Control Engineering Practice 67 (2017) 21–30

Contents lists available at ScienceDirect

Control Engineering Practice

journal homepage: www.elsevier.com/locate/conengprac

Practices of detecting and removing nuisance alarms for alarm overloading in thermal power plants[☆]

Jiandong Wang[a,*], Zijiang Yang[a], Kuang Chen[b], Donghua Zhou[a]

[a] College of Electrical Engineering and Automation, Shandong University of Science and Technology, Qingdao, Shandong Province, China
[b] College of Engineering, Peking University, Beijing, China

ARTICLE INFO

Keywords:
Industrial alarm systems
Alarm overloading
Nuisance alarms
Delay timers
Power plants

ABSTRACT

Alarm overloading refers to the most noticeable phenomenon in existing alarm systems: there are far too many alarms to be promptly handled by industrial plant operators. A large number of occurred alarms are nuisance alarms that are not associated with any actual abnormalities and are extremely detrimental to important roles of industrial alarm systems. This paper presents long-term industrial applications of three techniques on detecting and removing nuisance alarms to a thermal power generation unit. By deploying these techniques, the severity of alarm overloading phenomenon has been significantly alleviated. The average number of alarm occurrences per day has been reduced from 18,280 to 359 in the year of 2015, so that about 98% alarm occurrences have been removed without affecting true alarms.

1. Introduction

The industrial standard ISA-18.2 (2009) states “an alarm system is the collection of hardware and software that detects an alarm state, communicates the indication of that state to operators, and records changes in the alarm state”. Alarm systems are indispensable parts of modern computerized monitoring systems such as distributed control systems, and play a critically important role for safety and efficiency of industrial plants such as oil refineries, petrochemical facilities and power plants (Bransby & Jenkinson, 1998; Hollifield & Habibi, 2010; Rothenberg, 2009).

Most existing alarm systems suffer from poor performance with the most noticeable phenomenon of alarm overloading, namely, there are far too many alarms to be promptly handled by industrial plant operators (Wang, Yang, Chen, & Shah, 2016). For instance, Table 1 lists the benchmarks of three basic performance metrics of alarms systems in the guideline EEMUA-191 (2013) and the counterparts from surveys on 39 industrial plants (Rothenberg, 2009). Clearly, the alarm overloading phenomenon is omnipresent in various industries. A large number of occurred alarms are the nuisance ones that are not associated with any actual abnormalities. As a result, nuisance alarms provide no useful information and are extremely detrimental to the important roles of alarm systems. A true alarm may be overlooked by operators due to distractions from nuisance alarms, and may be distrusted by operators due to “cry wolf” effects of nuisance alarms. Therefore, the very first objective for an intelligent alarm system is to reduce the number of alarms presented to operators (Kirschen & Wollenberg, 1992).

Some techniques have been proposed for nuisance alarms in recent years. Bransby and Jenkinson (1998), Hollifield and Habibi (2010) and Rothenberg (2009) presented some common methods to handle chattering alarm and repeating alarms, including filters, deadbands, delay timers and shelving mechanisms. Ahnlund, Bergquist, and Spaanenburg (2003) developed filters to reduce the number of nuisance alarms for different classes of process variables. Xu, Wang, Izadi, and Chen (2012) proposed optimal design approaches for alarm thresholds and delay timers to remove false and missed alarms. Naghoosi, Izadi, and Chen (2011), Kondaveeti, Izadi, Shah, Black, and Chen (2012) and Kondaveeti, Izadi, Shah, Shook, Kadali, and Chen (2013) identified chattering alarms and designed alarm limits and deadbands based on a chattering index related to alarm run lengths. Gupta, Giridhar, Venkatasubramanian, and Reklaitis (2013) integrated wavelet analysis, principal component analysis and qualitative trend analysis to rationalize alarm thresholds for drug manufacturing. Wang and Chen (2014) formulated online methods to remove chattering and repeating alarms by adjusting alarm thresholds and using delay timers. Zhu, Shu, Zhao, and Yang (2014) designed dynamic alarm thresholds depending upon

☆ This work was supported by the National Natural Science Foundation of China under Grant No. 61433001 and the Research Fund for the Taishan Scholar Project of Shandong Province of China.

* Corresponding author.
E-mail addresses: jiandong@sdust.edu.cn (J. Wang), yangzj@sdust.edu.cn (Z. Yang), chenkuang@pku.edu.cn (K. Chen), zdh@mail.tsinghua.edu.cn (D. Zhou).

http://dx.doi.org/10.1016/j.conengprac.2017.07.003
Received 22 February 2017; Received in revised form 1 July 2017; Accepted 11 July 2017
0967-0661/

Table 1
Cross-industry study (Rothenberg, 2009).

	EEMUA	Oil–Gas	PetroChem	Power
Average alarms/day	144	1200	1500	2000
Peak alarms/10 min	10	220	180	350
Average alarms/10 min	1	6	9	8

multiple steady states and their transitions in order to avoid nuisance alarms from using constant alarm thresholds. O'Donoghue, Phillips, and Nicell (2015) carried out alarm duration analysis and prevented an average 26816 nuisance alarms per week by applying 5-minute alarm delay timers for water industries. Han, Gao, Xu, and Zhu (2016) optimized alarm thresholds for multivariate alarm systems based on joint probability densities of process variables to minimize probabilities of false and missed alarms. Xu, Li, Song, Wen, and Xu (2016) gave fuzzy alarm thresholds based on evidence theory to reduce numbers of false and missed alarms caused by uncertainties of process variables. Soares, Pinto, and Souza (2016) applied alarm prioritization and correlation analysis techniques to design and reduce the size of alarm sets for natural gas processing plants.

The main contribution of this paper is to present long-term industrial applications of three techniques to a thermal power generation unit, in order to reduce the number of nuisance alarms and alleviate the severity of alarm overloading. The techniques are mainly based on our previous studies in Xu et al. (2012) and Wang and Chen (2014), and have been revised according to industrial practices for a large number of alarm variables. All the three techniques are composed by two steps: first, nuisance alarms in historical data sets are detected by exploiting specific characteristics of nuisance alarms; second, delay timers are designed to reduce the number of nuisance alarms in the future, by exploiting statistical information of alarm durations, alarm intervals or corresponding process variables. The power generation unit involves more than 2000 alarm variables, and the current alarm system suffers from a severe phenomenon of alarm overloading. By deploying these techniques, the average number of alarm occurrences per day has been reduced from 18 280 to 359 in the year of 2015, i.e., about 98% alarm occurrences of nuisance alarms have been removed. Hence, the severity of the alarm overloading phenomenon has been greatly alleviated.

The rest of the paper is organized as follows. Section 2 introduces the main techniques to detect and remove nuisance alarms. Section 3 presents the detailed results of applying these techniques to the power generation unit. Some concluding remarks are given in Section 4.

2. Main techniques

This section introduces the basic information of alarm variables and main techniques to detect and remove nuisance alarms.

2.1. Basics of alarm variables

Let $x_a(t)$ represent the value of an alarm variable x_a at the time instant th, where t is the integer-valued sample index and h is the real-valued sampling period. In the sequel, $h = 1$ s is assumed to ease notations. The alarm variable x_a usually takes the value of '1' ('0') for the alarm (non-alarm) state. The very basic alarm generation mechanism is to set x_a into an alarm state when a process variable x overpasses a constant high (low) alarm threshold x_{tp}, i.e.,

$$x_a(t) = \begin{cases} 1, & \text{if } x(t) \geq (\leq)\ x_{tp} \\ 0, & \text{if } x(t) < (>)\ x_{tp}. \end{cases} \quad (1)$$

The alarm occurrence (alarm clearance) is the event that x_a switches from '0' to '1' (from '1' to '0'). Three metrics are involved in describing alarm occurrences and clearances, namely, the alarm duration, alarm interval and alarm run length. The alarm duration T_1 is the time duration of one alarm occurrence, i.e.,

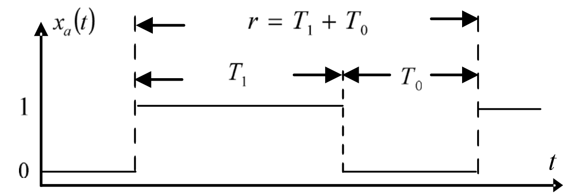

Fig. 1. Definitions of T_1, T_0 and r.

$$T_1 := t_2 - t_1 + 1, \quad (2)$$

where

$$x_a(t_1 - 1) = 0, \qquad x_a(t_2 + 1) = 0, \qquad \sum_{t=t_1}^{t_2} x_a(t) = t_2 - t_1 + 1, \quad \text{for } t_2 > t_1.$$

The alarm interval T_0 is the time interval from an alarm clearance to the next alarm occurrence, i.e.,

$$T_0 := t_2 - t_1 + 1, \quad (3)$$

where

$$x_a(t_1 - 1) = 1, \qquad x_a(t_2 + 1) = 1, \qquad \sum_{t=t_1}^{t_2} \left(1 - x_a(t)\right) = t_2 - t_1 + 1,$$

for $t_2 > t_1$.

The alarm run length r is the time distance between an alarm occurrence to the next one, i.e.,

$$r = T_1 + T_0. \quad (4)$$

Fig. 1 illustrates the definitions of T_1, T_0 and r.

Nuisance alarms refer to the occurred alarms that are not associated with any actual abnormalities. Hence, even if these alarms are ignored by operators, they would have no negative effects on industrial processes. By contrast, a true alarm must require a prompt operator response (ISA-18.2, 2009). By the definitions, nuisance and true alarms should be distinguished by looking at whether an alarm occurrence is accompanied by operator actions. Yuki (2002) and Noda, Higuchi, Takai, and Nishitani (2011) did so by focusing on the balance between alarm occurrences and operator actions. However, doing so is often infeasible for two reasons. First, operator actions are not recorded in a way being associated with alarm occurrences. Their connections have to be done manually, which is time-consuming and is confined by personal experiences subject to errors. Second, many operator actions are not recorded in computerized monitoring systems, e.g., checking the running condition of a sensor by visual inspection. Therefore, new ways of detecting nuisance alarms need to be developed. This is feasible for some specific types of nuisance alarms that have their own characteristics. Thus, these nuisance alarms can be detected effectively, as given in the following subsections.

2.2. Chattering alarms

Chattering alarms are the ones that make fast transitions between alarm and non-alarm states, where the transitions are not owing to corrective actions from operators. The industrial standard ISA-18.2 (2009) defines a chattering alarm as "an alarm that repeatedly transitions between the alarm state and the normal state in a short period of time". Repeating alarms are often taken as synonyms of chattering alarms. The industrial guide EEMUA-191 (2013) refers to an repeating alarm as "the same alarm raising and clearing repeatedly over a period of time". As special classes of chattering alarms, fleeting alarms are similar short-duration alarms that do not immediately repeat (ISA-18.2, 2009).