

NOAA - An Expert System managing the Telephone Network

*R. M. Goodman and B. E. Ambrose
California Institute of Technology, Pasadena, CA91125, USA
Ph: (818)3956811 Fax: (818)5688670
email: rogo@micro.caltech.edu*

*H. W. Latin and C. T. Ulmer
AGL Systems*

Abstract

A report is given on an expert system called NOAA, Network Operations Analyzer and Assistant, that manages the Pacific Bell Californian telephone network. Progress towards automatic implementation of expansive controls is complete. Progress towards restrictive controls is partially complete. Comments are made on current research including the use of neural networks for Time Series Prediction.

Keywords

Network Management, Telephone Network, Expert Systems, Expansive Controls, Restrictive Controls, Neural Networks

1 INTRODUCTION

Pacific Bell and Caltech have for several years been working on a real-time traffic management/expert system (Goodman, 1992, 1993). This project is called NOAA, Network Operations Analyzer and Assistant. The task of NOAA is to take information from the Pacific Bell network management computer, use it to isolate and diagnose exceptional events in the network and then recommend the same corrective advice as network management staff would in the same circumstances. A new company called AGL Systems has started up to continue the NOAA project and market it to all the Regional Bell telephone companies.

The NOAA project has several unique features:

- Provides expert system capability for complex decision making.
- Runs in real-time managing the whole of California's telephone traffic.
- Implements network controls automatically, 24 hours a day.
- Has a real-time earthquake information interface.
- Incorporates neural networks for time series prediction.

The rest of the paper gives a description of the Pacific Bell telephone network and the architecture of the Network Operations Analyzer and Assistant (NOAA) system. This is followed by sections on Expert Systems, Restrictive Controls, CUBE (Broadcast of Earthquakes), Research Aspects, and Conclusions.

2 PACIFIC BELL TELEPHONE 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 network in California is divided into North and South regions. Each region has a network management center associated with it. The network provides service to at least 4 million subscribers. The following network description is simplified for clarity.

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 (Bellcore, 1986). In the network as a whole, there are 15 tandems to be managed and over 400 end offices. The south is responsible for 6 of these tandems and about 200 end offices. The north is responsible for 9 tandems and about 200 end offices.

There are two types of trunk groups. *High usage* trunk groups 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 will therefore overflow onto the *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 home tandem. It is these final routes that provide the backbone of the network. The final routes are therefore closely monitored by the network managers. If such a final overflows then a customer gets an 'all circuits busy - please try again later' recording. It is the goal of network managers and NOAA to eliminate such messages as much as possible.

3 NOAA ARCHITECTURE

The Architecture of NOAA is shown in Figure 1. The Pacific Bell network management system is called NTMOS (officially NetMinder/NTM OS from AT&T). NOAA is connected over an ethernet data link and appears as an ordinary operators terminal to NTMOS. NOAA then runs on a Sun workstation under UNIX. Other operations systems interfaces are planned.

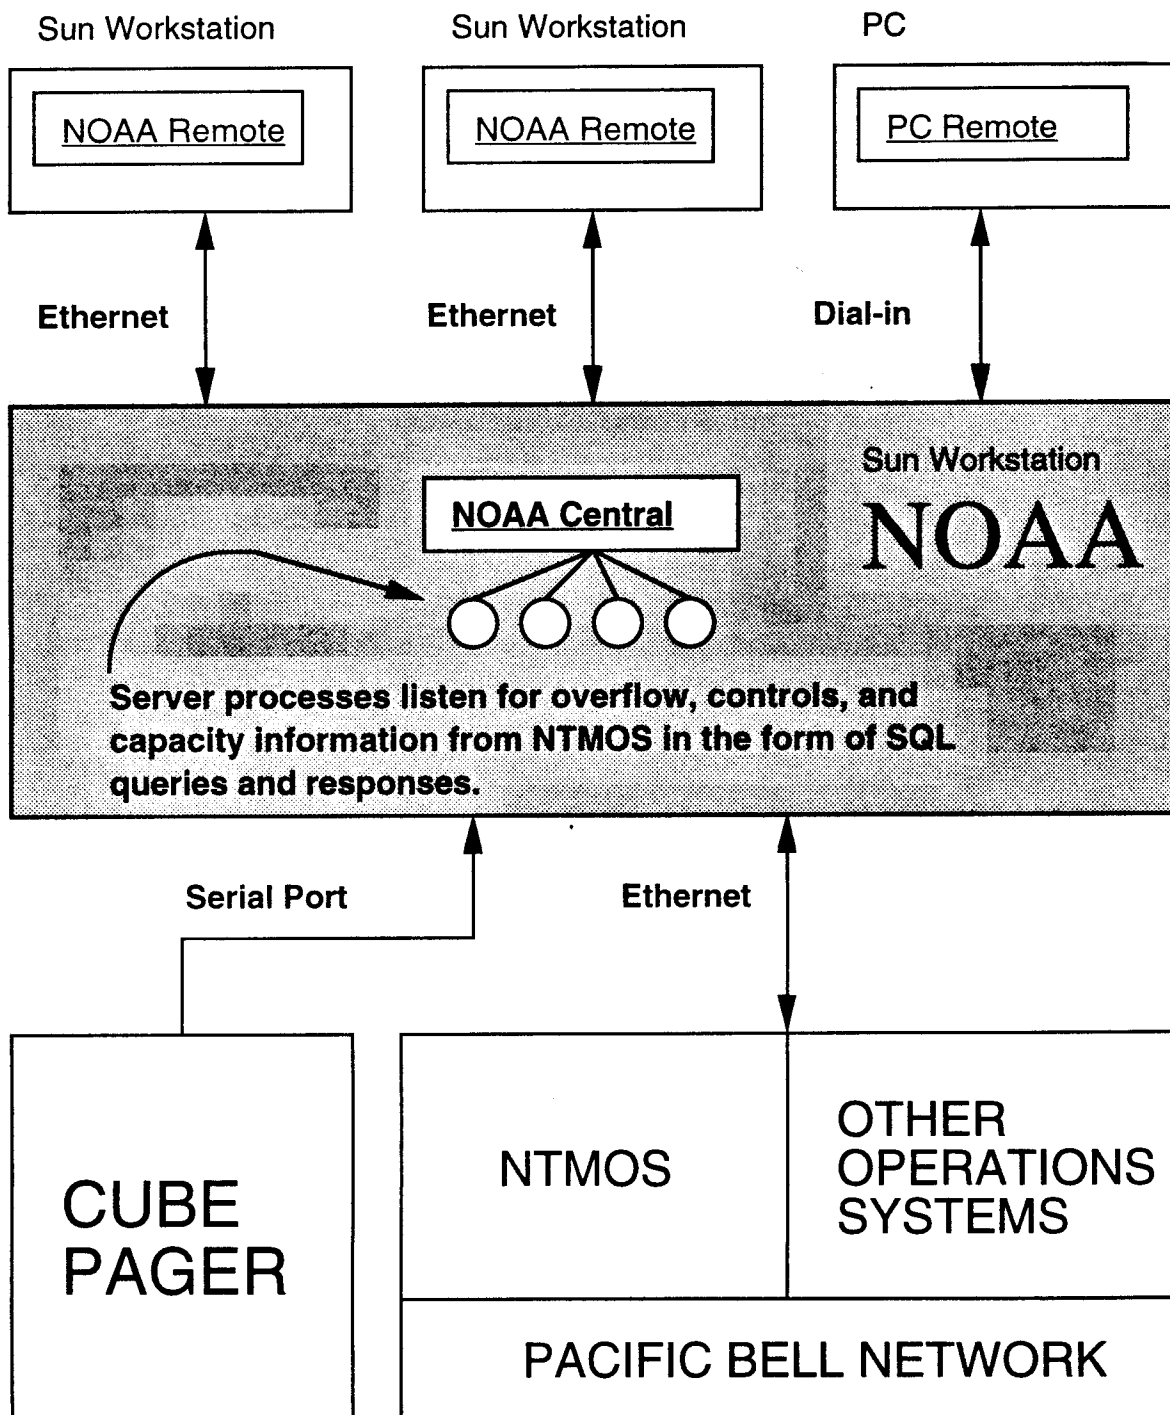


Figure 1 Architecture of NOAA.

4 EXPERT SYSTEMS

There have been other applications of expert systems to telephone traffic operations and management. For example (Sloman, 1994) lists the following among others. MAX from NYNEX and AMF from BT do fault isolation. NETTRAC from GTE and NEMESYS from AT&T do traffic management. However not all the features listed in the introduction are found in these products.

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 (trunks) failures may mean that overflow shoots up on related trunk groups. Occasionally maintenance operations may interfere with the data gathering and unreliable data is returned. Random overflows can occur on individual trunk groups. Most significantly, earthquakes can cause catastrophic overflows in a metropolitan area such as Los Angeles as people instinctively try to call loved ones after a moderate quake. The demand for dial tone can exceed normal operating loads by orders of magnitude, and bring the whole network to its knees.

After diagnosing the network problem, network management staff may choose to reroute traffic elsewhere (*expansive controls*) or cut the traffic off at its source (*restrictive controls*). Currently NOAA handles expansive controls and also restrictive controls to a lesser degree.

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

- rules that indicate which exceptions can be safely ignored. For example overflow on high usage routes is ignored;
- rules that indicate which routes can be used as candidate re-routes;
- rules that map a suggested re-route into a list of controls to effect the re-routes. E.g. certain other routes may have to be finalized first to prevent a round-robin situation. When a route is finalized, it no longer overflows onto a final route. A round robin situation is essentially a routing loop.

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. In addition, automated rule acquisition using our ITRULE algorithm has been used to extract rules. NOAA currently contains approximately 120 rules and this number is expected to grow as interfaces to other operations systems are added.

The automatic installation of controls raises questions about how the system fares in situations that are outside the rule base. In the short term, a button is available that marks a route as a special case. Also configuration files can be tailored to prevent NOAA from dealing with certain routes. For a more permanent fix, a suggestion screen is available to the operator, and based on the operators suggestions additions are made to the rule base to allow NOAA to deal with new situations.

As with any rule based system, including a good coverage of rules in the rule base has the advantage that any rarely seen special cases are immediately recognized as special cases and appropriately dealt with. In contrast, a human operator dealing with a rarely seen special case may need to refer to handbooks and reference material before implementing a control. 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.

5 RESTRICTIVE CONTROLS

The work of automating expansive controls is completed to the point where NOAA is capable of automatically implementing expansive controls and indeed this feature of NOAA is taken advantage of by the network operators. The next major goal is to provide the functionality to allow restrictive controls to be automatically implemented in the same fashion.

Restrictive controls are appropriate for call-in conditions, where most of the traffic has a low probability of completion, but its presence interferes with the normal network operations. Restrictive controls are also used for earthquake situations. In an earthquake situation, 10 times the traffic that the network is dimensioned for is typically present.

Interviews have been conducted with the network management operators in an attempt to find out the action of the network management operators in response to these and other failure possibilities.

For each event, the following questions were asked:

Awareness - How does the NM operator first become aware of the problem? What NTMOS statistics might be give-aways?

Decisions - During an event, what decisions have to be made? What control options are available? Is there coordination of actions with other personnel?

Decision Support Information - What information is needed to support each of the above decisions.

The following list of failure events was considered:

- Signalling System Failure
- Transmission Cable Cut
- Switch Office Failure
- Earthquake
- Call-in Event
- Weather Event

5.1 Signalling System Failure

Signalling information is used to set up and clear calls. More recently Signalling System No. 7 (SS7) signalling has allowed more flexible routing and number translation features.

With the older Multi-Frequency (MF) signalling, the signalling information is sent on the trunk carrying the call. If the signalling runs into problems, the individual trunk group will show problems and this will be detected by NTMOS.

With the newer SS7, the signalling is carried on a separate network to a special processing node called an STP. This makes it easy to install new signalling features by changing the software at the STP. If an STP were to fail, it would be a disaster. Redundancy is therefore supplied. Each office is linked to two STPs and each STP is loaded at a maximum of 50% so that if one STP fails, the other can take over.

The exact symptoms of a signalling system problem depend on switch type. In general increased ineffective call attempts, and low holding time of calls are observed. The appropriate action is for the signalling people to fix the STP.

5.2 Transmission Cable Cut

If a cable is cut, there may be enough capacity in the network to route around the point of failure. The cable may be carrying from tens to thousands of conversations. The main indication of a cable cut would be overflow on trunk groups. However even this information may not be available if the traffic levels are low. It may be that a single cable cut can halve the capacity of the trunk group. Thus with low traffic levels, no indication of any problem may be seen.

The appropriate action is to try to reroute any overflow around the failure. If no reroute paths remain intact nothing can be done.

5.3 Switch Office Failure

A switch office failure can be caused by a number of events. There may be a fire in the location, or the power supply equipment may fail, or the switch software may perform poorly in high load situations. The tandem switches are especially important to the health of the network, because of the volume of traffic that they carry.

'Discretes' from NTMOS are a good indicator of switch problems. Discretes are updated on a 30-second interval and hence provide early warning of switch malfunctions. The machine congestion discrete and the dial tone delay discrete indicate switch problems. It may be that the problem is temporary, in which case the appropriate action is to do nothing.

With SS7, congestion limiting controls may be automatically put in place if a problem is detected in sending traffic to a particular switch. The SS7 controls need to be augmented by manual controls if there is a switch failure. The manual controls would restrict traffic entering the network if the traffic probably would not complete. The manual controls would also reroute traffic to avoid heavily congested parts of the network.

Once the situation is diagnosed and controls put in place, the next action is to call people located near to the switch to check on the state of the switch. They have the decision power for removal of the controls.

5.4 Earthquake

The magnitude of the earthquake and the closeness to populated areas make a big difference in the severity of the event from the network managers point of view. A magnitude 5.0 in Los Angeles may be more serious than a magnitude 7.0 in the Mojave desert.

For serious earthquakes, say 6.0 or more in a populated area, there are many indications of problems. The discretes will indicate machine congestion and dial tone delay from switches whose load has increased. There will be lots of trunk group overflow from all over the region as every one picks up the phone to call their in-state and out-of-state friends and relations.

The Caltech CUBE broadcast of earthquake information should provide an indication of the magnitude of the quake and the location of the epicenter.

If the network is functioning ok, the appropriate action may be to partially directionalize the trunk groups to favor outgoing calls. In this case, outgoing call attempts are favored in the battle for the available resources. Any existing reroutes are taken out. 10 times more call attempts than the network is dimensioned for are typically present.

It is the experience of the network managers that the tandem exchanges win the battle for trunk group resources more often than the end offices. If this is seen to be the case, restrictive controls are put in at the tandems to allow both tandems and end-offices equal access to the trunks. Fairness of access to limited facilities is the guiding principle.

It may be useful to implement reroutes in less affected areas.

5.5 Call-in Event

If concert tickets go on sale at 10:00am on Monday morning and there is a lot of publicity about the event, a sharp increase of network traffic may be seen. Similarly if a cable TV company suffers a cable cut and goes off the air during primetime viewing, many people will call the cable operator at the same time to complain. Most of these calls will not complete, and customers may re-dial using auto-diallers. This volume of ineffective traffic may interfere with

the regular traffic by overloading the switches' and signalling systems' call processing capabilities.

This traffic is characterized by a large number of call attempts per circuit and low holding time. The tandem exchanges can provide an indication of when restrictive controls are appropriate through a hard to reach (htr) indication. This provides NOAA with information about an area code and telephone number prefix to which congestion is being experienced. NOAA can then do a table lookup to find the business that is associated with the telephone number, and place a restrictive control in all the offices in the network to cut down traffic whose destination is this number.

If a number is identified, it can do no harm to call gap the number. This won't affect calls to the number, provided the call volume is low, since its only action is to limit the number of calls accepted per 5 minute period. Even with call gaps in place, the office may be still overloaded by calls coming in from the long-distance network.

5.6 Weather Event

A weather event typically is a storm or blizzard. Weather events are characterized by a higher than usual level of traffic in the network. However the resources should still be available to handle the traffic. Thus although traffic may be 10 to 20% higher than usual, for the network managers, the only difference is a larger number of overflow exceptions to be handled. No special procedures are needed.

6 CUBE

CUBE is the Caltech / U S Geological Survey Broadcast of Earthquakes system. It provides epicenter and magnitude information of any earthquake occurring in California. In the event of a major earthquake NOAA applies a special set of rules to either scale back its recommendations or enter protective controls. Although CUBE only applies to California, the same type of system could conceivably be used to access information about other types of natural disaster, such as the National Hurricane Center's early warning system and tornado watch data.

Indications of an earthquake are first received on sensors distributed throughout California. This data is relayed to Caltech in Pasadena, where it is processed to provide epicenter location and magnitude information. Pager messages are then sent on the standard paging system to NOAA, and a data interface to the CUBE pager allows the message to be read and processed by NOAA.

Earthquake information received in real-time is displayed on NOAA's map in the form of a circle around the epicenter along with a numerical indication of the magnitude of the quake on the Richter scale. The map interface allows the operator an immediate identification of quake location and magnitude and well as identification of end-offices that may be impacted by the quake.

7 RESEARCH ASPECTS

During the course of developing NOAA, there have been opportunities for research. The involvement of the California Institute of Technology has been invaluable for investigation of these issues. Examples of the research issues that have been investigated are:

- the use of neural networks for time series prediction.
- the use of simulation to verify call saved metrics.
- the use of automated knowledge acquisition to generate rules describing correlation of exceptions.

7.1 Neural Networks

NOAA performs traffic prediction using *neural networks*. This comes under the heading of trend analysis. It can be applied to many time series found in NOAA's data structures to improve network capacity analysis and indicate potential network equipment shortages.

Neural networks have been used in applications ranging from pattern classification to associative memories. One of their main features is the ability to learn an arbitrary mapping between the network inputs and the outputs. In contrast to artificial intelligence algorithms, the learning is based on memorizing example patterns 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 (Hertz, 1991).

In this case, to aid in traffic management, the neural network was used to predict a future value of trunk occupancy on a route, based on previous readings. This provides a better indication of spare capacity for rerouting purposes and can 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).

The Quickprop (Fahlman, 1988) program for network training was used as it was advertised as having faster convergence than standard backprop. The quickprop program incorporates a weight decay factor which avoids overtraining. We modified it to include linear outputs since squashing functions on the output units will not aid function fitting.

A plot of hidden unit activations gave valuable insight into the features of the data. The features that were recognized in the training set by the hidden units were traffic level and rate of change of traffic level. In particular occasional traffic spikes showed strong activation for two of the hidden units. We are researching this feature as a means of signaling unusual conditions, e.g. the start of earthquake activity. This can then be used to automatically initiate restrictive controls.

7.2 Simulation for IRR metrics

NOAA displays a running total of the number of calls saved by network controls during the course of the day.

For ORR controls, which reroute calls that overflow from a problem route, the number of calls saved during a 5 minute period is simply equal to the number of calls that overflowed from the trunk group. A correction is made for any calls that were rerouted but still failed.

For IRR controls, which reroute calls before they even attempt the problem trunk group, the number of calls saved is not so easy to derive. Instead the number depends on (i) the number of trunks in the problem route (ii) the number of trunks in high usage routes that are overflowing to the problem route (iii) the holding time of calls and (iv) the number of call attempts on the problem route. A formula was derived which gave the number of calls saved assuming a knowledge of quantities (i), (iii) and (iv). In general, quantity (ii) is difficult to obtain. Simulations showed the formula accurately estimated the calls saved over a wide range of conditions. The formula itself is based on the Erlang Blocking formula that network planners use to find the number of trunks required for a given level of traffic.

7.3 Automated Knowledge Acquisition

In the knowledge acquisition process, we have been faced with the problem of developing rules via the traditional techniques of knowledge acquisition from human experts. This is a very time consuming process in terms of human resources, particularly expert availability. We have therefore investigated various automated knowledge acquisition techniques aimed at speeding up this process. In particular we have been concerned with the automated induction of rules from network management databases. These databases include trouble ticket databases, alarms databases, and topology databases. This area of learning from examples is referred to as machine learning, and a number of statistical and neural network algorithms exist that enable rules or correlations between data to be learned. We have developed our own algorithm ITRULE (Information Theoretic RULE Engineering) (Goodman, 1992). The ITRULE algorithm possesses a number of significant advantages over other algorithms in that the rules that are generated are ranked in order of informational priority or utility. It is thus an easy matter to directly load the rules into a standard expert system shell (such as NEXPERT), utilize an inferencing scheme based on these rule priorities, and have a working expert system performing inference in a matter of minutes. We have implemented the ITRULE suite of programs on a number of platforms (Sun, Mac, PC), and linked these into a number of expert system shells (NEXPERT, KES). This approach means that the expert system developer can 'instantly' generate and run a tentative expert system with little domain expertise. This 'bootstrap' expert system can then be used to refine the rules in conjunction with the domain expert in a fraction of the time of traditional 'cold' question and answer knowledge acquisition techniques.

8 CONCLUSIONS

Over the past three years, much work has been done in interfacing NOAA to the Pacific Bell network management computer and building the infrastructure for an expert system. The rules implemented in the program have been tested by running the program on live data. The loop has been closed and NOAA now carries out controls autonomously. Clearly considerations of reliability and robustness had to be taken into account when this step was carried out. Confidence in NOAA is very high, and NOAA is regarded by network management staff as a

valuable tool. In one case, where a switch had temporary problem, NOAA was able to implement 70 controls to route traffic around the switch in 15 minutes giving a much faster response than a human operator.

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 NTMOS. Two other sources being considered at present are NetMinder/NTP from AT&T which provides information about seizures of trunks, and a separate system which provides information about the SS#7 (Signalling System No. 7) signaling network.

The events of interest to the network managers are characterized by a sharp increase in traffic level or a sharp reduction in network resources. In some cases the increase in traffic level may be such that no network management controls are effective in managing the network throughput. In other cases, the scale of the event is smaller allowing re-routes or restrictive controls to bypass or reduce the problem.

There is plenty of scope for the rule-base of NOAA to be augmented to recognize these situations and take appropriate action. Some of the information to start doing this is already available from NTMOS. As interfaces to more Operations Systems become available, NOAA can begin to correlate event indications, and more effectively diagnose events.

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 NTMOS become available. As in any network management application, standardization of data formats between applications that share the data is an important requirement. The Bellcore GR495 (Bellcore, 1993) specification of network management information transmission should go some way to filling this gap.

To summarize, network management advice is currently being generated and controls automatically implemented for the whole of the California telephone network. As the rules that generate this advice were tuned, a robust network management application was developed that relieves network management staff of most of the need to supervise the day to day running of the telephone network.

9 REFERENCES

Bellcore, Network Management Intra-LATA Network Fundamentals, *BR 780-150-122*, Issue 1, December 1986.

Bellcore, Network Management Information Transmission Requirements, *BR GR-495-CORE*, Issue 1, November 1993.

Fahlman, S. E., Faster-Learning Variations on Back-Propagation: An Empirical Study in *Proceedings of the 1988 Connectionist Models Summer School*, Morgan Kaufman, 1988.

Goodman, R. M., Smyth, P., Higgins, C. M., Miller, J. W., Rule-Based Neural Networks for Classification and Probability Estimation, *Neural Computation*, Vol. 4, No. 6, November 1992.

Goodman, R. M., Ambrose, B., Latin, H., Finnell, S., Network Operations Analyzer and

Assistant (NOAA): A real-time traffic rerouting expert system, *Globecom*, Florida, December, 1992.

Goodman, R. M., Ambrose, B., Latin, H., Finnell, S., Network Operations Analyzer and Assistant (NOAA): A hybrid Neural Network / Expert System for Traffic Management, *IFIP*, San Francisco, April, 1993.

Hertz, J., Krogh, A., Palmer, R. G., *Introduction to the Theory of Neural Computation*, Lecture Notes Vol. I, Addison-Wesley, New York, 1991.

Sloman, M., *Network and Distributed Systems Management*, Addison-Wesley, New York, 1994.