

***Detecting and Reducing Redundancy in Alarm  
Networks***

Butters, Timothy D. and Güttel, Stefan and Shapiro,  
Jonathan L.

2015

MIMS EPrint: **2015.19**

Manchester Institute for Mathematical Sciences  
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary  
School of Mathematics  
The University of Manchester  
Manchester, M13 9PL, UK

ISSN 1749-9097

# Detecting and Reducing Redundancy in Alarm Networks

Timothy D. Butters<sup>1,2</sup>, Stefan Güttel<sup>2</sup>, and Jonathan L. Shapiro<sup>3</sup>

**Abstract**—Alarm systems are vital for the safe operation of almost all large-scale industrial and technical installations, such as chemical plants or power stations. The optimization of alarm systems has great potential to improve the safety of these installations, and also to increase their profitability through the reduction of automated shut-downs and suboptimal operation modes.

In this work we present a new approach to alarm system optimization through the identification of redundant alarms. Our approach is based on a ranking of alarms by their connectivity in the alarm network. We also propose an overall redundancy measure for the alarm system which can be used to monitor performance improvements after redundant alarms have been removed. We present an example demonstrating that our ranking technique provides operational staff with useful information, allowing them to enhance the effectiveness of their existing alarm systems.

## I. INTRODUCTION

Alarm systems are a key component in the safe operation of large-scale industrial installations. They are responsible for alerting plant operators to abnormal behavior and, if necessary, performing automated shut-downs to prevent serious damage. However, it is not uncommon for alarm systems to exhibit suboptimal performance, as documented in industrial catastrophes such as the partial nuclear meltdown at Three Mile Island and the Texas City refinery explosion, both having root causes in poor alarm management [1], [2].

When discussing system optimization it is important to identify a suitable quantitative measure that indicates the effectiveness of the system in performing its intended task. Perhaps the most commonly used measure in the context of alarm systems is “operator load,” which is defined as the number of alarms each operator has to address in a given time interval. The global standard for alarm management configuration, EEMUA 191, provides recommended limits on the operator load of at most 1 alarm per operator every 10 minutes [3].

Operator load not only measures the frequency of alarms, but also takes into account the feasibility of acknowledging and actioning each alarm in the appropriate manner. In the event of an alarm cascade, in which many alarms are triggered in a short period of time due to a single physical

event, it may be very difficult or impossible for the operators to identify the root cause of the problem [4]. In some reported cases, the manual suppression of noisy alarms by operators has had severe consequences. It is therefore important to suppress these noisy alarms and alarm cascades without losing safety-critical information. This can be achieved by identifying and removing redundancies in the alarm system.

With sensors becoming cheaper it is now commonplace to install a large number of these across a plant, ensuring that every physical component is monitored. Although this may be necessary for some aspects of operation, it is usually unnecessary to configure an alarm for each single sensor, a practice that has also become common, particularly during initial configuration. Because these sensors often measure physically linked systems (e.g., temperature conduction between components), a single physical event can trigger a large number of alarms. The identification of such redundancies would allow the alarm management system to be re-configured in a way that operator load is reduced, while maintaining safety of operation.

When undertaking alarm optimization exercises it is common practice to prioritize alarms based on their arrival frequency, aiming to reduce operator load through the direct suppression of the most active alarms. Although there is some merit to this approach it does not take into account the pathways of activation for each high-frequency alarm. This means that the removal or suppression of these alarms may remove important information from the operator. It is also possible that an alarm ranked lower by frequency but often followed by an alarm cascade has a bigger impact on operator load than a high-frequency alarm. If instead the alarms with highest *redundancy* were targeted this would not only lead to a reduction in operator load, but also ensure that safety-critical information is retained.

Various methods to optimize alarm system performance have been discussed in the literature [5], [6], [7], but these mainly focus on real-time optimization through automated intervention. These methods require detailed computational models of the plant, which are not only difficult to construct, but also often rely on the manufacturer’s equipment parameters, rather than reflecting the actual behavior of the equipment in operation. Commercialized solutions of this kind are both expensive and require time-consuming set-up.

With many industrial installations such as chemical plants and power stations being operational for many years, there are large amounts of alarm data available for analysis. Using this data would allow for the reliable identification of alarm dependencies. This would ensure that the results obtained are relevant to the specific installation, rather than being

\*This work was supported by an Innovate UK Knowledge Transfer Partnership (KTP-9315).

<sup>1</sup>T.D. Butters is with the Research & Development group at Sabisu Ltd., Sale, Manchester, M33 3SJ, UK. [tim.butters@sabisu.co](mailto:tim.butters@sabisu.co)

<sup>2</sup>S. Güttel and T.D. Butters are with the Department of Mathematics, The University of Manchester, Manchester, M13 9PL, UK. [stefan.guettel@manchester.ac.uk](mailto:stefan.guettel@manchester.ac.uk)

<sup>3</sup>J.L. Shapiro is with the Department of Computer Science, The University of Manchester, Manchester, M13 9PL, UK. [jonathan.l.shapiro@manchester.ac.uk](mailto:jonathan.l.shapiro@manchester.ac.uk)

suggested by an idealized computational model.

Here we present a novel approach that utilizes historical and streamed alarm data to identify and prioritize sources of redundancy within an alarm system. We also present an alarm system health measure that can be used to track the progress of a plant optimization exercise, and which allows for differential analysis with varying operation parameters. This is an important requirement in effective alarm system optimization as complex plants are likely to behave differently under varying operation modes.

## II. METHODS

To identify redundancies in an alarm system we analyse historical alarm data that was acquired during the operation of an industrial petrochemicals plant. Once redundant alarms have been identified the plant operations team can then decide to eliminate these alarms, or deploy suppression rules in the alarm management system. Our data-driven approach removes the need for complex computational models of the industrial process required by other alarm management solutions [8], [9], [10]. To access, aggregate, and process the large datasets required our method was built in the Sabisu platform (<http://www.sabisu.co>).

### A. Constructing the Alarm Transition Matrix

Our approach is based on a network interpretation of the alarm system, where connected alarms are modelled as links in a directed graph. For a comprehensive overview of the graph-theoretical tools used in this paper we refer to [11]. A more applied introduction to graphs and related computational methods can be found in [12]. In our notation, bold printed capitals stand for matrices (e.g.,  $\mathbf{A}$  and  $\mathbf{T}$ ), and bold lower-case letters denote vectors (e.g.,  $\mathbf{c}$  and  $\mathbf{r}$ ). The notation  $\mathbf{c}^T$  stands for transposition of the vector  $\mathbf{c}$ .

Our aim is to construct an  $n \times n$  transition matrix  $\mathbf{T}$  of the alarm system, in which each entry  $t_{i,j}$  represents the probability that an alarm  $A_i$  is followed by alarm  $A_j$  within a fixed time window of, say, 30 s. These probabilities can be obtained empirically from historical alarm data, as is best explained at a simple example.

Assume we are given a system of  $n = 3$  alarms  $A_1$ ,  $A_2$ ,  $A_3$ , sounding in the following sequence:




Fig. 1. A simple alarm system with  $n = 3$  alarms. Time goes from left to right. The under-braces symbolize a fixed time window of 30 s, and a link  $A_i \rightarrow A_j$  indicates that alarm  $A_j$  followed alarm  $A_i$  within 30 s.

In the first step we will unfold this alarm sequence into an ensemble of alarm pairs and single alarms. By *alarm pair* we mean a pair  $(A_i, A_j)$  representing an occurrence of

alarm  $A_j$  within 30 s of alarm  $A_i$ . By *single alarm* we mean a singleton  $(A_i)$  representing an occurrence of alarm  $A_i$  which was not followed by any other alarm for at least 30 s. The alarm sequence in Fig. 1 is transformed into the following alarm ensemble:




Fig. 2. The alarm ensemble corresponding to our simple alarm system.

The motivation for decoupling an alarm sequence into an alarm ensemble is that transitions on alarm ensembles are compatible with elementary linear algebra operations. Effectively we are splitting the alarm sequence into statistically independent events such that the resulting ensemble can be modelled as a Markov chain with transition matrix  $\mathbf{T}$ . As a step towards obtaining this matrix  $\mathbf{T}$ , let us use the ensemble representation to construct an  $n \times n$  matrix  $\mathbf{A}$  and an  $n$ -vector  $\mathbf{c}$  counting the number of alarm pairs and singles. The matrix  $\mathbf{A}$  has only zeroes on the diagonal and each non-diagonal entry  $a_{i,j}$  corresponds to the number of pairs  $(A_i, A_j)$  in the alarm ensemble. The  $i$ th entry in  $\mathbf{c}$  corresponds to the sum of the number of single alarms  $(A_i)$  and alarm pairs  $(A_i, A_j)$  in the ensemble. For the example in Fig. 2 we have

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} 3 \\ 1 \\ 0 \end{bmatrix}. \quad (1)$$

Finally, the  $n \times n$  alarm transition matrix  $\mathbf{T}$  is defined as

$$\mathbf{T} = \text{diag}(\mathbf{c})^\dagger \mathbf{A},$$

where  $\text{diag}(\mathbf{c})^\dagger$  corresponds to the  $n \times n$  diagonal matrix with entries  $c_i^{-1}$  (set to zero in the case  $c_i = 0$ ). For example, with  $\mathbf{A}$  and  $\mathbf{c}$  given in (1) we obtain

$$\mathbf{T} = \begin{bmatrix} 0 & 1/3 & 1/3 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}.$$

By construction, the matrix  $\mathbf{T}$  is guaranteed to be non-negative and row sub-stochastic. This matrix can be used to retrieve statistics regarding the alarm ensemble. For example, with the vector  $\mathbf{1} = [1, \dots, 1]^T$ , the  $i$ th entry in the matrix-vector product  $\mathbf{T}\mathbf{1}$  corresponds to the probability that alarm  $A_i$  triggers any other alarm in the ensemble. Likewise, the  $j$ th entry in the product  $\mathbf{c}^T \mathbf{T}$  returns the number of times alarm  $A_j$  followed an *alarm state* represented in  $\mathbf{c}$ . In our example we have

$$\mathbf{c}^T \mathbf{T} = [0 \quad 1 \quad 2],$$

which means that starting from an alarm state  $\mathbf{c}^T = [3 \ 1 \ 0]$  we are more likely to hear alarm  $A_3$  than alarm  $A_2$ , and we are unlikely to hear alarm  $A_1$ . Note how this interpretation corresponds to a transition from the left of Fig. 2 to the right.

State transitions are consistent with matrix products, and therefore we can perform longer transitions by multiplication with higher powers of  $\mathbf{T}$ . For example, starting again from state  $\mathbf{c}$  we obtain after two transitions

$$\mathbf{c}^T \mathbf{T}^2 = [0 \ 0 \ 1],$$

i.e., we are likely to hear alarm  $A_3$  alone, and finally

$$\mathbf{c}^T \mathbf{T}^3 = [0 \ 0 \ 0]$$

means that the alarm system eventually becomes silent.

### B. Alarm Ranking

Probably the most common ranking technique used in alarm management is to simply count the number of times each alarm was triggered over a sufficiently long time period. Although this approach gives some information about the alarm system it does not account for cascades of alarms where a less frequent alarm triggers a series of other incident alarms, which in turn trigger a large number of alarms. Such a cascade may have a much bigger impact on operator load than a single faulty alarm that sounds frequently.

To account for the connectivity of alarms we propose to rank them using higher-order powers of the transition matrix  $\mathbf{T}$ . From algebraic graph theory it is known that the  $(i, j)$  entry in  $\mathbf{T}^k$  can be interpreted as the weighted number of walks of length  $k$  from alarm  $A_i$  to  $A_j$  (see, e.g., [11], [13]). Summing weighted powers  $w_k \mathbf{T}^k$  hence gives a measure of *communicability* between alarm  $A_i$  and  $A_j$  via walks of nonzero length (i.e., loops are not allowed). The weighting coefficients  $w_k = 1/k!$  have been found to produce meaningful results, i.e., we use the matrix function

$$f(\mathbf{T}) := \sum_{k=1}^{\infty} \frac{\mathbf{T}^k}{k!} = e^{\mathbf{T}} - \mathbf{I},$$

which corresponds to the matrix exponential of  $\mathbf{T}$  minus the identity matrix  $\mathbf{I}$ . The entries of the matrix exponential are a popular communicability measure in network analysis (see, e.g., [14]), but other functions can give similar results [15].

We now define the vector

$$\mathbf{r}^T := \mathbf{c}^T f(\mathbf{T})$$

to be the *redundancy ranking* of the alarm system. The entries of this vector relate to the frequencies at which each alarm can be expected to be triggered by the initial alarm state  $\mathbf{c}$  (which is obtained as an average alarm state by counting as explained in the previous section). Redundant alarms are those which are often preceded by incident alarms and hence do not contribute much information by themselves. These alarms will have large entries in the ranking vector  $\mathbf{r}$ . This vector can be efficiently approximated via Taylor expansion of  $f(z) = e^z - 1$ . A sketch of the computational method used is as follows:

```
v = c; r = 0;
for k = 1 to m:
    v = v * T / k;
    r = r + v;
```

where  $\mathbf{c}$  stores the row vector  $\mathbf{c}^T$  and  $m$  is an integer chosen so that accuracy requirements are met (in our computations we used  $m = 10$ ).

We note that there are other ways to sum the terms in the Taylor expansion of  $f$  (such as, e.g., the Horner scheme [16]). However, it is useful to keep track of the vector  $\mathbf{v}$  as it can be expected to align with a dominant left eigenvector of  $\mathbf{T}$ , which essentially is the PageRank vector<sup>1</sup> [12]. Indeed, the PageRank vector can also be obtained as a limit of  $\mathbf{c}^T f_k(\mathbf{T})$  as  $k \rightarrow \infty$ , where the functions  $f_k(z)$  have Taylor coefficients  $(0, \dots, 0, \rho(\mathbf{T})^{-k}, 0, \dots)$  with the only nonzero in the  $k$ th position. Here,  $\rho(\mathbf{T})$  denotes the spectral radius of  $\mathbf{T}$ . We will now argue that  $\rho(\mathbf{T})$  can be viewed as a measure for redundancy in the alarm system.

### C. Measuring Redundancy

Recall that the powers of  $\mathbf{T}$  correspond to longer-range dependencies in the alarm system. In an optimal alarm system, a single physical event (such as, e.g., “Overheat warning in compressor X”) will trigger exactly one alarm and then become silent, but in practice incident alarms will follow (e.g., “High pressure in compressor X”). Therefore, one measure of redundancy in alarm systems is the limiting rate at which the matrices  $\mathbf{T}^k$  decay. More precisely, let  $\|\cdot\|$  be an arbitrary matrix norm, then we consider the spectral radius

$$\rho(\mathbf{T}) = \lim_{k \rightarrow \infty} \|\mathbf{T}^k\|^{1/k} \geq 0 \quad (2)$$

as a measure of *alarm redundancy*. The equality in (2) is independent of the matrix norm used and typically referred to as *Gelfand’s formula*. Note that  $\rho(\mathbf{T}) \leq 1$  because  $\mathbf{T}$  is sub-stochastic. A highly redundant alarm system will have  $\rho(\mathbf{T}) \approx 1$ , meaning that alarm cascades may potentially decay off very slowly (resulting in an alarm cascade).

It should be emphasized that the spectral radius  $\rho(\mathbf{T})$  may be approached by the decay rate only for very large  $k$  in (2). No single number can capture the whole information about  $\|\mathbf{T}^k\|^{1/k}$ , but nevertheless we found  $\rho(\mathbf{T})$  to at least be a good indicator of redundancy. Indeed, if  $\|\cdot\|_1$  denotes the vector 1-norm (which is just the sum of the entries in a non-negative vector), then

$$\|\mathbf{c}^T \mathbf{T}^k\|_1 \leq \|\mathbf{c}\|_1 \|\mathbf{T}^k\|_1 \approx \|\mathbf{c}\|_1 \rho(\mathbf{T})^k. \quad (3)$$

In words: the number of alarms after  $k$  transitions (which is  $\|\mathbf{c}^T \mathbf{T}^k\|_1$ ) is approximately proportional to the number of alarms in the initial state (which is  $\|\mathbf{c}\|_1$ ) multiplied by  $\rho(\mathbf{T})^k$ . We will demonstrate this behavior at a practical example in the following section.

<sup>1</sup>Strictly speaking, more conditions are required, e.g., the assumption that the spectral radius  $\rho(\mathbf{T}) > 0$  is the eigenvalue of largest modulus and it is a simple eigenvalue, and  $\mathbf{c}$  is not orthogonal to the associated eigenvector. Some of these requirements can be guaranteed by the Perron–Frobenius theory for non-negative matrices (see, e.g., [12, Section 15.2]).

### III. RESULTS

We analyzed industrial alarm data from a SABIC UK petrochemicals plant. The alarm log ranges over approximately two weeks (from 2014/04/01 to 2014/04/16) and collects alarm signals from  $n = 1433$  alarms.

A sparsity plot of the  $n \times n$  transition matrix  $\mathbf{T}$  is shown in Fig. 3. With 34948 nonzero elements this matrix has a sparsity of 1.7%. The matrix rows and columns have been reordered by the reverse Cuthill–McKee algorithm [17] so that nonzero elements are brought close to the diagonal. One can see from the resulting sparsity pattern that some alarms appear in clusters. The alarm redundancy was found to be

$$\rho(\mathbf{T}) = 0.856.$$




Fig. 3. Sparsity pattern of the transition matrix  $\mathbf{T}$  corresponding to alarm data from a large-scale petrochemicals plant.

In Fig. 4 we visualize the entries of the vectors  $\mathbf{c}^T$ ,  $\mathbf{c}^T\mathbf{T}$ , and  $\mathbf{c}^T\mathbf{T}^2$ . Recall from Section II-A that the  $i$ th entry in  $\mathbf{c}^T$  corresponds to the sum of the number of single alarms ( $A_i$ ) and alarm pairs ( $A_i, A_j$ ). In total, we found 89372 alarm singles and pairs (this corresponds to the vector norm  $\|\mathbf{c}\|_1$ ). One multiplication by  $\mathbf{T}$  reduces the total alarm number to 76262, and a further multiplication reduces it to 66155. We find that the reduction factor is very well approximated by our alarm redundancy measure  $\rho(\mathbf{T})$ , which is in agreement with (3).

We now apply the exponential-based ranking as described in Section II-B to determine the ten alarms with the highest redundancy. The effect of consecutively removing 1, 2, ..., 10 of these alarms on the redundancy measure is shown in Table I. This table also shows the alarm index (which is the unique identifier for each alarm corresponding to the




Fig. 4. Entries in the vectors  $\mathbf{c}^T\mathbf{T}^k$  for  $k = 0, 1, 2$  (from top to bottom).

row/column index in  $\mathbf{T}$ ), and the count (which is the number of times each alarm appeared in the alarm ensemble). We find that the redundancy of the alarm system is reduced significantly by eliminating the five alarms of highest rank. After this elimination the redundancy measure continues to decrease, but at a slower rate.

TABLE I  
EXPONENTIAL-BASED ALARM RANKING AND THE REDUCTION IN REDUNDANCY BY REMOVING THE HIGHEST-RANKED ALARMS.

Alarms Removed	Index	Count	Redundancy Measure
0	—	—	0.856
1	1265	5902	0.841
2	1266	3334	0.833
3	1279	863	0.822
4	1272	1152	0.818
5	1275	971	0.815
10	1013	269	0.808

By comparing with Table II, which shows the index and counts of the most frequent alarms, it becomes clear that a simple ranking based on alarm frequency will give different results. Although four of the alarms are present in both tables, they appear in a different order even though there is a significant difference in their respective counts. Alarm 1279 appears only in the exponential-based ranking, whereas alarm 1282 appears only in the ranking by frequency.

TABLE II  
ALARM RANKING BY FREQUENCY.

Rank	Index	Count
1	1265	5902
2	1266	3334
3	1272	1152
4	1282	1010
5	1275	971

Using a simple load calculation it was possible to estimate the reduction in operator load that would have been attained if the highest-ranked alarms in Table I were removed for the analysed time period. It was found that the average load per operator was reduced by approximately 12%.

#### IV. DISCUSSION

The results in Section III, and other tests not reported here, indicate that our approach to alarm ranking can be used effectively to reduce the load on plant operators. We found that the exponential-based ranking may differ significantly from the ranking of alarms by frequency. This means that there is dependency information in the alarm network which is worth factoring into the analysis. The frequency-based ranking fails to include alarm cascades, in which a less frequent alarm triggers an alarm with a large number of incident alarms. Such alarm cascades may result in much higher operator load than individual alarms that sound frequently, and the exponential-based ranking proposed in this paper takes this into account.

Having a reliable ranking method and the possibility to simulate the effect of alarm system modifications is particularly important in real-world applications as real changes to a running alarm system can take considerable time to enact. This is because any proposed change must first be approved by the alarm management team due to its safety-critical nature. It is therefore vital that alarms with the largest overall redundancy can be identified reliably and processed first.

The proposed alarm redundancy measure provides a useful quantitative metric to assess the current state of an alarm system, and it also allows to track improvements as the alarm system is optimized. Moreover, this measure can be used to compare alarm systems in different operation modes, or between similar industrial installations.

It is common for manufacturing plants to change their mode of operation due to environmental factors or varying customer demand. For example, this could mean running a chemical process on a different set of raw materials, or running the plant with modified parameters to produce a higher or lower grade of end product. With the plant operating in these distinct modes it is likely that different alarm behavior will be observed, possibly causing a change in the redundancy measure. Tracking this change in redundancy allows problematic operation modes to be identified and documented. Alarms being particularly active

in a problematic operation mode could be tagged and dealt with separately.

Redundant alarms can be handled in several possible ways in order to reduce the operator load. The simplest solution is to remove the alarm from the system. This is a possible course of action if the alarm has a high redundancy in all operation modes, which suggests that it rarely gives any meaningful information to the operators. However, if an alarm is highly redundant in only a subset of operation modes, removal is clearly not the best course of action. As the alarm provides useful information in some modes, it should be suppressed only in modes in which it has been identified as redundant. An alternative to these two methods is trigger-based suppression. Here a redundant alarm  $A_2$  is suppressed for 30s when it appears in combination with another alarm  $A_1$ , and both alarms have been identified to appear very likely together (so only in combination with alarm  $A_1$ , alarm  $A_2$  can be considered as redundant). This allows alarm  $A_2$  to act as normal when it appears by itself. In practice probably all three methods of alarm load reduction will be utilized.

#### V. CONCLUSIONS

The alarm ranking approach presented here provides an effective way to improve the performance of existing alarm systems. Due to its data-driven nature, this approach it is very easy to implement and versatile. The alarm ranking will automatically adapt to the specific alarm installation being analysed. Moreover, any changes the plant may go through, whether intended or incidental, will eventually be reflected in the alarm data, allowing for an automatic re-ranking of the alarms.

Our results show that in a real-world installation our ranking-based approach may significantly reduce operator load while making sure that important alarms are retained. The approach is currently being trialled by a number of Sabisu's customers, who have provided a positive initial response. In the future, with the permission of these customers, a comprehensive analysis of the method could be performed to quantify the reduction in operator load throughout a variety of operational environments.

This approach does not require any complex configuration, and can easily be implemented to any industrial installation using an alarm management system. We believe that its simplicity and versatility make it an extremely effective solution to alarm optimization problems.

#### ACKNOWLEDGMENT

The authors would like to thank T. J. Sharpe and S. Toner from Sabisu Ltd. for their help in obtaining data, and many discussions regarding the analysis and interpretation. They would also like to thank SABIC UK for permitting the use of their historical alarm logs for this study.

## REFERENCES

- [1] L. M. Toth, "The three mile island accident diagnosis and prognosis," in *189th Meeting of the American Chemical Society*. American Chemical Society, 1986, Conference Proceedings.
- [2] D. Holmstrom, F. Altamirano, J. Banks, G. Joseph, M. Kaszniak, C. Mackenzie, R. Shroff, H. Cohen, and S. Wallace, "CSB investigation of the explosions and fire at the BP Texas City refinery on March 23, 2005," *Process Safety Progress*, vol. 25, no. 4, pp. 345–349, 2006.
- [3] EEMUA 191, *Alarm Systems - A Guide to Design, Manufacturing and Procurement*, 3rd ed. The Engineering Equipment and Materials Users' Association, 2013.
- [4] R. B. Whittingham, *The Blame Machine: Why Human Error Causes Accidents*. Elsevier Butterworth-Heinemann, 2004.
- [5] J. Liu, K. W. Lim, W. K. Ho, K. C. Tan, R. Srinivasan, and A. Tay, "The intelligent alarm management system," *IEEE Software*, vol. 20, no. 2, pp. 66–71, 2003.
- [6] Y. Xu and Q. Zhu, "Alarm optimization for process industry based on matter-element analysis," *Journal of Chemical Industry and Engineering (China)*, vol. 7, p. 2, 2008.
- [7] J. Liu, J. Zhou, S. L. R. Li, K. Tan, W. K. Ho, R. Srinivasan, and A. Tay, "Intelligent alarm management through suppressing nuisance alarms and providing operator advice," in *Proceedings of the 4th IFAC Workshop On-Line Fault Detection and Supervision*, 2001, pp. 81–86.
- [8] M. Lind, "Modeling goals and functions of complex industrial plants," *Applied Artificial Intelligence: An International Journal*, vol. 8, no. 2, pp. 259–283, 1994.
- [9] P. Kämpjärvi, M. Sourander, T. Komulainen, N. Vatanski, M. Nikus, and S.-L. Jämsä-Jounela, "Fault detection and isolation of an on-line analyzer for an ethylene cracking process," *Control Engineering Practice*, vol. 16, no. 1, pp. 1–13, 2008.
- [10] A. Guasch, J. Quevedo, and R. Milne, "Fault diagnosis for gas turbines based on the control system," *Engineering Applications of Artificial Intelligence*, vol. 13, no. 4, pp. 477–484, 2000.
- [11] N. Biggs, *Algebraic Graph Theory*. Cambridge University Press, 1993.
- [12] A. N. Langville and C. D. Meyer, *Google's PageRank and Beyond: The Science of Search Engine Rankings*. Princeton University Press, 2011.
- [13] M. E. Newman, "Analysis of weighted networks," *Physical Review E*, vol. 70, no. 5, p. 056131, 2004.
- [14] M. Benzi, E. Estrada, and C. Klymko, "Ranking hubs and authorities using matrix functions," *Linear Algebra and its Applications*, vol. 438, no. 5, pp. 2447–2474, 2013.
- [15] M. Aprahamian, D. J. Higham, and N. J. Higham, "Matching exponential-based and resolvent-based centrality measures," Manchester Institute for Mathematical Sciences, The University of Manchester, UK, MIMS EPrint 2015.3, 2015.
- [16] W. G. Horner, "A new method of solving numerical equations of all orders by continuous approximations," *Philosophical Transactions of the Royal Society*, vol. 109, pp. 308–335, 1819.
- [17] E. Cuthill and J. McKee, "Reducing the bandwidth of sparse symmetric matrices," in *Proceedings of the 24th National Conference of the ACM*, 1969, pp. 157–172.