1

PM-1 3 1/2" x 4" PHOTOGRAPHIC MICROCOPY TARGET
NBS 1010a ANSI/ISO #2 EQUIVALENT

National Library
of Canada

Acquisitions and
Bibliographic Services Branch

395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale
du Canada

Direction des acquisitions et
des services bibliographiques

395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Produit en France

Produit en France

NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30, and subsequent amendments.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30, et ses amendements subséquents.

Canada

**Detection of Abnormal Router Behaviour
In A Wide Area Network**

by

Alan D. Farrell, B. Eng., P. Eng.

**A thesis submitted to the
Faculty of Graduate Studies and Research
in partial fulfilment of the requirements
of the degree of**

Master of Engineering

**Ottawa - Carleton Institute for Electrical Engineering
Faculty of Engineering
Department of Systems and Computer Engineering
Carleton University
Ottawa, Ontario, Canada K1S 5B6**

March 5, 1993

National Library
of Canada

Acquisitions and
Bibliographic Services Branch

395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale
du Canada

Direction des acquisitions et
des services bibliographiques

395, rue Wellington
Ottawa (Ontario)
K1A 0N4

Your file / Votre référence :

Our file / Notre référence :

The author has granted an irrevocable non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of his/her thesis by any means and in any form or format, making this thesis available to interested persons.

L'auteur a accordé une licence irrévocable et non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de sa thèse de quelque manière et sous quelque forme que ce soit pour mettre des exemplaires de cette thèse à la disposition des personnes intéressées.

The author retains ownership of the copyright in his/her thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without his/her permission.

L'auteur conserve la propriété du droit d'auteur qui protège sa thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

ISBN 0-315-84118-4

Canada

**Masters Theses
PUBLISH ABSTRACT ONLY
AGREEMENT**

Abstract No. _____

INDICATE YOUR APPROVAL OF THIS AGREEMENT BY SIGNING ITEM #11

PERSONAL DATA (Please type or print)

1. Full name FARRELL
(last) ALAN DOUGLAS
(first) (middle)
2. Year of birth 1965
3. Country of citizenship CANADA
4. Present mailing address 39 Bentworth Cres
Napan, ON
K2G 3X1
- Future mailing address _____

- Effective date _____ Telephone number (416) 596 3166

DO NOT WRITE IN THIS SPACE

Vol/Issue _____
School Code _____
Advisor _____

Abstract Length _____
TP Date _____

MASTERS DEGREE DATA

5. Full name of university conferring degree CARLETON UNIVERSITY
6. Abbreviation for degree awarded M. ENG 7. Date degree awarded SPRING 1993

TITLE/SUBJECT DATA

8. Attach a copy of your thesis title page and abstract to this Agreement Form.
9. Subject Category for your thesis. Enter 4-digit code from list on opposite side.

10157414

ABSTRACT REPRINTS

10. Please send _____ reprints of my abstract. (Minimum order 100 copies)
First set of 100	\$ 30.00
Additional sets of 100 @ \$5.00 each set	\$ _____
Subtotal	\$ _____
Sales tax for state of ship-to	\$ _____
10% shipping and handling	\$ _____
Enclosed please find my check for	TOTAL \$ _____

This is an agreement between the author and University Microfilms Inc., Ann Arbor, Michigan, providing for the publication of an abstract (of not more than 150 words) of the thesis in *Masters Abstracts International*, which is distributed on a subscription basis. If more than 150 words are contained in the abstract, it will be returned to the author for abridging. Mathematical formulas, diagrams, and other illustrative materials are not recommended for the printed abstract.

11. I agree to the provisions of this agreement.

Signature

Alan Farrell

Date

APRIL 6, 1993

The undersigned recommend to the Faculty of Graduate Studies and Research
acceptance of the thesis

**Detection of Abnormal Router Behaviour
In A Wide Area Network**

submitted by Alan Farrell, B.Eng., P. Eng.

in partial fulfilment of the requirements for
the degree of Master of Engineering

Thesis Supervisor

Thesis Supervisor

Chair, Department of Systems and Computer Engineering

Carleton University

February 17, 1993

Abstract

The detection of faults in large communications networks is a difficult problem. This thesis presents an algorithm which analyzes large quantities of statistical information to detect the abnormal behaviour of routers in a Wide Area Network. Factors affecting the utility of the algorithm are determined. The accuracy of the algorithm is established. Some preliminary investigation of improvements to the algorithm is done.

Acknowledgements

I wish to acknowledge the support and direction of my research supervisors, Dr. Nick Dawes and Dr. Bernard Pagurek. I would also like to thank Bell Northern Research for providing me the opportunity to further my education in this field and to Telecommunications Research Institute of Ontario (TRIO) for funding the project.

Table of Contents

CHAPTER 1.0	Introduction	1
<hr/>		
CHAPTER 2.0	Background - Pegasus Project and BNR WAN	3
	2.1 The BNR Wide Area Network	3
	Topology of the BNR WAN	4
	Current BNR Management and Diagnostic Elements	7
	Faults in the BNR Network	9
	2.2 The Pegasus Project	10
	The Pegasus Blackboard Architecture	10
	The Dragon Network Simulator	12
	The BANES (Basic Analysis in Nested Evidence Spaces) K.S.	13
	The END (Expert Network Diagnostician) K.S.	13
	The LEAPR (Line Element Anomaly Pattern Recognition) K.S.	14
	The NPATH and MEDIAN Knowledge Sources	14
	Other Knowledge Sources Under Development	15
<hr/>		
CHAPTER 3.0	State of the Art in Network Fault Diagnosis	16
	3.1 Rule Based Diagnostic Systems	16
	A Diagnostic Expert System for Wide Area Networks	17

- Automated Knowledge Acquisition from Network Management Databases 17
- 3.2 Pattern Recognition Techniques For Proactive Network Analysis 18
 - Anomaly Detection in a Local Area Network - R.A. Maxion 19
 - The LEAPR Knowledge Source 20
 - ATM Communications Network Control by Neural Networks - A. Hiramatsu 21
- 3.3 Probabilistic Technique for Reactive Break Fault Detection 22
- 3.4 The Role of Cooperating Processes in Diagnosis 23

CHAPTER 4.0 The Problem of Detecting Abnormal Router Behavior 26

- 4.1 Detecting Abnormal Router Behaviour: Problem Definition 26
- 4.2 There is a Need For an Algorithm to Detect Abnormal Router Behaviour 28
- 4.3 An Algorithm To Detect Abnormal Router Behaviour Will Be Beneficial 31

CHAPTER 5.0 The Development of an Algorithm to Detect Abnormal Router Behaviour 33

- 5.1 Phase One: Studying the Effect of a Router Delay 36
- 5.2 Phase Two: Can A Router With a Delay be Detected and Discriminated From Its Neighbours? 42
 - Step 1: Router With A Delay, Two Neighbours, 1 File Transfer 43
 - Step 2: Router with a Delay, Two Neighbours, Many File Transfers 48
- 5.3 Investigating an Algorithm To Detect a Router With a Delay 56
 - Step 1: Algorithm Definition 57
 - Step 2: Role of Traffic Mix in Algorithm Performance 58
 - Step 3: Impact of a Router's Delay in Various Zones 62
 - Step 3: Long Running Simulations 65
- 5.4 Alarm Thresholds and Algorithm Performance 79
 - Step 1: Establish Base Statistics 80
 - Step 2: No Faults in Network, Determine Algorithm's False Positive Rate 82
 - Faults in Network, Determine Algorithm's Correct Diagnosis and False Positive Rates 88
- 5.5 Preliminary Work in Future Research Areas 97
 - A Second Method of Defining Alarm Thresholds 98
 - Investigate Determination of Base Statistics - The Model of Correct Behaviour 104
 - Over What Time Frame Should Base Statistics Be Gathered ? 110

CHAPTER 6.0

Conclusions, Future Research, Summary of Contributions 115

- 6.1 Conclusions 115**
- 6.2 Future Research 117**
- 6.3 Summary of Contributions 118**

References 120

List of Tables

TABLE	DESCRIPTION	PAGE
5.0	Mean Bytes In For nRTPc2 Under Different Conditions	38
5.1	Ratios of Node In-Mean-Bytes Per Minute for Different States	45
5.2	(Standard Deviation/Mean In Bytes)*100% For Different States	47
5.3	Analyzing Node Level Traffic For Delay in nRTPc2	51
5.4	Analyzing Line Traffic Within a Node with Delay in nRTPc2	52
5.5	Router contribution levels, snmp traffic, varying delays in nRTPc2	59
5.6	Change in Router Contribution to Zone, nRTPc2 Delay = 0.1 second	63
5.7	Change in Router Contribution to Zone, nRTPc2 Delay = 0.5 second	63
5.8	Change in Router Contribution to Zone, nRTPc2 Delay = 1 second	63
5.9	Change in Router Contribution to Zone, nRTPc2 Delay = 2 seconds	64
5.10	Router Base Statistics Gathered Over 24 Hours (144 - 10 minute intervals)	82
5.11	Correct Diagnosis and False Positives As Delays are Inserted in Routers	92
5.12	Routers Contributing > 20% To Zone, Highlighting Those With High Standard Deviations	93
5.13	Router Base Statistics - Collected Under Two Different Conditions	106
5.14	Means, Standard Deviations For Routers' Percent Contribution, and Mean Bytes In (over 24 hours)	112

List of Figures

FIGURE	DESCRIPTION	PAGE
2.0	Basic Topology of the BNR Wide Area Network	5
2.1	Router and Ethernet Connectivity	6
2.2	Key Elements of BNR's Network Management Structure	8
2.3	The Pegasus Blackboard Architecture	12
3.0	LEAPR Knowledge Source Structure	20
5.0	Research Phases	34
5.1	Research Phases - Currently in Phase One	36
5.2	File transfer path through nRTpC2	39
5.3	Research Phases - Currently in Phase Two	42
5.4	Topology: Router with a delay, two neighbours, one file transfer	44
5.5	Topology: Router with a Delay, Two Neighbours, Many File Transfers	49
5.6	Zone Around a Router	55
5.7	Research Phases - Currently in Phase 3	56
5.8	Research Phases - Currently in Phase 4	79
5.9	Research Phases - Currently in Future Research Phase	97

List of Graphs

GRAPH	DESCRIPTION	PAGE
5.1	Effect of 0.5 Second Delay in nRTPc2 as FTP/SNMP Traffic Ratio Varies	60
5.2	Histogram of % Deviation From Mean Traffic Contribution - nRTPc2	70
5.3	Histogram of % Deviation From Mean Traffic Contribution - nRTPc1	70
5.4	Histogram of % Deviation From Mean Traffic Contribution, nPARc2	71
5.5	Histogram of % Deviation From Mean Traffic Contribution, bRTPc2	71
5.6	Histogram of % Deviation From Mean Traffic Contribution, nATLc1	72
5.7	Percent Contribution of Each Node to Zone as nRTPc2 Delay Varies	73
5.8	Percent Change in % Contribution of Each Node To Zone as nRTPc2 Delay Varies	74
5.9	Percent Contribution of Each Node to Zone as bRCHc1 Delay Varies	75
5.10	Percent Change in % Contribution of Each Node To Zone as bRCHc1 Delay Varies	76
5.11	Percent Contribution of Each Node to Zone as nMTHc1 Delay Varies	77
5.12	Percent Change in % Contribution of Each Node To Zone as nMTHc1 Delay Varies	78
5.13	Histogram - % Decrease in Zone Contribution, Routers 0 - 5 % Contribution over a 24 hour period	84
5.14	Histogram - % Decrease in Zone Contribution, Routers 5 - 10 % Contribution over a 24 hour period	84
5.15	Histogram - % Decrease in Zone Contribution, Routers 10 - 15 % Contribution over a 24 hour period	85
5.16	Histogram - % Decrease in Zone Contribution, Routers 15 - 20 % Contribution over a 24 hour period	85
5.17	Histogram - % Decrease in Zone Contribution, Routers > 20 % Contribution over a 24 hour period	86
5.18	Percentage False Positives For Each Router Category	87
5.19	Histogram - % Decrease in Zone Contr., Router > 20% Contr., nRTPc2 with a one second delay	89
5.20	Histogram - % Decrease in Zone Contribution, nRTPc2 with 1 second delay	89
5.21	Histogram - % Decrease in Zone Contribution, nRTPc2 with zero delay	90
5.22	Drop in Zone Contribution over 24 hours, Measured in "Percent" Drop and "Number of Std. Deviations Drop" , no fault in network	100
5.23	Drop in Zone Contributions for Healthy Routers, Measured in "Percent" Drop and "Number of Std. Deviations Drop"	101
5.24	Drop in Zone Contributions for Faulty Routers, Measured in "Percent" Drop and "Number of Std. Deviations Drop"	101
5.25	Drop in Contribution For Non-Faulty Nodes (Base Stats Do Not Include Faulty Behaviour In The Network)	108

List of Graphs

5.26	Drop in Contribution For Faulty Nodes (Base Stats Do Not Include Faulty Behaviour In The Network)	108
5.27	Drop In Contribution For Non-Faulty Nodes (Base Stats Include Faulty Behaviour In The Network)	109
5.28	Drop In Contribution For Faulty Nodes (Base Stats Include Faulty Behaviour In The Network)	109
5.29	nRTPc2's % Contribution to Zone over 72 - 20 minute Intervals (24 hours)	113
5.30	nRTPc2's Mean Bytes Into Node over 72 - 20 minute Intervals (24 hours)	113

Glossary of Common Acronyms

ATM:	Asynchronous Transfer Mode
BANES:	Basic Analysis in Nested Evidence Spaces
BISDN:	Broadband Integrated Services Digital Network
DCD:	Data Collector Device
END:	Expert Network Diagnostician
FTP:	File Transfer Protocol
KS:	Knowledge Source
LAN:	Local Area Network
LEAPR:	Line Element Anomaly Pattern Recognition
MAGNET:	BNR Network Monitoring Package
MAN:	Metropolitan Area Network
NOC:	Nework Operations Center
NPATH:	Operations Research Path Finding Algorithm
OSI:	Open Systems Interconnection
PING:	Packet INternet Groper
PC:	Personal Computer
SNMP:	Simple Network Management Protocol
WAN:	Wide Area Network

CHAPTER 1.0 Introduction

This thesis presents an algorithm used to analyze large amounts of network performance statistics to detect abnormal router behaviour in a Wide Area Network.

Faults in a Wide Area Network can manifest themselves as a broken component or as a malfunctioning component. The latter of these two faults can be very difficult to trace as there is often no clear indication of where the fault lies. The redundancy and fault tolerance of communication networks and protocols can hide these faults or lessen their impact on traffic flow. Human operator analysis of the large amounts of network traffic information is for the most part unfeasible. A disrupted file transfer may be the only observable indication of the presence of a fault somewhere in the network. As a result, these faults can be present in a network for long periods of time before they are diagnosed or result in component failure that is more directly observable. An algorithm that can identify a router that is not functioning properly (but is still functioning) would be an invaluable tool in network diagnosis.

The algorithm developed in this thesis monitors the traffic in a router relative to the routers around it. A drop in a router's traffic level compared to those around it is an indication that there is a fault in the router. The algorithm is capable of detecting a packet processing delay inserted in a router. It can be used in a proactive manner or, cooperating with other Knowledge Sources in a reactive manner. The factors which affect its utility have been documented.

The remainder of this thesis is organized as follows:

Chapter 2 presents background information on the Pegasus Fault Diagnosis project. A description of the BNR Wide Area Network (WAN), which is the target network for our diagnosis techniques, is given, including a discussion of network topology, current network management tools, and the types of faults encountered in the BNR WAN. The Pegasus project architecture and knowledge sources are briefly presented.

Chapter 3 reviews current research in the field of fault diagnosis and pattern analysis. Some of the Pegasus diagnostic techniques are examined in more detail, and other relevant research papers are discussed.

Chapter 4 provides a complete discussion of the diagnosis problem tackled in this thesis.

Chapter 5 documents the development of the diagnostic algorithm that will detect abnormal router behaviour in a WAN.

Chapter 6 presents the thesis conclusions, future research areas, and a summary of contributions of this research.

CHAPTER 2.0 Background - Pegasus Project and BNR WAN

The Pegasus Project incorporates a multiparadigm approach to fault detection in large communications networks. The Bell-Northern Research Wide Area Network is a very large communications network that is used as an example network for the Pegasus Project. This chapter provides an overview of the BNR network and the Pegasus Project.

2.1 The BNR Wide Area Network

The BNR WAN is a large, private communications network with many thousands of devices spanning four continents - Europe, Asia, Australia and North America. This section provides a brief overview of the BNR WAN and is divided into three subsections. First, an overview of the Network Topology is presented. Second, a summary of BNR's current management structure and diagnostic elements is given. Last, an overview of the types of problems encountered by network managers is given.

2.1.1 Topology of the BNR WAN

The basic components and general connectivity in the BNR WAN can be seen in Figure 2.0[1]. Up to 16 workstations are connected to primary hubs. In turn, up to 16 primary hubs can be connected to a secondary hub. The hierarchical structure of workstations, primary hubs, and the secondary hub is known as a workgroup LAN in BNR terminology. The workgroup LAN hierarchy provides connectivity between workstations located under the same secondary hub, consolidating local traffic which need not go to the LAN backbone. Important devices, such as mainframes and Fastpaths for Appletalk devices are connected directly to several LAN backbones.

Workgroup LANs are interconnected by a LAN Backbone through bridges. Bridges forward packets that are destined to other workgroups from the secondary hub to the ethernet and pick packets off the ethernet that are destined to its own workgroup and forwards them to its secondary hub.

Routers are used to interconnect LAN Backbones into Metropolitan Area Networks (MAN) and to interconnect MANs into a Wide Area Network (WAN). Figure 2.1 shows the interconnection of routers [1]. A router is usually connected to two ethernets, and using serial lines (56 kb to 1.5444Mb) it can connect to several other routers. Routers pick up packets from ethernet LANs destined for another ethernet LAN and forwards them to the appropriate router. The Router also accepts packets from other routers and forwards them to the next router along their path or puts them on the appropriate ethernet. Each router is equipped with a routing algorithm to dynamically choose the best route to another LAN. While this flexibility improves load balancing, performance, and error recovery abilities, it does complicate diagnosis of network faults because the robustness in the presence of a fault tends to mask the existence of the fault.

FIGURE 2.0 Basic Topology of the BNR Wide Area Network

FIGURE 2.1 Router and Ethernet Connectivity

2.1.2 Current BNR Management and Diagnostic Elements

Figure 2.2 shows the key elements of the BNR network management structure [2].

The BNR Wide Area Network is monitored from the Network Operations Centre (NOC) at one of BNR's labs. Network operators use a tool called Magnet to monitor the WAN. Magnet is a workstation based graphical monitoring system that presents the network status to operators in the NOC. Network operators can view different parts of the network at different levels of detail - from high level site status down to specific device status in specific sites. Colour coded icons are used to indicate the state of components in the network. Magnet has input in the form of files describing network topology and status information forwarded from Data Collector Devices (DCDs) throughout the network.

DCDs are polling devices spread throughout the network. Each is responsible for monitoring a number of network components. Only the more important components in the network are monitored - workstations, for example are not monitored while servers, mainframes, routers and bridges are. DCDs poll devices using PING (Packet InterNet Groper) and SNMP (Simple Network Management Protocol) messages. Status messages are forwarded from the DCDs to the Magnet workstation for analysis and presentation to the network operator.

The Magnet monitoring system is supplemented by diagnostic tools such as the Remote Diagnostic Access (RDA) Tool used to monitor network traffic. The RDA tool consists of a matrix switch connected to a secondary hub and a personal computer (PC) equipped with network analyzer software. The operator can use the PC in the NOC to analyze traffic throughout the network, using the matrix switch to tap into the network at specific points to analyze network traffic - for fragmentation, types of traffic, monitoring certain traffic types, etc. The RDA tool eliminates the need for network operators to physically leave the NOC to connect a datascopes to network equipment in different parts of the

FIGURE 2.2 Key Elements of BNR's Network Management Structure

building - city or country for that matter. The existence of tools such as the RDA allows BNR to consolidate its Network Management personnel in one location.

It is worth noting that not all network faults are detected using Magnet. Some faults only become known when a user calls the NOC to complain of a performance problem with his or her workstation or workgroup.

2.1.3 Faults in the BNR Network

Faults in the BNR WAN range from break faults, where a component is not performing at all, to non-break faults, where a component is functioning improperly. In this section, an overview of some of the causes of network faults is given, followed by a short discussion on the way in which these faults may manifest themselves. This section is not intended to be an in-depth analysis of communications system faults. It is intended to give the reader an appreciation of the wide range of issues network managers deal with.

There are a large number of devices in a large communications network. Devices are generally supplied by a variety of vendors. Devices supplied by the same vendor may have different vintage hardware or software. There are also a large number of protocols in use in a given network. Topology changes take place frequently. An environment such as this is understandably prone to failure from a variety of causes. Incompatibility of software versions or hardware versions between components of one or more vendors frequently causes faults. Protocols may not be fault-free, and it is possible that different vendors interpret or implement protocols in an inconsistent manner, occasionally resulting in abnormal behaviour. Improper connection of devices, such as a pulled cable or unclear connection may cause a fault. Bugs in communication software or failure of micro-chips and other hardware components cause faults. Misconfiguration of devices, unreported topology changes and other “finger” problems also cause faults in computer networks.

Faults in communication networks can manifest themselves in a variety of ways. Components may be rendered non-operational, or may work correctly but with poor performance, or may appear to work correctly when they are not. Faults may be steady or intermittent. Faults may be immediately seen by monitoring devices such as Magnet or

may remain unseen until a user telephones the NOC complaining of poor network performance or the inability to connect to a device. The robust and fault resistant topology of a network and its protocols means that some faults may exist for a long period of time before they are discovered. Faults may appear in an asymmetrical manner - workstation X can connect to workstation Y but workstation Y cannot connect to workstation X. Faults may masquerade as other faults or remain hidden by a second fault. For example, the failure of a DCD gives the appearance of the complete failure of a large portion of the network.

The environment described above is clearly a challenge to troubleshoot and maintain. Network operators have at their fingertips vast quantities of network status information which may or may not be useful in detecting a fault. Developing algorithms to process this information and summarize it in a useful manner provides the motivation for this thesis.

2.2 The Pegasus Project

The Pegasus system is designed to diagnose faults in communication networks [3]. It consists of a multiparadigm approach organized on a Blackboard Architecture. A full scale network simulator is used for system development and testing. The Pegasus System is built on a UNIX based Sun SPARCstation 2. The main languages used in its development are C, C++, and Pascal.

This section of the thesis gives a brief description of the Blackboard Architecture, the network simulator, and the knowledge sources currently implemented in Pegasus. Some of the knowledge sources are described in more detail in Chapter 3.

2.2.1 The Pegasus Blackboard Architecture

Blackboard architectures have been recognized as being good for complex, ill-defined problems by providing dynamic control, flexibility and modularity [4,5]. The Pegasus Blackboard provides the basis for a multiparadigm system with the ability to investigate

more than one fault at the same time through concurrent execution of knowledge sources. The different diagnostic agents can communicate and interact using the blackboard to operate in a cooperative fashion.

The major components of the Pegasus Blackboard are shown in Figure 2.3, [6]. The knowledge sources are a group of processes working together to diagnose faults in the network. The blackboard database is a global database used by the knowledge sources to communicate and interact. It is subdivided into six areas: Status Messages area, Traffic Messages area, Problems area, Hypotheses area, Knowledge Source Activation Record area, and the Requests and Replies area. The blackboard handling functions and control data structures allow knowledge sources to monitor the blackboard, locate changes easily and add their own changes. Using the blackboard, knowledge sources can cooperate, share information and build on each other's diagnoses. This aspect of using blackboards is discussed further in Chapter 3.

FIGURE 2.3 The Pegasus Blackboard Architecture

2.2.2 The Dragon Network Simulator

The Dragon Network Simulator provides a means by which knowledge sources can be evaluated in their ability to detect network faults [7]. It is unreasonable to expect to have unlimited access to real network data and be able to insert faults in a real network to evaluate the performance of a knowledge source in detecting them. The network simulator was built to allow the development and evaluation of different diagnostic methods. Once a diagnostic method is proven to work on the simulator, evaluation of the method using real network information is more justified.

The Dragon Simulator supports most of the communications functionality of very large WANs and ATM (Asynchronous Transfer Mode) / B-ISDN (Broadband Integrated Services Digital Network) Networks and operates at near real time. Network topology is obtained from three files. The SNMP, PING, and Open Systems Interconnection (OSI) management protocols are included. Both alarm and status information is available to the user. File transfers and interactive file sessions are supported. Network traffic varies with the time of day. A variety of network faults (randomly generated or manually inserted) can be introduced to the network and the effects of these faults can be observed. Both manual and program control levels are supported.

2.2.3 The BANES (Basic Analysis in Nested Evidence Spaces) K.S.

BANES was developed within the Pegasus group and uses a model of the system being diagnosed, and symptoms associated with the fault to identify the faulty device(s) [8,14,15]. BANES relies on the propagation of beliefs in an internal structure known as Nested Evidence Spaces. The nested evidence spaces are created from the network topology and the beliefs are generated from the propagation of status messages in the network.

BANES is analyzed in more detail in Chapter 3.

2.2.4 The END (Expert Network Diagnostician) K.S.

END is a rule based knowledge source [9]. Rule based systems are well established in represented knowledge, directly translating human experts' knowledge into rules a computer understands. Network experts were consulted and the diagnostic procedures for a number of faults were incorporated as rules in the END knowledge source. The COSMIC CLIPS shell was used to develop END.

END is used to investigate communication problems reported by users through phone calls. END is also used to verify the hypotheses generated by knowledge sources such as BANES.

END is discussed in more detail in Chapter 3.

2.2.5 The LEAPR (Line Element Anomaly Pattern Recognition) K.S.

The LEAPR knowledge source was developed to proactively detect network performance degradation before network users or operators detected it [10]. LEAPR uses a back propagation neural network implemented in the C language.

LEAPR monitors the line traffic statistics on a router to determine the health of the router. The neural network is trained to classify traffic into normal and abnormal by exposing it to examples classified by network communications experts.

LEAPR is discussed in more detail in Chapter 3.

2.2.6 The NPATH and MEDIAN Knowledge Sources

Operations Research techniques are used to analyze the information in network topology files [11].

The NPATH knowledge source is used to find the shortest path between nodes in the communications network. A specified number of shortest paths between two nodes are determined, as well as the most critical arcs and nodes between the source and destination. In the event that communication between two network components is disrupted, NPATH can be used to identify the critical arcs and nodes between these components. Once identified, the critical arcs and nodes can be examined in detail to determine if they are faulty.

The Median Knowledge Source is used to determine the optimum location of Data Collector Devices in the Wide Area Network. Optimum placement of DCDs results in a minimization of monitoring traffic in the network while maximizing its effectiveness

2.2.7 Other Knowledge Sources Under Development

Other Pegasus team members are continuing research into other fault diagnosis techniques. Machine learning techniques to analyze traffic patterns in a router are being investigated [16]. Case based reasoning techniques correlating attributes (traffic patterns, alarms, etc.) to specific types of faults are also being investigated [17].

CHAPTER 3.0 State of the Art in Network Fault Diagnosis

The field of diagnosis extends over a broad range of applications. This chapter presents a review of some important diagnostic approaches that can be used in network fault detection. The chapter is organized into four sections. First, the use of Rule Based Diagnostic Systems is discussed. Second, Pattern Recognition techniques applied to network analysis are discussed. Third, a novel technique for object classification and evidence accumulation is presented. Last, the role of cooperating processes in diagnosis is discussed.

3.1 Rule Based Diagnostic Systems

When one thinks of artificial intelligence and diagnosis, rule based “expert systems” are often the first to come to mind. Rule based systems capture a human expert’s knowledge in a set of rules which are stored in a knowledge base and manipulated by an appropriate inference engine. Two articles involving the use of rule based systems were reviewed. The first discusses the END knowledge source used to diagnose faults in com-

munication networks. The second addresses some of the limitations of rule based systems, and presents an algorithm to alleviate some of the shortcomings of rule based systems.

3.1.1 A Diagnostic Expert System for Wide Area Networks

The Expert Network Diagnostician (END) Knowledge Source [9] was developed as part of the Pegasus WAN fault diagnosis project and uses Expert Systems techniques to capture part of the knowledge and fault diagnosis procedures used by BNR network operators. Network diagnosticians were consulted and a portion of their expertise formalized in sets of rules. END is divided into four separate modules.

The Veri-line and Veri-element modules are both used to validate the BANES generated hypotheses regarding broken station or router elements and broken lines. The General Network Problem Locator module investigates communication problems reported by users and posted to the blackboard. The Network Overload Tracer module investigates more complex routing problems, including determining the source of data storms in the network.

The END modules illustrate the value of expert systems technology in automating human diagnostic procedures and knowledge used in a diagnosis environment. It also indicates that a large quantity of knowledge can be effectively divided into small, more manageable task specific knowledge sources.

3.1.2 Automated Knowledge Acquisition from Network Management Databases

Rule based systems have a number of limitations, including bottlenecks in knowledge acquisition and updating, difficulty in scaling up, and limitation to certain kinds of domains. Goodman and Latin [13] use machine learning techniques to alleviate the knowledge acquisition bottleneck in developing rule based expert systems. The ITRULE (Information Theoretic Rule) Induction Algorithm was developed to automatically generate useful rules from trouble ticket and alarms databases. These rules can be fine-tuned and

loaded into an expert system shell to provide an instant rule-based system, or they can be used as a basis for a knowledge acquisition session with a domain expert. An interesting aspect of this work is that many types of rules (depending on the type of database it is run on) can be generated, not just symptom class rules.

The ITRULE algorithm selects a right-hand side of interest and constructs first-order left-hand sides. The algorithm searches for more information and creates higher-order left-hand sides. A “J-measure” of utility is calculated for each rule, based on the probability that the rule will fire and the amount of information gained from the rule if it fires. This utility measure is used to rank the rules. This structure allows the knowledge base to deal with variable inputs and is also useful in conflict-resolution.

The ITRULE algorithm was implemented on a trouble ticket database and was found to be a useful data summarizing tool, recording how specific problems should be closed out and how long it should take to do so. It was also implemented on an alarm database, outputting a real-time prioritized list of the most important network alarms.

This paper by Goodman and Latin has some valuable contributions to the area of fault diagnosis. Most rule-based systems are of the symptom-class variety. The ITULE algorithm does not necessarily result in a diagnosis type of expert system, and it can be used as a data summarizing tool or detection type of system. The article shows that machine learning techniques can be used to overcome the difficulty of knowledge acquisition and make information available in the convenient format of rules. Another valuable addition is the idea of utility ranking of rules. In a real time system, when faced with the problem of conflict resolution, it is beneficial to know the cost of and value in firing the different rules.

3.2 Pattern Recognition Techniques For Proactive Network Analysis

The ITRULE algorithm has demonstrated the use of machine learning techniques in analyzing large quantities of information. Pattern recognition techniques have been devel-

oped to analyze network information in real time, allowing proactive detection of network faults. Two proactive network diagnostic methods employing pattern recognition have been reviewed. The first, by Maxion [12] develops a model of normalcy for a network's behaviour. Deviation from the model is indicative of a fault. The second, the LEAPR Diagnostic System [10] uses a Backpropagation Neural Network to characterize normal and abnormal traffic load on a router's lines. These two techniques are discussed below. A third article has been reviewed which uses pattern recognition techniques to control service quality in an asynchronous transfer mode (ATM) communication network.

3.2.1 Anomaly Detection in a Local Area Network - R.A. Maxion

R.A. Maxion describes an Ethernet anomaly detector. LAN traffic data was collected and curve fitting techniques were used on the traffic data to abstract a background model of normal behaviour. Weighting factors defined a learning and forgetting period to track traffic pattern drifts. Tolerance envelopes around the background were established. Real time data is collected and compared to the background model. Performance outside the tolerance envelope is indicative that a fault has occurred.

The diagnostic technique was trialed on the Carnegie Mellon University (CMU) Andrew network. It was applied conservatively and the performance was quite good - 100% of human detected faults were detected, it exhibited noise resistance, sensitivity and reliability when detecting anomalies. The Andrew network experienced an average of 12.7 anomalies per day during October 1989, 4.7 of which were alarmed using this diagnostic technique.

Maxion's work shows that volumes of data, unmanageable by humans, can be represented in a simple, meaningful and productive way, suitable for use in real time analysis. Technicians were freed from tedious data analysis activities. Maxion demonstrates that traffic analysis on a real time perpetual basis is possible.

Maxion has also presented a good characterization of communication networks - they are highly dynamic, noisy, and without global control. The nature of network protocols, designed to recover from faults, makes the detection of faults in a computer network very difficult.

3.2.2 The LEAPR Knowledge Source

The LEAPR (Line Element Anomaly Pattern Recognition) Knowledge Source was built as part of the Pegasus WAN fault diagnosis project and uses Neural Network techniques to classify the traffic on a router's lines as normal or abnormal.

LEAPR uses as input the traffic statistics for four of a router's lines - the top and bottom two performers. A back-propagation Neural Network classifies the statistics as being normal or abnormal. An adjustable Repetitiveness Counter (RC) is used to increase the accuracy of the Neural Network output. A diagram of LEAPR's structure is shown in Figure 3.0.

FIGURE 3.0 LEAPR Knowledge Source Structure

The LEAPR Knowledge Source has been found to be quite flexible - it is capable of monitoring different routers with different numbers of lines. The RC threshold is variable and allows fine tuning of the fault detection rate (ie: increase fault detection sensitivity at the risk of generating false alarms or decrease fault detection sensitivity while decreasing the risk of false alarms). An RC value of 1.91 yields 100% fault detection accuracy with an average of 3.54 false alarms per day for 10 different routers during normal operations. At the other extreme an RC value of 5.22 yields 34% detection accuracy with 0 false alarms per day. The LEAPR Knowledge Source can be used proactively - detecting abnormal network behaviour, perhaps before component degradation leads to complete component failure. The LEAPR Knowledge Source has demonstrated that neural network techniques can be used to establish traffic patterns in communications networks and detect faults that cause these patterns to be altered.

3.2.3 ATM Communications Network Control by Neural Networks - A. Hiramatsu

Hiramatsu [18] investigates a learning method using neural networks to control service quality in an asynchronous transfer mode (ATM) communication network. More specifically, back propagation neural networks are used to learn the relationship between the offered traffic and service quality. While his study is not directly in the area of fault diagnosis the problem is similar enough that his ideas are relevant.

ATM networks are highly flexible networks requiring flexible control of service quality. Different users and services of ATM have different requirements, and the traffic characteristics and quality requirements of the network change over time. As an example, the quality control parameters for cell transmission delay and cell loss rate are characterized by nonlinear functions with many inputs including: network component status, cell arrival rate, call generation rate, component utilization, and bit rate fluctuation. Hiramatsu divides ATM control into three areas, cell transfer level, call control level and network control level. He addresses the call control level (ie: acceptance or rejection of calls) in this paper,

and uses neural network techniques to determine decision values for acceptance or rejection of calls.

Hiramatsu provides a well-thought analysis of the problems ATM network control systems must deal with. When developing network management techniques for ATM networks it is likely these same problems will be encountered, hence the relevance of his work.

Hiramatsu observes that the call control portion of ATM control has a large number of inputs which are constantly changing. He suggests that traditional methods of call control cannot handle this. On the other hand, neural networks can learn a nonlinear function with many inputs and outputs with short processing speed. Hiramatsu implements a back-propagation neural network to learn the relations between the offered traffic and service quality. Hiramatsu tested his network on single and multi bit rate simulators with positive results. Hiramatsu suggests that in a real environment neural networks could be trained off-line using simulators to get initial weight values, and trained on-line for fine tuning.

While Hiramatsu was concerned with network control, his work has relevance to the field of fault detection. Some interesting issues must be considered when developing fault diagnosis techniques for ATM networks. The wide-ranging assortment of traffic types will result in a large variation of background "noise" from which to pick out faulty behaviour. Neural networks (or other pattern recognition techniques) must be able to pick out faulty patterns from a background which varies, suggesting a need for multiple models of normal behaviour. This will complicate the training and adaptability of neural networks. Which training algorithms will be the best?

3.3 Probabilistic Technique for Reactive Break Fault Detection

As discussed previously, analyzing the large amounts of information available can be a very difficult task. A complicating factor is the complex interactions and relationships

between components in a network. BANES (Basic Analysis in Nested Evidence Spaces) is a diagnostic technique which is used to classify objects and provide a framework upon which to collect evidence regarding the state of the components. BANES is used to pinpoint breakfaults in a monitored system [8]. BANES uses nested evidence spaces to classify objects and provide a framework for evidence accumulation.

Nested Evidence Spaces is a novel technique to classify objects according to their attributes. In an ordinary classification system, one attribute is generally subordinated to another, for example, colour may be subordinated to shape. Nested Evidence Spaces avoids this problem by allowing both AND and OR (conjunctive and disjunctive) relationships in the attribute hierarchy. This hierarchy is used to collect and propagate beliefs regarding the status of the hierarchy's components.

In BANES' application to the diagnosis of communication networks, the network topology is used to build the hierarchy and SNMP status records from the network's DCDs initiate the propagation of evidence. Components are marked as OK, Possibly Broken, or Broken, with an accompanying belief measure. Testing results indicate BANES has a 99.93% accuracy in diagnosing broken network components.

BANES is a novel diagnostic technique that seems to have overcome the difficulty diagnostic systems had in handling beliefs in large scale. The complexity of calculations increases linearly with the increase in topology. BANES has also established a meaningful method of representing the relationships between the attributes of components.

3.4 The Role of Cooperating Processes in Diagnosis

Some of the diagnostic applications discussed previously in this chapter are stand-alone systems while others are used in conjunction with other diagnostic techniques. Some of the recent work investigating the role of cooperating processes in diagnosis has been reviewed. This includes papers by Nii [4,5] discussing the use of blackboard systems

in general and as platforms for the interaction of cooperating processes. The Pegasus Project is also discussed from the viewpoint of the interaction of cooperating knowledge sources.

Nii [4,5] presents a well-thought out analysis of Blackboard Architectures. A blackboard system is defined as a relatively complex problem-solving model prescribing the organization of knowledge, data, and the problem solving behaviour within the overall organization. It provides a flexible approach to problem decomposition and knowledge acquisition.

Further, and in more general terms, blackboards allow the development of independent, specialized knowledge sources. These knowledge sources can be used together to build solutions to difficult problems that on their own, they could not reach. The use of cooperating processes is a powerful feature of blackboard systems. Cooperating knowledge sources can also be coordinated with planning strategies [19,20,21].

In the Pegasus Project, the END [9] and BANES [8] knowledge sources are used in a cooperating manner in two instances. First, one KS can be used to diagnose a situation that the other KS is unable to diagnose correctly. Second, one KS can be used to double check the validity of the hypotheses generated by the other KS.

The BANES knowledge source can be used to overcome the weaknesses of the END knowledge source. END has problems correlating multiple apparent faults to one real fault. Elements close to a data collector carry status information for elements further from the data collector. If an element close to the data collector is broken, then from the data collector's perspective all elements beneath that component are broken as well. END has difficulty making the determination that only one component is truly broken. END's weakness is the strength of BANES, which, using its nested evidence spaces has no difficulty correlating the multiple apparent faults to the one single true fault.

The END modules Veri-line and Veri-element were developed to verify the hypotheses output from BANES, and they were found to increase the accuracy of BANES; END validated BANES' correct diagnoses and corrected the diagnoses that were in error.

Using knowledge sources in a cooperating manner can result in fewer incorrect diagnoses and a wider range of diagnosability.

CHAPTER 4.0 The Problem of Detecting Abnormal Router Behavior

This chapter presents the problem tackled in this thesis: the detection of abnormal router behavior. It is divided into three main sections: First, a definition is given of the problem that this thesis tackles. Second, with reference to state of the art diagnostic procedures discussed in Chapter 3, it is shown that a new approach to this problem is warranted. Last, it is shown that solving this problem will be an advance in the area of network fault diagnosis.

4.1 Detecting Abnormal Router Behaviour: Problem Definition

The purpose of this thesis is to develop an algorithm that can process large amounts of network performance data and from this detect the abnormal behaviour of key network components such as routers. The term abnormal behaviour implies that a router is functioning, but is not functioning correctly, as opposed to a situation where the router is not functioning at all. The algorithm will extract useful information from a much larger volume of information and present it to the network operator in a concise manner. It is intended that this algorithm could be used in a proactive or a reactive manner, and in dif-

ferent types of communication networks. The factors affecting the utility of this algorithm will be determined. These points are expanded below.

Break faults in communication networks, where a component is rendered inoperable, are relatively easy to detect - their absence in the network is quite apparent and alarmed accordingly. Non-break faults, where components misbehave but are still operational are, on the other hand, quite difficult to detect. There may be no explicit alarms to indicate this kind of misbehaviour in the network, and a fault may remain undetected and possibly worsening until a user calls the network operations centre to complain of poor network performance. Network operators are inundated with information - network alarms, status messages and performance statistics. It is typically not possible to analyze all of the information at hand, particularly when there is a crisis somewhere in the network. Tools that can digest the information available and present it in a clear, concise and correct format to a network operator are valuable. Network operators work to determine the source of known problems, as well as to detect a problem and determine its source before it becomes noticeable to users. It is beneficial to have a diagnostic technique that can be used both proactively and reactively. The term reactive implies that a problem exists in a network and is quite apparent to users or network operators before a knowledge source is used to determine the source of the problem. The term proactive implies that a KS will detect a problem in its early stages of development before it is noticeable to network users.

Complex systems like a communication network can exhibit many types of faults and in many different ways. Some diagnostic techniques may detect some types of problems but be unable to detect others. By understanding the factors affecting the utility of an algorithm, it is easier to understand the boundary of its diagnostic domain.

4.2 There is a Need For an Algorithm to Detect Abnormal Router Behaviour

The following analysis of current research in the field of network fault diagnosis provides the justification for the work done in this thesis.

Rule based diagnostic systems, such as the END Knowledge Source [9] have been shown to be effective techniques for the automation of human diagnostic procedures. They too, however, have difficulties associated with them. As an example, the knowledge acquisition process can be quite cumbersome; it is difficult to arrange meetings with the busy human diagnostic experts and tedious to extract the proper information from them. Goodman and Latin have presented the ITRULE algorithm as a means of alleviating some of the problems associated with rule based systems, but difficulties remain. Rule based systems may be difficult to maintain, especially for telecommunications networks which are constantly evolving, and whose diagnostic techniques may be easily outdated. More generic diagnosis techniques may be desirable. Rule based systems are also generally reactive - the problem already exists before analysis begins. It would be beneficial to develop more proactive algorithms. Also, rule based systems primarily automate procedures that humans are very capable of doing on their own - it would be useful to develop diagnostic techniques that extend the capability of a human operator (ie: value added techniques), such as the near real time analysis of vast quantities of statistical information. To summarize, although rule based diagnostic systems are effective in some areas of diagnosis, there is still a need for the development of more generic, proactive and value added diagnostic techniques.

Pattern recognition techniques have been shown to be able to manipulate large quantities of network statistics and make meaningful determinations from them. Maxion's work [12] and the LEAPR Knowledge Source [10] presented in this thesis can be considered to be successful proactive error detection techniques; however, they are not without limitations. The success of pattern recognition techniques relies to some extent on the way in

which they are trained. The LEAPR knowledge source, for example, is trained off-line by presenting to it both abnormal and normal cases. The difficulty lies in defining which cases are normal and which are abnormal, and how one makes that distinction. Pattern recognition techniques such as Maxion's work are also affected by a learning and forgetting period. Maxion's algorithm uses weighting factors to continually adjust its model of normal behaviour. It is possible that gradual degradation may be incorporated into the model of normal behaviour and not necessarily identified as a problem. Both Maxion's work and the LEAPR K.S. focus on a router as an independent entity. The topology around the router and the behaviour of neighbouring components is not directly considered. For example, the Top-Two Bottom-Two Algorithm in the LEAPR K.S. selects the busiest and least busy lines on a router to provide a measure of the router's performance, thus making the algorithm independent of the number of lines in the router being monitored. It is not unreasonable to suggest that in some cases particular lines are always the top performers and other particular lines are always the lowest performers. A fault causing a shift in network traffic may cause these role to be reversed and if the net traffic measure is the same in both cases, then the algorithm will not flag the change in traffic pattern. Perhaps some traffic analysis should be done that is more directly tied with the topology of the network. This is not to say that Maxion's work and the LEAPR knowledge source are deficient in some way, just that there are other frames of reference upon which network performance data can be analyzed and that it is worth considering other algorithms to analyze network performance data. Hiramatsu's work using neural networks in ATM network control supports the idea that machine learning techniques can be used to analyze information in real time and that work in developing these types of algorithms is of value.

The BANES probabilistic diagnosis technique discussed in chapter 3 has been shown to be very effective at detecting break faults in communication networks, through the analysis of network status messages. BANES has not been used to analyze performance statistics from the network. This is an important source of information and algorithms that can

make meaningful observations from it are very valuable. The BANES algorithm has not been applied to detect the abnormal behaviour of network components and is not readily suited for it. Many communication problems in a network are caused not by broken components but by components that are not behaving properly. These conditions are not always “alarmed”, and the status messages that BANES analyzes may contain no indication that this type of problem exists. To summarize, BANES has been demonstrated as an effective algorithm to analyze network status messages, but algorithms to analyze network statistics (eg: performance data) are still in demand.

The work done by Goodman and Latin [13] in using machine learning techniques to generate rules from databases has demonstrated the feasibility and effectiveness of data summarizing tools. Traffic information is collected from communication networks in large quantities and is difficult to analyze in “real time”. The information typically is gathered in a database and may be analyzed at a later date to evaluate and make improvements to the network topology. Data summarizing tools like the ITRULE algorithm developed by Goodman and Latin, which can analyze traffic information and make concise summations in real time would tap a very valuable but for the most part previously untouched resource.

In Chapter 3, the role of cooperating processes in diagnosis was discussed. Many knowledge sources can work cooperatively on the same problem - each with a different angle on the problem, a different piece to the same puzzle, the end result being better than if the knowledge sources had worked independently. Each algorithm or knowledge source has certain limitations in its applicability that may be overcome by another type of knowledge source. It is therefore wise to develop a variety of algorithms for the analysis of communication network management data. The different views and cross-sections of data in the different algorithms may prove to be complementary and result in more comprehensive fault detection.

4.3 An Algorithm To Detect Abnormal Router Behaviour Will Be Beneficial

An algorithm that can analyze large quantities of data and detect the abnormal behaviour of network components is a valuable addition to a network operator's toolset for many reasons, as outlined below.

Fault detection systems can be developed to automate current procedures used by diagnosticians or to add new procedures to a diagnostician's toolset. The algorithm developed in this thesis falls into the latter category and is intended to supplement the toolset used by network diagnosticians. As such, it is valuable because it increases the diagnostician's ability to detect network problems that were previously difficult to detect, specifically the detection of abnormal router behaviour. It also allows diagnosticians to "tap" new resources such as the large quantities of network traffic statistics that are difficult for humans to analyze manually. By analyzing large quantities of information and representing it in a meaningful way, the algorithm developed in this thesis could be used to supplement the network diagnostic toolset and increase the network operator's ability to detect faults.

The algorithm developed in this thesis is intended to focus on the detection of non-break faults. The detection of break faults, through the use of tools like BANES have become much more easily diagnosable than non-break faults. Non-break faults typically do not have alarms or status messages that pinpoint their location. Abnormal component behavior is often made apparent by poor file transfer performance. The development of techniques to detect these types of faults, that otherwise could exist in a network indefinitely, is very valuable.

It is also important to develop tools that can be used in an operative manner as well as a reactive manner. Detecting the presence of a fault before it is noticeable to the network user reduces network down-time and increases user's productivity.

It is hoped that the type of algorithm developed in this thesis will be generalizable- i.e.: it will not rely on specific types of performance statistics, and be applicable to different types of communication networks.

To summarize, the problem tackled in this thesis is worthwhile because: it will analyze large quantities of information that humans have difficulty analyzing, extract the useful information and present it to network operators in a much reduced and more meaningful form. Network operators will find it easier to detect faults that previously were difficult or impossible to detect. They will be provided with a useful summary of a large amount of information that cannot realistically be analyzed "by hand". The algorithm developed will be able to operate in both a proactive and reactive manner, and it may be applicable to different types of networks.

CHAPTER 5.0 The Development of an Algorithm to Detect Abnormal Router Behaviour

This chapter documents the development of an algorithm to detect abnormal router behaviour in a Wide Area Network. The type of abnormal router behaviour studied in this thesis is the insertion of a packet processing (or throughput) delay in a router. Each packet that passes through the router has an additional delay added to its total processing time. Through discussion with several industry groups, the Pegasus group learned that this throughput delay symptom is often a primary indicator of a fault, and one that would be recognized by a network user. There may be other fault indicators, but they would not be as obvious to the user. A fault producing this type of symptom can be generated by many varieties of software bugs in routers, by misconfiguration and by overuse.

There were four main research phases in tackling the problem of detecting abnormal router behaviour. These phases are shown in Figure 5.0. The “fifth” phase in the figure

represents some investigative work into future research related to the algorithm developed in this thesis.

FIGURE 5.0 Research Phases

The four research phases are summarized as follows:

Phase 1: Effects of a Fault (Router Delay) on File Transfer and Router Performance:
The first phase of research is meant to establish, in simple terms, how a fault like a router

delay manifests itself in the performance of a router, and in the performance of file transfers passing through the router.

Phase 2: Can A Router With a Fault (Delay) be Discriminated From Its Fault-Free Neighbours? Given the manner in which a fault manifests itself, the second phase of research establishes how a faulty router can be discriminated from its healthy neighbours. The term neighbours refers to the routers connected to a given router (via serial lines or an ethernet).

Phase 3: Define And Evaluate an Algorithm For Router Fault Detection: Based on the observations of the first two phases of research, an algorithm to detect faulty router behaviour is defined. During this third phase, the utility of the algorithm is explored under different conditions.

Phase 4: Establish Alarm Thresholds and Monitor The Algorithm's Success At Detecting Router Faults: This last phase of research focuses on evaluating the performance (detection rate, false positive rate) of the algorithm defined in the third phase of research. Alarm thresholds are established and a large number of routers are monitored for long periods of time. The success rate of the algorithm at detecting delays inserted in the different routers is established.

Future Research: Some preliminary investigation of improvements to the algorithm are presented.

The remainder of this chapter is divided into five sections to describe the four phases of research and the preliminary investigation into future research areas. At the start of each section, Figure 5.0 is redrawn to illustrate how each section fits together. In some cases the figure is expanded to show the major steps involved in each particular phase of research. In each section, the experimental work and important conclusions of each phase (which form the basis for each succeeding phase of research) are presented. Any upgrades

made to the Pegasus Simulator to facilitate the research done in this thesis are also described.

5.1 Phase One: Studying the Effect of a Router Delay

FIGURE 5.1 Research Phases - Currently in Phase One

The first step in the process of developing an algorithm to detect abnormal router behaviour was to study the effect of an abnormal router delay on a file transfer and the performance impact on the router where the delay was inserted. Identification of the ways in

which the delay manifested itself provided the basis for developing a technique to detect and pinpoint the source of the abnormal behaviour. Figure 5.1 above shows this phase of research in relation to the other 3 phases of research.

Since the initial focus of research was on the effect of abnormal router behaviour on network traffic conditions, including file transfers, the following changes were made to the simulator:

The first major change to the Pegasus simulator was the addition of file transfer capability. Using the simulator interface, users can invoke (and cancel) file transfers between workstations. The user specifies source, destination, file size, packet size, resend timeout, and maximum retry parameters. An "odometer" output informs the user of packets transferred and the completion of the file transfer.

The second major change to the Pegasus simulator was the ability to introduce abnormal behaviour to a router. The kind of abnormal router behaviour studied in this thesis is the insertion of a packet processing delay in a router. Through the Pegasus simulator operator interface, the user selects a router and introduces a time delay composed of a base delay (in 5 micro-second units) and a random scaler (a small number, 1-4) that introduces a slight variability in the delay.

The above two simulator upgrades were sufficient for the initial stages of this thesis work. As the research became more focused, more upgrades were made to the simulator. These are presented where appropriate.

It is worth noting that the routing algorithm used in the Dragon simulator employs the functionality of the more common industry routing algorithms. The reaction to the insertion of a packet processing delay observed in the Dragon network simulator is deemed to be consistent with that of a real network.

A knowledge source was written that would collect the line statistics for a chosen router (nRTPc2) every minute. The knowledge source polled the Pegasus simulator for the bytes in and bytes out for each of a router's lines. These line traffic measurements were used to calculate the total bytes in and bytes out of a router for a one minute period. Every five minutes the knowledge source would calculate the average number of bytes in and out of a router in one minute. This would be done for a specified number of 5 minute intervals. A file transfer (10 megabytes, 512 byte packet size, 2 seconds before timeout, 2 retries) was invoked that would traverse router nRTPc2. See Figure 5.1 for a diagram of the file transfer path. The average bytes per minute into the router (over 5 - 1 minute intervals) was printed for the following states of nRTPc2:

- nRTPc2 with background snmp traffic only
- nRTPc2 with background traffic, and 1.25 second delay added to nRTPc2
- nRTPc2 with file transfer active
- nRTPc2 with file transfer active and 1.25 second delay added to nRTPc2

The information calculated by the knowledge source can be seen in Table 5.0 below.

Type of traffic through nRTPc2	Mean bytes in per minute for nRTPc2 (5 - 1 minute intervals)
normal traffic (background snmp)	5076.8
normal traffic + delay	4777.7
ftp traffic	67411.5
ftp traffic + delay	17747.0

TABLE 5.0 Mean Bytes In For nRTPc2 Under Different Conditions

FIGURE 5.2 File transfer path through nRTPc2

As can be seen from Table 5.0, the addition of a large file transfer through nRTPc2 greatly increased the traffic level compared to the normal background snmp traffic level (from 5076.8 bytes per minute to 67411.5 bytes per minute, a 1228% increase). The addition of a delay in router nRTPc2 decreased the amount of traffic processed by the router. The delay in nRTPc2 had a small effect in the low traffic level case (traffic level drops from 5076.8 bytes per minute to 4777.7 bytes per minute, a 5.9% decrease), and a very significant effect in the high traffic level (ftp) case (traffic level drops from 67411.5 bytes per minute to 17747.0 bytes per minute, a 74% decrease).

Qualitative observations during the data gathering include:

The file transfer was considerably slower when there was a delay inserted in node nRTPc2.

When the file transfer was inserted in the network the traffic increase observed in router nRTPc2 took place on the lines the ftp was using (ie: nRTPc2's lines to nRTPc1 and bRTPc2 which are also on the file transfer path). It was thus possible to trace the trajectory or path of the file transfer.

Key Observations and Direction Setting:

Two key observations can be made from this phase of experimentation:

1. A delay inserted in a router slows down a file transfer passing through that router. This is something a network user would observe.
2. A delay inserted in a router decreases the amount of traffic a router processes in a given interval. This is more readily observed if the router is loaded heavily. This observation is something a network operator would observe.

The above two observations provide the basis for two directions to pursue in creating a router delay fault detection system:

1. A file transfer that appears slow may be the victim of a delay in a router along the file transfer path. It seems possible to isolate the router that has the delay by systematically bisecting the path and timing small file transfers over the path segments. The segment that exhibits poor file transfer performance likely contains the faulty router. This technique is rather intrusive, and also could only be used in a reactive manner. It is worth noting that this technique is used by network operators to solve ftp problems in live networks. They observe the same slow ftp behaviour (as do the users who complain).
2. Given that a delay in a router decreases the amount of traffic that router can process, careful monitoring of router statistics (average bytes into / out of a router in a given time period under varying loads and varying node delays), will provide an indication as to the health of the router. The next step in developing an algorithm to detect abnormal router behaviour is to determine how a faulty router can be discriminated from its healthy neighbours.

Because of the limitations of the first direction (it is intrusive and can only be used in a reactive manner), it was felt that the second direction was the most promising, and it provides the motivation for the second phase of research. The first phase of research will not be investigated in the remainder of this thesis.

5.2 Phase Two: Can A Router With a Delay be Detected and Discriminated From Its Neighbours?

FIGURE 5.3 Research Phases - Currently in Phase Two

The focus of this second phase of research was to determine, by collecting and analyzing the routers' statistics, whether or not a router with a fault could be detected and discriminated from its healthy neighbours. Figure 5.3 above shows this phase of research in

relation to the other 3 phases of research. As indicated in the figure, there were two main steps in this phase of research.

First, the performance of a router with a delay was compared to two of its neighbours, one which shared a common file transfer, and the other which did not share the common file transfer.

Second, the performance of a router with a delay was compared to two of its neighbours, where a large number of file transfers were criss-crossing through all or some of the routers, simulating a more randomly generated traffic pattern.

The results of these two studies provided a platform on which to build an algorithm to detect a router which has a fault (packet processing delay) associated with it.

5.2.1 Step 1: Router With A Delay, Two Neighbours, 1 File Transfer

A knowledge source was developed that would probe a router for its statistics every minute. At the end of 15 one minute intervals the mean bytes into the node per minute was calculated and printed, as well as the Standard Deviation from this mean over the 15 intervals. The knowledge source would do this for a specified router or routers for the specified number (N) of 15 - one minute intervals. The measures used in the calculations are:

$$M = Mean = \frac{\sum_{n=1}^N (Xbytes/minute)}{N}$$

$$StandardDeviation = \sqrt{\frac{\sum_{n=1}^N (M - Xbytes/minute)^2}{N}}$$

FIGURE 5.4 Topology: Router with a delay, two neighbours, one file transfer

The topology shown in Figure 5.4 was chosen for the study. Router nRTPc2 was the router that has a delay inserted in it. A file transfer was set up through nRTPc2 and neighbouring node nRTPc1. The neighbouring router nATLc1 did not share a file transfer with nRTPc2. The file transfer was set up between routers bRTPx282 and nNSHx31 as indicated in Figure 5.4. A file transfer of this size would typically boost traffic levels by a factor of 10 (eg: from 6000 bytes per minute to 60000 bytes per minute).

Traffic statistics (router's mean bytes in per minute) were gathered for nRTPc2, nRTPc1, and nATLc1 for the following states:

- normal background (SNMP only) traffic, no delays or file transfers
- normal background (SNMP only) traffic, 1.25 second delay on nRTPc2
- file transfer through nRTPc2, nRTPc1, no delay on nRTPc2
- file transfer through nRTPc2, nRTPc1, 1.25 second delay on nRTPc2

To compare the performance of the different routers, the ratio of mean bytes in per minute between nodes was calculated for each state. Fluctuation in these ratios would give an indication of the impact of a delay on a node relative to its neighbours. These ratios are included in Table 5.1. The shaded columns on the right hand side of the table are the inverse of the ratios in the unshaded columns on the left hand side of the table, and are included for completeness. To establish an appreciation for the amount of noise in the signal, the ratio of the Standard Deviation to the Mean for each node in the different states was calculated; this information is included in Table 5.2.

Type of Traffic	Ratio of nodes' mean bytes in per minute					
	$\frac{nK/Pc2}{nRTPc1}$	$\frac{nRTPc2}{nATLc1}$	$\frac{nRTPc1}{nATLc1}$	$\frac{nRTPc1}{nRTPc2}$	$\frac{nATLc1}{nRTPc2}$	$\frac{nATLc1}{nRTPc1}$
normal	2.15	1.55	0.72	0.47	0.65	1.39
normal + delay	2.36	1.73	0.73	0.42	0.58	1.36
ftp	1.05	20.1	19.2	0.96	0.05	0.05
ftp + delay	1.2	5.5	4.6	0.83	0.18	0.22

TABLE 5.1 Ratios of Node In-Mean-Bytes Per Minute for Different States

From the Table 5.1, it can be seen that the addition of a delay in nRTPc2 did not have a significant impact on the traffic ratios under "normal" "SNMP only" traffic conditions (very little change between rows one and two in the table).

The addition of a file transfer increased the traffic flow in nRTPc2 and nRTPc1 in relation to nATLc1, which is not on the file transfer path (From rows 1 to 3 the nRTPc2:nATLc1 and nRTPc1:nATLc1 ratios increased from 1.55 to 20.1 and 0.72 to 19.2 respectively). Both of the routers on the file transfer path (nRTPc2 and nRTPc1) had approximately the same rate of traffic flow (as indicated by the row 3 nRTPc2:nRTPc1 ratio of 1.05).

The addition of a delay in nRTPc2 while the file transfer was in progress dropped the traffic flow in both nRTPc2 and nRTPc1 by a very similar amount (from rows 3 to 4 the ratio nRTPc2:nRTPc1 changes from 1.05 to 1.02 with the addition of the delay while the nRTPc2:nATLc1 and nRTPc1:nATLc1 ratios drop from 20.1 to 5.5 and 19.2 to 4.6 respectively). Again, the row 4 nRTPc2:nRTPc1 ratio of 1.20 shows that the routers still had nearly equal traffic levels.

It would appear that a delay in a router throttles the file transfer traffic, dropping the traffic throughput of the router with the delay. This throttling effect seems to be passed on to other routers on the file transfer path (as seen in router nRTPc1), but this experiment does not exhaustively investigate the extent to which the throttling is passed along the file transfer path.

The throttling of traffic is not surprising. The processing delay in the router creates a bottleneck along the file transfer path. For the router with the processing delay, throughput is decreased as packets are queued up waiting for their turn to be processed (passed to another router or ethernet). As the delay in the router increases, fewer packets are processed in a given period of time.

For the file transfer on whose path the delayed router lies, router performance drops. In most file transfer protocols, a limited number of packets will be sent to the destination. Once acknowledgment for these packets is received from the destination, more packets are sent. A delay along the file transfer path slows the send - acknowledge sequence, and the file transfer takes longer to complete.

Type of Traffic	(Standard Deviation/ MEAN)*100%		
	nRTPc2	nRTPc1	nATLc1
normal	10.3	13.5	21.7
normal + delay	17.3	23.8	44.6
ftp	2.0	1.5	25.1
ftp + delay	7.4	5.9	30.4

TABLE 5.2 (Standard Deviation/Mean In Bytes)*100% For Different States

Table 5.2 above shows the ratio of the Root Mean Square Standard Deviation to the Mean Bytes into the router per minute over the 15 one minute intervals. This was calculated for each node in the four different states. It can be seen that the deviation from the mean dropped as the traffic level was increased (between rows 1 and 3 the ratio of Deviation : Mean for nRTPc2 drops from 10.3 to 2.0 with the addition of an ftp through nRTPc2 while the ratio for nATLc1, which is at a lower traffic level, actually jumps a bit from 21.7 to 25.1). The fact that the ftp traffic is larger and less variable than the SNMP traffic is the likely cause of this effect.

Step 1 Key Observations and Direction Setting:

Three key observations can be made from this work:

1. A packet processing delay inserted in a router causes the file transfer traffic throughput to drop in that node. The delay is also reflected, to a degree, in the next node along the file transfer path.

2. The delay creates a bottleneck along the file transfer path, slowing the file transfer.
3. The traffic level Standard Deviation from the Mean decreases as the traffic level in a router increases. This will be considered in more detail in section 5.3.

Based on the above observations, further study is warranted in the following area:

1. Investigate the traffic levels of a router with a delay and its neighbouring routers with many criss-crossing file transfers providing more random traffic generation. Determine if the traffic level of the router with a delay drops and demonstrates different behaviour than the “healthy” routers around it.

5.2.2 Step 2: Router with a Delay, Two Neighbours, Many File Transfers

This step studied the effect that a delay inserted in a router had on that router and two of its neighbouring routers. A larger number of file transfers (7) were set up to create a pseudo-random high density traffic pattern. Some of the file transfers are common to all three of the routers, while others pass through only one or two of the routers. The simulator version used in this experiment was not able to run with more than 7 of these file transfers.

A knowledge source was written that would probe a router's lines for its traffic statistics every minute. At the end of 15 minutes the mean bytes into the node per minute and the standard deviation from the mean were calculated. As well, the mean bytes in and out for each line in a router, and the Standard Deviations from these means were calculated. The knowledge source would collect these statistics for a specified number of routers for a specified number of 15 - one minute intervals.

Figure 5.5 shows the topology used for this study. Router nRTPc2 was chosen to have a delay inserted in it. Routers nRTPc1 and nATLc1 are two of its neighbours. Seven file transfers were set up randomly around and through these routers to create a pseudo-random traffic flow. The file transfer parameters are indicated in the figure.

FIGURE 5.5 Topology: Router with a Delay, Two Neighbours, Many File Transfers

Traffic statistics were gathered for nRTPc2, nRTPc1, and nATLc1 for the following two states:

- 7 ftp's running, no delay on nRTPc2
- 7 ftp's running, 1.25 second delay on nRTPc2

The performance of the three routers was studied in two ways: Firstly, the total traffic for each router was studied and compared. Secondly, the traffic on a router's individual lines was studied.

The bytes into each router was determined for 15 one minute intervals both before and after the delay was added to nRTPc2. From this each router's "mean bytes in per minute" traffic was calculated. A ratio of this measure (i.e.: (after a delay was added to nRTPc2) / (before the delay was added to nRTPc2)) was calculated and can be seen in Table 5.3.

The total mean bytes in per minute for all 3 nodes was also calculated before and after the delay was added to nRTPc2. Each node's percentage contribution to the total traffic of all 3 nodes was calculated - both before and after a delay was inserted in nRTPc2. The ratio (% contribution before the delay) / (% contribution after the delay) was calculated. This information can also be seen in Table 5.3. Discussion of these results follows the table.

For each line in a router, the "mean bytes in per minute" over the 15 one minute intervals was calculated both before and after the delay was added to nRTPc2. A ratio of this measure (after a delay was added to nRTPc2) / (before the delay was added to nRTPc2) was calculated and can be seen in Table 5.4. The percent contribution of each line in a router to the total traffic in that router was calculated - both before and after a delay was inserted in nRTPc2. A ratio of this measure (after the delay was added to nRTPc2) / (before the delay was added to nRTPc2) can also be seen in Table 5.4. For nodes nRTPc1

and nATLc1, the shaded cells indicate the lines that are connected to nRTPc2, the router with the delay.

	Ratio of Node Mean Bytes In Traffic: (after delay) (before delay)	Node % contribution to total of all nodes (mean bytes in) for that router		
		before delay	after delay	ratio of after before
nRTPc2	0.21	31	16	0.52
nRTPc1	0.43	30	31	1.03
nATLc1	0.55	39	53	1.36

TABLE 5.3 Analyzing Node Level Traffic For Delay in nRTPc2

In the second column of Table 5.3 above, the ratio of mean bytes in (after delay / before delay) indicates that all three of the routers studied experienced a drop in traffic level after a delay was inserted in nRTPc2. The router with the delay, nRTPc2 experienced the most severe drop (0.21 of its previous level) while nRTPc1 and nATLc1 experienced less severe drops (0.43 and 0.55 respectively). This makes sense because all of the ftp's through nRTPc2 would be throttled by the delay, whereas in the neighbouring nodes, only the ftp's in common with nRTPc2 would be throttled by the delay in nRTPc2. In more general terms, all of nRTPc2's ftp traffic would be throttled to some extent by nRTPc2's delay while only a portion of the traffic in nRTPc1 and nATLc1 (that portion which is common with nRTPc2) would be throttled by the delay in nRTPc2.

The third, fourth and fifth columns of Table 5.3 above, (showing the node percentage contribution to the total traffic of all nodes) yields some interesting results as well. The node with the delay, nRTPc2, contributed a lower percentage to the total traffic of the three nodes after a delay was added to it (nRTPc2's contribution dropped from 31% to 16%). The contribution levels of the other two nodes nRTPc1 and nATLc1 rose (from 30% to

31% and 39% to 53% respectively as seen in the second and third rows, third and fourth columns). This phenomena is explained as follows. When a delay is inserted in a router, a bottleneck is created, and file transfer traffic through the router slows. The traffic throughput of the router decreases. The throttled file transfers decrease the neighbouring routers' throughput as well. However, in a situation where there are many file transfers active, only a portion of the neighbouring routers' traffic is dependent on the health of the delayed router. Thus, they will not experience the same performance loss as the delayed router. When comparing the percentage contribution of each router to the total throughput of all of the routers, the delayed router's contribution will be lower.

Node	Line, destination	Ratio Mean Bytes (after delay) (before delay)		Ratio of Line's contribution to total node traffic (after delay) (before delay)	
		In	Out	In	Out
nRTPc2	0 - e0.nRTP	0.16	0.24	0.77	1.16
	1 - e1.nRTP	0	0.15	0	0.71
	2 - nATLc1	0.22	0.14	1.07	0.67
	3 - nPARc2	0.22	0.27	1.05	1.28
	4 - hRTPc2	1.2	0.28	5.6	1.36
nRTPc1	0 - e0.nRTP	0.61	0.25	1.4	0.58
	1 - e1.nRTP	0.14	0	0.33	0
	2 - nANNc2	0.61	.21	1.44	0.49
	3 - nNSHc2	0.99	1	2.25	2.33
nATLc1	0 - e0.nATL	1.01	1.01	1.8	1.84
	1 - e1.nATL	0.22	0.14	0.40	0.25
	2 - nRTPc2	0.14	0.22	0.25	0.40
	3 - nLAKc2	1.01	1.01	1.82	1.83

TABLE 5.4 Analyzing Line Traffic Within a Node with Delay in nRTPc2

The third and fourth columns of Table 5.4 above show, for each line, the ratio of mean bytes in and mean bytes out (i.e.: (after the delay) / (before the delay)). The node with the delay, nRTPc2, experienced large traffic losses on all lines (after/before ratios range from 0.14 to 0.28, with the exception of line 4 in ratio of 1.2 - this is a low traffic volume, noisy line). It is significant that on the other two nodes nRTPc1 and nATLc1, their lines connected to the node with the delay (nRTPc2) experienced a drop in traffic after the delay was inserted in nRTPc2. The cells in the table corresponding to these lines are shaded. The other lines in nRTPc1 and nATLc1 did not experience a significant decrease in traffic level, or experienced a decrease that mirrored the decrease in the lines connected to nRTPc2. This latter phenomenon may be explained by the fact that the traffic generation using ftp's is not entirely random, and the throttling along the ftp path caused by the delay in nRTPc2 affected not only the lines connected to nRTPc2, but followed the trajectory of the ftp's through the nodes, affected some of the other lines not connected to nRTPc2.

The fifth and sixth columns of Table 5.4 above shows the ratios of a line's contribution to total node traffic (i.e.: (before delay added) / (after delay added)) for each line in the three routers. The ratios of nRTPc2's line contributions to the node total do not fluctuate as widely as those for nodes nRTPc1 and nATLc1. An exception to this is line 4 (to bRTPc2), which is a low traffic volume, noisy line. For routers nRTPc1 and nATLc1, the lines connected to nRTPc2 (the corresponding table cells are shaded) show a decrease in contribution to the node. Their other lines show an increase in contribution to the node, or mirror the decrease in contribution shown by the lines connected to nRTPc2.

Key Observations and Direction Setting

Four key observations can be made from this work:

1. When considering a router and its neighbours, if the router has a delay (fault) inserted in it then it will usually contribute a smaller percentage of processed traffic to the total of all of the routers than it did when it was healthy. The neighbouring routers tend to demonstrate an increase in the percentage contribution of processed traffic to the total processed traffic of all of the routers.
2. Further, looking at the lines of the router with the delay, the loss in traffic seems to be evenly distributed over all of the router's lines. This particular type of fault seems to affect all of the router's lines to the same degree.
3. Further to the second observation, looking at the lines of the neighbouring routers, the lines destined to the router with the delay contribute less to their node's total traffic. The other lines on these routers contribute a higher percentage of the router's total traffic. Some exceptions can be seen, but these may be attributed to the non-random nature of the traffic generation used in the experiment - a finite number of large ftp's.
4. Using file transfers to generate traffic is rather crude and non-random, since only a small, finite number can be invoked at a time (ie: less than 10). A better method of increasing traffic will be investigated in the next phase of research.

Based on the above observations, two main hypotheses come to mind:

1. Consider a router R and all routers connected to it to a depth of, say, 1. These routers make up a "zone". Each router in the zone contributes a percentage of the total traffic in the zone. If a delay is inserted in the router R, then it will not contribute the same percentage of traffic to the zone - in fact it will contribute less, essentially ending up in a "well" compared to its neighbours. The neighbouring nodes will experience a slight rise in percentage contribution, unless they are heavily dependent on the router R for a large portion of its traffic, in which case it may drop slightly as well.

2. In addition to the above, focusing on the neighbouring nodes, the lines connected to the router R with the delay will contribute less traffic to their node total. The other lines on the neighbouring routers will contribute a higher percentage to the node total. The poor performing lines (connected to the router with the delay) will together act as "vectors" pointing to the faulty node. Figure 5.6 below illustrates the concept of a zone around a router.

The first hypothesis is the simplest of the two, and computationally less expensive. It will form the basis of the fault detection algorithm defined in section 5.1. The second hypothesis will not be examined during the remainder of this thesis.

FIGURE 5.6 Zone Around a Router

5.3 Investigating an Algorithm To Detect a Router With a Delay

FIGURE 5.7 Research Phases - Currently in Phase 3

The third phase of research focused on defining and evaluating an algorithm to detect faults in routers. Figure 5.7 above shows this phase of research in relation to the other phases of research. As seen in the figure, there were four main steps in this phase of

research. First, the algorithm is defined. Secondly, the role of traffic mix (snmp type vs ftp type) on the utility of the algorithm was investigated. Thirdly, the impact of a router's delay in (i) the zone where it is in the center, and (ii) the other zones in which it is a border node, was studied. Lastly, long running simulations were done to determine what size of router delay the algorithm can detect and to get a firm grasp of the signal to noise ratio in the signal.

5.3.1 Step 1: Algorithm Definition

In the previous section, two hypotheses were presented regarding the detection of a misbehaving router. The first hypothesis (i.e.: comparing a router's performance to that of the routers in a zone around it) forms the basis of the algorithm to be investigated.

Algorithm: Consider a router R and all routers connected to it to a depth of, say, 1. These routers make up a "zone". The traffic throughput (defined as total bytes into the router) for each router in the zone is calculated for a given period of time. The total traffic throughput for the zone is defined as the sum of the individual throughputs for all routers in the zone. Each router in the zone contributes a percentage of the total traffic throughput in the zone. This percentage contribution is the critical measure for the algorithm, and it is used to demonstrate a faulty node's drop in performance relative to its neighbours. If a delay is inserted in the router R, then it will not contribute the same percentage contribution to the zone, in fact it will contribute less, and essentially end up in a "well" compared to its neighbours. The neighbouring nodes will experience a slight rise in percentage contribution; however, if they are heavily dependent on the router R for a large portion of its traffic, their percentage contribution may drop slightly. The algorithm will monitor routers in the network, specifically their contribution levels to their respective zones. If a router's contribution level drops from a base reference level by a given threshold, an alarm will be generated indicating the presence of a fault in that router.

5.3.2 Step 2: Role of Traffic Mix in Algorithm Performance

In the simulations done in section 5.2, it became apparent that using file transfers was not the best way to boost traffic levels in the network. Since only a finite number of file transfers can be invoked in the simulator, the traffic generated is not very random or evenly distributed, it is sort of an "all or nothing" way of generating traffic.

While investigating other ways of increasing the traffic level in the network, it became apparent that the type of traffic in the network had an impact on how sensitive the network was to the addition of a delay in a router.

A simulation was done in which network traffic was increased by increasing the rate at which snmp polling occurred. A zone around router nRTPc2 was selected for monitoring (the border nodes in the zone are nATLc1, nPARc2, bRTPc2 and nRTPc1), and traffic in the zone increased to approximately the same level as if a large number of file transfers had been invoked (approximately 8000 bytes/second for the entire zone). This was done by increasing the frequency with which Data Collector Devices (DCDs) polled network components, resulting in more SNMP packets in the network. Delays were inserted into router nRTPc2 (0 second delay, 1 second, 2 second, 6 second). Zone statistics were monitored (in a manner consistent with previous experiments - calculating a node's percent contribution to the zone) as the increasing delays were added. It was expected that router nRTPc2 would fall into a well in relation to its neighbours as the delay was increased. Instead, router traffic contribution levels remained consistent at the percentages shown in Table 5.5.

Size of Delay in nRTPc2	Contribution level of router to the zone				
	nRTPc2	nATLc1	nPARc2	bRTPc2	nRTPc1
0 seconds	27.8%	15.7%	6.9%	35.4%	14.2%
2 seconds	28.1%	15.6%	6.7%	35.3%	14.3%
4 seconds	28.0%	15.8%	6.7%	35.5%	14.0%
6 seconds	27.6%	15.9%	7.1%	35.2%	14.2%

TABLE 5.5 Router contribution levels, snmp traffic, varying delays in nRTPc2

The results in Table 5.5 indicate that snmp traffic is apparently not affected by a delay inserted in a router (there was very little change in router contribution levels to the zone when a delay was added). This is probably because the base frequency of snmp traffic (approximately three snmp polls per minute per device) was much greater than the delay. If a DCD polls a device once every 20 seconds then a 2 or 3 second delay in getting a response back does not affect the next poll to be sent out.

A study of the ratio of delay sensitive traffic (ftp) to delay insensitive traffic (snmp) and the corresponding impact of a router delay was deemed useful.

A zone around router nRTPc2 was selected. With no file transfers active in the network, a set of traffic statistics were gathered to determine the background level of traffic (mean node traffic contribution of nRTPc2 to its zone over 15 - one minute periods). Traffic levels were increased by adding finite numbers of file transfers. For each set of file transfers added the ratio of file transfer traffic to background SNMP traffic was calculated (these ratios were: 0,~2.5,~3.5,~4.7). For each set of traffic ratios the mean node traffic contribution of nRTPc2 to its zone over 15-one minute periods was measured before and after a half second delay was added to nRTPc2. The percentage drop in nRTPc2's zone traffic contribution was calculated and can be seen in Graph 5.1

GRAPH 5.1 Effect of 0.5 Second Delay in nRTPc2 as FTP/SNMP Traffic Ratio Varies

From the above graph it can be seen that as the ratio of sensitive to insensitive traffic was increased, the more a delay in nRTPc2 manifested itself, and the more nRTPc2 dropped into a well compared to the other routers in the zone. For example, with a traffic ratio of ~2.5, a half second delay dropped nRTPc2's traffic contribution by ~11% but with a traffic ratio of ~4.7, the delay dropped the contribution by ~25%.

Key Observations and Direction Setting

Two key observations can be made from this work:

1. In the simulation network there are two types of traffic: snmp traffic and ftp traffic.

SNMP traffic is not affected by a delay in a router. FTP traffic is affected by a delay in a router. SNMP monitoring devices send out polling packets at specific intervals to devices in the network. The speed with which these packets reach their destination and the speed with which the answer packets return to the monitoring devices has little effect on when the next SNMP poll is sent out. The polls are also sent out at a very low frequency. Thus, the insertion of a fault such as a router delay has little effect on SNMP traffic levels. FTP traffic is sensitive to router delays because of the concept of a window in file transfers and the high frequency with which packets are sent. The file transfer source will send a specific number of packets to the destination and will stop sending until some of these packets get acknowledged from the destination. As a router delay is increased the source has to wait longer for acknowledgements and the rate of file transfer decreases, resulting in a decrease in FTP traffic due to a throttling of the file transfer.

2. The ratio of sensitive to insensitive traffic determines to what degree a specific type of fault manifests itself, and in turn affects the ability of monitoring algorithms to detect that fault.

Based on the above observations, an important point becomes apparent when considering the problem of monitoring a real network for faults:

1. Traffic in a network is composed of many different types. Some types of traffic may be sensitive to some types of faults while being insensitive to others. A fault in a network may manifest itself differently depending on the mix of traffic in the network and their sensitivity to that type of fault. The fault, at one extreme, may not be very noticeable, or, at the other extreme, may be quite blatant.

5.3.3 Step 3: Impact of a Router's Delay in Various Zones

Implicit in the hypothesis under investigation, if a delay is inserted in a router, its performance level will be in a "well" compared to the routers in the zone around it. A router, however, will belong to more than one zone. It will belong to the zone where it is the center node, and it will also belong to the zones of the routers around it (where it will be a border node). It is worthwhile to observe the impact of a delay in a router in the zone around the router, and also in the zones where the router is a "border node". The concept of center nodes and border nodes is illustrated in Figure 5.6.

At this point it is worth noting that a new traffic generation feature was added to the Pegasus Simulator. At startup the user specifies a number of workstations to be added randomly throughout the simulator network, and a number of interactive file sessions to be invoked between these workstations. This feature provides a random generation of ftp-like traffic throughout the simulated network.

The router nRTPc2 was selected to have varying delays inserted in it (from 0.1 seconds to 2 seconds long). The zone around nRTPc2 was set to be monitored, as well as the zones centred around the border nodes in nRTPc2's zone. A large number of workstations and interactive file sessions were added to the network (3300 workstations, 1650 file sessions). With this level of traffic in the network and no delay inserted in nRTPc2 the mean percentage of nRTPc2's contribution to its own zone and in the zones of its border nodes was calculated. The percentage contribution of nRTPc2's border nodes was also calculated for each of their own zones, and nRTPc2's zone. The simulations and zone contribution calculations were repeated, each time a different delay was added to router nRTPc2. The change in each router's contribution level, as the delay in nRTPc2 was increased, can be seen in Tables 5.6 through 5.9

Router:	% Change in Router Contribution in Zone Centred Around:				
	nRTPc2	nATLc1	nPARc2	bRTPc2	nRTPc1
nRTPc2	-6.3%	-12	+1.3	-9.3	-3.4
nATLc1	+2.5	+3.1			
nPARc2	-8.3		-1.6		
bRTPc2	+7.8			+2.5	
nRTPc1	-2				+5.5

TABLE 5.6 Change in Router Contribution to Zone, nRTPc2 Delay = 0.1 second

Router:	% Change in Router Contribution in Zone Centred Around:				
	nRTPc2	nATLc1	nPARc2	bRTPc2	nRTPc1
nRTPc2	-7.1	-15	-5	-5.8	-2.6
nATLc1	+11	+2.6			
nPARc2	-2.1		-1.3		
bRTPc2	+6.3			-3.3	
nRTPc1	-12.3				+3.5

TABLE 5.7 Change in Router Contribution to Zone, nRTPc2 Delay = 0.5 second

Router:	% Change in Router Contribution in Zone Centred Around:				
	nRTPc2	nATLc1	nPARc2	bRTPc2	nRTPc1
nRTPc2	-15	-28	-15	-13	-0.4
nATLc1	+23	+7.7			
nPARc2	0		0		
bRTPc2	+12			-1.3	
nRTPc1	-14				+9

TABLE 5.8 Change in Router Contribution to Zone, nRTPc2 Delay = 1 second

Router:	% Change in Router Contribution in Zone Centred Around:				
	nRTPc2	nATLc1	nPARc2	bRTPc2	nRTPc1
nRTPc2	-18	-3.1	-12	-21	-4.7
nATLc1	+25	+9.2			
nPARc2	+0.8		-3.3		
bRTPc2	+13			-2.5	
nRTPc1	-15				+9

TABLE 5.9 Change in Router Contribution to Zone, nRTPc2 Delay = 2 seconds

In Tables 5.6 - 5.9, it can be seen that as the delay in nRTPc2 was increased, nRTPc2 contributed less to its own zone, and to the other zones where it was a border node (observe the continued drop in contribution in row 1 of the tables). The border nodes around nRTPc2 either increased their contributions to their own zone, or decreased slightly, but not nearly to the same extent as nRTPc2 did. In nRTPc2's zone, the border nodes exhibited fluctuations in both directions.

Key Observations and Direction Setting

Two key observations can be made from this work.

1. When a delay is inserted in a router the contribution to its own zone and the zones around it will drop. The larger the delay is, the more significant the drop.
2. When a delay is inserted in a router the contributions of the border nodes may increase or decrease in the delay node's zone, and will increase or decrease slightly in their own zones. Thus, a node that has no delay in it may exhibit a bit of a "well" effect due to circumstances in the network around it; however, the well effect seems less noticeable than if there were a delay in the node.

These two observations are supportive of the algorithm stated at the start of this section. They lead to the following steps in further verification of the algorithm:

1. To date, only short simulations have been done. Since nodes with no delay can exhibit a bit of a "well" effect (albeit a small one), it is desirable to do some long-running simulations to observe traffic fluctuations over a longer interval.
2. It is conceivable that normal network traffic fluctuations may cause nodes to display small "well" effects. Long running simulations will help to establish thresholds to classify "well" effects as normal or abnormal.

5.3.4 Step 3: Long Running Simulations

Implicit in the hypothesis under investigation, if a delay is inserted in a router, its performance level will be in a well compared to the routers in the zone around it. Traffic in a communications network has normal fluctuations that should not be misconstrued as indications of faults. Also, the ability to detect a delay (fault) in a router will depend on the size of the delay - the larger the delay, the more easily it is detected. The purpose of this section was to get a firm grasp on both of these factors by performing some long running simulations.

Two sets of long running simulations were done. The first set was done to get an idea of traffic fluctuations in routers (ie: mean, standard deviation over a long period of time) fluctuations in routers, and the long-running affects of a router delay in the router's zone. The second set of simulations was done to study a large number of zones in more detail, leading to the establishment of alarm thresholds.

For the first series of long-running simulations, a knowledge source was built that would monitor the statistics of a specified zone, as follows: Every two minutes the knowledge source would gather traffic statistics for the nodes. Five of these two minute periods composed an interval. After every interval the mean bytes into and out of each node was

calculated. The mean bytes in and out of each zone was calculated as the sum of the mean bytes in and out of each node in the zone. Each node's percentage contribution to the total zone traffic was also calculated.

After 100 intervals (16 hours and 40 minutes) of data was collected, the mean traffic contributions and standard deviations from this measure were calculated for each node.

The percent deviation from the mean traffic contribution was also calculated for each of the 100 intervals - to be used for a histogram of % deviation from the mean

Simulations were done for each of the following cases:

- node nRTPc2 with 0 delay (reference set representing normal behaviour)
- node nRTPc2 with 0.1, 0.2, 0.4, 0.5, 0.7, and 1 second delays.
- node bRCHc1 with 0 delay (reference set representing normal behaviour)
- node bRCHc1 with 0.1, 0.2, 0.4, 0.5, 0.7, and 1 second delays
- node nMTHc1 with 0 delay (reference set representing normal behaviour)
- node nMTHc1 with 0.1, 0.2, 0.4, 0.5, 0.7, and 1 second delays

Graphs 5.2 - 5.6 are histograms tabulating, for all of the 100 intervals monitored, the percent deviation of a router's traffic contribution (from the mean) to nRTPc2's zone, with zero delay in router nRTPc2. Each router's percent contribution to the zone and the Standard Deviation of the measure is indicated on the graph. Note that although this information was collected for the zones around bRCHc1 and nMTHc1, graphs are not shown for the sake of brevity (similar results were achieved).

Referring to these graphs (5.2 - 5.6), it can be seen that the distributions are well behaved. The Standard Deviation varies for each router. These two points have some significance in that there is some fluctuation in traffic contributions of routers over time, and the amount of fluctuation is different for each router.

Graph 5.7 presents each router's mean percentage contribution (over 100 intervals) to the total processed traffic nRTPc2's zone, as the delay in nRTPc2 is increased. The error bars in the graph indicate the first Standard Deviation from the mean. Graph 5.8 presents the percentage change in each router's (percentage) contribution to total zone traffic as the delay in nRTPc2 is increased.

Graph 5.9 presents each router's mean percentage contribution over 100 intervals to the total traffic in bRCHc1's zone as the delay in bRCHc1 is increased. The error bars in the graph indicate the first Standard Deviation from the mean. Graph 5.10 presents the percentage change in each router's (percentage) contribution to total zone traffic as the delay in bRCHc1 is increased.

Graph 5.11 presents each router's mean percentage contribution over 100 intervals to the total traffic in nMTHc1's zone as the delay in nMTHc1 is increased. The error bars in the graph indicate the first Standard Deviation from the mean. Graph 5.12 presents the percentage change in each router's (percentage) contribution to total zone traffic as the delay in nMTHc1 is increased.

Refer to graphs 5.7, 5.9, and 5.11 showing the percent contribution of each node to the total zone traffic as the delay inserted in the center router is increased (for zones around nRTPc2, bRCHc2, and nMTHc1 respectively). In each case, the center router experienced a decrease in contribution when a delay was inserted in it - the larger the delay, the greater the decrease. The effect of a delay seemed to manifest itself in router nRTPc2 slightly more than in bRCHc2 and nMTHc1. The border nodes in the zones increased their contributions for the most part.

It is interesting to note that some nodes contributed a very small percentage to a zone while others contributed a very large percentage to a zone (in the 3 zones studied, it ranged from ~1% to ~74%). This will be studied further in section 5.4.

Refer to graphs 5.8, 5.10, 5.12 showing the percentage change in percentage contribution for each node as the delay inserted in the center node is increased. Again, it can be seen that the center nodes (nRTPe2, bRCHc2, nMTHc1) all experienced a decrease in zone contribution as their delay was increased. They do not appear to have decreased at the same rate. It is interesting to note that for the border nodes in the zones, the only ones that exhibited a decrease in contribution as the center node's delay was increased, are nodes that have low contributions (<10%) to the zones. Other nodes that have higher contribution percentages exhibited increases in zone contributions.

Key Observations and Direction Setting:

Based on the above work four key observations can be made:

1. The percentage contribution of each router to a zone varies to some degree over time the extent to which it fluctuates is unique to each router.
2. When a delay is inserted in a router, its contribution to its own zone will decrease. The larger the delay, the larger the decrease. The rate at which the contribution decreases is unique to each router.
3. It is not necessary to monitor the percent contribution of all nodes in each zone. Only the percent contribution of the center nodes in each zone will be explicitly monitored. The border nodes are, in effect, being monitored in the zone in which they are the centre.
4. The percent contribution of routers to a zone ranges from very low to very high. The size of contribution a router makes to a zone may be a factor that affects the utility of the algorithm being developed in this thesis.

The observations above directed the next phase of research in the following way:

1. A large number of zones should be monitored for a longer period of time - 24 hours.

2. The simulator has a feature whereby traffic is varied with the time of day - this should be used in further simulations.
3. the percentage contribution of the center nodes in each zone should be monitored (the border nodes are, in effect, monitored in the zone in which they are the centre).
4. Collect a set of base statistics for each zone over a 24 hour period, with no delays inserted in any router. Insert delays in various routers (simulate for a 24 hour period) and by comparing the collected statistics to the base statistics, establish thresholds where the exceeding of the given threshold indicates a true fault exists in a network component.

GRAPH 5.2 Histogram of % Deviation From Mean Traffic Contribution - nRTPc2

GRAPH 5.3 Histogram of % Deviation From Mean Traffic Contribution - nRTPc1

GRAPH 5.4 Histogram of % Deviation From Mean Traffic Contribution, nPARc2

GRAPH 5.5 Histogram of % Deviation From Mean Traffic Contribution, bRTPc2

GRAPH 5.6 Histogram of % Deviation From Mean Traffic Contribution, nATLc1

Graph 5.8 Percent Change in % Contribution of Each Node To Zone as nRTPc2 Delay Varies

Graph 5.9 Percent Contribution of Each Node To Zone as bRCHc1 Delay Varies

Graph 5.10 Percent Change in % Contribution of Each Node To Zone as bRCHc1 Delay Varie:

Graph 5.12 Percent Change In % Contribution of Each Node To Zone as nMTHc1 Delay Vari

Delay Inserted in nMTHc1 - seconds

5.4 Alarm Thresholds and Algorithm Performance

FIGURE 5.8 Research Phases - Currently in Phase 4

This last phase of research, illustrated in Figure 5.4 above, focused on establishing thresholds of performance for routers - exceeding the thresholds would be an indication of abnormal behaviour, and hence the presence of a fault. This was done in two steps: First a

set of base measurements were established. Second, 24 hour simulations were done with no delay in the network, and with a one second delay inserted in various routers in turn.

A large number of routers (41) were selected for zone monitoring. These routers were randomly chosen, and represent routers that contribute low, medium and high percentages of traffic to their respective zones.

5.4.1 Step 1: Establish Base Statistics

To establish performance thresholds, a set of reference statistics (mean percentage contribution of traffic to zone, Standard Deviation, etc.) were calculated over a long running period (24 hours). The simulator was set up so that traffic varied with the time of day. Simulations were repeated, for situations where no delay was inserted in any node, and situations where delays were inserted in different routers in turn. By comparing the traffic contributions of the 41 routers under the different conditions, a feel for the best way to define thresholds was established. The remainder of this section outlines the steps taken in this experiment, and the results achieved.

A knowledge source was set up to monitor the zones around 41 routers. Every two minutes the knowledge source would gather traffic statistics for all of the nodes in the 41 zones. Five of these two minute periods composed an interval. After every interval the mean bytes into and out of a node was calculated, and from these measurements the node's traffic contribution to the total zone traffic was found. After 24 hours (144 intervals) of data collection, the mean traffic contributions and the Standard Deviation from this mean were calculated for each node. Table 5.10 shows the mean traffic contribution (and Standard Deviation of same) for each node to its zone, and the mean bytes into (and Standard Deviation of same) each router in a two minute period. Note that the routers are grouped according to their mean percentage contribution to their zone:

Router Number	Router Name	Percentage Contribution to Zone	Standard Deviation of % Contr.	Mean Bytes Into Router	Standard Deviation of Mean Bytes
0	nRTPc2	22.953	0.637	604485.25	47227.97
1	bRCHc1	27.109	0.671	499585.06	36584.56
2	nPARc2	29.839	0.857	624790.31	36584.56
3	nANNe2	29.813	0.398	1060320.25	70629.16
4	nNSHe2	26.649	1.089	483950.844	46180.04
5	bRTPc2	23.465	0.494	554419.81	35959.82
6	nMTHc1	27.670	1.052	751118.69	134701.28
7	bGRGe1	20.265	0.233	126917.57	2673.60
8	nLAKc2	21.798	0.427	609108.06	33878.75
9	nATLc2	22.295	0.588	541237.81	21275.54
10	nLAUc2	25.496	0.336	438168.56	31593.83
11	nPARc1	27.83	1.032	427788.44	34742.70
12	nMLVc1	29.358	0.737	563985.63	16679.58
13	nBVWc1	26.177	2.019	263582.25	37137.86
14	nBPDc1	35.565	2.788	323267.66	18113.93
15	bMERc1	24.77	0.784	1439078.88	146923.13
16	bCARc2	22.554	0.332	1409046.88	89581.68
17	bCARc4	22.757	1.030	1616546.13	179097.11
18	nLAUc1	27.913	3.211	655468.38	54665.40
19	bCRKc1	32.088	0.188	719766.94	43711.41
20	bKANc2	20.009	0.486	464247.75	29117.37
21	nRTPc1	16.150	1.380	412953.47	33813.953
22	nATLc1	19.809	0.793	433564.72	26296.03
23	bRTPc1	18.403	0.457	1179637.13	77157.84
24	nMTVe3	15.840	0.612	426680.38	33679.34
25	nBDCc2	19.385	1.912	439571.56	69244.11
26	bMERc2	19.809	2.762	1187763.5	387220.66
27	nMGRc1	14.722	1.611	184300.14	32363.35
28	nTORc1	13.179	2.225	362543.25	90778.84
29	bMERc3	13.81	0.445	596180.44	92616.17
30	bCARc5	12.529	0.258	665025.81	41965.02
31	nNCXc1	6.865	0.105	36841.88	1821.09
32	nLDNc1	9.059	1.094	73256.070	5894.031
33	nMISc1	5.223	1.383	59175.23	26601.57
34	bMERc5	9.138	0.365	438389.03	53701.03
34	bRCHc2	0.371	0.029	1859.44	174.94

Router Number	Router Name	Percentage Contribution to Zone	Standard Deviation of % Contr.	Mean Bytes Into Router	Standard Deviation of Mean Bytes
36	hRTPcA	0.329	0.218	5952.5	4333.80
37	hRTPcB	0.331	0.213	5998.28	4235.41
38	hRTPcZ	0.334	0.192	6028.11	3846.38
39	hCARcB	0.179	0.103	7267.17	4714.71
40	hCARcY	0.187	0.101	7551.72	4684.81

TABLE 5.10 Router Base Statistics Gathered Over 24 Hours (144 - 10 minute intervals)

From Table 5.10 above, it can be seen that a wide range of routers were studied:

- Routers 0 - 20 contributed more than 20% of the traffic in their zone.
- Routers 21 - 26 contributed between 15% and 20% of the traffic in their zone.
- Routers 27 - 30 contributed between 10% and 15% of the traffic in their zone.
- Routers 31 - 34 contributed between 5% and 10% of the traffic in their zone.
- Routers 35 - 40 contributed less than 5% of the traffic in their zone.

The base statistics calculated over a 24 hour period and displayed in Table 5.10 above were assumed to represent normal conditions in the network, and they formed the benchmark for comparison in further simulations. The router fault detection algorithm would make use of these while monitoring routers in the WAN. If a router's percentage contribution were to drop from the above levels by a threshold amount, an alarm would be issued, indicating the likelihood of a fault associated with that router.

5.4.2 Step 2: No Faults in Network, Determine Algorithm's False Positive Rate

The algorithm being evaluated in this thesis is based on the premise that if a router has a fault (delay) associated with it, then its performance will be in a well compared to its neighbours, i.e. its percentage contribution to its own zone will decrease. However, there are normal fluctuations in the network such that the router's contribution to the zone will

2 OF/DE 2

PM-1 3 1/2" x 4" PHOTOGRAPHIC MICROCOPY TARGET
NBS 1010a ANSI/ISO #2 EQUIVALENT

rise above and fall below the calculated benchmark mean. The next step in establishing thresholds was to study fluctuations from the mean base measurements over a 24 hour period while no delay was present in the network. Thresholds must take into consideration the normal fluctuation of traffic levels in the network, and they should not result in the generation of false alarms (false positives) when there is no fault present in the network.

A knowledge source was constructed to monitor the 41 routers as follows: Every two minutes the bytes into each router were collected, and the contribution of each router to its own zone was calculated. At the end of every 20 minutes interval, the mean contribution over the last ten 2-minute periods was calculated. The percentage change of this mean contribution from the base stats was calculated. This was done for a specified number of 20 minute intervals. Thus, a snapshot of router performance compared to the mean was available every 20 minutes.

A simulation was done as follows: The 41 routers were monitored for 72 - 20 minute intervals (24 hours). The percentage change from the mean base contribution was calculated at the end of each of the 20 minute intervals - resulting in a "diagnosis" every 20 minutes. The results of this can be seen in Graphs 5.13 - 5.17 which are histograms showing the number of times the routers have recorded a specific percentage drop in contribution over the 24 hour period. (Anytime the router has recorded an increase in contribution it is grouped under a "less than 1% decrease"). The 41 routers were grouped according to their "load" into five categories (hence the five graphs). The value of their "load" was taken as their percentage contribution to their own zone in Table 5.10 Graph 5.13 shows the results for routers contributing 0 - 5% of their zone's traffic (the low load category), Graph 5.14 shows the results for routers contributing 5% to 10% of their zone's traffic, etc.

GRAPH 5.13 Histogram - % Decrease in Zone Contribution, Routers 0 - 5 % Contribution over a 24 hour period

GRAPH 5.14 Histogram - % Decrease in Zone Contribution, Routers 5 - 10 % Contribution over a 24 hour period

GRAPH 5.15 Histogram - % Decrease in Zone Contribution, Routers 10 - 15 % Contribution over a 24 hour period

GRAPH 5.16 Histogram - % Decrease in Zone Contribution, Routers 15 - 20 % Contribution over a 24 hour period

GRAPH 5.17 Histogram - % Decrease in Zone Contribution, Routers > 20 % Contribution over a 24 hour period

Referring at graphs 5.13 - 5.17 , it can be seen that the traffic of routers contributing small percentages of traffic to their zone is more variable than the traffic of routers contributing larger percentages of traffic to their zone. Routers with higher zone contributions do not vary from the base measurements as much as routers with lower contributions. If one were to choose threshold levels to apply to all categories of routers, more false alarms would be found in routers with lower traffic contributions. This was done and is seen in Graph 5.18.. This graph plots the percentage of false positives (of the 72 intervals) as the threshold is varied from a 10% drop in contribution to a 16% drop in contribution

GRAPH 5.18 Percentage False Positives For Each Router Category

Looking at Graph 5.18 above, it would not seem a good idea to establish thresholds that would serve for all categories of routers. For example, a 10% cutoff would create a false positive diagnosis rate of 2.3% for those routers contributing > 20% to their zones and a false positive diagnosis rate of 50.5% for those routers contributing < 5% to their zones. A 16% cutoff value creates a false positive diagnosis rate of 0% for those routers contributing > 20% to their zones and a false positive diagnosis rate of 47% for those routers contributing < 5% to their zone.

For the time being, research focused on establishing thresholds for those routers which contribute > 20% to their zone. It was felt that once this was done, it would be possible to

establish thresholds for the other categories of routers, and make a decision on the applicability of this algorithm to them.

5.4.3 Faults in Network, Determine Algorithm's Correct Diagnosis and False Positive Rates

Section 5.4.2 studied the establishment of algorithm thresholds that would limit the number of false positives due to natural fluctuation of traffic levels. The purpose of this section was to study the algorithm's performance (correct diagnoses and false positives for a given threshold) when delays are inserted in various routers throughout the network. The 21 routers with more than 20% contribution to their zones were selected for monitoring. The base statistics gathered in section 5.4.1 (see Table 5.10) are still valid and were used as the benchmark for comparison in this experiment. The measurements taken in section 5.4.2 (see Graphs 5.13 - 5.18) influenced the thresholds chosen in this experiment and will be referred to again in this section.

Routers nRTPc2, nPARc2 and nMTHc1 were selected to have delays inserted in them. In turn, one second delays were inserted in each of them, and their performance statistics were gathered for 72 - twenty minute periods (24 hrs) using the same knowledge source and technique used to gather the statistics in section 5.4.2.

Graph 5.19 presents a histogram showing the instances of percent drop in zone contribution for the routers contributing more than 20% to their zone when a one second delay was inserted in nRTPc2. Graph 5.20 is a subset of Graph 5.18 and presents a histogram showing the instances of percent drop in zone contribution for router nRTPc2. This can be compared to Graph 5.21 which is a histogram from nRTPc2 under normal conditions. A discussion of these graphs follows.

GRAPH 5.19 Histogram - % Decrease in Zone Contr., Router > 20% Contr., nRTPc2 with a one second delay

GRAPH 5.20 Histogram - % Decrease in Zone Contribution, nRTPc2 with 1 second delay

GRAPH 5.21 Histogram - % Decrease in Zone Contribution, nRTPc2 with zero delay

Comparing Graph5.17 (nodes contributing > 20%, no delay) to Graph 5.20 (nodes contributing > 20%, delay in nRTPc2), it can be seen that the insertion of a delay has caused a minor shift in node-zone contributions ie: more nodes experienced decreases in zone contributions in excess of 10%, some as high as 25%. At the same time, comparing Graph 5.21 (nRTPc2 with no delay) to Graph5.20 (nRTPc2 with delay), it can be seen that the insertion of a delay in nRTPc2 caused a major shift (decrease) in nRTPc2's node contribution. If one were to select a threshold of 10%, we would achieve a 40.3% detection of the fault in nRTPc2, and a 4% false positive rate in the remaining routers.

These simulations were repeated for routers nPARc2 and nMTHc1. Histograms (like those in Graphs 5.19 - 5.21) for the simulations are not included in the thesis for the sake

of brevity; however, the essential data for the simulations is summarized in Table 5.11. Table 5.11 presents the critical information for the simulations where:

- there was no delay in any router
- router nRTPc2 had a 2 second delay inserted in it
- router nPARc2 had a 2 second delay inserted in it
- router nMTHc1 had a 2 second delay inserted in it

Table 5.11 includes for a given threshold value: the percentage of correct diagnosis of a problem, and the percentage of false positives generated. It also includes the names of the specific routers generating the false positives.

Name of Router With Delay (if any)	Percentage of Correct Diagnosis		Percentage of False Positives		
	Threshold = 10 Percent	Threshold = 15 Percent	Router Giving False Positive	Threshold = 10 Percent	Threshold = 15 Percent
no delay	N/A	N/A	nNSHc2	1.4%	0%
			nPARc1	1.4%	0%
			nBVWc1	8.3%	0%
			nBPDc1	2.8%	0%
			nLAUc1	33%	2.8%
			ALL ROUTERS	2.25%	0.01%
nRTPc2	40.3%	5.6%	nNSHc2	5.6%	0%
			nBVWc1	6.9%	0%
			bCARc4	28%	26%
			bCRKc1	1.4%	0%
			ALL ROUTERS	2.1%	1.3%
nPARc2	36.1%	2.8%	nRTPc2	12.5%	0%
			nNSHc2	11.1%	0%
			nLAUc2	1.4%	0%
			nLAUc1	37%	36%
			nBVWc1	9.7%	0%
			bCARc4	33%	30.6%
			nPARc1	1%	0%
			ALL ROUTERS	10%	4.9%
nMTHc1	22%	0%	nBVWc1	43%	26%
			bCARc4	60%	31%
			nLAUc1	69%	1.4%
			nNSHc2	2.8%	0%
			ALL ROUTERS	8.8%	3.2%

TABLE 5.11 Correct Diagnosis and False Positives As Delays are Inserted in Routers

Referring to Table 5.11 above, a threshold of a 10% drop in router - zone contributions yields what could be considered a satisfactory fault detection percentage for the three routers (between 22% and 40.3% in column 2). Unfortunately the percentage of false positives for all other routers (between 2.25% and 10% in column 5) would be considered unsatisfactory. Raising the threshold to 15% decreases the false positives for all routers (in column 6), but very much at the expense of the correct diagnosis percentage (in column 3).

Examining the routers which generate false positives, it turns out that they are all either neighbours of the routers with delays, or are “noisy routers” where the Standard Deviation of the percentage contribution to the zone is quite high (i.e.: a large percentage of the mean contribution to the zone). Table 5.12 below contains the routers whose contributions to their zone are $> 20\%$. The darkened rows indicate those routers which could be considered to have high Standard Deviations (ie: greater than 1.0 arbitrarily).

Router Number	Router Name	Percentage Contribution to Zone	Standard Deviation of % Contr.
0	nRTPc2	22.953	0.637
1	bRCHc1	27.109	0.671
2	nPARc2	29.839	0.857
3	nANNc2	29.813	0.398
4	nNSHc2	26.649	1.089
5	bRTPc2	23.465	0.494
6	nMTHc1	27.670	1.052
7	bGRGc1	20.265	0.233
8	nLAKc2	21.798	0.427
9	nATLc2	22.295	0.588
10	nLAUc2	25.496	0.336
11	nPARc1	27.83	1.032
12	nMLVc1	29.358	0.737
13	nBVWc1	26.177	2.019
14	nBPDc1	35.565	2.788
15	bMERc1	24.77	0.784
16	bCARc2	22.554	0.332
17	bCARc4	22.757	1.030
18	nLAUc1	27.913	3.211
19	bCRKc1	32.088	0.188
20	bKANc2	20.009	0.486

TABLE 5.12 Routers Contributing $> 20\%$ To Zone, Highlighting Those With High Standard Deviations

Looking first at the case where no delay is in the network, the routers generating false positives (nNSHc2, nPARc1, nBVWc1, nBPDc1, and nLAUc1) all have high Standard Deviations.

Looking at the case where a delay was inserted in router nPARc2, the routers which are neighbours of nPARc2 and generated false positives are: nRTPc2, nPARc1, and nLAUc2. The remaining routers which generated false positives all have high Standard Deviations.

For the cases where delays were inserted in nRTPc2 and nMTHc1, all of the routers which generate false positives have high Standard Deviations.

The situation where routers with high Standard Deviations are generating false positives could be remedied by tying in the value of Standard Deviation to the setting of the threshold. It may not be possible to remedy the situation where routers neighbouring a faulty node generate false positives. However, if the routers are close to the source of the problem, then their false alarms may not be considered serious.

MAJOR OBSERVATIONS AND FUTURE RESEARCH:

Four major observations are summarized from the work done in section 5.4, establishing alarm thresholds and monitoring the algorithm's success at detecting router faults.

1. The contributions of routers to their zones varies significantly between routers. The routers which contribute a higher percentage contribution to their zone demonstrate more stable behaviour in that their Standard Deviations are a smaller percentage of their mean contribution.
2. One - dimensional thresholds, applied to all categories of routers (those that have large contributions to their zones vs. those that have small contributions to their zones) are difficult to establish.
3. Routers can be classified in many different ways: the percentage contribution to their zone, the size of the Standard Deviation of that measure, or perhaps even by such things as the topology around the router.

4. Focusing on the category of routers which contribute more than 20% to their respective zones:

- The detection rate for misbehaving routers is quite good (between 22% and 40.3% of all diagnoses). The percentage of false positives (between 2.25% and 10%) needs to be decreased.
- False positives are generated by routers with high standard deviations of their percentage contribution and by routers which are neighbours to a router with a fault. Higher thresholds for routers with high Standard Deviations should eliminate them as a source of false positives. The false positives generated by routers which are neighbours to a router with a fault are not as serious since they direct the attention of the network operators to the area of the fault.

Based on the above observations, two areas for future research (which are out of the scope of this thesis) are apparent. Some preliminary investigation into these research areas has been done and can be seen in section 5.5. The two areas for future research are:

1. Look more closely at the establishment of thresholds. Rather than setting “across the board” thresholds of X% drop in percentage contribution to the zone, it may be worth considering the size of a router’s standard deviation when establishing a threshold. For the category of routers with more than 20% contribution to their zone, this strategy may reduce the false positive rate. It also may allow extension of the algorithm to cover all categories of routers, even those which were considered too “noisy” to include before.
2. Consider the way in which the base statistics (which form the benchmark for comparison when monitoring for errors) are determined. In the research done in this thesis, the base statistics were calculated as the mean contribution in a two minute interval over 24 hours. There are many ways of calculating the base statistics including:
 - Have multiple “models” of base statistics, each geared to a specific time frame in the day. This would theoretically reduce any error due to time of day fluctuations in traffic patterns.

- When monitoring a network for a given time period, use the statistics collected in the previous time period as the basis for comparison in the current time period - in effect, a constantly changing set of base statistics that “follows” the time of day fluctuations.

There are many possible ways to establish the set of measurements which define the normal behaviour of components in the network. It would be worth investigating some of the different ways, and determine the effect they have on the performance of the diagnostic algorithm.

5.5 Preliminary Work in Future Research Areas

FIGURE 5.9 Research Phases - Currently in Future Research Phase

The purpose of this section is to present some preliminary investigative work into areas that are out of the scope of the thesis. These future research areas are intended to improve the performance of the algorithm developed in this thesis. Figure 5.9 above illus-

trates how this work ties in with the four main phases of research, and it illustrates the three main steps in this section.

In the research presented in section 5.4, an unacceptable number of false positives were generated by the algorithm (actually, any false positives are unacceptable). It is worthwhile to identify ways of improving the accuracy of the algorithm. Although in-depth study of this is beyond the scope of the thesis, some preliminary investigation is warranted.

Two readily identified areas of improvement are firstly the manner in which alarm thresholds are set and secondly how the set of base reference statistics are established. This phase of research is divided into three sections to investigate the above mentioned areas of improvement. The first section investigates a second method of defining alarm thresholds. The second and third sections investigate the manner in which reference statistics defining “normal behaviour” are established.

5.5.1 A Second Method of Defining Alarm Thresholds

In the research presented in Section 5.4, the measure of health for a router was the percent contribution of a router’s processed traffic to its zone’s total. A base set of measurements characteristic of “no fault behaviour” were established. Every 20 minutes the knowledge source would calculate each router’s contribution to their zone. If the current contribution dropped from the base measure by a threshold percentage (e.g. 10%) then this was presumed to indicate a fault with the router. This was found to be suitable for the category of routers which contributed more than 20% to their respective zones. It was not suitable for routers which were “noisy”, and had highly variable traffic fluctuations as indicated by a high Standard Deviation from mean traffic contributions. It was also unsuitable for routers which were small contributors to their zones (0-5%, 5-10%, 10-15%, 15-20% categories).

The purpose of this section is to establish a threshold that will decrease false positives generated by the algorithm and make it applicable to all (or as many as possible) categories of routers. The threshold to be investigated could be considered “two-dimensional”. A router’s drop in contribution from the “normal” base level is measured in two ways: firstly, as a percentage of the base percentage contribution (% drop, as was used in section 5.4) and secondly, in terms of the number of standard deviations it has dropped from the mean.

Experimentation was done as follows:

The base statistics defining normal behaviour are the same ones used in section 5.4, and can be seen in Table 5.10 in Section 5.3. The knowledge source from Section 5.4 was used in a modified form to monitor the 41 routers and their zones for 72 - 20 minute intervals (24 hours). At the end of each 20 minute interval the knowledge source would calculate each router’s mean contribution to their zone for that interval. The drop from each router’s base percent contribution (if any) was measured in two ways:

- “percentage” drop from mean base contribution
- “number of standard deviations” drop from mean base contribution

These two measures provided the data for a series of graphs (X axis = “number of standard deviations drop”, Y-axis = “percentage drop”). From the graphs, thresholds were established. Discussion of the simulations performed and resulting graphs follows.

Two sets of simulations were done.

In the first set, the network was monitored for 24 hours with no fault inserted in the network. The data gathered for this simulation can be seen in Graph 5.22. Referring to this graph, it can be seen that while no delay was in the network, the drop in router contributions ranged from very small to, in a few isolated cases, as much as 60%. It must be noted that the routers with large “percentage” drops in contribution were routers in the category that contribute between 0 and 5% to their zone. This delineation is made on the graph.

Although the “percentage” drop in contribution was at times high, the “number of standard deviations” drop was less than 2.5 in all cases.

GRAPH 5.22 Drop in Zone Contribution over 24 hours, Measured in “Percent” Drop and “Number of Std. Deviations Drop” , no fault in network

The second set of simulations involved monitoring the network for 5 - 24 hour periods. In each of the 24 hours periods, a delay was inserted in a specified router, and data gathered in the same manner as in the first set of simulations. The five routers which, in turn, had delays inserted in them are bMERc3, nRTPc2, bCARcb, bRCHc1, and bCRKc1. The data gathered for the non-faulty nodes in each simulation was combined and plotted on Graph 5.23. The data gathered for each of the 5 faulty nodes was combined and plotted on Graph 5.24

GRAPH 5.23 Drop in Zone Contributions for Healthy Routers, Measured in “Percent” Drop and “Number of Std. Deviations Drop”

GRAPH 5.24 Drop in Zone Contributions for Faulty Routers, Measured in “Percent” Drop and “Number of Std. Deviations Drop”

Referring to Graph 5.23 showing the change in contributions for non-faulty nodes, it can be seen that many routers throughout the network are affected by a fault inserted in a router. The contribution of several routers has fallen, as indicated by the spread of data on the graph - an increase along the Y-axis (percent change from the normal behaviour mean) and an increase along the X-axis (number of std. deviation change from the normal behaviour mean). Router bCRKc1 stands out in particular, with a drop in zone contribution which is many standard deviations from the normal behaviour mean, though not a large percentage drop.

Since the routers in Graph 5.23 are not faulty, thresholds must be established that will not identify them as faulty (false positives). Two lines have been drawn on the graph that together approximate a threshold that will not generate false positives.

The threshold established in Graph 5.23 has been redrawn on Graph 5.24 which shows the change in contributions for faulty nodes. Referring to Graph 5.24, any point on the graph that lies beyond the two threshold lines represents a correct fault diagnosis. It can be seen that the fault detection rate would range from approximately 0% to 50% for the different routers.

Important Observations and Future Directions:

The following two observations can be made from this work:

1. When comparing sets of data to a reference set, there are different thresholds that can be established to perform the comparison. The threshold used in this section has led to some improvement in diagnosis compared to the threshold used in section 5.4. Specifically, the number of false positives has been decreased and the algorithm has been made applicable to all categories of routers. More research would be required to determine the diagnosis rate for the different categories of routers.

2. When performing a diagnosis, it is necessary to have a model of correct behaviour for the basis of comparison. In the research completed in this section, fault thresholds were established based on Graph 5.23 which represented the behaviour of non-faulty routers when faults existed in the network. If thresholds had been established based on Graph 5.22, which represented the behaviour of non-faulty routers when no fault existed in the network, then many false positives would have been generated on Graph 5.24. This brings to light an important consideration: When establishing a model of correct behaviour, how does one define what correct is? Is the correct behaviour of a router defined by its behaviour when there are no faults in the network or does it include its behaviour when there are faults present elsewhere in the network? The research done in this section indicates that this is an important consideration.

Based on the above observations, the following two areas of future research are apparent:

1. Consider the collection of a set of base statistics that are meant to represent normal behaviour of routers in the network. Investigate the effect of collecting the base statistics under conditions: firstly, where no faults are present in the network, and secondly, where faults are present elsewhere in the network (ie: the normal behaviour of a router will include its behaviour when routers around it have faults inserted in them).
2. Further to the first area of future research, investigate the impact of collecting base statistics over different time frames. Is a set of statistics calculated over a 24 hour mean suitable, or is better performance achieved by have multiple models of correct behaviour representing network conditions during different times of the day?

5.5.2 Investigate Determination of Base Statistics - The Model of Correct Behaviour

The purpose of this step in research was to investigate different ways in which base statistics could be established: Should the model of correct behaviour for a router be based on its performance when there are no faults in the network, or should it include the router's performance when there are faults in the routers around it? Does it matter?

Two sets of simulations were done in this experiment.

Set 1:

In the first set, base statistics were gathered (for 5 - 20 minute intervals) for 41 routers being monitored in the WAN while there were no faults anywhere in the network. The base statistics included each router's mean percentage contribution to its zone and the standard deviation of that measure. The base statistics for this set of simulations can be seen in the unshaded columns of Table 5.13.

A series of five simulations were done (for 5 - 20 minute intervals each) for the following cases:

- simulation 1 - no delay in the network
- simulation 2 - one second delay in node nRTPc2
- simulation 3 - one second delay in node nMTVc3
- simulation 4 - one second delay in node bCRKc1
- simulation 5 - one second delay in node nLAUc1

At the end of each 20 minute interval, the change in each router's contribution in terms of percent change from the base mean and the number of standard deviations from the base mean, was calculated. The data for the non-faulty routers in all five simulations was combined and can be seen in Graph 5.25. The data for the four faulty routers in all the simulations was combined and can be seen in Graph 5.26. A threshold was established based on the information in Graph 5.25, and is shown on both graphs.

Router Number	Router Name	Step 1: Base Stats With No Faults in the Network		Step 2: Base Stats With Faults Elsewhere in the Network	
		Percentage Contribution to Zone	Standard Deviation of % Contr.	Percentage Contribution to Zone	Standard Deviation of % Contr.
0	nRTPc2	23.63	0.27	23.391	0.433
1	bRCHc1	20.462	0.396	21.009	0.626
2	nPARc2	29.602	0.822	29.340	1.297
3	nANNc2	29.106	0.085	29.229	0.219
4	nNSHc2	20.491	0.819	21.283	1.140
5	bRTPc2	26.228	0.407	26.719	0.397
6	nMTHc1	28.192	0.315	27.940	0.338
7	bGRGc1	20.885	0.119	20.902	0.103
8	nLAKc2	21.570	0.276	21.490	0.310
9	nATLc2	22.252	0.653	22.069	0.595
10	nLAUc2	25.273	0.300	25.553	0.637
11	nPARc1	27.774	1.113	28.077	1.369
12	nMLVc1	30.234	0.226	30.479	0.239
13	nBVWc1	23.743	0.288	23.921	0.549
14	nBPDc1	32.243	0.481	34.560	0.657
15	bMERc1	23.544	0.111	23.462	0.230
16	bCARc2	22.272	0.092	22.458	0.217
17	bCARc4	24.152	0.159	24.184	0.244
18	nLAUc1	24.061	0.159	24.049	0.196
19	bCRKc1	32.026	0.163	32.104	0.143
20	bKANc2	19.457	0.122	19.677	0.206
21	nRTPc1	16.193	1.050	14.984	1.298
22	nATLc1	19.339	0.595	19.715	0.924
23	bRTPc1	15.986	0.309	15.820	0.309
24	nMTVc3	15.226	0.461	14.892	0.455
25	nBDCc2	16.570	0.317	16.682	0.487
26	bMERc2	23.276	0.155	23.333	0.410
27	nMGRc1	13.244	0.162	13.372	0.399
28	nTORc1	12.285	0.279	12.122	0.243
29	bMERc3	13.998	0.097	14.051	0.125
30	bCARc5	12.414	0.249	12.161	0.527
31	nNCXc1	2.853	0.068	2.858	0.100
32	nLDNc1	7.948	0.130	8.226	0.226
33	nMISc1	6.888	0.145	7.057	0.244

Router Number	Router Name	Step 1: Base Stats With No Faults in the Network		Step 2: Base Stats With Faults Elsewhere in the Network	
		Percentage Contribution to Zone	Standard Deviation of % Contr.	Percentage Contribution to Zone	Standard Deviation of % Contr.
34	bMERc5	8.669	0.234	8.798	0.161
34	bRCHc2	0.683	0.052	0.668	0.041
36	bRTPcA	0.201	0.054	0.210	0.054
37	bRTPcB	0.213	0.045	0.218	0.041
38	bRTPcZ	0.242	0.038	0.230	0.051
39	bCARcB	0.165	0.042	0.115	0.045
40	bCARcY	0.176	0.054	0.121	0.050

TABLE 5.13 Router Base Statistics - Collected Under Two Different Conditions

Set 2:

In the second set of simulations, base statistics were gathered for 5 - 20 minute intervals while different faults existed in the network. In the first 20 minute interval, there were no faults in the network. In the second 20 minute interval, there was a one second delay in router nRTPc2 (since nRTPc2 was faulty, the calculations for nRTPc2's base statistics did not include this particular interval). In the third 20 minute interval, there was a one second delay in router nMTVc3 (since nMTVc3 was faulty, the calculations for nMTVc3's base statistics did not include this particular interval). The fourth and fifth intervals had a delay in bCRKc1 and nLAUc1 respectively. The base statistics for this set of simulations can be seen in the shaded columns of Table 5.13.

A series of five simulations were done (for 5 - 20 minute intervals each) for the following cases:

- simulation 1 - no delay in the network
- simulation 2 - one second delay in node nRTPc2
- simulation 3 - one second delay in node nMTVc3
- simulation 4 - one second delay in node bCRKc1
- simulation 5 - one second delay in node nLAUc1

At the end of each 20 minute interval, the change in each router's contribution in terms of percent change from the base mean and the number of standard deviations from the base mean, was calculated. The data for the non-faulty routers in all five simulations was combined and can be seen in Graph 5.27. The data for the four faulty routers in all the simulations was combined and can be seen in Graph 5.28. A threshold was established based on the information in Graph 5.27, and drawn on both graphs.

Referring to Table 5.13, it can be seen that when base statistics include the performance of routers when faults exist in the network, there is a slight shift in the routers' mean percentage contributions to their zones. The Standard Deviation of the measure also, for the most part increases, indicating that traffic is more variable when there are faults in the network.

Referring to Graphs 5.25 and 5.27 showing the drop in contribution for non faulty routers, it can be seen that when base statistics include faulty behaviour in the network (Graph 5.27), then it is easier to establish lower thresholds.

Referring to Graphs 5.26 and 5.28 showing the drop in contribution for faulty routers, it can be seen that when base statistics include faulty behaviour in the network, the algorithm has a higher fault detection rate. As seen in Graph 5.26, where base statistics do not include faulty behaviour in the network, four of the routers have 100% fault detection rates while nMTVc3 has a 0% fault detection rate (the points for nMTVc3 are on the left and below the threshold, while the points for the other three routers are to the right or above the threshold. In Graph 5.28, where base statistics include faulty behaviour in the network, four of the routers have 100% fault detection rates while nMTVc3 has an 80% fault detection rate.

GRAPH 5.25 Drop in Contribution For Non-Faulty Nodes (Base Stats Do Not Include Faulty Behaviour In The Network)

GRAPH 5.26 Drop in Contribution For Faulty Nodes (Base Stats Do Not Include Faulty Behaviour In The Network)

GRAPH 5.27 Drop In Contribution For Non-Faulty Nodes (Base Stats Include Faulty Behaviour In The Network)

GRAPH 5.28 Drop In Contribution For Faulty Nodes (Base Stats Include Faulty Behaviour In The Network)

Important Observations and Future Research:

The following observation can be made from this work:

1. When implementing a diagnostic algorithm, the performance of the algorithm relies on the quality of the “base measurements” that define normal behaviour for the system being studied, and are the basis of comparison for the algorithm. In the Wide Area Network being studied in this thesis, it was found that routers in the network were very interdependent, and that algorithm performance was improved when a router’s normal behaviour included not just fault-free, correct behaviour, but its behaviour when faults existed elsewhere in the network.

Based on the above observation, the following area of research is apparent:

1. Diagnostic algorithms require a model of correct behaviour to be used as a basis of comparison. In live Wide Area Networks, the term “correct behaviour” does not likely mean fault-free behaviour, rather it would be interpreted as normal behaviour that includes the presence of faults somewhere in the network. The establishment and maintenance of a model of “correct behaviour” for a network is a complex task, and it is worth researching in more detail.

5.5.3 Over What Time Frame Should Base Statistics Be Gathered ?

The previous section investigated the collection of base statistics under different network conditions: firstly a fault free network and secondly a network with faults in it. Another consideration is the timeframe in which reference statistics are gathered. The research done in Sections 5.3 and 5.4 used a set of base statistics which were determined as a mean over a 24 hour period. It is easy to appreciate that traffic conditions in a network vary with the time of day, day of the week, and from week to week. How does one establish and maintain a set of base statistics under these conditions? There are many different timeframes over which base statistics may be gathered. For example, Maxion [12] used a

model of correct behaviour that was determined over the most recent two week period, and it was set up to account for variations in traffic flow with the time of day and day of the week. The LEAPR Knowledge Source [10] used as a model of correct behaviour the set of measurements gathered during the 20 minute interval preceding the current interval.

The purpose of this section was to do a preliminary investigation into the algorithm performance impact of choosing a set of base statistics calculated as a mean over 24 hours. Specifically, this investigation studied the time of day variations in the measured signal distorting the accuracy of the algorithm. To do this, a knowledge source was constructed that would monitor specified routers in the following manner: The knowledge source was run for 72 - 20 minute intervals (24 hours). At the end of each 20 minute period, for each monitored router, the mean percentage contribution to its zone was calculated. Also, to provide a comparison, the mean bytes into each monitored router (for the 20 minute period) was calculated and. The mean percentage contribution and mean bytes into each router, and the standard deviation of the two measures were calculated for the entire 24 hour period.

The simulation was done monitoring routers nRTPc2, nANNc2, and nMTHc1. Graphs 5.29 and 5.30 show, for router nRTPc2, the percent contribution of the router to its zone, and the mean bytes into the router over the 72 - 20 minute periods. Graphs for routers nANNc2 and nMTHc1 were not included for the sake of brevity.

Table 5.14 contains for all three routers:

- the mean percentage contribution of each router to their zones over 24 hours, the standard deviation of this measure, and the standard deviation as a percentage of the mean
- the mean bytes into each router over 24 hours, the standard deviation of this measure, , and the standard deviation as a percentage of the mean.

Referring to Graphs 5.29 and 5.30, it can be seen that both of the measures (a router's contribution to its zone, and a router's traffic load as measured in bytes into the router) vary throughout the day. It is possible that using a mean calculated over 24 hours for the

base measurements representing normal behaviour is not as accurate as a series of base measurements calculated throughout the day, which may better reflect time of day variation in the signal being monitored. It would appear that the measure of a “router’s percent-age contribution to its zone” is more variable than the measure of “it’s mean bytes into the node”. This is hard to distinguish on graphs that have different scales, and is better investigated using standard deviations, as done in Table 5.14.

Router Name	Measure: % Contribution to Zone			Measure: Mean Bytes Into Router		
	Mean over 24 hours	Standard Deviation	Std. Dev. as a % of the mean	Mean over 24 hours	Standard Deviation	Std. Dev. as a % of the mean
nRTPc2	22.958	0.583	2.54%	604528.9	42579.9	7%
nANNc2	29.813	0.387	1.3%	1060232.6	69787.4	6.6%
nMTHc1	27.67	1.027	3.71%	750924.31	133922.6	18%

TABLE 5.14 Means, Standard Deviations For Routers’ Percent Contribution, and Mean Bytes In (over 24 hours)

Refer to Table 5.14 above. Refer to the measure of a router’s “% contribution to its zone” in columns 2, 3, and 4. The standard deviation as a percentage of the mean is quite small, between 1.3% and 3.71%. Referring to the measure of a router’s “Mean Bytes In”, the standard deviation as a percentage of the mean is quite large: between 6.6% and 18%. This indicates that the first measure (% contribution to the zone) varies less throughout the day than the second measure (mean bytes in). This is quite significant in that the first measure was derived from the second measure, and provides a more stable signal.

GRAPH 5.29 nRTPc2's % Contribution to Zone over 72 - 20 minute Intervals (24 hours)

GRAPH 5.30 nRTPc2's Mean Bytes Into Node over 72 - 20 minute Intervals (24 hours)

Major Observations and Future Research

The following observations are summarized from the work done in this section:

1. When monitoring components in a Wide Area Network, any given measure or signal is likely to have natural variation with the time of day, day of the week, etc. and this will complicate the establishment of a benchmark performance value representative of healthy component behaviour. The maintenance of multiple models of behaviour, neural networks, and machine learning techniques can be used to overcome this problem.
2. Some measures or signals have more natural variation than others. Given a measure or signal of a component's performance, it may be possible to use an algorithm (such as the algorithm developed in this thesis) to manipulate the data and make the signal less variable (ie: smooths out the signal). This is valuable because it simplifies the establishment of a model of correct component behaviour. Ideally, one would want to tune a signal so that it did not vary unless a fault was introduced to the network.

Given the above observations, the following areas of future research are apparent:

1. Monitor the variation in the measure "percentage contribution of a router to its zone" over both time of day, and day of the week.
2. Study, in more detail, the establishment of a model of normal behaviour using the measure "percentage contribution of a router to its zone". For example,
 - use the measure as input to a neural network
 - establish and maintain multiple models (ie: a model for 12:00 am - 7:00 am, 7:00 am - 12:00 pm, etc.) of normal behaviour
3. Given a measure or signal of network performance that has a degree of variability in it, it may be possible to decrease the variability by "massaging" the data. Investigate other algorithms to massage the signal. Study the algorithms under fault-free and faulty conditions.

CHAPTER 6.0 Conclusions, Future Research, Summary of Contributions

Research in this thesis has focused on creating an algorithm to detect the abnormal behavior of a router in a Wide Area Network. The type of abnormal behaviour studied in this thesis is a packet processing delay which is inserted in a specified router.

This chapter is divided into three sections. The first section presents the conclusions reached in this thesis. The second section recommends future research areas related to the work done in this thesis. The last section summarizes the contributions this thesis has made in the area of communications network diagnosis.

6.1 Conclusions

The research done in this thesis has resulted in the four following conclusions:

1. An algorithm to detect abnormal router behaviour in a wide area network has been defined:

Consider a router *R* and all routers connected to it to a depth of 1. These routers make up a “zone”. The traffic throughput for each router in the zone is calculated for a given period of time. The total traffic throughput for the zone is the sum of the individual throughputs for all routers in the zone. Each router in the zone contributes a percentage of the total traffic throughput in the zone. This percentage contribution is the critical measure for the algorithm, and it is used to demonstrate a faulty node’s drop in performance relative to its neighbours. If a delay is inserted in the router *R*, then it will not contribute the same percentage contribution to the zone, in fact it will contribute less, essentially ending up in a “well” compared to its neighbours. The neighbouring nodes will experience a slight rise in percentage contribution; however if they are heavily dependent on the router *R* for a large portion of its traffic, their contribution may drop slightly. The algorithm will monitor routers in the network, specifically it will monitor the router’s contribution levels to their respective zones. If a router’s contribution level drops from a base level by a specified threshold, an alarm will be generated indicating the presence of a fault in the network.

2. The measure or signal established by the algorithm “percentage contribution of a router to its zone” has been found to have less natural variation than the measure or signal from which it was derived “mean bytes out of a router”, making it a better measure or signal.
3. Factors affecting the utility of the algorithm have been determined, including:
 - (i) traffic mix in the network
 - (ii) the inherent “noisiness” of a router being monitored
 - (iii) the magnitude of the fault present in the network
 - (iv) the quality of the base measurements which define “normal” behaviour of a router, and the alarm thresholds which identify “abnormal” behaviour
4. The performance of the algorithm has been measured as follows:

- (i) With the threshold defined as a 10% decrease in a router's contribution to its zone, the detection rate for router faults is between 22% and 40.3%. The false positive rate is between 2.25% and 10%.
- (ii) False positives are generated by routers which are neighbours to faulty routers and routers which are considered noisy. The term noisy, in this context, means that given a router's mean percentage contribution to its zone, the standard deviation of that measure is a large percentage of the mean percentage contribution to the zone.
- (iii) False positives created by the noisy routers can be eliminated by raising the threshold for those particular routers.

6.2 Future Research

The following identify areas of future research related to the work completed in this thesis:

1. The algorithm developed in this thesis could be used, in cooperation with the NPATH path finding algorithm [11] in a reactive manner to identify the likely cause of file transfer problems which users complain about (refer to section 2.2.6 of this thesis). If a user complained about poor file transfer performance between two nodes, NPATH could be used to identify the routers along the file transfer path, and the algorithm developed in this thesis could be used to determine whether they have a fault associated with them. The alarm thresholds used by the algorithm in a reactive mode will likely be different (i.e. lower) than the thresholds used by the algorithm in a proactive mode (which are geared to eliminating false positives at the expense of a lower diagnosis rate).
2. Further investigation into the establishment of a set of base statistics which define "normal" behaviour for a router, and act as the basis for comparison when monitoring for faults, is warranted. For example, rather than determining base statistics from a mean over 24 hours, should there be a number of models of normal behaviour - one for each interval of the day - to account for any time-of-day variation in network load? How does one determine what correct behaviour is?

3. Further investigation into the assigning of alarm thresholds is warranted. For example, rather than defining thresholds as a percentage drop from the base measurements, should the threshold take into consideration information such as the router's RMS Std. deviation, or whether the router is a high or low contributor to its zone? The use of these "adaptive" thresholds could be investigated in more detail.
4. The algorithm could be tested on different types of statistics. For example, rather than using a router's percentage contribution of processed traffic in a zone, would a router's percentage contribution of lost packets in a zone be meaningful measure?
5. The algorithm could be tested on other types of faults - router corruption for example.
6. Investigate other algorithms that manipulate a signal to smooth out the inherent variability and establish a more stable signal.
7. There were two hypotheses presented at the end of Section 5.2.2. The first hypothesis formed the basis of the rest of the work done in the thesis. It would be desirable to investigate the performance of the second hypothesis. A possible algorithm would be: For each router in the network, establish base statistics for the lines on the router - the percentage contribution of each line's traffic to the total traffic in the router. Monitor each router to isolate the line whose performance has dropped the most in a given interval. Get the name of the router at the other end of the line and put it in a list. Examine the list - if one router appears in the list a disproportionate number of times, it may be faulty - the poor performing lines of other routers have acted as "vectors" pointing toward the faulty node.

6.3 Summary of Contributions

The following contributions have been made by this thesis:

1. An algorithm has been defined which detects abnormal performance of routers in near real time.

2. The algorithm has demonstrated that the near real time analysis of network performance data is feasible and meaningful.
3. Some insight into the factors influencing the diagnosability of networks has been developed.
4. The algorithm has shown that the inherent variability in a measure or signal of network performance can be smoothed, to a degree, by “massaging” it to establish a second less variable signal.

References

- [1] BNR Information Technology, "BNR Local Area Network Environment: An Introduction for the Network User", Ottawa, Canada, July 1990.
- [2] Bell-Northern research, "MAGNet: Network management tools", Release 3.0, Ottawa, Canada, June 1990.
- [3] B. Pagurek, N. Dawes, A. Kaye, "A Multiple Paradigm Diagnostic System For Wide Area Communication Networks", Carleton University Technical Report SCE-91-37, 1991
- [4] H. Nii, "Part One - Blackboard Systems: The Blackboard Model of Problem Solving and the Evolution of Blackboard Architectures", The AI Magazine, pp. 38-53, Summer 1986.
- [5] H. Nii, "Part Two - Blackboard Application Systems, Blackboard Systems from a Knowledge Engineering Perspective". The AI Magazine, pp. 82 - 106, August 1986.
- [6] S. Iqneibi, "A Blackboard Architecture to Support Network Fault Diagnosis", Master of Engineering Thesis, Carleton University, 1992.
- [7] N. Dawes, "Dragon: A very high speed, very large scale network simulator", Carleton University Technical Report SCE-91-31, August 1991.
- [8] N. Dawes, B. Pagurek, J. Altoft, "Effective Network Diagnosis By Reasoning In Uncertain Nested Evidence Spaces", 12'th International Conference on Artificial Intelligence, Avignon, 1992.
- [9] R. Dussault, "A Diagnostic Expert System for Wide Area Networks", Master of Engineering Thesis, Carleton University, March 1992.
- [10] F. Viens, "A Neural Network Approach to Detect Traffic Anomalies in a Communication Network", Master of Engineering Thesis, Carleton University, 1992.
- [11] P. Kalab, "Operations Research Techniques Applied to Network Fault Management", Master of Engineering Thesis, Carleton University, 1992.

- [12] R.A. Maxion, "Anomaly Detection for Diagnosis", 20th IEEE International Conference on Fault Tolerant Computing, 1990, pp 20-27.
- [13] R.M. Goodman and H. Latin, "Automated Knowledge Acquisition from Network Management Databases", Integrated Network Management II, 1991, pp. 541-549
- [14] N. Dawes, J. Altoft, B. Pagurek, "Effective Network Diagnosis by Reasoning in Uncertain Nested Evidence Space", Carleton University Technical Report SCE-91-26; submitted to IEEE conference on Artificial Intelligence Applications (July 1991).
- [15] J. Altoft, N. Dawes, "RUNES: Dempster Shafer Reasoning in Uncertain Nested Evidence Spaces", Carleton University Technical Report SCE-91-09; submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence (1991).
- [16] Q. Bo , Masters student, Systems and Computer Engineering, Carleton University - In Conversation
- [17] L. Zary, Masters student, Systems and Computer Engineering, Carleton University - In Conversation
- [18] A. Hiramatsu, "ATM Communications Network Control by Neural Networks", IEEE Transactions on Neural Networks, Vol. 1 No. 1, March 1990, pp. 122-130.
- [19] M. Stefik, "Planning and Meta-Planning (MOLGEN: Part 2)", Artificial Intelligence, Vol. 16, 1981, pp. 141 - 170.
- [20] P.S. Rosenbloom and A. Newell, "The Chunking of Goal Hierarchies: A Generalized Model of Practice", R.S. Michalski, T.G. Carbonell and T.M. Mitchell (Editors) Machine Learning: An Artificial Intelligence Approach, Morgan-Kaufman, Los Altos, CA, 1986, pp. 247 - 287.
- [21] B. Hayes-Roth, F. Hayes-Roth, "A Cognitive Model of Planning", Cognitive Science, Volume 3, 1979, pp. 275 - 310.

The following book and paper were also referenced in preparing this report but were not explicitly summarized.

- [22] P. H. Winston, "Artificial Intelligence Volume 4" Addison-Wesley Publishing Company, 1984.
- [23] L.N. Cassel, C. Partridge, J. Westcott, "Network Management Architectures and Protocols: Problems and Approaches", IEEE Journal on Selected Areas in Communications, Volume 7, No. 7, September, 1989, pp. 1104 - 1113.

END

15-1 2-93

FIN