Proactive Network Maintenance Using Machine Learning

R Sasisekharan V Seshadri S M Weiss

AT&T Bell Laboratories, HR2K033

480 Red Hill Road

Middletown, N J 07748

sesh@eagle.hr.att.com

ABSTRACT

A new approach to proactively maintain a massively interconnected communications network is described. This approach has been applied to the detection and prediction of chronic transmission faults in AT&T's digital communications network. A windowing technique was applied to large volumes of diagnostic data, and these data were analyzed by machine learning methods. A set of conditions has been found that is highly predictive of chronic circuit problems, that is, problems that are likely to continue in the immediate future without diagnosis and repair. In addition, a few conditions have been found that are predictive of problems that affect multiple circuits. Such analyses over the complete network are helpful in proactively maintaining the network and in spotting trends for circuit problems. Proactive maintenance of the network can help in greatly improving the quality and reliability of a network by identifying potentially serious problems before they occur.

KEYWORDS

communications network, rule induction, proactive maintenance, machine learning, database mining

INTRODUCTION

With the increasing complexity of modern communications networks, there is a commensurate need for intelligent systems to help manage and maintain them. Most desirable are systems that can analyze and resolve problems in the network automatically, thus greatly improving the reliability and quality of the network. Artificial intelligence(AI) techniques have proven useful in building network operation systems[1]. An area in which such techniques can have significant impact is in the proactive maintenance of networks. Proactive maintenance of the network can help in greatly improving the quality and reliability of a network by identifying potentially serious problems before they degrade. This can be accomplished by monitoring network performance over time and spotting trends in problems. In addition, monitoring network performance can help in prioritizing network problems. The ability to prioritize network problems can significantly improve the quality of the network since those problems that have the most significant impact on the operations of the network can be addressed first.

Monitoring network performance involves analyzing extremely large amounts of diagnostic data that varies with time. Looking for patterns of behavior in such large volumes of data can only be accomplished by computer analysis. In this paper, we present an approach to do such an analysis using machine learning techniques. We illustrate the approach by specifically looking at detection and prediction of chronic transmission faults in AT&T's world-wide digital communications network.

Even though this paper focuses on telecommunications networks where many homogeneous network components exist, our work is applicable to other types of networks as well. In the following sections, we describe the application domain, explain the approach that we used to enable the machine to learn from time-dependent problems in the network, apply alternative methods of machine learning to our problem, and report the results we have obtained.

PROBLEM DESCRIPTION

Network operation systems (NOS) exist in the network to support provisioning, maintenance, operation, administration and management functions for the network and for individual network components. A circuit can be considered as a path in the network, which contains network components and links. Transmission problems on a circuit are seen by several of the network components through which the circuit connects. In a large network, such as AT&T's communications network, the ratio, of diagnostic data generated by various network components to the root problem that is responsible for them, is very large.

The different types of problems in the network can be broadly categorized into two classes, transient and non-transient. Transient transmission problems are very common in the network, yet their behavior and causes are not completely understood. Part of the difficulty in understanding them is related to separating the wheat from the chaff, that is, in learning to ignore glitches that will not be repeated and focussing instead on those transient problems that will recur (chronics). Chronics not only affect the quality of communications while they recur but also indicate degradation and potential future failures in the network. Thus it is an important and challenging problem to identify these chronics and isolate their causes.

Diagnostic procedures that attempt to resolve transient problems must rely on large volumes of historical information and a more complex analysis of patterns of behavior. One novel approach to diagnosing transient faults is found in an AT&T system known as SCOUT[2]. Using historical and topological information, SCOUT finds specific related circuits that share common patterns of faulty behavior. Typically, these are difficult transient problems, multiple circuit problems, or even forms of chronic faulty behavior. In this paper, we consider a related form of analysis of chronic behavior. We consider the performance of the *complete* AT&T network over *time*. Our objective is to determine whether there are patterns of behavior over the network such that the following predictions can be made:

- The faulty behavior will continue in the immediate future.
- The faulty behavior involves multiple circuits.

We do not attempt to solve specific problems. Instead, our objective is to determine whether there are any signature characteristics for those network problems that do not get fixed. We may not have enough data to determine the exact nature of the problem, but we may be able to predict accurately that the problem will not be fixed quickly during the normal course of operations. In order to achieve maximum reliability of the network, problems with these signatures must be considered, and those problems that are deemed critical must be addressed quickly.

From a computer science perspective, interest in this area centers around a number of issues related to mapping time-dependent events into a standard classification framework. These include the following:

- Defining classes to convey the concept of "chronic."
- Defining features that summarize historical events.

In addition, the sample size in this analysis numbers in the tens of thousands of cases. Looking for patterns of behavior in such large volumes of data can only be accomplished by computer analysis using machine learning techniques, possibly resulting in new information that cannot be obtained by typical human experience.

THE MACHINE LEARNING APPROACH

In order to identify chronics that have the potential to degrade the performance of the communications network, we adopted a computer-based approach of learning from historical data. The machine learning methodology is described in this section.

Describing the Goals and Measurements

The circuit-related questions that we have outlined in the previous sections need to be posed in a standard classification format, so that a number of interesting analytical techniques that are available can be applied. Prediction models that can be applied to a standard classification problem include decision trees, decision rules, statistical linear discriminants, neural nets, and nearest neighbor methods. In the standard classification format, samples of cases are obtained. For each case, identical measurements are taken, and at least one of these measures is the class label. Methods are applied which attempt to find patterns for one class that differ from other classes.

For our first problem, the class label is chronic failure on a circuit, a concept that has been defined in previous sections. The goal is to predict that current failures will continue to occur. We must also take into account that these failures are often transient, and that failures will likely not occur continuously in the future. Instead, a failure may occur in the future, but the occurrence may also be transient. Periods of time that are reasonably close to the current period are of interest.

The measurements that are used for prediction must summarize historical information. These measurements are recorded each time a fault occurs. Not all measurements are recorded for every fault, only those that directly measure the fault process. Faults are often transient, so the trends for a period of time must be measured. It is quite possible that many faults will occur for a short time, but these faults are not necessarily chronic. They can be fixed and do not reappear in the immediate future. Measurements must be specified that are useful in predicting the target concept, that is, future failures on the circuit.

This time-dependent problem was mapped into a standard classification format by the use of fixed time windows. Historical information for circuits was examined over a consecutive period of time, and this time period was divided into two windows, W_a and W_b. The objective was to use the measurements made in Wa to predict that problems will occur in Wb. We considered both our knowledge of the application and experimental data to arrive at reasonable sizes for each of the two windows. The windows were also divided into sub-units based on time, which we will refer to as a ***time unit***. We will refer to the size of W_a as T_a time units, and the size of W_b as T_b time units.

There are many reasonable measurements that can be taken over time. Assuming a fault occurs, an alarm or exception is noted. Included in the possibilities of measurements are the number of times such an event occurs, the average number of times an event occurs during a time unit, or the number of time units during which the event occurs. We defined around 30 performance features for this problem, based on the variation of diagnostic data over time and variation over space.

In addition to the time periods for the windows, another factor must be considered in defining the conditions for each window. This is the degree of chronic failure. For the two windows, W_a and W_b, the conditions that were chosen were different, as follows:

- W_a: any fault during the time period T_a.
- W_b: faults during at least half the time units of the following time period T_b.

The rationale for these periods is as follows. We must consider a prior period sufficiently long that prediction of continuing chronic behavior is feasible. We considered all circuits where a fault has occurred during Wa. Because faults are often transient, we must specify a reasonably long period for W_b. A fault that recurs over at least half the time units during W_b would be of interest because it indicates a clearly continuing and unresolved problem.

In our learning experiments, we examined all circuits with problems during predefined intervals. It must be emphasized that the predictions to be made are those for continuing chronic failures. These are not necessarily the most obvious or acute problems, which are often diagnosed and fixed quickly.

The second question that was examined is whether any patterns of measurements are indicative of multiple circuit involvement. We refer to failures, occurring on multiple circuits due to a common cause, as multiple circuit problems. We rely on SCOUT's analysis to determine if the problem on a circuit is a multiple circuit problem. The question that we hoped to answer is whether there are certain types of faults within the complete network that consistently suggest multiple circuit problems.

Learning from Data

Once sample data are obtained, several computer-based techniques can be used to make predictions. Among the more prominent techniques are:

- statistical linear discriminants[5],
- neural nets[6],
- decision rules or trees[7][8], and
- nearest neighbor methods[9].

Both neural nets and linear discriminants make predictions based on weighted functions.The linear discriminant uses a simple additive scoring function, for example

$$\text{If } a*X + b*Y > 3, \text{ choose Class 1.} \tag{1}$$

The neural net, on the other hand, can model more complex decision functions, typically non-linear functions.

Decision trees and decision rules pose solutions in the form of true or false conditions. For example,

$$\text{If } A>10 \text{ AND } B<20\text{, choose Class 1.} \tag{2}$$

A decision tree is a stylized model of decision rules, where every decision flows from the root node of the tree, and each path is mutually exclusive. More information about these techniques can be found in Weiss and Kulikowski[4].

For our application, we had a large number of samples obtained from the operating communications network. It was not feasible to try every variation of these methods on the entire set of samples. Instead, we performed some smaller experiments to see whether one approach offered an advantage over the others. Overall, for this application, nearest neighbor methods and statistical linear discriminants performed poorly. Neural nets and decision rules or trees were competitive, with a small edge for decision rules.

In experiments on the complete data sets, representing all circuit problems in the network over a fixed period, we relied mostly on rule induction. Rule and tree induction methods have been extensively described in published works[4]. In our study, we emphasized a rule induction technique called SWAP-1[3]. Rule induction methods attempt to find a compact covering rule set that completely separates the classes. The covering set is found by heuristically searching for a single best rule that covers cases for only one class. Having found a best conjunctive rule for a class C, the rule is added to the rule set and the cases satisfying it are removed from further consideration. This process is repeated until no cases remain to be covered. Unlike decision tree induction programs and other rule induction methods, SWAP-1 has the advantage that it uses optimization techniques to revise and improve its decisions. Once a covering set is found that separates the classes, the induced set of rules is further refined by either pruning or statistical techniques. Using train and test evaluation methods, the initial covering rule set is then scaled back to the most statistically accurate subset of rules.

It is quite possible that tuning many of the alternative methods could result in somewhat improved results for each method. However, based on our knowledge of the application, there are a number of reasons why the rule induction method appears most appropriate:

- The objective is to extract new information from the data. The hope is that we can gain insight into the performance of the network. Decision rules have the strongest explanatory capabilities of the cited models.
- We know in advance that this is a noisy environment. Perfect classification can be achieved on all samples only when chronic behavior is entirely consistent. This is not likely with all the efforts toward high reliability and the transient nature of many problems. Thus the expectation is to find a subset of conditions that are highly reliable predictors of chronic failure.
- Most of the measurements are counts of the number of time units during which an event occurred over a predefined time period. Thus these measurements are ordered discrete variables. They are not continuous. Patterns of these types of measurements are usually effectively described in terms of the greater than or less than operators which are used by the rule induction model.

- The minimum error solution is not necessarily the best solution. Because we are trying to extract new information, the preferred solution consists of the highest predictive rules even if they cover fewer cases. The preference is also for simple rules that enhance our understanding of network performance.

Given that the expectation is for predictions that cover only a partial number of chronic problems, decision rules most naturally model the partitioning of data. The efficacy of the individual rules can then be tested on independent test data.

Testing the Decision Model

The central method for building a predictive model is to learn from samples and test on independent data. In many applications there is a relative shortage of data. In these situations, a compromise is made by randomly partitioning the data into training and test sets. In our application, we had a large number of samples. Training was performed on a random subset of data for a time period, and some preliminary testing was done on the remaining data. Once a solution was found, further rigorous testing was performed by testing the solution on additional data from subsequent time periods. While it is traditional to train on two-thirds of the data and test on the remaining data, the goal of finding simpler, more predictive rules may sometimes be achieved by training with fewer cases. These are rules that cover fewer cases, but are more predictive for the cases that they cover.

RESULTS

In this section, we describe the results that were obtained, both in identifying chronic problems and in classifying problems associated with multiple circuits. We have also provided the results of our comparison of various learning methods as applied in this case.

Alternative Methods for Learning

Although we concluded that rule induction was the preferred learning method for this application, we performed several experiments to evaluate its competitiveness with alternative learning models. Table 1 summarizes the results. One third of the cases were randomly selected for testing, and the error rates in the table are based on the test case performance.

Simply choosing the largest class gives an error rate of 6.4%. The linear discriminant (Fisher's) that we used was the standard parametric discriminant found in statistical packages. We used it with feature selection. It is not unusual that this method does not perform well in a low prevalence situation.

Nearest neighbor methods are greatly affected by noisy variables. This result is for k=1 and euclidean distance.

The three remaining techniques were relatively close for this experiment. The decision tree was induced by CART, and the decision rules by Swap-1. The neural net was a standard backpropagation network, with a single hidden layer. Configurations were considered with from 0 to 6 hidden units, with the best test error rate for 0 hidden units.

Method	Error rate (%)
Prior	6.4
Linear discriminant	6.4
Nearest neighbor	5.6
Neural net	4.7
Decision tree	4.5
Decision rules	4.4

Table 1: Comparative Results For Alternative Learning Methods

Sample Size

We examined the historical records for several months during late 1992 and 1993. These samples were taken from the complete AT&T network, and covered a significant number of all transmission problems encountered. When compared to the billions of transmissions during a month and the size of the network, the number of problems is quite small. However, from a sampling perspective, we had a large sample, consisting of tens of thousands. Of these circuits, between 5 and 10% fulfilled our definition of chronic, that is, they had faults during at least half the time units during W_b.

Predicting Chronic Circuit Problems

We now return to the central task, namely, can we predict chronic problems, that is, problems that will continue in the immediate future? We were able to generate rule sets that were predictive. The rule set reduces to a set of conditions. If any of the conditions hold, the problem is very likely to be chronic. The conditions were of the form

- $X\_Feature_i > n_i$

where $X\_Feature_i$ is a performance feature based on the number of time units during which an event of a type *i* occurs and n_i is an integer.

We will refer to a set of five conditions, of the type described above that were generated, as *Ruleset0*. Another condition of a different type was also generated. It was of the form

- $Y\_Feature_a > m_a$ & $Y\_Feature_b > m_b$ & $Y\_Feature_c > m_c$

Where Y_Feature is a performance feature that is not based on the count of time units and m is an integer.

This condition is weaker than the other conditions. We will refer to the rule set that includes all the conditions in *Ruleset0* and this last condition as *Ruleset1*. While the predictive value of *Ruleset0* is better than *Ruleset1*, *Ruleset1* covers a larger portion of the chronic problem. This is illustrated in Figures 2 and 3.

Figure 2 plots the performance of these two rule sets over the course of several time periods. The rule set was induced during a time period when the prevalence of chronic problems was somewhat lower than during other time periods. Thus any solution induced for that time period necessarily would be highly predictive to overcome the odds of the larger class. Figure 3 plots the percentage of each time period's chronic problems covered by the rule sets.

Figure 4 plots the change in predictive value versus the number of time units with faults for a representative time period. Specifically, as the definition of chronic problems that are to be detected is made more selective, the performance improves in that the number of false alarms is reduced. At the same time, of course, coverage decreases, meaning that the number of transient problems, that are no longer classified as chronic but have the potential to result in performance degradation, increases.

Multiple Circuit Detection

Using the same measurements, we consider the detection of multiple circuit problems. Whether a problem on a circuit is a multiple circuit problem or not is determined by means of data obtained from SCOUT. For this application, we trained on data from one time period and tested on data from another time period. The classes are close to equal in size with a 45% - 55% split. Two conditions emerged as particular strong predictors of multiple circuit involvement. They were of the type:

- $X\_Feature_i > n_i$

Although they cover a relatively small percentage of all circuits with problems, when either of these conditions occur the likelihood that multiple circuits are involved is estimated at 93%.

We have identified certain conditions that are common to both chronic problem detection and multiple circuit detection. Thus, we have come up with a rule that predicts when a problem with a circuit is likely to be chronic as well as affect multiple circuits.

CONCLUSIONS

If all problems in a communication network were either transient or quickly repaired, it would not be necessary to detect chronic problems. However, chronic problems do occur. Identifying patterns of these problems should be helpful in characterizing problems that are not detected and repaired quickly. In our analysis, we found that the number of time units over which events occurred was critical in determining the likelihood that the problem would continue in the future. These rules suggest a form of momentum or inertia for chronic fault problems. There are a number of rationales for the validity of this form of analysis:

- Not all faults have this momentum to the same degree. We have identified those measurements that are predictive along with the corresponding thresholds, that is, the number of time units of faults during a window for which they are predictive.
- Of particular importance, any circuit that exhibits this behavior will likely continue this behavior. Thus, if the goal is to maximize reliability, circuits exhibiting these characteristics should be given priority in diagnosis and repair.

The most immediate application of these results is their use in reordering trouble tickets. Beyond that, though, we have provided a methodology to proactively maintain and monitor the performance of a network. The determination of trends can demonstrate whether progress is made in reducing chronic problems or whether chronic problems are increasing in the network. The best results for network performance occur when no patterns emerge or when they cover a smaller percentage of the problems.

Clearly, we are limited by the types of recorded measurements. If the causes of circuit problems were eventually determined and recorded, it might be possible to explore hypotheses directly related to the cause and repair of a problem. Such information is not currently available, and because of the complexity of a network with transient problems, such records may never be fully available.

The second issue that we have addressed is the prediction of multiple circuit involvement. We found that certain types of faults are good indicators of multiple circuit problems. Resolving multiple circuit problems is particularly useful in reducing the overall number of problems in the network.

We have considered a highly complex communications network, and analyzed its behavior over time. To facilitate proactive maintenance, we have developed measurements that are sampled for the complete network during regular time periods. While our results are bounded by the predictive capabilities of these measurements, this form of analysis did produce reasonable predictors. The analysis involves intensive computer processing of very large volumes of data. In both objectivity and pattern matching capability, such efforts are beyond the capabilities of human processing and experience.

REFERENCES

[1] P. A. Corn, R. Dube, A. F. McMichael and J. L. Tsay, "An Autonomous Distributed Expert System For Switched Network Maintenance", *Proc. of IEEE GLOBECOM '88*, pages 1530-1537, 1988.

[2] R. Sasisekharan, Y-K. Hsu, and D.Simen, "SCOUT: An Approach To Automate Diagnoses Of Faults In Large Scale Networks", to appear in *Proc. of IEEE GLOBECOM' 93*, Houston, 1993.

[3] S. Weiss and N. Indurkhya, Reduced Complexity Rule Induction, Proceedings of IJCAI-91, pages 678-684, Sydney, 1991.

[4] S. Weiss and C. Kulikowski, *Computer Systems That Learn: Classification And Prediction Methods From Statistics, Neural Nets, Machine Learning, And Expert Systems*, Morgan Kaufmann,1991.

[5] M. James, *Classification Algorithms*, Wiley, 1985

[6] J. McClelland and D. Rumelhart, *Explorations in Parallel Distributed Processing*, MIT Press, 1988

[7] L. Breiman, J. Friedman, R. Olshen and C. Stone, *Classification and Regression Trees*, Wadsworth, 1984.

[8] J. Quinlan, "Induction of Decision Trees", *Machine Learning*, vol. 1, pages 81-106, 1986.

[9] D. Aha and D. Kibler, "Noise Tolerant Instance-Based Learning Algorithms", *Proc. of IJCAI-89*, pages 794-799, Detroit, 1989.

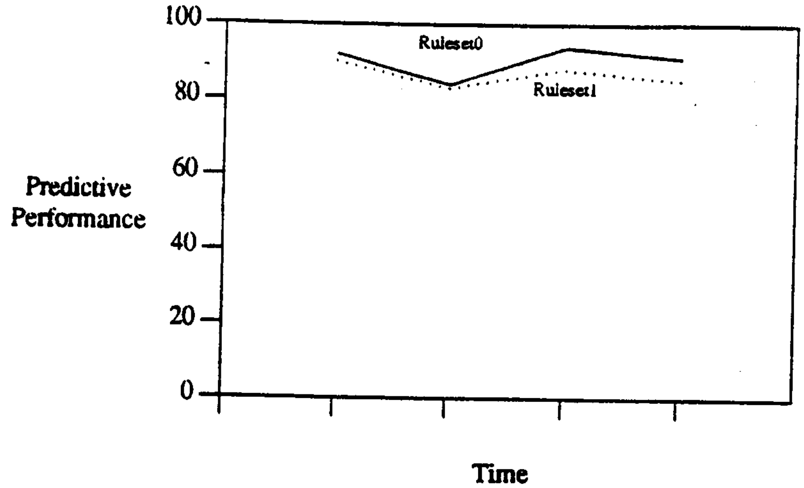

Figure 2: Predictive performance of rule sets.

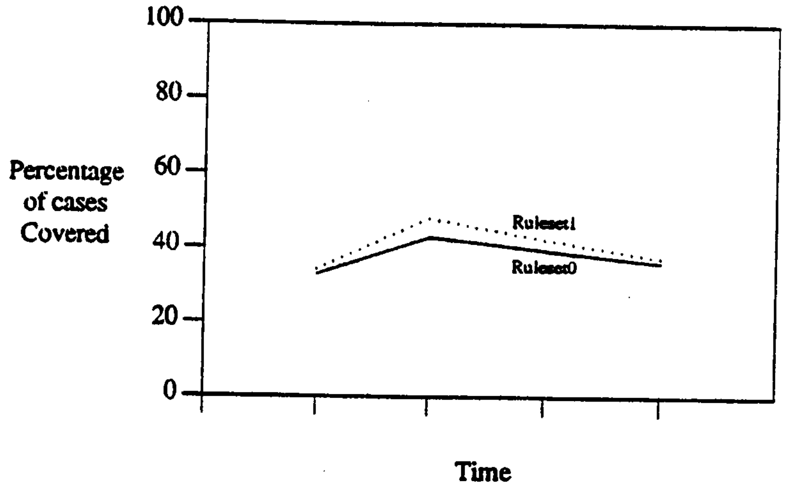

Figure 3: Percentage of Chronic problems covered by each rule set.

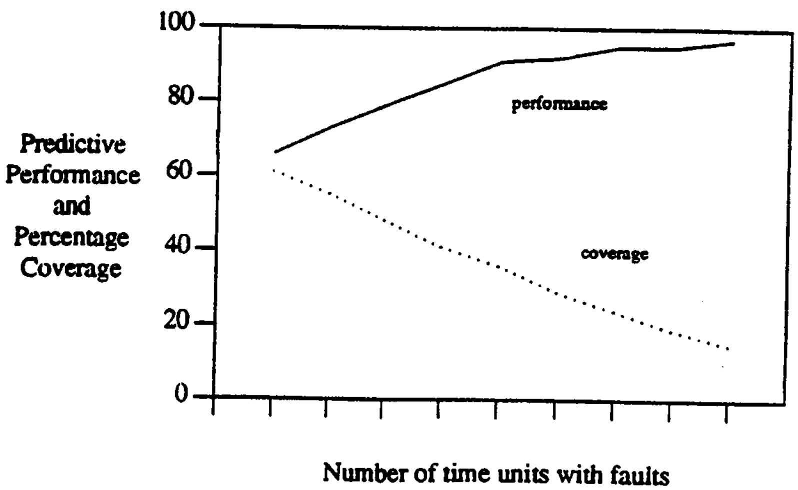

Figure 4: Performance of Ruleset0.