



Why Alarm Flooding Still Overwhelms Operators Despite Decades of Guidance

Jihoon Lim^{1*}, Takeshi Dlamini², Nare Harutyunyan³

¹Department of Engineering, Limpopo Institute of Engineering, Croatia

²Department of Engineering, Ural Center for Engineering, Philippines

³Department of Engineering, Achankovil School of Engineering, Japan

Article Info

Volume 1, Issue 1, 2026

doi: [10.62715/ijcepc.2026.v1.i1.pp27-31](https://doi.org/10.62715/ijcepc.2026.v1.i1.pp27-31)

Received: 18 October 2025

Accepted: 03 December 2025

Published: 03 February 2026

Abstract

Alarm floods, periods in which alarms arrive faster than a control room operator can read them, have been named as a contributing factor in major process accidents for more than forty years. Guidance on alarm management has been available for most of that time and is now codified in an international standard with a full lifecycle model. Yet the research literature and industry commentary continue to describe flooding as an unsolved problem. This perspective argues that the persistence is not mainly a failure of compliance. The guidance was shaped by the task of reducing nuisance alarms in normal operation, and its central method, the rationalisation of alarms one at a time against a fixed design condition, is poorly matched to a phenomenon that arises from the interaction of many individually justified alarms during a plant upset. Four further reasons are examined: alarms cost almost nothing to add and a great deal to remove, performance is judged on averages that hide the worst minutes, the engineering needed for state-dependent alarming is expensive and has to be kept current, and ownership of the alarm system is divided. Counter-arguments are considered. The article concludes that floods should be treated as a design case in their own right, defined by upset scenarios and tested as such.

Keywords: Alarm management, Alarm flood, Process safety, Human factors, Control room operators, IEC 62682

© 2026 Jihoon Lim et al. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

1. Introduction

The inquiry into the accident at Three Mile Island in 1979 recorded that more than a hundred alarms sounded in the control room in its first minutes and that the operators had no means of telling the important ones from the rest (Kemeny, 1979). Fifteen years later the investigation of the explosion at the Texaco refinery in Milford Haven reported that in the last eleven minutes before the explosion the two operators were presented with 275 alarms (HSE, 1997). In both cases the alarm system worked as designed. Each alarm was genuine. Together they were useless.

The response has been substantial. Research commissioned by the British safety regulator surveyed


* Corresponding author: Jihoon Lim, Department of Engineering, Limpopo Institute of Engineering, Croatia. E-mail: jihoon.lim@limpopo.edu

alarm practice in the chemical and power industries (Bransby and Jenkinson, 1998) and led to the EEMUA 191 guide, now in its third edition (EEMUA, 2013). The ISA-18.2 standard followed in the United States (ISA, 2016) and was taken up internationally as IEC 62682 (IEC, 2014). Books, vendor tools and consultancies have grown around them (Hollifield and Habibi, 2011; Rothenberg, 2009).

Even so, a survey of the research field as late as 2016 listed alarm flooding among its open problems (Wang *et al.*, 2016), and a review the following year reached a similar view (Goel *et al.*, 2017). This perspective asks why. Its argument is that the guidance, sound as it is, addresses flooding indirectly, and that the features of plants and organisations that generate floods are left largely untouched by it.

2. What the Guidance Asks For

The documents share a core. An alarm is an indication to the operator of an abnormal condition that requires a timely response. If no response is required, the event is not an alarm and belongs in a log. From that definition follows rationalisation: each configured alarm is examined, justified by a consequence and an operator action, given a priority and a setpoint, and recorded in a master alarm database. ISA-18.2 places this within a lifecycle that starts with an alarm philosophy and runs through identification, rationalisation, design, implementation, operation, maintenance, monitoring, management of change and audit (ISA, 2016). Figure 1 sets the main documents beside the accidents that prompted them.



The guidance also sets performance targets. EEMUA 191 suggests that a long-term average of fewer than one alarm in ten minutes is very likely to be acceptable to an operator, and that after a major plant upset fewer than ten alarms should appear in the first ten minutes (EEMUA, 2013). ISA-18.2 treats ten or more alarms in ten minutes at one operator position as the usual threshold of a flood. These numbers are reference points for assessment. They are not derived from a model of the operator, and the guides say as much.

3. Why Floods Persist

Five reasons are proposed here. They are drawn from the logic of the guidance and from the human factors and reliability literature, and are summarised in Table 1.

3.1. A flood is a property of the system

Rationalisation asks whether an alarm is justified. In a flood, most of the alarms are. A compressor trip in a tightly coupled plant produces real low flows, real high pressures and real level deviations in the units connected to it, each of which would merit an alarm if it occurred alone. Perrow (1984) described exactly this propagation as the signature of tightly coupled, interactively complex systems. A method that inspects alarms one at a time, against the steady state for which the setpoints were chosen, cannot see the flood, because the flood exists only in the combination. Woods (1995) made the cognitive side of the argument: an alarm is useful to the extent that it directs attention to what matters now, and that depends

on everything else competing for attention at the same moment.

Table 1: Reasons proposed for the persistence of alarm floods, and what each implies

Reason	Why current practice does not remove it	What would address it
Floods are a system property	Rationalisation judges alarms singly, at normal operation	Rationalise by upset scenario as well as by tag
Asymmetric cost of alarms	Adding an alarm is free; removing one needs justification	A budget of alarms per operator and per scenario
Averages hide the peaks	Monthly rates improve while worst minutes do not	Report flood count, duration and peak rate first
State-based logic is costly	Suppression rules need engineering and upkeep	Fund it as part of the safety case
Divided ownership	Operations, control and safety each hold a part	One accountable owner with authority over changes

3.2. Alarms are cheap to add and costly to remove

On a panel of hardwired annunciators every alarm occupied a window and had to earn it. In a distributed control system an alarm is a configuration entry, and several are available on every tag by default. The alarm management texts trace the growth in alarm counts to this change (Hollifield and Habibi, 2011). The asymmetry is reinforced by the way incidents are investigated. A hazard study or an incident review that recommends an additional alarm incurs no objection. A proposal to delete one invites the question of who will answer if it is later missed. Reason (1997) observed that defences added after each event can themselves become a source of opacity and risk. Alarm systems grow by exactly that mechanism, and a lifecycle standard slows the growth without reversing the incentive.

3.3. Performance is judged on averages

The metric most often reported is the average alarm rate, and sites that follow the guidance commonly bring it down by eliminating chattering and standing alarms with deadbands, delay timers and repairs to instruments (Izadi et al., 2009). This is worth doing. It improves the working day of the operator. It does little for the worst ten minutes of the year, which contribute almost nothing to an average. The standards do list flood measures, and graphical tools exist to display them (Kondaveeti et al., 2012), but an organisation that sets its target on the average will meet the target and keep its floods.

3.4. State-dependent alarming is expensive to engineer

The technical remedies for floods are known: alarms whose settings or enabling depend on the operating state of the plant, suppression of consequential alarms when their cause is already annunciated, first-out indication, and grouping. Each requires someone to work out, for each upset, which alarms are consequences of which causes, to encode that logic, to test it, and to revise it whenever the plant changes. A suppression rule that is wrong hides an alarm that was needed, so the engineering has to be done to the standard of a safety function. Research on data-driven methods, such as correlation analysis and the matching of flood sequences to earlier ones, aims to reduce this effort (Wang et al., 2016), but such methods learn from past floods and are weakest on the upset that has not happened before.

3.5. Ownership is divided

Alarm settings are touched by process engineers, control engineers, safety specialists, operations and maintenance, and by projects that come and go. The lifecycle model assumes an owner who enforces the philosophy through management of change. Where that role is weak or part-time, rationalisation becomes a project with an end date, and the system drifts back afterwards. Bainbridge (1983) pointed out long ago that automation leaves the operator with the tasks the designer could not automate, usually the hardest ones, and with fewer chances to practise them. A flood is the moment at which that irony is felt, and it is

seldom anyone's assigned job to design for it.

4. Counter-Arguments

The most direct objection is that the guidance works where it is applied and that persisting floods only show incomplete adoption. There is truth in this. The standards do provide for advanced alarming, the practitioner texts describe how to carry it out, and a document cannot be blamed for work not done. But the objection concedes the point at issue. If the part of the guidance that deals with floods is the part most often left undone, the reasons for that, cost, upkeep and ownership, are the real explanation and deserve to be addressed as such.

A second objection is that floods are unavoidable in a tightly coupled plant and that the proper response lies elsewhere: in automatic protection that does not rely on the operator, and in training. Automatic protection is indeed the right defence for the fastest events. Yet the operator remains responsible for diagnosis and recovery, and the theory of situation awareness holds that perception of the relevant elements comes before comprehension and projection (Endsley, 1995). An interface that defeats perception at the critical moment undermines everything built on it. The human factors literature on alarm design has said so consistently (Stanton, 1994).

A third objection is that better displays and analytics will make the problem disappear. They may help, but they share the dependence on engineering effort and maintenance described above, and they add a layer whose own failure modes have to be understood.

5. Implications

If the argument is right, three changes of emphasis follow. First, the flood should be a design case. For each operator position, a short list of credible upsets, such as loss of a utility, trip of a major machine or failure of a control system component, should be defined, and the alarm sequence each produces should be predicted, counted and reviewed, in the way that relief loads are reviewed for the same scenarios. Dynamic simulators used for operator training make this practicable.

Second, measurement should lead with the peaks. A site report that opens with the number of floods, their duration and their peak rate, and gives the average afterwards, directs effort differently from one that does the reverse. Third, the asymmetry of cost should be countered deliberately, by giving the alarm system an accountable owner and something like a budget, so that an addition during an upset scenario has to be traded against what is already there. None of this contradicts the existing standards. It asks that their least practised parts be treated as the centre of the subject.

6. Conclusion

Alarm flooding has outlasted decades of guidance because the guidance was built around the individual alarm and the ordinary day, while the flood belongs to the plant as a whole and to its worst minutes. The remedies exist but are costly to engineer and to maintain, the common metrics do not reward them, and no one clearly owns the result. Treating the upset as a design case, reporting peaks before averages and fixing accountability would align practice with what the accident reports have been saying since 1979.

Conflict of interest

To be declared by the author(s) before publication.

Funding

To be declared by the author(s) before publication.

References

- Bainbridge, L. (1983). *Ironies of automation*. *Automatica*, 19(6): 775-779.
[https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- Bransby, M.L. and Jenkinson, J. (1998). *The Management of Alarm Systems*. HSE Contract Research

- Report 166/1998. HSE Books, Sudbury.
- EEMUA (2013). *Alarm Systems: A Guide to Design, Management and Procurement*. EEMUA Publication 191, 3rd edition. Engineering Equipment and Materials Users Association, London.
- Endsley, M.R. (1995). *Toward a theory of situation awareness in dynamic systems*. *Human Factors*, 37(1): 32-64. <https://doi.org/10.1518/001872095779049543>
- Goel, P., Datta, A. and Mannan, M.S. (2017). *Industrial alarm systems: Challenges and opportunities*. *Journal of Loss Prevention in the Process Industries*, 50: 23-36.
- Hollifield, B.R. and Habibi, E. (2011). *Alarm Management: A Comprehensive Guide*, 2nd edition. International Society of Automation, Research Triangle Park, NC.
- HSE (1997). *The Explosion and Fires at the Texaco Refinery, Milford Haven, 24 July 1994*. Health and Safety Executive, HSE Books, Sudbury.
- IEC (2014). *IEC 62682:2014 Management of alarm systems for the process industries*. International Electrotechnical Commission, Geneva.
- ISA (2016). *ANSI/ISA-18.2-2016 Management of Alarm Systems for the Process Industries*. International Society of Automation, Research Triangle Park, NC.
- Izadi, I., Shah, S.L., Shook, D.S. and Chen, T. (2009). *An introduction to alarm analysis and design*. *Proceedings of the 7th IFAC Symposium on Fault Detection, Supervision and Safety of Technical Processes*, Barcelona.
- Kemeny, J.G. (chairman) (1979). *Report of the President's Commission on the Accident at Three Mile Island. The Need for Change: The Legacy of TMI*. President's Commission on the Accident at Three Mile Island, Washington, DC.
- Kondaveeti, S.R., Izadi, I., Shah, S.L., Black, T. and Chen, T. (2012). *Graphical tools for routine assessment of industrial alarm systems*. *Computers & Chemical Engineering*, 46: 39-47.
- Perrow, C. (1984). *Normal Accidents: Living with High-Risk Technologies*. Basic Books, New York.
- Reason, J. (1997). *Managing the Risks of Organizational Accidents*. Ashgate, Aldershot.
- Rothenberg, D.H. (2009). *Alarm Management for Process Control*. Momentum Press, New York.
- Stanton, N.A. (ed.) (1994). *Human Factors in Alarm Design*. Taylor & Francis, London.
- Wang, J., Yang, F., Chen, T. and Shah, S.L. (2016). *An overview of industrial alarm systems: Main causes for alarm overloading, research status, and open problems*. *IEEE Transactions on Automation Science and Engineering*, 13(2): 1045-1061. <https://doi.org/10.1109/TASE.2015.2464234>
- Woods, D.D. (1995). *The alarm problem and directed attention in dynamic fault management*. *Ergonomics*, 38(11): 2371-2393. <https://doi.org/10.1080/00140139508925274>