LAN Performance Analysis

Product Note

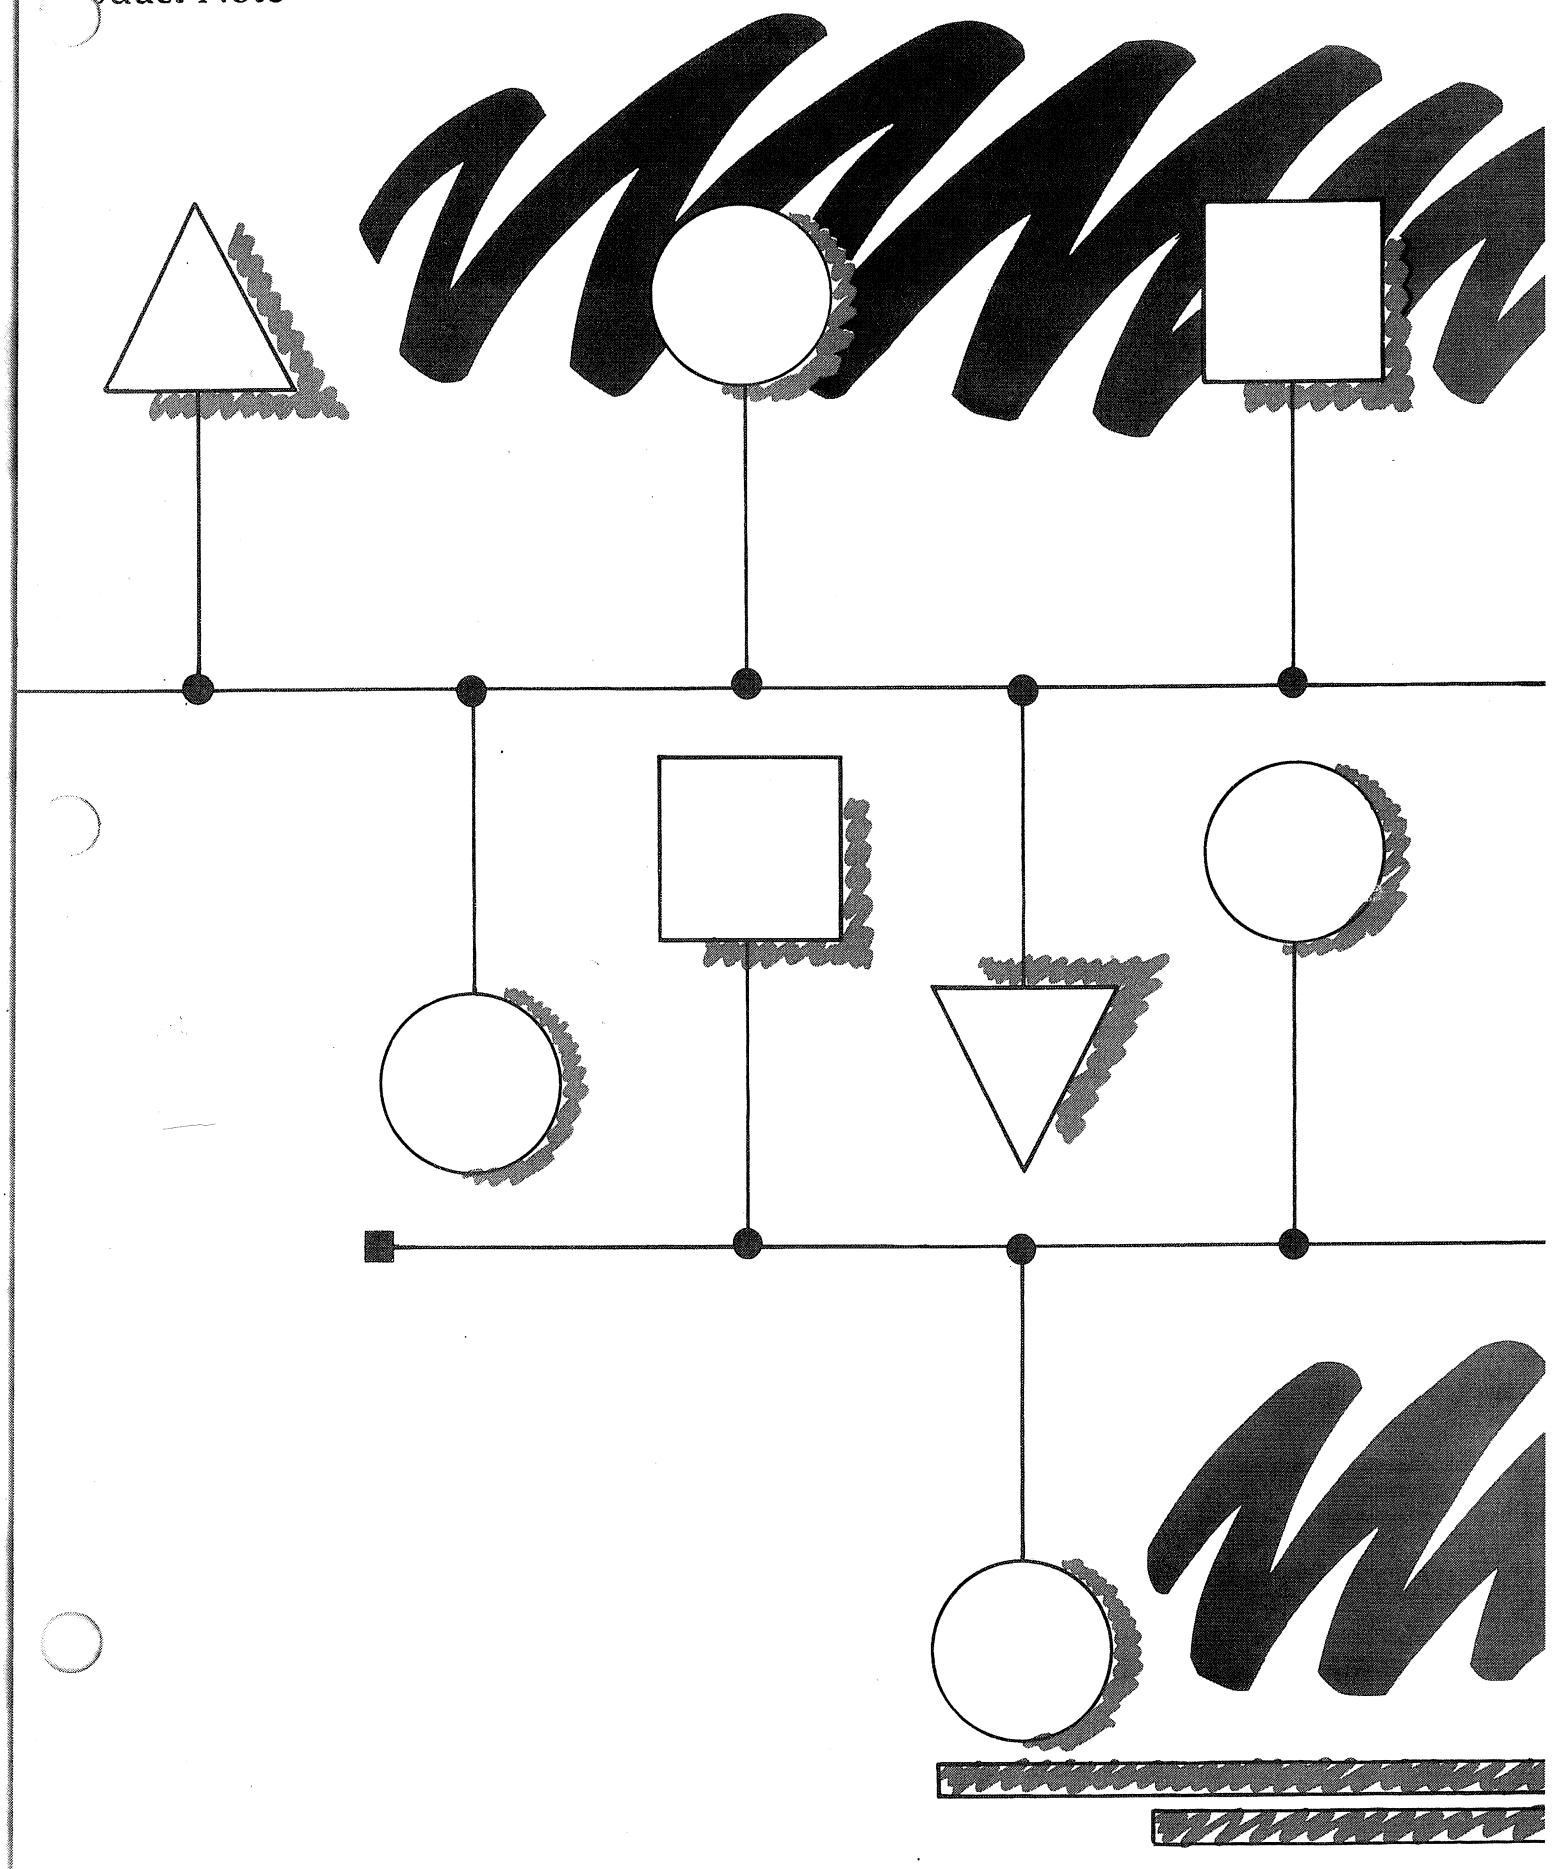

LAN Performance Analysis Product Note

Stalking the Renegade Node
- Five Steps to Network Bliss

Copyright

NOTICE

The information contained in this document is subject to change without notice.

HEWLETT-PACKARD PROVIDES THIS MATERIAL "AS IS" AND MAKES NO WARRANTY OF ANY KIND, EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. HEWLETT-PACKARD SHALL NOT BE LIABLE FOR ERRORS CONTAINED HEREIN OR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS) IN CONNECTION WITH THE FURNISHING, PERFORMANCE OR USE OF THIS MATERIAL WHETHER BASED ON WARRANTY, CONTRACT, OR OTHER LEGAL THEORY.

Some states do not allow the exclusion of implied warranties or the limitation or exclusion of liability for incidental or consequential damages, so the above limitation and exclusions may not apply to you. This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Contents

Chapter

The Layered Approach to Network Performance Analysis ..... 1.

Understanding the Basic Performance of Your Network ..... 2.

To Measure Your Network's Basic Activity ..... 3.

Determining the Sources of Network Traffic ..... 4.

Timing Measurements on the Network..... 5.

Archiving Network Performance Information ..... 6.

Testing for Changes in The Network Load ..... 7.

Network Utilization Calculation ..... Appendix A.

IEEE 802.2, 802.3 and Ethernet Specification ..... Appendix B.

Introduction

Stalking renegade nodes in a network jungle. Even the best big game hunter equips himself for the pursuit of his game. Network managers are no different. When you're seeking the source of a problem, you need every resource available to make your hunt successful. Your task is made even more difficult because the complexity of managing local area networks increases as networks expand. Compounding the issue is the proliferation of mixed-vendor environments among these growing networks. In order to effectively manage the growth and activities of a network, it is important to have a basic understanding of the daily performance and activities of the network.

Existing solutions for the evaluation of network performance include many vendor specific network management systems. These systems were developed to interact with the vendor's software residing at the nodes to obtain statistics information about each node. Since more than one vendor's logical network can coexist on a single cable, the user requires more than one network management system to obtain performance information for the entire network. Furthermore, these systems seldom provide an unbiased view of network activities as they obtain information by polling nodes, not by observing the actual activities on the network.

The HP LAN protocol analyzer and the HP LAN performance analysis application software is designed to provide unbiased information about network activities because it functions independently of any vendors' hardware and software. The analyzer derives network performance information by observing the actual network traffic.

This product note discusses performance analysis of Ethernet and IEEE 802.3 networks. It proposes a five step methodology for analyzing and managing network performance. In particular, it discusses how the HP LAN protocol analyzer and the HP LAN performance analysis system can be used to take those five easy steps to network bliss. Happy hunting.

Chapter 1

The Layered Approach to Network Performance Analysis

When looking at performance issues on a local area network, it is easiest to segment the communication system into the levels of the OSI model since each level uses different parameters to indicate system performance. For example, at the lower levels, the occurrences of collisions and misaligned frames are of interest, and in the upper levels the percentage of protocol overhead or traffic at a logical connection may be more critical in affecting the network's efficiency.

The OSI Model and LAN Architecture

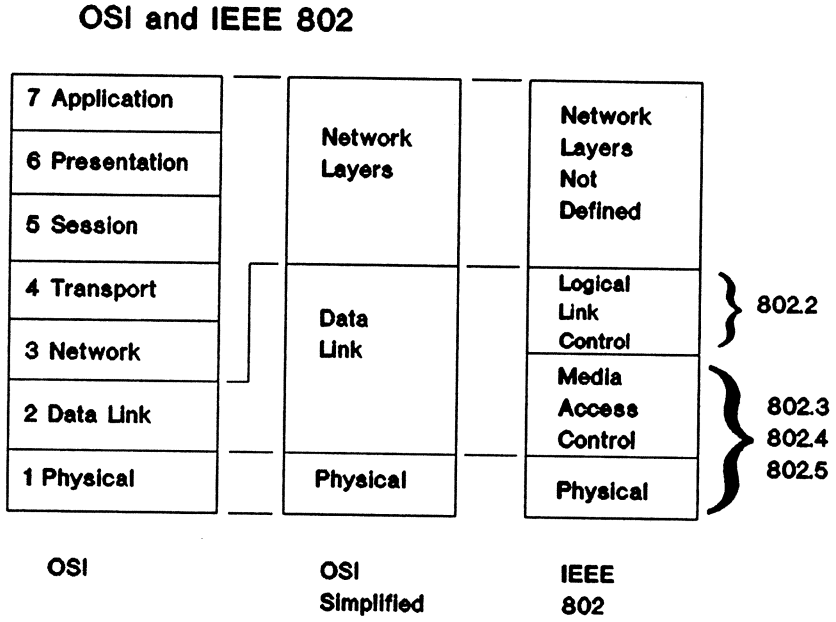

The OSI and IEEE models are both layered approaches to network architecture. In a simplified OSI model, there are three major levels: the physical (level 1), data link (level 2) and network layers (levels 3-7). In the IEEE 802 specification, the data link layer has been divided into two sublayers. These are called the Logical Link Control and Media Access Control. The IEEE 802 standard has subsections that define the protocol involved in each of these layers. The Logical Link Control layer is governed by the IEEE 802.2 standard and there are various standards that deal with the Media Access Control and Physical layers. Specifically the 802.3 is Ethernet-like, 802.4 is token passing on a bus and 802.5 is token passing on a physical ring. The 802 standard does not deal with the network layers of the OSI model. Having an IEEE 802 compatible network does not automatically mean that the individual devices connected to that network can communicate. The devices can be physically connected and can transmit frames, but the meaning of those frames cannot be established until the actual network layers of the two nodes are using identical or peer protocols.

The OSI Layers

The individual layers of the OSI model begin at level 1 and proceed through level 7. There is an implied level 0 which is the actual transmission media. Each level has a dedicated function in the communication model. For the purpose of this product note, we will limit discussion of the layer protocols to those that pertain to Ethernet and IEEE 802.3 local area networks.

Level 0 Physical Media - The physical media in an Ethernet or IEEE 802.3 network includes twisted pair and coaxial cable. The 50-ohm coaxial cable is more commonly used.

Level 1 Physical Layer - The physical layer defines the electrical and mechanical interfaces by which devices are physically connected and data is transmitted. In Ethernet/IEEE 802.3 networks, the physical layer includes a Medium Attachment Unit (MAU) or a transceiver, and the Attachment Unit Interface (AUI) cable or tranceiver cable.

Level 2 Data Link Layer - The data link layer is responsible for moving data reliably across the physical link. The Ethernet, IEEE 802.2 and IEEE 802.3 standards address the data link protocol requirement in the local area network.

Level 3 Network Layer - The network layer provides the means to establish, maintain, and terminate connections between systems as well as switching and routing information. In Ethernet and IEEE 802.3 networks, the Internet Protocol (TCP/IP), Internetwork Datagram Protocol (XNS) and DECnet Routing Protocol are three commonly used network layer protocols.

Level 4 Transport Layer - The transport layer is responsible for the end-to-end data integrity between systems. The Transmission Control Protocol (TCP/IP), Internet Transport Protocol (XNS) and DECnet NSP Protocol are three transport layer protocols commonly used in Ethernet and IEEE 802.3 networks.

Level 5 Session Layer - The session layer sets up and terminates sessions as well as coordinates the interaction between end-application processes. Examples of session layer protocols are ISO 8072 and CCITT X.214.

Level 6 Presentation Layer - The presentation layer is responsible for the character set and data code used and the data display format on a screen or printer. An example of a presentation layer protocol is ISO 8822.

Level 7 Application Layer - The application layer is responsible for the high level functions which provide support to the application or system activities. Examples of the application layer include Common Application Service Elements (CASE), ISO 8571 File Transfer Access and Management (FTAM), and CCITT X.400 Message Handling System (MHS).

Some networks that are based on Ethernet and IEEE 802.3 are DECnet, HP AdvanceNet, and Technical/Office Protocol (TOP). This product note focuses on the performance issues at the lower layers of the OSI model, levels 1 (physical) and 2 (data link).

Chapter 2

Understanding the Basic Performance of Your Network

The first step in the analysis of your network is to understand its basic performance. Key parameters such as the daily utilization (both continuous and peak), its variation over time, and the occurrences of different types of low level errors are all important in determining your network's basic performance.

It is important to remember that every network's performance level is different due to the different devices connected, the types of traffic (terminal versus file transfer traffic), and user's work habits. The best way to understand your network is to make a series of performance measurements and to interpret the information collected in a useful manner.

Network Utilization

What is network utilization ?

Network utilization is defined as the actual number of bits transmitted on the network at any instant divided by the maximum possible number of bits that can be transmitted at the same instant. Therefore, 100% network utilization per second is defined as sending as many bits as possible in one second without violating the network specifications. (See appendix for detail calculation.)

The HP LAN performance analysis system provides the means to gather network utilization information over a day, a week, a month or longer. The information can be plotted, printed or stored on a disc and then compared periodically.

Network utilization can vary drastically depending on the activities of the nodes on the network. If the network has a large number of Computer-Aided-Design (CAD) work stations connected to it, then the traffic tends to be bursty and the frames are long as images are transferred from one work station to another or to storage. Another factor that influences utilization is user's work habits. On networks that support a large number of users doing interactive work, peak traffic times will probably occur in mid-morning and mid-afternoon. Additional utilization peaks may occur if network backups or file synchronizations take place in the evening. It is always useful to know the largest instantaneous peak within a minute throughout a day. This information can be used to correlate trouble reports or user complaints of slow response time.

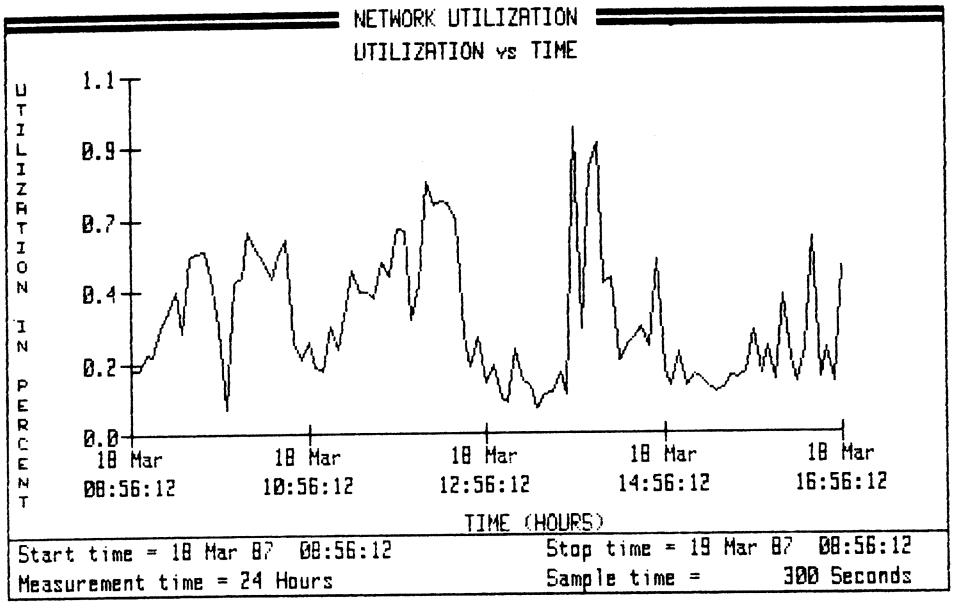

This graph shows utilization vs. time on a network. The sample time is set to 5 minutes. The duration of the measurement is 8 hours. You can see the average utilization during an 8-hour work day, for a network with mostly terminal traffic, is 0.2% with peaks around 10:15 am, 12:30 pm and 2:00 pm. Longer utilization cycles may occur on administrative networks. Peaks may occur on financial cycles, at the end of a month, a quarter or a year. On engineering networks, the utilization may increase close to project deadlines. On manufacturing networks, utilization cycles may correspond to the start/stop of a process.

What do these peaks indicate?

The utilization cycle information, especially the peaks, can be used to correlate trouble reports and poor response time on the network. With this information, you can take measures to improve response time during peak times by altering the utilization cycle or by splitting the network. If network troubles are reported around the peak time, the analysis system can be used to further diagnose the performance problem.

Measuring low level errors

The number of errors occurring on your network also indicates its performance. In Ethernet and IEEE 802.3 networks, errors that occur in levels one and two are bad frame check sequences, misaligned frames, and jabbers. Other parameters of interest include runts and collisions.

Bad Frame Check Sequences (Bad FCS) - A frame check sequence is used for error checking, to ensure that the bits of a frame are transmitted correctly across the network. The frame check sequence is calculated by the source node and transmitted with the frame. The receiving node then recalculates the FCS on the data received and compares it with the original FCS to determine if an error exists. In Ethernet and IEEE 802.3 networks, the FCS is the last 4 bytes of a frame.

Misaligned Frames - These are frames that have bad frame check sequences and the total number of bits in the frame is nondivisible by 8. Misaligned frames are sometimes caused by repeaters or media access units (MAU) that add or subtract bits from a frame. It is sometimes difficult to distinguish between a frame with a bad FCS and one that is misaligned because the error detection is dependent upon the MAU, which sometimes contributes to the error.

Jabbers -Jabber frames exceed the maximum allowable length for a frame on the network and are usually an indication of transmitter failure, but loose connections and other types of level 1 problems are probably the leading cause for jabber frames. When a transmitter continuously sends out bits, it is the MAU's responsibility to cut off the illegal transmission. If a MAU has to exercise its jabber control circuitry, it has to be power reset to function again. It is sometimes difficult to isolate the faulty node, even though the jabber frames are logged to disc, because the frames may not contain any information indicating where they came from. A jabber frame may contain all 1's even in the address fields. In that situation, you may have to resort to more tedious troubleshooting methods, such as a binary cable search, to find the culprit.

Runts - Runt frames are shorter than the minimum allowable length for a frame and are usually the result of collisions on the network. A runt frame is a frame with less than 64 bytes. Runt frames generated by collisions vary in size (all less than 64 bytes) depending on many factors. A runt is usually so small that it does not contain enough information to indicate its origin. Therefore, it is not usually possible to determine which nodes participated in the collision. Even if the information is there, it is not possible to determine which node caused the collision. The content of a runt frame is really not as important as the number of runt frames on the network. An excessive number of runts indicates a problem. The number of runt occurrences differ for each network.

Collisions - Collisions are not errors, but a part of the normal media access operation. Surprisingly enough, most networks do not have many collisions because the CSMA/CD access method used in Ethernet and IEEE 802.3 networks works very well under most network conditions. It is a good idea to measure network utilization as well as collisions over a period of time to see if there is a correlation.

Correlating network utilization and collision counts

Having high network utilization does not mean having a large number of collisions. Network utilization and collisions were measured using the HP LAN protocol analyzer and the HP LAN performance analysis system at an HP facility over a period of time. Daytime utilization was measured at approximately 2%, and nighttime utilization was up at 5%. Moreover, there were more collisions during the day than at night. It was discovered that during the day, there was more random user interactive traffic. At night, the traffic was mainly generated by multiple computer systems doing file backup. Even though the utilization was higher at night, the traffic was predictable. The systems were sending large frames for file backup. They managed to synchronize with one another so that the chance of a collision occurring was minimized.

Chapter 3

To Measure Your Network's Basic Activity

Using the HP LAN Performance Analysis System, you can easily measure your network's basic performance. The Network Summary reports network activities from the start of the measurement. This information is updated every second giving you a snapshot of the network activities.

NETWORK SUMMARY

27 Mar 87 15:14:49

Utilization and Throughput

Current	Average	Peak	
0.38	0.32	0.94	%
38	31	93	kbits/s
32	27	68	frms/s

Frame Parameters

Average Size	178 bytes
Maximum Size	1,510 bytes
Minimum Size	50 bytes
Total Frames	1.588E+6
Total Bytes	3.021E+8

Errors and Collisions

	Bad FCS/Misalign	Runts	Jabbers	Collisions	
Total Count	31	2	0	724	
Average	0.000E+0	0.000E+0	0.000E+0	1.865E-4	Cnt/frm
Peak	2.406E-3	1.552E-4	0.000E+0	9.091E-3	Cnt/frm

Start time = 18 Mar 87 08:56:12 Stop time = 19 Mar 87 08:56:12
Measurement time = 24 Hours Sample time = 300 Seconds

Utilization and Throughput - The utilization measurement is categorized into the current utilization, the average percentage and the peak percentage since the start of the measurement. Throughput is measured as the number of frames, or bits that have passed through the network.

Errors and Collisions - The analyzer keeps track of errors such as misaligned/bad FCS frames, and jabbers. It also counts the occurrences of runts and collisions.

Frame Size - This measurement is displayed as the average frame size and the maximum frame size since the start of test. Typically, terminal traffic tends to generate 64-byte frames as do most devices on the network. Large frames only appear when a file transfer or an image transfer occurs. These measurements are displayed in graphical as well as in tabular format.

To gather performance data over time, the Network Statistics menu should be selected. In this menu, the user can select parameters used in the measurements such as sample time, duration of each measurement and display boundaries.

Measuring Network Performance: Utilization, Errors/Collisions, Frame Timing, Frame Length

In the Network Statistics measurement, you can simultaneously collect information on network utilization, errors, collisions, interframe timing, and frame size. By varying the duration of the measurement, you can collect information over a long period of time. This information can be logged to disc, or output to a printer or plotter.

The following measurements are taken simultaneously over the same period of time.

- **Network Utilization**

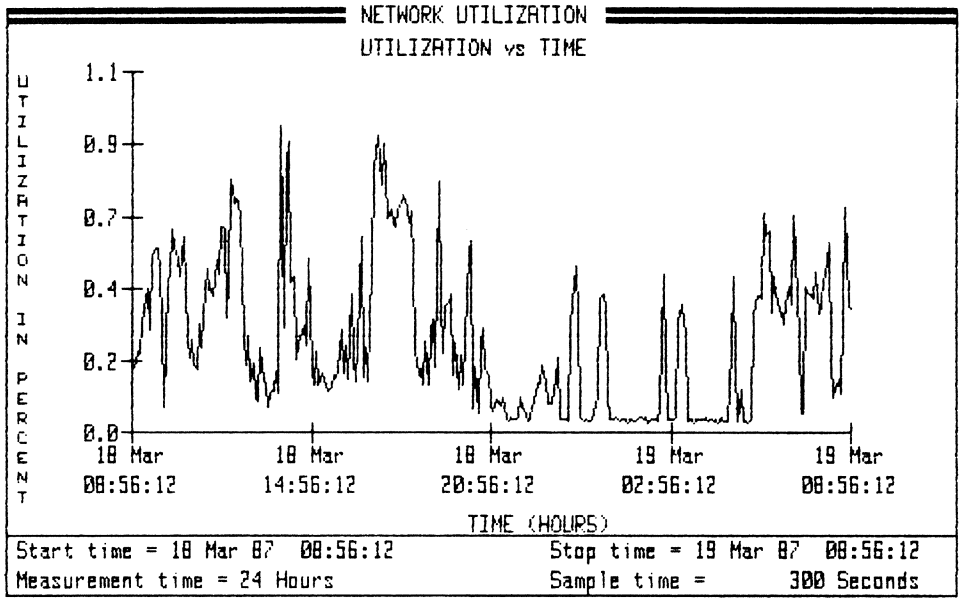

Knowing normal utilization characteristics, you can correlate the information with the error occurrences during the same period. The above graph shows network utilization vs. time measured over a 24 hour period. The analysis system allows for 360 samples with sample time from one second to four hours. The measurement value for each sample time is the average value during that period of time. One timing setup applies to all the network level measurements. If you want to make a measurement over a 24 hour period, a sample time of 5 minutes is recommended. Statistical data can also be stored at every sample time for review later on. If you need to make a measurement over a longer period of time, for example over 24 hours, a sample time of one hour is recommended. This allows for a measurement duration of 360 hours (15 days). A sample time longer than one hour may not give enough resolution for meaningful results. A measurement period greater than 15 days using a sample time of one hour can be obtained by using the autosequence features and menu described in Chapter 6.

- Network Errors and Collisions

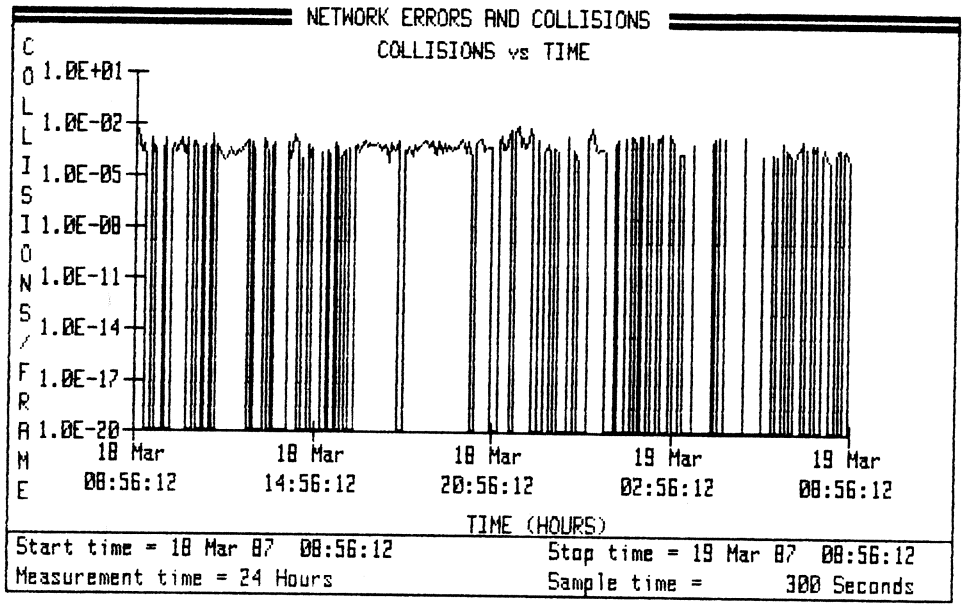

Collision occurrences are measured over the same 24 hour period. Since this network has mostly terminal traffic, it is characterized by a good deal of collisions. Other network errors, such as bad FCS/misaligned frames, runts and jabbers are also measured in the same time period.

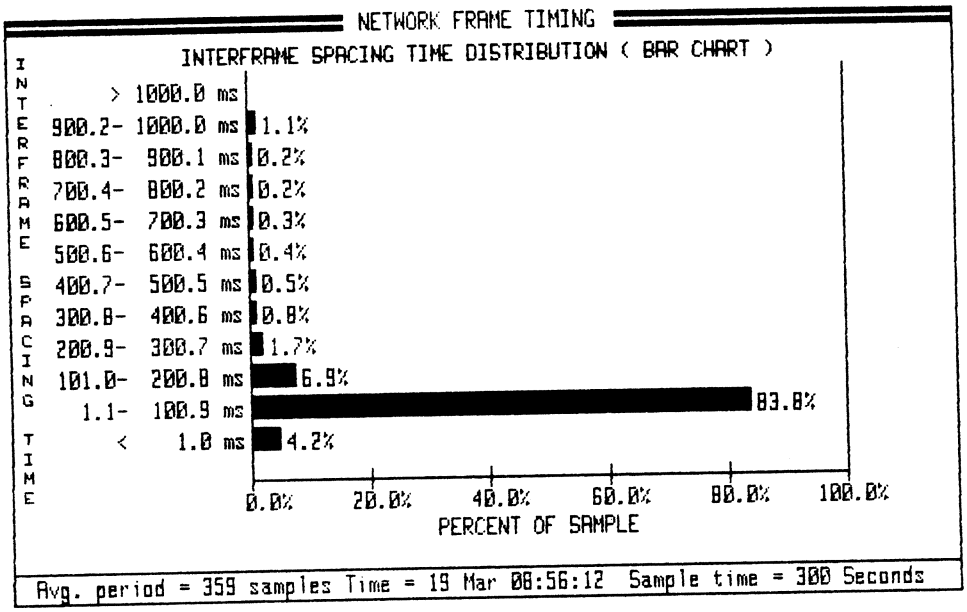

This measurement, taken over the same 24 hour period, shows the percentage of frames with different interframe spacing. On a busy network, the frames appear very close together and on a lightly loaded network, frames are further apart. On the other hand, if a lightly loaded network has a small interframe spacing for most of the frames, it shows that the frames appear in bursts on the network.

A parameter that influences the efficiency of the network is the burstiness of the traffic. Some nodes, particularly personal computers, can miss some frames destined for them if they arrive too close together. This occurs when the node or its network interface is not able to process the incoming frames fast enough. These errors are often masked from the user by the upper layer protocols. The retransmission takes place until the frames are properly received by the node. Retransmission of messsages wastes network capacity and affects node response time.

- **Network Frame Length**

Another important network parameter in the understanding of basic network characteristics is the frame size. A network with very small frame size is not fully utilizing the bandwidth of the network. Some devices may be configured for larger frame sizes so that the network bandwidth can be efficiently used. Of course, there is a balance between response time and network efficiency. Through experimentation, you can determine what is most appropriate for your network.

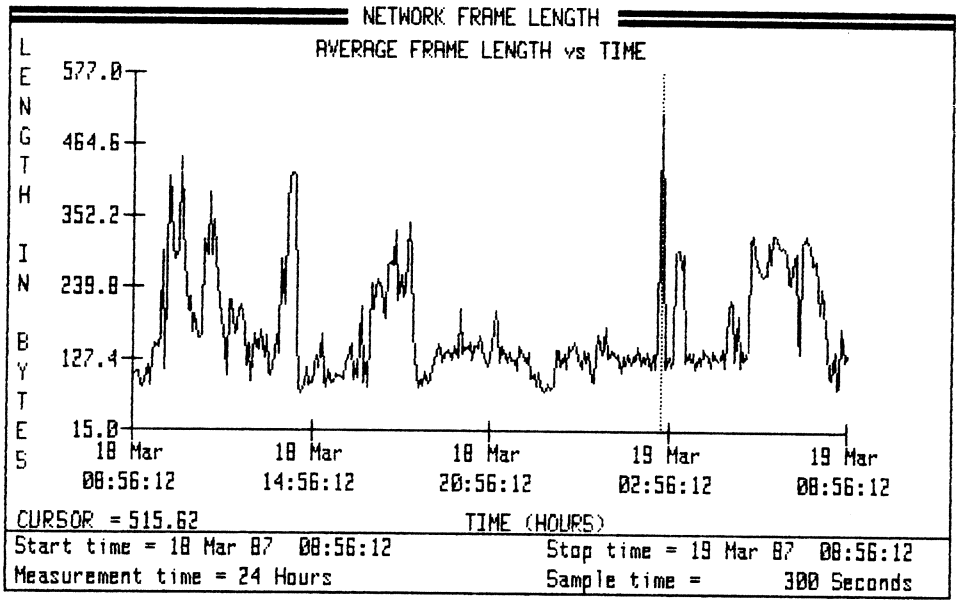

The first graph shows the average frame length vs. time. The measurement is taken over a 24 hour period as in network utilization. In this graph, the important information is not so much the actual values, but the changes in the values and the time of their occurrence. The changes, especially the peaks, indicate large file transfers taking place.

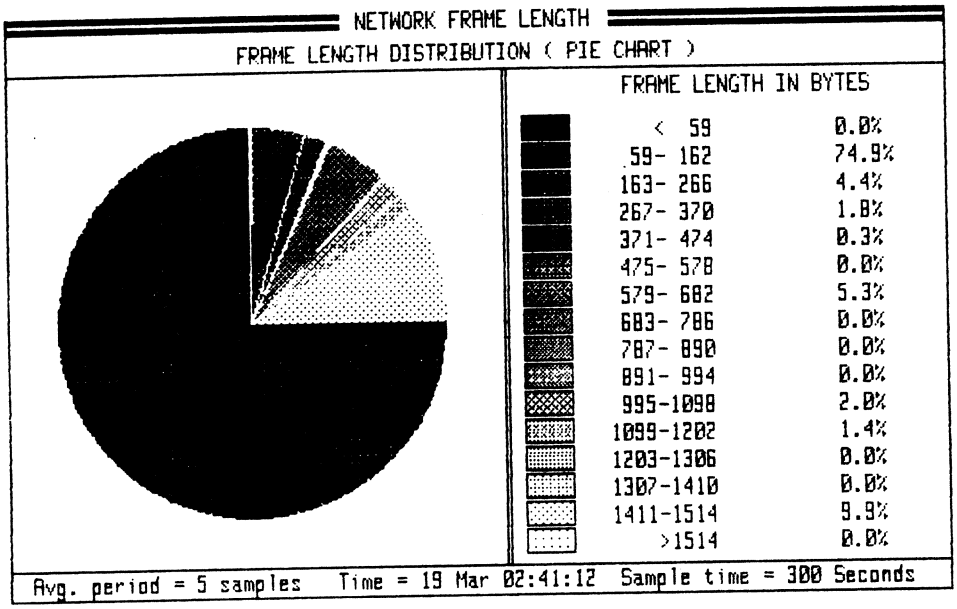

The second graph shows the detail frame length distribution during one of the peak times (marked by the cursor on the average frame length graph). In this sample time, there are a number of large frames due to the file transfer. Typical frame size on this network is between 64 to 200 bytes.

Many parameters affect the frame size in a network, among them are protocol used, type of end-user function supported, and internet connections. Protocols such as Xerox Network System (XNS) transport protocols limit the data field to less than 600 btyes while IEEE 802.3/Ethernet supports a maximum data field of 1500 bytes. Node configuration parameters are another factor that dictate frame size. End-user devices such as terminal servers tend to generate small, minimum length frames which allow the network to be very responsive but wastes capacity. File transfer applications generate maximum length frames using the network efficiently, but can block network access under some conditions. Asynchronous terminal servers communicating to host computers tend to generate small frames whereas host computers tend to generate larger frames since larger amounts of information are returned to the terminal. Finally, internet traffic also affects frame size. If an internet connected to the LAN supports short frames, a message may be split into several small frames before being forwarded to its destination node.

Using the guidelines provided in this chapter, measure your network's activities and fill in the following chart.

Your Network's Activities

	8 Hours	24 Hours	1 Week
Average Utilization			
Percentage			
kbits/sec			
Frames/sec			
Peak Utilization			
Time of occurrence			
kbits/sec			
Frames/sec			
Frame Parameters			
Average size			
Maximum			
Minimum			
Errors & Collisions			
Bad FCS/misaligned frames			
Runts			
Jabbers			
Collisions			

Questions To Think About...

What are the utilization characteristics of your network ? When are the peak times ? Are they what you expected ? What are they caused by ?

What is the average frame size on your network ? What type of traffic is on your network ?

What types of errors are on your network ? How does the collision rate correspond to the utilization characteristics on your network ?

Chapter 4

Determining the Sources of Network Traffic

Once you have established the baseline performance of your network, the next step is to determine the sources of network traffic. It is important to know which nodes are generating most of the traffic on the network because when a problem occurs, these nodes are the most likely to be involved. Knowing the sources of traffic on your network allows you to configure and segment the network based on network load. You also need to identify the nodes that are traffic concentrators because they may be critical devices on the network such as routers, bridges, file servers and printers. The overloading of such devices degrades network performance.

Node List Statistics

The easiest way to determine activity of the nodes on the network is by using the NODE LIST STATISTICS. In this measurement, the analyzer creates a node list by observing the network and collecting all the active addresses. At the same time, it also monitors node traffic activities. Node activities monitored include: frames transmitted and received, error count, error rate, and average frame size. The node list allows the user to assign a meaningful name to the Ethernet/IEEE 802.3 6-byte hex address.

NODE LIST STATISTICS

1 Sep 86 01:45:03

Node #	Node Name or Address	Last Sample Frame Cnt.		Frame Count	KByte Count	Error Count	Error Rate	Avg.Frm Size
1	hpfcrg	XMT	14	268	22	0	0.00E+0	71
		RCV	14	1,200	235	0	0.00E+0	184
2	hpfclq	XMT	14	644	55	0	0.00E+0	74
		RCV	14	625	50	0	0.00E+0	67
3	hpcnoa	XMT	10	536	45	0	0.00E+0	72
		RCV	10	500	36	0	0.00E+0	60
4	08-00-09-00-35-4A	XMT	1	520	51	0	0.00E+0	86
		RCV	2	465	36	0	0.00E+0	66
5	hpfcde	XMT	2	460	36	0	0.00E+0	66
		RCV	0	430	31	0	0.00E+0	60
6	hpfcdq	XMT	0	73	18	0	0.00E+0	234
		RCV	0	715	164	0	0.00E+0	217

Start time = 13 Jan 87 14:14:30 Stop time = 13 Jan 87 14:15:30
Measurement time = 60 Seconds Sample time = 2 Seconds

Printing

Abort
Print

If you want to study the traffic from a selected set of nodes, the analyzer can monitor only those specified addresses. If you have a standard node list for your network and your network is large and dynamic, the measurement can be set up to add new addresses to an existing list. While the measurement is running, the nodes are sorted by the total number of frames transmitted and received with the busiest node appearing on top of the list. The node list can accomodate 1000 node addresses.

It is also very convenient to be able to sort by error count and error rate as well. You should take note of the busiest nodes and high error rate nodes because these are the most likely problem nodes on the network.

Connections-to-node Measurement

Having identified the busiest nodes, you may want to do further traffic analysis on these nodes. Some nodes, such as file servers, gateways, routers and bridges, need regular monitoring. For bridges, the number of frames handled every second is critical. As the number of frames approach the performance limit of that device, steps need to be taken to optimize the situation. For file servers, the number of requests handled at any time is critical. If the number of requests processed continually exceeds its capability, then additional file servers may be needed.

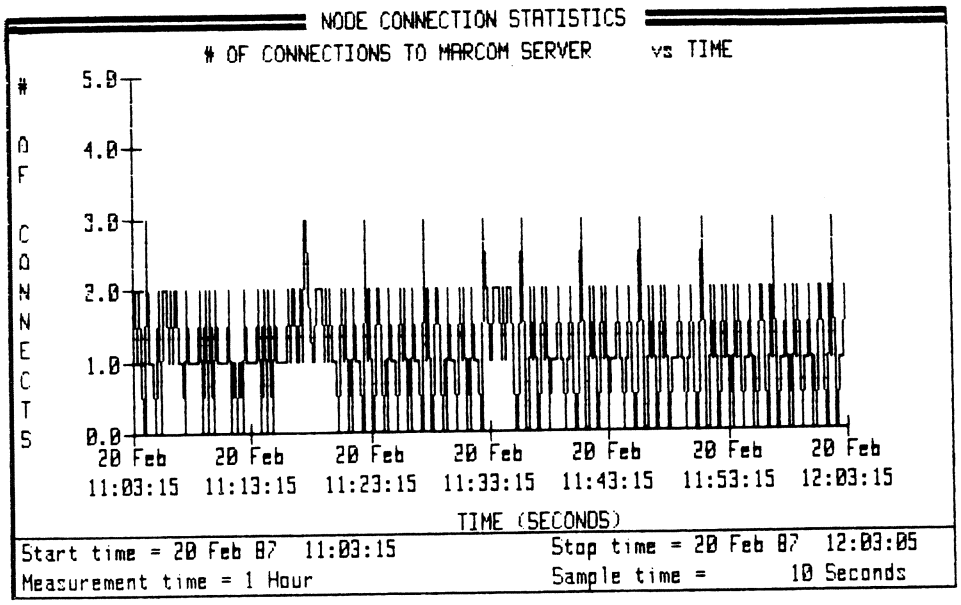

Using the connection-to-a-node measurement, you can determine the number of connections made to a node at any given time over the test period. In the graph shown, the number of level 2 (Ethernet or IEEE 802.3) addresses talking with a file server is monitored over a one hour period. The sample time was 10 seconds. Since this is a small network, with four nodes, the server has no problem supporting multiple nodes at the same time.

Connection Summary Measurement

Having identified the busiest nodes, it is also useful to know the busiest connection on your network. Having this information about your network helps to segment traffic when you need to reconfigure a busy network segment.

CONNECTION SUMMARY

19 Jan 87 11:26:19

Connection Node Name or Address	Last Sample Frame Cnt.		Frame Count	KByte Count	Error Count	Error Rate	Avg. Frm Size
hpdcdb	XMT	0	142	13	0	0.00E+0	77
hpfcrt	XMT	0	192	14	0	0.00E+0	60
hpfcla	XMT	6	136	10	0	0.00E+0	60
hpfcrj	XMT	6	149	35	0	0.00E+0	224
08-00-09-00-2E-54	XMT	4	100	13	0	0.00E+0	118
08-00-09-00-96-D0	XMT	7	141	10	0	0.00E+0	60
hpfcly	XMT	1	105	10	0	0.00E+0	81
08-00-09-00-7E-F5	XMT	1	103	7	0	0.00E+0	60

Start time = 13 Jan 87 14:03:22 Stop time = 13 Jan 87 14:04:22
Measurement time = 60 Seconds Sample time = 1 Second

Printing

Abort
Print

The connection summary measurement determines the busiest level 2 (Ethernet or IEEE 802.3) connections on your network, monitors the number of frames passed between the two nodes, the error count or error rate, as well as the frame size.

Communication Matrix

To examine connection level traffic pictorially, the communication matrix is the ideal tool. This measurement, making use of the top 30 addresses collected from the node list statistics, monitors the traffic among these busiest nodes.

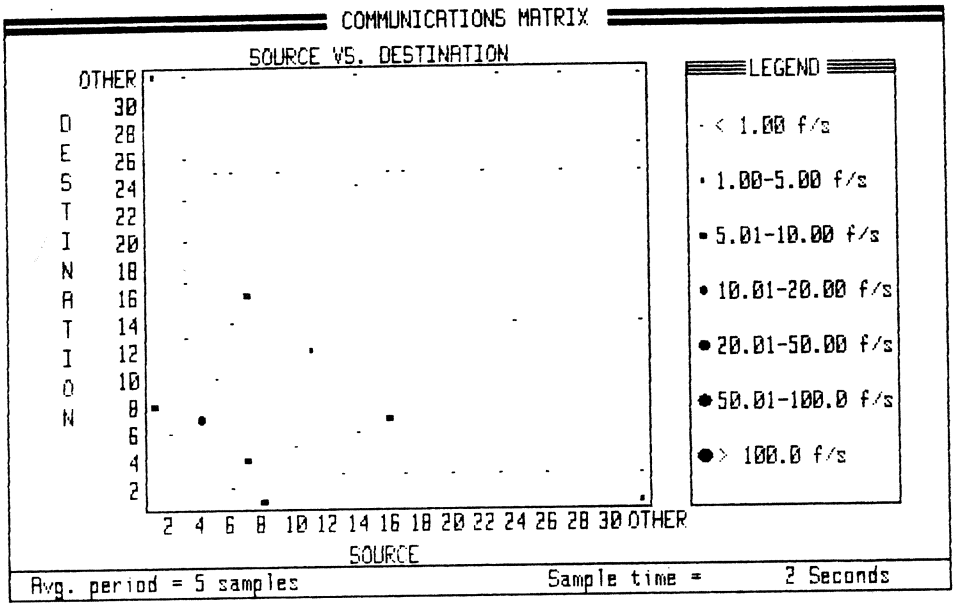

There are 31 addresses for the vertical and horizontal axis. Addresses on the vertical axis are destinations; addresses on the horizontal axis, sources. Frames having a particular source and destination address will show up on the matrix as a dot. The size and the shape of the dot denote the rate of frames transmitted on the network with that pair of addresses. The key information from this matrix is monitoring the traffic pattern of the nodes. If the node is a server or router, a row of dots will show up either horizontally or vertically showing the device talking to several other devices at the same time. For pairs of nodes conversing with each other, the dots will reflect across the diagonal of the matrix. The actual addresses of the nodes can be determined by moving the cursor to a dot. The addresses and the traffic level (frames/second) represented by the dot will be shown.

The communication matrix provides an easy graphical way to get a "feel" for your traffic characteristics.

Other Network Traffic Characterisitcs

Network Frame Analysis

If your network uses more than one upper level protocol, you may want to know which type is used more often. This indicates the utilization of a particular group of devices over other groups sharing the network. Further studies can then be made to ensure that these devices are not overloaded.

Using the filters in the HP 4971S, the user can define any paramenter within a frame for statistical analysis. In this example, the TYPE field is specified to analyze the different types of protocols used on the network.

TYPE field	Protocol
06-00	XNS
08-00	TCP/IP
60-01	DECMOP
60-03	DECNSP
60-04	DECLAT

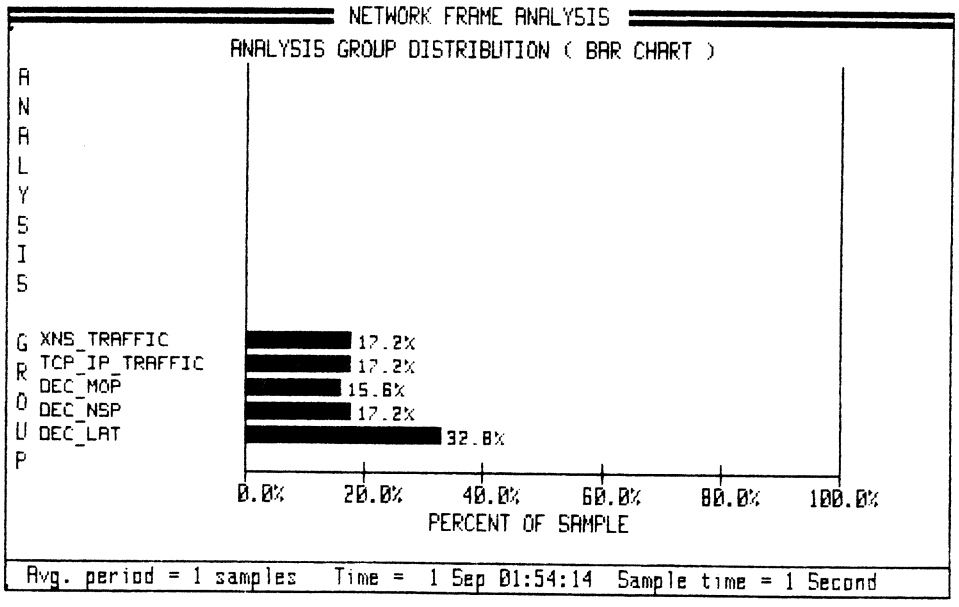

In the frame analysis measurement, the percentage of that type of protocol used is measured. In the same manner, any information within the data field can be statistically analyzed.

Types of traffic that waste network capacity and affect node response time are excess broadcast and multicast frames. Broadcast and multicast frames are addressed to all nodes or groups of nodes and must always be processed by the receiving nodes. A node can be overloaded by broadcast or multicast frames if its network interface has to continually process them.

To monitor the number of broadcast and multicast frames, you can define two filters specifying FF-FF-FF-FF-FF-FF for broadcast address and a multicast address in the destination address field. In the Network Level Frame Analysis measurement, these two filters are used to statistically analyze broadcast and multicast traffic on the network, reporting the percentage of traffic due to broadcast and multicast frames.

FILTER #	FILTER LABEL	DESTINATION ADDRESS	SOURCE ADDRESS	LENGTH OF FRAME: MINIMUM	MAXIMUM
0	XNS_TRAFFIC	XX-XX-XX-XX-XX-XX	XX-XX-XX-XX-XX-XX	15	2022
1	TCP_IP_TRAFFIC	XX-XX-XX-XX-XX-XX	XX-XX-XX-XX-XX-XX	15	2022
2	DEC_MOP	XX-XX-XX-XX-XX-XX	XX-XX-XX-XX-XX-XX	15	2022
3	DEC_NSP	XX-XX-XX-XX-XX-XX	XX-XX-XX-XX-XX-XX	15	2022
4	DEC_LAT	XX-XX-XX-XX-XX-XX	XX-XX-XX-XX-XX-XX	15	2022
5	BROADCAST	FF-FF-FF-FF-FF-FF	XX-XX-XX-XX-XX-XX	15	2022
6	MULTICAST	AB-00-04-XX-XX-XX	XX-XX-XX-XX-XX-XX	15	2022

46 Filter hardware bytes available

Add Filter | Delete Filter | Show Node Names * Show Hex Addresses | Edit all Fields | OTHER CHOICES | EXIT

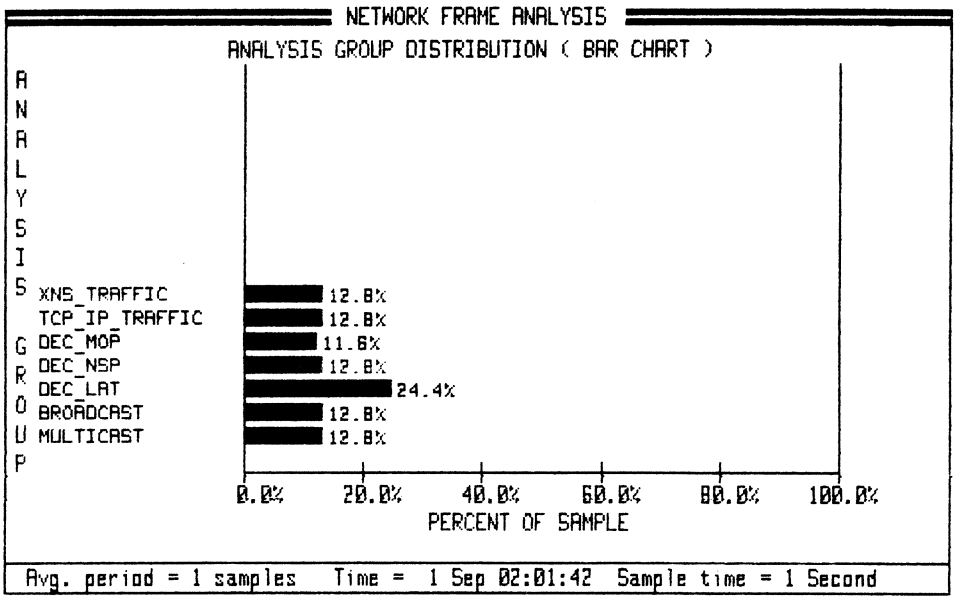

In the above example, the broadcast address is FF-FF-FF-FF-FF-FF, and the multicast address specified, AB-00-04-XX-XX-XX, is for DECnet. The percentage of traffic on the network due to these two addresses is monitored.

Using the guidelines discussed in this chapter, determine the sources of traffic on your network and fill in the following chart.

Sources of Your Network Traffic

	Frames transmitted	Frames received	Avg. frame size
The 5 busiest nodes			
1.			
2.			
3.			
4.			
5.			
The 5 busiest connections			
1.			
2.			
3.			
4.			
5.			

Questions To Think About...

What devices are at your busiest nodes ? Which device generate the most errors ?

Which pairs are your busiest connections ? If you have made the

What percentage of your network traffic is broadcast or multicast frames ?

If your network uses more than one type of protocol, what percentage of the traffic is each type ?

Chapter 5 Timing Measurements On The Network

As your network grows in number of nodes, traffic level, or geographical dispersion, delay measurements become important because network delay affects the response time of the network. The important timing measurements on the network are channel acquisition time and network response time.

Channel Acquisition Time

Channel acquisition time is measured from the time a message is ready to be sent to the time the frame appears on the network. This is a measurement of the delay in transmission due to the physical network. In a CSMA/CD network, such as Ethernet or IEEE 802.3 networks, the channel acquisition time is measured from the time the network interface is presented with the message for transmission to the time the message is successfully placed on the network. This takes into account all aspects of the specific media access control system, such as possible collisions, deferrals and backoffs. Measuring the variation of the channel acquisition time over a period of time indicates the response of the network to different types and amounts of network load.

CHANNEL ACQUISITION TIME

8 Jan 87 16:05:01

ABSOLUTE TIME	ACQUISITION TIME	% DEFERRED	COLL/MSG	% ABORTED
8 Jan 87 16:03:25	51 us	13.86	0.00	0.00
8 Jan 87 16:03:26	72 us	12.87	0.00	0.00
8 Jan 87 16:03:27	14 us	7.00	0.00	0.00
8 Jan 87 16:03:28	14 us	6.00	0.00	0.00
8 Jan 87 16:03:29	52 us	10.89	0.00	0.00
8 Jan 87 16:03:30	32 us	6.00	0.00	0.00
8 Jan 87 16:03:31	30 us	10.00	0.00	0.00
8 Jan 87 16:03:32	31 us	6.93	0.00	0.00
8 Jan 87 16:03:33	26 us	11.00	0.00	0.00
8 Jan 87 16:03:34	43 us	11.88	0.00	0.00
8 Jan 87 16:03:35	23 us	4.00	0.00	0.00
ALARMS (MAX) :	(100,000 us)	(50.00)	(8.00)	(0.00)

Start time = 8 Jan 87 16:03:24 Stop time = 8 Jan 87 16:04:24
Measurement time = 60 Seconds Sample time = 1 Second Msg/Second = 100

Printing

Abort
Print

The performance analysis system allows you to measure the channel acquisition time of the analyzer. In the example above, the analyzer sends out 100 messages/second with a sample time of one second. The ACQUISITION TIME is the average time it took the analyzer to transmit one message onto the network taking into account collisions, deferrals and backoffs. It also keeps track of the percent of messages it had to defer, i.e. the analyzer was ready to send the frame but the network was busy. As a result, the analyzer waited for the end of the transmission and 9.6 microseconds passed before sending out the frame. The COLL/MSG indicates the number of collisions encountered per 100 messages. The % ABORTED indicates the percent of messages aborted after 16 attempts stopped by collisions. The

channel acquisition time measurement result gives an indication of the time required for a node on the same segment to send a frame onto the network.

Network Response Time

Network response time measures the time it takes a message to travel to its destination and back. If two networks are connected through a satellite link or any wide area network, then it is useful to know the delays due to these connections. If a user is establishing a session with a computer across town and is having response time problems, the network response time measurement will be able to help you determine whether the delay is caused by the long distance connection. By comparing the normal delay with the measured delay, you can determine your next course of action.

```
MESSAGE #0 in format of XNS_Echo_Request        MESSAGE Label: XNS_Echo_Request
                                                FRAME Length:     80 Bytes
   FIELD LABEL                    HEX DATA                      CHARACTER DATA
----------------  ----------------------------------------------  --------------
DESTINATION       11-11-11-11-11-11    Node Name: -- Not Defined --
SOURCE            08-00-09-00-0A-2A    Node Name: -- Not Defined --
TYPE              06-00
Echo_Req_Chksum   FF-FF                                           ■■
IDP_Length        00-42                                           %B
IDP_Transport_Cnt 00                                              %
IDP_Packet_Type   02                                              %
IDP_Dest_Network  00-00-00-00                                     %%%%
IDP_Dest_Host     11-11-11-11-11-11                               %%%%%%
IDP_Dest_Socket   00-02                                           %%
IDP_Source        00-00-00-00                                     %%%%
IDP_Source_Host   08-00-09-00-0A-2A                               %%%%%%*
IDP_Source_Socket 00-02                                           %%
ER_Operation      00-01                                           %%
```

```
                        Katakana character data entry

JASCII 7  JASCII 7  JASCII 7  JASCII 7     JIS 8     EBCDIK                EXIT
 Par = 0   Par = 1   Odd Par   Even Par    Katakana*
```

For example, the XNS (Xerox Network Service) protocol defines a level 4 loopback protocol. To measure the response time of a node that employs the XNS protocol, the loopback message is created and then sent to the destination node. The node, upon receiving the message, will send back a responding message. The analysis system measures the total time it takes for the loopback message to reach the destination node and come back.

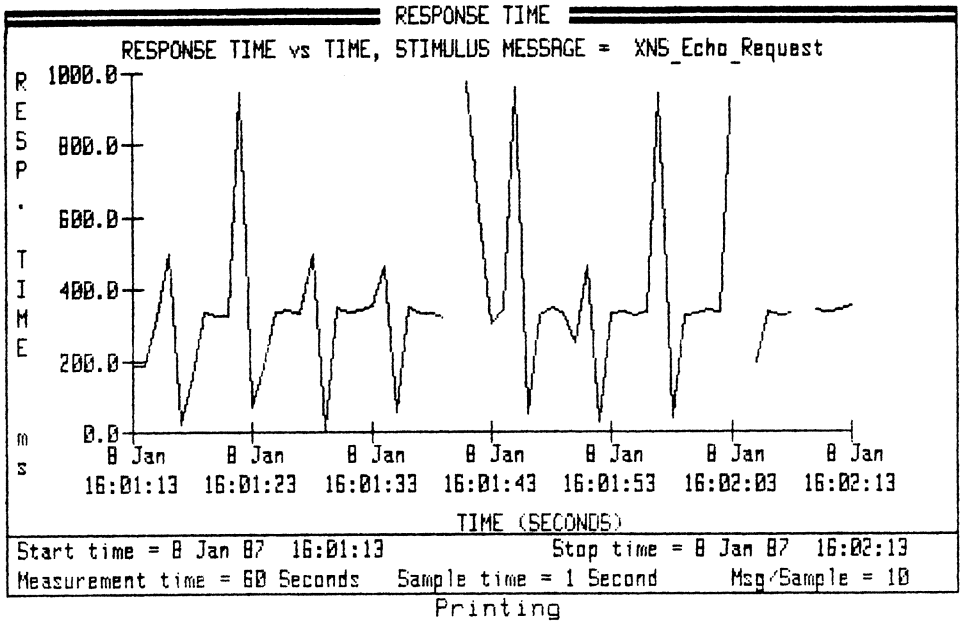

Abort
Print

The loopback message is sent up to 10 frames/second. This rate may be limited by the amount of time it takes the responding node to reply to the message. The average delay time is plotted on the graph. The variations of the response time in the graph shown are due to the low priority for servicing loopback messages at the node. The XNS loopback protocol is one of the many loopback protocols defined for local area networks. Other examples of loopback protocol includes the IEEE 802.2 XID and Ethernet loopback protocol.

The network response time measurement can also be used for troubleshooting the protocol stack since it is possible to implement loopback capability at every level of the stack. If a protocol problem is suspected, this measurement can be used to test the different layers starting from lowest layer of the stack. You can then easily identify the layer that is not functioning properly.

Using different loopback protocols defined for each layer in the protocol stack, the network response time measurement can also be used to measure protocol processing time. For example, a level 4 loopback takes longer than a level 2 loopback. The delay in a level 4 loopback is mainly due to protocol processing time whereas a level 2 loopback is mainly signal propagation time. This measurement can give some insight into the efficiencies of the different protocols.

Chapter 6

Archiving Network Performance Information

In order to manage the growth of your network and be aware of changes over time, you need to build up a database of your network's performance information. The objective is to collect the performance data of your network over a long period of time so that you can use the information to project growth and to study the changes of the network.

Autosequence/ Alarms Testing

The HP LAN network performance analysis system can be easily set up for gathering network performance information over a long period of time. These measurements are also easily repeatable. Using autosequence testing, the network can be monitored and performance information collected without operator interference. Autosequence testing also allows logging of unusual network activities, such as higher than normal utilization and excessive error occurrences, based on user-definable alarms. Alarms are thresholds for each measurement set by the user so that when any one is exceeded, it causes the system to log the information to a disc, output the information to a printer or execute another measurement.

Here is an example to help you understand how the autosequence testing and alarms can be used to your advantage. For example, your network utilization is typically 3% between 8 am and 10 am, however, on a particular day, it went up to almost 10% . At the same time, trouble calls came in. How do you go about diagnosing the problem ? The first step would be to investigate sources of the extra traffic, that is, which node or nodes were responsible for it. An autosequence and alarms menu such as the following would help to determine the problem.

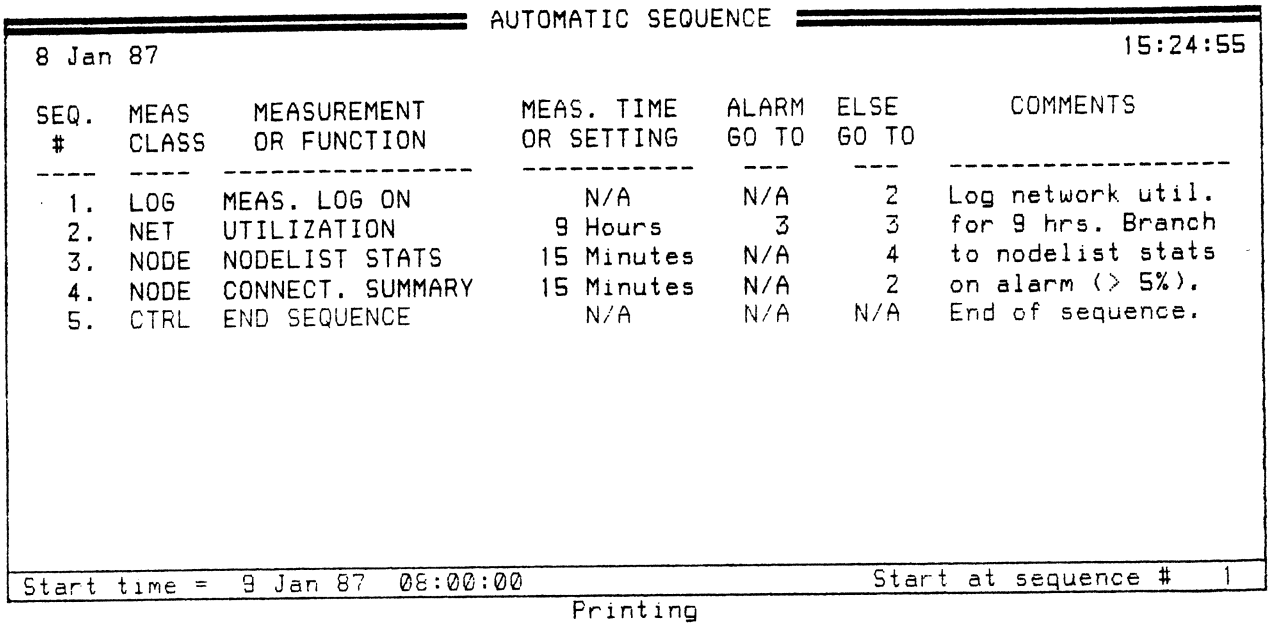

AUTOMATIC SEQUENCE

8 Jan 87 15:24:55

SEQ. #	MEAS CLASS	MEASUREMENT OR FUNCTION	MEAS. TIME OR SETTING	ALARM GO TO	ELSE GO TO	COMMENTS
1.	LOG	MEAS. LOG ON	N/A	N/A	2	Log network util.
2.	NET	UTILIZATION	9 Hours	3	3	for 9 hrs. Branch
3.	NODE	NODELIST STATS	15 Minutes	N/A	4	to nodelist stats
4.	NODE	CONNECT. SUMMARY	15 Minutes	N/A	2	on alarm (> 5%).
5.	CTRL	END SEQUENCE	N/A	N/A	N/A	End of sequence.

Start time = 9 Jan 87 08:00:00 Start at sequence # 1

Printing

Abort
Print

In the above autosequence, logging is activated, meaning performance information will be logged to the disc at every sample time. Next, the sequence is to start gathering network performance data at 8:00 am everyday. The first measurement is network utilization, the measurement duration is 9 hours. Since the the analysis system collects all network performance data at the same time, you will have collected 9 hours' worth of network utilization, errors and collisions, frame timing, frame length and any frame analysis information that you have specified. You can set alarms for the network utilization as well as error and collision counts so that when the thresholds are exceeded, the analysis system can branch to another measurement which can provide more information on the network anomaly. In this example, tripping the alarm causes the system to branch to the node list statistics. By doing so, the node list statistics will be able to provide information regarding the sources of the extra traffic and errors.

Other than logging all the network information at every sample time, the analysis can be set up to log the network information only when a threshold is exceeded. The same autosequence in the example can be set up the following way to log only when a threshold is exceeded:

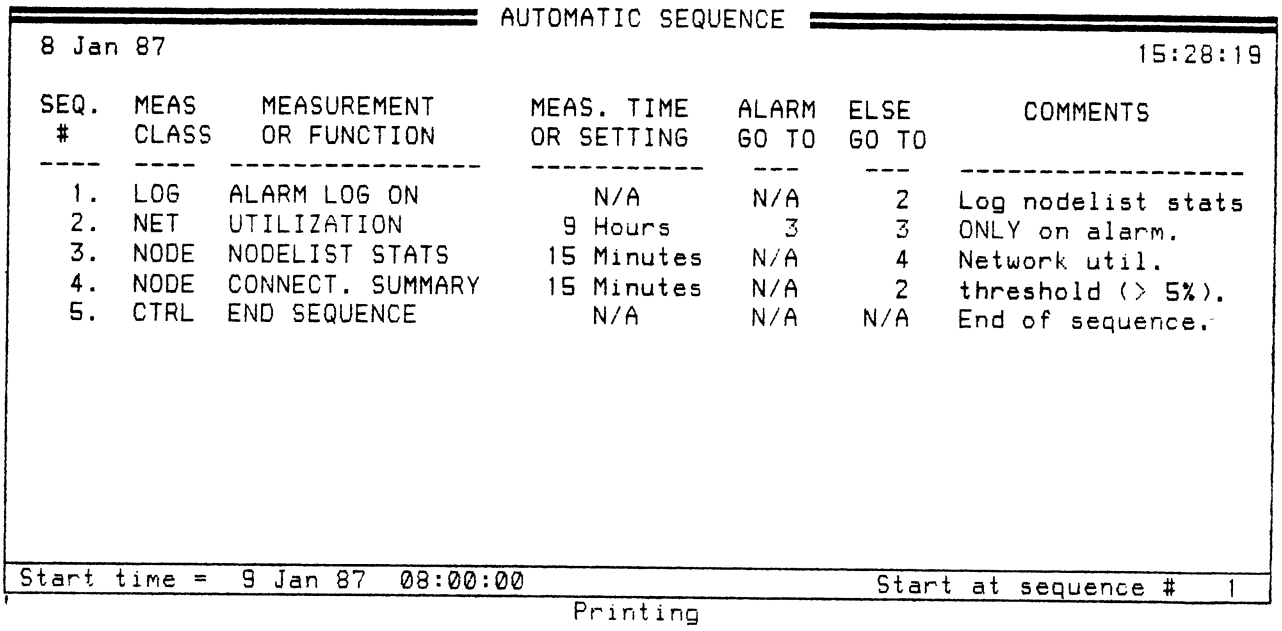

In this autosequence, alarm logging is activated. If any network level thresholds are exceeded while the system is in the network utilization category, it will branch to the node list statistics. At the same time, logging will be activated and the node list statistics will be logged to disc. Using the autosequence and event-triggered recording, important performance data will never be missed. Event-triggered logging is also supported for the printer and plotter.

The alarms menu can be activated to alert the operator as measurements are made. Alarms are available for network utilization as percentage, frames/second and kbit/second. It can also be set for error and collision counts, the number of connections to a node, channel acquisition and network response timing.

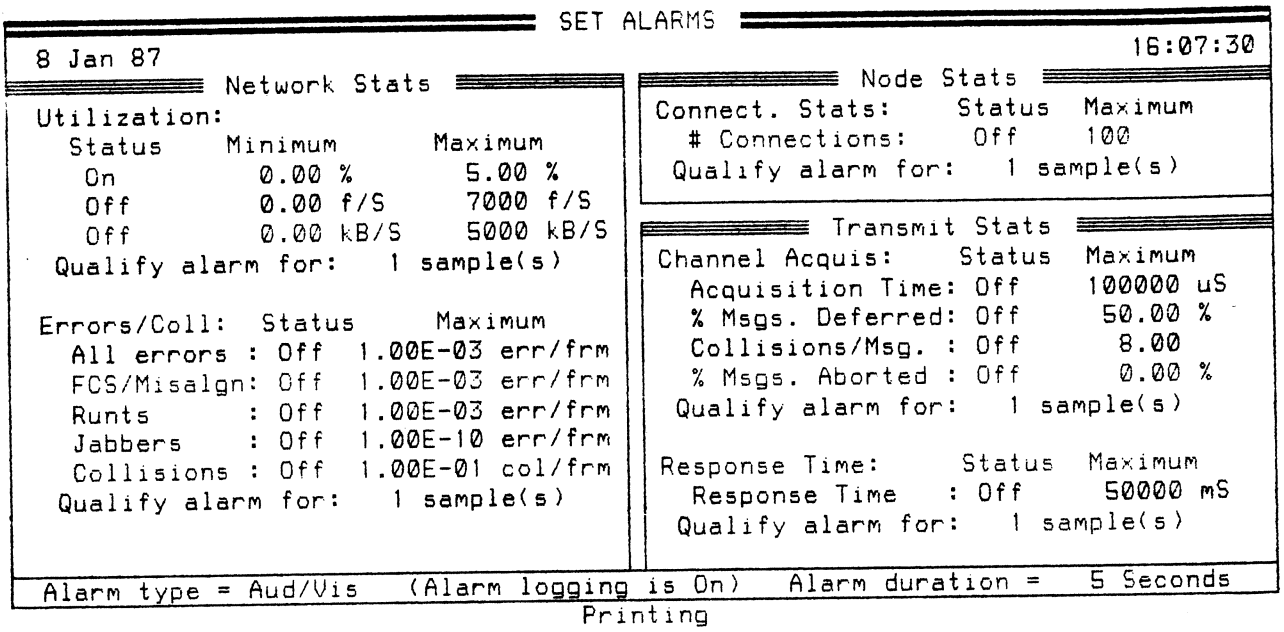

The "qualify alarm for" fields allow the user to require that the alarm condition be present for some number of measurement samples before sending the alert. The alert comes in three formats - audio, visual or both.

Plotting and Printing Network Performance Information

Any tabular and graphic screen can be ouput to an inkjet printer. All graphical screens can be output to a plotter. The user can specify the output device and then simply use the "SHIFT/PRINT" keys on the keyboard to output the screen. The performance analysis system supports printing to an HP 2225A ThinkJet Printer. Plotters supported are the HP 7475A, HP 7550A, HP 7440A.

Chapter 7

Testing for Changes in the Network Load

Establishing baseline performance information serves many purposes, one of which is looking at the effects of future growth. When evaluating the growth of a network, modeling tools are commonly used. One method is to introduce an independent controlled network load to observe the effects on network performance. Adding this simulated traffic and measuring the effects can provide a great deal of information on the impact of network growth and changes. When traffic simulation is in operation, it is important that the characteristic of the expected traffic growth is maintained. This includes the distribution of frame sizes and burst rates as well as the percentage increase.

The traffic generator in the LAN performance analysis system provides controlled amounts of simulated traffic on the network. It allows the user flexibility in defining the characteristics of the traffic. While the traffic generator is in operation, you can simultaneously make all the other received based measurements including network, connection, or node level performance measurements to determine the effects of the added traffic.

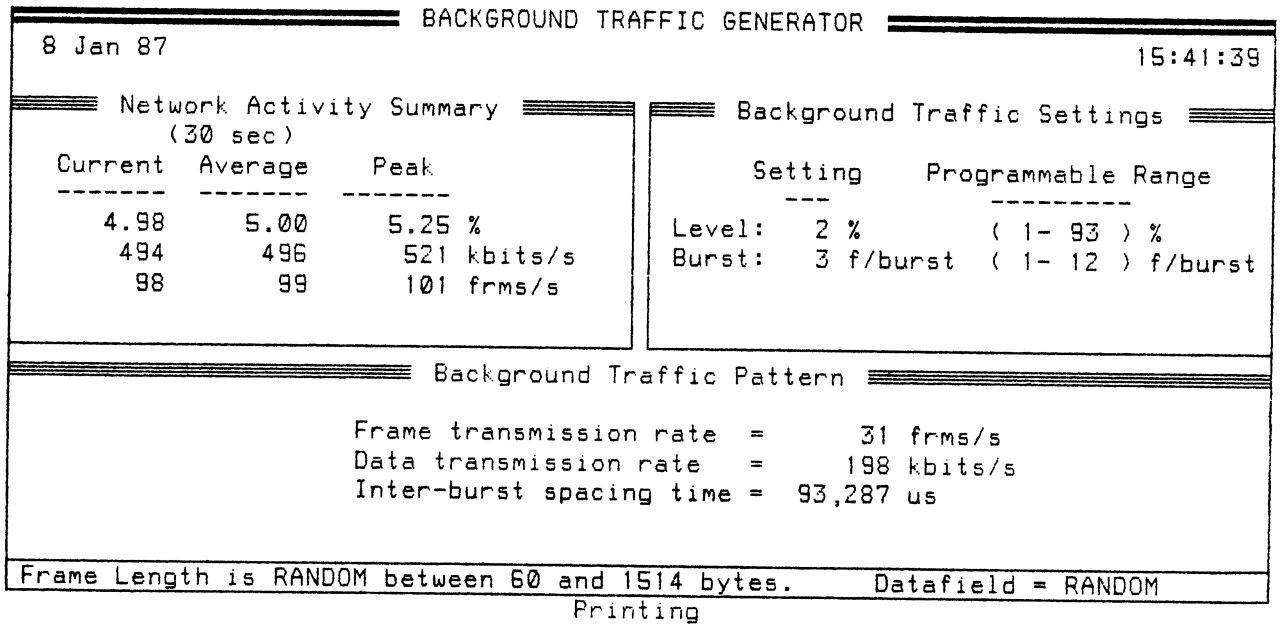

In the chart shown, the performance analysis system reports the current and average network utilization which includes the generated traffic. The background level utilization shows the amount of traffic added. You can also specify the level of traffic, the burst size and the type of traffic. The types of traffic include random length, fixed length, and user defined messages. The flexibility in the traffic generator allows you to accurately characterize the simulated traffic to predict future growth. While the traffic generator is operating on the network, the effect on the network can be shown by the network, connection and node level measurements.

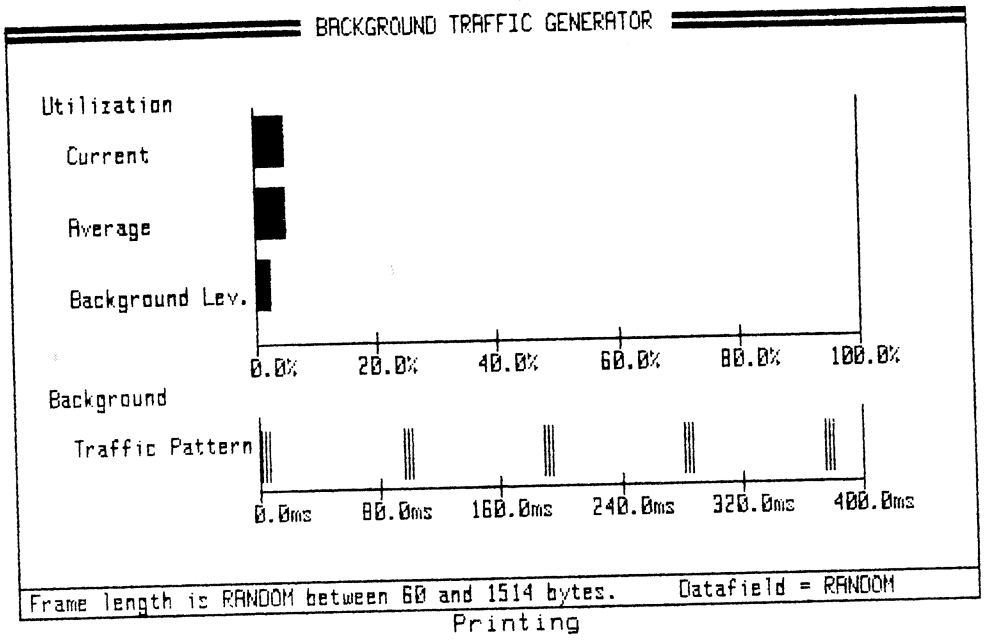

Abort
Print

Traffic generator information is also shown graphically. The first graph shows the current and average traffic level, on the network including the added traffic. It also shows the amount of traffic generated by the analysis system. The second graph shows the "bunching" of the frames sent on the network and the interburst time of the groups of frames.

Conclusion

Network performance analysis plays an important role in complex and growing networks. Performance information can be used to project growth and to understand network changes. Although there are many complex network management issues involved, the typical network manager needs to understand the fundamental measures to manage and control his network before tackling more complex problems. The simple guidelines in this product note can help provide good basic reliable service to the user group.

Appendix A

Network Utilization Calculation

The definition of 100% utilization in one second is being able to send as many bits as possible in one second without violating the frame length specification. Therefore we calculate using maximum length frames.

Max. length:

9.6 usec + 8 bytes preamble + 14 bytes header + 1500 bytes data + 4 byte FCS

9.6 usec + 1526 bytes => 9.6 usec + 12208 bits * 100 nsec/bit = 1.23 msec.

In one second there can be 1 sec/1.23 msec per max. frame = 812.74 max. frames.

There are also 812 slots of dead time (9.6 usec). We need to calculate the actual number of bits transmitted during 1 second by subtracting the dead time, converted to bits ((9.6 usec/100 nsec) * 812), from 10 million bits.

10000000 - (9.6 usec/100 nsec per bit) * 812 = 9922048 bits.

This is the definition for 100% utilization: transmitting 9922048 bits in one second without violating the frame length spec. If the network truly has 10000000 bits on it in one second, then the utilization would be over 100% (in fact 100.78%).

Let's calculate the utilization for 100 byte frames in one second:

8 bytes preamble + 100 bytes + 4 bytes (FCS) = 112 bytes

If the network has 50 frames of 100 byte frames (avg. length in one second), then the utilization of the network is

(112 * 8)* 50 = 44800 bits

Utilization = 44800/9922048 * 100% = 0.451%

Appendix B IEEE 802.2, 802.3 and Ethernet Specifications

Concise Specifications

Level 1: Physical layer

The physical layer consists of the physical signalling interface, Attachment Unit Interface (AUI) cable and the Medium Attachment Unit (MAU). The physical signalling interface is part of the Data Terminal Equipment (DTE) that converts the data from the upper layers to Manchester encoded data. The AUI cable has two 15-pin connectors on each end. The maximum length for this cable is 50 meters. The AUI cable connects the DTE to the MAU. Every node connects to the network through a MAU. MAUs, sometimes called transceivers, can only be placed in 2.5 meter increments apart. The MAU also detects collision while it is transmitting.

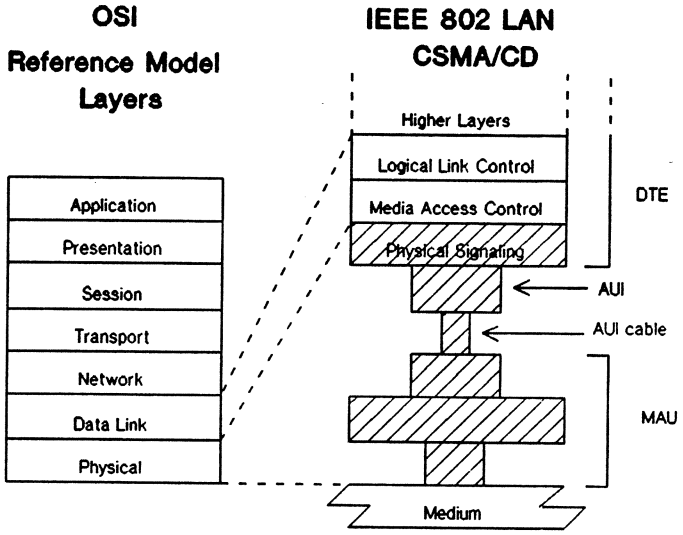

Physical Signalling Interface: This layer converts the data from the upper layers to Manchester encoded data before passing it on. The encoded data has the clock signal embedded in it. Since a Manchester encoded signal has a transition in the middle of every bit, it prevents the receive clock from losing synch when it encounters long strings of zeroes or ones. The second half of the bit time in the encoded signal is the true value of the bit. The following shows a Manchester encoded signal:

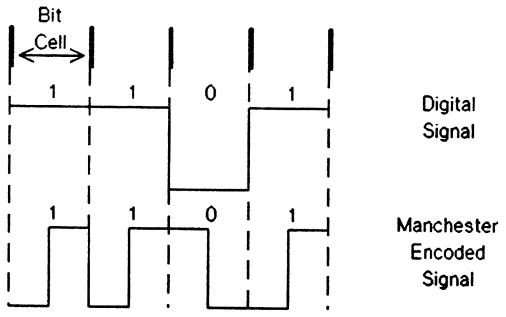

Physical Signalling Interface/AUI Interface: The interface supports one or more of the specified data rates in the IEEE 802.3 standard. (We are only interested in the 10 Mbps data rate.) It is capable of driving up to 50 meters of AUI cable. It also permits the DTE to test the AUI, AUI cable, MAU and the medium itself.

Attachment Unit Interface (AUI): The AUI carries encoded control and data signals. The signals are on balanced circuits. There are a total of four balanced circuits, data and control going out and data and control coming in. (Note: Control out is seldom implemented.) The timing information is encoded in the signals thus eliminating the need for separate timing circuits. The AUI operates in two modes: normal and monitor modes. In normal mode, the AUI is logically connected to the medium. The DTE is required to follow the media access algorithms, in this case CSMA/CD, to send data over

the AUI through the MAU onto the medium. The MAU always sends back the data it receives from the medium. In monitor mode, the MAU is logically disconnected from the medium and only functions as an observer on the medium.

Medium Attachment Unit (MAU): The MAU supports traffic rates at 10 Mbps. It will also drive up to 500 meters of coaxial trunk cable without the use of a repeater. It permits the DTE to test the MAU and the medium itself. It supports the CSMA/CD access mechanism defined with baseband signalling in the IEEE 802.3 standard. MAU functions are

Transmit function: To transmit serial data streams on the baseband medium from the local DTE entity to one or more remote DTE entities on the same network. To enable the collision presence function while transmitting.

Receive function: To receive serial data streams over the baseband medium.

Collision presence function: To detect the presence of two or more stations' concurrent transmissions. To detect collision while it is transmitting.

Monitor function: To inhibit the transmit function, but at the same time have the receive and collision presence functions operational.

Jabber function: To automatically shut off the transmitter if it exceeds the maximum allowable transmitting time - typically between 20 to 150 milliseconds.

Level 2: Data Link layer - Media Access Control IEEE 802.3/Ethernet

The media access control (MAC) sublayer interfaces with the physical layer and provides service for the Logical Link Control (LLC) to exchange data units with peer LLC sublayer entities.

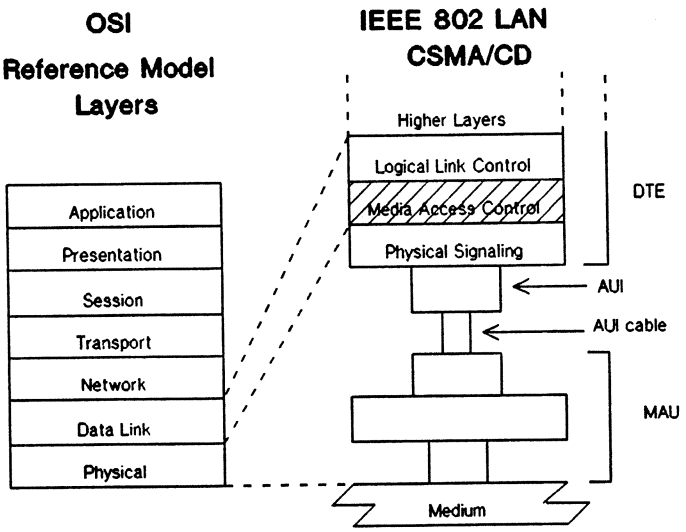

The MAC sublayers perform the following functions:

Data Encapsulation: This includes providing frame boundary delimitation and frame synchronization for transmitting and receiving data. It also handles source and destination addresses and error detection.

Media Access Management: The MAC sublayer accesses the medium using the CSMA/CD access method. It has the responsibility for avoiding and resolving collisions.

Carrier Sense Multiple Access with Collision Detect (CSMA/CD): CSMA/CD is the specified media access method for IEEE 802.3. When the MAC sublayer is ready to transmit a frame, it monitors the carrier sense signal from the physical layer to determine whether the media is being used in order to avoid a collision. When the medium is clear, it initiates frame transmission. The data is passed onto the physical layer for Manchester encoding and actual transmission. While transmitting, the physical layer also simultaneously monitors the medium for possible collision should other stations decide to transmit at the same time. When the transmission is completed without collision, the MAC sublayer informs the LLC sublayer and waits for the next request for frame transmission. In the event of a collision, the transmitting station's physical layer passes a collision detect signal to the MAC sublayer, The MAC sublayer then initiates the collision handling procedure by sending a jam sequence. This ensures that the duration of the collision is sufficient to be noticed by other stations involved in the collision. The MAC sublayer then terminates the transmission and waits for a randomly selected time interval before attempting to transmit again.

IEEE 802.3/Ethernet Frame Format

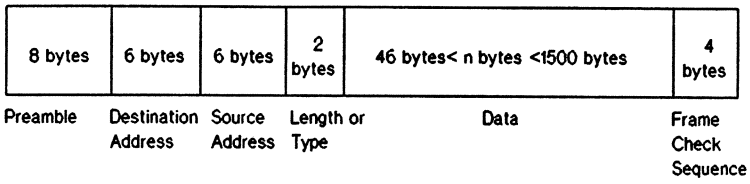

Each frame is viewed as a sequence of 8-bit bytes. The least significant bit (LSB) of each byte is always transmitted first except the frame check sequence (FCS).

Preamble: The preamble is a 64-bit synchronization pattern containing alternating 1s and 0s ending with two consecutive 1s. The last 1011 is called Start Frame Delimiter (SFD).
The preamble is

10101010101010101010101010101010101010101010101010101010101010101011

Destination Address: 6 bytes are designated for the destination address which specifies the station(s) to which the packet is being transmitted. The bit transmitted indicates the type of address.

0 - unique address of the one destination station
1 - a logical group of recipients
all 1s - broadcast address (all stations)

Source Address: 6 bytes are designated for the station that is transmitting the frame.

Type Field or Length Field: 2 bytes are designated to be the type field for Ethernet or length field for IEEE 802.3.

Type field (Ethernet) - This field is used to identify the higher level protocol type associated with the frame. It determines how the data field is to be interpreted. Examples of type fields are 0600H for XNS, 0800H for TCP/IP, and 9000H for Ethernet loopback protocol.

Length field (IEEE 802.3) - This field gives the number of data bytes in the data field. If the number of data bytes is less than the minimum required for proper operation of the protocol, a PAD field will be added at the end of the data field prior to the FCS field to make up a minimum size frame.

Data Field: This field contains an integral number of bytes ranging from 46 to 1500. If there aren't enough data bytes in the data field to meet the minimum requirement, pads will be added to make up for it. The reason for a minimum size packet is to distinguish valid packets from collision fragments

Frame Check Sequence: There are 4 bytes dedicated for the frame check sequence which uses a Cyclic Redunancy Check (CRC) code defined by the generating polynomial:

$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^{8} + x^{7} + x^{5} + x^{4} + x^{2} + x + 1$$

The CRC covers the address (destination/source), type or length and data field (including the pad field).

Maximum Frame Size: 1526 bytes (8 byte preamble + 14 byte header + 1500 data bytes + 4 byte CRC)

Minimum Frame Size: 72 bytes (8 byte preamble + 14 byte header + 46 data bytes + 4 byte CRC)

Level 2: Data Link layer - Logical Link Control IEEE 802.2

The Logical Link Control (LLC) sublayer provides the functionality for interfacing the data link layer (level 2) with the network layer (level 3). The LLC sublayer itself interfaces with the MAC sublayer, the network layer, and the LLC network management function. The LLC sublayer supports two types of data link control operations, connectionless and connection-oriented services.

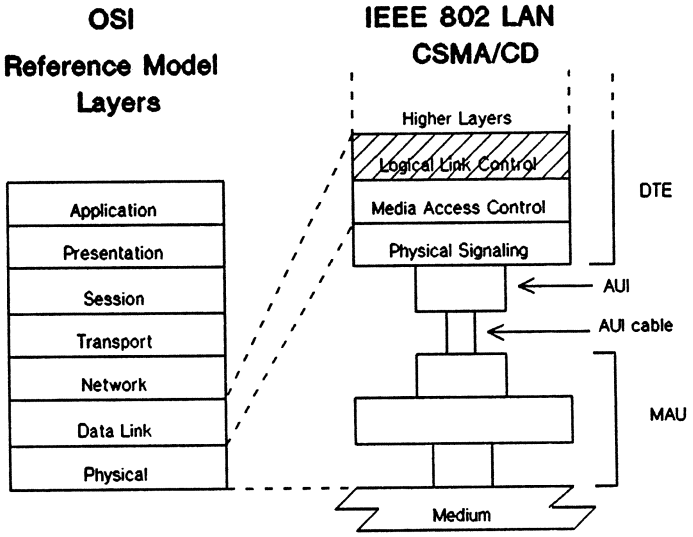

Unacknowledged Connectionless Service (datagram service): The data transfer service allows network entities to exchange link service data units without the establishment of a data link level connection. The data transfer can be point-to-point, multicast, or broadcast.

Connection-oriented Service (virtual circuit): This service provides the means for a network entity to request, or be notified of, the establishment of data link layer connections. It also provides the service for a network entity to send or receive link service data units over a data link layer connection.

Logical Link Control Frame Format

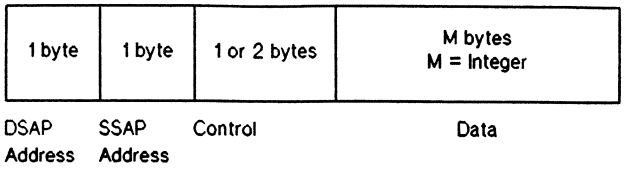

DSAP Address : Destination Service Access Point address is a one byte field which identifies the one, or more, service access points for which the LLC data field is intended. Only seven bits contain the actual address, the least significant bit indicates whether the address is a group or individual address.

SSAP Address : Source Service Access Point address is a one byte field which identifies the origin of the LLC data. Only seven bits contain the actual address, the least significant bit identifies whether the data information in the LLC frame is a command or response.

Control Field : This field contains one or two bytes used to designate command and response functions or sequence numbers for connectionless and connection-oriented data link services.

Data Field : The data field consists of an integral number of bytes.

Addresses and control sequences are delivered to or received from the media access control with the least significant bit first. The data sequence is delivered to the MAC in the same bit order as received from the network layer and vice versa.

Implementations of IEEE 802.3

There are four implementations of the IEEE 802.3 standard: 10BASE5, 10BASE2, 1BASE5, and 10BROAD36. Only the first two have been approved by the standards committe at this point. The last two implementations are still under discussion.

The notation used, for example, 10BASE5 denotes:

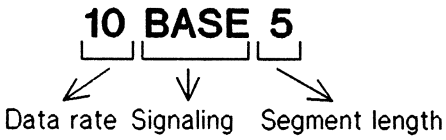

Therefore 10BASE5 stands for 10 Mbps, baseband signalling, 500 meters per segment of cable before a repeater is needed.

10BASE5 : This is the first implementation approved by the IEEE 802.3 standard committee. This is also the most common implementation for networking you will see today. Some of the 10BASE5 vendors are DEC (DECnet), Ungermann Bass (NetOne), Xerox (Ethernet), Interlan, 3Com.

10BASE2 : This implementation is sometimes called Cheapernet or Thin LAN. It also has a data rate of 10 Mbps and baseband signalling, but since the cable is an RG-58 coaxial cable, each segment can only be 250 meters. As the name indicates, it is less expensive to implement this network than the 10BASE5 network. The primary vendors for Thin LAN are 3Com's Cheapernet and HP's Officeshare network.

> *Note :* The HP LAN protocol analyzer supports both 10BASE5 and 10BASE2 (ThinLAN)

1BASE5 : This implemetation is sometimes called StarLAN. The configuration is a star/bus. The bus portion of the LAN is on a PC board. The star portion will be twisted pair. The primary vendor is AT&T**. This implementation is still under consideration by the IEEE 802.3 standards committee.

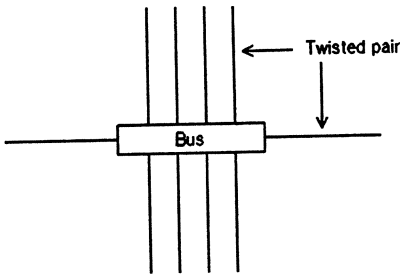

Starlan Configuration

10BROAD36 : This is a broadband implemntation of CSMA/CD. This proposal is under consideration of the IEEE 802.3 standards committee. The primary vendor is DEC*.

* DEC is a U.S. registered trademark of Digital Equipment Corp.
** AT&T is a U.S. registered trademark of American Telephone & Telegraph

HP Sales and Support Offices

For more information, call your local HP sales office listed in your telephone directory or an HP regional office listed below for the location of your nearest sales office.

United States:
Hewlett-Packard Company
4 Choke Cherry Road
Rockville, MD 20850
(301) 670-4300

Hewlett-Packard Company
5201 Tollview Dr.
Rolling Meadows, IL 60008
(312) 255-9800

Hewlett-Packard Company
5161 Lankershim Blvd.
No. Hollywood, CA 91601
(818) 505-5600

Hewlett-Packard Company
2015 South Park Place
Atlanta, GA 30339
(404) 955-1500

Canada:
Hewlett-Packard Ltd.
6877 Goreway Drive
Mississauga, Ontario L4V1M8
(416) 678-9430

Japan:
Yokogawa-Hewlett-Packard Ltd.
29-21, Takaido-Higashi 3-chome
Suginami-ku, Tokyo 168
(03) 331-6111

Latin America:
Latin American Region Headquarters
Monte Pelvoux Nbr 111
Lomas De Chapultepec
11000 Mexico, D.F. Mexico
(905) 596-79-33

Australia/New Zealand:
Hewlett-Packard Australia Ltd.
31-41 Joseph Street
Blackburn, Victoria 3130
Melbourne, Australia
(03) 895-2895

Far East:
Hewlett-Packard Asia Ltd.
22/F Bond Centre
West Tower
89 Queensway
Central, Hong Kong
(5) 8487777

Germany:
Hewlett-Packard GmbH
Vertriebzentrale Deutschland
Hewlett-Packard-Strasse
Postfach 1641
6380 Bad Homburg v.d.H.
Federal Republic of Germany
06172/400-0

France:
Hewlett-Packard France
Parc d'activité du Bois Briard
2, avenue du Lac
91040 EVRY Cedex, France
01/60 77 83 83

United Kingdom:
Hewlett-Packard Limited
Enquiry Group
Customer Support Centre
Eskdale Road
Winnersh Triangle
Wokingham
Berkshire RG11 5DZ
0734/69 66 22

Italy:
Hewlett-Packard Italiana S.p.A.
Via G. di Vittorio, 9
20063 Cernusco Sul Naviglio (MI)
Milano
02/923691

European Multi Country Region:
Hewlett-Packard S.A.
Route du Nant d'Avril 150
1217 Meyrin 2–Geneva
Switzerland
(41) 22/83 81 11

Or write to:

United States:
Hewlett-Packard Company
P.O. Box 10301,
Palo Alto, CA 94303-0890

Europe/Middle East/Africa:
Hewlett-Packard Company
Central Mailing Department
P.O. Box 529
1180 AM Amstelveen
The Netherlands

For all other areas:
Hewlett-Packard Company
Intercontinental Headquarters
3495 Deer Creek Rd.
Palo Alto, CA 94304
U.S.A.

HEWLETT PACKARD

Data subject to change

Printed in U.S.A. 2/89
5952-5103