conspicuous) colours. Similarly, unaccepted alarms should be more conspicuous than acknowledged or reset alarms. Thus:

> **Conspicuity of colour coding of alarms should reflect alarm priority and state.**

Note that the conspicuity of colours depends on the background colour on which they are displayed[24]. Also, whilst use of large numbers of colours may be very beneficial, e.g. for including TV images in schematics, this should be done with care to ensure that important alarm information does not lose conspicuity.

Flashing and blinking on schematics should only be used for unaccepted alarms, and all other movement on the graphic should be minimised.

Facilities to allow acceptance of alarms from schematics should be provided.

Navigation between screens should be intuitive and quick.

3.3.5 Audible Warnings

Audible warnings are normally generated in conjunction with the display of new alarms whether on annunciators, lists or schematics. However, control rooms generally include a number of other devices which can generate audible warnings, e.g. phones, radios. Often these warnings need to be directed at particular individuals. Thus, to achieve effective use of sound:

> **An integrated design should be developed for all audible warnings in the control room.**

This is an area where specialist advice may be desirable. The general guidance (49) is that a limited number of sounds should be used which should be easily

[24] Note that, because of the meaning commonly associated with it, red is often used to code the highest priority alarms without it actually being the most conspicuous colour. Slight tone change, for example, making the colour more orange can often increase conspicuity.

discernible. It is important that an operator should be able to identify audible warnings from alarms from the plant under the operator's control and distinguish these from alarms directed at other operators. It is desirable that the operator should be able to identify the priority of alarms, and this may be achievable using tone frequency or modulation frequency. With careful design it may also be possible to code the type of alarm using sound patterns or tone characteristics.

The HSE has published guidance on audible warnings (27). Relevant points are:

- so that they can be heard, acoustic signals should be set at a level considerably higher than the ambient noise at the signal frequency (e.g. 10dB above the level of ambient noise at the signal frequency). However, the level should neither startle the operator nor be painful;
- acoustic signals should be easily recognisable, particularly in terms of pulse length and the interval between pulses or groups of pulses;
- the same acoustic signal should not be used for more than one function;
- acoustic signals may be constant frequency or variable frequency (this includes an intermittent signal operating on a discrete frequency). Where both types of signal are used, the variable frequency signal should indicate a higher level of danger or a more urgent need for intervention or action.

Continual audible warnings from alarms can be extremely distracting, particularly during upsets. Palliatives, such as reducing volume during alarm floods, may be adopted. However, as discussed in Section 2.6 and elsewhere, it is preferable to tackle the basic problem and reduce the alarm load to a manageable level.

3.4 Training

The training of operators is a large topic, and this discussion will be limited to issues related to alarms systems.

All operators should be trained in the use of the alarm systems that they actually work with. This should comprise initial training, refresher training and training in any new alarm system facilities.

Training should be designed to ensure that the operator becomes familiar with the functionality of the alarm system and knows how it should be used. Training should also cover the diagnosis of faults in the alarm system itself and the operator response to such faults. Training methods include 'classroom' training, on-the-job training, simulator training, etc. Processes should be in place to regularly review the training content and identify the need for refresher training.

A key principle is that:

Training should cover all realistic operational usage of the alarm system.

For example, if simulator training is used, the operator should learn how to deal with spurious alarms and alarm floods. Simulations should not just represent idealised situations where every alarm is valid and meaningful unless this is what prevails on the plant. Training should also cover the diagnosis of plant faults from alarms and other information.

As indicated in Section 1.3, every alarm should have a defined response. Furthermore, the operator should know what the response should be. To achieve this, in addition to general training in use of the alarm system, the operator should be given specific training in how to respond to every important alarm and in the procedures that should be followed. The training should cover the consequences of failing to complete critical steps in the procedure.

3.5 Procedures

Each and every alarm should be covered by a written (or on screen) 'alarm response procedure' which should assist the operator in identifying and carrying out the necessary response. Many alarms may have a very similar response and may be covered by a general procedure. However, for critical alarms an individual procedure per alarm is generally justified.

The content and presentation of procedures can have a significant impact on the effectiveness of the operator's response. Procedures should thus be written in a clear and easy to use way which supports the operator and promotes human reliability.

Many organisations (particularly in the nuclear power industry) have produced in-house guides on writing procedures. Published guidance includes (21), (22), (24), (42), (44), (45), (46) and (50).

An alarm response procedure has two functions:

- **informative** - i.e. to provide information on how the alarm has been generated and what might have caused it. This information should include:
  - brief description of the alarm;
  - identification of the operating conditions in which the alarm is relevant;
  - list of probable causes;
  - values of alarm settings;
  - any special precautions or limitations relating to the alarm;
    - identification of any trips or automatic actions related to the alarmed variable;
  - reference number of appropriate plant detail schematic;
  - tag number and location reference of field alarm sensor;
- **instructive** - i.e. to provide step-by-step instructions on actions to be taken by the operator. These should clearly identify variables to be monitored, checks to be made and actions to be taken. Checks should be laid out in a clear fashion, e.g. "**IF** level X < 1% **THEN** ensure valve Y has tripped".

The procedure should also show:

- the issue number and date of the procedure;
- an explanation of the latest revision to the procedure;
- modification history and approvals for changes to the procedure.

On-screen alarm procedures should provide similar information to that provided in written procedures. However, there will be more flexibility for layering information or presenting alternative access routes to it.

The adequacy of procedures should be audited. Checks should be made that they do actually assist the operators in correctly performing tasks that they are unfamiliar with and that they are actually used by the operators.

3.6 Testing of Alarms

3.6.1 Management of Testing

A strategy should be developed for the testing of alarms.

In particular, the strategy should address the testing of safety related alarms to assure their reliability, where the test interval should be calculated to achieve the required target PFD_{avg}. Testing of other higher priority alarms may be required where there is a financial or environmental justification. Testing is unlikely to be necessary if the correct functioning of the alarm is regularly demonstrated in normal operation or where the effects of failure of the alarm does not justify testing.

There should be written test procedures. These may be generic for a number of devices or specific to the individual device. The test procedures should specify realistic tolerances on the point at which the alarm should become active (typically within ± 2.5% span of the alarm setting). This should be done to ensure that results do not depend on the subjective judgement of the person carrying out the test.

Testing should be carried out by suitably trained competent individuals.

The operator may need to take an active part in the test. Whether the operator does so or not, they should be kept aware of which alarms are being tested. It may be appropriate to divert alarms from the normal operator display while testing (see Appendix 8).

Results of the tests should be recorded, and these should be the results as found. Corrective actions should be recorded. The status and results of individual tests should be monitored. An overall review of the results of testing should be carried out periodically. It is good practice to review test results over time as it may be possible to amend test frequencies.

Testing should be carried out on the equipment as found. Any necessary maintenance, e.g. clearing of impulse lines, should be carried out following, not before testing.

Ideally, faults should be rectified at the time of testing. Where this is not appropriate, rectification should be initiated with the appropriate priority. The operator should be made aware of any outstanding defects.

3.6.2 Test Methodology

Where it can be done safely and without significant economic loss, and provided that it can be carried out in an acceptably short period of time, the test should be carried out by driving the alarmed process variable into the alarm state. This may be especially appropriate for some flow and level alarms.

Where simulation of a measurement is necessary, this should be done by injecting a signal into the primary side of the transmitter via the impulse piping and ensuring that the alarm operates at the appropriate point.

It is emphasised that alarms should **not** be tested by altering the alarm setting; this does not prove that the transmitter is capable of achieving the appropriate output. Similarly, alarms from smart instruments should not be tested by artificially overwriting the instrument output.

Where blockage of the impulse lines to an instrument is credible, the test should include a check that the impulse lines are clear.

Where there are alarms and trips on the same measurement, trips should be tested at the same time.

Batch plants may require different alarm settings for different products. Consideration should be given to testing before the first batch of each different product.

Different parts of the loop may be tested at different times, and, if appropriate, at different intervals, provided that, for safety related alarms, the required PFD_{avg} is achieved.

4. Measuring Performance

Section overview

This section discusses ways of measuring the performance of alarm systems - and how these measures may be used as targets in an alarm system construction or improvement exercise (see also Appendix 14 and Appendix 15).

4.1 Performance Metrics

A number of measurements of the performance of an alarm system are discussed in detail in Appendix 12. These measurements can be used:

- as performance targets for acceptability of a new alarm system;
- to assess the adequacy of an existing alarm system;
- as management tools for assessing the effectiveness of an on-going improvement programme;
- to identify specific nuisance alarms;
- to demonstrate to an independent auditor or regulator the performance of the alarm system.

Performance measurement with operator questionnaires and usability surveys are discussed in Appendix 14 and Appendix 15.

Appendix 12 also presents benchmark values for some of the measurements. These are empirical values based on industrial experience rather than fundamental theory.

Design benchmarks include the number of alarms per control valve/per analogue measurement/per digital measurement and the distribution of alarm priorities. They provide average figures that may point to potential future problems.

Usability metrics include the average alarm rate in steady operation, the number of alarms in 10 minutes after a plant upset, the average number of standing alarms and the average number of shelved[25] alarms.

The usability benchmarks may help in the assessment of whether the operator will find the alarm system easy to work with. Table 6 gives some benchmarks for usability. Several of them relate to whether the alarm workload imposed on the operator is one that the operator can cope with and they thus represent indicators of ergonomic acceptability. If these benchmarks were achieved it is suggested that the operator would find the alarm system extremely manageable. They are not absolute numbers and less demanding numbers could still produce manageable systems. However they should be regarded as target values to aim for.

[25] See glossary in Appendix 1 for definition of the term 'shelved'.

Usability metric	Benchmark value
usefulness questionnaire	nuisance score of 2.0 or less
average alarm rate in steady operation	less than one per 10 minutes
alarms in 10 minutes after plant upset	under 10
average number of standing alarms	under 10
average number of shelved alarms	under 30

Table 6 Benchmark values for usability metrics

It is possible that the benchmark values are not achievable in the short term or may not be financially viable in the long term. Nevertheless, setting targets and goals from which improvements can be seen will give major benefits not just to the loading on the operator, but also financial benefits in terms of less plant outage and better performance.

4.1.1 Key Performance Indicators (KPIs)

In order to define performance levels for an alarm system, it is necessary to define a set of quantitative key performance indicators (KPIs). The KPIs should relate to the basic usability benchmarks defined in Appendix 12 and be calculated over a reasonably long period of time (e.g. a week) – the measurement period. The main benchmark figure is concerned with average alarm rates and is given as: "A long term average alarm rate in steady operation of less than 1 per 10 minutes is very likely to be acceptable".

The KPIs can be expressed per 10 minute time period to match this or, because this is a long term average, per hour, as a more familiar time period. The calculations are then performed using the appropriate time period. The functionality of the tools being used to collect, collate and analyse the data may influence the time period used.

Three KPIs are suggested:

- Average Alarm Rate;
- Maximum Alarm Rate;
- % of time Alarm Rates are outside of acceptability target.

These not only characterise the alarm system performance in a meaningful and powerful manner, but also are simple to calculate and can be generated automatically.

These apply to each operator. They are described generically and then illustrated using a 10 minute time period.

Average alarm rate

This is a simple measure of the average level of interruption imposed on the operator by the alarm system. It is calculated, over the measurement period, by:

> total number of alarms annunciated to the operator / total number of time periods

Any periods where the alarm history was unavailable are excluded from the calculation.

Maximum alarm rate

This is the worst case load during any **ten minute** timeslice. It is calculated by splitting the alarm journal into consecutive ten minute timeslices, and recording the maximum number of alarms which were annunciated to the operator during any of the ten minute timeslice. It is then expressed per required time period.

Percentage of time alarm rates are outside of acceptability target

This is a simple measure of the proportion of the time that the alarm system is outside the alarm rate target. This performance measure is useful for showing improvements made to an alarm system during alarm rationalisation. The target is based on the "Benchmarks for assessing average alarm rates" table (Figure 21) in Appendix 12. The initial target value under consideration is the "over-demanding state" rate of 1 per 2 minute period.

The KPI is calculated by splitting the alarm journal into timeslices (based on the required time period) and calculating the number of alarms which were annunciated to the operator during any of these timeslices. The proportion of timeslices where the number of alarms exceeded the target is normally expressed as a percentage. Any timeslices where alarm history was unavailable are excluded from the calculation.

Improvements to the alarm system should be reflected by this percentage figure reducing. Once this percentage has reduced to a low figure, the target can be reduced to the "manageable state" rate (i.e. 1 per 5 minute period), again with the objective of reducing this percentage, and finally to the "acceptable state" rate (i.e. 1 per 10 minute period).

For example, using a 10 minute time period, the KPIs would be:

Average number of alarms per 10 minute period:

> total number of alarms annunciated to the operator over the measurement period / the number of 10 minute periods

Maximum number of alarms per 10 minute period:

> This is calculated directly as specified in the generic description above.

Percentage of time alarm rates are outside target:

> The alarm journal is split into 10 minute periods and the number of alarms annunciated in each 10 minute period is calculated. The initial target value is 5 alarms per 10 minute period. The subsequent targets are 2 alarms per 10 minute period and then 1 alarm per 10 minute period. The proportion of periods where the number of alarms exceeded target is easily calculated as a percentage.

Note – the base calculation for the maximum alarm rate KPI is always based on a 10 minute timeslice and multiplied by 6 to give the per hour figure (rather than using a 1 hour timeslice).

There are also lower value secondary metrics that can be defined that show other aspects of the management of the alarm system, some of which are defined in Appendix 12 but have been included again here. It is important that a review of the system is undertaken prior to the setting of individual targets for improvement so that achievable targets can be set. The suggested targets listed below should be set on a plant by plant basis; the objective is to show improvement.

Number of periods of intense alarm activity

This metric captures the **peaks of alarm activity.** It is calculated by splitting the alarm journal into consecutive 10 minute timeslices, and calculating the number of alarms which were annunciated to the operator during each of these timeslices. The number of timeslices where the load exceeds the target of 100 alarms is counted. This number can then be expressed as a percentage of the whole measurement period. The aim is for this percentage figure to reduce.

Shelved alarms

There are two parts to this metric. The first part is the calculation of the number of shelved alarms. The suggested target for this is under 30. The second part is to measure the duration of each shelved alarm. The alarm management strategy should define what reviews are necessary for shelved alarms and how often they should be reviewed.

Standing alarms

There are two parts to this metric. The first part is the calculation of the number of standing alarms. The suggested target for this is under 10. The second part is to measure the length of time each standing alarm is active. It is necessary to define what a 'Standing Alarm' is for calculation purposes, as effectively any alarm active in the system could be deemed to be a standing alarm. One definition is any alarm active for a full operating shift or longer.

Top 10 load percentage

A simple figure that gives an indication if there is a good distribution of alarms or if the system is loaded by a few 'bad actors'. It is calculated by expressing the total number of occurrences of the top ten most frequent alarms as a percentage of the total number of alarm occurrences, over a set measurement period.

4.1.2 Performance Levels

As noted in 4.1 above and Appendix 12, the benchmark figures given are for alarm systems which are considered extremely manageable from the operator's point of view. In moving an alarm system towards this state, it will pass through a number of performance levels and indeed for some plants, depending on various factors, a lesser performance level may be acceptable. In general however achieving a higher performance level will deliver higher plant availability and safety.

A five level model has been devised which formalises this approach. This can be used to define an appropriate target for new systems or as a way of measuring where a system currently stands and where it is seeking to move to.

The five levels of alarm system performance range from 'Overloaded' at the bottom end of the scale, through 'Reactive', 'Stable' and 'Robust' to 'Predictive' as the highest level of performance. This is illustrated in Figure 7.

Each performance level is defined according to the set of three Primary KPIs defined above, and is also described in qualitative terms from the perspective of the control room operator. This structure allows a fuller debate of what performance level might be appropriate under different circumstances – and suggests which alarm improvement techniques might be applicable at each level.

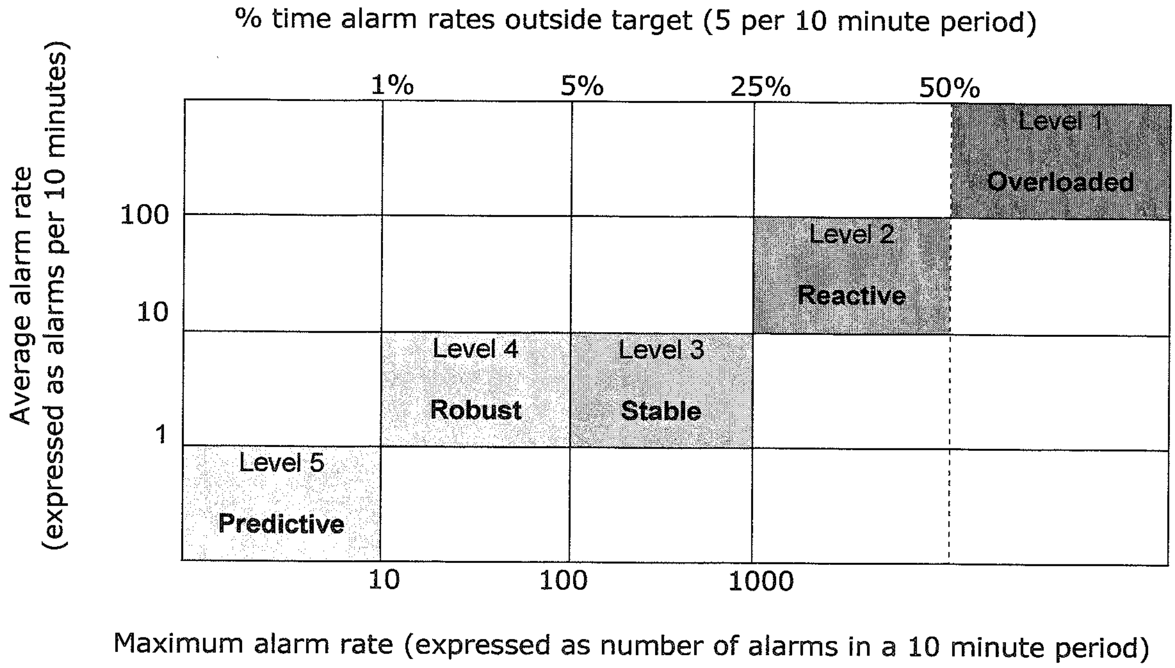

Figure 7 Performance Levels (on 10 minute time base)

The five levels of Alarm System performance are described fully in (12) and in Appendix 13, along with criteria that can be used to set an appropriate level for a particular asset. The levels can be summarised as follows:

level 1. **Overloaded.** At this level the Alarm System is subject to a continuously high rate of alarms, and deteriorates rapidly during process upset. Unfortunately this is typical of many modern installations with DCS systems.

level 2. **Reactive.** This could be considered the minimum 'entry level' for most plants. It is, typically, representative of a new DCS that has been implemented with the minimum of best practice, or an existing system that has received some initial attention – particularly with regard to the 'bad actors', those few alarms that contribute consistently with no real meaning. Some improvement has been made to the average alarm rate, by comparison with Level 1, but the peak rate during upset is still unmanageable and the alarm system will continue to represent an unhelpful distraction to the operator for long periods.

level 3. **Stable.** Typically, by careful selection of which variables to alarm, either via a rationalisation exercise or via robust engineering of alarms up-front during a project phase, improvements have now been made to both the average alarm and peak alarm rates, by comparison with Level 2. Problems due to 'bad actors' have been kept under control by regular review and continuous improvement, but there still remains a problem with the burst alarm rate. In general the alarms have been well defined for normal operation, but the system is less useful during plant upset.

level 4. **Robust.** Possibly at the limit of what is achievable with commercially available technology today, this level of performance represents a realistic target for most plants. Both the average and the peak alarm rates are under control, the latter under the full range of foreseeable plant operating scenarios. The use of dynamic techniques to improve the real time performance of the alarm system is likely to be extensive.

level 5. **Predictive.** For many plants this may not be achievable today with commercially available control technology. Even when achievable, it may not be justified for all plants. It will require fully adaptive alarming, whereby the alarm system predicts the future state of the plant and adjusts its configuration to meet the needs of the moment.

A number of processes and provisions of hardware and/or software are suggested in the performance level tables in Appendix 13 as a mechanism for improving from one performance level to the next, for a plant of typical size. Whilst it is recognised that powerful results can be obtained from parallel implementation of improvement methods from different levels, it is expected that each performance level will only be fully achievable if all elements of the preceding levels have been accomplished. For example, advanced alarm dynamic methods are unlikely on their own to achieve a robust system (Level 4) unless rationalisation improvements are complete (Level 3). Advanced alarm handling techniques are almost certain to be inappropriate as a solution for poor basic practices, since improving alarm system performance is a strongly hierarchical task.

4.2 Data analysis tools

When starting on any alarm rationalisation project, it is necessary to understand the scale of any problems that exist within the alarms system and one of the first acquisitions that is likely to be beneficial is an alarm logging and analysis tool.

A number of the metrics identified in Appendix 12 involve statistical analysis of the alarms that occur, i.e.:

- number of alarms over a defined period or after a defined event;
- most frequent alarms over a period/after an event;
- counts of standing alarms at defined times;
- counts of shelved alarms at defined times;
- identification of longest standing alarms over a defined period;
- proportion of alarms at each priority during a defined period;
- measurements of operator acceptance times;
- auto- and cross-correlation of alarm records.

For some annunciator systems the collection of data may be difficult and may preclude all but the most basic statistical analysis. For other systems the data analysis can be done manually from a printed alarm log, but again this tends to be laborious. It is recommended that:

Tools should be provided for routine statistical analysis of alarms.

There are two main components that form these data analysis tools, the first being the data gathering and storage of alarm and event data, and the second is the analysis of the stored data. Some systems effectively undertake this analysis on line as the alarm data is being gathered and do not rely on historically stored data.

The logging of alarms by the alarm analysis tools can be undertaken in a number of ways. Three typical methods are:

analysis within the alarm processor - many alarm systems store a record of the alarm log, so this is an obvious place to carry out the data analysis if suitable tools are available;

analysis using standard PC software - many alarm systems allow the alarm log to be exported as data files that can be loaded into a standard PC database. Statistics can then be computed using suitable PC tools;

analysis using proprietary PC software - many proprietary software packages are available. These packages tend to be connected via OPC (open connectivity), but in their basic form can be a simple printer replacement. They take the output that is directed to the alarm and system printers and sort the data so it can be analysed.

It is important to note that some metrics cannot be calculated just from a simple log of alarm changes. For example, if the number of standing alarms is calculated by examining a section of such a log, this count will exclude alarms which remained standing throughout the period. To overcome this, the log has to be enhanced (e.g. by printing 'midnight snapshots' of all standing alarms). This may imply changes to the alarm logging software which may be difficult or costly on some systems.

It should also be noted that, if users want to look back several weeks or months (e.g. to search for occurrences of specific alarms), then the alarm archive must be big enough to hold the necessary data. This may require large storage devices.

4.2.1 Specification of Alarm Logging and Analysis Tools

The alarm data logging tool should provide all the necessary information that is required to generate the performance indicators specified above, either providing a mechanism to export the data to a software package capable of performing the calculation (such as a spreadsheet), or ideally the data logger will automatically generate the key statistics itself.

The following is a list of some of the facilities that could be considered when purchasing an alarm logging tool:

- simple to connect to existing system;
- ability to integrate diverse systems;
- the duration of data collection and archiving;
- provide reports as standard (without engineering effort);
- present data in graphical and tabular form;
- ability to export data to Excel and other packages;
- tools to sort or filter data on any pattern or time frame;
- analysis tools to provide statistical information on any pattern. As minimum:
    - frequency analysis;
    - alarm rate;
    - pattern distribution (i.e. priority, alarm type, section, operator etc.);
    - operator response time (time form alarm to acknowledgment);
    - reaction time (time from High to High High alarm and similar).

The alarm system should provide long term storage of alarms and events. The storage should be based on a database and should as minimum include storage of the following parameters for each alarm:

- timestamp;
- activation timestamp (i.e. the alarm timestamp if this was an acknowledgement);
- tag name / object name (unique reference that identifies the alarm);
- description (descriptive text related to the tag name)
- condition and state (["Alarm", "On"], ["Valve", "Opened"] etc.);
- event/alarm type ("Event", "Alarm", "System alarm" etc.);
- priority level or severity;
- process sectioning.

The activation timestamp will make it much easier to calculate Alarm Management parameters like:

- alarm duration;
- operator response time.

As part of selecting the alarm logging and analysis facilities consideration should be given to a number of metrics that may be required to be reported. If these metrics can be generated automatically by the system, then this will save engineering effort at each reporting period. A typical analysis report for a defined time period might include:

- alarm rates (including KPIs);
- nuisance alarms
    - top 10-50;
    - percentage of total;
    - comparisons day, month, year;
- shelved alarms
    - count per period;
- standing alarms
    - count of active alarms for period;
    - duration of each alarm;
    - identification of longest standing alarms over a defined period;
- measurements of operator acceptance times;
- operator acceptance times;
- prioritisation
    - alarm configured per priority (often provided from the control system);
    - alarm rates per priority;
- alarm configuration enable/disable status;
- alarm configuration types usage (e.g. numbers of PVHI (process variable high), PVLO (process variable low) etc.).

Further, more advanced analysis might include:

- number of alarms over a defined period or after a defined event;
- analysis of frequent events and pattern recognition;
- most frequent alarms after an event;
- auto- and cross-correlation of alarm records;
- alarm database "enforcement" history including exceptions found.

5. Managing an Improvement Programme

Section overview

The previous section has identified a number of metrics that may be used to assess the performance of an operational alarm system. As indicated, these can be used as performance targets to guide an improvement process. In this section the improvement activity is considered in more detail. Key points are:

- *success depends on real and continued management commitment;*
- *operating staff should be deeply involved in the improvement process;*
- *the improvement process should be structured and driven by performance metrics;*
- *the improvement programme should address both normal and upset operation;*
- *simple techniques can eliminate many nuisance alarms, but these have to be applied by individuals who understand plant operation;*
- *alarm system changes should be controlled under a graded set of plant modification procedures.*

5.1 Culture of Improvement

A key aspect to achieving improvement in alarm systems, as in many other things, is for the organisation to propagate a culture which encourages improvement. This requires a real commitment by the senior management of the plant. The objective of improving the alarm systems needs to be made clear to all staff involved, and they need to be helped and encouraged to develop a coherent and co-ordinated strategy for achieving this. In addition a single individual should be given overall responsibility for managing the alarm system (e.g. control of changes, management of records and documentation) to ensure consistent standards are set and maintained.

The results of the survey carried out for the HSE (5) indicated that there was a positive correlation between management commitment to improvement and achievement of good results.

A number of techniques exist for improving alarm system, some of which are quite sophisticated using complex logic, knowledge-based processing, etc. However, a key message is that, typically, many of the problems with existing alarm systems are relatively basic and can be dealt with using simple techniques. For success, these should be implemented by people who fully understand the actual operational practices on the plant and the engineering implications of change. Examples of these techniques are given in Appendix 5 and Appendix 6 and in Section 5.4.

There is no technological universal remedy that can be applied to alarm systems which will provide instant and universal improvement in performance. Experience has shown that, first, there needs to be a management determination to achieve improved performance, second, there needs to be thorough application of the basic improvement techniques, and third, if performance is still not good enough, the more advanced techniques should be applied (see Table 7).

IMPROVEMENT STRATEGY
• management determination to achieve improvement • thorough application of basic techniques • if performance is still not adequate, apply advanced techniques

Table 7 Elements of an improvement strategy

5.2 Team Composition

A powerful tool for achieving alarm system improvement is to assign the task to specific teams or individuals. The primary end users of an alarm system are the operators and supervisors; the secondary users are maintenance personnel, engineering staff, etc. The alarm system is there to help these people do their jobs better. Whilst these staff may not be readily available, it is strongly recommended that:

> **The operators and supervisors should be involved deeply in any programme to improve an alarm system.**

A typical team composition might be:

- one or more operators from each shift;
- supervisor;
- instrumentation engineer;
- alarm processing system engineer;
- process/mechanical/safety engineer.

The activities of such teams may be clearly focused by setting performance targets such as "Reduce the long term average alarm rate to one per 5 minutes". Such performance measures are excellent for assessing performance and monitoring progress. However, care should be taken not to over-emphasise these (e.g. by linking them to bonus pay), so that they are achieved regardless of the broader operational and safety implications.

It is recommended that procedures are introduced whereby operations staff are positively encouraged to report operational errors. This should be on a 'no-blame' basis. Human error is an increasing contributor to accidents in many industries. 'No-blame' reporting assists in the identification of the reasons for these errors, which are mostly due to limitations in system design or operator training, so that steps can be taken to stop them happening again. The value of 'no-blame' reporting by nuclear power plant operators and aircraft pilots is very well proven.

5.3 Alarm Review

A key part of the process of improving the alarm system is for competent individuals or teams to review the alarms. The review should cover every alarm in the system, however it may be beneficial if it is organised such that the most important work is tackled first, e.g.:

- use performance measures to identify the nuisance alarms (e.g. most frequent, most likely after a trip) and progressively work through these, or;
- start with the safety related alarms and work down through the priorities[26].

The review should identify changes to be made to the alarms (and possibly the alarm system) and is likely to involve several of the techniques described in this Guide. The review should also identify individuals to do the work and set timescales and priorities for implementation (see Table 8). These decisions should be recorded in writing.

THE OBJECTIVES OF ALARM REVIEW

- minimise the number of alarms consistent with proper protection of people, plant and the environment;
- ensure that all alarms are relevant, truthful and understandable at all times;
- ensure that alarm rates are manageable at all times;
- ensure that all alarms have defined responses (this thus includes the separation out of event messages);
- ensure that alarms are properly prioritised.

Table 8 The objectives of alarm review

A key point about an alarm review process is that it should address both the alarms that occur in normal operation and those that occur following a major plant upset. These two activities require different approaches as shown in Table 9. The review of normal operation is generally easier, but upsets should not be neglected as they often involve significantly greater risks.

CHARACTERISTICS OF ALARM REVIEWS	
NORMAL OPERATION	PLANT UPSETS
• considerable data available to be collected and analysed • similar situations tend to recur • improvements can be monitored in metrics such as reduced alarm rates	• rare events (but, generally, the most risky) • each upset tends to be different • each upset needs in-depth analysis • sometimes difficult to demonstrate that improvements have worked

Table 9 Comparison of characteristics of alarm reviews for normal operation and plant upset

The alarm review should be structured to ensure it is thorough. This may be done by working through a list of questions for each alarm to examine, e.g. the purpose of the alarm, the implications of it being missed, the values of parameters associated with generating it. A suggested list of questions is given in Appendix 2.

Ideally, the answers to many of these questions relate to design intent that should have been formally documented when the plant was originally designed

[26] Alternatively fault studies, design specifications and HAZOP reviews should be re-examined to ensure that all events requiring operator intervention do have appropriate alarms.

and constructed (see Section 6). In addition, as part of the alarm system modification procedures, this design documentation should have been updated during the life of the plant whenever modifications were made (see Section 5.5). In this ideal situation, the design documentation should be further updated and extended as a result of the alarm review.

In practice, if a review is carried out some time into the operating life of the plant, the review team may find it hard to determine the original design intent behind many alarms. In this case, the answers from the questions asked in the review should be formally recorded and used as the basis for establishing an information database for the alarm system.

5.4 Effective Techniques

Alarm reviews will identify a variety of problems. This will include nuisance alarms which need to be modified in some way to stop them being a nuisance, as well as more general difficulties, such as incorrect settings, inadequate alarm messages, repeating alarms. This section provides a list of things that can be done to improve alarms. In doing this, it brings together techniques that have been mentioned in several different places in this Guide.

Often it is found that many alarms can be greatly improved by doing quite simple things. To reflect this, the techniques have been very crudely ranked in Table 10 according to the likelihood of getting benefit from the technique versus the effort in applying it. However, this crude ranking should not be taken to mean that the lower ranked techniques should not be applied. It simply means that if there are limited resources, the higher ranked techniques are the ones that are generally worth trying first. It is reiterated that review of performance during upsets is particularly important, which is why it has been placed top of the list.

5.5 Control of Modifications

Since the alarm systems are part of the defence of the plant against hazard, any changes resulting from alarm reviews need to be carried out in a responsible way. Thus all proposed changes should be fully analysed, their consequences should be determined, and agreed changes should be recorded with reasons. There should be a formal change control procedure. Changes should be recorded on an appropriate form and approved by identified suitably competent persons, e.g. the alarm systems engineer, the operations supervisor and the safety manager. Thus:

> **There should be defined procedures to control changes to the alarm system.**

In practice, certain changes to the alarm system have low safety significance. For example, changes to settings or deadbands on low priority alarms may not merit stringent review by several individuals, and relaxed procedures may enable the system to be optimised more quickly and with less resource expenditure. By contrast, all changes to safety related alarms should be carefully considered, perhaps including off-site approval. Consequently, a graded modification procedure is appropriate. This should ensure that changes to safety related and other important alarms are identified and carefully controlled.

TECHNIQUES FOR IMPROVING ALARM SYSTEMS

High benefit

- review alarm behaviour following all upset incidents to confirm usability
- tune alarm settings on nuisance alarms
- adjust deadbands on alarms which often repeat
- eliminate alarms which have no defined operator response
- ensure critical and high priority is allocated to appropriate alarms
- review alarm messages which operators do not understand or know how to respond to
- introduce an alarm shelving facility
- introduce single line annunciation of repeating alarms on alarm list displays

Medium benefit

- suppress alarms from out of service plant
- replace nuisance absolute alarms on controlled variables with deviation alarms
- apply filtering, transient suppression and de-bounce timers to repeating alarms
- replace digital alarm sensors causing nuisance with analogue sensors
- install automatic control/protection to reduce the operational importance of alarms
- redesign actuator discrepancy alarms causing nuisance
- re-engineer alarms from 'bad' signals so that they do not cause nuisance
- introduce logic to combine and simplify redundant sets of alarms
- introduce logic for eclipsing multi-level alarms (e.g. high and high-high)

Other

- introduce automatically adjusted alarm settings
- introduce operator set alarms
- apply counters and auto-shelving to repeating alarms
- introduce logic to dynamically re-prioritise alarms
- group alarms which all need the same operator response
- automatically suppress alarms according to the operating mode of the plant item
- develop intelligent logic for identifying the most important alarms

Table 10 Effective techniques for improving alarm systems

The Site Management Strategy document (see Table 4) should include a description of alarm modification procedures within it. This would define how priorities are set, the procedures for changing settings and which alarms need special review prior to being changed, etc.

6. Buying a New Alarm System

Section overview

This section discusses issues relating to procurement of a new alarm system plus associated engineering. It is relevant both to procurement for a new plant and for an alarm system replacement exercise.

6.1 Investment Appraisal

Investment justification is a critical issue in alarm systems. Methods are needed for determining how much is worth investing in alarm systems. This applies to new plant build, to alarm system replacement and to alarm system improvement. Few companies are going to make investment without some demonstration that it is worth it in terms of improved safety or profitability.

Unfortunately, it has to be recognised that it is, generally, difficult to develop rigorous cost-benefit analyses for proposed investments in alarm systems. Operating experience shows that a consequence of this tends to be under-investment of money and design effort, leading to avoidable risks to people and significant lifetime financial losses. There are major benefits to be gained at many plants from improving the alarm systems and ensuring that operators are not overloaded with nuisance alarms, and are able to properly investigate and respond to all the alarms that arise.

Appendix 16 presents some data on incidents involving alarm system shortcomings. This provides a general pointer to the size of the safety and financial benefits of ensuring good alarm system performance. However, it does not provide a measure of the difference in benefit between a 'good' alarm system and a 'bad' one, nor does it give information on the benefits likely to be achieved at a particular site.

To estimate the site-specific benefits from alarm system investment the following approaches may be used.

Collect records of incidents

Incident data should be collected, recorded and analysed in a structured manner by nominated personnel. This allows qualitative and quantitative analysis to be performed, which in turn allows key areas of interest to be identified and the result of any improvements to be monitored. Incident data may be collected:

- At the site of interest. Analysis of this data should identify the contribution that alarm systems make to these incidents and calculate statistics on the financial losses and hazards involved;
- Across sites within the same company. This will provide statistically better estimates of, e.g. the frequency of large incidents;
- Across the industry. Cross-company incident reporting provides a long term driver to alarm systems and overall plant safety improvement. Examples of such activities can be found in the chemical and nuclear industries where national and international data exchange occurs.

Analyse avoidable loss

- Assess the size of the total avoidable operational loss (see refinery example in Appendix 16).

Measure alarm system performance

- Measure the usability of the alarm system. This may be done by informal discussion with operating staff or more formally by use of operator surveys or other performance metrics. Appendix 12 lists some metrics and provides some benchmark figures.

Predict likely cost of ownership

- Cost out the time spent handling alarms and curing alarm problems. In addition, cost production losses and losses due to inefficient operation.

It will take some significant effort plus a long term company commitment to collecting data to use the above approaches to develop a rigorous justification for investment. Conversely, experience shows that the typical consequence of current custom and practice is poor alarm system performance, risks to people and avoidable financial losses. Thus, if it is not practical to carry out a rigorous analysis, the approach needs to be more systematic and thorough than current practice.

6.2 Contractual Implications

There are numerous contractual options for the procurement of plant equipment such as alarm systems, and a full discussion of the pros and cons of the different approaches is outside the scope of this document. However, there are a number of strategic issues specifically relevant to alarm systems that are worth exposing.

As was indicated in Table 2, the development of a new alarm system involves both the purchase and installation of equipment plus considerable engineering design and configuration. Some of the key contractual issues are:

Allocation of activities

The contract strategy should ensure that all of the above activities are carried out and are allocated to parties with the appropriate skill and experience. Expertise required will include:

- plant process design;
- field equipment design;
- alarm system configuration;
- ergonomics;
- plant maintenance; and
- plant operation.

Scheduling of activities

The engineering practices recommended in this Guide are considered to be the best available in terms of minimising total lifetime costs and, thus, should be adopted when developing a new alarm system. However, there may be options as to when some activities are carried out. For example, it would be possible to initially install a minimal alarm system and then extend and optimise it early in the operational life of the plant. Such an alarm system is likely to cost more in total, but in some cases it may have cash flow and contractual benefits that may make it commercially attractive. If contractual approaches like this are chosen, care should be taken to ensure that any interim alarm system is fully effective and usable in its handling of safety-related alarms.

Contractual tests of acceptability

Many types of contract need a clear test of acceptability of the delivered product. This can be difficult to define for an alarm system. In addition, it generally needs a reasonable period of plant operation for the performance of the alarm system to be optimised. For both these reasons, it can be hard to sustain acceptance tests on alarm system performance within the context of a complete plant construction contract, i.e. they are too vague, and performance is still being improved when the contractors are pressing for plant handover. Contract strategies should take account of these difficulties. For example, post-commissioning alarm system optimisation might be purchased as a defined package outside the main plant construction contract.

6.3 Specifying Alarm Functionality

The procurement of a new alarm system often involves the purchaser in writing a specification of required functionality or in comparing the functionality of different systems that are offered. A checklist of required alarm functionality has been prepared which has categorised particular functions as 'essential', 'valuable' or 'possible'. This is included in Appendix 17.

This checklist can be used as a basis for writing a procurement specification. Alternatively, purchasers can require those providing tenders to bid against the checklist and identify any functions in the 'essential' and 'valuable' categories they do not provide.

Purchasers in the power industry are also referred to European standard EN 45510 Part 8-1 (6) which provides guidance on the writing of functional specifications for power station control and instrumentation equipment (including alarm systems).

6.4 Specifying Engineering

The checklist referred to above covers the functionality of the alarm handling system. However, it is possible to buy an alarm system with excellent functionality, and end up with very poor performance if the system and the individual alarms within it are not engineered properly. Consequently, in competitive tender situations, care should be taken to ensure that good engineering is not diluted by the competitive process. This may require the purchaser to find some way of specifying what engineering is required.

There are two approaches to this:

Specification of design procedures

Here the procedures to be followed are specified. For example, the purchaser might specify that the supplier shall carry out a safety study, each alarm shall be documented using the headings identified in Appendix 2, certain of the techniques defined in Appendix 8 and Appendix 9 shall be implemented, and defined quality control assurance procedures shall be applied. Detailed information may be given on what is involved in each activity. The risk here is that the supplier will follow all the steps, but only as a minimum and so will be able to argue that the contract has been met; but the overall performance will fall short of what was wanted.

Specification of required performance

In this approach the purchaser specifies some required performance tests that the alarm system is required to pass. These might involve some of the metrics identified in Appendix 12. It is then left to the supplier to find the most economical way of meeting this performance. The risk here is that the performance tests will be passed, but the real usability of the system will be low. This is because usability is hard to specify in a quantifiable way.

It is recommended that, for effective procurement, a mixture of both approaches should be used. Both the performance required and the steps to be followed to meet that performance should be specified.

This will involve more effort than is commonly expended by procurers at present. However, many existing alarm systems are unsatisfactory in that too little has been invested in initial engineering compared with what is lost in the plant lifetime in avoidable incidents. It follows that it will be worthwhile for users to specify more fully what they want to be supplied and to accept that their alarm systems will have a higher initial capital cost than they do at present.

6.5 Ensuring Usability

The alarm system should support the user in his tasks. Therefore, it should be designed to meet the user's needs and operate within operator capabilities. However, the alarm system is only one part of the operator interface provided in the control room. Thus, the complete interface should to be designed to best ergonomics practices as an integrated and usable system. To help to achieve this, the purchaser should specify that the supplier follows a user-centred design methodology. This might include elements such as those given in Table 11.

ELEMENTS IN A USER-CENTRED DESIGN

- design review. The supplier and the purchaser set up a working party comprising designers, operators and ergonomics specialists that reviews the operator interface design through the various stages of design, construction and commissioning;
- task analysis of all operator activities to identify the tasks the operator performs and the information and controls required to carry them out;
- development of design standards for the operator interface covering areas such as alarm message structure, design of graphics, use of colour, etc.;
- operator involvement in the testing and evaluation of prototype designs;
- early ergonomics evaluation of prototype designs by the purchaser;
- pre-commissioning operator training in use of the operator interface;
- formal demonstration of the usability of the commissioned operator interface;
- on-going monitoring and improvement of performance.

Table 11 Elements in a user-centred operator interface design

Appendices

Appendix 1 Glossary

This Appendix provides a glossary of some of the terms used in the Guide. It is not intended to be comprehensive, but to concentrate on those terms that are likely to be unclear or cause confusion.

Definitions	*Description*
Abnormal situation	A disturbance or series of disturbances in a process that cause plant operations to deviate from their normal operating state.
Abnormal Situation Management (ASM)®	Abnormal Situation Management and ASM® are U.S. trademarks of Honeywell International.
Accepted	Alarm state: an alarm is accepted when the operator has indicated awareness of its presence (usually by push button or mouse click). It is unaccepted until this has been done.
Acknowledge	The operator action that indicates recognition of a new alarm.
Active Alarm	An alarm condition which is on (i.e. limit has been exceeded and condition continues to exist).
Alarm	An audible or visible means of indicating to the operator an equipment or process malfunction or abnormal condition.
Alarm Deadband	The range through which an input must be varied from the alarm limit necessary to clear the alarm.
Alarm Flood Alarm Overload	The situation where more alarms are received than can be physically addressed by a single console operator.
Alarm Limit Alarm Threshold Alarm Trip Point Alarm Set Point	The threshold value or discrete state of a process variable that triggers the alarm.
Alarm Management	The processes and practices for determining, documenting, designing, monitoring, and maintaining alarm systems.
Alarm Message	Text information presented to the operator that describes the alarm condition.
Alarm Priority	The ranking of alarms by severity and response time.
Alarm Processor	Refers to the part of the system for processing and displaying alarms.
Alarm Rationalisation	A process whereby a multi-function team determines what alarm configuration (priority and settings) is required for individual parameters in the control system.
Alert	A lower priority notification than an alarm, that has no serious consequence if ignored or missed. In some industries also referred to as a Prompt or Warning.
Alarm Response Time	The time between the process condition becoming abnormal and the initiation of the alarm state.
Alarm System	Refers to the complete system for generating and handling alarms including field equipment, signal conditioning and transmission, alarm processing and alarm display (it also includes hardware, software and supporting information (e.g. alarm response procedures, management controls)).
Cleared	Alarm state: an alarm is cleared when the condition has returned to normal.
Competency	The sufficient knowledge and skill required to effectively perform an activity or task.

Definitions	*Description*
Console	The interface for a single operator to monitor the process.
Control Room Layout	The physical organisation of equipment (consoles, tables, radios, etc.) in a control room.
Critical Alarm	The highest level of alarm priorities - immediate operator action is required or a serious plant incident will occur.
Deadband	When a deadband is applied, then the alarm is arranged to be raised at one level but cleared at a different level.
Dependent Failure	A failure which is in some way related to or dependent on another failure. Terms that were used in the past to describe this were 'common mode failure' or 'common cause failure'.
Disable	An alarm is disabled when the system is configured such that the alarm will not be generated even though the base alarm condition is present.
Emergency Shutdown (ESD) System	An automatic protection system which will act to shut down the plant if it enters a potentially dangerous state. In some countries this is called a Safety Instrumented System (SIS).
Event	A change in plant or equipment condition.
Grouping	A single grouped alarm may be used to display a number of different initiating events from a plant system.
Hazard and Operability Studies (HAZOP)	A structured analysis technique to assess the hazards and operability of a process design.
Health and Safety Executive (HSE)	The GB's Health and Safety Executive (HSE) is responsible for the regulation of almost all of the risks to health and safety.
Inhibit	To manually or automatically prevent the transmission of the alarm message to the operator.
Initiating event	A failure or other condition that can cause an alarm.
Alarm Log	The historical record of all alarm messages.
Nuisance Alarm	Alarms which do not generate a specific action or response from the operator.
Operating Procedure	A set of explicit guidelines and instructions to be followed by the operator.
Operator	A member of the operations team who is assigned to monitor and control a portion of the process and is working at the control system's console.
Prompt	A request from the control system that the operator perform some process action that the system cannot perform or that requires operator authority to perform.
Operator Response Time	The time between the annunciation of the alarm and when action is required to prevent the consequences of the alarm related event.
PFD_{avg}	The average probability of a system failing to perform its design function on demand. For an alarm system the 'design function' would be to generate an appropriate alarm, and 'on demand' would be on those occasions on which it should be generated (also see IEC 61508 (29)).
Prioritisation	The process of assigning to an alarm a level of importance that can be implemented within the alarm system.
Raised	An alarm is raised or initiated when the condition creating the alarm has occurred.
Remote Alarm	An alarm from a remotely operated facility.

Definitions	*Description*
Reset	An alarm is reset when it is in a state that it can be removed from the displayed list.
Safety Related Alarm	An alarm which is claimed to provide significant risk reduction from hazards to people and which is implemented independently from the process control system.
Shelving	Shelving is a facility where the operator is able to temporarily prevent an alarm from being displayed to him when it is causing him nuisance. A shelved alarm will be removed from the list and will not re-annunciate until un-shelved.
Standing	An alarm is standing whilst the condition persists (raised and standing are often used interchangeably).
Suppress	An alarm is suppressed when logical criteria are applied to determine that the alarm should not occur, even though the base alarm condition (e.g. alarm setting exceeded) is present.
Unaccepted	An alarm is accepted when the operator has indicated awareness of its presence (usually by push button or mouse click). It is unaccepted until this has been done.
Workstation	A computer station with a display (CRT or LCD), keyboard, and pointing device (mouse, trackball, etc.).
£	UK pound
$	US dollar

Acronyms

ASM®	Abnormal Situation Management®
BPCS	Basic Process Control System
CRT	Cathode Ray Tube
DCS	Distributed Control System
ESD System	Emergency Shutdown System
HAZOP	Hazard and Operability Studies
ISO	International Standards Organization
KPI	Key Performance Indicator
LCD	Liquid Crystal Display
LOPA	Layers of Protection Analysis
MOC	Management of Change
P&ID	Piping and Instrumentation Diagram
PFD	Probability of Failure on Demand
PFD_{avg}	Average Probability of Failure on Demand
PLC	Programmable Logic Controller
SCADA	Supervisory Control And Data Acquisition
SIF	Safety Instrumented Function
SIL	Safety Integrity Level
SIS	Safety Instrumented System
VDU	Visual Display Unit

Appendix 2 Design of Individual Alarms

This Appendix provides a checklist of issues that need to be resolved in the design of each alarm. This information should be recorded when the design takes place to provide a design database for use during the lifetime of the plant.

Often on existing plants there are many problems with individual alarms. For example:

- the alarm may be using an unsuitable source signal;
- the settings may be inappropriate, the priority may be wrong;
- the alarm message may be unclear;
- the operator may not know what to do about the alarm.

Such problems often indicate that insufficient thought was expended when the alarm was first conceived. With modern computer-based alarm systems, it is very easy to introduce alarms with a tendency for them to be configured without sufficient design effort and without considering the operational benefits which they will confer.

There is also the tendency for designers of individual plant systems to propose alarms from their particular systems, without a proper regard for the operational importance of their alarms compared against alarms from other areas. Similarly, HAZOP reviews can tend to result in the proliferation of alarms as extra lines of defence against potential hazards. However, as each extra alarm is introduced, the chances of overloading the operator with alarms increases, and the alarm system overall becomes less effective as a line of defence.

To overcome these problems, a formal design process should be followed so that all the key aspects of an alarm are considered from the outset and at a time when changes can be made at minimum cost. If this is done, the up-front costs of designing the alarm system will tend to rise, but there are likely to be long term savings in system commissioning and operation.

One way to formalise the design process is to require full justification and documentation of each alarm. This can be done by use of checklists, proforma or, preferably, by building up a computer database that can then be updated as changes occur during the life of the plant. This Appendix provides a checklist of some of the information that may need to be recorded.

It will undoubtedly be time consuming to answer all of these questions for the several thousand alarms that may be required on a large plant. However, design issues should be resolved for each and every alarm, and if the design process is well managed, these decisions should be recorded.

This same checklist can be used as the basis for an alarm review. The key questions that are very important and should be answered for every alarm are as follow. The answers to these can be considered the '**minimum design documentation**':

- what is the purpose of the proposed alarm?
- what response is the operator required to make to the alarm?
- what are the likely consequences if the operator does not respond to the alarm?
- what time is available for the operator to respond to the alarm?
- how likely is it that the operator response will be effective?

A2.1 Risk Assessment/Purpose of Alarm

What is the purpose of the proposed alarm? What hazard or process risk is it intended to provide warning of? What will be the consequences if the alarm fails or is missed?

> If the proposed alarm is purely informative then it almost certainly should not be an alarm.

What is the severity of the risk in terms of potential loss of life or injury, environmental impact, plant damage or economic loss?

> Any hazards to people should be linked into the formal risk assessment for the plant. For economic risks it is advantageous if the size of the potential plant damage or loss can be expressed financially. One major chemical company suggests the cost figures in Table 12 for repairs to damaged pumps, compressors, etc.:

Item	average repair cost
small pump e.g. injection pump	£2,000
medium pump e.g. reflux pump	£5,000
large pump e.g. amine circulator	£10,000
major pump e.g. MOL booster	£25,000
compressor	£200,000
gas turbine	£1,000,000
heat exchanger	£50,000

Table 12 Typical equipment repair costs

The cost of any lost production or business loss should be added to any equipment repair costs. These are often many times greater than repair costs. Similarly, uninsured losses are often many times greater than insured losses. The HSE has published guidance on the cost of accidents at work (23).

How frequently is the risk likely to occur?

> If it is difficult to be specific, it may be appropriate to select from:

- once a week;
- once a month;
- several times a year;
- once a year, once in 3 years;
- once in 5 years;
- once in 10 years;
- less than once in 10 years.

Is the plant also protected against the risk by protective systems (either mechanical systems - such as relief valves - or instrument-based systems)? If not, should a protective system be used rather than, or as well as, an alarm?

Are any reliability claims made in the plant Safety Case for the protection supplied by this alarm? Are these such that the alarm should be classified as safety related? If the alarm is not safety related, is the economic/environmental risk such that there would be benefits in implementing the alarm outside the process control system?

> As discussed in Section 2.3, safety related alarms should be implemented outside the process control systems (unless that is itself safety related). There may also be ergonomic advantages (and possibly reliability advantages) in displaying some important environmental/economic alarms in particularly obvious ways, e.g. on individual annunciators.

What will be the implications of failure of the instrument generating the alarm? Can such failures be detected? Is it possible to validate the input signal? Should the instrument be duplicated?

How likely is it that the operator response will be effective? What is the likely stress level of the operator when required to respond?

> If the operator cannot do anything to prevent the risk indicated by the alarm, then it is providing little benefit and should not be an alarm.

A2.2 Prioritisation

What are the likely consequences if the operator does not respond to the alarm?

> As discussed in Section 2.5.1 and Appendix 5, safety, environmental and economic consequences should be assessed. The direct consequences of not responding to the alarm should be considered, and also the possibility of equipment failing and causing injuries or damage. This requires some estimation of equipment reliability.

What time is available for the operator to respond to the alarm? Is the alarm 'time critical'?

> Some alarms, such as approach to trip alarms, may require urgent response if a large economic impact is to be avoided (however, this may depend on how quickly the alarmed variable is moving in the dangerous direction). Other alarms, e.g. an alarm indicating low thermal efficiency on a power plant, may indicate that costs are small but steadily accumulating. In Appendix 5, an example is given in which an alarm is classified as time critical if making no response within 3 minutes is likely to be too late. Other definitions of what is time critical may be used if different prioritisation rules are defined.

What priority should be allocated to the alarm?

> This will depend on the answers given above on the severity of the consequences of missing the alarm and on how quickly the operator needs to respond. Further guidance on prioritisation is given in Section 2.5.1 and Appendix 5.

Should the priority be automatically changed according to operating conditions?

A2.3 Operator Information/Response

What response is the operator required to make to the alarm?

> As indicated in Section 1.3, the response may be an action, a conditional action or a cognitive switch. It is important that the response is clearly

defined for each alarm. Note that it is useful to identify if this response should change according to the circumstances, e.g. whether the same response should be taken in normal operation and in the middle of a plant upset. If a response cannot be defined then this signal should not be used as an alarm.

What is the alarm message?

There is benefit in using alarm messages that are easy to read and understand. To this end it is desirable if guidelines are developed on the standard format of alarm messages and on the terms and abbreviations to be used within them. For example:

- superheater should not be abbreviated in one message as s/heater, in another as S/Htr, and in yet another as S/H;
- messages should not be formatted in some place as "process variable, state" - e.g. "S/H outlet pressure high" - and in other places as "state, process variable" - e.g. "high S/H outlet pressure";
- text descriptors of process variables are preferable to tag numbers.

What information does the operator need to decide how to respond? Is additional messaging (text or graphic) required under different conditions, or to clarify rare/complex alarms?

There may be a need to develop special displays that collect together this information and present it in an easily accessible form. On a programmable display system it may also be useful to have 'shortcuts' from the alarm list display to these graphics.

How long will it take for the operator to respond to the alarm and for the plant to respond to the corrective action?

If the diagnosis of the cause of the alarms or the implementation of the response is very drawn out then there may be a need for improved operator support tools. Assessing likely response times may also be of value in estimating the acceptable rate of presenting alarms - see Section 1.3.

What form of alarm response procedure should be provided?

A2.4 Alarm Setting

What is the normal value of the alarmed process variable? What is the value at which economic loss, hazard or environmental damage will occur?

What is the alarm setting? Does there need to be more than one alarm, e.g. a high and a high-high? Does the alarm setting need to change according to operating conditions? Does the alarm need to be a combination of time and extent over the limit (e.g. as in overheating of a motor due to high current)?

How much does the alarmed process variable fluctuate in normal operation? How much does it fluctuate in plant disturbances which do not bring the alarmed plant item into a hazardous state? What deadband would be suitable in generation of the alarm?

A2.5 Suppression

Will the alarm come up in a large plant disturbance or trip? Should logic be used to suppress the alarm?

Are there other circumstances when the process variable will exceed the alarm setting and not represent a risk (e.g. when starting up or shutting down the plant item)? Should the alarm be suppressed in these circumstances?

Should this alarm be suppressed if other more significant alarms occur?

What signals should be used to trigger the suppression?

> Alarm suppression is discussed in Appendix 8.

Will the process variable ever become invalid (e.g. go out of range, become faulty)? What effect will this have on the alarm?

> Out of range variables can be a major cause of repeating alarms (see Appendix 9). If a signal is hovering around the maximum instrument range and going in and out of the invalid state, then this repeating alarm may not be eliminated by use of a deadband. This is because most existing systems apply deadband to alarm settings, but not to invalid settings. Appendix 10 includes discussion of signal validation.

A2.6 Management Control

What procedures are required to change the alarm settings? Should operators or supervisors be able to do this?

> Site procedures should address the authorisation of alarm setting changes.

Is it acceptable for an operator to temporarily shelve this alarm?

> It is generally recommended that alarm shelving facilities should be available. However, on some sites the operator may not be authorised to shelve alarms, or special procedures may have to be followed for shelving certain important alarms. Site procedures should address the authorisation of alarm shelving.

Will the alarm require testing? How will the alarm be tested? How will the alarm be maintained?

> Guidance on testing and how to decide if testing is appropriate is given in Section 3.6, and on choice of sensor in Appendix 10. These are very important practical questions. For example, how do you test a high level float switch installed directly into a vessel containing material at a particular temperature/pressure, except by bringing the level up to that limit? Is it safe or practicable to do this, and is it acceptable to disturb plant production? The answers to these questions may result in the provision of a float chamber or in the replacement of the float switch by an alternative level measurement. This may be more expensive to purchase, but cheaper to test and maintain.

Appendix 3 Quantitative and Qualitative Risk Assessment

This Appendix provides an introduction to the topic of risk assessment. It introduces quantitative risk assessment and gives an example of a fault tree including an alarm. It also discusses qualitative risk assessment and the role of HAZOP reviews.

Risk assessments can either be quantitative or qualitative.

Quantitative assessments are normally appropriate where the hazard is greatest, e.g. where there are explosive risks where fatalities are possible. When applied to alarm systems they are used to demonstrate that the protective system (in this case the alarm plus the operator response to it, plus all the other protective devices) is adequate to reduce the risk to a defined target.

There are a variety of useful tools for quantitative risk assessment including Reliability Data Analysis, Fault Tree Analysis, Event Tree Analysis and Human Error Quantification (38). However, all quantitative assessment methods require specialist skills and should be carried out by individuals who have received specific training in the techniques.

The second part of Appendix 4 provides an example of the calculation of failure rates for a safety related alarm system and illustrates the interaction between the reliability of the hardware and the operator.

Figure 8 shows an example of a quantitative risk assessment using Fault Tree Analysis for a system including an alarm. The analysis estimates a PFD_{avg} of 0.1 for the operator failing to notice/respond appropriately to the alarm. A separate PFD_{avg} figure of 0.01 is given for the failure of the alarm hardware to generate an alarm when required to do so.

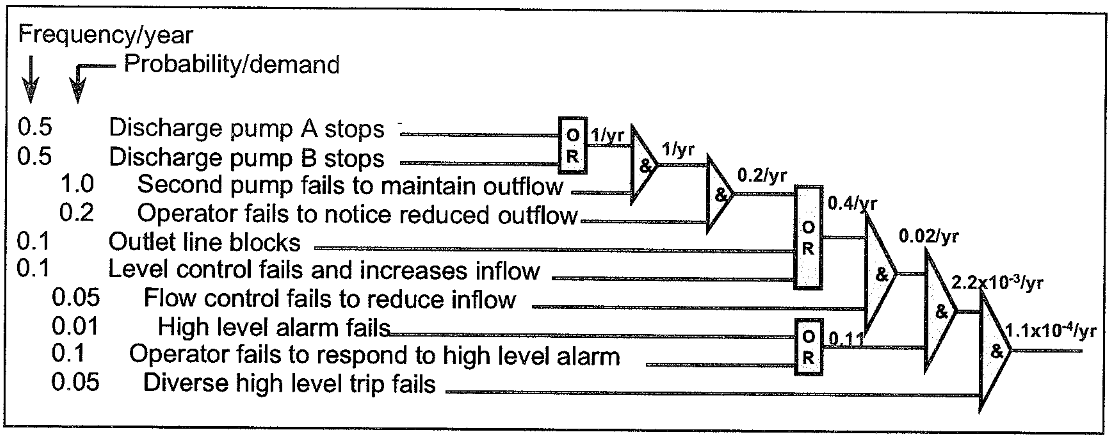
Frequency/year
Probability/demand
0.5 Discharge pump A stops
0.5 Discharge pump B stops
1.0 Second pump fails to maintain outflow
0.2 Operator fails to notice reduced outflow
0.1 Outlet line blocks
0.1 Level control fails and increases inflow
0.05 Flow control fails to reduce inflow
0.01 High level alarm fails
0.1 Operator fails to respond to high level alarm
0.05 Diverse high level trip fails
OR
1/yr
&
1/yr
&
0.2/yr
OR
0.4/yr
&
0.02/yr
OR
0.11
&
2.2x10-3/yr
&
1.1x10-4/yr

Figure 8 Example fault tree including alarm

Qualitative risk assessment is an alternative technique that can be used for assessing risks and deciding what integrity level is required of the alarm system. It involves making some assessment of the severity of the consequences of the operator not responding to the alarm and of the time available to the operator to respond. This is used to decide whether the alarm should be implemented within the control system or independent of it in a stand-alone system. A qualitative assessment can be thought of as the first stage in the process of prioritising the alarm (see Section 2.5.1 and Appendix 5). A qualitative assessment may identify some alarms as potentially safety related. These might be subjected to further quantitative assessment.

Figure 9 shows an example of a qualitative approach to deciding whether it is appropriate to use a standard alarm implemented within the process control system or to use a safety related alarm[27]. If the approach suggests implementation in a stand-alone system, consideration should be given to using alternative forms of protection, e.g. independent safety interlocks.

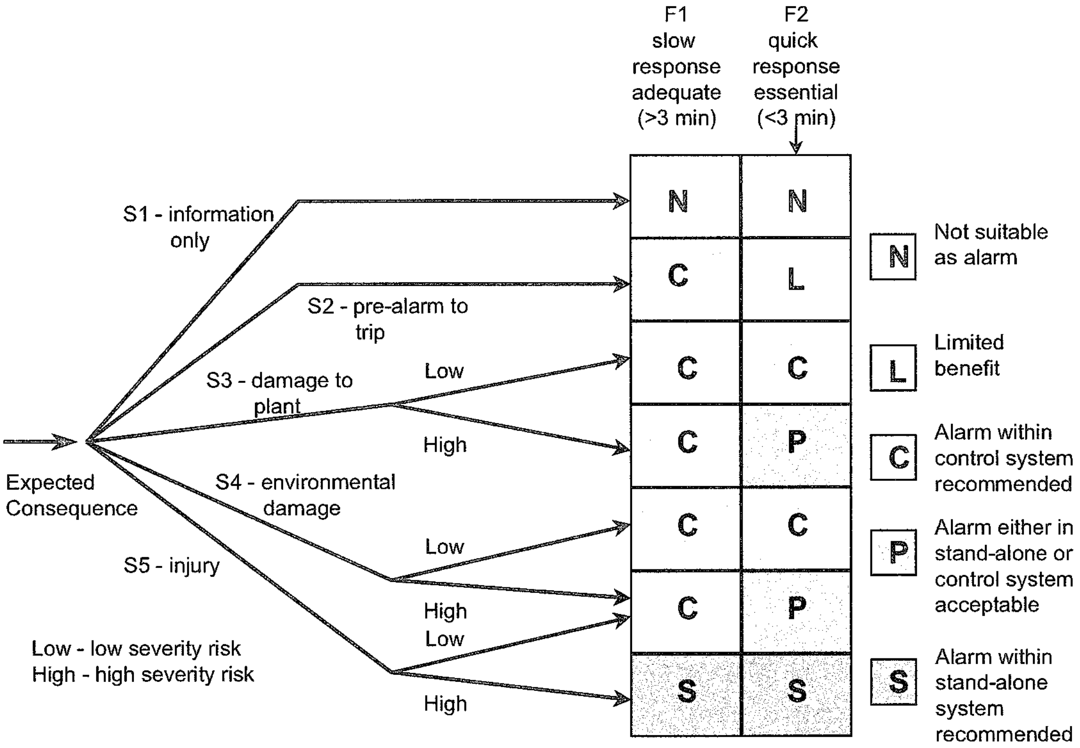
F1
slow response adequate (>3 min)
F2
quick response essential (<3 min)
S1 - information only
S2 - pre-alarm to trip
S3 - damage to plant
S4 - environmental damage
S5 - injury
Expected Consequence
Low
High
Low
High
Low
High
Low - low severity risk
High - high severity risk
N N
C L
C C
C P
C C
C P
S S
N Not suitable as alarm
L Limited benefit
C Alarm within control system recommended
P Alarm either in stand-alone or control system acceptable
S Alarm within stand-alone system recommended

Figure 9 Example of risk graph for qualitative risk assessment

The chart shows how the decision depends on the size and type of risk and the required speed of response.

Appendix 4 gives an example of a qualitative risk assessment for one alarm.

The risk assessment process may include Hazard and Operability Studies (HAZOPs). A HAZOP is a detailed and systematic review of the plant design and outline operating and maintenance procedures to identify the consequences of deviation from design intent (15), (34). It is carried out when the design has been finalised and when line diagrams (e.g. P&IDs) and operating procedures are available.

A HAZOP is a procedure for reviewing a design, not for carrying out design. If it identifies shortcomings, then the designers should be required to re-examine their design, modify it as necessary and resubmit it for a further HAZOP. Design decisions such as installing additional alarms should not be made during the HAZOP review itself.

[27] This is in effect equivalent to assessing whether the alarm has a claimed probability of failure on demand (for operator plus system) of more or less than 0.1 (see Section 2.3.4).

Appendix 4 Examples of Risk Assessment

This Appendix provides an example of a qualitative risk assessment being applied to a high current pre-trip alarm for an electric motor. It also provides an example of a quantitative risk calculation showing how the hardware failure rate, the hardware test rate and the operator reliability all contribute to the overall reliability for an alarm.

A4.1 Qualitative Risk Assessment - Example

This example illustrates a risk assessment for a high current alarm on an electric motor driving a large pump. The example is complex and illustrates a typical case of an alarm providing protection against only some potential faults. It also shows the interaction between the design of a protection system, a control system and an alarm.

A4.1.1 Identification of Risks

A potential hazard exists if the motor current gets too high. This could cause over-heating of the motor windings. Insulation could be damaged and a short circuit could occur. High currents would be drawn from the electrical supplies which would normally be detected by a suitable form of overcurrent and/or thermal overload protection relay(s).

Depending on the nature of the fault it is possible that pre-trip alarms and actual plant tripping may occur at the same instant, e.g. in the case of a short circuit. In this instance pre-alarms offer little benefit.

In the case of a short circuit, the overcurrent/earth fault protection will act along with the associated circuit breaker to clear the fault from the system. Alternatively, in fuse-contactor arrangements, the fuse would act to remove the short circuit.

In the case of a thermal overload, the motor protection would be expected to operate to remove the overload from the motor, i.e. via tripping of the associated circuit breaker or suitably rated contactor.

In the case of a short circuit, the motor may sustain extensive damage, possibly beyond economic repair. This is true even with the most modern, reliable and fast acting protection system.

In the case of a thermal overload, it would normally be expected that the thermal overload protection would remove the overload prior to motor damage occurring.

Experience shows that there are only very minor hazards to people associated with the faults. Such failures do not usually affect the integrity of the earthing of the motor body, so the risk of electric shock is low. This assumes the earthing system is correctly designed, installed, checked and maintained. There is also negligible chance of objects being projected through the motor casing and causing injury. It is extremely unlikely that the fault would cause a noise or a flash sufficient to result in injury to people. It is also assumed in this example that the motor is not located in a flammable atmosphere, so there is negligible chance of explosion.

The fault carries negligible risk of environmental damage.

A4.1.2 Protection

To protect against the minor potential hazard to people and the potential economic loss due to plant damage, motor overload protection is implemented within the motor switchgear. This may use thermocouples buried within the windings of the motor and/or measurement of current to the windings to detect when the temperature is becoming excessive. This is used to trip the motor. The motor manufacturers have carried out a risk assessment for the motor and recorded it in writing. This risk assessment would have identified all potential causes of trip (see below) and assessed the likelihood of them occurring.

A4.1.3 Automatic Control

The pump operates under automatic control. If the demand on the pump is high and the working fluid temperature is low, then an excessive current could be drawn by the motor and the protection could operate and trip the motor. This would lead to production losses. To avoid these and reduce the requirement for continuous operator surveillance, a current limiter has been installed within the automatic control.

A4.1.4 Alarm

There are a range of circumstances that could cause the motor protection to operate:

- fast-acting plant hardware or control system faults that could cause a trip within seconds. There is little benefit in protecting against these with a pre-trip alarm;
- slower acting hardware or control system faults. A pre-trip alarm would provide protection against these;
- normal plant disturbances (e.g. changes in fluid temperature). These would be contained by the limiter in the control system if it was operating on automatic. However, if the control was on manual, these disturbances could cause a motor protection trip. In some cases the disturbances will take several minutes to cause a trip so they can be protected against by a high current or thermal pre-trip alarm.

It is decided to install a pre-trip alarm. This is based on high motor current and includes filtering to provide an approximate model of the heating dynamics in the motor. It is accepted that this alarm will not prevent all trips, but it will provide a worthwhile reduction in trip frequency. It is assumed that the risk assessment for the protection has shown that that there is no need to make any safety claims for the pre-trip alarm. Hence, its purpose is purely economic and it can be implemented as a standard alarm within the control system (see Table 5). In order to protect against both slow control system faults and normal plant disturbances, it is necessary to use a different current measurement from that used by the automatic control.

A4.1.5 Alarm Setting and Priority

In this example, priority will be assessed using the procedure illustrated in Appendix 5. This means that an assessment must be made of the severity of consequence of missing the alarms and the time available for response to the alarm.

It is assumed that if the operator misses the alarm the pump would trip and cause a loss of production with an average cost £800. There is a possibility which is estimated to have a probability of 10^{-2} of the automatic trip not working. As discussed earlier, the expected safety and environmental consequences of this are considered negligible. It is estimated that the financial consequence of the trip not working and the motor being written off is £10,000. Thus, on a risk basis, the expected consequence of relying on the trip is £100 (£10,000x10^{-2}). The total economic consequence is therefore £900 (£800 + £100).

The faults being protected by the alarm were discussed above. The majority of these are assessed as relatively slow acting and thus the alarm is not categorised as 'time critical'. Based on the example figures given in Appendix 5 and shown in Figure 6 and Figure 10, this puts the alarm just above the boundary between low and medium priority. This alarm may, therefore, be provisionally categorised as medium priority. However, it would be one of the first alarms to be moved down to low priority if, after all alarms were provisionally prioritised, it was found that there were too many medium priority alarms occurring.

Given the range of faults that may be experienced, the operator should be given the maximum time to respond in order to catch the maximum number of faults. The alarm should be set a small margin above the limit of largest normal operational fluctuations (see Figure 4) to avoid spurious annunciations.

A4.2 Quantitative Risk Assessment - Example

This second example assumes that a safety related alarm is generated by taking an analogue measurement, converting it to an alarm in a trip amplifier, and wiring this to a discrete alarm annunciator on the control desk. It is assumed that this annunciator is very obviously positioned, and all the other requirements for safety related alarms given in Table 5 are fully satisfied.

Reasonable failure rate figures for the hardware components might be as shown in Table 13:

Component	Failure rate per year
instrument	0.02 (1 in 50 years)
trip amplifier	0.02
annunciator (including bulbs)	0.05 (1 in 20 years)

Table 13 Typical failure rates for hardware components

Adding these together gives a total hardware failure rate of 0.09 per year. It is assumed that one third of these are failures to danger, i.e. a rate of 0.03 per year.

If the system is tested twice per year, then the interval between tests is 0.5 year. This implies that the probability of the alarm being in a dangerous failed state at any time, i.e. its Average Probability of Failure on Demand is:

$$\begin{aligned} PFD_{avg}(\text{hardware}) &= 0.5 \text{ x (fail to danger rate) x (test interval)} \\ &= 0.5 \text{ x } 0.03 \text{ x } 0.5 \\ &= 0.0075 \end{aligned}$$

Note that if the alarm were never tested, its PFD_{avg} would tend asymptotically to 1.

If it is also assumed that the operator reliability would not normally be better than 0.1 so:

$$PFD_{avg}(\text{operator}) = 0.\ 1$$

This gives the overall risk reduction for the operator and system as:

$$\begin{aligned} PFD_{avg}(\text{overall}) &= PFD_{avg}(\text{hardware}) + PFD_{avg}(\text{operator}) \\ &= 0.1075 \end{aligned}$$

A lower failure rate could be achieved by using a higher integrity hardware arrangement

Appendix 5 Setting of Priority

This Appendix discusses the philosophy of alarm prioritisation, including the prioritisation of safety related alarms. It also gives example methods for setting the priority of alarms based on the severity of the consequences of the operator failing to respond appropriately to the alarm and on the time available for response.

A5.1 Severity of Consequences

As indicated in Section 2.5.1, priority should depend on severity of consequence that will follow from the operator not responding appropriately to the alarm. Suppose, for example, that two alarms occurred together, both of which were providing warning of potential plant damage which the operator could prevent. However, if the operator did not respond, the damage would be expected to cause a financial loss of £100 for one alarm and of £10,000 for the other. It is quite clear that the second alarm is more important than the first so this should be given the higher priority. Thus:

> **The prioritisation of an alarm should be based on the expected consequences that the operator can prevent by responding appropriately to it.**

When assessing severity of expected consequence, account should be taken of other systems that will act to mitigate risk if the operator fails to respond. This can be illustrated by comparing the expected consequences in terms of safety of two different alarms, i.e.:

- **pre-trip alarm:** this occurs, say, once per year. If the operator fails to respond to it to correct the disturbance, automatic protection will operate to prevent a dangerous event (e.g. a release of hazardous material, an explosion). If the automatic protection fails then, taking into account the severity of the event and the likelihood of people being near that item of plant, it is estimated that on average there would be 2 injuries following this dangerous event.
- **'final-warning' alarm:** for this to occur there have to be some very unlikely events, such as multiple failures of protection systems or extremely disturbed and abnormal operational conditions[28]. When the final-warning alarm does occur, there is immediate danger and if the operator does nothing the statistical expectation is that there will be 0.2 injuries.

[28] They are many practical examples of final-warning alarms, for example:
an alarm from a monitor installed on a sea discharge. If this comes up, pollution is flowing into the sea and, if nothing is done, an environmental limit could be breached;
a toxic (or inflammable) gas release alarm on a chemical plant. If this has occurred there may real potential for injury;
an alarm on a gas-cooled nuclear reactor indicating the failure of the primary and secondary protection systems. The simultaneous failure of these systems is virtually impossible and is calculated to occur significantly less than once every 10^6 years. However, in the extremely unlikely event that it does happen, the operator can manually initiate the injection of boron beads into the reactor to shut it down and prevent the possibility of release of radioactivity into the atmosphere.

As demonstrated in this example, the severity of consequence for an alarm depends on mitigating systems 'downstream' of the alarm, and not on those that had to fail for the alarm to occur.

The system containing the pre-trip alarm is potentially associated with a more serious end event, since it could result in 2 injuries, whereas the system containing the final warning alarm has an end event of 0.2 injuries. However, the final warning-alarm should be given a higher priority than the pre-trip alarm, because the final warning alarm has occurred and the expected consequence is 0.2 injuries, whereas the expected consequence for the pre-trip alarm is no injury due to the automatic protection equipment operating. If the two alarms were to occur together it would be much more important in terms of safety to respond to the final-warning alarm.

It should be noted that operation of the automatic protection system following a lack of operator response to the pre-trip alarm will have a financial consequence that may result in the pre-trip alarm being assigned a higher priority than would be assigned for the safety consequence alone.

A5.2 Time Available

As indicated in Section 2.5.1, priority may also need to take account of the time available compared with the time required for the corrective action to be performed and to have the desired effect. Thus, if there are two alarms of similar consequence, but one needs fast action to prevent the consequence and the other does not, then there may be benefit in prioritising the first alarm higher than the second so that it gets dealt with first. To take an everyday example, a brake pad wear alarm on a car has more serious potential consequences than an alternator alarm if it is not dealt with, but the latter may need immediate response to prevent the alternator being destroyed.

The value of weighting priority according to time available depends on the typical alarm load. If the load is low there should be adequate time for the operator to deal properly with all alarms. It may then be useful to use priority to emphasise the time critical alarms so that they are dealt with more quickly. If the alarm load is high, then there may not be time to deal with all alarms, and a significant proportion may be neglected for many minutes. At this point, it is important to ensure that the alarms with the greater consequences are not the ones that are ignored. Thus, there should be a lesser weighting on time available.

Because many alarm systems do, in practice, suffer from some alarm overloads it is suggested that the weighting given to time available should be limited. It is suggested that time available should at the most increase the priority of an alarm by one priority band.

A5.3 Priority Distribution of Alarms

The primary purpose of prioritisation is ergonomic, i.e. to make it easier for the operator to identify important alarms when a number occur together. Consequently, to be an effective discriminator, the relative frequency of occurrence of alarms of different priority should reduce with increased priority. This concept is illustrated in Table 15. This suggests approximate figures for target maximum rates of occurrence of alarms of different priorities, as shown in Table 14. It is seen that the occurrence rate reduces by a factor of around 5 for each increase in priority.

Priority band	Target maximum occurrence rate
safety related/critical	very infrequently
high	less than 5 per shift
medium	less than 2 per hour
low	less than 10 per hour

Table 14 Target maximum occurrence rates of alarms of different priorities

Note that the total rate implied by Table 15 is around 12 per hour and is, therefore, consistent with the "Manageable" benchmark rate of Table 25. If the benchmark rate of "Very Likely to be Acceptable" is appropriate, then the target maximum rates of Table 15 will need to be lower.

An important point of principle is that, to be an effective discriminator, the frequency of occurrence of high priority alarms should not be reduced to a minimum. If there are very, very few high priority alarms, say one per month, then the ones that do occur will get particularly special attention. However, if the definition of what is a high priority alarm is reduced so that there are, for instance, 2 or 3 high priority alarms per shift, then these should still be obvious and get the operator's attention. This will be a more effective use of prioritisation.

This argument implies that allocation of priority should be an iterative process. Thus, the allocation of alarms to priority bands should be adjusted based on operating experience until it becomes most effective as a discriminator of relatively important alarms. However, it is to be emphasised that this activity should be carried out only after it has been confirmed that all alarms should be alarms, that all low value nuisance alarms have been eliminated and that the overall alarm rate is reduced to an acceptable level. Additionally, if there are lots of higher priority alarms, a first step should be to question whether the level of automatic control/protection is too low and too much is being demanded of the operator before re-assigning priorities to get optimum discrimination.

Priority band	Alarms configured during system design
critical	about 20 altogether
high	5% of total
medium	15% of total
low	80% of total

Table 15 Priority distribution during system configuration

At the time of design, it will probably be difficult to predict what the alarm occurrence rate will be in practice. As a guide it is suggested that during design alarms should be configured in the approximate ratios shown in Table 15. Performance should then be reviewed during commissioning and early operation, and priorities should be adjusted to achieve a performance similar to that shown in Table 14.

It is strongly emphasised that the numbers in Table 15 and Figure 10 should be taken as approximate indicators of effective discrimination between priorities rather than exact targets. In particular, the priority distribution is expected to be dependent on the type of plant and the speed of response required. On plants with fast dynamic responses, there are likely to be a higher proportion of higher-priority alarms.

A5.4 Prioritisation of Safety Related Alarms

The prioritisation of a safety related alarms will be explored with an example. Suppose the system containing the pre-trip alarm described in Section A5.1 was not considered to be safe enough, i.e. the risks of the control, alarm and protection system failing and the likely injuries had been calculated to be unacceptably high. Suppose that, to make this system safer, it was decided to make the pre-trip alarm safety related. If suitably engineered, this could bring the PFD_{avg} that could be claimed for the alarm down from say 0.1 to 0.01. To justify this claim, as discussed in Section 2.3.3, the pre-trip alarm would have to be implemented outside the control system and given its own highly conspicuous individual annunciator.

Suppose, however, that the final-warning alarm has been implemented as a high priority alarm in the control system. Now, if both the pre-trip and the final-warning alarms occur together, the operator is likely to deal with the pre-trip alarm before the final-warning alarm, as it has been made more conspicuous and appears more important. However, in terms of expected consequences, the final-warning alarm is much more important than the pre-trip alarm and should be dealt with first. Thus, in taking the relatively low consequence pre-trip alarm and engineering it as safety related, it has been given an importance that it does not deserve.

The general conclusion from this is that:

> **Only alarms which would be implemented as highest priority in the control system should be considered as candidates to be safety related alarms.**

There is also a cost-benefit aspect. To make an alarm safety related will almost certainly cost rather more than to leave it as a normal alarm; it will have to be engineered to the same standard as an automatic protection system of similar safety integrity level. Rather than making an alarm safety related, it may be more cost-effective to consider using a normal alarm and installing an additional and, preferably, diverse form of plant protection.

Wherever possible, the plant should be designed so that there is time for the operator to correct the majority of faults before they escalate into emergencies. Consequently, wherever possible, there should be a precursor alarm associated with each safety related alarm.

A5.5 Example Procedures for Setting Priority

This section presents three examples of algorithms for setting priority based on the principles given above. This is intended to illustrate the concepts, and the numbers used are only illustrative. A variety of other algorithms for setting priority has been proposed and may be equally effective in practice.

A5.5.1 Method 1: Summating Consequences

In this example the priority is based on the sum of the safety, environmental and financial consequences. Theoretically this is a logical and rational approach.

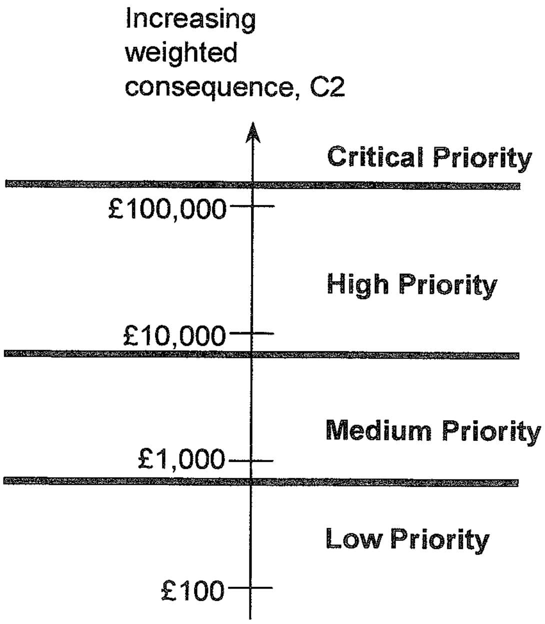

Weighted total consequence, C2	Priority
C2 < 900	Low
900 < C2 < 6,000	Medium
6,000 < C2 < 150,000	High
C2 > 150,000	Critical

Figure 10 Possible priority bands to achieve required priority distribution using summation

The proposed steps in setting priority are:

1. Estimate the safety, environmental and financial consequences of missing the alarm. This will require some estimation of the size of risks and of the likely failure rates of equipment. In doing this it may be acceptable initially to use rough approximations, and refine them if there is a possibility that the alarm may be high or critical priority;
2. Convert the consequences into common units. In this example the conversion factors assumed are:
   CS = 10^6 x (safety consequence in terms of risk of injury)
   CE = 10^6 x (environmental consequence in terms of risk of environmental incident)
   CF = 1 x (financial consequence in pounds);
3. Add the consequences together, i.e.:
   C1 = CS + CE + CF;
4. Assess whether the alarm is 'time critical'. In this example an alarm will be categorised as time critical if it is likely to be too late if the operator fails to make the appropriate response to the alarm within 3 min;

5. Increase the weighting on time critical alarms, in this example by a factor of three, i.e.:

   ```
   IF (alarm is time critical) THEN
   C2 := 3 * C1
   ELSE
   C2 := C1
   ```

   It will be seen after the next stage that this means that time critical alarms have their priority increased by about half a priority band (though the amount varies depending what band they are in);
6. On the basis of the time available weighted total consequence, C2, rank all the alarms in the system in order. Adjust each of the priority bands up or down to achieve the priority distribution given in Table 14. This might end up with prioritisation bands something like those shown in Figure 10. What this could mean for an alarm with a purely economic consequence is shown in Figure 6 Section 2.5.1.

Consequence	Priority	Value when time available more than 3 min	Value when time available less than 3 min
Safety	Low	Risk of injury < 9×10^{-4}	Risk of injury < 3×10^{-4}
	Medium	Risk of injury > 9×10^{-4}	Risk of injury > 3×10^{-4}
	High	Risk of injury > 6×10^{-3}	Risk of injury > 2×10^{-3}
	Critical	Risk of injury > 1.5×10^{-1}	Risk of injury > 5×10^{-2}
Environment	Low	Risk of incident < 9×10^{-4}	Risk of incident < 3×10^{-4}
	Medium	Risk of incident > 9×10^{-4}	Risk of incident > 3×10^{-4}
	High	Risk of incident > 6×10^{-3}	Risk of incident > 2×10^{-3}
	Critical	Risk of incident > 1.5×10^{-1}	Risk of incident > 5×10^{-2}
Financial	Low	Expected loss < £900	Expected loss < £300
	Medium	Expected loss > £900	Expected loss > £300
	High	Expected loss > £6,000	Expected loss > £2,000
	Critical	Expected loss > £150,000	Expected loss > £50,000

Table 16 Priority break points for alarms with just safety, or just environmental, or just financial consequences

Table 16 is expanded out in Figure 11 to show what the prioritisation break points would be for an alarm with only safety or environmental or financial consequences. Note that the actual figures will vary from plant to plant. A large economic loss on one plant may be quite small on another, etc.

A5.5.2 Method 2: Taking Maximum Consequence

Whilst logical and rational, the summation of consequences approach is based on the assumption that consequences can be numerically estimated with some precision. This may be difficult. It may be quicker and more practical to use some heuristic estimates of consequences such as 'small', 'medium', 'large' and 'very large'. These may then be combined by taking the maximum priority given by considering safety, environmental and financial consequences alone.

The procedure is illustrated in Table 17. First the safety consequences are considered, and based on some heuristic rules, the alarm is given a safety priority, P_S of low, medium, high or critical priority.

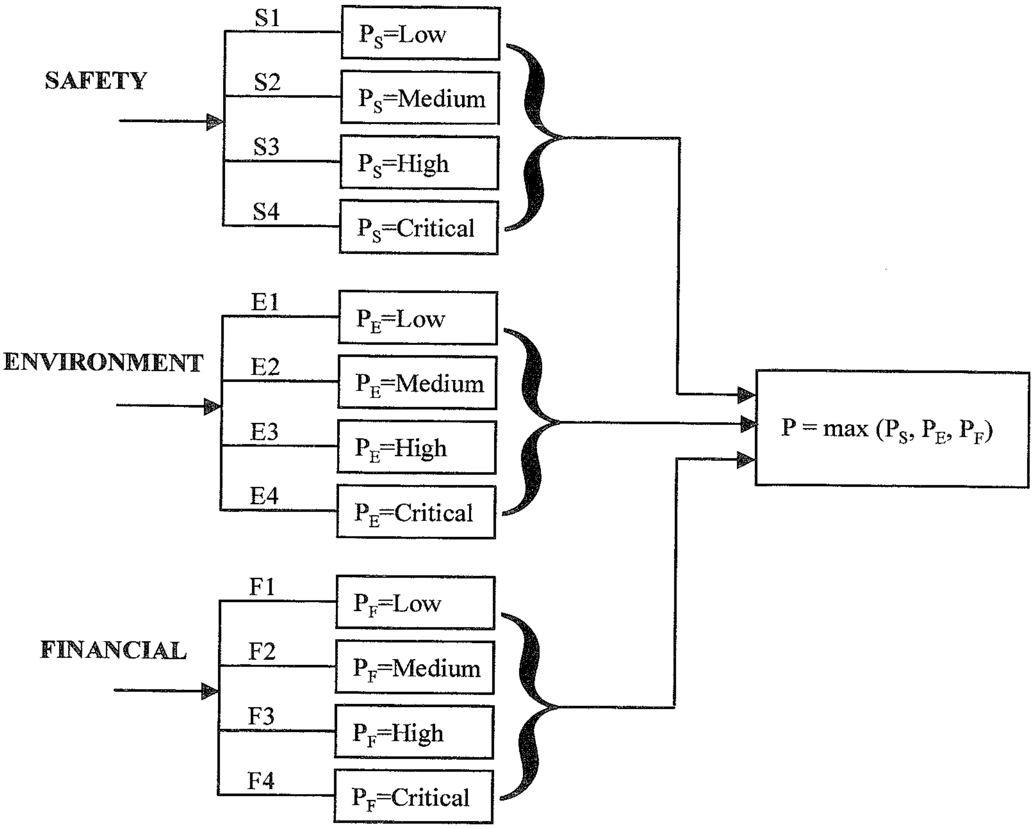

Figure 11 Prioritisation using maximum of individual priorities

A typical set of heuristic rules for classifying the safety consequence are shown in the top table in Table 17. In this table the time available assessment is ignored and the safety consequence is simply estimated as 'small', 'medium', 'large' and 'very large'. Approximate numeric equivalents for these terms are also shown in the table (using different figures from those given in Method 1). A similar process is followed to assign an environmental priority, P_E, and a financial priority, P_F, and the tables relating to these are also shown in Table 17. Then the actual priority given to the alarm is the highest of the safety, environmental and financial priority.

Expected Safety Consequences

Factor	Size	Heuristic measure	Numeric measure of risk of injury
S1	Small	Negligible risk of failure to respond to alarm resulting in a situation likely to cause injury	Below 10^{-3} (i.e. less than 1 in 1000 chance of injury)
S2	Medium	Remote possibility of injury	10^{-3} - 10^{-2}
S3	Large	Potentially dangerous situation with some possibility of putting people at risk of injury	10^{-2} - 10^{-1}
S4	Very large	Dangerous situation with real potential for injury/death	Above 10^{-1} (i.e. greater than 1 in 10 chance of injury)

Expected Environmental Consequences			
Factor	**Size**	**Heuristic measure**	**Numeric measure of risk of breach of limits**
E1	Small	Negligible risk of failure to respond to alarm resulting in any breach of environmental limits	Below 10^{-3} (i.e. less than 1 in 1000 chance of incident)
E2	Medium	Remote possibility of breach of environmental limits	10^{-3} - 10^{-2}
E3	Large	Situation with some possibility of breach of environmental limits	10^{-2} - 10^{-1}
E4	Very large	Situation with real potential for serious breach of environmental limits	Above 10^{-1} (i.e. greater than 1 in 10 chance of incident)

Expected Financial Consequences			
Factor	**Size**	**Heuristic measure**	**Numeric measure of risk of financial loss**
F1	Small	No immediate likelihood of plant damage but the possibility of this has increased. Minor loss in productivity or efficiency	Below £1000
F2	Medium	Some chance of minor plant damage. Significant reduction in plant output, e.g. 10% reduction for 1 hour	£1000 - £10,000
F3	Large	High chance of minor plant damage or low chance of serious plant damage. Significant loss of production, e.g. loss of an hour of total plant output	£10,000 - £100,000
F4	Very large	High chance of serious plant damage. Serious and prolonged output loss, e.g. loss of one day of complete plant output	£100,000+

Table 17 Heuristic rules for allocation of safety, environmental and financial priority

The use of heuristic rules can be a quick way of assigning priority to alarms. However, it has two disadvantages:

- care has to be taken that the heuristic rules for categorising consequences as 'small', 'medium', 'large' and 'very large' are consistent across the whole plant. In practice this may be achieved only by assigning some approximate numeric values to the terms;
- if only heuristic values recorded then it may be hard to later decide which alarms should be upgraded or down-graded if the priority distribution is found to give poor discrimination.

A5.5.3 Method 3: General Alarm Assessment

The following flowcharts (Figure 12, Figure 13, Figure 14 and Figure 15) have been developed as an example of how a team may review the priority assessment of process related alarms. The criteria in the relevant boxes can be customised for a particular site.

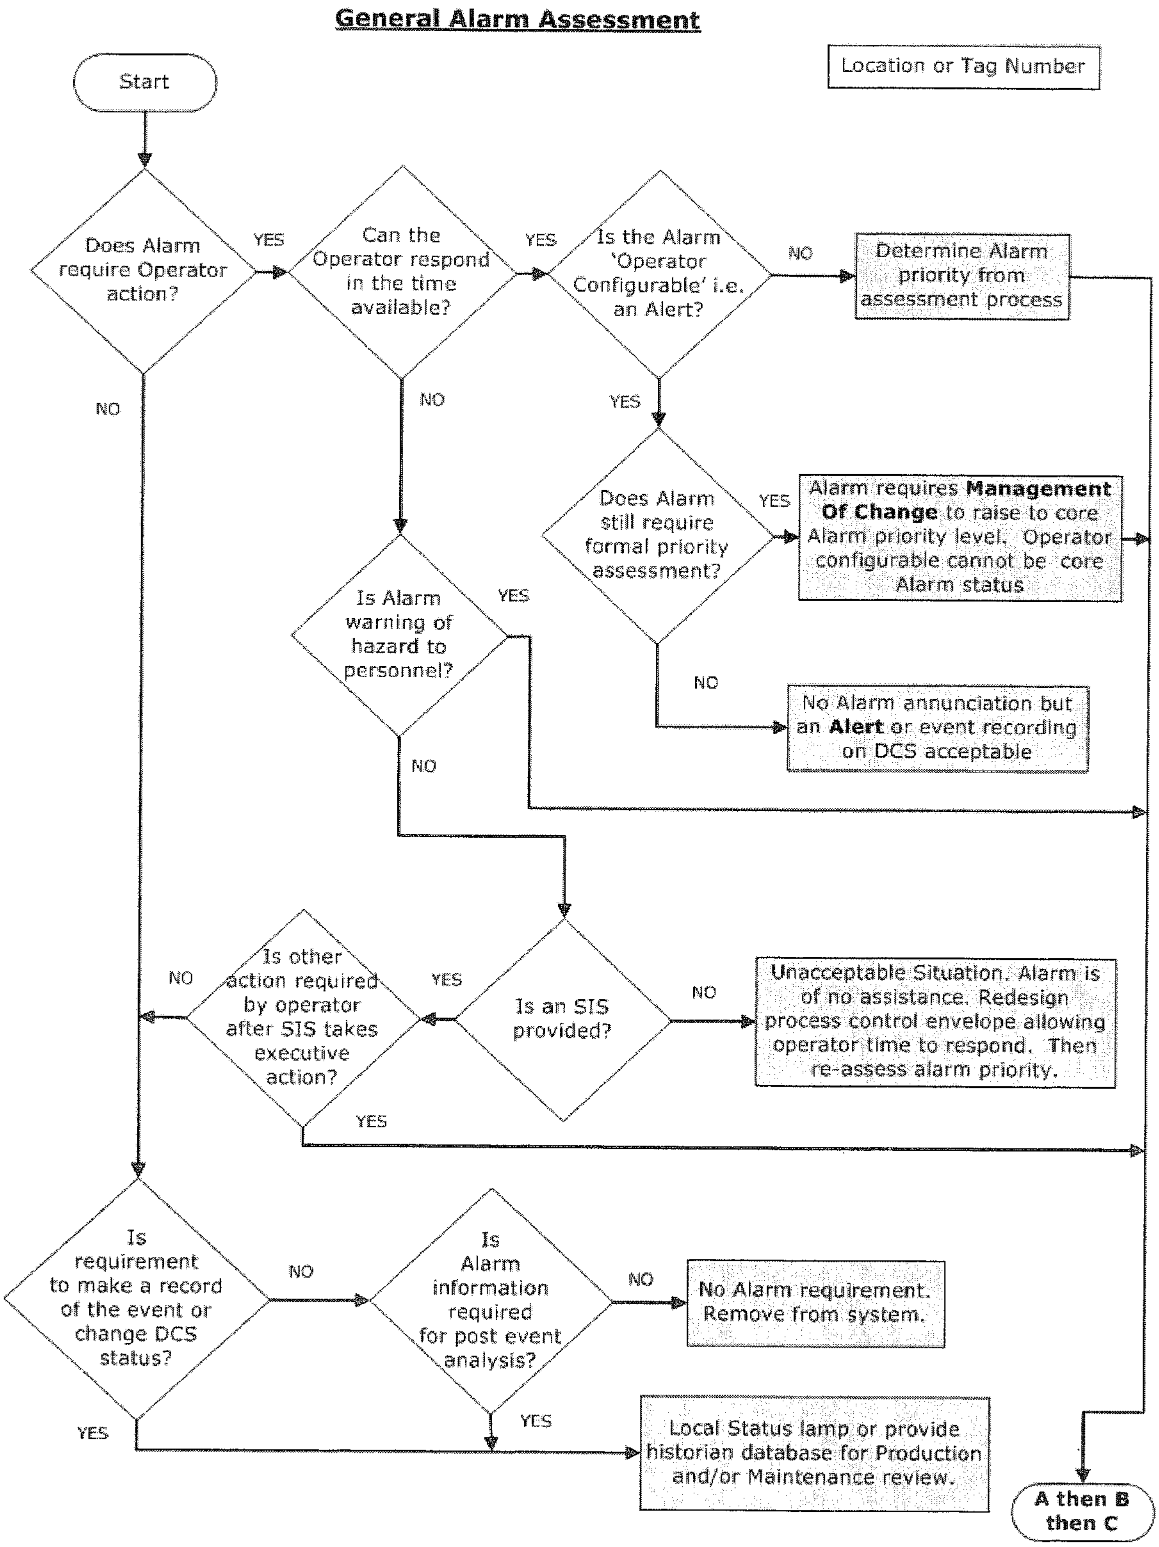
General Alarm Assessment
Location or Tag Number
Start
Does Alarm require Operator action?
YES
NO
Can the Operator respond in the time available?
YES
NO
Is the Alarm 'Operator Configurable' i.e. an Alert?
NO
YES
Determine Alarm priority from assessment process
Does Alarm still require formal priority assessment?
YES
NO
Alarm requires Management Of Change to raise to core Alarm priority level. Operator configurable cannot be core Alarm status
No Alarm annunciation but an Alert or event recording on DCS acceptable
Is Alarm warning of hazard to personnel?
YES
NO
Is an SIS provided?
YES
NO
Unacceptable Situation. Alarm is of no assistance. Redesign process control envelope allowing operator time to respond. Then re-assess alarm priority.
Is other action required by operator after SIS takes executive action?
NO
YES
Is requirement to make a record of the event or change DCS status?
NO
YES
Is Alarm information required for post event analysis?
NO
YES
No Alarm requirement. Remove from system.
Local Status lamp or provide historian database for Production and/or Maintenance review.
A then B then C

Figure 12 General Alarm Assessment

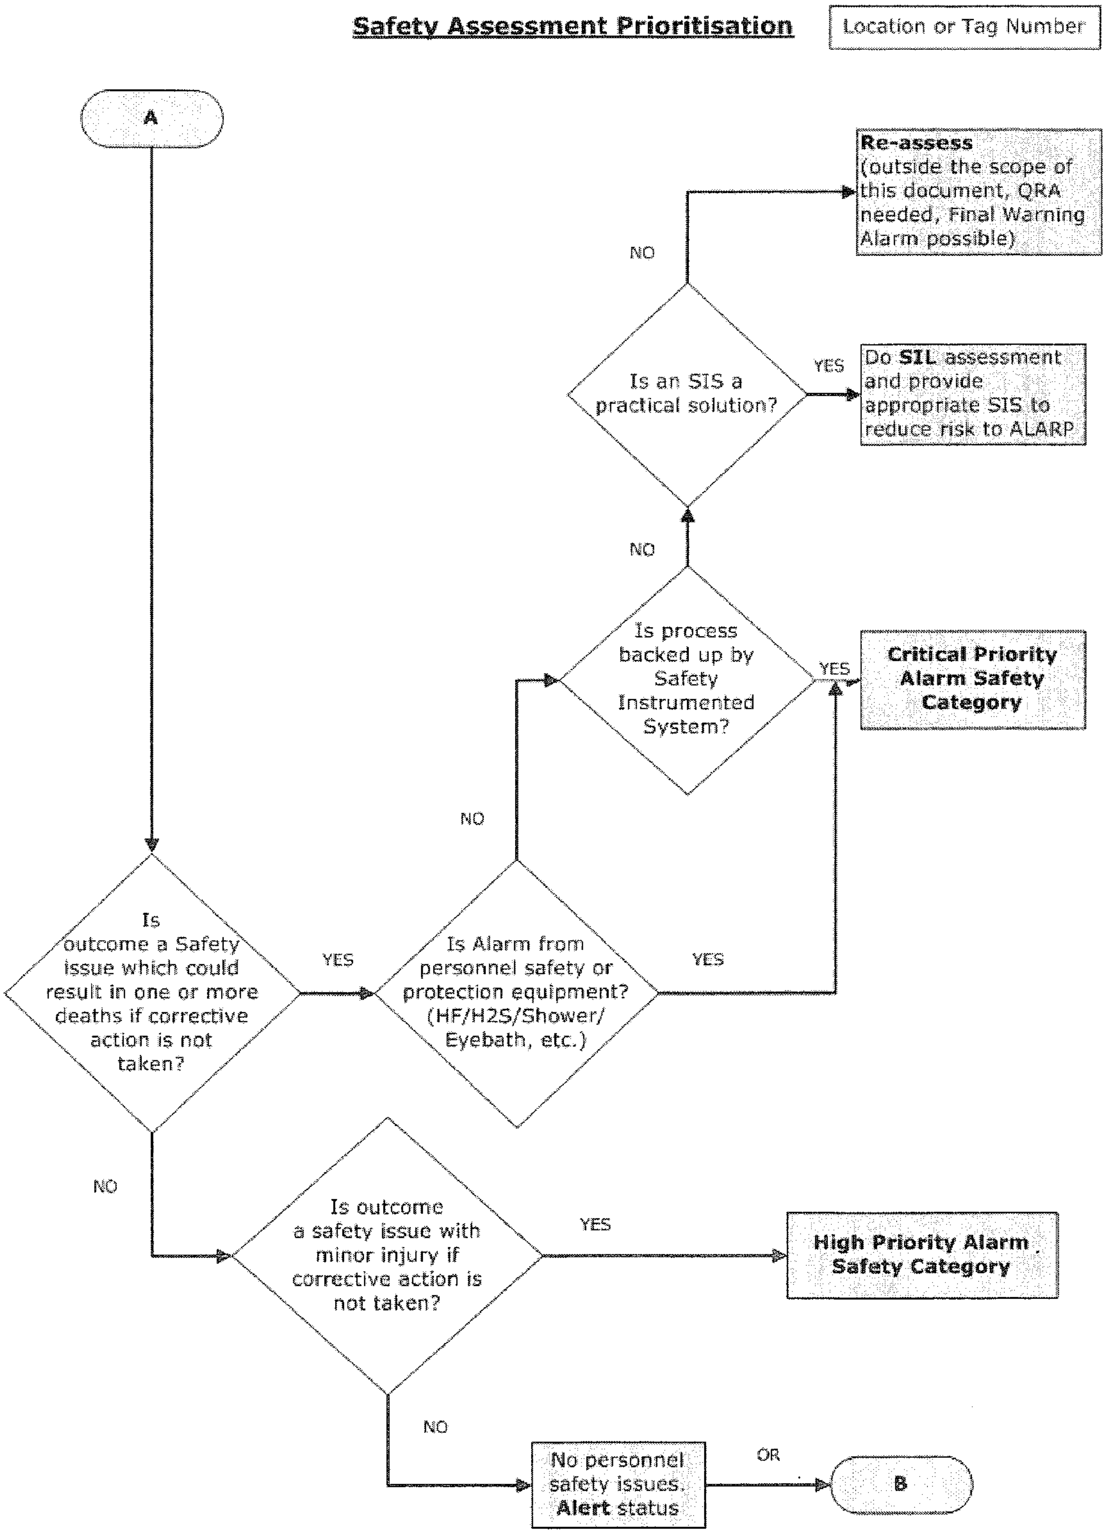
Safety Assessment Prioritisation
Location or Tag Number
A
Is outcome a Safety issue which could result in one or more deaths if corrective action is not taken?
YES
NO
Is Alarm from personnel safety or protection equipment? (HF/H2S/Shower/Eyebath, etc.)
YES
NO
Is process backed up by Safety Instrumented System?
YES
NO
Critical Priority Alarm Safety Category
Is an SIS a practical solution?
YES
NO
Do SIL assessment and provide appropriate SIS to reduce risk to ALARP
Re-assess (outside the scope of this document, QRA needed, Final Warning Alarm possible)
Is outcome a safety issue with minor injury if corrective action is not taken?
YES
NO
High Priority Alarm Safety Category
No personnel safety issues. Alert status
OR
B

Figure 13 Safety Assessment Prioritisation

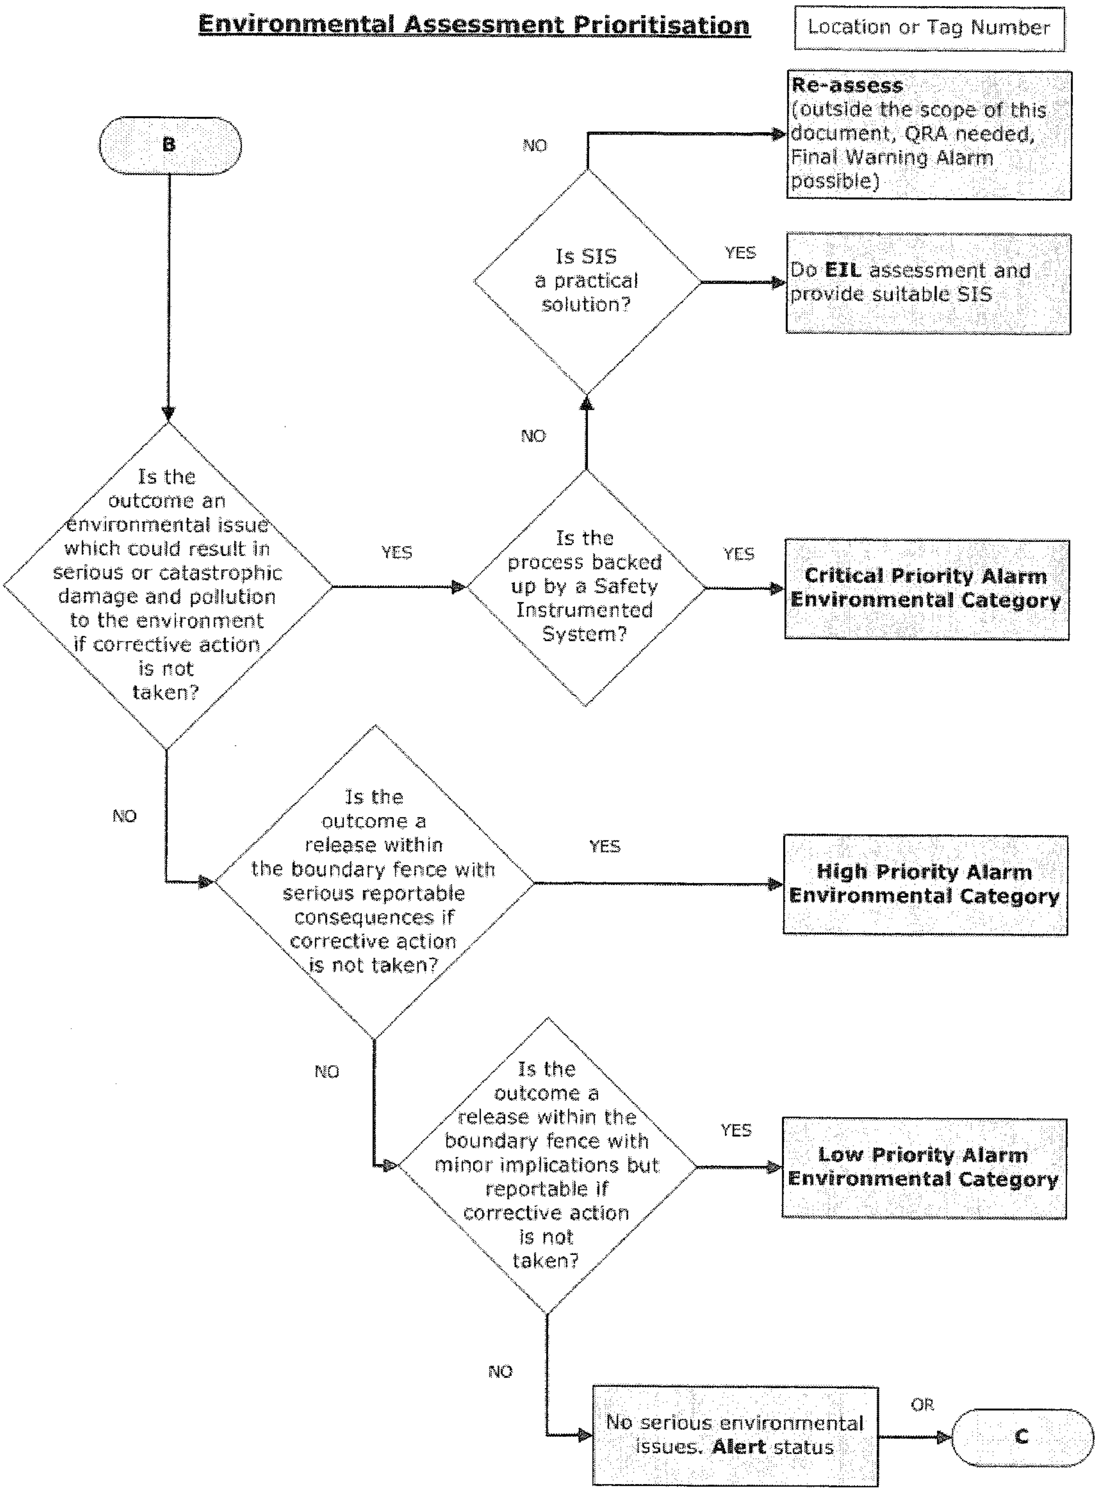
Environmental Assessment Prioritisation
Location or Tag Number
B
Is the outcome an environmental issue which could result in serious or catastrophic damage and pollution to the environment if corrective action is not taken?
YES
NO
Is the process backed up by a Safety Instrumented System?
YES
NO
Critical Priority Alarm Environmental Category
Is SIS a practical solution?
NO
YES
Re-assess (outside the scope of this document, QRA needed, Final Warning Alarm possible)
Do EIL assessment and provide suitable SIS
Is the outcome a release within the boundary fence with serious reportable consequences if corrective action is not taken?
YES
NO
High Priority Alarm Environmental Category
Is the outcome a release within the boundary fence with minor implications but reportable if corrective action is not taken?
YES
NO
Low Priority Alarm Environmental Category
No serious environmental issues. Alert status
OR
C

Figure 14 Environmental Assessment Prioritisation

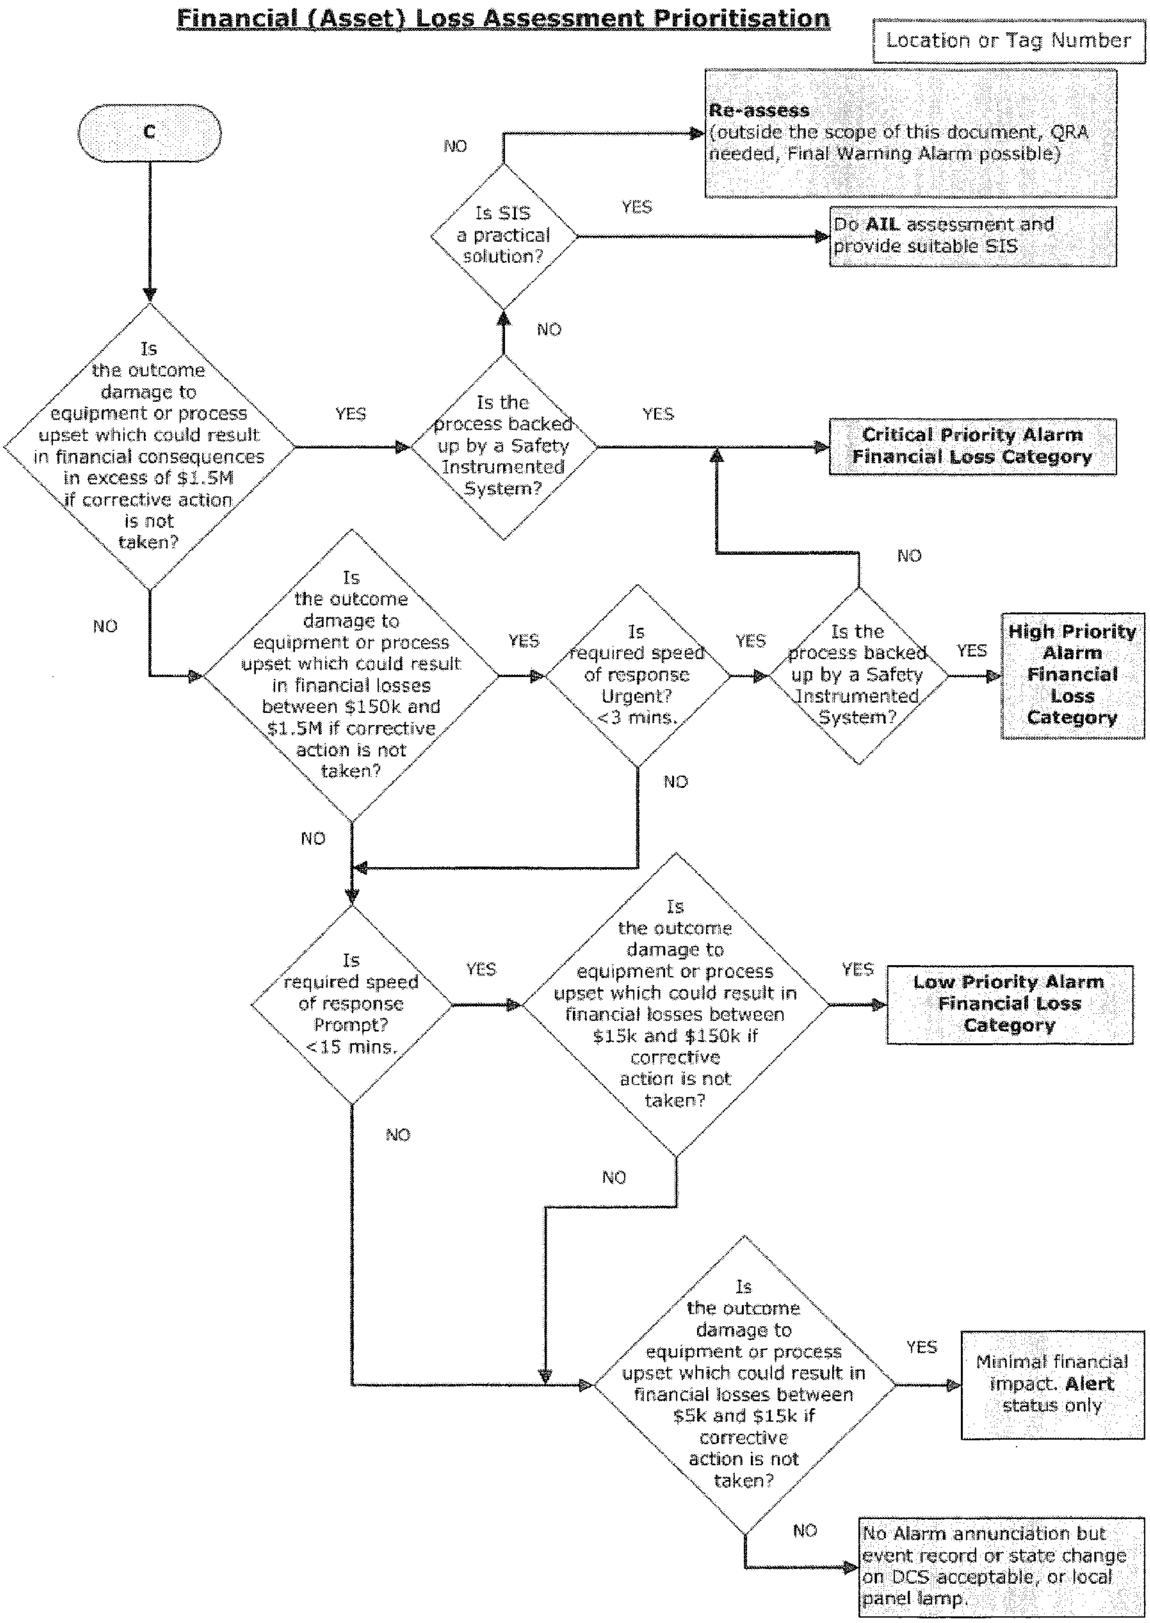
Financial (Asset) Loss Assessment Prioritisation
Location or Tag Number
C
Is the outcome damage to equipment or process upset which could result in financial consequences in excess of $1.5M if corrective action is not taken?
YES
NO
Is the process backed up by a Safety Instrumented System?
YES
NO
Critical Priority Alarm Financial Loss Category
Is SIS a practical solution?
NO
YES
Re-assess (outside the scope of this document, QRA needed, Final Warning Alarm possible)
Do AIL assessment and provide suitable SIS
Is the outcome damage to equipment or process upset which could result in financial losses between $150k and $1.5M if corrective action is not taken?
YES
NO
Is required speed of response Urgent? <3 mins.
YES
NO
Is the process backed up by a Safety Instrumented System?
YES
NO
High Priority Alarm Financial Loss Category
Is required speed of response Prompt? <15 mins.
YES
NO
Is the outcome damage to equipment or process upset which could result in financial losses between $15k and $150k if corrective action is not taken?
YES
NO
Low Priority Alarm Financial Loss Category
Is the outcome damage to equipment or process upset which could result in financial losses between $5k and $15k if corrective action is not taken?
YES
NO
Minimal financial impact. Alert status only
No Alarm annunciation but event record or state change on DCS acceptable, or local panel lamp.

Figure 15 Financial (Asset) Loss Assessment Prioritisation

A5.6 Record keeping

When prioritisation of alarms is carried out, a record should be kept of how the calibration was done. It is recommended that a form should be developed for this purpose. Information to record might include:

- reference to calibration rules used;
- alarm identifier/description;
- description of the defined response to the alarm;
- verbal descriptions of expected safety, environmental and economic consequences if the appropriate response is not made to the alarm;
- measures of expected safety, environmental and economic consequences;
- measure of time criticality;
- priority allocated;
- name of person doing the allocation;
- date when allocation done;
- record of any independent review of the allocation;
- modification approval/implementation record.

This information is a subset of the information listed in Appendix 2 which should be recorded when each alarm is designed.

Table 18 could be a typical method of recording information and data.

Unit XXX Alarms Review

Tag Number:	Device location number or other recognised reference
Description:	Description as appears on the window or faceplate for the alarm.
PID Number:	
Keyword:	For search purpose or other key descriptor.
Range:	Lower limit - Upper limit Engineering measuring units etc.
Category:	SAFETY, ENVIRONMENTAL or FINANCIAL
Agreed Settings:	

	Low		High	
	Setting	**Priority**	Setting	**Priority**
PV Alarm:		CRITICAL, HIGH MEDIUM or LOW		CRITICAL, HIGH MEDIUM or LOW
Dev. Alarm:		CRITICAL, HIGH MEDIUM or LOW		CRITICAL, HIGH MEDIUM or LOW
Alarm State:	False or True	CRITICAL, HIGH MEDIUM or LOW		
Off Norm Prty				
State 1:				
State 2:				

Condition being monitored: Actual process description not the alarm faceplate description.
Basis of setting: Engineering calculation, hold up time in vessel, physical limit of metallurgy etc.
Consequence of failure to respond: Consequence should the alarm be overlooked or missed in alarm flood situation. Do not consider any ESD back up. **Time to Respond:** This is the time available for the production operator to recognise the alarm, determine what action needs to be taken and then to bring the function back to below the alarm state.
Operator action to be taken: Action or activities necessary to bring the process into a controlled condition after the alarm has been annunciated. Adjusting set point is not adequate. There will be other observations etc. that should all be annotated.
Applicable IP/ MOC?: Any management of change documantation required to make adjustements to settings or physical changes.

	Production	**Process Engineer**	**I/E Maintenance**
Name:			
Signature:			
Date:			
Revision:			

Table 18 Example of a typical method of recording information and data

Appendix 6 Types of Alarms

This Appendix lists the mechanisms for detecting alarms and gives examples of how they might typically be used.

Some commonly used mechanisms for alarm detection are:

- **Absolute alarms:** generated by comparison of an analogue signal against a defined alarm setting. Typically used to warn of the parameter approaching an absolute limit such as a trip or safety-valve setting. These are simple to engineer, but tend to cause difficulty in abnormal operational situations. They may need conditioning to take account of plant operating state. The setting of absolute alarms is discussed further in Section 2.4.2.

- **Bit-pattern alarms:** alarms generated when a pattern of digital signals matches a predetermined pattern.

- **Calculated alarms:** generated in software from applications such as modulating controls, batch and sequential controls, efficiency calculations, etc. and often based on logical combinations of several signals. Ideally it should be possible for the application software to dynamically alter alarm settings, priorities, deadbands, etc.

- **Control and instrumentation system alarms:** generated from faults within the control and instrumentation system hardware or software. These can be useful, but care should be taken that they are relevant to the operator and easily understood by the operator. These are commonly misused. Note that it may be desirable to be able to separate the display of these system alarms from the process alarms.

- **Deviation alarms:** generated if the difference between two analogue signals exceeds a certain size. Typically used to compare a controlled variable against the controller set point and to warn that the control system is failing to operate effectively. Care should be taken that the alarm is not generated during large plant disturbances when the control system is doing the best it can, but the control error is large. To avoid this, they may need to incorporate time windows or time delays. The alarms may also need to take account of whether a control loop is selected to automatic or manual operation.

- **Discrepancy alarms:** generated by comparing an expected plant state against an actual plant state.

  Discrepancy alarms are often applied to actuators. This requires some model of the expected movement of the actuator in response to movement commands from the operator or from automatic controls or sequences. The model may be fairly crude, e.g. the actuator should always show the correct state (i.e. open or closed) within 10 seconds of a command being sent to it, or may be more sophisticated with actuator slew rate parameters and deadbands in it. A good discrepancy alarm check will identify faults such as stuck actuators or actuator runaway.

  In practice, actuators tend to degrade or not perform as modelled and spurious discrepancy alarms can be generated. For example, an actuator may get sluggish and take 12 seconds to close rather than the assumed 10 seconds. Spurious discrepancy alarms can be a nuisance, particularly if there is a large upset affecting the whole plant and many actuators simultaneously generate spurious alarms. Consequently, discrepancy alarms should be designed to be as robust as possible. It is also desirable to have simple procedures for widening tolerance bands in discrepancy alarms on actuators that are slightly degraded and are awaiting maintenance.

- **Rate-of-change alarms:** generated by the rate of change of an analogue signal exceeding a defined setting. These can often provide an early indication of an upset. However, they should be used with great care since noise on the signal tends to be amplified and can result in spurious alarms. Thus, they may need to incorporate time windows or time delays.

- **Recipe-driven alarms:** alarms that are turned on or off in different plant states, typically by some type of sequence controller. An example might be an alarm for a batch plant that is made active in particular operational phases.

Some more advanced alarm detection mechanisms are:

- **Adaptive alarms:** generated using the 'rate-of-change' or 'deviation' principle in combination with absolute thresholds. Where the parameter is well clear of the absolute threshold, the allowable rate-of-change or deviation is relaxed; as the parameter approaches the absolute threshold, the limits on rate-of-change or deviation are tightened.

- **Adjustable alarms:** absolute alarms in which the alarm settings are adjusted to suit operating conditions. This can be done automatically according to some predefined logic. For example, vibration alarm settings on a rotating machine might be made a function of speed, or boiler steam temperatures alarms might be made a function of operating load.

- **Operator-set alarms:** alarms in which the settings may be manually adjusted by the operator to suit their needs. Also included are temporary alarms on variables and settings chosen by the operator. Mechanisms need to be provided to ensure that operators do not use alarms that they set up, but which have since been changed by other operators without their knowledge.

- **Re-triggering alarms:** alarms which are automatically re-annunciated to the operator in certain conditions. This might be when the alarm has been standing for more than a predefined time setting, or when the alarmed variable had moved significantly above its setting. Thus, a high pressure alarm might be set to be raised at 100 bar and to re-annunciate for every 5 bar increase above this (i.e. at 105 bar, 110 bar, etc.).

- **Statistical alarms:** generated by some statistical process to filter out significant changes from process noise.

For the purposes of clarity another term is explained:

- **First-up alarms:** these are not strictly a type of alarm, but are a facility for examining the order of occurrence of alarms. Typically a first-up system would be used to identify the cause of plant item trips. For example, a pump might be tripped by a low flow switch or by loss of power to the motor. If the pump trips, the motor will stop and the flow will drop very quickly, so both a 'motor stopped' and a 'low flow' alarm will be generated. The time difference between them will be fractions of a second, and this will be difficult to resolve with many commercial alarm systems. Fast scanning logic allows the first alarm after the trip to be identified, and hence the cause of the trip to be determined. Some commercial alarm annunciator systems have facilities for displaying first-up alarms built into their hardware functionality. For example, a number of alarms may be put into a first-up group. The first alarm raised in the group might be indicated by a flashing light, the subsequent alarms by a steady light.

Appendix 7 Alerts

Operators today are constantly bombarded with a multitude of process-related tasks; often they must simultaneously:

- monitor multiple process units;
- diagnose and mitigate abnormal situations;
- recall and perform ongoing tasks.

In addition, an operator must integrate alarm system interrupts, radio requests and process control requirements with normal unit monitoring. Such heavy demands can strain, and may even overload, an operator's attention. Process incidents can further tax even the most experienced operators, requiring them to put ongoing tasks on hold until the event is mitigated.

In an effort to deal with this complexity and to manage multiple tasks, operators often use their alarm screens to notify them of pre-alarm conditions and to remind them of process events that need attention. Using the alarm displays in this manner may overload and clutter the system. Instead of resolving a difficult situation, this method often has the opposite effect – creating further difficulties. This situation has developed because operators have not been given tools that adequately keep pace with increasing process unit complexity.

A direct result of the efforts of the Abnormal Situation Management Consortium, the "Alert" concept was identified to deliver the automatic notification capabilities that are so seriously needed.

The concept has major similarities with that of an alarm – and an important difference. The *similarities* are:

- process conditions requiring attention are defined;
- alerts are built using a configuration tool according to the identified requirements (e.g. a condition based on a tag value or difference in tag values; a timer to "alert" when the time expires; a combination of multiple conditions, etc.);
- alert conditions are continually checked – and "annunciated" on a display (but ***not*** the standard alarm display) or some other notification device (e.g. a pager, text message, email, etc.);
- the operator responds accordingly.

The important *difference* (between an alarm and an alert) is:

- alerts are always of lower priority than alarms and ignoring an alert does not have serious consequences.

Alerts, therefore, help the operator to run the plant more efficiently – they should never be safety-related (in the IEC 61508 sense) (29) or related to some other condition that has a serious impact on the plant or its surroundings – since such conditions are properly dealt with in the alarm or other protective system. During upsets, many alerts might be generated (so "alert floods" can occur) – but, if alerts have been properly used, the operator knows that they can ignore them until all alarms have been dealt with. Notice also, that since alerts would often replace some of the low priority alarms that would otherwise have been used, the size of alarm floods are reduced, and can then be much more manageable.

Appendix 8 Logical Processing of Alarms

This Appendix describes a number of different techniques that can be used for processing signals from alarm sensors to generate more meaningful alarms for display to the operator. They include:

- *grouping of alarms;*
- *suppression of redundant alarms;*
- *eclipsing of several alarms on the same variable;*
- *suppression of alarms from out of service plant;*
- *suppression of alarms according to plant operating mode;*
- *suppression of alarms following major events;*
- *intelligent fault detection;*
- *automatic alarm load shedding;*
- *handling alarms from equipment under test.*

A8.1 General Comments

The logical processing methods discussed in this Appendix can be extremely useful in improving the operational value of alarms. However, it should be stressed that:

- the operator should be kept informed when logical processing is removing alarms from the display, e.g. by automatic suppression;
- any logical processing should be done in a manner which minimises the possibility of error. For example, it may be preferable to implement Boolean logic using configuration tools that can display the logic in a clear and unambiguous way to non-technical users and allow it to be easily checked, rather than using general programming languages;
- if the alarm is safety related, the implementation of the logical processing should comply with the requirements of IEC 61508 (29).

A8.2 Grouping

A single grouped alarm may be used to display a number of different initiating events from a plant system. For example, there might be one alarm 'Water treatment plant fault' that is annunciated by a range of different faults. This technique is valid only where all the constituent alarms in a group are of the same priority and require the same initial response from the recipient. Typically this response might be to send an operator to investigate the situation local to plant, which then might require an individual response based upon the additional information obtained locally.

It is good practice to design group alarms so they re-flash and need re-acceptance if a second initiating event comes up whilst the group alarm is already standing. Because of this, the use of group alarms does not generally reduce the total number of times that alarms are annunciated to the operator.

One difficulty with group alarms is that the operator will often need to obtain more information about the particular cause of the alarm before deciding what to do about it. If the grouping is implemented out on the plant (as is often done) this will require investigation by the operator or their assistant at the plant, which can be time consuming. Consequently, in many modernisation exercises in which old annunciator-based alarms are being replaced by DCS/SCADA systems, group alarms are being eliminated and the signals that were used to generate the group alarm are being wired into the alarm processor as individual alarms.

The conclusion of the above is that grouping of alarms is not often an effective technique for increasing usability. However, it can be effective where there is redundant instrumentation which generates multiple alarms from the same common cause or where a group of alarms have a similar operator response. Group alarms can also be used to flash format selection buttons or highlight items on overview schematics (see Section 3.3.4).

A8.3 Suppression

The following suppression techniques result in alarm signals from equipment being assessed as not appropriate for display to the operator. These techniques can be very valuable, but should be applied with care. Safety problems have arisen from inappropriate use of suppression.

Some users design the suppression to reduce the alarm priority rather than eliminating the alarm altogether. It may either be set to the lowest normally displayed priority, or to a priority that the operator has to select for display.

A8.3.1 Redundancy Logic

Often multiple measurements are made of the same process variable (e.g. so that they can be used for majority voting in safety systems). If alarms are generated from these individual measurements then there will be multiple alarms all indicating the same thing. Suppression logic can ensure that only a single alarm is displayed to the operator.

Note that sometimes the same alarms are generated by different systems (e.g. a control system and a protection system). It is advantageous to suppress duplicates if this can be done with no loss of integrity.

A8.3.2 Eclipsing

Sometimes there will be several alarms generated from a single process variable. For example, there might be a high alarm set at one setting and a high-high alarm set at a slightly higher setting. Logic can be used to suppress the alarms of lower operational significance when the more significant alarms are raised. For example, a high alarm might be suppressed when a high-high alarm is raised. Note that eclipsing may reduce alarms on display, but may not necessarily reduce the number of alarms the operator has to accept.

A8.3.3 Out of Service Plant

Some alarms are of operational significance when a plant item is running, but not when it is out of service. For example, a low outlet flow alarm from a pump will not be relevant when the pump is not running. Logic can be used to suppress such alarms (or reduce their priority to a band which is not normally displayed to the operator). It can be useful if the system still allows the operator to view such suppressed alarms if necessary, and also to override the suppression.

The computation of the plant running logic flags should be done carefully and must take account of all the various ways of operating the plant. This is particularly true when computing flags representing the running status of very large plant systems or of the complete plant. For example, when the plant is not producing output, there may be still circumstances when particular plant sub-systems are run and need their alarms to be active.

Plant running flags should also take account of the detail of the start up sequence for the plant item. For example, when starting a large machine, the lubricating oil systems come into service before the machine starts to rotate, and their alarms need to be made active at this stage. Different logic may be required in shut down. It takes careful thought and debate between engineering and operations staff to get such logic right.

Despite these difficulties, alarms from out of service plant are often quite a severe cause of nuisance. Serious consideration should be given to designing logic to suppress them.

An example of 'out of service plant' mode suppression of alarms at BP Oil Grangemouth Refinery is given in (10).

A8.3.4 Operating Mode

Certain alarms are only relevant in particular plant operating modes. For example, when a plant is starting up it may be difficult to avoid transiently exceeding the limits on smoke emissions that apply in steady operation. Thus, it may be appropriate to suppress alarms or change alarm settings according to the plant operating mode.

The number of modes chosen will depend on the type of plant, but might typically include:

- shut down;
- starting up;
- steady operation ;
- different recipes or process stages in a batch process;
- plant maintenance.

Often the transition between these modes is not clearly defined and varies from occasion to occasion. Also different alarms may become relevant at different stages in the transition and also depending on the direction of the transition. Thus the design of the suppression logic does require considerable care. The mode may be calculated from plant measurements and automatically changed, or may be selected manually by the operator. A good option is for the alarm system to make a calculation of mode and recommend when mode switches should be made, and for the operator to then confirm these.

A8.3.5 Major Event

Typically the biggest alarm load on the operator is after a major plant upset. Such disturbances are often particularly stressful for the operator, and can also be considered as relatively hazardous periods of operation, as many plant items are expected to change state and, thus, there are more things to go wrong. It is particularly important, therefore, to try to improve alarm performance in this period. This has to take some priority in an alarm improvement programme (see Section 5). It is the most important thing to do once the basic work on alarm settings, messages, repeating alarms, etc. has been done. It may involve other techniques, e.g. operating mode suppression.

Many of the alarms occurring after a major upset will relate to events that are expected to happen. For example, if a total plant shut down is initiated then there will be much inter-tripping of plant items, and many parameters will go outside their normal operating ranges. The use of logic to suppress these expected alarms offers significant benefit.

Also, the use of logic to identify missing events is operationally important. For example, in a plant shut down many trips will operate. The operator wants to know only about the trips that do not operate, or the valves that do not shut. This generally represents an inversion of the alarms required in normal operation. Take, for example, a nuclear reactor. In normal operation one wants an alarm if any control rod becomes fully inserted into the reactor. After a trip one wants an alarm on any control rod that is not fully inserted.

A8.4 Intelligent Fault Detection

The majority of industrial alarm systems are configured as 'one input - one output'. As discussed above, logical processing of alarms can be used to increase their relevance to the operator, particularly in fault situations.

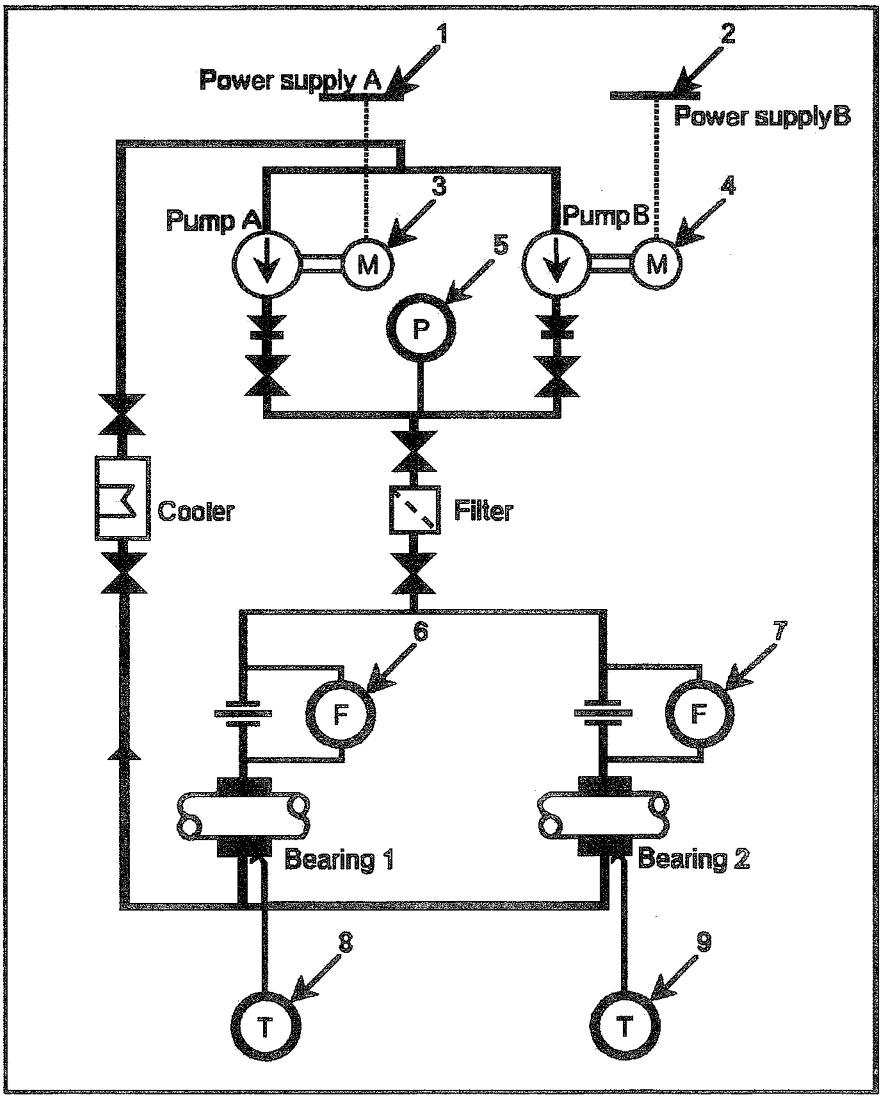

Figure 16 Plant system with redundant equipment

Intelligent fault detection is a term which covers a range of methods for logically processing alarms to reduce the amount of displayed information and increase its relevance. For example, in a complex system there may be several alarms that will be produced following a single fault. Some sophisticated processing systems are able to identify the root cause of a fault from the pattern of resulting alarms. A full literature review is given in (5).

A8.4.1 Pattern Recognition

Early attempts to address the problem of excessive numbers of alarms using computer-based systems adopted approaches which sought to 'analyse' alarms, either by establishing logical cause-consequence links between sets of alarms, or by attempting to identify standard patterns of alarms which could be identified with particular faults on plant items or complete systems.

The nature of the problem may be illustrated with an example (32) which will be referred to in the following discussions. Figure 16 shows the simplified arrangement of two electrically powered pumps for supplying oil to bearings on a large rotating plant. Normally one pump is in service, the other is in auto-standby. There are 9 alarms on the system.

In practice, an oil supply system would often have more alarms than shown here, e.g. with discrepancy alarms on the valves, differential pressure alarms on the filter, etc. Also many systems are more complex than the one shown here, and the relationships between alarms more complicated.

Table 19 below shows the pattern of alarms that might occur on the example system. A number of such alarm analysis systems were developed in the late 1960s and early 1970s for nuclear power plants, etc.

Fault	Alarms								
	1	2	3	4	5	6	7	8	9
Supply A failed	□	-	□	-	□	-	-	-	-
Pump A failed	-	-	□	-	□	-	-	-	-
Supply B failed	-	□	-	□	□	-	-	-	-
Pump B failed	-	-	-	□	□	-	-	-	-
Filter DP high	-	-	-	-	-	□	□	□	□
Bearing 1 oil flow lost	-	-	-	-	-	□	-	□	-
Bearing 2 oil flow lost	-	-	-	-	-	-	□	-	□

Table 19 Alarm pattern array

Although they are apparently appealing, experience world-wide has shown that such alarm analysis approaches have limitations, i.e.:

- in practice, there are a large number of possible faults on a large plant, so the store of patterns and the matching process can become cumbersome;
- combinations of faults are not easily catered for;
- the number of faults which can be reliably analysed before they have occurred is finite and not large, particularly if a full-scope plant simulator is not available;
- many alarm analysis systems have been shown to be resource-intensive, both to create and maintain, and may not reliably identify root-cause faults.

Consequently the technique is of limited application, but may be useful for restricted applications.

A8.4.2 Neural Networks

Early alarm analysis systems used formal computer-language declarations of relationships or patterns, often requiring significant amounts of computer code. More recently 'neural networks' techniques have received much attention. These programs are loosely modelled after the operation of the human brain. They allow plant fault analysis to be attempted without any formal logic modelling or rule definitions.

A neural network contains a very large number of interconnected 'nodes'. Each node can acquire a scalar value, depending on the signals fed to it via the interconnections. The combination of the individual node values produces an 'output' signal. The neural network is repeatedly presented with data representing the values of individual parameters and states on the plant. At each presentation, the network 'learns' the relationship of input signals to the resulting output. The network is 'trained' by presenting data representing normal operation and plant faults, until a reliable set of programme outputs is obtained, which can then be associated with particular plant faults. The technique has been used in numerous studies, but is still in its infancy in terms of working commercial alarm applications.

A8.4.3 Fuzzy Logic

A further approach to computerised fault detection is to use probabilistic methods based on fuzzy logic. Here the pattern recognition method is extended by assigning a probability of occurrence to each specified condition. The current alarm state is then used to calculate a probability value which is compared with the stored pattern values and a match sought. Again fuzzy logic has yet to find widespread commercial application for alarm handling.

A8.4.4 Knowledge-based Reasoning

The advent of effective symbolic processing and object-oriented programming techniques in the 1980s, together with the development of cost-effective powerful computing platforms on which to run such software, made it possible to approach the question of process fault detection in a more competent way. This made possible the development of so-called 'knowledge-based reasoning', where plant behaviour is modelled in logical and symbolic form. Through suitable interface design, plant knowledge is entered into the programme in near-plain English text form. A typical logic expression might be:

```
IF  the-output-pressure-of-pump-A is FALLING
AND
IF  the-state-of-circuit-breaker-A is CLOSED
THEN CONCLUDE
    pump-A has a-pump-fault
AND SEND MESSAGE TO OPERATOR
    "Pump A - Output Pressure fault".
```

A8.4.5 Model-based Reasoning

A more sophisticated approach to complex plant fault detection is that of model-based reasoning in which a mathematical model of a plant, typically comprising sufficient differential equations to describe the behaviour of all the relevant 'healthy' plant parameters, is run in real-time in a computer. The model produces a set of data describing the instantaneous value of various plant parameters, pressures, temperatures, etc., at various nodes in the process. These data are compared, again in real-time, with the equivalent actual plant measurements. If a plant fault exists, this will be reflected in the plant measured values. By detecting differences between model parameters and plant parameters, and by applying suitable computer logic, the existence and location of a variety of faults can be identified, often well before conventional alarms have detected any measurable change. The approach has produced good results in controlled studies.

A8.4.6 Overview of Intelligent Fault Detection Methods

Common to all the approaches described as intelligence fault detection, is that they derive or synthesise higher order statements about the plant from lower order information, e.g. process measurements, event information, or alarms. They must all be seen as add-ons which complement an existing good quality basic process alarm system. They will produce results if the basic information system is sound. None of the approaches will cure fundamental faults in the basic alarm system, and should not be considered as doing so.

The major problem with all the computerised fault detection techniques described is that, even with efficient implementation tools, they require considerable engineering analysis of plant behaviour. Some of this can be done 'on paper' from the plant design information, but generally considerable post-commissioning tuning is also required. Applying these techniques also demands some 'failure mode analysis' to be performed to ensure missing or incorrect input data can not cause false conclusions. Questions also remain with the artificial intelligence and expert systems techniques about demonstrating that a procedure that is developed on a limited range of plant transients will be effective in unexpected situations.

The development and application of any of the techniques described requires specialist knowledge. In all of the approaches, not insignificant computing and process engineering resources are required for any credible industrial process size. Pattern recognition approaches require large computing machines. Neural networks require significant amounts of plant data for training and are relatively unproved compared with the other techniques. Knowledge-based approaches require a great deal of data extraction and data engineering, together with specialist software. Mathematical models require an adequate model to be constructed and proved.

The more advanced methods are an active area of research and development. The HSE survey (5) provides a comprehensive literature review. Unfortunately the reports of large scale practical applications on working plants are few and far between. A key point is that all the advanced analysis methods work best when there are a minimum of spurious or low value alarms. Thus, priority should be given to applying other, more basic methods described in this Guide to eliminate the simple problems, and only then to invest in the more advanced methods.

A8.5 Automatic Alarm Load Shedding

As discussed in Section 2.6, there are fundamental limits on the amount of information that any human operator can assimilate and the number of actions the operator can perform. In very simplistic terms, a human is like a computer with limited processing power. A plant operator has many demands on their time, and consequently, unless the operator is to be overloaded, only a proportion of this available 'processing power' can be allocated to reading and responding to alarms.

On the other hand, on most process plants there will be a large variety of plant upset scenarios that could occur and result in many alarm signals going into their alarm state within a short period. The logical consequence from this is, that there is almost always a potential for the alarm load that the alarm system can generate to exceed that which the operator can handle. This potential exists even on a plant with well designed alarm systems with all spurious alarms eliminated. This implies that meaningful alarms are sometimes being generated with which the operator cannot be expected to deal.

If the operator is overloaded with alarms the operator will have to find some way of coping with the situation. The evidence is that the strategy taken is often very crude, e.g. to accept alarms without reading them, or to abandon use of the alarm system altogether, or to look only at selected alarms (e.g. high priority alarms or only the alarms shown on a particular overview screen). Alarm overloads are often associated with times of particular operator stress, and the operator should not be expected always to deal with the overload in the way that appears most sensible in hindsight.

It appears, therefore, that there could be considerable benefit if the alarm system were automatically to limit the alarm display rate to a rate with which the operator could cope reasonably comfortably. To do this, the alarm system would have to make some selection of the operationally most appropriate alarms to display. Some suggestions for algorithms for doing this are given in (4). The selection algorithm does not have to be optimal in its selection; it just has to be better in dealing with overloads than the average stressed operator. It is better for the algorithm to display most of the meaningful alarms in a way that the operator can understand and reject a few alarms that might potentially have been useful, than for the operator to be overloaded and ignore all alarms.

At the time of publication, this form of automatic alarm load shedding remains a research concept rather than a proven practical method. However, it is seen as an important direction for future development. Note that if this method were to be applied to select or suppress safety related alarms, it would need to be assessed as part of the safety related system.

A8.6 Alarms from Equipment Under Test

It is common for numerous alarms to be generated from plant and equipment when it is undergoing maintenance or testing. Routine testing of automatic protection systems can be a particular problem. Logic can, in principle, be used to automatically suppress these alarms, but this is difficult because:

- testing is generally carried out during plant shut downs, and the maintenance work which is often being done in parallel with the testing could invalidate the status signals needed by the suppression logic;
- the suppression logic needs signals to detect the start and end of test activities.

An alternative approach is to provide facilities so that operators can manually demote predefined groups of alarms to a priority at which they are displayed but do not generate an audible warning or require acceptance. This approach is relatively simple to implement but does require responsible and systematic use by the operator to avoid alarms being left demoted when testing is completed. (The use of a time-out to ensure that the priorities are automatically reset should be considered). It is also important to recognise that equipment under maintenance is a common source of hazards and care should be taken in the suppression of alarms indicating these hazards.

Appendix 9 Repeating and Fleeting Alarms

This Appendix discusses the source of repeating and fleeting alarms and describes a number of techniques for dealing with them. These include:

- *engineering noise-free alarm signals*
- *low pass filtering to eliminate noise*
- *transient suppression of fleeting alarms*
- *deadband*
- *de-bounce timer*
- *counter*
- *shelving*
- *release ('one-shot' shelving)*
- *auto-shelving*
- *single line annunciation*
- *logging of repeating alarms*

A9.1 Why Repeating Alarms Occur

Repeating alarms, i.e. the same alarm raising and clearing repeatedly over a period of time, are the most common form of nuisance alarm. On a typical plant, repeating alarms may account for around 50% of the alarm annunciations (5), (7), and (51). They cause nuisance to the operator because the operator will (depending on the alarm system design) have to silence the alarm hooter and/or accept the alarm each time it occurs. In addition, in some designs, repeating alarms can cause the alarm list display to fill up with the same alarm message, obscuring other messages.

This Appendix deals with ways in which the nuisance of repeating alarms may be reduced or eliminated. Note that much of the discussion is also applicable to fleeting alarms, i.e. alarms which are raised and clear shortly afterwards.

Repeating and fleeting alarms can be generated in several ways, e.g.:

- noise on a process variable when it is near an alarm setting;
- real high frequency fluctuations on a process variable;
- repeated action of on-off control loops.

Figure 17 shows an example of how noise (at a fixed frequency) on a process signal can result in repeating alarms, if the alarm is generated by comparing a process variable against a fixed setting.

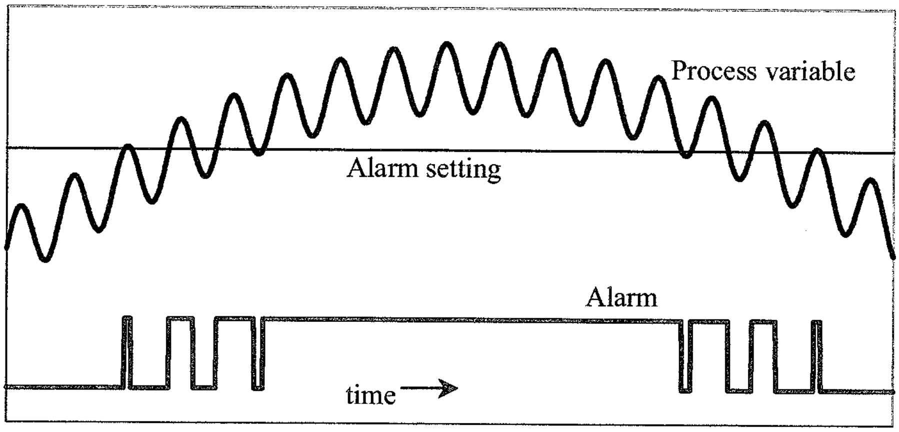

Figure 17 Repeating alarms generated by noise on a process signal

Since repeating alarms are often associated with unexpected signal fluctuations, they are often generated during plant upsets when alarm loads tend to be high and additional nuisance alarms are particularly unwelcome. Hence, it is important to provide protection against them.

In most circumstances, repeated annunciations of the same alarm are of little significance to the operator. When the alarm first occurs the operator will take the action they think appropriate. This may be action that will result in the process variable being moved back from the alarm setting, and in this case, the alarm may repeat a few times then go away. Alternatively, the operator may take a considered decision to note the situation, but take no corrective action. In this case the operator will tend to assume that further annunciations of the alarm are simply indicating that the process variable is moving around near the alarm setting. Note however, that if the process variable generating a repeating alarm moves further into alarm and the alarm becomes steady, then this may be of great operational importance. Consequently the suppression of repeating alarm should always be done with care.

There are a number of ways of dealing with repeating alarms as described below. For effective treatment, a combination of several of the methods should be used. Some techniques involve identifying effects on individual alarms that may cause them to repeat and eliminating these effects. However, it has to be recognised that in practice it is virtually impossible to totally eliminate all sources of repeating alarms, and some backstop methods - such as single line annunciation - are also needed.

A9.2 Engineering of Signals

Noise on analogue plant signals can cause nuisance in controllers, on pen recorders, on graphic displays and can also result in spurious repeating alarms. Process measurements should, therefore, be designed to be as 'clean' as possible. For example, this means not locating instruments at points where the process flow is turbulent, and avoiding measurements which are the difference between two noisy signals. Many problems can be avoided just by applying established good instrumentation practices

A9.3 Filtering

Modern instrumentation can be very responsive and can often follow the high frequency noise on process signals. Low pass filtering is frequently required to reduce noise from such signals.

Low pass filtering may be applied using digital algorithms within the alarm processor. However, because of aliasing effects, any noise above the digital input sampling frequency should be removed by analogue filtering before digital filtering is attempted[29]. This filtering can be achieved either by choosing an instrument with an in-built adjustable low pass filter, or by having filters in the alarm processor analogue input cards.

[29] Digital filtering in the DCS puts load on the DCS processor. This can be a reason for performing all necessary noise removal with analogue filters.

Table 20 gives recommended default values for filter time constants.

Signal type	Time constant
flow	2 sec
level	2 sec
pressure	1 sec
temperature	0 sec

Table 20 Table of filter time constants

It is recommended that time constants should be individually tuned for signals used in controllers. Filter time constants should be recorded in documentation. In some cases the filtering needs to be specifically developed for the particular signal. An example of this would be a high current alarm for an electric motor. High current will cause damage if it persists for some time and the motor components heat up, but a short transient over-current will not necessarily cause damage. The motor current signal is typically very noisy. If it is simply passed through a low pass filter and a limit detector, then it can be hard to find a filter time constant that avoids fleeting or repeating alarms and still generates the alarm before damage occurs. Use of an alarm based on a low-pass filtered version of the deviation above the current limit can reduce the number of spurious annunciations and still provide warning of damage.

A9.4 Transient Suppression

Some processes are known to transiently pass through an alarm state. For example, during start up, the current drawn by an electric motor may transiently reach levels well above the steady state high current limit. Filtering based on a plant running signal may be used to briefly suppress such expected transient alarms, e.g. for 3 seconds following the starting of a pump.

A9.5 Deadband

When a deadband is applied then, as shown in Figure 18, the alarm is arranged to be raised at one level but cleared at a different level.

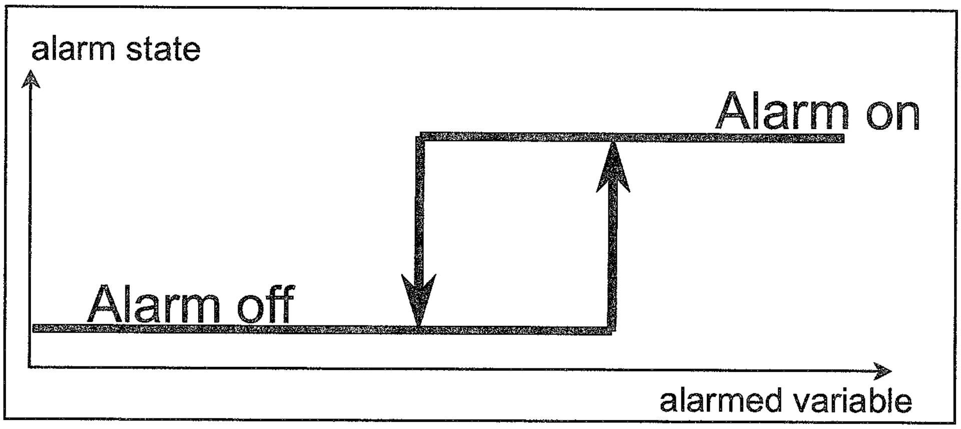

Figure 18 Transfer function of deadband

An example of its effect is shown in Figure 19. Notice that the size of the deadband has to be greater than the process fluctuations for it to be effective.

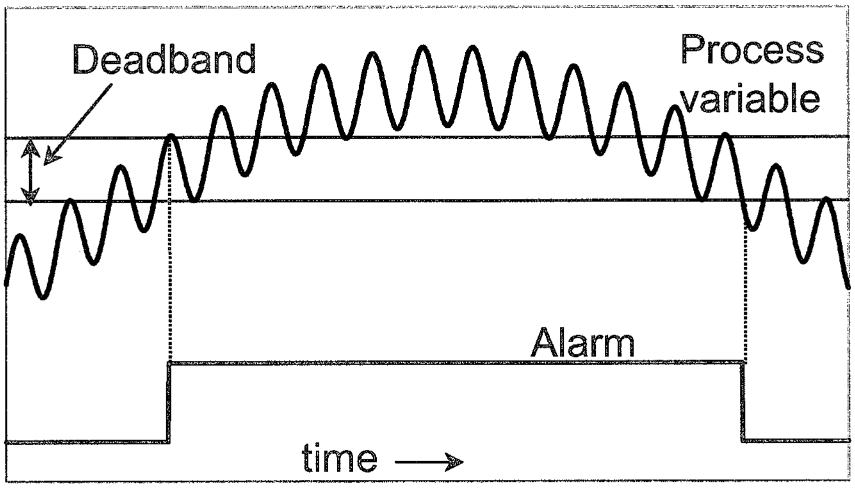

Figure 19 The elimination of repeating alarms using deadband

For alarms generated from analogue measurements, the use of a deadband is often very effective in eliminating repeating alarms. Table 21 shows recommended default values for deadband settings:

Signal type	Deadband
flow	5%
level	5%
pressure	2%
temperature	1%

Table 21 Table of default deadband settings

To get best performance the size of the deadband will need to be tuned to match the characteristics of the process signal. This should be part of the commissioning of the alarm system. In addition, deadbands may need to be optimised during the subsequent operation of the plant if particular alarms are found to be causing nuisance by repeating frequently[30].

Digital alarm sensors such as limit switches often include a small mechanical deadband that cannot be adjusted. If the process variable is noisy, this inherent deadband is often insufficient to stop repeating alarms being generated. Then some other method of repeating alarm suppression will need to be used or the alarm should be changed to one derived from an analogue measurement.

[30] Repeating alarms are often annunciated during plant shut downs and start ups as process variables move through alarm settings. Optimisation of deadbands tends to concentrate on the nuisance alarms occurring during normal operation, and these less usual repeating alarms can get neglected. They then cause a high operator load in a time of particular high plant activity and (unless the control and alarm systems are well engineered) associated high operator stress.

A9.6 De-bounce Timer

Since many repeating alarms are ones that are occurring frequently, one way of reducing the nuisance that they cause is to use a timer to eliminate cleared/raised signal pairs from the alarms displayed to the operator. For example, suppose after a long quiet period an alarm is raised, one second later it clears, and one second later it is raised again and remains raised. The operator needs to see the alarm raised only when it happens initially and if it stays raised.

An example of an algorithm for doing this is as follows:

- when an alarm is first raised it is annunciated immediately. When the alarm clears there is initially no change in the display to the operator and the change is 'held' for a period of time. If the alarm is raised again before this period has elapsed, then no messages are sent to the operator. If, however, the alarm is not raised again in the hold period, it is decided that it has truly cleared, and the appropriate change is made on the operator display.

Note that the precise details of implementation of this algorithm depend on how the alarm processor displays alarms to the operator.

A key aspect in the success of this method is the choice of the hold time period. For alarms on slow-acting systems (e.g. tank levels, temperatures) operators may be happy to use a timer as long as a minute. However, for an alarm where the operator would expect a fast response to a corrective action that the operator takes, e.g. a high current alarm on a motor that they can manually control, the operator may not accept a timer longer than a few seconds. This may limit the effectiveness of the method for eliminating repeating nuisance alarms.

Table 22 gives recommended default values for de-bounce timers.

Signal type	Hold time
flow	15 sec
level	60 sec
pressure	15 sec
temperature	60 sec
other	5 sec

Table 22 Table of default de-bounce timer parameters

A9.7 Counter

This method is based on detecting repeating alarms by counting the number of times the particular alarm repeats. This is done by having an up/down counter which is incremented (up to a maximum limit) every time the alarm is raised. If the alarm is repeating, the counter will eventually increment up above a 'repeating limit'. Up to this point, all alarms are displayed to the operator in the normal way. After the repeating limit is reached, the alarm is displayed to the operator as standing, and all subsequent raising or clearing of the alarm is hidden from the operator. However, at the same time the counter is being decremented down towards zero at regular intervals (e.g. once per minute). Once the counter goes below the repeating limit, the alarms are displayed to the operator in the normal way.

A key aspect in this method is choosing the values of the counter maximum limit, repeating limit and decrement rate. In practice 'repeating alarms' have a wide range of characteristics. Some raise and clear very quickly over short periods of time, others at slower rates and over longer periods. Some repeat intermittently. Consequently, it is difficult to choose universal values for the three parameters, and they require individual tuning.

A9.8 Shelving

Shelving is a facility where the operator is able to temporarily prevent an alarm from being displayed to him when it is causing him nuisance. A shelved alarm will be removed from the list and will not re-annunciate until un-shelved.

Shelving is a very useful and powerful operational tool for the management of an alarm system. The reason for this is that, even with excellent system management, it is inevitable that some alarms will sometimes become temporarily of no value. This may be due to instrument malfunction, lack of maintenance or noise causing repeating alarms. The usability of the alarm system is increased by allowing an operator to eliminate the distraction from alarms that the operator knows are of no operational relevance.

Because of its power to 'hide' alarms, shelving should be used only in the following circumstances:

- the operator has quick and easy access to view the list of shelved alarms and can print it out;
- the operator can very easily un-shelve an alarm;
- the operating procedures require the operators at shift changeover to check the list of shelved alarms and the reasons for them being there;
- the operators are fully trained in the implications of shelving and are seen to be using it responsibly;
- practices are in place such that one operator controlling a plant area cannot shelve an alarm without another operator controlling the same plant area (if there is one) being aware of it.

Shelving is normally performed from an alarm list. Each alarm to be shelved should have to be individually selected. Once selected, it is acceptable to allow several alarms to be shelved as a group.

It is desirable if shelved alarms (and released alarms) are displayed on plant mimic graphics in some fairly low contrast coding that is less obvious than normal alarms[31].

Whilst not essential, there can be advantages in requiring the operator (or the supervisor) to fill in a form when alarms are shelved. This might record:

- the alarm shelved;
- the person shelving the alarm;
- the reason for shelving;
- the time of shelving;

[31] Note that this feature implies that the shelved and released alarms are processed within the alarm handling software and not blocked out at the input to it (as is done in some implementations of alarm shelving).

- the planned time of un-shelving;
- additional information (e.g. the number of the job card raised to cure the fault).

It may be beneficial to present this form to the operator on the alarm display to ensure it is filled it in when the shelving is done. Then, at the planned time of un-shelving, either the operator might be prompted to un-shelve the alarm, or the alarm might be automatically un-shelved.

A further feature that is often provided is a default 'maximum shelve time allowed' after which the alarm is automatically un-shelved. Some users consider this a core safety feature.

Careful consideration should be given to whether every alarm should be shelveable, or whether some alarms, e.g. critical alarms, should be defined as un-shelveable. However, in practice faulty critical alarms can be just as much of a nuisance as other faulty alarms and, therefore, should be allowed to be shelved unless there is some other method (e.g. single line annunciation) that would limit the nuisance caused in this situation.

It should be the case that critical alarms are only shelved on a temporary basis, whilst all efforts are made to remedy the fault. Formal authorisation from a technically competent authority should be required for such shelving to be instigated. Ideally, automated reminders of the situation should be generated on a regular basis, to prevent the shelved status becoming permanent.

Thus, the restriction should be applied only if there is some other method, such as single line annunciation, that limits the nuisance from faulty un-shelveable alarms.

A9.9 Release ('One-shot' Shelving)

A 'release' is a facility that can be applied to a standing alarm in a similar way to which shelving is applied. A released alarm is temporarily removed from the alarm list and put on the shelf. There is no indication to the operator when the alarm clears, but it is taken off the shelf. Hence, when the alarm is raised again it appears on the alarm list in the normal way. Thus, the release is effectively 'one-shot' shelving. This facility is useful when there is an alarm which the operator fully understands (e.g. because it is from a plant under maintenance) and expects to stand for some time.

A9.10 Auto-shelving

A paper (9) describes use of an algorithm for automatically shelving alarms which are detected as repeating frequently. The algorithm works as follows:

- if more than 9 occurrences of an alarm occur in 5 minutes or less, then the 10th alarm is marked on the VDU screen in a colour to indicate it is a 'repeating' alarm. When this is accepted by the operator, the alarm is automatically shelved for 20 minutes. After 20 minutes it is put back on the alarm list 'on trial'. If it does not repeat more than 9 times in any 5 minute period during the next 20 minutes, then it ceases to be 'on trial' and becomes 'normal'. However, if repeating does recur when the alarm is on trial, then the alarm is automatically re-shelved for twice the original period (i.e. for 40 minutes). This process of doubling up the shelve time can continue up to a limit of 640 minutes. When the alarm has been automatically shelved, the operator can un-shelve it manually if the operator wishes (though this does not reset the on-trial timer).

Note that this method differs from the single line annunciation method in that, once an alarm has been shelved, it ceases to sound the audible warning.

This algorithm has been used in one power station for a number of years and is well-liked by the operators. Although the shelving process is done automatically, the operators are aware of what is going on, still feel in control, and have the benefit of fewer disturbances. It would appear to be a method that might usefully be provided in other systems.

A9.11 Single Line Annunciation

As mentioned above, repeating alarms can be a particular nuisance if they are displayed on alarm lists and allowed to fill up the display shown to the operator.

As discussed in Appendix 11, different manufacturers display alarm lists in different ways. In the method detailed in that Appendix, the alarms fill the page downwards, and when the page is full, the operator manually brings fresh alarms onto view. In other systems, the latest alarm is at the top of the displayed page, and when a new alarm comes in, it pushes the older alarms down the page.

Despite all the variants in alarm list displays, one key principle is that the alarm list displayed to the operator should be designed so that it cannot be swamped with repeats of the same alarm. The importance of this has been demonstrated in simulated replays of incidents (18), and has been taken into account in the design of many (but not all) proprietary DCS/SCADA systems.

The key principle means that, if there is a page of an alarm list on display to the operator, then a repeating alarm should only be shown once in this page.

If an alarm is already displayed on the page shown to the operator there are two options for dealing with a repeat occurrence of it:

- either the existing alarm message is left where it is, but it is updated to show its current state and time, or;
- the alarm message is moved to the position of the newest alarm, the old message is deleted, and the list is shuffled up/down as appropriate to eliminate the space.

Each method has disadvantages. In the first method the page on display loses its chronological order. In the second method there is some 'unnecessary' shuffling of the list that can make it harder to read. Despite these disadvantages, either method is preferable to allowing the displayed alarm page to be flooded by a repeating alarm.

It might be argued that this method of dealing with alarms is not needed if there is proper setting of deadbands and de-bounce timers, and there is a facility for the operators to shelve alarms. This argument is not valid for the following reasons:

- there will be many hundreds of alarms on a large plant. The deadbands and timers are unlikely to be set on all of them to cope with every unusual circumstance that may arise. It only needs one plant signal to go into an unexpectedly large oscillation for repeating alarms to occur and for the operator display to become flooded;

- experience shows that operators tend not to be prepared to devote effort to shelving repeating alarms in the operational stress of a significant plant upset. It is desirable for the alarm system to automatically do what it can to avoid the alarm list becoming unusable[32].

It should be considered an essential requirement that any system providing an alarm list display to the operator, should not be swamped by a repeating alarm.

A9.12 Logging of Repeating Alarms

All the methods described in this section relate to improvement of the display of repeating alarms to the operator. Some consideration may need to be given to the logging of repeating alarms. If all repeating alarm annunciations are printed or recorded to magnetic media, this can consume a considerable amount of paper/disc space. For example, some proprietary DCSs only have capacity for a few thousand alarms in their history buffer. If there is a major plant disturbance which also produces many repeating alarms, this buffer can quickly become full resulting in the loss of records of what caused the incident.

The best solution to this is, firstly, to eliminate repeating alarms at source and, secondly, to install a magnetic store large enough that all alarms can be logged. However, if this is not practical, storage can be minimised by only putting an alarm in the log if:

- an alarm that had been both cleared and accepted became raised, or;
- an alarm that had been cleared was accepted by the operator, or;
- an alarm that had been accepted became cleared.

This can be very effective in reducing the disc space taken up by repeating alarms and is certainly better than allowing incident records to be overwritten. However, it has the disadvantage that the log is no longer 'complete', and, for example, it becomes harder to analyse it to identify repeating alarms. Consequently it is preferable to install bigger history buffers wherever this is possible.

A9.13 Summary

The methods described in this section can be very effective in reducing the nuisance from repeating alarms. If they are properly set up, the deadband, de-bounce timer, counter and single line annunciation methods provide benefit with little downside. There is benefit in using several of these methods in

[32] In an incident that occurred some years ago on a large industrial plant, a severe plant upset caused alarms to be presented to the operator at a rate of 150 alarms per minute for a period of 12 minutes. This completely swamped the alarm system. 90% of this alarm activity was due to two repeating alarms. Without these two alarms the average alarm rate would have been reduced to about 15 per minute; still unmanageably high but rather better. The operators were very active in controlling the upset, but did find time to shelve the two repeating alarms 12 minutes after the start of the incident, which is why the rate fell at that point. The message from this incident was that even very good and well-trained operators should not be expected to manage repeating alarms under stress, and the alarm system should do what it can to automatically minimise the disturbance they cause.

combination. The shelving and auto-shelving methods can be very useful but require the operators to be aware of their implications and behave responsibly.

None of the methods provides the operator with a good warning of the alarm that repeats for a period and then moves significantly into the alarm region. More sophisticated alarming, e.g. use of high and high-high alarms may be required here. Alternatively, some sort of adaptive filtering of the underlying trend in the process signal might prove effective.

Appendix 10 Design of Field Alarm Sensors

This Appendix deals with issues relating to the field sensors used for generating alarms. This is an important topic since many alarm system problems derive from shortcomings in the sensors. Topics dealt with include:

- *the choice of analogue sensors or switches;*
- *the location of sensors;*
- *the choice of sensor range;*
- *the validation of signals from alarm sensors;*
- *the transmission of alarm signals;*
- *dependent failures in alarm sensors.*

A10.1 Choice of Alarm Sensor

There is frequently a choice of either generating an alarm by comparison of the output of an analogue sensor against a setting, or by installing an alarm sensor incorporating a switch (e.g. a movement limit switch or a fluid level switch).

Relevant issues in this choice are:

- **testing** - to test a switch type alarm it is often necessary to bring the process variable up to the alarm setting and confirm that the alarm is generated. This may involve disturbance to operation and bringing the plant close to an unsafe state. For thorough testing the same procedure will be required for an analogue alarm. However, it may often be acceptable to carry out a partial test of an analogue alarm by, firstly, checking that the measurement is calibrated and functioning correctly in normal operation, and, secondly, checking that a simulated measurement above the alarm setting does generate an alarm. Care should be taken to ensure that testing in this manner does not leave any element, e.g. impulse lines or part of the transmission cable, untested. Adjusting the alarm setting to initiate an alarm at operating conditions is not a valid method since this does not check that the transmitter and receiver are able to achieve the alarm setting;
- **adjustment of alarm setting** - for an analogue alarm, generally, the adjustment of the setting at which the alarm operates can be simply achieved by changing a value in software or adjusting a setting on a trip amplifier. A few switch type alarms may include the means for easily adjusting the point at which they operate. However, for others, e.g. a fluid level switch in a tank, the point at which the alarm operates may depend on the physical location of the device. Hence, once fixed and mounted, it may be impossible to change the alarm setting without major mechanical work.
- **avoidance of repeating alarms** - repeating alarms can be produced by noise on the process variable (see Appendix 9). For analogue alarms, it is often relatively easy to remove this noise by filtering in the instrument or by introduction of deadband in the trip amplifier or software. Switch type alarm devices often have a fixed built-in deadband and limited inherent filtering. If these are too small compared with the normal fluctuations on the process variable, then repeating alarms can be generated that can be hard to eliminate;
- **covert failure rate** - the generation of an analogue alarm, generally, will involve more hardware than a switch type alarm, and consequently the hardware failure rate will tend to be higher. However, a significant proportion of faults in analogue alarm hardware are likely to be detected if the analogue measurement is validated and is also displayed to the operator. Consequently,

the covert failure rate of analogue alarms tends to be lower than switch type alarms, and they thus tend to have an overall lower PFD_{avg};

- **smart transmitters** - these offer a number of advantages and some disadvantages. They ease changes in configuration and have self-diagnostic features. Their configuration is, however, vulnerable to inadvertent change. If used for safety related alarms they should be from established manufacturers with quality approval;
- **cost** - switch type alarms tend to have a lower initial capital cost than analogue alarms.

Thus, analogue alarms are usually easier to test, easier to adjust, less prone to repeating, and have a lower covert failure rate than switch type alarms. They tend to have higher initial cost, but experience shows that in many cases the overall balance will favour the installation of an analogue alarm sensor. Switches may be appropriate where a diverse sensing mechanism is required to avoid dependent failures. They may also be justified in certain safety related applications.

If a designer is considering using a switch, the following points should be checked:

- the alarm setting is well defined and will not need to be changed;
- noise on the process signal will not cause the switch to generate spurious repeating alarms;
- a defined procedure exists for testing the alarm. It may be required that this is demonstrated as part of the plant commissioning;
- the covert failure rate of the digital alarm is acceptable;
- if operationally necessary, the operator has a means of observing how far the process variable is above or below the alarm setting.

A10.2 Measurement of Actuator Position

Alarms are often generated from actuator position signals. Where this is done, generally it is preferable to use a measurement of the final actuation element position rather than some measurement of the demanded position sent to the actuator. It is noted that both in the Three Mile Island accident (39) and the Milford Haven explosion (23) operators were misled by indications of demanded actuator positions that did not match actual positions.

Often, it may also be preferable to base alarms on measurements of process flows through valves or dampers rather than on actuator positions.

A10.3 Location of Sensors

Every alarm sensor should be located in a position where it will work reliably for a long period, generate a good signal and be easy to maintain.

Some specific points worth highlighting are:

- **location of switches** - see comments above on adjustment of alarm setting;
- **avoidance of noisy signals** - sensor locations should be chosen where the process variable being measured is stable. For example, if a pressure or flow instrument is located close to a bend or other obstruction in the pipework, then the signal is likely to be noisy and generate repeating alarms;

- **location for maintenance** - alarm sensors should be installed with suitable access for maintenance. In general, on continuously-operating plant, alarm sensors should be installed in such a way that they can be isolated from the process flow path for testing or replacement. This may require external chambers, bypass piping, etc. In special circumstances when in-production maintenance is not possible, it may be necessary to duplicate the alarm sensors.

A10.4 Sensor Range

In order to decide how to respond when an alarm occurs, the operator will often need to make some check of the plant item concerned. This may require an analogue measurement of the alarmed variable to be available for display to the operator, e.g. on a plant schematic display, so that the operator can see how close it is getting to the unsafe state.

In terms of alarm sensor range this implies:

- the range should be wide enough to allow any required alarm setting;
- if the measurement from the alarm sensor is displayed to the operator, a further margin should be provided so that the severity of the condition can be assessed.

In some cases, despite the dependent failure disadvantages indicated later in this Appendix, an alarm will be generated within a control loop by comparing the controlled variable against an alarm setting. When this is done, care should be taken in choosing a range for the measurement of the controlled variable which is suitable both for control and for alarm.

For control purposes, it is often desirable to use a narrow range instrument which extends only, for example, 50% beyond the required range of adjustment of the control set point. This will tend to increase control accuracy by reducing the effects of instrument drift, noise, calibration errors, etc. However, this narrow range may not cover the extreme limits that the controlled variable can reach safely during a plant perturbation. The optimum alarm sensor range (see above) may extend even further outside this narrow range.

If a suitable compromise cannot be found, then both a wide range alarm sensor and a narrow range control sensor should be installed, and both signals should be made available for display to the operator. The narrow range display enables him to check functioning of the controller.

A10.5 Validation of Measurements

Wherever possible, signals from alarm sensors should be validated on-line to identify faulty instrumentation. When instrument faults are detected:

- all alarms from the faulty instrument should be automatically suppressed;
- some form of 'bad data' notification should be given (possibly an alarm) so that maintenance work can be initiated[33].

[33] It may be desirable for notification of this fault to be directed towards the maintenance engineers rather than displayed as an operator alarm. However, the process variable should be shown as 'bad' on all graphic displays.

In practice it is often difficult to provide on-line validation of alarms from switch devices; more is possible with alarms from analogue measurements. One simple and common approach is to include a check on the range of the analogue measurement. This should be designed so that real instrument faults can be distinguished from out of range process variables. This is important because process variables are most likely to go out of range during plant upsets, which is a time when spurious bad data alarms are particularly unwelcome.

To take an illustration, suppose that a 4-20 mA signal is derived from a pressure transducer with normal operating range 0-5 bar and that there is a high pressure alarm at 4 bar. Suppose also that the instrument has been chosen so that, if the pressure goes above 5 bar, the instrument loses linearity and its output saturates at 22 mA. Then, when a 22 mA signal is detected, it can be assumed that the pressure is high and the high pressure alarm can remain raised. On the other hand, if the pressure transducer fails and its output goes to 25 mA, then this should be detected as an invalid signal and the high pressure alarm should automatically be suppressed and the bad data notification should be raised. To achieve this functionality will require careful matching of the instrument range with the alarm processor analogue scanner range, and careful design of the fault detection algorithm[34]. For example, it will be impossible to distinguish between the two different situations if the scanner range is limited to 4-20 mA.

Problems of spurious bad data alarms can arise if fault detection is improperly designed. In particular, repeating patterns of high and bad data alarms may be generated if the process variable hovers around the bad data limit. A normal deadband algorithm, which is the most commonly used technique for eliminating repeating alarms, will not be effective on this type of alarm. It is recommended, therefore, that validity checking algorithms should be designed such that they include the same deadband as applied to the alarm settings.

The above discussion has considered the most common method for validating individual analogue measurements. Other methods that can be used include rate of change checks and analysis of sensor signal characteristics to identify blocked impulse lines. It is also possible to do cross validation between sets of signals. For example, one faulty measurement may be detected in a set of similar measurements by voting or 'best-mean' type algorithms. If well engineered, these techniques can help to reduce the likelihood of low value alarms. More sophisticated model based validation techniques have also been applied (see the literature review in (5)).

A10.6 Signal Transmission

Transmission of an alarm signal from the field to the annunciation device should be considered equal in importance to the field sensing device and the annunciation system.

Where an alarm is part of a control or indication loop, then the integrity of the transmission medium should be to the level required by the most critical function as determined by the risk assessment. Where the alarm is independent from any

[34] One method of dealing with out of range signals is 'PV clamping', i.e. clamping the measured process variable at a limit if it goes outside say 2% to 98% of the normal instrument range. Whilst this avoids some bad data problems, the need to identify faulty sensors remains.

other loop, the integrity should be as determined by the risk assessment. For example, a safety related alarm for which a PFD_{avg} of 0.01 is claimed for the hardware, may require a two-out-of-three sensor vote, hard-wired cabling back to the annunciation facility, the cables from each field sensor to be separate and run in different cable ducts/trunking/tray/raceway, etc. In another case where the claimed PFD_{avg} is 0.1, it may be acceptable that the alarm signal be field multiplexed with single pair wires back to the annunciation system. In this case the probability of losing all signals, should a multiplexer fail or the cable be damaged, may well be acceptable.

A10.7 Dependent Failures

When carrying out risk assessments for alarm systems it is important to pay particular attention to dependent failures (sometimes referred to as common cause or common mode failures). Three issues relating to alarm sensors are:

- **sharing of measurements** - if alarm systems share measurements with protection or control systems, then failure of the measurement will disable the common functions. Take, for example, a tank where a level transmitter provides a signal used both for control and for a high level alarm. If the level transmitter fails and indicates a level below the controller set point, then the controller will act to drive the level until the tank overflows and the high alarm will be inoperative. Such dependent failures are important since often the reliability of instrumentation is lower than that of the alarm, control or protection equipment. It is noted, however, that sharing of a measurement between an alarm and a control or protection function does offer some benefit, as the alarm will indicate failures of the non-common elements (e.g. failure of the control system itself), so it may sometimes be acceptable;
- **sharing of instrument connections** - instrument connections into the process fluid can have a tendency to block or become severely restricted. Examples of this are on power stations where instrument lines can block with coal dust and on hydrocarbon processing plants where lines can wax up. If alarms, protection or controls have separate instruments but share connection points, then dependent failures can occur;
- **shared services** - shared services, such as transducer power supplies, nitrogen purge or steam tracing can cause dependent failures to occur.

The guidance here is that:

> **Where shared measurement, shared services or shared process connections are used, the associated risk of dependent failures should be carefully assessed.**

Appendix 11 Design of Alarm List Displays

This Appendix considers the design of alarm list displays. It is arguable whether there is one best way to display alarm lists, since many different ways have been found reasonably workable in practice. In general, if alarm rates are low, there are many acceptable ways of displaying them to the operator on an alarm list; if alarm rates are high, then it becomes progressively harder to display them satisfactorily on a list, and eventually all methods become unacceptable.

This Appendix describes one set of design features for alarm list display which has proved successful on a number of large industrial plants. Variations in approach are not necessarily unacceptable, but the implications of any differences should be carefully assessed. Many vendors supply systems that conform to most of these guidelines, and it may be best to review their default configuration before making changes. Small stand-alone alarm systems may be acceptable with reduced functionality.

This Appendix is intended to be helpful to designers of alarm processing systems or to users who are comparing alternative manufacturer's offerings.

A11.1 Alarm States

In order to provide recommendations on how display lists should work, it is necessary to define a terminology for the states that a displayed alarm may have. These are:

- an alarm is **raised** or initiated when the condition creating the alarm has occurred;
- an alarm is **standing** whilst the condition persists (raised and standing are often used interchangeably);
- an alarm is **cleared** when the condition has returned to normal;
- an alarm is **accepted** when the operator has indicated awareness of its presence (usually by push button or mouse click). It is **unaccepted** until this has been done;
- an alarm is **reset** when it is in a state that it can be removed from the displayed list.

Figure 20 shows the actions which cause transitions between these states and will help to explain some of the following guidance.

A11.2 Content of Alarm Entry

Each alarm entry in the alarm list should show:

- the alarm state marker (unaccepted, accepted, standing, clear, reset);
- the alarm priority marker;
- the alarm message;
- time and date (the latter is necessary to identify long standing alarms).

The state and priority markers may make use of a combination of symbol, text, brightness, colour and position coding with the most important markers being the most conspicuous.

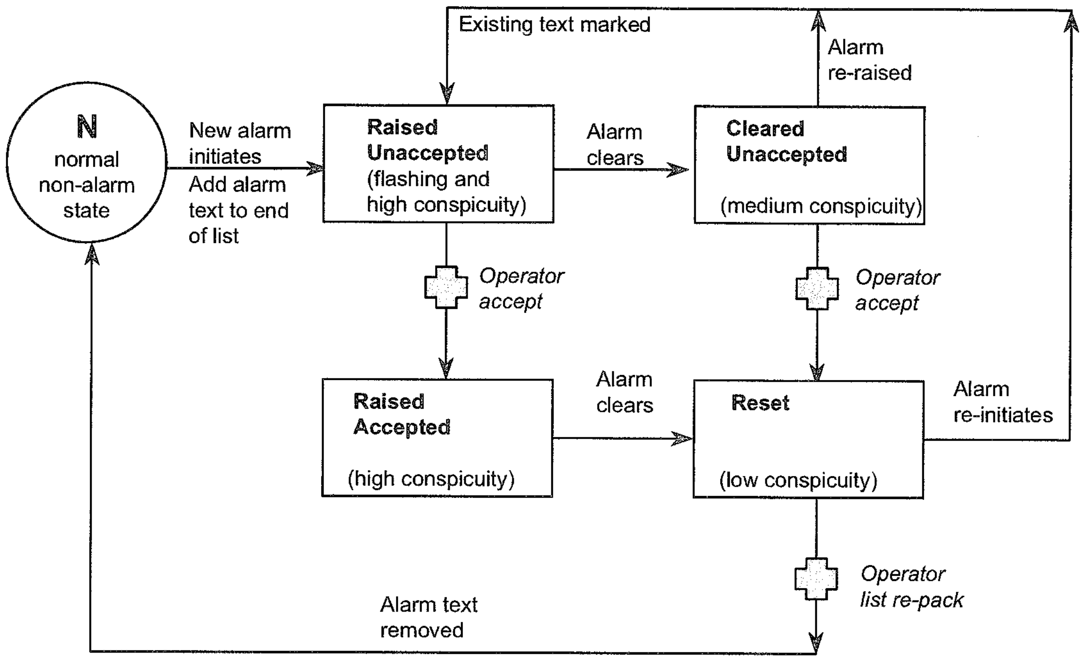

Figure 20 Alarm state transition diagram

Due to the possibility of the operator having impaired colour vision, colour coding should not be used on its own. Optionally, the alarm entry in the list may also show:

- value of the setting transgressed;
- current value of the alarmed variable;
- peak value of the alarmed variable;
- alarm tag number.

The design of the alarm messages should be performed with care. It is often found that there are some messages that the operators do not understand or to which they do not know how to respond. Table 23 provides guidance for avoiding this.

ATTRIBUTES OF A GOOD ALARM MESSAGE

- clearly identifies the condition that has occurred;
- uses terms that the operator is familiar with;
- uses consistent abbreviations from a standard site dictionary of abbreviations;
- has a consistent message structure[35];
- does not rely on the learning of tag names or numbers;
- has been checked for usability during actual plant operation.

Table 23 Attributes of a good alarm message

[35] This might mean, for example, that the message was always constructed in the order: variable, qualifier, status, e.g. boiler pressure A1 high.

Care should be taken when designing alarm messages to highlight the information which is most valuable to the operator[36].

Typically, an operator will receive many alarms and have to read many alarm messages over their working shift. Consequently the font size used in messages should be chosen to allow very easy legibility. The font size should be significantly larger than the minimum size ergonomically recommended for legibility from the operator's 'normal' working position. This will allow the operator to move back from their normal working position at times and also to discuss alarms with someone else standing further away from the screen.

It is also helpful to design the layout of the page of alarm messages in a way which makes it easy to remember the position of a specific alarm. This could be done, e.g. by placing a dividing line or gap after every fifth alarm.

New entries on the alarm list should be drawn to the operator's attention by a flashing indicator. The actual alarm entry message text should not flash. On acceptance the flashing indicator should be removed or steadied.

A11.3 Positioning of New Alarms

Experience shows that, to achieve a satisfactory alarm list display, the operator's view of the list should only change when actioned by the operator. Automatic scrolling, removal of alarms, or page movement are, therefore, to be avoided. Thus, if an alarm appears on the screen, it remains in a fixed position until some action is taken by the operator.

This implies that a page orientated alarm list display should be used[37]. The page should progressively fill downwards as new alarms come in (this matches normal reading practice). When a page becomes full, an indication should be given that following pages contain un-accepted alarms, and this should also indicate the highest priority of these alarms. However, a new page should only be displayed when a 'page on' command (or some other forward movement command) is made by the operator. This prevents the list changing whilst the operator is

[36] An example of this is as follows. Underground stations will have a number of escalators for carrying passengers to and from the trains. A modern escalator might typically be fitted with 120 alarms. A significant proportion of these indicate that the escalator has been automatically tripped by its protection. There are a range of different equipment faults that can cause this to happen. The primary operator response is the same for any of these alarms. The operator assesses the impact on congestion and takes steps to deal with any problem (e.g. by reversing the running direction of an escalator, starting another, holding people at the entry gates). The secondary response is to inform the maintenance people of the problem. It is only when it comes to the detail of this task that the operator needs the information about what the exact fault is.
In designing the message for one of these alarms it is most important to convey the fact that the escalator has tripped. The cause is secondary. Ways of achieving this might be:

- structure the message in two quite distinct parts (possibly in different colours) the first saying, e.g. "Escalator 7 tripped", the second saying "kink link (RHS) failed";
- show the primary message, e.g. "Escalator 7 tripped" on the alarm list and provide an easy way (such as a pop-up window) for accessing the diagnostic information.

Many similar examples of a group of alarms having the same primary response can be found across other industry sectors.

[37] This page orientated design is the preferred design of alarm list display. However, in practice, fully scrollable lists are often used, in which the alarms are added at the top and the list continuously scrolls downwards as new alarms come in. Whilst not preferred, some users may find this design acceptable.

reading previous alarms and also makes it easier for the operator to keep track of alarms that have occurred and to identify new alarms.

The list should be in chronological order (though this can be transgressed by repeating alarm suppression - see below).

A11.4 Alarm Acceptance

Alarm acceptance should be available on the list both on an individual alarm basis and on an 'all unaccepted alarms on view' basis. The acceptance of a single alarm should involve only one operator action (e.g. double-clicking on the alarm).

The acceptance of an alarm involves both silencing of the audible warning and acceptance of the alarm on the display. Separate silence and display accept buttons should be provided so that the operator can eliminate the nuisance from the audible warning and then accept the alarm on the display later when the operator has investigated it. These separate buttons also help the operator to manage a heavy alarm load. A combined silence and display accept button should also be provided as it allows further options for managing the list.

Cleared alarms should be displayed in medium conspicuity and annotated to indicate the clear state.

A11.5 Removal of Alarms

Except when shelving alarms, it should be possible to remove alarms from the list only when they are in the reset state. Typically alarms become reset when they are both cleared and accepted[38].

Reset alarms should be displayed in a low conspicuity. Reset alarm entries should be removed from the list only when requested by the operator. This is done by the operator performing a re-pack, either of the whole list or of the page on display.

A11.6 Movement through the List

As indicated above, movement through the list should be fully under operator control. Controls should be provided to:

- move to the start or end of the list;
- move to view the first unaccepted alarm;
- move up or down the list by varying amounts (e.g. manual scrolling, movement by one line, movement by a section, e.g. half a page, or movement by one page).

[38] Some plant owners install 'ring-back' alarms that require that the clearance, as well as the raising of certain alarms, must be accepted by the operator. Thus ring-back alarms do not become reset until their clearance has been accepted. This design of resetting is not represented on Table 23. A significant disadvantage of ring-back alarms is that they increase operator workload, so if used, the number should be limited.

A11.7 Dealing with Nuisance Alarms

As discussed in Appendix 9, operators are often faced by low value nuisance alarms. Standing alarms of low operational value can cause nuisance, but much more serious on alarm list displays are the problems caused by repeating alarms. They can cause the list to become overloaded and make it extremely difficult for the operator to view other alarms. Consequently alarm list displays should be designed such that repeating alarms do not cause them to become unusable.

Appendix 9 describes various methods for dealing with repeating alarms. As a minimum the following is recommended:

- effort should be made during the design of each alarm to ensure that it does not generate spurious repeating alarms;
- adjustable deadband should be included in all alarms derived from analogue inputs;
- manual shelving should be provided;
- single line annunciation should be provided to ensure that any repeating alarms that do occur do not make the alarm list display unusable.

A11.8 Display Filtering

A single 'All-Alarm' alarm list display can make identification of faults difficult when large numbers of alarms occur. A filtering capability allows smaller, more manageable lists to be selected for view. The following types of list may be useful:

- **priority** - this allows different priorities of alarms to be filtered out from the All-Alarm list. It can be used to make important alarms more accessible and obvious;
- **category** - one advantage of alarm annunciators is that the spatial organisation of alarms can make it easy to recognise patterns of related alarms. This is more difficult with alarm list displays. This can, to some extent, be alleviated by the use of category filters. Grouping alarms into categories, e.g. based round plant areas, and providing facilities to select alarm lists filtered on these categories is a highly desirable feature;
- **named lists** - this provides for a useful tool when performing specific tasks. Pre-configured filters can be defined so that only alarms associated with the given task are displayed when the list is selected;
- **modes** - some conditions that are alarmed during one mode of plant operation may not need to be alarmed during a different mode. Automatic detection of operating mode and suppression of appropriate alarms is one possibility (see Appendix 8), but care needs to be taken that the mode selection is robust. Manual selection by the operator of lists filtered according to operating mode can provide a simpler alternative;
- **unaccepted** - for systems where repeated alarm suppression methods can cause the list to be out of chronological order, it is useful to be able to view just those alarms that need accepting;
- **shelved** - the shelved list is, essentially, a filtered list.

Appendix 12 Performance Metrics

This Appendix describes a number of ways of measuring the performance of an alarm system. These include:

- *operator questionnaires (see also Appendix 14);*
- *alarm usefulness surveys (see also Appendix 15);*
- *incident recording;*
- *assessment of numbers of alarms in a system;*
- *measurement of average alarm rate;*
- *identification of frequent alarms;*
- *measurement of number of alarms following a major plant upset;*
- *measurement of operator response time;*
- *measurement of number of standing alarms;*
- *analysis of the priority distribution of alarms configured and occurring;*
- *correlation techniques.*

This Appendix also presents some benchmark values that may be used to guide a design or improvement programme. Some of these benchmark values are collected together in the main body of the report in Table 6.

A12.1 Operator Questionnaires

The key performance requirement of an alarm system can be summed up in the word 'usability'. Whilst this is difficult to define in quantifiable terms, an overall assessment of usability can be made by surveying operator opinion using questionnaires. This is relatively 'low cost' as operators can fill in the questionnaires during quiet periods whilst working at the control desk. Questionnaires can be of a general nature or can be used to follow up specific points in more detail.

A general operator opinion questionnaire is provided in Appendix 14. This questionnaire was used in the HSE-sponsored survey on 96 operators from 13 sites. If this questionnaire is used elsewhere, the results can be compared against the results given in the HSE report.

A12.2 Usefulness of Alarms

One measure of the performance of an alarm system is how 'useful' the operator finds the alarms. Are a high proportion of alarms of no or little value to the operator, or are they all indicating events which need action from the operator or, at least, provide information of real operational value?

Assessing the usefulness of alarms requires judgement rather than just making a simple quantitative measurement. One method of doing this in normal operation is by using a questionnaire survey (see Appendix 15)[39]. Results from this may be analysed to give a 'nuisance score'. It is suggested that a **nuisance score under 2.0** should be a benchmark figure at which to aim.

[39] Note that a survey of this sort provides operator views only on the usefulness of alarms, and may not reflect the designer's or the safety specialist's intent. For example, in a certain fault condition which an operator has not encountered, a certain alarm which is generally considered to be a nuisance, may be crucial. However, this does not invalidate the usefulness survey, but a proper examination needs to be made of any alarm that is rated of little use or as a nuisance, before any decision is taken to eliminate or change it.

A12.3 Incident Recording

As mentioned previously in Section 6.1, records of incidents involving failures of operators to respond to alarms, etc. can be used to provide a measure of alarm system performance. However, the problem with these measures is that such incidents occur relatively infrequently, so a long period is required to make confident estimates of average rates, etc. Hence they are not very useful for driving an improvement exercise.

Incident statistics are valuable for making longer term decisions, e.g. for use in investment appraisals. If incidents are recorded, it is desirable to identify and record alarm systems incidents with minor consequences, as well as major incidents. A general finding when considering injuries to people is that organisations with a high frequency of minor accidents tend also to have a high frequency of major accidents. The same principle is likely to apply to alarm systems incidents.

A12.4 Number of Alarms

In general, the more alarms installed per operator, the greater the likelihood of problems with spurious alarms, high alarm rates, etc. Conversely, a small number of alarms should not be viewed as necessarily desirable, as this may mean that the operator will be unaware of important events. A balance has to be struck, and it is inevitable that on a large complex plant a large numbers of alarms will be required.

Firstly, the sophistication of the alarm system and the effort put in to the optimisation of the individual alarms should increase as the size of the alarm systems increases. Some guidance based on industry experience is given in Table 25.

Secondly, guidance that can be given is on the number of alarms that should be expected to be configured per plant subsystem. The metrics given in Table 26 are only indicative, but may give some pointer to whether designers are likely to have installed too many or too few alarms on a plant. The first row gives the typical number of alarms deriving from each control actuator (e.g. valve, damper). The second and third rows relate to additional measurements and sensors not associated with actuators.

A12.5 Average Alarm Rate

Average alarm rates provide a good and simple indication of the workload imposed on the operator by the alarm system. Measurements made over reasonably long periods, e.g. every week, provide useful performance indicators. Rates should be calculated on a per operator basis and should count all alarm messages requiring acceptance by the operator[40].

[40] If alarm return to normal events have to be accepted by the operator, as is the case on a few systems, these should be included.

Number of installed alarms per operator	Approach to design and management
under 100	Simple technology (e.g. annunciators) likely to be acceptable. Operational problems may often be solved by individually tailored solutions, e.g. special electronics. One significant risk is that the operator and the alarm systems may have a safety role that is not fully appreciated. In addition, poor alarm system performance may not be recognised and given management priority.
100-300	If annunciator based alarms are used on their own they need to be very carefully designed to ensure good usability. Mixed annunciator and computer-based alarm displays can be effective. If purely computer-based displays are used, then the alarm system may not have to be sophisticated, but does need good tools for dealing with basic problems, e.g. repeating alarms, alarms from out of service plant. Effort should be put into identifying key alarms and ensuring that they are clearly displayed. Management commitment needs to be made to solve operational problems, but much may be achievable with relatively little effort.
301-1000	Sophisticated computer-based alarm handling with powerful tools for logical processing and suppression of alarms is very desirable, possibly in combination with individual displays of critical alarms. Significant effort has to be expended on ensuring a consistent and considered approach to the design of individual alarms. The optimisation of operational performance is likely to require involvement of a team of engineering and operational staff for several months.
over 1000	A major alarm system, which will require significant investment throughout its life cycle and which needs to be developed and maintained to the best practices available within industry, if it is to work well.

Table 24 How system size affects design and management

Metric	Low	Average	High
Alarms per control actuator	1	4	6
Alarms per analogue measurement	0.5	1	2
Alarms per digital measurement	0.2	0.4	0.6

Table 25 Guidance on alarms per plant sub-system

As has been discussed in Sections 1.3 and 2.6, the acceptable alarm rate will depend on the workload that the alarm system should be allowed to impose on the operator and the time that should be allowed for the operator to implement a response.

In normal operation, the majority of alarms should require the operator to carry out an action. This will, typically, require the operator to select a schematic, make a check of some operating conditions, and then carry out an action, such as raising a job card, phoning a plant assistant or changing a control set point. On the basis of user experience it is suggested that 1 minute is the minimum average time to allow for this per alarm. However, the operator will have other plant supervisory tasks to carry out and the alarm system should not take all their attention.

The benchmark figures given in the Table 26 result from an analysis of operating experience in a number of industries. Note that these figures may appear low compared with current industry averages. However, many current systems generate a high proportion of spurious alarms which, generally, can be dealt with more quickly than useful alarms.

Long term average alarm rate in steady operation	Acceptability
more than 1 per minute	Very likely to be unacceptable
one per 2 minutes	Likely to be over-demanding (industry average in HSE survey)
one per 5 minutes	Manageable
less than one per 10 minutes	Very likely to be acceptable

Table 26 Benchmarks for assessing average alarm rates

The operator is likely to find the alarm system harder to use if the alarm rate is variable rather than steady. Thus, the operator will find it difficult if there are no alarms for a long time and then there is a burst of many alarms. It might be reasonable to expect the operator to be able to cope if the alarm system takes their full attention for a period of, e.g. 10 minutes, but not if it does it for longer than this. This means that, on the basis of a 1 minute average response time, there should be very few 10 minute periods in which more that 10 alarms occur. Statistical measurements can be made to check if this is the case.

A12.6 Identification of Frequent Alarms

In an alarm improvement exercise, it is extremely useful to identify the alarms which are occurring most often. These can be given priority in review and correction work. In addition, statistics, such as the number of occurrences of the 10 most frequent alarms in a given period, can provide a powerful performance measurement.

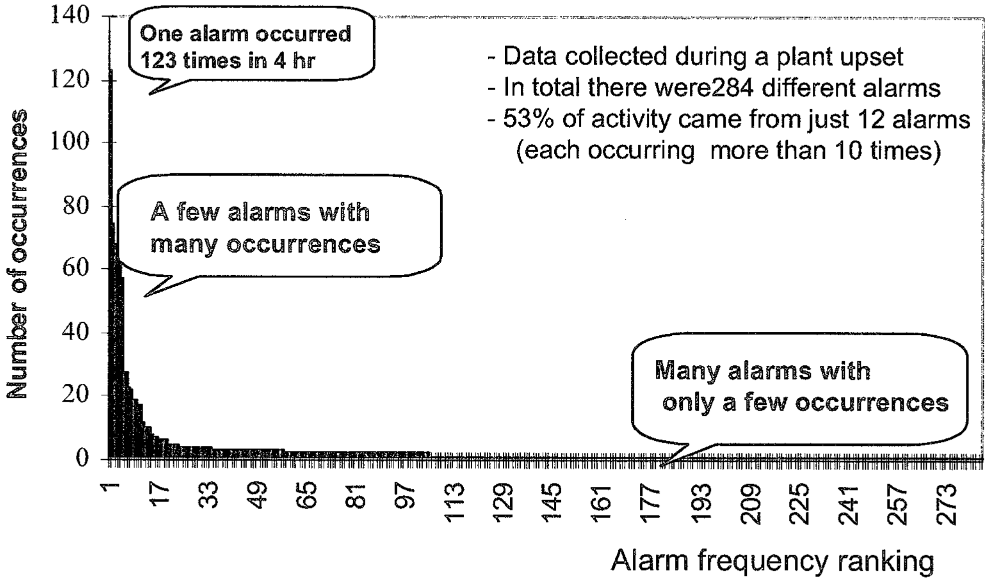

Figure 21 Example of alarm frequency analysis

Figure 21 provides an example of the frequency analysis of alarms at a petrochemical plant.

In the period analysed there were 284 different alarms, but just 12 alarms contributed 53% of the alarm activity. This sort of frequency distribution, where there are a few alarms generating much of the activity, is quite common both in normal operation and following upsets.

A12.7 Number of Alarms Following a Major Plant Upset

Plant upsets are periods when the performance of the operator is particularly important. During major upsets many plant sub-systems will be called upon to change operating point, to shut down or to start up. Consequently, there is a significantly greater likelihood of these sub-systems failing to work properly than in normal operation. Actions by the operator can have a very major impact in terms of avoiding production loss, plant damage and hazards to people.

In addition to the number of alarms in a period, it may also be useful to identify those alarms that were fleeting, i.e. were active only for a short period of time. Counts of the number of times that an alarm 'fleeted' may be an indication that it is, or could become, a repeating alarm.

Unfortunately, operators quite often find alarm systems very difficult to manage following a major plant upset, due to the very large number of alarms that can get displayed.

Figure 22 shows an example of the load that can be experienced. The stress of the situation makes the operator even less tolerant to such problems than in normal operation.

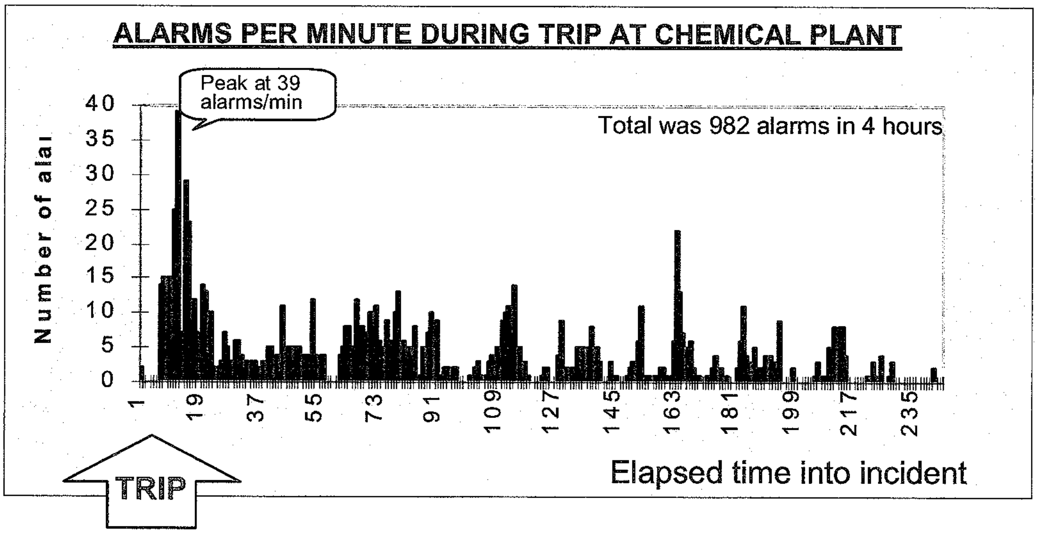

Figure 22 Example plot of alarms per minute

Measurement of the number of alarms occurring following major upsets is an important metric of the usability of the alarm system[41]. Various measures may be used, e.g. the number in the first minute, the first 10 minutes or the first hour after the upset. The most appropriate measure may depend on the plant dynamics and for how long typical upsets last.

Table 27 gives general guidance, based on experience, on the acceptability of alarm rates in the first 10 minutes following an upset.

Number of alarms displayed in 10 minutes following a major plant upset	Acceptability
more than 100	Definitely excessive and very likely to lead to the operator abandoning use of the system
20-100	Hard to cope with
under 10	Should be manageable - but may be difficult if several of the alarms require a complex operator response

Table 27 Guidance on alarm rate following an upset

Whilst the above metrics provide an indication of usability of the alarm system, in themselves they do nothing to improve performance. To do this, it is necessary to examine the post-upset alarms and identify ones of low operational value. Tools such as frequency analysis may be useful (see above).

A12.8 Number of Standing Alarms

There are several reasons why an alarm may be standing:

- the alarm is indicative of an operational problem that the operator should do something about or of which the operator should be aware. The alarm stands until such time as the operator is able to resolve the problem, and it should be very unusual for this period to extend beyond a maximum of a few hours. Such standing alarms are valuable as a reminder of things that the operator should be dealing with and which still need the operator's attention;
- the alarm is indicative of a problem about which others (typically maintenance staff) should do something. When the alarm comes up, the operator should take action to inform these people either verbally or using a work request form. Subsequently, the alarm provides a reminder that the work is outstanding. The operator may need to know about the implications of the alarm, e.g. that certain items of plant are unavailable, but this is little different

[41] Note that the reduction in the number of alarms is not the only way of improving operator performance following a major upset or plant trip, nor necessarily the most effective. For example, in the nuclear power industry, considerable emphasis is placed on having clear well-defined operating procedures. Thus, following a reactor trip the operator would refer to the book of post-trip procedures and carry out a highly structured series of checks and actions to ensure that the shut down proceeded properly. A number of hard-wired and computer-based displays are provided to assist in this task. This approach provides confidence that the operator will identify any critical failure in the shut down sequence. However, there are noticeable additional benefits if the alarm system can be designed to properly support the operator in these circumstances. This will also help in any plant upsets where the plant behaviour is not covered by procedures.

from many other plant status indicators. A 'maintenance shelf' (see below) should be provided to handle this type of alarm;

- the alarm is indicative of a problem that cannot be resolved in the short term. For example, it may have to wait until an annual overhaul. The alarm no longer serves any value in terms of making the operator aware of something requiring attention. Again a 'maintenance shelf' should be provided for this type of alarm;
- the alarm is spurious in the sense that it requires no action and does not indicate an operational problem (i.e. it is not really an alarm). One such example is alarms generated by maintenance staff from equipment under test (e.g. instruments being calibrated).

A high number of standing alarms indicates either that the plant is being badly operated or maintained, or that there are a lot of alarms being generated from things that do not require operator attention (and which possibly need elimination, suppression or re-prioritisation). Thus, the total number of standing alarms provides some indication of alarm system performance and may be used as a performance target.

A quick insight into the performance of the alarm system may be obtained by talking through the list of standing alarms with the operator. This form of 'snap survey' can also lead to the identification of alarms which are standing for no good reason.

More formal monitoring of standing alarms is often valuable. This should involve a periodical count of the number of standing alarms, e.g. once per week and should lead on to an identification and review of those alarms which are frequently found to be standing. This may be done automatically or manually. It is also be useful to identify **'long standing' alarms**[42] and, e.g. generate printouts of the ten longest standing alarms.

To ease the management of standing alarms, the alarm system should have a facility to put standing alarms on to a 'maintenance shelf' until such time as the problem is resolved. Maintenance shelved alarms should be displayed (in some suitably toned down format) on plant graphic displays, but should be eliminated from the operator's normal alarm list display.

It is recommended that a site with a large alarm system (i.e. containing greater than 1000 alarms) should be targeting to achieve on average **fewer than 10 standing alarms plus fewer than 30 shelved alarms** (excluding maintenance shelved alarms).

A12.9 Priority Distribution

The prioritisation of alarms is discussed at length in Section 2.5.1 and Appendix 5. The effective use of priority as a discriminator of important information was discussed. Information was presented in Table 15 on the approximate distribution of alarms into priority bands that should be aimed for during system configuration, and Table 14 showed what this should mean in terms of the occurrence rate of alarms during operation.

[42] To do this manually requires that the DCS allows the date at which alarms became active to be observed.

Useful information can be gained by measuring the rates and patterns of occurrence of alarms of different priority. For example, if critical alarms are occurring every 10 minutes on average, then this suggests that either:

- the plant is grossly unsafe, or;
- some critical alarms relate to events of low operational significance.

In either case review would be necessary. Similarly, if high priority alarms occurred only once per week on average, then it can be argued that the use of priority as an 'attention getter' might be more effective if high priority was assigned to more frequent events. In addition, if high priority alarms occur in groups, i.e. there tends to be a burst of several high priority alarms in a short space of time, then this would suggest benefit can be gained by more logical processing to combine these alarms into a single more meaningful alarm.

A12.10 Operator Response Time

If a computerised record of the alarm log and of the operator actions is available, it becomes possible to calculate the time between an alarm occurring and the operator accepting it. Long delays may be an indication of operator overload. Short acceptance times may be an indication that the operator is finding that many alarms are of little operational value and is accepting them without significant investigation.

A12.11 Correlation

Auto and cross-correlation are statistical techniques for analysing time varying signals to identify underlying patterns that may be hidden in noise. Consequently they have potential value in the analysis of alarm logs. To apply correlation analysis to alarm occurrences, a continuous analogue signal has to be generated, e.g. by giving it a value of 1.0 for the period when the alarm is active and a value of 0.0 when it is inactive. This signal can then be auto-correlated with itself or cross-correlated with the analogue signal derived from other alarms.

One exercise has been reported in which the alarm log was analysed to calculate the cross-correlation between each alarm and every other alarm. This required some significant processing, but gave results such as "about 70% of the time, Alarm A occurs 15 seconds after Alarm B". Such information was used to drive a review process, e.g. to consider whether Alarm A could be eliminated.

It is also suggested that auto-correlation might assist in the identification of nuisance alarms. For example, if there was a non-obvious repeating alarm occurring every 40 minutes, this would be shown by a sinusoidal component within the graph of its auto-correlation function.

Another option which may be considered is to correlate alarms and state changes, e.g. to identify that Alarm A is closely followed by Alarm B whenever Valve C is shut.

Note that operators may well be able to point out patterns in alarm occurrences based on their operating experience, and this may be more cost-effective, certainly for initial investigations - than using statistical methods.

Appendix 13 Performance Levels

This section expands the concepts defined in Appendix 12, Section 4.1.2 and also gives guidance on defining an appropriate performance level.

13.1 Defining an Appropriate Performance Level

It is important that the performance level that is appropriate for an individual asset is set as a target. This target will be affected by a range of factors, including:

- The range of tasks performed by the control room operator. In some cases this role is restricted to management of the process, using the DCS and its alarm system, with no requirement to leave the control room. In other cases, however, the control room operator may also be responsible for coordinating traffic (e.g. marine, aviation) around the plant, managing communications (e.g. radio, tannoy), authorising work activities and even touring the plant to take local measurements and performing manual operations. Where the operator's role is more focused (the former case), the absolute performance level that is appropriate is likely to be lower, since (simplistically) the operator can afford to spend a larger proportion of time interacting with the alarm system.

- Complexity. It is likely that increased complexity will drive a need for higher levels of alarm system performance. Simplistically, if it takes the operator longer to understand the implications and correct response to each alarm, then in any period of time the operator can deal with fewer alarms if the operator is to remain on top of the situation.

- The consequences of failure to act. For all assets there are a range of severities associated with the operator failing to respond to an alarm, as characterised by the prioritisation of the alarm. Application of standards such as IEC 61508 (29) and underlying risk criteria will tend to normalise these between assets, but in some cases, particularly on older plants or unusual processes, there may be a greater reliance on operator action in response to alarms in order to avoid significant hazards. In such cases, it is clear that a higher performance level for the alarm system would be appropriate.

- The required speed of response. As with the consequences of failure to act (see above), there is also a range of periods in which the operator has to react in order to avoid the consequences of missing the alarm altogether. Some assets will have generally faster process dynamics than others (e.g. due to process instability, exothermic reactions, high gas velocity processes) and in these cases a higher performance level for the alarm system would also be appropriate.

- 'Centrality' of alarmed plant. Even where the consequence of failure to act is a safe shut down of the plant, the relationship of alarmed plant to connected plants (with separate alarm systems) may also be important. The shut down of a water treatment facility may be relatively tolerable for a short period because of buffer storage capacity, whereas the loss, even temporary, of a core process unit (such as an ethylene cracker) may have altogether wider implications in terms of impact on associated units, as well as a higher commercial penalty. The performance level of the alarm system on the latter would be expected to be higher than for the former.

- Level of automation and fallback strategies. Plants with a higher level of automation will, typically, require less manual intervention and could therefore potentially manage with an alarm system exhibiting a lower performance level. However, what is probably more important is the abruptness of the transition to lower fall-back controls when these higher levels of automation fail for some reason. If the transition is abrupt (i.e. nothing between full multivariable predictive control and direct manipulation of valve positions), then a higher alarm system performance level would be indicated (i.e. the alarm system actually has to be designed to be usable at the lower level of automation).

- The cost of implementing higher performance levels. Particularly for plants with an older DCS, the cost of improving the alarm system performance level may be very high or the task may even be impractical. In these cases, provided the appropriate risk criteria are met, investments elsewhere to improve plant performance may deliver greater value (economic, safety and environmental) than efforts to raise the alarm system to a higher performance level.

In theory, the choice of performance level is independent of plant size and the number of operators. This is because the definition of each level in this vision is based on the qualitative 'feel' to the 'average' operator (in relation to their area of the plant), and because the quantitative metrics are expressed in absolute numbers (e.g. alarms per hour) per operator.

The implication of a larger plant scope under the control of each operator (measured, e.g. by the number of control loops) is that more effort will have to be put into the alarm system to achieve a given performance level. The measures to achieve each performance level are, therefore, quoted for a typical plant, e.g. a refinery or chemical plant operator managing around 200 control loops, or an oil production facility operator managing around 75 control loops (with, typically, a correspondingly larger number of monitoring points).

The choice of which alarm system performance level is appropriate will vary between assets, but as a general rule higher alarm system performance levels will deliver higher plant availability and safety.

	Performance Level: 1 - Overloaded
Characteristics	• Alarm system is difficult to use during normal operation and in practice ignored during plant upset as it becomes unusable • Low operator confidence in the alarm system, which is often ignored for long periods • Important alarms are difficult or impossible to discriminate from less important ones, and the alarm system gives little or no advance warning of plant upset • Many alarms are meaningless or of little value • Many alarms are 'commoned' together before reaching the operator, with no drill-down detail available • Alarms are often disabled by the operator because they represent a nuisance, and are frequently then forgotten about (i.e. never re-enabled)

Typical KPIs	• Av. Alarms / 10 min > 100 • Max alarms / 10 min >1000 • % hours when there were more than 30 alarms > 50%
Typical operator interface	• Diverse uncoordinated interfaces for alarms, including DCS and hard wired annunciators • Basic (default) DCS interface for alarms, typically a text-based only alarm display • Some basic alarm representation on process schematics • A significant proportion of information needed by operator workstation to interpret alarms is not available in the control room • Alarms are delivered to a printer but archiving is haphazard and printer is out of service for significant periods
Typical alarm system functionality	• Basic DCS alarm system with no supporting alarm response manual • Alarms are presented on a simple 'first-up' basis, with uncertainties introduced by variable system time delays • All alarms have to be acknowledged by the operator, although return-to-normal does not require any operator intervention • Many alarm setpoints are inappropriate for even the normal operating mode • Extensive use of default settings for alarm tuning with little customisation • Alarm journal of unpredictable duration available only on DCS history module, with no electronic archiving
Typical ancillary processes	• No site- or project-specific alarm management philosophy exists • No clear or agreed understanding of the purpose of the alarm system • No measurement of alarm system performance • No control of alarms disabled by the operator
Comments	• Represents a typical DCS during initial plant commissioning (or re-commissioning following plant re-instrumentation) where alarm management has not been subject to explicit focus by the project team
Typical focus for further improvement	• Establish a site-specific alarm philosophy document • Establish a well-defined change control process for alarms, linked to the agreed alarm philosophy • Analyse alarm journals to identify 'bad actors' and address these as a priority • Invest in software/hardware for electronic alarm journal archiving • Survey alarm tuning parameters (deadband, etc.) and implement generic improvements • Establish minimum (e.g. paper-based) control mechanism for alarms disabled by the operator • Improve alarm representation on process schematics, particularly for critical alarms

	Performance Level: 2 - Reactive
Characteristics	• Alarm system is more stable and useful during normal operation, but is often unusable in practice during plant upset • The operator reacts more to the rate of alarm generation rather than to the detail of the alarms themselves • Some heed paid by operators to alarm prioritisation, but known to be unreliable • The alarm system gives some early warning of plant upset • Some alarms are still meaningless or of little value • Alarms are often disabled by the operator because they represent a nuisance, and are sometimes forgotten about
Typical KPIs	• Av. alarms / 10 min < 100 but > 10 • Max alarms / 10 min >1000 • % hours when there were more than 30 alarms < 50% but > 25%
Typical operator interface	• Alarms are consistently delivered via the DCS, with drill-down detail available, albeit subject to the operator knowing where to look • Basic DCS interface for alarms, with a text-based alarm display and consistent identification (e.g. using colour/flashing etc) of at least the critical alarms on process schematics • Alarm annunciator display, if available, is limited to default functionality • Past alarms journals are not available to control room operator
Typical alarm system functionality	• Basic DCS alarm system with no supporting alarm response manual • Alarms are presented accurately on a 'first-up' basis within a manually selectable scope of plant (area or unit, etc.), with the uncertainties introduced by variable system time delays removed or well understood by the operator • Settings for alarm tuning (e.g. deadbands, time delays) are consistent with best practice, with some customisation for alarms that have historically caused problems • Alarm journal of unpredictable duration is stored on DCS historian for immediate history • Alarm journal is archived electronically for longer term availability, accessible from a non-DCS interface
Typical ancillary processes	• Site- or project-specific alarm management philosophy exists, but may not be widely understood • All significant changes to alarms are subject to formal control processes • Alarm system performance, including determination of 'bad actors' is measured periodically, but involves significant manual input • Some control mechanism exists for alarms disabled by the operator, possibly paper-based, but its use may not be consistent and there is no automated mechanism to ensure that alarms are not disabled and forgotten
Comments	• Some effort has been put into alarm management, but impact is limited and patchy

Typical focus for further improvement	• Reinforce alarm management philosophy and ensure site-wide adoption • Establish automated analysis and delivery of alarm system performance metrics (together with an ongoing 'bad actors' list) • Implement grouping of alarms with an identical operator action, and discrepancy alarming to identify associated actions • Carry out basic alarm rationalisation to reduce the content of the alarm system to only what is meaningful (as determined by the site alarm management philosophy) and identify the correct alarm setpoints • Implement software alarm shelving to support control of alarms disabled by the operator

	Performance Level: 3 - Stable
Characteristics	• Alarm system is reliable during normal operation, providing early warning of impending plant upset, but is less useful during plant upset • Operators are confident in the appropriateness of the alarm prioritisation and react consistently and fast to all critical alarms • All alarms are meaningful and have a defined response, although during process upset this may no longer be relevant
Typical KPIs	• Av. alarms / 10 mins < 10 but > 1 • Max alarms / 10 mins < 1000 but > 100 • % hours when there were more than 30 alarms < 25% but > 5%
Typical operator interface	• DCS interfaces for alarms include a text-based alarm presentation and a spatial-cognitive annunciator display, together with consistent identification of alarms (e.g. using colour/flashing, etc.) on process schematics • Critical alarms are always on view, using dedicated displays (or display areas) • All alarms are accessible with a single operator action and the system provides one-step links from each to the operating schematic appropriate to resolving the related problem
Typical alarm system functionality	• Basic DCS alarm system has been enhanced with key elements to underpin effective alarm management • Alarm setpoints are all appropriate for normal operation • Use is made of grouping for alarms with identical operator actions and discrepancy alarming to identify associated actions • Software alarm shelving is available for the operators to disable alarms as necessary, but with an enforcement function to ensure that alarms are not disabled and then forgotten • Bulk masking of alarms is possible, e.g. as a coarse filter based on priority during plant upset, but options are limited and little used by operators • Some dynamic treatment is carried out for individual alarms, but this is of limited functionality and constrained to single variable triggering
Typical ancillary processes	• Site-specific alarm management philosophy exists and is consistently applied

	• An alarm response manual that defines the action associated with all alarms is available and maintained. This is available at-line (i.e. in the control room, but not integrated with the DCS) for the operator to consult • Alarm system performance, including determination of 'bad actors' is analysed and delivered at least weekly by a fully automated process
Comments	• Significant effort has been put into alarm management, with demonstrable impact • Burst alarm rate is still a problem
Typical focus for further improvement	• Implement automatic dynamic alarm management for logical blocks of alarms • Improve usability of manually-initiated alarm masking features • Implement adaptive alarm tuning, e.g. to automatically suppress bouncing alarms • Integrate the alarm response manual into the DCS alarm system interface • Implement model-based multivariate alarming to provide early warning and avoid multiple single variable alarms

	Performance Level: 4 (Robust)
Characteristics	• Alarm system is reliable during all plant modes, including normal operation and plant upset • Operators have a high degree of confidence in the alarm system and have time to read and understand all alarms
Typical KPIs	• Av. alarms / 10 mins < 10 but > 1 • Max alarms / 10 mins < 100 but > 10 • % hours when there were more than 30 alarms < 5% but > 1%
Typical operator interface	• The alarm response manual, containing the expected operator action in response to each alarm and the likely consequence if this being ineffective, is available on-line (i.e. integrated into the DCS alarm system interface) for the operator to consult as necessary • The alarm system adjusts automatically according to plant operating mode, displaying only the alarms that are relevant under the current conditions • Priority safety and priority production alarms are always displayed in the same location on the operator interface to facilitate pattern recognition
Typical alarm system functionality	• DCS alarm system is fully enhanced for optimal alarm management • A large proportion of the alarms are treated dynamically, so that they can be annunciated to the operator only when they have a response that is appropriate for the current operating mode • Manually-initiated bulk masking of alarms is possible, based on a range of criteria (including priority- and equipment-based), and this is consistently used by the operators • Adaptive alarm tuning is applied consistently, e.g. to automatically suppress bouncing alarms • Some use is made of model-based alarming, to warn of deviation from multivariate relationships, before single measurements become significantly upset

Typical ancillary processes	- A full process of continuous improvement is established and running for the alarm system, with identified responsibilities and accountabilities. Key performance indicators are published at a high level in the organisation
Comments	- This possibly represents the limit of performance with currently-available technology
Typical focus for further improvement	- Implement automatic event diagnosis, combining pattern matching with surveillance of analogue variables in order to diagnose critical events that give rise to multiple alarms - Implement advanced alarm filtering, to remove predictable secondary alarms - Implement procedure monitors, to provide procedural support during critical operations, including identification of 'the next most important alarm/action' relevant to this task - Implement model-based intelligent operator support systems both (a) for individual alarms and (b) to guide the operator towards proactive intervention during normal operation rather than relying on reaction to alarms towards the edge of the operating envelope

	Performance Level: 5 (Predictive)
Characteristics	- The alarm system is stable at all times and provides the operator with the right information at the right time, in order to avoid process upset or minimise the impact of any upset that does occur - The operator actively 'patrols' the process schematics and corrects deviations before they are significant enough to cause an alarm
Typical KPIs	- Av. alarms / 10 mins < 1 - Max alarms / 10 mins < 10 - % hours when there were more than 30 alarms < 1%
Typical operator interface	- Extensive use is made of pattern recognition for alarm interpretation, both in the alarm annunciator display and on the operating schematics - Alarms are provided with individual intelligent support, giving an estimated remaining time to the associated consequence (e.g. time remaining to vessel overfill)
Typical alarm system functionality	- Automatic diagnosis, combining pattern matching with surveillance of analogue variables, provided to diagnose critical events that would otherwise give rise to multiple alarms - Advanced alarm filtering is used extensively to alarm only root causes - Procedure monitors are comprehensively available and used by the operators to ensure that all critical tasks follow the best practice and do not omit any key steps - Intelligent (expert) operator support systems prompt the operator to intervene proactively before an alarm is raised
Typical ancillary processes	
Comments	- Full realisation of the EEMUA 191 targets - This may not be achievable with currently-available technology
Typical focus for further improvement	- Not Applicable - this represents the best level of performance for currently available operator/DCS technologies

13.2 Validation of Metrics

The ASM Consortium has carried out an assessment of performance levels actually achieved in a number of ASM Consortium member plants. A summary of the results were given in a conference presentation (16). The abstract of that presentation is reproduced below.

"The Abnormal Situation Management Consortium has completed a series of studies related to effective alarm management practices for the refining and petrochemicals industry. These studies related directly to the alarm system performance guidelines published in the Engineering Equipment and Materials User Association's (EEMUA) Publication No. 191.

Results from 37 unique operator consoles indicate that the EEMUA recommendation for average alarm rate during normal operations (i.e., less than one alarm per 10 minutes), while not universally demonstrated, is achievable today. Our study found that about one-third of the consoles surveyed were able to achieve this recommended alarm rate guideline for normal operations and about one-quarter more consoles were achieving the EEMUA "manageable" level of 1 to 2 alarms per 10 minute period. However, the EEMUA recommendation for peak alarm rates following a major plant upset (i.e. not more than 10 alarms in the first 10 minutes) appears to be a challenge, given today's practices and technology. Only 2 of the 37 consoles came close to achieving the alarm rate guideline for upset conditions. This suggests that to achieve alarm system guidelines for upset conditions, more advanced site practices and alarm-handling technology (e.g., dynamic or mode-based alarming) are required. In studying the relationships between the observed alarm rate performance and other metrics collected, along with anecdotal information gathered (a subset of which is included here), we conclude that there is no "silver bullet" for achieving the EEMUA alarm system performance recommendations. Rather, a metrics-focused continuous improvement program that addresses key lifecycle management issues is most appropriate."

Appendix 14 Operator Questionnaire

This Appendix shows a questionnaire that may be used to assess operator's views on their alarm systems. This questionnaire was completed by 96 operators at 13 sites in the HSE-sponsored survey (5). If it is used at other sites, then comparisons can be drawn with the 'industry average' results given in the report of that survey.

Questionnaire for Plant Operators

Insert introductory text explaining:

- why the questionnaire survey is being carried out;
- what is intended to be obtained from the survey;
- whether the survey is confidential, and if so how confidentiality will be maintained;
- who is carrying out the survey;
- who to approach with questions regarding the survey;
- who to return the questionnaire to.

1. What is your job, and on what plant/unit?			
2. How long have you worked with the present alarm systems?			
	Years	months	
3. How well do the alarm systems support you in normal steady operations?			
very good	OK	poor	very poor
4. How well do the alarm systems support you during plant faults or trips?			
very good	OK	poor	very poor
5. What about the number of alarms in the system?			
too many alarms	many but necessary	few but adequate	too few alarms

Normal Steady Operation

6. How many alarms do you get in normal steady operation?

per hour

7. How often do you find that an alarm that comes up is a repeat of an alarm you have already seen in the last 5 minutes?

70-100% of alarms	40-70% of alarms	20-40% of alarms	less than 20% of alarms

8. Do you suffer from the following 'nuisance' alarms?

	often	sometimes	rarely
Alarms which are wrongly prioritised			
Alarms from plant that is shut down			
Two or more alarms occurring at the same time that mean the same			
Alarms occurring in a trip which are only relevant in steady operation			

9. What proportion of alarms is really useful to you in operating the plant?

all essential	most useful	few useful	very few useful

10. Do you fully understand each alarm message and know what to do about it?

always	mostly	sometimes

11. Consider a normal operating situation and 10 typical alarms. How many of the 10 alarms:-

Require you to take positive action, e.g. operate a valve, speak to an assistant	
Cause you to bring up a format and monitor something closely	
Are noted as useful information	
Are read and quickly forgotten	

Plant faults and trips

12. How many alarms would you get during a large plant fault or trip?

in the first minute	in the next 10 minutes	in the next hour

13. Do you keep an alarm list display on permanent display during a large plant fault or trip?

Yes		No	

14. How often do you look through the alarm list display during a large plant fault or trip?

several times a minute	once every couple of minutes	once every 10 minutes	less than once every 10 minutes

15. How often in a large plant fault or trip do the alarms come too fast for you to take them in?

mostly	sometimes	rarely

16. How often in a large plant fault or trip are you forced to accept alarms without having time to read and understanding them?

always	quite often	sometimes	never

17. Does the alarm system help you to pick out key safety related events during a large plant fault or trip?

very well	some help	little help	a nuisance

General

18. What do you think of the procedures for getting changes made to alarm settings etc.?

over-restrictive and cumbersome	strict but safe	easy to use - but you have to be careful what you do	sloppy and uncontrolled

19. Compared with the other things they do to improve your control systems, do your site engineers put enough effort into improving the alarm systems?

too much	about right	too little

20. What features of the alarm systems do you like best?

21. What features of the alarm systems do you like least?

22. If you could change any part of the alarm systems what features would you add to help you run the plant?

23. What features would you remove because they do not help or you do not like them?

24. Can you add any other comments which might help us improve alarm systems?

Please continue overleaf

Appendix 15 Usefulness Questionnaire

This Appendix gives an example of a questionnaire that may be used during quiet periods of operation to assess the operator's opinion of the usefulness of alarms which are displayed. Results may be analysed to produce a 'nuisance score' for the alarm system.

One way of measuring the usefulness of alarms is to ask operators to express a view on the value of each of them. This can be done with the questionnaire given on the next page, though it should be used only during in quiet periods of operation. It is seen that the operators are asked to write in the alarms that occurred over a period, and tick one of five columns which are labelled 'Action', 'Check', 'Noted', 'Little use' and 'Nuisance', respectively.

When filling in the form the operators are asked to fill in the time when they start the exercise and the time when they finish it and to identify each alarm. This may be useful if particular nuisance alarms identified on the forms are to be further investigated.

The form asks operators to enter repeating alarms only once. Repeating and fleeting alarms might, typically, represent half of the alarm load and the usefulness of these is worth identifying. However, statistical analysis of alarm logs is better for obtaining accurate estimates of the average percentage of fleeting and repeating alarms.

Once several forms have been completed, an overall 'nuisance score' may be calculated. This was done in the HSE survey (5) by applying a weighting to the alarms in each category according to a judgement of the 'spuriousness' of the alarm. These weightings are given in the second row of Table 28. It is seen that an alarm categorised as 'Action' gets a weighting of 0, and an alarm categorised as 'Nuisance' gets a weighting of 10. Thus if 100% of alarms for a site were in the 'Nuisance' category, the score for that site would be 10.

	Action	Check	Noted	Little use	Nuisance	Weighted score
Total alarms in each category	113	89	128	61	58	
Percentage alarms in each category	25%	20%	28%	14%	13%	
Weighting	0	1	3	6	10	3.18

Table 28 Example of weighting of results of usefulness questionnaires

Table 28 shows an example of the weightings being used to calculate the overall 'nuisance score'. The weightings and figures are those used for analysing data from the HSE-sponsored survey (5). The average nuisance score across 11 plants was 3.18. It is suggested that a score of under 2.0 represents a plant with alarms which are generally useful.

How many alarms are useful?

This form is being used to measure what proportion of alarms are found useful by operators.

As each alarm occurs we want you to write its title in the table below. This only needs to be a very brief abbreviation, e.g. "Boiler Press Hi".

We then want you to put a tick in one of the five columns depending on how useful you think the alarm is. An example is shown in the table.

Please do this for 10 consecutive alarms that occur during normal operation. Choose a time when you are not too busy so it will not distract you from running the plant.

The five column headings mean the following:

> **'Action'** Tick this column if you took a positive action like operating a valve, changing a control set point, phoning a plant attendant, etc.
>
> **'Check'** Tick this column if you made some check of the plant status, e.g. checked a measurement already on display, called up a new VDU graphic.
>
> **'Noted'** Tick this column if you did not take any action or make any plant check, but you were still glad that the alarm was displayed to you.
>
> **'Little use'** Tick this column if the alarm was of no real use to you.
>
> **'Nuisance'** Tick this column if the alarm was a complete waste of your time.

If you have already entered the alarm once in the table and it occurs again, please just put a tick in the repeat column.

Start time .... Date .... Plant/unit ....

Alarm Title	**Action**	**Check**	**Noted**	**Little use**	**Nuisance**	**Repeats**
Boiler Press Hi		✓				

Finish time ....

Appendix 16 The Costs of Poor Alarm Performance

This Appendix provides some illustrative examples of incidents where poor alarm system performance has contributed to financial loss, risk to people or environmental damage. This information might be used in a justification for investment in alarm systems and may increase awareness of the importance of alarm system failings.

A16.1 Introduction

Financial losses and accidents occur for a multitude of different reasons on process plant. Often, because designers have tried hard to prevent such incidents, the ones that do happen are due to combinations of several unexpected events occurring together. In addition, it is hard to collect data about all incidents involving financial loss or risk to people or the environment, especially the smaller incidents and the near misses. Because of these difficulties it is hard to make statements like "there have been 'x' incidents on process plants in the UK in the last year due to alarm system failures that have cost a total of 'y' million pounds". It is possible, however, to list some of the known incidents that have occurred involving alarm system failures and to quote estimates of the losses involved. Such information is collected together in this Appendix.

A16.2 Three Mile Island

The accident that occurred at the Three Mile Island nuclear power station in 1979 caused major damage to the plant that resulted in it being permanently shut down. The financial losses were about $1 billion. There were no fatalities or serious injuries as a direct result of the incident, but there was some minor release of radioactive material into the environment. There were various causes for the accident, but with hindsight it is clear that if the operators had fully appreciated what was going on, they would have been able to prevent the accident. There were a number of shortcomings in the operator interface. One of the problems was that the operators were loaded with numerous alarms, and that several key alarms were misleading (39).

A16.3 Milford Haven Refinery

The explosion and fires at the Texaco Refinery, Milford Haven in the UK in 1994 (23), (40) resulted in plant damage which cost £48 million to repair and loss of production which significantly affected the UK refining capacity. There were 26 minor injuries but, due to a fortuitous combination of circumstances, there were no serious injuries or fatalities. The plant owners were fined £200,000 plus costs. The accident was caused by equipment failure coupled with poor design of a modification. With hindsight, the plant operators had ample time to recognise and prevent the accident. They were hampered by a lack of good overview graphics on their VDU displays and by the fact that alarms were being presented at an estimated rate of one every 2-3 seconds in the 5 hours leading up to the accident. There were 275 alarms in the 10.7 minutes before the explosion.

A16.4 Channel Tunnel Fire

The Channel Tunnel fire that occurred in November 1996 resulted in nearly £200 million loss including the damage and losses in operating revenue. A number of passengers suffered shock and were affected by smoke. Alarm system problems were one factor that affected the management of the incident. Design concerns about false alarms had resulted in a philosophy of unconfirmed and confirmed fire alarms being adopted. A consequence of this was that the fire detection system reacted but did not give immediate warning of what was a significant developing fire. Furthermore, during the first minutes of the incident the Rail Control Centre operators were submerged in an overload of information and alarms. One of the recommendations of the inquiry into the accident was that Eurotunnel must develop and install an alarm management system (14).

A16.5 IChemE Accident Database

The Institution of Chemical Engineers (IChemE) maintains an accident database that includes records of accidents that have been made known to it on a confidential basis or which have been reported publicly. The IChemE also runs an international Safety and Loss Prevention Group.

A16.6 Nuclear Power Incident Databases

A number of owners of nuclear power plants contribute to various national and international databases which records incidents which have a potential safety implication. Incidents are categorised according to severity and the great majority of those recorded are of a minor nature. These databases can be large, containing records from plants all over the world; however, access to the databases is restricted.

A16.7 HSE Accident Analysis

The HSE has published general information on the cost of accidents at work (25). Whilst not specific to alarm systems, it does contain guidance of relevance.

A16.8 HSE Alarm Survey

The recent survey of alarm systems in the chemical and power industries carried out on behalf of the HSE (5) identified the following incidents involving alarm system failures at the 15 plants visited:

- 4 incidents resulting in plant damage, namely:
  - damage to a compressor costing £1 million, plus bringing forward of an outage costing £12 million;
  - destruction of a pump costing £250,000 to replace;
  - destruction of another pump costing £250,000 to replace;
  - £20,000 of plant damage due to fire plus a potential £250,000 production loss.
- 3 incidents resulting in production loss, namely:
  - an unplanned trip due to a missed alarm costing £250,000;
  - trips about once per day on a new plant, of which many could possibly have been prevented if there had been fewer nuisance alarms occurring;
  - a trip resulting in loss of 5 days production with a selling cost around £3 million.
- 3 incidents causing excessive environmental discharge, namely:
  - release of a gas resulting in serving of an Improvement Notice;
  - overflow of a vessel containing hazardous fluid;
  - three incidents at another plant involving overflow of vessels containing hazardous fluids.

A number of other incidents at plants not visited were also reported. It is noted that all the incidents identified in this survey were based on unstructured conversations with individual engineers rather than any systematic or rigorous study. Consequently it is likely that many incidents were not identified in the survey.

A16.9 Abnormal Situation Management Consortium

In 1999, a major project in the USA was undertaken by the Abnormal Situation Management Consortium involving a number of suppliers, major petrochemical companies and research organisations. Its objective was to improve plant performance in 'abnormal situations' (1). These situations encompass a range of events outside the 'normal' plant operating modes, e.g. trips, fires, explosions or toxic releases. The ASM Consortium has carried out a survey in the US petrochemical industry and estimated that there are losses of perhaps $10-20 billion per year from abnormal situations (1), (36), (48). This is approximately equal to the total annual profits of that industry. The following evidence led to this estimate:

- plant surveys showed that incidents were frequent with typical costs ranging from $100,000 to well in excess of $1 million per year. For example, one plant surveyed had 240 shut downs per year at a total cost of $8 million. Many of these shut downs were preventable;
- it was found that refineries on average suffer a major incident once every three years costing on average $80 million;
- one insurance company's statistics showed that the industry was claiming on average over $2.2 billion per year due to equipment damage. It is likely that actual total losses to the companies would be significantly higher than that claimable.

Personnel injuries or fatalities were also associated with some of the more serious incidents. Whilst alarm system failures were only implicated in a proportion of the above losses, the surveys did show that they were a major contributor, and the loss incidents frequently involved the operator being overloaded with alarm floods.

A16.10 Refinery Study

Much of the information given so far in this section relates to the major and more obvious loss incidents. However, poor performance of alarm systems and other operator support tools can result in smaller and less obvious losses. Experience suggests that these incidents are very much more frequent that the large incidents. They are reflected in numerous small deviations from optimum operating conditions. These deviations will be reflected in various key performance indicators (KPI) for the plant such as plant throughput, plant product quality, plant efficiency, etc.

These ideas may be illustrated by an example, and Figure 23 shows a typical plot of a key performance indicator over time. Poor performance results in lost opportunity costs, lost profit, unnecessary shut downs and equipment damage.

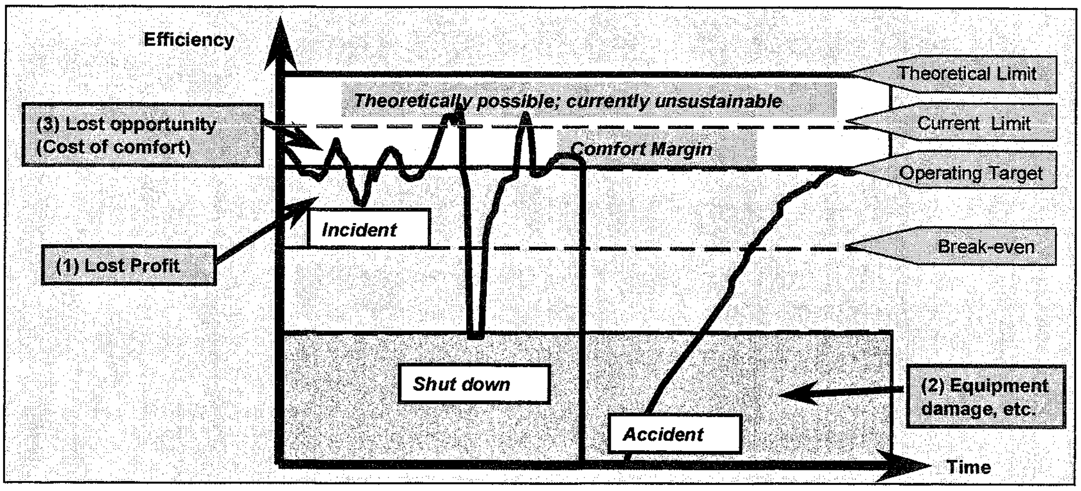

Figure 23 Typical plot of Key Performance Indicator

Plant data of this sort may be plotted in histogram form and this is illustrated in Figure 24. The ideal performance would be represented by a narrow histogram centred on the operating target, and the difference between this and the actual productivity histogram indicates the avoidable losses. The financial loss associated with this difference can be calculated.

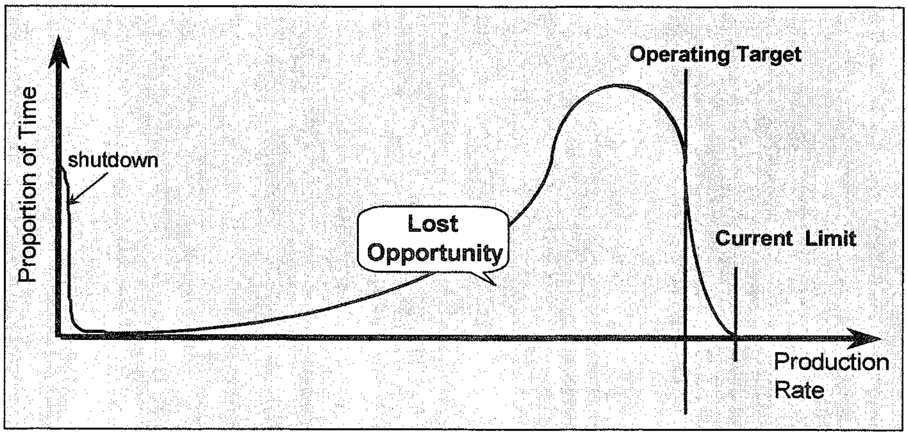

Figure 24 Productivity histogram

Using one year's data from three hydrocarbon processing plants, it has been estimated that small disturbances from optimal production account for 3-8% of plant throughput (11). For a typical oil refinery this equates to an annual cost of £3-10 million. Not all this loss will be recoverable just from installing better alarm systems, but some part of it should be. More powerful decision support tools that integrate better with information and planning systems should also help to reduce these losses.

One important point of note is that plotting a productivity histogram from actual plant data can give a good statistical estimate of the losses due to the relatively frequent incidents such as product quality deviations or unplanned shut downs. However, it cannot give a good statistical estimate of the losses due to events such as very large accidents as these will be particularly infrequent (see Figure 25). If alarm system investment is made on the basis of the small incidents alone then there will be under-investment. Furthermore, the total costs of these large accidents can be very high indeed, in the worst cases running into billions of pounds and involving multiple loss of life. It should be recognised that good alarms systems can play a significant part in reducing the likelihood of these rare accidents.

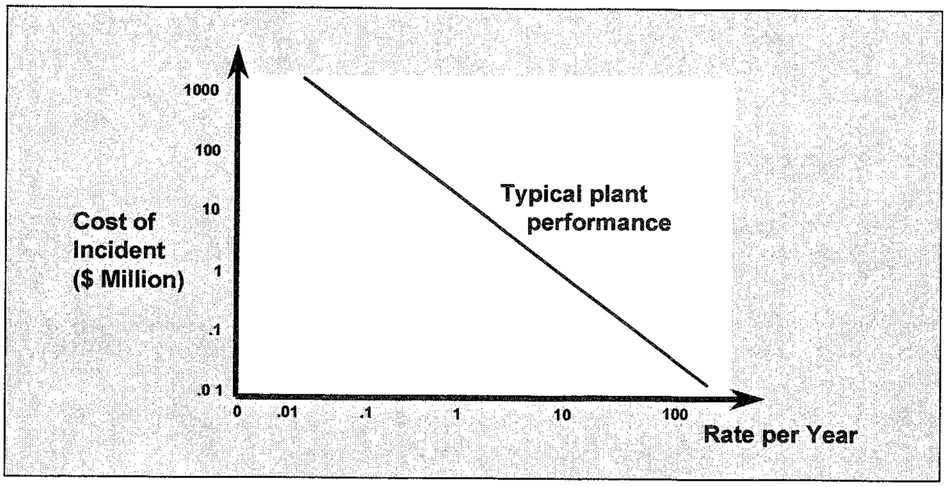

Figure 25 The frequency of loss incidents

Appendix 17 Specification Checklist - Large System

This Appendix collects together information from other sections of this Guide into a summary list of the features that should be included in a specification for a new large alarm system.

The checklist given in this Appendix provides information to assist in the writing of a specification for the procurement of a new alarm system. It can also be used for assessing the functionality of alarm systems offered by suppliers.

Requirements have been graded as follows:

- **Essential** - functionality that it is recommended should be provided. Much of this has been discussed in more detail elsewhere in this Guide. These requirements are written in **bold**;
- **Valuable** - functionality that is often found useful, but may not always be needed. These requirements are written in normal text and identified as 'valuable';
- **Possible** - options that some users may wish to consider. These requirements are written in *italic* and identified as 'possible'.

The checklist has been written on the assumption that a large system is being procured, e.g. in excess of 500 alarms. Some requirements may possibly be relaxed for smaller systems. The checklist also assumes use of programmable display devices and not everything will be relevant to an annunciator-based system.

It should be noted that the checklist does not cover qualitative features of an alarm system, e.g. how easy it is to configure the system or maintain it. These factors are very important in terms of the overall performance of the alarm system. However, they are highly system dependent and hence difficult to specify. It is recommended that the purchaser should request the supplier to provide information to allow an assessment to be made.

Inputs

The alarm system should be able to generate alarms from:

- digital inputs to the alarm processor (volt free or powered);
- analogue inputs to the alarm processor from transmitters (e.g. 4-20 mA, 0-1V, 0-5V, 0-10V), thermocouples, resistance thermometers, etc.;
- systems connected to the alarm processor over communication lines;
- automatic controls or other calculations made in the alarm processor system;
- faults in the alarm processor hardware or software.

(Valuable) There should be timing logic to detect fast repeating digital input alarms and, hence, suppress them.

(Valuable) There should be loop resistance checking on digital inputs.

Alarms from digital inputs should be time stamped to within 1 second generally, but may need to be within 100 millisecond or better for applications such as sequence of event logging when this is incorporated into the alarm logging system.

(Valuable) Digital input alarms should be time-tagged at source.

(Valuable) Alarms received over communications lines should be time-tagged at source.

Fast, sequence of event logging of particular alarms or events is required. (Comment: Typically for perhaps 10% of the inputs and maybe at 10-100 millisecond time resolution. For some special cases resolution down to 1 millisecond may be needed).

Alarms from analogue inputs should be time stamped to within 1 second generally, but may need to be within 100 millisecond for applications such as sequence of event logging.

The following alarm types should be available for configuration on analogue inputs:

- low and high;
- low-low and high-high;
- (Valuable) rate of change;
- (Valuable) deviation from a set point.

There should be an adjustable deadband associated with all alarms derived from analogue values.

Alarms should be generated if analogue signals go faulty.

(Valuable) Discrepancy alarms should be available which indicate the difference between the expected state of a system and its actual state (e.g. discrepancy on a control loop that has tripped from auto into manual, or discrepancy between expected actuator position and actual measured actuator position). Note that, as discussed in Appendix 6, discrepancy alarms should be designed to take account of expected dynamics and to be robust against slight changes from expected behaviour.

Processing of Alarms

(Valuable) All alarm settings, deadbands, messages, etc. should be stored in a system-wide database. This database should include information on when data items within it are changed.

(Valuable) Operators should be able to change the alarm settings on some alarms (i.e. those marked in the database as operator-adjustable).

(Valuable) Supervisors should be able to change the alarm settings on some alarms (i.e. those marked in the database as supervisor-adjustable).

(Possible) Facilities should be provided to the operators for defining their own alarms.

Alarms should be prioritised. (Comment: There should be 3 priorities within any one system for normal display of alarms, plus other priorities for alarms not normally displayed):

- it should not be possible for the operator to change priorities of any alarms;
- *(Possible) the alarm system should be capable of automatically changing the priority of alarms according to operating conditions.*

It should be possible to shelve individual alarms.

- (Valuable) this facility should be available to the operator;
- (Valuable) this facility should be available to the shift supervisor;
- the operator should be able to observe which alarms are shelved;
- (Valuable) facilities should be provided for the operators to document the reason for shelving alarms (Comment: It can be as important to know why an alarm is shelved, as it is to know that it is shelved);
- (Valuable) there should be an in-built software limit on the period for which alarms can be shelved;
- (Valuable) the operator should be prevented from shelving specified alarms;
- the occurrences of shelved alarms should be logged.

The following facilities for automatically suppressing alarms from appearing on the operator's display should be provided:

- suppression according to plant operating mode (e.g. shut down, starting up, full load);
- suppression according to the operating state of particular plant items (e.g. suppression of alarms from a pump which is out of service);
- *(Possible) suppression of alarms from plant under test;*
- (Valuable) suppression of normally expected alarms in a short period after a major event (e.g. a plant trip, a loss of electrical power);
- (Valuable) suppression of related alarms in cause-consequence groups;
- *(Possible) suppression of alarms using expert systems or other similar techniques;*
- the operator should be provided with facilities for observing alarms which have been automatically suppressed.

'Alarm coalescing' should be provided (e.g. the merging of multiple alarms from 2-out-of-3 voting systems to generate a single alarm). (Valuable) It should be possible to synthesise an alarm based on the states of various other alarms.

Display of Alarms

(Valuable) It should be possible for some alarms to be generated and displayed independently of the alarm processor.

(Possible) There is a requirement for alarms to be displayed on individual annunciators.

(Possible) There is a requirement for alarm annunciators driven from the alarm processor.

The following audible functions are required:

- a sounder on all (or only some) alarm processor alarms;
- different sounds for different priorities of alarms;
- *(Possible) different sounds for alarms and for discrepancies;*
- (Valuable) different sounds for alarms on annunciators and on screens;
- *(Possible) voice output of alarm messages (Comment: This is not yet widely established so must be used with great care).*

The alarm status should be indicated on objects (e.g. valves, measurements) drawn on schematic and control faceplate graphics.

It should be possible for operators to accept alarms from schematics and faceplates.

(Valuable) Plant objects displayed on schematics should be hierarchically connected in terms of alarms (e.g. a pump displayed on a schematic should go into alarm if one of the lower level alarms associated with it, such as "pump vibration high", goes into alarm).

An alarm list display should be provided:

- it should be possible to accept alarms from the alarm list (i.e. without referring to the detailed plant display);
- (Valuable) it should be possible to accept alarms on the alarm list display individually or by page by a single operator action;
- *(Possible) for certain alarms it should be necessary for the operator to accept the alarm when it is raised and also when it clears;*
- **the list should show alarm occurrences in strict chronological order according to their time tag** (Comment: Single line annunciation of repeating alarms may cause loss of order of the display);
- **it should be possible for the operator to select a historical log of alarms in strict chronological order** (Comment: This should go back certainly one shift. Many users think at least 1 week, some think even longer);
- (Valuable) it should be possible for the operator to select a filtered list which shows only alarms of a particular priority;
- (Valuable) it should be possible for the operator to select a filtered list which shows only alarms from particular plant areas;
- *(Possible) it should be possible for the operator to view a list of all standing alarms.*

(Valuable) Alarm messages relating to analogue limit transgressions should show the alarm setting.

(Possible) Alarm messages relating to analogue limit transgressions should be continuously updated to show the value of the analogue signal.

Alarms of different priorities should be displayed in different colours. (Comment: This is especially important on graphics. Colour coding is discussed in Sections 3.3.1 and 3.3.4.).

(Possible) Guidance on the actions to be taken after an alarm should be available on the alarm display.

(Possible) Information on the likely cause of the alarm should be available on the alarm display.

There should be a facility to enable the operator to bring up a detailed graphical display (e.g. a plant detail simulation) relating to the latest alarm with a single action.

(Possible) There should be facilities for causing important alarms to be re-annunciated after a certain delay time.

Logging of Alarms[43]

Alarms should be continuously logged to magnetic or optical disc storage.

The disc storage should be large enough to hold 1 year of alarm records. It should be possible to copy this into long term archive. (Comment: On large alarm systems it may too costly to store 1 year of records on disc and some alternative may have to be taken.)

(Possible) facilities should also be provided to continuously log alarms to paper printers.

(Valuable) Operator acceptances of alarms should be logged.

(Valuable) Every alarm occurrence should be logged even if it repeating at a high frequency. (Comment: If there is limited memory for storing the log, this requirement may have to be relaxed.)

Facilities should be provided for exporting alarm logs to off-line management information systems or to PCs.

The following facilities should be provided for analysis of alarm logs:

- analysis of total numbers of alarms in a given period;
- (Valuable) searches for/counts of occurrences of specific alarms in a given period;
- identification of the most frequent alarms in a given period;
- (Valuable) identification of repeating alarms;
- (Valuable) an entry against each alarm clearance in the log to indicate the time period for which it stood;
- (Valuable) capability for the operator to annotate alarm/event logs with observations to help in evaluating operation.

(Possible) There should be facilities to replay an alarm log through an operator display to be able to simulate what the operator saw.

[43] Facilities for logging alarms are often combined with facilities for logging plant events, e.g. control mode changes, plant status changes, operator actions, etc. Requirements for event logging are not included here.

Engineering of Alarm Systems

During alarm system design, formal documentation should be prepared giving the reasons for every alarm, the consequences of not responding to it, the time of response, the required operator action, etc. (see Appendix 2)

Risk assessments should be carried out to identify all alarms which are required to support the safety of the plant.

Alarm response procedures should be written for all safety related alarms.

(Possible) Alarm response procedures should be written for all alarms. (Comment: needed on some sites.)

Alarm settings etc. should be subject to formal change control during plant commissioning.

(Valuable) In plant commissioning there will be a formal requirement to measure the alarm rates during steady operation or after plant trips and meet defined targets.

Tools should be provided to audit the current alarm database and compare against a recorded approved database to produce an exception report and optionally, to automatically reset 'unauthorised' changes.

Appendix 18 Specification Checklist - Small System

This Appendix contains a checklist of functionality that might, typically, be expected to be found in a small dedicated alarm system driving annunciators.

A18.1 Where Dedicated Alarm Systems are Used

Appendix 17 provides a checklist of the functionality that might be specified for a large alarm system driving in excess of 500 alarms. There is often a need for smaller alarm systems, say, on small batch chemical plants or in local control rooms for utilities on large plant. In some cases the required alarm functionality can be provided as part of an integrated plant alarm system. However, in other cases it will be more cost-effective to install a dedicated annunciator-based system. This Appendix provides a checklist of the minimum functionality that should be sought in such a system. Because they generally contain fewer alarms, dedicated systems tend to have less problems than large systems. Nevertheless, the design principles given in this Guide should be followed to ensure that they are effective and usable systems.

A18.2 Input Handling

- **digital and analogue capability** - analogue inputs can be transmitters (e.g. 4-20 mA, 0-1V, 0-5V, 0-10V), thermocouples, resistance thermometers, etc. Digital inputs can be volt free or powered;
- **voltage isolation** - up to 5kV isolation between digital inputs, 2kV between analogue inputs;
- **sequence of event recording** - time stamping of events to 100 millisecond resolution should be provided, but many systems achieve time stamping to 1 millisecond;
- **first-up discrimination** - again 1 millisecond resolution between first and subsequent alarms is often achieved, though coarser resolution may often be acceptable;
- **programmable delays** - can be set on each input for eliminating repeating alarms;
- **individual input inhibit and group inhibit** - this can be useful for nuisance alarms and for maintenance of field devices;
- **modular build** - can be designed to suit single input to unlimited number of inputs.

A18.3 Alarm Processing

- **integral and remote logic** - can be used for managing annunciation applications;
- **alarm management** - typical features would be Boolean logic with group output and cascade functions for the more complex alarm management applications;
- **software programmable** - providing all alarm annunciation sequences as per ISA RP18-1 (31);
- **recipe-driven alarms** - which can be used in process start up conditions etc.

A18.4 Alarm Display/Output

- **multiple alarm annunciation groups** - can be used to drive several independent alarm annunciation panels, each with their own operator controls;
- **shelving features** - both automatic and manual shelving of nuisance and repeating inputs;
- **test functions** - for lamp and input status condition;
- **indications** - both filament type and light emitting diode type illumination;
- **programmable group outputs** - both hardware and software programmable outputs;
- **serial communications** - available using, as a minimum, Modbus RTU communications.

A18.5 Environmental

- **EMC capability** - suitable for the environment in which it will be installed, e.g. installation near 132kV and 400kV switchgear or near radio transmitters, including portable radios;
- **IP rating** - enclosure should have a rating appropriate to the environment in which it is installed to prevent dust or water ingress;
- **hazardous area mounting** - equipment and enclosures certified for specified applications.

Appendix 19 Alarm Suppression Hazard Study

Suppression of alarms involves a potential risk of depriving the operator of important possibly safety-related information. This Appendix gives an example of procedures for the review of alarm suppression proposals.

Section 5 has described how an alarm improvement exercise should be carried out on an existing alarm system. This should follow some structured process of reviewing alarms and eliminating or re-engineering those which are of low value. Significant improvement will be generally be achievable from simple things such as adjustment of settings or deadbands, but to fully optimise the value of all alarms, more sophisticated logical processing techniques (see Appendix 8) become necessary.

All modifications made in such an improvement programme should be carried out responsibly and carefully to ensure that the operator is not deprived of operationally important alarms. As the suppression techniques become more sophisticated the risks of making mistakes probably increase - as does the implications of these mistakes - and the need for control of modifications becomes more important. For example, operating mode or major event suppression will involve the suppression of large numbers of different alarms, so if it is incorrectly implemented, it could potentially deprive the operator of considerable numbers of meaningful alarms.

The procedures used by different companies to control modifications vary and a general discussion of these is outside the scope of this Guide. What is provided in this Appendix are some notes on the procedures followed by one company for the review of operating mode suppression modifications. These are provided as an example rather than as definite guidance. Some companies may need to follow more stringent procedures, e.g. including review by independent parties, others may find less restrictive procedures to be quite acceptable given the safety and commercial risks that they are involved with.

The alarm hazard study methodology developed by the company is intended to be applied when using software-based operating mode suppression of alarms. It follows procedures somewhat similar to those used in HAZOP reviews (15), (34), though in no way replaces those or any other safety or operability reviews.

The study should be carried out at a time when the design of the suppression is completely specified, i.e. when the software functional specification, the suppression algorithms and the alarms suppressed are all defined.

The core team to carry out the hazard study should comprise:

- chairman - with experience in the suppression methodology and if possible also familiar with the specific plant;
- site control room operator - with experience in operating the plant;
- site process or operations engineer - familiar with the design of the plant.

Additional team members may include:

- engineer with alarm processing system expertise;
- site control/instrument engineer;
- site safety engineer.

If not a member of the above team, the designer of the suppression rules should be available to provide advice. One team member should be appointed to record notes of the study meetings.

The main documents required for the study are:

- functional design specification for the alarm suppression system which includes details of the operating modes for suppression, the operating parameters selected to identify the modes and the lists of alarms to be suppressed;
- up-to-date piping and instrumentation diagrams (P&IDs) for the plant;
- cause and effect charts for the plant;
- HAZOP and any other control, safety and operability reviews[44] for the plant.

The steps in the study are:

Familiarisation

The suppression proposals should be presented to the team and studied to ensure that all team members understand them in detail. Team members should also familiarise themselves with the process and control of the plant from drawings and documentation.

Review of operating modes

Taking each operating mode in turn, the team should discuss the operating parameters selected to identify the mode. The team should agree if the proposed operating parameters uniquely identify the mode or if other operating conditions would also be identified. Consideration should be given to the 'default' mode (which, in most cases, would be expected to enable all alarms) and the team should discuss all the conditions under which that mode would be selected, e.g. normal operation of the plant, failure or false signals on the input to the alarm system, failure of the alarm suppression software itself.

Review of alarms to be suppressed

The team should systematically work though each mode and each alarm to be suppressed in that mode. For each alarm:

- the alarm should be identified on P&IDs and its basic purpose(s) be agreed, e.g. it is a high temperature alarm to warn that a product rundown temperature is getting too high and could cause a problem in the storage tank;
- the HAZOP and other control, safety and operability review reports for the plant should be checked for any specific requirements for this alarm;
- the team should consider the effect of the alarm being suppressed under the mode being studied. The chairman should lead the team by a structured series of questions as well as by allowing free-ranging team discussion of the

[44] A conventional HAZOP review may not be sufficient for a plant with complex interactive control systems. In this case the review must be specifically structured to consider the potential for multiple simultaneous faults. The review must ensure that the operator is fully familiar with the control system purpose and functionality, that he is able to adequately monitor what it is doing in all situations, and that safeguards exist to maintain process safety and operability in the event of a controller fault. Some companies have developed formal procedures for carrying out such control, safety and operability reviews.

impact of suppressing the alarm. The structured questions to the team should include:

- for the selected operating mode, is the loss of the agreed basic purpose of the alarm likely to create a hazard or lead to an operational difficulty?
- is the alarm used for a purpose other than the agreed basic purpose, i.e. is it used to infer a problem elsewhere, and, if so, does loss of the alarm for the inferred purpose create a potential hazard or operational difficulty?
- is there another alarm which will provide similar information (e.g. a pump stopped alarm and a pump discharge low flow alarm could, in many circumstances, provide the same information to the control operator) and, if so should one, other or both be suppressed?
- is there any other potential hazard or operability problem created by suppressing this alarm?
- will the suppression of this alarm be unacceptable if certain other alarms are not displayed to the operator, e.g. because they are shelved or disabled?
- if any potential hazards or operability problems are identified by the team a record should be made on the log sheet (see below) to identify the potential hazard or operability problem and to make a recommendation for change.

Reporting

A log sheet should be filled in for each alarm which should identify:

- the alarm identifier;
- the function of the alarm;
- the operating mode being considered;
- the implication on the plant if the alarm is suppressed in the mode considered;
- any additional function which is inferred from the alarm;
- any other alarm from which the function of the alarm being considered can be inferred;
- any potential hazard or operability problem identified;
- any recommendation or comment.

In addition, an overall report of the complete alarm hazard study should be written. This should include all the log sheets plus a brief report of the study. The recommendations and conclusions of the study team should be given including a statement whether, subject to satisfactory resolution of the recommendations contained in the report, the suppression scheme can be put into service safely.

Follow-up

The recommendations identified in the overall report should be followed-up by the plant management and implemented on the alarm system by appropriately skilled people prior to the suppression scheme being put into operation. This may involve some iterations and further meeting of the review team. Once the scheme is in operation it should be carefully tested, and this may involve identification of further changes to the suppression software.

Appendix 20 Alarm Management in Batch Plants

This section has been added to assist those working with batch operated production units. It identifies areas of alarm management specific to the batch and semi-batch environment that may need to be considered in the operation of the alarm system. The core principles as defined throughout this Guide are as appropriate to batch processes as they are to continuous processes, but the emphasis in some areas may be different:

- *batch plants normally do not suffer from large alarm floods and can usually manage plant upsets by holding a plant state;*
- *many alarms are generated directly from control logic (phases, sequences etc.). These are a form of intelligent alarm handling, they are only generated at a specific time, they can contain specific information and can direct the operator in the action they should take;*
- *many batch sequences proceed with a high level of operator interaction via 'operator prompts' (alerts) on the computer interface.*

Topics discussed include:

- *batch plant operation;*
- *operator information;*
- *alarm rates;*
- *application of alarm priorities;*
- *design of alarms.*

A20.1 Batch Plant operation

There are a number of differences in the way batch plants are operated as opposed to the way continuous processes are operated. Therefore there is a need to consider how information is presented and how the alarm system is designed for each operation.

There can be large variations in the way batch plants and individual processes are operated and are manned. Often the control rooms are not continuously populated, with operators performing manual activities as a regular part of the batch process.

Examples of variations can be:

- single product or multiple product;
- single path or multiple path;
- automatic and or manual control;
- sequence of operations;
- manual additions;
- manual confirmations;
- mobile operators moving in and out of control room;
- operators waiting for the control system to inform of the next step;
- multiple operators covering same plant area;
- mobile alarm information;
- instruction from control system.

Typical considerations in the development of the Alarm System for a batch plant may be:

- what happens when control room is unmanned?;
- how is information passed on and to whom?;
- can critical alarms be missed?;
- understanding failure modes of equipment and what alarms they raise;
- alarms need to be configured to take account of which section of the plant the process is running;
- how are batch events, such as end of phase X notified to operator if not through alarm system?

There can be a large variation in the way control equipment is managed and therefore how the alarm system is controlled:

- central control rooms;
- satellite control rooms;
- remote workstations.

Typical considerations include:

- who is in overall control;
- who accepts alarms;
- can alarms be accepted by one operator and therefore missed by another;
- segregation of duties – do operator only see alarms that are relevant to them and the activities they are undertaking.

A20.1.1 Alarm Thresholds

The alarm system may need to provide flexibility to allow:

- alarm limits setting dependant on the type of plant are very different;
- alarm limits change with recipe;
- alarm limits change with plant state;
- alarms enabled change with plant state;
- minor alarm floods from services;
- alarms directly from sequences.

Typical considerations include:

- can alarms be left active or inhibited when the batch is finished?;
- what happens to alarms when batch/phases are held or aborted?;
- are safety alarms still active even when product control is finished?;
- is there segregation of equipment alarm, product alarms and operator prompts?;
- testing of alarms before the first batch of each different product.

A20.1.2 Replacement of Batch with Semi-Batch or Continuous Processes

From an inherent safety viewpoint, semi-batch processes should be operated wherever possible in contrast to the all-in batch type of process, in order to reduce the potential for reaction runaway / over-temperature / over-pressure / containment loss etc. The HSE publication HSG143 on designing and operating safe chemical reaction processes (26) provides guidance.

A20.2 Operator Information

As part of batch operation, different types of information are output for the operator's attention, these are in the form of prompts, messages and alarms.

Typical operator prompts and messages may include:

- confirmation of operator actions;
- informing the operator of their next process action;
- confirmation of recipe information.

Messages are generated so that production records can be maintained and so that the operator knows what has been achieved. For pharmaceutical plants these messages are essential as they may form part of the GMP (Good Manufacturing Practice) batch record.

The difference between alarms and operator prompts:

- an **alarm** is an **unexpected** event that requires **timely operator action**;
- an operator **prompt** is an **expected event** that may require operator action but is **not necessarily time critical**.

An operator prompt is a request from the control system that the operator perform some process action that the system cannot, or that requires operator authority to perform. It is normally not time critical and, typically, requires a positive confirmation from the operator that the action has been completed. The initial operator prompt and the response of the operator should be saved to a log. This may be the event log.

Ideally operator prompts should be displayed on a separate display so that they can be differentiated from alarms; they should have a separate mechanism from the alarms for drawing the attention of the operator.

A20.3 Alarm Rates

When calculating the performance of the alarm system, allowance needs to be made for operator prompts and messages. The time taken for the operator to respond to an alarm may be increased because of these other activities. Reducing the number of operator prompts may be difficult because of the nature of the process that may require intense manual activity at certain times.

In a complex batch process, operator prompts may easily exceed current 'normal state' alarm targets (i.e. 1 per 10 minutes). There is a need to consider the definition of 'alarm'. Operator prompts in a batch sequence are a normal part of the process operations, they are not signalling an abnormal event.

Where clear differentiation between operator prompts and alarms has been achieved, the prompts should be excluded from any counts of alarm rates; they should be considered within the overall 'human factors' assessment as part of the operator base workload.

A20.4 Application of Alarm Priorities in Batch Processes

Given that alarm floods are not generally an issue, then a different approach may be taken when prioritising alarms. Rather than trying to apply the normal 5% : 15% : 80% ratio between high : medium : low, it may be worth considering prioritising alarms by their function, e.g. safety, equipment protection and operational alarms generated from sequences/ phase logic.

It may also be worth considering splitting these even further, by separating operational alarms to dedicated operator displays for each type of alarm and/or each section of the plant/operator's area. It is recognised that the special requirements for batch plants means that this is somewhat inconsistent with guidance given elsewhere in this document.

A20.5 Design of the Alarms

Figure 26 is typical of the structure of the control of a batch process, the diagram is based on the ISA S88.01 Batch Standard model. The diagram shows two structures, the physical structure, i.e. the plant equipment, and the procedural structure, i.e. control of materials.

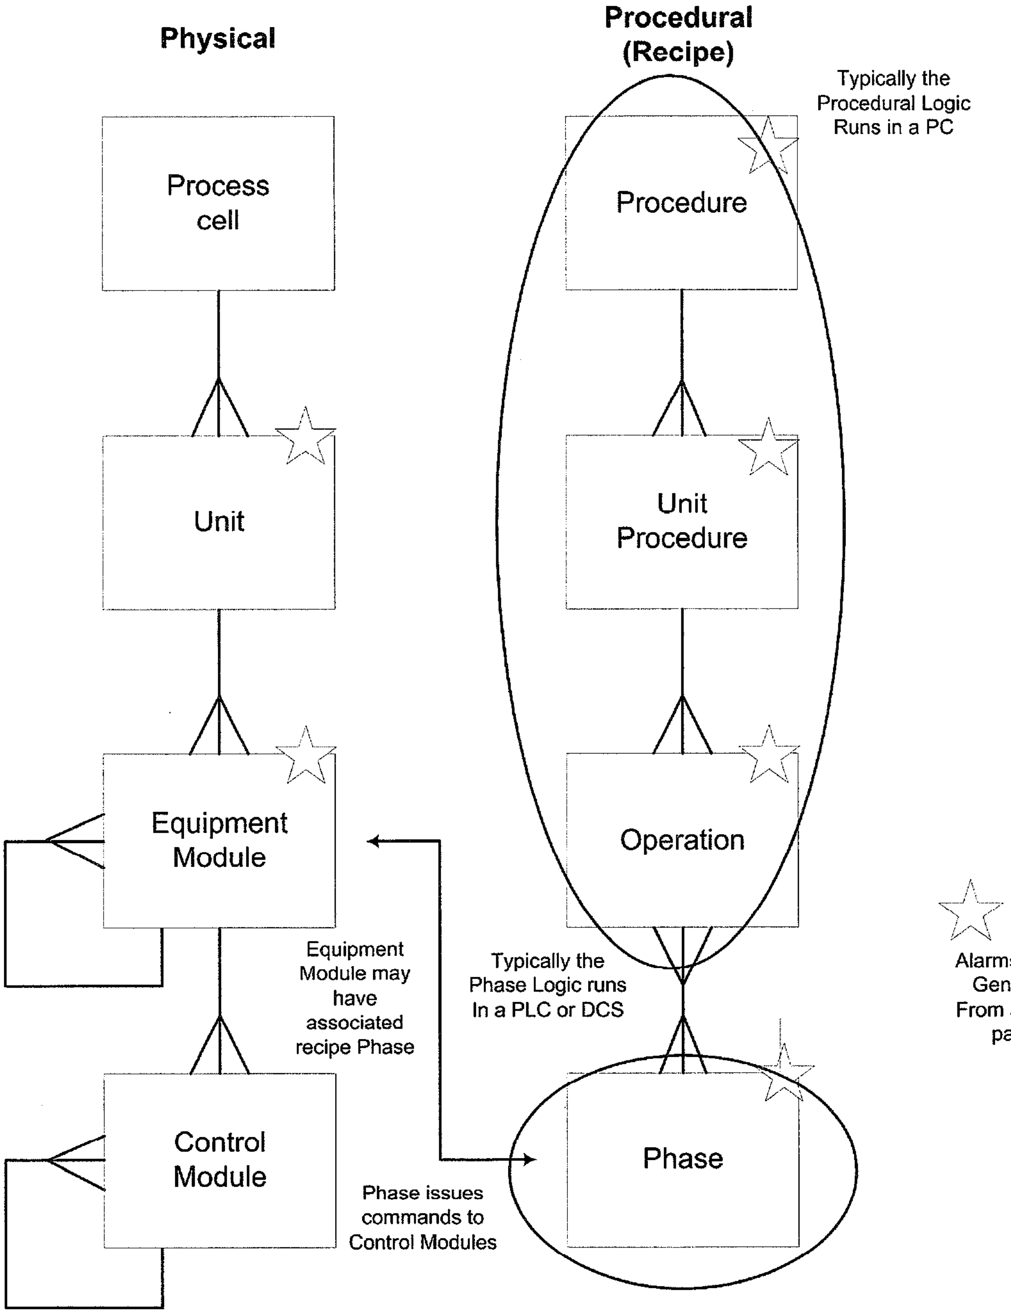
Physical
Procedural (Recipe)
Typically the Procedural Logic Runs in a PC
Process cell
Procedure
Unit
Unit Procedure
Equipment Module
Operation
Alarms can be Generated From all these palces
Equipment Module may have associated recipe Phase
Typically the Phase Logic runs In a PLC or DCS
Control Module
Phase
Phase issues commands to Control Modules

Figure 26 S88 Model

A20.6 Generation of Alarms

Alarm can be generated from a number of sources as described in Figure 27. Where alarms are initiated directly from the phase logic or procedure the alarm message can be more specific because logic can easily be applied. Alarms can be enabled and disabled depending on the state of the process, thus providing a form of 'intelligent' alarm handling.

It is important that standard information such as phase name and phase step is meaningful and easily interpreted by the operator.

Equipment alarms need to be continuously monitored regardless of the product under control.

Safety related alarms need to be separated from the procedural requirements of the product production.

A20.7 Management of Alarm Limits Thresholds

Where alarm threshold values vary depending upon the type of process or product, a formal robust procedure should be implemented to ensure that alarm threshold values are set appropriately for the product sequence(s) to be monitored and controlled. Sufficient safeguards to prevent inappropriate legacy values remaining operational within the alarm system from previous product runs should be put in place.

Extreme caution is urged with respect to inhibiting alarms as hazards may be generated post completion of the batch. Such conditions may include, e.g. when process heat is inadvertently applied to the final product post completion, thereby potentially leading to high temperature, overpressure and containment loss via pressure relief or even vessel or system failure etc.

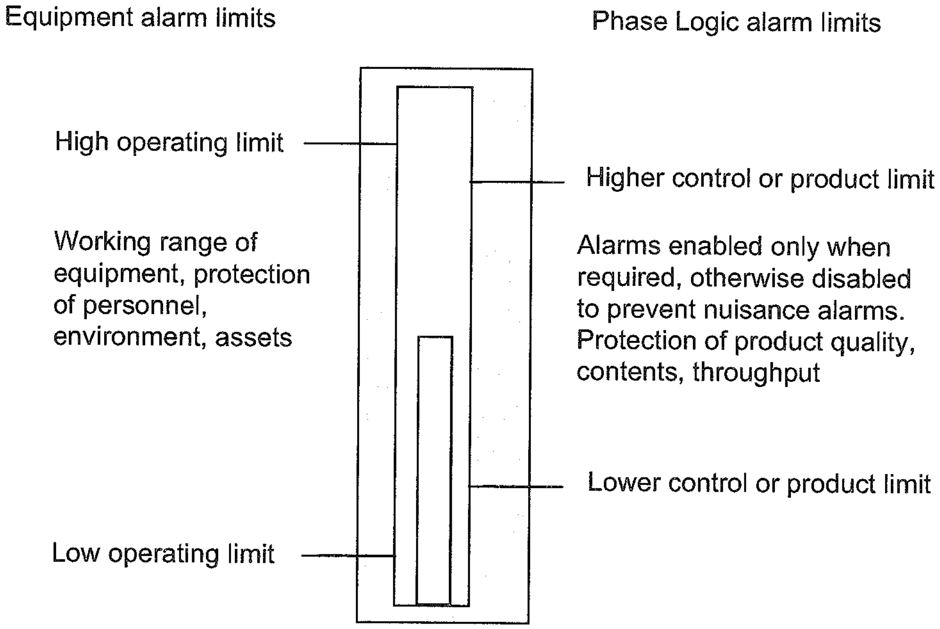
Equipment alarm limits
Phase Logic alarm limits
High operating limit
Higher control or product limit
Working range of equipment, protection of personnel, environment, assets
Alarms enabled only when required, otherwise disabled to prevent nuisance alarms. Protection of product quality, contents, throughput
Lower control or product limit
Low operating limit

Figure 27 Ideal – Separate set of alarm thresholds for control and equipment

Figure 27 describes an example of the split of alarm thresholds where a single input is used for equipment and procedural alarm generation. Good segregation is necessary so that there is time for operators to respond to safety related aspects, to ensure the equipment does not exceed its normal operating envelope, as well as maintaining the product to its specification.

As with continuous production processes, it may be difficult to maintain segregation between alarm thresholds as assets are operated near to limits.

A20.8 Sources of Nuisance Alarms

The generation and management of nuisance alarms are no different in batch activities from other types of processes, but there are at least two additional areas that are worth mentioning, they are:

- alarm enabled in an inappropriate phase – e.g. not turning off alarm on completion;
- product changes without consideration of alarm setting.

A20.9 Sources for Alarm Floods

Alarm floods are often less of an issue for batch production. There is less interconnection between vessels, disturbances tend not to propagate, and when a hazardous condition is identified it can usually be contained by shutting down the affected reactor. Most high alarm loads are generated when services and common items fail, such as air, nitrogen or other utility supplies.

Minor floods can occur when the control room is not continuously manned and the operator returns. Segregating alarms as described in Section A20.4 may help the operator manage these better.

Considerations include:

- review consequential alarms for each service;
- use alarm grouping to reduce number of alarms;
- review operator overload from a combination of alarms, prompts and messages;
- put alarms, prompts and messages into separate lists and selection keys.

A20.10 Getting Information to Remote Operators

A number of different tools are available that can be utilised to pass information to the operator while outside the control room, these include:

- **pagers** – carried by mobile operators, key alarms can be directed to the pager, allows all operators to be made aware of all critical alarms. Use for key alarms only;
- **remote alarm banners** – displays can be installed at key locations around production area, key alarms can be scrolled;
- **remote indication lights** – these could be simple indicator lights that identify when an alarm group is active. Simple but can be effective;
- **remote operator workstations** – located on plant rather than in the control room;

- **variable tone alarm klaxons** – audible sounds that direct the operator to view alarms on remote terminal or return to control room, different tones for different types of alarm, may include different coloured flashing lights for noisy areas;
- **fixed annunciator panels or fixed alarm displays** – allows operator to quickly view status of key alarm when returning to control room. Ideally suited for safety related alarms as they will always be in view regardless of the control room operation.

Appendix 21 Alarm System Improvement Process

Experience has shown that in order to progress with alarm systems improvements, a formal and methodical process is necessary if the best results are to be obtained.

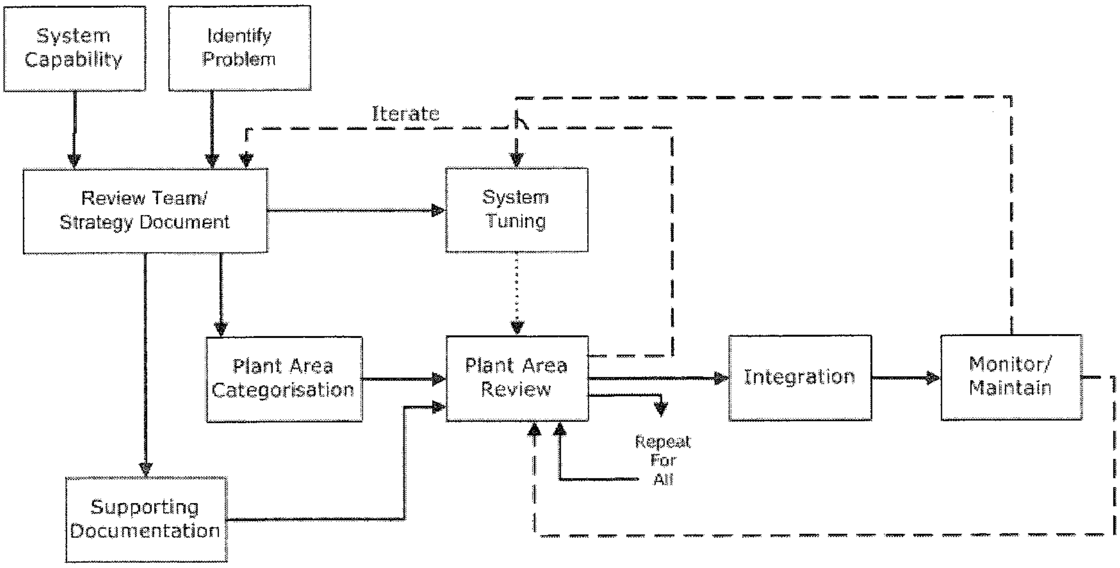

Figure 28 Alarm System Improvement Process

The overall process necessary to be achieved is illustrated in Figure 28. This sets out the major items that comprise an alarm systems review and their interlinking. Whilst the actual method of performing each of the steps may be specific to a particular site or company, the overall process is generic. Tools and techniques to assist with each step can be developed as pertinent to a particular site. Those described here are illustrative of the types of methods that can be applied to achieve the overall result.

It has to be recognised that a full alarm systems review is a resource intensive and time consuming process. A pragmatic approach can be to first conduct an alarm system tuning exercise. Whilst not giving the total advantages of a full review, this can often produce a substantial increase in performance for much less effort and is often a good way of starting to bring the alarm system into the manageable situation. A full system review can then be conducted. The techniques required for the tuning exercise still apply to an overall review and the process below includes a section on system tuning.

Pre-cursors

There are two pre-cursors to a full system review:

- to confirm and identify the alarm system problems;
- to understand the existing alarm system capability.

Problem Identification

A full alarm system review should only be undertaken if there is a fundamental problem with the alarm system performance. Whilst this is usually the case, it is wise to confirm this and determine what the major problems are. It would be unusual for a system with performance problems not to exhibit the common problems of:

- too high alarm rates in both normal and abnormal situations;
- the number of standing and/or shelved alarms;
- the incorrect priority of most alarms.

As noted in Section 4.2, an alarm analysis tool can be extremely helpful in exposing the limitations of a system.

A necessary step at this point is also to talk to the operators. They will undoubtedly be able to pinpoint the difficulties of using the system and of course any improvement is aimed at easing their task.

System Capability

Whilst not a reflection on the performance of the alarm system, it is necessary to understand the capability and functionality of the alarm system before undertaking a review, so as to be aware of what improvement techniques can be supported by the system.

This survey would include such items as types of alarms supported, ability to use logic, suppression techniques, whether shelving is supported, number of priorities supported and categorisation. The display capability would also need to be considered.

Review team / Strategy Document

Having determined that an alarm review is necessary, the first step is to set up a review team. This should be composed of a variety of individuals who can all bring their own individual knowledge to the proceedings. This will, typically, involve operations staff, C&I staff, process engineers and safety engineers.

The first task is to produce the Alarm Design Strategy Document (see Table 3). This is the guiding document for the whole process and ensures that the objectives and methods of the review are spelt out. This will allow for consistency in the review process.

An important part of this process is ensuring that guidance documentation is produced on a number of issues, so that when the review is undertaken the principles and checksheets are in place which will speed the process and also ensure the consistency of approach.

Supporting Documentation

This Guide, both in the main document and the Appendices, gives much information and guidance on issues such as apportioning hazard severity, operator response time, alarm frequency, prioritisation as well as the items that need considering for individual alarm design (Appendix 2).

These criteria have to be determined for individual sites before the alarm review can take place and should be used as supporting documentation when considering each individual alarm. It is also useful, as *aide memoires*, to prepare checklists of other relevant techniques that need to be considered as part of the review process. For instance, what type of alarms can be used, examples of logical alarm processing, message formats and types of improvement techniques. The site Alarm Management Strategy Document should also be produced.

System Tuning

Before undertaking a full alarm system review it can be useful to perform some system tuning, so as to make a quick improvement in day-to-day operation of the alarm system for substantially less effort. The techniques used and the methods employed will also be applicable to the full alarm system review. Often in a system tuning exercise, generic improvements which can be applied wholesale to large numbers of alarms can be identified – hence the opportunity to make a substantial improvement for a less resource intensive effort.

In performing a system tuning it is useful to have access to an alarm analysis tool which can help identify the current problem alarms. These can be dealt with on an individual basis (often a maintenance problem can cause a single alarm to be the current nuisance alarm) to the immediate benefit of the operations staff.

It is often useful in a system tuning exercise to focus attention on one plant area at a time, especially if this allows a large number of similar alarms to be looked at in the same manner (e.g. boiler metal temperatures).

The types of techniques to be used in a tuning exercise are those of ensuring that the basics of the existing alarms are correct, hence the focus on correct setpoints, deadbands and scan rates. When considering equipment maintenance regimes, care must also be taken to ensure that nothing is done that results in an under active alarm (e.g. moving a setpoint too far from a marginal setting). Similarly, the maintenance regime can help to uncover failed sensors.

This exercise can also consider whether groups of alarms need simple logic applied (e.g. timer delays on motor start alarms), whether there are many 'alarms' generated by the software that are irrelevant to the operator and whether many are duplicates (e.g. from both analogue and digital inputs).

Plant Area Categorisation

A full alarm system review process will have many alarms to consider. In order to make this manageable it is necessary to consider the whole system in smaller segments. For most plant the most obvious way of doing this is to categorise the alarms by plant area. Most plants will already have a set of plant areas; these should be reviewed to ensure that they are appropriate for the alarm system.

Each alarm should be assigned to a plant area category. The alarm review can then take place on a plant area category by plant area category basis.

Plant Area Review

It is usual to review all the existing alarms in a plant area category. An alternative approach, sometimes favoured when the plant is relatively old and has a number of enhancements, is to start with a blank piece of paper and determine what alarms should be assigned to that area of plant. In practice, the review is often approached with a combination of both methods, reviewing what exists and examining plant to ensure that the correct set of alarms is present after the review.

Each alarm is examined in turn, using the documentation checksheets/criteria, etc., previously prepared. The existing alarm data is compared against the strategy requirements. This is the crucial task of ensuring that only **true** alarms are included in the final alarm list. The 'data sheet' (see below) should clearly define why the alarm has been approved, what is its role and what is the appropriate operator action.

Alarm Data Sheet / Database

To facilitate the process, various computer orientated tools can be applied. This will depend upon how the existing alarm system is managed. One approach is to export the alarm list (data) into a spreadsheet or database (this will probably have already been done in order to create the plant area lists). A 'data sheet' should be created for each alarm (again, this can be as part of a spreadsheet or database), recording all the relevant information pertinent to that alarm that is agreed upon by the review process (again, Appendix 2 is relevant here).

A semi-automatic system can be created if required. For example, utilising a spreadsheet to calculate the priority based on criteria setting of other alarm properties.

Integration

The above processes will produce a base set of alarms with appropriate limits, deadbands, etc. to optimise the performance of the system for each individual alarm.

A further optimisation of the alarm system can also be applied by looking at the interaction of plant and situations. The two main areas to consider are:

- operating mode of plant;
- trips.

Considering the operating mode of plant and the relationship of alarms in this situation, can produce further rationalisation due to alarms not being relevant in a given mode or other alarms giving the required information

Trip situations can reveal alarm dependencies which can allow suppression of alarms, expected events in this situation which are therefore not needed to be alarmed, and a role reversal where it is the 'missing' event that needs to be alarmed.

Monitor / Maintain

Whilst not strictly part of the initial alarm system review process, an ongoing monitor / maintain role has to be instigated to ensure that the alarm system continues to benefit from the improvement process and does not degrade (see the Site Alarm Management Strategy (Table 4)). This includes ensuring that the alarm system is correctly identifying alarm situations, as well as preventing overload.

The two main tools here are to continue to monitor the performance using an alarm analysis tool, with a weekly top ten approach a useful mechanism, and to have a set of alarm system performance targets to measure the performance against.

The ongoing monitor/maintenance role will make use of both the tuning techniques approach and, as required, the deeper reviews as in the main body of the process.

Appendix 22 References

1. Andow, P., "Abnormal situation management: a major US programme to improve management of abnormal situations", IEE Colloquium on Stemming the Alarm Flood, London, 17 June 1997

2. Bliss, J. P., "The cry wolf phenomenon and its effect on alarm responses (false alarms)" PhD dissertation, University of Central Florida, USA, 1993

3. Bransby, M. L., "Advanced control desks for power stations", IBC Symposium on Human Factors in the Electricity Industries, London, 17/18 Oct 1995

4. Bransby, M. L., "Matching alarms systems to people", UKACC Conf. Control 98, Swansea, 1-4 Sept 1998

5. Bransby, M. L. and Jenkinson, J., "Survey of alarm systems in the chemical and power industries", HSE Research Report CRR 166, 1998

6. British Standards Institution, "Guide for the procurement of power station equipment. Control and Instrumentation", BS EN 45510-8-1:1998

7. British Standards Institution, "Basic and safety principles for man-machine interface, marking and identification. Coding principles for indicators and actuators", BS EN 60073:2002

8. Broadhead, N., "Optimisation of the Sizewell B alarm system", IEE Colloquium on Stemming the Alarm Flood, London, 17 June 1997

9. Burnell, E. & Dicken, C. R., "Handling of repeating alarms", IEE Colloquium on Stemming the Alarm Flood, London, 17 June 1997

10. Campbell Brown, D. & O'Donnell, M., "Too much of a good thing? Alarm management experiences in BP Oil", IEE Colloquium on Stemming the Alarm Flood, London, 17 June 1997

11. Campbell Brown, D., "Alarm design and performance - an industry perspective", IBC Seminar on Safe & Reliable Control Room Operations, London, 1 Dec 1997

12. Campbell Brown, D., "Horses for courses – A vision for alarm management", IBC seminar on Alarm Systems, London, Informa House, June 26/27, 2002

13. Carey, M. S., "Safety management of process faults: a position paper on human factors approaches for the design of operator interfaces to computer-based control systems", HSE Contract Research Report no 60, 1993

14. Channel Tunnel Safety Authority, "Inquiry into the fire on the Heavy Goods Vehicle Shuttle 7539 on 18 November 1996", 1997

15. Chemical Industries Association, "A Guide to Hazard and Operability Studies", London, 1997

16. Reising, D. V. C. & Montgomery, T., "Achieving effective alarm system performance: Results of ASM Consortium benchmarking against the EEMUA

Guide for Alarm Systems", 20th Annual CCPS International Conference, Atlanta, GA, 11-13 April 2005

17. Proceedings of the 20th Annual CCPS International Conference, Atlanta, GA, USA 11-13 April 2005

18. Dicken, C. R. & Marsland, C. R., "Improvements in alarm processing to reduce the effect of repeating alarms in an existing computer based system", IEE Colloquium on The application of computers to alarm detection, analysis and presentation, London, 25 Nov 1987

19. Electric Power Research Institute, "Human factors guide for nuclear power plant control room development", EPRI-NP-3659, Project 1637-1, August 1984

20. EEMUA Publication 201, "Process plant control desks utilising human-computer interfaces - A guide to design, operational and human interface issues", 2002, ISBN 0 85931 136 8

21. Hartley, J., "Designing instructional text", Kogan Page, New York, 1977

22. Hartley, J., "Eighty ways of improving instructional text", IEEE Trans. on Professional Communication, Vol PC-24, No 1, March 1981

23. Health & Safety Executive, "The explosion and fires at Texaco Refinery, Milford Haven, 24 July 1994", HSE, 1997

24. Health & Safety Executive, "Reducing error and influencing behaviour", HS(G)48, http://www.hse.gov.uk/humanfactors/comah/procedures.htm

25. Health & Safety Executive, "The cost of accidents at work", HS(G)96 2nd edition, ISBN 0 7176 1343 7, 1997

26. Health & Safety Executive, "Designing and operating safe chemical reaction processes", HS(G)143, ISBN 0 7176 1051 9, 2000

27. Health & Safety Executive, "Guidance on the Health and Safety (Safety Signs and Signals) Regulations 1996", Document L64, ISBN 0 7176 0870 0, 1996

28. Hopkins, Andrew, "Lessons from Longford: the Esso Gas Plant Explosion", CCH Australia Ltd, 2000, ISBN 1 8646 8422 4

29. International Electrotechnical Commission, IEC 61508: "Functional safety: Safety related systems Parts 1-7"

30. International Electrotechnical Commission, IEC 61511: "Functional Safety – Safety instrumented systems for the process industry sector Parts 1-3"

31. Instrument Society of America, "Annunciator sequences and specifications", ISA-RP 18.1, 1979 (rev. 1992)

32. Jenkinson, J., "Modern power station practice" Volume F, Chapter 6 "Central control rooms", Pergamon Press, 1991

33. Kirwan, B. & Ainsworth, L. K., "A guide to task analysis", Taylor and Francis, London, 1992

34. Kletz, T., "Hazop and Hazan - Identifying and assessing process industry hazards", 3rd ed., Institution of Chemical Engineers, Rugby, 1992

35. Mogford, J., "Fatal accident investigation report, isomerization unit explosion final report" – BP plc, Texas City USA, December 9th 2005

36. Nimmo, I., "Abnormal situations management - adequately address abnormal operations", American Institute of Chemical Engineers, Vol 91, No 9, 36-45, Sept 1995

37. Reason, J., "Human error", Cambridge University Press, 1990

38. Redmill, F. & Rajan, J., "Human factors in safety critical systems", Butterworth Heinemann, London, 1997

39. Rubinstein, E. & Mason, J. F., "An analysis of Three Mile Island. The accident that shouldn't have happened. The technical blow-by-blow." IEEE Spectrum, 33-42, November 1979

40. Powell-Price, M., "The explosion and fires at the Texaco Refinery, 24 July 1994", IChemE Loss Prevention Bulletin, 138, 3-10, 1997

41. Swain, A. D. & Guttermann, H. E., "Handbook of human reliability analysis with emphasis on nuclear power plant applications", NUREG/CR 1278, Albuquerque, NM, Sandia National Laboratories, 1983

42. Shepherd, A., et al., "Developing best operating procedures", SRD, Culcheth, 1991

43. U.S. Nuclear Regulatory Commission, "Human-system interface design review guideline", NUREG-0700 Rev 1, June 1996

44. U.S. Nuclear Regulatory Commission, "Guidelines on the preparation of emergency operating procedures", NUREG-0899, 1982

45. U.S. Nuclear Regulatory Commission, "Checklist for evaluating emergency operating procedures used in nuclear power plants", NUREG/CR-2005, 1983

46. U.S. Nuclear Regulatory Commission, "Techniques for preparing flowchart-format emergency operating procedures", NUREG/CR-5228, 1989

47. Williams, J. C., "A data-based method for assessing and reducing human error performance", IEEE Conference on Human Factors in Nuclear Power, Monterey, CA, 436-450, 6-9 June 1988

48. Windhorst, J., "Abnormal situation management and process safety", IPSG Meeting, Terneuzen, May 1998

49. Woodson, W. E., Tillman, B. & Tillman, P., "Human factors design handbook" McGraw-Hill, 1992

50. Wright, P., "Presenting technical information: a survey of research findings", Elsevier, Amsterdam, 1977

51. Zwaga, H. J. G. & Hoonhout, H. C. M., "Supervisory control behaviour and the implementation of alarms in process control", in "Human Factors in Alarm Design", N Stanton (ed.), Taylor & Francis, 1994

Appendix 23 Bibliography

1. BS EN/ISO 11064-1:2000 "Ergonomic design of control centres - Part 1: Principles for the design of control centres"

2. HSE Information Sheet, "Better Alarm Handling", Chemicals Sheet No. 6. http://www.hse.gov.uk/pubns/chis6.pdf

3. "Guide to the application of IEC 61511 to safety instrumented systems in the UK process industries" Draft for committee and industry review November 2006, EIC, EEMUA, UKOOA (draft document available from EEMUA)

EEMUA Publication: Feedback Form

An electronic version of this Form is available at **www.eemua.org/publications.htm**

Date feedback provided: (Dd/mm/yyyy)	EEMUA Publication: 191 2nd Edition Title: Alarm Systems: A Guide to Design, Management and Procurement
Feedback provided by: Contact name: Phone: E-mail:	Feedback provided on behalf of which organisation? Organisation name: Location:

EEMUA ref (EEMUA use only)	Page no.	Paragraph Figure/ Table concerned	Type of comment (General/ Technical/ Editorial)	COMMENTS	Proposed change (Please provide alternative wording, or propose removal of, or addition to, existing wording.)	OBSERVATIONS OF THE EXECUTIVE (EEMUA use only) on each comment submitted
	4	Fig 2	E	This is an example only.	No change to wording required.	Example only

When you have completed this form with all your comments, please save and e-mail to info@eemua.org or post a paper copy to: Publications Department, EEMUA, 10-12 Lovat Lane, London EC3R 8DN.

EEMUA PUBLICATIONS CATALOGUE

All EEMUA Publications can be purchased on-line. To order a publication contact EEMUA via the website at www.eemua.org, by e-mail, fax or phone.

ELECTRICAL

186 A Practitioner's Handbook - Electrical Installation, Inspection and Maintenance in Potentially Explosive Atmospheres

181 A Guide to Risk Based Assessments of In-situ Large Ex 'e' and Ex 'N' Machines

133 Specification for Underground Armoured Cable Protected against Solvent Penetration and Corrosive Attack

INSTRUMENTATION AND CONTROL

201 Process Plant Control Desks Utilising Human-Computer Interfaces – A Guide to Design, Operational and Human Interface Issues

191 Alarm Systems - A Guide to Design, Management and Procurement

189 A Guide to Fieldbus Application for the Process Industry

187 Analyser Systems - A Guide to Maintenance Management

178 A Design Guide for the Electrical Safety of Instruments, Instrument / Control Panels and Control Systems

175 Code of Practice for Calibration and Checking Process Analysers

155 Standard Test Method for Comparative Performance of Flammable Gas Detectors against Poisoning

138 Design and Installation of On-Line Analyser Systems

138 S1 Design and Installation of On-Line Analyser Systems. A Guide to Technical Enquiry and Bid Evaluation

MECHANICAL PLANT AND EQUIPMENT

204 Piping and the European Pressure Equipment Directive: Guidance for Plant Owners / Operators

203 Guide to the Application of ISO 3183 Parts 2 (1996) and 3 (1999) Petroleum and Natural Gas Industries - Steel Pipes for Pipelines - Technical Delivery Conditions

200 Guide to the Specification, Installation and Maintenance of Spring Supports for Piping

199 On-Line Leak Sealing of Piping - Guide to Safety Considerations

196	Valve Purchasers' Guide to the European Pressure Equipment Directive
192	Guide for the Procurement of Valves for Low Temperature (Non-cryogenic) Service
188	Guide for Establishing Operating Periods of Safety Valves
185	Guide for Hot Tapping on Piping and other Equipment
184	Guide to the Isolation of Pressure Relieving Devices
182	Specification for Integral Block and Bleed Valve Manifolds for Direct Connection to Pipework
179	A Working Guide for Carbon Steel Equipment in Wet H_2S Service
173	Specification for Production Testing of Valves - Part 4 Butterfly and Globe Valves
172	Specification for Production Testing of Valves - Part 3 Gate Valves
171	Specification for Production Testing of Valves - Part 2 Plug Valves
170	Specification for Production Testing of Valves - Part 1 Ball Valves
168	A Guide to the Pressure Testing of In-Service Pressurised Equipment
164	Seal-less Centrifugal Pumps: Class 1
153	EEMUA Supplement to ASME B31.3-1996 Edition, Process Piping
151	Liquid Ring Vacuum Pumps and Compressors
143	Recommendations for Tube End Welding: Tubular Heat Transfer Equipment, Part 1 - Ferrous Materials

OFFSHORE

197	Specification for the Fabrication of Non-Primary Structural Steelwork for Offshore Installations
194	Guidelines for Materials Selection and Corrosion Control for Subsea Oil and Gas Production Equipment
176	Specification for Structural Castings for Use Offshore
158	Construction Specification for Fixed Offshore Structures in the North Sea
146	90/10 Copper Nickel Alloy Piping for Offshore Applications - Specification: Fittings
145	90/10 Copper Nickel Alloy Piping for Offshore Applications - Specification: Flanges Composite and Solid
144	90/10 Copper Nickel Alloy Piping for Offshore Applications - Specification: Tubes Seamless and Welded

STORAGE TANKS AND VESSELS

190 Guide for the Design, Construction and Use of Mounded Horizontal Cylindrical Bulk Storage Vessels for Pressurised LPG at Ambient Temperatures

183 Guide for the Prevention of Bottom Leakage from Vertical Cylindrical Steel Storage Tanks

180 Guide for Designers and Users on Frangible Roof Joints for Fixed Roof Storage Tanks

159 Users' Guide to the Inspection, Maintenance and Repair of Above-ground Vertical Cylindrical Steel Storage Tanks

154 Guidance to Owners on Demolition of Vertical Cylindrical Steel Storage Tanks and Storage Spheres

147 Recommendations for the Design and Construction of Refrigerated Liquefied Gas Storage Tanks

NOISE

161 Guide to the Selection and Assessment of Silencers and Acoustic Enclosures

141 Guide to the Use of Noise Procedure Specification

140 Noise Procedure Specification

104 Noise: A Guide to Information required from Equipment Vendors

GENERAL

206 Risk Based Inspection: A Guide to Effective Use of the RBI Process

195 Compendium of EEMUA Information Sheets on Topics Related to Pressure Containing Equipment

193 Recommendations for the Training, Development and Competency Assessment of Inspection Personnel

149 Code of Practice for the Identification and Checking of Materials of Construction in Pressure Systems in Process Plants

148 Reliability Specification - Model clauses for inclusion in purchasing specifications for equipment items and packages

105 Factory Stairways, Ladders and Handrails (Including Access Platforms and Ramps)

101 Lifting Points - A Design Guide

STORAGE TANKS AND VESSELS

190 Guide for the Design, Construction and Use of Mounded Horizontal Cylindrical Bulk Storage Vessels for Pressurised LPG at Ambient Temperatures

183 Guide for the Prevention of Bottom Leakage from Vertical Cylindrical Steel Storage Tanks

180 Guide for Designers and Users on Frangible Roof Joints for Fixed Roof Storage Tanks

159 Users' Guide to the Inspection, Maintenance and Repair of Above-ground Vertical Cylindrical Steel Storage Tanks

154 Guidance to Owners on Demolition of Vertical Cylindrical Steel Storage Tanks and Storage Spheres

147 Recommendations for the Design and Construction of Refrigerated Liquefied Gas Storage Tanks

NOISE

161 Guide to the Selection and Assessment of Silencers and Acoustic Enclosures

141 Guide to the Use of Noise Procedure Specification

140 Noise Procedure Specification

104 Noise: A Guide to Information required from Equipment Vendors

GENERAL

206 Risk Based Inspection: A Guide to Effective Use of the RBI Process

195 Compendium of EEMUA Information Sheets on Topics Related to Pressure Containing Equipment

193 Recommendations for the Training, Development and Competency Assessment of Inspection Personnel

149 Code of Practice for the Identification and Checking of Materials of Construction in Pressure Systems in Process Plants

148 Reliability Specification - Model clauses for inclusion in purchasing specifications for equipment items and packages

105 Factory Stairways, Ladders and Handrails (Including Access Platforms and Ramps)

101 Lifting Points - A Design Guide