INTERNATIONAL JOURNAL OF COMMUNICATION SYSTEMS
Int. J. Commun. Syst. **12**, 217–228 (1999)

Robust event correlation scheme for fault identification in communication networks

Chi-Chun Lo* and Shing-Hong Chen

Institute of Information Management, National Chiao-Tung University, 1001 Ta Hsueh Road, Hsinchu, Taiwan 300, Republic of China

SUMMARY

The complexity of communication networks and the amount of information transferred in these networks have made the management of such networks increasingly difficult. Since faults are inevitable, quick detection, identification, and recovery are crucial to make the systems more robust and their operation more reliable. This paper proposes a novel event correlation scheme for fault identification in communication networks. This scheme is based on the algebraic operations of sets. The causality graph model is used to describe the cause-and-effect relationships between network events. For each disorder, and each manifestation, a unique prime number is assigned. The use of the greatest common devisor (GCD) makes the correlation process simple and fast. A simulation model is developed to verify the effectiveness and efficiency of the proposed scheme. From simulation results, we notice that this scheme not only identifies multiple disorders at one time but also is insensitive to noise. The time complexity of the correlation process is close to a function of n, where n is the number of observed manifestations, with order $O(n^2)$; therefore, the on-line fault identification is easy to achieve. Copyright © 1999 John Wiley & Sons, Ltd.

KEY WORDS: event correlation; fault identification; algebraic operation of set; causality graph model; greatest common devisor (GCD)

1. Introduction

Faults are unavoidable in large and complex communication networks, but quick detection and identification can significantly improve network reliability. Network faults are often the result of underlying problems such as hardware or software failures, performance bottleneck, configuration inconsistency. Since a single fault in one resource often causes alarms in other related resources, operational staff must be able to correlate the observed alarms to identify and localize underlying problems. However, this manual process does not scale to the growing speed, complexity, and size of today's communication networks. Computer automation of this manual process becomes increasingly desirable.

Although the OSI management standard provides a framework for managing faults in heterogeneous open systems, it does not address the methodology used to detect and diagnose faults. To fill this gap, various theoretical approaches have been suggested. Rule-based expert systems so far have been the major approach for solving the alarm correlation problem.[1,2] This approach suits well-defined problems where the environment is not very dynamic. Diagnostic

* Correspondence to: Chi-Chun Lo, Institute of Information Management, National Chiao-Tung University, 1001 Ta Hsueh Road, Hsinchu, Taiwan 300, ROC. E-mail-cclo@cc.nctu.edu.tw

CCC 1074–5351/99/030217–12$17.50

Received March 1998
Accepted December 1998

reasoning provides another basis for developing expert systems with which it can find the solutions to multiple simultaneous problems.[3] In Reference 3, the set covering model is proposed and the concepts of 'disorder' and 'manifestation' are described. Case-based reasoning offers potential solutions to the problems of adaptation and knowledge acquisition.[4] Using finite state machines (FSMs) to detect fault is another approach.[5–7] This method is able to cope with incomplete information and unforeseen faults. However, all approaches mentioned above are very sensitive to the 'noise' (e.g. lost, delayed, etc.) in the alarm stream; therefore their use for the real-world network is limited. The application of probabilistic reasoning is another well-known approach.[8–12] Using Bayesian network to identify faults in the linear lightwave networks has been presented in Reference 8. Wang and Schwartz[12] use *a priori* knowledge and probabilistic estimates of link failures to pick out links that are likely to be faulty. For newly installed systems, however, such information is not available. Recent study of fault identification has been focused on event correlation with coding approach.[13] The complete set of events caused by a problem (or disorder) is represented by a 'code' that identifies the problem. Correlation is simply the process of 'decoding' the set of observed symptoms (or manifestations) by determining which problem matches its code. The causality graph model is used to describe the causal relations between events in the coding method. Nonetheless, code length needs to be extended when new events are created; consequently, computing complexity increases drastically due to redundant codes.

In this paper, we propose a novel event correlation scheme that has its origin of the algebra of sets.[14] This scheme not only identifies multiple disorders at one time but also is insensitive to noise. In the following section, the causality graph model is described. In Section 3, the proposed scheme is detailed. Simulation results and analyses are given in Section 4. Section 5 describes the limitations of the proposed scheme and future works. Finally, Section 6 concludes this paper.

2. Causality graph model[13]

Network operations management consists mainly of monitoring, interpreting, and handling events, where an event is defined as an exceptional condition in the operation of the network. A disorder is an event that can be handled directly. Some disorders are directly observable, while others can be observed only indirectly by observing their manifestations. Manifestations are defined as events that are observable; for example, degraded application performance is a manifestation of the faulty interface problem. Manifestations cannot be handled instantly; to make a manifestation go away, it needs to handle its root cause disorder. Relationships are essential components of correlation, because disorders and manifestations propagate from one object to another along relationships.

A natural candidate for representing disorder domain is the causality graph model. Causality is a partial order relation between events. The notation $d \rightarrow m$ is used to illustrate the causality of event m by event d. The relation $\rightarrow$ can be described by a causality graph whose nodes represent events and whose directed edges represent causality. For example, Figure 1(a) depicts the causality graph of a network consisting of 11 nodes. As shown in Figure 1(b), nodes of a causality graph may be labeled as either a disorder (d) or a manifestation (m).

Causality graph may include information that does not contribute to correlation analysis. Certain manifestations are not directly caused by any disorder but only by other manifestations; for instance, manifestation 7 in Figure 1(b). These indirect manifestations may be eliminated without loss of information. Events may form cyclic relation; e.g. $d_1 \rightarrow d_2 \rightarrow d_3 \rightarrow d_1$, many-to-one relation; e.g. $d_1, d_2, d_3 \rightarrow d_4$, or inference relation; e.g. $d_1 \rightarrow d_2 \rightarrow d_3$. All these relations represent

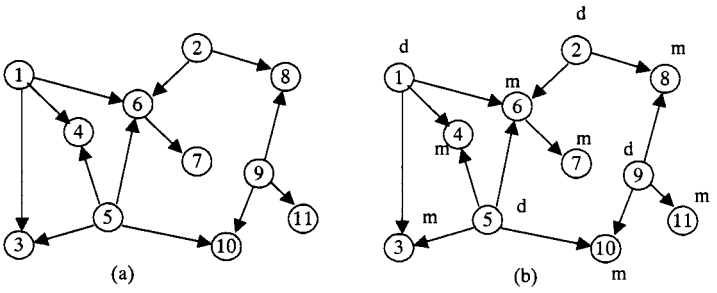

Figure 1. (a) A causality graph; (b) its labelling

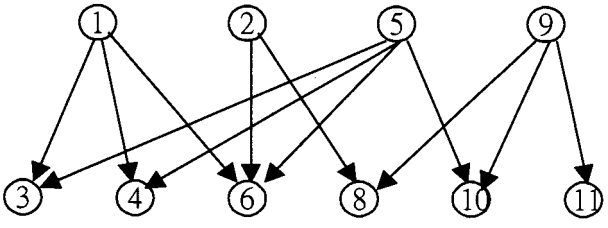

Figure 2. A correlation graph

causal equivalence; consequently, all involved events can be aggregated into a single event.[13] In this paper, we assume that all causality graphs are properly pruned and there are no cyclic, many-to-one, and inference relations. On the basis of these assumptions, the causality graph shown in Figure 1(a) can be reduced to the correlation graph illustrated in Figure 2.

3. Proposed event correlation scheme

3.1. Notation

In order to use the algebraic operations of sets for correlation, the information contained in the correlation graph is grouped into different sets. The following set notation are defined the same as those defined by Reggia *et al.*[3] Two discrete finite sets, **D** and **M**, are defined where **D** represents all possible disorders d_i that can occur, and **M** represents all possible manifestations m_j that may occur when one or more disorders are presented. We will assume that $\mathbf{D} \cap \mathbf{M} = \phi$. To capture the intuitive notion of causation, we assume knowledge of a correlation $\mathbf{C} \subseteq \mathbf{D} \times \mathbf{M}$, where $\langle d_i, m_j \rangle \in \mathbf{C}$, represents '$d_i$ can cause m_j'. Given **D**, **M** and **C**, the following two sets are defined:

$$\text{man}(d_i) = \{m_j | \langle d_i, m_j \rangle \in \mathbf{C}\}, \quad \forall d_i \in \mathbf{D}$$

and

$$\text{causes}(m_j) = \{d_i | \langle d_i, m_j \rangle \in \mathbf{C}\}, \quad \forall\ m_j \in \mathbf{M}$$

These sets represent all possible manifestations caused by d_i, and all possible disorders that cause m_j, respectively.

The transformation from the correlation graph to algebraic sets can be best illustrated by examples. Consider the correlation graph shown in Figure 2, where the set of disorders, **D**, is equal to $\{d_1, d_2, d_5, d_9\}$ and the set of manifestations, **M**, is equal to $\{m_3, m_4, m_6, m_8, m_{10}, m_{11}\}$. According to the above set notation, the set of manifestations of each disorder in **D** can be derived as follows:

$$\begin{aligned}\text{man}(d_1) &= \{m_3, m_4, m_6\}\\ \text{man}(d_2) &= \{m_6, m_8\}\\ \text{man}(d_5) &= \{m_3, m_4, m_6, m_{10}\}\\ \text{man}(d_9) &= \{m_8, m_{10}, m_{11}\}, \text{ respectively}\end{aligned}$$

The set of disorders of each manifestation in **M** can be derived as follows:

$$\begin{aligned}\text{causes}(m_3) &= \{d_1, d_5\}\\ \text{causes}(m_4) &= \{d_1, d_5\}\\ \text{causes}(m_6) &= \{d_1, d_2, d_5\}\\ \text{causes}(m_8) &= \{d_2, d_9\}\\ \text{causes}(m_{10}) &= \{d_5, d_9\}\\ \text{causes}(m_{11}) &= \{d_9\}, \text{ respectively.}\end{aligned}$$

3.2. Event correlation

Since the correlation scheme is based on the operations in the algebra of sets, the event correlation procedure is then to find a set of disorders whose man(d_i) is optimally included in the set of observed manifestations.

The event correlation procedure is done by the selection process, followed by the identification process. Figure 3 illustrates this procedure. The selection process is simply a series of intersection and union operations on causes(m_j) of observed manifestations. The identification process uses the inclusion relation to identify real disorders.

The purpose of the selection process is to reduce the number of inferred disorders. In the selection process, an elimination rule is used to remove those unlikely disorders which are included in causes(m_j) of observed manifestations. In essence, this rule eliminates all of the

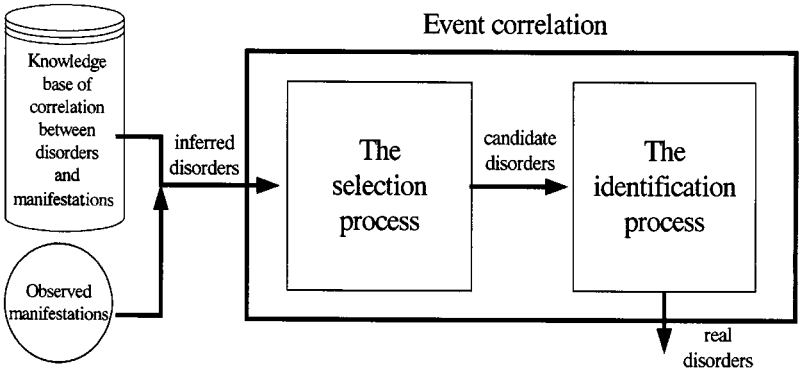

Figure 3. The event correlation procedure

disorders which have only one active manifestation. For any inferred disorder, at least one active manifestation should be observed. Therefore, disorders with only one active manifestation could be good candidates to be eliminated so as to reduce the search space. However, there is a potential problem associated with the elimination rule. It will remove those disorders which truly have only one manifestation. This problem can be fixed as follows: for every one-manifestation disorder in the knowledge base, a unique 'dummy' manifestation is assigned. In the real world, both real and dummy manifestation will be triggered once the one-manifestation disorder occurs. For a set of k observed manifestations, the elimination rule can be expressed as follows:

$$(\text{causes}(m_1) \cap \text{causes}(m_2)) \cup (\text{causes}(m_1) \cap \text{causes}(m_3)) \cup \cdots \cup$$
$$(\text{causes}(m_{k-2}) \cap \text{causes}(m_k)) \cup (\text{causes}(m_{k-1}) \cap \text{causes}(m_k)) \quad (1)$$

In the identification process, the primary goal is to identify those disorders that have been selected in the selection process, whose man(d_i) is a subset of the set of observed manifestations. Since information loss is unavoidable in communication networks, the set of observed manifestations usually does not contain all of the manifestations included in causes(m_j) of those selected disorders. Thus it is necessary to set a level of tolerance for evaluating the possibility of each selected disorder. The tolerance level specifies the maximal number of manifestation loss allowed. The tolerance level, t, satisfies the following inequality:

$$\|A - A'\| \leqslant t \quad (2)$$

where A represents the set man(d_i) of the selected disorder, A' represents the intersection between A and the set of observed manifestations, $\|A - A'\|$ represents the number of elements in set $(A - A')$.

If t is equal to 0, then A and A' are identical; i.e. manifestation loss is not allowed. If t is equal to 1, then A' may have one less element than A; i.e. only one missing manifestation is allowed.

Correlation Scheme. The correlation scheme can be stated as follows:

Step 1: Select candidate disorders by using the elimination rule given in equation (1).

Step 2: Set the tolerance level according to equation (2), and then use the inclusion relation to identify real disorders.

Example. Suppose that two disorders d_1 and d_2 in the correlation graph of Figure 2 have occurred. Two cases are considered. In the first case, there is no manifestation loss. In the second case, there are missing manifestations.

Case 1: For the correlation graph shown in Figure 2, the set of observed manifestations, $\mathbf{M}'$, is equal to $\{m_3, m_4, m_6, m_8\}$. In the selection process, the elimination rule can be expressed as follows:

$$(\text{causes}(m_3) \cap \text{causes}(m_4)) \cup (\text{causes}(m_3) \cap \text{causes}(m_6)) \cup \cdots \cup$$
$$(\text{causes}(m_4) \cap \text{causes}(m_8)) \cup (\text{causes}(m_6) \cap \text{causes}(m_8)) \quad (3)$$

After applying the elimination rule of equation (3), we get the set of candidate disorders, $\mathbf{D}'$, which is equal to $\{d_1, d_2, d_5\}$. In the identification process, the set man(d_i) of each candidate

disorder in $\mathbf{D}'$ is checked whether it is contained in $\mathbf{M}'$; two disorders d_1 and d_2 are identified as real disorders. Since we assume that two disorders d_1 and d_2 have occurred; therefore, the proposed scheme did find the exact solution.

Case 2: Assume that manifestation m_6 in *Case 1* has been lost; therefore, the set of observed manifestations, $\mathbf{M}'$, is equal to $\{m_3, m_4, m_8\}$. After applying the elimination rule of equation (1), we get the set of candidate disorders, $\mathbf{D}'$, which is equal to $\{d_1, d_5\}$. In the identification process, we set the tolerance level to be equal to 1. First, we examine d_1. By applying equation (2), we get the following:

$$A' = \text{man}(d_1) \cap \mathbf{M}' = \{m_3, m_4\} \quad \text{and} \quad \| \text{man}(d_1) - A' \| = 1$$

Thus we conclude that disorder d_1 is a real disorder. Second, we examine d_5 and find that it is not a real disorder. Note that the proposed scheme found disorder d_1, but failed to identify disorder d_2.

3.3. Implementation issues

For the purpose of computer simulation, we need to assign a numeric value to each disorder, and each manifestation, respectively. Also, in order to convert the symbolic set operations into numeric operations, we need to revise the correlation scheme.

Event representation. The following notation is defined:

$d\text{ID}_i$: disorder identifier; it is a unique prime number assigned to disorder d_i,
$m\text{ID}_j$: manifestation identifier; it is a unique prime number assigned to manifestation m_j,
$Md\text{ID}_i$: the product of the identifiers of all manifestations caused by disorder d_i,
$Dm\text{ID}_j$: the product of the identifiers of all disorders that cause the same manifestation m_j,
OM: the product of the identifiers of all observed manifestations.

According to the above notation, each disorder d_i can be represented by an order pair ($d\text{ID}_i$, $Md\text{ID}_i$), and each manifestation m_j can be represented by an order pair ($m\text{ID}_j$, $Dm\text{ID}_j$).

Take Figure 2 as an example. First, by assigning prime number to each disorder and each manifestation, we have the following:

$$d\text{ID}_1 = 2; \quad d\text{ID}_2 = 3; \quad d\text{ID}_5 = 5; \quad d\text{ID}_9 = 7 \quad m\text{ID}_3 = 2; \quad m\text{ID}_4 = 3; \quad m\text{ID}_6 = 5;$$
$$m\text{ID}_8 = 7; \quad m\text{ID}_{10} = 11; \quad m\text{ID}_{11} = 13$$

Then, we can derive the representation of each disorder and each manifestation as follows:

$$d_1 = (d\text{ID}_1, Md\text{ID}_1) = (d\text{ID}_1, m\text{ID}_3 \times m\text{ID}_4 \times m\text{ID}_6) = (2, 2 \times 3 \times 5) = (2, 30)$$
$$d_2 = (3, 35), d_5 = (5, 330), d_9 = (7, 1001)$$
$$m_3 = (m\text{ID}_3, Dm\text{ID}_3) = (m\text{ID}_3, d\text{ID}_1 \times d\text{ID}_5) = (2, 2 \times 5) = (2, 10)$$
$$m_4 = (3, 10); m_6 = (5, 6); m_8 = (7, 1001); m_{10} = (11, 35); m_{11} = (13, 7)$$

Revised scheme. Three numeric functions, f, g, and h are defined. The f function factorizes a number into the set of its prime factors; e.g. $f(30)$ is equal to $\{2, 3, 5\}$. The g function extracts the greatest common divisor of two numbers; e.g. $g(6, 9)$ is equal to 3. The h function makes sure that

there is only one instance of a prime number existing in a prime number set; e.g. $h(\{2, 2, 3, 3, 5\})$ is equal to $\{2, 3, 5\}$.

The correlation scheme given in Section 3.2 can be rewritten as follows:

Step 1: Select candidate disorders by using the following elimination rule:

$$h(f(g(Dm\mathrm{ID}_1, Dm\mathrm{ID}_2), f(g(Dm\mathrm{ID}_1, Dm\mathrm{ID}_3)), \ldots, f(g(Dm\mathrm{ID}_{k-1}, \quad Dm\mathrm{ID}_k))) \tag{4}$$

where k is the number of observed manifestations.

Step 2: Set the tolerance level t, and then identify those disorders d_i, selected in step 1, that satisfy

$$\| f(Md\mathrm{ID}_i/g(Md\mathrm{ID}_i, \mathrm{OM})) \| \leqslant t \tag{5}$$

Note that $\| A \|$ denotes the number of elements in set A and

$$\| f(1) \| = \| \emptyset \| = 0$$

Example. Suppose that disorders d_1 in the correlation graph of Figure 2 have occurred and there is no manifestation loss. It is assumed that manifestation m_3, m_4, and m_6 are observed. It is necessary to identify which disorders have occurred. In step 1, the elimination rule of equation (4) can be expressed as follows:

$$h(f(g(Dm\mathrm{ID}_3, Dm\mathrm{ID}_4)), f(g(Dm\mathrm{ID}_3, Dm\mathrm{ID}_6)), f(g(Dm\mathrm{ID}_4, Dm\mathrm{ID}_6))) \tag{6}$$

By reducing equation (6), we have the following:

$$h(f(10), f(2), f(2)) = \{2, 5\} = \{d\mathrm{ID}_1, d\mathrm{ID}_5\}$$

Therefore, d_1 and d_5 are candidate disorders.

In step 2, suppose that the tolerance level t is set to 0. Thus, it can be shown that

$$\| f(Md\mathrm{ID}_1/g(Md\mathrm{ID}_1, \mathrm{OM})) \| = \| f(Md\mathrm{ID}_1/30) \| = \| f(1) \| = 0 \leqslant t \tag{7}$$

$$f(Md\mathrm{ID}_5/g(Md\mathrm{ID}_5, \mathrm{OM})) \| = \| f(Md\mathrm{ID}_5/30) \| = \| f(11) \| = 1 > t \tag{8}$$

where

$$\mathrm{OM} = m\mathrm{ID}_3 \times m\mathrm{ID}_4 \times m\mathrm{ID}_6 = 2 \times 3 \times 5 = 30,$$

$$g(Md\mathrm{ID}_1, \mathrm{OM}) = g(30, 30) = 30$$

$$g(Md\mathrm{ID}_5, \mathrm{OM}) = g(330, 30) = 30$$

From equations (7) and (8), we conclude that only disorder d_1 has occurred. Since we assume that disorders d_1 have occurred; therefore, the proposed scheme did find the exact solution.

4. Simulation results and analyses

Simulation experiments have been conducted to evaluate the effectiveness and efficiency of the proposed event correlation scheme. Test cases include some 6000 manifestations and 9500 disorders. The proposed scheme is coded in C and is run on an Intel Pentium 133 processor. Events are randomly generated. The benchmark model makes two conservative assumptions. It assumes an underinstrumented system where the number of observed manifestations is much

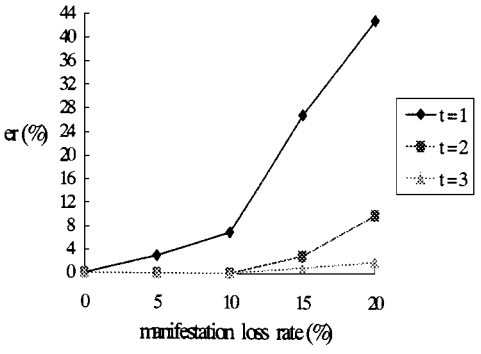

Figure 4. Correlation error ratio of the proposed scheme: er—correlation error ratio; t—tolerance level

smaller than the number of disorders; typical systems are overinstrumented. It also assumes a sparse propagation model where only a small number of manifestations are caused by a typical disorder; in real-world systems with complex dependencies, disorders tend to propagate very widely.

4.1. Effectiveness

Two factors, the correlation ratio (CR) and the hit ratio (HR), are used to measure the effectiveness of the proposed scheme. The correlation ratio shows the robustness of the proposed scheme. The hit ratio demonstrates the accuracy of the solution; i.e., real disorders, obtained from the proposed scheme. The effectiveness, E, is defined as follows:

$$E = \mathrm{CR} \times \mathrm{HR} = (1 - \mathrm{er}) \times \frac{\mathrm{ce}}{\mathrm{te}} \tag{9}$$

where er represents the percentage of exact disorders† not identified by the proposed scheme; ce represents the number of exact disorders identified by the proposed scheme; te represents the number of real disorders identified by the proposed scheme.

For different manifestation loss rates, Figure 4 presents the correlation error ratio for the proposed scheme. For the purpose of comparison, we have also simulated the coding scheme given by Yemini *et al.*[13] Figure 5 depicts the simulation results of the coding scheme.

By examining Figures 4 and 5, we give a detailed comparison between the effectiveness of the proposed scheme and that of the coding scheme in Table I. For the manifestation loss rates 5 and 10, the proposed scheme performs better than the coding scheme. It is worth mentioning that the proposed scheme did not have correlation error while the coding scheme did, when the tolerance level was relaxed to 2; which is equivalent to the radius of 1 specified in the coding scheme. For the manifestation loss rates 15 and 20, the proposed scheme performs better than the coding scheme when t is 3; which is equivalent to the radius of 1·5 specified in the coding scheme. The above

† Exact disorders are those disorders which are defined in the knowledge base and whose occurrence cause the observed manifestations.

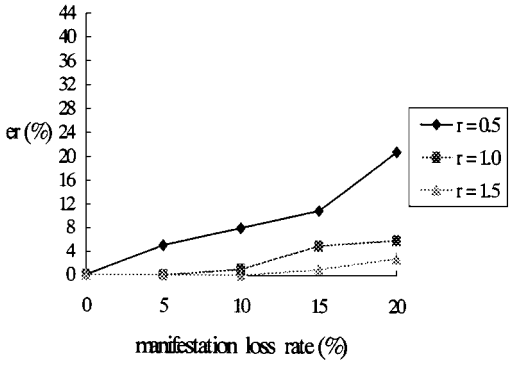

Figure 5. Correlation error ratio of the coding scheme: er—correlation error ratio; *r*—codebook radius

observations demonstrate that the proposed scheme can converge to a better solution than the coding scheme by properly setting the tolerance level.

4.2. Efficiency

Time complexity is used to evaluate the efficiency of the proposed scheme. The solution of the proposed scheme can be obtained in

$$\binom{n}{2} + m$$

steps, where n represents the number of observed manifestations, m represents the number of candidate disorders.

Since it can be shown that

$$\binom{n}{2} + m = \frac{n(n-1)}{2!} + m = O(n^2) \tag{10}$$

thus, the time complexity of the correlation scheme is close to a function of n with order $O(n^2)$.

For the coding scheme, which consists of the codebook selection phase and the decoding phase, its time complexity can be shown to be equal to

$$\binom{n}{k} \times M = O(M \cdot n^k) \tag{11}$$

where n represents the number of observed manifestations, M represents the number of disorders in a managed domain and M is much greater than n, k is equal to the minimum of the following set: $\{\text{number of manifestations caused by disorder } d_i | i = 1, 2, \ldots, M\}$.

By comparing equation (10) with equation (11), we observe that the time complexity of the proposed scheme is much better than that of the coding scheme.

Table I. Comparison between the effectiveness of the proposed scheme and that of the coding scheme

scheme / symptoms loss rate (%)	The proposed scheme									The coding scheme								
	$t = 1$			$t = 2$			$t = 3$			$r = 0{\cdot}5$			$r = 1{\cdot}0$			$r = 1{\cdot}5$		
	CR	HR	E	CR	HR	E	CR	HR	E	CR	HR	E	CR	HR	E	CR	HR	E
5	0·97	1	0·97	1	1	1	1	1	1	0·95	1	0·95	1	1	1	1	0·99	0·99
10	0·93	1	0·93	1	1	1	1	1	1	0·92	1	0·92	0·99	1	0·99	1	0·99	0·99
15	0·73	1	0·73	0·93	1	0·93	0·99	1	0·99	0·89	1	0·89	0·95	1	0·95	0·99	1	0·99
20	0·57	1	0·57	0·87	1	0·87	0·98	1	0·98	0·79	1	0·79	0·94	1	0·94	0·97	1	0·97

Note: CR—the correlation ratio; HR—the hit ratio; E—effectiveness; t—tolerance level; r—codebook radius.

5. Limitations and future works

5.1. Limitations

The proposed scheme is based on the causality graph model, which is deterministic in nature. In order to identify real disorders, enough manifestations have to be collected over a period of time. Thus, this scheme is not suited for real-time diagnosis. The second limitation associated with the proposed scheme is the one-manifestation disorder problem as described in Section 3.2. The 'dummy' manifestation used to remedy this problem causes overheads in both creating the knowledge base and triggering alarms.

5.2. Future works

The proposed scheme does not take into account the uncertainty relationship between disorder and manifestation. Therefore, further investigation into this problem is required. Probabilistic reasoning[8–12] and artificial intelligent[1–3] techniques could be incorporated into both the selection and identification processes.

6. Conclusions

In this paper, an event correlation scheme for fault identification in communication networks is proposed. It is composed of two processes, the selection process and the identification process. The selection process selects candidate disorders. The identification process identifies real disorders. This scheme is based on the algebraic operations of sets. The causality graph model is used to describe the cause-and-effect relationships between network events. Prime numbers are used to represent disorders and manifestations. A simulated model was built to evaluate the proposed scheme. The efficiency and effectiveness of the proposed scheme can be easily verified by simulation results. This scheme has the following merits: it can identify multiple disorders at one time, it is robust to noise, and its time complexity is close to a function of n, where n is the number of observed manifestations, with order $O(n^2)$.

References

1. M. Frontini, J. Griffin and S. Towers, 'A knowledge-based system for fault localization in wide area networks', in *Integrated Network Management II*, North-Holland, Amsterdam, pp. 519–530, 1991.
2. P. Hong and P. Sen, 'Incorporating non-deterministic reasoning in managing heterogeneous network faults', in *Integrated Network Management II*, North-Holland, Amsterdam, pp. 481–492, 1991.
3. J. R. Reggia, D. S. Nau and P. Y. Yang, 'Diagnostic expert systems based on a set covering model', *Int. J. Man-Mach Stud.*, **19**, 437–460 (1983).
4. L. Lewis, 'A case-based reasoning approach to management of fault in communication networks', *Proc. Conf. AI Appl., IEEE*, pp. 114–120, 1993.
5. A. T. Bouloutas, G. Hart and M. Schwartz, 'Identification of finite state machine using unreliable partially observed data sequences', *IEEE Trans. Commun.*, **42**, 523–533 (1992).
6. I. Rouvellou and G. W. Hart, 'Automatic alarm correlation for fault identification', *Proc. IEEE INFOCOM 95 The Conf. on Computer Commun.*, 553–561, 1995.
7. C. Wang and M. Schwartz, 'Fault detection with multiple observers', *IEEE/ACM Trans. Networking*, **1**(1), 48–55 (1993).
8. R. H. Deng, A. A. Lazar and W. Wang, 'A probabilistic approach to fault diagnosis in linear lightwave networks', *IEEE J. Select. Areas Commun.*, **11**(9), 1438–1448 (1993).
9. G. Jakobson and M. D. Weissman, 'Alarm correlation', *IEEE Network*, 11, 52–59 (1993).

10. I. Katzela and M. Schwartz, 'Schemes for fault identification in communication networks', *IEEE Trans. Commun.*, **3**, 753–764, 1995.
11. J. Pearl, *Probabilistic Reasoning in Intelligent Systems*: *Networks of Plausible Inference*, Morgan Kaufman Publishers, Los Altos, CA, 1997.
12. C. Wang and M. Schwartz, 'Identification of faulty links in dynamic-routed networks', *IEEE J. Select. Areas Commun.*, **11**, 1449–1460 (1993).
13. S. A. Yemini *et al.*, 'High speed and robust event correlation', *IEEE Commun. Mag.*, **34**(5), 82–90 (1996).
14. H. B. Enderton, *Elements of Set Theory*, Academic Press, California, 1977.

Authors' biographies:

Chi-Chun Lo was born in Taipei, Taiwan, Republic of China, on 22 August 1951. He received the BS Degree in mathematics from the National Central University, Taiwan, in 1974, the MS Degree in computer science from the Memphis State University, Memphis, TN, in 1978, and the PhD Degree in computer science from the Polytechnic University, Brooklyn, NY, in 1987. From 1981 to 1986, he was employed by Bell Laboratories, Holmdel, NJ, as a Member of Technical Staff. From 1986 to 1990, he worked for the Bell Communications Research as a Member of Technical Staff. Since 1990, he has been with the Institute of Information Management, National Chaio-Tung University, Taiwan, and is now an Associate Professor. He served as the Director of the Institute from 1994 to 1996. His major current research interests include network design algorithm, network management, network security, and multimedia system.

Shing-Hong Chen was born in Tainan, Taiwan, Republic of China, on 17 July 1963. He received the BS Degree in applied mathematics from the Chung-Cheng Institute of Technology, Taiwan, in 1986, and the MS Degree in resource management from the National Defense Management College, Taiwan, in 1992. Currently, he is working on the PhD Degree, in the area of fault management for communications network, at the National Chiao-Tung University, Taiwan. His research interests include network management, algorithm, and data compression.

