

Continuous Network Monitoring for Fast Detection of Performance Problems

Hassan Hajji and Behrouz Hodayoun Far
 Department of Information and Computer Science
 Saitama University,
 Saitama 338-8570, JAPAN
 e-mail: {hajji, far}@cit.ics.saitama-u.ac.jp

Keywords — Performance problems detection, network model, residual generation, online change detection, MIB-II information base

Abstract — This paper addresses the problem of network monitoring, for fast detection of performance problems. The network behavior is modeled as clusters of dependent objects of the Management Information Base-II (MIB-II). Each cluster is modeled as finite mixture of simple regression models. Network baseline parameters are identified from routine operation data. An online residual generation method, based on successive parameter identification, is introduced. Residuals are shown to be stationary, with mean zero under normal operation. Performance problems are characterized by sudden jumps in the mean. Detection is formulated as an online change point problem, where the task is to process residuals sequentially and raise alarms as soon as anomalies occur. An analytical expression of false alarm rate allows us to choose the threshold, automatically. Experimental results on a real network showed that the monitoring agent is able to detect even slight changes in the characteristics of the network, while maintaining a low alarm rate.

1 INTRODUCTION

Networks and distributed processing systems have become an important substrate of modern information technology. The rapid growth of these systems throughout the workplace has given rise to a discontinuity in expertise of human operators to manage them. There is a need for automating the management functions to reduce network operations and management cost.

Detection of network problems is a crucial step in automating network management. It has a direct impact on the accuracy of fault, performance and security management functions. From a control viewpoint, well designed fault and performance problems detection algorithms enhance the network control capability, by providing timely indication of network incipient problems. The possibility of early detection of performance degradation can alleviate the constant fire-fighting of network managers. Early warnings from the monitoring agent can trigger preventive actions, and serious and expensive outages can be avoided.

Most of the existing research assumed that the alarm generating mechanism is accurate, and network problems are given a priori [6, 5, 25]. Current practice in network management rely on user-defined thresholds for detection. Alarms are generated when some variable of interest crosses a prede-

defined threshold. Generally, the predefined value of the threshold is no more than an estimation of the normal range within which the measured feature is believed to operate. Not only there is little objective insights on how to choose these thresholds, but also there is a risk of missing subtle changes in the network state [4]. In addition, the complexity and size of current network systems makes them vulnerable to novel faults and performance degradation patterns.

The main difficulty in network anomaly detection is the lack of a general definition of what constitutes normal behavior [8]. The dynamics of the network normal operations need to be identified from routine operation data. To this end, [11] characterizes the normal behavior by different templates, obtained by taking the standard deviations of observations (typically Ethernet load and packets count), at different operating times. An observation is declared abnormal if it exceeds the upper bound of the envelope. Given the bursty nature of network traffic, the standard deviation estimates are likely to be distorted, making subtle changes in the network state go undetected. To mitigate the effect of the non-stationary nature of network traffic, [4] considered the model formed by segmenting time series obtained from Simple Network Management Protocol (SNMP) Management Information Base (MIB) objects. Observations are declared abnormal if they do not fit an auto-regressive model of the traffic inside segments. In [20] the observation are declared abnormal after a statistical test with the mean of 24-hour period sample. In these approaches, the assumption of piecewise constancy of the traffic is questionable, since the traffic burst is not sustained long enough to allow accurate estimation, and the combining scheme for correlating different alarms from the objects is ad-hoc.

In this paper, we address the problem of performance problems detection in IP-Networks, where the knowledge about the problems to be detected is not required. The emphasis is on fast detection with minimal human supervision – an important requirement for reducing potential impact of problems on network services users. We propose a model of the network operations in terms of MIB objects dependencies, and we show that the parametric characterization of this dependency can be described as a finite mixture of simple regression models. Model parameters are identified from routine operation data, using the Expectation Maximization (EM) algorithm.

A new method for residual generation, based on successive parameter identification, is introduced. The residuals are shown to be approximately multivariate Normal, with mean zero under normal operations, and sudden jumps in this mean are characteristics of abnormal conditions. The

Table 1: clusters of variables for performance problems detection

Dependent Variable (Y)	Independent Variable (X)
ifInNUcastPkts	
ipInReceives	ifInPkts for all interfaces
ipInDelivers	ipInReceives + ifInPkts (for loopback interface)
ipForwDatagrams	ipInReceives - ipInDelivers
ifOutPkts for all interfaces	ipForwDatagrams + ipOutRequests
ifOutNUcastPkts	

detection problem is formulated as a *change point problem*. A real-time online change detection algorithm is designed to process, sequentially, the residuals and raise an alarm as soon as the anomaly occurs. We motivate this formulation through a real problem scenario that occurred in Saitama university network. The proposed approach requires neither the set of faults and performance degradation nor the thresholds to be supplied by the user. Experimental results showed the effectiveness of the method on real data. A very low alarm rate and a high detection capability has been demonstrated.

This paper is arranged as follows: Section 2 introduces our proposed model of the network normal behavior, and the learning algorithm for parameter identification from routine operation data. Section 3 introduces our proposed approach for residual generation, and the formulation of the network problem detection. In section 4, we present results of our experiments in a real network. We conclude in section 5.

2 NORMAL OPERATIONS BASELINING

The goal of this section is to characterize network normal behavior using MIB-II [12] variables, and to identify network model parameters from routine operation data.

A Network Model

Network performance problems can be characterized as level-2 and level-3 problems [7]. Our proposed model is to define the network normal behavior model in terms of the relationships between objects at these levels. That is, instead of studying individual objects, the normal behavior of the network is defined as the parametric characterization of clusters of dependent variables in the interface and network layers. The node's view of the of the network behavior is, then, the aggregation of these models. Table 1 shows object dependencies at the network and interface level, as can be extracted from the functional requirement of TCP/IP protocol stack, or case diagrams [19].

This approach has three main benefits. First, examining the relationship between dependent variables provides a robust method for detecting network anomalies: it allows the model to interpret individual variable values in conjunction with dependent variable values. For example, the number of interface errors alone is not a clear-cut indication of network problems unless studied in relation to the total amount of inbound traffic [9]. Second, the statistical characteristics of individual variables is non-stationary, and better characterization of normal behavior can be obtained by examining explicitly variables dependency. Segmentation of individual variables, as proposed in [4, 20], overlooks the effect of variables on each other. Clearly, the segmentation of the variable

ipInReceives that counts the number of received IP packets, for instance, could be better done if all incoming packets *ifInPkts* are taken into account. Third, the specific nature of the MIB-II shows that structural change in the dependency between these variables is a symptom of an abnormal behavior. For example, if the relation between *ifInPkts* and *ipInReceives* changes, it is because of errors caused by lack of buffer space, line noise or unknown protocols packets and so on. In this sense, selected objects form a sufficient subset that needs to be continuously monitored. In addition, it is guaranteed that these variables statistics are available across most operating systems kernels. Currently, there is little support for detailed interface error statistics.

Most of operating systems share the same dependency as in Table 1, except for some systems that put loop-back packets directly in the IP input queue, rather than requiring network interface cards to read their own transmissions. We can account for this case by adding the loop-back packets to the *ipInReceives* as shown in Table 1.

B Objects Dependency Parametric Model

To be able to identify the network operation parameters from operation data, we have to define a parametric model for network operations. Our approach to network model parameterization is to view each cluster (X, Y) as switching between different regimes, where each regime is a simple linear regression model. This is a form of what referred to in the literature as *switching regression* [15, 13]. The observations (x_i, y_i) are generated by one of the K linear regression, as shown in the following equations:

$$y_i = b'_k x_i + \epsilon_{ki} \quad k = 1, \dots, K \quad (1)$$

$$p(y_i | x_i) = \sum_{k=1}^K \frac{\pi_k}{\sqrt{2\pi}\sigma_k} \exp \frac{-(y_i - b'_k x_i)^2}{2\sigma_k^2} \quad (2)$$

For the case of single variables in Table 1, such as *IfInNUcastPkts*, we have:

$$y_i = m_k + \epsilon_{ki} \quad k = 1, \dots, K \quad (3)$$

$$p(y_i) = \sum_{k=1}^K \frac{\pi_k}{\sqrt{2\pi}\sigma_k} \exp \frac{-(y_i - m_k)^2}{2\sigma_k^2} \quad (4)$$

The errors ϵ_k are assumed to be Gaussian, with mean 0 and variance σ_k . The column vector b_k is made of the slope and the intercept for the regime k . Abusing the notation slightly, the integer K denotes the number of regimes, for both the mixture of regression and Normal distributions. Each regime k has a mixing probability, denoted by π_k .

Finite mixture models for network operation baselining captures both clusters that may be described *parsimoniously* by linear relationships, and more generally non-linear dynamics of any given cluster. In fact, finite Gaussian mixture models are general enough to approximate any continuous function with a finite number of discontinuities, under appropriate regularity conditions [23]. For any cluster of variable (X, Y) that are linearly related, or even locally linear with slowly time-varying parameters, an adaptive algorithm with suitably chosen *forgetting factor* can track the model parameters. In general, however, breaks in the linear relationship are normal, and the idea then is to explicitly accommodate these breaks in the network model. The resulting model is, then, a

mixture of regimes, where each regime describe a given mode of network operations.

The network normal behavior is then characterized by the parameters of the finite mixture model. In the next subsection, we show how these parameters are identified.

C Learning Model Parameters

Identifying the network normal operations from routine operation data amounts to estimate the parameter θ of the switching regression. The vector θ consists of the vectors b_1, \dots, b_K , the variances $\sigma_1, \dots, \sigma_K$ and the mixing probabilities π_1, \dots, π_K . Given a training set of N independent and identically distributed data points (x_i, y_i) , the Maximum Likelihood (ML) estimator is the vector $\hat{\theta}$ that maximizes the likelihood function $L(\theta)$, given by:

$$\hat{\theta} = \arg \max_{\theta} L(\theta) \quad (5)$$

$$L(\theta) = \sum_{i=1}^N \sum_{k=1}^K \pi_k f_{ik} \quad (6)$$

$$f_{ik} = \frac{1}{\sqrt{2\pi}\sigma_k} \exp\left(-\frac{(y_i - b'_k x_i)^2}{2\sigma_k^2}\right) \quad (7)$$

Viewing the separation variable that assigns each observation to its corresponding regime as missing, the estimation can be formally identified with the Expectation Maximization (EM) algorithm [3]. The EM algorithm estimate the unknown parameter θ by iteratively maximizing the expected log likelihood function. For each iteration, the Estimation Step (E-Step) calculates the weights of each observation, with respect to each of the regimes, and the Maximization Step (M-Step) calculates a refined estimates of the parameters [2, 16].

The EM algorithm is numerically stable, in the sense that the incomplete data likelihood is increased at each iteration [2]. Also, since the Maximization Step (M-Step) relies on the complete data, the maximization in this step is analytically solvable. In fact, in our case the M-Step is simply the weighted least squares for the case of switching regression. Similar reasoning can be done for the case of finite Normal distributions mixtures.

3 ONLINE PROBLEMS DETECTION

In this section we discuss how the residuals are generated, and our formulation of the problem of anomaly detection.

A Residual Generation

The residual generation method we propose is based on the algorithm for parameter estimation. After convergence of the learning algorithm (Section C), we keep updating the slopes of the models, assuming that only one iteration of the EM algorithm is used for each new observation (x_n, y_n) . It can be shown, that:

$$b_k^n = b_k^{n-1} + (X_n^T W_{kn} X_n)^{-1} w_{kn} x_n (y_n - b_k^{n-1} x_n) \quad (8)$$

For the case of mixture of Normal distributions, the mean m_k of each regime k is updated as follows:

$$m_k^n = m_k^n + \left(\sum_{i=1}^n w_{ki}\right)^{-1} w_{kn} (y_n - m_k^{n-1}) \quad (9)$$

Where b_k^n and m_k^n are the estimate of the slope parameter b_k and mean m_k at iteration n , respectively. X_n is a $n \times 1$




Figure 1: comparison of the behavior of the drifts $(b_k^n - b_k^{n-1})$ and $(b_k^n - \hat{b}_k)$ under the same network conditions

vector made of the regressor x_i , and X_n^T is its transpose vector. W_{kn} is the $n \times n$ diagonal matrix, made of the weights w_{ki} . The initial parameters b^0 and μ^0 are the estimates obtained using the batch EM (Section C). It is worth nothing to note that similar results to Equations (8) and (9) can be obtained following the recursive parameter estimation proposed in [21], where parameters are updated online for each new observation.

Under normal conditions, the difference between the successive values b^n and b^{n-1} is expected to fluctuate around zero. This difference should not drift constantly in a fixed direction. On the other hand, if this difference drifts systematically over long duration, then the new observations are generated by a different model, and the recursion in Equation (8) will alter the parameter b to its new value. The idea, then, is to generate the residuals based on the K-variate random variable $(b^n - b^{n-1})$. The mean value of this difference is a good indicator of the health of the network. The same arguments can be done for the parameter m .

There is two major advantages of the residuals generated this way. First, successive identification of the parameters allows the model to adaptively track local changes in its parameters. It is unrealistic to assume that the model parameters will remain exactly the same over all the operating times of the network. Second, the difference $(b^n - b^{n-1})$ does not depend on the "true" value of the parameter b . This is very important since, in practice, we do not know this "true" value, and the only available information is the value \hat{b} , estimated from the data. Approximating b with \hat{b} , and studying the difference $(b^n - \hat{b})$ is possible, but our experiments showed that this approach is inefficient. Figure 1 compares both differences for a duration of one hour under the same network conditions. Results are shown only for one of the two regimes of the cluster, denoted by k . It can that the difference $(b_k^n - \hat{b}_k)$ is not symmetric around zero, while the difference $(b_k^n - b_k^{n-1})$ is both symmetric and very close to zero under normal conditions.

Under appropriate regularity conditions (see Appendix), we can show that the K-variate residuals e_n given by:

$$e_n = (b^n - b^{n-1})^T \Lambda^{-1} (b^n - b^{n-1}) \quad (10)$$

$$\Lambda = \text{diag} \left(\frac{\sqrt{w_{kn} x_n \hat{\sigma}_k}}{X_n^T W_{kn} X_n} \right) \quad (11)$$

For the case of mixture of Normal distributions, the residuals e_n are given by:

$$e_n = (m^n - m^{n-1})^T \Lambda^{-1} (m^n - m^{n-1}) \quad (12)$$

$$\Lambda = \text{diag} \left(\frac{\sqrt{w_{kn} \hat{\sigma}_k}}{\sum_{i=1}^n w_{ki}} \right) \quad (13)$$

are approximately Normal, with mean zero under network normal conditions. Approximating the variance matrix of e_n with Λ involves an informal argument about the asymptotic distribution of the univariate residuals e_{kn} (see Appendix). Note that e_n given in Equation (10) (respectively Equation (12)) is simply the difference $(b^n - b^{n-1})$ (respectively $(m^n - m^{n-1})$), scaled such that its variance-covariance matrix becomes Identity. Figure 2 shows the behavior of the residuals e_{kn} under the same network conditions as in Figure 1. It can be seen that these residuals are stable, and their mean is very close to 0.

In summary, network operations are characterized by the distribution the residuals e_n . We showed that, under normal operations, the residuals e_n are approximately Normal with mean zero and variance Identity matrix. The next section shows the behavior of the residuals under abnormal conditions, and how we formulate and solve the detection problem.

B Anomaly Detection

Anomaly detection is determining the discrepancy between the normal behavior and the predicted behavior. Figure 3 shows the behavior of the residuals generated by the model under a real abnormal condition that affected Saitama university network, due to badly formatted packets. As shown in Figure 3-a, this abnormal condition causes a sudden jump in the mean of the residuals. Figure 3-b shows the behavior of the residuals just before the sudden jump in the mean. Interestingly, we notice that the sudden jump is preceded by a slight change in the mean of residuals. If the detection approach is designed to be sensitive to slight changes in the operating characteristics of the network, we could have predicted the problem of Figure 3 at least 19 minutes before it became serious. The problem could have been avoided, or at least addressed immediately after its occurrence. In general, however, not all problems presents signs to allow their prediction. In this case, we require our detection method to raise alarm as soon as change in the mean occurs.

Consider the residuals E_c^n obtained by observing sequentially the residuals e_i from time point c to n . Under the normal operations of the network, the sample of e_n follows a K-variate Normal distribution with mean θ_0 as 0 and Identity variance-covariance matrix (Section A). At some unknown time point c , a change happens in the model, and the new generated residuals shift to a new distribution, with a different mean, denoted by θ_1 . The goal is to find a *decision function* and a *stopping rule* that detects this change and raise an alarm as soon as possible, under a controlled false alarm rate. This formulation is known in sequential analysis literature as the *disruption problem*. The main difference with classical hypothesis testing is that the sample size is a function of the observations made so far (i.e. not fixed a priori), and the distribution of the residuals is known, when the process being monitored, is in control. The goal is to achieve fast detection of change, by using no more than the sufficient sample size to decide whether an alarm is to be raised or not.

It is well-known that for known probability distribution after change, Page-Lorden cumulative sum (CUSUM) [14, 10] test is optimal, in the sense that it minimizes the delay to detection, among all tests with a given false alarm rate. However, in the present case of network anomaly detection, we do not have a priori knowledge about the probability distribution after change P_{θ_1} , and the change point c . The common extension of Page-Lorden CUSUM test consists of estimating the post-change distribution mean, and the change point from the data. This approach is known as the Generalized Likelihood Ratio (GLR) test [1]. That is, for the unknown parameter θ_1 of the distribution of $P_{\theta_1}(e_i)$ after change, and the change point c are estimated from data, using the maximum likelihood estimator. The resulting decision function is given by:

$$R_n = \sup_{1 \leq c \leq n} \sup_{\theta_1} \ln \frac{P(E_c^n | \theta_1, c)}{P(E_c^n | \theta_0)} \quad (14)$$

$$T_n = \inf \{n : R_n > \lambda\} \quad (15)$$

In our case, where pre-change and post-change distributions are Normal, the maximization problem of Equation (14) can be worked out explicitly. It has a simple form, given by:

$$S_0 = (0, \dots, 0)^T \quad S_n = \sum_{i=1}^n e_i \quad (16)$$

$$T_n = \inf \{n : \max_{0 \leq c < n} \frac{\|S_n - S_c\|}{\sqrt{n-c}} > \lambda\} \quad (17)$$

The equation assumes that after change, the distribution of the residuals is still Normal, but with different mean. For the abnormal case, it is hard to obtain an unbiased fit of the post-change distribution $P_{\theta_1}(e_i)$. Fortunately such accurate estimation is not crucial. What is needed is that, when the an anomaly occurs, the closest Normal distribution, obtained by maximum likelihood estimation, has a mean significantly different from zero.

C Tuning the Threshold λ

So far we have introduced the decision function and the stopping rule used for online detection of network faults and performance degradation. The remainder of our problem setup concerns the choice of the design threshold λ .

It can be shown that the expectation of the stopping rule, under no change denoted by $E_\infty(T)$, is given by [17]:

$$E_\infty(T) \sim \frac{\Gamma(K/2) 2^{K/2} \exp(\lambda^2/2)}{\lambda^K \int_0^\lambda x v^2(x) dx} \text{ as } \lambda \rightarrow \infty \quad (18)$$

$$v(x) = 2x^2 \exp\left(-2 \sum_1^\infty n^{-1} \Phi\left(\frac{-xn^{1/2}}{2}\right)\right), x > 0 \quad (19)$$

Where Φ denotes the Normal distribution function. For calculation, see [17, 18] for an approximation of $v(x)$. Not surprisingly, Equation (18) turns out to be the mean time between false alarms. It follows that, given a desired false alarm rate, we can *recover* the design threshold λ , by solving Equation (18).

4 EVALUATION AND RESULTS

The network monitoring algorithms described earlier has been implemented in a real network. MIB-II [12] describes




Figure 2: Residuals plots under the same network conditions as in Figure (1): (a) Univariate residuals e_{kn} for regime k of the switching regression, as given by Equation (10) (b) e_{kn} for finite Normal mixtures, as given by Equation (12)




Figure 3: Residuals generated by one of the two regimes of $(ifInPkts, IpInReceives)$, under abnormal conditions: (a) abrupt change in the mean of the residuals (b) slight change in the mean of the residuals just before the sudden disruption

information objects for traffic monitoring in TCP/IP networks, and operating system kernels normally keep counters corresponding to these objects. The agent monitors the network by accessing the kernel statistics, for each of the monitored variables. Currently, the agent runs on our network file server, under Solaris 2.6 operating system. To test the agent detection capability, we implemented a program to access the underlying data-link layer for fault injection. The fault injection program runs on another machine, and it is set to inject faults, by assembling and injecting appropriate packets. The goal is to prove that alarms generated by the agent are, effectively, due to abnormal network conditions. The remainder of this section presents the results that show the monitoring agent detection capability.

A Implementation Aspects

To allow our model to track time varying parameters, we introduce an exponential forgetting factor $0 < \zeta \leq 1$, that reduces the effect of old observations, much in the same way as proposed in [24]. Evaluating the sum $\sum_{i=1}^n w_{ki} x_i^2$ (Equation (8), Section A) is then replaced by $\sum_{i=1}^n \zeta^i w_{ki} x_i^2$. Similar modification is done for finite Normal mixture parameter up-

dating (Equation (9), Section A).

Unfortunately, Equation (17) can not be written recursively. Consequently, the number of residuals to be inspected can grow large. To circumvent this difficulty, we use a moving horizon of fixed length, where the starting point of the horizon moves one step forward as new observation are made available to inspection. For the mean time between false alarms (Section C), it is fixed to 8640, which corresponds to 24-hours period, given that samples are collected every 10 seconds. Finally, for the number of components K , it was found empirically that at most 4 regimes are enough to describe the data satisfactorily. Work is underway to infer the number of components automatically from data.

B Detection Capability

The agent detection capability is first illustrated using the network problem, introduced earlier in Section B. The problem showed up as a streaming network interface card sending excessively badly formatted packets. Figure 4 shows the behavior of the residuals and test statistic as detected by the cluster $(ifInPkts, ipInReceives)$. As shown in the figure, it takes approximately 30 samples (5 minutes) to detect the




Figure 4: slight change in the mean of the residuals is detected 19 minutes before the sudden disruption

slight change in the residuals. In this particular case, the threshold is crossed 19 minutes before the sudden disruption. It could have been possible to address *proactively* the problem before it became serious, or at least draw the attention of network operators earlier, before the impact of the problem is felt by all network users. Detection capability of the agent is tested with other problems. The obtained results are shown in Figure 5. For IP packet loss and excessive outbound broadcasts, appropriately assembled packets are sent every two seconds. For the remaining problems, assembled packets are injected every one second. It can be seen that the agent could detect all of these problems, with reasonably short delay.

We note that, apart from reasons given in Section A for defining network model as the parametric characterization of dependent objects, there is no MIB counter for ARP operations. Also, for the problem of packet loss induced by assembling Ethernet packets with non-existent protocol type, some kernels do not have any entry for this type of errors, even if this object is part of the standard MIB-II information base. For IP operations, one can easily check that the difference between *ipInDelivers* and *ipInReceives* can be sometimes very large, without noticing any change in IP related errors. This is true even when we take into consideration the fact that loopback packets are to be added to *IpInReceives*. With current under-instrumented networks, our approach to network modeling and performance problems detection can provide useful insights about incipient problems.

In all the results reported here, we actually stress-tested the agent detection capability. The total amount of all injected packets is smoothed over the whole test duration. In practice, network problems are expected to induce large magnitude changes in the model residuals. The detection is expected to even work better.

C Alarm rate

The final aspect we investigate in our proposed approach is the false alarms rate. Ideally, we would like to estimate this rate, given that the network is operating normally. Unfortunately, it is difficult to gain perfect knowledge about all the subtle changes in the network behavior. Instead, Table 2 shows the average alarm rate per hour, for data collected during a period of 24-hours.

Table 2: Average number of alarms per hour for each of the clusters of the network model

Clusters	Average alarm rate per hour
<i>ifInNUcastPkts</i>	0.04
<i>ipInReceives, ifInPkts</i>	0.50
<i>ipInReceives, ipInDelivers</i>	0.08
<i>ipForwDatagrams, ipInReceive - ipInDelivers</i>	0.04 *
<i>ifOutPkts, ipForwDatagrams + ipOutRequests</i>	0.16
<i>ifOutNUcastPkts</i>	0.00

Some comments are now in order. First, we note that *ifOutNUcastPkts* is very stable, recording a zero alarm rate per 24-hours. The same comment is also true for *ifInNUcastPkts*. This raises questions about whether the switching regression model adds any accuracy to the model. It is probably enough to model all the variable by a finite mixture model. We plan to investigate this question in detail in the future. Note that even under this conditions, our method compares favorably to results in [4, 20]. Second, note that results in the Table 2 are not the minimum alarm rate that could be achieved. In times the network is almost idle, the alarm rate can be extremely low, even for long duration. In fact during experiments conducted in summer vacation (August 2000), only 8 alarms are raised for the 28956 samples collected (80 hours approximately), for the cluster *ifInPkts-ipInReceives*. This makes an average of 0.1 alarms per hour.

5 CONCLUSION

In this paper, we developed an online technique for fast detection of performance problems in IP-Networks. We proposed a model of the network operations in terms of MIB objects dependencies, and we showed that the parametric

*Results are from an internal router of Saitama University network

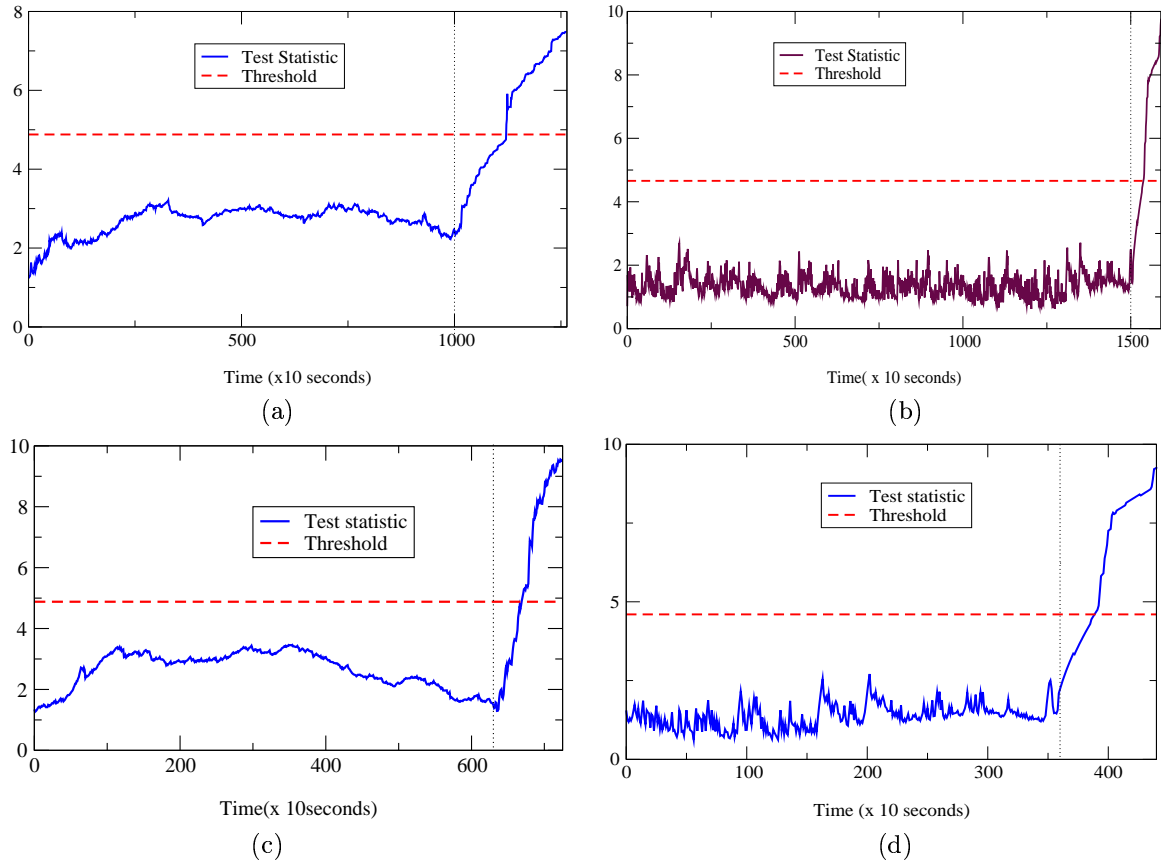


Figure 5: Behavior of the test statistic corresponding to excessive ARP packets, excessive inbound broadcasts, excessive outbound broadcasts, and IP packet loss problems: (a) IP packet discards (b) Outbound broadcast packets (c) Inbound broadcast packets (d) Inbound ARP packets

characterization of this dependency is amenable to a finite mixture model. Model parameters are identified from routine operation data, using the EM algorithm. A new method for residual generation, based on successive parameter identification, is introduced. The residuals are shown to be approximately Normal, with mean zero under normal operations, and sudden jumps in this mean are characteristics of abnormal conditions. A real-time online change detection algorithm processes, sequentially, the residuals and raise an alarm as soon as the anomaly occurs. The proposed approach requires neither the set of faults and performance degradation nor the thresholds to be supplied by the user. Experimental results showed the effectiveness of the method on real data. A low alarm rate and a high detection capability have been demonstrated.

APPENDIX

Empirical Verification of Residuals Distribution

We studied the assumption of residuals normality using multivariate quantiles, calculated after enforcing the matrix to be diagonal, and the mean vector to be 0. Figure 6 shows the plots of the ranked Mahanobolis distance d_i^2 versus the expected corresponding χ_K^2 value for the residuals. It can be concluded that the fits are sufficiently linear, and do pass through the origin, which validates the residuals normality. Results are shown only for the cluster (*ifInPkts, ipInReceives*)

and the variable *ifInNUcastPkts*

There is some theoretical evidences that our conjecture about the distribution of the residuals e_n may hold in general. First, If the number of components is one, e_n are effectively Normal with mean 0 and variance given in Equation (11). Second, in the general case, if we note by I_c and I the Fisher information matrix of the complete and observed sample, respectively, then it can be shown (see [21, 22]) that the variable b^n is asymptotically normal, given that the smallest eigenvalue of $I_c^{-1}I$ is bigger than $1/2$. It follows, then, that the mean $E(e_n) = 0$.

To prove the variance in the general case, we conjecture that the regimes are separated enough, such that each data point (x_n, y_n, w_{kn}) is described with the regime k , where the weights w_{kn} is treated as measure of appropriateness. The goal here is to support our claim that $\sqrt{w_{kn}}(y_n - b_k x_n)$ follows a Normal distribution with mean 0 and variance σ_k^2 , much in the same way as weighted linear regression. If the weight w_{ki} is zero, no residual is generated for that regime. We also assume that each component k are updated independently of other components. Asymptotically, one can reasonably assume that as the algorithm converges, the regimes gets, effectively, separated.

Under this conjecture, it easy to show that $\sqrt{w_{kn}}(y_n - b_k^{n-1} x_n)$ follows a Normal distribution with mean 0 and vari-

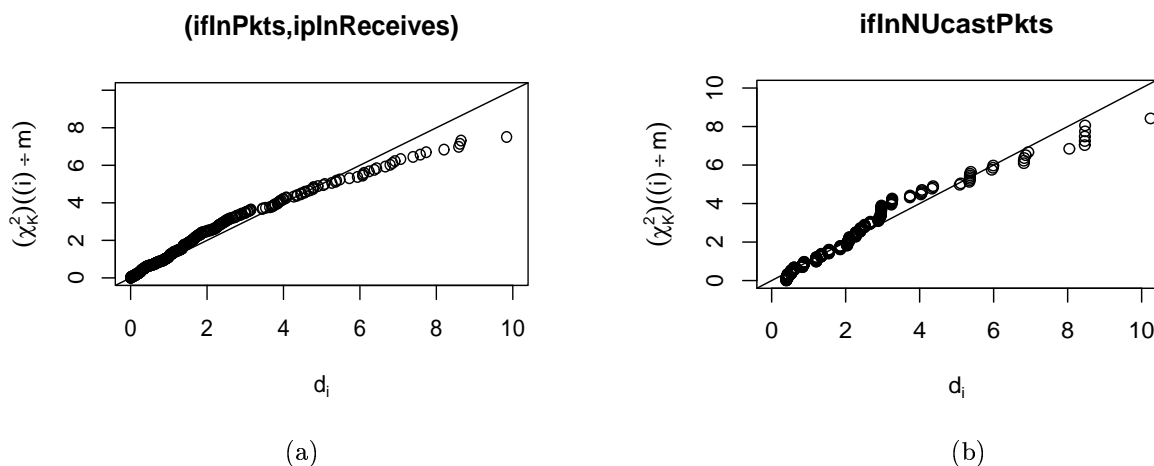


Figure 6: Ranked Mahalanobis distance (d_i^2 , where i is i -th smallest distance) versus the expected corresponding χ^2 value for the residuals (a) case of $\{iflnPkts,iplnReceives\}$, (b) case of the $iflnNUcastPkts$. All fits are both linear and pass through the origin

ance:

$$\sigma_k^2(1 + w_{kn}x_n(X'_{n-1}W_{k(n-1)}X_{n-1})^{-1}x_n) \quad (20)$$

The term $w_{kn}x_n(X'_{n-1}W_{k(n-1)}X_{n-1})^{-1}x_n$ tends quickly to 0, so it is safe to ignore it. We see, from Equation (8), that the variance of $(b_k^n - b_k^{n-1})$ is given by:

$$Var(b_k^n - b_k^{n-1}) = (X'_n W_{kn} X_n)^{-2} w_{kn} x_n^2 \sigma_k^2 \quad (21)$$

It follows that e_n are Normal with mean 0 and variance Identity matrix. Similar arguments can be made for the case of Normal mixtures.

References

- [1] M. Basseville and I. V. Nikiforov. 1993. *Detection of Abrupt Changes : Theory and Application*, Prentice-Hall.
- [2] A. Dempster, N. Laird and D. Rubin. 1977. "Maximum Likelihood from Incomplete Data via the EM Algorithm", *J. R. Statist. Soc. Series B*, Vol. 39: 1-38.
- [3] M. J. Hartley. 1978. Comment on "Estimating mixtures of Normal Distributions and Switching Regressions", *Journal of the American Statistical Association*, Vol. 73: 738-741.
- [4] C. S. Hood and C. Ji. 1997. "Proactive Network Fault Detection," in *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, (Kobe, Japan), April, 1147-1155.
- [5] G. Jakobson and M. D. Weissman. 1993. "Alarm Correlation", in *IEEE Network*, Vol.7, No.6, November: 52-59.
- [6] I. Katzela and M. Schwarz. 1995. "Schemes For Fault Identification in Communication Networks", *IEEE/ACM Transactions on Networking*, Vol. 3, No. 6: 753-764.
- [7] S. Keshav and R. Sharma. 1998. "Achieving Quality of Service through Network Performance Management", in *Proceeding of NOSSDAV '98*.
- [8] L. LaBarre. 1991. "Management by Exception: OSI event generation, reporting, and logging", in *Proceedings of Second International Symposium on Integrated Network Management*.
- [9] A. Leinwand, K. Fang Conroy. 1996. *Network management, a practical perspective*. 2nd Edition. Addison-Wesley.
- [10] G. Lorden. 1971. "Procedures for reacting to a change in distribution", in *Annals of Mathematical Statistics*, Vol.42: 1897-1908.
- [11] R. A. Maxion and F. E. Feather. 1990. "A Case Study of Ethernet Anomalies in a Distributed Computing Environments", *IEEE Transactions on Reliability*, Vol. 39, No. 4: 433-443.
- [12] K. McCloghrie, M. Rose. 1991. "Management Information Base for Network Management of TCP/IP-based internets: MIB-II", RFC 1213.
- [13] G. J. McLachlan, and K. E. Basford. 1988. *Mixture Models: Inference and Application to Clustering*. New York: Dekker.
- [14] E. S. Page. 1954. "Continuous Inspection Schemes", in *Biometrika*, Vol. 41: 100-115.
- [15] R. E. Quandt. 1972. "A New Approach to Estimating Switching Regressions", in *Journal of the American Statistical Association*, Vol. 67, No.338: 306-310.
- [16] R. A. Redner and H. F. Walker. 1984. "Mixture Densities, Maximum Likelihood and the EM Algorithm", in *SIAM Review*, Vol. 26, No. 2: 195-239.
- [17] D. Siegmund and E. S. Venkatraman. 1995. "Using the Generalized Likelihood Ratio Statistic for Sequential Detection of a Change Points", in *The Annals of Statistics*, Vol. 23, No.1: 255-271.
- [18] D. Siegmund. 1985. *Sequential Analysis: Tests and Confidence Intervals*. Springer, New York.
- [19] W. Stalling. 1999. *The SNMPv1, SNMPv2 and RMON*. Addison-Wesley, Reading, MA, USA.
- [20] M. Thottan and C. Ji. 1998. "Proactive Anomaly Detection Using Distributed Intelligent Agents", *IEEE International Workshop on Systems Management*.
- [21] D. M. Titterton. 1984. "Recursive Parameter Estimation using Incomplete", in *J. R. Statist. Soc. Series B*, Vol.46, No.2: 257-267.
- [22] D. M. Titterton, and J-M, Jiang. 1983. "Recursive Estimation Procedures for Missing-data problems", in *Biometrika*, Vol.70, No.3: 613-624.
- [23] N. A. Vlassis. 1999. "A Kurtosis-Based Dynamic Approach to Gaussian Mixture Modeling", in *IEEE Transactions On Systems, Man, And Cybernetics, Part A*, Vol.29 No.4: 393-399.
- [24] E. Weinstein, M. Feder and A. V. Oppenheim. 1990. "Sequential algorithms for parameter estimation based on Kullback-Leibler information measure", in *IEEE Trans. Acous., Speech, Signal Processing*, Vol.38, No.9: 1652-1654.
- [25] S. Yemini, S. Kliger, E. Mozes, Y. Yemini, D. Ohsie. 1996. "High speed and robust event correlation", *IEEE communication Magazine*, May: 82-90.