

A Hybrid Expert System / Neural Network Traffic Advice System

Rodney M. Goodman¹, Barry Ambrose¹, Hayes Latin² and Sandee Finnell³

¹ Electrical Engineering Dept., California Institute of Technology, Pasadena, CA 91125.

² Systems Technology, Pacific Bell.

³ Southern California Network Management Center, Pacific Bell.

Abstract

A hybrid Expert System / Neural Network Traffic Advice system has been developed to aid traffic management in the Southern Californian telephone network. The system has been working on live data since September 1991 and generates rerouting advice that agrees with that generated by the present network management procedures. A modular software design was adopted to allow for evolution. A graphics interface allows the user to easily navigate through the display of exception conditions and advice. Exceptions are shown highlighted on a map of Southern California. A neural network telephone traffic predictor was tested prior to incorporation in the system.

Keyword Codes: I.2.1, I.2.6, C.2.3

Keywords: Applications and Expert Systems, Learning, Network Operations

Introduction

In September 1990, a project was started by CalTech and Pacific Bell to investigate the implementation of a real-time autonomous expert system which would aid network management of the Pacific Bell Southern Californian telephone network [GOOD92]. The ultimate aim of the project was to produce and automatically implement recommendations for network controls. To do this, the knowledge in everyday use by the network management staff had to be written down, and coded for use in a computer system while retaining flexibility and openness to change.

The situations that the network management staff are exposed to are complex. Every day new problems arise. Examples are single random overloads of traffic on a given route, high volume phone-in conditions, disasters such as earthquakes, unusual calling patterns such as occur on Mother's Day, killer trunks, which are trunks that seize a call attempt and immediately release and other events. The introduction of new facilities, such as advanced signalling systems, adds additional problems that have to be diagnosed.

While many implementations exist that embed Expert Systems in existing operations software

[CALL88, CRON88, ERIC89, FRON91, MARQ88], the constraint of having to operate in real time and therefore prioritize exception conditions made this an interesting project. In addition, to our knowledge, the incorporation of neural networks to aid in time series prediction of telephone traffic statistics had not been previously attempted.

To date, a number of software modules have been developed which use data gathered in real time from the network. The software can produce a graphic analysis of the network, describe how the network is functioning in terms of exception conditions and recommend controls. The controls are given in the format that the network manager is accustomed to. Each module has been extensively tested.

The rest of the paper gives a description of the Pacific Bell telephone network and the architecture of the Network Operators Advice and Assistance (NOAA) system. This is followed by the details of how NOAA works presented logically in the order (i) input information available to NOAA (ii) decision making and decision support for NOAA (iii) controls recommended by NOAA. The map based Graphics Interface developed for NOAA is described separately followed by the neural network experiments and finally an indication is given of future work.

The Pacific Bell Network

In order to gain some appreciation of the network management tasks, one must have a description of the network to be managed. The Pacific Bell Southern Californian telephone network provides service to at least 1 million subscribers and has a hierarchical structure. For administrative and legal purposes it is divided into Local Access and Transport Areas (LATAs) [BELL86a] and the network management center in Pasadena manages LATA's 5 and 6. Geographically LATA's 5 and 6 extend to San Diego in the south and include all the Los Angeles urban area.

The network is hierarchical. End offices are the exchanges that serve customers, and tandems are the exchanges used for traffic between end offices that are not directly connected [BELL86a]. There are 6 tandems to be managed and over 200 end offices.

There are two types of trunk groups. High usage trunk groups exist between offices that are sufficiently close and have sufficient traffic flow between them. For economic reasons, these high usage trunk groups, which are the most common type of trunk group, are dimensioned to be lossy, i.e. during the busy hour they are not guaranteed to have enough capacity to carry all offered traffic. Traffic (call attempts) will therefore overflow onto a smaller number of Final trunk groups which are dimensioned to provide a good Grade of Service. In general there will be a final route between each end office and its parent tandem and between the tandems. It is these final routes that provide the backbone of the network. The final routes are therefore closely monitored by the network managers.

Architecture of NOAA

The Architecture of NOAA is shown in Figure 1. The Pacific Bell network management system is called EADAS [BELL86b]. NOAA is connected over a 4800 baud data link and appears as an ordinary operators terminal to EADAS. NOAA then runs on a Sun workstation.

There are four software modules that make up NOAA, each made up of about 3000 lines of C code. These are:

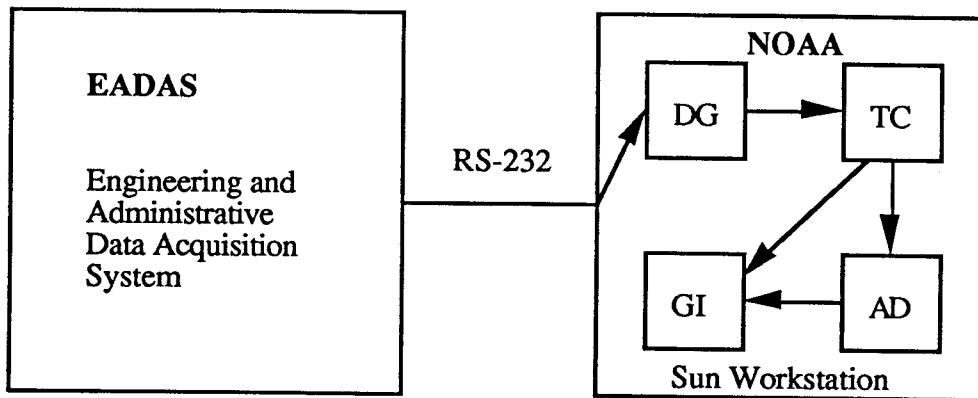


Figure 1 Architecture of NOAA

- data gathering (DG).

This module is responsible for logging into the EADAS system and maintaining the basic dialog with EADAS. The UNIX tip command is used for this purpose. A significant effort was put into making this module robust in the face of communication difficulties, such as loss of synchronization in the EADAS-NOAA dialogue.

- tandem calculations (TC).

This module takes the data which has been downloaded from EADAS and carries out the calculations needed to transform the data into usable statistics. The data is downloaded in register format. These registers, also known as peg counts, are integer counts of events such as trunk seizures, that have happened in a five minute interval. The calculations that are carried out are well defined and produce more useful statistics such as seizures per circuit per hour. In the next version of NOAA, neural networks will be used to predict future traffic statistics given the previous readings. This should improve the accuracy of the advice.

- Advice (AD)

Given the statistics on the various trunk groups, exception situations are noted and a severity measure associated with each exception. A diagnosis based on observed symptoms is carried out according to rules supplied by network management staff. Finally advice is given about the response to the exception. If a control is advised, it is given in the format expected by the network operator. This can be then approved by the operator and implemented through EADAS.

- Graphics Interface (GI)

The graphics interface uses a map of Southern California. Exceptions are indicated by red, suggested controls by green and actual controls already in

place by blue. Additional information is given when the operator clicks on the map. This information includes calculated statistics on the route and any advice generated for that route.

Statistics

NOAA makes its diagnosis on the basis of available information about the trunk groups. This information is available to NOAA in the form of counts of certain events for each trunk group during a 5 minute period. These events are for instance (i) a trunk being seized (ii) a seizure that is not successful i.e. overflow (iii) a survey of trunk occupancy at a given instant. On the basis of these counts, the following standard statistics are calculated:

| | |
|-----|--|
| ACH | Attempts per circuit per hour |
| CCH | Connections per circuit per hour |
| OFL | Percentage of attempts that overflowed |
| USG | Percentage of trunks occupied on average |
| HT | Holding Time of Calls |

The 5 minute period is a compromise. The shorter the period between data collections, the faster the response to any potential problems, but the larger the amount of information to be processed, requiring more expensive hardware. The 5 minute period means that problems can be cleared relatively rapidly. It is known that mean traffic increases do occur on a time scale of the holding time whereas traffic fluctuations occur on a time scale of the inter-arrival rate of calls [FIL190].

Since a large amount of information is available, filtering is carried out. Firstly exceptions on high usage routes are ignored and only final route data are examined. Secondly exceptions on special purpose final routes are ignored. For example, we expect a trunk group for the time announcement to have a short holding time. Thirdly overflows are treated as having higher priority than other exceptions, since an overflow on a final may represent lost traffic.

Decisions

When an exception condition has been noted on a trunk route, there could be many possible explanations for it. Typically phone-ins to radio stations and TV stations may generate excess call attempts. Facilities failures may mean that overflow shoots up on related trunk groups. Occasionally the data gathering may interfere with maintenance operations and unreliable data is returned. Finally random overflows can occur on individual trunk groups.

For any reroutes, a search for spare capacity is carried out. If there is an overflow problem between A and B where A is an end office and B is a tandem, then the program first looks for capacity on other trunk groups going from A to B, and then carries out a search for A to B via C possibilities. Here C has to be one of the 6 tandem exchanges, because of the hierarchical nature of the network. The candidate reroutes are then sorted according to available capacity and controls are suggested that make use of the minimum number of reroutes that achieves the required capacity.

A severity measure is attached to each exception. This measure is biased to give overflow exceptions a high severity figure, since overflows can result in lost traffic. The severity measure is also made proportional to the size of the trunk group and the amount by which the statistic is above threshold. Thus a 10% overflow on 10 circuits is considered equivalent to 1% overflow on 100 circuits. A log scale is then used to scale the figure between 0 and 100. This

severity measure has proved useful for sorting exceptions prior to displaying them. In a full scale emergency with the loss of a lot of capacity, a sort by severity would indicate which trunk groups to deal with first. Note that 911 (emergency) calls have their own reserved trunks from end offices and are not supervised by NOAA.

The rules used in the program are of three separate types:

- rules which indicate which exceptions can be safely ignored. For example overflow on high usage routes is ignored;
- rules which indicate which routes can be use as candidate reroutes;
- rules that map a suggested reroute into a list of controls to effect the reroutes. E.g. certain other routes may have to be "finalized" first to prevent a round-robin situation.

DISREGARD ANY EXCEPTIONS ON TRUNK GROUP COMMON LANGUAGES (CLLIS) ENDING WITH "MD" (EX: LSANCA02AMD; LSANFDRCCMD). EXCLUDE SAME WHEN SEARCHING FOR VIA ROUTE CANDIDATES.

DISREGARD ANY EXCEPTIONS ON CLLIS INDICATING "PB" IN THE STATE DESIGNATION (EX: OKLDPB0349T). EXCLUDE SAME WHEN SEARCHING FOR VIA ROUTE CANDIDATES.

DISREGARD ANY EXCEPTIONS ON THE FOLLOWING HIGH VOLUME CALL-IN CLLIS:

HLWDCA01520 SNANCA01977 COTNCA1143A
SIMICA11629 SNDGCA0157X

Table 1 Typical Network Management Rules

Some of the above rules were already written down in operators handbooks. Others were supplied by the network management staff. Examples of the rules are given in Table 1.

As with any rule based system, including the rules in a computer program has the advantage that any special cases that only occur rarely are immediately recognized and the operators handbook does not have to be referred to. However for complete trust in the system, the rule base has to be extensively tested and compared with the experts analysis in a wide range of cases.

Controls

It was mentioned above that there are many causes of exceptions on the trunk groups. Similarly there are many possible controls that can be put in to improve the network throughput. In general, controls can be divided into protective controls and expansive controls [FILI90, BELL88].

Protective controls are used to cut down the amount of call attempts entering the network. For example in the event of a phone-in, controls can be put into each end office to block call-

attempts to that telephone number. Alternatively a certain percentage can be let through. This limits the amount of congestion in the network. Clearly the calls destined for a phone-in are "bad" traffic as there is a high possibility of non completion, and therefore they should be blocked at source and not rerouted.

On the other hand, expansive controls allow trunk groups that are normally separate to share their capacity. If there is a single random overload, then expansive controls are called for. Two alternatives here are overflow rerouting, which reroutes overflow traffic only, and an immediate reroute which puts all call attempts onto the alternative route first. Each is appropriate for certain cases. Finally spray rerouting is possible, whereby a number of alternatives are given and each is tried for a given call. Clearly today's telephone exchanges have a lot more flexibility than the hardwired mechanical exchanges that were common thirty years ago.

Graphics Interface

It was found that this application was eminently suitable for graphical interfacing. Initially a scrolling list of exception conditions was presented to the operator in a window. Buttons allowed the operator to pull up information about the statistics on a trunk group and the recommended advice.

This was then modified to allow the operator to choose from a number of lists of exceptions, each list corresponding to a different statistic over threshold.

Finally a map interface was developed (see Figure 2). Without doubt, this was the easiest way to (i) display an analysis of the network, showing the exceptions clearly (ii) pull up information about individual exceptions. The map, with its translation and zoom features, could be mastered by any user in minutes and added a "fun" element to navigation through the program. Coordinates of most of the cities in Southern California were available from a public database.

Neural Networks

Neural networks have been used in applications ranging from pattern classification to associative memories. One of their main features is the ability to learn a mapping between the network inputs and the outputs. In contrast to artificial intelligence algorithms, the learning is based on memorizing examples by the process of *adjusting weights* in the network, rather than looking up rules. Much progress has been made on the algorithms used to train neural networks [HKP91].

Figure 3 gives an example of a two input single output neural network. The two input single output linear predictor can be considered as a simple form of neural network.

In our case, to aid in traffic management, we tried to use the neural network to predict a future value of trunk occupancy on a route, based on six previous readings over a thirty minute period. The prediction would provide a better indication of spare capacity for rerouting purposes and could also be used for extrapolation in the event of data not being available. The advantage of using a neural network for this application is that it can implement non-linear mappings between the inputs (in this case the previous occupancy readings) and the output (the predicted occupancy reading). Figure 4 shows a graph of occupancy on a trunk route over the course of a day.



Network Operators' Assistant
 quit search overflow status

cmdtool (CONSOLE) - /bin/csh
 netrcxx cd
 netrcxx iscreen
 screendump NOAA_SCREEN

Jan 14
 3:25:33 pm

Overflow Exception List

| Route | Severity | nccts | Type | service | pc |
|---------------------------|----------|-------|------|---------|----|
| Currently No Routes Match | | | | | |

Tue Jan 14 1992 15:10

Overflow Exc

| Route | Severity | nccts | Type | service | pc |
|-------------------------------|----------|-------|------|---------|-----|
| shokca0589t - (lmrca11d50-af) | | 240 | 2way | fi | 462 |
| shokca0589t - brnca13d50-ss) | | 240 | 2way | fi | 313 |
| shokca0589t - hlwda01d50-ss) | | 456 | 2way | fi | 277 |
| shokca0589t - nhwda02611-df) | | 72 | out | fi | 272 |
| shokca0589t - nhwda02g90-ct) | | 312 | 2way | fi | 253 |
| shokca0589t - brnca11c90-ss) | | 720 | 2way | fi | 230 |
| shokca0589t - nhllca01d50-af) | | 888 | 2way | fi | 217 |

Tue Jan 14 1992 15:15


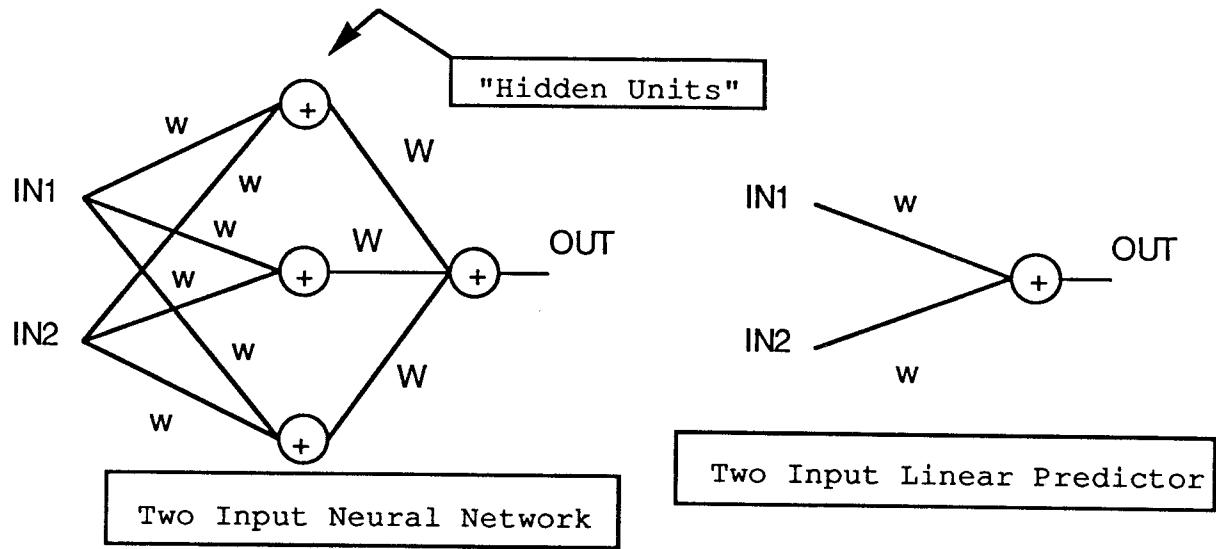


Figure 2 Graphical Interface for NOAA



w indicate weights and + represents addition which may be followed by a non-linearity in the case of the neural network. The non-linearity may be for example $\text{Tanh}(x)$.

Figure 3 Neural Network and Linear Predictor

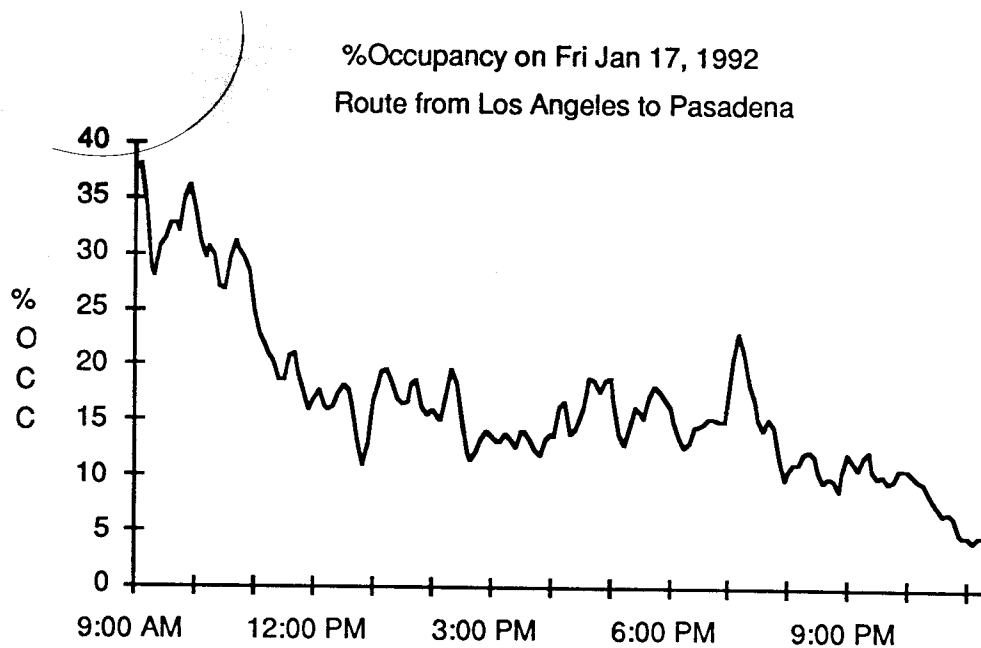


Figure 4 Traffic on a Trunk Route

After trying a number of simple examples, we choose a network with a small number of hidden units (3-5) for faster learning. We scaled the inputs to be between 0 and 1, although this was

not essential. We used non-linear hidden units but linear output units. We used a training set consisting of 893 examples and a test set of 609 examples. We used a linear predictor (also known as a perceptron in the neural network literature) as a standard for comparison. The results are shown in Table 2.

| | Training Set | Test Set |
|------------------|-----------------------|-----------------------|
| Linear Predictor | 2.39×10^{-4} | 3.63×10^{-4} |
| Neural Network | 2.18×10^{-4} | 8.20×10^{-4} |

Table 2 Mean Square Error for Two Prediction Methods

The table indicates that the neural network was better able to memorize the training examples but less well able to generalize on the test set. Given the added computation required to learn the non-linear mapping, it would seem that a linear predictor is the most suitable method for prediction of telephone traffic. This can also be justified theoretically. It is well known that a linear predictor works well for stationary processes whose probability densities are approximately normally distributed as in the case of telephone traffic [CHRIS91].

Future Work

Over the past 18 months, much work has been done in interfacing NOAA to the Pacific Bell network management computer and building the infrastructure for an expert system. Over the coming months the rules implemented in the program will be further tested by continuing to run the program on live data. When sufficient cases have been tested, it is intended to close the loop and allow NOAA to carry out controls autonomously. Clearly considerations of reliability and robustness have also to be taken into account when this step is to be carried out.

The ability of NOAA to diagnose problems correctly and to take the correct actions will be enhanced if the system has other information sources besides EADAS. Two other sources being considered at present are NSCS which provides information about seizures of trunks, and a separate system which provides information about the SS#7 signalling network.

Experiments with neural networks have shown that at little cost, linear predictors can be used to extrapolate statistics and generate more accurate assessments of available route capacity. A short term goal is to build the linear prediction capability into the TC module.

Looking at the long term future for NOAA, the definition of a standard data format for exceptions and for statistical information about trunk group performance would help in minimizing the cost of upgrade of NOAA, as new versions of EADAS become available. As in any network management application, standardization of data formats between applications that share the data is an important requirement. The ISO standards work on fault management and performance management should go some way to filling this gap.

To summarize, network management advice is currently being generated for the Southern California telephone network. As the rules that generate this advice are tuned, a robust network management application will have been developed that relieves network management staff of some of the need to supervise the day to day running of the telephone network.

References

- BELL86a Bellcore, Network Management Intra-LATA Network Fundamentals, BR 780-150-122, Issue 1, December 1986.
- BELL86b Bellcore, Network Management Technical Overview, BR 780-150-130, Issue 2, December 1986.
- BELL88 Bellcore, Network Management Control Descriptions, BR 780-150-128, Issue 3, June 1988.
- CALL88 Callahan, P., Expert Systems for AT&T Switched Network Maintenance, pp. 93-103, AT&T Technical Journal, Vol. 67, No. 1, January 1988.
- CHRIS91 Christensen, R., Linear Models for Multivariate, Time Series, and Spatial Data, Springer-Verlag, New York, 1991.
- COOP81 Cooper, R. B., Introduction to Queueing Theory, 2nd ed., North Holland, New York, 1981.
- CRON88 Cronk, R., Callahan, P., Bernstein, L., Rule Based Expert Systems for Network Management and Operations: An introduction, pp. 7-21, IEEE Network Magazine, Vol. 4, No. 5, September 1988.
- ERIC89 Ericson, E., Ericson, L., Minoli, D., (eds), Expert Systems Applications in Integrated Network Management, Norwood MA, Archtech House, 1989.
- FELL50 Feller, W., An Introduction to Probability Theory and its Applications, Volume I, ed. 3, Wiley and Sons, New York, 1950.
- FILI90 Filipiak, J., Real Time Network Management, Draft Copy (not published), April, 1990.
- FRON91 Frontini, M., Griffin, J., Towers, S., A Knowledge-Based System for Fault Localization in Wide Area Networks, pp. 519-530, IFIP Conf. on Integrated Network Management II, Washington, 1991.
- GOOD92 Goodman, R. M. et al, Network Operators Advice and Assistance (NOAA): A Real-Time Traffic Rerouting Expert System, Globecom, Orlando, December 1992.
- HKP91 Hertz, J., Krogh, A., Palmer, R.G., Introduction to the theory of neural computation, Lecture Notes Volume I, Santa Fe Institute Studies in the Sciences of Complexity, Addison Wesley, 1991.
- MARQ88 Marques, T., A symptom driven Expert System for Isolating and Correcting Network Faults, pp. 6-14, IEEE Communications, Vol. 26, No. 3, March 1988.